



Geometry & Topology

Volume 30 (2026)

Issue 3 (pages 835–1201)



GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITORS

Robert Lipshitz University of Oregon
lipshitz@uoregon.edu
András I Stipsicz Alfréd Rényi Institute of Mathematics
stipsicz@renyi.hu

BOARD OF EDITORS

Mohammed Abouzaid	Stanford University abouzaid@stanford.edu	Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Dan Abramovich	Brown University dan_abramovich@brown.edu	Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu	Sándor Kovács	University of Washington skovacs@uw.edu
Arend Bayer	University of Edinburgh arend.bayer@ed.ac.uk	Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Agnès Beaudry	University of Colorado Boulder agnes.beaudry@colorado.edu	Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Mark Behrens	University of Notre Dame mbehren1@nd.edu	Jianfeng Lin	Tsinghua University linjian5477@mail.tsinghua.edu.cn
Mladen Bestvina	University of Utah bestvina@math.utah.edu	Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Martin R Bridson	University of Oxford bridson@maths.ox.ac.uk	Haynes R Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Tobias H Colding	Massachusetts Institute of Technology colding@math.mit.edu	Aaron Naber	Institute for Advanced Studies anaber@ias.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk	Peter Ozsváth	Princeton University petero@math.princeton.edu
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu	Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
David M Fisher	Rice University davidfisher@rice.edu	Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
Mike Freedman	Microsoft Research michaelf@microsoft.com	Roman Sauer	Karlsruhe Institute of Technology roman.sauer@kit.edu
David Gabai	Princeton University gabai@princeton.edu	Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de	Natasa Sesum	Rutgers University natasas@math.rutgers.edu
Cameron Gordon	University of Texas gordon@math.utexas.edu	Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Jesper Grodal	University of Copenhagen jg@math.ku.dk	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Misha Gromov	IHES and NYU, Courant Institute gromov@ihes.fr	Kirsten Wickelgren	Duke University kirsten.wickelgren@duke.edu
Mark Gross	University of Cambridge mgross@dpms.cam.ac.uk	Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2026 is US \$905/year for the electronic version, and \$1275/year (+\$80, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY



mathematical sciences publishers
nonprofit scientific publishing

<http://msp.org/>

© 2026 Mathematical Sciences Publishers

The level structure in quantum K -theory and mock theta functions

YONGBIN RUAN, MING ZHANG AND YAOXIONG WEN

This is the first in a sequence of papers to develop the theory of levels in quantum K -theory and study its applications. Here, we give an adelic characterization of the range of the J -function in quantum K -theory with level structure. As an application, we prove a mirror theorem for permutation-equivariant quantum K -theory with level structures of toric varieties. In the study of the mirrors of some simple examples, we see the surprising appearance of Ramanujan's mock theta functions.

1 Introduction

1.1 Overview

More than a decade ago, quantum K -theory was introduced in [Givental 2000; Lee 2004] as the K -theoretic analog of quantum cohomology. Its recent revival stems partially from a physical interpretation of quantum K -theory as a 3D-quantum field theory in a 3-manifold of the form $S^1 \times \Sigma$ (see [Kapustin and Willett 2013; Jockers and Mayr 2020]). Because of this mysterious physical connection, the B-model counterpart of quantum K -theory is q -hypergeometric series, itself a classical subject. This connection was recently confirmed by Givental [2015e] as the mirror of the so-called J -function of the permutation-equivariant quantum K -theory.

Classically, K -theory is more closely related to representation theory, comparing to cohomology theory. It is natural to revisit quantum K -theory from the representation-theoretic point of view. In fact, a variant of quantum K -theory was already studied in [Okounkov 2017; Aganagic et al. 2018; Aganagic and Okounkov 2017; Pushkar et al. 2020; Koroteev et al. 2021; Okounkov and Smirnov 2022; Su et al. 2020] in relation to quantum groups. One of the predominant features of representation theory is the existence of an additional parameter, called the level. A natural question is whether it is possible to extend the current version of quantum K -theory to include this notion of level. In this paper, we answer the question affirmatively in the context of quasimap theory. This is the first in a sequence of papers to develop the theory of levels in quantum K -theory and study its applications.

Our motivating example is an old physical result [Witten 1995], which relates the quantum cohomology ring of the Grassmannian to the Verlinde algebra. Early explicit physical computations [Gepner 1991; Vafa 1992; Intriligator 1991] indicate that they are isomorphic as algebras, but have different pairings. Witten [1995] gave a conceptual explanation of the isomorphism, by proposing an equivalence between the quantum field theories which govern the Verlinde algebra and the quantum cohomology of the Grassmannian. His physical derivation of the equivalence naturally leads to a mathematical problem that

these two objects are conceptually isomorphic (without referring to the detailed computations). A great deal of work has been done in [Agnihotri 1995; Marian and Oprea 2007a; 2007b; 2010; Belkale 2008]. However, to the best of our knowledge, a complete conceptual proof of the equivalence is missing.

Assuming a basic knowledge of quantum K -theory, a key and yet more or less trivial observation is that the Verlinde algebra is a K -theoretic invariant. To be more precise, let X be a smooth projective variety. Suppose that $\overline{\mathcal{M}}_{g,k}(X, \beta)$ is the moduli space of stable maps to X . Quantum cohomology studies integrals of the form

$$\int_{[\overline{\mathcal{M}}_{g,k}(X, \beta)]^{\text{vir}}} \alpha,$$

where $[\overline{\mathcal{M}}_{g,k}(X, \beta)]^{\text{vir}}$ is the so-called *virtual fundamental cycle* and α is a “tautological” cohomology class. In quantum K -theory, we replace the virtual fundamental cycle by the *virtual structure sheaf* $\mathcal{O}_{\overline{\mathcal{M}}_{g,k}(X, \beta)}^{\text{vir}}$. We also replace the integral by the holomorphic Euler characteristic

$$\chi(\overline{\mathcal{M}}_{g,k}(X, \beta), E \otimes \mathcal{O}_{\overline{\mathcal{M}}_{g,k}(X, \beta)}^{\text{vir}}),$$

where E is some natural K -theory class on $\overline{\mathcal{M}}_{g,k}(X, \beta)$. For the Verlinde algebra, the relevant moduli space is the moduli space of semistable parabolic $U(n)$ -bundles $\mathcal{M}_{U(n)}(\alpha_1, \dots, \alpha_k)$ on a fixed genus- g marked curve $(C, p_1, \dots, p_k)$ with parabolic structure at p_i indexed by $\alpha_i \in V_l(U(n))$. Here, l is a nonnegative integer and $V_l(U(n))$ denotes the level- l Verlinde algebra. A new ingredient is a certain determinant line bundle, denoted by $\det$, over the moduli space $\mathcal{M}_{U(n)}(\alpha_1, \dots, \alpha_k)$. The level- l Verlinde algebra calculates the holomorphic Euler characteristic

$$\langle \alpha_1, \dots, \alpha_k \rangle_{g, \text{Verlinde}}^l = \chi(\mathcal{M}_{U(n)}(\alpha_1, \dots, \alpha_k), \det^l).$$

Based on the above description, the Verlinde algebra is clearly a K -theoretic object, and we should compare it with the quantum K -theory of the Grassmannian (with an appropriate notion of levels). Let $\mathfrak{Bun}_G$ be the moduli stack of principal bundles over curves. Let $\pi : \mathfrak{E}_{\mathfrak{Bun}_{g,k}} \rightarrow \mathfrak{Bun}_G$ be the universal curve and let $\mathfrak{P} \rightarrow \mathfrak{E}_{\mathfrak{Bun}_{g,k}}$ be the universal principal bundle. Given a finite-dimensional representation R of G , we consider the inverse determinant of cohomology

$$\det_R := (\det R\pi_*(\mathfrak{P} \times_G R))^{-1}.$$

This is a line bundle over $\mathfrak{Bun}_G$. Suppose $X = Z//G$ is a GIT quotient. Let $\mathcal{Q}^\epsilon$ be the moduli stack of ϵ -stable quasimaps to $X = Z//G$ (see Section 2.2). Then there is a natural forgetful morphism $\mu : \mathcal{Q}^\epsilon \rightarrow \mathfrak{Bun}_G$. We define the level- l determinant line bundle as

$$\mathcal{D}^{R,l} = \mu^*(\det_R)^l.$$

We will often refer to $\mathcal{D}^{R,l}$ as the *level structure*. In general, when X is a smooth complex projective variety but not a GIT quotient, one can still define determinant line bundles over the moduli space of stable maps $\overline{\mathcal{M}}_{g,k}(X, \beta)$ as follows. Let $\mathcal{R}$ be a vector bundle over X . Let $\pi : \mathcal{C} \rightarrow \overline{\mathcal{M}}_{g,k}(X, \beta)$ be the

universal curve and let $\text{ev} : \mathcal{C} \rightarrow X$ be the universal evaluation map. We define the level- l determinant line bundle as

$$\mathcal{D}^l := (\det R\pi_*(\text{ev}^* \mathcal{R}))^{-l}.$$

This definition agrees with the previous one when X is a GIT quotient (see Definition 2.4 and Remark 2.5).

With the above definition of the level- l determinant line bundle $\mathcal{D}^{R,l}$, we can define the level- l quantum K -invariants and quasimap invariants by twisting with $\mathcal{D}^{R,l}$ (see Section 3). The ordinary quantum K -theory corresponds to the case $l = 0$. One of the main purposes of this article is to verify that the level- l quantum K -theory satisfies all the axioms of the ordinary quantum K -theory (see Section 2.3).

One can use irreducible $\text{GL}_n(\mathbb{C})$ -representations to define two sets of invariants: Verlinde invariants from the theory of (semi)stable parabolic vector bundles, and level- l quantum K -theory invariants of the Grassmannian $\text{Gr}(n, n + l)$. The *K-theoretic Verlinde/Grassmannian correspondence* describes a conjectural formula between these two sets of invariants (see [Ruan and Zhang 2023] for the precise formula).

1.2 Mirror theorem and mock theta functions

The proof of the Verlinde/Grassmannian correspondence will be discussed in different papers. The other main results of this paper are various mirror theorems, in the same style as the recent work of Givental [2017a; 2015a; 2015b; 2015c; 2015d; 2015e; 2015f; 2015g; 2017b; 2020; 2017c]. In the study of quantum K -theory with level structures, we see the surprising appearance of Ramanujan's mock theta functions in some of the simplest examples.

Let X be a smooth complex projective variety and let $\mathcal{R}$ be a vector bundle over X . When $X = Z // G$ is a GIT quotient, we assume $\mathcal{R}$ is of the form $(Z \times R) // G$, with R a finite-dimensional representation of G . By abuse of terminology, we refer to the vector bundle $\mathcal{R}$ as the “representation” R . Let Q be the Novikov variables. We fix a λ -algebra Λ which is equipped with Adams operations Ψ^i for $i = 1, 2, \dots$. Let $\{\phi_a\}$ be a basis of $K^0(X) \otimes \mathbb{Q}$ and let $\{\phi^a\}$ be the dual basis with respect to the twisted pairing (5). Let q be a formal variable and let $\mathbf{t}(q)$ be a Laurent polynomial in q with coefficients in $K^0(X) \otimes \mathbb{Q}$. The *permutation-equivariant K-theoretic J-function* $\mathcal{J}_{S_\infty}^l(\mathbf{t}(q), Q)$ of level l and representation R is defined by

$$\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t}(q), Q) := 1 - q + \mathbf{t}(q) + \sum_a \sum_{\beta \neq 0} Q^\beta \phi^a \left\langle \frac{\phi_a}{1 - qL}, \mathbf{t}(L), \dots, \mathbf{t}(L) \right\rangle_{0,k+1,\beta}^{R,l,S_k}.$$

Here $\langle \cdot \rangle_{0,k+1,\beta}^{R,l,S_k}$ denotes the permutation-equivariant quantum K -invariants of level l and L denote the cotangent line bundles. The J -function $\mathcal{J}_{S_\infty}^{R,l}$ can be viewed as elements in the loop space $\mathcal{K}$ defined by

$$\mathcal{K} := [K^0(X) \otimes \mathbb{C}(q)] \otimes \mathbb{C}[[Q]],$$

where $\mathbb{C}(q)$ denotes the field of complex rational functions in q . There is a natural Lagrangian polarization $\mathcal{K} = \mathcal{K}_+ \oplus \mathcal{K}_-$, where $\mathcal{K}_+$ consists of Laurent polynomials and $\mathcal{K}_-$ consists of reduced rational functions

that are regular at $q = 0$ and vanishing at $q = \infty$. We denote by $\mathcal{L}_{S_\infty}^{R,l}$ the range of $\mathcal{J}_{S_\infty}^{R,l}$, i.e.,

$$\mathcal{L}_{S_\infty}^{R,l} = \bigcup_{t(q) \in \mathcal{K}_+} \mathcal{J}_{S_\infty}^{R,l}(t(q), Q) \subset \mathcal{K}.$$

Due to the stacky structure of the moduli space of stable maps, the J -function $\mathcal{J}_{S_\infty}^{R,l}(t(q), Q)$, as a function in q , has poles at roots of unity. The main technical tool is a generalization of Givental and Tonita's *adelic characterization* of points on the cone $\mathcal{L}_{S_\infty}^{R,l}$ with the presence of level structure, i.e., we describe the Laurent expansion of $\mathcal{J}_{S_\infty}^{R,l}$ at each primitive root of unity in terms of a certain twisted *fake* quantum K -theory. The precise statement is rather technical, and we present it in Theorem 4.5. When $l = 0$, the determinant line bundle $\mathcal{D}^{R,l}$ is trivial, and we recover the ordinary quantum K -theory. The adelic characterization of the cone in the ordinary quantum K -theory was introduced in [Givental and Tonita 2014], and its generalization to the permutation-equivariant theory is given in [Givental 2015b]. The proofs of all these results are based on application of the virtual Kawasaki Riemann–Roch formula to moduli spaces of stable maps (see Section 4).

Let $\mathcal{L}_{S_\infty}$ denote the range of the permutation-equivariant big J -function in ordinary quantum K -theory (i.e., with trivial level structure). As an application of Theorem 4.5, we prove that certain “determinantal” modifications of points on $\mathcal{L}_{S_\infty}$ lie on the cone $\mathcal{L}_{S_\infty}^{R,l}$ of quantum K -theory of level l .

Theorem 1.1 *If*

$$I = \sum_{\beta \in \text{Eff}(X)} I_\beta Q^\beta$$

lies on $\mathcal{L}_{S_\infty}$, then the point

$$I^{R,l} := \sum_{\beta \in \text{Eff}(X)} I_\beta Q^\beta \prod_i (L_i^{-\beta_i} q^{(\beta_i+1)\beta_i/2})^l$$

lies on the cone $\mathcal{L}_{S_\infty}^{R,l}$ of permutation-equivariant quantum K -theory of level l . Here $\text{Eff}(X)$ denotes the semigroup of effective curve classes on X , L_i are the K -theoretic Chern roots of $\mathcal{R}$, and $\beta_i := \int_\beta c_1(L_i)$.

In Proposition 5.1, we give explicit formulas for level- l (torus-equivariant) I -functions of toric varieties. Moreover, we prove the following toric mirror theorem:

Theorem 1.2 *Assume that X is a smooth quasiprojective toric variety. The level- l torus-equivariant I -function $(1-q)I_X^{R,l}$ of X lies on the cone $\mathcal{L}_{S_\infty}^{R,l}$ in the permutation- and torus-equivariant quantum K -theory of level l of X .*

In the study of toric mirror theorems for quantum K -theory with level structure, a remarkable phenomenon is the appearance of Ramanujan's mock theta functions. We first establish some notation. Denote the standard representation of G by St and its dual by $\text{St}^\vee$. When $G = \mathbb{C}^*$, any n -dimensional representation of G is determined by a *charge vector* $(a_1, \dots, a_n)$ with $a_i \in \mathbb{Z}$: a $\mathbb{C}^*$ -action on $\mathbb{C}^n$ can be explicitly described by

$$\lambda \cdot (x_1, \dots, x_n) \rightarrow (\lambda^{a_1} x_1, \dots, \lambda^{a_n} x_n), \quad \text{where } \lambda \in \mathbb{C}^*.$$

In the following propositions, we consider GIT quotients $\mathbb{C}^n // \mathbb{C}^*$, and we refer to the $\mathbb{C}^*$ -actions by their associated charge vectors.

Proposition 1.3 Consider $X = \mathbb{C} // \mathbb{C}^* = [(\mathbb{C} \setminus 0) / \mathbb{C}^*]$, where the $\mathbb{C}^*$ -action is the standard action by multiplication. The $\mathbb{C}^*$ -equivariant K -ring $K_{\mathbb{C}^*}(X)$ is isomorphic to the representation ring $\text{Repr}(\mathbb{C}^*)$. Let $\lambda \in K_{\mathbb{C}^*}(X)$ be the equivariant parameter corresponding to the standard representation. For the $\mathbb{C}^*$ -representations St and $\text{St}^\vee$, we have the explicit formulas of the equivariant small I -functions

$$I_X^{\text{St}, l}(q, Q) = 1 + \sum_{n \geq 1} \frac{q^{n(n-1)l/2}}{(1 - \lambda^{-1}q)(1 - \lambda^{-1}q^2) \cdots (1 - \lambda^{-1}q^n)} Q^n,$$

$$I_X^{\text{St}^\vee, l}(q, Q) = 1 + \sum_{n \geq 1} \frac{q^{n(n+1)l/2}}{(1 - \lambda^{-1}q)(1 - \lambda^{-1}q^2) \cdots (1 - \lambda^{-1}q^n)} Q^n.$$

By choosing certain specializations of the parameters, we obtain Ramanujan's mock theta functions of order 3,

$$I_X^{\text{St}, l=1}(q^2, Q)|_{\lambda=-1, Q=q} = 1 + \sum_{n \geq 1} \frac{q^{n^2}}{(1 + q^2)(1 + q^4) \cdots (1 + q^{2n})},$$

$$I_X^{\text{St}, l=1}(q^2, Q)|_{\lambda=q, Q=q} = 1 + \sum_{n \geq 1} \frac{q^{n^2}}{(1 - q)(1 - q^3) \cdots (1 - q^{2n-1})},$$

$$I_X^{\text{St}, l=1}(q^2, Q)|_{\lambda=-q, Q=1} = 1 + \sum_{n \geq 1} \frac{q^{n(n-1)}}{(1 + q)(1 + q^3) \cdots (1 + q^{2n-1})},$$

and Ramanujan's mock theta functions of order 5,

$$I_X^{\text{St}, l=2}(q, Q)|_{\lambda=-1, Q=q} = 1 + \sum_{n \geq 1} \frac{q^{n^2}}{(1 + q)(1 + q^2) \cdots (1 + q^n)},$$

$$I_X^{\text{St}, l=2}(q^2, Q)|_{\lambda=q, Q=q^2} = 1 + \sum_{n \geq 1} \frac{q^{2n^2}}{(1 - q)(1 - q^3) \cdots (1 - q^{2n-1})},$$

$$I_X^{\text{St}^\vee, l=2}(q, Q)|_{\lambda=-1, Q=1} = 1 + \sum_{n \geq 1} \frac{q^{n(n+1)}}{(1 + q)(1 + q^2) \cdots (1 + q^n)},$$

$$I_X^{\text{St}^\vee, l=2}(q^2, Q)|_{\lambda=q, Q=1} = 1 + \sum_{n \geq 1} \frac{q^{2n^2+2n}}{(1 - q)(1 - q^3) \cdots (1 - q^{2n-1})}.$$

Proposition 1.4 Let a_1 and a_2 be two positive integers which are coprime. We consider the target $X_{a_1, a_2} = [(\mathbb{C}^2 \setminus 0) / \mathbb{C}^*]$ with charge vector (a_1, a_2) and a line bundle $p = \{(\mathbb{C}^2 \setminus 0) \times \mathbb{C}\} / \mathbb{C}^*$ with charge vector $(a_1, a_2, 1)$. Let λ_1 and λ_2 be the equivariant parameters. For the $\mathbb{C}^*$ -representations St

and $\text{St}^\vee$, we have the explicit formulas of the equivariant small I -functions

$$I_{X_{a_1, a_2}}^{\text{St}, l}(q, Q) = 1 + \sum_{n \geq 1} \frac{p^{nl} q^{n(n-1)l/2}}{(1 - p^{a_1} \lambda_1^{-1} q) \cdots (1 - p^{a_1} \lambda_1^{-1} q^{a_1 n}) (1 - p^{a_2} \lambda_2^{-1} q) \cdots (1 - p^{a_2} \lambda_2^{-1} q^{a_2 n})} Q^n,$$

$$I_{X_{a_1, a_2}}^{\text{St}^\vee, l}(q, Q) = 1 + \sum_{n \geq 1} \frac{p^{nl} q^{n(n+1)l/2}}{(1 - p^{a_1} \lambda_1^{-1} q) \cdots (1 - p^{a_1} \lambda_1^{-1} q^{a_1 n}) (1 - p^{a_2} \lambda_2^{-1} q) \cdots (1 - p^{a_2} \lambda_2^{-1} q^{a_2 n})} Q^n.$$

By choosing $(a_1, a_2) = (1, 1)$ and certain specializations of the parameters, we obtain the following four Ramanujan's mock theta functions of order 3:

$$I_{X_{1,1}}^{\text{St}, l=2}(q, Q)|_{p=1, \lambda_1=\lambda_2=-1, Q=q} = 1 + \sum_{n \geq 1} \frac{q^{n^2}}{((1+q)(1+q^2) \cdots (1+q^n))^2},$$

$$I_{X_{1,1}}^{\text{St}, l=2}(q, Q)|_{p=1, \lambda_1=(1+\sqrt{3}i)/2, \lambda_2=(1-\sqrt{3}i)/2, Q=q} = 1 + \sum_{n \geq 1} \frac{q^{n^2}}{(1-q+q^2)(1-q^2+q^4) \cdots (1-q^n+q^{2n})},$$

$$\frac{1}{(1-q)^2} I_{X_{1,1}}^{\text{St}^\vee, l=2}(q^2, Q)|_{p=1, \lambda_1=\lambda_2=q^{-1}, Q=1} = \sum_{n \geq 0} \frac{q^{2n^2+2n}}{((1-q)(1-q^3) \cdots (1-q^{2n+1}))^2},$$

$$\frac{1}{1+q+q^2} I_{X_{1,1}}^{\text{St}^\vee, l=2}(q^2, Q)|_{p=1, \lambda_1=(-1+\sqrt{3}i)/2, \lambda_2=(-1-\sqrt{3}i)/2, Q=1}$$

$$= \sum_{n \geq 0} \frac{q^{2n^2+2n}}{(1+q+q^2)(1+q^3+q^6) \cdots (1+q^{2n+1}+q^{4n+2})}.$$

Proposition 1.5 Let a and b be two positive integers which are coprime. We consider the target $X_{a,-b} = [(\mathbb{C} \setminus 0) \times \mathbb{C}] / \mathbb{C}^*$ with charge vector $(a, -b)$ and a line bundle $p = [(\mathbb{C} \setminus 0) \times \mathbb{C} \times \mathbb{C}] / \mathbb{C}^*$ with charge vector $(a, -b, 1)$. Let λ and μ be the equivariant parameters of the standard $(\mathbb{C}^*)^2$ -action on $X_{a,-b}$. For the $\mathbb{C}^*$ -representation St , we have the explicit formula for the equivariant small I -function

$$I_{X_{a,-b}}^{\text{St}, l}(q) =$$

$$1 + \sum_{n \geq 1} (-1)^{bn} \frac{p^{nl-b^2n} q^{n(n-1)l-bn(bn-1)/2} \mu^{-bn} (1-p^b \mu) (1-p^b \mu q) \cdots (1-p^b \mu q^{bn-1})}{(1-p^a \lambda^{-1} q) (1-p^a \lambda^{-1} q^2) \cdots (1-p^a \lambda^{-1} q^{an})} Q^n.$$

In particular, we have order 7 mock theta functions

$$I_{X_{2,-1}}^{\text{St}, l=3}(q, Q)|_{p=1, \lambda=1, \mu=q, Q=-q^2} = 1 + \sum_{n \geq 1} \frac{q^{n^2}}{(1-q^{n+1}) \cdots (1-q^{2n})},$$

$$\frac{q}{1-q} I_{X_{2,-1}}^{\text{St}, l=3}(q, Q)|_{p=1, \lambda=q^{-1}, \mu=q, Q=-q^4} = \sum_{n \geq 1} \frac{q^{n^2}}{(1-q^n) \cdots (1-q^{2n-1})},$$

$$\frac{1}{1-q} I_{X_{2,-1}}^{\text{St}, l=3}(q, Q)|_{p=1, \lambda=q^{-1}, \mu=q, Q=-q^3} = \sum_{n \geq 1} \frac{q^{n^2-n}}{(1-q^n) \cdots (1-q^{2n-1})}.$$

Remark 1.6 The targets that we consider in Propositions 1.4 and 1.5 are, in general, orbifolds. The full I -functions have components corresponding to the twisted sectors of the (rigidified) inertia stacks of the orbifold targets. However, to match with mock theta functions, we only consider the components of I -functions corresponding to the untwisted sectors. See the discussion on the orbifold case in Remarks 4.14 and 4.15.

It is interesting that we can recover Ramanujan's mock theta functions using only very simple targets.

One of the attractive features of quantum K -theory is the appearance of q -hypergeometric series as mirrors of K -theoretic J -functions. Recall the definition of the q -Pochhammer symbol

$$(a; q)_n := (1 - a)(1 - aq) \cdots (1 - aq^{n-1}) \quad \text{for } n > 0,$$

and $(a; q)_0 := 1$. A general q -hypergeometric series can be written as

$${}_r\phi_s = \sum_{n \geq 0} \frac{(\alpha_1; q)_n \cdots (\alpha_r; q)_n}{(\beta_1; q)_n \cdots (\beta_s; q)_n} \frac{z^n}{(q; q)_n} [(-1)^n q^{n(n-1)/2}]^{1+s-r}.$$

For the quantum K -theory of level 0, i.e., Givental and Lee's quantum K -theory, we only see special q -hypergeometric series of the form

$$\sum_{n \geq 0} \frac{(\alpha_1; q)_n \cdots (\alpha_r; q)_n}{(\beta_1; q)_n \cdots (\beta_s; q)_n} \frac{z^n}{(q; q)_n}.$$

The level structure naturally introduces the term

$$[(-1)^n q^{n(n-1)/2}]^{1+s-r}.$$

Proposition 1.7 Consider the target $X_{1,-1} := O(-1)_{\mathbb{P}^s}^{\oplus r} = [\{(\mathbb{C}^{s+1} \setminus 0) \times \mathbb{C}^r\} / \mathbb{C}^*]$ with the charge vector $(1, 1, \dots, 1, -1, -1, \dots, -1)$. Let $p = [\{(\mathbb{C}^{s+1} \setminus 0) \times \mathbb{C}^r \times \mathbb{C}\} / \mathbb{C}^*]$ be a line bundle with charge vector $(1, \dots, 1, -1, \dots, -1, 1)$. Let $\lambda_1, \dots, \lambda_{s+1}, \mu_1, \dots, \mu_r$ be the equivariant parameters of the standard $(\mathbb{C}^*)^{s+r+1}$ -action on $X_{1,-1}$. Then the equivariant small I -function has the explicit form

$$I_{X_{1,-1}}^{\text{St}, l=1+s}(q) = 1 + \sum_{n \geq 1} (-1)^{nr} \prod_{i=1}^r (p\mu_i)^{-n} p^{(1+s)n} \frac{(p\mu_1, q)_n \cdots (p\mu_r, q)_n}{(p\lambda_1^{-1}q; q)_n \cdots (p\lambda_s^{-1}q; q)_n} \frac{Q^n}{(p\lambda_{s+1}^{-1}q; q)_n} (q^{n(n-1)/2})^{1+s-r}.$$

Hence we can recover the general q -hypergeometric series by setting $p = 1$, $\lambda_i^{-1}q = \beta_i$ for $i = 1, \dots, s$, $\lambda_{s+1}^{-1} = 1$ and $\mu_j = \alpha_j$ for $j = 1, \dots, r$, $Q = (-1)^{1+s} z \prod_{i=1}^r \mu_i$.

Recall that Gromov–Witten theory (of Calabi–Yau varieties) is related to quasimodular forms. Mock modular forms are another class of modular objects, which are different from the quasimodular forms. Yet, they share some common properties. The above mirror theorems suggest an exciting possibility: that the natural geometric home of mock modular forms is quantum K -theory with nontrivial level structures. We certainly would like to investigate this surprising connection further.

Plan of the paper

In Section 2, we introduce the notion of level in the K -theoretic quasimap theory and establish its main properties. In Section 3, we define the K -theoretic quasimap invariants with level structure and their permutation-equivariant version. In Section 4, we characterize the cone $\mathcal{L}_{S_\infty}^{R,l}$ of level l in terms of the cone $\mathcal{L}_{S_\infty}$ of level 0. In Section 5, we compute the K -theoretic toric I -functions of level l , and prove a toric mirror theorem for quantum K -theory with level structure.

2 Level structure

In this section, we assume, unless otherwise indicated, that X is given by a geometric invariant theory (GIT) quotient. As mentioned in the introduction, the level structure is defined by determinant line bundles. We first recall the definition of determinant line bundles and then define the notion of level in quasimap theory.

2.1 Determinant line bundles

In this subsection, we briefly review the construction of determinant line bundles.

Let $\mathcal{X}$ be a Deligne–Mumford stack. Let $\mathcal{E}$ be a locally free, finitely generated $\mathcal{O}_{\mathcal{X}}$ module. We define the determinant line bundle of $\mathcal{E}$ as

$$\det(\mathcal{E}) := \bigwedge^{\operatorname{rank}(\mathcal{E})} \mathcal{E},$$

where $\bigwedge^i$ denotes the i -th wedge product. In general, let $\mathcal{F}^\bullet$ be a complex of coherent sheaves on $\mathcal{X}$ which has a bounded locally free resolution, i.e., there exists a bounded complex of locally free, finitely generated $\mathcal{O}_{\mathcal{X}}$ modules $\mathcal{G}^\bullet$ and a quasi-isomorphism

$$\mathcal{G}^\bullet \rightarrow \mathcal{F}^\bullet.$$

We define the determinant line bundle associated to $\mathcal{F}^\bullet$ by

$$\det(\mathcal{F}^\bullet) := \bigotimes_n \det(\mathcal{G}^n)^{(-1)^n}.$$

We summarize some basic properties of this construction in the following proposition:

Proposition 2.1 *Let $\mathcal{F}$ be a complex of coherent sheaves which has a bounded locally free resolution. Then:*

- (1) *The construction of $\det(\mathcal{F}^\bullet)$ does not depend on the locally free resolution.*
- (2) *For every short exact sequence of complexes of sheaves which have bounded locally free resolutions*

$$0 \rightarrow \mathcal{F}^\bullet \xrightarrow{\alpha} \mathcal{G}^\bullet \xrightarrow{\beta} \mathcal{H}^\bullet \rightarrow 0,$$

we have a functorial isomorphism

$$i(\alpha, \beta) : \det(\mathcal{F}^\bullet) \otimes \det(\mathcal{H}^\bullet) \xrightarrow{\sim} \det(\mathcal{G}^\bullet).$$

- (3) The operator $\det$ commutes with base change. To be more precise, for every (representable) morphism of Deligne–Mumford stacks

$$g : \mathcal{X} \rightarrow \mathcal{Y},$$

we have an isomorphism

$$\det(Lg^*) \xrightarrow{\sim} g^* \det.$$

In the case when $\mathcal{X}$ is a scheme, the above proposition is proved in [Knudsen and Mumford 1976]. Note that these properties are preserved under flat base change. Therefore they hold for stacks as well.

2.2 Level structure in quasimap theory

In this subsection, we first recall the quasimap theory for nonsingular GIT quotients introduced in [Ciocan-Fontanine et al. 2014]. Then we define the level structure in this setting and discuss its generalizations in orbifold quasimap theory.

Let $Z = \operatorname{Spec}(A)$ be a complex affine algebraic variety in $\mathbb{C}^n$ and let G be a reductive group acting on it. Let $\theta : G \rightarrow \mathbb{C}^*$ be a character determining a G -equivariant line bundle $L_\theta := Z \times \mathbb{C}$. Let $Z^s(\theta)$ and $Z^{ss}(\theta)$ be the stable and semistable loci, respectively. Throughout the paper, we assume $Z^s(\theta) = Z^{ss}(\theta)$ is nonsingular. Furthermore, we assume that G acts freely on $Z^s(\theta)$. It follows that the GIT quotient $Z//_\theta G$ is nonsingular and quasiprojective. For simplicity, we drop θ from the notation of the GIT quotient. The unstable locus is defined as $Z_{\text{us}} := Z - Z^s(\theta)$. Recall that we can identify the G -equivariant Picard group $\operatorname{Pic}^G(Z)$ with the Picard group $\operatorname{Pic}([Z/G])$ of the quotient stack $[Z/G]$ by sending an G -equivariant line bundle L to $[L/G]$. Let $\beta \in \operatorname{Hom}_{\mathbb{Z}}(\operatorname{Pic}^G(Z), \mathbb{Z})$.

Definition 2.2 [Ciocan-Fontanine et al. 2014] A quasimap is a tuple $(C, p_1, \dots, p_k, P, s)$ where

- $(C, p_1, \dots, p_k)$ is a connected, at most nodal, k -pointed projective curve of genus g ,
- P is a principal G -bundle on C ,
- s is a section of the induced fiber bundle $P \times_G Z$ on C such that (P, s) is of class β , i.e., the homomorphism

$$\operatorname{Pic}^G(Z) \rightarrow \mathbb{Z}, \quad L \mapsto \deg_C(s^*(P \times_G L)),$$

is equal to β .

We require that there be only finitely many basepoints, i.e., points $p \in C$ such that $s(p) \in Z_{\text{us}}$. An element $\beta \in \operatorname{Hom}_{\mathbb{Z}}(\operatorname{Pic}^G(Z), \mathbb{Z})$ is called L_θ -effective if it can be represented as a finite sum of classes of quasimaps. We also refer to β as a curve class. Denote by E the semigroup of L_θ -effective (curve) classes.

A quasimap $(C, p_1, \dots, p_k, P, s)$ is called *prestable* if the basepoints are disjoint from the nodes and marked points on C . Given a rational number $\epsilon > 0$, a prestable quasimap is called ϵ -stable if it satisfies the following conditions:

- (1) $\omega_{C,\log} \otimes \mathcal{L}_\theta^\epsilon$ is ample, where $\omega_{C,\log} := \omega_C(\sum_{i=1}^k p_i)$ is the twisted dualizing sheaf of C and

$$\mathcal{L}_\theta := u^*(P \times_G L_\theta) \cong P \times_G \mathbb{C}_\theta.$$

- (2) $\epsilon l(x) \leq 1$ for every point x in C , where

$$l(x) := \text{length}_x(\text{coker}(u^* \mathcal{I} \rightarrow \mathcal{O}_C)).$$

Here $\mathcal{I}$ is the ideal sheaf of the closed subscheme $P \times_G Z_{\text{us}}$ of $P \times_G Z$.

Let $\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta) = \{(C, p_1, \dots, p_k, P, s)\}$ be the moduli stack of ϵ -stable quasimaps. Ciocan-Fontanine et al. [2014] show that this stack is a separated Deligne–Mumford stack of finite type and it is proper over the affine quotient $Z/\text{aff} G := \text{Spec}(A^G)$. When Z has only local complete intersection singularities, the ϵ -stable quasimap space $\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)$ admits a canonical perfect obstruction theory.

Remark 2.3 There are two extreme chambers for the stability parameter ϵ .

- (1) $(\epsilon = \infty)$ -stable quasimaps: One can check that when $(g, k) \neq (0, 0)$ and $\epsilon > 1$, the quasimap space $\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)$ is isomorphic to the moduli space of stable maps $\overline{\mathcal{M}}_{g,k}(Z//G, \beta)$. When $(g, k) = (0, 0)$, the same holds with $\epsilon > 2$. Therefore, when ϵ is sufficiently large, we denote the ϵ -stable quasimap space by

$$\mathcal{Q}_{g,k}^\infty(Z//G, \beta) = \overline{\mathcal{M}}_{g,k}(Z//G, \beta)$$

and refer to it as the $(\epsilon = \infty)$ -theory.

- (2) $(\epsilon = 0+)$ -stable quasimaps: Fix $\beta \in E$. For each $\epsilon \in (0, 1/\beta(L_\theta)]$, the ϵ -stability is equivalent to the condition that the underlying curve C of a quasimap does not have rational tails and, on each rational bridge, the line bundle $\mathcal{L}_\theta$ has strictly positive degree. Since we need to consider different β at the same time, we reformulate the stability condition as

$$\omega_{C,\log} \otimes \mathcal{L}_\theta^\epsilon \text{ is ample for all } \epsilon \in \mathbb{Q}_{>0}.$$

Quasimaps which satisfy the above stability condition are referred to as $(\epsilon = 0+)$ -stable quasimaps.

To define the level structure, we introduce some notation first. Let $\mathfrak{M}_{g,k}$ be the algebraic stack of prestable nodal curves and $\mathfrak{Bun}_G$ be the relative moduli stack

$$\mathfrak{Bun}_G \xrightarrow{\phi} \mathfrak{M}_{g,k}$$

of principal G -bundles on the fibers of the universal curve $\mathfrak{C}_{g,k} \rightarrow \mathfrak{M}_{g,k}$. The morphism ϕ is smooth. There is a forgetful morphism, which forgets the section s ,

$$\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta) \xrightarrow{\mu} \mathfrak{Bun}_G.$$

Let $\tilde{\pi} : \mathfrak{C}_{\mathfrak{Bun}_{g,k}} \rightarrow \mathfrak{Bun}_{g,k}$ be the universal curve which is the pullback of $\mathfrak{C}_{g,k}$ along ϕ . Let $\tilde{\mathfrak{P}} \rightarrow \mathfrak{C}_{\mathfrak{Bun}_{g,k}}$ be the universal principal G -bundle. We denote by $\pi : \mathcal{C}_{g,k} \rightarrow \mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)$ the universal curve on the quasimap space. Let $\mathcal{P} \rightarrow \mathcal{C}_{g,k}$ be the universal principal bundle given by the pullback of $\tilde{\mathfrak{P}} \rightarrow \mathfrak{C}_{\mathfrak{Bun}_{g,k}}$.

Definition 2.4 Given a finite-dimensional representation R of G , we define the level- l determinant line bundle over $\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)$ as

$$(1) \quad \mathcal{D}^{R,l} := (\det R\pi_*(\mathcal{P} \times_G R))^{-l}.$$

Alternatively, one can define $\mathcal{D}^{R,l}$ to be the pullback of the determinant line bundle $(\det R\pi_*(\widetilde{\mathfrak{P}} \times_G R))^{-l}$ on $\mathfrak{Bun}_{g,k}$ via μ .

Remark 2.5 The definition mentioned in the introduction is the second one. It is conceptually better in the sense that it does not depend on the different moduli spaces over $\mathfrak{Bun}_{g,k}$. In our case, these moduli spaces are the ϵ -stable quasimap spaces $\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)$ for different ϵ . However, $\mathfrak{Bun}_{g,k}$ is an Artin stack, and it is technically more difficult to work with it. Formally, we will use the first definition as the working definition.

Remark 2.6 Note that in Definition 2.4, the bundle $\mathcal{P} \times_G R$ is the pullback of the vector bundle $[Z \times R/G] \rightarrow [Z/G]$ along the evaluation map to the quotient stack $[Z/G]$. Therefore, given a vector bundle $\mathcal{R}$ on X , we can use (1) to define a determinant line bundle over the moduli space of stable maps $\overline{\mathcal{M}}_{g,k}(X, \beta)$, even when X is not a GIT quotient. To be more precise, let $\pi : \mathcal{C} \rightarrow \overline{\mathcal{M}}_{g,k}(X, \beta)$ be the universal curve and let $\text{ev} : \mathcal{C} \rightarrow X$ be the universal evaluation map. We define the level- l determinant line bundle as

$$\mathcal{D}^{R,l} := (\det R\pi_*(\text{ev}^* \mathcal{R}))^{-l}.$$

We will often abuse notation by referring to the vector bundle $\mathcal{R}$ as the “representation” R .

The above construction can be easily generalized to orbifold quasimap theory. To be more precise, suppose the target can be written as $[Z^s/G]$. Now, we do not assume G acts freely on the stable locus $Z^s(\theta)$. Therefore $[Z^s/G]$ is in general a Deligne–Mumford stack. The quasimap theory for such orbifold GIT targets is established in [Cheong et al. 2015]. According to [ibid., Section 2.4.5], we still have universal curves and universal principal G -bundles over moduli spaces of ϵ -stable orbifold quasimaps. Therefore the level- l determinant line bundle can still be defined using (1).

2.3 Properties of the level structure in quasimap theory

In this subsection, we study the level- l determinant line bundle in the case $\beta = 0$ and its pullbacks along some natural morphisms between moduli spaces of quasimaps. An important property of the level structure $\mathcal{D}^{R,l}$ is that it splits “correctly” among nodal strata (see Proposition 2.9). In the following discussion, we assume $X = Z//G$ is a GIT quotient, so that moduli spaces of quasimaps are defined. When X is a smooth projective variety but not a GIT quotient, the same results hold for the determinant line bundles defined in Remark 2.6. In fact, the arguments used in the proofs are identical for both cases.

2.3.1 Mapping to a point Assume that $\beta = 0$. Then any quasimap is a constant map and the morphism

$$\mathcal{Q}_{g,k}^\epsilon(X, 0) \xrightarrow{\text{stab} \times \text{ev}} \overline{\mathcal{M}}_{g,k} \times X$$

is an isomorphism. Here $\text{stab} : \mathcal{Q}_{g,k}^\epsilon(X, 0) \rightarrow \overline{M}_{g,k}$ denotes the stabilization morphism of source curves of quasimaps, and $\text{ev} : \mathcal{Q}_{g,k}^\epsilon(X, 0) \rightarrow X$ denotes the constant evaluation morphism. Let P be the principal G -bundle $Z^s \rightarrow X = Z//G$.

Lemma 2.7 *The universal bundle $\mathcal{P}$ over the universal curve $\mathcal{C} = \overline{C}_{g,k} \times X$ is equal to $\pi_2^*(P)$, where $\overline{C}_{g,k}$ is the universal curve over $\overline{M}_{g,k}$ and $\pi_2 : \overline{C}_{g,k} \times X \rightarrow X$ is the second projection.*

Proof In general, there is an evaluation map from the universal curve $\mathcal{C}$ to the quotient stack $[Z/G]$ and $\mathcal{P}$ is the pullback of P along this map. The lemma follows from the observation that the evaluation map is given by the second projection π_2 in this case. $\square$

Corollary 2.8 *Let $\mathcal{R} := P \times_G R$ be the associated vector bundle on X and let $\pi : \overline{C}_{g,k} \rightarrow \overline{M}_{g,k}$ be the canonical morphism. We have*

$$\mathcal{D}^{R,l} = (\wedge^{\text{rk}(R)g} (R^1 \pi_* \mathcal{O}_{\overline{C}_{g,k}} \boxtimes \mathcal{R}))^l \otimes (\wedge^{\text{rk}(R)} \mathcal{R})^{-l}.$$

Proof By Lemma 2.7, we have $\mathcal{P} \times_G R = \pi_2^*(\mathcal{R})$. Therefore the pushforward $R\pi_*(\pi_2^*(\mathcal{R}))$ is equal to $\mathcal{O}_{\overline{M}_{g,k}} \boxtimes \mathcal{R} - R^1 \pi_* \mathcal{O}_{\overline{C}_{g,k}} \boxtimes \mathcal{R}$ via the projection formula. Note that $\text{rk}(R^1 \pi_* \mathcal{O}_{\overline{C}_{g,k}}) = g$. $\square$

2.3.2 Cutting edges For $i = 1, 2$, we denote by $\text{ev}_{k_i} : \mathcal{Q}_{g_i, k_i}^\epsilon(X, \beta_i) \rightarrow X$ the evaluation morphism at the last marking. Consider the cartesian diagram

$$\begin{array}{ccc} \mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times_X \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2) & \xrightarrow{\Phi} & \mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2) \\ \downarrow & & \downarrow \text{ev}_{k_1} \times \text{ev}_{k_2} \\ X & \xrightarrow{\Delta} & X \times X \end{array}$$

where Δ is the diagonal embedding of X . Let $\mathcal{C}$ and $\mathcal{C}'$ denote the universal curves over

$$\mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times_X \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2) \quad \text{and} \quad \mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2),$$

respectively. Let $\mathcal{P}$ and $\mathcal{P}'$ be the universal principal G -bundles over $\mathcal{C}$ and $\mathcal{C}'$, respectively. We can define level structures

$$\mathcal{D}_{\mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times_X \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2)}^{R,l} = \mathcal{D}_{\mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1)}^{R,l} \boxtimes \mathcal{D}_{\mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2)}^{R,l} \quad \text{and} \quad \mathcal{D}_{\mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2)}^{R,l}$$

using (1). The following proposition shows that the level structure splits “correctly” among nodal strata:

Proposition 2.9 *Let $x : \mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times_X \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2) \rightarrow \mathcal{C}$ be the section corresponding to the node. Then*

$$(2) \quad \Phi^*(\mathcal{D}_{\mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1)}^{R,l} \boxtimes \mathcal{D}_{\mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2)}^{R,l}) = \mathcal{D}_{\mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times_X \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2)}^{R,l} \otimes \det(x^*(\mathcal{P} \times_G R))^{-l}.$$

Proof Consider the commutative diagram

$$\begin{array}{ccc} \Phi^* \mathcal{C}' & \xrightarrow{p} & \mathcal{C} \\ & \searrow \pi' & \downarrow \pi \\ & & \mathcal{Q}_{g_1, k_1}^\epsilon(X, \beta_1) \times_X \mathcal{Q}_{g_2, k_2}^\epsilon(X, \beta_2) \end{array}$$

Notice that $\mathcal{C}$ is obtained by gluing along two sections of marked points of $\Phi^*\mathcal{C}'$. For any locally free sheaf $\mathcal{F}$ on $\mathcal{C}$, we have a short exact sequence (the normalization exact sequence)

$$0 \rightarrow \mathcal{F} \rightarrow p_*p^*\mathcal{F} \rightarrow x_*x^*\mathcal{F} \rightarrow 0.$$

which induces the natural isomorphism

$$(3) \quad \det(R\pi_*(p_*p^*\mathcal{F}))^{-1} \cong \det(R\pi_*(\mathcal{F}))^{-1} \otimes \det(R\pi_*(x_*x^*\mathcal{F}))^{-1}.$$

Note that

$$\det(R\pi_*(p_*p^*\mathcal{F}))^{-1} = \det(R\pi'_*(p^*\mathcal{F}))^{-1} \quad \text{and} \quad \det(R\pi_*(x_*x^*\mathcal{F}))^{-1} = \det(x^*\mathcal{F})^{-1}.$$

Now take $\mathcal{F} = \mathcal{P} \times_G R$ and $\mathcal{F}' = \mathcal{P}' \times_G R$. Finally, the lemma follows from the isomorphism $p^*\mathcal{F} \cong \Phi^*\mathcal{F}'$, equation (3), and cohomology and base change. $\square$

2.3.3 Contractions Fix g_1, g_2 and k_1, k_2 such that $g = g_1 + g_2$ and $k = k_1 + k_2$. We denote the basic gluing maps by

$$r : \overline{M}_{g_1, k_1+1} \times \overline{M}_{g_2, k_2+1} \rightarrow \overline{M}_{g, k}, \quad q : \overline{M}_{g-1, k+2} \rightarrow \overline{M}_{g, k}.$$

Let k' be a nonnegative integer. Let $\underline{k}' = (k'_1, \dots, k'_{m+1})$ and $\underline{\beta} = (\beta_1, \dots, \beta_{m+1})$ be partitions of k' and β , respectively. For simplicity, we denote by $\mathcal{Q}_{m, \underline{k}', \underline{\beta}}^\epsilon$ the fiber product

$$\mathcal{Q}_{g_1, k_1+k'_1+1}^\epsilon(X, \beta_1) \times_X \underbrace{\mathcal{Q}_{0, 2+k'_3}^\epsilon(X, \beta_3) \times_X \cdots \times_X \mathcal{Q}_{0, 2+k'_{m+1}}^\epsilon(X, \beta_{m+1})}_{m-1 \text{ factors}} \times_X \mathcal{Q}_{g_2, k_2+k'_2+1}^\epsilon(X, \beta_2).$$

Let $\text{st} : \mathcal{Q}_{g, k+k'}^\epsilon(X, \beta) \rightarrow \overline{M}_{g, k}$ be the morphism defined by forgetting the (quasi)map and the last k' markings, then stabilizing the source curve. Consider the commutative diagram

$$\begin{array}{ccc} \bigsqcup \mathcal{Q}_{m, \underline{k}', \underline{\beta}}^\epsilon & \xrightarrow{\sqcup r_{m, \underline{k}', \underline{\beta}}} & \mathcal{Q}_{g, k+k'}^\epsilon(X, \beta) \\ \downarrow & & \downarrow \text{st} \\ \overline{M}_{g_1, k_1+1} \times \overline{M}_{g_2, k_2+1} & \xrightarrow{r} & \overline{M}_{g, k} \end{array}$$

Here the disjoint union is over partitions $\underline{k}'$ of k' and partitions $\underline{\beta}$ of β . The above commutative diagram induces a morphism

$$\Psi_m : \bigsqcup \mathcal{Q}_{m, \underline{k}', \underline{\beta}}^\epsilon \rightarrow (\overline{M}_{g_1, k_1+1} \times \overline{M}_{g_2, k_2+1}) \times_{\overline{M}_{g, k}} \mathcal{Q}_{g, k+k'}^\epsilon(X, \beta).$$

Using the same argument as in the proof of [Lee 2004, Proposition 11], one can show that the virtual structure sheaves satisfy

$$\sum_m (-1)^{m+1} \Psi_{m*} \sum_{\underline{k}', \underline{\beta}} \mathcal{O}_{\mathcal{Q}_{m, \underline{k}', \underline{\beta}}^\epsilon}^{\text{vir}} = r^! \mathcal{O}_{\mathcal{Q}_{g, k+k'}^\epsilon(X, \beta)}^{\text{vir}}.$$

Let $\mathcal{C}_{g,k+k'}$ and $\mathcal{C}_{m,\underline{k}',\underline{\beta}}$ be the universal curves on $\mathcal{Q}_{g,k+k'}^\epsilon(X, \beta)$ and $\mathcal{Q}_{m,\underline{k}',\underline{\beta}}^\epsilon$, respectively. Let $\mathcal{P}_{g,k+k'}$ and $\mathcal{P}_{m,\underline{k}',\underline{\beta}}$ be the corresponding universal principal G -bundles. The level structures

$$\mathcal{D}_{\mathcal{Q}_{g,k+k'}^\epsilon(X, \beta)}^{R,l} \quad \text{and} \quad \mathcal{D}_{\mathcal{Q}_{m,\underline{k}',\underline{\beta}}^\epsilon}^{R,l}$$

can be defined using (1). To prove that quantum K -theory with level structure satisfies the same axioms as Givental and Lee's quantum K -theory, we need the following proposition:

Proposition 2.10 $\mathcal{D}_{\mathcal{Q}_{m,\underline{k}',\underline{\beta}}^\epsilon}^{R,l} = (r_{m,\underline{k}',\underline{\beta}})^* \mathcal{D}_{\mathcal{Q}_{g,k+k'}^\epsilon(X, \beta)}^{R,l}.$

Proof The proposition follows from the following cartesian diagram and cohomology and base change:

$$\begin{array}{ccc} \mathcal{C}_{m,\underline{k}',\underline{\beta}} & \longrightarrow & \mathcal{C}_{g,k+k'} \\ \downarrow & & \downarrow \\ \mathcal{Q}_{m,\underline{k}',\underline{\beta}}^\epsilon & \xrightarrow{r_{m,\underline{k}',\underline{\beta}}} & \mathcal{Q}_{g,k+k'}^\epsilon(X, \beta) \end{array}$$

□

3 The K -theoretic quasimap theory with level structure

In this section, we first define the K -theoretic quasimap invariants with level structure and their permutation-equivariant version. For most of the discussion, we assume that the GIT quotient $Z//G$ is nonsingular and projective. The cases when the target $[Z^s/G]$ is noncompact or an orbifold are discussed at the end of this section.

3.1 K -theoretic quasimap invariants with level structure

In this subsection, we first briefly recall Givental and Lee's quantum K -theory. Then we define K -theoretic quasimap invariants with level structure.

The quantum K -theory or K -theoretic Gromov–Witten theory was introduced in [Givental 2000; Lee 2004]. Let X be a smooth projective variety and let $\overline{\mathcal{M}}_{g,k}(X, \beta)$ be the moduli space of stable maps to X . The moduli space is known to be a proper Deligne–Mumford stack (see for example [Behrend and Manin 1996]). In particular, for any coherent sheaf $\mathcal{E}$ on $\overline{\mathcal{M}}_{g,k}(X, \beta)$, we can consider its K -theoretic pushforward to the point $\text{Spec } \mathbb{C}$, i.e., we can take its Euler characteristic

$$\chi(\mathcal{E}) = \sum_i (-1)^i h^i(\mathcal{E}),$$

where $h^i(\mathcal{E}) := \dim_{\mathbb{C}} H^i(\overline{\mathcal{M}}_{g,k}(X, \beta), \mathcal{E})$.

From perfect obstruction theory, Lee [2004] constructs a virtual structure sheaf $\mathcal{O}^{\text{vir}} \in K_0(\overline{\mathcal{M}}_{g,k}(X, \beta))$, where $K_0(\overline{\mathcal{M}}_{g,k}(X, \beta))$ denotes the Grothendieck group of coherent sheaves on $\overline{\mathcal{M}}_{g,k}(X, \beta)$. The virtual structure sheaf $\mathcal{O}^{\text{vir}}$ has the following properties:

- (1) If the obstruction sheaf is trivial and hence $\overline{\mathcal{M}}_{g,k}(X, \beta)$ is smooth, then $\mathcal{O}^{\text{vir}}$ is the structure sheaf of $\overline{\mathcal{M}}_{g,k}(X, \beta)$.

- (2) If the obstruction sheaf Obs is locally free, then $\mathcal{O}^{\text{vir}} = \sum_i (-1)^i \wedge^i \text{Obs}^\vee$. Here $\wedge^i \text{Obs}^\vee$ denotes the i -th wedge product of the dual of the obstruction bundle.

Since we assume X to be smooth, the Grothendieck group of locally free sheaves on X , denoted by $K^0(X)$, is isomorphic to the Grothendieck group of coherent sheaves $K_0(X)$. We denote both of them by $K(X)$. Suppose that E_i are K -theory elements in $K(X)$ and let L_i denote the i -th cotangent line bundles. The K -theoretic Gromov–Witten invariants are defined by

$$\langle E_1 L_1^{l_1}, \dots, E_k L_k^{l_k} \rangle_{g,k,\beta} = \chi \left(\overline{\mathcal{M}}_{g,k}(X, \beta), \prod_i \text{ev}_i^* E_i \otimes L_i^{l_i} \otimes \mathcal{O}^{\text{vir}} \right),$$

where $\text{ev}_i : \overline{\mathcal{M}}_{g,k}(X, \beta) \rightarrow X$ are the evaluation morphisms at the i -th marking. Let $E \subset H_2(X, \mathbb{Z})$ be the semigroup generated by effective curve classes on X . We define the *quantum K -potential of genus 0* by

$$\mathcal{F}(t, Q) := \frac{1}{2}(t, t) + \sum_{k=0}^{\infty} \sum_{\beta \in E} \frac{Q^\beta}{k!} \langle t, \dots, t \rangle_{0,k,\beta},$$

where $t \in K(X)_{\mathbb{Q}} := K(X) \otimes \mathbb{Q}$ and $(t, t) := \chi(t \otimes t)$ is the Mukai pairing. Let $\phi_0 = \mathcal{O}_X, \phi_1, \phi_2, \dots$ be a basis of $K(X)_{\mathbb{Q}}$. One can define the “quantized” pairing on $K(X)_{\mathbb{Q}}$ by

$$((\phi_i, \phi_j)) := \mathcal{F}_{ij} = \partial_{t_i} \partial_{t_j} \mathcal{F}(t, Q).$$

Lee [2004] shows that quantum K -theory satisfies all the usual axioms of cohomological Gromov–Witten theory except the flat identity axiom.

In the following discussion, we assume X can be represented as a GIT quotient $Z // G$. As mentioned before, the moduli space of stable maps $\overline{\mathcal{M}}_{g,k}(Z // G, \beta)$ can be identified with $\mathcal{Q}_{g,k}^\epsilon(Z // G, \beta)$ for large ϵ . According to [Ciocan-Fontanine et al. 2014], for general ϵ , the ϵ -stable quasimap space $\mathcal{Q}_{g,k}^\epsilon(Z // G, \beta)$ is proper and it admits a two-term perfect obstruction theory, assuming Z has only lci singularities. Hence, by [Lee 2004], one can construct a virtual structure sheaf $\mathcal{O}^{\text{vir}}$ on $\mathcal{Q}_{g,k}^\epsilon(Z // G, \beta)$.

Definition 3.1 The K -theoretic quasimap invariants of level l are defined by

$$\langle E_1 L_1^{l_1}, \dots, E_k L_k^{l_k} \rangle_{g,k,\beta}^{Z // G, R, l, \epsilon} = \chi \left(\mathcal{Q}_{g,k}^\epsilon(Z // G, \beta), \prod_i \text{ev}_i^* E_i \otimes L_i^{l_i} \otimes \mathcal{O}^{\text{vir}} \otimes \mathcal{D}^{R, l} \right) \in \mathbb{Z},$$

where $E_i \in K(Z // G) \otimes \mathbb{Q}$.

We shall usually suppress $Z // G$ from the notation if there is no confusion. Note that these invariants are all integers.

3.2 Quasimap graph space and $J^{R, l, \epsilon}$ -function

In this subsection, we first recall the definition and properties of the ϵ -stable quasimap graph space. Then we define an important generating series $\mathcal{J}^{R, l, \epsilon}$ of K -theoretic quasimap invariants of level l .

Given a rational number $\epsilon > 0$, the *quasimap graph space*, denoted by $\mathcal{QG}_{g,k}^\epsilon(Z//G, \beta)$, is introduced in [Ciocan-Fontanine et al. 2014]. It is the moduli space of the tuples

$$((C, x_1, \dots, x_k), P, u, \varphi),$$

where $((C, x_1, \dots, x_k), P, u)$ is a prestable quasimap satisfying $\epsilon l(x) < 1$ for every point x on C , and the new data φ is a degree-1 morphism from C to $\mathbb{P}^1$. The curve C has a unique rational component C_0 such that $\varphi|_{C_0} : C_0 \rightarrow \mathbb{P}^1$ is an isomorphism and the complement C/C_0 is contracted by φ . The ampleness condition imposed on the tuples is modified to

$$\omega_{\overline{C \setminus C_0}} \left(\sum x_i + \sum y_j \right) \otimes \mathcal{L}_\theta^\epsilon \quad \text{is ample,}$$

where x_i are marked points on $\overline{C \setminus C_0}$ and y_j are the nodes $\overline{C \setminus C_0} \cap C_0$. Ciocan-Fontanine et al. [2014] show that the quasimap graph space is also a separated Deligne–Mumford stack which is proper over the affine quotient. Moreover, when Z has only lci singularities, the canonical obstruction theory on the graph space is perfect. Similarly, we can define the level- l determinant line bundle $\mathcal{D}_{\mathcal{QG}}^{R,l}$ on $\mathcal{QG}_{g,k}^\epsilon(Z//G, \beta)$ using the universal principal G -bundles over its universal curve.

There is a natural $\mathbb{C}^*$ -action on the graph spaces. Let $[x_0, x_1]$ be homogeneous coordinates on $\mathbb{P}^1$, and set $0 := [1, 0]$ and $\infty := [0, 1]$. We consider the standard $\mathbb{C}^*$ -action on $\mathbb{P}^1$,

$$(4) \quad t \cdot [x_0, x_1] = [tx_0, x_1] \quad \text{for all } t \in \mathbb{C}^*.$$

This induces an action on the ϵ -stable quasimap graph space $\mathcal{QG}_{g,k}^\epsilon(Z//G, \beta)$ by rescaling the parametrized rational component. According to [Ciocan-Fontanine and Kim 2014, Section 4.1], the $\mathbb{C}^*$ -fixed locus can be described as

$$(\mathcal{QG}_{g,k}^\epsilon(Z//G, \beta))^{\mathbb{C}^*} = \coprod F_{g_2, k_2, \beta_2}^{g_1, k_1, \beta_1},$$

where the disjoint union is over all possible splittings

$$g = g_1 + g_2, \quad k = k_1 + k_2, \quad \beta = \beta_1 + \beta_2$$

with $g_i, k_i \geq 0$ and β_i effective. In the stable cases, an ϵ -stable parametrized quasimap

$$((C, x_1, \dots, x_k), P, u, \varphi) \in F_{g_2, k_2, \beta_2}^{g_1, k_1, \beta_1}$$

is obtained by gluing two ϵ -stable quasimaps of types (g_1, k_1, β_1) and (g_2, k_2, β_2) to a constant map $\mathbb{P}^1 \rightarrow p \in Z//G$ at 0 and ∞ , respectively. Therefore, the component $F_{g_2, k_2, \beta_2}^{g_1, k_1, \beta_1}$ is isomorphic to the fiber product

$$\mathcal{Q}_{g_1, k_1 + \bullet}^\epsilon(Z//G, \beta_1) \times_{Z//G} \mathcal{Q}_{g_2, k_2 + \bullet}^\epsilon(Z//G, \beta_2)$$

over the evaluation maps at the special marked points $\bullet$. When one of the components at 0 or ∞ is unstable, we use the following conventions:

- (1) For the unstable cases $(g_1, k_1, \beta_1) = (0, 0, 0)$ or $(0, 1, 0)$ (and likewise for (g_2, k_2, β_2)), we define

$$\mathcal{Q}_{0,0+\bullet}^\epsilon(Z//G, 0) := Z//G, \quad \mathcal{Q}_{0,1+\bullet}^\epsilon(Z//G, 0) := Z//G, \quad \text{ev}_\bullet = \text{Id}_{Z//G}.$$

- (2) For the unstable cases $(g_1, k_1, \beta_1) = (0, 0, \beta_1)$ with $\beta_1 \neq 0$ and $\epsilon \leq 1/\beta_1(L_\theta)$, we denote by

$$\mathcal{Q}_{0,0+\bullet}(Z//G, \beta_1)_0$$

the moduli space of quasimaps $(C = \mathbb{P}^1, P, u)$ such that $u(x) \in P \times_G Z^s$ for $x \neq 0 \in \mathbb{P}^1$ and $0 \in \mathbb{P}^1$ is a basepoint of length $\beta_1(L_\theta)$. Similarly, we define $\mathcal{Q}_{0,0+\bullet}(Z//G, \beta_2)_\infty$ to be the moduli space of quasimaps whose only basepoint is of length $\beta_2(L_\theta)$ and located at ∞ . Using these definitions, we have

$$F_{0,0,\beta_2}^{g,k,\beta_1} \cong \mathcal{Q}_{g,k+\bullet}^\epsilon(Z//G, \beta_1) \times_{Z//G} \mathcal{Q}_{0,0+\bullet}(Z//G, \beta_2)_\infty$$

for $k \geq 1$ and $\epsilon \leq 1/\beta_2(L_\theta)$. Similarly,

$$F_{g,k,\beta_2}^{0,0,\beta_1} \cong \mathcal{Q}_{0,0+\bullet}(Z//G, \beta_1)_0 \times_{Z//G} \mathcal{Q}_{g,k+\bullet}^\epsilon(Z//G, \beta_2)$$

for $k \geq 1$ and $\epsilon \leq 1/\beta_1(L_\theta)$. When $g = k = 0$ and $\epsilon \leq \min\{1/\beta_1(L_\theta), 1/\beta_2(L_\theta)\}$, we have

$$F_{0,0,\beta_2}^{0,0,\beta_1} \cong \mathcal{Q}_{0,0+\bullet}(Z//G, \beta_1)_0 \times_{Z//G} \mathcal{Q}_{0,0+\bullet}(Z//G, \beta_2)_\infty.$$

We denote by $\mathcal{R}$ the vector bundle $Z \times_G R \rightarrow [Z/G]$ and its restriction to $Z//G$. We define the twisted pairing on $K(Z//G)_\mathbb{Q}$ by

$$(5) \quad (u, v)^{R,l} := \chi(u \otimes v \otimes (\det \mathcal{R})^{-l}), \quad \text{where } u, v \in K(Z//G)_\mathbb{Q}.$$

Let $\{\phi_a\}$ be a basis of $K(Z//G)_\mathbb{Q}$ and let $\{\phi^a\}$ be the dual basis with respect to the above twisted pairing $(\cdot, \cdot)^{R,l}$. Let $t = \sum_i t^i \phi_i \in K(Z//G)_\mathbb{Q}$. We define the $J^{R,l,\epsilon}$ -function of level l to be¹

$$(6) \quad \mathcal{J}^{R,l,\epsilon}(t, Q) = 1 - q + t + \sum_a \sum_{(k,\beta) \neq (0,0), (1,0)} Q^\beta \phi^a \left\langle \frac{\phi_a}{1 - qL}, t, \dots, t \right\rangle_{0,k+1,\beta}^{R,l,\epsilon}.$$

In the above summation, the quasimap moduli spaces are empty when $k = 0$, $\beta \neq 0$, $\beta(L_\theta) \leq 1/\epsilon$, and the unstable terms are defined by $\mathbb{C}^*$ -localization on the graph space $\mathcal{QG}_{0,0}^\epsilon(Z//G, \beta)$. To be more precise, we consider the fixed-point locus $F_{0,\beta} := \mathcal{Q}_{0,0+\bullet}(Z//G, \beta)_0$ of the $\mathbb{C}^*$ -action. The unstable terms in (6) are defined to be

$$(1 - q) \sum_a \sum_{\substack{\beta \neq 0 \\ \beta(L_\theta) \leq 1/\epsilon}} Q^\beta (\text{ev}_\bullet)_* \left(\mathcal{O}_{F_{0,\beta}}^{\text{vir}} \otimes \left(\frac{\text{tr}_{\mathbb{C}^*} \mathcal{D}^{R,l}}{\text{tr}_{\mathbb{C}^*} \wedge^* N_{F_{0,\beta}}^\vee} \right) \right),$$

where $\text{ev}_\bullet$ is the evaluation map at $\infty \in \mathbb{P}^1$, $N_{F_{0,\beta}}$ is the virtual normal bundle of the fixed locus $F_{0,\beta}$ in $\mathcal{QG}_{0,0}^\epsilon(Z//G, \beta)$ and $\wedge^* N_{F_{0,\beta}}^\vee := \sum_i (-1)^i \wedge^i N_{F_0}^\vee$ is the K -theoretic Euler class of $N_{F_{0,\beta}}$. Here,

¹In Ciocan-Fontanine and Kim's convention, the J -function starts at 1. The definition given here agrees with Givental's convention in which the J -function starts at the dilaton shift $1 - q$.

the trace of a $\mathbb{C}^*$ -equivariant bundle V , when restricted to the fixed-point locus, is a virtual bundle defined by the eigenspace decomposition with respect to the $\mathbb{C}^*$ -action, i.e., we have

$$\mathrm{tr}_{\mathbb{C}^*}(V) := \sum_i q^i V(i),$$

where $t \in \mathbb{C}^*$ acts on $V(i)$ as multiplication by t^i .

For $1 < \epsilon \leq \infty$, i.e., the $(\epsilon = \infty)$ -theory, (6) defines $J^{R,l,\infty}$ -function (or $J^{R,l}$ -function for short) in the quantum K -theory of level l . In this case, we use $\langle \cdot \rangle^{R,l,\infty}$ or simply $\langle \cdot \rangle^{R,l}$ to denote quantum K -invariants of level l . Following [Givental and Tonita 2014], we introduce the *symplectic loop space formalism*. Recall that E denotes the semigroup of L_θ -effective curve classes on $Z//_\theta G$. The Novikov ring $\mathbb{C}[[Q]]$ is defined as

$$\mathbb{C}[[Q]] := \overline{\left\{ \sum_{\beta \in E} c_\beta Q^\beta \mid c_\beta \in \mathbb{C} \right\}}.$$

Here the completion is taken with respect to the $\mathfrak{m}$ -adic topology, where $\mathfrak{m}$ denotes the maximal ideal generated by nonzero elements of E . We define the *loop space* as

$$\mathcal{K} := [K(Z//G) \otimes \mathbb{C}(q)] \otimes \mathbb{C}[[Q]],$$

where $\mathbb{C}(q)$ is the field of complex rational functions in q . By viewing the elements in $\mathbb{C}(q) \otimes \mathbb{C}[[Q]]$ as the coefficients, we extend the twisted pairing $(\cdot, \cdot)^{R,l}$ to $\mathcal{K}$ via linearity. There is a natural symplectic form Ω on $\mathcal{K}$ defined by

$$(7) \quad \Omega(f, g) := [\mathrm{Res}_{q=0} + \mathrm{Res}_{q=\infty}](f(q), g(q^{-1}))^{R,l} \frac{dq}{q}, \quad \text{where } f, g \in \mathcal{K}.$$

With respect to Ω , there is a *Lagrangian polarization* $\mathcal{K} = \mathcal{K}_+ \oplus \mathcal{K}_-$, where

$$\mathcal{K}_+ = [K(Z//G) \otimes \mathbb{C}[q, q^{-1}]] \otimes \mathbb{C}[[Q]] \quad \text{and} \quad \mathcal{K}_- = \{f \in \mathcal{K} \mid f(0) \neq \infty, f(\infty) = 0\}.$$

As before, let $\{\phi_a\}$ be a basis of $K(Z//G)_{\mathbb{Q}}$ and let $\{\phi^a\}$ be the dual basis with respect to the twisted pairing $(\cdot, \cdot)^{R,l}$. Let $\mathbf{t}(q) = \sum_{i,j} t_j^i \phi_i q^j \in \mathcal{K}_+$ be an arbitrary Laurent polynomial. We define the *big $J^{R,l}$ -function of level l* to be the function $\mathcal{J}^{R,l}(\mathbf{t}(q), Q) : \mathcal{K}_+ \rightarrow \mathcal{K}$ given by

$$\mathcal{J}^{R,l}(\mathbf{t}(q), Q) = 1 - q + \mathbf{t}(q) + \sum_a \sum_{(k,\beta) \neq (0,0), (1,0)} \frac{Q^\beta}{k!} \phi^a \left\langle \frac{\phi_a}{1 - qL}, \mathbf{t}(L), \dots, \mathbf{t}(L) \right\rangle_{0,k+1,\beta}^{R,l,\infty}.$$

Define the *genus-0 K -theoretic descendant potential of level l* by

$$(8) \quad \mathcal{F}^{R,l}(\mathbf{t}, Q) := \sum_{k,\beta} \frac{Q^\beta}{k!} \langle \mathbf{t}(L), \dots, \mathbf{t}(L) \rangle_{0,k,\beta}^{R,l,\infty}.$$

We can identify the cotangent bundle $T^*\mathcal{K}_+$ with the symplectic loop space $\mathcal{K}$ via the Lagrangian polarization and the *dilaton shift* $f \rightarrow f + (1 - q)$. Then $\mathcal{J}^{R,l}$ coincides with the differential of the

descendant potential up to the dilaton shift, i.e., we have

$$\mathcal{J}^{R,l} = 1 - q + t(q) + d_t \mathcal{F}^{R,l}(t, Q).$$

In the case $l = 0$, the above fact is proved in [Givental and Tonita 2014, Section 2]. The same argument works for arbitrary l .

For $(\epsilon = 0+)$ -stable quasimap theory, the definition (6) gives the $I^{R,l}$ -function of level l of $Z//G$,

$$\begin{aligned} \mathcal{I}^{R,l}(t, Q) := \frac{\mathcal{J}^{R,l,0+}(t, Q)}{1 - q} &= 1 + \frac{t}{1 - q} + \sum_a \sum_{\beta \neq 0} Q^\beta (\text{ev}_\bullet)_* \left(\mathcal{O}_{F_{0,\beta}}^{\text{vir}} \otimes \left(\frac{\text{tr}_{\mathbb{C}^*} \mathcal{D}^{R,l}}{\text{tr}_{\mathbb{C}^*} \wedge^* N_{F_{0,\beta}}^\vee} \right) \right) \\ &+ \sum_a \sum_{k \geq 1, (k, \beta) \neq (1, 0)} \frac{Q^\beta}{k!} \phi^a \left\langle \frac{\phi_a}{(1 - q)(1 - qL)}, t, \dots, t \right\rangle_{0, k+1, \beta}^{R, l, \epsilon=0+}, \end{aligned}$$

where $t \in K(Z//G)_{\mathbb{Q}}$.

3.3 The permutation-equivariant quasimap K -theory with level structure

Givental [2017a; 2015a; 2015b; 2015c; 2015d; 2015e; 2015f; 2015g; 2017b; 2020; 2017c] introduced the *permutation-equivariant* quantum K -theory, which takes into account the S_n -action on the moduli spaces of stable maps by permuting the marked points. The definition can be easily generalized to incorporate the level structure.

Let Λ be a λ -algebra, i.e., an algebra over $\mathbb{Q}$ equipped with abstract Adams operations Ψ^k for $k = 1, 2, \dots$. Here $\Psi^k : \Lambda \rightarrow \Lambda$ are ring homomorphisms which satisfy $\Psi^r \Psi^s = \Psi^{rs}$ and $\Psi^1 = \text{id}$. We often assume that Λ includes the Novikov variables, the algebra of symmetric polynomials in a given number of variables, and the torus-equivariant K -ring of a point. We also assume that Λ has a maximal ideal Λ_+ and is equipped with the Λ_+ -adic topology. For example, we can choose

$$\Lambda = \mathbb{Q}[[N_1, N_2, \dots]][[Q]][\Lambda_0^\pm, \dots, \Lambda_N^\pm],$$

where N_i are the Newton polynomials (in infinitely or finitely many variables) and Q denotes the Novikov variable(s). The parameters Λ_i denote the torus-equivariant parameters. The Adams operations Ψ^r act on N_m and Q by $\Psi^r(N_m) = N_{rm}$ and $\Psi^r(Q^\beta) = Q^{r\beta}$, respectively. We assume their actions on the torus-equivariant parameters are trivial.

Similar to the “ordinary” quasimap K -theory with level structure, we define the loop space by

$$\mathcal{K} := [K(Z//G) \otimes \Lambda] \otimes \mathbb{C}(q).$$

As before, this is equipped with a symplectic form defined by (7), and it has a Lagrangian polarization

$$\mathcal{K} = \mathcal{K}_+ \oplus \mathcal{K}_-,$$

where $\mathcal{K}_+$ is the subspace of Laurent polynomials in q and $\mathcal{K}_-$ is the subspace of reduced rational functions which are regular at $q = 0$ and vanish at $q = \infty$.

Consider the natural S_k -action on the quasimap moduli space $\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)$ by permuting the k marked points. Notice that the virtual structure sheaf $\mathcal{O}_{\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)}$ and the determinant line bundle $\mathcal{D}^{R,l}$ are invariant under this action. Therefore, we have the S_k -module

$$[\mathbf{t}(L), \dots, \mathbf{t}(L)]_{g,k,\beta} := \sum_m (-1)^m H^m \left(\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta), \mathcal{O}_{\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)}^{\text{vir}} \otimes \mathcal{D}^{R,l} \otimes \bigotimes_{i=1}^k \mathbf{t}(L_i) \right)$$

for $\mathbf{t}(q) \in \mathcal{K}_+$.

Definition 3.2 The correlators of the permutation-equivariant quasimap K -theory of level l are defined by

$$\langle \mathbf{t}(L), \dots, \mathbf{t}(L) \rangle_{g,k,\beta}^{R,l,\epsilon,S_k} := \pi_* \left(\mathcal{O}_{\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)}^{\text{vir}} \otimes \mathcal{D}^{R,l} \otimes \bigotimes_{i=1}^k \mathbf{t}(L_i) \right),$$

where π_* is the K -theoretic pushforward along the projection

$$\pi : [\mathcal{Q}_{g,k}^\epsilon(Z//G, \beta)/S_k] \rightarrow [\text{pt}].$$

Remark 3.3 When the λ -algebra Λ is chosen to be $\mathbb{Q}[[Q]]$, we refer to the invariants as the *symmetrized* invariants. The pushforward map in Definition 3.2 carries information only about the dimensions of the S_k -invariant parts of sheaf cohomology $[\mathbf{t}(L), \dots, \mathbf{t}(L)]_{g,k,\beta}$. We refer to [Givental 2017a, Example 4] for more details.

For the permutation-equivariant quasimap K -theory, we also consider the $J^{R,l,\epsilon}$ -function, and define the cone $\mathcal{L}_{S_\infty}$ to be the range of the $J^{R,l,\infty}$ -function.

Definition 3.4 The permutation-equivariant K -theoretic $J^{R,l,\epsilon}$ -function of $Z//G$ of level l is defined by

$$(9) \quad \mathcal{J}_{S_\infty}^{R,l,\epsilon}(\mathbf{t}(q), Q) := 1 - q + \mathbf{t}(q) + \sum_a \sum_{(k,\beta) \neq (0,0), (1,0)} Q^\beta \left\langle \frac{\phi_a}{1 - qL}, \mathbf{t}(L), \dots, \mathbf{t}(L) \right\rangle_{0,k+1,\beta}^{R,l,\epsilon,S_k} \phi^a,$$

where the unstable terms in the summation are the same as those in (6).

Note that, in the above definition of permutation-equivariant $J^{R,l,\epsilon}$ -function, we do not need to divide each term by $k!$.

Definition 3.5 We define the Givental cone $\mathcal{L}_{S_\infty}^{R,l}$ as the range of $\mathcal{J}_{S_\infty}^{R,l,\infty}$, i.e.,

$$\mathcal{L}_{S_\infty}^{R,l} := \bigcup_{\mathbf{t}(q) \in \mathcal{K}_+} \mathcal{J}_{S_\infty}^{R,l,\infty}(\mathbf{t}(q), Q) \subset \mathcal{K}.$$

Remark 3.6 In the ordinary — i.e., permutation-*nonequivariant* — quantum K -theory, the range of the J -function is a cone which coincides with the differential of the descendant potential (up to the dilaton shift). Therefore the range of the ordinary K -theoretic J -function is a Lagrangian cone in the loop space $\mathcal{K}$. However, in the permutation-equivariant theory, it is explained in [Givental 2015f] that the cone $\mathcal{L}_{S_\infty}^{R,l}$ is not Lagrangian.

3.4 The level structure in equivariant quasimap theory and orbifold quasimap theory

When $Z//G$ is not proper, one can still define the *equivariant* quasimap invariants if $Z//G$ has an additional torus action such that the fixed-point loci in the quasimap moduli spaces are proper. Ciocan-Fontanine et al. [2014, Section 6.3] explain how to define cohomological quasimap invariants via virtual localization. Similarly, one can define equivariant K -theoretic quasimap invariants (with level structure) for noncompact GIT targets using the K -theoretic virtual localization formula (see [Qu 2018, Section 3.2]). With this understood, we define the $J^{R,l,\epsilon}$ -function of equivariant K -theoretic quasimap invariants of level l using (6). Its permutation-equivariant generalization is straightforward.

If we do not assume G acts freely on the stable locus Z^s , then the target $X := [Z^s/G]$ is naturally an orbifold. For such orbifold GIT targets, a quasimap is a tuple $((C, x_1, \dots, x_k), [u])$, where $(C, x_1, \dots, x_k)$ is a k -pointed, genus- g twisted curve (see [Abramovich et al. 2008, Section 4]) and $[u]$ is a representable morphism from $(C, x_1, \dots, x_k)$ to X . We refer the reader to [Cheong et al. 2015, Section 2.3] for the details of the ϵ -stability imposed on those tuples. We denote by $\mathcal{Q}_{g,k}^\epsilon(X, \beta)$ the moduli stack of ϵ -stable quasimaps to the orbifold X . Cheong et al. [2015, Theorem 2.7] show that this moduli stack is Deligne–Mumford and proper over the affine quotient. Furthermore, if Z only has lci singularities, then $\mathcal{Q}_{g,k}^\epsilon(X, \beta)$ has a canonical perfect obstruction theory. Let $\mathcal{C} \rightarrow \mathcal{Q}_{g,k}^\epsilon(X, \beta)$ be the universal curve. The universal principal G -bundle $\mathcal{P} \rightarrow \mathcal{C}$ is defined as the pullback of the principal G -bundle $Z \rightarrow [Z/G]$ via the universal morphism $[u]: \mathcal{C} \rightarrow [Z/G]$. In the orbifold setting, we can still define the level- l determinant line bundle $\mathcal{D}^{R,l}$ using (1).

According to [Cheong et al. 2015, Section 2.5.1], there are natural evaluation morphisms

$$\mathrm{ev}_i: \mathcal{Q}_{g,k}^\epsilon(X, \beta) \rightarrow \bar{I}_\mu X, \quad ((C, x_1, \dots, x_k), [u]) \mapsto [u]_{|x_i} \quad \text{for } i = 1, \dots, k.$$

Here $\bar{I}_\mu X$ denotes the rigidified cyclotomic inertia stack of X which parametrizes representable maps from gerbes banded by finite cyclic groups to X . Let L_i be the universal cotangent line bundle whose fiber at $((C, x_1, \dots, x_k), [u])$ is the cotangent space of the coarse curve $\underline{C}$ of C at the i -th marked point x_i . For nonnegative integers l_i and classes $E_i \in K^0(\bar{I}_\mu X) \otimes \mathbb{Q}$, we define the K -theoretic quasimap invariants of level l as

$$\langle E_1 L_1^{l_1}, \dots, E_k L_k^{l_k} \rangle_{g,k,\beta}^{X,R,l,\epsilon} = \chi \left(\mathcal{Q}_{g,k}^\epsilon(X, \beta), \prod_i \mathrm{ev}_i^* E_i \otimes L_i^{l_i} \otimes \mathcal{O}^{\mathrm{vir}} \otimes \mathcal{D}^{R,l} \right).$$

When $\epsilon = \infty$ and $l = 0$, this definition recovers the K -theoretic Gromov–Witten invariants of X defined in [Tonita and Tseng 2013].

In the orbifold setting, one can still define the quasimap graph space $\mathcal{QG}_{g,k}^\epsilon(X, \beta)$ (see [Cheong et al. 2015, Section 2.5.3]). The definition of the determinant line bundle $\mathcal{D}^{R,l}$ over the graph space is straightforward. We choose a basis $\{\phi_a\}$ of $K^0(\bar{I}_\mu X) \otimes \mathbb{Q}$. Let $\{\phi^a\}$ be the dual basis with respect to the twisted pairing $(\cdot, \cdot)^{R,l}$ on $K^0(\bar{I}_\mu X) \otimes \mathbb{Q}$ given by

$$(u, v)^{R,l} := \chi(\bar{I}_\mu X, u \otimes \bar{t}^* v \otimes \det^{-l}(\bar{I}_\mu \mathcal{R})).$$

Here $\bar{\iota}$ is the involution induced by $(x, g) \mapsto (x, g^{-1})$ and $\bar{\iota}_\mu \mathcal{R}$ is a vector bundle over $\bar{\iota}_\mu X$ such that the fiber over (x, H) , with $H \subset \text{Aut } x$, is the H -fixed subspace of $\mathcal{R}_x$. With all the notation understood, it is straightforward to adapt the definition of cohomological orbifold quasimap $\mathcal{J}^\epsilon$ -function [Cheong et al. 2015, Definition 3.1] to the (permutation-equivariant) K -theoretic setting.

4 Adelic characterization in quantum K -theory with level structure

In this section, we focus on quantum K -theory, i.e., $(\epsilon = \infty)$ -quasimap theory. We first recall the (virtual) Lefschetz–Kawasaki Riemann–Roch formula. This is the main tool in analyzing the poles of the $J^{R,l}$ -function. We give an adelic characterization of points on the cone $\mathcal{L}_{S_\infty}^{R,l}$ in Theorem 4.5. As an application of the adelic characterization, we prove that certain “determinantal” modifications of points on the cone $\mathcal{L}_{S_\infty}$ of level 0 lie on the cone $\mathcal{L}_{S_\infty}^{R,l}$ of level l . This result will be used in the proof of the toric mirror theorem in Section 5. We assume in this section that X is a smooth projective variety which is not necessarily a GIT quotient. The determinant line bundle $\mathcal{D}^{R,l}$ is defined as in Remark 2.6.

4.1 Virtual Lefschetz–Kawasaki Riemann–Roch formula

To understand the poles of the generating series of the permutation-equivariant quantum K -invariants, we recall the *Lefschetz–Kawasaki Riemann–Roch formula* in [Givental 2017b].

Let h be a finite-order automorphism of a holomorphic orbibundle E over a compact smooth orbifold $\mathcal{M}$. The (super)trace of h on the sheaf cohomology $H^*(\mathcal{M}, E)$ can be computed as an integral over the h -fixed-point locus $I\mathcal{M}^h$ in the inertia orbifold $I\mathcal{M}$:

$$(10) \quad \text{tr}_h H^*(\mathcal{M}, E) = \chi^{\text{fake}} \left(I\mathcal{M}^h, \frac{\text{tr}_{\tilde{h}} E}{\text{tr}_{\tilde{h}} \wedge^* N_{I\mathcal{M}^h}^\vee} \right) := \int_{[I\mathcal{M}^h]} \text{td}(T_{I\mathcal{M}^h}) \text{ch} \left(\frac{\text{tr}_{\tilde{h}} E}{\text{tr}_{\tilde{h}} \wedge^* N_{I\mathcal{M}^h}^\vee} \right).$$

We explain the ingredients of this formula as follows. By definition, we can choose an atlas of local charts $U \rightarrow U/G(x)$ of $\mathcal{M}$. The local description of the inertia orbifold $I\mathcal{M}$ near $x \in \mathcal{M}$ is given by $[\coprod_{g \in G(x)} U^g/G(x)]$, where $U^g \subset U$ denotes the fixed-point locus of g . The automorphism h can be lifted to an automorphism $\tilde{h}$ of the chart U^g . More precisely, the orbibundle lifts to a $G(x)$ -equivariant bundle over U^g . Then the transformation h can be lifted to an automorphism $\tilde{h}$ of the chart U^g and that of the $G(x)$ -equivariant bundle in $|G(x)|$ many ways. We denote by $(U^g)^{\tilde{h}}$ the fixed-point locus of $\tilde{h}$ in U^g . Then the local description of the orbifold $I\mathcal{M}^h$ is given by $[\coprod_g (U^g)^{\tilde{h}}/G(x)]$. We refer to the connected components of $I\mathcal{M}^h$ as *Kawasaki strata*. Near a point $(x, [g]) \in I\mathcal{M}^h$, the tangent and normal orbifold bundles $T_{I\mathcal{M}^h}$ and $N_{I\mathcal{M}^h}$ are identified with the tangent bundle and normal bundle to $(U^g)^{\tilde{h}}$ in U , respectively. In the denominator of the right side of (10), $\wedge^* N_{I\mathcal{M}^h}^\vee := \sum_{i \geq 0} (-1)^i \wedge^i N_{I\mathcal{M}^h}^\vee$ is the K -theoretic Euler class of the normal bundle $N_{I\mathcal{M}^h}$. The trace bundle $\text{tr}_{\tilde{h}} F$ is the virtual orbifold bundle

$$\text{tr}_{\tilde{h}} F := \sum_{\lambda} \lambda F_{\lambda},$$

where F_λ are the eigenbundles of $\tilde{h}$ corresponding to the eigenvalues λ . Finally, td and ch denote the *Todd class* and *Chern character*.

By choosing h to be the identity map, we obtain *Kawasaki's Riemann–Roch formula* [1979] from (10):

$$(11) \quad \chi(\mathcal{M}, E) = \chi^{\text{fake}}\left(I\mathcal{M}, \frac{\text{tr}_g E}{\text{tr}_g \wedge^* N_{I\mathcal{M}}^\vee}\right).$$

When $\mathcal{M}$ is no longer smooth, Tonita [2014b] proved a *virtual* Kawasaki formula: under the assumption that $\mathcal{M}$ has a perfect obstruction theory and admits an embedding into a smooth orbifold which has the resolution property, Kawasaki's formula still holds true if we replace the structure sheaves and the tangent and normal bundles in the formula by their virtual counterparts. According to [Abramovich et al. 2007], the moduli stacks of stable maps to smooth projective varieties satisfy the assumptions of Tonita's theorem. In the next subsection, we apply the virtual Kawasaki Riemann–Roch (KRR) formula to $\overline{\mathcal{M}}_{g,k}(X, \beta)/S_k$ to study the poles of the J -function.

4.2 Adelic characterization

In this subsection, we first recall the adelic characterization [Givental 2015b] of the cone $\mathcal{L}_{S_\infty}$ in the level-0 permutation-equivariant quantum K -theory. Then we generalize it to describe points on the cone $\mathcal{L}_{S_\infty}^{R,l}$ of level l .

4.2.1 The level-0 case In the level-0 case, i.e., Givental and Lee's quantum K -theory, the permutation-equivariant invariants are defined as

$$\langle \mathbf{t}(L), \dots, \mathbf{t}(L) \rangle_{g,k,\beta}^{S_k} := \chi\left(\overline{\mathcal{M}}_{g,k}(X, \beta)/S_k, \mathcal{O}_{\overline{\mathcal{M}}_{g,k}(X, \beta)}^{\text{vir}} \otimes \bigotimes_{i=1}^k \mathbf{t}(L_i)\right),$$

where $\mathbf{t}(q)$ is a Laurent polynomial in q with coefficients in $K^0(X) \otimes \mathbb{Q}$. This is a special case of Definition 3.2 with ϵ sufficiently large and $l = 0$. The J -function is defined as

$$\mathcal{J}_{S_\infty}(\mathbf{t}(q), Q) := 1 - q + \mathbf{t}(q) + \sum_a \sum_{(k,\beta) \neq (0,0), (1,0)} Q^\beta \left\langle \frac{\phi_a}{1 - qL}, \mathbf{t}(L), \dots, \mathbf{t}(L) \right\rangle_{0,k+1,\beta}^{S_k} \phi^a.$$

Here $\{\phi_a\}$ and $\{\phi^a\}$ are bases of $K^0(X) \otimes \mathbb{Q}$ dual with respect to the Mukai pairing

$$(\phi_a, \phi_b) := \chi(\phi_a \otimes \phi_b).$$

Recall the definition of the loop space $\mathcal{K}$ from Section 3.2,

$$\mathcal{K} := [K^0(X) \otimes \mathbb{C}(q)] \otimes \mathbb{C}[[Q]].$$

With respect to the symplectic form (7) with $l = 0$, there is a Lagrangian polarization $\mathcal{K} = \mathcal{K}_+ \oplus \mathcal{K}_-$, where $\mathcal{K}_+$ consists of Laurent polynomials in q and $\mathcal{K}_-$ consists of reduced rational functions in q . The Givental cone $\mathcal{L}_{S_\infty} \subset \mathcal{K}$ is defined as the image of $\mathcal{J}_{S_\infty} : \mathcal{K}_+ \rightarrow \mathcal{K}$.

To study the poles of the series $\mathcal{J}_{S_\infty}(q)$, we apply the virtual KRR formula to $\overline{\mathcal{M}}_{0,k+1}(X, \beta)/S_k$, where the symmetric group acts on the last k markings. By the virtual KRR formula, each term in the J -function can be written as a summation of fake Euler characteristics over the Kawasaki strata. Note that the Kawasaki strata parametrize stable maps with prescribed automorphisms, i.e., equivalence classes of pairs (C, f, h) , where (C, f) is a stable map to X and h is an automorphism of the map. Here, h is allowed to permute the last k markings, but it has to preserve the first marking (with the insertion $1/(1-qL)$). Denote by η the eigenvalue of h on the *cotangent* line to the curve at the first marking.

There are two types of Kawasaki strata. Over the Kawasaki strata with $\eta = 1$, the input $\mathrm{tr}_h(1/(1-qL))$ in the fake Euler characteristics becomes $1/(1-q\bar{L})$, where $1-\bar{L}$ is nilpotent. From the finite expansion

$$\frac{1}{1-q\bar{L}} = \sum_{i \geq 0} \frac{q^i (\bar{L}-1)^i}{(1-q)^{i+1}},$$

we see that the contributions to the J -function from the Kawasaki strata with $\eta = 1$ have poles at $q = 1$.

Over the Kawasaki strata where $\eta \neq 1$ is a primitive m -th root of unity, the insertion $\mathrm{tr}_h(1/(1-qL))$ in the fake Euler characteristics becomes $1/(1-q\eta\bar{L})$, where $1-\bar{L}$ is nilpotent. By considering the finite expansion

$$\frac{1}{1-q\eta\bar{L}} = \sum_{i \geq 0} \frac{(q\eta)^i (\bar{L}-1)^i}{(1-q\eta)^{i+1}},$$

we see that they contribute terms with possible poles at the root of unity η^{-1} to the J -function. We refer the reader to [Givental 2015b] for a nice diagram cataloging all the strata.

For each primitive m -th root of unity η , we denote by $\mathcal{J}_{S_\infty}(\mathbf{t})_{(\eta)}$ the Laurent expansion of the J -function in $1-q\eta$ and regard it as an element in the loop space of power Q -series with vector Laurent series in $1-q\eta$ as coefficients:

$$\mathcal{K}^\eta := K^0(X) \left[\frac{1}{1-q\eta}, 1-q\eta \right] \otimes \mathbb{C}[[Q]].$$

The contributions from the untwisted sector of $\overline{\mathcal{M}}_{0,k}(X, \beta)/S_k$ in the virtual KRR formula are called the *fake* K -theoretic Gromov–Witten (GW) invariants. More precisely, they are defined by

$$\langle \mathbf{t}(L), \dots, \mathbf{t}(L) \rangle_{0,k,\beta}^{\mathrm{fake}} := \int_{[\overline{\mathcal{M}}_{0,k}(X, \beta)]^{\mathrm{vir}}} \prod_{i=1}^k \mathrm{ch}(\mathbf{t}(L_i)) \mathrm{Td}(T^{\mathrm{vir}}),$$

where $[\overline{\mathcal{M}}_{0,k}(X, \beta)]^{\mathrm{vir}}$ is the virtual fundamental class of the moduli space and T^{vir} is the virtual tangent bundle of $\overline{\mathcal{M}}_{0,k}(X, \beta)$. We define the J -function in the fake quantum K -theory, $\mathcal{J}_{\mathrm{fake}} : \mathcal{K}_+^1 \rightarrow \mathcal{K}^1$, by

$$\mathcal{J}_{\mathrm{fake}}(\mathbf{t}(q), Q) := 1 - q + \mathbf{t}(q) + \sum_a \sum_{(k,\beta) \neq (0,0), (1,0)} \frac{Q^\beta}{k!} \phi^a \left\langle \frac{\phi_a}{1-qL}, \mathbf{t}(L), \dots, \mathbf{t}(L) \right\rangle_{0,k+1,\beta}^{\mathrm{fake}}.$$

Here the input $\mathbf{t}(q)$ belongs to

$$\mathcal{K}_+^1 := K^0(X) [[1-q]] \otimes \mathbb{C}[[Q]].$$

Denote by $\mathcal{L}_{\text{fake}} \subset \mathcal{K}^1$ the range of the series $\mathcal{J}_{\text{fake}}$. The negative space $\mathcal{K}_-^1$ of the polarization is spanned by $\{\phi^a q^k / (1-q)^{k+1} \mid a = 1, \dots, \dim K^0(X)_{\mathbb{Q}}, k = 0, 1, \dots\}$. The input $\mathbf{t}(q)$ of $\mathcal{J}_{\text{fake}}$ can be obtained from the projection of $\mathcal{J}_{\text{fake}}$ to $\mathcal{K}_+^1$ along $\mathcal{K}_-^1$.

Givental [2015b] gives the following adelic characterization of the values of $\mathcal{J}_{S_{\infty}}(\mathbf{t})$:

Theorem 4.1 [Givental 2015b] *The values of $\mathcal{J}_{S_{\infty}}(\mathbf{t})$ are characterized by the following requirements:*

- (1) $\mathcal{J}_{S_{\infty}}(\mathbf{t})$ has possible poles only at 0, ∞ and roots of unity.
- (2) $\mathcal{J}_{S_{\infty}}(\mathbf{t})_{(1)}$ lies on $\mathcal{L}_{\text{fake}}$.
- (3) For every primitive root of unity η of order $m \neq 1$,

$$\mathcal{J}_{S_{\infty}}(\mathbf{t})_{(\eta)} \left(\frac{q^{1/m}}{\eta} \right) \in \sqrt{\frac{\lambda_{-1}(T_X)}{\lambda_{-1}(\Psi^m T_X)}} \exp \sum_{i \geq 1} \left(\frac{\Psi^i T_X^{\vee}}{i(1 - \eta^{-i} q^{i/m})} - \frac{\Psi^{im} T_X^{\vee}}{i(1 - q^{im})} \right) \mathcal{T}_m(\mathcal{J}_{S_{\infty}}(\mathbf{t})_{(1)}),$$

where Ψ^m is the m -th Adams operation on $K^0(X)_{\mathbb{Q}}$, which acts on line bundles as $L \mapsto L^m$; $\lambda_{-1}(E) := \sum_i (-1)^i \wedge^i E^{\vee}$ is the K -theoretic Euler class of a vector bundle E ; and $\mathcal{T}_m(f)$ is the space described in Definition 4.2 below.

We recall the following definition from [Tonita 2018]:

Definition 4.2 Let f be a point on $\mathcal{L}_{\text{fake}}$ and let $T(f)$ be the tangent space to $\mathcal{L}_{\text{fake}}$ at f , considered as the image of a map $S(q, Q) : \mathcal{K}_+^1 \rightarrow \mathcal{K}$. We extend the Adams operations from $K^0(X)_{\mathbb{Q}}$ to $\mathcal{K}^1$ by $\Psi^m(q) = q^m$ and $\Psi^m(Q) = Q^m$. Let $\Psi^{1/m}$ be the inverse of Ψ^m , acting as $q \mapsto q^{1/m}$ and $Q \mapsto Q^{1/m}$. We define the space $\mathcal{T}_m(f)$ to be the image of the conjugate of $S(q, Q)$,

$$\Psi^m \circ S(q, Q) \circ \Psi^{1/m} : \mathcal{K}_+^1 \rightarrow \mathcal{K}^1.$$

Remark 4.3 As explained in [Tonita 2018, Remark 5.6], the explicit operator in condition (3) of Theorem 4.1 can be written as a composition $\square_{\eta} \square_m^{-1}$. The definitions of the operators $\square_{\eta}$ and $\square_m$ are given in Proposition 4.10.

4.2.2 The general case of level l Let R be a vector bundle over X . According to Remark 2.6, the level structure $\mathcal{D}^{R,l}$ is defined as

$$\mathcal{D}^{R,l} = (\det \mathcal{R}_{k,\beta})^{-l},$$

where $\mathcal{R}_{k,\beta} := R\pi_*(\text{ev}^* \mathcal{R})$ is the *index bundle*, $\pi : \mathcal{C} \rightarrow \overline{\mathcal{M}}_{g,k}(X, \beta)$ is the universal curve, and $\text{ev} : \mathcal{C} \rightarrow X$ is the universal evaluation morphism. To state the adelic characterization theorem of the cone $\mathcal{L}_{S_{\infty}}^{R,l}$ in the permutation-equivariant quantum K -theory with level structure, we introduce the fake invariants of level l ,

$$\langle \mathbf{t}(L), \dots, \mathbf{t}(L) \rangle_{0,k,\beta}^{\text{fake}, R, l} := \int_{[\overline{\mathcal{M}}_{0,k}(X, \beta)]^{\text{vir}}} \prod_{i=1}^k \text{ch}(\mathbf{t}(L_i)) \text{Td}(T^{\text{vir}}) \text{ch}(\mathcal{D}^{R,l}),$$

and the fake J -function of level l ,

$$\mathcal{J}_{\text{fake}}^{R,l}(\mathbf{t}(q), Q) := 1 - q + \mathbf{t}(q) + \sum_a \sum_{(k,\beta) \neq (0,0), (1,0)} \frac{Q^\beta}{k!} \phi^a \left\langle \frac{\phi_a}{1-qL}, \mathbf{t}(L), \dots, \mathbf{t}(L) \right\rangle_{0,k+1,\beta}^{\text{fake}, R, l}.$$

Here $\{\phi^a\}$ is the dual basis of $\{\phi_a\}$ with respect to the twisted pairing

$$(u, v)^{R,l} = \chi(u \otimes v \otimes (\det \mathcal{R})^{-l}).$$

Denote by $\mathcal{L}_{\text{fake}}^{R,l} \subset \mathcal{K}^1$ the range of the series $\mathcal{J}_{\text{fake}}^{R,l}$. Based on the relationship between gravitational descendants and ancestors of fake quantum K -theory [Givental 2004], one can show that $\mathcal{L}_{\text{fake}}^{R,l}$ is an overruled Lagrangian cone. We refer the reader to [Givental and Tonita 2014, Section 3] for more details.

Convention 4.4 We will consider various twisted theories, in which the pairings are usually different. In particular, the dual bases $\{\phi^a\}$ which appear in the definitions of various J -functions may not be the same. To relate J -functions in different theories, we need to regard them as elements of the same loop space. This is achieved by rescaling the elements in loop spaces. For example, there is a rescaling map

$$(\mathcal{K}^1, (\cdot, \cdot)^{R,l}) \rightarrow (\mathcal{K}^1, (\cdot, \cdot)), \quad E \mapsto E \otimes (\det \mathcal{R})^{-l/2},$$

which identifies the loop space in fake quantum K -theory of level l with that in fake quantum K -theory of level 0.

One of our main results is the following adelic characterization of values of the big J -function in quantum K -theory of level l , generalizing Theorem 4.1:

Theorem 4.5 *The values of $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})$ are characterized by the following requirements:*

- (1) $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})$ has possible poles only at 0, ∞ and roots of unity.
- (2) $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(1)}$ lies on $\mathcal{L}_{\text{fake}}^{R,l}$.
- (3) For every primitive root of unity η of order $m \neq 1$,

$$\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(\eta)} \left(\frac{q^{1/m}}{\eta} \right) \in \sqrt{\frac{\lambda_{-1}(T_X)}{\lambda_{-1}(\Psi^m T_X)}} \exp \sum_{i \geq 1} \left(\frac{\Psi^i T_X^\vee}{i(1 - \eta^{-i} q^{i/m})} - \frac{\Psi^{im} T_X^\vee}{i(1 - q^{im})} \right) \mathcal{T}_m^{R,l}(\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(1)}),$$

where the space $\mathcal{T}_m^{R,l}(\mathbf{f})$ is defined as in Definition 4.2 but starting with a point $\mathbf{f} \in \mathcal{L}_{\text{fake}}^{R,l}$.

4.2.3 The proof of Theorem 4.5 We follow the proofs of [Givental and Tonita 2014; Tonita 2018]. The first requirement in Theorem 4.5 is obviously satisfied. For the second condition, we apply the virtual KRR formula to the J -function. Let $\tilde{\mathbf{T}}(q)$ be the sum of $\mathbf{t}(q)$ and all contributions in $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})$ which are regular at $q = 1$. According to Proposition 2.9, the level structure $\mathcal{D}^{R,l}$ splits “correctly” over nodal strata. With this understood, the following proposition follows from an argument identical to the one given in [Tonita 2018, Proposition 5.2]:

Proposition 4.6 We have

$$\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t}(q))_{(1)} = \mathcal{J}_{\text{fake}}^{R,l}(\tilde{\mathbf{T}}(q))$$

as elements in $\mathcal{K}^1$. In particular, this shows that $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(1)}$ lies on the cone $\mathcal{L}_{\text{fake}}^{R,l}$.

Before we move on to prove the third condition in Theorem 4.5, we characterize the cone $\mathcal{L}_{\text{fake}}^{R,l}$ of fake quantum K -theory of level l in terms of the cone $\mathcal{L}_{\text{fake}}$ of level 0. This characterization will be needed later in the proof of Theorem 4.5.

Note that the fake quantum K -theory is a version of twisted cohomological Gromov–Witten theory. The machinery of twisted cohomological Gromov–Witten invariants was introduced in [Coates and Givental 2007], and generalized in various directions in [Tseng 2010; Tonita 2014a]. Let $\mathcal{H}$ be the loop space of the cohomological GW theory of X ,

$$\mathcal{H} := H^{\text{even}}(X, \mathbb{C})[z^{-1}, z][[Q]].$$

This is equipped with a natural symplectic form Ω_H on $\mathcal{H}$ given by

$$\Omega_H(\mathbf{f}, \mathbf{g}) = \text{Res}_{z=0}(\mathbf{f}(-z), \mathbf{g}(z)) dz,$$

where the pairing $(\cdot, \cdot)$ is the Poincaré pairing. With respect to Ω_H , there is a Lagrangian polarization $\mathcal{H} = \mathcal{H}_+ \oplus \mathcal{H}_-$, where

$$\mathcal{H}_+ := H^{\text{even}}(X, \mathbb{C})[[z]][[Q]], \quad \mathcal{H}_- := \frac{1}{z} H^{\text{even}}(X, \mathbb{C})[z^{-1}][[Q]].$$

Inside $\mathcal{H}$, one can define an overruled Lagrangian cone $\mathcal{L}_H$ by the image of the cohomological big J -function. Since we will not use the explicit description of $\mathcal{L}_H$ in this paper, we refer the reader to [Coates and Givental 2007] for the basic definitions in the cohomological GW theory.

Convention 4.7 Throughout this subsection, we identify $\mathcal{K}^1$ with $\mathcal{H}$ via the Chern character

$$\text{qch} : \mathcal{K}^1 \rightarrow \mathcal{H}, \quad E \mapsto \text{ch}(E), \quad q \mapsto e^z.$$

Hence, K -theoretic insertions (e.g., ϕ_a and L) in the correlators of twisted cohomological theories should be understood as their Chern characters (e.g., $\text{ch } \phi_a$ and $\text{ch } L$).

By definition, the fake quantum K -invariants are obtained from the cohomological invariants by inserting the Todd class $\text{Td}(T^{\text{vir}})$ of the virtual tangent bundle T^{vir} . According to [Coates 2003], the virtual tangent bundle can be written as

$$(12) \quad T^{\text{vir}} = \pi_*(\text{ev}^* T_X - 1) - \pi_*(L_{k+1}^\vee - 1) - (\pi_* i_* \mathcal{O}_{\mathcal{Z}})^\vee$$

in $K^0(\overline{\mathcal{M}}_{0,k}(X, \beta))$. Here, $i : \mathcal{Z} \rightarrow \mathcal{C}$ is the embedding of the nodal locus. The three parts correspond respectively to

- (i) deformations of maps to X of a fixed source curve,
- (ii) deformations of complex structure and configuration of markings, and
- (iii) smoothing the nodes.

Coates and Givental [2006] prove that the cone $\mathcal{L}_{\text{fake}}$ (of level 0) is given explicitly in terms of $\mathcal{L}_H$ as

$$\text{qch}(\mathcal{L}_{\text{fake}}) = \Delta \mathcal{L}_H,$$

where the *loop group transformation* Δ is the Euler–Maclaurin asymptotics of the infinite product

$$\Delta \sim \prod_{\text{Chern roots } x \text{ of } T_X} \prod_{r=1}^{\infty} \frac{x - rz}{1 - e^{-x+rz}}.$$

Here, Δ acts on $\mathcal{H}$ by pointwise multiplication, and it is determined only by the Todd class of the first summand in the expression of T^{vir} . The second and third summands in (12) are respectively responsible for the changes of the dilaton shifts and the polarizations between $\mathcal{H}$ and $\mathcal{K}^1$. We refer to [Coates 2003] for the details.

The fake quantum K -invariants of level l are obtained from those of level 0 by inserting one more class $\text{ch}(\mathcal{D}^{R,l})$. Its effect on the Lagrangian cone is described in the following proposition:

Proposition 4.8 *Under the identification $z = \log q$, we have the identity*

$$\mathcal{L}_{\text{fake}}^{R,l} = \exp\left(-l\left(\frac{\text{ch}_2 \mathcal{R}}{z}\right)\right) \mathcal{L}_{\text{fake}}$$

in the loop space $(\mathcal{K}^1, (,))$.

Proof Recall that the level structure $\mathcal{D}^{R,l}$ is defined as a certain power of the determinant of the index bundle $\mathcal{R}_{k,\beta} = R\pi_*(\text{ev}^* \mathcal{R})$. Note that

$$(13) \quad \text{ch}(\mathcal{D}^{R,l}) = \exp(-l \cdot \text{ch}_1(\mathcal{R}_{k,\beta})).$$

According to [Coates and Givental 2007], the cone of a theory twisted by a general multiplicative characteristic class of the form

$$\exp\left(\sum_{i \geq 0} s_i \text{ch}_i(\mathcal{R}_{k,\beta})\right)$$

is obtained from the cone of the untwisted theory by applying the operator

$$\exp\left(\sum_{m,i \geq 0} s_{2m-1+i} \frac{B_{2m}}{(2m)!} \text{ch}_i(\mathcal{R}) \cdot z^{2m-1}\right).$$

Here we set $s_{-1} = 0$ and the Bernoulli numbers B_{2m} are defined by

$$\frac{t}{1 - e^{-t}} = \frac{t}{2} + \sum_{m \geq 0} \frac{B_{2m}}{(2m)!} t^{2m},$$

and the operator acts on $\mathcal{H}$ by pointwise multiplication. For the twisting class (13), we have $s_1 = -l$ and $s_i = 0$ if $i \neq 1$. By applying the above result, we obtain the corresponding loop group transformation,

$$\exp\left(-l\left(\frac{\text{ch}_2(\mathcal{R})}{z} + \frac{\text{ch}_0(\mathcal{R}) \cdot z}{12}\right)\right).$$

Note that the cone $\mathcal{L}_{\text{fake}}$, being overruled, is invariant under multiplication by functions of z . Therefore, we can ignore the second summand in the exponent of the above operator. $\square$

Now let us prove the third condition in Theorem 4.5. Let $\eta \neq 1$ be a primitive root of unity of order m . The Kawasaki strata in $\overline{\mathcal{M}}_{0,k+1}(X, \beta)/S_k$ which contribute terms with poles at $q = \eta^{-1}$ to the J -function are called the *stem spaces* in [Givental and Tonita 2014]. We give a brief description of stem spaces here, and we refer the reader to [ibid., Section 8] for more details. Let (C', f, h) be a point in these strata. A node is *balanced* if the eigenvalues of h on the two branches are inverse. Let $C_+ \subset C'$ denote the unique maximal subcurve, containing the first marked point, that consists of a chain of balanced $\mathbb{P}^1$ such that the m -th power h^m acts as the identity. There are only two smooth points on C_+ which are fixed by h : the first marking on the first component, and one more on the last component. The second point is called the *butt* in [loc. cit.]. The butt can be a regular point, a marking or a node in C' . The automorphism h acts on the cotangent space at the butt by η^{-1} . The other marked points and unbalanced nodes on C_+ are cyclically permuted by h . We denote by C the quotient of C_+ by the $\mathbb{Z}_m$ -symmetry generated by h . The quotient curve C together with the induced quotient stable map is called a *stem* in [loc. cit.]. Note that a stem curve can carry: unramified marked points, coming from symmetric configurations of m -tuples of markings on the cover; or nodes, coming from m -tuples of symmetric nodes on the cover, where further components of C' , cyclically permuted by h , are attached.

One of the key observations in [loc. cit.] is that the data (C_+, C, f) also represents a stable map to the orbifold $X/\mathbb{Z}_m = X \times B\mathbb{Z}_m$ in the sense of [Chen and Ruan 2002; Abramovich et al. 2002]. Therefore, the contributions with poles at $q = \eta^{-1}$ in the KRR formula for the J -function can be expressed as cohomological integrals over the moduli space of stable maps to $X \times B\mathbb{Z}_m$, twisted by the Todd classes of the traces of the virtual tangent and normal bundles of the Kawasaki strata, and the Chern class of the trace of the level structure $\mathcal{D}^{R,l}$. To be more precise, we introduce some notation first. Let $\overline{\mathcal{M}}_{0,k+2}^{X,\beta}(\eta)$ denote the stem space. It parametrizes stems of degree β , which are quotient maps by the $\mathbb{Z}_m$ -symmetry generated by g . Here g acts by η and η^{-1} on the cotangent lines at the first and last markings of the covering curve, respectively. The only markings on the covering curve fixed by h are the first and last markings. Note that the stem space is a Kawasaki stratum in $\overline{\mathcal{M}}_{0,mk+2}(X, m\beta)$. According to [Givental and Tonita 2014, Proposition 5], the stem space $\overline{\mathcal{M}}_{0,k+2}^{X,\beta}(\eta)$ is isomorphic to the moduli space $\overline{\mathcal{M}}_{0,k+2}^{X/\mathbb{Z}_m,\beta}(g, 1, \dots, 1, g^{-1})$ of stable maps to the orbifold $X/\mathbb{Z}_m$. Here the sequence $(g, 1, \dots, 1, g^{-1})$ indicates the sectors where the evaluation maps land. We also consider the stem space $\overline{\mathcal{M}}_{0,k+1}^{X,\beta}(\eta)$ parametrizing stems whose butts are regular points. Similarly, we have an isomorphism between $\overline{\mathcal{M}}_{0,k+1}^{X,\beta}(\eta)$ and $\overline{\mathcal{M}}_{0,k+1}^{X/\mathbb{Z}_m,\beta}(g, 1, \dots, 1)$.

For simplicity, we denote the stem space $\overline{\mathcal{M}}_{0,k+1}^{X,\beta}(\eta)$ by $\overline{\mathcal{M}}$. Modeling on the contributions in the virtual KRR formula applied to the stack $\overline{\mathcal{M}}_{g,k+1}(X, \beta)/S_k$, we define the correlators in the stem theory of level l by

$$\left\langle \frac{\phi}{1 - qL^{1/m}}, t(L), \dots, t(L) \right\rangle_{0,k+1,\beta}^{\text{stem}, R, l} := \int_{[\overline{\mathcal{M}}]^{\text{vir}}} \text{td}(T_{\overline{\mathcal{M}}}) \text{ch} \left(\frac{\text{ev}_1^* \phi \cdot \prod_{i=2}^k \text{ev}_i^* t(L_i) \cdot \text{tr}_g \mathcal{D}^{R,l}}{(1 - qL_1^{1/m}) \cdot \text{tr}_g (\wedge^* N_{\overline{\mathcal{M}}}^\vee)} \right).$$

Here $[\overline{\mathcal{M}}]^{\text{vir}}$ is the virtual fundamental class of the moduli space $\overline{\mathcal{M}}_{0,k+1}^{X/\mathbb{Z}_m,\beta}(g, 1, \dots, 1)$ of stable maps to $X/\mathbb{Z}_m$, and $T_{\overline{\mathcal{M}}}$ and $N_{\overline{\mathcal{M}}}$ are, respectively, the virtual tangent and normal bundles to $\overline{\mathcal{M}}$, considered as a Kawasaki stratum in $\overline{\mathcal{M}}_{0,mk+1}(X, m\beta)$. The line bundle L_1 is formed by the cotangent spaces of stem curves at the first markings, while $L_1^{1/m}$ corresponds to the cotangent line bundle of the covering curves (see [ibid., Section 7] for the explanation). From the definition, we see that the stem theory of level l is a type of twisted cohomological GW theory of $X \times B\mathbb{Z}_m$.

Before we investigate the stem theory further, let us recall some basic facts about the GW theory of the orbifold $X \times B\mathbb{Z}_m$. In this case, the Lagrangian cone of the cohomological GW theory of $X \times B\mathbb{Z}_m$ is the product of m copies of the Lagrangian cone of the GW theory of X . It lies inside the product of m copies of the Fock space $\mathcal{H}$. We refer to each copy of the Lagrangian cone as a *sector*. These sectors are labeled by elements of $\mathbb{Z}_m = \{1, g, \dots, g^{m-1}\}$.

The following proposition relates the Laurent expansion of $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})$ at $q = \eta^{-1}$ to generating series in stem theory:

Proposition 4.9 *Let $\delta\mathbf{t}(q)$ be the contributions in $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})$ which are regular at $q = \eta^{-1}$, i.e.,*

$$\delta\mathbf{t}(q) = 1 - q + \mathbf{t}(q) + \tilde{\mathbf{t}}(q),$$

where $\tilde{\mathbf{t}}(q)$ is the sum of all the contributions from Kawasaki strata $\overline{\mathcal{M}}_{0,k+1}^{X,\beta}(\xi)$ with $\xi \neq \eta$. Then

$$\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(\eta)} = \delta\mathbf{t}(q) + \sum_a \sum_{(k,\beta) \neq (0,0)} \frac{Q^{m\beta}}{k!} \phi^a \left\langle \frac{\phi_a}{1 - q\eta L^{1/m}}, \mathbf{T}(L), \dots, \mathbf{T}(L), \delta\mathbf{t}\left(\frac{L^{1/m}}{\eta}\right) \right\rangle_{0,k+2,\beta}^{\text{stem}, R, l},$$

where:

- (1) The evaluation morphisms at the marked points land in the twisted sector of $B\mathbb{Z}_m$ labeled by the sequence $(g, 1, \dots, 1, g^{-1})$.
- (2) $\mathbf{T}(L) = \Psi^m \tilde{\mathbf{T}}(L)$, where Ψ^m acts on cotangent line bundles $L \mapsto L^m$, elements of $K^0(X)_{\mathbb{Q}}$, and Novikov variables $Q^\beta \mapsto Q^{m\beta}$.
- (3) $\tilde{\mathbf{T}}(q)$ is the input point of $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(1)}$, i.e., it is determined by

$$1 - q + \tilde{\mathbf{T}}(q) = (\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(1)})_+,$$

where $(\cdot)_+$ denotes the projection along $\mathcal{K}_-^1$ to $\mathcal{K}_+^1$.

Proof Proposition 2.9 shows that the determinant line bundle $\mathcal{D}^{R,l}$ factorizes “nicely” over nodal strata. With this in mind, the argument of [ibid., Proposition 2] applies here with one slight change: when determining $\tilde{\mathbf{T}}(q)$, we do not impose the same condition $\mathbf{t}(q) = 0$ as in [loc. cit.]. This is because in the permutation-equivariant theory, we are allowed to permute marked points. $\square$

Proposition 4.9 shows that the Laurent expansion $\mathcal{J}_{S_\infty}^{R,l}(\mathbf{t})_{(\eta)}(\eta^{-1}q^{1/m})$ of the J -function around $q = \eta^{-1}$ can be identified with a tangent vector to the cone of stem theory of level l ,

$$\delta \mathcal{J}^{\text{st},R,l}(\delta \mathbf{t}, \mathbf{T}') := \delta \mathbf{t}(q^{1/m}) + \sum_a \sum_{(k,\beta) \neq (0,0)} \frac{Q^\beta}{k!} \phi^a \left\langle \frac{\phi_a}{1 - q^{1/m} L^{1/m}}, \mathbf{T}'(L), \dots, \mathbf{T}'(L), \delta \mathbf{t}(L^{1/m}) \right\rangle_{0,k+2,\beta}^{\text{stem},R,l},$$

after replacing Q^β with $Q^{m\beta}$ (but not in $\delta \mathbf{t}$). Note that $\mathbf{T}'$ belongs to the sector labeled by 1 and the tangent vector $\delta \mathcal{J}^{\text{st},R,l}(\delta \mathbf{t}, \mathbf{T})$ belongs to the sector labeled by g^{-1} .

Now let us study the stem theory of level l using the formalism of twisted cohomological GW theory of $X \times B\mathbb{Z}_m$. It follows from the definition that the stem theory of level l is obtained from the untwisted cohomological Gromov–Witten theory of $X \times B\mathbb{Z}_m$ by twisting the classes

$$(14) \quad \text{td}(T_{\overline{\mathcal{M}}})/\text{ch}(\text{tr}_g(\wedge^* N_{\overline{\mathcal{M}}}^\vee))$$

and

$$(15) \quad \text{ch}(\text{tr}_g \mathcal{D}^{R,l}).$$

The trace in the first twisting class (14) is computed in [ibid., Section 8], and the effects of this twisting class on the Lagrangian cone and the genus zero potential of the untwisted theory are also studied in [loc. cit.]. We summarize them in the following proposition:

Proposition 4.10 [Givental and Tonita 2014] *The effects of the twisting class (14) on the Lagrangian cone and the genus zero potential of the untwisted theory are described as follows:*

- (i) *The sectors labeled by 1 and g^{-1} are rotated by the operators $\square_m$ and $\square_\eta$, respectively. These two operators are the Euler–Maclaurin asymptotics of the infinite products*

$$\begin{aligned} \square_m &\sim \prod_i \left(\sqrt{\frac{x_i}{1 - e^{-mx_i}}} \prod_{r=1}^{\infty} \frac{x_i - rz}{1 - e^{-mx_i + rmz}} \right), \\ \square_\eta &\sim \prod_i \left(\sqrt{\frac{x_i}{1 - e^{-x_i}}} \prod_{r=1}^{\infty} \frac{x_i - rz}{1 - \eta^{-r} e^{-x_i + rz/m}} \right), \end{aligned}$$

where x_i are the Chern roots of the tangent bundle T_X .

- (ii) *The dilaton shift changes from $-z$ to $1 - q^m$.*
- (iii) *There are changes of polarizations of symplectic loop spaces. More precisely, in the sector labeled by 1, the negative space of the polarization is spanned by*

$$\phi^a \Psi^m \frac{q^k}{(1 - q)^{k+1}},$$

whereas in the sector labeled by g^{-1} , it is spanned by

$$\phi^a \frac{q^{k/m}}{(1 - q^{1/m})^{k+1}}.$$

We compute the second twisting class (15), and describe its effect on the Lagrangian cone. Let p be the universal family of stem curves. By abuse of notation, we still use ev to denote the universal evaluation morphism from the universal family of quotient curves to $X/\mathbb{Z}_m$. Let $\mathbb{C}_{\eta^i}$ be the topologically trivial line bundle on $X/\mathbb{Z}_m$ on which g acts as multiplication by η^i . According to a simple argument in [Tonita 2018], the trace of the index bundle $R\pi_*(\text{ev}^* \mathcal{R})$ can be expressed as

$$\text{tr}_g(R\pi_*(\text{ev}^* \mathcal{R})) = \sum_{i=0}^{m-1} \eta^i R p_*(\text{ev}^* \mathcal{R} \otimes \mathbb{C}_{\eta^i}).$$

For simplicity, we denote $R p_*(\text{ev}^* \mathcal{R} \otimes \mathbb{C}_{\eta^i})$ by $\overline{\mathcal{R}}_i$. Then

$$\begin{aligned} (16) \quad \text{ch tr}_g \mathcal{D}^{R,l} &= \text{ch} \left(\det \sum_{i=0}^{m-1} \eta^i \overline{\mathcal{R}}_i \right)^{-l} \\ &= \text{ch} \left(\prod_{i=0}^{m-1} (\eta^i)^{\text{ch}_0 \overline{\mathcal{R}}_i} \prod_{i=0}^{m-1} \det \overline{\mathcal{R}}_i \right)^{-l} \\ &= \prod_{i=0}^{m-1} \exp(-l(i \log(\eta) \text{ch}_0 \overline{\mathcal{R}}_i + \text{ch}_1 \overline{\mathcal{R}}_i)). \end{aligned}$$

Here, to take the logarithm of η , we may assume that η is taken in the principal branch of the logarithm.

Proposition 4.11 *Twisting by the class (15) rotates the sector labeled by the identity of the Lagrangian cone of $X \times B\mathbb{Z}_m$ by*

$$D_m := \exp \left(-ml \left(\frac{\text{ch}_2 \mathcal{R}}{z} \right) \right).$$

The sector labeled by g^{-1} is rotated by the same operator D_m .

Proof The proof is based on the orbifold quantum Riemann–Roch theorem developed in [Tseng 2010]. Let E be an orbifold vector bundle over $X \times B\mathbb{Z}_m$. Consider a general twisting class

$$\exp \left(\sum_{j \geq 0} s_j \text{ch}_j p_*(\text{ev}^* E) \right).$$

According to [Tseng 2010, Theorem 1], this corresponds to rotation by the operator

$$\exp \left(\sum_{j \geq 0} s_j \left(\sum_{n \geq 0} \frac{(A_n)_{j+1-n} z^{n-1}}{n!} + \frac{\text{ch}_j E^{(0)}}{2} \right) \right).$$

Here A_n is an operator which acts on all sectors. The restriction $(A_n)|_{X, g^i}$ of A_n to the sector labeled by g^i is defined by

$$(A_n)|_{(X, g^i)} = \sum_{r=0}^{m-1} B_n\left(\frac{r}{m}\right) \text{ch } E_i^{(r)},$$

where $E_i^{(r)}$ (respectively $E^{(0)}$) is the subbundle of the restriction of E to (X, g^i) on which g^i acts with eigenvalue $e^{2\pi\sqrt{-1}r/m}$ (respectively 1). The notation $(A_n)_j$ denotes the degree- j component of the operator A_n . The Bernoulli polynomials are defined by

$$\sum_{n \geq 0} B_n(x) \frac{t^n}{n!} = \frac{te^{tx}}{e^t - 1}.$$

In our case, the twisting class is given by (16). Let $\tilde{D}_{\eta, i}$ denote the symplectic transformations corresponding to the i -th factor of the twisting class (16), restricted to (X, g^{-1}) . For each $i \in \{0, \dots, m-1\}$, the operator A_n in the definition of $\tilde{D}_{\eta, i}$ is given by

$$(A_n)|_{(X, g^{-1})} = B_n\left(\frac{i}{k}\right) \text{ch } \mathcal{R}.$$

By the orbifold quantum Riemann–Roch theorem, the operator $\tilde{D}_{\eta, i}$ equals

$$\exp\left(-l \log(\eta) \left(\frac{\text{ch}_1 \mathcal{R}}{z} + B_1\left(\frac{i}{m}\right) \text{ch}_0 \mathcal{R}\right) - l \left(\frac{\text{ch}_2 \mathcal{R}}{z} + B_1\left(\frac{i}{m}\right) \text{ch}_1 \mathcal{R} + \frac{B_2(i/m) \text{ch}_0 \mathcal{R}}{2} z\right) - l \frac{\text{ch}_1 \mathcal{R}}{2}\right).$$

Let $\tilde{D}_\eta = \prod_{i=0}^{m-1} \tilde{D}_{\eta, i}$. To simplify the expression of $\tilde{D}_\eta$, we use the fact that $n \log(\eta) = 0$ if n is an integer divisible by m . Keeping this in mind, we obtain

$$\tilde{D}_\eta = \exp\left(-l \left(\frac{m \text{ch}_2 \mathcal{R}}{z} + \frac{\text{ch}_0 \mathcal{R}}{12m} z + \frac{\text{ch}_0 \mathcal{R}}{6} \log(\eta)\right)\right).$$

Note that the factor $\exp(-l(z \text{ch}_0 \mathcal{R}/(12m) + \log(\eta) \text{ch}_0 \mathcal{R}/6))$ in $\tilde{D}_\eta$ is a scalar z -series and thus it preserves the overruled Lagrangian cone. We can drop it and obtain the operator D_m .

For the sector labeled by the identity, we denote by $\tilde{D}_i$ the restriction of the operator corresponding to the i -th factor of the twisting class (16). It is easy to check that

$$(A_n)|_{(X, 1)} = B_n(0) \text{ch } \mathcal{R},$$

and the operator $\tilde{D}_i$ equals

$$\exp\left(-il \log(\eta) \left(\frac{\text{ch}_1 \mathcal{R}}{z}\right) - ml \left(\frac{\text{ch}_2 \mathcal{R}}{z} + \frac{\text{ch}_0 \mathcal{R}}{12} z\right)\right).$$

Let $\tilde{D} = \prod_{i=0}^{m-1} \tilde{D}_i$. Again by using the fact that $n \log(\eta) = 0$ if $m|n$, we can simplify the operator $\tilde{D}$ to

$$\exp\left(-ml \left(\frac{\text{ch}_2 \mathcal{R}}{z} + \frac{\text{ch}_0 \mathcal{R}}{12} z\right)\right).$$

We can drop the second term in the exponent because it is a constant z -series. □

Recall the Lagrangian cone of $X \times B\mathbb{Z}_m$ is m copies of the Lagrangian cone of X . The tangent vector $\delta\mathcal{J}^{\text{st}, R, l}(\delta t, T)$ has application point in the sector labeled by 1 but is tangent in the direction labeled by g^{-1} . Using the above discussion, we locate $\text{qch} \delta\mathcal{J}^{\text{st}, R, l}(\delta t, T)$ as follows:

Proposition 4.12 $\text{qch} \delta\mathcal{J}^{\text{st}, R, l}(\delta t, T)$ lies in the tangent space $\square_\eta \square_m^{-1}(\mathcal{T}_{\mathcal{I}^{tw}} \square_m D_m \mathcal{L}_H)$ to the cone of the stem theory of level l at a certain point $\mathcal{I}^{tw}$. The input T satisfies

$$\text{qch}(1 - q^m + T(q)) = [\mathcal{I}^{tw}]_+,$$

where $[\cdot]_+$ denotes the projection along the negative space of the polarization of the sector labeled by 1.

Proof The argument of [Tonita 2018, Proposition 5.8] applies here. We briefly explain the relation between the application point $\mathcal{I}^{tw}$ and the input T . Note that the application point $\mathcal{I}^{tw}$ lies on the Lagrangian of the stem theory of level l in the sector labeled by 1. According to Proposition 4.10(ii), the new dilaton shift is $1 - q^m$. This explains the equality in the proposition. $\square$

To prove the third condition in Theorem 4.5, we need to identify $\mathcal{T}_{\mathcal{I}^{tw}} \square_m D_m \mathcal{L}_H$ with $\mathcal{T}_m(\mathcal{J}_{S_\infty}(t)_{(1)})$. We first show that:

Proposition 4.13 We have

$$\text{qch}^{-1}(\square_m D_m \mathcal{L}_H) = \tilde{\Psi}^m \mathcal{L}_{\text{fake}}^{R, l}.$$

Here, the Adams operation $\tilde{\Psi}^m$ acts on K -theory classes of X and q by $\Psi^m(q) = q^m$, but not on the Novikov variables Q .

Proof It is proved in [Givental and Tonita 2014, Proposition 9] that $\text{qch}^{-1}(\square_m \mathcal{L}_H) = \Psi^m \mathcal{L}_{\text{fake}}$. In that proof, one needs to extend the action of the Adams operator Ψ^m on cohomology classes via the Chern isomorphism:

$$\text{ch}(\Psi^m(\text{ch}^{-1} a)) = m^{\deg(a)/2} a.$$

By Proposition 4.8, we have $\mathcal{L}_{\text{fake}}^{R, l} = D_1 \mathcal{L}_{\text{fake}}$. We conclude the proof by noticing that $\Psi^m(D_1) = D_m$. $\square$

Let $\tilde{T}(q)$ be the input point of $\mathcal{J}_{\text{fake}}^{R, l}(\tilde{T}(q))$ determined by

$$1 - q + \tilde{T}(q) = (\mathcal{J}_{S_\infty}^{R, l}(t)_{(1)})_+,$$

where $(\cdot)_+$ denotes the projection along $\mathcal{K}_-^1$ to $\mathcal{K}_+^1$. Let $T'(q)$ be the input point of $\mathcal{I}^{tw}$ such that

$$(17) \quad \tilde{\Psi}^m(\mathcal{J}_{\text{fake}}^{R, l}(\tilde{T}(q))) = \mathcal{I}^{tw}(T'(q)).$$

We claim that $\tilde{\Psi}^m(\tilde{T}(q)) = T'(q)$. This equality holds because, according to Propositions 4.10 and 4.13, the operation $\tilde{\Psi}^m : \mathcal{K}^1 \rightarrow \mathcal{K}^1$ identifies the cone $\mathcal{L}_{\text{fake}}^{R, l}$ with the cone $\square_m D_m \mathcal{L}_H$, the polarization of the fake quantum K -theory with the polarization in the sector labeled by 1 of the stem theory, and the old dilaton shift $1 - q$ with the new one $1 - q^m$. Therefore $\tilde{\Psi}^m$ must also map the input point $\tilde{T}(q)$ of the fake J -function to the input point $T'(q)$ of $\mathcal{I}^{tw}$. Recall from Proposition 4.9 that $\Psi^m(\tilde{T}(q)) = T(q)$.

Then it follows from the definitions of Ψ^m and $\tilde{\Psi}^m$ that $T'(q)$ is obtained from $T(q)$ by replacing Q^β with $Q^{\beta/m}$.

By differentiating the relation (17), we get

$$\begin{aligned} \tilde{\Psi}^m \left(f(q) + \sum \frac{Q^\beta}{k!} \phi^a \left\langle \frac{\phi_a}{1-qL}, \tilde{T}(L), \dots, \tilde{T}(L), f(L) \right\rangle_{0,k+2,\beta}^{\text{fake},R,l} \right) \\ = \tilde{\Psi}^m f(q) + \sum \frac{Q^\beta}{k!} \tilde{\Psi}^m \phi^a \left\langle \frac{\tilde{\Psi}^m \phi_a}{1-q^m L^m}, T'(L), \dots, T'(L), \tilde{\Psi}^m f(L) \right\rangle_{0,k+2,\beta}^{\text{stem},R,l}. \end{aligned}$$

The right-hand side is a tangent vector in $\mathcal{T}_{\text{Tw}} \square_m D_m \mathcal{L}_H$ along the direction of $\delta t' := \tilde{\Psi}^m f(q)$. The left-hand side becomes $\Psi^m \circ S(q, Q) \circ \Psi^{1/m}(\delta t')$ after we replace Q^β with $Q^{m\beta}$ (including such a change in $\tilde{T}$ but excluding it in $f(q)$). This concludes the proof of Theorem 4.5.

Remark 4.14 When the target X is an orbifold, the adelic characterization of points on the cone $\mathcal{L}$ of the ordinary — i.e., permutation-*nonequivariant* — quantum K -theory is developed in [Tonita and Tseng 2013]. In this case, the Lagrangian cone $\mathcal{L}$ has different sectors, and each sector corresponds to a connected component of the rigidified inertia stack $\bar{I}_\mu X$. Let $f : C \rightarrow X$ be an orbifold stable map. Here C is an orbifold curve with possible orbifold structures at the marked points and nodes. Let $\underline{f} : \underline{C} \rightarrow \underline{X}$ be the map between coarse moduli spaces. There is a short exact sequence

$$1 \rightarrow K \rightarrow \text{Aut}(f) \rightarrow \text{Aut}(\underline{f}) \rightarrow 1.$$

The kernel K consists of automorphisms of $C \rightarrow X$ that fix $\underline{C} \rightarrow \underline{X}$. These automorphisms are referred to as “ghost automorphisms” in [Abramovich et al. 2003], and they arise from the stacky nodes of the source curve.

To analyze poles of K -theoretic J -functions in the orbifold setting, we still apply the virtual KRR formula to the moduli space of orbifold stable maps. In this case, there are extra contributions from twisted sectors corresponding to ghost automorphisms. The key observation in [Tonita and Tseng 2013] is that once we add the appropriate contributions from ghost automorphisms in the definition of fake K -theoretic GW invariants, the formalism of adelic characterizations carries over to the orbifold setting. We refer the reader to [ibid., Definition 3.1] for the precise definition of fake K -theoretic invariants and [ibid., Theorem 4.1] for the adelic characterization in the orbifold and permutation-*nonequivariant* setting. We only mention that if we restrict to the untwisted sector of the cone $\mathcal{L}$, the main theorem in [loc. cit.] specializes to Theorem 4.1 above.

The generalization of [ibid., Theorem 4.1] to the permutation-equivariant setting is straightforward: we only need to change the application point of the tangent space in [ibid., Definition 4.3] from $\mathcal{J}_1(0)$ to the Laurent expansion $(\mathcal{J}_{S_\infty}(t))_1$ of the J -function at $q = 1$. Since the determinant line bundle splits “correctly” among nodal strata, we can also generalize [ibid., Theorem 4.1] to permutation-equivariant quantum K -theory with level structure. In this paper, we focus on recovering examples of mock theta functions. For this purpose, we only need to consider the untwisted sector of Lagrangian cones of orbifold

targets. Once we make this restriction, the statement of the adelic characterization is the same as in Theorem 4.5.

4.3 Determinantal modification

In this subsection, we use the adelic characterization to prove Theorem 1.1, which gives us a way to obtain points on $\mathcal{L}_{S_\infty}^{R,l}$ by making certain “determinantal” modifications to points on the level-0 cone $\mathcal{L}_{S_\infty}$.

Let us restate Theorem 1.1:

Theorem 1.1 *If*

$$I = \sum_{\beta \in \text{Eff}(X)} I_\beta Q^\beta$$

lies on $\mathcal{L}_{S_\infty}$, then the point

$$I^{R,l} := \sum_{\beta \in \text{Eff}(X)} I_\beta Q^\beta \prod_i (L_i^{-\beta_i} q^{(\beta_i+1)\beta_i/2})^l$$

lies on the cone $\mathcal{L}_{S_\infty}^{R,l}$ of permutation-equivariant quantum K -theory of level l . Here $\text{Eff}(X)$ denotes the semigroup of effective curve classes on X , L_i are the K -theoretic Chern roots of $\mathcal{R}$, and $\beta_i := \int_\beta c_1(L_i)$.

Proof Suppose $I = \sum_{\beta \in \text{Eff}(X)} I_\beta Q^\beta$ is a point on $\mathcal{L}_{S_\infty}$. Let $I^{R,l}$ be its “determinantal” modification $\sum_\beta I_\beta Q^\beta \prod_i (L_i^{-\beta_i} q^{\beta_i(\beta_i-1)/2})^l$. According to Convention 4.4, we compare different cones in the same loop space $\mathcal{K}^1$. In particular, $\mathcal{L}_{\text{fake}}^{R,l}$ and the tangent space in Theorem 4.5 are viewed as subspaces of $\mathcal{K}^1$. Therefore, to show $I^{R,l}$ lies on $\mathcal{L}_{S_\infty}^{R,l}$, we need to work with the series after the rescaling

$$\tilde{I}^{R,l} := (\det \mathcal{R})^{-l/2} I^{R,l}.$$

Note that the square root of the determinant line bundle may not exist. Instead, we treat it as a formal variable that rescales the elements in the loop space. This rescaling modifies the pairing in the loop space, which, nonetheless, remains well defined.

We denote by $I_{(\eta)}$ and $\tilde{I}_{(\eta)}^{R,l}$ the Laurent expansions of I and $\tilde{I}^{R,l}$ in $1 - q\eta$, respectively. Let $Q_1, \dots, Q_n$ be the Novikov variables. Let p_i be the degree-2 cohomology classes corresponding to Q_i , and let $P_i = e^{-p_i} \in K^0(X)$. We denote by L_i the K -theoretic Chern roots of $\mathcal{R}$. In other words, $\text{ch } L_i = e^{l_i}$, where l_i are cohomological Chern roots of $\mathcal{R}$. We write l_i as a linear function $f_i(p_1, \dots, p_n)$ in terms of the basis $p_1, \dots, p_n$.

It is clear that $\tilde{I}^{R,l}$ satisfies the first condition in Theorem 4.5. Now, we check the second condition. It follows from the lemma in the proof of [Coates and Givental 2007, Theorem 2] that the operator

$$\Phi := \prod_{i=1}^n \exp \left(l \left(\frac{(f_i(p_j - z Q_j \partial_{Q_j}))^2}{2z} + \frac{f_i(p_j - z Q_j \partial_{Q_j})}{2} \right) \right)$$

preserves $\mathcal{L}_{\text{fake}}$. Define $d_j = -\langle c_1(P_j), \beta \rangle$ to be the components of degree β . By a simple computation, one can show that

$$\begin{aligned} \Phi(Q^\beta) &= \prod_{i=1}^n \exp\left(l\left(\frac{(f_i(p_j - zd_j))^2}{2z} - \frac{f_i(p_j - zd_j)}{2}\right)\right) Q^\beta \\ &= \exp\left(l\left(\frac{\text{ch}_2 \mathcal{R}}{2z} - \frac{\text{ch}_1 \mathcal{R}}{2}\right)\right) Q^\beta \\ &\quad \cdot \prod_{i=1}^n \exp\left(l\left(\frac{(f_i(p_j - zd_j))^2}{2z} - \frac{f_i(p_j - zd_j)}{2}\right)\right) \exp\left(-l\left(\frac{\text{ch}_2 \mathcal{R}}{2z} - \frac{\text{ch}_1 \mathcal{R}}{2}\right)\right) \\ &= \exp\left(l\left(\frac{\text{ch}_2 \mathcal{R}}{z} - \frac{\text{ch}_1 \mathcal{R}}{2}\right)\right) Q^\beta \prod_i (L_i^{-\beta_i} q^{\beta_i(\beta_i+1)/2})^l. \end{aligned}$$

It follows that

$$(18) \quad \Phi(I_{(1)}) = \exp\left(l\left(\frac{\text{ch}_2 \mathcal{R}}{z}\right)\right) \tilde{I}_{(1)}^{R,l}.$$

Since the left-hand side lies on $\mathcal{L}_{\text{fake}}$, we conclude that the Laurent expansion of $\tilde{I}^{R,l}$ at $q = 1$ lies on the cone $\mathcal{L}_{\text{fake}}^{R,l} = \exp(-l(\text{ch}_2 \mathcal{R}/z)) \mathcal{L}_{\text{fake}}$.

Now we check the third condition. Suppose the tangent space to $\mathcal{L}_{\text{fake}}$ at $I_{(1)}$ is given as the image of a map

$$S(q, Q) : \mathcal{K}_+^1 \rightarrow \mathcal{K}^1.$$

Then, by (18), the tangent space to $\mathcal{L}_{\text{fake}}^{R,l}$ at $\tilde{I}_{(1)}^{R,l}$ is given as the image of a map

$$S'(q, Q) = \exp\left(-l\left(\frac{\text{ch}_2 \mathcal{R}}{z}\right)\right) \Phi \circ S(q, Q) : \mathcal{K}_+^1 \rightarrow \mathcal{K}^1.$$

Here we use the fact that the Novikov variables are contained in the λ -algebra and hence they preserve tangent spaces.

Recall from Definition 4.2 that the space $\mathcal{T}_m(I_{(1)})$ is defined as the image of a map

$$\Psi^m \circ S(q, Q) \circ \Psi^{1/m} : \mathcal{K}_+^1 \rightarrow \mathcal{K}^1.$$

Then $\mathcal{T}_m^{R,l}(\tilde{I}_{(1)}^{R,l})$ is given as the image of

$$\begin{aligned} \Psi^m \circ S'(q, Q) \circ \Psi^{1/m} &= \Psi^m \circ \exp\left(-l\left(\frac{\text{ch}_2 \mathcal{R}}{z}\right)\right) \Phi \circ S(q, Q) \circ \Psi^{1/m} \\ &= \exp\left(-ml\left(\frac{\text{ch}_2 \mathcal{R}}{z}\right)\right) \Psi^m \circ \Phi \circ S(q, Q) \circ \Psi^{1/m} \\ &= D_m \Phi^m(\Psi^m \circ S(q, Q) \circ \Psi^{1/m}), \end{aligned}$$

where $\Phi^m := \Psi^m(\Phi)$ is given by

$$\Psi^m(\Phi) = \prod_{i=1}^n \exp\left(l\left(\frac{(f_i(mp_j - zQ_j\partial Q_j))^2}{2mz} + \frac{f_i(mp_j - zQ_j\partial Q_j)}{2}\right)\right).$$

Here we use the fact that the Adams operation Ψ^m acts on the degree-2 classes z and p_j as multiplication by m , and its action on the differential operator $zQ_j\partial Q_j$ is trivial.² This shows that $\tau_m^{R,l}(\tilde{I}_{(1)}^{R,l}) = D_m\Phi^m(\tau_m(I_{(1)}))$.

By assumption, we have

$$I_{(\eta)}\left(\frac{q^{1/m}}{\eta}\right) \in \square_\eta \square_m^{-1} \tau_m(I_{(1)}).$$

Then

$$\begin{aligned} (19) \quad \tilde{I}_{(\eta)}^{R,l}\left(\frac{q^{1/m}}{\eta}\right) &= \sum_{\beta} \exp\left(-\frac{1}{2} \text{ch}_1 \mathcal{R}\right)(I_{\beta})_{(\eta)}\left(\frac{q^{1/m}}{\eta}\right) Q^{\beta} \prod_i (L_i^{-\beta_i} q^{(\beta_i+1)\beta_i/(2m)} \eta^{-(\beta_i+1)\beta_i/2})^l \\ &= \exp\left(-l \text{ch}_2 \mathcal{R} \cdot \left(\frac{z}{m} - \log \eta\right)^{-1}\right) (\Phi(I_{(\eta)})) \left(\frac{q^{1/m}}{\eta}\right) \\ &= \exp\left(-ml \left(\frac{\text{ch}_2 \mathcal{R}}{z}\right)\right) (\Phi(I_{(\eta)})) \left(\frac{q^{1/m}}{\eta}\right). \end{aligned}$$

Here we use the fact that $m \cdot \log(\eta) = 0$.

By an elementary computation, one can show that, for any series f in q and Q , we have

$$(20) \quad (\Phi(f))\left(\frac{q^{1/m}}{\eta}\right) = \Phi^m \mathcal{D}_{\eta} f\left(\frac{q^{1/m}}{\eta}\right),$$

where the operator $\mathcal{D}_{\eta}$ is defined by

$$\mathcal{D}_{\eta} = \prod_{i=1}^n \exp\left(l\left(-\log(\eta) f_i(Q_j\partial Q_j)^2 + \frac{f_i(\log(\eta)Q_j\partial Q_j)}{2} - \frac{m-1}{m} \frac{f_i(mp_j - zQ_j\partial Q_j)}{2}\right)\right).$$

Here the substitution $q \mapsto q^{1/m}/\eta$ corresponds to the change $z \mapsto z/m - \log(\eta)$ in the expression of Φ .

It follows from (20) that (19) equals

$$(21) \quad \left(D_m\Phi^m\mathcal{D}_{\eta}\left(I_{\eta}\left(\frac{q^{1/m}}{\eta}\right)\right)\right) \in D_m\Phi^m\mathcal{D}_{\eta} \square_{\eta} \square_m^{-1} \tau_m(I_{(1)}).$$

Since all the operators above have constant coefficients (i.e., independent of Q), they commute. We claim that $\mathcal{D}_{\eta}$ preserves $\tau_m(I_{(1)})$. This is because, by definition, we have

$$\mathcal{D}_{\eta}\tau_m(I_{(1)}) = \mathcal{D}_{\eta}\Psi^m \circ S(q, Q) \circ \Psi^{1/m}\mathcal{K}_{+}^1.$$

Let

$$\mathcal{D}_{\eta,1} = \prod_{i=1}^n \exp\left(l\left(-\log(\eta) f_i(Q_j\partial Q_j)^2 + \frac{1}{2} f_i(\log(\eta)Q_j\partial Q_j)\right)\right)$$

²This is because $\Psi^m(zQ_j\partial Q_j) = mzQ_j^m\partial Q_j^m = zQ_j\partial Q_j$.

and

$$\mathcal{D}_{\eta,2} = \prod_{i=1}^n \exp \left(l \left(- \frac{(m-1) f_i (mp_j - z Q_j \partial_{Q_j})}{2m} \right) \right)$$

be the factors of $\mathcal{D}_{\eta}$. Then it is easy to check that the first factor $\mathcal{D}_{\eta,1}$ commutes with $\Psi^m \circ S(q, Q) \circ \Psi^{1/m}$, and hence preserves $\mathcal{T}_m(I_{(1)})$. The second factor $\mathcal{D}_{\eta,2}$ satisfies the commutation relation

$$\mathcal{D}_{\eta,2} \Psi^m = \Psi^m \prod_{i=1}^n \exp \left(l \left(- \frac{(m-1) f_i (p_j - z Q_j \partial_{Q_j})}{2m} \right) \right).$$

According to [Givental and Tonita 2014, Corollary 1], the second operator on the right-hand side preserves the tangent space $S(q, Q) \circ \Psi^{1/m} \mathcal{K}_+^1$. Therefore, we have shown that the space $\mathcal{T}_m(I_{(1)})$ is $\mathcal{D}_{\eta}$ -invariant.

We can further simplify the space on the right-hand side of (21) as

$$\begin{aligned} D_m \Phi^m \mathcal{D}_{\eta} \square_{\eta} \square_m^{-1} \mathcal{T}_m(I_{(1)}) &= D_m \Phi^m \square_{\eta} \square_m^{-1} \mathcal{D}_{\eta} \mathcal{T}_m(I_{(1)}) \\ &= D_m \Phi^m \square_{\eta} \square_m^{-1} \mathcal{T}_m(I_{(1)}) \\ &= D_m \Phi^m \square_{\eta} \square_m^{-1} (\Phi^m)^{-1} D_m^{-1} (\mathcal{T}_m^{R,l}(\tilde{I}_{(1)}^{R,l})) \\ &= \square_{\eta} \square_m^{-1} (\mathcal{T}_m^{R,l}(\tilde{I}_{(1)}^{R,l})). \end{aligned} \quad \square$$

Remark 4.15 Suppose the target is an orbifold. As explained in Remark 4.14, the adelic characterization of points on the untwisted sector of the Lagrangian cone of X is the same as the one given in Theorem 4.5. Using the same proof as above, we can show that, if I is a point on the untwisted sector of $\mathcal{L}_{S_{\infty}}$, then the determinantal modification $I^{R,l}$ lies on the untwisted sector of $\mathcal{L}_{S_{\infty}}^{R,l}$.

5 Toric mirror theorem and mock theta functions

In this section, we first recall a mirror theorem proved in [Givental 2015d] for ordinary quantum K -theory of toric varieties. This corresponds to Theorem 1.2 in the level 0 case. By combining Givental's result with Theorem 1.1, proved in the previous section, we obtain a toric mirror theorem for level- l quantum K -theory. In some simple cases, we recover Ramanujan's mock theta functions from toric I -functions with nontrivial level structure.

5.1 K -theoretic I -function and mock theta function

In this subsection, we first explicitly compute the (torus-equivariant) small I -functions with level structures for toric varieties, using quasimap graph spaces. Then we use torus localization to prove a toric mirror theorem (Theorem 1.2), following [Givental 2015d]. In the study of quantum K -theory with nontrivial level structures, a remarkable phenomenon is the appearance of Ramanujan's mock theta functions.

Let $M \cong \mathbb{Z}^n$ be a n -dimensional lattice and let N be its dual lattice. To every complete nonsingular fan $\Sigma \subset N_{\mathbb{R}}$, we can associate a n -dimensional smooth projective variety X_{Σ} . We denote by $\Sigma(1)$ the set

of 1-dimensional cones in Σ . Let $m = |\Sigma(1)|$. Each $\rho \in \Sigma(1)$ determines a Weil divisor D_ρ on X_Σ and the Picard group of X_Σ is determined by the short exact sequence

$$(22) \quad 0 \rightarrow M \rightarrow \mathbb{Z}^{\Sigma(1)} \rightarrow \text{Pic}(X_\Sigma) \rightarrow 0.$$

Here the inclusion is defined by $m \mapsto \sum_\rho \langle m, \rho \rangle D_\rho$. Now let us describe the quotient construction of X_Σ . Since $\text{Pic}(X_\Sigma)$ is torsion free, we choose an integral basis $\{L_1, \dots, L_s\}$ of it, where $s = m - n$. Then the surjective map in (22) is given by the transpose of an integral $s \times m$ matrix $Q = (Q_{a\rho_b})_{a=1, \dots, s; b=1, \dots, m}$ which is called the charge matrix of X_Σ . Applying $\text{Hom}(-, \mathbb{C}^*)$ to the exact sequence (22), we get an exact sequence.

$$1 \rightarrow G \rightarrow (\mathbb{C}^*)^{\Sigma(1)} \rightarrow N \otimes \mathbb{C}^* \rightarrow 1,$$

where $G := \text{Hom}(\text{Pic}(X_\Sigma), \mathbb{C}^*) \cong (\mathbb{C}^*)^s$. The first map in the above short exact sequence defines the G -action on $\mathbb{C}^{\Sigma(1)}$

$$(23) \quad \mathbf{t} \cdot (z_{\rho_1}, \dots, z_{\rho_m}) = \left(\prod_{a=1}^s t_a^{Q_{a\rho_1}} z_{\rho_1}, \dots, \prod_{a=1}^s t_a^{Q_{a\rho_m}} z_{\rho_m} \right),$$

where $\mathbf{t} = (t_1, \dots, t_s) \in (\mathbb{C}^*)^s$. By choosing an appropriate linearization of the trivial line bundle on $\mathbb{C}^{\Sigma(1)}$ (see, e.g., [Dolgachev 2003, Chapter 12]), the semistable and stable loci are equal. We denote this linearized trivial line bundle by L_Σ and the stable loci by $U(\Sigma)$. Let z_ρ be the coordinates in $\mathbb{C}^{\Sigma(1)}$. We define a subvariety

$$Z(\Sigma) = \left\{ (z_\rho) \in \mathbb{C}^{\Sigma(1)} \mid \prod_{\rho \notin \sigma} z_\rho = 0, \sigma \in \Sigma \right\}.$$

Then

$$U(\Sigma) = \mathbb{C}^{\Sigma(1)} \setminus Z(\Sigma).$$

The toric variety X_Σ is the geometric quotient $U(\Sigma)/G$. Let P be the principal G -bundle $\mathbb{C}^{\Sigma(1)} \rightarrow [\mathbb{C}^{\Sigma(1)}/G]$. Let $\pi_i : G \rightarrow \mathbb{C}^*$ be the projection to the i -th component and let R_j be the characters given by $\mathbf{t} = (t_1, \dots, t_s) \mapsto \prod_{a=1}^s t_a^{Q_{a\rho_j}}$ for $1 \leq j \leq m$. Then the line bundles L_i and $\mathcal{O}(-D_{\rho_j})$ are the restrictions of the associated line bundles of P with the characters π_i and R_j , respectively, to X_Σ .

Note that X_Σ admits a $T^m := (\mathbb{C}^*)^{\Sigma(1)}$ -action, i.e., T^m acts by diagonal matrix on $\mathbb{C}^{\Sigma(1)}$. We denote by P_i and U_ρ the T^m -equivariant line bundles corresponding to L_i and $\mathcal{O}(D_\rho)$, respectively. In the T^m -equivariant K -group $K_{T^m}^0(X_\Sigma) \otimes \mathbb{Q}$, we have the multiplicative relation

$$U_\rho = \prod_{i=1}^s P_i^{\otimes Q_{i\rho}} \Lambda_\rho^{-1},$$

where Λ_ρ are the generators of $\text{Repr}(T^m)$ corresponding to the projection to the component labeled by ρ . See [Givental 2015d] for more details.

Now let us compute the (T^m -equivariant) small I -function of X_Σ with level structures, using the quasimap graph space. Let $\beta \in \text{Hom}_{\mathbb{Z}}(\text{Pic}^{T^m}(\mathbb{C}^{\Sigma(1)}), \mathbb{Z})$ be an L_Σ -effective class. According to

[Ciocan-Fontanine and Kim 2010, Lemma 3.1.8], a point in the quasimap graph space $QG_{0,k}^{\epsilon=0+}(X_\Sigma, \beta)$ is specified by the data

$$((C, p_1, \dots, p_k), \{\mathcal{P}_i \mid i = 1, \dots, s\}, \{u_\rho\}_{\rho \in \Sigma(1)}, \varphi),$$

where

- $(C, p_1, \dots, p_s)$ is a connected, at most nodal, curve of genus 0 and p_i are distinct nonsingular points of C ,
- $\mathcal{P}_i$ are line bundles on C of degree $f_i := \beta(L_i)$,
- $u_\rho \in \Gamma(C, \mathcal{L}_\rho)$, where $\mathcal{L}_\rho$ is defined by

$$\mathcal{L}_\rho := \bigotimes_{i=1}^s \mathcal{P}_i^{\otimes Q_{i\rho}},$$

- $\varphi : C \rightarrow \mathbb{P}^1$ is a regular map such that $\varphi_*[C] = [\mathbb{P}^1]$.

The stability conditions are discussed in Section 3.2. When $(g, k) = (0, 0)$, we have $C \cong \mathbb{P}^1$ and $\mathcal{P}_i \cong \mathcal{O}_{\mathbb{P}^1}(f_i)$. The line bundles $\mathcal{L}_\rho$ are isomorphic to $\mathcal{O}_{\mathbb{P}^1}(\sum_{i=1}^s f_i Q_{i\rho}) = \mathcal{O}_{\mathbb{P}^1}(\beta_\rho)$, where $\beta_\rho := \beta(\mathcal{O}(D_\rho))$. Therefore, a point on $QG_{0,0}^{\epsilon=0+}(X_\Sigma, \beta)$ is specified by sections $\{u_\rho \in \Gamma(\mathbb{P}^1, \mathcal{O}_{\mathbb{P}^1}(\beta_\rho)) \mid \rho \in \Sigma(1)\}$. We choose coordinates $[x_0, x_1]$ on $\mathbb{P}^1$ and consider the standard action $\mathbb{C}^*$ -action defined by (4). Let F_0 be the distinguished fixed-point locus parametrizing quasimaps whose degrees are concentrated only at 0. According to [Ciocan-Fontanine and Kim 2010, Section 7.2], we have the identification

$$(24) \quad F_0 \cong \bigcap_{\{\rho \mid \beta_\rho < 0\}} D_\rho \subset X_\Sigma, \quad (z_\rho x_0^{\beta_\rho}) \mapsto (z_\rho),$$

where (z_ρ) are the coordinates on X_Σ .

Let R be a character of $G = (\mathbb{C}^*)^s$ defined by $t \cdot z = \prod_{i=1}^s t_i^{r_i} z$ where $r_i \in \mathbb{Z}$. Recall that the small I -function of X_Σ of level l and representation R is defined by

$$I^{R,l}(q) = 1 + \sum_a \sum_{\beta \neq 0} Q^\beta \chi \left(F_0, \text{ev}^*(\phi_a) \otimes \left(\frac{\text{tr}_{\mathbb{C}^*} \mathcal{D}^{R,l}}{\text{tr}_{\mathbb{C}^*} \wedge^*(N_{F_0/QG}^{\text{vir}})^\vee} \right) \right) \phi^a.$$

It is not difficult to check that, under the identification (24), we can identify the virtual normal bundle $N_{F_0/QG}^{\text{vir}}$ in $K^0(F_0)$ with

$$(25) \quad N_{F_0/QG}^{\text{vir}} = \sum_{\{\rho \mid \beta_\rho > 0\}} \sum_{i=1}^{\beta_\rho} \mathcal{O}(D_\rho)|_{F_0} \otimes \mathbb{C}_{-i} - \sum_{\{\rho \mid \beta_\rho < 0\}} \sum_{i=1}^{-\beta_\rho-1} \mathcal{O}(D_\rho)|_{F_0} \otimes \mathbb{C}_i,$$

where $\mathbb{C}_a$ denotes the representation of $\mathbb{C}^*$ on $\mathbb{C}$ with weight $a \in \mathbb{Z}$.³ Let $\mathcal{P}$ be the universal principal G -bundle on $F_0 \times \mathbb{P}^1 \subset X_\Sigma \times \mathbb{P}^1$. Then the associated line bundle $\mathcal{P} \times_G R$ can be identified with

³Here, the $\mathbb{C}^*$ is the torus acting on the graph space.

$\bigotimes_{i=1}^s L_i^{r_i} \otimes \mathcal{O}_{\mathbb{P}^1}(\beta_R)$, where $\beta_R := \sum_{i=1}^s r_i f_i$. We denote the line bundle $\bigotimes_{i=1}^s L_i^{r_i}$ by $\mathcal{R}$. Let $\pi : F_0 \times \mathbb{P}^1 \rightarrow F_0$ be the projection. When $\beta_R \geq 0$, we have

$$(26) \quad \mathcal{D}^{R,l} = \det^{-l} R\pi_*(\mathcal{R} \otimes \mathcal{O}_{\mathbb{P}^1}(\beta_R)) = \det^{-l} (\mathcal{R} \otimes R^0\pi_*(\mathcal{O}_{\mathbb{P}^1}(\beta_R))) = \mathcal{R}^{-l(\beta_R+1)} \otimes \mathbb{C}_{l\beta_R(\beta_R+1)/2}.$$

When $\beta_R < 0$, a similar calculation derives the same formula $\mathcal{D}^{R,l} = \mathcal{R}^{-l(\beta_R+1)} \otimes \mathbb{C}_{l\beta_R(\beta_R+1)/2}$.

We give the explicit formulas of the (torus-equivariant) small I -functions of toric varieties in the following proposition:

Proposition 5.1 *The small I -function of a toric variety X_Σ of level l and character R is given by*

$$I^{R,l}(q) = 1 + \sum_{\beta \in \text{Eff}(X)} Q^\beta \mathcal{R}^{-l\beta_R} q^{l\beta_R(\beta_R+1)/2} \prod_{\rho \in \Sigma(1)} \frac{\prod_{j=-\infty}^0 (1 - \mathcal{O}(-D_\rho)q^j)}{\prod_{j=-\infty}^{\beta_\rho} (1 - \mathcal{O}(-D_\rho)q^j)},$$

and its equivariant version is given by

$$I^{R,l,\text{eq}}(q) = 1 + \sum_{\beta \in \text{Eff}(X)} Q^\beta \tilde{\mathcal{R}}^{-l\beta_R} q^{l\beta_R(\beta_R+1)/2} \prod_{\rho \in \Sigma(1)} \frac{\prod_{j=-\infty}^0 (1 - U_\rho q^j)}{\prod_{j=-\infty}^{\beta_\rho} (1 - U_\rho q^j)}.$$

Here $\mathcal{R} := \bigotimes_{i=1}^s L_i^{r_i}$ is the line bundle associated to the character R , and $\tilde{\mathcal{R}} \otimes \bigotimes_{i=1}^s P_i^{r_i}$ and U_ρ are the equivariant line bundles corresponding to $\mathcal{R}$ and $\mathcal{O}(-D_\rho)$, respectively.

Proof Note that the evaluation map $\text{ev}_\bullet$ makes F_0 as a closed subvariety of X_Σ , whose normal bundle is given by $\bigoplus_{\rho: \beta_\rho < 0} \mathcal{O}(D_\rho)$. Then, the proposition follows easily from (25) and (26). Note that one factor $\mathcal{R}^{-l}$ in (26) disappears due to the change of pairings (see Convention 4.4). $\square$

By extending Givental's localization argument [2015d] to the setting with level structure, we prove a toric mirror theorem. Now let us restate Theorem 1.2.

Theorem 1.2 *Assume that X_Σ is a smooth quasiprojective toric variety. Let $I^{R,l,\text{eq}}(q)$ be the level- l torus-equivariant small I -function given in Proposition 5.1. Then the series $(1-q)I^{R,l,\text{eq}}(q)$ lies on the cone $\mathcal{L}_{S_\infty}^{R,l,\text{eq}}$ in the symmetrized torus-equivariant quantum v -theory of level l of X_Σ .*

Proof For simplicity, we denote by I the I -function $I^{R,l,\text{eq}}$ of X_Σ . Let $\{\phi_\alpha\}_{\alpha \in X_\Sigma^{T^m}}$ be the fixed-point basis of $K_{T^m}^0(X_\Sigma)$ and let $\{\phi^\alpha\}$ be the dual basis with respect to the pairing (5). For each fixed point α , we denote by $J(\alpha) \subset \Sigma(1)$ the cardinality- s subset such that α equals the intersection $\bigcap_{\rho \notin J(\alpha)} D_\rho$. Write $I = \sum_\alpha I^{(\alpha)} \phi_\alpha$. We denote by $U_\rho(\alpha)$ and $\mathcal{R}(\alpha)$ the restrictions of U_ρ and $\mathcal{R}$ to the fixed point α , respectively. For $\rho \in J(\alpha)$, we have $U_\rho(\alpha) = 1$. Hence $I^{(\alpha)}$ can be explicitly written as

$$I^{(\alpha)}(q) = 1 + \sum_{\beta \in \text{Eff}'(X_\Sigma)} Q^\beta \frac{\mathcal{R}(\alpha)^{-l\beta_R} q^{l\beta_R(\beta_R+1)/2}}{\prod_{\rho \in J(\alpha)} \prod_{j=1}^{\beta_\rho} (1 - q^j)} \prod_{\rho \notin J(\alpha)} \frac{\prod_{j=-\infty}^0 (1 - U_\rho(\alpha)q^j)}{\prod_{j=-\infty}^{\beta_\rho} (1 - U_\rho(\alpha)q^j)}.$$

Here $\text{Eff}'(X_\Sigma)$ denotes the semigroup of effective curve classes β such that $\beta_\rho \geq 0$. The terms with $\beta_\rho < 0$ disappear because there is a factor $(1 - q^0)$ in the numerators.

We first observe that, for the point target, the cone $\mathcal{L}_{S_\infty}^{\text{pt}, R, l}$ of level l coincides with the cone $\mathcal{L}_{S_\infty}^{\text{pt}}$ of level 0. This is because, in the case of the point target, the determinant line bundle $\mathcal{D}^{R, l}$ is always topologically trivial (with possible equivariant weights). Combining this observation with the fact that the level structure $\mathcal{D}^{R, l}$ splits “nicely” among nodal strata, we can extend the argument in [Givental 2015a] to prove that a point $f = \sum_{\alpha} f^{(\alpha)} \phi_{\alpha}$ lies on $\mathcal{L}_{S_\infty}^{R, l, \text{eq}}$ if and only if the following are satisfied:

- (1) When expanded as meromorphic functions with poles only at roots of unity, $f^{(\alpha)}$ lie on the cone $\mathcal{L}_{S_\infty}^{\text{pt}}$ in the permutation-equivariant quantum K -theory of the point target space.
- (2) Away from $q = 0, \infty$ and roots of unity, $f^{(\alpha)}$ may have at most simple poles at $q = U_{\rho}(\alpha)^{-1/m}$ with $\rho \notin J(\alpha)$ and $m = 1, 2, \dots$ for generic values of $\Lambda_1, \dots, \Lambda_m$. The residues satisfy the recursion relations

$$\text{Res}_{q=U_{\rho}(\alpha)^{-1/m}} f^{(\alpha)}(q) \frac{dq}{q} = -\frac{\phi^{\alpha} Q^{md_{\alpha\rho}}}{m \cdot C_{\alpha\rho}(m)} f^{(\rho)}(U_{\rho}(\alpha)^{-1/m}).$$

Here $C_{\alpha\rho}(m) = \lambda_{-1}(T_p \overline{\mathcal{M}}_{0,2}(X_{\Sigma}, md_{\alpha\rho})) \cdot (\mathcal{R}(\alpha)^{-md_{\alpha\rho}^R} q^{md_{\alpha\rho}^R(md_{\alpha\rho}^R+1)/2})^l$, where:

- (a) T_p denotes the virtual tangent space to the moduli space at the point p represented by the m -multiple cover of the one-dimensional orbit connecting α and ρ . The explicit formula of the equivariant weights of the K -theoretic Euler class $\lambda_{-1}(T_p)$ is given in [Givental 2015d].
- (b) $d_{\alpha\rho}$ denotes the degree of the one-dimensional orbit connecting α and ρ .
- (c) $d_{\alpha\rho}^R := \langle d_{\alpha\rho}, c_1(\mathcal{R}) \rangle$.

We want to show that $(1-q)I$ satisfies (1) and (2). It is proved in the main theorem⁴ of [Givental 2015d] that the series

$$\tilde{I}^{(\alpha)} = 1 + \sum_{\beta \in \text{Eff}^r(X_{\Sigma})} \frac{Q^{\beta}}{\prod_{\rho \in J(\alpha)} \prod_{j=1}^{\beta_{\rho}} (1-q^j)} \prod_{\rho \notin J(\alpha)} \frac{\prod_{j=-\infty}^0 (1-U_{\rho}(\alpha)q^j)}{\prod_{j=-\infty}^{\beta_{\rho}} (1-U_{\rho}(\alpha)q^j)}$$

represents a value of $\mathcal{J}_{S_\infty}^{\text{pt}}(t(q), Q)/(1-q)$, i.e., $(1-q)\tilde{I}^{(\alpha)}$ lies on the cone $\mathcal{L}_{S_\infty}^{\text{pt}}$. Note that $I^{(\alpha)}$ is obtained from $\tilde{I}^{(\alpha)}$ by a “determinantal” modification. Therefore, it follows from Theorem 1.1 that $(1-q)I^{(\alpha)}$ lies on $\mathcal{L}_{S_\infty}^{\text{pt}, R, l} = \mathcal{L}_{S_\infty}^{\text{pt}}$.

To prove that $(1-q)I^{(\alpha)}$ satisfies the second condition, we rewrite $\mathcal{R}(\alpha)^{-l\beta_R} q^{l\beta_R(\beta_R+1)/2}$ as

$$(27) \quad \frac{\prod_{j=-\infty}^{\beta_R} (\mathcal{R}(\alpha)^{-1} q^j)^l}{\prod_{j=-\infty}^0 (\mathcal{R}(\alpha)^{-1} q^j)^l}.$$

Note that, for all j , we have $\mathcal{R}(\alpha) = \mathcal{R}(\beta) \lambda^{-d_{\alpha\rho}^R}$, where $\lambda = U_{\rho}(\alpha)$. Hence, at $q = \lambda^{-1/m}$,

$$\mathcal{R}(\alpha)^{-1} q^j = \mathcal{R}(\beta)^{-1} q^{j-md_{\alpha\rho}^R}.$$

⁴Givental [2015d] defines the I -function to sum over all $\beta \in \mathbb{Z}^m$. However, the same argument works if we restrict the summation to curve classes in the semigroup $\text{Eff}(X_{\Sigma})$.

The formula (27) is equivalent to

$$\mathcal{R}(\alpha)^{-ml d_{\alpha\rho}^R} q^{ml d_{\alpha\rho}^R (md_{\alpha\rho}^R + 1)/2} \frac{\prod_{j=-\infty}^{\beta_R - md_{\alpha\rho}^R} (\mathcal{R}(\beta)^{-1} q^j)^l}{\prod_{j=-\infty}^0 (\mathcal{R}(\beta)^{-1} q^j)^l}$$

at $q = \lambda^{-1/m}$. Combining the above equivalence with the result in [Givental 2015d, page 10], we obtain

$$\text{Res}_{q=U_\rho(\alpha)^{-1/m}} (1-q) I^{(\alpha)}(q) = -\frac{\phi^\alpha Q^{md_{\alpha\rho}}}{m \cdot C_{\alpha\rho}(m)} (1 - U_\rho(\alpha)^{-1/m}) I^{(\rho)}(U_\rho(\alpha)^{-1/m}),$$

which is equivalent to the residue formula in condition (2).

Since I is defined over the λ -algebra $\mathbb{Z}[\Lambda_1^\pm, \dots, \Lambda_m^\pm][[Q]]$, it takes values in the symmetrized theory (see Remark 3.3). $\square$

When X_Σ is projective, we may pass to the nonequivariant limit in Theorem 1.2 to obtain:

Corollary 5.2 *When X_Σ is a smooth projective toric variety, the level- l small I -function $I^{R,l}$ given in Proposition 5.1 lies on the cone $\mathcal{L}_{S_\infty}^{R,l}$ in the symmetrized quantum K -theory of level l .*

We denote by St and $\text{St}^\vee$ the standard representation and its dual representation of $\mathbb{C}^*$. As corollaries of Proposition 5.1, we give proofs for Propositions 1.3–1.5 and 1.7.

Proof of Proposition 1.3 Let the target be $X = (\mathbb{C} \setminus 0)/\mathbb{C}^*$, where the action is the standard action. Then the proposition follows directly from Proposition 5.1. $\square$

Proof of Proposition 1.4 Let the target be $X_{a_1, a_2} = (\mathbb{C}^2 \setminus \{(0, 0)\})/\mathbb{C}^*$ with charge vector (a_1, a_2) . As mentioned in Remarks 1.6, 4.14 and 4.15, we only consider the untwisted component of the orbifold I -function, and its formula is given by Proposition 5.1. $\square$

Proof of Proposition 1.5 For positive integers a, b , we consider $X_{a, -b} = \{(\mathbb{C} - 0) \times \mathbb{C}\}/\mathbb{C}^*$ with charge vector $(a, -b)$. Let λ, μ be the generators of $\text{Repr}((\mathbb{C}^*)^2)$ corresponding to the first and second projections of $(\mathbb{C}^*)^2$ onto its factors. From Proposition 5.1, the untwisted component of the orbifold I -function is given by

$$\begin{aligned} & I_{X_{a, -b}}^{\text{St}, l}(q) \\ &= 1 + \sum_{n \geq 1} \frac{p^{nl} q^{n(n-1)l/2} (1-p^{-b} \mu^{-1}) (1-p^{-b} \mu^{-1} q^{-1}) \cdots (1-p^{-b} \mu^{-1} q^{1-bn})}{(1-p^a \lambda^{-1} q) (1-p^a \lambda^{-1} q^2) \cdots (1-p^a \lambda^{-1} q^{an})} Q^n \\ &= 1 + \sum_{n \geq 1} (-1)^{bn} \frac{p^{nl-b^2n} q^{(n(n-1)l-bn(bn-1))/2} \mu^{-bn} (1-p^b \mu) (1-p^b \mu q^1) \cdots (1-p^b \mu q^{bn-1})}{(1-p^a \lambda^{-1} q) (1-p^a \lambda^{-1} q^2) \cdots (1-p^a \lambda^{-1} q^{an})} Q^n. \quad \square \end{aligned}$$

Proof of Proposition 1.7 We consider the target $O(-1)_{\mathbb{P}^{s-1}}^{\oplus r} = X_{\mathbf{1}, -\mathbf{1}} = \{(\mathbb{C}^s - 0) \times \mathbb{C}^r\} / \mathbb{C}^*$ with the charge vector $(1, 1, \dots, 1, -1, -1, \dots, -1)$. It follows from Proposition 5.1 that

$$\begin{aligned} I_{X_{\mathbf{1}, -\mathbf{1}}}^{\text{St}, l=1+s}(q) &= 1 + \sum_{n \geq 1} Q^n p^{nl} q^{n(n-1)l/2} \frac{(p^{-1}\mu_1^{-1}, q)_n \cdots (p^{-1}\mu_r^{-1}, q)_n}{(p\lambda_1^{-1}q; q)_n \cdots (p\lambda_s^{-1}q; q)_n} \\ &= 1 + \sum_{n \geq 1} (-1)^{nr} \prod_{i=1}^r (p\mu_i)^{-n} p^{(1+s)n} \frac{(p\mu_1, q)_n \cdots (p\mu_r, q)_n}{(p\lambda_1^{-1}q; q)_n \cdots (p\lambda_s^{-1}q; q)_n} Q^n (q^{n(n-1)/2})^{1+s-r}. \quad \square \end{aligned}$$

Remark 5.3 Traditionally, there are two approaches to prove genus-zero mirror theorems: an older approach of Givental and Tonita using adelic descriptions of cones, and a more recent wall-crossing approach of Ciocan-Fontanine and Kim using quasimap theory. In the first version of our paper, we used the wall-crossing approach. Recently, we discovered a gap in the proof. Roughly speaking, the K -theoretic version of the polynomiality property [Ciocan-Fontanine and Kim 2014, Lemma 7.6.1] is not strong enough to determine the coefficients of the K -theoretic S -operators and J -functions recursively. In the current version, we switch back to Givental and Tonita’s older technique. It is still an interesting question if we can improve the wall-crossing approach to prove the mirror theorem in the K -theory setting.

Acknowledgements

Ruan would like to thank Aaron Bertram, Prakash Belkale and Jie Zhou for interesting discussions. Zhang would like to thank Felix Janda, Feng Qu, Valentin Tonita, Rachel Webb and Fenglong You for helpful discussions. We also thank the referee for valuable comments and suggestions on a previous draft.

Ruan is partially supported by the NSF grant DMS 1405245 and the NSF FRG grant DMS 1159265.

References

- [Abramovich et al. 2002] **D Abramovich, T Graber, A Vistoli**, *Algebraic orbifold quantum products*, from “Orbifolds in mathematics and physics” (Madison, WI, 2001) (A Adem, J Morava, Y Ruan, editors), Contemp. Math. 310, Amer. Math. Soc., Providence, RI (2002) 1–24 [MR](#)
- [Abramovich et al. 2003] **D Abramovich, A Corti, A Vistoli**, *Twisted bundles and admissible covers*, Comm. Algebra 31:8 (2003) 3547–3618 [MR](#)
- [Abramovich et al. 2007] **D Abramovich, T Graber, M Olsson, H-H Tseng**, *On the global quotient structure of the space of twisted stable maps to a quotient stack*, J. Algebraic Geom. 16:4 (2007) 731–751 [MR](#)
- [Abramovich et al. 2008] **D Abramovich, T Graber, A Vistoli**, *Gromov–Witten theory of Deligne–Mumford stacks*, Amer. J. Math. 130:5 (2008) 1337–1398 [MR](#)
- [Aganagic and Okounkov 2017] **M Aganagic, A Okounkov**, *Quasimap counts and Bethe eigenfunctions*, Mosc. Math. J. 17:4 (2017) 565–600 [MR](#)
- [Aganagic et al. 2018] **M Aganagic, E Frenkel, A Okounkov**, *Quantum q -Langlands correspondence*, Trans. Moscow Math. Soc. 79 (2018) 1–83 [MR](#)
- [Agnihotri 1995] **S Agnihotri**, *Quantum cohomology and the Verlinde algebra*, PhD thesis, University of Oxford (1995)

- [Behrend and Manin 1996] **K Behrend, Y Manin**, *Stacks of stable maps and Gromov–Witten invariants*, Duke Math. J. 85:1 (1996) 1–60 MR
- [Belkale 2008] **P Belkale**, *Quantum generalization of the Horn conjecture*, J. Amer. Math. Soc. 21:2 (2008) 365–408 MR
- [Chen and Ruan 2002] **W Chen, Y Ruan**, *Orbifold Gromov–Witten theory*, from “Orbifolds in mathematics and physics” (Madison, WI, 2001) (A Adem, J Morava, Y Ruan, editors), Contemp. Math. 310, Amer. Math. Soc., Providence, RI (2002) 25–85 MR
- [Cheong et al. 2015] **D Cheong, I Ciocan-Fontanine, B Kim**, *Orbifold quasimap theory*, Math. Ann. 363:3–4 (2015) 777–816 MR
- [Ciocan-Fontanine and Kim 2010] **I Ciocan-Fontanine, B Kim**, *Moduli stacks of stable toric quasimaps*, Adv. Math. 225:6 (2010) 3022–3051 MR
- [Ciocan-Fontanine and Kim 2014] **I Ciocan-Fontanine, B Kim**, *Wall-crossing in genus zero quasimap theory and mirror maps*, Algebr. Geom. 1:4 (2014) 400–448 MR
- [Ciocan-Fontanine et al. 2014] **I Ciocan-Fontanine, B Kim, D Maulik**, *Stable quasimaps to GIT quotients*, J. Geom. Phys. 75 (2014) 17–47 MR
- [Coates 2003] **T H Coates**, *Riemann–Roch theorems in Gromov–Witten theory*, PhD thesis, University of California, Berkeley (2003) MR Available at <https://www.proquest.com/docview/305341155>
- [Coates and Givental 2006] **T Coates, A Givental**, *Quantum cobordisms and formal group laws*, from “The unity of mathematics” (P Etingof, V Retakh, IM Singer, editors), Progr. Math. 244, Birkhäuser, Boston, MA (2006) 155–171 MR
- [Coates and Givental 2007] **T Coates, A Givental**, *Quantum Riemann–Roch, Lefschetz and Serre*, Ann. of Math. 165:1 (2007) 15–53 MR
- [Dolgachev 2003] **I Dolgachev**, *Lectures on invariant theory*, London Mathematical Society Lecture Note Series 296, Cambridge Univ. Press (2003) MR
- [Gepner 1991] **D Gepner**, *Fusion rings and geometry*, Comm. Math. Phys. 141:2 (1991) 381–411 MR
- [Givental 2000] **A Givental**, *On the WDVV equation in quantum K -theory*, Michigan Math. J. 48 (2000) 295–304 MR
- [Givental 2004] **A B Givental**, *Symplectic geometry of Frobenius structures*, from “Frobenius manifolds” (C Hertling, M Marcolli, editors), Aspects Math. E36, Friedr. Vieweg, Wiesbaden (2004) 91–112 MR
- [Givental 2015a] **A Givental**, *Permutation-equivariant quantum K -theory, II: Fixed point localization*, preprint (2015) arXiv 1508.04374
- [Givental 2015b] **A Givental**, *Permutation-equivariant quantum K -theory, III: Lefschetz’ formula on $\overline{\mathcal{M}}_{0,n}/S_n$ and adelic characterization*, preprint (2015) arXiv 1508.06697
- [Givental 2015c] **A Givental**, *Permutation-equivariant quantum K -theory, IV: D_q -modules*, preprint (2015) arXiv 1509.00830
- [Givental 2015d] **A Givental**, *Permutation-equivariant quantum K -theory, V: Toric q -hypergeometric functions*, preprint (2015) arXiv 1509.03903
- [Givental 2015e] **A Givental**, *Permutation-equivariant quantum K -theory, VI: Mirrors*, preprint (2015) arXiv 1509.07852
- [Givental 2015f] **A Givental**, *Permutation-equivariant quantum K -theory, VII: General theory*, preprint (2015) arXiv 1510.03076
- [Givental 2015g] **A Givental**, *Permutation-equivariant quantum K -theory, VIII: Explicit reconstruction*, preprint (2015) arXiv 1510.06116
- [Givental 2017a] **A Givental**, *Permutation-equivariant quantum K -theory, I: Definitions: Elementary K -theory of $\overline{\mathcal{M}}_{0,n}/S_n$* , Mosc. Math. J. 17:4 (2017) 691–698 MR
- [Givental 2017b] **A Givental**, *Permutation-equivariant quantum K -theory, IX: Quantum Hirzebruch–Riemann–Roch in all genera*, preprint (2017) arXiv 1709.03180
- [Givental 2017c] **A Givental**, *Permutation-equivariant quantum K -theory, XI: Quantum Adams–Riemann–Roch*, preprint (2017) arXiv 1711.04201

- [Givental 2020] **A Givental**, *Permutation-equivariant quantum K -theory, X: Quantum Hirzebruch–Riemann–Roch in genus 0*, Symmetry Integrability Geom. Methods Appl. 16 (2020) art. id. 031 MR
- [Givental and Tonita 2014] **A Givental, V Tonita**, *The Hirzebruch–Riemann–Roch theorem in true genus-0 quantum K -theory*, from “Symplectic, Poisson, and noncommutative geometry” (T Eguchi, Y Eliashberg, Y Maeda, editors), Math. Sci. Res. Inst. Publ. 62, Cambridge Univ. Press (2014) 43–91 MR
- [Intriligator 1991] **K Intriligator**, *Fusion residues*, Modern Phys. Lett. A 6:38 (1991) 3543–3556 MR
- [Jockers and Mayr 2020] **H Jockers, P Mayr**, *A 3d gauge theory/quantum K -theory correspondence*, Adv. Theor. Math. Phys. 24:2 (2020) 327–457 MR
- [Kapustin and Willett 2013] **A Kapustin, B Willett**, *Wilson loops in 3-dimensional $N = 6$ supersymmetric Chern–Simons theory and their string theory duals*, preprint (2013) arXiv 1302.2164
- [Kawasaki 1979] **T Kawasaki**, *The Riemann–Roch theorem for complex V -manifolds*, Osaka Math. J. 16:1 (1979) 151–159 MR
- [Knudsen and Mumford 1976] **F F Knudsen, D Mumford**, *The projectivity of the moduli space of stable curves: I: Preliminaries on “ $\det$ ” and “ Div ”*, Math. Scand. 39:1 (1976) 19–55 MR
- [Koroteev et al. 2021] **P Koroteev, P P Pushkar, A V Smirnov, A M Zeitlin**, *Quantum K -theory of quiver varieties and many-body systems*, Selecta Math. 27:5 (2021) art. id. 87 MR
- [Lee 2004] **Y-P Lee**, *Quantum K -theory, I: Foundations*, Duke Math. J. 121:3 (2004) 389–424 MR
- [Marian and Oprea 2007a] **A Marian, D Oprea**, *Counts of maps to Grassmannians and intersections on the moduli space of bundles*, J. Differential Geom. 76:1 (2007) 155–175 MR
- [Marian and Oprea 2007b] **A Marian, D Oprea**, *The level-rank duality for non-abelian theta functions*, Invent. Math. 168:2 (2007) 225–247 MR
- [Marian and Oprea 2010] **A Marian, D Oprea**, *GL Verlinde numbers and the Grassmann TQFT*, Port. Math. 67:2 (2010) 181–210 MR
- [Okounkov 2017] **A Okounkov**, *Lectures on K -theoretic computations in enumerative geometry*, from “Geometry of moduli spaces and representation theory” (R Bezrukavnikov, A Braverman, Z Yun, editors), IAS/Park City Math. Ser. 24, Amer. Math. Soc., Providence, RI (2017) 251–380 MR
- [Okounkov and Smirnov 2022] **A Okounkov, A Smirnov**, *Quantum difference equation for Nakajima varieties*, Invent. Math. 229:3 (2022) 1203–1299 MR
- [Pushkar et al. 2020] **P P Pushkar, A V Smirnov, A M Zeitlin**, *Baxter Q -operator from quantum K -theory*, Adv. Math. 360 (2020) art. id. 106919 MR
- [Qu 2018] **F Qu**, *Virtual pullbacks in K -theory*, Ann. Inst. Fourier (Grenoble) 68:4 (2018) 1609–1641 MR
- [Ruan and Zhang 2023] **Y Ruan, M Zhang**, *Verlinde/Grassmannian correspondence and rank 2 δ -wall-crossing*, Peking Math. J. 6:1 (2023) 217–306 MR
- [Su et al. 2020] **C Su, G Zhao, C Zhong**, *On the K -theory stable bases of the Springer resolution*, Ann. Sci. Éc. Norm. Supér. 53:3 (2020) 663–711 MR
- [Tonita 2014a] **V Tonita**, *Twisted orbifold Gromov–Witten invariants*, Nagoya Math. J. 213 (2014) 141–187 MR
- [Tonita 2014b] **V Tonita**, *A virtual Kawasaki–Riemann–Roch formula*, Pacific J. Math. 268:1 (2014) 249–255 MR
- [Tonita 2018] **V Tonita**, *Twisted K -theoretic Gromov–Witten invariants*, Math. Ann. 372:1-2 (2018) 489–526 MR
- [Tonita and Tseng 2013] **V Tonita, H-H Tseng**, *Quantum orbifold Hirzebruch–Riemann–Roch theorem in genus zero*, preprint (2013) arXiv 1307.0262
- [Tseng 2010] **H-H Tseng**, *Orbifold quantum Riemann–Roch, Lefschetz and Serre*, Geom. Topol. 14:1 (2010) 1–81 MR
- [Vafa 1992] **C Vafa**, *Topological mirrors and quantum rings*, from “Essays on mirror manifolds” (S-T Yau, editor), Int. Press, Hong Kong (1992) 96–119 MR
- [Witten 1995] **E Witten**, *The Verlinde algebra and the cohomology of the Grassmannian*, from “Geometry, topology, & physics” (S-T Yau, editor), Conf. Proc. Lecture Notes Geom. Topology IV, Int. Press, Cambridge, MA (1995) 357–422 MR

YONGBIN RUAN ruanyb@zju.edu.cn

Department of Mathematics, University of Michigan, Ann Arbor, MI, United States

Current address: Institute of Advanced Study for Mathematics, Zhejiang University, Hangzhou, China

MING ZHANG zhangmsq@umich.edu

Department of Mathematics, University of Michigan, Ann Arbor, MI, United States

YAOXIONG WEN y.x.wen.math@gmail.com

School of Mathematics, Korea Institute for Advanced Study, Seoul, South Korea

Proposed: András I Stipsicz

Received: May 17, 2018

Seconded: Mark Gross, Arend Bayer

Revised: February 11, 2025

Springer theory for symplectic Galois groups

KEVIN MCGERTY AND THOMAS NEVINS

A classical and beautiful story in geometric representation theory is the construction by Springer of an action of the Weyl group on the cohomology of the fibres of the Springer resolution of the nilpotent cone. We establish a natural extension of Springer’s theory to arbitrary symplectic resolutions of conical symplectic singularities. We analyse features of the action in the case of affine quiver varieties, constructing Weyl group actions on the cohomology of ADE quiver varieties, and also consider “symplectically dual” examples arising from slices in the affine Grassmannian. Along the way, we document some basic features of the symplectic geometry of quiver varieties.

1. Introduction	883
2. Symplectic resolutions	886
3. Basics of symplectic Springer theory	892
4. Symplectic geometry of quiver varieties	902
5. Quiver varieties of type ADE	907
6. Bionic symplectic varieties and Springer theory	912
7. Slices in the affine Grassmannian	916
8. Tensor products of representations	923
Acknowledgements	925
References	925

1 Introduction

A complex *symplectic variety* is a (quasiprojective) variety X with at worst rational, Gorenstein singularities and with a symplectic form ω , ie a closed, nondegenerate algebraic 2-form, on its smooth locus. A resolution of singularities $\pi : Y \rightarrow X$ is a *symplectic resolution* if ω extends to a symplectic form on all of Y .

A celebrated example is the cotangent bundle to a flag variety of a reductive group, which, as was first observed by a number of people including Springer and Steinberg, is a symplectic resolution of the nilpotent cone of the corresponding Lie algebra. In recent years, much effort has been focused on discovering to what extent one can generalise known representation-theoretic phenomena from Springer resolutions to all symplectic resolutions.

It is a beautiful observation of Markman [2010] in the projective case — shortly thereafter extended to the conical affine case by Namikawa [2010] — that if $\pi : Y \rightarrow X$ is a symplectic resolution, then the Poisson deformation base of Y is a Galois covering of the corresponding deformation base of X . The corresponding “symplectic Galois group”, in the case of the cotangent bundle of the flag variety, is

MSC2010: 14B07, 32S30.

© 2026 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

precisely the Weyl group, and in general it is the product of reflection groups attached to the geometry of slices to the resolution in codimension two.

In the present paper, we show that the cohomologies of the fibres of a symplectic resolution $\pi : Y \rightarrow X$ carry representations of its symplectic Galois group, generalising the seminal work of Springer [1976]. We also show that these representations can be constructed in a number of ways, mirroring the various constructions in the classical case: in particular, one can use both small resolutions and a version of the nearby cycles construction over a higher-dimensional base. Along the way, we establish some fundamental but not well-documented features of the symplectic geometry of Nakajima quiver varieties, and compute their symplectic Galois groups in the finite-type case.

A more precise description of our results is the following. Suppose that X is a conical symplectic variety (Definition 2.7) with symplectic resolution $Y \xrightarrow{\pi} X$. Let $\mathcal{Y} \rightarrow B_Y$ and $\mathcal{X} \rightarrow B_X$ be the corresponding versal Poisson deformations, with morphism $\mathcal{Y} \xrightarrow{\tilde{\pi}} \mathcal{X}$, as constructed in [Markman 2010; Namikawa 2010] (and reviewed in Section 2.2), with symplectic Galois group W . Fixing a coefficient field k , we define constructible complexes

$$\mathrm{HC} := \tilde{\pi}_! k_{\mathcal{Y}}[\dim(\mathcal{Y})] \quad \text{and} \quad \mathrm{Spr} := \pi_! k_Y[\dim(Y)],$$

the *symplectic Harish-Chandra sheaf* and *symplectic Springer sheaf*. Let $\mathrm{Perv}_{\mathrm{symp}}(X)$ denote the category of perverse sheaves on X (with coefficients in k) which are smooth along the stratification by symplectic leaves. We also define, in Section 3.2, a nearby cycles sheaf $\mathcal{P}$ for the family $\mathcal{X} \rightarrow B_X$ (again, notation as in Section 2.2). This sheaf carries a natural action of the fundamental group of the complement of the discriminant locus, which we call the *braid group* of the symplectic resolution.

Theorem 1.1 (Theorem 3.2, Corollary 3.12, Theorem 3.21, Proposition 3.24) (1) *The complexes HC and Spr are perverse sheaves, which are semisimple if $\mathrm{char}(k) = 0$. Moreover HC is the intersection cohomology extension of the local system on $\mathcal{X}^{\mathrm{reg}}$ given by $k[W]$, the regular representation of the symplectic Galois group W .*

(2) *We have $i_X^* \mathrm{HC} \cong \mathrm{Spr}$, where $i_X : X \hookrightarrow \mathcal{X}$ is the inclusion.*

(3) *We have a natural algebra isomorphism*

$$k[W] \xrightarrow{a} \mathrm{End}(\mathrm{HC}).$$

(4) *We get an adjoint pair of functors*

$$(- \otimes_{k[W]} \mathrm{Spr}) : k[W]\text{-mod} \rightleftarrows \mathrm{Perv}_{\mathrm{symp}}(X) : \mathrm{Hom}(\mathrm{Spr}, -).$$

(5) *The group W acts faithfully on the sheaf Spr, yielding, for each $x \in X$, an action of W on $H^*(\pi^{-1}(x))$.*

(6) *We have $\mathcal{P} \cong \mathrm{Spr}$ and the action of the braid group factors through the symplectic Galois group W . Moreover, this action coincides with the action of W induced by the intermediate extension functor.*

Assuming that k is coprime to $|W|$, the W action on Spr decomposes it into W -isotypic components: $\text{Spr} \cong \bigoplus_{\rho \in \widehat{W}} \mathcal{E}_\rho \otimes \rho$, for certain perverse sheaves $\mathcal{E}_\rho$. In the classical case, the perverse sheaves $\mathcal{E}_\rho$ are irreducible, and hence one obtains¹ the “Springer correspondence” between irreducible W -representations and certain local systems on nilpotent orbits. In Section 3.3 we discuss some examples which show that in our more general setting the situation is more subtle: the sheaves $\mathcal{E}_\rho$ need not be irreducible and indeed may be zero or have nonisomorphic irreducible constituents. It would be interesting to better understand the structure of the sheaves $\mathcal{E}_\rho$.

Turning to Nakajima’s quiver varieties, we review several features of their geometry, most likely known to experts, but not all of which appear yet to be documented in the literature. Fixing a quiver Q as in Section 4.1, let $Y(\mathbf{v}, \mathbf{w})$, respectively $X(\mathbf{v}, \mathbf{w})$, be the smooth symplectic quiver variety and affine quiver variety associated to the dimension vectors $\mathbf{v}$ and $\mathbf{w}$.

Theorem 1.2 [Bellamy and Schedler 2021, Propositions 4.6 and 5.2] (1) $X(\mathbf{v}, \mathbf{w})$ is a conical affine symplectic variety.

(2) $Y(\mathbf{v}, \mathbf{w})$ is a symplectic resolution of some conical affine symplectic variety.

(3) If Q is a Dynkin quiver of finite type, then for each $\mathbf{v}$ and $\mathbf{w}$ there exists $\mathbf{v}'$ so that $X(\mathbf{v}', \mathbf{w}) \cong X(\mathbf{v}, \mathbf{w})$ and the natural projective morphism $\pi_{\mathbf{v}'} : Y(\mathbf{v}', \mathbf{w}) \rightarrow X(\mathbf{v}', \mathbf{w})$ is a symplectic resolution.

We caution the reader that in general $Y(\mathbf{v}, \mathbf{w})$ need *not* be a resolution of $X(\mathbf{v}, \mathbf{w})$; cf Example 5.3. Bellamy and Schedler [2021], as well as establishing part (1) of the above theorem, also determine precisely when the varieties $X(\mathbf{v}, \mathbf{w})$ have symplectic resolutions.

Furthermore, we calculate the symplectic Galois group of quiver varieties for finite-type Dynkin quivers. More precisely, associated to the finite-type Dynkin quiver Q and dimension vectors $\mathbf{v}$ and $\mathbf{w}$ is a parabolic subroot system (denoted by Φ_μ in Section 5) with associated Weyl group W_μ . We show:

Theorem 1.3 (Theorem 5.8) *If Q is a Dynkin quiver of finite type, the symplectic Galois group of the quiver variety $X(\mathbf{v}, \mathbf{w})$ is W_μ .*

Combining Theorem 1.3 with our general theory, we obtain Weyl group actions on the cohomology of ADE quiver varieties. Weyl group actions have previously been obtained by Lusztig [2000], Nakajima [2003] and Maffei [2002].

Outside of finite type, as we show in Example 5.12, the naive generalisation of Theorem 1.3 fails. It is not clear to the authors what the correct general statement should be. It is plausible that one needs to take into account the existence of actions by larger algebras than the Kac–Moody algebras on topological invariants of quiver varieties. Bellamy and Schedler [2021, Section 7] establish some important results in the general setting, but they work in the unframed setting, so a direct comparison with our results is not immediate. See Remark 5.13 for more details.

We should note that much has already been written about categorical actions of braid groups on (derived categories of) quiver varieties and their quantisations. For example, if $\mathcal{D}$ is the discriminant

¹Since a simple perverse sheaf is the intermediate extension of a local system.

locus in $H^2(Y, \mathbb{C})$ (see Theorem 2.12) and $\mathcal{D}^c$ denotes its complement, then it is shown in [Braden et al. 2016] that there is an action of $\pi_1(\mathcal{D}^c/W)$ on $D^b(\mathbb{O}_g)$, the bounded derived category of the geometric category $\mathbb{O}$ defined in [Braden et al. 2016]. By taking Hochschild cohomology in an appropriate sense, this action should “deategorify” to an action of the braid group on the cohomology of the symplectic resolution, and compatibility with supports should ensure one also obtains actions on the cohomology of the fibres of the resolution also. These actions should then descend to the Weyl group actions constructed here. We hope to explore this relation in more detail in future work. However, even if one views our actions, at a lower categorical level, as a poor relation of categorical actions, our approach does work uniformly, independent of coefficient field. It is also quite concrete and topological; one important feature of Springer theory is the interplay amongst various approaches.

Convention 1.4 Throughout the paper, all varieties are over the field $\mathbb{C}$ of complex numbers, unless otherwise stated. We will also need the analytic topology, and in that case we will, by abuse of notation, identify a variety X with its analytification X^{an} . We write k to denote a coefficient field for sheaves, cohomology, etc.

2 Symplectic resolutions

2.1 Symplectic varieties and Poisson structures

Symplectic varieties were introduced by Beauville [2000].

A *symplectic variety* X is a normal variety which admits an (algebraic) symplectic form ω on its smooth locus X^{reg} such that, for some (or equivalently, every) resolution of singularities $\pi : Y \rightarrow X$, the pullback of ω extends to a regular (though not necessarily nondegenerate) 2-form on all of Y . In the case where the pullback of ω does extend to a symplectic form on all of Y , we say that π is a *symplectic resolution*. It is not necessarily the case that a symplectic variety admits a symplectic resolution (see the work of Fu [2003] on nilpotent orbit closures for examples).

We restrict our attention to affine symplectic varieties and their symplectic resolutions. These have been studied by a number of people. We will rely heavily on Kaledin’s lovely paper [2006], in which he provides a basic structure theory for symplectic varieties which we now review. We begin by noting that, since X is normal, the Poisson bracket on functions on X^{reg} given by ω extends to a Poisson bracket on all of X . A subvariety of a Poisson variety X is said to be Poisson if it is the vanishing locus of a (radical) Poisson ideal sheaf.

Definition 2.1 If X is a Poisson variety, then, considering it as an analytic space, we may define the *symplectic leaves* of X to be the equivalence classes of the equivalence relation generated by Hamiltonian flows on X . Thus $x, y \in X$ are equivalent if there is a finite sequence of Hamiltonian flows

$$\mathcal{H}_t^i : [0, a_i] \rightarrow \text{Symp}(X)$$

for $1 \leq i \leq n$ such that if $x_0 = x$ and $x_i = \mathcal{H}_{a_i}^i(x_{i-1})$ then $x_n = y$. Note that, in general, symplectic leaves need not be algebraic.

Definition 2.2 We say that a Poisson variety Z is *generically nondegenerate* if the bivector giving the Poisson bracket yields a generically nondegenerate alternating two-form. A Poisson variety is said to be *holonomic* if the restriction of the Poisson bracket to any Poisson subvariety is generically nondegenerate.

It is shown in [Kaledin 2006] that any symplectic variety is a holonomic Poisson variety. Note however that the converse is false (already in dimension 2, for example for simple elliptic singularities).

Recall that the *singularity stratification* of a reduced scheme X of finite type is the stratification whose closed subsets $X = X_0 \supset X_1 \supset X_2 \supset \cdots \supset X_k$ have the property that $X_k^\circ := X_k \setminus X_{k+1}$ equals X_k^{reg} , the smooth locus of X_k . This stratification is canonical and thus is preserved by all automorphisms of X and vector fields on X . The strata of this stratification however need not be connected. We will normally refine it so as to obtain connected strata. By abuse of notation we refer to such a refinement as the singularity stratification. We note that these strata are still clearly preserved by all vector fields on X .

Theorem 2.3 [Kaledin 2006] *Let X be a complex symplectic variety. The locally closed strata of the singularity stratification of X are locally closed smooth symplectic subvarieties. The closures X_i of the X_i° are irreducible Poisson subvarieties and every Poisson subvariety of X is the closure of a stratum. Moreover, the normalisation of any Poisson subvariety $Y \subset X$ is also a symplectic variety.*

Remark 2.4 The singularity stratification is, in general, too coarse, and many desirably properties of a stratification will only hold once it is refined. One virtue of symplectic varieties, however, is that the singularity stratification already possesses many of these properties. The most basic of these is for the strata to satisfy the “axiom of closure”, that is, the closure of a stratum of the partition is a union of strata. The following lemma, which is implicit in the proof of Proposition 3.1 of [Kaledin 2006], shows this is indeed the case for the singularity stratification of a symplectic variety.

Lemma 2.5 *The strata of the singularity stratification of a holonomic Poisson variety X (hence, in particular, a symplectic variety) are given by its symplectic leaves and thus they satisfy the axiom of closure.*

Proof We use induction on $\dim(X)$. The smooth locus of X^{reg} is a smooth holonomic variety, and hence, by Lemma 1.4 of [Kaledin 2006] is in fact smooth symplectic, and thus is a single symplectic leaf. The singular locus X^{sing} is a proper closed subscheme preserved by all automorphisms of X and all vector fields. In particular if $C^1, \dots, C^k$ are the irreducible components of X^{sing} , then each C^i is preserved by all Hamiltonian vector fields, and is therefore a Poisson subscheme that is necessarily holonomic. By induction it follows that the singularity stratification of each $\overline{C}^i$ is given by its symplectic leaves, and since by definition the symplectic leaves of $\overline{C}^i$ are precisely the symplectic leaves of X contained in $\overline{C}^i$, the stratifications of the components agree on their overlaps, which completes the inductive step. $\square$

Remark 2.6 We thank the referee for simplifying our original proof of the previous lemma.

2.2 Deformation theory and the symplectic Galois group

While standard deformation theory for nonprojective or singular varieties is generally rather intractable, it turns out that one often obtains a more tractable theory of *Poisson deformations* for algebraic varieties

equipped with a Poisson bracket. This deformation theory has been studied by Namikawa [2011] and Ginzburg and Kaledin [2004]; here we will largely follow Namikawa. Given a Poisson variety X there is a Poisson deformation functor PD_X from the category of local Artinian rings with residue field $\mathbb{C}$ to sets. This assigns to an Artinian ring R the set of pairs $(\mathcal{X}, \phi)$ consisting of a flat morphisms² $\mathcal{X} \rightarrow T = \mathrm{Spec}(R)$ equipped with a relative Poisson bracket $\{.,.\}$ and an isomorphism $\phi : \mathcal{X} \rightarrow \mathcal{X} \times_T \mathbb{C}$, taken up to the natural notion of isomorphism. We write $\mathbb{C}[\epsilon]$ for the ring of dual numbers over $\mathbb{C}$.

Definition 2.7 A *conical symplectic variety* is an affine symplectic variety X equipped with a $\mathbb{G}_m$ -action with nonnegative weights on $\mathbb{C}[X]$ and $\mathbb{C}[X]^{\mathbb{G}_m} = \mathbb{C}$, and for which the symplectic form transforms with weight $\ell > 0$. When X has a symplectic resolution $\pi : Y \rightarrow X$, the $\mathbb{G}_m$ -action lifts uniquely to Y .

Now suppose that X is an affine symplectic variety and $\pi : Y \rightarrow X$ is a symplectic resolution, we thus have Poisson deformation functors for both X and Y . If $p : \mathcal{Y} \rightarrow T$ is a deformation of Y , then it can be checked that $\mathrm{Spec}(p_*(\mathcal{O}_{\mathcal{Y}})) \rightarrow T$ is a deformation of X , so that one obtains a natural transformation $\pi_* : \mathrm{PD}_Y \rightarrow \mathrm{PD}_X$. One then has the following theorem of Namikawa:

Theorem 2.8 [Lehn et al. 2012, Theorem 2.1] *In the above setting both PD_X and PD_Y are unobstructed and are pro-representable. More precisely, one has a $\mathbb{G}_m$ -equivariant commutative diagram:*

$$(2-1) \quad \begin{array}{ccc} \mathcal{Y} & \xrightarrow{\theta_Y} & B_Y \\ \downarrow \tilde{\pi} & & \downarrow q \\ \mathcal{X} & \xrightarrow{\theta_X} & B_X \end{array}$$

where $B_Y = \mathrm{PD}_Y(\mathbb{C}[\epsilon]) \cong \mathbb{A}^d$ and $B_X = \mathrm{PD}_X(\mathbb{C}[\epsilon]) \cong \mathbb{A}^d$ with $d = \dim(H^2(Y, \mathbb{C}))$. The deformations $\mathcal{Y} \rightarrow B_Y$ and $\mathcal{X} \rightarrow B_X$ are formally universal Poisson deformations of Y and X respectively at 0. The map q is a finite Galois with $q(0) = 0$.

The action of $\mathbb{G}_m$ on the affine spaces $\mathbb{A}^d$ is of weight ℓ by construction. There is a natural period map $B_Y \rightarrow H^2(Y, \mathbb{C})$ for the deformation $\mathcal{Y}$ defined, for example, in step 2 of the proof of Theorem 1.1 in [Namikawa 2010]; see also [Kaledin and Verbitsky 2002] and [Ginzburg and Kaledin 2004]. Roughly, this map assigns to each fibre of the deformation the class of its symplectic form (for this to make sense one must show that the H^2 of the fibres can be canonically identified, which follows from the topological triviality of the deformation — see Proposition 2.9 below). The period map identifies B_Y with $H^2(Y, \mathbb{C})$. It follows from this that the $\mathbb{G}_m$ -action in the above theorem has weight ℓ on $B_Y \cong H^2(Y, \mathbb{C})$. The symplectic Galois group W arises as the covering group for the map q at 0.

Note that the affinisation $\mathcal{Y}^{\mathrm{aff}}$ is given by $\mathcal{Y}^{\mathrm{aff}} = \mathcal{X} \times_{B_X} B_Y$. It is well known to experts that the family $\mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ is topologically trivial. For example, it is noted by Namikawa [2010] that the arguments of Slodowy [1980, Section 4.2] for slices in the nilpotent cone generalise, as we discuss in some more detail below. The resulting triviality of cohomology is alternatively a consequence of Theorem 1.13 of [Ginzburg and Kaledin 2004] or Theorem 10.2.2 of [Dobrovolska et al. 2016].

²In this section we must work with schemes and not just varieties.

Proposition 2.9 *The morphism $\mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ is trivial as a topological fibration.*

Proof We may choose a contracting linear $\mathbb{G}_m$ -action on $\mathbb{A}^N$ for some positive integer N , and a $\mathbb{G}_m$ -equivariant embedding of the affinisation $\mathcal{Y}^{\text{aff}} = \mathcal{X} \times_{B_X} B_Y$ in $\mathbb{A}^N \times B_Y$ for some N , compatibly with projection to B_Y . Pull the function $\bar{\phi}(z_1, \dots, z_N) = \sum_i |z_i|^2$ on $\mathbb{A}^N$ back to a function $\phi : \mathcal{Y} \rightarrow \mathbb{R}$ that is proper on each fibre of $\mathcal{Y} \rightarrow B_Y$. This is an exhaustion function, in the sense that it is bounded below (by zero) and proper: for any $a > 0$ the preimage $\phi^{-1}([0, a])$ has compact fibres over the compact closed hyperdisk $\{(z_1, \dots, z_N) : \sum |z_i|^2 \leq a\}$.

It follows that $\phi : \mathcal{Y} \rightarrow \mathbb{R}$ is a proper map and hence the subset $\mathcal{Y}_{\leq r} := \phi^{-1}([0, r]) \subset \mathcal{Y}$ is thus compact for each $r > 0$. When r is sufficiently small, the boundary of $\mathcal{Y}_{\leq r}$ is a smooth manifold. This follows from the transversality of the affinisation map to $\phi^{-1}(r)$. To see this, note that because the $\mathbb{G}_m$ -action has positive weights, near 0 the boundary of the disk $\{z \in \mathbb{A}^N : \phi(z) \leq r\}$ is transverse to $\mathbb{R}^\times$ -orbits, and the affinisation map $\mathcal{Y} \rightarrow \mathcal{Y}^{\text{aff}}$ is equivariant. This argument also shows that $\mathcal{Y}_{\leq r}$ intersects $Y = \theta_Y^{-1}(0)$, where $0 \in B_Y$, transversely. It follows by continuity that the boundary of $\mathcal{Y}_{\leq r}$ intersects the fibres $\mathcal{Y}_\lambda = \theta_Y^{-1}(\lambda)$, where $\lambda \in B_Y$, transversely for small λ . It follows that, near $0 \in B_Y$, the map $\mathcal{Y}_{\leq r} \rightarrow B_Y$ is a proper submersion and hence by the Ehresmann fibration lemma (for manifolds with boundary) it is a locally trivial fibration over U some open ball around $0 \in B_Y$. Explicitly, we have shown that for sufficiently small $r > 0$ there is an open ball U about 0 such that θ_Y restricted to $\mathcal{Y}_{\leq r} \cap \theta_Y^{-1}(U)$ is a locally trivial fibration.

To extend this to the entire family $\theta_Y : \mathcal{Y} \rightarrow B_Y$, one must again use the $\mathbb{G}_m$ -action: Fix some r as above so that if $K = \mathcal{Y}_{\leq r} \cap \theta_Y^{-1}(U)$ then $\theta_{Y|K}$ is a locally trivial fibration. Then using the action of $t \in [1, \infty) \in \mathbb{R}^\times$ we see that the boundary of tK intersects the fibres of the family $\mathcal{Y} \rightarrow B_Y$ transversely. Using a Hermitian metric compatible with the projection to B_Y one can then project the radial vector field given by the infinitesimal $\mathbb{R}^\times$ -action to obtain a nonvanishing vector field along the fibres, which is complete by the properness of ϕ . Standard arguments in Morse theory (cf Theorem 3.1 of [Milnor 1963]) now show that the rescaled downward flow of this vector field defines a diffeomorphism $\mathcal{Y}_{\geq r} \cap \phi_Y^{-1}(U) \cong (\mathcal{Y}_r \cap \phi_Y^{-1}(U)) \times [r, \infty)$. It follows that $\phi_Y^{-1}(U) \rightarrow U$ is a topologically trivial fibration. (In fact with appropriate arguments one can make it C^∞ trivial.) One can then again use the $\mathbb{R}^\times$ action to expand from U to all of B_Y and conclude that θ_Y is a locally trivial (and hence trivial as B_Y is contractible) fibration. $\square$

2.3 The universal conic deformation

The deformations $\mathcal{Y} \rightarrow B_Y$ and $\mathcal{X} \rightarrow B_X$, in addition to being formally Poisson universal at the origin 0, are also universal once one takes the $\mathbb{G}_m$ -action into account:³ Given a symplectic variety Y with a $\mathbb{G}_m$ -action $\rho : \mathbb{G}_m \rightarrow \text{Aut}(Y)$ which scales the symplectic form ω_Y with weight ℓ , we will need to consider flat $\mathbb{G}_m$ -equivariant families over a linear base $\mathfrak{a}$, that is, families $f : \mathcal{Y}_\mathfrak{a} \rightarrow \mathfrak{a}$ of algebraic symplectic varieties over a vector space $\mathfrak{a}$ with central fibre Y , where $\mathbb{G}_m$ acts on the vector space $\mathfrak{a}$ by the ℓ -th power of the scaling action and the action on $\mathcal{Y}_\mathfrak{a}$ induces the action ρ on the central fibre Y , with moreover

³Everything here is presumably well known to the experts, but we failed to find a suitable reference, so for the convenience of the reader we outline the details here.

$t.\tilde{\omega} = t^\ell \tilde{\omega}$, where $\tilde{\omega}$ is the relative symplectic form of the family $\mathcal{Y}_a$ over a . We shall call such a family a *weight- ℓ conic deformation*. The following result is a rephrasing of the work of Namikawa [2008]:

Proposition 2.10 *The family $\mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ is universal among weight- ℓ conic deformations, that is, any such family $\mathcal{Y}_a \rightarrow a$ is obtained by pullback from a linear map $\alpha : a \rightarrow H^2(Y, \mathbb{C})$. Moreover, the family $\mathcal{X} \rightarrow B_X = H^2(Y, \mathbb{C})/W$ is also universal among weight- ℓ conic deformations of X , that is, any such family $\mathcal{X}_b \rightarrow b$ is obtained by pullback along a $\mathbb{G}_m$ -equivariant morphism $\beta : b \rightarrow B_X$.*

Proof We sketch a proof. As already noted above in Theorem 2.8, the deformation $\mathcal{Y} \rightarrow B_Y$ is formally Poisson-universal at the origin 0. The space B_Y is obtained from the pro-representing object by taking $\mathbb{G}_m$ -finite vectors, and one can check that the period map p is compatible with the $\mathbb{G}_m$ -actions (with weight ℓ on $H^2(Y, \mathbb{C})$).

Now if we have any weight- ℓ conic deformation $\mathcal{Y}_a \rightarrow a$, then by universality its formal completion at 0 is pulled back from the universal deformation $(\hat{\mathcal{Y}}, \hat{\omega})$ by a map $\tilde{f}$, say. The $\mathbb{G}_m$ -actions then allow us to spread this out to a map $f : a \rightarrow H^2(Y, \mathbb{C})$, which by homogeneity (everything in sight has weight ℓ for the $\mathbb{G}_m$ -action) must be linear. The claim for $\mathcal{X} \rightarrow B_X$ is established in an entirely similar manner. $\square$

2.4 The symplectic Galois group

It is shown in [Namikawa 2010] that q is a Galois covering.

Definition 2.11 We call the covering group W of q the *Galois group* of our symplectic resolution, or simply the *symplectic Galois group*.

Namikawa's work shows that W can be described quite explicitly in terms of the exceptional divisors of the resolution π : Let Σ denote the singular locus of X and Σ_0 denote the union of the strata of X of codimension 4 or higher, and set $U = X \setminus \Sigma_0$. Let $\mathcal{B}$ denote the set of connected components of $\Sigma \setminus \Sigma_0$, the singular locus of U . For each B in $\mathcal{B}$, a transverse slice S_B through a point $b \in B$ yields a two-dimensional symplectic singularity — that is, a rational surface singularity, whose minimal resolution is given by the restriction of π . Let W'_B denote the Weyl group of type ADE attached to this singularity. As the point $b \in B$ varies, the type of this singularity is locally constant, and thus one obtains an action of $\pi_1(B)$ on the components of the exceptional divisor of $\pi^{-1}(S_B)$. As these components can be identified with a set of simple roots for the simply laced root system attached to W_B , this monodromy action thus induces a diagram automorphism of W'_B and we set $W_B \subset W'_B$ to be its fixed-point subgroup. The main theorem of [Namikawa 2010] then shows that the symplectic Galois group is

$$W = \prod_{B \in \mathcal{B}} W_B,$$

and so $B_X = H^2(Y, \mathbb{C})/W$. Moreover, the space $H^2(Y, \mathbb{C})$ can be identified with the direct sum of the reflection representations of the W_B .

Note that in diagram (2-1) each square is commutative, but the right-hand square is *not* Cartesian, since for example the fibre product $\mathcal{X}' := \mathcal{X} \times_{B_X} B_Y$ is affine. Following Namikawa [2015], we let $\Pi : \mathcal{Y} \rightarrow \mathcal{X}'$

be the natural map and, for $t \in H^2(Y, \mathbb{C})$, let $\Pi_t : \mathcal{Y}_t \rightarrow \mathcal{X}'_t$ be the restriction of Π to the preimage $\mathcal{Y}_t = \psi_Y^{-1}(t)$.

We now recall a result of Namikawa. Let $\mathcal{D} \subset H^2(Y, \mathbb{C})$ be the discriminant locus: that is, $\mathcal{D}$ is defined to be the closure of the locus where Π_t is not an isomorphism. It can be shown that this is also the locus where $\mathcal{Y}_t$ is not affine, or equivalently, where $\mathcal{X}_t$ is singular; see [Namikawa 2011] for details. For a $\mathbb{Q}$ -vector space H write $H_{\mathbb{C}} = \mathbb{C} \otimes_{\mathbb{Q}} H$.

Theorem 2.12 [Namikawa 2015] *There are finitely many linear subspaces $\{H_i\}_{i \in I}$ of $H^2(Y, \mathbb{Q})$ of codimension 1 such that $\mathcal{D} = \bigcup_{i \in I} (H_i)_{\mathbb{C}}$. Moreover, the open chambers in $H^2(Y, \mathbb{R})$ complementary to the arrangement of real hyperplanes $\{(H_i)_{\mathbb{R}}\}_{i \in I}$ are given by $\{w(\text{Amp}(\pi_k))\}$, where $w \in W$ and π_k are the projective crepant resolutions of X , and $\text{Amp}(\pi_k)$ denotes the ample cone of π_k .*

Thus, in particular, the discriminant locus of the map Π is an algebraic hypersurface: more precisely, a finite union of hyperplanes.

2.5 (Semi)smallness and Poisson deformations

Definition 2.13 A morphism of varieties $f : Y \rightarrow X$ is said to be *semismall* if the fibre product $Z = Y \times_X Y$ has dimension at most $d = \dim(Y)$ (and hence exactly d). If moreover any irreducible component of Z which has dimension d maps to a dense subset of X then f is said to be *small*.

Kaledin shows that any symplectic resolution $Y \rightarrow X$ is semismall. We now show that suitable families of resolutions form small maps.

Definition 2.14 Let $f : \mathcal{X} \rightarrow B$ be a family of affine symplectic varieties such that for Zariski-generic $b \in B$ the variety $X_b = f^{-1}(b)$ is smooth and $\dim(X_b)$ is independent of $b \in B$. Let $\pi : \mathcal{Y} \rightarrow \mathcal{X}$ be a family of fibrewise symplectic resolutions, that is, we have a commutative diagram

$$\begin{array}{ccc} \mathcal{Y} & \xrightarrow{\pi} & \mathcal{X} \\ & \searrow g & \swarrow f \\ & B & \end{array}$$

such that $\pi_b : Y_b \rightarrow X_b$ is a symplectic resolution, where $Y_b = (f \circ \pi)^{-1}(b)$. We will call f a *smoothing family*. Note that by assumption, for Zariski-generic $b \in B$, the resolution π_b is an isomorphism.

Proposition 2.15 *Suppose that $f : \mathcal{X} \rightarrow B$ is a smoothing family of d -dimensional affine symplectic varieties and $\pi : \mathcal{Y} \rightarrow \mathcal{X}$ a family of fibrewise symplectic resolution as above. Then π is a small morphism.*

Proof Let D denote the closure of the locus of $b \in B$ for which X_b is singular. By assumption, D is a proper subvariety of B . Let $\mathcal{X}_{\text{sing}} = f^{-1}(D)$, a proper subvariety of $\mathcal{X}$. To show that π is small it suffice to show that $\dim(\mathcal{Y} \times_{\mathcal{X}_{\text{sing}}} \mathcal{Y}) < \dim(\mathcal{Y})$.

To see this, fix $b \in B$. By Lemma 2.11 of [Kaledin 2006], since $\pi_b : Y_b \rightarrow X_b$ is a projective birational morphism from the smooth variety Y_b to the symplectic variety X_b , it is semismall. It follows that $\dim(Y_b \times_{X_b} Y_b) = \dim(Y_b) = d$. Considering the natural map $\mathcal{Y} \times_{\mathcal{X}} \mathcal{Y} \rightarrow B$ we see that

$$\dim(\mathcal{Y} \times_{\mathcal{X}_{\text{sing}}} \mathcal{Y}) \leq d + \dim(D) < \dim(\mathcal{Y}).$$

□

Corollary 2.16 Suppose that $\mathcal{Y}_{\mathfrak{a}} \rightarrow \mathfrak{a}$ is a conic deformation of a conic symplectic resolution $\pi : Y \rightarrow X$ with the generic fibre of $\mathcal{Y}_{\mathfrak{a}}$ affine. Then if $\mathcal{X}_{\mathfrak{a}}$ is the affinisation of $\mathcal{Y}_{\mathfrak{a}}$, the natural morphism $\mathcal{Y}_{\mathfrak{a}} \rightarrow \mathcal{X}_{\mathfrak{a}}$ is small.

Proof By Proposition 2.10, the family $\mathcal{Y}_{\mathfrak{a}}$ is a pullback of the universal family by a linear map $\alpha : \mathfrak{a} \rightarrow H^2(Y, \mathbb{C})$. Since the generic fibre of $\tilde{\pi}$ is affine, the image of α is not contained in $\mathcal{D}$ the discriminant locus of the universal deformation, hence the locus in $\mathfrak{a}$ for which the fibres of the map $\mathcal{X}_{\mathfrak{a}} \rightarrow \mathfrak{a}$ are singular is a proper closed subset, hence we may apply the previous proposition. $\square$

Remark 2.17 It might arguably be more natural to consider, in the context of a conic symplectic resolution $\pi : Y \rightarrow X$, the morphism $\mathcal{Y} \rightarrow \mathcal{X}$ between respective universal deformations (over the deformation base B_X rather than $H^2(Y, \mathbb{C})$). Since the map from $H^2(Y, \mathbb{C}) \rightarrow B_X$ is finite, however, it follows immediately from the above that this map is also small.

3 Basics of symplectic Springer theory

We lay out some general algebraic and topological features of symplectic Springer theory. We let k denote a coefficient field for constructible sheaves (and their cohomologies). Unless otherwise stated we will assume that $\text{char}(k) \nmid |W|$.

3.1 Harish-Chandra and Springer sheaves

We maintain the notation of diagram (2-1).

Definition 3.1 The constructible complex

$$\text{HC} := \tilde{\pi}_! k_{\mathcal{Y}}[\dim(\mathcal{Y})]$$

is the *symplectic Harish-Chandra sheaf* on $\mathcal{X}$. The constructible complex

$$\text{Spr} := \pi_! k_Y[\dim(Y)]$$

is the *symplectic Springer sheaf* on X . When we wish to emphasise the particular symplectic resolution relevant, we will write Spr_X , and similarly when we wish to emphasise which coefficient field is being used, we will write $\text{Spr}(k)$.

We collect some basic properties below. Let $\text{Perv}_{\mathcal{G}}(X)$ denote the category of perverse sheaves on X (with coefficients in k) which are smooth along the stratification by symplectic leaves.

Theorem 3.2 Let W denote the symplectic Galois group.

- (1) The complexes HC and Spr are perverse sheaves, which are semisimple when k has characteristic zero. Moreover, HC is the intersection cohomology complex on $\mathcal{X}$ extending the local system on $\mathcal{X}^{\text{reg}}$ given by the regular representation of W the symplectic Galois group.
- (2) We have $i_X^* \text{HC} \cong \text{Spr}$, where $i_X : X \hookrightarrow \mathcal{X}$.

(3) We have a natural algebra isomorphism

$$k[W] \xrightarrow{a} \text{End}(\text{HC}).$$

(4) We get an adjoint pair of functors

$$(- \otimes_{k[W]} \text{Spr}) : k[W]\text{-mod} \rightleftarrows \text{Perv}_{\text{symp}}(X) : \text{Hom}(\text{Spr}, -).$$

Proof Statement (1) is immediate from the semismallness of $\pi : Y \rightarrow X$ and the smallness of the map $\tilde{\pi}$ established in Proposition 2.15. Note that the semisimplicity statement about the local system will follow from the proof of (3) below. Statement (2) follows from proper base change.

For statement (3), let B_X^{reg} denote the open subset of the base B_X of the versal deformation over which $\mathcal{X}$ has smooth fibres. Write $\mathcal{X}^{\text{reg}}$ for $\mathcal{X} \times_{B_X} B_X^{\text{reg}}$ and let $j : \mathcal{X}^{\text{reg}} \hookrightarrow \mathcal{X}$ denote the inclusion. Note that when restricted to the regular locus $\tilde{\pi}^{-1}(\mathcal{X}^{\text{reg}}) \rightarrow \mathcal{X}^{\text{reg}}$ is an unramified Galois W -cover. Now, observe that smallness of $\tilde{\pi}$ implies that $\text{HC} = j_{!*} \tilde{\pi}_* k_{\tilde{\pi}^{-1}(\mathcal{X}^{\text{reg}})}[-\dim(\mathcal{X})]$. To complete the proof of (3) we use Lemma 3.3 below. This shows that

$$\text{End}(\text{HC}) \cong \text{End}(\tilde{\pi}_* k_{\tilde{\pi}^{-1}(\mathcal{X}^{\text{reg}})}) \cong k[W].$$

This defines the isomorphism a .

Finally, the existence of the adjunction claimed in (4) is a version of the standard tensor-Hom adjunction, provided we prove that the sheaf Spr lies in $\text{Perv}_{\mathcal{G}}(X)$, that is, it is smooth along the stratification of X by symplectic leaves. We will give a proof of this using the nearby cycles description of the Springer sheaf which we will give in Section 3.2, since the nearby cycles sheaf is constructible with respect to a Whitney stratification of the exceptional fibre that satisfies the a_f condition. $\square$

Lemma 3.3 For a dense open subset $U \xrightarrow{\iota} X$ and a local system $\mathcal{L}$ on U ,

$$\text{End}_X(\iota_{!*} \mathcal{L}) = \text{End}_U(\mathcal{L}).$$

Proof We have homomorphisms

$$\text{End}_U(\mathcal{L}) \rightarrow \text{End}_X(\iota_{!*} \mathcal{L}) \rightarrow \text{End}_U(\mathcal{L}),$$

the first given by $(-)\iota_*$ and the second by restriction to U . The composite is the identity; hence it suffices to show that the latter map is injective. If $\phi \in \text{End}_X(\iota_{!*} \mathcal{L})$ and $\phi|_U = 0$, it follows that $\phi(\iota_{!*} \mathcal{L})$ is supported on $X \setminus U$. But the intermediate extension $\iota_{!*} \mathcal{L}$ is the unique extension that has no sub- or quotient object supported on $X \setminus U$, so $\text{im}(\phi) = 0$, as required. $\square$

Remark 3.4 In the case of HC this is a direct consequence of the definitions — the pushforward is, by the smallness assumption, an intersection cohomology sheaf, and the local systems that arise will be semisimple provided $\text{char}(k)$ is coprime to $|W|$.

On the other hand, for Spr , although its perversity is an immediate consequence of semismallness, the semisimplicity of Spr is not automatic if one only assumes that $\text{char}(k) \nmid |W|$. Indeed, [de Cataldo and Migliorini 2002] shows that the decomposition theorem for semismall maps can be related to the nondegeneracy of certain intersection forms, and the symplectic Galois group is not sufficient to determine these forms in general. We will give an explicit example of this later in the paper.

Definition 3.5 If P is an irreducible constituent of Spr , then, since Spr is constructible with respect to the stratification of X by its symplectic leaves, $P \cong \text{IC}(\mathcal{L})$ for some local system $\mathcal{L}$ on a stratum S , and the pair $(S, \mathcal{L})$ uniquely determine P (it is the “middle-extension” of $\mathcal{L}$ to X). We call such a local system a *Springer local system*.

Remark 3.6 From the point of view of $\text{Perv}_{\mathcal{G}}(X)$, an immediate deficiency in the theory described above is that $k[W]\text{-mod}$ does not detect all of the category $\text{Perv}_{\mathcal{G}}(X)$. This is of course already true in the classical setting of Springer theory when X is the nilpotent cone of a reductive group. In that setting, however, Lusztig [1984] shows that one can understand all of the simple objects in $\text{Perv}_{\mathcal{G}}(X)$ by a Harish-Chandra-style classification in which each object is associated to an irreducible representation of a finite Coxeter group, which arises by studying the endomorphism algebra of complexes induced from what Lusztig calls *cuspidal local systems*. In this context, the simple perverse sheaves arising in the Springer correspondence are those induced from the trivial local system on a maximal torus. (This is a cuspidal local system only because a torus has no proper parabolic subgroups.) Indeed more turns out to be true—one in fact gets an orthogonal decomposition of corresponding derived category of $\text{Perv}_{\mathcal{G}}(X)$, where each block of the decomposition can be labelled by a cuspidal local system on a Levi subalgebra. The objects of $\text{Perv}_{\mathcal{G}}(X)$ one can understand using the Springer sheaf might thus be thought of as the “principal block” of $\text{Perv}_{\mathcal{G}}(X)$.

Perhaps the easiest examples of cuspidal local systems arise from the fact that, when the group is SL_n , the equivariant fundamental group of the regular nilpotent orbit is cyclic of order n . The nontrivial representations of the cyclic group then correspond to local systems on the regular orbit which are cuspidal. In any case where the open symplectic leaf of our symplectic variety X has a nontrivial étale fundamental group, the nontrivial local systems on X^{reg} cannot be Springer local systems, ensuring $\text{Perv}_{\mathcal{G}}(X)$ has more simple objects than the “principal block”.

It is not immediately obvious how to even frame a generalisation to the setting of symplectic resolutions of Lusztig’s generalisation of Springer’s results. One reason for this is of course that it is difficult to define an appropriate analogue of a parabolic subgroup. At the very least it seems one must require X to possess additional torus symmetries which lift to the symplectic resolution.

Remark 3.7 As we will see in examples later in this section, the stronger version of assertion (3), which readers familiar with the classical theory might have expected to see, namely that the composition of the map a with the natural map $\text{End}(\text{HC}) \rightarrow \text{End}(\text{Spr})$ is an isomorphism, is not true in general. Neither, then, is it true that the induction functor from $k[W]\text{-modules}$ to $\text{Perv}_{\text{symp}}(X)$ is faithful. Indeed we will see later in this section examples which show that the map from $k[W]$ to $\text{End}(\text{Spr})$ need neither be injective nor surjective. The following lemmas will, however, imply a weaker statement, namely that W embeds into the automorphisms of Spr .

Next we note that the deformation $\tilde{\pi} : \mathcal{Y} \rightarrow \mathcal{X}$ gives us a more elementary way of constructing an action of W on the cohomology of Y . We first remark that the endomorphism algebra of the sheaf Spr has a natural geometric interpretation. Let $Z = X \times_Y X$ be the fibre product of X with itself over Y (the “symplectic

Steinberg variety”). Its Borel–Moore homology has a natural convolution structure preserving the middle-dimensional piece $H_{\text{BM}}^{\text{mid}}(Z)$; the latter has a natural basis given by the irreducible components of Z .

Lemma 3.8 $\text{End}(\text{Spr}) \cong H_{\text{BM}}^{\text{mid}}(Z).$

Proof The argument is standard: see, for example, Section 3.4 of [Chriss and Ginzburg 1997]. $\square$

In terms of the identification $\text{End}(\text{Spr}) \cong H_{\text{BM}}^{\text{mid}}(Z)$, the map from the symplectic Galois group W to $H_{\text{BM}}^{\text{mid}}(Z)$ can be seen as a specialisation-of-cycles map, similar to the construction of Springer representations in [Chriss and Ginzburg 1997].

Since we are assuming that our symplectic resolution $\pi : Y \rightarrow X$ is conic, the fixed point locus $F = Y^{\mathbb{G}_m}$ is a smooth projective variety contained in $L = \pi^{-1}(0)$ (where as usual we write 0 for the cone point of X). Moreover, since this $\mathbb{G}_m$ action contracts Y to F , it follows by an argument of Slodowy [1980, Section IV] that L is a homotopy retract of Y , and so we have isomorphisms in cohomology $H^*(Y) \cong H^*(L)$. (Note however that in general, L is not smooth and $F \subsetneq L$.)

Definition 3.9 At a regular point $b \in B_Y$, that is, at a point outside the discriminant locus, the map $\tilde{\pi}$ is W -sheeted local isomorphism and we get an isomorphism between each fibre of θ_Y over the orbit $W \cdot b \subset B_Y$ and $\theta_X^{-1}(\pi(b))$. This induces an action of W on the cohomology of the fibres (for example induced by the action of W on the complex $(\theta_Y)_!(k_{\mathcal{O}_Y})$). But the deformation $\mathcal{O}_Y$ is topologically trivial (by Proposition 2.9), so the cohomologies of the fibres of θ_Y are canonically identified, and hence we get an action on the cohomology of Y . We will denote this action by $\rho : W \rightarrow \text{GL}(H^*(Y))$. Note that the action on the fibres at regular points is precisely the one inducing the morphism a as in Theorem 3.2.

Lemma 3.10 The restriction of ρ to $H^2(Y)$ agrees with the action of W on $B_Y = H^2(Y)$.

Proof This follows from the construction of the period map, which is given by taking the homology class of the symplectic form of the fibres of ψ_Y . See for example Step 2 in the proof of Theorem 1.1 in [Namikawa 2010] for details. $\square$

Lemma 3.11 The action of W on $H^*(Y)$ via ρ coincides with the action of W on $\pi_!(k_Y)_0 = H^*(L)$ via the isomorphism $H^*(L) \cong H^*(Y)$ described in Definition 3.9.

Proof We have the diagram

$$\begin{array}{ccccc} \mathcal{O}_Y & \xrightarrow{\theta_Y} & B_Y & & \\ \tilde{\pi} \downarrow & & \downarrow \pi & & \\ \mathcal{X} & \xrightarrow{\theta_X} & B_X & \xrightarrow{=} & B_Y/W \end{array}$$

As already noted, the $\mathbb{G}_m$ -action on X extends to the deformations $\mathcal{X}$ and $\mathcal{O}_Y$, and hence the sheaf $\tilde{\pi}_!(k_{\mathcal{O}_Y})$ is equivariant with respect to the $\mathbb{G}_m$ -action (since the constant sheaf $k_{\mathcal{O}_Y}$ clearly is, and the morphism $\tilde{\pi}$ is equivariant). It follows that the natural morphism $H^*(\mathcal{X}, \tilde{\pi}_!(k_{\mathcal{O}_Y})) \rightarrow \tilde{\pi}_!(k_{\mathcal{O}_Y})_0$ is an isomorphism; see for example [Kaledin 2006, Lemma 5.6]. Now $H^*(\mathcal{O}_Y, k) = H^*(\mathcal{X}, \tilde{\pi}_*(k_{\mathcal{O}_Y})) = H^*(\mathcal{X}, \tilde{\pi}_!(k_{\mathcal{O}_Y}))$. But the

map θ_Y is a topologically trivial fibration, so the restriction-to-fibres maps $H^*(\mathcal{Y}, k) \rightarrow H^*(\mathcal{Y}_t, k)$ are all isomorphisms. It follows readily that the composition

$$H^*(Y, k) = H^*(\mathcal{Y}_0, k) \cong H^*(\mathcal{Y}, k) \rightarrow \tilde{\pi}_!(k)_0 \cong \pi_!(k)_0$$

is W -equivariant, as required. $\square$

Corollary 3.12 *The symplectic Galois group W acts faithfully on the sheaf Spr . Moreover this action yields, for any $x \in X$, a natural action of W on the cohomology $H^*(\tilde{\pi}^{-1}(x))$.*

Proof The second assertion follows immediately from the action of W on Spr induced by the restriction of the action on HC , since the cohomology in question is just the stalk of Spr at $x \in X$. For the first part, note that it is enough to check that the induced action on some stalk cohomologies of Spr is faithful; but this follows for the stalk at $0 \in X$ by the preceding lemmas. Indeed, since $\pi^{-1}(0) = L$, these show that action of W on $H^2(\pi^{-1}(0))$ is just that of W on B_Y , which is faithful by construction. $\square$

3.2 Stratifications and Thom's a_f condition for Poisson deformations

Let $X \subset U \subset \mathbb{C}^n$ be an analytic subset of U an open set in $\mathbb{C}^n$, and suppose $f : X \rightarrow \mathbb{C}$ is an analytic map, and $x \in f^{-1}(0)$. For $\delta \ll \varepsilon \ll 1$ the restriction of f to $f^{-1}(\dot{D}_\delta \setminus \{0\}) \cap B_\varepsilon(x)$ is a topological fibration. Here $\dot{D}_\delta$ is the open disc of radius δ centred at 0, and $B_\varepsilon(x)$ denotes the closed ball of radius ε around x (with respect to some Hermitian metric say). This was first proved by Lê Dũng Tráng [1977], generalising results of Milnor, who proved the result when X is a hypersurface in $\mathbb{C}^n$. The fibre of the resulting fibration is therefore known as the *Milnor fibre*. If we take $f : X \rightarrow Y$ where $\dim(Y)$ is greater than 1, the corresponding statement is false.

Example 3.13 Let $f : \mathbb{C}^3 \rightarrow \mathbb{C}^2$ be given by $f(x, y, z) = (x^2 - y^2z, y)$. It can be checked that f is flat, however it is not a topological fibration in any ball around $0 \in \mathbb{C}^3$: indeed it is easy to see that some fibres are connected and some are not. (See the introduction to [Sabbah 1983] for a detailed discussion.)

However, one can impose a condition, going back to R Thom, on the map f which ensures that an analogue of Lê's result holds. In order to describe this we will need some basic material from stratification theory, as discussed for example in [Mather 2012]. First note that it follows from Theorem 2.3 that any symplectic variety is holonomic in the sense of Definition 2.2, and has finitely many symplectic leaves.

Lemma 3.14 *Let X be a symplectic variety. Then the canonical singularity stratification of X consists of the symplectic leaves of X and moreover it satisfies the Whitney conditions.*

Proof We have already seen in Theorem 2.3 that the (connected components of) the singularity stratification are the symplectic leaves of X , so it remains to check the Whitney conditions. To see this we use the basic fact which establishes the existence of Whitney stratifications: given any analytic variety X with a partition $X = \bigsqcup_{i \in I} X_i^0$, and strata X_i^0, X_j^0 satisfying $X_j^0 \subset X_i = \overline{X_i^0}$, the set of points $x \in X_j^0$ which do not satisfy the Whitney a and b condition is an analytic set of dimension strictly smaller than X_j^0 . We call such points irregular.

But now if $x \in X_j^0 \subset X_i$ is a point of irregularity for a pair of strata X_i^0 and X_j^0 , then its orbit under any Hamiltonian flow (which preserves both X_i^0 and X_j^0) will consist entirely of points of irregularity. It follows that the set of points of irregularity is a union of symplectic leaves in X . But then if the irregular locus is nonempty it must be all of X_j^0 , contradicting the fact that it has smaller dimension than $\dim(X_j^0)$. $\square$

Remark 3.15 The iterative procedure of taking (a component of) the smooth locus, removing it, and repeating until only smooth components remain yields, in the case of a holonomic variety X , the symplectic leaves of that variety, and the above shows that these form a Whitney stratification of X . The same procedure can of course be applied to an arbitrary singular variety, but in general it will only yield a partition of the variety into locally closed pieces, which need not satisfy the Whitney conditions.

Definition 3.16 Let M be a connected smooth complex analytic variety of dimension n and suppose that $f : M \rightarrow V$ is an analytic map, where $V \subset \mathbb{C}^r$ is a neighbourhood of $0 \in \mathbb{C}^r$. Let V^{reg} be the set of regular values of f and V^{sing} the set of nonregular values. We assume that U^{sing} is a proper analytic subvariety of V . Consider $M^0 = f^{-1}(U^{\text{reg}})$, a smooth submanifold of M . Suppose that $E = f^{-1}(0)$ is equipped with a Whitney stratification $\mathcal{S}$. For $x \in M$, let $\text{Fib}_x = f^{-1}(f(x))$ be the fibre of the map f which contains x . We will say that $(f, M, \mathcal{S})$ satisfies Thom's A_f condition if for each stratum $S \in \mathcal{S}$, whenever a sequence of points $(x_i)_{i=1}^\infty$ in M^0 converges to $y \in S$, is such that the limit

$$T = \lim_{i \rightarrow \infty} T_{x_i} \text{Fib}_{x_i} \subset T_y M$$

exists, it contains the tangent space to S at y , that is, we have $T \supseteq T_y S$.

When (f, M, S) satisfies the A_f condition, then Grinberg [1998] (following an earlier observation of Lê Dũng Tráng [1977]) showed that it makes sense to talk about the nearby cycles sheaf of f , and the higher-dimensional base results in a richer monodromy action.

The construction is as follows: Let D be a neighbourhood of zero in $\mathbb{C}$, and let $\gamma : D \rightarrow U$ be an embedded analytic arc, such that $\gamma(0) = 0$ and $\gamma(t) \in M^0$ for $t \neq 0$. Form the fibre product $M_\gamma = M \times_\gamma D$, and let $f_\gamma : M_\gamma \rightarrow V$ be the obvious map. We may consider the nearby cycles sheaf $\psi_{f_\gamma}(k_{M_\gamma}[d])$ (where d shifts appropriately to ensure perversity). Then we have the following:

Proposition 3.17 [Grinberg 1998, Proposition 2.4] *Let $(f, M, \mathcal{S})$ be as in Definition 3.16 and let $\gamma : D \rightarrow U$ be an embedded analytic arc as above.*

- (i) *The sheaves $\psi_{f_\gamma}(k_{M_\gamma}[d])$ for different γ are all isomorphic. We may therefore omit the subscript γ , and call the sheaf $\mathcal{P}_f$ the nearby cycles of f . It is a perverse sheaf on E , constructible with respect to $\mathcal{S}$.*
- (ii) *The local fundamental group $\pi_1((U \cap B_\epsilon))$, where B_ϵ is a small ball around the origin, acts on $\mathcal{P}_f$ by monodromy. We denote this action by $\mu : \pi_1(U \cap B_\epsilon) \rightarrow \text{Aut}(\mathcal{P}_f)$.*
- (iii) *The sheaf $\mathcal{P}_f$ is Verdier self-dual.*

Proof The monodromy action of $\pi_1(U \cap B_\epsilon)$ is the assertion of the proposition which is not standard, and we include Grinberg's proof for the convenience. If $x \in E$ is a point in the exceptional fibre lying in a stratum S , take a slice N at x transverse to $\mathcal{S}$ and take $0 < \epsilon \ll \delta \ll 1$ (chosen in decreasing order). Then for λ a regular value with $\|\lambda\| < \epsilon$ the *Milnor fibre* of f is $M_{f,N,x} = f^{-1}(\lambda) \cap N \cap B_{x,\delta}$, where $B_{x,\delta}$ is the ball of radius $\delta > 0$ centred at x . The stalk cohomology of $P_{f,y}$ at x is $H^*(M_{f,N,x})$ and these can be shown to vary as a local system using the argument in Section 1 of [Lê Dũng Tráng 1977]. $\square$

Example 3.18 In the case of $T^*\mathcal{B}$, the resolution of the nilpotent cone, the symplectic deformation space is the Grothendieck resolution $\tilde{\mathfrak{g}}$ equipped with its canonical map to the universal Cartan $\mathfrak{h}$, which can be identified with $H^2(T^*\mathcal{B}, \mathbb{C})$ via Borel's description of $H^*(\mathcal{B})$. The symplectic Galois group W is the (usual) Weyl group of the algebraic group G , and the universal Poisson deformation of $\mathcal{N}$, the nilpotent cone, is $\mathfrak{g}$ the Lie algebra of the group G , equipped with the adjoint quotient map from $\mathfrak{g}$ to $\mathfrak{g} // G \cong \mathfrak{h} / W$.

Let $\mathcal{A}$ denote the union of the root hyperplanes in $\mathfrak{h}$. Then $U^{\text{sing}} = \mathcal{A} / W \subset \mathfrak{h} / W = B_X$ (and hence U^{reg} is its complement). The orbit stratification of $\mathcal{N}$ satisfies the A_f condition for the adjoint quotient map.

We will now show that this theory can be applied in our symplectic setting. We keep the notation of diagram (2-1) and of Theorem 2.12.

Definition 3.19 Let $B_X^{\text{reg}} = (H^2(Y, \mathbb{C}) \setminus \mathcal{D}) / W$, and let $U = \theta_X^{-1}(B^{\text{reg}})$.

Lemma 3.20 *The symplectic leaves of X satisfy Thom's A_{θ_X} condition with respect to the regular locus U of θ_X .*

Proof That the set U is the regular locus of the morphism θ_X follows from [Namikawa 2015]; see Theorem 2.12.

Let $x \in S$ be a point in a symplectic leaf $S \subseteq X$, and suppose $(y_n)_{n \geq 1}$ is a sequence in U which converges to $x \in S$ with the tangent spaces $T_{y_i} \text{Fib}_{y_i}$ converging in a suitable Grassmannian to T . We must show that $T_x S \subseteq T$.

To show this, it suffices to produce, for $\eta \in T_x S$, a sequence of tangent vectors in $T_{y_i} \text{Fib}_{y_i}$ which converge to η . Since S is an integral Poisson subscheme, we may write the derivation given by η via the Poisson structure as $g \mapsto \{f_\eta, g\}(x)$ for some function f_η . Now picking any extension of this function to a function $\tilde{f}_\eta$ on the variety $\mathcal{X}$ we see that if (y_n) is a sequence of points in U tending to x , since the fibre of $\tilde{\pi}$ at each y_n is smooth symplectic, the derivation $\eta_n = \{\tilde{f}_\eta, \cdot\}(y_n)$ is tangent to Fib_{y_n} and clearly

$$\lim_{n \rightarrow \infty} \eta_n = \eta,$$

so that the tangent space $T_x S$ lies in the limit of the tangent spaces of the fibres of $\tilde{\pi}$ as required. $\square$

Combining the previous lemma with Proposition 3.17 we obtain a perverse sheaf $\mathcal{P} = \mathcal{P}_{\theta_X}$ on X equipped with a monodromy action $\mu : \pi_1(B^{\text{reg}} \cap \mathbb{B}_\epsilon)$, where $\mathbb{B}_\epsilon$ is a small ball centred at $0 \in B_X$. By the conic nature of our setting, restricting to this ball does not change the fundamental group.

Theorem 3.21 We have $\mathcal{P} \cong \pi_!(k_Y)$, and the action of the braid group $\pi_1(U)$ factors through the symplectic Galois group W .

Proof Let us first see that $\mathcal{P}$ is isomorphic to $\pi_!(k_Y)$: pick a line $K \subseteq H^2(Y, \mathbb{C})$ which intersects $\mathcal{D}$ only at $0 \in H^2(Y, \mathbb{C})$ (this is possible since $\mathcal{D}$ is a finite union of hyperplanes) and let $C = q(K)$ be the W -orbit of K , that is, its image under $q : H^2(Y, \mathbb{C}) \rightarrow H^2(Y, \mathbb{C})/W$ the quotient map. We may construct $\mathcal{P}$ by forming the (one-dimensional) nearby cycles for the restriction $\theta_C : \mathcal{X}_C \rightarrow C$ of the map θ_X to $\theta_X^{-1}(C)$. Let $\theta_K : \mathcal{Y}_K \rightarrow K$ denote the restriction of θ_Y to K , so that we have a commutative diagram

$$\begin{array}{ccc} \mathcal{Y}_K & \xrightarrow{\tilde{\pi}_K} & \mathcal{X}_C \\ \theta_K \downarrow & & \downarrow \theta_C \\ K & \xrightarrow{q|_K} & C \end{array}$$

where $\tilde{\pi}_K$ denotes the restriction of $\tilde{\pi}$ to $\theta_Y^{-1}(K)$. Since the map $\tilde{\pi}_K$ is proper and the nearby cycles construction commutes with pushforward by $\tilde{\pi}$, we have

$$\psi_{\theta_C}(\tilde{\pi}_{K!}(k_{\mathcal{Y}_K})) \cong \tilde{\pi}_{K!}(\psi_{\theta_K}(k_{\mathcal{Y}_K})).$$

Now $\psi_{\theta_K}(k_{\mathcal{Y}_K}) \cong k_Y$, since, by Proposition 2.9, $\tilde{\pi} : \mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ is a topologically trivial fibre bundle. On the other hand, since $\tilde{\pi}$ is a local isomorphism over U^{reg} (see Definition 3.19) it follows that the nearby cycles sheaves $\psi_{\theta_C}(\tilde{\pi}_{K!}(k_{\mathcal{Y}_K}))$ and $\psi_{\theta_C}(k_{\mathcal{X}_C})$ are isomorphic, and hence $\mathcal{P} \cong \pi_!(k_Y) = \text{Spr}$. Finally, the topological triviality of the fibre bundle θ_Y shows that the action of the fundamental group on $\mathcal{P}$ must factor through W as required. $\square$

Corollary 3.22 The Springer sheaf Spr is constructible with respect to the stratification of X by its symplectic leaves. Moreover, the action of the monodromy action of the symplectic Galois group on the cohomology of the fibres of π respects the ring structure on cohomology.

Proof This follows immediately from the isomorphism $\text{Spr} \cong \mathcal{P}$ and the general fact that the nearby cycles sheaf is constructible with respect to a Whitney stratification satisfying the A_f -condition; see for example [Goresky and MacPherson 1988]. For the final sentence, the compatibility with the ring structure on cohomology comes from the fact that the nearby cycles construction shows that monodromy action actually acts on the homotopy type of the fibres of π . This is essentially contained in Slodowy's work, and explained in detail in the case of classical Springer theory in [Treumann 2009]. $\square$

Remark 3.23 As the referee pointed out to us, it should be possible to deduce this using the Hamiltonian flows on a leaf.

Finally, we show that the monodromy action of the symplectic Galois group given by the nearby cycles construction is the same as that given by taking the intermediate extension.

Proposition 3.24 The monodromy action of W on $\pi_!(k_Y)$ induced by the isomorphism

$$\mathcal{P} \cong \pi_!(k_Y[\dim(Y)])$$

coincides with the action of W induced by the intermediate extension.

Proof By Lemmas 3.10 and 3.11 we see that the action of W on $H^*(\pi^{-1}(0)) = H^*(L)$ is faithful, where $L = \pi^{-1}(0)$. It thus suffices to compare the two actions on $H^*(L)$. But by Lemma 3.11 again, the intermediate extension action is given by the “naive action” of Definition 3.9 and it is easy to see that Grinberg’s monodromy action also coincides with this action on $\pi^{-1}(0)$. $\square$

3.3 Examples and a question

We continue with a discussion of some examples. For simplicity, we take $k = \mathbb{C}$.

Example 3.25 The simplest cases of symplectic singularities are the Kleinian singularities. Let Γ be a finite subgroup of $\mathrm{SL}_2(\mathbb{C})$, let $S = \mathbb{C}^2/\Gamma$ be the corresponding quotient singularity, and let $\pi : Y \rightarrow S$ be the minimal resolution. Now $\mathrm{Spr} = \pi_!(\mathbb{C}_Y[2])$ is a semisimple perverse sheaf. The only sheaves which can occur are the constant sheaf (which is perverse in this case because S is rationally smooth) and the skyscraper sheaf supported at the cone point. Using proper base change to compute the cohomology of the fibres, this shows that $\mathrm{Spr} \cong \mathbb{C}_S[2] \oplus \mathbb{C}_0^{\oplus r}$, where r is the number of components in the exceptional fibre. The action of the Weyl group is then via the trivial representation on $\mathbb{C}_S[2]$ and via the reflection representation on the summand supported on singular point. The action on $\mathbb{C}_0^{\oplus r}$ is just the action of W on $H^2(\widetilde{\mathbb{C}^2/\Gamma})$, and by Corollary 3.22 the W respects the intersection form, hence we see this is the reflection representation of W .

This example shows that in general one cannot expect that the map from $\mathbb{C}[W]$ to $\mathrm{End}(\mathrm{Spr})$ is injective (even though the map $W \rightarrow \mathrm{Aut}(\mathrm{Spr})$ is injective). Indeed as soon as W has more than two conjugacy classes (which is the case unless $W = Z_2$) the map from $\mathbb{C}[W]$ to $\mathrm{End}(\mathrm{Spr})$ will fail to be injective. (We thank the referee for suggesting we set out how stark the failure of injectivity is here.)

Example 3.26 Let $Y = T^*\mathbb{P}^n$ be the cotangent bundle of n -dimensional projective space. This can be described explicitly (viewing $\mathbb{P}^n$ as the space of lines in $\mathbb{C}^{n+1}$) as

$$Y = \{(\alpha, \ell) \in \mathfrak{gl}_{n+1} \times \mathbb{P}^n : \mathrm{im}(\alpha) \subseteq \ell \subseteq \ker(\alpha) \subseteq \mathbb{C}^{n+1}\}.$$

The natural map π given by the restriction of the projection to the first factor is the moment map for the natural GL_{n+1} -action on Y , and it realises Y as a symplectic resolution of its image $X = \overline{\mathbb{O}_{\min}}$ the closure of the minimal nonzero nilpotent orbit in $\mathfrak{gl}_{n+1}$. It is well-known and easy to see that $\dim(X) = 2n$, and moreover, if $\alpha \neq 0$ then $\ell = \mathrm{im}(\alpha)$, so that π is an isomorphism over $\mathbb{O}_{\min}$. Thus for $n > 1$ we see that X is smooth in codimension 2, so that its symplectic Galois group is trivial. On the other hand, for any coefficient field k , we have the Springer sheaf $\mathrm{Spr}_k = \pi_!(k_Y[2n])$ and, as noted in Lemma 3.8, its endomorphism algebra $\mathrm{End}(\mathrm{Spr})$ is isomorphic to the middle-dimensional Borel–Moore homology of $Z_Y = Y \times_X Y$, and it is easy to see that this has two components, C_1 and C_2 say, where C_1 is the fibre of the map $Z_Y \rightarrow Y$ over $0 \in \mathfrak{gl}_{n+1}$ and C_2 is the closure of the locus $\{((\alpha, \ell), (\alpha, \ell)) \in Y \times Y : \alpha \neq 0\}$, so that $\mathrm{End}(\mathrm{Spr}_k)$ is 2-dimensional. However, a straightforward computation using the techniques of [Chriss and Ginzburg 1997] shows that, if we write $[C_i]$ for the class in Borel–Moore homology associated to the component C_i , then $[C_1] * [C_1] = (n+1)[C_2]$ (where $n+1 = \chi(\mathbb{P}^n)$) and so $\mathrm{Spr}(k)$ is semisimple if and only if p and $n+1$ are coprime.

This example also shows that the map $\mathbb{C}[W] \rightarrow \text{End}(\text{Spr})$ need not be surjective. (Note that the case $n = 1$ recovers the ordinary Springer sheaf for $\mathfrak{sl}_2$, where the Weyl group does of course control the semisimplicity of Spr .)

Example 3.27 Let $X = \text{Sym}^n(\mathbb{A}^2)$ be the symmetric product of the affine plane $\mathbb{A}^2$, equipped with the usual scaling action of $\mathbb{G}_m$. It is well-known that a resolution of singularities is given by the Hilbert–Chow morphism $\pi : \text{Hilb}^n(\mathbb{A}^2) \rightarrow \text{Sym}^n(\mathbb{A}^2)$, and viewing $\mathbb{A}^2 = T^*\mathbb{A}^1$, it becomes a symplectic resolution.

The symmetric product X is naturally stratified into locally closed pieces X_ν given by partitions ν of n (the partition measures the multiplicities of points) and these are evidently integral Poisson subvarieties. Moreover, as in [Göttsche and Soergel 1993] for example, we have $\dim(X_\nu) = 2n - 2\ell(\nu)$, where $\ell(\nu)$ is the length of the partition ν (ie the number of parts in the partition). It follows that the only codimension-2 stratum is that labelled by $(2, 1^{n-2})$. In order to calculate the symplectic Galois group it is thus enough to calculate a slice through this stratum (and check monodromies). However, in our case since we know it must be a Kleinian singularity, and moreover the description of the fibres of the Hilbert–Chow morphism show that the resolution of the slice has exceptional fibre consisting of a single $\mathbb{P}^1$ (the punctual Hilbert scheme of two points) the slice must be a type A_1 singularity and hence $W = S_2$, the symmetric group on two letters.

On the other hand, it is shown in [Göttsche and Soergel 1993] that the Hilbert–Chow morphism is semismall with every stratum being relevant, and moreover

$$\pi_!(\mathbb{C}_Y[\dim(Y)]) = \bigoplus_{\lambda} \text{IC}(X_\lambda),$$

where λ runs over the partitions of n . Thus the Springer sheaf Spr in this case is a direct sum of the simple perverse sheaves on X associated to the constant local systems on each stratum by the intermediate extension functor. Thus if $p(n)$ denotes the number of partitions of n , we see that for $n \geq 2$ there are thus $p(n)$ simple constituents in the Springer sheaf. On the other hand, the symplectic Galois group is S_2 for all n . It follows that $\text{End}(\text{Spr}_X) \cong \mathbb{C}^{p(n)}$, which is larger than $\mathbb{C}[S_2]$ provided $n \geq 2$, so that we obtain another example where the homomorphism from the group algebra to $\text{End}(\text{Spr}_X)$ fails to be surjective — indeed as $n \rightarrow \infty$ one can ensure that the failure is in some sense arbitrarily bad!

The last example also has another interesting consequence: Suppose for simplicity that we work over an algebraically closed field k for which the Springer sheaf is semisimple. Then it can be decomposed into a direct sum

$$(3-1) \quad \text{Spr} \cong \bigoplus_{i \in I} \text{IC}(S_i, \mathcal{L}_i) \otimes V_i,$$

where $i \in I$ runs over a finite set of local systems supported on symplectic leaves of X , and each V_i is a finite-dimensional vector space (which can be identified as usual with $\text{Hom}(\text{IC}(S_i, \mathcal{L}_i), \text{Spr})$. The V_i thus become representations of W . In the classical theory, this yields a map from the irreducible representations of W to the set of local systems supported on nilpotent orbits, because each irreducible W -representation occurs uniquely as some V_i .

Notice that Example 3.27 shows that in our general setting, an irreducible representation of W the symplectic Galois group may be associated to more than one local system: Indeed when $n > 2$ there are $p(n) \geq n$ many local systems (the constant local systems on the various strata) but only two irreducible representations of S_2 , the symplectic Galois group. Thus at least one of its two irreducible representations must be associated to more than one local system.

The following question was suggested to us by Gwyn Bellamy.

Question 3.28 When are the W -representations V_i in the isotypic decomposition (3-1) of Spr irreducible?

From the study of resolutions of slices in the affine Grassmannian in Section 7, we will obtain a large supply of examples where the irreducibility should be false in some sense “generically”. We are very grateful to the referee for pointing out the following example, which also shows that a satisfactory answer to Bellamy’s question will need to identify some natural conditions which can ensure the irreducibility of the representations V_i .

Example 3.29 Similarly to Example 3.27, consider $\widetilde{\text{Hilb}^n(\mathbb{C}^2/\Gamma)}$, where $\Gamma < \text{SL}_2(\mathbb{C})$ is a nontrivial finite subgroup and $n \geq 2$. The symplectic Galois group is $W_\Gamma \times S_2$ (there are precisely two codimension-2 strata once $n \geq 2$, with slices equal to the slices we saw in Examples 3.25 and 3.27, respectively; for more details see Example 5.12).

If c denotes the number of conjugacy classes in W , the middle-dimensional homology of $\widetilde{\text{Hilb}^n(\mathbb{C}^2/\Gamma)}$ has dimension

$$\dim(H^{\text{mid}}(\widetilde{\text{Hilb}^n(\mathbb{C}^2/\Gamma)})) = \binom{n+c-2}{c-2}.$$

Moreover, by the bounds on the stalk cohomologies of intersection cohomology sheaves, this cohomology is precisely V_0 , where 0 denotes the stratum of the cone point of $\text{Sym}^n(\mathbb{C}^2/\Gamma)$ (indeed we saw in Example 3.25, the case of the minimal resolution of a Kleinian singularity, that the middle cohomology of the smooth surface was attached to the skyscraper sheaf at the singularity). Clearly, provided $|W| > 2$ so that $c > 2$, this grows with n , and so must eventually become a reducible representation of W .

4 Symplectic geometry of quiver varieties

We investigate symplectic features of Nakajima’s quiver varieties. While much of this section is likely well understood by experts, we document the details where we could not find suitable references.

4.1 Basic set-up

We briefly recall the setting, using [Nakajima 1998] as our principal reference. Let $Q = (I, \Omega)$ be a quiver with vertex set I and edge set Ω , such that the edges $h \in \Omega$ are directed with starting vertex $s(h)$ and terminating vertex $t(h)$. The *adjacency matrix* of a quiver Q is the matrix $A = A_Q$ with entries $a_{ij} = |\{h \in \Omega : s(h) = i, t(h) = j\}|$. Let $\bar{Q}$ denote the doubled quiver, ie the quiver with vertex set I but edge set $H = \Omega \sqcup \bar{\Omega}$, where $\bar{\Omega}$ is in bijection with Ω via a map $h \mapsto \bar{h}$ so that $s(\bar{h}) = t(h)$, and $t(\bar{h}) = s(h)$.

For convenience we treat $(\overline{\cdot})$ as an involution on H by setting $(\overline{\overline{h}}) = h$. Note that $A_{\overline{Q}} = A_Q + A_Q^T$ is symmetric.

For a doubled quiver $\overline{Q}$ we define the associated (symmetric) *Cartan matrix* $C = C_{\overline{Q}} = (c_{ij})$ to be $2I - A_{\overline{Q}}$. Let $R = \mathbb{Z}[I] = \{f : I \rightarrow \mathbb{Z}\}$ and $P = \text{Hom}(R, \mathbb{Z})$ be its dual. For $i \in I$ we write $\check{\alpha}_i$ for the indicator function of $i \in I$, and Λ_i for the corresponding element of the basis of P dual to $\{\check{\alpha}_i : i \in I\}$. Now R carries a symmetric bilinear form $\langle \cdot, \cdot \rangle$ given by $\langle \check{\alpha}_i, \check{\alpha}_j \rangle = c_{ij}$, and using this form one defines “reflections” $s_i : R \rightarrow R$, where $s_i(x) = x - \langle x, \check{\alpha}_i \rangle \check{\alpha}_i$ for all $x \in R$. The group $W(\overline{Q})$ is then defined to be the subgroup of $\text{Aut}(R)$ generated by the $\{s_i : i \in I\}$. Note that $W(\overline{Q})$ acts on P in the usual way via $w(\lambda)(x) = \lambda(w^{-1}(x))$. The bilinear form $\langle \cdot, \cdot \rangle$ yields a map $\theta_R : R \rightarrow P$ given by $\theta_R(x)(y) = \langle x, y \rangle$, and we will write α_i for $\theta_R(\check{\alpha}_i)$. We will also write $\Lambda_{\mathbf{w}} = \sum_{i \in I} w_i \Lambda_i$ and $\alpha_{\mathbf{v}} = \sum_{i \in I} v_i \alpha_i$. If $v = \sum_{i \in I} v_i \alpha_i \in R$ we write $\text{supp}_R(v) = \{i \in I : v_i \neq 0\}$, and similarly for $\mu = \sum_{i \in I} \mu_i \Lambda_i \in P$ we write $\text{supp}_P(\mu) = \{i \in I : \mu_i \neq 0\}$.

A quiver variety is defined by a pair $\mathbf{v}, \mathbf{w} \in \mathbb{Z}_{\geq 0}^I$. Given $\mathbf{v}, \mathbf{w}$, choose I -graded vector spaces V and W with $\dim(V_i) = v_i$, $\dim(W_i) = w_i$ for each $i \in I$, and let

$$M(\mathbf{v}, \mathbf{w}) = \bigoplus_{h \in H} \text{Hom}(V_{s(h)}, V_{t(h)}) \oplus \bigoplus_{i \in I} \text{Hom}(W_i, V_i) \oplus \bigoplus_{i \in I} \text{Hom}(V_i, W_i).$$

We will write an element of this space as a triple (x, p, q) . Picking a function $\varepsilon : H \rightarrow \{\pm 1\}$ such that $\varepsilon(h) + \varepsilon(\overline{h}) = 0$, the formula

$$\omega((x, p, q), (x', p', q')) = \sum_{h \in H} \text{tr}(\varepsilon(h) x_h x'_h) + \sum_{i \in I} \text{tr}(p'_i q_i + p_i q'_i)$$

defines a symplectic form on $M(\mathbf{v}, \mathbf{w})$ for which the action of the group $G_{\mathbf{v}} = \prod_{i \in I} \text{GL}(V_i)$ is Hamiltonian, with moment map

$$\mu(x, p, q) = \left(\sum_{h: s(h)=i} \varepsilon(h) x_{\overline{h}} x_h + p q \right)_{i \in I} \in \mathfrak{g}_{\mathbf{v}}^* = \bigoplus_{i \in I} \mathfrak{gl}(V_i)$$

(where we identify $\mathfrak{gl}_V$ with its dual using the trace pairings). We may then consider two kinds of quotients of $M(\mathbf{v}, \mathbf{w})$ by the action of $G_{\mathbf{v}}$. The first is the affine quotient, which we will denote by $X(\mathbf{v}, \mathbf{w})$. That is,

$$X(\mathbf{v}, \mathbf{w}) = \text{Spec}(\mathbb{C}[\mu^{-1}(0)]^{G_{\mathbf{v}}}).$$

Since it is known since Hilbert that the invariant functions separate closed orbits in this setting, the points of $X(\mathbf{v}, \mathbf{w})$ thus correspond to the closed $G_{\mathbf{v}}$ -orbits in $\mu^{-1}(0)$. Notice that the Poisson bracket associated to the symplectic form ω is clearly preserved by the action of $G_{\mathbf{v}}$, so that it induces a Poisson structure on $\mathbb{C}[\mu^{-1}(0)]^{G_{\mathbf{v}}}$, and hence $X(\mathbf{v}, \mathbf{w})$ is a Poisson variety.

The second quotient is the GIT quotient determined by the linearisation of the trivial line bundle given by the determinant character χ (more generally, one could take any sufficiently general character of $G_{\mathbf{v}}$):

$$Y(\mathbf{v}, \mathbf{w}) = \text{Proj} \left(\bigoplus_{n \geq 0} \mathbb{C}[\mu^{-1}(0)]^{\chi^n} \right).$$

In [Nakajima 1998] it is shown that $Y(\mathbf{v}, \mathbf{w})$ coincides with the quotient of $\mu^{-1}(0)^s$ by the action of $G_{\mathbf{v}}$, where $\mu^{-1}(0)^s$ is the set of stable points of $\mu^{-1}(0)$, and $G_{\mathbf{v}}$ acts freely there, so the quotient is again an algebraic symplectic manifold. The description of $Y(\mathbf{v}, \mathbf{w})$ immediately shows that there is a projective morphism $\pi : Y(\mathbf{v}, \mathbf{w}) \rightarrow X(\mathbf{v}, \mathbf{w})$. It is easy to check also that the variety $Y(\mathbf{v}, \mathbf{w})$, if nonempty, has dimension $2\mathbf{v}^t \mathbf{w} - \mathbf{v}^t C \mathbf{v}$, where $C = C_{\bar{Q}}$ is the Cartan matrix of the doubled quiver.

Remark 4.1 Our notation, as is common in the subject, labels quiver varieties by the dimension vectors $\mathbf{v}$ and $\mathbf{w}$ despite the definition using a choice of vector spaces of those dimensions. Since different choices of vector spaces yield isomorphic varieties however, the imprecision is harmless.

4.2 Affine quiver varieties are symplectic

Let (Q, I, H) be a quiver and $\mathbf{v}, \mathbf{w}$ be dimension vectors. Let $Y = Y(\mathbf{v}, \mathbf{w})$, respectively $X = X(\mathbf{v}, \mathbf{w})$, be the smooth symplectic quiver variety and affine quiver variety associated to the dimension vectors $\mathbf{v}$ and $\mathbf{w}$ and let $\pi : Y \rightarrow X$ denote the natural projective morphism.

We first investigate the affine symplectic variety X . Note that since the morphism $\pi : Y \rightarrow X$ is not in general birational, the following is not obvious.

Theorem 4.2 [Bellamy and Schedler 2021, Theorem 1.2] *An affine quiver variety X is a symplectic variety.*

As mentioned in the introduction, Bellamy and Schedler [2021] also give explicit criteria for when an affine quiver variety admits a symplectic resolution.

It follows from the above theorem that every affine quiver variety is holonomic. In fact, one can give an explicit description of the symplectic leaves. Recall from [Nakajima 1998] that X is stratified by stabiliser-type (in fact it is an example of a Luna stratification): a point $z \in X$ corresponds to a closed orbit in $\mu^{-1}(0)$. Let $(x, p, q) \in \mu^{-1}(0)$ be a point lying in this orbit and let H denote its stabiliser in $G_{\mathbf{v}}$. Since the orbit of (x, p, q) is closed the subgroup H is a reductive subgroup of $G_{\mathbf{v}}$. If (H) denotes the conjugacy class of the subgroup H in $G_{\mathbf{v}}$ we say that (H) is the stabiliser type of $z \in X$. Conversely, given a conjugacy class (H) of reductive subgroups of $G_{\mathbf{v}}$ we let $X_{(H)}$ be the subset of all points of X of stabiliser type (H) . This yields a finite⁴ stratification $X = \bigsqcup_{(H)} X_{(H)}$.

Remark 4.3 This stratification is a special case of a stratification introduced by Luna [1973]. In the context of Hamiltonian reductions in the differentiable setting it was first studied by Sjamaar and Lerman [1991], who establish in that context its description in terms of symplectic leaves.

Lemma 4.4 *The variety X is a holonomic variety in the sense of Kaledin. Its symplectic leaves are precisely the strata of the orbit stratification in the sense of [Nakajima 1998].*

Proof As already noted, the variety X carries a natural Poisson variety structure. Martino [2008] showed that the strata of the orbit stratification of X are precisely the symplectic leaves of X , and hence X has only finitely many symplectic leaves. It follows from the discussion in Section 2.1 that X is holonomic. $\square$

⁴There are finitely many conjugacy classes of such reductive subgroups, as can be seen by root combinatorics.

Remark 4.5 As noted in [Bellamy and Schedler 2021], the argument of [Martino 2008] incorrectly invokes a result of Schwarz to show that the orbit-type strata are connected. However, this gap is readily filled using the connectedness of quiver varieties as is explained in [Bellamy and Schedler 2021, Corollary 3.25].

4.3 Smooth quiver varieties are symplectic resolutions

While it is not, in general, true that an affine quiver variety has a symplectic resolution, it is nonetheless the case that a smooth symplectic quiver variety $Y(\mathbf{v}, \mathbf{w})$ is always a symplectic resolution — although not necessarily of the associated affine quiver variety $X(\mathbf{v}, \mathbf{w})$.

Proposition 4.6 *The smooth quiver variety $Y(\mathbf{v}, \mathbf{w})$ is a symplectic resolution.*

Although this seems to be known to experts, it does not seem to be documented anywhere. We thank the referee for suggesting the current organisation of our proof, and in particular for pointing out the virtue of first establishing the following lemma, which appears to be folklore, but we could not find a reference.

Lemma 4.7 *Suppose that Y is a smooth symplectic quasiprojective variety. Then Y is a symplectic resolution if and only if it is birational to its affinisation.*

Proof If $\pi : Y \rightarrow X$ exhibits Y as a symplectic resolution of an affine symplectic variety X , it follows from Grauert–Reimenschneider that $\pi_*(\mathcal{O}_Y) = \mathcal{O}_X$; hence Y is the affinisation of X , as required.

For the converse, suppose that the morphism $a : Y \rightarrow Y^{\text{aff}}$ is birational. Let $V \subseteq Y^{\text{aff}}$ denote the locus over which a is not an isomorphism. Since the affinisation of a smooth variety is automatically normal, we may apply Zariski’s main theorem to deduce that the complement of V in N is codimension 2 in Y^{aff} . Now the Poisson bivector associated to the symplectic form of Y induces, through a the affinisation map, a Poisson structure on V . Since Y^{aff} is normal, and the complement of V is codimension 2 in Y^{aff} it follows that the Poisson structure extends to all of Y^{aff} , and hence a is a symplectic resolution. $\square$

With the lemma in hand, Proposition 4.6 will follow if we can show that $Y = Y(\mathbf{v}, \mathbf{w})$ is birational to its affinisation. In general, this need *not* be $X(\mathbf{v}, \mathbf{w})$, but this can be rectified using reflection functors, as we now briefly review:

We have defined quiver varieties, following [Nakajima 1998], using the notion of stability provided by determinant character, but any choice of character θ of $G_{\mathbf{v}}$ yields varieties $Y_{\theta}(\mathbf{v}, \mathbf{w})$ and $X_{\theta}(\mathbf{v}, \mathbf{w})$. Now the existence of polystable representations for θ will in general mean that $Y_{\theta}(\mathbf{v}, \mathbf{w})$ need not be smooth, but for all suitably generic θ will possess no such representations and thus for all such generic θ the variety $Y_{\theta}(\mathbf{v}, \mathbf{w})$ is smooth. Note that we may identify the abelian group of characters of $G_{\mathbf{v}}$ with $R = \mathbb{Z}[I]$. Explicitly, if θ is a character of $G_{\mathbf{v}}$ and $g = (g_i)_{i \in I} \in G_{\mathbf{v}}$ then $\theta(g) = \prod_{i \in I} \det(g_i)^{n_i}$ and $\theta \mapsto \sum_{i \in I} n_i \check{\alpha}_i$.

Reflection functors for quiver varieties have been studied by Lusztig [2000], Nakajima [2003] and Maffei [2002]. In the algebrogeometric (as opposed to hyper-Kähler) setting these provide, for each vertex $i \in I$ an isomorphism of varieties $s_{i,w} : Y_{\theta}(\mathbf{v}, \mathbf{w}) \rightarrow Y_{s_i(\theta)}(s_i * \mathbf{v}, \mathbf{w})$. Here $s_i(\theta)$ is defined by

viewing θ as an element of R , and $s_i * v$ is defined so that, if $\alpha_v = \sum_{i \in I} v_i \alpha_i$ and $\Lambda_w = \sum_{i \in I} w_i \Lambda_i$, then $\Lambda_w - \alpha_{s_i * v} = s_i(\Lambda_w - \alpha_v)$.

Proposition 4.8 *Let $Y = Y(v, w)$. The affinisation map $a : Y \rightarrow Y^{\text{aff}}$ is a birational morphism.*

Proof Using the isomorphisms induced by reflection functors, it is shown in Proposition 46 of [Maffei 2002] that any quiver variety $Y_\theta(v, w)$ is isomorphic to a quiver variety $Y_{\theta'}(v', w)$ with the property that $(\Lambda_w - \alpha_v)(\tilde{\alpha}_i) \geq 0$ is, in the terminology of that paper, *dominant*. Moreover, the proof of Lemma 54 of that paper shows that, in this case, $\mu^{-1}(0)^{\theta'}$ is open dense in $\mu^{-1}(0)$, hence $Y_{\theta'}(v', w)$. Since we work with generic θ (the action of $W(\bar{Q})$ preserves the locus of generic stability conditions) the action of $G_{v'}$ on $\mu^{-1}(0)^{\theta'}$ is free, and hence it follows that $\pi_{v'} : Y_{v'}(v', w)$ is birational. Since $\pi_{v'}$ is also projective, and thus its image is closed, it follows that $\pi_{v'}$ is a surjective birational morphism. Since $X(v', w)$ is affine, it is clear that the affinisation map a factors through $\pi_{v'}$, and since $\pi_{v'}$ is birational, it follows a fortiori that a is birational. $\square$

Remark 4.9 The preceding proof shows slightly more than is asserted in the statement: the smooth symplectic variety $Y(v, w)$ is a symplectic resolution of the an affine quiver variety. The authors thank the referee for alerting them to the results of [Maffei 2002], which simplified considerably our argument in the proof of the preceding proposition.

Finally, we note that a smooth quiver variety is always a *conical* symplectic resolution. Indeed, if we let $\mathbb{G}_m$ act on the vector space $M(v, w)$ with weight 1, we obtain $\mathbb{G}_m$ -actions on $X(v, w)$ and $Y(v, w)$ for which π is equivariant. The action has weight 2 on the symplectic form on $Y(v, w)$. The following is then clear:

Lemma 4.10 *The induced $\mathbb{G}_m$ -action on $Y(v, w)$ gives $Y(v, w)$ the structure of a conical symplectic resolution.*

Remark 4.11 The $\mathbb{G}_m$ -action need not be unique: in general, the automorphism group of a quiver variety may contain a higher-rank torus, and many of its one-parameter subgroups may yield conical structures on the quiver variety. Such group actions (or *flavour symmetries* of supersymmetric gauge theories) play a central role in the beautiful story of symplectic duality discovered by Braden, Licata, Proudfoot and Webster [Braden et al. 2016].

It should also be noted that Nakajima [1998] uses a different $\mathbb{G}_m$ -action, which depends on a choice of orientation Ω of the quiver. For quivers with no oriented loops it is a conical action. See [Nakajima 1998, Proposition 3.22] for more details.

4.4 Universal Hamiltonian reduction and the universal Poisson deformation

The moment map for the action of G_v on $M(v, w)$ takes values in $\mathfrak{g}_v^*$. We have

$$\mathfrak{z}_v = (\mathfrak{g}_v / [\mathfrak{g}_v, \mathfrak{g}_v])^* \subseteq \mathfrak{g}_v^*.$$

We may thus consider the flat family of reductions $\mu^{-1}(\mathfrak{z}_v) //_{\chi} G_v$ over $\mathfrak{z}_v$, and it follows from the universal property that this deformation has an associated period map $\kappa : \mathfrak{z}_v \rightarrow H^2(Y(v, w))$. It is known that the map κ is actually given by the Chern character map: for details see [Losev 2012].

Theorem 4.12 [McGerty and Nevins 2018, Theorem 1.1] *For any quiver variety the map*

$$H_{G_v}^*(\star) \cong H_{G_v}^*(E_V) \rightarrow H_{G_v}^*(\mu^{-1}(0)^{\text{ss}}) \rightarrow H^*(Y(v, w))$$

is surjective.

Corollary 4.13 *The period map κ is surjective.*

Thus, up to identifying the kernel of κ , we obtain a realisation of the universal graded Poisson deformation.

5 Quiver varieties of type ADE

In this section, we provide refined information on the quiver varieties associated to finite-type simply laced diagrams. Thus, from now on, we will assume that the underlying graph of our quiver Q is a simply laced Dynkin diagram associated to a Lie algebra $\mathfrak{g} = \mathfrak{g}_Q$. As we noted before, the results in this subsection are presumably well known to experts. We use the notation introduced in Section 4.1.

Let $X(v, w)^{\text{reg}}$ be the locus of points in $X(v, w)$ corresponding to closed G_v -orbits in $\mu^{-1}(0)$ with trivial stabiliser. (This is an open but possibly empty subset of $X(v, w)$.) Note that, if we fix w , then if v' has $v'_i \leq v_i$ we may choose V'_i as a direct summand of V_i , and hence extending by zero yields a natural map $X(v', w) \rightarrow X(v, w)$ which is injective.

Proposition 5.1 *For an affine quiver variety X attached to a simply laced finite-type Dynkin diagram, the orbit strata (and hence the symplectic leaves) are given by*

$$X(v, w) = \bigsqcup_{v'} X(v', w)^{\text{reg}},$$

where the v' that occur as nonempty strata are those for which both $v'_i \leq v_i$ for all $i \in I$ and such that $\Lambda_w - \alpha_{v'} \in P^+$.

Proof This is almost contained in [Nakajima 1998]. Indeed it is shown in Remark 3.28 of that paper that for quiver of type ADE, the symplectic leaves are all of the form $X(v', w)^{\text{reg}}$, and moreover Corollary 10.8 therein shows that $X(v, w)^{\text{reg}}$ is nonempty if and only if $\Lambda_w - \alpha_v \in P^+$, and $\Lambda_w - \alpha_v$ is a weight of the irreducible representation $L(\Lambda_w)$ of highest weight Λ_w . But it is known that the weights $\Pi(\Lambda_w)$ of the irreducible representation $L(\Lambda_w)$ are, in the sense of [Humphreys 1972, Section 13.4], a saturated set with highest weight Λ_w , and so, by Lemma B of [loc. cit.] a weight μ lies in $P^+ \cap \Pi(\Lambda_w)$ if and only if $\Lambda_w - \mu \in Q^+$. Applying this to $\mu = \Lambda_w - \alpha_v$ gives the result. $\square$

Finally we note that in the finite-type case, $X(v, w)$ always has a symplectic resolution.

Proposition 5.2 *Let $X(\mathbf{v}, \mathbf{w})$ be a finite-type affine quiver variety. Then there is a $\mathbf{v}' \in \mathbb{Z}^I$ with $\alpha_{\mathbf{v}'} \leq \alpha_{\mathbf{v}}$ such that $X(\mathbf{v}', \mathbf{w}) \cong X(\mathbf{v}, \mathbf{w})$ and $\pi_{\mathbf{v}'} : Y(\mathbf{v}', \mathbf{w}) \rightarrow X(\mathbf{v}', \mathbf{w})$ is a symplectic resolution.*

Proof By Proposition 5.1 there is a $\mathbf{v}'$ with $X(\mathbf{v}', \mathbf{w})^{\text{reg}}$ dense and open in $X(\mathbf{v}, \mathbf{w})$. It follows that the natural injective morphism $i : X(\mathbf{v}', \mathbf{w}) \rightarrow X(\mathbf{v}, \mathbf{w})$ is surjective, and hence bijective. Since by the main result of [Crawley-Boevey 2003] the variety $X(\mathbf{v}, \mathbf{w})$ is normal, it follows that i is an isomorphism.

By Lemma 4.7, we must show that $\pi_{\mathbf{v}'}$ is the affinisation morphism of $Y(\mathbf{v}', \mathbf{w})$. Now [Nakajima 1998, Corollary 10.11] shows that the projective morphism $\pi_{\mathbf{v}}$ is a semismall resolution of its image, and moreover all strata contained in the image are relevant. Since $\dim(X(\mathbf{v}', \mathbf{w})) = \dim(Y(\mathbf{v}', \mathbf{w}))$, it follows that $\pi_{\mathbf{v}'}$ is dominant, and so the image is all of $X(\mathbf{v}', \mathbf{w})$. Since $X(\mathbf{v}', \mathbf{w})$ is affine, the morphism $\pi_{\mathbf{v}'}$ factors through the affinisation morphism $a : Y(\mathbf{v}', \mathbf{w}) \rightarrow Y(\mathbf{v}', \mathbf{w})^{\text{aff}}$, and the induced map $\pi'_{\mathbf{v}'} : Y(\mathbf{v}', \mathbf{w})^{\text{aff}} \rightarrow X(\mathbf{v}', \mathbf{w})$ is birational. Moreover, since $X(\mathbf{v}', \mathbf{w})$ is normal, the fibres of $\pi'_{\mathbf{v}'}$ are connected. Since they must also be proper and $Y(\mathbf{v}', \mathbf{w})^{\text{aff}}$ is affine, they must be points. It follows that $\pi'_{\mathbf{v}'}$ is a bijective proper morphism, so again by the normality of $X(\mathbf{v}', \mathbf{w})$, it is an isomorphism, hence $\pi_{\mathbf{v}'}$ is the affinisation morphism for $Y(\mathbf{v}', \mathbf{w})$, as required. $\square$

The above highlights one elementary criterion for a projective morphism to be a resolution of singularities, namely that it must at least be surjective! Despite its trivial nature, this criterion is essentially the only way the natural morphisms $\pi : Y(\mathbf{v}, \mathbf{w}) \rightarrow X(\mathbf{v}, \mathbf{w})$ can fail to be a resolution, and examples where this happens are not difficult to find even in low dimensions:

Example 5.3 Take Q to be the quiver with one node, and let $v = w = 2$. Then $Y(2, 2)$ has dimension 0 (and so is a point) while $X(2, 2)$ is the rational singularity $\mathbb{C}^2/\{\pm 1\}$, and thus π is not surjective. An explicit calculation shows that the zero fibre $\mu^{-1}(0)$ of the moment map has three irreducible components, two of which are picked out by stability conditions. These components each contain a free $G_{\mathbf{v}}$ -orbit however, so that the invariant functions are constant on each, while on the third component there are no free orbits, and there are nonconstant $G_{\mathbf{v}}$ -invariant functions on it.

Remark 5.4 In fact, the argument of the previous proposition can be modified to show that provided the map $\pi : Y(\mathbf{v}, \mathbf{w}) \rightarrow X(\mathbf{v}, \mathbf{w})$ is surjective, it will be a symplectic resolution. Moreover, since $X(\mathbf{v}, \mathbf{w})$ is irreducible and π is proper, it suffices to check that $\dim(Y(\mathbf{v}, \mathbf{w})) = \dim(X(\mathbf{v}, \mathbf{w}))$. We thank the referee for this clarifying remark.

If, on the other hand, π fails to be surjective, then a second difficulty arises: It is not clear (at least to the authors) that, in general, the Poisson subvarieties of $X(\mathbf{v}, \mathbf{w})$ (one of which must be the image of π) must be normal. Indeed this is false for arbitrary symplectic resolutions, as it is well-known that nilpotent orbit closures outside of type A can fail to be normal.

5.1 Weyl groups and symplectic Galois groups of quiver varieties of type ADE

We now compute the symplectic Galois groups for type ADE quiver varieties. As before, we let $\mathfrak{g}_Q$ denote the semisimple Lie algebra attached to the simply laced Dynkin diagram underlying the quiver Q . Nakajima [1998] has shown how to construct the irreducible highest-weight representations of $\mathfrak{g}_Q$ in the

cohomology of quiver varieties, where if $\Lambda_{\mathbf{w}}$ is the highest weight, the variety $Y(\mathbf{v}, \mathbf{w})$ corresponds to the weight space of weight $\Lambda_{\mathbf{w}} - \alpha_{\mathbf{v}}$. It turns out that if $W = W_Q$ is the Weyl group attached to the Lie algebra $\mathfrak{g}_Q$, then the symplectic Galois group of the symplectic resolution $Y(\mathbf{v}, \mathbf{w})$ is simply the parabolic subgroup of the W_Q which stabilises this weight. The corresponding assertion can fail for simply laced quivers Q of affine type; see Example 5.12 for more details.

Let us introduce some notation: recall that the roots of $\mathfrak{g}$ are denoted by Φ . If $\lambda \in P$, then we write

$$\Phi_{\lambda} = \{\alpha \in \Phi : (\lambda, \alpha) = 0\}$$

for the parabolic subroot system associated to λ , and Φ_{λ}^+ for the positive roots in Φ_{λ} . Let us denote by $\Phi_{\lambda}^{\max}$ the roots in Φ_{λ} which are maximal in Φ_{λ} with respect to the partial order on R^+ .

Lemma 5.5 *Let $X(\mathbf{v}, \mathbf{w})$ be a quiver variety such that $X(\mathbf{v}, \mathbf{w})^{\text{reg}} \neq \emptyset$. Then if $\mu = \Lambda_{\mathbf{w}} - \alpha_{\mathbf{v}}$, the codimension two strata of $X(\mathbf{v}, \mathbf{w})$ are indexed by the roots in $\Phi_{\mu}^{\max}$.*

Proof Recall that $X(\mathbf{v}', \mathbf{w})^{\text{reg}} \neq \emptyset$ if and only if $\Lambda_{\mathbf{w}} - \alpha_{\mathbf{v}'} \in P^+$, and this is a stratum of $X(\mathbf{v}, \mathbf{w})$ if $\alpha_{\mathbf{v}'} \leq \alpha_{\mathbf{v}}$. Let $\alpha_{\mathbf{v}} = \alpha_{\mathbf{v}'} + \beta$, and let $\mu = \Lambda_{\mathbf{w}} - \alpha_{\mathbf{v}}$. Then if $\Lambda_{\mathbf{w}} - \alpha_{\mathbf{v}'} \in P^+$, the dimension formula for quiver varieties shows that $X(\mathbf{v}', \mathbf{w})^{\text{reg}}$ is a codimension two stratum precisely when

$$2(\mu, \beta) + (\beta, \beta) = 2.$$

Since $\mu \in P^+$ and $\beta \in R^+$, we must have $(\mu, \beta) = 0$ and $(\beta, \beta) = 2$, that is, $\beta \in \Phi_{\mu}$, the parabolic subsystem of Φ corresponding to μ . Thus the codimension two strata are labelled by roots $\beta \in \Phi_{\mu}^+$ for which $\mu + \beta \in P^+$.

We now claim that this set is exactly $\Phi_{\mu}^{\max}$. Indeed if β is such that $\mu + \beta \in P^+$, then for any $\alpha \in \Phi_{\mu}^+$ we have $0 \leq (\mu + \beta, \alpha) = (\beta, \alpha)$, and hence (as our root system is simply laced, so root strings have length at most 2) it follows that $\beta \in \Phi_{\mu}^{\max}$.

Conversely, if $\beta \in \Phi_{\mu}^{\max}$, and $\alpha_i \in \Phi$ is a simple root, consider $(\mu + \beta, \alpha_i)$. If $\alpha_i \in \Phi_{\mu}$ then $(\beta, \alpha_i) \geq 0$ since β is maximal (otherwise $s_{\alpha_i}(\beta) = \beta + \alpha_i \in \Phi_{\mu}$) and so $(\mu + \beta, \alpha_i) \geq 0$. If $\alpha_i \notin \Phi_{\mu}$, then since $\mu \in P^+$ we must have $(\mu, \alpha_i) \geq 1$, and since our root system is simply laced, $(\alpha_i, \beta) \geq -1$, and hence $(\mu + \beta, \alpha_i) \geq 0$ in this case also. $\square$

Lemma 5.6 *Suppose that $X(\mathbf{v}, \mathbf{w})$ is as above, and $\beta \in \Phi_{\mu}^{\max}$. Then the slice through the stratum associated to β by Lemma 5.5 is isomorphic to the Kleinian singularity indexed by the subroot system $\Phi(\beta) = \{\alpha \in \Phi : \alpha \leq \beta\}$, an indecomposable subsystem of Φ .*

Proof Nakajima [2001] constructs an (analytic) slice through a stratum $X(\mathbf{v}', \mathbf{w})^{\text{reg}}$ which is isomorphic to $X(\mathbf{v} - \mathbf{v}', \mathbf{w} - C\mathbf{v}')$. In our case it follows that the slices are isomorphic to $X(\mathbf{v}_0, \mathbf{w}_0)$ where $\alpha_{\mathbf{v}_0} = \beta$ and $\Lambda_{\mathbf{w}_0} = \mu + \beta$. It is known that the support $\text{supp}_R(\beta)$ of a root β (see Section 4.1) is always a connected subdiagram of the Dynkin diagram. The condition that $(\beta, \mu) = 0$ shows that $X(\mathbf{v}_0, \mathbf{w}_0)$ is isomorphic to the quiver variety $X(\mathbf{v}'_0, C'(\mathbf{v}'_0))$ for the quiver associated to the subdiagram Q' given by $\text{supp}_R(\beta)$. (Here C' is the Cartan matrix attached to Q' and $\mathbf{v}'_0$ is just $\mathbf{v}_0$ viewed as a dimension vector

for Q' .) It is well known that this quiver variety is the Kleinian singularity associated to the Weyl group of type Q' , or equivalently, the root system $\Phi(\beta)$. $\square$

Remark 5.7 The arguments above can also be used to show that the two-dimensional quiver varieties associated to an ADE quiver Q are precisely the Kleinian singularities attached to Dynkin diagrams which occur as subdiagrams of Q .

Theorem 5.8 Suppose $\mathbf{v}$ and $\mathbf{w}$ are as above. Let W_μ be the Weyl group of the subroot system Φ_μ . The symplectic Galois group associated to the quiver variety $X(\mathbf{v}, \mathbf{w})$ is W_μ .

Proof We have seen that the codimension-two strata are labelled by the set $\Phi_\mu^{\max}$. Each $\beta \in \Phi_\mu^{\max}$ is associated to an indecomposable subroot system $\Phi(\beta)$ of Φ_μ . Lemma 5.6 shows that the slice through the associated stratum is the Kleinian singularity associated to the subroot system, which has symplectic Galois group $W(\Phi(\beta))$, the Weyl group of the root system $\Phi(\beta)$. Now Namikawa's description of the symplectic Galois groups shows that the group attached to $X(\mathbf{v}, \mathbf{w})$ is

$$W = \prod_{\beta \in \Phi_\mu^{\max}} W(\Phi(\beta))^{\sigma_\beta},$$

where σ_β is a diagram automorphism of the Dynkin diagram induced by monodromy coming from the stratum. Thus it remains to show that this monodromy must be trivial. We show this in the following proposition. $\square$

Proposition 5.9 Let $\mathbf{v}$ and $\mathbf{w}$ be such that $X(\mathbf{v}, \mathbf{w})^{\text{reg}} \neq \emptyset$. The monodromy automorphism σ_β for every codimension-two stratum in $X(\mathbf{v}, \mathbf{w})$ is trivial.

Proof Let $\eta = \Lambda_{\mathbf{w}} - \alpha_{\mathbf{v}}$, and suppose that $X(\mathbf{v}_1, \mathbf{w})^{\text{reg}}$ is a codimension-two stratum of $X(\mathbf{v}, \mathbf{w})$. Then by Lemma 5.5, there is some $\beta \in \Phi_\eta^{\max}$ such that $\alpha_{\mathbf{v}} = \alpha_{\mathbf{v}_1} + \beta$.

Let $Z(\mathbf{v}, \mathbf{v}, \mathbf{w}) = Y(\mathbf{v}, \mathbf{w}) \times_{X(\mathbf{v}, \mathbf{w})} Y(\mathbf{v}, \mathbf{w})$, and let $\tilde{\pi} : Z(\mathbf{v}, \mathbf{v}, \mathbf{w}) \rightarrow X(\mathbf{v}, \mathbf{w})$ be the canonical map. By [Nakajima 1998, Theorem 7.2], $Z(\mathbf{v}, \mathbf{v}, \mathbf{w})$ is a Lagrangian subvariety of $Y(\mathbf{v}, \mathbf{w}) \times Y(\mathbf{v}, \mathbf{w})$.

Let $H_{\text{top}}(Z(\mathbf{v}, \mathbf{v}, \mathbf{w}))$ denote the top-dimensional Borel–Moore homology of $Z(\mathbf{v}, \mathbf{v}, \mathbf{w})$ and write $[C]$ for the fundamental class of a component C of $Z(\mathbf{v}, \mathbf{v}, \mathbf{w})$. Let

$$H_{\text{top}}(Z(\mathbf{v}, \mathbf{v}, \mathbf{w}))_{\mathbf{v}_1} = \text{span}_{\mathbb{K}} \{[C] : C \text{ a component of } Z(\mathbf{v}_1, \mathbf{v}_2, \mathbf{w}) \cap \overline{\tilde{\pi}^{-1}(X(\mathbf{v}_1, \mathbf{w})^{\text{reg}})}\}.$$

For the irreducible highest-weight representation $L(\lambda)$ of the Lie algebra $\mathfrak{g}_Q$ associated to our quiver, let $L(\lambda)_\mu$ denote its μ -weight space. It is shown in [Nakajima 1998, Theorem 10.15] that there is a natural isomorphism

$$H_{\text{top}}(Z(\mathbf{v}_1, \mathbf{v}_1, \mathbf{w}))_{\mathbf{v}} \cong L(\eta + \beta)_\eta^* \otimes L(\eta + \beta)_\eta.$$

But now the number s of components in $H_{\text{top}}(Z(\mathbf{v}, \mathbf{v}, \mathbf{w}))_{\mathbf{v}_1}$ is just the number of orbits of the monodromy action on the components of $\tilde{\pi}^{-1}(x)$ for $x \in X(\mathbf{v}_1, \Lambda_{\mathbf{w}})^{\text{reg}}$. Now clearly there are r^2 such components, where r is the number of components of $\pi_{\mathbf{v}}^{-1}(x)$. Since Theorem 10.2 of [Nakajima 1998] shows that the dimension of $L(\eta + \alpha)_\eta$ is exactly r , it follows that $s = r^2$, and hence the monodromy action must be trivial, as required. $\square$

Example 5.10 It will be convenient in this example to write $X(\alpha_v, \Lambda_w)$ rather than $X(v, w)$. Suppose that Q is the A_3 quiver and $\Lambda_w = \Lambda_1 + \Lambda_2 + \Lambda_3$, $\alpha_v = \alpha_1 + \alpha_2 + \alpha_3$. Then $X(\alpha_v, \Lambda_w)$ is 4-dimensional, and has four strata, labelled by $0, \alpha_1 + \alpha_2, \alpha_2 + \alpha_3$ and $\alpha_v = \alpha_1 + \alpha_2 + \alpha_3$. Here the stratum corresponding to 0 is a point, and each of $X(\alpha_1 + \alpha_2, \Lambda_w)^{\text{reg}}$ and $X(\alpha_2 + \alpha_3, \Lambda_w)^{\text{reg}}$ is two-dimensional. The slice through the stratum $X(\alpha_1 + \alpha_2, \Lambda_w)$ is isomorphic to $X(\alpha_3, 2\Lambda_3)$, while that through $X(\alpha_2 + \alpha_3, \Lambda_w)$ is isomorphic to $X(\alpha_1, 2\Lambda_1)$. These are both Kleinian singularities of type A_1 , and so the symplectic Galois group in this case is $S_2 \times S_2$, which is the parabolic subgroup of S_4 stabilizing the weight $\Lambda_w - \alpha_v = \Lambda_2$.

Our calculation of the symplectic Galois group for finite type quiver varieties, along with the general theory of Section 3, have the following immediate consequence:

Corollary 5.11 *Let $\pi : Y(v, w) \rightarrow X(v, w)$ be as above, and let $\lambda = \Lambda_w - \alpha_v$. If $x \in X(v, w)$, then the cohomologies $H^*(\pi^{-1}(x))$ carry an action of the group W_λ .*

Example 5.12 We give an example of a symplectic Galois group for a quiver variety of affine type. Let Q_Γ be the affine quiver associated by the McKay correspondence to the finite subgroup Γ of $\text{SL}_2(\mathbb{C})$. Then if we take $\alpha_v = n\delta$, where δ is a generator for the imaginary roots, and $\Lambda_w = \Lambda_0$, where 0 denotes the extending node of the affine quiver, then it is known that $X(v, w)$ is isomorphic to $\mathbb{C}^{2n}/\Gamma_n$, where $\Gamma_n = \Gamma \wr S_n$. Now for the quotient of a symplectic vector space by a symplectic reflection group such as Γ_n it is known that the stratification by stabiliser type coincides with that of the symplectic leaves. It follows that there are two codimension two strata: one labelled by the group Γ and the other by a subgroup of order two generated by a simple reflection in S_n (of course if $n = 1$ then this stratum does not exist). It can be checked that the symplectic Galois group is therefore either $W(\Gamma)$, the finite Weyl group associated to Γ if $n = 1$, or $S_2 \times W(\Gamma)$ for $n > 1$.

Let W_{aff} denote the affine Weyl group, that is, the Weyl group associated to the affine Lie algebra $\mathfrak{g}_Q$ whose Dynkin diagram underlies the affine quiver Q . Note that W_{aff} fixes the imaginary root δ , and thus $\Lambda_0 - \delta$ and $\Lambda_0 - n\delta$ for $n > 1$ have the same stabiliser in W_{aff} , thus the naive generalisation of Theorem 5.8 fails to hold.

Remark 5.13 If one is willing to work with all quivers and arbitrary dimension vectors, the distinction between “framed” and “unframed” quivers disappears thanks to the “Crawley-Boevey trick”. In essence this exchanges framing vertices for “framing edges” all connected to a single framing vertex, denoted by ∞ , equipped with a 1-dimensional vector space. The “diagonal” copy of $\mathbb{G}_m$ in the unframed quiver can be identified with the gauge group at the vertex ∞ . Bellamy and Schedler [2021, Section 7] calculate the codimension-two strata in an arbitrary unframed affine quiver variety in terms of the root system combinatorics of the associate Kac–Moody Lie algebra $\mathfrak{g}_Q$ and establish the isomorphism type of a transverse slice through each of these strata. They thus give the analogues of our Lemmas 5.5 and 5.6 in the unframed setting. It is however not clear, at least to the authors, how to interpret the results of [Bellamy and Schedler 2021] in terms of the Kac–Moody algebra $\mathfrak{g}_Q$ (rather than $\mathfrak{g}_{Q'}$).

Remark 5.14 The proof that the monodromy action on the components of the fibres of π_v should work in a somewhat more general context than that of finite-type quivers: Nakajima’s results [1998, Section 10]

produce an isomorphism from a modified form of the enveloping algebra to the “regular” part of the convolution algebra

$$\mathcal{X} = \bigoplus_{\mathbf{v}^1, \mathbf{v}^2} H_{\text{top}}(Z(\mathbf{v}^1, \mathbf{v}^2, \mathbf{w})^{\text{reg}}),$$

where the superscript “reg” means that one restricts to the open subset which is the complement of the closure of points in $Z(\mathbf{v}_1, \mathbf{v}_2, \mathbf{w})$ whose image under $\tilde{\pi}$ does not lie in any locus $X(\mathbf{v}', \mathbf{w})^{\text{reg}}$. In the finite-type case the regular set exhausts $Z(\mathbf{v}_1, \mathbf{v}_2, \mathbf{w})$, but in general it is a proper subset. The proof of Proposition 5.9 immediately extends to show that the monodromy action is trivial on the components of the fibres of codimension-two strata which are regular.

Remark 5.15 Many authors, including Nakajima [2003], Lusztig [2000] and Maffei [2002], have studied actions of the Weyl group on the cohomology of quiver varieties. It would be interesting to compare the actions constructed here with these actions. A novelty of our approach is that the action of the group W_λ is constructed without the use of a presentation of W_λ : previous approaches define the action of simple reflections and then must explicitly check the braid relations to deduce the existence of a Weyl group action. On the other hand, these approaches give a more explicit description of the action of the simple reflections.

A first step in such a comparison might be to exhibit an explicit homomorphism from the Weyl group to the symplectic Galois group. As the previous example shows, we should not expect this to be surjective in general, but it seems reasonable to expect that there should be a natural embedding. Evidence for this is provided, for example, by the generic isomorphisms constructed in [Maffei 2002]. We hope to address the existence of such an embedding in future work.

6 Bionic symplectic varieties and Springer theory

In this section we consider a special class of conic symplectic resolutions, which we call *bionic*. These are symplectic resolutions $\pi : Y \rightarrow X$ which, in addition to having a conical action $\lambda : \mathbb{G}_m \rightarrow \text{Aut}(Y)$ acting with positive weight on the symplectic form, have another action $\mathbb{G}_m$ -action $\rho : \mathbb{G}_m \rightarrow \text{Aut}(Y)$, commuting with the λ -action, which is required to be Hamiltonian (and thus preserves the symplectic form on Y). This class of resolutions has been considered by a number of authors; see for example [Losev 2017] and [Braden et al. 2016].

Lemma 6.1 *The actions λ and ρ naturally extend to an action on the deformation $\pi : \mathcal{Y} \rightarrow \mathcal{X}$. Moreover, the Hamiltonian action acts trivially on the base of the deformation of Y , thus it induces an action on each deformation Y_t for $t \in H^2(Y, \mathbb{C})$.*

Proof The fact that the action λ extends is proved in [Namikawa 2011]. The arguments in [Namikawa 2008, Lemma 20] show that the Hamiltonian action naturally extends to a Hamiltonian action on the formal deformation. That argument moreover shows that the action preserves the base of the deformation. Since ρ commutes with λ , the action on the formal deformation extends to the algebraicisation, and moreover again fixes the base of the deformation, and hence induces a Hamiltonian action on each deformation. $\square$

Though we will focus on the case where Y^ρ is finite, it is checked in [Losev 2017] that the connected components of the fixed point locus are always themselves conical symplectic resolutions.

Proposition 6.2 [Losev 2017, Proposition 5.1] *Each connected component of the fixed point locus Y^ρ for the ρ -action on Y is a conical symplectic resolution.*

Lemma 6.3 *Let $\rho_c : U_1 \rightarrow \text{Aut}(Y)$ be the restriction of the $\mathbb{G}_m$ action to the circle U_1 . The $\mathbb{C}^\infty$ -trivial fibration $\mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ can be made ρ_c -equivariantly topologically trivial. In particular, the family $\mathcal{Y}^\rho \rightarrow H^2(Y, \mathbb{C})$ is topologically locally trivial.*

Proof That the trivialisation can be done U_1 -equivariantly follows immediately from Thom's second isotopy theorem (see [Mather 2012] for example) that the $\mathbb{C}^\infty$ -trivial fibration $\mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ can be made ρ_c -equivariantly topologically trivial. The local triviality of the fixed-point locus follows immediately. $\square$

Lemma 6.4 *Suppose that ρ has finitely many fixed points on Y . Then $\mathcal{Y}^\rho$ is a trivial étale covering of $B_Y = H^2(Y, \mathbb{C})$.*

Proof The map $q_Y : \mathcal{Y} \rightarrow B_Y$ is smooth, hence it is surjective on tangent spaces. The action of ρ is trivial on B_Y , and if $x \in Y^\rho$ then dq_Y must induce an isomorphism on the ρ -fixed subspace of $T_x \mathcal{Y}$ (it must be surjective, and the kernel of dq_Y is the subspace $T_x Y$, so since x is an isolated fixed point of the ρ -action on Y , it must also be injective). The dimension of the ρ -fixed subspace in the fibres of the restriction of $T\mathcal{Y}$ to $\mathcal{Y}^\rho$ is locally constant. Moreover, since λ commutes with ρ , it will contract $\mathcal{Y}^\rho$ to a subvariety of Y^ρ , hence the rank calculation for fixed points on Y imply that $dq_Y|_{\mathcal{Y}^\rho}$ is an isomorphism at every point $y \in \mathcal{Y}^\rho$. Thus q_Y restricts to an étale map on $\mathcal{Y}^\rho$.

Next we take an equivariant compactification $\bar{\mathcal{Y}}$ of $\mathcal{Y}$, using for example [Brion 2017, Theorem 1.2]. Applying the same argument used above we see that the fixed locus on the compactification is proper and étale, and hence finite. By taking an equivariant compactification, one sees by the same argument that the fixed locus on the compactified locus is proper and étale, hence finite. If the restriction to the original locus was not finite, it can only be because the fibre multiplicities are nonconstant. But by Lemma 6.4, we know $\mathcal{Y}^\rho \rightarrow H^2(Y, \mathbb{C})$ is a locally trivial family, hence the fibre multiplicities are locally constant, and so, as $H^2(Y, \mathbb{C})$ is connected, constant. It follows that the fibre multiplicities are constant, and it follows that q_Y is finite, as required. Since $H^2(Y, \mathbb{C})$ is simply connected, it follows immediately that it is a trivial covering. $\square$

Example 6.5 Let $Y = T^*\mathcal{B}$ be the flag variety for a reductive group G and $X = \mathcal{N}$ the nilpotent cone. It is well known that the moment map $\mu : Y \rightarrow X$ is a symplectic resolution. If T is a maximal torus of G and $\rho : \mathbb{G}_m \rightarrow T$ is a one-parameter subgroup, then the induced action of T on Y is Hamiltonian, so ρ gives a Hamiltonian $\mathbb{G}_m$ -action on Y . Let $L = G^\rho$ be the subgroup of G fixed by the action of ρ . It is a Levi subgroup of G . The fixed-point locus of the ρ -action is a disjoint union of varieties isomorphic to $T^*\mathcal{B}_L$, where $\mathcal{B}_L$ is the flag variety of L . The connected components are in bijection with W/W_L , where W is the Weyl group of G and W_L is the Weyl group of L . In this case the restriction of the moment map is in fact a symplectic resolution.

We now restrict again to the case where ρ has finitely many fixed points. We will also, for simplicity, assume that all sheaves have $\mathbb{Q}$ -coefficients. In this case we have a permutation action of W on the fixed point loci Y_t^ρ due to the trivialisation given by Lemma 6.4. Let us denote by P the corresponding permutation representation of W . We claim that the representation of the symplectic Galois group on $H^*(\pi^{-1}(0))$ is isomorphic to P . To show this, we will use equivariant cohomology, and hence we need to lift all of our constructions to the setting of the equivariant derived category. This is completely standard; see [Bernstein and Lunts 1994] for example. As we will only need to work with the group U_1 , which is connected, the forgetful functor from equivariant perverse sheaves to perverse sheaves is fully faithful and hence equivariance, roughly speaking, is a property and not extra structure. The constant sheaf on any U_1 -variety is equivariant, and where functors are compatible with the U_1 action, they lift to the equivariant setting, hence Spr lifts to a U_1 -equivariant perverse sheaf. Since $H_{U_1}^*(\star) = H^*(B\mathbb{G}_m)$, in the equivariant setting all cohomologies are modules over $H^*(B\mathbb{G}_m)$, which is a polynomial ring in one variable.

We will show that $H^*(\pi^{-1}(0))$ is isomorphic to the permutation representation P using localisation and the nearby cycles description of Spr . The equivariant derived category has a six functor formalism allowing us to lift the nearby cycles functor. In the equivariant cohomology of the fixed points, the family $\mathcal{Y}^\rho \rightarrow H^2(Y, \mathbb{C})$ is trivial, with the generic action of W being given by the permutation action.

Fix a Hermitian metric on $\mathcal{X}$ which is invariant under the action of U_1 the maximal compact subgroup of ρ , the Hamiltonian $\mathbb{G}_m$ -action (since X is affine this can be done globally), and a W -invariant metric on $H^2(Y, \mathbb{C})$, so that it descends to give a metric on $B_X = H^2(Y, \mathbb{C})/W$. The *Milnor fibre* of the map $\theta_X : \mathcal{X} \rightarrow B_X$ at a point $x \in X$ is $F_{x,\epsilon} = B(x, \epsilon) \cap \theta_X^{-1}(v)$, where $\delta = d(0, v) \ll \epsilon \ll 1$. We will write $\mathbb{B}_\epsilon$ for the ball of radius ϵ centred at $0_X \in \mathcal{X}$, and $B_X(\delta)$ for the ball of radius δ centred at $0 \in B_X$, and $B_X(\delta)^{\text{reg}} = B_X(\delta) \cap B_X^{\text{reg}}$ for the regular values of θ_X in $B_X(\delta)$ (cf Definition 3.19).

Lemma 6.6 *The stalk at $0 \in X$ of the U_1 -equivariant lift of the Springer sheaf Spr_{U_1} is $H_{U_1}^*(F_{\epsilon,v})$, where $F_{\epsilon,v} = \mathbb{B}_\epsilon \cap \theta_X^{-1}(v)$ for $0 < \|v\| < \delta \ll \epsilon \ll 1$.*

Proof Note that since we have chosen our Hermitian metric to be U_1 -invariant, the spaces $F_{x,\epsilon}$ are preserved by U_1 , and so it makes sense to speak their U_1 -equivariant cohomology. By Theorem 3.21, the Springer sheaf is given by a nearby cycles construction. Now it follows from the construction of the nearby cycles functor that for sufficiently small ϵ and δ , the spaces F_ϵ are all homeomorphic, and the stalk of the nearby-cycles sheaf $\mathcal{P}_{\theta_X,0}$ is given by the cohomology of F_ϵ , and hence working in the equivariant setting, this is $H_{U_1}^*(F_\epsilon)$. $\square$

Theorem 6.7 *Suppose that the Hamiltonian $\mathbb{G}_m$ acts with finitely many fixed points. Then the symplectic Galois group W permutes the fixed points in a generic deformation $\mathcal{X}_t$ and the cohomology $H^*(X_t^\rho)$ has a filtration whose associated graded is isomorphic as a W -representation to $H^*(Y)$.*

Proof If the Hamiltonian $\mathbb{G}_m$ -action ρ has finitely many fixed points on Y , then we have seen that the fixed-point locus $\mathcal{Y}^\rho$ is a trivial finite cover of $H^2(Y, \mathbb{C})$. Moreover, since the action of ρ is compatible with the deformation it follows that the symplectic Galois group action commutes with the action of ρ

and so acts on $\mathcal{Y}^\rho$, hence the triviality of the covering $\mathcal{Y}^\rho \rightarrow H^2(Y, \mathbb{C})$ yields a permutation action of W on the sheets of the covering.

Now the action of the symplectic Galois group on the stalks of the Springer sheaf is, via the nearby cycles description, given by the monodromy action arising from the fact that the construction of ψ_{θ_X} is locally constant as one varies the analytic arc used in the construction of ψ_{θ_X} ; see Proposition 3.17. To prove the theorem we use localisation theory for the U_1 -action, following the strategy of [Treumann 2009]. Let $F_\epsilon = \tilde{\pi}^{-1}(\mathbb{B}_\epsilon \cap \theta_X^{-1}(B_X(\delta)^{\text{reg}}))$ and F_ϵ^ρ the fixed points of ρ in F_ϵ . Then if ψ is the restriction of $\theta_X \circ \tilde{\pi}$ to F_ϵ and $\iota : F_\epsilon^\rho \rightarrow F_\epsilon$ the inclusion map, and $\phi = \psi \circ \iota$, we have diagrams

$$\begin{array}{ccc} F_\epsilon^\rho & \xrightarrow{\iota} & F_\epsilon \\ & \searrow \phi & \swarrow \psi \\ & B_X^{\text{reg}}(\delta) & \end{array} \quad \text{and} \quad \begin{array}{ccc} (F_\epsilon^\rho)_{hU_1} & \xrightarrow{(\iota)_{hU_1}} & (F_\epsilon)_{hU_1} \\ & \searrow \tilde{\phi} & \swarrow \tilde{\psi} \\ & BU_1 \times B_X^{\text{reg}}(\delta) & \end{array}$$

where if M is any topological space with a U_1 -action we write M_{hU_1} for the Borel construction $EU_1 \times_{U_1} M$, and the top row of the diagram on the right is obtained from that on the left by applying the Borel construction to the inclusion map ι . The maps with the maps $\tilde{\phi}$ and $\tilde{\psi}$ induced by ϕ and ψ since the U_1 -action preserves the fibres of those maps. In more detail, we have a natural map

$$\tilde{\psi} : (F_\epsilon)_{hU_1} = EU_1 \times_{U_1} F_\epsilon \rightarrow BU_1 \times B_X(\delta)^{\text{reg}}, \quad \tilde{\psi}([e, x]) = (U_1 \cdot e, \psi(x)),$$

with fibre

$$\tilde{\psi}^{-1}(p, v) = F_{\epsilon, v} = \psi^{-1}(v) = \mathbb{B}_\epsilon \cap \theta_X^{-1}(v) \quad \text{for } v \in B_X^{\text{reg}}(\delta), p \in BU_1.$$

The map $\tilde{\phi}$ is defined in the same way.

Now let $\pi_2 : BU_1 \times B_X(\delta)^{\text{reg}} \rightarrow B_X(\delta)^{\text{reg}}$ be the projection to the second factor and let

$$\mathcal{F}_\rho = \bigoplus_{i \in \mathbb{Z}} R^i(\pi_2 \circ \tilde{\phi})_*(\mathbb{Q}) \quad \text{and} \quad \mathcal{F} = \bigoplus_{i \in \mathbb{Z}} R^i(\pi_2 \circ \tilde{\psi})_*(\mathbb{Q}).$$

Then $\mathcal{F}_\rho$ and $\mathcal{F}$ are $\mathbb{Z}$ -graded locally constant sheaves with fibres at $v \in B_X(\delta)^{\text{reg}}$ isomorphic to the U_1 -equivariant cohomology of the fibres $\phi^{-1}(v)$ and $\psi^{-1}(v)$, respectively. This shows that $\tilde{\psi}^{-1}(v) = (\psi^{-1}(v))_{hU_1}$ so that we may view p as a fibration relative to the B_X^{reg} , ie as a fibration in the category of locally trivial fibre bundles over B_X^{reg} : As a result, the associated Leray spectral sequence for the fibration p can be treated as a spectral sequence of locally constant sheaves over B_X^{reg} , that is, as a spectral sequence of $\pi_1(B_X^{\text{reg}})$ -representations.

Fix a ring isomorphism of the rational cohomology $H^*(BU_1, \mathbb{Q})$ with $\mathbb{Q}[t]$ (so that t has degree 2). By Lemma 6.4, the map $\mathcal{Y}^\rho \rightarrow H^2(Y, \mathbb{C})$ is a trivial covering, ie it is isomorphic to the product $Y^\rho \times H^2(Y, \mathbb{C})$. Moreover, U_1 acts trivially on this space, hence the Borel construction for F_ϵ^ρ is just $BU_1 \times Y^\rho \times B_X^{\text{reg}}(\delta)$. It follows that the graded local system $\mathcal{F}_\rho$ is just the constant local system with fibre $H^*(BU_1) \otimes H^0(Y^\rho) = \text{Ind}_1^W(\mathbb{Q}[t] \otimes H^0(Y^\rho))$, where Ind_1^W denotes induction from the trivial subgroup of W .

First consider the map $\tilde{\psi}$. If $\pi_2 : BU_1 \times B_X(\delta)^{\text{reg}} \rightarrow B_X(\delta)^{\text{reg}}$ denotes the projection to the second factor, we therefore have a commutative diagram

$$\begin{array}{ccc} (F_\epsilon)_{hU_1} & \xrightarrow{p} & BU_1 \times B_X^{\text{reg}}(\delta) \\ & \searrow \tilde{\psi} & \swarrow \pi_2 \\ & B_X^{\text{reg}}(\delta) & \end{array}$$

The stalks of the sheaves on its E_2 -page are $E_2^{ij} = H^i(BU_1, \mathbb{Q}) \otimes H^j(F_\epsilon, \mathbb{Q})$, with the sequence converging to $H_{U_1}^*(F_\epsilon)$. However, it is known that by [Kaledin 2006, Theorem 2.12] the cohomology of the fibres of a symplectic resolution $\pi : Y \rightarrow X$ are nonzero only in even degrees, hence $H^i(BU_1) \otimes H^j(F_\epsilon)$ is zero unless $i = 2k$ and $j = 2l$, in which case it is $t^k \otimes H^{2l}(F_\epsilon)$. Thus the differential on the E^2 page vanishes for degree reasons and so the spectral sequence collapses at the E_2 -page. In particular, it follows that $H_{U_1}^*(F_\epsilon)$ is a free $\mathbb{Q}[t]$ -module, and that, if we set $F_k(H_{U_1}^*(F_\epsilon)) = t^k H_{U_1}^*(F_\epsilon)$, then $\text{gr}^F(H_{U_1}^*(F_\epsilon)) \cong \mathbb{Q}[t] \otimes H^*(F_\epsilon)$ as graded rings.

Now consider $\iota^* : H_{U_1}^*(F_\epsilon) \rightarrow H_{U_1}^*(F_\epsilon^\rho)$. By the localisation theorem in equivariant cohomology (see for example [Quillen 1971]) this map is injective, and becomes an isomorphism after localising. In view of the degeneration of the Leray spectral sequence for $H_{U_1}^*(F_\epsilon)$, it suffices to invert $t \in \mathbb{Q}[t]$. We may therefore identify $H_{U_1}^*(F_\epsilon)$ with its image in $\mathbb{Q}[t] \otimes H^0(Y^\rho)$. As $\mathbb{Q}[t]$ is a Euclidean domain, there is a basis $B = \{e_1, \dots, e_n\}$ of $\mathbb{Q}[t] \otimes H^0(Y^\rho)$ and $f_1, \dots, f_n \in \mathbb{Q}[t]$ for which $\{f_1 e_1, \dots, f_n e_n\}$ is a basis of $H_{U_1}^*(F_\epsilon)$. Moreover, since $\mathbb{Q}[t] \otimes H^0(Y^\rho)$ is a graded $\mathbb{Q}[t]$ -module, and $H_{U_1}^*(F_\epsilon)$ a graded submodule, we may assume that the basis $\{e_1, \dots, e_n\}$ and the f_i are homogeneous elements. It follows that B is a basis of $H^0(Y^\rho) = H_{U_1}^0(Y^\rho)$ and that we may assume $f_i = t^{n_i}$ for some $n_i \in \mathbb{Z}_{\geq 0}$. But now if we define $G_k(H^0(Y^\rho)) = \text{span}_{\mathbb{Q}}\{e_i : n_i \geq k\}$, we have

$$\text{gr}^G(H^0(Y^\rho)) \cong H_{U_1}^*(F_\epsilon)/t \cdot H_{U_1}^*(F_\epsilon) \cong H^*(F_\epsilon) \cong H^*(Y). \quad \square$$

Remark 6.8 The action of W on the fixed-point locus of generic fibres of $\mathcal{X}$ can also be described as follows: fix $s_0 \in B_X^{\text{reg}}$ and pick $t_0 \in H^2(Y, \mathbb{C})$ with $q(t_0) = s_0$. Now $\mathcal{Y}_{t_0} \cong \mathcal{X}_{s_0}$ via the restriction $\tilde{\pi}_{t_0}$ of $\tilde{\pi}$, so that $\mathcal{X}_{s_0}^\rho$ is identified with $\mathcal{Y}_{t_0}^\rho$. But W acts on $\mathcal{Y}$ compatibly with the ρ -action (by the universality of the deformation $\mathcal{Y}$ for example), where $w \in W$ satisfies $w(\mathcal{Y}_t) = \mathcal{Y}_{w(t)}$, where $\tilde{\pi}_{w(t_0)} : \mathcal{Y}_{w(t_0)} \rightarrow \mathcal{X}_{q(w(t_0))}$ is an isomorphism and $q(w(t_0)) = q(t_0) = s_0$. It follows that we obtain an action of W on $\mathcal{X}_{q(t)}$ by setting $w(x_0) = \tilde{\pi}_{w(t)} \circ w \circ \tilde{\pi}_t^{-1}(x_0)$ for all $w \in W$ and $x_0 \in \mathcal{X}_{s_0}^\rho$.

Since $\mathcal{Y}^\rho \rightarrow H^2(Y, \mathbb{C})$ is a trivial covering, we may identify Y^ρ with the fixed point locus $\mathcal{Y}_t^\rho$ for any t , thus we may also view this as an action of W on Y^ρ .

7 Slices in the affine Grassmannian

The idea that slices in the affine Grassmannian should have a kind of Springer theory associated to them goes back to work of Ngô [1999, Section 2], and was studied in the context of knot invariants by Kamnitzer [2011]. A detailed study of the Poisson geometry of these slices was initiated in [Kamnitzer

et al. 2014], and we will review briefly some results of that paper below. As should be clear to the reader, the main contribution of this section is to compare the Beilinson–Drinfel’d deformation with the universal conic deformation and obtain some basic insight into the symplectic Springer theory for these symplectic varieties which, following the insights of Braverman, Finkelberg and Nakajima [Braverman et al. 2019] should perhaps be thought of as the “symplectic duals” of the finite-type quiver varieties studied in Section 5.

7.1 Background on slices

In this section we review some facts about slices in the affine Grassmannian. An excellent introduction to the Geometric Satake correspondence and the affine Grassmannian is [Zhu 2017], and more details on the geometry of slices are given in [Kamnitzer et al. 2014].

Let G be a reductive algebraic group over $\mathbb{C}$. We write $G[[t^{\pm 1}]]$ or $G((t^{\pm 1}))$ for the $\mathbb{C}[[t^{\pm 1}]]$ or $\mathbb{C}((t^{\pm 1}))$ points of G , and similarly we write $G[t]$ for the $\mathbb{C}[t]$ -points of G . The *thick affine Grassmannian* is the scheme $\mathbf{Gr} = G((t^{-1}))/G[t]$, which contains the thin affine Grassmannian Gr via the maps

$$G((t))/G[[t]] \cong G[t, t^{-1}]/G[t] \hookrightarrow G((t^{-1}))/G[t].$$

A coweight λ for G may be viewed as an element of $G[t, t^{-1}]$ and hence yields a point in the thin affine Grassmannian, which we denote by t^λ . Set, for λ and μ dominant coweights of G ,

$$\mathrm{Gr}^\lambda = G[t].t^\lambda \quad \text{and} \quad \mathbf{Gr}_\mu = G_1[[t^{-1}]].t^{w_0\mu},$$

where $G_1[[t^{-1}]]$ is the kernel of the natural homomorphism $G[[t^{-1}]] \rightarrow G$ given by evaluation at $t^{-1} = 0$, and as usual w_0 denotes the longest element of the Weyl group of G .

For λ and μ as above with $\mu \leq \lambda$, let $\mathrm{Gr}_\mu^{\bar{\lambda}} = \overline{\mathrm{Gr}^\lambda} \cap \mathbf{Gr}_\mu$. This is an affine slice to the orbit Gr^μ through the point $t^{w_0\mu}$, since $\mathbf{Gr}_\mu$ is transverse to every Gr^ν and intersects Gr^μ precisely at the point $t^{w_0\mu}$. It is known to have dimension $2\langle \rho, \lambda - \mu \rangle$ (where ρ as usual denotes the half-sum of positive roots). There is a natural action of $\mathbb{G}_m$ on Gr by “loop rotation”. This action preserves both $G[t]$ - and $G((t^{-1}))$ -orbits, and hence preserves each Gr_μ^λ .

Lemma 7.1 *Let λ and μ be dominant coweights of G such that $\mu \leq \lambda$. Then the variety $\mathrm{Gr}_\mu^{\bar{\lambda}}$ is a conical symplectic variety in the sense of Definition 2.7. Moreover, its symplectic leaves are $\mathrm{Gr}_\mu^\lambda = \mathrm{Gr}^\lambda \cap \mathbf{Gr}_\mu$.*

Sketch of proof The details of the proof of these statements are contained in Theorems 2.5 and 2.7 of [Kamnitzer et al. 2014]. Here we outline the key ideas. The Poisson structure is constructed from the Manin triple $(\mathfrak{g}((t^{-1})), \mathfrak{g}[t], t^{-1}\mathfrak{g}[[t^{-1}]])$, where the pairing between $\mathfrak{g}[t]$ and $t^{-1}\mathfrak{g}[[t^{-1}]]$ being given by the residue of the Killing form (note that $\mathrm{Lie}(G_1[[t^{-1}]]) = t^{-1}\mathfrak{g}[[t^{-1}]]$). The loop rotation action preserves $\mathrm{Gr}_\mu^{\bar{\lambda}}$ and contracts it to the unique fixed point $t^{w_0\mu}$, and its compatibility with the Manin triple described above ensures it makes $\mathrm{Gr}_\mu^{\bar{\lambda}}$ into a conic symplectic variety as claimed (acting with weight -1 on the Poisson bracket). The claim about the symplectic leaves follows from a calculation of the rank of the Poisson structure which is due to Lu and Yakimov [2008]. $\square$

Remark 7.2 If X is an arbitrary conical symplectic variety, it may not have a symplectic resolution. It is known, however, using the minimal model programme, that X will always have a $\mathbb{Q}$ -factorial terminalisation, that is, a crepant birational morphism $\pi : Y \rightarrow X$ such that Y is $\mathbb{Q}$ -factorial and possesses only terminal singularities. The former condition means that if D is a Weil divisor then there is some integer N such that $N.D$ is Cartier, while the latter condition means that any smooth resolution of Y will have positive discrepancy (so that Y is as smooth as a crepant (partial) resolution can get). Namikawa [2010] shows that, if $\pi : Y \rightarrow X$ is such a morphism, then Y is sufficiently smooth that it possesses a universal conical Poisson deformation $\mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ to which the conclusions of Theorem 2.8 apply. But then if X has a smooth crepant resolution, Y' say, Theorem 2.8 for the universal deformation of Y' shows that the generic Poisson deformations of X will be smooth, but then since $\tilde{\Pi} : \mathcal{Y} \rightarrow \mathcal{X} \times_{B_X} H^2(Y, \mathbb{C})$ induces an isomorphism $\mathcal{Y}_t \cong \mathcal{X}_{q(t)}$ for generic $t \in H^2(Y, \mathbb{C})$, we see that the generic deformation of Y must also be smooth. But Theorem 5.5 of [Namikawa 2011] shows that $\mathcal{Y} \rightarrow H^2(Y, \mathbb{C})$ is a locally trivial deformation (as a flat family of varieties without Poisson structure), so that $Y = \mathcal{Y}_0$ is smooth if $\mathcal{Y}_t$ is smooth. Thus remarkably, the existence of a single smooth crepant resolution implies that *any* $\mathbb{Q}$ -factorial terminalisation of X will in fact be smooth!

The previous remark implies that, given a conical symplectic variety X , if one can produce an explicit $\mathbb{Q}$ -factorial terminalisation $\pi : Y \rightarrow X$, then X has a symplectic resolution if and only if Y is in fact smooth, in which case of course, π is such a resolution! Perhaps surprisingly, [Kamnitzer et al. 2014] provide such a Y when $X = \text{Gr}_{\mu}^{\bar{\lambda}}$. To describe it we must first recall the convolution product for $G[t]$ -invariant subsets of $\mathbf{Gr}$. Let $p : G((t^{-1})) \rightarrow \mathbf{Gr}$ be the natural map. Then if $A, B \subseteq \mathbf{Gr}$ are $G[t]$ -invariant, we let $A \tilde{\times} B = p^{-1}(A) \times_{G[t]} B$.

Now consider $\text{Gr}_{\mu}^{\bar{\lambda}}$: given λ , assuming (as we shall) that G is of adjoint type, we may pick a sequence $\bar{\lambda} = (\lambda_1, \dots, \lambda_n)$ of fundamental coweights such that $\lambda = \lambda_1 + \dots + \lambda_n$. Then we have open and closed convolutions

$$\begin{array}{ccc} \text{Gr}^{\bar{\lambda}} = \text{Gr}^{\lambda_1} \tilde{\times} \dots \tilde{\times} \text{Gr}^{\lambda_n} & & \overline{\text{Gr}}^{\bar{\lambda}} = \overline{\text{Gr}}^{\lambda_1} \tilde{\times} \dots \tilde{\times} \overline{\text{Gr}}^{\lambda_n} \\ \downarrow m & \text{and} & \downarrow \bar{m} \\ \text{Gr}^{\lambda} & & \overline{\text{Gr}}^{\lambda} \end{array}$$

where m and $\bar{m}$ denote the convolution maps. If we set $\text{Gr}_{\mu}^{\bar{\lambda}} = \bar{m}^{-1}(\text{Gr}_{\mu})$ then by [Kamnitzer et al. 2014, Theorem 2.7], $\text{Gr}_{\mu}^{\bar{\lambda}}$ is a $\mathbb{Q}$ -factorial terminalisation of $\text{Gr}_{\mu}^{\lambda}$, and therefore, by Remark 7.2, a symplectic resolution of $\text{Gr}_{\mu}^{\lambda}$ whenever such a resolution exists. Provided the fundamental coweights λ_i are minuscule as weights for the dual group $G^{\vee}$, the orbits Gr^{λ_i} are actually closed, and hence the convolution product $\text{Gr}^{\bar{\lambda}}$ is smooth. A refinement of this observation shows the following.

Theorem 7.3 [Kamnitzer et al. 2014, Theorem 2.9]. *The variety $\text{Gr}_{\mu}^{\bar{\lambda}}$ has a symplectic resolution precisely when there do not exist coweights $\nu_1, \dots, \nu_k$ such that $\nu_1 + \dots + \nu_n = \mu$ and for all k , ν_k is a weight of $V(\lambda_k)$ (the irreducible representation of highest weight λ_k for the dual group $G^{\vee}$) and for some k , ν_k is not an extremal weight of $V(\lambda_k)$.*

Remark 7.4 We shall assume from now on that our variety $\mathrm{Gr}_{\mu}^{\bar{\lambda}}$ possesses a symplectic resolution. In this case, one can check (see Theorem 2.9 in [Kamnitzer et al. 2014]) that the preimage of $\mathrm{Gr}_{\mu}^{\bar{\lambda}}$ lies in the open convolution space $\mathrm{Gr}^{\lambda_1} \widetilde{\times} \cdots \widetilde{\times} \mathrm{Gr}^{\lambda_n}$, hence we will write $\mathrm{Gr}_{\mu}^{\bar{\lambda}}$ for the resolution.

Remark 7.5 The theorem shows that we must take λ to be a sum of minuscule coweights, if we wish to ensure $\mathrm{Gr}_{\mu}^{\bar{\lambda}}$ always has a symplectic resolution. When G is of type A , every fundamental weight is minuscule, hence this is automatic, but in general we see that the theorem gives a large supply of symplectic varieties which do not possess a symplectic resolution.

Namikawa's recipe for computing the symplectic Galois group can be carried out in these cases thanks to the work of [Malkin et al. 2005, Theorem 5.2]:

Lemma 7.6 *The codimension-two strata of the symplectic variety $\mathrm{Gr}_{\mu}^{\bar{\lambda}}$ are those $\mathrm{Gr}_{\nu}^{\lambda}$ with ν dominant and $\nu = \lambda - \check{\alpha}_i \geq \mu$ for some simple coroot $\check{\alpha}_i$. Moreover, a transverse slice through this stratum is isomorphic to a Kleinian singularity of the form A_{p_i-1} , where $p_i = \langle \check{\alpha}_i, \lambda \rangle$.*

Proof This follows immediately from the dimension formula $\dim(\mathrm{Gr}_{\mu}^{\bar{\lambda}}) = 2\langle \rho, \lambda - \mu \rangle$ and Theorem 5.2 of [Malkin et al. 2005]. $\square$

Once one has computed the codimensional 2 strata, in order to calculate the symplectic Galois group one must compute the monodromy action of the fundamental group of the strata on the exceptional fibres of minimal resolutions of the codimension two strata. However, it is asserted in [Kamnitzer et al. 2014] that these strata are simply connected, and hence the monodromy is trivial. For simplicity, in the remainder of this section we will fix $\mu = 0$.

Let $\{\check{\omega}_i\}_{i=1}^r$ be the set of fundamental coweights for G .

Proposition 7.7 *Let $Y = \mathrm{Gr}_0^{\bar{\lambda}}$, and assume that the conditions of Theorem 7.3 hold for $(\lambda, \mu) = (\lambda, 0)$ (and hence Y has a symplectic resolution). If $\lambda = \sum_{i=1}^r p_i \check{\omega}_i$ then its symplectic Galois group is isomorphic to $\prod_{p_i > 1} S_{p_i}$.*

Proof By Lemma 7.6, the slices through codimension two strata are Kleinian singularities of type A_{p_i-1} , and thus the associated reflection group is the symmetric group on p_i letters. To be more precise, given λ as in the statement of the proposition, we obtain a codimension-two stratum corresponding to $\nu = \lambda - \check{\alpha}_k$ provided ν is dominant. But if $C = (c_{ij})$ denotes the Cartan matrix of dual group $G^{\vee}$, then

$$\nu(\alpha_j) = p_j - c_{kj},$$

which is at least p_j if $j \neq k$, while if $j = k$ then $\nu(\alpha_k) = p_k - 2$ and hence $p_i \geq 2$. It follows that for each k , $1 \leq k \leq r$ we have a codimension 2 stratum in $\mathrm{Gr}_0^{\bar{\lambda}}$ provided $p_k > 1$, and in that case a slice through that stratum is a Kleinian singularity of type A_{p_k-1} . It follows that the symplectic Galois group is $\prod_{p_i > 1} S_{p_i}$ as required. $\square$

Remark 7.8 If we take $\mu \neq 0$ then the symplectic Galois group is in general smaller: one needs in addition, for each fundamental weight λ_i , that we have $\mu \leq \lambda - \alpha_i$.

Next we recall the geometric Satake correspondence, discovered by Lusztig, whose full categorical meaning was elucidated in the beautiful paper [Mirković and Vilonen 2007]. This establishes an equivalence of Tannakian categories between the representations of the Langlands dual group $G^\vee$ and the category of $G[[t]]$ -equivariant perverse sheaves on Gr . Under this equivalence, the convolution product of sheaves (recalled above for subsets of Gr) corresponds to tensor product of representations. As such, we have the following:

Proposition 7.9 *If $m : \mathrm{Gr}_0^{\bar{\lambda}} \rightarrow \mathrm{Gr}_0^{\bar{\lambda}}$ is a symplectic resolution, then the associated Springer sheaf $\mathrm{Spr}_{\bar{\lambda}}$ corresponds to the tensor product $V(\lambda_1) \otimes \dots \otimes V(\lambda_n)$ of representations of $G^\vee$. If i denotes the inclusion of our slice into the affine Grassmannian, it follows that*

$$\mathrm{Spr}_{\bar{\lambda}} \cong \bigoplus_v i^*(\mathrm{IC}(\mathbb{C}_v)) \otimes M_{\bar{\lambda}, v},$$

where $M_{\bar{\lambda}, v}$ is the multiplicity space, a vector space of dimension equal to the multiplicity of the irreducible V_v in the tensor product representation $V_{\lambda_1} \otimes \dots \otimes V_{\lambda_n}$.

Proof The irreducible representations $V(\lambda_i)$ correspond under the Geometric Satake Correspondence to the simple perverse sheaves IC_{λ_i} , the intersection cohomology extension of the constant local system $k_{\mathrm{Gr}^\lambda}[(2\rho, \lambda)]$. The tensor product of representations corresponds to the pushforward via $m_!$ of the twisted product

$$(7-1) \quad \mathrm{IC}_{\lambda_1} \tilde{\boxtimes} \dots \tilde{\boxtimes} \mathrm{IC}_{\lambda_n}$$

of the sheaves IC_{λ_i} . On the other hand, we have

$$\mathrm{Spr}_{\bar{\lambda}} = m_!(k_{\mathrm{Gr}_0^{\bar{\lambda}}}[(2\rho, \lambda)]).$$

However, it is known that the smooth locus of $\mathrm{Gr}^{\bar{\mu}} = \mathrm{Gr}^\mu$ for all μ and hence $\mathrm{Gr}_{\mu_0}^{\bar{\lambda}}$ is the smooth locus of $\mathrm{Gr}_0^{\bar{\lambda}}$. But by our hypotheses $\mathrm{Gr}_0^{\bar{\lambda}}$ is smooth, and hence it is exactly equal to the smooth locus of $\mathrm{Gr}_\mu^{\bar{\lambda}}$. It follows that on this locus the twisted product (7-1) is just the constant sheaf (with appropriate shift). The claim about the multiplicity spaces thus follows from the Satake equivalence. $\square$

Remark 7.10 These tensor product multiplicities are computable, via a number of combinatorial tools, eg crystal bases, Littelmann paths, etc. In order for $M_{\bar{\lambda}, v}$ to be an irreducible representation of the symplectic Galois group W for all the irreducible constituents of $\mathrm{Spr}_{\bar{\lambda}}$ one would need the group algebra $k[W]$ to be the full commutator algebra of the action of $\mathfrak{g}$ in $\mathrm{End}(V_{\lambda_1} \otimes \dots \otimes V_{\lambda_n})$. While this is true, famously, for the n -th tensor power of the vector representation of $\mathfrak{sl}_n$, where it is known as Schur–Weyl duality, in general one should expect kW to be too small to capture all of the commutant algebra. We will give an explicit example of this in Corollary 8.4.

7.2 Comparing the universal deformation and the BD deformation

The Beilinson–Drinfel’d deformation family arises by considering the affine Grassmannian spread over $\mathbb{P}^1$. We equip $\mathbb{P}^1$ with a $\mathbb{G}_m$ -action, fixing points $\{0, \infty\} \subseteq \mathbb{P}^1$ and pick a local coordinate $t \in \Gamma(\mathbb{P}^1, \mathcal{H}_{\mathbb{P}^1})$

at 0 having a pole at $\infty \in \mathbb{P}^1$. We will require our $\mathbb{G}_m$ -action to have weight 1 on our coordinate t . The basic observation is the following: let E^0 denote the trivial G -bundle on $\mathbb{P}^1$ and consider the set

$$\{(E, \phi) : E \text{ is a principal } G\text{-bundle on } \mathbb{P}^1 \text{ and } \phi : E_{|\mathbb{P}^1 \setminus \{0\}} \rightarrow E_{|\mathbb{P}^1 \setminus \{0\}}^0 \text{ is an isomorphism}\},$$

where E^0 denotes the trivial G -bundle. Using our local coordinate t at 0 one can readily obtain a bijection between this set and the affine Grassmannian Gr . Since any principal G -bundle on $\mathbb{P}^1 \setminus \{0\}$ is necessarily trivial, they can all be obtained by a modification of the trivial bundle at 0, and hence we see that Gr uniformises the moduli space of G -bundles on $\mathbb{P}^1$. Indeed observing that a trivialisation on $\mathbb{P}^1 \setminus \{0\}$ is just a G -valued function on $\mathbb{P}^1 \setminus \{0\}$, that is, an element of $G[t^{-1}]$, we see⁵ that

$$\text{Bun}_G(\mathbb{P}^1) \cong G[t^{-1}] \backslash \text{Gr}.$$

Definition 7.11 If (E, ϕ) is a pair consisting of a principal G -bundle on $\mathbb{P}^1$ together with a trivialisation $\phi : E_{|\mathbb{P}^1 \setminus \{a\}} \rightarrow E_{|\mathbb{P}^1 \setminus \{a\}}^0$ of E on $\mathbb{P}^1 \setminus \{a\}$ (thus, as before, we let E^0 denote the trivial bundle). The bundle E therefore corresponds to point on the affine Grassmannian $G((s))/G[[s]]$, where $s = t - a$, and we say E has *Hecke type* λ and write $\text{Inv}(\phi) = \lambda$ if this point lies in Gr^λ , and say E has Hecke type $\leq \lambda$ if $\text{Inv}(\phi) = \nu$ with $\text{Gr}^\nu \subseteq \text{Gr}^\lambda$.

The action of $G_1[[t^{-1}]]$ yield changes of trivialisation which preserve ∞ , and as a result one can describe the slices Gr_μ in terms of the isomorphism type of the bundle E and the parabolic reduction that it yields at ∞ . We will restrict our attention, however, to the case where $\mu = 0$, which corresponds in the moduli interpretation simply to the triviality of the bundle E . The importance of the moduli description is that it allows us to deform the affine Grassmannian over $\mathbb{A}^n$: we write $\text{Gr}_{\mathbb{A}^n}$ for the family over $\mathbb{A}^n$ whose fibre at $\underline{a} = (a_1, \dots, a_n)$ is just

$$\text{Gr}_{\underline{a}} =$$

$$\{(E, \phi) : E \text{ is a principal } G\text{-bundle on } \mathbb{P}^1 \text{ and } \phi : E_{|\mathbb{P}^1 \setminus \{a_1, \dots, a_n\}} \rightarrow E_{|\mathbb{P}^1 \setminus \{a_1, \dots, a_n\}}^0 \text{ is an isomorphism}\}.$$

Now fix a coweight λ with a decomposition $\lambda = \lambda_1 + \dots + \lambda_n$ as a sum of minuscule coweights. We think of the coweights λ_i as giving a “colouring” the divisor $\sum_{i=1}^n a_i$, that is, we set $D = \sum_{i=1}^n \lambda_i \cdot a_i$ and view it as a finite-supported coweight-valued function on $\mathbb{A}^1$: if $x \in \mathbb{A}^1$ then we set $D(x) = \sum_{i: a_i = x} \lambda_i$, and define

$$\text{Gr}_{0, \underline{a}}^\lambda = \{(E, \phi) \in \text{Gr}_{0, \underline{a}} : (E, \phi) \text{ has Hecke type } D(x) \text{ for all } x \in \mathbb{P}^1\},$$

$$\text{Gr}_{0, \underline{a}}^{\bar{\lambda}} = \{(E, \phi) \in \text{Gr}_{0, \underline{a}} : (E, \phi) \text{ has Hecke type } \leq D(x) \text{ for all } x \in \mathbb{P}^1\}.$$

Letting $\underline{a} = (a_1, \dots, a_n)$ vary over $\mathbb{A}^n$ we thus obtain a family of symplectic varieties $\text{Gr}_{0, \mathbb{A}^n}^\lambda$ (from a Poisson structure defined by a global version of the Manin triple used above) with central fibre $\text{Gr}_0^{\bar{\lambda}}$.

⁵Though clearly we have only argued at the level of $\mathbb{C}$ -points!

To deform our resolution $\text{Gr}_\mu^{\vec{\lambda}}$ of $\text{Gr}_0^{\vec{\lambda}}$ we use the twisted products of Gr^{λ_i} (which, since we are assuming each λ_i is minuscule, are smooth projective varieties). In more detail, we set

$$\begin{array}{ccc} \widetilde{\text{Gr}}_{0,\vec{a}}^{\vec{\lambda}} & \equiv \{ (E^1, E^2, \dots, E^n, \phi_1, \dots, \phi_n) \mid \phi_i : E_{|\mathbb{P}^1 \setminus \{a_i\}}^i \xrightarrow{\cong} E_{|\mathbb{P}^1 \setminus \{a_i\}}^{i-1}, \text{Inv}(\phi_i) \leq \lambda_i \} \\ \downarrow m_\lambda & & \downarrow \\ \text{Gr}_{0,\vec{a}}^{\vec{\lambda}} & & (E^n, \phi_1 \circ \dots \circ \phi_n) \end{array}$$

Here, by our convention that E^0 denotes the trivial bundle, ϕ_1 gives a trivialisation of E_1 over $\mathbb{P}^1 \setminus \{a_1\}$. Note that the composition $\phi_1 \circ \dots \circ \phi_n$ is defined on $\mathbb{P}^1 \setminus \{a_1, \dots, a_n\}$. Now again letting $\vec{a}$ vary we thus obtain a family $\widetilde{\text{Gr}}_{0,\mathbb{A}^n}^{\vec{\lambda}}$ which has $\text{Gr}_0^{\vec{\lambda}}$ as its central fibre. Let n_i be the multiplicity with which the fundamental weight $\tilde{\omega}_i$ occurs in the sequence $(\lambda_1, \lambda_2, \dots, \lambda_n)$, and let $\Sigma_{\vec{\lambda}} = \prod_{i=1}^\ell S_{n_i}$ (where ℓ is the rank of G).

Because we have coloured our divisor, our affine deformation no longer lives over $\text{Sym}^n \mathbb{A}^1$. Instead it lives over $\mathbb{A}^n / \Sigma_{\vec{\lambda}}$. We thus obtain a diagram

$$\begin{array}{ccc} \widetilde{\text{Gr}}_{0,\mathbb{A}^n}^{\vec{\lambda}} & \longrightarrow & \mathbb{A}^n \\ \pi_{\text{BD}} \downarrow & & \downarrow q_{\text{BD}} \\ \text{Gr}_{0,\mathbb{A}^n}^{\vec{\lambda}} & \longrightarrow & \mathbb{A}^n / \Sigma_{\vec{\lambda}} \end{array}$$

Note that since we are assuming the λ_i are all minuscule, the families $\text{Gr}_{0,\mathbb{A}^n}^{\vec{\lambda}}$ and $\text{Gr}_{0,\mathbb{A}^n}^{\bar{\lambda}}$ coincide. Recall that we gave our choice of coordinate t weight 1 for our $\mathbb{G}_m$ -action on $\mathbb{P}^1$. This action on $\mathbb{P}^1$ induces an action of $\mathbb{G}_m$ on the diagram above, and thus we have a deformation of a conical symplectic resolution over a linear base $\mathbb{A}^n$ (for the smooth deformation $\widetilde{\text{Gr}}_{0,\mathbb{A}^n}^{\vec{\lambda}}$).

Now as in Proposition 2.10, Namikawa has shown that his deformation spaces are universal for conic Poisson deformations; see Proposition 2.10 for a careful statement. By universality, there is a unique linear map $\alpha : \mathbb{A}^n \rightarrow H^2(\text{Gr}_0^{\vec{\lambda}})$ such that the families $\widetilde{\text{Gr}}_{0,\mathbb{A}^n}^{\vec{\lambda}} \rightarrow \mathbb{A}^n$ and $\text{Gr}_{0,\mathbb{A}^n}^{\vec{\lambda}} \rightarrow \mathbb{A}^n / \Sigma_{\vec{\lambda}}$ are obtained by pulling back the universal families.

Lemma 7.12 *The map $\alpha : \mathbb{A}^n \rightarrow H^2(\text{Gr}_0^{\vec{\lambda}}, \mathbb{C})$ is an isomorphism, and thus the symplectic Galois group action on $H^2(\text{Gr}_0^{\vec{\lambda}})$ is that of $\Sigma_{\vec{\lambda}}$.*

Proof First note that our $\mathbb{G}_m$ -action gives $\mathbb{A}^n$ the structure of a vector space, with 0-vector the fixed point of the action. Now the twisted products forming $\text{Gr}_0^{\vec{\lambda}}$ are in fact isomorphic to a product, and hence we have $\text{Pic}(\text{Gr}_0^{\vec{\lambda}}) = \prod_{i=1}^n \text{Pic}(\text{Gr}_0^{\lambda_i}) = \mathbb{Z}^n$, so that $\dim(H^2(\text{Gr}_0^{\vec{\lambda}})) = n$ (as the cohomology of a symplectic resolution is spanned by algebraic cycles). Since α is a linear map between two vector spaces of the same dimension, it follows that it is an isomorphism provided it is injective. But if $K = \ker(\alpha) \neq \{0\}$, then $(\widetilde{\text{Gr}}_{0,\mathbb{A}^n}^{\vec{\lambda}})_K = K \times Y$, and hence the $\mathbb{G}_m$ action must have weight 0 on K . But the $\mathbb{G}_m$ action has weight 1 on all of $\mathbb{A}^n$, hence α must be injective and hence an isomorphism as required. The final sentence of the statement of the lemma is immediate. $\square$

Corollary 7.13 *Every symplectic resolution of $\mathrm{Gr}_0^{\bar{\lambda}}$ is of the form $\mathrm{Gr}_0^{\bar{\lambda}}$ for some decomposition of λ .*

Proof This follows from Namikawa's results: any conic symplectic resolution $\pi : Y \rightarrow X$ allows one to build its associated universal deformation, and the universality implies that its different movable cones will exhaust all symplectic resolutions of X . $\square$

To identify $\Sigma_{\bar{\lambda}}$ with the symplectic Galois group we will use the following lemma, which was first observed in [Kamnitzer 2011, Proposition 2.8]. We outline a proof for the reader's convenience.

Lemma 7.14 *The monodromy action of Σ_{λ} on Spr corresponds, under the geometric Satake correspondence, to the action by permutation of tensor factors.*

Proof This is essentially a consequence of the fusion construction of the commutativity constraint in [Mirković and Vilonen 2007]. In particular, it is shown in Section 5 of that paper that the convolution product can be interpreted via the fusion product, and in the proof of Lemma 6.1 of Section 6 that the pushforward to $\mathbb{A}^1 \times \mathbb{A}^1$ of the fusion product sheaf has constant cohomologies. These facts imply the required compatibility. $\square$

It follows that we can reduce the calculation the action of the symplectic Galois group on the Springer sheaf $\mathcal{S}_{\bar{\lambda}}$ to the question of calculating the action of $\prod_{i \in I} S_{n_i}$ on the multiplicity spaces $M_{\bar{\lambda}, v}$ (in the notation of the preceding subsection). We use this in the next section to show that the multiplicity spaces in the Springer sheaf need not always be irreducible.

8 Tensor products of representations

The irreducible representations of $\mathfrak{sl}_{n+1}$, as for any semisimple Lie algebra, are indexed by highest weights. In the case of $\mathfrak{sl}_{n+1}$ however, these can naturally be identified with partitions $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_k)$, where $k \leq n$. Indeed, the root lattice of $\mathfrak{sl}_{n+1}$ is $Q = \{\alpha \in \mathbb{Z}^{n+1} : \sum_{i=1}^{n+1} \alpha_i = 0\}$, and hence the weight lattice P is naturally identified with $\mathbb{Z}^{n+1}/\mathbb{Z}\mathbf{1}$, where $\mathbf{1} = (1, 1, \dots, 1) \in \mathbb{Z}^{n+1}$. The set of dominant weights P^+ consists of the elements $\lambda \in P$ whose representatives $\tilde{\lambda} = (\tilde{\lambda}_1, \dots, \tilde{\lambda}_{n+1}) \in \mathbb{Z}^{n+1}$ satisfy $\tilde{\lambda}_i \geq \tilde{\lambda}_{i+1}$ for $1 \leq i \leq n$. Clearly given $\lambda \in P$ it has a unique representative $\tilde{\lambda}$ with $\tilde{\lambda}_{n+1} = 0$. Thus any dominant weight λ of $\mathfrak{sl}_{n+1}$ has a partition with at most n parts associated to it, and we will write this partition as $(\lambda_1, \dots, \lambda_n)$. By abuse of terminology, we will refer to the dominant weight and its associated partition interchangeably. Write $V(\lambda)$ for the irreducible highest-weight representation with highest weight λ , though if we specify λ as a partition, eg $\lambda = (2, 1^2)$, we will write $V(2, 1^2)$ rather than $V((2, 1^2))$.

Definition 8.1 A partition of a positive integer n is a weakly decreasing sequence $\lambda = (\lambda_1, \dots, \lambda_l)$ of nonnegative integers (so $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_l \geq 0$) with $\sum_{i=1}^l \lambda_i = n$. The *length* of λ is defined to be $\ell(\lambda) = \max\{k : \lambda_k > 0\}$. A *Young diagram of shape λ* is a left-justified arrangement of boxes with λ_i boxes in row i , where the rows are arranged in descending order. By convention we take $\lambda_k = 0$ if $k > l = \ell(\lambda)$.

Now fix n a positive integer. A *tableau* T is a Young diagram of length at most n , where each box is assigned an integer between 1 and n . It is *standard* if the rows and columns are strictly increasing, and *semistandard* if its columns are strictly increasing but its rows are only weakly increasing. The *weight* $w(T)$ of a semistandard (skew)-tableau T is the sequence $(\mu_1, \mu_2, \dots, \mu_n)$, where μ_i is the number of boxes in the tableau with entry i . We write $\mathcal{S}(\lambda)$ for the set of semistandard tableaux of shape λ (ie whose underlying Young diagram is λ). Given a (skew)-tableau T we will write $T_{\geq j}$ for the subtableau of T formed by the entries in columns $j, j+1, \dots$.

Thus, for example, the Young diagram below is the diagram associated to the partition $(2, 2, 1)$ of 5, and the tableau is semistandard with shape $(4, 3, 2, 1)$ and weight $(3, 3, 3, 1)$:

		1	1	1	2
		2	2	3	
		3	3		
		4			

Theorem 8.2 (Littlewood–Richardson rule [Stembridge 2002]) *If $\mathfrak{g} = \mathfrak{sl}_n$ and λ, μ, ν are dominant weights for $\mathfrak{g}$, then the simple representation $V(\nu)$ occurs in the tensor product $V(\lambda) \otimes V(\mu)$ with multiplicity $C_{\lambda, \mu}^{\nu}$, where $C_{\lambda, \mu}^{\nu}$ is the number of tableaux T of shape μ which satisfy*

- (i) $\lambda + w(T_{\geq j}) \in P^+$ for each $j \geq 1$, and
- (ii) $\lambda + w(T) = \nu$.

The fundamental representations of $\mathfrak{sl}_{n+1}$ are the exterior powers of V , the vector representation. They are the representations associated to the partitions (1^k) for $k = 1, \dots, n$. They are minuscule representations (that is, their weights form a single Weyl-group orbit).

Corollary 8.3 *Let $\omega_k = (1^k)$.*

- (1) *We have $c_{\mu, \omega_k}^{\lambda} \in \{0, 1\}$.*
- (2) *If $n \geq 9$ and $\lambda = \omega_3 = (1^3)$, then*


$$V(1^3) \otimes V(1^3) = V(1^6) \oplus V(2, 1^5) \oplus V(2^2, 1^4) \oplus V(2^3),$$

and $V(1^3)^{\otimes 3}$ contains $V(2^3, 1^3)$ with multiplicity 4.

Proof For part (1), note that since $\omega_k = (1^k)$ has only 1 column, Stembridge’s version of the Littlewood–Richardson rule shows that the multiplicity of $V(\lambda)$ in $V(\mu) \otimes V(\omega_k)$ is just the number of semistandard tableaux of shape 1^k and weight $\nu - \lambda$. But since (1^k) has only one column, the weight of the semistandard tableau uniquely determines the tableau.

For part (2), to calculate $V(1^3) \otimes V(1^3)$ we must find all (semi)standard tableaux T of shape (1^3) such that $(1^3) + w(T)$ is a partition. Now if $(1^3) + w(T) = (a_1, a_2, \dots, a_9)$, each $a_i \in \{0, 1, 2\}$. Let $m = \max\{a_i : 0 \leq i \leq 9\}$. If $m = 2$, then we must have $w_1 = 1$ and $w_2 \in \{2, 4\}$ and $w_3 \in \{3, 4\}$ or $\{5\}$, respectively, while if $m = 1$ then $(w_1, w_2, w_3) = (4, 5, 6)$.

Finally, it is easy to see that if we tensor with $V(1^3)^{\otimes 2}$ with $V(1^3)$, and let $\nu = (2^3, 1^3)$, then for each of the 4 summands $V(\mu)$ of $V(1^3)^{\otimes 2}$ the difference $\nu - \mu$ is a partition:



Hence $V(2^3, 1^3)$ occurs in $V(1^3)^{\otimes 3}$ with multiplicity 4 as required. $\square$

Recall from Definition 3.5 that we say a local system $\mathcal{E}$ on a symplectic leaf is a Springer local system if $\mathrm{IC}(\mathcal{E})$ occurs as a constituent of the Springer sheaf.

Corollary 8.4 *Let $\lambda = 3\varpi_3$ and $\mu = 0$, and let $\vec{\lambda} = (\varpi_3, \varpi_3, \varpi_3)$. We may decompose the Springer sheaf $\mathcal{S}$ for $\pi : \mathrm{Gr}_0^{\vec{\lambda}} \rightarrow \mathrm{Gr}_0^{\vec{\lambda}}$ as*

$$\mathcal{S} \cong \bigoplus_{\nu \leq \lambda} \mathrm{IC}(\mathrm{Gr}_0^{\nu}) \otimes \mathcal{M}_{\nu},$$

where each $\mathcal{M}_{\nu}$ carries a natural action of $W = S_3$ the symplectic Galois group. When $\nu = (2^3, 1^3)$, this representation is reducible. Thus the symplectic Springer correspondence does not necessarily yield a function from Springer local systems to irreducible representations of the symplectic Galois group.

Proof By Proposition 7.9, the dimension of the multiplicity space M_{ν} is the multiplicity of $V(2^3, 1^3)$ in $V(1^3)^{\otimes 3}$, and Corollary 8.3 shows that this multiplicity is 4. It follows that M_{ν} cannot be an irreducible representation of S_3 . $\square$

Acknowledgements

We are grateful to many people for helpful conversations and comments, in particular Gwyn Bellamy, Tom Braden, Victor Ginzburg, Mark Goresky, Dmitry Kaledin, Joel Kamnitzer and Oded Yacobi. Both authors are grateful to MSRI for excellent working conditions during the initial work on this project. McGerty was supported by a Royal Society University Research Fellowship and the EPSRC Programme Grant on motivic invariants and categorification, and by a Fisher Visiting Professorship at the University of Illinois at Urbana-Champaign. Nevins was supported by NSF grants DMS-1159468, DMS-1502125 and DMS-1802094 and NSA grant H98230-12-1-0216, and by an All Souls Visiting Fellowship. Both authors were supported by MSRI.

References

- [Beauville 2000] **A Beauville**, *Symplectic singularities*, Invent. Math. 139:3 (2000) 541–549 MR
- [Bellamy and Schedler 2021] **G Bellamy, T Schedler**, *Symplectic resolutions of quiver varieties*, Selecta Math. 27:3 (2021) art. id. 36 MR

- [Bernstein and Lunts 1994] **J Bernstein, V Lunts**, *Equivariant sheaves and functors*, Lecture Notes in Mathematics 1578, Springer (1994) MR
- [Braden et al. 2016] **T Braden, A Licata, N Proudfoot, B Webster**, *Quantizations of conical symplectic resolutions, II: Category $\mathcal{O}$ and symplectic duality*, Astérisque 384, Soc. Math. France, Paris (2016) 75–179 MR
- [Braverman et al. 2019] **A Braverman, M Finkelberg, H Nakajima**, *Coulomb branches of $3d \mathcal{N} = 4$ quiver gauge theories and slices in the affine Grassmannian*, Adv. Theor. Math. Phys. 23:1 (2019) 75–166 MR
- [Brion 2017] **M Brion**, *Algebraic group actions on normal varieties*, Trans. Moscow Math. Soc. 78 (2017) 85–107 MR
- [de Cataldo and Migliorini 2002] **M A A de Cataldo, L Migliorini**, *The hard Lefschetz theorem and the topology of semismall maps*, Ann. Sci. École Norm. Sup. 35:5 (2002) 759–772 MR
- [Chriss and Ginzburg 1997] **N Chriss, V Ginzburg**, *Representation theory and complex geometry*, Birkhäuser, Boston (1997) MR
- [Crawley-Boevey 2003] **W Crawley-Boevey**, *Normality of Marsden–Weinstein reductions for representations of quivers*, Math. Ann. 325:1 (2003) 55–79 MR
- [Dobrovolska et al. 2016] **G Dobrovolska, V Ginzburg, R Travkin**, *Moduli spaces, indecomposable objects and potentials over a finite field*, preprint (2016) arXiv 1612.01733
- [Fu 2003] **B Fu**, *Symplectic resolutions for nilpotent orbits*, Invent. Math. 151:1 (2003) 167–186 MR
- [Ginzburg and Kaledin 2004] **V Ginzburg, D Kaledin**, *Poisson deformations of symplectic quotient singularities*, Adv. Math. 186:1 (2004) 1–57 MR
- [Goresky and MacPherson 1988] **M Goresky, R MacPherson**, *Stratified Morse theory*, Ergebnisse der Math. (3) 14, Springer (1988) MR
- [Göttsche and Soergel 1993] **L Göttsche, W Soergel**, *Perverse sheaves and the cohomology of Hilbert schemes of smooth algebraic surfaces*, Math. Ann. 296:2 (1993) 235–245 MR
- [Grinberg 1998] **M Grinberg**, *A generalization of Springer theory using nearby cycles*, Represent. Theory 2 (1998) 410–431 MR
- [Humphreys 1972] **J E Humphreys**, *Introduction to Lie algebras and representation theory*, Graduate Texts in Mathematics 9, Springer (1972) MR
- [Kaledin 2006] **D Kaledin**, *Symplectic singularities from the Poisson point of view*, J. Reine Angew. Math. 600 (2006) 135–156 MR
- [Kaledin and Verbitsky 2002] **D Kaledin, M Verbitsky**, *Period map for non-compact holomorphically symplectic manifolds*, Geom. Funct. Anal. 12:6 (2002) 1265–1295 MR
- [Kamnitzer 2011] **J Kamnitzer**, *The Beilinson–Drinfeld Grassmannian and symplectic knot homology*, from “Grassmannians, moduli spaces and vector bundles” (D A Ellwood, E Previato, editors), Clay Math. Proc. 14, Amer. Math. Soc., Providence, RI (2011) 81–94 MR
- [Kamnitzer et al. 2014] **J Kamnitzer, B Webster, A Weekes, O Yacobi**, *Yangians and quantizations of slices in the affine Grassmannian*, Algebra Number Theory 8:4 (2014) 857–893 MR
- [Lê Dũng Tráng 1977] **Lê Dũng Tráng**, *Some remarks on relative monodromy*, from “Real and complex singularities” (Oslo, 1976) (P Holm, editor), Sijthoff & Noordhoff, Alphen aan den Rijn (1977) 397–403 MR
- [Lehn et al. 2012] **M Lehn, Y Namikawa, C Sorger**, *Slodowy slices and universal Poisson deformations*, Compos. Math. 148:1 (2012) 121–144 MR
- [Losev 2012] **I Losev**, *Isomorphisms of quantizations via quantization of resolutions*, Adv. Math. 231:3-4 (2012) 1216–1270 MR
- [Losev 2017] **I Losev**, *On categories $\mathcal{O}$ for quantized symplectic resolutions*, Compos. Math. 153:12 (2017) 2445–2481 MR
- [Lu and Yakimov 2008] **J-H Lu, M Yakimov**, *Group orbits and regular partitions of Poisson manifolds*, Comm. Math. Phys. 283:3 (2008) 729–748 MR
- [Luna 1973] **D Luna**, *Slices étales*, from “Sur les groupes algébriques”, Mém. Soc. Math. France 33, Soc. Math. France, Paris (1973) 81–105 MR

- [Lusztig 1984] **G Lusztig**, *Intersection cohomology complexes on a reductive group*, Invent. Math. 75:2 (1984) 205–272 MR
- [Lusztig 2000] **G Lusztig**, *Quiver varieties and Weyl group actions*, Ann. Inst. Fourier (Grenoble) 50:2 (2000) 461–489 MR
- [Maffei 2002] **A Maffei**, *A remark on quiver varieties and Weyl groups*, Ann. Sc. Norm. Super. Pisa Cl. Sci. 1:3 (2002) 649–686 MR
- [Malkin et al. 2005] **A Malkin**, **V Ostrik**, **M Vybornov**, *The minimal degeneration singularities in the affine Grassmannians*, Duke Math. J. 126:2 (2005) 233–249 MR
- [Markman 2010] **E Markman**, *Modular Galois covers associated to symplectic resolutions of singularities*, J. Reine Angew. Math. 644 (2010) 189–220 MR
- [Martino 2008] **M Martino**, *Stratifications of Marsden–Weinstein reductions for representations of quivers and deformations of symplectic quotient singularities*, Math. Z. 258:1 (2008) 1–28 MR
- [Mather 2012] **J Mather**, *Notes on topological stability*, Bull. Amer. Math. Soc. 49:4 (2012) 475–506 MR
- [McGerty and Nevins 2018] **K McGerty**, **T Nevins**, *Kirwan surjectivity for quiver varieties*, Invent. Math. 212:1 (2018) 161–187 MR
- [Milnor 1963] **J Milnor**, *Morse theory*, Annals of Mathematics Studies 51, Princeton Univ. Press (1963) MR
- [Mirković and Vilonen 2007] **I Mirković**, **K Vilonen**, *Geometric Langlands duality and representations of algebraic groups over commutative rings*, Ann. of Math. 166:1 (2007) 95–143 MR
- [Nakajima 1998] **H Nakajima**, *Quiver varieties and Kac–Moody algebras*, Duke Math. J. 91:3 (1998) 515–560 MR
- [Nakajima 2001] **H Nakajima**, *Quiver varieties and finite-dimensional representations of quantum affine algebras*, J. Amer. Math. Soc. 14:1 (2001) 145–238 MR
- [Nakajima 2003] **H Nakajima**, *Reflection functors for quiver varieties and Weyl group actions*, Math. Ann. 327:4 (2003) 671–721 MR
- [Namikawa 2008] **Y Namikawa**, *Flops and Poisson deformations of symplectic varieties*, Publ. Res. Inst. Math. Sci. 44:2 (2008) 259–314 MR
- [Namikawa 2010] **Y Namikawa**, *Poisson deformations of affine symplectic varieties, II*, Kyoto J. Math. 50:4 (2010) 727–752 MR
- [Namikawa 2011] **Y Namikawa**, *Poisson deformations of affine symplectic varieties*, Duke Math. J. 156:1 (2011) 51–85 MR
- [Namikawa 2015] **Y Namikawa**, *Poisson deformations and birational geometry*, J. Math. Sci. Univ. Tokyo 22:1 (2015) 339–359 MR
- [Ngô 1999] **B C Ngô**, *Faisceaux pervers, homomorphisme de changement de base et lemme fondamental de Jacquet et Ye*, Ann. Sci. École Norm. Sup. 32:5 (1999) 619–679 MR
- [Quillen 1971] **D Quillen**, *The spectrum of an equivariant cohomology ring, I*, Ann. of Math. 94 (1971) 549–572 MR
- [Sabbah 1983] **C Sabbah**, *Morphismes analytiques stratifiés sans éclatement et cycles évanescents*, from “Analysis and topology on singular spaces, II, III” (Luminy, 1981), Astérisque 101–102, Soc. Math. France, Paris (1983) 286–319 MR
- [Sjamaar and Lerman 1991] **R Sjamaar**, **E Lerman**, *Stratified symplectic spaces and reduction*, Ann. of Math. 134:2 (1991) 375–422 MR
- [Slodowy 1980] **P Slodowy**, *Four lectures on simple groups and singularities*, Communications of the Mathematical Institute, Rijksuniversiteit Utrecht 11, Rijksuniversiteit Utrecht, Mathematical Institute, Utrecht (1980) MR
- [Springer 1976] **T A Springer**, *Trigonometric sums, Green functions of finite groups and representations of Weyl groups*, Invent. Math. 36 (1976) 173–207 MR
- [Stembridge 2002] **J R Stembridge**, *A concise proof of the Littlewood–Richardson rule*, Electron. J. Combin. 9:1 (2002) Note 5, 4 MR
- [Treumann 2009] **D Treumann**, *A topological approach to induction theorems in Springer theory*, Represent. Theory 13 (2009) 8–18 MR
- [Zhu 2017] **X Zhu**, *An introduction to affine Grassmannians and the geometric Satake equivalence*, from “Geometry of moduli spaces and representation theory” (R Bezrukavnikov, A Braverman, Z Yun, editors), IAS/Park City Math. Ser. 24, Amer. Math. Soc., Providence, RI (2017) 59–154 MR

KEVIN MCGERTY mcgerty@maths.ox.ac.uk

Mathematical Institute, University of Oxford, Oxford, United Kingdom

<https://www.maths.ox.ac.uk/people/kevin.mcgerty>

THOMAS NEVINS

Department of Mathematics, University of Illinois at Urbana-Champaign, Urbana, IL, United States

Proposed: András I Stipsicz

Received: June 18, 2019

Seconded: Dan Abramovich, David Fisher

Revised: March 27, 2024

Unipotent morphisms

DANIEL BRAGG, JACK HALL AND SIDDHARTH MATHUR

We introduce the theory of unipotent morphisms of algebraic stacks and prove a surprising local to global principle for a class of vector bundles. Two sample applications of our methods are a unipotent analogue of Gabber’s theorem for torsion $\mathbf{G}_m$ -gerbes, and that smooth Deligne–Mumford stacks with quasiprojective coarse spaces satisfy the resolution property in positive characteristic. Our main tool is a descent result for flags, which we prove using results of Schäppi.

1 Introduction

In the early 1990s, Gabber answered a question of Grothendieck for quascompact schemes admitting an ample line bundle: every cohomological Brauer class on such a scheme is represented by an Azumaya algebra; see Gabber [14] and de Jong [22]. This result has since been reinterpreted in terms of the *resolution property* for algebraic stacks (i.e., every quasicoherent sheaf is a quotient of a direct sum of vector bundles) and gerbes; see Edidin, Hassett, Kresch and Vistoli [12], Gross [17], Kresch and Vistoli [24], Mathur [27] and Totaro [45]. In this language, Gabber’s result can be restated as follows:

Theorem (Gabber) *Let X be a quascompact scheme that admits an ample line bundle. Let $\mathcal{X} \rightarrow X$ be a $\mathbf{G}_m$ -gerbe. If the cohomology class $[\mathcal{X}] \in H^2(X, \mathbf{G}_m)$ of $\mathcal{X}$ is torsion, then $\mathcal{X}$ is a global quotient stack (see Definition 2.1) and has the resolution property.*

Using quasiprojective methods and Gabber’s theorem, Kresch and Vistoli show that smooth, separated and generically tame Deligne–Mumford stacks with quasiprojective coarse moduli spaces have the resolution property [24]. It has long been a challenge to produce nontrivial vector bundles without such hypotheses.

Here we introduce new methods to construct nontrivial vector bundles on schemes, algebraic spaces and algebraic stacks. Our key idea is to leverage the abundance of cohomology in unipotent settings. A sample application of our results is the following additive analogue of Gabber’s result (also see Theorem 6.3):

Theorem A *Let X be a quascompact algebraic stack with affine diagonal and the resolution property. If $\mathcal{G} \rightarrow X$ is a $\mathbf{G}_a$ -gerbe, then $\mathcal{G}$ has the resolution property.*

In particular, this implies that $\mathbf{G}_a$ -gerbes over smooth separated schemes have the resolution property. Note that the analogous multiplicative statement is still unknown, even in dimension 3. We can also extend Kresch and Vistoli’s result to the generically wild setting:

Theorem B *Let $\mathcal{X}$ be a smooth and separated Deligne–Mumford stack of finite type over a field k . Let $\pi: \mathcal{X} \rightarrow X$ be the associated coarse moduli space. If X is quasiprojective over $\mathrm{Spec} k$, then $\mathcal{X}$ has the resolution property.*

MSC2020: primary 14A20; secondary 14G17.

© 2026 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

We expect the key new technical input for the above results to be of general interest: it states that one may descend certain vector bundles along flat morphisms after locally taking a direct sum (see Theorem 4.8 for a more general statement):

Theorem C *Let $\mathcal{X} \rightarrow X$ be a morphism of quasicompact and quasiseparated algebraic stacks. Suppose that X has affine diagonal and satisfies the resolution property. If there is a faithfully flat morphism $X' \rightarrow X$, where X' is affine, and a vector bundle V' on $\mathcal{X}' = \mathcal{X} \times_X X'$ that admits a filtration*

$$0 = V'_0 \subseteq V'_1 \subseteq \cdots \subseteq V'_n = V'$$

*with $V'_i/V'_{i-1} \simeq \mathcal{O}_{\mathcal{X}'}$ for all i (i.e., V' is **trivially graded**), then there is a vector bundle V on $\mathcal{X}$ and a split surjection $V_{\mathcal{X}'} \rightarrow V'$.*

In fact, our results apply in the broader setting of *unipotent morphisms*, a notion that we develop and characterize in Section 5. Roughly speaking, the theory is modeled on gerbes with unipotent stabilizers and algebraic stacks of the form $[Z/U_n]$, where Z is an algebraic space and U_n is the unipotent subgroup of GL_n consisting of unitriangular matrices. Compare with [12; 17; 45], where quotients of the form $[Z/\mathrm{GL}_n]$ and the resolution property are considered. In fact, we present a unipotent enrichment of their results (see Theorem 5.4) and, better still, we prove a striking descent statement. Indeed, Theorem C reveals the following local to global principle: *a locally unipotent morphism over a base with enough flags is globally unipotent* (see Theorem 5.5).

These notions are closely related to many foundational questions of independent interest. For example, over a field k a *unipotent group scheme* is a closed subgroup of $U_{n,k}$ (see Demazure and Grothendieck [10, Theorem XVII.3.5] and Remark 5.14 for several other characterizations). It is natural to ask: to what extent is this true over a general base?

Question 1.1 Let $G \rightarrow S$ be a group scheme with unipotent geometric fibers. Is there an embedding $G \rightarrow U_{n,S}$ for some $n \geq 1$? Is this true locally on S ?

As a consequence of our methods, if S is a regular separated scheme or admits an ample line bundle, then the existence of a flat-local embedding in U_n implies the existence of a Zariski-local embedding of G in U_n (see Corollary 5.13 for a more general statement). In Section 5.4, we refine Question 1.1 and explain some cases where the answer is positive (see Examples 5.22 and 5.23) and one case, in positive characteristic, where it is negative (Example 5.21(1)). One may also compare Question 1.1 with a question of Conrad on when smooth affine group schemes can be embedded into GL_n ; see Conrad [7].

Unipotent morphisms are closely related to the notion of a faithful moduli space, due to Alper (personal communication, 2017). This is a morphism of algebraic stacks $f: \mathcal{X} \rightarrow X$ such that

- (1) the natural map $\mathcal{O}_X \rightarrow f_*\mathcal{O}_{\mathcal{X}}$ is an isomorphism (i.e., f is Stein), and
- (2) if F is quasicoherent on $\mathcal{X}$ and $f_*F = 0$, then $F = 0$.

Faithful moduli spaces should help with the analysis of moduli stacks with only unipotent stabilizers (compare with good or adequate moduli spaces of Alper [2; 3]) that appear in the unstable locus of GIT

quotients. In contrast to good moduli spaces, however, faithful moduli spaces are rarely compatible with base change. This makes them difficult to reduce to simple cases where we can apply deformation theory. Faithful moduli spaces are also stable under composition.

We show that faithful moduli spaces and locally unipotent morphisms have unipotent stabilizers (Propositions 5.15 and 7.10) and that faithful moduli spaces are generically gerbes for a unipotent group (Proposition 7.12). We also show that certain Stein unipotent morphisms are faithful moduli spaces (Proposition 7.1).

2 Vector bundles, global quotients, and the resolution property

We briefly recall a circle of ideas from [12; 17; 45]. We also make some minor refinements, which will be useful. To agree with the conventions of [17, Section 5], a *vector bundle* on X will mean a locally free $\mathcal{O}_X$ -module of finite constant rank.

Let X be an algebraic stack and let n be a nonnegative integer. If V is a vector bundle on X of rank n , there is an associated *frame bundle* $\mathrm{Fr}(V) = \underline{\mathrm{Isom}}_X(\mathcal{O}_X^{\oplus n}, V)$. Then $\mathrm{Fr}(V) \rightarrow X$ is a GL_n -torsor and sits in the following 2-cartesian diagram:

$$\begin{array}{ccc} \mathrm{Fr}(V) & \longrightarrow & \mathrm{Spec} \mathbf{Z} \\ \downarrow & & \downarrow \\ X & \longrightarrow & B\mathrm{GL}_n \end{array}$$

It is well known that this association induces an equivalence of categories between vector bundles on X of rank n and GL_n -torsors.

Definition 2.1 Let X be an algebraic stack. We say that X is a *global quotient* if $X \cong [Z/\mathrm{GL}_n]$, where Z is a quasicompact and quasiseparated algebraic space with an action of GL_n for some nonnegative integer n . If V denotes the vector bundle of rank n associated to the GL_n -torsor $Z \rightarrow X$, we say that V is *faithful*. A morphism of algebraic stacks $f: X \rightarrow S$ is a *global quotient* if there is a vector bundle V of rank n on X with frame bundle that is quasicompact, quasiseparated and representable over S . In this case, we say that V is *f -faithful*.

Remark 2.2 A vector bundle V on X is easily seen to be f -faithful for $f: X \rightarrow S$ if and only if the relative stabilizer groups of f act faithfully on V at every geometric point of X . In particular, the relative stabilizer groups are affine [8, Corollary VI_B.1.4.2]. Moreover, if $X \xrightarrow{f} S \xrightarrow{g} T$ are quasicompact and quasiseparated morphisms of algebraic stacks and V is $(g \circ f)$ -faithful, then V is f -faithful. Also, X is a global quotient if and only if the morphism $X \rightarrow \mathrm{Spec} \mathbf{Z}$ is a global quotient.

Remark 2.3 It is natural to ask if algebraic stacks such as $X_1 = \coprod_{n \geq 1} B\mathrm{GL}_{n,\mathbf{Z}}$ or $X_2 = B\mathbf{Z}$, where $\mathbf{Z}$ denotes the constant group scheme associated to the integers, should be considered global quotients. There is certainly a locally free sheaf on X_1 with a faithful action, but it has nonconstant and unbounded rank; also, X_2 admits a vector bundle of rank 2 with a faithful action whose frame is quasicompact but

not quasiseparated. We exclude these as we are unable to shed much light on them and also to adhere to the conventions of [17, Section 5].

Let $\mathcal{A}$ be an abelian category. Let Λ be a set of objects of $\mathcal{A}$. Recall that Λ *generates* $\mathcal{A}$ (or Λ is *generating*) if the functor

$$H_\Lambda: \mathcal{A} \rightarrow \mathbf{Ab}: x \mapsto \prod_{a \in \Lambda} \operatorname{Hom}_{\mathcal{A}}(a, x)$$

is faithful. If $\mathcal{A}$ is closed under small coproducts, then Λ is generating if and only if every $x \in \mathcal{A}$ admits a presentation of the form $\bigoplus_{a \in \Lambda} (a^{\oplus I_a}) \twoheadrightarrow x$.

Remark 2.4 Let $\mathcal{A}$ be an abelian category and Λ a set of objects of $\mathcal{A}$. Then H_Λ is *conservative* (i.e., if $f: x \rightarrow x'$ in $\mathcal{A}$ and $H_\Lambda(f)$ is bijective, then f is an isomorphism) if and only if Λ generates $\mathcal{A}$. Note that while H_Λ conservative implies that it is *zero-reflecting* (i.e., $H_\Lambda(x) = 0$ implies that $x = 0$), the converse does not hold in general — it does when the $a \in \Lambda$ are all $\mathcal{A}$ -projective. For a counterexample, consider the affine line with the doubled origin; then $H_{\{\mathcal{O}\}}$ is zero-reflecting but not conservative.

Example 2.5 If X is an algebraic stack, then the abelian category $\mathbf{QCoh}(X)$ always admits a set of generators [43, Tag 0781]. If X is quasicompact and quasiseparated, then a set of generators can always be found amongst the $\mathcal{O}_X$ -modules of finite presentation [39]. Note this result is much simpler if X is noetherian [25, Proposition 15.4] or has quasifinite and separated diagonal [37].

The *resolution property*, which we refine and recall below, is about algebraic stacks that have a generating set of vector bundles.

Definition 2.6 Let X be an algebraic stack. Let $\mathbf{V}$ be a set of vector bundles on X . We say that X has the *$\mathbf{V}$ -resolution property* if it is quasicompact, quasiseparated and $\mathbf{V}$ generates $\mathbf{QCoh}(X)$. Following [17, Definition 5.1], we say X has the *resolution property* if it has the $\mathbf{V}$ -resolution property for some set $\mathbf{V}$.

Remark 2.7 Fix a noetherian algebraic stack X and let $\mathbf{V}$ denote the set of isomorphism classes of vector bundles on X . Then X has the $\mathbf{V}$ -resolution property if and only if every coherent sheaf on X is the quotient of a locally free sheaf of finite rank; see [17, Remark 5.2].

Just as in [17], we will be interested in a relative version of the resolution property, so we recall the following definition:

Definition 2.8 [17, Definition 2.7] Let $f: X \rightarrow S$ be a morphism of algebraic stacks. Let $\mathbf{V}$ be a set of finitely presented quasicoherent sheaves on X . We say that $\mathbf{V}$ is *f -generating* if f is quasicompact and quasiseparated and there is a set of (not necessarily finitely presented) quasicoherent $\mathcal{O}_S$ -modules $\mathbf{W}$ such that $\mathbf{V} \otimes_{\mathcal{O}_X} f^* \mathbf{W} = \{V \otimes_{\mathcal{O}_X} f^* W : V \in \mathbf{V}, W \in \mathbf{W}\}$ generates $\mathbf{QCoh}(X)$. We say that $\mathbf{V}$ is *universally f -generating* if for every morphism of algebraic stacks $s: S' \rightarrow S$, the set $s'^* \mathbf{V} = \{s'^* V : V \in \mathbf{V}\}$, where $s': X' = X \times_S S' \rightarrow X$ and $f': X' \rightarrow S'$ denote the projections, is f' -generating.

Remark 2.9 Note that $\mathbf{V}$ is universally f -generating if and only if it is so on an fpqc covering of the target; see [17, Proposition 2.8 (iii)]. Also, if $\mathbf{V}$ is f -generating, then it remains so after quasiffine base change

(combine [17, Proposition 2.8 (ii), (v) and (vi)]). In particular, if S has quasiaffine diagonal, then $\mathbf{V}$ is f -generating if and only if it is universally f -generating; see [17, Corollary 2.10]. Ignoring set-theoretic issues, Example 2.5 implies that we can take $\mathbf{W} = \mathbf{QCoh}(S)$ in the definition of f -generating.

Definition 2.10 Let $f: X \rightarrow S$ be a morphism of algebraic stacks. If $\mathbf{V}$ is a set of vector bundles on X , we say f has the $\mathbf{V}$ -resolution property if $\mathbf{V}$ is universally f -generating. Following [17, Definition 5.1], we say that f has the resolution property if f has the $\mathbf{V}$ -resolution property for some $\mathbf{V}$. If such a $\mathbf{V}$ arises as a set of submodules of the polynomials in a fixed vector bundle V and its dual that are split by restriction to the frame bundle of V , then we say that V is an f -tensor generator [17, Definition 6.1]. Such a V is f - R -faithful if it is also f -faithful (note that f -tensor-generators are automatically f -faithful if f has relatively affine stabilizers by [17, Theorem 6.4]).

Remark 2.11 An algebraic stack X has the $\mathbf{V}$ -resolution property if and only if the morphism $X \rightarrow \mathrm{Spec} \mathbf{Z}$ has the $\mathbf{V}$ -resolution property. In general, if Y has the $\mathbf{V}'$ -resolution property and $f: X \rightarrow Y$ is a morphism of algebraic stacks with the $\mathbf{V}$ -resolution property, then X has the $(\mathbf{V} \otimes_{\mathcal{O}_X} f^* \mathbf{V}')$ -resolution property [17, Proposition 2.8(v)].

Example 2.12 Let A be an abelian scheme of positive dimension over a field k . If $g: \mathrm{Spec} k \rightarrow BA$ denotes the standard covering, then $\mathcal{O}_{\mathrm{Spec} k}$ is g -generating but not universally g -generating. Also, g^* induces an equivalence $\mathbf{QCoh}(BA) \simeq \mathbf{QCoh}(\mathrm{Spec} k)$, so $f: BA \rightarrow \mathrm{Spec} k$ has the resolution property with $\mathcal{O}_{BA}$ as a f -tensor-generator. However, BA admits no f -faithful vector bundles.

Example 2.13 If X admits an ample family $L_1, \dots, L_r$ of line bundles [43, Tag 0FXR] (e.g., X is quasiprojective over a field or is noetherian, normal and $\mathbf{Q}$ -factorial with affine diagonal), then it satisfies the $\mathbf{V}$ -resolution property [17, Example 5.9(i)], where $\mathbf{V} = \{L_i^m : m \leq 0 \text{ for } i = 1, \dots, r\}$. In dimensions 2 and above, there are proper schemes with no nontrivial line bundles. However, 2-dimensional separated algebraic spaces always have enough vector bundles; see [42, Theorem 2.1; 16, Theorem 5.2; 27, Theorem 41]. Nothing is known for smooth separated algebraic spaces or normal separated schemes over a field in dimension ≥ 3 , except that they enjoy the resolution property after removing a closed subset of codimension ≥ 3 [29].

Example 2.14 For algebraic stacks the situation is more complicated. Kresch and Vistoli showed that a Deligne–Mumford stack which is smooth, generically tame and separated over a field satisfies the resolution property if its coarse moduli space is quasiprojective [24, Theorem 1.3]. The case of 2-dimensional normal tame algebraic stacks with finite diagonal over a field was settled in [27, Theorem 1]. There are also nontorsion (and hence, nonregular) $\mathbf{G}_m$ -gerbes which do not have the resolution property [18, Remark 1.11b]. As far as the authors are aware, there is no known example of a separated or smooth algebraic stack with affine diagonal which does not have the resolution property.

Surprisingly, Totaro and Gross (see [45, Theorem 1.1; 17, Theorems 1.1, and 6.10]) were able to relate quotient stacks and the resolution property. We have the following refinement, whose proof is the same as Gross', but we just keep track of the generating set:

Theorem 2.15 Let $f: X \rightarrow S$ be a morphism of quasicompact and quasiseparated algebraic stacks. Assume that f has affine stabilizers.

- (1) Let $\mathbf{V}$ be a set of vector bundles on X that is closed under direct sums. If $\mathbf{V}$ is universally f -generating, then there is a vector bundle $V \in \mathbf{V}$ of rank n on X whose frame bundle is quasiffine over S .
- (2) Let V be a vector bundle on X . Then V has frame bundle quasiffine over S if and only if V is an f -tensor generator.

In particular, a quasicompact and quasiseparated algebraic stack with affine stabilizers at closed points has the resolution property if and only if it can be written as $[U/\mathrm{GL}_n]$, where U is a quasiffine scheme.

Proof Claim (2) is [17, Theorem 6.4]. For (1), consider the natural inverse system of X -stacks formed by

$$F_J: \prod_{W \in J} \mathrm{Fr}(W) \rightarrow X,$$

where $J \subset \mathbf{V}$ is a finite subset. Since each frame bundle is affine over X , the transition maps in the inverse system are affine. Hence the inverse limit is an algebraic stack $p: F \rightarrow X$, which is affine over X [37, Theorem C(i)]. Since $\mathbf{V}$ is generating for f and every vector bundle in $\mathbf{V}$ becomes trivial when restricted to F , it follows that $\mathcal{O}_F$ is a tensor generator for the morphism $f \circ p: F \rightarrow X \rightarrow S$. Thus the morphism $f \circ p$ is quasiffine [17, Proposition 4.1]. By absolute approximation [37, Theorem C(i)], there is a finite subset $J \subset \mathbf{V}$ such that $F_J \rightarrow S$ is quasiffine. By [35, Lemma 1.1], the frame bundle of $V = \bigoplus_{W \in J} W \in \mathbf{V}$ is quasiffine. For the final claim, it suffices to show that an algebraic stack with the resolution property and affine stabilizers at closed points, has affine stabilizers everywhere. This is just [17, Lemma 5.15]. $\square$

If a quasicompact and quasiseparated scheme has an ample line bundle, then it can be written in the form $[U/\mathrm{GL}_1]$, where U is a quasiffine scheme. From this perspective, it is natural to view the resolution property as a higher-dimensional analogue of quasiprojectivity. Note that the converse does not hold (e.g., take $U = \mathbb{A}^2 - \{0\}$ with GL_1 acting with weights $(1, -1)$; then $[U/\mathrm{GL}_1]$ is the line with the doubled origin, which is not separated, so does not admit an ample line bundle—but it does admit an ample family of line bundles).

Definition 2.16 Let S be an algebraic stack. We say that a morphism $G \rightarrow S$ is a *group* if it is a group object in representable morphisms over S . If $G \rightarrow S$ is a group that is flat and of finite presentation, then we say that it is

- (1) *embeddable* if it admits a group monomorphism to $\mathrm{GL}(E)$ for some vector bundle E of rank n on S ,
- (2) *R-embeddable* if it admits a group monomorphism to $\mathrm{GL}(E)$ for some vector bundle E of rank n on S such that the quotient $\mathrm{GL}(E)/G$ is quasiffine over S (note that this implies that $G \rightarrow \mathrm{GL}(E)$ is a closed immersion).

Example 2.17 An embedded group is not necessarily R-embedded, even over a field. Indeed, if $B_{2,\mathbb{C}} \subset GL_{2,\mathbb{C}}$ is a Borel subgroup, then $GL_{2,\mathbb{C}}/B_{2,\mathbb{C}} \simeq \mathbf{P}_{\mathbb{C}}^1$.

Remark 2.18 If $G \rightarrow S$ is a group that is flat and of finite presentation, then $G \rightarrow S$ is embeddable if and only if the structure morphism of the classifying stack $BG \rightarrow S$ is a global quotient. Indeed, if $BG \rightarrow S$ is a global quotient, then there is a vector bundle E on BG whose total space is representable over S . This gives a representable S -morphism $BG \rightarrow BGL_{\text{rank } E}$; the induced map on inertia groups gives a monomorphism $G \hookrightarrow GL(E')$ for some vector bundle E' on S . The converse is similar. In particular, $G \rightarrow S$ is embeddable if and only if there is a vector bundle E on S with a faithful action of G . We call these *faithful* G -representations. If, in addition, $G \rightarrow S$ has affine fibers, then it is R-embeddable if and only if $BG \rightarrow S$ has the resolution property (Theorem 2.15). In particular, such a $G \rightarrow S$ is R-embeddable if and only if there is a vector bundle E on S with a faithful action such that the quotient $GL(E)/G \rightarrow S$ is a quasiaffine morphism. We will call these G -representations *R-faithful*.

We expect the following result to be well known to experts. Over a field, it is due to Rosenlicht [34, Theorem 3].

Proposition 2.19 *Let S be an algebraic stack.*

- (1) *The quotient morphism $GL_{n,S}/U_{n,S} \rightarrow S$ is quasiaffine.*

Let $G \subseteq U_{n,S}$ be a closed subgroup that is flat and of finite presentation over S .

- (2) *If S is a normal scheme and the quotient $U_{n,S}/G$ is representable by a scheme (e.g., if S is Dedekind or the spectrum of a field, see [8, Remark VI_B.9.3(b)]), then the quotient $U_{n,S}/G \rightarrow S$ is quasiaffine.*
- (3) *If $G \rightarrow S$ is finite, then the quotient $U_{n,S}/G \rightarrow S$ is affine.*

Proof Since the inclusion $U_{n,S} \subset GL_{n,S}$ is defined over $\text{Spec } \mathbf{Z}$, in (1) it suffices to assume $S = \text{Spec } \mathbf{Z}$. In particular, we may assume that S is Dedekind. We will now prove that both of the homogenous spaces $U_{n,S}/G$ and $GL_{n,S}/U_{n,S}$ are quasiaffine over S using [33, Theorem VII.2.1], thereby establishing both (1) and (2). Since $GL_{n,S}$ and $U_{n,S}$ are both smooth with connected fibers, it suffices to show that the abelian groups $\text{Pic}(U_{n,k(s)}/G_{k(s)})$ and $\text{Pic}(GL_{n,k(s)}/U_{n,k(s)})$ are torsion. To this end, first note that the cohomology of the complex

$$\text{Hom}(K, \mathbf{G}_m) \rightarrow \text{Pic}(H/K) \rightarrow \text{Pic}(H)$$

is torsion for any inclusion of finite-type group schemes $K \subset H$ over a field by [33, Theorem VII.1.5]. Moreover, by [9, Proposition XVII.2.4(ii)] we always have $\text{Hom}(K, \mathbf{G}_m) = 0$ for every unipotent group scheme K over a field. Further, the Picard groups of $U_{n,k(s)}$ and $GL_{n,k(s)}$ both vanish. Indeed, in the former case the underlying scheme is isomorphic to affine space, and in the latter case it is a principal open of affine space. Thus the two outer terms in the complex are zero in both of our cases and it follows that $\text{Pic}(U_{n,k(s)}/G_{k(s)})$ and $\text{Pic}(GL_{n,k(s)}/U_{n,k(s)})$ are torsion, as desired.

For (3), the morphism $U_{n,S} \rightarrow U_{n,S}/G$ is finite flat. Since $U_{n,S} \rightarrow S$ is affine, the same is true for $U_{n,S}/G$; e.g., [37, Theorem 8.1]. $\square$

Proposition 2.20 *Let $f: X \rightarrow S$ be a quasicompact and quasiseparated morphism of algebraic stacks.*

- (1) *Let W and V be vector bundles on X . If the frame bundle F_W is representable (resp. quasiaffine) over S , then the frame bundle of $V \oplus W$ is also representable (resp. quasiaffine) over S .*
- (2) *If f is adequately affine (e.g., a gerbe banded by a reductive group scheme, a good/adequate moduli space morphism or a coarse moduli space morphism [3, Definition 4.1.1]), then every f -faithful vector bundle V on X has a frame bundle which is affine over S .*

Proof For the representability part of (1), the hypothesis implies that the relative stabilizers act faithfully on the fibers of W and therefore they must act faithfully on the fibers of $V \oplus W$, as desired. For the latter claim, see [35, Lemma 1.1]. For (2), $F_V \rightarrow S$ is representable and adequately affine, and so is affine [3, Theorem 4.3.1]. $\square$

3 Schäppi's theorem

Let A be a ring and let M be a flat A -module. Lazard's theorem [26] tells us that M is a filtered colimit of free A -modules of finite rank. Extending this result to the nonaffine situation is surprisingly subtle [13]. Nonetheless, Schäppi recently proved a remarkable Lazard-type theorem for a restricted but very useful class of flat modules that arise in algebraic geometry [40, Theorem 1.3.1].

Schäppi approaches his result through comodules over flat Hopf algebroids. Since Schäppi's result is crucial for our article and ought to be better known amongst algebraic geometers, we have translated his category-theoretic proof into a direct proof for algebraic stacks. We have made some simple but algebro-geometrically natural generalizations to his hypotheses. Note that while nonaffine schemes tend to not have interesting projective objects, there are interesting algebraic stacks (e.g., those with affine tame [1] or good [2] moduli spaces) that do.

Theorem 3.1 (Schäppi) *Let $f: Y \rightarrow X$ be a flat morphism of quasicompact and quasiseparated algebraic stacks. Let $\mathbf{V}$ be a set of isomorphism classes of vector bundles on X . Let M be a vector bundle on Y that is a projective object of $\mathbf{QCoh}(Y)$. If X has the $\mathbf{V}$ -resolution property, then f_*M is a filtered colimit of finite direct sums of objects of $\mathbf{V}$ and their duals. In particular, if Y is cohomologically affine (e.g., an affine scheme), then $f_*\mathcal{O}_Y$ is a filtered colimit of vector bundles.*

Proof We may replace $\mathbf{V}$ by the set of all finite direct sums of objects of $\mathbf{V}$ and their duals. Let G be a quasicoherent $\mathcal{O}_X$ -module. Consider the category $\mathbf{V}_G$ whose objects are pairs (H, η) , where $H \in \mathbf{V}$ and $\eta: H \rightarrow G$ is an $\mathcal{O}_X$ -module homomorphism. A morphism $h: (H, \eta) \rightarrow (H', \eta')$ in $\mathbf{V}_G$ is an $\mathcal{O}_X$ -module homomorphism $h: H \rightarrow H'$ such that $\eta = \eta' \circ h$. Consider the functor $\mu_G: \mathbf{V}_G \rightarrow \mathbf{QCoh}(X)$ that sends (H, η) to H . It remains to establish the following two claims:

Claim 3.1.1 *Every pair of objects in $\mathbf{V}_G$ has an upper bound and $\text{colim}(\mu_G) \simeq G$.*

Proof If (H_1, η_1) and $(H_2, \eta_2) \in \mathbf{V}_G$, then $(H_1 \oplus H_2, \eta_1 \oplus \eta_2) \in \mathbf{V}_G$. This proves the existence of upper bounds for every pair. We now prove that $\text{colim}(\mu_G) \simeq G$. To do this, consider a quasicoherent

$\mathcal{O}_X$ -module N together with $\mathcal{O}_X$ -module homomorphisms $v_{(H,\eta)}: H \rightarrow N$ for every $(H, \eta) \in \mathbf{V}_G$, such that if $h: (H, \eta) \rightarrow (H', \eta')$ is a morphism in $\mathbf{V}_G$, then $v_{(H,\eta)} = v_{(H',\eta')} \circ h$. It suffices to show that there is a uniquely induced morphism $v: G \rightarrow N$ such that $v_{(H,\eta)} = v \circ \eta$ for every $(H, \eta) \in \mathbf{V}_G$. Since X has the $\mathbf{V}$ -resolution property, G admits a presentation

$$\bigoplus_{j \in J} P_j \xrightarrow{\bigoplus p_j} \bigoplus_{(H,\eta) \in \mathbf{V}_G} H \xrightarrow{\bigoplus_{(H,\eta) \in \mathbf{V}_G} \eta} G \rightarrow 0,$$

where P_j belongs to $\mathbf{V}$ for all $j \in J$. Let $j \in J$ and note that the morphism $p_j: P_j \rightarrow \bigoplus_{(H,\eta) \in \mathbf{V}_G} H$ factors as

$$P_j \xrightarrow{\tilde{p}_j} \bigoplus_{(H,\eta) \in I_j} H \subseteq \bigoplus_{(H,\eta) \in \mathbf{V}_G} H,$$

where $I_j \subseteq \mathbf{V}_G$ is finite. Let $\tilde{Q}_j = \bigoplus_{(H,\eta) \in I_j} H$ and let $\tilde{q}_j = \bigoplus_{(H,\eta) \in I_j} \eta: \tilde{Q}_j \rightarrow G$ be the resulting morphism. Then $\tilde{q}_j \circ \tilde{p}_j = 0$ and so $\tilde{p}_j: (P_j, 0) \rightarrow (\tilde{Q}_j, \tilde{q}_j)$ is a morphism in $\mathbf{V}_G$. Hence $v_{(\tilde{Q}_j, \tilde{q}_j)} \circ \tilde{p}_j = v_{(P_j, 0)} = v_{(P_j, 0)} \circ 0 = 0$. But this means $(\bigoplus_{(H,\eta) \in \mathbf{V}_G} v_{(H,\eta)}) \circ (\bigoplus_j p_j) = 0$. By the universal property of cokernels, there is a unique morphism $v: G \rightarrow N$ such that $v \circ \eta = v_{(H,\eta)}$ for all $(H, \eta) \in \mathbf{V}_G$. $\square$

Claim 3.1.2 *If M is a vector bundle on Y that is a projective object in $\mathbf{QCoh}(Y)$, then $\mathbf{V}_{f_*M}$ is filtered.*

Proof Consider a pair of morphisms $h_1, h_2: (H, \eta) \rightarrow (H', \eta')$; we must show that these can be coequalized in $\mathbf{V}_{f_*M}$. Take the duals of h_1 and h_2 , which results in morphisms $h_1^\vee, h_2^\vee: H'^\vee \rightarrow H^\vee$. Let E be their equalizer in $\mathbf{QCoh}(X)$. If $(F, \rho) \in \mathbf{V}_E$, then taking duals of everything results in a commutative diagram:

$$\begin{array}{ccccc} & & \eta' & & \\ & & \curvearrowright & & \\ H & \xrightarrow{h_1} & H' & \xrightarrow{\quad} & F^\vee \\ & \searrow h_2 & \downarrow & \searrow & \downarrow \\ & & H' & \xrightarrow{\quad} & F^\vee \\ & & \eta' & & \end{array}$$

Thus we just need to produce $(F, \rho) \in \mathbf{V}_E$ that admits a compatible morphism $F^\vee \rightarrow f_*M$. By taking adjoints in the above diagram, we see that it is sufficient to produce a compatible morphism $f^*F^\vee \rightarrow M$. But M is a vector bundle, so it is sufficient to produce a compatible morphism $M^\vee \rightarrow f^*F$. Dualizing their defining diagrams we obtain

$$\begin{array}{ccccccc} & & & & f^*H'^\vee & & \\ & & & & \uparrow & & \\ & & & & f^*h_1^\vee & & \\ & & & & \downarrow & & \\ M^\vee & \xrightarrow{f^*F} & f^*E & \xrightarrow{f^*\rho} & f^*H'^\vee & \xrightarrow{f^*h_1^\vee} & f^*H^\vee \\ & \searrow & \downarrow & \searrow & \uparrow & \searrow & \\ & & f^*H'^\vee & \xrightarrow{f^*h_2^\vee} & f^*H^\vee & & \end{array}$$

But f is flat, so f^*E is also the equalizer of $f^*h_1^\vee$ and $f^*h_2^\vee$; hence there is a compatibly induced morphism $M^\vee \rightarrow f^*E$. By Claim 3.1.1, $E \simeq \text{colim}(\mu_E)$; therefore $f^*E \simeq f^*\text{colim}(\mu_E)$. Now

$M^\vee$ is a vector bundle and Y is quasicompact and quasiseparated, so the functor $\mathrm{Hom}_{\mathcal{O}_Y}(M^\vee, -) = \Gamma(Y, M \otimes_{\mathcal{O}_Y} -)$ preserves filtered colimits of quasicoherent sheaves [20, Lemma 1.2(iii)]. Since M is projective in $\mathbf{QCoh}(Y)$, $M^\vee$ is projective¹ in $\mathbf{QCoh}(Y)$. Hence $\mathrm{Hom}_{\mathcal{O}_Y}(M^\vee, -)$ is also exact and so commutes with *all* colimits of quasicoherent sheaves. Thus we obtain

$$\mathrm{Hom}_{\mathcal{O}_Y}(M^\vee, f^* \mathrm{colim}(\mu_E)) \simeq \mathrm{Hom}_{\mathcal{O}_Y}(M^\vee, \mathrm{colim}(f^* \mu_E)) \simeq \mathrm{colim}_{(F, \rho) \in \mathbf{V}_E} \mathrm{Hom}_{\mathcal{O}_Y}(M^\vee, f^* F).$$

In particular, $M^\vee \rightarrow f^* E$ factors through some morphism $f^* \rho: f^* F \rightarrow f^* E$, where $(F, \rho) \in \mathbf{V}_E$ [43, Tag 09WR]. The claim follows. $\square$

Example 3.2 In Theorem 3.1, it is necessary to include the duals of the generating set $\mathbf{V}$. Indeed, let $X = \mathbf{P}_{x,y}^1$ over a field k with coordinates x and y . Let $Y = D_+(x)$ and take $f: Y \rightarrow X$ to be the inclusion. Then $L = \mathcal{O}(1)$ is ample and X has the $\mathbf{V} = \{L^m\}_{m \leq 0}$ -resolution property (Example 2.13). But $f_* \mathcal{O}_Y \simeq \varinjlim_n L^n$ cannot be written as a filtered colimit of direct sums of objects from $\mathbf{V}$.

The converse to Theorem 3.1 is an earlier result of Hovey [21, Proposition 1.4.4]:

Proposition 3.3 *Let $f: Y \rightarrow X$ be a faithfully flat affine morphism of quasicompact and quasiseparated algebraic stacks, where Y is quasiaffine. Let $\mathbf{V}$ be a set of vector bundles on X . If $f_* \mathcal{O}_Y$ is a filtered colimit of duals of finite direct sums of objects of $\mathbf{V}$, then X has the $\mathbf{V}$ -resolution property.*

Proof It suffices to prove that $\mathbf{V}$ generates $\mathbf{QCoh}(X)$. Write $f_* \mathcal{O}_Y = \mathrm{colim}_{\lambda \in \Lambda} H_\lambda^\vee$, where $H_\lambda = \bigoplus_{V \in J_\lambda} V^{\oplus n_\lambda(V)}$ for $J_\lambda \subseteq \mathbf{V}$ a finite subset and $n_\lambda(V)$ finite. Let $M \in \mathbf{QCoh}(X)$; then there are natural isomorphisms

$$\begin{aligned} \Gamma(Y, f^* M) &= \mathrm{Hom}_{\mathcal{O}_X}(\mathcal{O}_X, f_* f^* M) \simeq \mathrm{Hom}_{\mathcal{O}_X}(\mathcal{O}_X, f_* \mathcal{O}_Y \otimes_{\mathcal{O}_X} M) && (f \text{ is affine}) \\ &\simeq \mathrm{Hom}_{\mathcal{O}_X}(\mathcal{O}_X, \mathrm{colim}_\lambda H_\lambda^\vee \otimes_{\mathcal{O}_X} M) \simeq \mathrm{colim}_\lambda \mathrm{Hom}_{\mathcal{O}_X}(\mathcal{O}_X, H_\lambda^\vee \otimes_{\mathcal{O}_X} M) \\ &\simeq \mathrm{colim}_\lambda \mathrm{Hom}_{\mathcal{O}_X}(H_\lambda, M) \end{aligned}$$

It follows that if $\phi: M \rightarrow N$ is a homomorphism of quasicoherent $\mathcal{O}_X$ -modules, then there is a commutative diagram

$$\begin{array}{ccc} \Gamma(Y, f^* M) & \xrightarrow{\simeq} & \mathrm{colim}_\lambda \mathrm{Hom}_{\mathcal{O}_X}(H_\lambda, M) \\ \Gamma(Y, f^* \phi) \downarrow & & \downarrow \mathrm{colim}_\lambda \mathrm{Hom}_{\mathcal{O}_X}(H_\lambda, \phi) \\ \Gamma(Y, f^* N) & \xrightarrow{\simeq} & \mathrm{colim}_\lambda \mathrm{Hom}_{\mathcal{O}_X}(H_\lambda, N) \end{array}$$

In particular, if $\mathrm{Hom}_{\mathcal{O}_X}(V, \phi): \mathrm{Hom}_{\mathcal{O}_X}(V, M) \rightarrow \mathrm{Hom}_{\mathcal{O}_X}(V, N)$ is the zero map for all $V \in \mathbf{V}$, then $\Gamma(Y, f^* \phi)$ is the zero map. Since Y is quasiaffine, $f^* \phi = 0$. But f is faithfully flat, so $\phi = 0$. That is, $\mathbf{V}$ generates $\mathbf{QCoh}(X)$. $\square$

Let S be an integral quasicompact quasiseparated algebraic stack with generic point ξ . We say that a quasicoherent $\mathcal{O}_S$ -module F is *torsion free* if $\ker(F \rightarrow (i_\xi)_* i_\xi^* F) = 0$, where $i_\xi: S_\xi \hookrightarrow S$ is the generic residual gerbe [36, Theorem B.2]. We record here the following folklore result:

¹This follows from three observations: a direct summand of a projective is projective, if M is projective then $M \otimes_{\mathcal{O}_Y} Q$ is projective for any vector bundle Q , and $M^\vee$ is a direct summand of $M^\vee \otimes_{\mathcal{O}_Y} M \otimes_{\mathcal{O}_Y} M^\vee$.

Corollary 3.4 *Let T be an integral quasicompact algebraic stack with affine diagonal. If every torsion free quasicoherent $\mathcal{O}_T$ -module of finite type is a vector bundle, then T has the resolution property.*

Proof Since T is quasicompact with affine diagonal, there is a smooth covering $\tau: \operatorname{Spec} A \rightarrow T$ and τ is affine. We may write $\tau_* \mathcal{O}_{\operatorname{Spec} A} = \varinjlim_{\lambda} N_{\lambda}$, where the N_{λ} are finite-type quasicoherent $\mathcal{O}_T$ -submodules of $\tau_* \mathcal{O}_{\operatorname{Spec} A}$ [38]. Then the N_{λ} are torsion free, and so are all vector bundles. By Proposition 3.3, the result follows. $\square$

Remark 3.5 Corollary 3.4 gives a simple generalization of [6, 1.4.5]. Indeed, if T is an integral quasicompact quasiseparated algebraic stack that admits a faithfully flat cover by a finite disjoint union of spectra of Dedekind or Prüfer domains (i.e., every finitely generated ideal is invertible), then every torsion free quasicoherent $\mathcal{O}_T$ -module of finite type is a vector bundle [23, Theorem 1].

4 Flags

Our goal is to construct interesting vector bundles on algebraic stacks. The main problem—even for schemes—is that vector bundles are glued from local data. The key idea here is to add the additional structure of a flag to a vector bundle. This ends up being surprisingly useful.

Let Y be an algebraic stack. Let $V_{\bullet}$ be a quasicoherent $\mathcal{O}_Y$ -module V with a finite filtration

$$0 = V_0 \subseteq V_1 \subseteq \cdots \subseteq V_{n-1} \subseteq V_n = V.$$

We call n the *length* of the filtration. Recall that $V_{\bullet}$ is a *flag* if the graded pieces $\operatorname{gr}_i(V_{\bullet}) = V_i/V_{i-1}$ are vector bundles for all $i = 1, \dots, n$. Note that this condition implies that the V_i are also vector bundles. A flag is *complete* if the graded pieces are line bundles. We say that $V_{\bullet}$ is *trivially graded* if $\operatorname{gr}_i(V_{\bullet})$ is a trivial vector bundle for all i . Note that a vector bundle which admits the structure of a trivially graded flag is sometimes referred to as a unipotent vector bundle in the literature; see [31, Section 1; 30, Definition 4.5]. We now have the following key definition:

Definition 4.1 Let $f: X \rightarrow S$ be a morphism of algebraic stacks. Let $V_{\bullet}$ be a flag on X . We say that $V_{\bullet}$ is *f -graded* (or *S -graded*) if for each i there is a vector bundle E_i on S and an isomorphism $\phi_i: \operatorname{gr}_i(V_{\bullet}) \simeq f^* E_i$. If the vector bundle V is f -(R-)faithful, then we say the same of the flag $V_{\bullet}$.

Note that trivially graded flags are graded by every morphism. Remarkably, we can establish that graded flags often descend (Theorem 4.8). We are not aware of any related result in the literature. We begin, however, with the universal example:

Example 4.2 Let S be an algebraic stack. Fix an ordered sequence of n line bundles $\mathcal{L} = (L_1, \dots, L_n)$ on S . Define a category fibered in groupoids $\operatorname{Flag}_{\mathcal{L}} \rightarrow \mathbf{Sch}/S$ as follows: Its objects over $f: X \rightarrow S$ are pairs $(V_{\bullet}, \{\phi_i\}_{i=1}^n)$, where

- (1) $V_{\bullet}$ is a flag of length n , and
- (2) $\phi_i: \operatorname{gr}_i(V_{\bullet}) \simeq f^* L_i$ are isomorphisms for $i = 1, \dots, n$.

That is, an X -point of $\text{Flag}_{\mathcal{L}}$ is an S -graded flag on X whose graded pieces are isomorphic to the f^*L_i . A morphism in $\text{Flag}_{\mathcal{L}}$ is an isomorphism of vector bundles that is compatible with the filtrations and isomorphisms to the L_i . Note that there is a distinguished object in $\text{Flag}_{\mathcal{L}}$ defined over S :

$$\ell = \left(L_1 \subseteq \cdots \subseteq \bigoplus_{j=1}^{n-1} L_j \subseteq \bigoplus_{j=1}^n L_j, \{\text{id}_{L_i}\}_{i=1}^n \right).$$

Every other object in $\text{Flag}_{\mathcal{L}}$ becomes isomorphic to ℓ after passing to a smooth covering. In particular, $\text{Flag}_{\mathcal{L}}$ is a gerbe and the section defined by ℓ induces an equivalence $\text{Flag}_{\mathcal{L}} \simeq BU(\mathcal{L})$, where $U(\mathcal{L}) = \text{Aut}(\ell)$. Clearly $U(\mathcal{L}) \subseteq \text{GL}(\bigoplus_{j=1}^n L_j)$. Moreover, the following diagram 2-commutes:

$$\begin{array}{ccc} X & \xrightarrow{\text{Fr}(V)} & B\text{GL}_{n,S} \\ \downarrow (V_\bullet, \{\phi_i\}) & & \uparrow \wr \text{Fr}(\bigoplus_j L_j) \\ BU(\mathcal{L}) & \longrightarrow & B\text{GL}(\bigoplus_{j=1}^n L_j) \end{array}$$

If $\mathcal{L} = (\mathcal{O}_S, \dots, \mathcal{O}_S)$, then there is a natural identification

$$\text{Aut}(\ell) \simeq U_{n,S}.$$

Thus for any choice of $\mathcal{L}$, $\text{Aut}(\ell)$ is locally isomorphic to $U_{n,S}$. In particular, $U(\mathcal{L}) \rightarrow S$ is a flat group of finite presentation and it follows from Proposition 2.19 that $\text{GL}(\bigoplus_{j=1}^n L_j)/U(\mathcal{L}) \rightarrow S$ is quasiaffine, and $U(\mathcal{L}) \rightarrow S$ is R-embeddable. Hence $BU(\mathcal{L})$ is an algebraic stack, the morphism $BU(\mathcal{L}) \rightarrow S$ has the resolution property and the universal flag $F_\bullet(\mathcal{L})$ on $BU(\mathcal{L})$ is a $(BU(\mathcal{L}) \rightarrow S)$ -tensor generator. In particular, if S has the resolution property, then so does $BU(\mathcal{L})$.

If $B_{n,\mathbb{Z}}$ denotes the Borel subgroup of upper triangular matrices in $\text{GL}_{n,\mathbb{Z}}$ and $B_{n,\mathbb{Z}} \rightarrow \mathbf{G}_m^n$ is the induced quotient, then the following diagram is 2-cartesian:

$$(4.2.1) \quad \begin{array}{ccc} BU(\mathcal{L}) & \longrightarrow & BB_{n,\mathbb{Z}} \\ \downarrow & & \downarrow \\ S & \xrightarrow[\mathcal{L}]{} & B\mathbf{G}_m^n \end{array}$$

Let $f: X \rightarrow S$ be a morphism of algebraic stacks and let $\mathbf{Vect}(X)$ denote the set of isomorphism classes of vector bundles on X . We let

$$\mathbf{Flag}_\emptyset \subseteq \mathbf{Flag}_f \subseteq \mathbf{Flag} \subseteq \mathbf{Vect}(X)$$

denote the sets of isomorphism classes of vector bundles that admit trivially graded complete flags, complete f -graded flags, and complete flags, respectively.

Example 4.3 Standard arguments show that $\mathbf{Flag}$, $\mathbf{Flag}_f$ and $\mathbf{Flag}_\emptyset$ are stable under finite direct sums, the taking of duals, finite tensor products and extensions.

Example 4.4 Let X be an algebraic stack. If

- (1) X is affine,
- (2) X is quasiaffine,
- (3) X is quasiprojective over affine,
- (4) X admits an ample family of line bundles,
- (5) $\mathbf{QCoh}(X)$ is generated by a set of line bundles (e.g., $X = B\mathbf{G}_{m,\mathbf{Z}}^n$), or
- (6) $X = BU_{n,k}$, where k is a field,

then X has the **Flag**-resolution property. These assertions are trivial. If X has the **Flag**-resolution property, then Sch\"appi's theorem (Theorem 3.1) implies that if $p: \mathrm{Spec} A \rightarrow X$ is flat, then $p_* \mathcal{O}_{\mathrm{Spec} A} \simeq \varinjlim_{\lambda} F_{\lambda}$, where the F_{λ} are complete flags on X . Note that $B\mathrm{GL}_{n,\mathbf{Z}}$ does not have the **Flag**-resolution property if $n > 1$.

The following result yields a useful sufficient condition for the existence of a flag structure on a vector bundle:

Lemma 4.5 *Let T be an integral quasicompact quasiseparated algebraic stack with generic point $\eta: \mathrm{Spec} k \rightarrow T$, where k is a field. Let $\gamma: G \rightarrow T$ be a flat group of finite presentation and let $f: BG \rightarrow T$ be the induced morphism. Let V be a vector bundle on BG .*

- (1) *If the adjunction $z: f^* f_* V \rightarrow V$ is an isomorphism after restriction along η , then it is an isomorphism.*
- (2) *If every torsion free quasicoherent $\mathcal{O}_T$ -module of finite type is a vector bundle (Remark 3.5) and G_{η} is unipotent, then V admits an f -graded flag.*

Proof By [36, Theorem B.2], η factors as $\mathrm{Spec} k \xrightarrow{\eta'} T_{\eta} \xrightarrow{i_{\eta}} T$, where T_{η} is the residual gerbe and η' is faithfully flat. For (1), since z_{η} is an isomorphism, $z_{T_{\eta}}$ is an isomorphism. Let $p: T \rightarrow BG$ be the usual section to f ; by descent, V is described by a group homomorphism $v: G \rightarrow \mathrm{GL}(p^* V)$. Note that z is an isomorphism if and only if $\ker(v) = G$ as z corresponds to the inclusion of G -invariants of V into V . Now $\ker(v) \subseteq G$ is a closed immersion as $\mathrm{GL}(p^* V) \rightarrow T$ is separated and contains $\gamma^{-1}(T_{\eta})$, which is dense since γ is open. Hence $\ker(v) = G$.

We prove (2) by induction on $\mathrm{rank}(V)$, the case $\mathrm{rank}(V) = 0$ being trivial; assume that $\mathrm{rank}(V) > 0$. If $z_{T_{\eta}}$ is an isomorphism, then (1) says that z_V is an isomorphism and so we would be done. Otherwise, if $z_{T_{\eta}}$ is not an isomorphism, then $(f^* f_* V)_{T_{\eta}}$ is nonzero as after restriction along η it is the inclusion of the invariants of a nontrivial representation of the unipotent group G_{η} [10, Theorem XVII.3.5]. In particular, there is a nontrivial quotient $V_{T_{\eta}} \rightarrow W$ on T_{η} . Now define $\overline{W}$ to be the image of the composition $V \rightarrow (i_{T_{\eta}})_* V_{T_{\eta}} \rightarrow (i_{T_{\eta}})_* W$. Since the map $V \rightarrow (i_{T_{\eta}})_* W$ is nonzero, $\overline{W}$ is a nonzero torsion free quasicoherent $\mathcal{O}_{BG}$ -module of finite type, and so is a vector bundle. Also, $\overline{W}_{T_{\eta}} = W$ and so $\mathrm{rank}(\overline{W}) = \mathrm{rank}(W) < \mathrm{rank}(V)$. We now have an exact sequence of vector bundles on BG :

$$0 \rightarrow K \rightarrow V \rightarrow \overline{W} \rightarrow 0.$$

The ranks of K and $\overline{W}$ are less than that of V , and so by induction they admit f -graded flags; consequently so does V . $\square$

Using Lemma 4.5, we can now prove the following:

Proposition 4.6 *Let $\mathcal{L}$ be a sequence of n line bundles on an algebraic stack S .*

- (1) *The morphism $f: BU(\mathcal{L}) \rightarrow S$ has the $\mathbf{Flag}_f$ -resolution property.*
- (2) *If S is affine, then f has the $\mathbf{Flag}_{\mathcal{O}}$ -resolution property.*

Proof It suffices to show (1) universally by [17, Proposition 2.8(iv)]. Therefore we may suppose $f: BU(\mathcal{L}) \rightarrow BG_{m,\mathbf{Z}}^n$, where $\mathcal{L}$ is the universal sequence of n line bundles. Since $BB_{n,\mathbf{Z}} \simeq BU(\mathcal{L})$ has the resolution property (Example 4.2), it suffices to show that every vector bundle V is a complete f -graded flag. Since $\mathrm{Spec} \mathbf{Z} \rightarrow BG_{m,\mathbf{Z}}^n$ is a smooth covering, Lemma 4.5 and Remark 3.5 imply that V admits an f -graded flag. But every vector bundle on $BG_{m,\mathbf{Z}}^n$ is a direct sum of line bundles, so (1) follows from Example 4.3.

For (2), it suffices to show that the vector bundle V underlying every f -graded complete flag $V_\bullet$ is a direct summand of a vector bundle W which admits a trivially graded complete flag structure. We proceed by induction on n , the length of $V_\bullet$. If $n = 0$, then the result is trivial. If $n > 0$, then we may write V as an extension

$$0 \rightarrow V_{n-1} \rightarrow V \rightarrow f^*L \rightarrow 0$$

for some line bundle L on S . Since S is affine, there is a split surjection $\gamma: \mathcal{O}_S^{\oplus m} \twoheadrightarrow L$. Then $f^*\gamma: \mathcal{O}_{BU(\mathcal{L})}^{\oplus m} \twoheadrightarrow f^*L$ is a split surjection. Now pull the extension back along $f^*\gamma$, so we may replace f^*L with a trivial bundle of rank m . By induction, there is a split surjection $\delta: U \twoheadrightarrow V_{n-1}$, where U admits a trivially graded complete flag. Pushing forward along a section of δ , V is a direct summand of a vector bundle that is an extension of trivially graded complete flags. Now apply Example 4.3. $\square$

The following concept will be useful. A *refinement* of a flag $V_\bullet$ of length n is another flag $V'_0 \subseteq \cdots \subseteq V'_m = V$ of length $m \geq n$ together with a strictly order-preserving function $\rho: \{0 < 1 < \cdots < n\} \rightarrow \{0 < 1 < \cdots < m\}$ such that $V_i = V'_{\rho(i)}$ for all i .

Example 4.7 Let Y be an algebraic stack. Let $V_\bullet$ be a flag on Y . Suppose that $\mathrm{gr}_i(V_\bullet)$ admits the structure of a complete flag for each $i \geq 0$; then $V_\bullet$ admits a complete refinement. To see this, use Example 4.3.

Theorem 4.8 *Consider a 2-cartesian diagram of algebraic stacks:*

$$\begin{array}{ccc} X' & \xrightarrow{q} & X \\ f' \downarrow & & \downarrow f \\ S' & \xrightarrow{p} & S \end{array}$$

Assume that:

- (1) *S is quasicompact with affine diagonal and has the resolution property.*
- (2) *f is quasicompact and quasiseparated.*
- (3) *p is faithfully flat and S' is affine.*

If X' admits a q -graded flag $V_\bullet$ of length n , then X admits a flag $W_\bullet$ of length n and a degree-preserving split surjection $q^*W_\bullet \twoheadrightarrow V_\bullet$. Moreover:

- (a) If $V_\bullet$ is f' -(R)-faithful, then $W_\bullet$ is f -(R)-faithful.
- (b) If $p_*\mathcal{O}_{S'}$ is a filtered colimit of vector bundles F_λ on X (see Theorem 3.1) and $\mathrm{gr}_i(V_\bullet) \simeq q^*E_i$, then there exists λ_i with $\mathrm{gr}_i(W_\bullet) \simeq f^*F_{\lambda_i} \otimes_{\mathcal{O}_X} E_i$.
- (c) If $V_\bullet$ is fq -graded, then $W_\bullet$ is f -graded.
- (d) Assume that S has the **Flag**-resolution property.
 - (i) If $V_\bullet$ is complete, then $W_\bullet$ admits a complete refinement.
 - (ii) If $V_\bullet$ is fq -graded, then $W_\bullet$ admits a complete f -graded refinement.

Proof Part (a) follows from Proposition 2.20(1), the main claim and [17, Proposition 2.8(iii)]. Part (c) follows from (b). Part (d) follows from (b) and Examples 4.7 and 4.4. Part (b) will follow from the construction of $W_\bullet$, which we prove by induction on $n \geq 0$. The base case is trivial so we assume $n \geq 1$. Choose a vector bundle E on X such that $V_1 = q^*E$; then $(V/V_1)_\bullet$ is a q -graded flag of length $n - 1$. By induction, there is a flag $W_\bullet$ of length $n - 1$ on X and a split surjection $q^*W_\bullet \rightarrow (V/V_1)_\bullet$ as in the statement of the theorem. We now pull back the defining short exact sequence

$$0 \rightarrow q^*E \rightarrow V \rightarrow V/q^*E \rightarrow 0$$

along the surjection $q^*W \twoheadrightarrow V/q^*E$. This results in a commutative diagram with exact rows, whose vertical arrows are easily checked to be split surjective:

$$\begin{array}{ccccccc} 0 & \longrightarrow & q^*E & \longrightarrow & V' & \longrightarrow & q^*W \longrightarrow 0 \\ & & \parallel & & \downarrow & & \downarrow \\ 0 & \longrightarrow & q^*E & \longrightarrow & V & \longrightarrow & V/q^*E \longrightarrow 0 \end{array}$$

Thus we may replace V by V' and assume that $(V/V_1)_\bullet = q^*W_\bullet$.

Since S' is affine and S has affine diagonal, it follows that p and q are affine and faithfully flat morphisms. It follows that we may push forward the top row in the above diagram to obtain an exact sequence of quasicoherent sheaves on X :

$$0 \rightarrow q_*q^*E \rightarrow q_*V' \rightarrow q_*q^*W \rightarrow 0.$$

Pulling this sequence back along the injective adjunction $W \hookrightarrow q_*q^*W$, we obtain a diagram with exact rows:

$$\begin{array}{ccccccc} 0 & \longrightarrow & q_*q^*E & \longrightarrow & W' & \longrightarrow & W \longrightarrow 0 \\ & & \parallel & & \downarrow & & \downarrow \\ 0 & \longrightarrow & q_*q^*E & \longrightarrow & q_*V' & \longrightarrow & q_*q^*W \longrightarrow 0 \end{array}$$

Since E is a vector bundle and q is affine, the projection formula and flat base change produce natural isomorphisms

$$q_* q^* E \simeq q_* \mathcal{O}_{X'} \otimes_{\mathcal{O}_X} E \simeq (f^* p_* \mathcal{O}_{S'}) \otimes_{\mathcal{O}_X} E.$$

Assume we are given that $p_* \mathcal{O}_{S'} \simeq \varinjlim_{\lambda} F_{\lambda}$, where the F_{λ} are vector bundles, or we produce such a description by Sch\"appi's theorem (Theorem 3.1). Since the adjunction $p^* p_* \mathcal{O}_{S'} \twoheadrightarrow \mathcal{O}_{S'}$ is split surjective, taking λ sufficiently large, we can assume that the induced morphisms $p^* F_{\lambda} \rightarrow \mathcal{O}_{S'}$ are also split surjective. But we also have natural isomorphisms

$$\begin{aligned} \mathrm{Ext}_{\mathcal{O}_X}^1(W, q_* q^* E) &\simeq H^1(X, W^{\vee} \otimes_{\mathcal{O}_X} q_* q^* E) \\ &\simeq H^1(X, \varinjlim_{\lambda} (f^* F_{\lambda} \otimes_{\mathcal{O}_X} W^{\vee} \otimes_{\mathcal{O}_X} E)) \\ &\simeq \varinjlim_{\lambda} H^1(X, f^* F_{\lambda} \otimes_{\mathcal{O}_X} W^{\vee} \otimes_{\mathcal{O}_X} E) & [20, \text{Lemma 1.2(iii)}] \\ &\simeq \varinjlim_{\lambda} \mathrm{Ext}_{\mathcal{O}_X}^1(W, f^* F_{\lambda} \otimes_{\mathcal{O}_X} E). \end{aligned}$$

Thus the extension W' is the pushforward of an extension

$$0 \rightarrow f^* F_{\lambda} \otimes_{\mathcal{O}_X} E \rightarrow W'_{\lambda} \rightarrow W \rightarrow 0$$

along the morphism $f^* F_{\lambda} \otimes_{\mathcal{O}_X} E \rightarrow f^*(p_* \mathcal{O}_{S'}) \otimes_{\mathcal{O}_X} E \simeq q_* q^* E$. Pulling this exact sequence back along q , we obtain an exact sequence

$$0 \rightarrow q^* f^* F_{\lambda} \otimes_{\mathcal{O}_{X'}} q^* E \rightarrow q^* W'_{\lambda} \rightarrow q^* W \rightarrow 0.$$

But the pushforward of this extension along the split surjection $q^* f^* F_{\lambda} \otimes_{\mathcal{O}_{X'}} q^* E \twoheadrightarrow q^* E$ is clearly isomorphic to V' and the resulting morphism $q^* W'_{\lambda} \twoheadrightarrow V'$ is split surjective. $\square$

5 Unipotent morphisms and groups


5.1 Unipotent morphisms

Motivated by the results of the previous section, we make the following definition:

Definition 5.1 Let $f: X \rightarrow S$ be a morphism of algebraic stacks. Then f is

- *(R-)unipotent* if X admits a complete f -graded f -(R-)faithful flag,
- *locally (R-)unipotent* if there is an fpqc covering $S' \rightarrow S$ [43, Tag 022B] such that $f \times_S S'$ is (R-)unipotent, and
- *geometrically (R-)unipotent* if for every algebraically closed field k and morphism $\mathrm{Spec} k \rightarrow S$, the induced morphism $X \times_S \mathrm{Spec} k \rightarrow \mathrm{Spec} k$ is (R-)unipotent.

Locally R-unipotent and unipotent morphisms have affine and quasiaffine diagonals, respectively. We have the following sequence of implications:



The solid arrows follow from the definitions. The dashed arrows are partial converses. The dashed vertical implications are valid only for certain groups over normal bases. The dashed horizontal implications are valid only in special cases.

Quasicompact and quasiseparated representable morphisms are always unipotent. There are two prototypical examples of R-unipotent morphisms:

- (1) $BU(\mathcal{L}) \rightarrow S$ (Example 4.2), and
- (2) quasiaffine morphisms $X \rightarrow S$.

We will return to these examples often. We begin this section with the following two trivial lemmas:

Lemma 5.2 *Let $f: X \rightarrow S$ be an (R-)unipotent morphism of algebraic stacks. If $S' \rightarrow S$ is a morphism of algebraic stacks, then $f \times_S S'$ is (R-)unipotent. Moreover, the same holds for the local and geometric versions.*

Lemma 5.3 *Let $X \xrightarrow{h} Y \xrightarrow{g} S$ be morphisms of algebraic stacks.*

- (1) *If g is unipotent and h is quasicompact, quasiseparated and representable, then $g \circ h$ is unipotent.*
- (2) *If g is R-unipotent and h is quasiaffine, then $g \circ h$ is R-unipotent.*

Moreover, the analogous properties hold for the local and geometric versions.

In general, unipotent morphisms are not stable under composition (Example 5.6). We are optimistic that locally unipotent morphisms have better stability properties under composition, but this appears to be surprisingly subtle — even when S is the spectrum of a field, Y is quasiaffine and X is a gerbe over Y .

We now have the following characterization of (R-)unipotent morphisms, which provides a unipotent enrichment of the Totaro–Gross theorem (Theorem 2.15).

Theorem 5.4 *Let $f: X \rightarrow S$ be a morphism of algebraic stacks with S quasicompact. Then the following are equivalent:*

- (1) *The morphism f is R-unipotent.*

- (2) For some integer $n \geq 0$, there is an ordered sequence of n line bundles $\mathcal{L}$ on S and a factorization of f as $X \xrightarrow{h} BU(\mathcal{L}) \rightarrow S$, where h is quasifinite.
- (3) The morphism f has the **Flag** $_f$ -resolution property and the relative inertia stack $I_f \rightarrow X$ has affine fibers.

Moreover, f is unipotent if and only if there is a quasicompact, quasiseparated and representable morphism $X \xrightarrow{h} BU(\mathcal{L})$ as in (2).

Proof The statement regarding unipotent morphisms will follow from our arguments from the equivalence (1) $\iff$ (2).

We have (2) $\implies$ (1) by Lemma 5.3 and the R -unipotence of $BU(\mathcal{L}) \rightarrow S$. As for (1) $\implies$ (2), by definition there is a complete f -graded flag $V_\bullet$ of length n on X that is f -(R -)faithful. That is, there is an ordered sequence of line bundles $\mathcal{L} = (L_1, \dots, L_n)$ on S together with isomorphisms $\phi_i: \mathrm{gr}_i(V_\bullet) \simeq f^* L_i$ for all i . By Example 4.2, we obtain a morphism $X \xrightarrow{h} BU(\mathcal{L})$. Postcomposing this with the quasifinite morphism $BU(\mathcal{L}) \rightarrow BGL(\bigoplus_{i=1}^n L_i) \simeq BGL_{n,S}$ yields a quasifinite morphism by Theorem 2.15(2), and the result follows from [43, Tag 054G].

For (2) $\implies$ (3), Proposition 4.6 and [17, Proposition 2.8(v)] imply that f has the **Flag** $_f$ -resolution property. Finally, for (3) $\implies$ (2), the Gross–Totaro theorem (Theorem 2.15) produces a $V_\bullet \in \mathbf{Flag}_f$ which is f - R -faithful. $\square$

A remarkable consequence of Theorem 4.8 is that local (R -)unipotence implies (R -)unipotence on bases with the **Flag**-resolution property.

Theorem 5.5 *Let $f: X \rightarrow S$ be a locally unipotent (resp. locally R -unipotent) morphism of algebraic stacks. Let S be quasicompact with affine diagonal.*

- (1) *If S has the resolution property, then X is a global quotient (resp. has the resolution property).*
- (2) *If S has the **Flag**-resolution property, then f is unipotent (resp. R -unipotent).*

Proof By assumption, there is a faithfully flat cover $p: S' \rightarrow S$ such that $f': X \times_S S' \rightarrow S'$ admits a complete f' -graded f' -faithful (resp. f' - R -faithful) flag $V_\bullet$. Passing to a smooth cover of S' , we may assume that S' is affine and that the flag is trivially graded, and so $V_\bullet$ is pf' -graded. Now apply Theorem 4.8 to

$$\begin{array}{ccc} X \times_S S' & \xrightarrow{q} & X \\ f' \downarrow & & \downarrow f \\ S' & \xrightarrow{p} & S \end{array}$$

Then we obtain a flag $W_\bullet$ on X and a split surjection $q^* W_\bullet \rightarrow V_\bullet$. Part (c) implies that $W_\bullet$ is f -graded. Part (a) says that $W_\bullet$ is f -faithful (resp. f - R -faithful), so X is a global quotient (resp. has the resolution property). Part (d) says that if S has the **Flag**-resolution property, then $W_\bullet$ admits a complete refinement, which is also f -graded. Hence f admits a complete f -graded f -(R -)faithful flag. That is, f is unipotent (resp. R -unipotent). $\square$

Example 5.6 Let S be a proper normal surface over $\mathbf{C}$ with no nontrivial line bundles [41]. Let $Y \subseteq S$ be the inclusion of the regular locus, which is an open subscheme whose complement $S \setminus Y$ is a finite set of closed points. Then Y admits a nontrivial very ample line bundle L . Let $\mathcal{U} = BU(\mathcal{O}, L^\vee) \cong BL$. Then

- (1) $\pi: \mathcal{U} \rightarrow Y$ and $j: Y \rightarrow S$ are R-unipotent,
- (2) the composition $p = j \circ \pi$ is locally R-unipotent but not unipotent, and
- (3) S has the resolution property but not the **Flag**-resolution property.

Indeed, (1) is clear from Theorem 5.4. Also, (3) follows from (2) and Theorem 5.5.

For (2), to see that p is locally R-unipotent, it suffices to restrict to affine opens $\text{Spec } A \subseteq S$; then $Y_A = Y \cap \text{Spec } A$ is quasiaffine and so $L_{Y_A}^\vee$ is globally generated. Hence there is a quasiaffine morphism $\mathcal{U} \rightarrow BG_{a, Y_A}^n$ for some n ; after composition with the open immersion $BG_{a, Y_A}^n \subseteq BG_{a, \text{Spec } A}^n$, we see that p is locally R-unipotent by Lemma 5.3(2).

If $p = j \circ \pi$ is unipotent, since S has no nontrivial line bundles, $\mathcal{U}$ must admit a faithful trivially graded flag. Hence it suffices to show that every trivially graded flag $V_\bullet$ on $\mathcal{U}$ is of the form $\pi^* W_\bullet$ for some trivially graded flag $W_\bullet$ on Y . A simple calculation shows that $\mathbf{R}^1 \pi_* \mathcal{O}_{\mathcal{U}} \simeq L^\vee$. Note that Y is not quasiaffine because $H^0(Y, \mathcal{O}_Y) = \mathbf{C}$ [43, Tag 01P9], and therefore L is nontrivial [43, Tag 01QE]. In turn, this implies $H^0(Y, L^\vee) = 0$ because by composing with a nonzero section of L , a nonzero section $L^\vee$ would yield a trivialization of $L^\vee$. Therefore, since $L^\vee$ has no nonzero global sections, if $W_\bullet$ is a trivially graded flag $W_\bullet$ on Y then $\text{Hom}_{\mathcal{O}_Y}(W, L^\vee) = 0$. Now let $V_\bullet$ be a trivially graded flag on $\mathcal{U}$. By induction on the length of $V_\bullet$, we may assume that $V_\bullet$ is an extension

$$e: 0 \rightarrow \pi^* \mathcal{O}_Y \rightarrow V_\bullet \rightarrow \pi^* W_\bullet \rightarrow 0,$$

where $W_\bullet$ is a trivially graded flag on Y . There is also an exact sequence

$$0 \rightarrow \text{Ext}_{\mathcal{O}_Y}^1(W, \mathcal{O}_Y) \rightarrow \text{Ext}_{\mathcal{O}_{\mathcal{U}}}^1(\pi^* W, \pi^* \mathcal{O}_Y) \rightarrow \text{Hom}_{\mathcal{O}_Y}(W, \mathbf{R}^1 \pi_* \mathcal{O}_{\mathcal{U}}).$$

Since $\text{Hom}_{\mathcal{O}_Y}(W, L^\vee) = 0$, the first map is bijective and so our extension e is pulled back from Y .

Two useful application of Theorem 5.5 are the following:

Corollary 5.7 *Let $f: X \rightarrow S$ be a locally (R-)unipotent morphism of algebraic stacks. If S is quasicompact, then there is a smooth surjection $S' \rightarrow S$ such that $f \times_S S'$ is (R-)unipotent.*

Proof Passing to a smooth cover of S , we may assume that S is an affine scheme and so has the **Flag**-resolution property. The result follows from Theorem 5.5. $\square$

Corollary 5.8 *Let $f: X \rightarrow S$ be a morphism of algebraic stacks.*

- (1) *If f is geometrically (R-)unipotent, then $f \times_S \text{Spec } k: X \times_S \text{Spec } k \rightarrow \text{Spec } k$ is (R-)unipotent for every field k and morphism $\text{Spec } k \rightarrow S$.*
- (2) *Let $S' \rightarrow S$ be surjective. If $f': X \times_S S' \rightarrow S'$ is geometrically (R-)unipotent, then f is geometrically (R-)unipotent.*

Proof In (2), we may assume that $S = \operatorname{Spec} k$ and $S' = \operatorname{Spec} k'$, where $k \subseteq k'$ is a field extension and k' is algebraically closed. It suffices to prove that if f' is (R-)unipotent, then f is (R-)unipotent, which also implies (1). But $S' \rightarrow S$ is an fpqc covering, so f is locally (R-)unipotent. Since S is affine, it has the **Flag**-resolution property. The result now follows from Theorem 5.5. $\square$

5.2 Unipotent groups

A powerful source of unipotent morphisms will be unipotent groups and gerbes. Let k be a field. Let G be an algebraic group over k . There are several characterizations of unipotence over k [9, Theorem XVII.3.5]. We take the following: G is *unipotent* if there is an embedding into the upper triangular unipotent matrices $G \hookrightarrow U_{n,k} \subset \operatorname{GL}_{n,k}$ for some n . Families of unipotent groups over general bases are more subtle, however, so we make the following definition:

Definition 5.9 Let S be an algebraic stack. A group $G \rightarrow S$ that is flat and of finite presentation is said to be *(R-)unipotent*, *locally (R-)unipotent* or *geometrically (R-)unipotent* if the corresponding gerbe $BG \rightarrow S$ is so.

Remark 5.10 There is a potential for confusion with this terminology: unipotence of the group $G \rightarrow S$ is not the same thing as unipotence of the morphism $G \rightarrow S$. For instance, every quasicompact, quasiseparated and representable morphism of algebraic stacks is unipotent, but an affine group scheme $G \rightarrow S$ need not be unipotent in the sense of Definition 5.9.

Remark 5.11 A group $G \rightarrow S$ that is flat and of finite presentation is unipotent (resp. R-unipotent) if and only if there exists an ordered sequence of line bundles $\mathcal{L}$ on S , an object $V_\bullet$ of $\operatorname{Flag}_{\mathcal{L}}(S)$ and a monomorphism (resp. with quasiaffine quotient)

$$G \hookrightarrow \operatorname{Aut}_{\operatorname{Flag}_{\mathcal{L}}}(V_\bullet) \subset \operatorname{GL}(V).$$

Thus (R-)unipotent groups are (R-)embeddable. Note that the above automorphism group is an inner form of $U(\mathcal{L})$. In particular, if S is a scheme, then G can be embedded in a Zariski form of $U_{n,S}$. Hence, if $S = \operatorname{Spec} k$ is the spectrum of a field, then $G \rightarrow \operatorname{Spec} k$ is unipotent (in the sense of Definition 5.9) if and only if there exists an embedding $G \subseteq U_{n,k}$ for some n , so both possible definitions of a unipotent group scheme agree over a field.

Example 5.12 By Proposition 2.19, the group $U_{n,S} \rightarrow S$ is R-unipotent, as is any flat and finitely presented closed subgroup $H \subset U_{n,S}$ when either S is the spectrum of a field or a Dedekind domain, or $H \rightarrow S$ is finite. In particular, if $G \rightarrow S$ is (locally) unipotent and either $G \rightarrow S$ is finite or S is the spectrum of a field, then it is (locally) R-unipotent.

Now translating Theorem 5.5 into groups we obtain the following:

Corollary 5.13 Let S be a quasicompact algebraic stack with affine diagonal. Let $G \rightarrow S$ be a locally (R-)unipotent group.

- (1) If S has the resolution property, then $G \rightarrow S$ is (R-)embeddable.
- (2) If S has the **Flag**-resolution property, then $G \rightarrow S$ is (R-)unipotent.

Note the stark contrast to the case of tori. Tori are always locally embeddable, but may not be embeddable even over a projective curve [9, X.1.6 and XI.4.6].

Remark 5.14 By Example 5.12 and Corollary 5.13, all the definitions in Definition 5.9 agree when S is the spectrum of a field. Thus one may speculate what the best definition of unipotence is over a general base. We argue that local R-unipotence is the best behaved as evidenced by Theorem 5.5, Examples 5.12, 5.16, 5.17 and 5.22, and Proposition 7.1.

We have the following proposition, which justifies our terminology:

Proposition 5.15 *Let $f: X \rightarrow S$ be a geometrically unipotent morphism of algebraic stacks. Then the relative geometric stabilizers of f are all unipotent groups.*

Proof We may assume that $S = \operatorname{Spec} l$, where l is an algebraically closed field. Let $\bar{x}: \operatorname{Spec} k \rightarrow X$ be a point of X , where k is an algebraically closed field. Let $G_{\bar{x}} = \operatorname{Aut}_X(\bar{x})$ be its automorphism group. Then we have a representable morphism $\tilde{x}: BG_{\bar{x}} \rightarrow X$. By Lemma 5.3(1), $BG_{\bar{x}}$ is unipotent. Now apply the last sentence of Remark 5.11. $\square$

Example 5.16 Let S be an algebraic stack. Let E be a locally free sheaf on S of rank n . Then the vector group $\mathbf{V}(E) \rightarrow S$ is locally R-unipotent. We may work locally on S , so we may assume that $E \simeq \mathcal{O}_S^{\oplus n}$. In this case, $\mathbf{V}(E) \simeq \mathbf{G}_{a,S}^n$. Now $\mathbf{G}_{a,\mathbf{Z}}^n \subseteq U_{n+1,\mathbf{Z}} \subseteq \operatorname{GL}_{n+1,\mathbf{Z}}$. Since $\mathbf{Z}$ is a Dedekind domain, it follows from Example 5.12 that $\mathbf{G}_{a,\mathbf{Z}}^n \rightarrow \operatorname{Spec} \mathbf{Z}$ is R-unipotent. Hence $\mathbf{V}(E) \rightarrow S$ is locally R-unipotent. If S is quasicompact with affine diagonal and the **Flag**-resolution property, then $\mathbf{V}(E) \rightarrow S$ is even R-unipotent (Corollary 5.13).

Example 5.17 Let $p > 0$ be a prime. Let S be an algebraic $\mathbf{F}_p$ -stack. Let $G \rightarrow S$ be a finite étale group scheme of degree p^d for some $d \geq 0$. Then $G \rightarrow S$ is locally R-unipotent. Indeed, smooth locally on S we may assume that $G \rightarrow S$ is a constant group scheme. We may thus assume that $S = \operatorname{Spec} \bar{\mathbf{F}}_p$ and G is a constant p -group. Standard facts from group theory show that G is unipotent, and the claim follows.

Example 5.18 If X denotes the affine line with doubled origin, then it is an R-unipotent scheme. Indeed if $x, y \in X$ denote the two origins, then the trivially graded flag $\mathcal{O}_X = I_{\{x,y\}} \subset I_{\{x\}} \oplus I_{\{y\}}$ is a tensor generator, where I_C is the ideal sheaf associated to a reduced closed subscheme $C \subset X$.

5.3 Representable R-unipotent morphisms

We now briefly discuss representable (locally) R-unipotent morphisms $f: X \rightarrow S$. Characterizing such morphisms seems subtle. For instance, consider for simplicity the case when $S = \operatorname{Spec} k$ is the spectrum of a field. It follows from Theorem 5.4 that representable R-unipotent morphisms $f: X \rightarrow \operatorname{Spec} k$ are exactly quotients of the form T/U_n , where T is a quas affine scheme over k acted on freely by U_n . In particular, a quas affine morphism over k is R-unipotent and hence geometrically R-unipotent. The converse is not true, however.

Example 5.19 (1) Let k be a field and denote by X the complement of a single point in the exceptional locus of $\mathbf{A}_k^2$ blown up at the origin. We claim that the morphism $f: X \rightarrow \mathbf{A}^2$ is geometrically R-unipotent but is not quasiaffine. Indeed, the geometric fibers are all affine and hence the morphism is geometrically R-unipotent. On the other hand, if f was quasiaffine then the natural map $g: X \rightarrow \operatorname{Spec} \Gamma(X, \mathcal{O}_X)$ would be an open immersion [43, Tag 01SM] and this cannot be true since $\Gamma(X, \mathcal{O}_X) = k[x, y]$.

(2) In [5, Example 3.16], Asok and Doran describe a free action of $U_1 = \mathbf{G}_a$ on $\mathbf{A}^5$ such that the quotient $\mathbf{A}^5/\mathbf{G}_a$ is not even a scheme. More generally, given a smooth affine scheme T over k equipped with a free and proper action of U_n , [5, Corollary 3.18] yields an effective criterion for determining when the quotient T/U_n will be affine or quasiaffine.

We have the following result, whose proof we defer until Section 7:

Proposition 5.20 *Let $f: X \rightarrow S$ be a representable and geometrically R-unipotent morphism of algebraic stacks. If f is proper, then f is finite.*

It follows from Proposition 5.20 that representable geometrically R-unipotent morphisms over a field do not contain any positive-dimensional proper subvarieties.

5.4 Geometrically unipotent vs locally unipotent groups

We conclude this section with some thoughts on the following natural refinement of Question 1.1: which geometrically unipotent groups are locally unipotent? Note that locally unipotent groups are *always* quasiaffine; in particular, they are separated and schematic. In positive characteristic, it is easy to produce geometrically unipotent group algebraic spaces that are not separated. There are also separated geometrically unipotent group algebraic spaces that are not schemes. In particular, these give examples of geometrically unipotent groups that are not locally unipotent.

Example 5.21 Let k be a field of characteristic 2. Let $S = \mathbf{A}_k^1$. Let $H \subseteq (\mathbf{Z}/2\mathbf{Z})_S = G$ be the étale subgroup obtained by deleting the nontrivial point over the origin.

(1) Let $Q = G/H \rightarrow S$ be the line with the doubled origin, viewed as a group scheme. Then $Q \rightarrow S$ is geometrically unipotent, but nonseparated. In particular, it is not locally unipotent.

(2) We have $H \subseteq G \subseteq \mathbf{G}_{a,S}$. Let $Q' = \mathbf{G}_{a,S}/H \rightarrow S$. Then $Q' \rightarrow S$ is smooth, geometrically unipotent, with connected fibers, nonseparated and not a scheme [8, VI_B.5.5]. This shows that the closed subgroup condition in Proposition 2.19(2) is necessary.

(3) A more sophisticated example was constructed in [33, X.14]: take $T = \mathbf{A}_k^2$; then there is a *closed* subgroup $N \subseteq \mathbf{G}_{a,T}^2$ with $N \rightarrow T$ étale. Taking $Q'' = \mathbf{G}_{a,T}^2/N$, we produce a smooth separated geometrically unipotent group of finite presentation with connected fibers. There it is shown that Q'' is not representable by a scheme. This shows that the condition that the quotient is representable by a scheme in Proposition 2.19(2) is necessary.

In characteristic 0, there are strong results in the existing literature, however:

Example 5.22 Let S be an algebraic stack of equicharacteristic 0. Let $G \rightarrow S$ be geometrically unipotent and flat-locally schematic. Then $G \rightarrow S$ is affine. Indeed, this is local on S , so we may assume that S is affine. By standard limit methods, we may assume that S is of finite type over $\operatorname{Spec} \mathbf{Z}$ and so excellent; in particular, S has noetherian normalization. If S is either normal or $G \rightarrow S$ is embeddable, then there is an isomorphism of groups

$$\exp: \mathbf{V}(\operatorname{Lie}(G)) \rightarrow G,$$

where we endow the left-hand side with the group structure coming from the Baker–Campbell–Hausdorff formula. In particular, when $G \rightarrow S$ is commutative it is a vector group, and so is locally R-unipotent. The statement involving the exponential map is established in [44, Section 1.3; 33, XV.3]. That $G \rightarrow S$ is affine follows from the normal case and Chevalley’s theorem [37, Theorem 8.1].

Example 5.23 Let $G \rightarrow S$ be an affine and geometrically unipotent group.

- (1) If $S = \operatorname{Spec} A$, where A is a Prüfer domain, then $G \rightarrow S$ is R-unipotent.
- (2) If S is equicharacteristic $p > 0$ and $G \rightarrow S$ is a finite locally free commutative group scheme with dual of height 1, then $G \rightarrow S$ is locally R-unipotent.

For (1), by Lemma 4.5 every vector bundle on BG admits the structure of a complete S -graded flag. In combination with Corollary 3.4, this implies that G is R-unipotent. For (2), the exact sequence of [4, Proposition 1.1] implies there is an affine morphism $BG \rightarrow B\mathbf{V}(\omega)$ for some vector group $\mathbf{V}(\omega) \rightarrow S$. Since $\mathbf{V}(\omega) \rightarrow S$ is locally R-unipotent (Example 5.16), (2) follows.

6 Applications

6.1 Unipotent gerbes and the resolution property

Theorem 5.5 immediately yields a proof of Theorem A.

Proof of Theorem A A $\mathbf{G}_a$ -gerbe is locally R-unipotent (Example 5.12). It follows from Theorem 5.5 that $\mathcal{G}$ has the resolution property. $\square$

Remark 6.1 Assume that X lives over a scheme S and $G \rightarrow S$ is locally R-unipotent: for example, $G = \mathbf{G}_a^n$, $G = U_{n,S}$, or S is the spectrum of a field and $G \subseteq U_{n,S}$ (see Example 5.12). Then the argument above for Theorem A also holds for G -gerbes $\mathcal{G} \rightarrow X$.

We next obtain the following corollary of Theorem 5.5, which is key for our proof of Theorem B:

Corollary 6.2 Let S be a quasicompact algebraic stack with affine diagonal over $\mathbf{F}_p$. Let $\pi: \mathcal{G} \rightarrow S$ be a gerbe that is relatively Deligne–Mumford, separated and with p -power order inertia.

- (1) If S has the resolution property, then $\mathcal{G}$ has the resolution property.
- (2) If S has the **Flag**-resolution property, then π is R-unipotent.

Proof Since π is relatively Deligne–Mumford and separated, there is a smooth covering $S' \rightarrow S$ such that $\mathcal{G} \times_S S' \simeq BG'$, where $G' \rightarrow S'$ is finite étale. The p -power order inertia assumption on π implies that $G' \rightarrow S'$ has degree p^d for some $d \geq 0$. Example 5.17 shows that $G' \rightarrow S'$ is locally R-unipotent. Now apply Theorem 5.5. $\square$

6.2 Deligne–Mumford gerbes

Let X be an algebraic stack. Then there is an injective group homomorphism

$$\mathrm{Br}(X) \rightarrow \mathrm{Br}'(X) = \mathrm{H}^2(X, \mathbf{G}_m)_{\mathrm{tors}}$$

from the geometric Brauer group of X to the cohomological Brauer group of X . However, it is still not known when the map is surjective (see, for instance, [28] for an introduction to this problem).

We are now in a position to prove one of our key results:

Theorem 6.3 *Let S be a connected quasicompact algebraic stack with affine diagonal and the resolution property over a field k . Let $\pi: \mathcal{G} \rightarrow S$ be a gerbe that is relatively Deligne–Mumford and separated. If $\mathrm{Br}(T) = \mathrm{Br}'(T)$ for all finite étale morphisms $T \rightarrow S$ (e.g., if S is a quasicompact scheme that admits an ample line bundle), then $\mathcal{G}$ has the resolution property.*

Proof The gerbe $\mathcal{G} \rightarrow S$ is locally banded by a finite constant group scheme $\underline{G}_S$. It follows that the sheaf of isomorphisms (in the category of Bands over S)

$$I = \underline{\mathrm{Isom}}(\underline{\mathrm{Band}}(\mathcal{G}), \underline{G}_S)$$

is a torsor under the outer automorphism group $\underline{\mathrm{Out}}(\underline{G}_S) \rightarrow S$. This map is finite étale, so we may replace S with I as the resolution property descends under such maps [17, Proposition 5.3(vii)]. Hence we can assume that $\mathcal{G} \rightarrow S$ is banded by $\underline{G}_S$, which is constant. Now there is a factorization

$$\mathcal{Z} \rightarrow \mathcal{G} \rightarrow S,$$

where $\mathcal{Z} \rightarrow S$ is a gerbe banded by the center of $\underline{G}_S$, which is a finite abelian constant group scheme and $\mathcal{Z} \rightarrow \mathcal{G}$ is finite étale and surjective. Indeed, define $\mathcal{Z}$ to be the stack of banded equivalences

$$\mathcal{Z} = \underline{\mathrm{HOM}}_{\mathrm{id}}(B\underline{G}_S, \mathcal{G}) \rightarrow S.$$

By [15, IV.2.3.2(iii)], $\mathcal{Z}$ is banded by $Z(\underline{G}_S) = Z(G)_S$ and the natural evaluation map $\mathrm{ev}: \mathcal{Z} \rightarrow \mathcal{G}$ induces a map on bands equal to the inclusion $Z(G)_S \subset \underline{G}_S$. Again, by [17, Proposition 5.3(vii)], we may now replace $\mathcal{G}$ by $\mathcal{Z}$. Thus $\mathcal{G}$ is banded by the constant group scheme associated to the finite abelian group

$$\mathbf{Z}/p_1^{n_1}\mathbf{Z} \times \cdots \times \mathbf{Z}/p_l^{n_l}\mathbf{Z},$$

where the p_i are primes. Hence $\mathcal{G} \cong \mathcal{R}_1 \times_S \cdots \times_S \mathcal{R}_l$, where $\mathcal{R}_i$ is a $\mathbf{Z}/p_i^{n_i}\mathbf{Z}$ -gerbe for all i . For each p_i that is prime to the characteristic of k , consider the finite and separable field extension $k \subseteq k'$ obtained by adding in all the $p_i^{n_i}$ -th roots of unity. As before, it suffices to prove the result after base changing to k' .

In particular, for those p_i we now have $\mathbf{Z}/p_i^{n_i}\mathbf{Z} \cong \mu_{p_i^{n_i}}$. Since $\mathrm{Br}(S) = \mathrm{Br}'(S)$, the resulting $\mathcal{R}_i$ have the resolution property. For the remaining p_i that are the characteristic of k , these $\mathcal{R}_i$ have the resolution property by Corollary 6.2. An easy calculation now shows that $\mathcal{G}$ has the resolution property. $\square$

6.3 Deligne–Mumford stacks in positive characteristic

Proof of Theorem B We may assume that $\mathcal{X}$ is connected. By [32, Proposition 2.1], the morphism $\mathcal{X} \rightarrow X$ can be rigidified. Hence it factors into a gerbe morphism and a coarse moduli space morphism:

$$\mathcal{X} \rightarrow \mathcal{X}^{\mathrm{rig}} \rightarrow X.$$

Here $\mathcal{X}^{\mathrm{rig}}$ has generically trivial stabilizers. Therefore $\mathcal{X}^{\mathrm{rig}}$ is a global quotient [12, Theorem 2.18]. By [24, Theorem 2.1] (see also [11, Corollary 4.5]), there is a finite flat cover $Z \rightarrow \mathcal{X}^{\mathrm{rig}}$, where Z is a quasiprojective scheme. Base changing along this cover, we may replace $\mathcal{X}^{\mathrm{rig}}$ by Z [17, Proposition 5.3(vii)]. Then $\mathcal{X} \rightarrow X$ is a separated Deligne–Mumford gerbe, where X is a quasiprojective scheme over $\mathrm{Spec} k$. The result now follows from Theorem 6.3 and Gabber’s theorem [22]. $\square$

7 Faithful moduli spaces

If V is a nonzero representation of the unipotent group $U_{n,k}$ over a field k , then the set of fixed points $V^{U_{n,k}}$ is nonzero [9, XVII, Proposition 3.2]. If X is a quasiasfine scheme and F is a nonzero quasicoherent sheaf on X , then $\Gamma(X, F)$ is nonzero. The following result shows that this property generalizes to any locally R -unipotent morphism:

Proposition 7.1 *Let $f: X \rightarrow S$ be a locally R -unipotent morphism of algebraic stacks. If F is quasicoherent on X and $f_*F = 0$, then $F = 0$.*

Proof By flat base change we may assume that f is R -unipotent and S is affine. Let N be a quasicoherent $\mathcal{O}_S$ -module; then $\mathrm{Hom}_{\mathcal{O}_X}(f^*N, F) = \mathrm{Hom}_{\mathcal{O}_S}(N, f_*F) = 0$. It follows that if V belongs to $\mathbf{Flag}_f$, then $\mathrm{Hom}_{\mathcal{O}_X}(V, F) = 0$. As f is R -unipotent, it has the $\mathbf{Flag}_f$ -resolution property (Theorem 5.4). Hence $F = 0$. $\square$

The following definition is due to Alper (personal communication, 2017).

Definition 7.2 Let $f: X \rightarrow S$ be a quasicompact and quasiseparated morphism of algebraic stacks. We say that f is a *faithful moduli space* if

- (1) the natural map $\mathcal{O}_S \rightarrow f_*\mathcal{O}_X$ is an isomorphism (i.e., f is Stein), and
- (2) if $F \in \mathbf{QCoh}(X)$ and $f_*F = 0$, then $F = 0$ (i.e., f_* is zero-reflecting).

Remark 7.3 Faithful moduli spaces are compatible with quasiasfine flat base change. Indeed, if $p: S' \rightarrow S$ is flat and quasiasfine and $f': X' = X \times_S S' \rightarrow S'$ is the induced morphism, then flat base change says that (1) is preserved. Let $p': X' \rightarrow X$ be the induced morphism; then p_* and p'_* are zero-reflecting on quasicoherent sheaves because they are quasiasfine. Hence $p_*f'_* \simeq f_*p'_*$ is zero-reflecting on

quasicoherent sheaves, and so f'_* is zero-reflecting on quasicoherent sheaves, which proves (2). In particular, if $f: X \rightarrow S$ is a faithful moduli space and S has quasiaffine diagonal, then f remains a faithful moduli space after arbitrary flat base change on S .

Example 7.4 A Stein and locally R-unipotent morphism is a faithful moduli space (Proposition 7.1). Also if S is a normal scheme, then any open immersion $U \subseteq S$ with complement of codimension ≥ 2 is Stein and R-unipotent. In particular, there is not necessarily an induced bijection on closed points, so some may find the term “moduli space” misleading.

Remark 7.5 A faithful moduli space f need not imply that the functor f_* is faithful. Indeed, f_* being faithful forces f to be quasiaffine (because the structure sheaf would be f -generating). Moreover, f being Stein implies that the quasiaffine morphism f is an open immersion. Thus the second example in Example 7.4 is essentially the only example of a faithful moduli space for which f_* is faithful.

Remark 7.6 The converse to Proposition 7.1 does not hold: if X denotes affine n -space with a doubled origin and $f: X \rightarrow \mathbf{A}^n$ is the natural map, f is a faithful moduli space but is not locally R-unipotent unless $n = 1$ (see Example 5.18). Indeed, it cannot be locally R-unipotent when $n \geq 2$, because Theorem 5.5 would imply that X has affine diagonal. On the other hand, f is Stein. To see that f_* is zero-reflecting, use Lemma 7.7.

Lemma 7.7 Let $f: X \rightarrow Y$ be a morphism of algebraic stacks that is representable and quasifinite. Then $f_*: \mathbf{QCoh}(X) \rightarrow \mathbf{QCoh}(Y)$ is zero-reflecting.

Proof By passing to a smooth cover of Y , we may assume that Y is an affine scheme and X is an algebraic space. Let F be a quasicoherent $\mathcal{O}_X$ -module with $f_*F = 0$; then the same is true for any of its finite-type subsheaves, so we may assume F is also of finite-type. If $F \neq 0$, then we may replace X with $V(\text{Ann}_{\mathcal{O}_X}(F))$ and then Y with $V(\ker(\mathcal{O}_Y \rightarrow f_*\mathcal{O}_X))$. Then $\text{Ann}_{\mathcal{O}_X}(F) = 0$ and f has schematically dense image. Since f is quasifinite, there is a nonempty open subset $V \subseteq Y$ such that $f^{-1}(V) \rightarrow V$ is finite. Then $F|_{f^{-1}(V)} = 0$ and so $F = 0$, which is a contradiction. $\square$

Example 7.8 Let k be a field. Let G be an algebraic group scheme over k . Then $BG \rightarrow \text{Spec } k$ is a faithful moduli space if and only if G is unipotent. One implication follows from Example 7.4. For the other, if V is a finite-dimensional representation of G over k , then $V^G \neq 0$. By induction, V is a trivially graded flag on BG . Taking V to be a faithful representation of G , we see that G admits a faithful flag. In particular, it is unipotent.

Example 7.9 Let $Y \xrightarrow{g} X \xrightarrow{f} S$ be quasicompact and quasiseparated morphisms of algebraic stacks.

- (1) If f and g are faithful moduli spaces, then $f \circ g$ is a faithful moduli space.
- (2) If f_* is zero-reflecting, then $X \rightarrow \underline{\text{Spec}}_S(f_*\mathcal{O}_X)$ is a faithful moduli space.

The following result is a version of Proposition 5.15 for faithful moduli spaces:

Proposition 7.10 Let $f: X \rightarrow S$ be a faithful moduli space with affine stabilizers. Then the relative geometric stabilizers of f are unipotent groups.

Proof We may assume that S is an affine scheme; then X is quasicompact and quasiseparated. Let $\bar{x}: \text{Spec } k \rightarrow X$ be a geometric point of X , where k is an algebraically closed field. Let $G = \text{Aut}(\bar{x})$, which is a group scheme of finite type over $\text{Spec } k$. There is an induced quasiaffine morphism $BG \rightarrow X$ [36, Theorem B.2] and so $BG \rightarrow \text{Spec } k$ is a faithful moduli space (Example 7.9(2)). By Example 7.8, the result follows. $\square$

The following lemma turns out to be very useful:

Lemma 7.11 *Let $f: X \rightarrow S$ be a quasicompact, quasiseparated and Stein morphism of algebraic stacks. If $I \subseteq \mathcal{O}_X$ is a quasicoherent sheaf of ideals, then $V(I) \subseteq f^{-1}(V(f_*I))$.*

Proof Indeed, the quasicoherent sheaf defining the closed substack $f^{-1}(V(f_*I)) \subseteq X$ is the image of $f^*f_*I \rightarrow \mathcal{O}_X$, which factors through I . $\square$

Proposition 7.12 *Let $f: X \rightarrow S$ be a faithful moduli space with affine stabilizers. If S is reduced and quasiseparated with affine stabilizers, then there is a nonempty open $U \subseteq S$ with $f^{-1}(U) \rightarrow U$ a locally R -unipotent gerbe.*

Proof We may assume that S is quasicompact; then [19, Proposition 2.6(i)] implies that there is a nonempty open quasicompact $U \subseteq S$ with affine diagonal. Let $\text{Spec } A \rightarrow U$ be a smooth cover; then the composition $\text{Spec } A \rightarrow U \rightarrow S$ is quasiaffine and smooth. It follows from Remark 7.3 that we may now assume that $S = \text{Spec } A$, where A is reduced.

Now there is a nonempty open immersion $\mathcal{G} \subseteq X_{\text{red}}$ such that $\mathcal{G}$ is a gerbe with the resolution property with affine coarse space $g: \mathcal{G} \rightarrow T$ (combine [43, Tags 06RC and 06NH] with [19, Proposition 2.6(i)]). Let $I \subseteq \mathcal{O}_X$ be a quasicoherent ideal sheaf defining the reduced closed complement Z of $\mathcal{G} \subseteq X$. If $Z \neq \emptyset$, then $\Gamma(X, \mathcal{O}_X) \rightarrow \Gamma(Z, \mathcal{O}_Z)$ is nonzero because it is a ring homomorphism and $\Gamma(Z, \mathcal{O}_Z) \neq 0$. In particular, $\Gamma(X, I) \subsetneq A$. Since $I \neq 0$, it follows that we may choose $a \in \Gamma(X, I) \setminus \{0\}$ that is not a unit. Since A is reduced, a is not nilpotent and so $\text{Spec } A_a \neq \emptyset$. But Lemma 7.11 implies that $Z \cap f^{-1}(\text{Spec } A_a) = \emptyset$, and so we may replace A by A_a and assume that $Z = \emptyset$. Hence $\mathcal{G} = X_{\text{red}} \subseteq X$ is a surjective closed immersion with defining ideal J . By Lemma 7.11 again, we may further shrink S so that $J = 0$ and $\mathcal{G} = X$ is a gerbe. But the composition $\mathcal{O}_S \rightarrow t_*\mathcal{O}_T \rightarrow f_*\mathcal{O}_X$ is an isomorphism, as is $\mathcal{O}_T \rightarrow g_*\mathcal{O}_X$; hence, $T \simeq S$ and $X \rightarrow S$ is of finite presentation. By standard limit methods and Theorem 5.4, we may replace S by the spectrum of the localization at a minimal prime of A , which is a field [43, Tag 00EU]. By Example 7.8 and Remark 5.14 the result follows. $\square$

Corollary 7.13 *Let k be a field. Let $X \rightarrow \text{Spec } k$ be a faithful moduli space. Then $X \rightarrow \text{Spec } k$ is an R -unipotent gerbe. In particular, if X is an algebraic space, then $X \rightarrow \text{Spec } k$ is an isomorphism.*

We conclude with the proof of our characterization of proper and representable geometrically R -unipotent morphisms:

Proof of Proposition 5.20 We may assume that S is affine. By Zariski's main theorem, it suffices to prove that f is quasifinite. Hence we may further reduce to the situation where S is the spectrum of an

algebraically closed field k and f is R -unipotent. Let $X \rightarrow S' \rightarrow S$ be the Stein factorization of f . Then $S' \rightarrow S$ is finite and the morphism $X \rightarrow S'$ is again R -unipotent, and so we may further reduce to the situation where X is geometrically connected and f is Stein. It follows from Example 7.4 that $X \rightarrow \operatorname{Spec} k$ is a faithful moduli space. By Corollary 7.13, $X \rightarrow \operatorname{Spec} k$ is an isomorphism, which gives the claim. $\square$

Acknowledgements

We would like to thank Jarod Alper, Brian Conrad, Andrea Di Lorenzo, Aise Johan de Jong, Andrew Kresch, Nikolas Kuhn, Max Lieblich, Martin Olsson, Stefan Schröer, Tuomas Tajakka, and Angelo Vistoli for helpful comments. We are especially thankful to David Rydh for carefully reading an earlier draft and providing many invaluable suggestions. We would also like to thank the referees for a number of helpful suggestions. Bragg was supported by the National Science Foundation under NSF postdoctoral research fellowship DMS-1902875. Hall was partially supported by the Australian Research Council DP210103397 and FT210100405. Mathur conducted part of this research in the framework of the research training group GRK 2240: Algebro-geometric Methods in Algebra, Arithmetic and Topology, which is funded by the Deutsche Forschungsgemeinschaft. Mathur was also supported by the Swedish Research Council under grant 2016-06596 while in residence at Institut Mittag-Leffler in Djursholm, Sweden during the Fall of 2021. Mathur was also supported by FONDECYT regular grant 1230402.

References

- [1] **D Abramovich, M Olsson, A Vistoli**, *Tame stacks in positive characteristic*, Ann. Inst. Fourier (Grenoble) 58:4 (2008) 1057–1091 MR
- [2] **J Alper**, *Good moduli spaces for Artin stacks*, Ann. Inst. Fourier (Grenoble) 63:6 (2013) 2349–2402 MR
- [3] **J Alper**, *Adequate moduli spaces and geometrically reductive group schemes*, Algebr. Geom. 1:4 (2014) 489–531 MR
- [4] **M Artin, J S Milne**, *Duality in the flat cohomology of curves*, Invent. Math. 35 (1976) 111–129 MR
- [5] **A Asok, B Doran**, *On unipotent quotients and some $\mathbb{A}^1$ -contractible smooth schemes*, Int. Math. Res. Pap. 2007:2 (2007) art. id. rpm005 MR
- [6] **F Bruhat, J Tits**, *Groupes réductifs sur un corps local, II: Schémas en groupes. Existence d’une donnée radicielle valuée*, Inst. Hautes Études Sci. Publ. Math. 60 (1984) 197–376 MR
- [7] **B Conrad**, *Smooth linear algebraic groups over the dual numbers*, MathOverflow question (2010) Available at <http://mathoverflow.net/q/22078>
- [8] **M Demazure, A Grothendieck**, *Schémas en groupes, Tome I: Propriétés générales des schémas en groupes, Exposés I–VII (SGA 3_I)*, Lecture Notes in Math. 151, Springer (1970) MR Annotated reprint in Documents Mathématiques (Paris) 7 (2011)
- [9] **M Demazure, A Grothendieck**, *Schémas en groupes, Tome II: Groupes de type multiplicatif, et structure des schémas en groupes généraux, Exposés VIII–XVIII (SGA 3_{II})*, Lecture Notes in Math. 152, Springer (1970) MR Annotated preprint available at <https://webusers.imj-prg.fr/patrick.polo/SGA3/>
- [10] **M Demazure, A Grothendieck**, *Schémas en groupes, Tome III: Structure des schémas en groupes réductifs, Exposés XIX–XXVI (SGA 3_{III})*, Lecture Notes in Math. 153, Springer (1970) MR Annotated reprint in Documents Mathématiques (Paris) 8 (2011)
- [11] **N Deshmukh, A Hogadi, S Mathur**, *Quasi-affineness and the 1-resolution property*, Int. Math. Res. Not. 2022:3 (2022) 2224–2249 MR
- [12] **D Edidin, B Hassett, A Kresch, A Vistoli**, *Brauer groups and quotient stacks*, Amer. J. Math. 123:4 (2001) 761–777 MR

- [13] **S Estrada, P A G Asensio, S Odabaşı**, *A Lazard-like theorem for quasi-coherent sheaves*, *Algebr. Represent. Theory* 16:4 (2013) 1193–1205 MR
- [14] **O Gabber**, *Some theorems on Azumaya algebras*, from “The Brauer group” (Les Plans-sur-Bex, 1980) (M Kervaire, M Ojanguren, editors), *Lecture Notes in Math.* 844, Springer (1981) 129–209 MR
- [15] **J Giraud**, *Cohomologie non abélienne*, *Grundle Math. Wissen. Band 179*, Springer (1971) MR
- [16] **P Gross**, *The resolution property of algebraic surfaces*, *Compos. Math.* 148:1 (2012) 209–226 MR
- [17] **P Gross**, *Tensor generators on schemes and stacks*, *Algebr. Geom.* 4:4 (2017) 501–522 MR
- [18] **A Grothendieck**, *Le groupe de Brauer, II: Théorie cohomologique*, from “Séminaire Bourbaki, 9”, *Soc. Math. France, Paris* (1995) Exposé 297, 287–307 MR
- [19] **J Hall, D Rydh**, *Algebraic groups and compact generation of their derived categories of representations*, *Indiana Univ. Math. J.* 64:6 (2015) 1903–1923 MR
- [20] **J Hall, D Rydh**, *Perfect complexes on algebraic stacks*, *Compos. Math.* 153:11 (2017) 2318–2367 MR
- [21] **M Hovey**, *Homotopy theory of comodules over a Hopf algebroid*, from “Homotopy theory: relations with algebraic geometry, group cohomology, and algebraic K-theory” (P Goerss, S Priddy, editors), *Contemp. Math.* 346, Amer. Math. Soc., Providence, RI (2004) 261–304 MR
- [22] **A J de Jong**, *A result of Gabber*, preprint (2003) Available at <https://www.math.columbia.edu/~dejong/papers/2-gabber.pdf>
- [23] **I Kaplansky**, *Modules over Dedekind rings and valuation rings*, *Trans. Amer. Math. Soc.* 72 (1952) 327–340 MR
- [24] **A Kresch, A Vistoli**, *On coverings of Deligne–Mumford stacks and surjectivity of the Brauer map*, *Bull. London Math. Soc.* 36:2 (2004) 188–192 MR
- [25] **G Laumon, L Moret-Bailly**, *Champs algébriques*, *Ergebnisse der Math.* (3) 39, Springer (2000) MR
- [26] **D Lazard**, *Sur les modules plats*, *C. R. Acad. Sci. Paris* 258 (1964) 6313–6316 MR
- [27] **S Mathur**, *The resolution property via Azumaya algebras*, *J. Reine Angew. Math.* 774 (2021) 93–126 MR
- [28] **S Mathur**, *Experiments on the Brauer map in high codimension*, *Algebra Number Theory* 16:3 (2022) 747–775 MR
- [29] **S Mathur, S Schröer**, *The resolution property holds away from codimension three*, *Trans. Amer. Math. Soc.* 376:2 (2023) 1041–1063 MR
- [30] **S Mukai**, *Semi-homogeneous vector bundles on an Abelian variety*, *J. Math. Kyoto Univ.* 18:2 (1978) 239–272 MR
- [31] **T Oda**, *Vector bundles on an elliptic curve*, *Nagoya Math. J.* 43 (1971) 41–72 MR
- [32] **M Olsson**, *A boundedness theorem for Hom-stacks*, *Math. Res. Lett.* 14:6 (2007) 1009–1021 MR
- [33] **M Raynaud**, *Faisceaux amples sur les schémas en groupes et les espaces homogènes*, *Lecture Notes in Math.* 119, Springer (1970) MR
- [34] **M Rosenlicht**, *On quotient varieties and the affine embedding of certain homogeneous spaces*, *Trans. Amer. Math. Soc.* 101 (1961) 211–223 MR
- [35] **D Rydh**, *Remarks on “The resolution property for schemes and stacks” by B Totaro* (2009) Available at <https://people.kth.se/~dary/resprop20090923.pdf>
- [36] **D Rydh**, *Étale dévissage, descent and pushouts of stacks*, *J. Algebra* 331 (2011) 194–223 MR
- [37] **D Rydh**, *Noetherian approximation of algebraic spaces and stacks*, *J. Algebra* 422 (2015) 105–147 MR
- [38] **D Rydh**, *Approximation of sheaves on algebraic stacks*, *Int. Math. Res. Not.* 2016:3 (2016) 717–737 MR
- [39] **D Rydh**, *Absolute noetherian approximation of algebraic stacks*, preprint (2023) arXiv 2311.09208
- [40] **D Schäppi**, *A characterization of categories of coherent sheaves of certain algebraic stacks*, preprint (2012) arXiv 1206.2764 To appear in *J. Pure Appl. Algebra*
- [41] **S Schröer**, *On non-projective normal surfaces*, *Manuscripta Math.* 100:3 (1999) 317–321 MR
- [42] **S Schröer, G Vezzosi**, *Existence of vector bundles and global resolutions for singular surfaces*, *Compos. Math.* 140:3 (2004) 717–728 MR
- [43] *The Stacks project* (2005–) Available at <http://stacks.math.columbia.edu>

- [44] **J Tong**, *Unipotent groups over a discrete valuation ring (after Dolgachev–Weisfeiler)*, from “Autour des schémas en groupes, III” (G Prasad, P Polo, editors), Panor. Synthèses 47, Soc. Math. France, Paris (2015) 173–225 MR
- [45] **B Totaro**, *The resolution property for schemes and stacks*, J. Reine Angew. Math. 577 (2004) 1–22 MR

DANIEL BRAGG bragg@math.utah.edu

Department of Mathematics, University of Utah, Salt Lake City, UT, United States

JACK HALL jack.hall@unimelb.edu.au

School of Mathematics and Statistics, The University of Melbourne, Parkville, Australia

SIDDHARTH MATHUR siddharth.mathur@uga.edu

Department of Mathematics, Pontificia Universidad Católica de Chile, Santiago, Chile

Department of Mathematics, University of Georgia, Athens, GA, United States

Proposed: Marc Levine

Received: February 16, 2022

Seconded: Mark Gross, Dan Abramovich

Revised: November 25, 2024

On the Donaldson–Scaduto conjecture

SAMAN HABIBI ESFAHANI AND YANG LI

Motivated by G_2 -manifolds with coassociative fibrations in the adiabatic limit, Donaldson and Scaduto conjectured the existence of associative submanifolds homeomorphic to a three-holed 3-sphere with three asymptotically cylindrical ends in the G_2 -manifold $X \times \mathbb{R}^3$, or equivalently similar special Lagrangians in the Calabi–Yau 3-fold $X \times \mathbb{C}$, where X is an A_2 -type ALE hyperkähler 4-manifold. We prove this conjecture by solving a real Monge–Ampère equation with a singular right-hand side, which produces a potentially singular special Lagrangian. Then, we prove the smoothness and asymptotic properties for the special Lagrangian using inputs from geometric measure theory. The method produces many other asymptotically cylindrical $U(1)$ -invariant special Lagrangians in $X \times \mathbb{C}$, where X arises from the Gibbons–Hawking construction.

1 Introduction

Simon Donaldson [2017] initiated a program to study G_2 -manifolds through coassociative K3 fibrations $\pi : M^7 \rightarrow B^3$ over a 3-dimensional base B^3 , in the adiabatic limit where the diameters of the K3 fibers shrink to zero. This program is expected to lead to large classes of new examples of compact torsion-free G_2 -manifolds.

Subsequent work of Donaldson and Scaduto [2022] provided a conjectural limiting description of certain associative submanifolds in the adiabatic setting. Roughly speaking, the generic part of the associative submanifolds is fibered over one-dimensional gradient flowlines inside the base B^3 . These flowlines are allowed to end on the discriminant locus of $\pi : M^7 \rightarrow B^3$, and generically, three flowlines can meet to form a triple junction point. Donaldson and Scaduto made several conjectures concerning the existence of the local model for the triple junction.

In the “global” version of the conjecture, let (X^4, I_1, I_2, I_3) be a hyperkähler K3 surface, and let $\alpha_1, \alpha_2, \alpha_3$ be (-2) -classes in $H_2(X^4; \mathbb{R})$, namely $\alpha_i^2 = -2$ with respect to the intersection product, such that $\alpha_1 + \alpha_2 + \alpha_3 = 0$. Each α_i determines a complex structure J_i in the S^2 -family of complex structures, so that for suitable choices of $v_i \in \mathbb{R}^3$, we have three cylindrical associative submanifolds $P_i := \Sigma_i \times (\mathbb{R}^+ v_i) \subset X^4 \times \mathbb{R}^3$ for $i \in \{1, 2, 3\}$.

Conjecture (Donaldson and Scaduto) *There is an associative submanifold homeomorphic to a three-holed 3-sphere $P \subset X^4 \times \mathbb{R}^3$ with three ends asymptotic to cylinders P_1, P_2 and P_3 .*

In the “local” version of the conjecture, the K3 surface is replaced with an A_2 -type ALE gravitational instanton $X_{A_2}^4$. In the Gibbons–Hawking ansatz, $X_{A_2}^4$ is defined as the completion of a $U(1)$ -bundle over $\mathbb{R}^3 \setminus \{p_1, p_2, p_3\}$. We make the genericity assumption that p_1, p_2, p_3 are three noncolinear points

MSC2020: 32Q25, 53C26, 53C29, 53C38, 53D20.

© 2026 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).




Figure 1: Donaldson-Scaduto conjecture.

in $\mathbb{R}^3$, which corresponds to the condition that the three (-2) -spheres Σ_i are holomorphic with respect to different complex structures. The relation to the global version is that when the K3 surface is the small desingularization (e.g., smoothing or resolution) of an orbifold with local A_2 -singularity, then the local version captures the metric behavior near the desingularization region.

We prove the local version in this note.

Theorem 1 (Donaldson–Scaduto conjecture, local version) *There exists a $U(1)$ -invariant associative submanifold $P \subset X^4_{A_2} \times \mathbb{R}^3$ homeomorphic to a three-holed 3-sphere, with three ends asymptotic to the half-cylinders $\Sigma_i \times (\mathbb{R}^+ v_i)$, where $i \in \{1, 2, 3\}$.*

In fact, since the vectors v_1 , v_2 , and v_3 lie in a plane, say $\mathbb{R}^2 \times \{0\} \subset \mathbb{R}^3$, the associative submanifold P can be equivalently interpreted as a special Lagrangian submanifold in $X^4 \times \mathbb{R}^2$ with an appropriate Calabi–Yau structure. Our method readily generalizes to the A_{n-1} -type ALE or ALF gravitational instantons X where $n \geq 3$ and the monopole points $p_1, \dots, p_n$ in the Gibbons–Hawking ansatz are in the “convex position”, namely they are the vertices of a convex polygon in a plane in $\mathbb{R}^3$, arranged in the counterclockwise orientation. We equip $X \times \mathbb{R}^2$ with a natural product Calabi–Yau 3-fold structure.

Theorem 2 (generalization) *There is an $(n-1)$ -dimensional family of $U(1)$ -invariant special Lagrangian submanifolds in the Calabi–Yau 3-fold $X \times \mathbb{R}^2$, each homeomorphic to an n -holed 3-sphere and with n asymptotically cylindrical ends, modeled on the product of $\Sigma_i \subset X$ and $\{(p_{i+1} - p_i) \cdot y = c_i\} \subset \mathbb{R}^2$, where the parameters $\{c_i\}_{i=1}^n$ satisfy one constraint, $\sum_{i=1}^n c_i = 0$.*

The $n - 1$ parameters geometrically correspond to the translation of the n asymptotic cylinder ends, subject to one constraint coming from the vanishing of the integral of $\text{Im}(\Omega)$. Specifically, two of these parameters account for global translations along the $\mathbb{R}^2$ direction, while the remaining $n - 3$ parameters yield geometrically distinct special Lagrangians. Moreover, these special Lagrangians remain rigid after fixing the asymptotic conditions, as studied in [Habibi Esfahani 2022, Theorem 43].

Remark 3 In [Bera et al. 2025, Theorem 1], it is shown that any special Lagrangian submanifold homeomorphic to an n -holed 3-sphere with n asymptotically cylindrical ends, satisfying the conditions of Theorem 2, belongs to the $(n-1)$ -dimensional family of $U(1)$ -invariant special Lagrangians constructed in this paper.

We expect these special Lagrangians to be useful as building blocks in the gluing construction of new special Lagrangians in “local Calabi–Yau 3-folds” admitting a fibration of A_{n-1} -type spaces over a Riemann surface.

Plan of the paper We focus on Theorem 2. In Section 2, we introduce the geometric structures on the ambient spaces. In Section 3.1, we dimensionally reduce the $U(1)$ -invariant special Lagrangian conditions to an equation for surfaces in the symplectic quotient. Under an additional graphical assumption, this leads to a 2-dimensional real Monge–Ampère equation for some potential φ over the convex polygon with vertices $p_1, \dots, p_n$. In Section 3.2, we solve the appropriate Dirichlet problem for φ , where the boundary data is given by affine linear functions on each edge of the polygon. The $U(1)$ -bundle L° over the gradient graph of the solution over the open solid convex polygon is an open $U(1)$ -invariant special Lagrangian in $X \times \mathbb{R}^2$.

In Section 4.1, we take the closure of L° , denoted by L . To prove L is a closed submanifold, we need two extra ingredients. First, the gradient of φ diverges to infinity near the edges of the convex polygon away from the vertices, and therefore, the edges give no contribution to the closure of the special Lagrangian; this uses some analysis on the real Monge–Ampère equation. Second, the vertex contributions introduce only smooth points to the special Lagrangian; this uses some geometric measure theory, as well as the classification of $U(1)$ -invariant special Lagrangian cones in $\mathbb{C}^3$, following an earlier idea of Joyce [2005]. In Section 4.2, we prove an exponential decay estimate and show that L has the expected asymptotic behavior. In Section 4.3, we prove L is homeomorphic to an n -holed 3-sphere. This concludes the proof of Theorem 2.

2 Preliminaries: ambient spaces

We recall the hyperkähler structure on the $U(1)$ -invariant gravitational instanton X , and describe the Calabi–Yau structure on $Z = X \times \mathbb{R}^2$ and the G_2 -structure on $M = X \times \mathbb{R}^3$.

Hyperkähler structure Let X be a complete noncompact $U(1)$ -invariant hyperkähler 4-manifold, given by the Gibbons–Hawking construction as follows. Let $n \geq 3$, and let $p_1, p_2, \dots, p_n$ be n distinct points in $\mathbb{R}^3$. We will assume that p_i are contained in the plane $\mathbb{R}^2 \times \{0\} \subset \mathbb{R}^3$, and in the “convex position”, namely they are the vertices of a convex polygon, arranged in counterclockwise order. In the $n = 3$ case, up to coordinate rotation, this amounts to the genericity assumption that p_1, p_2, p_3 are noncolinear.

Let u_1, u_2, u_3 denote the coordinates on $\mathbb{R}^3$. Let

$$\pi : X^\circ \rightarrow \mathbb{R}^3 \setminus \{p_1, p_2, \dots, p_n\}$$

be a principal $U(1)$ -bundle, with Chern class 1 on small S^2 around each point p_i . Define a positive harmonic function $V : \mathbb{R}^3 \setminus \{p_1, p_2, \dots, p_n\} \rightarrow \mathbb{R}$ by

$$V(u) = A + \sum_{i=1}^n \frac{1}{2|u - p_i|}, \quad \text{where } A = \text{constant} \geq 0,$$

and let θ be a $U(1)$ -connection on X° with curvature 2-form $-*dV$. The Gibbons–Hawking ansatz describes a hyperkähler structure on X° given by the symplectic forms

$$\omega_1 = \theta \wedge du_1 + V du_2 \wedge du_3, \quad \omega_2 = \theta \wedge du_2 + V du_3 \wedge du_1, \quad \omega_3 = \theta \wedge du_3 + V du_1 \wedge du_2,$$

and the metric $g = V^{-1}\theta^2 + V \sum_{i=1}^3 du_i^2$.

The coordinates u_1, u_2, u_3 are the moment maps with respect to the symplectic forms $\omega_1, \omega_2, \omega_3$, respectively. The manifold X is obtained by adding a point above each p_i , and the hyperkähler structure extends smoothly to X , with the corresponding complex structures I_1, I_2, I_3 . In fact, for each $(a_1, a_2, a_3) \in S^2 \subset \mathbb{R}^3$, we obtain a complex structure $\sum_{i=1}^3 a_i I_i$ on X . For $A = 0$, the hyperkähler manifold X is an A_{n-1} -type ALE space, and for $A > 0$, it is an A_{n-1} -type ALF space.

Let $\Sigma_i = \pi^{-1}[p_i, p_{i+1}]$ be the preimage of $[p_i, p_{i+1}]$, the line segment from p_i to p_{i+1} , and for convenience write $p_{n+1} = p_1$. Each Σ_i is a 2-sphere, which is holomorphic with respect to the complex structure associated with the vector v_i ,

$$v_i = \frac{(p_{i+1} - p_i)}{|p_{i+1} - p_i|} \in S^2 \cap (\mathbb{R}_{(u_1, u_2)}^2 \times \{0\}) \subset \mathbb{R}_{(u_1, u_2, u_3)}^3.$$

Calabi–Yau structure Let $Z = X \times \mathbb{R}_{(y_1, y_2)}^2$. The 6-dimensional manifold Z can be equipped with the Calabi–Yau structure

$$g_Z = g_X + g_{\mathbb{R}^2}, \quad \omega = \omega_3 + dy_2 \wedge dy_1, \quad \Omega = (\omega_1 + i\omega_2) \wedge (dy_2 + i dy_1),$$

where y_1, y_2 denote the coordinates on $\mathbb{R}^2$. Note that with our convention,

$$\omega^3 = \frac{3\sqrt{-1}}{4} \Omega \wedge \bar{\Omega}.$$

We extend the $U(1)$ -action on X to a $U(1)$ -action on Z by $e^{it} \cdot (q, y) \rightarrow (e^{it} \cdot q, y)$ for any $q \in X$ and $y \in \mathbb{R}^2$.

Let $\tilde{v}_i = Rv_i$, where $R : \mathbb{R}_{(u_1, u_2)}^2 \rightarrow \mathbb{R}_{(y_1, y_2)}^2$ is the linear transformation given by the 90-degree rotation,

$$R(\partial_{u_1}) = -\partial_{y_2}, \quad R(\partial_{u_2}) = \partial_{y_1}.$$

Let L_i be $\Sigma_i \times (\mathbb{R}^+ \cdot \tilde{v}_i) \subset X \times \mathbb{R}_{(y_1, y_2)}^2$ translated along some vector in the $\mathbb{R}_{(y_1, y_2)}^2$, so that L_i is contained inside

$$\Sigma_i \times \{y \in \mathbb{R}^2 \mid (p_{i+1} - p_i) \cdot y = c_i\} \subset X \times \mathbb{R}^2, \quad \text{for } i \in \{1, 2, \dots, n\}.$$

A direct computation shows that the L_i are $U(1)$ -invariant special Lagrangians in Z ,

$$\omega|_{L_i} \equiv 0 \quad \text{and} \quad \text{Im}(\Omega)|_{L_i} \equiv 0 \quad \text{for } i \in \{1, 2, \dots, n\}.$$

Remark 4 The 90-degree rotation is an artifact of our choice of holomorphic volume form Ω . This choice will be useful when we later derive the real Monge–Ampère equation.

G_2 structure Let $M = X \times \mathbb{R}^3$. The 7-dimensional manifold M can be equipped with a torsion-free G_2 -structure

$$\begin{aligned}\phi = & dy_1 \wedge dy_2 \wedge dy_3 + dy_1 \wedge du_1 \wedge \theta - V dy_1 \wedge du_2 \wedge du_3 \\ & + dy_2 \wedge du_2 \wedge \theta - V dy_2 \wedge du_3 \wedge du_1 \\ & + dy_3 \wedge du_3 \wedge \theta - V dy_3 \wedge du_1 \wedge du_2,\end{aligned}$$

where y_1, y_2 and y_3 denote the coordinates on $\mathbb{R}^3$. The associated G_2 -metric is given by $g_M = g_X + g_{\mathbb{R}^3}$.

Let $P_i = \Sigma_i \times (\mathbb{R}^+ \cdot v_i) \subset X \times \mathbb{R}^3$ for $i \in \{1, 2, 3\}$, where we regard v_i as a vector in the space $\mathbb{R}^3_{(y_1, y_2, y_3)} \cong \mathbb{R}^3_{(u_1, u_2, u_3)}$. The cylinders P_i are $U(1)$ -invariant associatives in M , namely $\phi|_{P_i} = \text{vol}_{P_i}$, where the volume form vol_{P_i} is defined with respect to the restriction of the Riemannian metric g_M to P_i .

Identifying M as $Z \times \mathbb{R}_{y_3}$, we have $\phi = -dy_3 \wedge \omega - \text{Im}(\Omega)$. For any submanifold L in $X \times \mathbb{R}^2$, we define $\tilde{L} = \{(q, y) \in X \times \mathbb{R}^2 \mid (q, R^{-1}y) \in L\}$. For any $a \in \mathbb{R}$, the submanifold $P = L \times \{a\}$ is an associative submanifold of M if and only if $\tilde{L}$ is a special Lagrangian submanifold in $X \times \mathbb{R}^2$. The upshot is that Theorem 1 follows from Theorem 2.

3 Dimensional reduction and real Monge–Ampère equation

We look for the dimensional reduction on the desired special Lagrangian submanifolds $L \subset X \times \mathbb{R}^2$ with cylindrical ends L_i in the symplectic quotient, where the case $n = 3$ has been conjectured by Donaldson and Scaduto. Under some heuristic assumptions, this leads to a certain real Monge–Ampère equation with a specific Dirichlet boundary condition. The rest of this section will rigorously establish the existence and the properties of the solution.

3.1 The setup

The $U(1)$ -invariance of the asymptotic ends L_i motivates us to search for L among $U(1)$ -invariant submanifolds. The $U(1)$ -moment map associated to the symplectic form ω on Z is u_3 , which must be constant on the $U(1)$ -invariant Lagrangian L , and the constant is fixed to be zero, since it is zero on the asymptotic cylinders. The special Lagrangian conditions $\omega|_L \equiv 0$ and $\text{Im}(\Omega)|_L \equiv 0$ for the $U(1)$ -invariant $L \subset X \times \mathbb{R}^2$ reduce to

$$V du_1 \wedge du_2 - dy_1 \wedge dy_2 = 0 \quad \text{and} \quad du_1 \wedge dy_1 + du_2 \wedge dy_2 = 0,$$

respectively. The dimensionally reduced Lagrangian is

$$L_{\text{red}} := L/U(1) \subset Z_{\text{red}} := u_3^{-1}(0)/U(1).$$

Topologically, Z_{red} can be identified with $\mathbb{R}^4$ with coordinates (u_1, u_2, y_1, y_2) , but equipped with a degenerate reduced Kähler structure. Let

$$\pi_1 : Z_{\text{red}} \rightarrow \mathbb{R}^2_{(u_1, u_2)} \quad \text{and} \quad \pi_2 : Z_{\text{red}} \rightarrow \mathbb{R}^2_{(y_1, y_2)}$$

be the natural projection maps.




Figure 2: $L_{i,\text{red}} \subset \mathbb{R}^4 = \mathbb{R}^2_{(u_1, u_2)} \times \mathbb{R}^2_{(y_1, y_2)}$.

The $U(1)$ -reduction of the cylindrical special Lagrangians L_i results in half-strips $L_{i,\text{red}}$ contained inside the cylinder

$$[p_i, p_{i+1}] \times \{(p_{i+1} - p_i) \cdot y = c_i\} \subset \mathbb{R}^2_{(u_1, u_2)} \times \mathbb{R}^2_{(y_1, y_2)},$$

where $[p_i, p_{i+1}]$ denotes the closed line segment in $\mathbb{R}^2 \times \{0\} \subset \mathbb{R}^3$ that connects the points p_i and p_{i+1} , and we will require the n parameters c_i to sum to zero; see the Appendix. Therefore,

$$\pi_1(L_{i,\text{red}}) = [p_i, p_{i+1}] \subset \mathbb{R}^2_{(u_1, u_2)} \quad \text{and} \quad \pi_2(L_{i,\text{red}}) = \text{a translation of } (\mathbb{R}^+ \cdot \tilde{v}_i) \subset \mathbb{R}^2_{(y_1, y_2)},$$

as shown in Figure 2. Therefore, $\pi_1(\bigcup_i L_{i,\text{red}})$ is the boundary of the convex polygon with vertices $p_1, \dots, p_n$, and $\pi_2(\bigcup_i L_{i,\text{red}})$ is the union of n rays, as shown in Figure 3 for the case $n = 3$.

Graphical case The asymptotic cylindrical requirement motivates us to look for the reduction L_{red} of the conjectural special Lagrangian L , as (the closure of) the graph of a map

$$F : U \subset \mathbb{R}^2_{(u_1, u_2)} \rightarrow \mathbb{R}^2_{(y_1, y_2)}, \quad \text{with } (y_1, y_2) = F(u_1, u_2),$$

where U is the interior of the convex polytope with vertices $p_1, \dots, p_n$.

The projection of L_{red} to $\mathbb{R}^2_{(y_1, y_2)}$ is expected to be a thickening of the union of the n rays, whose shape resembles an amoeba, as illustrated in Figure 4. We will study these shapes more in Section 3.4.





Figure 3: $\pi_1(\bigcup L_{i,\text{red}})$, left, and $\pi_2(\bigcup L_{i,\text{red}})$, right.

Figure 4: L_{red} in $Z_{\text{red}} = \mathbb{R}^4$, in the case $n = 3$.

The equation $du_1 \wedge dy_1 + du_2 \wedge dy_2 = 0$ is equivalent to $\partial y_2 / \partial u_1 = \partial y_1 / \partial u_2$. This implies that we can define a function $\varphi : U \rightarrow \mathbb{R}$ such that $F = \nabla \varphi$, namely

$$y_1 = \frac{\partial \varphi}{\partial u_1} \quad \text{and} \quad y_2 = \frac{\partial \varphi}{\partial u_2}.$$

The equation $V du_1 \wedge du_2 - dy_1 \wedge dy_2 = 0$ reduces to the following real Monge–Ampère equation for φ :

$$(1) \quad \det D^2 \varphi = V = A + \sum_{i=1}^n \frac{1}{2|u - p_i|}.$$

Note that V has singularities at the vertices of the convex polygon U .

Dirichlet boundary condition We now use the expected asymptote of the special Lagrangian L to heuristically motivate the Dirichlet boundary condition on the real Monge–Ampère equation. The rest of the paper will start from the Dirichlet problem and construct the conjectured special Lagrangian L . In Section 4.2, we will see that the Dirichlet boundary condition results in the correct asymptotic behavior.

Notice that the vector $\tilde{v}_i \in \mathbb{R}^2$ is normal to the edge $[p_i, p_{i+1}]$. Since L_{red} is the graph of $\nabla \varphi$ and its asymptote at infinity is the union of L_i , we expect that when $u \in U$ approaches the open line segment (p_i, p_{i+1}) on the boundary of the convex polygon, the normal derivative tends to infinity, while the tangential derivative tends to a constant, $\nabla \varphi \cdot (p_{i+1} - p_i) \rightarrow c_i$. Since $\sum c_i = 0$, we can write $c_i = b_{i+1} - b_i$, and notice that adding a global constant to b_i will be inconsequential. Thus the boundary value of φ on the edge is expected to be *affine linear*, namely

$$(2) \quad \varphi(p_i) = b_i \quad \text{and} \quad \varphi(tp_i + (1-t)p_{i+1}) = tb_i + (1-t)b_{i+1} \quad \text{for all } t \in [0, 1].$$

Remark 5 The real Monge–Ampère equation is invariant under adding an affine linear function. Geometrically, adding a constant to φ has no effect on the special Lagrangian L , and adding a linear function amounts to translating L along some vector in $\mathbb{R}^2_{(y_1, y_2)}$. In particular, for the triangle case $n = 3$, we can reduce to the zero boundary data.

The real Monge–Ampère equation is also invariant under the Euclidean motion of the convex polygon U , with the corresponding change to V . Later, when we analyze the local behavior of φ near an open edge, we will sometimes reduce to the “standard position”, where the open edge is contained in the u_2 -axis, and U lies inside the right half-plane, to simplify notation.

3.2 Solving the Dirichlet problem

Theorem 6 (Dirichlet problem) *Let $b_1, \dots, b_n \in \mathbb{R}$. There is a unique continuous convex function $\varphi : \bar{U} \rightarrow \mathbb{R}$, with boundary data (2), which is smooth in U and solves the real Monge–Ampère equation (1).*

The remainder of this section is dedicated to the proof of this theorem. We use an approximation strategy to deal with the failure of strict convexity of the domain.

Proof Step 1 (approximate solutions): We take a smooth 1-parameter expanding family of strictly convex smooth domains $U_t \subset U$, where $t \in (0, 1)$, converging to $U_1 = U$ as $t \rightarrow 1$, as in Figure 5. The piecewise linear boundary data (2) can be extended to some Lipschitz function $\bar{\phi}$ on $\bar{U}$. We consider the following Dirichlet problem for $\varphi_t : U_t \rightarrow \mathbb{R}$ for each $t \in (0, 1)$:

$$(3) \quad \begin{cases} \det D^2 \varphi_t = V & \text{on } U_t, \\ \varphi_t = \bar{\phi} & \text{on } \partial U_t. \end{cases}$$

Lemma 7 [Rauch and Taylor 1977] *Let $\Omega \subset \mathbb{R}^2$ be a strictly convex domain, $g : \partial\Omega \rightarrow \mathbb{R}$ a continuous function, and μ a nonnegative Borel measure on Ω with $\mu(\Omega) < \infty$. Then, there exists a unique convex function $f \in C(\bar{\Omega})$ such that*

$$(4) \quad \begin{cases} \det D^2 f = \mu & \text{on } \Omega, \\ f = g & \text{on } \partial\Omega. \end{cases}$$

In our case, for every $t \in (0, 1)$, we let $\Omega = U_t$ and define $\mu = V du_1 du_2$. Notice here that V is strictly positive within U_t and uniformly bounded in L^1 ,

$$\int_{U_t} V du_1 du_2 \leq \int_U V du_1 du_2 < +\infty \quad \text{for all } t \in (0, 1).$$




Figure 5: Approximating domains $U_t \rightarrow U$, in the case $n = 3$.

Therefore, by Lemma 7, there is a unique convex function $\varphi_t \in C(\bar{U}_t)$ satisfying (3).

Step 2 (uniform bounds and the limit): By the convexity of φ_t , and the Lipschitz property of $\bar{\phi}$, we obtain the uniform upper bound

$$\varphi_t(u) \leq \bar{\phi}(u) + C \operatorname{dist}(u, \partial U_t) \quad \text{for all } u \in U_t.$$

We now derive a uniform lower bound.

Lemma 8 *There is a uniform constant, independent of t , such that*

$$\varphi_t(u) \geq \bar{\phi}(u) - C \operatorname{dist}(u, \partial U)^{1/2} \quad \text{for all } u \in U_t.$$

Proof As in Remark 5, we put U into the standard position, namely that the boundary edge of interest $[p_i, p_{i+1}]$ lies on the u_2 -axis, and U lies in the right half-plane, as in Figure 6, and without loss of generality, $\bar{\phi} = 0$ on this boundary edge. It suffices to show $\varphi_t(u) \geq -Cu_1^{1/2}$ for some constant C independent of t .

By the Lipschitz property of $\bar{\phi}$, we know $|\bar{\phi}(u)| \leq C_1 u_1$ on U , for some constant C_1 . The boundary data of $\varphi_t + C_1 u_1$ is nonnegative, and it solves the same real Monge–Ampère equation (1).

By the Alexandrov estimate,

$$\max(0, -\varphi_t(u))^2 \leq C_2 \operatorname{diam}(U_t) \cdot \operatorname{dist}(u, \partial U_t) \cdot \int_{U_t} V \, du_1 \, du_2 \leq C_3 u_1 \quad \text{for all } u \in \bar{U}_t,$$

as required. $\square$

Combining the upper and lower bounds, we obtain a uniform bound on the C^0 -norm of φ_t . By the convexity of φ_t , this implies a uniform Lipschitz bound on any fixed compact subset of U , as $t \rightarrow 1$. By Arzelà–Ascoli, we can extract a subsequence of φ_t which C^0 -converges to some continuous function φ on any compact subset of U , which is a viscosity solution to the real Monge–Ampère equation. By passing the upper and lower bounds to the limit as $t \rightarrow 1$, we obtain

$$(5) \quad \bar{\phi}(u) - C \operatorname{dist}(u, \partial U)^{1/2} \leq \varphi(u) \leq \bar{\phi}(u) + C \operatorname{dist}(u, \partial U).$$

Thus φ extends continuously to $\bar{U}$ and achieves the boundary data (2), which agrees with $\bar{\phi}$ on ∂U . The uniqueness of the solution is a standard consequence of the maximum principle.

Step 3 (interior smoothness): Notice V is smooth and strictly positive in U . As a standard fact about real Monge–Ampère equation in dimension two, the solution φ to the real Monge–Ampère equation is smooth in the interior domain U . This fact is the consequence of two standard results: Caffarelli [1990] proved that the singular set must propagate to the boundary along some line segment, while the partial regularity of Mooney [2015] showed that the singular set has codimension one Hausdorff measure zero. $\square$

3.3 Gradient divergence near the edges

In this section, we study the behavior of φ near an open edge of the convex polygon U . The following theorem is essential in proving the smoothness of the special Lagrangian L .

Proposition 9 Let u_* be any point on a boundary edge (p_i, p_{i+1}) in $\partial U \setminus \{p_1, \dots, p_n\}$. Then, as $u \in U$ tends to u_* , the normal and tangential gradient components satisfy

$$\nabla^{\text{normal}} \varphi(u) \rightarrow +\infty \quad \text{and} \quad \nabla \varphi(u) \cdot (p_{i+1} - p_i) \rightarrow c_i.$$

Proof The tangential component converges to a constant by the convexity of φ , the affine linearity of the boundary data, and the boundary continuity estimate (5) by considering the convex function restricted to line segments parallel to the boundary edge.

We now focus on the normal gradient component and place the convex polygon into the standard position by Remark 5, so the open boundary edge (p_i, p_{i+1}) containing u_* lies on the u_2 -axis, the domain is contained in the right half-plane, and the boundary data is zero on this edge. The normal gradient component is just $-\partial_{u_1} \varphi$. Using convexity, $\partial_{u_1} \varphi$ is bounded from the above near u_* .

We suppose for contradiction that $\partial_{u_1} \varphi(u)$ stays bounded for some sequence $u \rightarrow u_*$, and therefore, the gradient at u stays bounded. Then there exists some subgradient for φ at u_* , which must be of the form $(-\Lambda, 0)$ for some $\Lambda \in \mathbb{R}$, because the tangential component is zero. Let

$$\Lambda_0 = \inf\{\Lambda \mid (-\Lambda, 0) \text{ is a subgradient of } \varphi \text{ at some point } v \in (p_i, p_{i+1})\},$$

so in particular $\varphi(u) \geq -\Lambda_0 u_1$ for any $u \in \bar{U}$. Thus $(-\Lambda_0, 0)$ is a subgradient at every point on the open edge, and $\partial_{u_1} \varphi \geq -\Lambda_0$ on U .

We fix a small constant $h > 0$ such that u_* has distance at least $2h$ to the vertices. Let $R(u_*, h, \varepsilon)$ be the rectangle with length h and width $\varepsilon \ll 1$, as shown in Figure 6. By the Monge–Ampère equation, and the strict positivity of V ,

$$(6) \quad C^{-1} h \varepsilon \leq \int_{R(u_*, h, \varepsilon)} V \, du_1 \, du_2 = \int_{\nabla \varphi(R(u_*, h, \varepsilon))} dy_1 \, dy_2.$$

By the convexity of φ and the bound $-\Lambda_0 u_1 \leq \varphi \leq C u_1$, we deduce

$$|\partial_{u_2} \varphi(u)| \leq C(\Lambda_0, h) u_1 \quad \text{for all } u \in R(u_*, h, \varepsilon).$$




Figure 6: Standard position in the case $n = 3$. $R(u_*, h, \varepsilon)$ is the shaded region.

Thus by considering the gradient image,

$$\left| \int_{\nabla \varphi(R(u_*, h, \varepsilon))} dy_1 dy_2 \right| \leq C(\Lambda_0, h) \varepsilon \sup_{R(u_*, h, \varepsilon)} (\partial_{u_1} \varphi(u) + \Lambda_0).$$

Contrasting with (6), for any small $\varepsilon > 0$,

$$\sup_{R(u_*, h, \varepsilon)} (\partial_{u_1} \varphi(u) + \Lambda_0) \geq C(\Lambda_0, h)^{-1}.$$

In the limit $\varepsilon \rightarrow 0$, we can extract some subgradient at some boundary point, which contradicts the minimality of Λ_0 . $\square$

3.4 Solutions near vertices

In this section, we examine the behavior of the gradient of φ near the vertices of the convex polygon, which will be important in studying the smoothness of L . The ideal picture to have in mind, which we justify in this section, is shown in Figure 7 for the case $n = 3$. The Map $F = \nabla \varphi$ takes the bounded gray solid convex polygon to the unbounded gray area. Denote the subgradient sets at the vertices by

$$C_{p_i} = \{y \in \mathbb{R}_{(y_1, y_2)}^2 \mid \varphi(u) - \varphi(p_i) \geq \langle y, u - p_i \rangle \text{ for all } u \in \bar{U}\}.$$

Lemma 10 *The sets C_{p_i} are disjoint convex closed subsets of $\mathbb{R}^2$ contained in the wedge region*

$$C_{p_i} \subset W_{p_i} = \{y \in \mathbb{R}^2 \mid y \cdot (p_{i+1} - p_i) \leq b_{i+1} - b_i\} \cap \{y \in \mathbb{R}^2 \mid y \cdot (p_{i-1} - p_i) \leq b_{i-1} - b_i\}.$$

Remark 11 The wedge region is a translated copy of the wedge

$$W'_{p_i} = \{y \in \mathbb{R}^2 \mid y \cdot (p_{i+1} - p_i) \leq 0\} \cap \{y \in \mathbb{R}^2 \mid y \cdot (p_{i-1} - p_i) \leq 0\},$$

in which the directions of its two extremal rays are specified by the vectors $\tilde{v}_{i-1}$ and $\tilde{v}_i$. For different i , the wedges W'_{p_i} can only intersect along boundary rays, and the intersections between different W_{p_i} have areas bounded by some constant depending only on $\{b_k\}_{k=1}^n$.




Figure 7: Mapping $F = \nabla \varphi : U \subset \mathbb{R}_{(u_1, u_2)}^2 \rightarrow \mathbb{R}_{(y_1, y_2)}^2$.




Figure 8: Adjacent wedges W_{p_i} and $W_{p_{i+1}}$ intersect along a common ray, but nonadjacent wedges may intersect nontrivially inside a compact set.

Proof of Lemma 10 Suppose C_{p_i} and C_{p_j} have a common vector. Then by convexity, this vector is a subgradient for all the points on the line segment $[p_i, p_j]$. If the open segment (p_i, p_j) is in the interior domain U , this would contradict the strict convexity of φ (which follows from the smoothness of the real Monge–Ampère solution), and if (p_i, p_j) is an edge in the boundary, this would imply the existence of subgradient at every point of (p_i, p_j) , contradicting Proposition 9. This proves the disjointness of the C_{p_i} . $\square$

Lemma 12 (image of the gradient) $\nabla\varphi(U) = \mathbb{R}^2 \setminus \left(\bigcup_i C_{p_i} \right)$.

Proof First we claim that $\mathbb{R}^2 = \bigcup_i C_{p_i} \cup \nabla\varphi(U)$. Given any $y \in \mathbb{R}^2$, we consider the graph of the affine linear function $a + \langle y, u \rangle$ on $\bar{U}$, where $a \in \mathbb{R}$ increases from negative infinity. There must be some a for which the graph first touches the graph of the convex function φ . This shows that y is a subgradient at some point on $\bar{U}$. But the divergence of the gradient on the edge (Proposition 9) shows that there is no subgradient at any point on the open edges, so y is a subgradient either at an interior point or at one of the vertices.




Figure 9: Subgradients C_{p_1} , C_{p_2} and C_{p_3} , in the case $n = 3$.

By the strict convexity of φ in U , we see

$$\nabla\varphi(U) \cap C_{p_i} = \emptyset.$$

Thus $\nabla\varphi(U) = \mathbb{R}^2 \setminus (\bigcup_i C_{p_i})$. $\square$

Lemma 13 For any sequence of points $u \in U$ converging to p_i , after passing to a subsequence, either $\nabla\varphi(u) \rightarrow \infty$, or $\nabla\varphi(u)$ converges to a point in ∂C_{p_i} .

Proof Suppose $|\nabla\varphi(u)|$ stays bounded. After passing to a subsequence, $|\nabla\varphi(u)|$ converges to some y . Then y is a subgradient at p_i , so lies in C_{p_i} . However, the limit y has to lie in the closure of $\nabla\varphi(U)$, which is disjoint from the interior of C_{p_i} . Thus the gradient lies on ∂C_{p_i} . $\square$

Lemma 14 gives some asymptotic decay bound on the gradient image $\nabla\varphi(U)$.

Lemma 14 Suppose $y \in \nabla\varphi(U)$ and $|y| \geq 1$. Then, for one of the boundary rays $\mathcal{R}$ of some wedge region W_{p_i} , we have $\text{dist}(y, \mathcal{R}) \leq C/|y|$, where the distance is measured in the Euclidean $\mathbb{R}^2$, and the constant C is independent of y .

In particular, Lemma 14 shows that a scenario similar to the one shown in Figure 10, where $\nabla\varphi(U)$ contains an infinite wedge, cannot happen.

Proof Without loss of generality, we assume that y is large compared to $\max_i |b_i - b_{i+1}|$. The rays $\mathbb{R}_{\geq 0}\tilde{v}_i$ partition $\mathbb{R}^2$ into wedge-shaped regions W'_{p_i} . We choose the direction $\tilde{v}_i$ which minimizes the angle with the direction of y . Notice $\tilde{v}_i$ is parallel to the boundary ray between W_{p_i} and $W_{p_{i+1}}$, and by the choice of $\tilde{v}_i$ and the largeness of y , we see that y must lie in either W_{p_i} or $W_{p_{i+1}}$, and without loss of generality we focus on $y \in W_{p_i}$.

Let q_i be the vertex point of W_{p_i} , and let $\mathcal{R} = \mathbb{R}_{\geq 0}\tilde{v}_i + q_i$, and let y' be the intersection point of the ray $\mathcal{R}$ with the ray $y + \mathbb{R}_{\leq 0}\tilde{v}_{i-1}$. If $y = y'$, then the distance to the ray is zero, and we are done. So without loss of generality $y \neq y'$, and we consider triangle $T(y)$ with vertices at y, y', q_i .





Figure 10: $\nabla\varphi(U)$ cannot contain an infinite wedge W .

Figure 11: Triangle $T(y)$.

By construction the triangle $T(y)$ is contained in W_{p_i} . By elementary geometry, the area of this triangle is comparable to $|y| \operatorname{dist}(y, \mathcal{R})$ when $|y|$ is much larger than $\max_i |b_i - b_{i+1}|$.

Using the partition $\mathbb{R}^2 = \nabla\phi(U) \cup \bigcup_j C_{p_j}$,

$$\begin{aligned}
 (7) \quad \operatorname{Area}(W_{p_i} \setminus C_{p_i}) &\leq \int_{\nabla\phi(U)} dy_1 dy_2 + \operatorname{Area}\left(W_{p_i} \cap \bigcup_{j \neq i} C_{p_j}\right) \\
 &\leq \int_{\nabla\phi(U)} dy_1 dy_2 + \sum_{j \neq i} \operatorname{Area}(W_{p_i} \cap W_{p_j}) \\
 &\leq \int_U V du_1 du_2 + C \leq C.
 \end{aligned}$$

Here the second line uses $C_{p_j} \subset W_{p_j}$, and the third line uses the Monge–Ampère equation and the fact that nonadjacent wedge regions W_{p_j} only intersect within a compact set depending only on $\{b_k\}$ and the direction of the rays. Since $T(y) \subset W_{p_i}$, in particular, $\operatorname{Area}(T(y) \setminus C_{p_i}) \leq C$.

We claim that $C_{p_i} \cap T(y)$ is empty. Suppose $y'' \in C_{p_i} \cap T(y) \subset W_{p_i}$. Then by the convexity of C_{p_i} , and the fact that $y \notin C_{p_i}$ because of $y \in \nabla\phi(U) \subset \mathbb{R}^2 \setminus \bigcup C_{p_i}$, we deduce that the ray starting from y in the direction $y - y''$ lies in the complement of C_{p_i} . Since the direction of $y - y''$ lies in the wedge region W'_{p_i} , by the convexity of C_{p_i} the set $W_{p_i} \setminus C_{p_i}$ contains some infinite wedge region with vertex at y . This is impossible by the finiteness of $\operatorname{Area}(W_{p_i} \setminus C_{p_i})$.

Combining the above,

$$|y| \operatorname{dist}(y, \mathcal{R}) \leq C \operatorname{Area}(T(y)) \leq C,$$

hence $\operatorname{dist}(y, \mathcal{R}) \leq C/|y|$ when $|y|$ is large. $\square$

4 Regularity, asymptotics and topology

This section aims to prove Theorem 2 by producing the desired special Lagrangian from the solution of the real Monge–Ampère equation, and establishing its smoothness, asymptotic properties and topology.

4.1 Smoothness of the special Lagrangian

Let $q : u_3^{-1}(0) \rightarrow Z_{\text{red}}$ be the quotient map, and $L^\circ = q^{-1}(\text{Graph}(\nabla\varphi|_U))$. By the smoothness of the real Monge–Ampère solution φ , clearly L° is a smooth special Lagrangian submanifold in Z , diffeomorphic to $\mathbb{D}^2 \times S^1$; however, it is not a closed subset of Z . Let L be the current of integration defined by L° , so that L contains also points in the closure of L° . The goal of this section is to prove the following theorem.

Theorem 15 *L is represented by a smooth submanifold of Z without boundary.*

The proof of Theorem 15 follows from Lemmas 16 and 20.

Lemma 16 *L is a closed integral current, $\partial L = 0$. Moreover, any point in the support of L either lies on L° or lies above some vertex p_i .*

Proof We first show that L has locally finite mass inside any given ball in Z . Since the gradient $\nabla\varphi$ diverges to infinity near the open edges (p_i, p_{i+1}) , we only need to show that mass cannot accumulate near the vertices p_i . Now by the special Lagrangian condition, for any $U(1)$ -invariant compact set $K \subset Z$,

$$\text{Mass}(L \cap K) = \int_{L^\circ \cap K} \text{Re}(\Omega) = 2\pi \int_{(L^\circ \cap K)/U(1)} (du_1 \wedge dy_2 - du_2 \wedge dy_1).$$

Since φ is a smooth strictly convex function, both integrands are positive, and the maps $(u_1, u_2) \mapsto (u_1, y_2)$ and $(u_1, u_2) \mapsto (y_1, u_2)$ are injective, so $\int_{(L^\circ \cap K)/U(1)} du_1 \wedge dy_2$ is simply the Lebesgue measure of the (u_1, y_2) projection of $(L^\circ \cap K)/U(1)$, and likewise with the second integrand. By the boundedness of the range of $(u_1, y_2), (u_2, y_1)$ on $L^\circ \cap K$, we deduce that the integral is finite.

The support of L is contained in the closure of L° . Since the gradient $\nabla\varphi$ diverges to infinity, there cannot be any limiting point in the support of L whose projection to $\mathbb{R}_{(u_1, u_2)}^2$ lies above the open edges (p_i, p_{i+1}) . Thus, the only points added in the closure lie above the vertices p_i , and by Lemma 13, they lie above $\{p_i\} \times \partial C_{p_i}$.

Let ψ be any smooth and compactly supported test 2-form on Z , and for small $r > 0$, let χ_r be a cutoff function supported in the union of $B_{X^4}(p_i, r) \times \mathbb{R}^2 \subset Z$, with value one on the union of $B_{X^4}(p_i, r/2) \times \mathbb{R}^2$, and whose g_X -gradient is bounded by Cr^{-1} . Since $(1 - \chi_r)\psi$ is supported away from the $U(1)$ -fixed locus, integration by parts gives

$$\int_L \chi_r d\psi = - \int_L d\chi_r \wedge \psi = - \int_{L^\circ} d\chi_r \wedge \psi,$$

hence

$$\left| \int_L \chi_r d\psi \right| \leq C(\psi)r^{-1} \sum_{i=1}^n \text{Mass}(L^\circ \cap \text{supp}(\psi) \cap B_{X^4}(p_i, r) \times \mathbb{R}^2).$$

Now y_1 and y_2 are bounded within the support of ψ , and from the Gibbons–Hawking ansatz we know $|u - p_i| \leq Cr^2$ on $B_{X^4}(p_i, r)$. The same argument as in the local finiteness of measure now gives a

bound $|\int_L \chi_r d\psi| \leq C(\psi)r$. Taking the limit $r \rightarrow 0$, we deduce that $\int_L d\psi = 0$ for any test 2-form, which means $\partial L = 0$. $\square$

Now, we prove the smoothness. L is a special Lagrangian integral current and, in particular, a minimal integral current. Therefore, there exists a tangent cone at each point x on L . The proof of smoothness is based on the following implication of Allard's regularity theorem.

Proposition 17 *A point $x \in \text{supp}(L)$ is a smooth point if and only if every tangent cone $N \subset \mathbb{C}^3$ at x is a 3-plane with multiplicity one.*

Let $q(x) \in \{p_i\} \times \partial C_{p_i}$. Any tangent cone $N \subset \mathbb{C}^3$ at x is a $U(1)$ -invariant tangent cone in $\mathbb{C}^3$. To prove every tangent cone of L is a 3-plane with multiplicity one, we employ Joyce's classification [2005] of $U(1)$ -invariant special Lagrangian cones in $\mathbb{C}^3$.

Proposition 18 [Joyce 2005; Haskins 2004] *Let N be a special Lagrangian cone without boundary in $\mathbb{C}^3$ invariant under the $U(1)$ -action given by*

$$e^{i\theta} : (z_1, z_2, z_3) \rightarrow (e^{i\theta} z_1, e^{-i\theta} z_2, z_3) \quad \text{for } e^{i\theta} \in U(1),$$

where $N \setminus \{0\}$ is connected. Then there exists $A \in [-1, 1]$ and functions $w : \mathbb{R} \rightarrow (-1, 1)$, and $\alpha, \beta : \mathbb{R} \rightarrow \mathbb{R}$ satisfying the system of differential equations

$$\begin{aligned} \left(\frac{dw}{dt}\right)^2 &= 4((1-w)^2(1+2w) - A^2), & \frac{d\alpha}{dt} &= \frac{A}{1-w}, \\ \frac{d\beta}{dt} &= \frac{-2A}{1+2w}, & (1-w)(1+2w)^{\frac{1}{2}} \cos(2\alpha + \beta) &= A, \end{aligned}$$

such that, away from points $(z_1, z_2, z_3) \in \mathbb{C}^3$ with $z_j = 0$ for some j , we may locally write N in the form $\Phi(r, s, t)$ with $r > 0$ and $s, t \in \mathbb{R}$, where

$$\Phi : (r, s, t) \mapsto (re^{i(\alpha(t)+s)} \sqrt{1-w(t)}, e^{i(\alpha(t)-s)} \sqrt{1-w(t)}, re^{i\beta(t)} \sqrt{1+2w(t)}),$$

and exactly one of the following holds:

- $A = 1$. Then, N is the $U(1)^2$ -invariant special Lagrangian $\mathbb{T}^2$ -cone


$$\{(re^{i\theta_1}, re^{i\theta_2}, re^{i\theta_3}) \mid r > 0, \theta_1, \theta_2, \theta_3 \in \mathbb{R}, \theta_1 + \theta_2 + \theta_3 = 0\}.$$

- $A = -1$. Then, N is the $U(1)^2$ -invariant special Lagrangian $\mathbb{T}^2$ -cone

$$\{(re^{i\theta_1}, re^{i\theta_2}, re^{i\theta_3}) \mid r > 0, \theta_1, \theta_2, \theta_3 \in \mathbb{R}, \theta_1 + \theta_2 + \theta_3 = \pi\}.$$

- $A = 0$. Then, for some $\phi \in (-\pi, \pi]$, either $N = \Pi_\phi^+$, or $N = \Pi_\phi^-$, or N is the singular union $\Pi_\phi^+ \cup \Pi_\phi^-$, where $\Pi_\phi^\pm$ are the special Lagrangian 3-planes

$$\Pi_\phi^+ = \{(z, ie^{-i\phi}\bar{z}, re^{i\phi}) \mid z \in \mathbb{C}, r \in \mathbb{R}\} \quad \text{and} \quad \Pi_\phi^- = \{(z, -ie^{-i\phi}\bar{z}, re^{i\phi}) \mid z \in \mathbb{C}, r \in \mathbb{R}\}.$$

Figure 12: The π_1 -projection of a tangent cone at p_i .

- $0 < |A| < 1$. Then, the function $w(t)$ may be written in terms of the Jacobi elliptic functions. It is nonconstant and periodic in t with period T depending only on A , and $2\alpha + \beta$ is also nonconstant and periodic in t with period T .

We proceed by ruling out every possibility on Joyce’s list except 3-planes. We do this using the following lemma.

Lemma 19 Let $N \subset \mathbb{C}^3$ be a special Lagrangian tangent cone of L at x , where $q(x) \in p_i \times \partial C_{p_i}$. Let $U = \pi_1(N/U(1)) \subset \mathbb{R}^2_{(u_1, u_2)}$. The set U is a subset of the infinite wedge with vertex p_i and two rays along the direction $\overrightarrow{p_i p_{i+1}}$ and $\overrightarrow{p_i p_{i-1}}$.

Proof The image of L in $\mathbb{R}^2_{(u_1, u_2)}$ is a convex polygon. Let W be the infinite wedge with vertex at p_i , and two boundary rays $\overrightarrow{p_i p_{i+1}}$ and $\overrightarrow{p_i p_{i-1}}$. In particular its opening angle is less than π . The π_1 -projections of all the special Lagrangians obtained by rescaling L around the base point x , are all contained in W , so by passing to the limit, the same holds for the projection of the tangent cone N to $\mathbb{R}^2_{(u_1, u_2)}$. $\square$

Lemma 20 Let $N \subset \mathbb{C}^3$ be a special Lagrangian tangent cone at x in L , where $q(x) \in \{p_i\} \times \partial C_{p_i}$. Then, N is a 3-plane with multiplicity 1.

Proof We apply Joyce’s classification to the connected components of the tangent cone, and rule out every other possibility of the list of Proposition 18.

Step 1 ($U(1)^2$ -invariant $\mathbb{T}^2$ -cone): The cases $A = 1$ and $A = -1$ are similar, so we focus on $A = 1$. In this scenario,

$$\pi_1(N/U(1)) = \{r^2 e^{-i\theta_3} \mid r > 0, \theta_3 \in \mathbb{R}\},$$

and $r^2 e^{-i\theta_3}$ can take any value in $\mathbb{C} = \mathbb{R}^2_{(u_1, u_2)}$; consequently, $\pi_1(N/U(1)) = \mathbb{R}^2$. In particular, $\pi_1(N)$ is not subset of a wedge with angle less than 180° , which contradicts Lemma 19.

Step 2 (union of two 3-planes): Suppose $\text{supp}(N)$ contains $\Pi_+^\phi \cup \Pi_-^\phi$. We have

$$u_1(z, i e^{-i\varphi} \bar{z}) = \text{Re}(z_1 z_2) = |z|^2 \sin(\varphi) \quad \text{and} \quad u_2(z, i e^{-i\varphi} \bar{z}) = \text{Im}(z_1 z_2) = |z|^2 \cos(\varphi).$$

Hence

$$\pi_1(\Pi_+^\phi) = \{r^2(\sin(\varphi), \cos(\varphi)) \mid r \in \mathbb{R}\} \quad \text{and} \quad \pi_1(\Pi_-^\phi) = \{-r^2(\sin(\varphi), \cos(\varphi)) \mid r \in \mathbb{R}\}.$$

Therefore, $\pi_1(N/U(1)) = \{R(\sin(\varphi), \cos(\varphi)) \mid R \in \mathbb{R}\}$ forms a line. In particular, it is not subset of a wedge with angle less than 180° , contradicting Lemma 19.

Step 3 (multiplicity and graphicality): The special Lagrangian L projects with degree one to the (y_1, u_2) -plane. As in Theorem 5.6. in [Joyce 2005], for any given cutoff function $\chi \geq 0$ on $\mathbb{R}^2$,

$$\int_L \chi(y_1, u_2)(\theta du_1 - V du_2 du_3) \wedge dy_1 = 2\pi \int_S \chi(y_1, u_2) dy_1 du_2,$$

where S is the projection of $L/U(1)$ on the (y_1, u_2) -plane. By passing to a tangent cone N of L at x , we get

$$\int_N \chi(\operatorname{Re}(z_3), \operatorname{Im}(z_1 z_2)) \operatorname{Re}(dz_1 \wedge dz_2) \wedge \operatorname{Re}(dz_3) \leq 2\pi \int_{\mathbb{R}^2} \chi(y_1, u_2) dy_1 du_2.$$

In a small open neighborhood around a generic point $(y_1, u_2) \in \mathbb{R}^2$, the tangent cone N is nonsingular and divides into k components, therefore for χ supported near this point we get

$$\int_N \chi(\operatorname{Re}(z_3), \operatorname{Im}(z_1 z_2)) \operatorname{Re}(dz_1 \wedge dz_2) \wedge \operatorname{Re}(dz_3) = 2\pi k \int_{\mathbb{R}^2} \chi(y_1, u_2) dy_1 du_2.$$

Therefore, $k = 0$ or 1 . Thus the preimage of $(y_1, u_2) \in \mathbb{R}^2$ contains at most one point counting multiplicity. The same conclusion holds for the projections to (y_2, u_1) , and indeed any choice of direction of y and the corresponding direction of u specified by the partial Legendre transform. This forces that there is at most one connected component in N , and it has multiplicity one.

Furthermore, following Proposition 5.5 in [Joyce 2005], in the Jacobi elliptic cone case $0 < |A| < 1$, the (y_1, u_2) projection has degree greater than one; hence this case is ruled out. The only remaining possibility is that N is a flat 3-plane with multiplicity one. $\square$

4.2 Asymptotics of the special Lagrangians

In this section, we prove that L has the expected asymptotic behavior.

Theorem 21 *Asymptotically near infinity, L is an exponentially small C^k graph over the model special Lagrangian cylinders $\bigcup_{i=1}^n L_i$.*

We divide the proof into a few steps.

Step 1 (Cauchy–Riemann type equation): The region of L close to spatial infinity must project to a small neighborhood of one of the edges. Moreover, Lemma 14 shows that the projection to the (y_1, y_2) plane is close to some ray $\mathcal{R}$, which we can without loss of generality take to be inside $\{(p_{i+1} - p_i) \cdot y = c_i\}$. Up to rotating and translating the coordinates, we reduce the problem to the standard position, so the edge lies in $\{u_1 = 0\}$, and the ray is simply $\{y_2 = 0, y_1 < 0\}$, and the region lies in $\{y_1 < 0, |y_1| \gg 1\} \cap \{u_2 \in [a, b]\}$, where $[a, b]$ describes the boundary edge.

Via the partial Legendre transform, we see the reduced Lagrangian L_{red} is graphical over the (u_2, y_1) variables, except at the vertices. The special Lagrangian condition can be rewritten as the Cauchy–Riemann type equation

$$\frac{\partial y_2}{\partial y_1} = -\frac{\partial u_1}{\partial u_2}, \quad \frac{\partial y_2}{\partial u_2} = V \frac{\partial u_1}{\partial y_1}.$$

Then Lemma 14 provides the preliminary decay estimate $|y_2| \leq C/|y_1|$. Moreover, since φ is a convex function with $C^{1/2}$ boundary modulus of continuity, we obtain

$$|y_1|u_1 = -y_1u_1 = -\partial_{u_1}\varphi(u_1, u_2)u_1 \leq \varphi(0, u_2) - \varphi(u_1, u_2) \leq Cu_1^{1/2},$$

so we have the preliminary estimate $u_1 \leq C|y_1|^{-2}$. Thus L_{red} is a C^0 -small graph over the (u_2, y_1) variables, with decay rate estimate $O(|y_1|^{-1})$.

Now the Cauchy–Riemann type equation is quasilinear elliptic, so away from the singular locus $V = \infty$ corresponding to $u_2 \in \{a, b\}$, we can bootstrap the smallness of the C^0 -norm to smallness of the C^k -norm. More geometrically, the asymptotic model is the half-cylinder $\Sigma_i \times \mathbb{R}^+$. For any given $\epsilon > 0$, we can find some $R \gg 1$ such that on the region $\{y_1 < -R, u_2 \in [a + \epsilon, b - \epsilon]\}$ away from the vertex, our special Lagrangian is a C^k -small perturbation of the model with C^1 -norm bounded by ϵ .

Step 2 (quantitative smoothness): We need to prove quantitative smoothness estimate for L_{red} near the vertex region, and in our coordinates this means u_2 close to the endpoints a, b of the edges. This is based on Allard’s regularity theorem. Notice the ambient manifold has bounded geometry in our region of interest.

Proposition 22 (Allard’s regularity) *There exists a universal constant $\epsilon_0 \ll 1$, and some small fixed r_0 depending on the ambient manifold such that the following holds. Let X be an n -dimensional multiplicity one stationary integral varifold inside the coordinate ball $B(p, r)$ with $r \leq r_0$. Assume that p lies on the support of X , and the volume $\mathcal{H}^n(X \cap B(p, r)) \leq (\omega_n + \epsilon_0)r^n$. Then $X \cap B(p, r/2)$ is a $C^{1,\alpha}$ graph over the tangent plane through p , with the $C^{1,\alpha}$ -norm bounded by $\frac{1}{100}$.*

Without loss of generality p lies in $a \leq u_2 \leq a + \epsilon$, where $\epsilon \ll r_0$ will be fixed later. We compute the volume on a small geodesic ball of radius $r \leq r_0$,

$$\text{Mass}(L_{\text{red}} \cap B(p, r)) = \int_{L_{\text{red}} \cap B(p, r)} \text{Re}(\Omega) = 2\pi \int_{L_{\text{red}} \cap B(p, r)} (du_1 \wedge dy_2 - du_2 \wedge dy_1).$$

On the integration region away from $a \leq u_2 \leq a + \epsilon$, by the smallness of the C^1 -norms of y_2 and u_1 , we see that the integral contribution is bounded by $\omega_3 r^3(1 + O(\epsilon))$. On the other hand, the contribution from the region $\{a \leq u_2 \leq a + \epsilon\} \cap B(p, r)$ can be estimated by the same idea as in Lemma 16: the integral of $dy_1 \wedge du_2$ is the Lebesgue area of the (y_1, u_2) projection, which gives a contribution bounded by $O(\epsilon r)$. The integral of $du_1 \wedge dy_2$ is the Lebesgue area of the projection to the (u_1, y_2) plane. Since u_1, y_2 are both $O(|y|^{-1})$, this contribution is bounded by $O(|y|^{-2})$. In total,

$$\text{Mass}(L_{\text{red}} \cap B(p, r)) \leq \omega_3 r^3(1 + C\epsilon) + C\epsilon r + C|y|^{-2}.$$

Now for fixed r , we can choose $\epsilon \ll 1$ and $R \gg 1$ such that for $y_1 < -R$, all the remainder terms can be dominated by $\epsilon_0 r^3$, so that

$$\text{Mass}(L_{\text{red}} \cap B(p, r)) \leq \omega_3(1 + \epsilon_0)r^3.$$

Thus we can apply Allard regularity to deduce the quantitative smoothness of L close to infinity. Combining with the C^0 -decay, it follows that L is a C^k -small graph over the model half-cylinder. The C^k -norms of u_1, y_2 are both bounded by $O(|y|^{-1})$. In particular,

$$\int_{\{y_1 < -1, u_2 \in [a, b]\} \cap L} |\nabla y_2|^2 < +\infty.$$

Step 3 (exponential decay): It remains to improve the C^k decay to exponential decay.

We first notice that y_2 defines a harmonic function on L . This is because y_2 has zero Hessian on the ambient manifold, and L is a minimal surface.

Lemma 23 *For any sufficiently large $R \geq R_0$,*

$$(8) \quad \int_{y_1 \leq -R} |\nabla y_2|^2 \leq Ce^{-\gamma R},$$

for a constant $\gamma > 0$ independent of R .

Proof We already know that the end of L is a small $C^{k,\alpha}$ -graph over the cylindrical model. We define

$$F(R) = \int_{\{y_1 \leq -R\} \cap L} |\nabla y_2|^2.$$

Since y_2 is harmonic, for any $c \in \mathbb{R}$, we can apply the divergence theorem to deduce

$$F(R) = \int_{L \cap \{y_1 = -R\}} (y_2 - c) \nabla y_2 \cdot \vec{n},$$

whence by Cauchy–Schwarz and the Poincaré inequality,

$$F(R) \leq \left(\int_{L \cap \{y_1 = -R\}} |\nabla y_2|^2 \right)^{1/2} \left(\int_{L \cap \{y_1 = -R\}} |y_2 - c|^2 \right)^{1/2} \leq C \int_{L \cap \{y_1 = -R\}} |\nabla y_2|^2.$$

Thus $-CF' \geq F$, which implies the exponential decay. $\square$

Since L is already C^k -regular, we can bootstrap this to C^k exponential decay for y_2 . Using the elliptic system, it is then easy to see that L is asymptotically an exponentially small graph over the model cylindrical special Lagrangian.

4.3 Topology of the constructed special Lagrangians

We conclude by proving the last component of Theorem 2, thereby confirming the Donaldson–Scaduto conjecture, Theorem 1.

Theorem 24 *L is homeomorphic to an n -holed 3-sphere.*




Figure 13: Boundary components C_i and sets U_0 and U_i , in the case $n = 3$.

Proof Let L' be the 3-manifold obtained by truncating the ends of L at a sufficiently large distance R , denoted by L_R , and sealing them by adding 3-balls, resulting in a closed 3-manifold. We show $L' \cong S^3$.

First argument (employing the Poincaré conjecture): We prove that L' is simply connected.

Note that L is fibered over L_{red} , where L_{red} is homeomorphic to $\mathbb{D} \setminus \{a_1, \dots, a_n\}$ for n distinct boundary points $a_1, \dots, a_n$. The fiber over any interior point $z \in L_{\text{red}}^\circ$ is a copy of $U(1)$, and the fibers collapse to a point when $z \in \partial L_{\text{red}}$.

Let $C_1, \dots, C_n$ be the boundary components of L_{red} . Recalling the projection map $q : u_3^{-1}(0) \rightarrow Z_{\text{red}}$, let $V_i := q^{-1}(U_i) \subset L$ be the preimage of an open neighborhood U_i of the boundary component C_i in L_{red} such that $U_i \cap U_j = \emptyset$ when $i \neq j$. Each V_i is homeomorphic to $\mathbb{D} \times \mathbb{R}$. Let U'_i be another open neighborhood of the boundary component C_i in L_{red} slightly smaller than U_i . Let V_0 be the open set in L defined as the preimage of $U_0 := L_{\text{red}} \setminus \bigcup_{i=1}^n \bar{U}'_i$. The set V_0 is homeomorphic to $\mathbb{D} \times U(1)$. The configuration of open sets in L_{red} is shown in Figure 13.

Let $x_0 \in V_0 \cap V_1$ be the base point. We have $\pi_1(V_0, x_0) \cong \mathbb{Z}$, with a generator presented by a curve encircling the $U(1)$ -fiber based at x_0 . Furthermore, $\pi_1(V_1, x_0) = \{0\}$ and $\pi_1(V_0 \cap V_1, x_0) \cong \mathbb{Z}$, and the inclusion map $V_0 \cap V_1 \rightarrow V_0$ takes the generator of $\pi_1(V_0 \cap V_1, x_0)$ to the generator of $\pi_1(V_0, x_0)$. Therefore, by Van Kampen's theorem, $\pi_1(V_0 \cup V_1, x_0) = \{0\}$. Applying Van Kampen's theorem again repeatedly and adding V_i inductively yields $\pi_1(V_0 \cup V_1 \cup \dots \cup V_n, x_0) = \{0\}$, namely $\pi_1(L) = 0$. Consequently, $\pi_1(L') = 0$, and therefore, by the Poincaré conjecture, L' is a 3-sphere.

Second argument (without employing the Poincaré conjecture): We can extend q to obtain $q : L' \rightarrow L'_{\text{red}}$, where $L'_{\text{red}} \cong \mathbb{D}^2$ is the truncated version of L_{red} capped off with n half-discs, where the preimage of any interior point under q is a copy of S^1 , and the preimage of any boundary point is a single point. In other words, L'_{red} is an S^1 -bundle over the interior of $\mathbb{D}^2$, which collapses to a point above each boundary point.

Let C be an embedded circle in L'_{red} which divides it into two regions: the interior D_1 and the exterior D_2 . Let $\Sigma = q^{-1}(C)$, and $\Sigma_i = q^{-1}(D_i)$ for $i \in \{1, 2\}$. The Heegaard surface $\Sigma = \mathbb{T}^2$, and

handlebodies $\Sigma_1 \cong \overline{\mathbb{D}}^2 \times S^1$ and $\Sigma_2 \cong S^1 \times \overline{\mathbb{D}}^2$, leading to a genus 1 Heegaard decomposition of L' , where the gluing map of this decomposition maps the meridian of Σ_1 to the longitude of Σ_2 and the longitude of Σ_1 to the meridian of Σ_2 . This description characterizes the genus 1 Heegaard decomposition of S^3 . $\square$

Appendix: Parameter count

Our main construction depends on the parameters $\{b_k\}_{k=1}^n$. Changing these parameters by the same additive constant amounts to adding a constant to the Dirichlet solution φ , which does not affect the special Lagrangian. In total, our construction depends on $n - 1$ real parameters. Furthermore, as shown in [Habibi Esfahani 2022], each of these special Lagrangians with fixed asymptotics are rigid.

We now provide an alternative perspective on why the deformations of the n asymptotically cylindrical ends are subject to one additional constraint. We suppose that the n asymptotic half cylinders are contained in $L_i = \Sigma_i \times l_i$, where we recall that $\Sigma_i = \pi^{-1}[p_i, p_{i+1}]$ and $l_i \subset \mathbb{R}^2$ is defined by

$$l_i = \{y \in \mathbb{R}^2 \mid y \cdot (p_{i+1} - p_i) = c_i\} \quad \text{for some } c_i \in \mathbb{R}.$$

Lemma 25 *Let L be an asymptotically cylindrical special Lagrangian in $X \times \mathbb{R}^2$ with asymptotes ends L_i for $i = 1, \dots, n$. Then, we have $\sum_{i=1}^n c_i = 0$.*

Proof We can find a primitive for $\text{Im}(\Omega)$,

$$\text{Im}(\Omega) = \omega_1 \wedge dy_1 + \omega_2 \wedge dy_2 = d\lambda, \quad \text{with } \lambda = y_1 \omega_1 + y_2 \omega_2.$$

We apply Stokes' theorem to L truncated at a very large distance R , which is a manifold L_R with boundary diffeomorphic to $\bigcup_{i=1}^n \Sigma_i$. Since L is a special Lagrangian,

$$0 = \int_{L_R} \text{Im}(\Omega) = \sum_{i=1}^n \int_{\Sigma_i} \lambda.$$

By definition, L is an exponentially small graph over L_i in the asymptotic regime. Consequently, we can evaluate the boundary integrals on L_i with an error of $O(e^{-cR})$, which disappears in the limit when $R \rightarrow +\infty$. We have

$$\left(\int_{\Sigma_i} \omega_1, \int_{\Sigma_i} \omega_2 \right) = 2\pi(p_{i+1} - p_i) \in \mathbb{R}^2,$$

hence

$$\sum_{i=1}^n \int_{\Sigma_i} \lambda = \sum_{i=1}^n (y \cdot (p_{i+1} - p_i) + O(e^{-cR})) = \sum_{i=1}^n (c_i + O(e^{-cR})),$$

and letting $R \rightarrow +\infty$, we deduce $\sum_{i=1}^n c_i = 0$. $\square$

Acknowledgements

We thank our common teacher Prof. Simon Donaldson for suggesting this problem to us, and the referee for useful comments on the previous draft. Habibi Esfahani thanks Mark Haskins and Rafe Mazzeo for discussions on related topics. Li was supported by a Clay Maths Research Fellowship.

References

- [Bera et al. 2025] **G Bera, S Habibi Esfahani, Y Li**, *Uniqueness in the local Donaldson–Scaduto conjecture*, Int. Math. Res. Not. 2025:16 (2025) art. id. rnaf245 MR
- [Caffarelli 1990] **L A Caffarelli**, *A localization property of viscosity solutions to the Monge–Ampère equation and their strict convexity*, Ann. of Math. (2) 131:1 (1990) 129–134 MR
- [Donaldson 2017] **S Donaldson**, *Adiabatic limits of co-associative Kovalev–Lefschetz fibrations*, from “Algebra, geometry, and physics in the 21st century” (D Auroux, L Katzarkov, T Pantev, Y Soibelman, Y Tschinkel, editors), Progr. Math. 324, Springer (2017) 1–29 MR
- [Donaldson and Scaduto 2022] **S Donaldson, C Scaduto**, *Associative submanifolds and gradient cycles*, from “Surveys in differential geometry 2019: differential geometry, Calabi–Yau theory, and general relativity, II” (H-D Cao, S-T Yau, editors), Surv. Differ. Geom. 24, Int., Boston, MA (2022) 39–65 MR
- [Habibi Esfahani 2022] **S Habibi Esfahani**, *Monopoles, singularities and hyperkahler geometry*, PhD thesis, State University of New York at Stony Brook (2022) MR Available at <https://www.proquest.com/docview/2719007490>
- [Haskins 2004] **M Haskins**, *Special Lagrangian cones*, Amer. J. Math. 126:4 (2004) 845–871 MR
- [Joyce 2005] **D Joyce**, *$U(1)$ -invariant special Lagrangian 3-folds, III: Properties of singular solutions*, Adv. Math. 192:1 (2005) 135–182 MR
- [Mooney 2015] **C Mooney**, *Partial regularity for singular solutions to the Monge–Ampère equation*, Comm. Pure Appl. Math. 68:6 (2015) 1066–1084 MR
- [Rauch and Taylor 1977] **J Rauch, B A Taylor**, *The Dirichlet problem for the multidimensional Monge–Ampère equation*, Rocky Mountain J. Math. 7:2 (1977) 345–364 MR

SAMAN HABIBI ESFAHANI saman_habibiesfahani@fas.harvard.edu
Department of Mathematics, Duke University, Durham, NC, United States
 Current address: *Center of Mathematical Sciences and Applications, Harvard University*
Cambridge, MA, United States

YANG LI yl454@cam.ac.uk
Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, United States
 Current address: *Department of Pure Mathematics and Mathematical Statistics, Cambridge University*
Cambridge, United Kingdom

Proposed: András I Stipsicz
 Seconded: Simon Donaldson, Ciprian Manolescu

Received: February 11, 2024
 Revised: April 14, 2025

Milnor’s invariants for knots and links in closed orientable 3-manifolds

RYAN STEES

In his 1957 paper, John Milnor introduced a collection of invariants for links in S^3 detecting higher-order linking phenomena by studying lower central quotients of link groups and comparing them to those of the unlink. These invariants, now known as Milnor’s $\bar{\mu}$ -invariants, were later shown to be topological link concordance invariants and have since inspired decades of consequential research. Milnor’s invariants have many interpretations, and there have been numerous attempts to extend them to other settings. In this paper, we extend Milnor’s invariants to topological concordance invariants of knots and links in general closed orientable 3-manifolds. These invariants unify and generalize all previous versions of Milnor’s invariants in dimension 3, including Milnor’s original invariants for links in S^3 .

1. Introduction	983
2. Preliminaries	986
3. Main results	989
4. n -basings	1003
5. Concordance invariance	1007
6. Characterization and realization for h_n	1009
7. Characterization and realization for θ_n and $\bar{\mu}_n$	1024
8. Relationships to previous work	1035
9. Future directions	1042
Appendix. Properties of $X_n(L)$ and $\iota_n(L)$	1046
Acknowledgements	1048
References	1048

1 Introduction

In his 1954 PhD thesis and subsequent paper, John Milnor [1957] introduced a collection of integer-valued invariants for links in S^3 , now known as *Milnor’s $\bar{\mu}$ -invariants*, which inductively compare lower central quotients of link groups to those of the unlink. The simplest $\bar{\mu}$ -invariants are the classical linking numbers first defined by C F Gauss [1877], and the higher invariants are higher-order linking numbers. Much is known about Milnor’s $\bar{\mu}$ -invariants; in particular, Casson [1975] used work of Stallings [1965] to show that Milnor’s invariants are topological link concordance invariants. Milnor’s invariants are ubiquitous, having interpretations in terms of Massey products [Turaev 1976; Porter 1980], homotopy-theoretic data [Orr 1989], surface systems in S^3 [Cochran 1990], the Kontsevich integral [Habegger and Masbaum 2000] and twisted Whitney towers in the 4-ball [Conant et al. 2014]. We summarize the characterizing property of these $\bar{\mu}$ -invariants in the following theorem.

MSC2020: primary 57K10, 57K31; secondary 57K30, 57K40, 57N65.

© 2026 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

Milnor's theorem [1957] *Let $L \subset S^3$ be an m -component link, and let U be the m -component unlink, with link groups π and F , respectively. Suppose the n -th lower central quotients of π and F are isomorphic. Then the following three statements are equivalent:*

- (1) *The $\bar{\mu}$ -invariants of L of length n agree with those of U (that is, they vanish).*
- (2) *The $(n+1)$ -st lower central quotients of π and F are isomorphic.*
- (3) *The $\bar{\mu}$ -invariants of L of length $n+1$ are well-defined.*

There have been numerous attempts to extend Milnor's $\bar{\mu}$ -invariants to other settings; see [Miller 1995; Heck 2011; Kuzbary 2024; Cha and Orr 2024]. This paper extends Milnor's invariants to concordance invariants of knots and links in any closed orientable 3-manifold. Our work unifies and generalizes all previous extensions of Milnor's invariants in dimension 3.

1.1 New contributions

To any link L in a closed orientable 3-manifold M , we associate a tower of spaces $\{X_n(L)\}_{n \geq 1}$, well-defined up to homotopy equivalence, which come equipped with canonical homotopy classes of maps $\iota_n(L) : M \rightarrow X_n(L)$. We have the following diagram:

$$\begin{array}{ccccccc}
 & & M & & & & \\
 & \swarrow \iota_{n+1}(L) & \downarrow \iota_n(L) & \searrow \iota_1(L) & & & \\
 \cdots & \longrightarrow & X_{n+1}(L) & \longrightarrow & X_n(L) & \longrightarrow & \cdots \longrightarrow X_1(L)
 \end{array}$$

We will show that a concordance between two links $L, L' \subset M$ induces homotopy equivalences $X_n(L') \xrightarrow{\cong} X_n(L)$ for all n and that the maps $\iota_n(L') : M \rightarrow X_n(L')$ and $\iota_n(L) : M \rightarrow X_n(L)$ correspond via these homotopy equivalences.

To define invariants, we will fix a link $L \subset M$ and compare other links to L inductively. Given some link $L' \subset M$, and assuming we have homotopy equivalences $X_k(L') \xrightarrow{\cong} X_k(L)$ for all $k \leq n$, we obstruct the existence of a homotopy equivalence $X_{n+1}(L') \xrightarrow{\cong} X_{n+1}(L)$ by comparing the maps $\iota_n(L') : M \rightarrow X_n(L')$ and $\iota_n(L) : M \rightarrow X_n(L)$ via the homotopy equivalence $X_n(L') \xrightarrow{\cong} X_n(L)$. Through this inductive comparison, we extract a sequence of topological concordance invariants h_n , which we call the *lower central homotopy invariants* of L' relative to L . These invariants are homotopy classes of maps in $[M, X_n(L)]_0 / \text{Aut}(\pi / \Gamma_n, \partial)$, where $\text{Aut}(\pi / \Gamma_n, \partial)$ is a group acting on $[M, X_n(L)]_0$ we discuss in Section 4.1.

An approach pioneered by K Orr [1989] inspired our approach to defining the lower central homotopy invariants; we define them in a similar manner to invariants of P Heck [Heck 2011]. We paraphrase the main results involving these invariants:

- *Invariance:* Theorem A, stated in Section 3.1 and proved in Section 5.1, asserts that the h_n are topological concordance invariants.

- *Characterization*: Theorem B, stated in Section 3.2 and proved in Section 6.2, states that the h_n characterize n -cobordism, a geometric relationship between links. Roughly speaking, L' and L'' are n -cobordant if they cobound a surface in $M \times [0, 1]$ which looks like a concordance to the n -th lower central quotients. Theorem B states that $h_n(L') = h_n(L'')$ if and only if L' is n -cobordant to L'' .
- *Realization*: Theorem C, stated in Section 3.3 and proved in Section 6.3, gives exact conditions under which a homotopy class in $[M, X_n(L)]_0$ is realized as the invariant of some link $L' \subset M$.

The invariants h_n also determine two additional new invariants we will discuss and explore:

- The invariants θ_n , which we call the *lower central homology invariants*, are homology classes in $H_3(X_n(L))/\text{Aut}(\pi/\Gamma_n, \partial)$. They are images of the fundamental class $[M]$ under the maps the h_n induce on homology.
- The invariants $\bar{\mu}_n$, which we refer to as *Milnor invariants*, are a further reduction of the h_n . Like Milnor's $\bar{\mu}$ -invariants, they determine the lower central quotients of link groups one step at a time.

The invariants θ_n and $\bar{\mu}_n$ are useful in the greatly extended setting of $\mathbb{Z}[G]$ -homology concordance for suitable groups G , where we allow the links we compare to live in different 3-manifolds. A $\mathbb{Z}[G]$ -homology concordance between $L \subset M$ and $L' \subset M'$ is a concordance in a $\mathbb{Z}[G]$ -homology cobordism between M and M' . We now summarize the main results involving these reductions of h_n :

- *Invariance*: Theorem A', stated in Section 3.5 and proved in Section 5.2, asserts that the θ_n and $\bar{\mu}_n$ are topological homology concordance invariants.
- *Characterization*: Theorem D, stated in Section 3.6 and proved in Section 7.1, states that the $\bar{\mu}_n$ characterize algebraic information we call an n -basing. Roughly speaking, an n -basing for L' relative to L is an isomorphism on n -th lower central quotients of link groups which preserves the peripheral structures of meridians and longitudes. We state a concise version of Theorem D below for the sake of direct comparison with Milnor's theorem. See Section 3.6 for the definition of the $\bar{\mu}_n$ and the full statement of Theorem D.
- *Lifting property*: Corollary D₂, stated and proved in Section 3.6, asserts that when comparing links in the same 3-manifold the invariant $\bar{\mu}_n$ vanishes if and only if the lower central homotopy invariant h_n lifts to a realizable homotopy class in $[M, X_{n+1}(L)]_0$.
- *Realization*: Theorem E, stated in Section 3.7 and proved in Section 7.2, gives exact conditions under which a homology class $\theta \in H_3(X_n(L))$ is realized as the integral homology concordance invariant $\theta = \theta_n(L')$ of some link L' in some 3-manifold M' .

Theorem D Fix an m -component link $L \subset M$. Let $L' \subset M$ be another m -component link, and let π and π' be the link groups of L and L' , respectively. Suppose L' admits an n -basing relative to L , i.e., suppose the n -th $\pi_1(M)$ -lower central quotients of π and π' are admissibly isomorphic. Then the following three statements are equivalent:

- (1) We have $\bar{\mu}_n(L') = \bar{\mu}_n(L)$.

- (2) The link L' admits an $(n+1)$ -basing relative to L , that is, the $(n+1)$ -st $\pi_1(M)$ -lower central quotients of π and π' are admissibly isomorphic.
- (3) The invariant $\bar{\mu}_{n+1}(L')$ is well-defined.

We emphasize a few features of the invariants h_n and their reductions θ_n and $\bar{\mu}_n$, which relate them to previous extensions of Milnor's $\bar{\mu}$ -invariants:

- The vanishing of the invariants h_n , taken in appropriate contexts, implies the vanishing of all previous versions of Milnor's invariants in dimension 3, including Milnor's original $\bar{\mu}$ -invariants.
- For links in S^3 and every $n \geq 2$, the invariant $\bar{\mu}_n$ relative to the unlink is equivalent to Milnor's (total) $\bar{\mu}$ -invariant of length $n + 1$.
- For links in S^3 and for $n \geq 2$, the invariants h_n relative to the unlink are equivalent to invariants of Orr [1989].
- The invariants h_n are defined for knots and can be nontrivial for knots in 3-manifolds $M \neq S^3$; see [Miller 1995; Heck 2011; Kuzbary 2024].
- For empty links, the invariants θ_n and $\bar{\mu}_n$ are precisely the Cha–Orr invariants [2024] of homology cobordism of 3-manifolds, “Milnor invariants of 3-manifolds”. Cha and Orr show that these invariants exhibit highly nontrivial behavior.

Notation Throughout this paper, $L \subset M$, $L' \subset M'$ and $L'' \subset M''$ will denote m -component links in the closed orientable 3-manifolds M , M' and M'' . All links are assumed to be oriented and ordered. We will usually assume the 3-manifolds are oriented, so that meridians of links and fundamental classes of the 3-manifolds are well-defined. The 3-manifolds M , M' and M'' will be equipped with surjective homomorphisms φ , φ' and φ'' from their respective fundamental groups onto a fixed group G . Often, we will have $M = M' = M''$, $G = \pi_1(M)$ and $\varphi = \varphi' = \varphi'' = \text{id}_{\pi_1(M)}$, and we will indicate when this is not the case. We will denote the regular neighborhood of a submanifold X in a manifold Y by νX and the exterior of X in Y by E_X . The link group of $L \subset M$, the fundamental group $\pi_1(E_L)$, will be denoted by π . We will let $\Gamma \trianglelefteq \pi$ be the normal subgroup $\Gamma = \ker(\pi \twoheadrightarrow \pi_1(M) \twoheadrightarrow G)$. We will use analogous notation for the links L' and L'' .

2 Preliminaries

2.1 Concordance in 3-manifolds

A non-simply-connected 3-manifold M allows for embedding phenomena which cannot occur in S^3 . In such 3-manifolds, we do not have an analogue of the unknot or unlink in the sense that there is no natural choice of knot or link to which all others can be reasonably compared. For instance, knots which are concordant must be freely homotopic; in other words, concordant knots represent the same element of $[S^1, M]$. A free homotopy between concordant knots is given by the composition of the concordance with projection to M :

$$S^1 \times [0, 1] \hookrightarrow M \times [0, 1] \xrightarrow{\text{proj}} M.$$

To develop concordance invariants, we will fix some link $L \subset M$ and compare other links relative to L . The fixed link L will play the role of the unlink.

The nontriviality of $\pi_1(M)$ also implies that knots $K \subset M$ may not all have preferred longitudes. A *longitude* for a knot K is a section of the sphere bundle of K , that is, a parallel of K in $\partial\nu K$. We also refer to the homology class of such a curve as a longitude. Removing a regular neighborhood of a concordance between two knots $K, K' \subset M$ yields a cobordism rel boundary between knot exteriors E_K and $E_{K'}$ which identifies $\partial\nu K$ and $\partial\nu K'$ and induces a correspondence between the meridians and longitudes of K and K' . If K is nullhomologous in M , then one can pick a preferred longitude by taking the pushoff of K which is nullhomologous in the exterior E_K . A concordance induces a correspondence between preferred longitudes if they exist; however, in the absence of preferred longitudes, there is *a priori* some ambiguity in how $\partial\nu K'$ and $\partial\nu K$ are to be identified when attempting to obstruct the existence of a concordance between K' and K . In other words, the absence of a preferred longitude prevents a canonical identification of $\partial\nu K$ with $S^1 \times S^1$. For knots representing torsion classes in $H_1(M)$, a rational-valued linking number remedies this problem (see [Friedl et al. 2019] or [Raoux 2020]), but it is not clear what to do in the case where $[K] \in H_1(M)$ has infinite order. The techniques of Section 3 resolve the aforementioned issue. To define the lower central homotopy invariants, we fix an identification of longitudes and then understand how to account for the dependence on this chosen identification.

In the last few decades, there has been renewed interest in extending the study of concordance to general 3-manifolds [Miller 1995; Schneiderman 2003; Heck 2011; Davis et al. 2018; Kuzbary 2024], as well as to generalized versions of concordance. These include homology concordance [Goldsmith 1978; Levine 2016; Zhou 2021; Hom et al. 2022] and almost-concordance [Celoria 2018; Yildiz 2018; Friedl et al. 2019; Nagel et al. 2019], concordance up to local knotting. Relatively little is known in these settings. For instance, the following conjecture of Celoria and Friedl–Nagel–Orson–Powell remains open in many cases:

Almost-concordance conjecture [Celoria 2018; Friedl et al. 2019] *Every free homotopy class $x \in [S^1, M]$ in a closed orientable 3-manifold $M \neq S^3$ which does not admit a dual 2-sphere contains infinitely many concordance classes of knots modulo local knotting.*

M Powell has suggested that the work in this paper could be useful for attacking many of the remaining cases; indeed, in [Steas 2025] we consider the Almost-concordance conjecture in a large class of open cases; namely, we consider nontrivial $x \in [S^1, M]$ in aspherical M .

2.2 Lower central series

Recall that the (classical) *lower central series* of a group π is the series of normal subgroups

$$\cdots \trianglelefteq \pi_{n+1} \trianglelefteq \pi_n \trianglelefteq \cdots \trianglelefteq \pi_2 \trianglelefteq \pi_1 = \pi,$$

where $\pi_2 = [\pi, \pi]$, the commutator subgroup of π , and $\pi_{n+1} = [\pi, \pi_n]$ in general. We call the quotients π/π_n the *lower central quotients* of π .

We recall a fundamental result of Stallings and Dwyer concerning lower central quotients and group homology, which will play an important role in proving some of the main results of this work. For an abbreviated proof, see [Cha and Orr 2024].

Stallings–Dwyer theorem [Stallings 1965; Dwyer 1975] *Suppose $f : A \rightarrow B$ is a group homomorphism inducing an isomorphism $H_1(A) \xrightarrow{\cong} H_1(B)$. For any positive integer n , f induces an isomorphism $f_n : A/A_n \xrightarrow{\cong} B/B_n$ if and only if f induces a surjection $H_2(A) \twoheadrightarrow H_2(B)/K_j(B)$ for all $j < n$, where $K_j(B) = \ker(H_2(B) \rightarrow H_2(B/B_j))$.*

Now fix a group G , and suppose a group π admits a surjective homomorphism $\varphi : \pi \twoheadrightarrow G$. We consider a finer lower central series, the G -lower central series, defined as

$$\cdots \trianglelefteq \Gamma_{n+1} \trianglelefteq \Gamma_n \trianglelefteq \cdots \trianglelefteq \Gamma_2 \trianglelefteq \Gamma \trianglelefteq \pi,$$

where $\Gamma = \ker \varphi$. This lower central series was also used in work of Heck [2011]. We recover the classical lower central series when $G = \{1\}$. The quotients π/Γ_n are the G -lower central quotients of π . The G -lower central quotients of the link group π of a link $L \subset M$, for suitable groups G , will play a central role in our study of concordance. Most often, we will take $G = \pi_1(M)$ and φ the surjective homomorphism $\pi \twoheadrightarrow \pi_1(M)$ induced by the inclusion $j : E_L \hookrightarrow M$.

2.3 Homology properties of concordance

Cappell and Shaneson [1974] observed that a concordance exterior is a $\mathbb{Z}[\pi_1(M)]$ -homology cobordism rel boundary. In particular, if L and L' are concordant, we have an isomorphism

$$H_*(E_{L'}; \mathbb{Z}[\pi_1(M)]) \xrightarrow{\cong} H_*(E_L; \mathbb{Z}[\pi_1(M)]),$$

along with an identification of $\partial E_{L'}$ and ∂E_L preserving meridians, orientations and orderings of components. We will call such an identification *admissible*.

Apart from the case of knots in S^3 , the $\mathbb{Z}[\pi_1(M)]$ -homology of E_L , along with boundary data, can provide nontrivial information about the concordance class of L . (On the other hand, E_K is an integral homology S^1 for any knot $K \subset S^3$, where the meridian and longitude of K represent 1 and 0, respectively, in $H_1(S^1) \cong \mathbb{Z}$.) Using the Stallings–Dwyer theorem, we see that the $\pi_1(M)$ -lower central quotients of link groups are concordance invariants.

Corollary 2.1 *A concordance C between L and L' induces isomorphisms over $\pi_1(M)$ on $\pi_1(M)$ -lower central quotients $\pi'/\Gamma'_n \xrightarrow{\cong} \pi/\Gamma_n$ for all n .*

Proof Let $\bar{E}_{L'}$, $\bar{E}_C$ and $\bar{E}_L$ denote the $\pi_1(M)$ -covers of $E_{L'}$, E_C , and E_L , respectively. Consider the inclusion-induced isomorphisms

$$H_*(E_{L'}; \mathbb{Z}[\pi_1(M)]) \xrightarrow{\cong} H_*(E_C; \mathbb{Z}[\pi_1(M)]) \xleftarrow{\cong} H_*(E_L; \mathbb{Z}[\pi_1(M)]),$$

which yield isomorphisms

$$H_*(\bar{E}_{L'}; \mathbb{Z}) \xrightarrow{\cong} H_*(\bar{E}_C; \mathbb{Z}) \xleftarrow{\cong} H_*(\bar{E}_L; \mathbb{Z}).$$

Let $\pi^C = \pi_1(E_C)$ and $\Gamma^C = \ker(\pi^C \twoheadrightarrow \pi_1(M \times [0, 1]))$. Then $\bar{E}_{L'}$, $\bar{E}_C$ and $\bar{E}_L$ have fundamental groups Γ' , Γ^C and Γ , respectively. By the Stallings–Dwyer theorem, we have isomorphisms on lower central quotients $\Gamma'/\Gamma'_n \xrightarrow{\cong} \Gamma^C/\Gamma_n^C \xleftarrow{\cong} \Gamma/\Gamma_n$ for all n . The result then follows from the five lemma. $\square$

Corollary 2.1 motivates the definitions of the lower central homotopy invariants h_n given in Section 3.1. The h_n will obstruct the existence of certain isomorphisms between $\pi_1(M)$ -lower central quotients of link groups. As in Milnor's original work, they will compare the $\pi_1(M)$ -lower central quotients of link groups to those of a fixed link $L \subset M$.

3 Main results

In this section, we present our extension of Milnor's invariants to links in closed orientable 3-manifolds. For the duration of Sections 3.1–3.3, all links will live in a fixed 3-manifold M , and for a link group π we will take $\Gamma = \ker(\pi \twoheadrightarrow \pi_1(M))$, so that the groups π/Γ_n are the $\pi_1(M)$ -lower central quotients of π . We will indicate when these assumptions change.

3.1 Lower central homotopy invariants

Recall from Section 1.1 that we will associate to a link $L \subset M$ a tower of spaces $\{X_n(L)\}_{n \geq 1}$, well-defined up to homotopy equivalence, with canonical homotopy classes of maps $\iota_n(L) : M \rightarrow X_n(L)$ as in the diagram

$$\begin{array}{ccccccc} & & M & & & & \\ & \swarrow \iota_{n+1}(L) & \downarrow \iota_n(L) & \searrow \iota_1(L) & & & \\ \cdots & \longrightarrow & X_{n+1}(L) & \longrightarrow & X_n(L) & \longrightarrow & \cdots \longrightarrow X_1(L) \end{array}$$

To define the concordance invariants h_n , we fix a link $L \subset M$ and compare other links to L . For another link $L' \subset M$, we inductively define homotopy classes of maps, depending on L' , from M into the tower $\{X_n(L)\}$. The n -th invariant h_n of L' relative to L is a homotopy class of maps $h_n(L', \phi) \in [M, X_n(L)]_0$ into the n -th level of this tower. If L' is concordant to L , we can define such a homotopy class for every n .

Consider the diagram

$$\mathcal{D}_n(L) = \begin{array}{c} E_L \xrightarrow{p_n} K(\pi/\Gamma_n, 1) \\ j \downarrow \\ M \end{array}$$

where $j : E_L \hookrightarrow M$ is the inclusion and p_n is induced by the canonical projection $\pi \twoheadrightarrow \pi/\Gamma_n$.

Define $X_n(L)$ to be the standard homotopy pushout of $\mathcal{D}_n(L)$, that is, the based space

$$X_n(L) = \frac{K(\pi/\Gamma_n, 1) \sqcup (E_L \times [0, 1])^\times \sqcup M}{p_n(x) \sim (x, 0), j(x) \sim (x, 1)},$$

where $(E_L \times [0, 1])^\times$ is the reduced cylinder on E_L . The standard homotopy pushout square equips the space $X_n(L)$ with a canonical based homotopy class of maps $\iota_n(L) \in [M, X_n(L)]_0$ represented by

the inclusion $M \hookrightarrow X_n(L)$. This inclusion is a canonical extension of $E_L \hookrightarrow M_{p_n}^\times$, where $M_{p_n}^\times$ is the reduced mapping cylinder of p_n . We may interpret the homotopy pushout $X_n(L)$ as:

- The space $M_{p_n}^\times \cup_{E_L} M$. In particular, $X_n(L)$ is obtained from M by attaching cells of dimensions 2 and higher.
- The space $M_{p_n}^\times \cup_{\partial E_L} \nu L$ obtained by gluing νL to the reduced mapping cylinder $M_{p_n}^\times$.

The canonical projections $p_{m,n} : \pi / \Gamma_m \twoheadrightarrow \pi / \Gamma_n$ for $m \geq n$ induce morphisms of diagrams $\mathcal{D}_m(L) \rightarrow \mathcal{D}_n(L)$ over $\pi_1(M)$, which induce based maps $\psi_{m,n}(L) : X_m(L) \rightarrow X_n(L)$. We will often suppress notation and write $\psi_{m,n}$. Via these maps $\psi_{m,n}$, the collection of spaces $\{X_n(L)\}$ forms a tower over $\pi_1(M)$:

$$\begin{array}{ccccccc} & & M & & & & \\ & \swarrow \iota_{n+1}(L) & \downarrow \iota_n(L) & \searrow \iota_1(L) & & & \\ \cdots & \xrightarrow{\psi_{n+2,n+1}} & X_{n+1}(L) & \xrightarrow{\psi_{n+1,n}} & X_n(L) & \xrightarrow{\psi_{n,n-1}} & \cdots \xrightarrow{\psi_{2,1}} X_1(L) \end{array}$$

For a list of well-definedness, naturality and compatibility properties the $X_n(L)$, $\iota_n(L)$ and $\psi_{m,n}$ satisfy, see the Appendix.

Now fix a link $L \subset M$. The *lower central homotopy invariants* relative to L will be maps $h_n(L', \phi) \in [M, X_n(L)]_0$. The following notion of *n-basing* gives a method of producing a homotopy class of maps depending on some other link $L' \subset M$ into the n -th level $X_n(L)$ of the tower $\{X_n(L)\}$ corresponding to the fixed link L .

Definition 3.1 Fix an m -component link $L \subset M$, and let $L' \subset M$ be another m -component link. An *n-basing* for L' relative to L is an ordered pair (ϕ, ϕ_∂) such that

- (1) $\phi_\partial : H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L)$ is an admissible isomorphism, and
- (2) $\phi : \pi' / \Gamma'_n \xrightarrow{\cong} \pi / \Gamma_n$ is an isomorphism over $\pi_1(M)$ such that ϕ_∂ is *compatible* with ϕ , that is, given any basing τ' for L' , there exists a basing τ for L such that the following diagram commutes:

$$\begin{array}{ccc} \pi_1(\partial E_{L'} \cup \tau') & \xrightarrow[\phi_{\partial*}]{\cong} & \pi_1(\partial E_L \cup \tau) \\ \downarrow \pi' & & \downarrow \pi \\ \pi' / \Gamma'_n & \xrightarrow[\phi]{\cong} & \pi / \Gamma_n \end{array}$$

Condition (2) in the above definition is satisfied for any basing τ' if and only if it is satisfied for a single basing τ' . Recall that an isomorphism $H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L)$ is *admissible* if it preserves meridians, orientations and orderings of components. We call a homeomorphism $\partial E_{L'} \rightarrow \partial E_L$ *admissible* if it induces an admissible isomorphism. We often suppress notation and write ϕ instead of (ϕ, ϕ_∂) . We discuss n -basings in further detail in Section 4. For now, we list some of their relevant properties:

- If L' is concordant to L , an n -basing exists for every n ; see Corollary 4.3.
- An n -basing induces a k -basing for all $k \leq n$.
- The existence of a 1-basing is equivalent to the statement that, when based appropriately, corresponding components of L' and L agree in $\pi_1(M)$; see the discussion at the beginning of Section 2.1.

Our terminology is inspired by the case of links in S^3 , where a choice of meridians for a link (a basing of the link), modulo the n -th lower central subgroup of its link group, is equivalent to a choice of isomorphism between the n -th lower central quotient of the link group and that of the unlink, assuming such an isomorphism exists. In Section 8.1, we prove that for links in S^3 the existence of an n -basing for a link relative to the unlink is equivalent to the statement that Milnor's $\bar{\mu}$ -invariants of length n vanish.

An n -basing for a link L' relative to L contains sufficient information to define the lower central homotopy invariant h_n . We postpone the proof of the following lemma until the end of this section.

Lemma 3.2 *An n -basing ϕ for L' relative to L determines a well-defined based homotopy class of maps of pairs $(E_{L'}, \partial E_{L'}) \rightarrow (M_{p_n}^\times, \partial E_L)$, restricting to an admissible homeomorphism $\partial E_{L'} \xrightarrow{\cong} \partial E_L$, which extends canonically to a based homotopy class of maps of pairs*

$$h_n(L', \phi) : (M, \nu L') \rightarrow (X_n(L), \nu L),$$

which may be taken to map L' homeomorphically to L , all over $\pi_1(M)$.

Definition 3.3 Fix an m -component link $L \subset M$, and let $L' \subset M$ be another m -component link. Suppose L' admits an n -basing (ϕ, ϕ_∂) relative to L . The n -th lower central homotopy invariant of the pair (L', ϕ) relative to L is the homotopy class $h_n(L', \phi) \in [M, X_n(L)]_0$ induced by the n -basing ϕ ; see Lemma 3.2.

With our notation, $h_n(L, \text{id}) = \iota_n(L)$, where $\text{id} = (\text{id}_{\pi/\Gamma_n}, \text{id}_{H_1(\partial E_L)})$. By Lemma 3.2, $h_n(L', \phi)$ is a well-defined homotopy class, but it may depend on the choice of n -basing ϕ . This dependence is measured by the group of *self- n -basings* of L , written $\text{Aut}(\pi/\Gamma_n, \partial)$, which consists of all n -basings of L relative to itself. The group $\text{Aut}(\pi/\Gamma_n, \partial)$ is a subgroup of $\text{Aut}(\pi/\Gamma_n) \times \text{Aut}(H_1(\partial E_L))$ and acts on the set $[M, X_n(L)]_0$ by post-composition with homotopy self-equivalences of $X_n(L)$. See Section 4.1 for more details. To remove the indeterminacy introduced by the choice of n -basing, we take the value of $h_n(L', \phi)$ in the orbit space $[M, X_n(L)]_0/\text{Aut}(\pi/\Gamma_n, \partial)$.

Definition 3.4 Let L and L' be as in the previous definition. The n -th lower central homotopy invariant of L' relative to L is the image $h_n(L')$ of the homotopy class $h_n(L', \phi)$ in the orbit space $[M, X_n(L)]_0/\text{Aut}(\pi/\Gamma_n, \partial)$.

We call n the *length* of the invariant h_n . For conciseness, we will often refer to $h_n(L')$ as the n -th *h-invariant* of L' . To compare the link L' to the fixed link L , we compare $h_n(L', \phi)$ and $h_n(L')$ to the distinguished classes $h_n(L, \text{id})$ and $h_n(L)$ associated to the fixed link L .

Definition 3.5 We say the n -th *h-invariant* for L' relative to L *vanishes* if $h_n(L', \phi) = h_n(L, \text{id})$ in $[M, X_n(L)]_0$ or $h_n(L') = h_n(L)$ in $[M, X_n(L)]_0/\text{Aut}(\pi/\Gamma_n, \partial)$.

Theorem A, which we prove in Section 5, asserts the concordance invariance of the lower central homotopy invariants h_n . Note that we may use h_n not only to compare other links L' to the fixed link L , but to compare links L' and L'' relative to the fixed link L .

Theorem A *The lower central homotopy invariants h_n are invariants of concordance. More precisely, fix an m -component link $L \subset M$. Suppose the m -component links $L', L'' \subset M$ are concordant. Then the following statements hold for all n :*

- (1) *There is a canonical n -basing (ϕ', ϕ'_∂) for L'' relative to L' induced by the concordance.*
- (2) *An n -basing for L' relative to L exists if and only if an n -basing for L'' relative to L exists, and $h_n(L')$ is defined if and only if $h_n(L'')$ is defined.*
- (3) *If (ϕ, ϕ_∂) is an n -basing for L' relative to L , then $h_n(L', \phi) = h_n(L'', \phi \circ \phi')$, where ϕ' is the n -basing in (1).*
- (4) *If $h_n(L')$ and $h_n(L'')$ are defined using any n -basings relative to L , then*

$$h_n(L') = h_n(L'') \in [M, X_n(L)]_0 / \text{Aut}(\pi / \Gamma_n, \partial).$$

In particular, if L' is concordant to the fixed link L , then $h_n(L')$ is defined and vanishes for all n .

By construction, if the invariant $h_n(L', \phi)$ is defined, then $\psi_{n,k} \circ h_n(L', \phi) = h_k(L', \phi_k)$ for all $k \leq n$, where ϕ_k is the k -basing induced by the n -basing ϕ . Thus, the n -th lower central homotopy invariant determines the k -th invariant for every $k \leq n$. We will eventually show in Corollary D₁ that the invariants h_n are defined inductively for $n \geq 2$. If the n -th invariant vanishes, then the $(n+1)$ -st invariant is defined. In fact, we will characterize precisely when the $(n+1)$ -st invariant is defined; see Section 3.6.

Proof of Lemma 3.2 Suppose L' admits an n -basing (ϕ, ϕ_∂) relative to L . Since ϕ_∂ is an admissible isomorphism, it corresponds to a unique (up to isotopy) admissible homeomorphism $h_\partial : \partial E_{L'} \rightarrow \partial E_L$. Choose any basing τ' for L' . Since ϕ_∂ is compatible with ϕ , there exists a basing τ for L such that h_∂ may be extended to a map $\partial E_{L'} \cup \tau' \rightarrow \partial E_L \cup \tau$ where the following diagram commutes:

$$\begin{array}{ccc} \pi_1(\partial E_{L'} \cup \tau') & \xrightarrow[\phi_{\partial*}]{\cong} & \pi_1(\partial E_L \cup \tau) \\ \downarrow & & \downarrow \\ \pi' / \Gamma'_n & \xrightarrow[\phi]{\cong} & \pi / \Gamma_n \end{array}$$

Postcomposing with the inclusion $\partial E_L \cup \tau \hookrightarrow M_{p_n}^\times$, we now have the extension problem over $\pi_1(M)$:

$$\begin{array}{ccccc} \partial E_{L'} \cup \tau' & \longrightarrow & M_{p_n}^\times & \longrightarrow & B\pi_1(M) \\ \downarrow & \nearrow ?\exists & \nearrow & & \\ E_{L'} & & & & \end{array}$$

The space $M_{p_n}^\times$ is a $K(\pi/\Gamma_n, 1)$ and ϕ is an isomorphism over $\pi_1(M)$, so there is no obstruction to extending to a map $h : E_{L'} \rightarrow M_{p_n}^\times$ inducing ϕ . The map h is the desired map of pairs $h : (E_{L'}, \partial E_{L'}) \rightarrow (M_{p_n}^\times, \partial E_L)$ over $\pi_1(M)$.

We will next show that the homotopy class of this homotopy equivalence h given above depends only on the n -basing (ϕ, ϕ_∂) . Suppose $h' : (E_{L'}, \partial E_{L'}) \rightarrow (M_{p_n}^\times, \partial E_L)$ is another map of pairs over $\pi_1(M)$ defined using the same n -basing. We first define the homotopy between h and h' on $\partial E_{L'}$. Since both h_∂ and $h'|_{\partial E_{L'}}$ are homeomorphisms $\partial E_{L'} \xrightarrow{\cong} \partial E_L$ which induce ϕ_∂ , they are homotopic via some homotopy H_∂ . To show h and h' are homotopic, we solve the extension problem over $\pi_1(M)$ below:

$$\begin{array}{ccc} (E_{L'} \times \{0, 1\}) \cup (\partial E_{L'} \times [0, 1]) & \xrightarrow{(h \sqcup h') \cup H_\partial} & M_{p_n}^\times \longrightarrow B\pi_1(M) \\ \downarrow & \nearrow \text{?} \exists & \\ E_{L'} \times [0, 1] & & \end{array}$$

Suppose τ'_0 and τ_0 are the basings for L' and L used to define the map h , as above, and τ'_1 and τ_1 are the basings used to define the map h' . The basings τ_0 and τ_1 for L differ by a wedge of circles α , so that $\alpha\tau_0 \simeq \tau_1$ rel endpoints. Similarly, τ'_0 and τ'_1 differ by a wedge of circles α' , so that $\alpha'\tau'_0 \simeq \tau'_1$. We have $\phi_*(\alpha') \simeq \alpha$ by construction. Let $\alpha_0 = h(\alpha')$. In $M_{p_n}^\times$, $\alpha_0 \simeq \alpha$.

We extend our map to the trace of a homotopy $H_\tau \subset E_{L'} \times [0, 1]$ between $\alpha'\tau'_0$ and τ'_1 over a homotopy from $h(\alpha'\tau'_0) = \alpha_0\tau_0$ to $h'(\tau'_1) = \tau_1$. Note that $\alpha_0\tau_0 \simeq \alpha\tau_0 \simeq \tau_1$ rel endpoints. We now have a based map $(E_{L'} \times \{0, 1\}) \cup (\partial E_{L'} \times [0, 1]) \cup H_\tau \rightarrow M_{p_n}^\times$. We extend to all of $E_{L'} \times [0, 1]$ as in the previous extension problem. There is no obstruction to extending since $M_{p_n}^\times$ is a $K(\pi/\Gamma_n, 1)$ and h and h' both induce the isomorphism ϕ , which is a homomorphism over $\pi_1(M)$. Thus, h and h' are based homotopic as maps of pairs.

To conclude the proof, we show h extends over $\pi_1(M)$ to a canonical homotopy class of maps of pairs $\bar{h} : (M, \nu L') \rightarrow (X_n(L), \nu L)$. We show that h and h' above both extend canonically and the canonical extensions are homotopic. The homotopy seen above between h and h' is a homotopy of pairs sending $\partial E_{L'} \times [0, 1] \rightarrow \partial E_L$. We may view this as a homotopy of maps from the sphere bundle of $L' \subset M'$ to the sphere bundle of $L \subset M$. There is a canonical extension of this homotopy to a homotopy of disk bundles $\nu L' \times [0, 1] \rightarrow \nu L$ given by coning on the fibers of the sphere bundles. Since $X_n(L) = M_{p_n}^\times \cup_{\partial E_L} \nu L$ and $M = E_{L'} \cup_{\partial E_{L'}} \nu L'$, we may extend the homotopy $E_{L'} \times [0, 1] \rightarrow M_{p_n}^\times$ constructed above to a homotopy $M \times [0, 1] \rightarrow X_n(L)$. This yields a canonical homotopy class of maps of pairs $\bar{h} : (M, \nu L') \rightarrow (X_n(L), \nu L)$, which is still a map over $\pi_1(M)$ because the 2-cells we have attached are glued along elements in Γ' and Γ . Just as h may be taken to restrict to a homeomorphism of sphere bundles $h_\partial : \partial E_{L'} \xrightarrow{\cong} \partial E_L$, $\bar{h}$ may be taken to restrict to a homeomorphism of disk bundles $h_\nu : \nu L' \xrightarrow{\cong} \nu L$ sending L' to L . $\square$

3.2 n -cobordism of links

The lower central homotopy invariants provide obstructions to the existence of surfaces which successively approximate a concordance between links in $M \times [0, 1]$. We call such surfaces n -cobordisms. Orr [1989] introduced and studied such surfaces for links in S^3 .

Definition 3.6 Two m -component links $L, L' \subset M$ are n -cobordant if there exists a properly embedded orientable m -component surface $\Sigma = \bigsqcup_i \Sigma_i \subset M \times [0, 1]$ such that for each i :

- (1) $\Sigma_i \cap (M \times \{0\}) = L_i$ and $\Sigma_i \cap (M \times \{1\}) = L'_i$.
- (2) For any basing of Σ_i determined by a basing of L_i , the images of $\pi_1(L_i)$ and $\pi_1(\Sigma_i)$ in $\pi_1(E_\Sigma)$ agree modulo the n^{th} $\pi_1(M)$ -lower central subgroup, that is,

$$\text{im}(\pi_1(L_i) \rightarrow \pi \rightarrow \pi^\Sigma \twoheadrightarrow \pi^\Sigma / \Gamma_n^\Sigma) = \text{im}(\pi_1(\Sigma_i) \rightarrow \pi^\Sigma \twoheadrightarrow \pi^\Sigma / \Gamma_n^\Sigma),$$

where $\pi^\Sigma = \pi_1(E_\Sigma)$ and $\Gamma^\Sigma = \ker(\pi^\Sigma \twoheadrightarrow \pi_1(M))$, and where the inclusions $\Sigma_i \hookrightarrow E_\Sigma$ and $L_i \hookrightarrow E_L$ are induced by a trivialization of $\nu\Sigma$.

In other words, an n -cobordism is a surface that looks like a concordance to the n -th $\pi_1(M)$ -lower central quotients. A concordance is an n -cobordism for all n . Condition (2) implies that for all i the images of the longitudes of L_i and L'_i , when based appropriately, agree in $\pi^\Sigma / \Gamma_n^\Sigma$: they cobound a copy of Σ_i on $\partial\nu\Sigma$ and are therefore homologous in E_Σ . Since the images of $\pi_1(L_i)$ and $\pi_1(\Sigma_i)$ in $\pi^\Sigma / \Gamma_n^\Sigma$ agree, and because the image of $\pi_1(L_i)$ is cyclic and therefore abelian, the longitudes agree in $\pi^\Sigma / \Gamma_n^\Sigma$. We prove the following property of n -cobordism in Section 6:

Proposition 3.7 An n -cobordism between the links L and L' induces a canonical n -basing of L' relative to L and a canonical isomorphism $\pi' / \Gamma'_{n+1} \xrightarrow{\cong} \pi / \Gamma_{n+1}$.

We will also prove the following theorem which states that the lower central homotopy invariants are complete invariants of n -cobordism. Note again that the invariants h_n can compare two links L' and L'' over the fixed link L . A comparison between the h -invariants of L' and L'' is still meaningful even if they do not agree with the h -invariant of L . Heck [2011] proved a version of this theorem for knots in prime 3-manifolds in the case where we compare some knot to the fixed knot.

Theorem B Fix an m -component link $L \subset M$. Let $L', L'' \subset M$ be m -component links, and suppose there exist n -basings (ϕ', ϕ'_∂) and $(\phi'', \phi''_\partial)$ for L' and L'' , respectively, relative to L . Then the following statements are equivalent:

- (1) We have $h_n(L', \phi') = h_n(L'', \phi'')$.
- (2) The links L' and L'' are n -cobordant via an n -cobordism inducing the n -basing

$$(\psi, \psi_\partial) = ((\phi')^{-1} \circ \phi'', (\phi'_\partial)^{-1} \circ \phi''_\partial)$$

of L'' relative to L' .

As a consequence, the following basing-independent statements are equivalent:

- (1') We have $h_n(L') = h_n(L'')$.
- (2') The links L' and L'' are n -cobordant.

In particular, $h_n(L', \phi')$ vanishes if and only if L' is n -cobordant to the fixed link L via a surface whose exterior induces the isomorphisms ϕ' and ϕ'_∂ , and $h_n(L')$ vanishes if and only if L' is n -cobordant to the fixed link L .

3.3 Realization of lower central homotopy invariants

Questions concerning the realizability of Milnor-type invariants date back to Milnor's original work [1957]. Leveraging a realization theorem, Orr [1989] computed the number of linearly independent $\bar{\mu}$ -invariants of length n for every n . Similar computations for knots and links in other 3-manifolds can provide answers to open questions. Recall the Almost-concordance conjecture from Section 2.1. Celoria [2018] and Friedl, Nagel, Orson and Powell [Friedl et al. 2019] conjecture that, aside from certain exceptional cases, there are infinitely many almost-concordance classes representing any free homotopy class in $[S^1, M]$, where M is any closed orientable 3-manifold. Cappell and Shaneson [1974] first studied concordance up to local knotting. More recently, Celoria coined the term *almost-concordance* to mean concordance up to local knotting.

As the lower central homotopy invariants are almost-concordance invariants, proving either of the following statements is sufficient for proving the Almost-concordance conjecture for some fixed M and some fixed free homotopy class $x \in [S^1, M]$:

- (1) For some knot $K \subset M$ representing x and some $n \in \mathbb{N}$, the set of equivalence classes of elements in $[M, X_n(K)]_0$ realized as nonvanishing invariants of knots in M is infinite.
- (2) For some knot $K \subset M$ representing x , there exists a family $\{K_n\}_{n \geq 1}$ of knots in M such that $h_i(K_n) = h_i(K)$ for $i < n$ but $h_n(K_n) \neq h_n(K)$.

The following theorem, which characterizes for a fixed $L \subset M$ which classes in $[M, X_n(L)]_0$ are realized by links in M , can be leveraged to provide far more information about concordance classes of knots in a fixed free homotopy class. We prove this realization theorem in Section 6.

Theorem C Fix an m -component link $L \subset M$, and let $f \in [M, X_n(L)]_0$. Then there exists an m -component link $L' \subset M$ which admits an n -basing ϕ relative to L such that $h_n(L', \phi) = f$ if and only if the following three conditions hold.

- (1) The composition $H_3(M) \xrightarrow{f_*} H_3(X_n(L)) \xrightarrow{\Delta} H_2(\partial E_L)$, where Δ is the connecting homomorphism in the Mayer–Vietoris sequence corresponding to the decomposition $X_n(L) = M_{p_n}^\times \cup_{\partial E_L} \nu L$, sends the fundamental class $[M]$ to $\sum_i [T_i]$, the sum of the fundamental classes of the tori in $\partial E_L = \bigsqcup_i T_i$.
- (2) On fundamental groups, we have $f_* = \iota_n(L)_*$.
- (3) The image of the composition

$$H^1(X_n(L)) \xrightarrow{\cap f_*[M]} H_2(X_n(L)) \xrightarrow{\text{proj}} H_2(X_n(L))/\text{im}(K_j(\pi/\Gamma_n))$$

contains the image of $H_2(\pi/\Gamma_n)/K_j(\pi/\Gamma_n)$ for all $j < n$, where

$$K_j(\pi/\Gamma_n) = \ker(H_2(\pi/\Gamma_n) \rightarrow H_2(\pi/\Gamma_j)),$$

and where all coefficients are in $\mathbb{Z}[\pi_1(M)]$.

Condition (3) above may be viewed as a Poincaré duality property imposed on the class $f_*[M]$. It holds for all $j < n$ if and only if it holds for $j = n - 1$. We will see in the proof of Theorem C that condition (3) follows from conditions (1) and (2) in the case where $\pi_2(M) = 0$ or, more generally, if the composition

$$H_2(E_L; \mathbb{Z}[\pi_1(M)]) \xrightarrow{j_*} H_2(M; \mathbb{Z}[\pi_1(M)]) \xrightarrow{\iota_n(L)_*} H_2(X_n(L); \mathbb{Z}[\pi_1(M)])$$

is zero; see Corollary 6.2.

3.4 G -lower central homotopy invariants

We may reduce the lower central homotopy invariants from Section 3.1 by changing from the group $\pi_1(M)$ to some quotient group G (which could be the trivial group). We will use the resulting G -lower central homotopy invariants to study $\mathbb{Z}[G]$ -homology concordance in Sections 3.5–3.7. For this section, we assume there is a fixed surjective homomorphism $\varphi : \pi_1(M) \twoheadrightarrow G$ onto some group G . For a link $L \subset M$, we now denote the kernel of the composition $\pi \twoheadrightarrow \pi_1(M) \twoheadrightarrow G$ by Γ , so that the groups π / Γ_n are now the G -lower central quotients of π . The homological properties of concordance which hold over $\pi_1(M)$ also hold over G . A concordance exterior is a $\mathbb{Z}[G]$ -homology cobordism rel boundary and induces isomorphisms over G on G -lower central quotients of link groups.

Just as in Section 3.1, we define a tower of spaces $\{X_n^G(L)\}$ with maps $\iota_n^G(L) : M \rightarrow X_n^G(L)$. The space $X_n^G(L)$ is the homotopy pushout of a diagram $\mathcal{D}_n^G(L)$ analogous to the diagram $\mathcal{D}_n(L)$ from Section 3.1 but with the group π / Γ_n the n -th G -lower central quotient of π . The map $\iota_n^G(L)$ is still the homotopy class $M \rightarrow X_n^G(L)$ induced by the homotopy pushout diagram. We have the diagram over G below:

$$\begin{array}{ccccccc} & & M & & & & \\ & \swarrow \iota_{n+1}^G(L) & \downarrow \iota_n^G(L) & \searrow \iota_1^G(L) & & & \\ \cdots & \xrightarrow{\psi_{n+2,n+1}^G} & X_{n+1}^G(L) & \xrightarrow{\psi_{n+1,n}^G} & X_n^G(L) & \xrightarrow{\psi_{n,n-1}^G} & \cdots \xrightarrow{\psi_{2,1}^G} X_1^G(L) \end{array}$$

We will often omit the group G from our notation when it is understood from context.

To define invariants in this setting, we fix a link $L \subset M$. For other links $L' \subset M$, we require a $\mathbb{Z}[G]$ -homology version of an n -basing which we will call an n -basing over G relative to L . When $G = \pi_1(M)$ and $\varphi = \text{id}_{\pi_1(M)}$, an n -basing over G is just an n -basing in the sense of Section 3.1. We will often refer to an n -basing over G simply as an n -basing. An analogue of Lemma 3.2, Proposition 4.2, implies that an n -basing ϕ over G for L' relative to L induces a well-defined based homotopy class of maps of pairs $(E_{L'}, \partial E_{L'}) \rightarrow (M_{p_n}^\times, \partial E_L)$, restricting to an admissible homeomorphism $\partial E_{L'} \xrightarrow{\cong} \partial E_L$, which extends canonically to a based homotopy class of maps of pairs $h_n^G(L', \phi) : (M, \nu L') \rightarrow (X_n^G(L), \nu L)$, all over G . We define the G -lower central homotopy invariant to be this homotopy class $h_n^G(L', \phi)$. This invariant agrees with the lower central homotopy invariant $h_n(L', \phi)$ when $G = \pi_1(M)$ and $\varphi = \text{id}_{\pi_1(M)}$. Note that $h_n^G(L, \text{id}_{\pi/\Gamma_n}) = \iota_n^G(L)$. We may again remove the dependence on the n -basing by defining $h_n^G(L')$, the image of $h_n^G(L', \phi)$ in the orbit space $[M, X_n^G(L)]_0 / \text{Aut}(\pi / \Gamma_n, \partial)$, where the group acting is now the group of self- n -basings over G . Analogues of Theorems A and B hold for the invariants h_n^G .

Any n -basing ϕ over $\pi_1(M)$ induces an n -basing ϕ^G over G . Assuming an n -basing for L' relative to L exists, we may consider the relationship between $h_n(L', \phi)$ and $h_n^G(L', \phi^G)$. The surjections of $\pi_1(M)$ -lower central quotients onto G -lower central quotients yield morphisms of homotopy pushout squares $\mathcal{D}_n(L) \rightarrow \mathcal{D}_n^G(L)$ inducing well-defined homotopy classes of maps $\rho_n^G : X_n(L) \rightarrow X_n^G(L)$. It is straightforward to show that

$$\rho_n^G \circ h_n(L', \phi) = h_n^G(L', \phi^G) \in [M, X_n^G(L)]_0.$$

3.5 Homology concordance invariants

We now consider the ideas from Sections 3.1 and 3.4 in the broader setting of $\mathbb{Z}[G]$ -homology concordance, concordance in $\mathbb{Z}[G]$ -homology cobordisms, for suitable groups G . We introduce two $\mathbb{Z}[G]$ -homology concordance invariants θ_n^G and $\bar{\mu}_n^G$ determined by the G -lower central homotopy invariant h_n^G from Section 3.4. As we will see in Theorem D, the invariants $\bar{\mu}_n^G$ inductively determine the G -lower central quotients of link groups.

For the remainder of Section 3, we allow the links L , L' , and L'' to live in different closed orientable 3-manifolds M , M' and M'' . We assume there are fixed surjective homomorphisms φ , φ' and φ'' from the fundamental groups of these 3-manifolds onto a common group G (which could be the trivial group). We consider these homomorphisms as fixed data accompanying the 3-manifolds and use them to define $\mathbb{Z}[G]$ -coefficient systems on the 3-manifolds. Denote the kernels of the surjective homomorphisms from the link groups onto G by Γ , Γ' and Γ'' , respectively.

Definition 3.8 A $\mathbb{Z}[G]$ -homology cobordism from (M, φ) to (M', φ') is an orientable 4-manifold W with boundary $M \sqcup M'$, along with a diagram

$$\begin{array}{ccc} M' & & \\ \downarrow & \searrow \varphi' & \\ W & \longrightarrow & BG \\ \uparrow & \nearrow \varphi & \\ M & & \end{array}$$

such that the inclusions $M \hookrightarrow W$ and $M' \hookrightarrow W$ induce isomorphisms on $H_*(-; \mathbb{Z}[G])$.

Definition 3.9 Two m -component links $L \subset M$ and $L' \subset M'$ are $\mathbb{Z}[G]$ -homology concordant if there exists a proper, locally flat embedding

$$C : \bigsqcup_{i=1}^m (S^1 \times [0, 1]) \hookrightarrow W$$

into some $\mathbb{Z}[G]$ -homology cobordism W between (M, φ) and (M', φ') which restricts on the ends to the links L and L' .

Note that we return to the setting of Section 3.1 if we take $M = M'$, $W = M \times [0, 1]$, $G = \pi_1(M)$ and $\varphi = \varphi' = \text{id}_{\pi_1(M)}$, as a concordance is a $\mathbb{Z}[\pi_1(M)]$ -homology concordance. Taking $G = \{1\}$ reduces Definition 3.9 to the more familiar notion of integral homology concordance.

The fact that L and L' live in different 3-manifolds does not prevent us from comparing the towers $\{X_n^G(L')\}$ and $\{X_n^G(L)\}$. Given an n -basing (ϕ, ϕ_∂) over G for L' relative to L , we may define the homotopy class $h_n^G(L', \phi)$; however, comparing the homotopy classes $h_n^G(L', \phi) \in [M', X_n^G(L)]_0$ and $h_n^G(L, \text{id}) \in [M, X_n^G(L)]_0$ is not straightforward. We would like to determine whether these two classes extend over some $\mathbb{Z}[G]$ -homology cobordism between (M, φ) and (M', φ') . This may be difficult to measure, so we instead ask if the images of the fundamental classes $[M]$ and $[M']$ agree in $H_3(X_n^G(L))$. This is the case if the homotopy classes $h_n^G(L', \phi)$ and $h_n^G(L, \text{id})$ extend over some $\mathbb{Z}[G]$ -homology cobordism. We are immediately led to a reduction of the lower central homotopy invariants, which we call the *lower central homology invariants* θ_n^G .

Definition 3.10 Fix an m -component link $L \subset M$, and let $L' \subset M'$ be another m -component link. Suppose L' admits an n -basing (ϕ, ϕ_∂) over G relative to L . Then the n -th G -lower central homology invariant of the pair (L', ϕ) relative to L is the homology class

$$\theta_n^G(L', \phi) = h_n^G(L', \phi)_*[M'] \in H_3(X_n^G(L)).$$

We will often suppress the group G from our notation when it is understood from context. To be concise, we will often refer to θ_n as the n -th θ -invariant. The n -th θ -invariant is well-defined because the homotopy class $h_n^G(L', \phi)$ is well-defined. As with the h -invariant, the θ -invariant may depend on the choice of n -basing. We account for this by taking the image of $\theta_n^G(L', \phi)$ in the orbit space $H_3(X_n^G(L))/\text{Aut}(\pi/\Gamma_n, \partial)$, where $\text{Aut}(\pi/\Gamma_n, \partial)$ now denotes the group of *self- n -basings of L over G* .

Definition 3.11 Let L and L' be as in the previous definition. The n -th G -lower central homology invariant of the link L' relative to L is the image $\theta_n^G(L')$ of the homology class $\theta_n^G(L', \phi)$ in the orbit space $H_3(X_n^G(L))/\text{Aut}(\pi/\Gamma_n, \partial)$.

As before, we will say the θ -invariant *vanishes* if $\theta_n^G(L', \phi) = \theta_n^G(L, \text{id}) \in H_3(X_n^G(L))$ or if $\theta_n^G(L') = \theta_n^G(L) \in H_3(X_n^G(L))/\text{Aut}(\pi/\Gamma_n, \partial)$. We summarize invariance properties of the θ -invariants in Theorem A', which we prove in Section 5.

Theorem A' The n -th G -lower central homology invariant θ_n^G is an invariant of $\mathbb{Z}[G]$ -homology concordance. More precisely, fix an m -component link $L \subset (M, \varphi)$. Suppose the m -component links $L' \subset (M', \varphi')$ and $L'' \subset (M'', \varphi'')$ are $\mathbb{Z}[G]$ -homology concordant. Then the following statements hold for all n :

- (1) There is an n -basing (ϕ', ϕ'_∂) over G for L'' relative to L' induced by the homology concordance.
- (2) An n -basing over G for L' relative to L exists if and only if an n -basing over G for L'' relative to L exists, and the homology class $\theta_n^G(L')$ is defined if and only if $\theta_n^G(L'')$ is defined.

(3) If (ϕ, ϕ_∂) is an n -basing over G for L' relative to L , then $\theta_n^G(L', \phi) = \theta_n^G(L'', \phi \circ \phi')$, where ϕ' is any n -basing from (1).

(4) If $\theta_n^G(L')$ and $\theta_n^G(L'')$ are defined using any n -basings over G relative to L , then

$$\theta_n^G(L') = \theta_n^G(L'') \in H_3(X_n^G(L))/\text{Aut}(\pi/\Gamma_n, \partial).$$

3.6 Determination of G -lower central quotients

We now consider the set of homology classes realized by the invariants θ_n^G . Not all classes are necessarily realizable. In fact, the same is true for the invariants from [Cha and Orr 2024] which the θ_n^G generalize. Define

$$\mathcal{R}_n^G(L) = \left\{ \theta \in H_3(X_n^G(L)) \left| \begin{array}{l} \theta = \theta_n^G(L', \phi) \text{ for some link } L' \text{ in some 3-manifold } (M', \varphi') \\ \text{with } \varphi' : \pi_1(M') \twoheadrightarrow G \text{ which admits some } n\text{-basing } \phi \text{ over } G \\ \text{relative to } L \end{array} \right. \right\}.$$

We will often suppress G from the notation and write $X_n(L)$, $\theta_n(L)$ and $\mathcal{R}_n(L)$. We have a well-defined function $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$ induced by the canonical map $\psi_{n+1,n}$. This is a restriction of the homomorphism $(\psi_{n+1,n})_* : H_3(X_{n+1}(L)) \rightarrow H_3(X_n(L))$.

As seen in [Cha and Orr 2024], the set $\mathcal{R}_n(L)$ is not necessarily a subgroup of $H_3(X_n(L))$, so the notion of the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$ is not well-defined. Nevertheless, we make the following definition:

Definition 3.12 A class $\theta \in \mathcal{R}_n(L)$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$ if θ lies in the image of the map $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$.

In other words, a realizable class $\theta \in H_3(X_n(L))$ vanishes in the cokernel if there exists a link $L' \subset (M', \varphi')$ which admits an $(n+1)$ -basing ϕ over G relative to L such that $(\psi_{n+1,n})_*(\theta_{n+1}(L', \phi)) = \theta$.

Likewise, we will say a class $\theta \in \mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial)$ vanishes in the cokernel if θ lies in the image of the composite

$$\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L) \rightarrow \mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial).$$

The orbit space $\mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial)$ is well-defined because the action of any self-basing on a realizable class yields another realizable class. The following theorem, which we prove in Section 7, characterizes the notion of vanishing in the cokernel.

Theorem D₀ Fix an m -component link $L \subset M$. Suppose the m -component link $L' \subset M'$ admits an n -basing ϕ over G relative to L for some $n \geq 2$. Then the following statements are equivalent:

- (1) The link L' admits an $(n+1)$ -basing $\tilde{\phi}$ over G relative to L which is a lift of ϕ .
- (2) The invariant $\theta_n(L', \phi)$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$.

As a consequence, the following basing-independent statements are equivalent:

- (1') The link L' admits some $(n+1)$ -basing over G relative to L .
- (2') The invariant $\theta_n(L')$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial)$.

The notion of vanishing in the cokernel extends to an equivalence relation $\sim$ on all of $\mathcal{R}_n(L)$. Rather than comparing a link L' to the fixed link L , we can compare two links L' and L'' over the fixed link L . Given $\theta \in \mathcal{R}_n(L)$, there exists a 3-manifold (M', φ') and a link $L' \subset M'$ which admits an n -basing ϕ relative to L such that $\theta_n(L', \phi) = \theta$. Observe that ϕ induces a bijection $\mathcal{R}_n(L') \xrightarrow{\cong} \mathcal{R}_n(L)$ via $\theta_n(L'', \phi') \mapsto \theta_n(L'', \phi \circ \phi')$. Let I_θ be the image of the composition $\mathcal{R}_{n+1}(L') \rightarrow \mathcal{R}_n(L') \xrightarrow{\cong} \mathcal{R}_n(L)$. We prove in Lemma 7.2 that the set $\{I_\theta \mid \theta \in \mathcal{R}_n(L)\}$ partitions $\mathcal{R}_n(L)$. Let $\sim$ be the associated equivalence relation.

Definition 3.13 Let $\theta, \theta' \in \mathcal{R}_n(L)$. We say that $\theta \sim \theta'$ if and only if $\theta' \in I_\theta$.

The invariant $\bar{\mu}_n^G$ is the equivalence class of the θ -invariant under this equivalence relation.

Definition 3.14 Fix an m -component link $L \subset M$, and let $L' \subset M'$ be another m -component link. Suppose L' admits an n -basing (ϕ, ϕ_∂) over G relative to L . The n -th G -Milnor invariant of the pair (L', ϕ) relative to L is $\bar{\mu}_n^G(L', \phi) = [\theta_n^G(L', \phi)] \in \mathcal{R}_n^G(L)/\sim$, the class of the n -th θ -invariant under the equivalence relation $\sim$.

Again, we sometimes suppress G when it is understood from context. As with the h - and θ -invariants, we can remove dependency on the choice of n -basing ϕ . We define a coarser equivalence relation $\approx$ on $\mathcal{R}_n(L)$ as follows:

Definition 3.15 Let $\theta, \theta' \in \mathcal{R}_n(L)$. We say that $\theta \approx \theta'$ if and only if there exists a self-basing $\psi = (\psi, \psi_\partial) \in \text{Aut}(\pi/\Gamma_n, \partial)$ of L such that $\psi \cdot \theta' \sim \theta$. In other words, $\theta \approx \theta'$ if, upon choosing $L' \subset M'$ which admits an n -basing ϕ over G relative to L with $\theta_n(L', \phi) = \theta$, the class $[\theta']$ is in the image of the composition

$$\mathcal{R}_{n+1}(L') \rightarrow \mathcal{R}_n(L') \xrightarrow[\phi_*]{\cong} \mathcal{R}_n(L) \rightarrow \mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial).$$

Definition 3.16 Let L and L' be as in Definition 3.14. The n -th G -Milnor invariant of the link L' relative to L is $\bar{\mu}_n^G(L') = [\theta_n^G(L', \phi)] \in \mathcal{R}_n^G(L)/\approx$, the class of the n -th θ -invariant under the equivalence relation $\approx$.

As with the h - and θ -invariants, we will say the $\bar{\mu}$ -invariant *vanishes* if we have $\bar{\mu}_n^G(L', \phi) = \bar{\mu}_n^G(L, \text{id})$ or $\bar{\mu}_n^G(L') = \bar{\mu}_n^G(L)$. Observe that $\bar{\mu}_n^G(L', \phi)$ vanishes if and only if $\theta_n^G(L', \phi)$ vanishes in the cokernel, and likewise for $\bar{\mu}_n^G(L')$ and $\theta_n^G(L')$.

We prove the following theorem, which strengthens Theorem D₀, in Section 7. This theorem reduces to Milnor's theorem when we take $M = M' = S^3$ and our fixed link L to be the unlink. It is both a justification for calling these new invariants $\bar{\mu}_n$ "Milnor invariants" and a way to show that *the vanishing of these invariants implies the vanishing of all previous versions of Milnor's invariants in dimension 3*. Note that we have already seen a version of this theorem in Section 1.1.

Theorem D Fix an m -component link $L \subset (M, \varphi)$. Suppose the m -component links $L' \subset (M', \varphi')$ and $L'' \subset (M'', \varphi'')$ admit n -basings ϕ' and ϕ'' over G relative to L for some $n \geq 2$. Then the following statements are equivalent:

- (1) The link L' admits an $(n+1)$ -basing $\tilde{\phi}'$ over G relative to L which is a lift of ϕ' .
- (2) The invariant $\bar{\mu}_n(L', \phi')$ vanishes.
- (3) The invariant $\bar{\mu}_{n+1}(L', \tilde{\phi}')$ is well-defined.

As a consequence, the following basing-independent statements are equivalent:

- (1') The link L' admits some $(n+1)$ -basing over G relative to L .
- (2') The invariant $\bar{\mu}_n(L')$ vanishes.
- (3') The invariant $\bar{\mu}_{n+1}(L')$ is well-defined.

Even if the $\bar{\mu}$ -invariant of L' does not vanish, we have the following equivalent statements concerning the links L' and L'' :

- (4) The n -basing $(\phi')^{-1} \circ \phi''$ of L'' over G relative to L' lifts to an $(n+1)$ -basing.
- (5) We have $\bar{\mu}_n(L', \phi') = \bar{\mu}_n(L'', \phi'')$.

As a consequence, the following basing-independent statements are equivalent:

- (4') The link L'' admits some $(n+1)$ -basing over G relative to L' .
- (5') We have $\bar{\mu}_n(L') = \bar{\mu}_n(L'')$.

Note that (1)–(3) and (1')–(3') in the above theorem are just the statement of Theorem D_0 .

Corollary D_1 For $n \geq 2$, the lower central homotopy invariants, G -lower central homology invariants, and G -Milnor invariants are defined inductively. If any of these invariants vanishes at level n , the invariant at level $n+1$ is well-defined.

Proof All that is required to define any of the invariants for L' relative to L at level $n+1$ is an $(n+1)$ -basing for L' relative to L . The $\bar{\mu}$ -invariants are determined by both the h - and θ -invariants. If any n -th invariant is defined and vanishes, then the n -th $\bar{\mu}$ -invariant is defined and vanishes. Theorem D then implies that an $(n+1)$ -basing exists. $\square$

We have just seen in Theorem D that the invariants $\bar{\mu}_n$ determine the lower central quotients one step at a time. Corollary D_2 below, which follows from Theorems C and D, provides an alternative perspective: The invariant $\bar{\mu}_n(L')$ is the obstruction to lifting the homotopy invariant $h_n(L')$ to the next level of the tower $\{X_n(L)\}$. Note that it concerns links L and L' in the same 3-manifold M .

Corollary D_2 Suppose $L' \subset M$ admits an n -basing ϕ over $\pi_1(M)$ relative to $L \subset M$ for some $n \geq 2$. Then $h_n(L', \phi)$ is defined. The following statements are equivalent:

- (1) The invariant $h_n(L', \phi)$ lifts to a class $f \in [M, X_{n+1}(L)]_0$ such that the image of the composition

$$H^1(X_{n+1}(L)) \xrightarrow{\cap f_*[M]} H_2(X_{n+1}(L)) \xrightarrow{\text{proj}} H_2(X_{n+1}(L))/\text{im}(K_j(\pi/\Gamma_{n+1}))$$

contains the image of $H_2(\pi/\Gamma_{n+1})/K_j(\pi/\Gamma_{n+1})$ for all $j < n+1$, where all coefficients are in $\mathbb{Z}[\pi_1(M)]$.

- (2) The invariant $\bar{\mu}_n(L', \phi)$ vanishes.

As a consequence, the following basing-independent statements are equivalent:

(1') For some n -basing ϕ' , the invariant $h_n(L', \phi')$ lifts to a class $f \in [M, X_{n+1}(L)]_0$ such that the image of the composition

$$H^1(X_{n+1}(L)) \xrightarrow{\cap f_*[M]} H_2(X_{n+1}(L)) \xrightarrow{\text{proj}} H_2(X_{n+1}(L))/\text{im}(K_j(\pi/\Gamma_{n+1}))$$

contains the image of $H_2(\pi/\Gamma_{n+1})/K_j(\pi/\Gamma_{n+1})$ for all $j < n+1$, where all coefficients are in $\mathbb{Z}[\pi_1(M)]$.

(2') The invariant $\bar{\mu}_n(L')$ vanishes.

Proof (1) $\Rightarrow$ (2): Suppose $h_n(L', \phi)$ lifts to such a map $f \in [M, X_{n+1}(L)]_0$. The map f automatically satisfies conditions (1) and (2) of Theorem C since $h_n(L', \phi)$ is a realizable class and therefore satisfies those conditions. The additional assumption on the composition $\text{proj} \circ (\cap f_*[M])$ allows us to apply Theorem C and conclude $f = h_{n+1}(L'', \tilde{\phi}')$ for some $L'' \subset M$ and some $(n+1)$ -basing $\tilde{\phi}'$. This implies the homology invariant $\theta_n(L', \phi) = (\psi_{n+1,n})_* \theta_{n+1}(L'', \tilde{\phi}') = \theta_n(L'', \phi')$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$, where ϕ' is the n -basing induced by $\tilde{\phi}'$. Equivalently, $\bar{\mu}_n(L', \phi)$ vanishes.

(2) $\Rightarrow$ (1): Now suppose $\bar{\mu}_n(L', \phi)$ vanishes. By Theorem D, ϕ lifts to an $(n+1)$ -basing $\tilde{\phi}$, so $h_{n+1}(L', \tilde{\phi})$ is defined. We have $\psi_{n+1,n} \circ h_{n+1}(L', \tilde{\phi}) = h_n(L', \phi)$, so $h_n(L', \phi)$ lifts to a realizable class for which the desired condition on the composition $\text{proj} \circ (\cap f_*[M])$ is satisfied by Theorem C.

(1') $\iff$ (2') follows similarly. $\square$

The assumption on the composition $\text{proj} \circ (\cap f_*[M])$ in Corollary D₂ is the statement of condition (3) from Theorem C for the class f , where n is replaced with $n+1$. Recall from our discussion in Section 3.3 that condition (3) of Theorem C follows from conditions (1) and (2) of Theorem C in many cases; see also Corollary 6.2. Thus, in many cases, Corollary D₂ states that the lower central homotopy invariant h_n lifts if and only if the Milnor invariant $\bar{\mu}_n$ vanishes.

3.7 Realization of lower central homology invariants

The following realization theorem for the invariants θ_n in the case where the group G is the trivial group is analogous to Theorem C from Section 3.3. We prove Theorem E in Section 7.

Theorem E *Let the fixed group G be the trivial group. Fix an m -component link $L \subset M$, and let $\theta \in H_3(X_n^G(L))$. Then $\theta \in \mathcal{R}_n^G(L)$ if and only if the following three conditions hold:*

(1) *The connecting homomorphism $H_3(X_n^G(L)) \xrightarrow{\Delta} H_2(\partial E_L)$ from the Mayer–Vietoris sequence corresponding to the decomposition $X_n^G(L) = M_{p_n}^\times \cup_{\partial E_L} \nu L$ sends θ to $\sum_i [T_i]$, the sum of the fundamental classes of the tori in $\partial E_L = \bigsqcup_i T_i$.*

(2) *The cap product*

$$\cap \theta : tH^2(X_n^G(L)) \rightarrow tH_1(X_n^G(L))$$

is an isomorphism.

(3) *The image of the composition*

$$H^1(X_n^G(L)) \xrightarrow{\cap \theta} H_2(X_n^G(L)) \xrightarrow{\text{proj}} H_2(X_n^G(L))/\text{im}(K_j(\pi/\pi_n))$$

contains the image of $H_2(\pi/\pi_n)/K_j(\pi/\pi_n)$ for all $j < n$.

We briefly compare this theorem to Theorem C. Conditions (1) and (3) of Theorem E are analogous to conditions (1) and (3) of Theorem C. Condition (3) holds for all $j < n$ if and only if it holds for $j = n - 1$. Analogous to Theorem C, condition (3) of Theorem E follows from conditions (1) and (2) in the case where the composition $H_2(E_L; \mathbb{Z}) \rightarrow H_2(M; \mathbb{Z}) \rightarrow H_2(X_n^G(L); \mathbb{Z})$ is zero.

An analogue of condition (2) of Theorem E is missing from Theorem C. In the proof of Theorem E, this condition allows us to leverage a 3-dimensional homology surgery result of V Turaev [1984] to produce a 3-manifold with the desired homological properties along with a link realizing the class θ . Condition (2) in Theorem E is necessary to apply Turaev's theorem. This is not needed in Theorem C, where we begin with a homotopy class of maps from M . Proving Theorem E in the case of other groups G would rely on extending Turaev's result to the setting of $\mathbb{Z}[G]$ -coefficients; this is a topic of future consideration.

4 n -basings

We require an n -basing over G for a link $L' \subset M'$ relative to the fixed link $L \subset M$ in order to define the n -th G -lower central homotopy invariant, the n -th G -lower central homology invariant and the n -th G -Milnor invariant of L' relative to L . Recall from Sections 3.1 and 3.4 that an n -basing over G is a pair (ϕ, ϕ_∂) , where $\phi : \pi' / \Gamma'_n \xrightarrow{\cong} \pi / \Gamma_n$ is an isomorphism over G on n -th G -lower central quotients and $\phi_\partial : H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L)$ is an *admissible* isomorphism which is *compatible* with ϕ . We now motivate this definition by showing that a $\mathbb{Z}[G]$ -homology cobordism induces an n -basing over G for all n .

Proposition 4.1 *For all n , a $\mathbb{Z}[G]$ -homology concordance between the links $L \subset M$ and $L' \subset M'$ induces a based homotopy equivalence of pairs $(M_{p_n}^\times, \partial E_{L'}) \xrightarrow{\cong} (M_{p_n}^\times, \partial E_L)$, restricting to an admissible homeomorphism $\partial E_{L'} \xrightarrow{\cong} \partial E_L$ and extending canonically to a based homotopy equivalence of pairs $(X_n^G(L'), \nu L') \xrightarrow{\cong} (X_n^G(L), \nu L)$. In the case of a concordance in $M \times [0, 1]$, where $M' = M$ and $G = \pi_1(M)$, the based homotopy equivalence $(M_{p_n}^\times, \partial E_{L'}) \xrightarrow{\cong} (M_{p_n}^\times, \partial E_L)$ is canonical.*

Proof Suppose C is a concordance between $L \subset M$ and $L' \subset M'$ in the $\mathbb{Z}[G]$ -homology cobordism W . We take the basepoint of the homology cobordism W to be the basepoint of M . Choose some path γ in W missing C from the basepoint of M to the basepoint of M' whose image in BG under the based map $W \rightarrow BG$ is nullhomotopic. An analogue of Corollary 2.1, with an identical proof, implies that the inclusions $E_{L'} \hookrightarrow E_C \hookleftarrow E_L$ induce isomorphisms $\pi' / \Gamma'_n \xrightarrow{\cong} \pi^C / \Gamma_n^C \xleftarrow{\cong} \pi / \Gamma_n$ for all n , where $\pi^C = \pi_1(E_C)$ and $\Gamma^C = \ker(\pi^C \twoheadrightarrow \pi_1(W) \twoheadrightarrow G)$.

We obtain the diagram below, where $X_n^G(C)$ is a standard homotopy pushout analogous to $X_n^G(L)$, and where $\partial \nu C$ denotes the intersection $\partial E_C \cap \nu C$, not the manifold boundary of $\partial \nu C$ which includes νL and $\nu L'$. Since γ is nullhomotopic in BG , this diagram is a diagram over G . The homotopy equivalences

of pairs on the second row are induced by the group isomorphisms, and we may extend to the homotopy equivalences on the bottom row by gluing in νC :

$$\begin{array}{ccccc}
 (E_{L'} \cup \gamma, \partial E_{L'}) & \xrightarrow{\quad} & (E_C, \partial \nu C) & \xleftarrow{\quad} & (E_L, \partial E_L) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 (M' \cup \gamma, \nu L') & \xrightarrow{\quad} & (W, \nu C) & \xleftarrow{\quad} & (M, \nu L) \\
 \downarrow \iota_n(L') & \searrow & \downarrow & \searrow & \downarrow \iota_n(L) \\
 (X_n^G(L') \cup \gamma, \nu L') & \xrightarrow{\quad} & (X_n^G(C), \nu C) & \xleftarrow{\quad} & (X_n^G(L), \nu L)
 \end{array}$$

$(M'_{p'_n}^\times \cup \gamma, \partial E_{L'}) \xrightarrow{\cong} (M'_{p'_n}, \partial E_{L'}) \xrightarrow{\quad} (M'_{p'_n}^\times, \partial \nu C) \xleftarrow{\quad} (M'_{p'_n}, \partial E_L) \xrightarrow{\cong} (M'_{p'_n}^\times, \partial E_L)$

In the case where C is a concordance in $M \times [0, 1]$ (where $M' = M$ and $G = \pi_1(M)$), we may take γ to be $* \times [0, 1]$, and all homotopy equivalences in the diagram are then canonical. $\square$

We now prove that such homotopy equivalences $(M'_{p'_n}^\times, \partial E_{L'}) \xrightarrow{\cong} (M'_{p'_n}, \partial E_L)$ as in Proposition 4.1 are in one-to-one correspondence with n -basings over G for L' relative to L . The following proposition is a refinement of Lemma 3.2 from Section 3.1.

Proposition 4.2 *The link L' admits an n -basing ϕ over G relative to L if and only if there exists a based homotopy equivalence of pairs*

$$h_\phi : (M'_{p'_n}^\times, \partial E_{L'}) \xrightarrow{\cong} (M'_{p'_n}, \partial E_L)$$

over G restricting to an admissible homeomorphism $\partial E_{L'} \xrightarrow{\cong} \partial E_L$. Additionally, any two homotopy equivalences realizing the same n -basing are based homotopic as maps of pairs. Furthermore, such a homotopy equivalence h extends canonically (up to based homotopy) to a based homotopy equivalence of pairs $\bar{h}_\phi : (X_n^G(L'), \nu L') \xrightarrow{\cong} (X_n^G(L), \nu L)$.

Proof We begin with the first statement. One direction requires little work: Suppose such a homotopy equivalence of pairs h exists. Then we may take ϕ and ϕ_∂ to be the isomorphisms

$$h_* : \pi_1(M'_{p'_n}^\times) \xrightarrow{\cong} \pi_1(M'_{p'_n}) \quad \text{and} \quad (h|_{\partial E_{L'}})_* : H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L),$$

respectively. By assumption, the homeomorphism $\partial E_{L'} \xrightarrow{\cong} \partial E_L$ is admissible, so ϕ_∂ is admissible. All that remains is to see that ϕ_∂ is compatible with ϕ in the sense of Definition 3.1. Take any basing τ' for L' . Its image $h(\tau')$ is a basing for L in $M'_{p'_n}$. We find a corresponding basing in E_L as follows: Take any basing τ^0 for L in E_L . The basings $h(\tau')$ and τ^0 differ up to homotopy in $M'_{p'_n}$ by some wedge of circles which represent elements of π / Γ_n . Adjust τ^0 by lifts of these elements to π to obtain a new basing τ in E_L . By construction, the desired diagram commutes.

The other direction is entirely similar to the proof of Lemma 3.2. We replace the pair $(E_{L'}, \partial E_{L'})$ in that proof with the pair $(M_{p_n}^\times, \partial E_{L'})$. The resulting map

$$h_\phi : (M_{p_n}^\times, \partial E_{L'}) \xrightarrow{\cong} (M_{p_n}^\times, \partial E_L)$$

is a homotopy equivalence of pairs because we construct it using the isomorphism ϕ .

We also follow the proof of Lemma 3.2 to prove the remainder of the proposition. Using similar arguments, the homotopy equivalence h_ϕ is unique up to based homotopy of maps of pairs and extends canonically (up to based homotopy) to a map of pairs

$$\bar{h}_\phi : (X_n^G(L'), \nu L') \xrightarrow{\cong} (X_n^G(L), \nu L)$$

over G . Because h_ϕ is a homotopy equivalence of pairs, $\bar{h}_\phi$ is, too. $\square$

By the proof of the above proposition, the map of pairs $(E_{L'}, \partial E_{L'}) \rightarrow (M_{p_n}^\times, \partial E_L)$ and its extension constructed in the proof of Lemma 3.2 factor through the homotopy equivalences h_ϕ and $\bar{h}_\phi$, respectively. Therefore, the composition $\bar{h}_\phi \circ \iota_n(L')$ is precisely the n -th h -invariant $h_n(L', \phi)$.

Corollary 4.3 *A $\mathbb{Z}[G]$ -homology concordance between the links $L \subset M$ and $L' \subset M'$ induces an n -basing over G for L' relative to L for all n . In the case of a concordance in $M \times [0, 1]$, where $M' = M$ and $G = \pi_1(M)$, the induced n -basing is canonical.*

4.1 On the choice of n -basing

The invariants h_n , θ_n^G and $\bar{\mu}_n^G$ defined in Sections 3.1, 3.5 and 3.6 at first depend on the choice of n -basing for L' . We therefore study how different choices of n -basing are related. Suppose (ϕ, ϕ_∂) and (ϕ', ϕ'_∂) are two n -basings for L' relative to L . We will show that

$$(\psi, \psi_\partial) = (\phi' \circ \phi^{-1}, \phi'_\partial \circ \phi_\partial^{-1}) \in \text{Aut}(\pi / \Gamma_n) \times \text{Aut}(H_1(\partial E_L))$$

constitutes an n -basing over G for L relative to itself. We will call such an n -basing a *self- n -basing* or *self-basing*. By Proposition 4.2, to each self-basing we may associate a canonical homotopy class of homotopy self-equivalences of the pair $(X_n^G(L), \nu L)$ over G . In Proposition 4.6, we show that the collection of all self- n -basings (ψ, ψ_∂) , denoted by $\text{Aut}(\pi / \Gamma_n, \partial)$, forms a subgroup of $\text{Aut}(\pi / \Gamma_n) \times \text{Aut}(H_1(\partial E_L))$ isomorphic to the group of homotopy classes of homotopy self-equivalences of the pair $(M_{p_n}^\times, \partial E_L)$ over G which restrict to admissible homeomorphisms $\partial E_L \rightarrow \partial E_L$. In fact, because every such ψ is an isomorphism over G and every such ψ_∂ is admissible, we will have that

$$\text{Aut}(\pi / \Gamma_n, \partial) \leq \text{Aut}_G(\pi / \Gamma_n) \times \mathbb{Z}^m \leq \text{Aut}(\pi / \Gamma_n) \times \text{Aut}(H_1(\partial E_L)),$$

where Aut_G denotes automorphisms over G , and where $\mathbb{Z}^m \leq \text{Aut}(H_1(\partial E_L))$ is the subgroup of automorphisms of $H_1(\partial E_L)$ which fix the meridians of L .

Lemma 4.4 *Suppose (ϕ, ϕ_∂) is an n -basing over G for L' relative to L . Then $(\phi^{-1}, \phi_\partial^{-1})$ is an n -basing over G for L relative to L' .*

Proof Since ϕ_∂ is admissible, so is ϕ_∂^{-1} . We next see that ϕ_∂^{-1} is compatible with ϕ^{-1} . Let τ be a basing for L . Choose any basing τ'_0 for L' . Since ϕ_∂ is compatible with ϕ , there exists a basing τ_0 for L such that the following square commutes:

$$\begin{array}{ccc} \pi_1(\partial E_{L'} \cup \tau'_0) & \xrightarrow[\phi_{\partial*}]{\cong} & \pi_1(\partial E_L \cup \tau_0) \\ \downarrow & & \downarrow \\ \pi'/\Gamma'_n & \xrightarrow[\phi]{\cong} & \pi/\Gamma_n \end{array}$$

Since τ and τ_0 are both basings for L , they differ up to homotopy by some wedge of circles in E_L . Consider the image of this wedge of circles in π/Γ_n . Mapping this to π'/Γ'_n via ϕ^{-1} and choosing any lift to the group π' yields a wedge of circles in $E_{L'}$. Adjust τ'_0 by this wedge of circles to obtain a new basing τ' for L' . Using ϕ_∂^{-1} , define an isomorphism $\pi_1(\partial E_L \cup \tau) \xrightarrow{\cong} \pi_1(\partial E_{L'} \cup \tau')$, which fits into the diagram

$$\begin{array}{ccc} \pi_1(\partial E_L \cup \tau) & \xrightarrow[\phi_{\partial*}^{-1}]{\cong} & \pi_1(\partial E_{L'} \cup \tau') \\ \downarrow & & \downarrow \\ \pi/\Gamma_n & \xrightarrow[\phi^{-1}]{\cong} & \pi'/\Gamma'_n \end{array}$$

Thus, ϕ_∂^{-1} is compatible with ϕ^{-1} , so $(\phi^{-1}, \phi_\partial^{-1})$ is an n -basing over G for L relative to L' . $\square$

Lemma 4.5 Suppose (ϕ', ϕ'_∂) is an n -basing over G for L'' relative to L' and (ϕ, ϕ_∂) is an n -basing over G for L' relative to L . Then the composition $(\phi \circ \phi', \phi_\partial \circ \phi'_\partial)$ is an n -basing over G for L'' relative to L .

Proof Since ϕ_∂ and ϕ'_∂ are both admissible, so is the composition $\phi_\partial \circ \phi'_\partial$. Let τ'' be a basing for L'' . Since ϕ'_∂ is compatible with ϕ' , we obtain the desired basing τ' for L' . This in turn yields the desired basing τ for L since ϕ_∂ is compatible with ϕ . We have the following diagram:

$$\begin{array}{ccccc} \pi_1(\partial E_{L''} \cup \tau'') & \xrightarrow[\phi'_{\partial*}]{\cong} & \pi_1(\partial E_{L'} \cup \tau') & \xrightarrow[\phi_{\partial*}]{\cong} & \pi_1(\partial E_L \cup \tau) \\ \downarrow & & \downarrow & & \downarrow \\ \pi''/\Gamma''_n & \xrightarrow[\phi']{\cong} & \pi'/\Gamma'_n & \xrightarrow[\phi]{\cong} & \pi/\Gamma_n \end{array}$$

Thus, $\phi_\partial \circ \phi'_\partial$ is compatible with $\phi \circ \phi'$, so $(\phi \circ \phi', \phi_\partial \circ \phi'_\partial)$ is an n -basing over G for L'' relative to L . $\square$

The preceding Lemmas 4.4 and 4.5 reveal that the set of self- n -basings is a group. The following proposition reveals some of the structure of this group.

Proposition 4.6 The collection $\text{Aut}(\pi/\Gamma_n, \partial)$ of all self- n -basings over G of L forms a subgroup of $\text{Aut}(\pi/\Gamma_n) \times \text{Aut}(H_1(\partial E_L))$ isomorphic to the group of homotopy classes of homotopy self-equivalences of the pair $(M_{p_n}^\times, \partial E_L)$ over G which restrict to admissible homeomorphisms $\partial E_L \xrightarrow{\cong} \partial E_L$. Furthermore, every self-basing gives rise to a canonical homotopy class of homotopy self-equivalences of the pair $(X_n(L), \nu L)$ over G .

Proof Any self- n -basing (ψ, ψ_∂) is an element of $\text{Aut}(\pi/\Gamma_n) \times \text{Aut}(H_1(\partial E_L))$. Taking $L'' = L' = L$ in Lemmas 4.4 and 4.5, we see that $\text{Aut}(\pi/\Gamma_n, \partial)$ is closed under inverses and the group operation. Thus, $\text{Aut}(\pi/\Gamma_n, \partial) \leq \text{Aut}(\pi/\Gamma_n) \times \text{Aut}(H_1(\partial E_L))$. By Proposition 4.2, a self-basing (ψ, ψ_∂) over G is equivalent to a homotopy class of homotopy self-equivalences of the pair $(M_{p_n}^\times, \partial E_L)$ over G which restricts to an admissible homeomorphism $\partial E_L \xrightarrow{\cong} \partial E_L$. Such a homotopy self-equivalence extends over G to a canonical homotopy class of homotopy self-equivalences of the pair $(X_n(L), \nu L)$, again by Proposition 4.2. $\square$

By Proposition 4.6, $\text{Aut}(\pi/\Gamma_n, \partial)$ acts (on the left) on both $[M, X_n(L)]_0$ and $H_3(X_n(L))$ by post-composition with homotopy self-equivalences of $X_n(L)$, so that the orbit space of each of these sets modulo the group action is well-defined.

5 Concordance invariance

In this section, we prove Theorems A and A'. See Sections 3.1 and 3.5, respectively, for the statements of these theorems. Theorem A asserts the invariance of h_n (hence θ_n and $\bar{\mu}_n$) under concordance. Theorem A', whose proof is similar to that of Theorem A, asserts the invariance of θ_n (hence $\bar{\mu}_n$) under $\mathbb{Z}[G]$ -homology concordance.

5.1 Proof of Theorem A

Proof Fix $L \subset M$, and assume $L', L'' \subset M$ are concordant.

Statement (1), which asserts the existence of a canonical n -basing (ϕ', ϕ'_∂) for L'' relative to L' for each n , follows directly from Corollary 4.3.

Next, we prove (2), which states that for each n an n -basing for L' relative to L exists if and only if an n -basing for L'' relative to L exists, and that the invariant $h_n(L')$ is defined if and only if $h_n(L'')$ is defined. Suppose (ϕ, ϕ_∂) is an n -basing for L' relative to L . By Lemma 4.5, the composition $(\phi \circ \phi', \phi_\partial \circ \phi'_\partial)$ constitutes an n -basing for L'' relative to L . Conversely, suppose (ϕ, ϕ_∂) is an n -basing for L'' relative to L . By Lemma 4.4, $((\phi')^{-1}, (\phi'_\partial)^{-1})$ is an n -basing for L' relative to L'' , so that $(\phi \circ (\phi')^{-1}, \phi_\partial \circ (\phi'_\partial)^{-1})$ is an n -basing for L' relative to L , again by Lemma 4.5. This proves (2), since the existence of an n -basing relative to L is all that is needed to define the invariant h_n .

We now prove (3), which gives the precise relationship between the h -invariants of L' and those of L'' . Again assume (ϕ, ϕ_∂) is an n -basing for L' relative to L . By Proposition 4.2, we obtain a homotopy equivalence of pairs $h_\phi : (M_{p_n}^\times, \partial E_{L'}) \xrightarrow{\cong} (M_{p_n}^\times, \partial E_L)$ restricting to an admissible homeomorphism $\partial E_{L'} \xrightarrow{\cong} \partial E_L$ and extending to a homotopy equivalence of pairs $\bar{h}_\phi : (X_n(L'), \nu L') \xrightarrow{\cong} (X_n(L), \nu L)$, all over $\pi_1(M)$. Since L' and L'' are concordant, from Proposition 4.1 we obtain a canonical homotopy equivalence of pairs $h' : (M_{p_n}^\times, \partial E_{L''}) \xrightarrow{\cong} (M_{p_n}^\times, \partial E_{L'})$ restricting to an admissible homeomorphism $\partial E_{L''} \xrightarrow{\cong} \partial E_{L'}$ and extending canonically to a homotopy equivalence of pairs $\bar{h}' : (X_n(L''), \nu L'') \xrightarrow{\cong} (X_n(L'), \nu L')$. In fact, using the language of Proposition 4.2, $h' = h_{\phi'}$ and $\bar{h}' = \bar{h}_{\phi'}$, where ϕ' is the canonical n -basing from statement (1). We have the following commutative diagram over $\pi_1(M)$ from

the proof of Proposition 4.1 (assume the concordance C misses $* \times [0, 1]$, and recall that $\partial \nu C$ denotes the intersection $\partial E_C \cap \nu C$):

$$\begin{array}{ccccc}
 (E_{L''}, \partial E_{L''}) & \xrightarrow{\quad} & (E_C, \partial \nu C) & \xleftarrow{\quad} & (E_{L'}, \partial E_{L'}) \\
 \downarrow & \searrow & \downarrow & \swarrow & \downarrow \\
 & (M_{p_n''}^\times, \partial E_{L''}) & \xrightarrow{\cong} & (M_{p_n^C}^\times, \partial \nu C) & \xleftarrow{\cong} & (M_{p_n'}^\times, \partial E_{L'}) \\
 & \downarrow & & \downarrow & & \downarrow \\
 (M, \nu L'') & \xrightarrow{\quad} & (M \times [0, 1], \nu C) & \xleftarrow{\quad} & (M, \nu L') \\
 \searrow \iota_n(L'') & & \searrow & & \searrow \iota_n(L') \\
 & (X_n(L''), \nu L'') & \xrightarrow{\cong} & (X_n(C), \nu C) & \xleftarrow{\cong} & (X_n(L'), \nu L')
 \end{array}$$

Thus, the composite homotopy equivalence $h_\phi \circ h_{\phi'} : (M_{p_n''}^\times, \partial E_{L''}) \xrightarrow{\cong} (M_{p_n'}^\times, \partial E_{L'})$ restricts to an admissible homeomorphism $\partial E_{L''} \xrightarrow{\cong} \partial E_{L'}$ and induces the n -basing $(\phi \circ \phi', \phi_\partial \circ \phi'_\partial)$. The map $h_\phi \circ h_{\phi'}$ extends to the homotopy equivalence of pairs $\bar{h}_\phi \circ \bar{h}_{\phi'}$. Thus,

$$h_n(L'', \phi \circ \phi') = \bar{h}_\phi \circ \bar{h}_{\phi'} \circ \iota_n(L'') = \bar{h}_\phi \circ \iota_n(L') = h_n(L', \phi),$$

where the first and last equalities follow by definition, and the middle equality follows from the commutativity of the above diagram.

Finally, we prove (4), which states that $h_n(L') = h_n(L'')$ when defined. Suppose (ϕ, ϕ_∂) and $(\phi'', \phi''_\partial)$ are n -basings for L' and L'' , respectively, relative to L . Then $(\phi \circ \phi', \phi_\partial \circ \phi'_\partial)$ is also an n -basing for L'' relative to L , and by (3) we have $h_n(L', \phi) = h_n(L'', \phi \circ \phi')$. Observe that by Lemmas 4.4 and 4.5, $(\psi, \psi_\partial) := (\phi'' \circ (\phi \circ \phi')^{-1}, \phi''_\partial \circ (\phi_\partial \circ \phi'_\partial)^{-1})$ is a self- n -basing for L . Then $h_n(L'', \phi'') = (\psi, \psi_\partial) \cdot h_n(L'', \phi \circ \phi') = (\psi, \psi_\partial) \cdot h_n(L', \phi)$, so $h_n(L') = h_n(L'') \in [M, X_n(L)]_0 / \text{Aut}(\pi / \Gamma_n, \partial)$. $\square$

Theorem A implies that the invariants θ_n and $\bar{\mu}_n$ defined in this context are also invariants of concordance.

5.2 Proof of Theorem A'

Proof We proceed in a manner similar to the proof of Theorem A. Fix $L \subset (M, \varphi)$, and suppose $L' \subset (M', \varphi')$ and $L'' \subset (M'', \varphi'')$ are $\mathbb{Z}[G]$ -homology concordant. Statement (1), which asserts the existence of an n -basing over G for L'' relative to L' for all n , is the statement of Corollary 4.3. The proofs of statements (2) and (4) are entirely similar to the proofs of the analogous statements (2) and (4) of Theorem A. Together, these state that $\theta_n(L')$ is defined if and only if $\theta_n(L'')$ is defined, in which case $\theta_n(L') = \theta_n(L'')$.

The proof of statement (3), which gives the precise relationship between the θ -invariants for L' and the θ -invariants for L'' , is nearly the same as the proof of (3) of Theorem A. Using the notation from the proof of Theorem A, we again have $h_n(L'', \phi \circ \phi') = \bar{h}_\phi \circ \bar{h}' \circ \iota_n(L'')$ and $h_n(L', \phi) = \bar{h}_\phi \circ \iota_n(L')$. By the commutativity of the diagram from the proof of Proposition 4.1 (which now involves a $\mathbb{Z}[G]$ -homology

cobordism W instead of $M \times [0, 1]$ as seen in the proof of Theorem A), we have

$$\begin{aligned}\theta_n(L'', \phi \circ \phi') &= h_n(L'', \phi \circ \phi')_*[M''] = (\bar{h}_\phi)_* \bar{h}'_* \iota_n(L'')_*[M''] \\ &= (\bar{h}_\phi)_* \iota_n(L')_*[M'] = h_n(L', \phi)_*[M'] = \theta_n(L', \phi),\end{aligned}$$

which is the claimed equality in statement (3). $\square$

Just as with Theorem A, Theorem A' immediately implies that the $\bar{\mu}_n$ are also invariants of $\mathbb{Z}[G]$ -homology concordance.

6 Characterization and realization for h_n

In this section, we prove characterization and realization properties of the lower central homotopy invariants h_n . We begin in Section 6.1 with a more detailed discussion of n -cobordism of links, including the proof of Proposition 3.7, which states that n -cobordisms induce n -basings. In Section 6.2, we prove Theorem B, which asserts that the invariants h_n are complete invariants of n -cobordism. We then prove Theorem C, which determines exactly which elements of $[M, X_n(L)]_0$ are realized as the h -invariants of links in M , in Section 6.3. For all of Section 6, all of our links live in M and the fixed group G is $\pi_1(M)$. Throughout this section, given a properly embedded surface $\Sigma \subset M \times [0, 1]$, we will use the notation $\partial \nu \Sigma$ to denote the intersection $\partial E_\Sigma \cap \nu \Sigma$, not the manifold boundary of $\nu \Sigma$, which also includes $(M \times \{0, 1\}) \cap \nu \Sigma$. Some of the proofs in this section are similar to proofs seen in [Heck 2011].

6.1 n -cobordism of links

Recall Definition 3.6: The links L and L' are n -cobordant if there exists a properly embedded orientable m -component surface $\Sigma = \bigsqcup_i \Sigma_i \subset M \times [0, 1]$ such that for each i :

- (1) $\Sigma_i \cap (M \times \{0\}) = L_i$ and $\Sigma_i \cap (M \times \{1\}) = L'_i$.
- (2) For any basing of Σ_i determined by a basing of L_i , the images of $\pi_1(L_i)$ and $\pi_1(\Sigma_i)$ in $\pi_1(E_\Sigma)$ agree modulo the n^{th} $\pi_1(M)$ -lower central subgroup, that is,

$$\text{im}(\pi_1(L_i) \rightarrow \pi \rightarrow \pi^\Sigma \twoheadrightarrow \pi^\Sigma / \Gamma_n^\Sigma) = \text{im}(\pi_1(\Sigma_i) \rightarrow \pi^\Sigma \twoheadrightarrow \pi^\Sigma / \Gamma_n^\Sigma),$$

where $\pi^\Sigma = \pi_1(E_\Sigma)$ and $\Gamma^\Sigma = \ker(\pi^\Sigma \twoheadrightarrow \pi_1(M))$, and where the inclusions $\Sigma_i \hookrightarrow E_\Sigma$ and $L_i \hookrightarrow E_L$ are induced by a trivialization of $\nu \Sigma$.

We prove some properties of n -cobordisms culminating with the fact that an n -cobordism induces an n -basing for L' relative to L .

Lemma 6.1 *Suppose Σ is an n -cobordism between L and L' . Then the image of the composition*

$$H_2(\partial E_L; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(\partial \nu \Sigma; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(\pi^\Sigma / \Gamma_n^\Sigma; \mathbb{Z}[\pi_1(M)])$$

equals the image of $H_2(\partial \nu \Sigma; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(\pi^\Sigma / \Gamma_n^\Sigma; \mathbb{Z}[\pi_1(M)])$. The same holds for L' .

Proof In this proof, we denote the $\pi_1(M)$ -cover of a space X by $\bar{X}$. Observe that $\partial\nu\Sigma \cong \Sigma \times S^1$ and $\overline{\partial\nu\Sigma} \cong \bar{\Sigma} \times S^1$, where Σ and $\bar{\Sigma}$ denote parallels of the cores of $\nu\Sigma$ and $\overline{\nu\Sigma}$, respectively. Since each component of Σ is a surface with boundary, it follows that a collection of generators of $H_2(\partial\nu\Sigma; \mathbb{Z}[\pi_1(M)])$ may be represented by embedded tori in $\overline{\partial\nu\Sigma}$. Each such torus has a symplectic basis consisting of a meridian to a component of $\bar{\Sigma}$, which we refer to as a meridian for the torus, and a curve lying on that component, which we call a longitude for the torus. Therefore, the image of any such generator of $H_2(\partial\nu\Sigma; \mathbb{Z}[\pi_1(M)])$ in $H_2(\pi^\Sigma/\Gamma_n^\Sigma; \mathbb{Z}[\pi_1(M)])$ is represented by a map $T^2 \rightarrow \overline{\partial\nu\Sigma} \rightarrow K(\Gamma^\Sigma/\Gamma_n^\Sigma, 1)$. It suffices to find a solution, up to homotopy, to the lifting problem

$$\begin{array}{ccc} & & \overline{\partial E_L} \\ & \nearrow \text{?}\exists & \downarrow \\ T^2 & \longrightarrow & K(\Gamma^\Sigma/\Gamma_n^\Sigma, 1) \end{array}$$

Base L (a parallel of L in ∂E_L) using some basing τ . This induces an equivariant basing $\bar{\tau}$ for $\bar{L}$ (a parallel of the core of $\nu\bar{L}$ in $\overline{\partial E_L}$), hence for $\bar{\Sigma}$. Extend this basing in the cover along the appropriate component of $\bar{\Sigma}$ to give a basing of the torus. The (based) meridian of T^2 is homotopic in $\overline{\partial\nu\Sigma}$ to a meridian for $\bar{L}$, so first define the image of the meridian of T^2 to be this meridian of $\bar{L}$. Next, note that the (based) longitude of T^2 represents an element of $\pi_1(\bar{\Sigma} \cup \bar{\tau}) \leq \pi_1(\Sigma \cup \tau)$. Since Σ is an n -cobordism, the image of this longitude in $\pi^\Sigma/\Gamma_n^\Sigma$ agrees with the image of some element $x \in \pi_1(L_i \cup \tau_i)$ for some i . But the image of the longitude is also an element of Γ^Σ . Since $E_L \hookrightarrow E_\Sigma$ is a map over $\pi_1(M)$, $x \in \Gamma \cap \pi_1(L_i \cup \tau_i) = \pi_1(\bar{L}_i \cup \bar{\tau}_i)$. Map the longitude of T^2 to a curve representing x . This gives a lifting of the 1-skeleton of T^2 up to homotopy.

Finally, we may choose any lift of the 2-cell of T^2 . The resulting composite

$$T^2 \rightarrow \overline{\partial E_L} \rightarrow K(\Gamma^\Sigma/\Gamma_n^\Sigma, 1)$$

is homotopic to the original map $T^2 \rightarrow \overline{\partial\nu\Sigma} \rightarrow K(\Gamma^\Sigma/\Gamma_n^\Sigma, 1)$ since the maps agree on π_1 . Thus, we have produced a map $T^2 \rightarrow \overline{\partial E_L}$ that represents an element of $H_2(\partial E_L; \mathbb{Z}[\pi_1(M)])$ which maps onto the image of the chosen generator of $H_2(\partial\nu\Sigma; \mathbb{Z}[\pi_1(M)])$ in $H_2(\pi^\Sigma/\Gamma_n^\Sigma; \mathbb{Z}[\pi_1(M)])$.

The same proof holds for L' . Recall from Section 3.2 that the condition

$$\text{im}(\pi_1(L_i) \rightarrow \pi \rightarrow \pi^\Sigma \twoheadrightarrow \pi^\Sigma/\Gamma_n^\Sigma) = \text{im}(\pi_1(\Sigma_i) \rightarrow \pi^\Sigma \twoheadrightarrow \pi^\Sigma/\Gamma_n^\Sigma)$$

implies the analogous condition for L' . □

We are now prepared to prove that n -cobordisms induce n -basings.

Proposition 3.7 *Suppose $\Sigma \subset M \times [0, 1]$ is an n -cobordism between L and L' . Then Σ induces a canonical n -basing (ϕ, ϕ_∂) for L' relative to L and a canonical isomorphism $\pi'/\Gamma'_{n+1} \xrightarrow{\cong} \pi/\Gamma_{n+1}$ over $\pi_1(M)$.*

Proof We first obtain the desired isomorphisms on $\pi_1(M)$ -lower central quotients.

Claim The inclusion $E_L \hookrightarrow E_\Sigma$ induces isomorphisms $\pi/\Gamma_k \xrightarrow{\cong} \pi^\Sigma/\Gamma_k^\Sigma$ for all $k \leq n+1$.

The claim is true for $k = 1$ because $E_L \hookrightarrow E_\Sigma$ is a map over $\pi_1(M)$. We use an induction argument along with the five lemma applied to the following diagram with exact rows:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \Gamma_{k-1}/\Gamma_k & \hookrightarrow & \pi/\Gamma_k & \twoheadrightarrow & \pi/\Gamma_{k-1} & \longrightarrow & 1 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 1 & \longrightarrow & \Gamma_{k-1}^\Sigma/\Gamma_k^\Sigma & \hookrightarrow & \pi^\Sigma/\Gamma_k^\Sigma & \twoheadrightarrow & \pi^\Sigma/\Gamma_{k-1}^\Sigma & \longrightarrow & 1 \end{array}$$

For the $k = 2$ base case, it suffices to prove $H_1(E_L; \mathbb{Z}[\pi_1(M)]) \rightarrow H_1(E_\Sigma; \mathbb{Z}[\pi_1(M)])$ induced by inclusion is an isomorphism since

$$\begin{aligned} H_1(E_L; \mathbb{Z}[\pi_1(M)]) &= \Gamma/\Gamma_2, \\ H_1(E_\Sigma; \mathbb{Z}[\pi_1(M)]) &= \Gamma^\Sigma/\Gamma_2^\Sigma, \\ \pi/\Gamma &= \pi^\Sigma/\Gamma^\Sigma = \pi_1(M). \end{aligned}$$

Consider the following diagram of exact sequences of pairs (all coefficients are in $\mathbb{Z}[\pi_1(M)]$, which we suppress):

$$\begin{array}{ccccccc} H_2(M) & \longrightarrow & H_2(M, E_L) & \twoheadrightarrow & H_1(E_L) & \longrightarrow & 0 \\ \parallel & & \downarrow & & \downarrow & & \\ H_2(M \times [0, 1]) & \longrightarrow & H_2(M \times [0, 1], E_\Sigma) & \twoheadrightarrow & H_1(E_\Sigma) & \longrightarrow & 0 \end{array}$$

It suffices to prove that $H_2(M, E_L; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(M \times [0, 1], E_\Sigma; \mathbb{Z}[\pi_1(M)])$ is an isomorphism. By excision,

$$\begin{aligned} H_2(M, E_L; \mathbb{Z}[\pi_1(M)]) &\cong H_2(vL, \partial E_L; \mathbb{Z}[\pi_1(M)]), \\ H_2(M \times [0, 1], E_\Sigma; \mathbb{Z}[\pi_1(M)]) &\cong H_2(v\Sigma, \partial v\Sigma; \mathbb{Z}[\pi_1(M)]). \end{aligned}$$

Note that $H_2(v\Sigma; \mathbb{Z}[\pi_1(M)]) = 0$ since the components of Σ are surfaces with boundary. Now consider the diagram of exact sequences of pairs below (again, all coefficients are in $\mathbb{Z}[\pi_1(M)]$):

$$\begin{array}{ccccccc} 0 & \longrightarrow & H_2(vL, \partial E_L) & \longrightarrow & H_1(\partial E_L) & \twoheadrightarrow & H_1(vL) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & H_2(v\Sigma, \partial v\Sigma) & \longrightarrow & H_1(\partial v\Sigma) & \twoheadrightarrow & H_1(v\Sigma) \longrightarrow 0 \end{array}$$

Even with coefficients in $\mathbb{Z}[\pi_1(M)]$, we have splittings

$$\begin{aligned} H_1(\partial E_L; \mathbb{Z}[\pi_1(M)]) &\cong H_1(L; \mathbb{Z}[\pi_1(M)]) \oplus H_1\left(\bigsqcup m_i; \mathbb{Z}[\pi_1(M)]\right), \\ H_1(\partial v\Sigma; \mathbb{Z}[\pi_1(M)]) &\cong H_1(\Sigma; \mathbb{Z}[\pi_1(M)]) \oplus H_1\left(\bigsqcup m_i; \mathbb{Z}[\pi_1(M)]\right), \end{aligned}$$

where the m_i denote both meridians of L and meridians of Σ . The homomorphism

$$H_1(\partial E_L; \mathbb{Z}[\pi_1(M)]) \rightarrow H_1(\partial v\Sigma; \mathbb{Z}[\pi_1(M)])$$

may be written as a direct sum with respect to these splittings. Furthermore, the kernels of the rightmost horizontal maps in the diagram are precisely $H_1(\bigsqcup m_i; \mathbb{Z}[\pi_1(M)])$. Thus, we have (with coefficients in $\mathbb{Z}[\pi_1(M)]$)

$$\begin{aligned} H_2(vL, \partial E_L) &\cong \ker(H_1(\partial E_L) \rightarrow H_1(vL)) \cong H_1\left(\bigsqcup m_i\right) \\ &\cong \ker(H_1(\partial v\Sigma) \rightarrow H_1(v\Sigma)) \cong H_2(v\Sigma, \partial v\Sigma). \end{aligned}$$

Therefore, the map $\Gamma/\Gamma_2 \rightarrow \Gamma^\Sigma/\Gamma_2^\Sigma$ induced by inclusion is an isomorphism. By the five lemma, $\pi/\Gamma_2 \rightarrow \pi^\Sigma/\Gamma_2^\Sigma$ is an isomorphism. This proves the base case $k = 2$.

Now suppose $\pi/\Gamma_k \xrightarrow{\cong} \pi^\Sigma/\Gamma_k^\Sigma$ is an isomorphism for some $2 \leq k \leq n$. We will obtain an isomorphism $\pi/\Gamma_{k+1} \xrightarrow{\cong} \pi^\Sigma/\Gamma_{k+1}^\Sigma$. We invoke the 5-term natural exact sequence of Stallings [1965]; see the proof of the Stallings–Dwyer theorem in [Cha and Orr 2024]. We have the following diagram, where all homology groups have coefficients in $\mathbb{Z}[\pi_1(M)]$:

$$\begin{array}{ccccccccc} H_2(E_L) & \twoheadrightarrow & H_2(\pi) & \longrightarrow & H_2(\pi/\Gamma_k) & \twoheadrightarrow & \Gamma_k/\Gamma_{k+1} & \longrightarrow & 0 \\ \downarrow & & \downarrow & & \downarrow \cong & & \downarrow & & \\ H_2(E_\Sigma) & \twoheadrightarrow & H_2(\pi^\Sigma) & \longrightarrow & H_2(\pi^\Sigma/\Gamma_k^\Sigma) & \twoheadrightarrow & \Gamma_k^\Sigma/\Gamma_{k+1}^\Sigma & \longrightarrow & 0 \end{array}$$

Note that $H_2(\pi; \mathbb{Z}[\pi_1(M)])$ is the second homology of the $\pi_1(M)$ -cover of a $K(\pi, 1)$, which is a $K(\Gamma, 1)$. Therefore, $H_2(\pi; \mathbb{Z}[\pi_1(M)]) = H_2(\Gamma; \mathbb{Z})$, and similarly for the other group homology terms appearing above. The surjectivity of the first map in each row follows from Hopf’s theorem, which gives an exact sequence

$$\pi_2(X) \rightarrow H_2(X) \twoheadrightarrow H_2(\pi_1(X)) \rightarrow 0$$

for any CW complex X ; see [Brown 1982, Section II.5]. The rest of each row is exact by Stallings [1965]. (Each Stallings 5-term exact sequence we use has an isomorphism between the final two terms.) The third vertical arrow is an isomorphism by the induction hypothesis. By the diagram above, we obtain an isomorphism $\Gamma_k/\Gamma_{k+1} \xrightarrow{\cong} \Gamma_k^\Sigma/\Gamma_{k+1}^\Sigma$ if we prove

$$\mathrm{im}(H_2(E_L; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(\pi/\Gamma_k; \mathbb{Z}[\pi_1(M)]))$$

is mapped isomorphically to

$$\mathrm{im}(H_2(E_\Sigma; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(\pi^\Sigma/\Gamma_k^\Sigma; \mathbb{Z}[\pi_1(M)])).$$

We then apply the five lemma to obtain the desired isomorphism $\pi/\Gamma_{k+1} \xrightarrow{\cong} \pi^\Sigma/\Gamma_{k+1}^\Sigma$.

We build a tower of spaces $X_j(\Sigma)$, analogous to $X_j(L)$, similarly to the proof of Proposition 4.1. The space $X_j(\Sigma)$ has the decomposition

$$X_j(\Sigma) = M_{p_j^\Sigma}^\times \cup_{\partial v\Sigma} v\Sigma.$$

For each j , there is a canonical morphism of diagrams $\mathcal{D}_j(L) \rightarrow \mathcal{D}_j(\Sigma)$ which induces a map $X_j(L) \rightarrow X_j(\Sigma)$ restricting to inclusions $vL \hookrightarrow v\Sigma$ and $E_L \hookrightarrow E_\Sigma$. Thus, we have the following commutative

square of Mayer–Vietoris sequences, where all coefficients are in $\mathbb{Z}[\pi_1(M)]$:

$$\begin{array}{ccccccc}
 H_2(\partial E_L) & \longrightarrow & H_2(E_L) & \longrightarrow & H_2(M) & \longrightarrow & H_1(\partial E_L) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 H_2(\partial E_L) & \longrightarrow & H_2(\pi/\Gamma_k) & \longrightarrow & H_2(X_k(L)) & \longrightarrow & H_1(\partial E_L) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 H_2(\partial v\Sigma) & \longrightarrow & H_2(E_\Sigma) & \longrightarrow & H_2(M \times [0, 1]) & \longrightarrow & H_1(\partial v\Sigma) \\
 \downarrow & \searrow & \downarrow & \searrow & \downarrow & \searrow & \downarrow \\
 H_2(\partial v\Sigma) & \longrightarrow & H_2(\pi^\Sigma/\Gamma_k^\Sigma) & \longrightarrow & H_2(X_k(\Sigma)) & \longrightarrow & H_1(\partial v\Sigma)
 \end{array}$$

We use this diagram to study the images of $H_2(E_L) \rightarrow H_2(\pi/\Gamma_k)$ and $H_2(E_\Sigma) \rightarrow H_2(\pi^\Sigma/\Gamma_k^\Sigma)$. Observe that $H_2(vL) = H_2(v\Sigma) = 0$ and $H_1(\partial E_L) \rightarrow H_1(\partial v\Sigma)$ is injective.

First, we observe that

$$\text{im}(H_2(E_L) \rightarrow H_2(\pi/\Gamma_k) \rightarrow H_2(\pi^\Sigma/\Gamma_k^\Sigma)) \subseteq \text{im}(H_2(E_\Sigma) \rightarrow H_2(\pi^\Sigma/\Gamma_k^\Sigma)).$$

The reverse inclusion is a diagram chase which we briefly describe. Let $x \in \text{im}(H_2(E_\Sigma) \rightarrow H_2(\pi^\Sigma/\Gamma_k^\Sigma))$. Denote the image of x in $H_2(X_k(\Sigma))$ by y . Then $y \in \text{im}(H_2(M \times [0, 1]) \rightarrow H_2(X_k(\Sigma)))$. We may then consider y as the image of some $z \in H_2(M)$. The image $\Delta(z)$ of z under the connecting homomorphism is 0 because $\Delta(y) = 0$ and $H_1(\partial E_L) \hookrightarrow H_1(\partial v\Sigma)$ is injective. Thus, z is the image of some $\tilde{z} \in H_2(E_L)$. Since the images of $\tilde{z}$ and x agree in $H_2(X_k(\Sigma))$, the image of $\tilde{z}$ in $H_2(\pi^\Sigma/\Gamma_k^\Sigma)$ differs from x by the image of some element in $H_2(\partial v\Sigma)$. By Lemma 6.1, since Σ is an n -cobordism, this difference is represented by an element in the image of $H_2(\partial E_L)$. Adjusting $\tilde{z}$ by this element, we obtain an element $\tilde{x} \in H_2(E_L)$ which maps to $x \in H_2(\pi^\Sigma/\Gamma_k^\Sigma)$. This completes the proof of the claim. The induction carries through the step from $k = n$ to $k = n + 1$, after which Lemma 6.1 no longer applies since an n -cobordism only guarantees the desired property through level n .

The same argument shows that the inclusion $E_{L'} \hookrightarrow E_\Sigma$ induces isomorphisms $\pi'/\Gamma'_k \xrightarrow{\cong} \pi^\Sigma/\Gamma_k^\Sigma$ for all $k \leq n + 1$. We will take ϕ , part of the data for the desired n -basing, to be the composite $\phi : \pi'/\Gamma'_n \xrightarrow{\cong} \pi^\Sigma/\Gamma_n^\Sigma \xrightarrow{\cong} \pi/\Gamma_n$. To complete the proof of the proposition, we must find a canonical admissible isomorphism $\phi_\partial : H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L)$ compatible with ϕ . Define ϕ_∂ as follows: Send each meridian m'_i of L' to the corresponding meridian m_i of L . A longitude λ'_i for the component L'_i of L' determines a corresponding copy of Σ_i in $\partial v\Sigma$, which we also denote by Σ_i . Send λ'_i to the longitude λ_i for the component L_i of L which is homologous to λ'_i in this pushoff. The isomorphism ϕ_∂ is admissible.

We now check that ϕ_∂ is compatible with ϕ . We may assume that $\Sigma \subset M \times [0, 1]$ misses $* \times [0, 1]$. To simplify notation, we work inside the reduced cylinder $(M \times [0, 1])^\times$. Choose any basing τ' for L' . Also choose any basing τ^0 for L . We will eventually modify τ^0 . We consider these basings in $E_{L'}$ and E_L , respectively, and assume for each i that the basings of the components L'_i and L_i meet $\partial v\Sigma_i$ in Σ_i .

For each i , choose a path γ_i on Σ_i connecting the basings of L'_i and L_i and forming a based loop α_i together with the two basings. The α_i represent elements in π^Σ . Let β_i denote a lift of the image

of α_i under the composite $\pi^\Sigma \twoheadrightarrow \pi^\Sigma / \Gamma_n^\Sigma \xrightarrow{\cong} \pi / \Gamma_n$ to an element of π . Modify the i -th factor of the basing τ^0 of L by β_i . Call this new basing $\tau = \beta\tau^0$. We claim that τ is the desired basing for L , that is, the following diagram commutes:

$$\begin{array}{ccc} \pi_1(\partial E_{L'} \cup \tau') & \xrightarrow[\phi_{\partial*}]{\cong} & \pi_1(\partial E_L \cup \tau) \\ \downarrow & & \downarrow \\ \pi' / \Gamma'_n & \xrightarrow[\cong]{} \pi^\Sigma / \Gamma_n^\Sigma \xleftarrow[\cong]{} & \pi / \Gamma_n \end{array}$$

It is enough to check that this commutes for (based) meridians and longitudes of L' . We have two basings for L in E_Σ : the basing $\tau = \beta\tau^0$ (which is a basing in E_L), and the basing $\tau'\gamma$ given by following the basing for L' with the paths on Σ connecting L' to L . The based meridians and longitudes of L given by these two different basings agree in $\pi^\Sigma / \Gamma_n^\Sigma$: In $\pi^\Sigma / \Gamma_n^\Sigma$, the images of β_i and α_i agree. But $\alpha_i\tau_i^0$ is homotopic rel endpoints to $\tau'_i\gamma_i$. It is therefore enough to show that the following diagram commutes:

$$\begin{array}{ccc} \pi_1(\partial E_{L'} \cup \tau') & \xrightarrow[\phi_{\partial*}]{\cong} & \pi_1(\partial E_L \cup \tau'\gamma) \\ \downarrow & & \downarrow \\ \pi' / \Gamma'_n & \xrightarrow[\cong]{} & \pi^\Sigma / \Gamma_n^\Sigma \end{array}$$

This commutes on meridians, as the based meridians $\tau'\gamma$ induces for L are just the based meridians τ' induces for L' dragged down Σ along γ . To see that this diagram also commutes on longitudes, we consider τ' as a basing for Σ and leverage the definition of n -cobordism. Consider the based longitudes $l'_i = \tau'_i\lambda'_i$ for L' and $l_i = \tau'_i\gamma_i\lambda_i$ for L as elements of $\pi_1(\Sigma_i)$ (with the basing τ'_i). Since the two longitudes cobound Σ_i , the product $l'_i l_i^{-1}$ is equal to $\prod_j [a_j^i, b_j^i]$, where $\{a_j^i, b_j^i\}_j$ is a symplectic basis for Σ_i with appropriate basings. Since Σ is an n -cobordism, the image of $\pi_1(\Sigma_i)$ in $\pi^\Sigma / \Gamma_n^\Sigma$ is cyclic, so $l'_i l_i^{-1} = \prod_j [a_j^i, b_j^i] \mapsto 1 \in \pi^\Sigma / \Gamma_n^\Sigma$. Thus, the above diagram commutes, and (ϕ, ϕ_∂) is a canonical n -basing for L' relative to L induced by Σ . $\square$

6.2 Proof of Theorem B

We now move toward the proof of Theorem B, which asserts the equivalence of n -cobordism and equality of h -invariants; see Section 3.2 for the statement of the theorem. We first prove Theorem B₀, which is just Theorem B in the case where one of the two links we compare is the fixed link L .

Theorem B₀ *Fix an m -component link $L \subset M$. Let $L' \subset M$ be another m -component link, and suppose there exists an n -basing (ϕ, ϕ_∂) for L' relative to L . Then the following statements are equivalent:*

- (1) *The invariant $h_n(L', \phi)$ vanishes.*
- (2) *The links L and L' are n -cobordant via an n -cobordism inducing the n -basing (ϕ, ϕ_∂) .*

As a consequence, the following basing-independent statements are equivalent:

- (1') *The invariant $h_n(L')$ vanishes.*
- (2') *The links L and L' are n -cobordant.*

Proof The equivalence of (1') and (2') will be a consequence of the equivalence of (1) and (2).

(1) $\implies$ (2): Suppose $h_n(L', \phi)$ vanishes. Then $h_n(L', \phi) = h_n(L, \text{id})$, so there is a based homotopy $H : M \times [0, 1] \rightarrow X_n(L)$ restricting to $h_n(L, \text{id})$ on $M \times \{0\}$ and $h_n(L', \phi)$ on $M \times \{1\}$. Take the transverse preimage of $L \subset X_n(L)$ under H to obtain a properly embedded surface $\Sigma \subset M \times [0, 1]$. Since L and L' are the respective preimages of L under $h_n(L, \text{id})$ and $h_n(L', \phi)$, the surface Σ restricts to L and L' on $M \times \{0\}$ and $M \times \{1\}$, respectively.

Consider the surface exterior E_Σ , and note that H restricts to a map $E_\Sigma \rightarrow M_{p_n}^\times$. This induces a homomorphism $\pi^\Sigma \rightarrow \pi / \Gamma_n$ over $\pi_1(M)$, which induces $\pi^\Sigma / \Gamma_n^\Sigma \rightarrow \pi / \Gamma_n$, where $\pi^\Sigma = \pi_1(E_\Sigma)$ and $\Gamma_n^\Sigma = \ker(\pi^\Sigma \twoheadrightarrow \pi_1(M))$. We also have the homomorphism $\pi \rightarrow \pi^\Sigma$ induced by inclusion, which induces a homomorphism $\pi / \Gamma_n \rightarrow \pi^\Sigma / \Gamma_n^\Sigma$. Since the homotopy H restricts to $h_n(L, \text{id})$ on $M \times \{0\}$, the composition $\pi / \Gamma_n \rightarrow \pi^\Sigma / \Gamma_n^\Sigma \rightarrow \pi / \Gamma_n$ is $\text{id}_{\pi / \Gamma_n}$. In particular, $\pi / \Gamma_n \rightarrow \pi^\Sigma / \Gamma_n^\Sigma$ is injective and $\pi^\Sigma / \Gamma_n^\Sigma \rightarrow \pi / \Gamma_n$ is surjective.

The surface Σ would be our candidate for an n -cobordism between the two links, except it may have too many components. In Claim 1, we perform surgeries on Σ to ensure the preimage of each component of L has only one component. More precisely, we surger $M \times [0, 1]$ by attaching canceling pairs of 5-dimensional 1- and 2-handles to the interior of $(M \times [0, 1]) \times \{1\} \subset (M \times [0, 1]) \times [0, 1]$. The resulting surgery cobordism contains a cobordism from Σ to a new surface in $M \times [0, 1]$ with one fewer component. The surjectivity of $\pi^\Sigma / \Gamma_n^\Sigma \twoheadrightarrow \pi / \Gamma_n$ is key for guaranteeing that $H : M \times [0, 1] \rightarrow X_n(L)$ extends over the surgery cobordism. In Claim 2, we argue that $\pi / \Gamma_n \hookrightarrow \pi^\Sigma / \Gamma_n^\Sigma$ is an isomorphism. We use an induction argument, proving $\Gamma_k / \Gamma_{k+1} \rightarrow \Gamma_k^\Sigma / \Gamma_{k+1}^\Sigma$ is an isomorphism for all $k < n$ and applying the five lemma.

Claim 1 We may modify the surface Σ and the homotopy H rel boundary so that the preimage of each component of L under H is a connected surface with boundary.

Suppose the preimage of some component L_i of L under H has more than one component. Choose some component Σ'_i of $H^{-1}(L_i)$, and let Σ_i be the union of the remaining components. Find a path α in E_Σ from the interior of $\partial\nu\Sigma_i$ to the interior of $\partial\nu\Sigma'_i$ whose boundary maps under H to a single point of $\partial\nu L_i$ and whose image in $M_{p_n}^\times$ is nullhomotopic. Such a path exists since $\pi^\Sigma / \Gamma_n^\Sigma \twoheadrightarrow \pi / \Gamma_n$ is surjective.

We now perform surgeries to connect Σ'_i to Σ_i , modifying the map H rel boundary. Connect α to Σ_i and Σ'_i by small intervals in D^2 -fibers of the regular neighborhoods of the surfaces whose images in $X_n(L)$ agree. We also call this extended path α . Its endpoints map to some point $p \in L_i \subset X_n(L)$. Parametrize the closure of a neighborhood of p as $D^1 \times D^2$, where $D^1 \subset L_i$ is the closure of a neighborhood of p in L_i . The endpoints of α have neighborhoods whose closures may be parametrized as $D^1 \times D^1 \times D^2$, which map under H to $D^1 \times D^2$ by projection onto the last two coordinates.

Attach a 5-dimensional 1-handle

$$D^1 \times D^4 = D^1 \times (D^1 \times D^1 \times D^2) \quad \text{to} \quad (M \times [0, 1]) \times \{1\} \subset (M \times [0, 1]) \times [0, 1]$$

along the two 4-balls $D^1 \times D^1 \times D^2$ to obtain a surgery cobordism Z from $M \times [0, 1]$ to a 4-manifold W which is itself a cobordism from M to M , the result of surgery on $M \times [0, 1]$ corresponding to the

attachment of the 1-handle. Extend the map $H \times [0, 1] : (M \times [0, 1]) \times [0, 1] \rightarrow X_n(L)$ to Z by projection of $D^1 \times D^4$ to D^4 followed by H , which agrees on $\{0\} \times D^4$ and $\{1\} \times D^4$. The map $Z \rightarrow X_n(L)$ restricts to a map $H' : W \rightarrow X_n(L)$ which further restricts to $h_n(L, \text{id})$ and $h_n(L', \phi)$ on $\partial W = M \sqcup M$. Define a surface $\Sigma'_i \subset W$ by removing from Σ_i and Σ'_i the D^2 -neighborhoods of the endpoints of α and gluing in $D^1 \times S^1 \subset D^1 \times D^4$. The surface Σ'_i has one fewer component than $\Sigma_i \cup \Sigma'_i$.

We now attach a 5-dimensional 2-handle to Z to cancel the 1-handle described above. The attaching circle for this 2-handle is the path α along with the core of the 1-handle pushed to the boundary of Z . The map $Z \rightarrow X_n(L)$ extends over the 2-handle with the image of the 2-handle missing $\nu L \subset X_n(L)$ since the image of the attaching circle in $M_{p_n}^\times$ is nullhomotopic. Because the handles cancel, the resulting space is homeomorphic to $(M \times [0, 1]) \times [0, 1]$. The resulting map $(M \times [0, 1]) \times [0, 1] \rightarrow X_n(L)$ restricts to a map $H'' : (M \times [0, 1]) \times \{1\} \rightarrow X_n(L)$ which is another homotopy from $h_n(L, \text{id})$ to $h_n(L', \phi)$.

By construction, $(H'')^{-1}(L)$ has one fewer component than $H^{-1}(L)$. We may continue inductively, reducing the number of components of $H^{-1}(L)$ until we reach an m -component surface in $M \times [0, 1]$ cobounded by L and L' which is the preimage of L under a homotopy $M \times [0, 1] \rightarrow X_n(L)$ from $h_n(L, \text{id})$ to $h_n(L', \phi)$. We continue to call this surface Σ and this homotopy H . This proves Claim 1.

Claim 2 *The new surface $\Sigma \subset M \times [0, 1]$ induces an isomorphism $\pi / \Gamma_n \xrightarrow{\cong} \pi^\Sigma / \Gamma_n^\Sigma$.*

The proof of Claim 2 is analogous to part of the proof of Proposition 3.7, so we will omit some details. Since we modified the homotopy H rel boundary, the new homotopy obtained through the proof of Claim 1, which we still call H , restricts to $h_n(L, \text{id})$ on $M \times \{0\}$. Thus, the composition $\pi / \Gamma_n \rightarrow \pi^\Sigma / \Gamma_n^\Sigma \rightarrow \pi / \Gamma_n$ remains $\text{id}_{\pi / \Gamma_n}$. In fact, the composition $\pi / \Gamma_k \rightarrow \pi^\Sigma / \Gamma_k^\Sigma \rightarrow \pi / \Gamma_k$ is $\text{id}_{\pi / \Gamma_k}$ for all $k \leq n$. So $\pi / \Gamma_k \hookrightarrow \pi^\Sigma / \Gamma_k^\Sigma$ is injective for all $k \leq n$.

The proof of Claim 2 for $n = 2$ is entirely analogous to part of the proof of Proposition 3.7. We now use induction. Suppose $\pi / \Gamma_k \rightarrow \pi^\Sigma / \Gamma_k^\Sigma$ is an isomorphism for some $2 \leq k < n$. We have already seen that $\pi / \Gamma_{k+1} \hookrightarrow \pi^\Sigma / \Gamma_{k+1}^\Sigma$ is injective. Consider the diagram of short exact sequences below:

$$\begin{array}{ccccccc} 1 & \longrightarrow & \Gamma_k / \Gamma_{k+1} & \hookrightarrow & \pi / \Gamma_{k+1} & \twoheadrightarrow & \pi / \Gamma_k \longrightarrow 1 \\ & & \downarrow & & \downarrow & & \downarrow \cong \\ 1 & \longrightarrow & \Gamma_k^\Sigma / \Gamma_{k+1}^\Sigma & \hookrightarrow & \pi^\Sigma / \Gamma_{k+1}^\Sigma & \twoheadrightarrow & \pi^\Sigma / \Gamma_k^\Sigma \longrightarrow 1 \end{array}$$

To prove $\pi / \Gamma_{k+1} \hookrightarrow \pi^\Sigma / \Gamma_{k+1}^\Sigma$ is also surjective, it suffices to show $\Gamma_k / \Gamma_{k+1} \hookrightarrow \Gamma_k^\Sigma / \Gamma_{k+1}^\Sigma$ is surjective. We again invoke the Stallings 5-term natural exact sequence [1965] in the following diagram:

$$\begin{array}{ccccccc} H_2(\Gamma) & \longrightarrow & H_2(\Gamma / \Gamma_k) & \twoheadrightarrow & \Gamma_k / \Gamma_{k+1} & \longrightarrow & 0 \\ \downarrow & & \downarrow \cong & & \downarrow & & \\ H_2(\Gamma^\Sigma) & \longrightarrow & H_2(\Gamma^\Sigma / \Gamma_k^\Sigma) & \twoheadrightarrow & \Gamma_k^\Sigma / \Gamma_{k+1}^\Sigma & \longrightarrow & 0 \end{array}$$

Since the composite $H_2(\Gamma / \Gamma_k) \xrightarrow{\cong} H_2(\Gamma^\Sigma / \Gamma_k^\Sigma) \twoheadrightarrow \Gamma_k^\Sigma / \Gamma_{k+1}^\Sigma$ is surjective, so is $\Gamma_k / \Gamma_{k+1} \rightarrow \Gamma_k^\Sigma / \Gamma_{k+1}^\Sigma$. This proves Claim 2.

We now have a diagram of groups induced by the surface $\Sigma \subset M \times [0, 1]$ and the homotopy H :

$$\begin{array}{ccc}
 \pi' / \Gamma'_n & & \\
 \downarrow & \searrow \phi & \\
 \pi^\Sigma / \Gamma_n^\Sigma & \xrightarrow{\cong} & \pi / \Gamma_n \\
 \uparrow \cong & \nearrow & \\
 \pi / \Gamma_n & &
 \end{array}$$

Thus, in the diagram we also have isomorphisms $\pi' / \Gamma'_n \xrightarrow{\cong} \pi^\Sigma / \Gamma_n^\Sigma$ and $\pi^\Sigma / \Gamma_n^\Sigma \xrightarrow{\cong} \pi / \Gamma_n$. In particular, since the map $M \times [0, 1] \rightarrow X_n(L)$ sends Σ to L , for each component Σ_i of Σ the image of $\pi_1(\Sigma_i)$ is sent to the image of $\pi_1(L_i)$ under the isomorphism $\pi^\Sigma / \Gamma_n^\Sigma \xrightarrow{\cong} \pi / \Gamma_n$. Because the inverse of this isomorphism is induced by the inclusion $E_L \hookrightarrow E_\Sigma$, the surface Σ is an n -cobordism by definition. Furthermore, Σ induces the isomorphisms ϕ and ϕ_∂ by construction. This proves (1) $\implies$ (2).

(2) $\implies$ (1): Now suppose $\Sigma \subset M \times [0, 1]$ is an n -cobordism between L and L' . By Proposition 3.7, Σ induces some n -basing, say, (ϕ, ϕ_∂) . This yields inclusions of the pairs $(E_L, \partial E_L)$ and $(E_{L'}, \partial E_{L'})$ into $(E_\Sigma, \partial \nu \Sigma)$, which induce the isomorphism ϕ_∂ on the boundary. We have a diagram of isomorphisms

$$\begin{array}{ccc}
 \pi' / \Gamma'_n & & \\
 \cong \downarrow & \searrow \phi & \\
 \pi^\Sigma / \Gamma_n^\Sigma & \xrightarrow{\cong} & \pi / \Gamma_n \\
 \cong \uparrow & \nearrow & \\
 \pi / \Gamma_n & &
 \end{array}$$

Our goal is a diagram

$$\begin{array}{ccc}
 (E_{L'}, \partial E_{L'}) & & \\
 \downarrow & \searrow & \\
 (E_\Sigma, \partial \nu \Sigma) & \dashrightarrow & (M_{p_n}^\times, \partial E_L) \\
 \uparrow & \nearrow & \\
 (E_L, \partial E_L) & &
 \end{array}$$

of pairs of spaces. We wish to construct the map $(E_\Sigma, \partial \nu \Sigma) \rightarrow (M_{p_n}^\times, \partial E_L)$. We first extend the maps on E_L and $E_{L'}$ to $\partial \nu \Sigma$. Choose a basing for L (which induces a basing on Σ), and choose a trivialization of νL (hence $\nu \Sigma$ and $\nu L'$). For each i , we have the diagram

$$\begin{array}{ccc}
 \pi_1(L_i) & \longrightarrow & \pi / \Gamma_n \\
 \downarrow & & \downarrow \cong \\
 \pi_1(\Sigma_i) & \longrightarrow & \pi^\Sigma / \Gamma_n^\Sigma
 \end{array}$$

where L_i and Σ_i are parallels of the cores of νL and $\nu \Sigma$, respectively, induced by the trivializations.

Since Σ is an n -cobordism, the image of $\pi_1(L_i)$ in $\pi^\Sigma/\Gamma_n^\Sigma$ is onto the image of $\pi_1(\Sigma_i)$. Because $\pi/\Gamma_n \xrightarrow{\cong} \pi^\Sigma/\Gamma_n^\Sigma$ is an isomorphism, the same is true in π/Γ_n . We extend the map to the 1-skeleton of Σ by sending each based loop in Σ to a based loop in L whose image agrees in π/Γ_n . There is no obstruction to defining a map from the 2-skeleton of Σ (i.e., all of Σ) to L . To complete the extension to $\partial\nu\Sigma$, we use the trivializations to map $\partial\nu\Sigma \cong S^1 \times \Sigma \rightarrow S^1 \times L \cong \partial E_L$. We now have a map of pairs $(E_L \cup \partial\nu\Sigma \cup E_{L'}, \partial\nu\Sigma) \rightarrow (M_{p_n}^\times, \partial E_L)$.

As $M_{p_n}^\times$ is an Eilenberg–Mac Lane space, there is no obstruction to extending to all of E_Σ given the diagram of group isomorphisms above. Thus, we obtain a map of pairs $(E_\Sigma, \partial\nu\Sigma) \rightarrow (M_{p_n}^\times, \partial E_L)$, which fits into the diagram above. This map restricts to a bundle map $\partial\nu\Sigma \rightarrow \partial E_L$ from the sphere bundle of Σ to the sphere bundle of $L \subset X_n(L)$. Extend to the corresponding disk bundles $\nu\Sigma$ and νL by coning on fibers. This yields a map $H : M \times [0, 1] \rightarrow X_n(L)$ which is a homotopy between $h_n(L, \text{id})$ and $h_n(L', \phi)$. Thus, the invariant $h_n(L', \phi)$ vanishes.

We have now proven the equivalence of (1) and (2). To complete the proof of the theorem, we show the corresponding basing-independent statements are equivalent.

(1') $\iff$ (2'): The invariant $h_n(L')$ vanishes if and only if $h_n(L', \phi)$ vanishes for some n -basing ϕ . By the equivalence of (1) and (2), this is true if and only if L and L' are n -cobordant via some n -cobordism inducing the n -basing ϕ , i.e., if and only if L and L' are n -cobordant. $\square$

Theorem B, which compares links L' and L'' relative to the fixed link L , follows readily from Theorem B₀.

Proof of Theorem B Fix a link $L \subset M$. Let $L', L'' \subset M$ be links which admit n -basings (ϕ', ϕ'_∂) and $(\phi'', \phi''_\partial)$, respectively, relative to L . Lemmas 4.4 and 4.5 imply $(\psi, \psi_\partial) = ((\phi')^{-1} \circ \phi'', (\phi'_\partial)^{-1} \circ \phi''_\partial)$ is an n -basing of L'' relative to L' . Let $\bar{h}_{\phi'} : (X_n(L'), \nu L') \xrightarrow{\cong} (X_n(L), \nu L)$ be the homotopy equivalence of pairs induced by the n -basing (ϕ', ϕ'_∂) (provided by Proposition 4.2). Observe that

$$\bar{h}_{\phi'} \circ h_n(L', \text{id}) = h_n(L', \phi') \quad \text{and} \quad \bar{h}_{\phi'} \circ h_n(L'', \psi) = h_n(L'', \phi'').$$

By Theorem B₀, L' and L'' are n -cobordant via an n -cobordism inducing the n -basing ψ if and only if $h_n(L'', \psi) = h_n(L', \text{id})$. Since $\bar{h}_{\phi'}$ is a homotopy equivalence, this holds if and only if $\bar{h}_{\phi'} \circ h_n(L'', \psi) = \bar{h}_{\phi'} \circ h_n(L', \text{id})$, that is, $h_n(L'', \phi'') = h_n(L', \phi') \in [M, X_n(L)]_0$.

The equivalence of the analogous basing-independent statements follows immediately from the equivalence of (1) and (2) in a similar manner to the proof of Theorem B₀. $\square$

6.3 Proof of Theorem C

Recall that Theorem C states precise conditions under which a homotopy class $f \in [M, X_n(L)]_0$ is realized as $h_n(L', \phi)$ for some link $L' \subset M$ and some n -basing ϕ ; see Section 3.3 for the full statement. Some elements of the proof of Theorem C mirror arguments from the proof of Theorem B₀.

Proof ($\implies$): Suppose $f \in [M, X_n(L)]_0$ is realizable. Then $f = h_n(L', \phi)$ for some m -component link $L' \subset M$ which admits the n -basing (ϕ, ϕ_∂) relative to L . By Proposition 4.2, the n -basing ϕ induces a

canonical homotopy equivalence of pairs

$$\bar{h}_\phi : (X_n(L'), \nu L') \xrightarrow{\simeq} (X_n(L), \nu L).$$

Since $f = h_n(L', \phi)$, we have by definition that $f \simeq \bar{h}_\phi \circ \iota_n(L')$. Recall that we obtain the diagram

$$\begin{array}{ccc} M & \xrightarrow{\iota_n(L')} & X_n(L') \\ & \searrow & \downarrow \bar{h}_\phi \\ & & B\pi_1(M) \\ & \swarrow & \uparrow \\ M & \xrightarrow{\iota_n(L)} & X_n(L) \end{array}$$

We can see that $f_* = (\bar{h}_\phi)_* \circ \iota_n(L')_* = \iota_n(L)_*$, that is, condition (2) holds, from the induced diagram on fundamental groups:

$$\begin{array}{ccc} \pi_1(M) & \xrightarrow{\cong} & \pi_1(X_n(L')) \\ & \searrow & \downarrow \cong \\ & & \pi_1(M) \\ & \swarrow & \uparrow \cong \\ \pi_1(M) & \xrightarrow{\cong} & \pi_1(X_n(L)) \end{array}$$

Next, since $\bar{h}_\phi$ is a map of pairs, the composition $\Delta \circ f_* = \Delta \circ (\bar{h}_\phi)_* \circ \iota_n(L')_*$ sends $[M]$ to $\sum_i [T_i]$, so condition (1) holds.

Finally, we show condition (3) holds, which concerns the homomorphism

$$\text{proj} \circ (\cap f_*[M]) : H^1(X_n(L); \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(X_n(L); \mathbb{Z}[\pi_1(M)]/\text{im}(K_j(\pi/\Gamma_n))).$$

Consider the diagram

$$\begin{array}{ccccccc} & & H^1(M) & \xleftarrow[\cong]{f_*} & H^1(X_n(L)) & & \\ & & \cap[M] \downarrow \cong & & \downarrow \cap f_*[M] & \searrow \text{proj} \circ (\cap f_*[M]) & \\ (*) & H_2(E_{L'}) & \xrightarrow{j'_*} & H_2(M) & \xrightarrow{f_*} & H_2(X_n(L)) & \xrightarrow{\text{proj}} H_2(X_n(L))/\text{im}(K_j(\pi/\Gamma_n)) \\ & \downarrow \cong & & & \uparrow & & \uparrow \\ & H_2(\pi') & \longrightarrow & H_2(\pi'/\Gamma'_n) & \xrightarrow[\cong]{\phi_*} & H_2(\pi/\Gamma_n) & \xrightarrow{\text{proj}} H_2(\pi/\Gamma_n)/K_j(\pi/\Gamma_n) \end{array}$$

where all coefficients are in $\mathbb{Z}[\pi_1(M)]$ and $j < n$. Recall that $K_j(\pi/\Gamma_n)$ is the kernel of $H_2(\pi/\Gamma_n) \rightarrow H_2(\pi/\Gamma_j)$. The map $\cap[M]$ is an isomorphism by Poincaré duality. The map ϕ_* is an isomorphism by the assumption that ϕ is an n -basing.

To prove that the map f^* on the top row is an isomorphism, we invoke the (cohomology) universal coefficient spectral sequence, a cohomology spectral sequence with

$$E_2^{p,q} = \text{Ext}_{\mathbb{Z}[\pi_1(M)]}^p(H_q(-; \mathbb{Z}[\pi_1(M)]), \mathbb{Z}[\pi_1(M)]),$$

which converges to $H^*(-; \mathbb{Z}[\pi_1(M)])$. See [Levine 1977, Theorem 2.3] or [Rotman 2009, Theorem 11.34] (note that the pages of this spectral sequence are the transposes of those seen in [Levine 1977], with the differentials modified accordingly). From this spectral sequence, we obtain the following natural short exact sequence, akin to the universal coefficient theorem for cohomology:

$$\begin{aligned} 0 \rightarrow \text{Ext}^1(H_0(-), \mathbb{Z}[\pi_1(M)]) \rightarrow H^1(-; \mathbb{Z}[\pi_1(M)]) \\ \rightarrow \ker(\text{Ext}^0(H_1(-), \mathbb{Z}[\pi_1(M)]) \rightarrow \text{Ext}^2(H_0(-), \mathbb{Z}[\pi_1(M)])) \rightarrow 0. \end{aligned}$$

(For convenience, we suppress the coefficient ring $\mathbb{Z}[\pi_1(M)]$ in some places.) The map $f : M \rightarrow X_n(L)$ induces a map of spectral sequences, and therefore a map of short exact sequences of the type seen above. By condition (2), f induces isomorphisms on $H_i(-; \mathbb{Z}[\pi_1(M)])$, $i = 0, 1$, which implies that $f^* : H^1(X_n(L); \mathbb{Z}[\pi_1(M)]) \rightarrow H^1(M; \mathbb{Z}[\pi_1(M)])$ is an isomorphism by the five lemma.

We now apply the Stallings–Dwyer theorem with $A = \Gamma'$ and $B = \Gamma/\Gamma_n$; see Section 2.2. By assumption, $\Gamma' \rightarrow \Gamma/\Gamma_n$ induces an isomorphism on n -th lower central quotients. Therefore, with $\mathbb{Z}[\pi_1(M)]$ coefficients, the map $H_2(\pi') \rightarrow H_2(\pi/\Gamma_n)/K_j(\pi/\Gamma_n)$, which is the composition of the three maps on the bottom row of diagram (*), is surjective: Note that $H_2(\pi'; \mathbb{Z}[\pi_1(M)])$ is the second homology of the $\pi_1(M)$ -cover of a $K(\pi', 1)$, which is a $K(\Gamma', 1)$, so we have $H_2(\pi'; \mathbb{Z}[\pi_1(M)]) = H_2(\Gamma'; \mathbb{Z})$. Similarly, $H_2(\pi/\Gamma_n; \mathbb{Z}[\pi_1(M)]) = H_2(\Gamma/\Gamma_n; \mathbb{Z})$. Combining with the surjectivity of $H_2(E_{L'}) \twoheadrightarrow H_2(\pi')$ and following the diagram counterclockwise beginning with $H_2(E_{L'})$, we see that $H_2(E_{L'})$ surjects onto the image of $H_2(\pi/\Gamma_n)/K_j(\pi/\Gamma_n)$ in $H_2(X_n(L))/\text{im}(K_j(\pi/\Gamma_n))$. Thus, the composition of maps following the diagram clockwise from $H_2(E_{L'})$ to $H_2(X_n(L))/\text{im}(K_j(\pi/\Gamma_n))$ also surjects onto the image of $H_2(\pi/\Gamma_n)/K_j(\pi/\Gamma_n)$ in $H_2(X_n(L))/\text{im}(K_j(\pi/\Gamma_n))$. This implies the image of $\text{proj} \circ (\cap f_*[M])$ contains the image of $H_2(\pi/\Gamma_n)/K_j(\pi/\Gamma_n)$. Thus, condition (3) holds for all $j < n$.

($\Leftarrow$): To prove the converse, suppose $f \in [M, X_n(L)]_0$ satisfies conditions (1)–(3). We may assume the image of f contains $L \subset X_n(L)$. Define $L' = f^{-1}(L)$, the transverse preimage of L under f . The restriction of f to $E_{L'}$ yields a map $E_{L'} \rightarrow M_{p_n}^\times$, hence a homomorphism $\pi' \rightarrow \pi/\Gamma_n$. Since $f_* = \iota_n(L)_*$ on fundamental groups by condition (2), this is a homomorphism over $\pi_1(M)$. In Claim 1 below, we will show that the induced homomorphism $\pi'/\Gamma'_n \rightarrow \pi/\Gamma_n$ is surjective. Claim 3 will upgrade this to an isomorphism. In between, we will need Claim 2 to arrange that L' has the same number of components as L .

Claim 1 *The restriction of f to $E_{L'}$ induces a surjective homomorphism $\pi'/\Gamma'_n \twoheadrightarrow \pi/\Gamma_n$.*

First, we prove $\Gamma'/\Gamma'_2 \rightarrow \Gamma/\Gamma_2$ is surjective. Consider the diagram of Mayer–Vietoris sequences below corresponding to the decompositions $M = E_{L'} \cup_{\partial E_{L'}} \nu L'$ and $X_n(L) = M_{p_n}^\times \cup_{\partial E_L} \nu L$; recall

that $M_{p_n}^\times$ is a $K(\pi/\Gamma_n, 1)$, and all coefficients are in $\mathbb{Z}[\pi_1(M)]$:

$$\begin{array}{ccc} H_1(\partial E_{L'}) \twoheadrightarrow H_1(E_{L'}) \oplus H_1(\nu L') & \longrightarrow & 0 \\ \downarrow & & \downarrow \\ H_1(\partial E_L) \twoheadrightarrow H_1(\pi/\Gamma_n) \oplus H_1(\nu L) & \longrightarrow & 0 \end{array}$$

The vertical map on the left is surjective by construction. The vertical map on the right splits as a direct sum of homomorphisms. Thus, by the diagram,

$$\Gamma'/\Gamma'_2 = H_1(E_{L'}; \mathbb{Z}[\pi_1(M)]) \rightarrow H_1(\pi/\Gamma_n; \mathbb{Z}[\pi_1(M)]) = \Gamma/\Gamma_2$$

is surjective. By an argument given in the proof of the Stallings–Dwyer theorem (see [Cha and Orr 2024]), the homomorphism $\Gamma'/\Gamma'_n \rightarrow \Gamma/\Gamma_n$ is surjective. Because $\pi'/\Gamma'_n \rightarrow \pi/\Gamma_n$ is a homomorphism over $\pi_1(M)$, the five lemma implies that $\pi'/\Gamma'_n \rightarrow \pi/\Gamma_n$ is surjective.

Claim 2 *We may reduce the number of components of L' without altering the homotopy class of the map f such that the resulting link, which we continue to call L' , has m components and maps homeomorphically to L .*

The proof of Claim 2 is similar to the proof of Claim 1 of (1) $\implies$ (2) from the proof of Theorem B_0 with one notable exception. When we finish reducing the number of components of L' , we want each component of the resulting link to map homeomorphically onto a component of L . Suppose some component L_i of L has a transverse preimage which is not connected, that is, $f^{-1}(L_i)$ is more than one component of the link L' . Similarly to the proof of Theorem B_0 , we attach 4-dimensional 1-handles to $M \times \{1\} \subset M \times [0, 1]$ in which we connect the components of $f^{-1}(L_i)$, along with canceling 4-dimensional 2-handles. We are left with a homotopy $M \times [0, 1] \rightarrow X_n(L)$ which is f on $M \times \{0\}$ and a map g on $M \times \{1\}$ such that the preimage of L_i under g is a single component L'_i . Just as in the proof of Theorem B_0 , we leverage the surjection $\pi'/\Gamma'_n \twoheadrightarrow \pi/\Gamma_n$. By condition (1), $\Delta g_*[M] = \Delta f_*[M] = \sum_i [T_i]$. Thus, by additional homotopy of g (which we may take to be the identity outside of $\nu L'_i$), we may arrange that g maps L'_i homeomorphically to L_i . We do the same for all components of L and obtain a map, which we continue to call g , such that $g \simeq f$ and $g^{-1}(L)$ is an m -component link. This proves Claim 2.

Claim 3 *We have an isomorphism $\pi'/\Gamma'_n \xrightarrow{\cong} \pi/\Gamma_n$.*

Observe that the surgeries performed in the proof of Claim 2 did not alter the surjectivity of the homomorphism $\pi'/\Gamma'_n \twoheadrightarrow \pi/\Gamma_n$.

To begin, we show that

$$\Gamma'/\Gamma'_2 = H_1(E_{L'}; \mathbb{Z}[\pi_1(M)]) \rightarrow H_1(E_L; \mathbb{Z}[\pi_1(M)]) = \Gamma/\Gamma_2$$

is now an isomorphism. For the decompositions

$$M = E_{L'} \cup_{\partial E_{L'}} \nu L', \quad X_n(L) = M_{p_n}^\times \cup_{\partial E_L} \nu L \quad \text{and} \quad M = E_L \cup_{\partial E_L} \nu L$$

consider the corresponding diagram of Mayer–Vietoris sequences

$$\begin{array}{ccccccc}
 H_2(M) & \longrightarrow & H_1(\partial E_{L'}) & \longrightarrow & H_1(E_{L'}) \oplus H_1(\nu L') & \longrightarrow & 0 \\
 \downarrow g_* & & \downarrow \cong & & \downarrow & & \\
 H_2(X_n(L)) & \longrightarrow & H_1(\partial E_L) & \longrightarrow & H_1(\pi/\Gamma_n) \oplus H_1(\nu L) & \longrightarrow & 0 \\
 \uparrow \iota_n(L)_* & & \parallel & & \uparrow \cong & & \\
 H_2(M) & \longrightarrow & H_1(\partial E_L) & \longrightarrow & H_1(E_L) \oplus H_1(\nu L) & \longrightarrow & 0
 \end{array}$$

where all coefficients are in $\mathbb{Z}[\pi_1(M)]$.

The rightmost vertical maps split as direct sums of homomorphisms, and $H_1(\nu L')$ maps isomorphically onto $H_1(\nu L)$. It suffices to show the composition of the second column of vertical maps, which are induced by restrictions of g and $\iota_n(L)$, sends $\text{im}(H_2(M) \rightarrow H_1(\partial E_{L'}))$ isomorphically to $\text{im}(H_2(M) \rightarrow H_1(\partial E_L))$. To show this, we prove the diagram below commutes:

$$\begin{array}{ccc}
 H^1(M) & \xrightarrow[\cong]{\cap[M]} & H_2(M) \xrightarrow{\Delta} H_1(\partial E_{L'}) \\
 g^* \uparrow \cong & & \downarrow \cong \\
 H^1(X_n(L)) & & H_1(\partial E_L) \\
 \iota_n(L)_* \downarrow \cong & & \parallel \\
 H^1(M) & \xrightarrow[\cong]{\cap[M]} & H_2(M) \xrightarrow{\Delta} H_1(\partial E_L)
 \end{array}$$

All coefficients are still in $\mathbb{Z}[\pi_1(M)]$.

The leftmost vertical arrows are isomorphisms by the same universal coefficient spectral sequence argument used in the forward direction of this proof, since for $i = 0, 1$, the maps g_* and $\iota_n(L)_*$ are isomorphisms on $H_i(-; \mathbb{Z}[\pi_1(M)])$. Consider the two commutative squares below, again with coefficients in $\mathbb{Z}[\pi_1(M)]$. We use these squares to see that the diagram above commutes. The commutativity of the squares

$$\begin{array}{ccc}
 H^1(X_n(L)) & \xrightarrow{\text{inc}^*} & H^1(\partial E_L) \\
 \cap(\iota_n(L)_*[M]) \downarrow & & \downarrow \cap(\Delta \iota_n(L)_*[M]) \\
 H_2(X_n(L)) & \xrightarrow{\Delta} & H_1(\partial E_L)
 \end{array}
 \quad \text{and} \quad
 \begin{array}{ccc}
 H^1(X_n(L)) & \xrightarrow{\text{inc}^*} & H^1(\partial E_L) \\
 \cap(g_*[M]) \downarrow & & \downarrow \cap(\Delta g_*[M]) \\
 H_2(X_n(L)) & \xrightarrow{\Delta} & H_1(\partial E_L)
 \end{array}$$

follows from the relationship between the connecting homomorphism in a Mayer–Vietoris sequence and the cap product; see [Spanier 1995].

We check that the large square commutes. Let $\alpha \in H^1(X_n(L); \mathbb{Z}[\pi_1(M)])$. Then

$$\begin{aligned}
 g_* \Delta(g^* \alpha \cap [M]) &= \Delta g_*(g^* \alpha \cap [M]) = \Delta(\alpha \cap g_*[M]) = \text{inc}^*(\alpha) \cap (\Delta g_*[M]) \\
 &= \text{inc}^*(\alpha) \cap (\Delta \iota_n(L)_*[M]) = \Delta(\alpha \cap \iota_n(L)_*[M]) \\
 &= \Delta \iota_n(L)_*(\iota_n(L)^* \alpha \cap [M]) = \iota_n(L)_* \Delta(\iota_n(L)^* \alpha \cap [M]).
 \end{aligned}$$

The third and fifth equalities follow from the commutative squares above. The fourth equality follows from condition (1) or by construction. Thus, the diagram commutes, and we obtain an isomorphism $\Gamma' / \Gamma'_2 \xrightarrow{\cong} \Gamma / \Gamma_2$.

We consider the diagram (*) again, but with g in place of f . Recall all coefficients are in $\mathbb{Z}[\pi_1(M)]$. The diagram commutes for all $j < n$. We require $j < n$ simply to ensure that $K_j(\pi / \Gamma_n)$ is defined. By assumption (3), the image of $\text{proj} \circ (\cap g_*[M]) = \text{proj} \circ (\cap f_*[M])$ contains the image of $H_2(\pi / \Gamma_n) / K_j(\pi / \Gamma_n)$ in $H_2(X_n(L)) / \text{im}(K_j(\pi / \Gamma_n))$. We claim that the composition

$$H_2(\pi') \rightarrow H_2(\pi / \Gamma_n) / K_j(\pi / \Gamma_n)$$

of the three arrows on the bottom row is onto. Claim 3 then follows from the Stallings–Dwyer theorem.

To prove this surjectivity, we consider the diagram of Mayer–Vietoris sequences seen previously. We copy a different part of this diagram below:

$$\begin{array}{ccccccc} H_2(\partial E_{L'}) & \longrightarrow & H_2(E_{L'}) & \longrightarrow & H_2(M) & \longrightarrow & H_1(\partial E_{L'}) \\ \downarrow \cong & & \downarrow & & \downarrow g_* & & \downarrow \cong \\ H_2(\partial E_L) & \longrightarrow & H_2(\pi / \Gamma_n) & \longrightarrow & H_2(X_n(L)) & \longrightarrow & H_1(\partial E_L) \\ \parallel & & \uparrow & & \uparrow \iota_n(L)_* & & \parallel \\ H_2(\partial E_L) & \longrightarrow & H_2(E_L) & \longrightarrow & H_2(M) & \longrightarrow & H_1(\partial E_L) \end{array}$$

All coefficients are in $\mathbb{Z}[\pi_1(M)]$.

A diagram chase using diagram (*) and the diagram above completes the proof of Claim 3. Let $x \in H_2(\pi / \Gamma_n)$. By condition (3), there exists $\alpha \in H^1(X_n(L))$ such that the images of x and α agree in $H_2(X_n(L)) / \text{im}(K_j(\pi / \Gamma_n))$. Let $\bar{x}$ and a denote the respective images of these elements in $H_2(X_n(L))$. Then $\bar{x} - a \in \text{im}(K_j(\pi / \Gamma_n))$, so there exists some $y \in K_j(\pi / \Gamma_n) \leq H_2(\pi / \Gamma_n)$ whose image in $H_2(X_n(L))$ is $\bar{x} - a$. The element $x - y \in H_2(\pi / \Gamma_n)$ maps to $a \in H_2(X_n(L))$, so a is in the image of $H_2(\pi / \Gamma_n)$. In particular, this implies that $\Delta(a) = 0$, where Δ is the connecting homomorphism in the appropriate Mayer–Vietoris sequence. We apply the two isomorphisms at the top of diagram (*) to α to produce an element $b \in H_2(M)$ such that $g_*(b) = a$. By the rightmost vertical isomorphisms in the Mayer–Vietoris diagram, $\Delta(b) = 0$, too. Thus, there exists $c \in H_2(E_{L'})$ which maps to b in $H_2(M)$. Following diagram (*) counterclockwise from $H_2(E_{L'})$ to $H_2(\pi / \Gamma_n)$, consider the image z of the element c . Since $x - y$ and z both map to $a \in H_2(X_n(L))$, the Mayer–Vietoris sequence implies that $z - (x - y)$ is in the image of $H_2(\partial E_L)$. Using the isomorphism at the upper-left of the Mayer–Vietoris diagram, subtract from $c \in H_2(E_{L'})$ the image of this element. This yields $c' \in H_2(E_{L'})$ whose image in $H_2(\pi / \Gamma_n)$ is $z - (z - (x - y)) = x - y$. Since $y \in K_j(\pi / \Gamma_n)$, $x - y$ and x have the same image in $H_2(\pi / \Gamma_n) / K_j(\pi / \Gamma_n)$.

Thus, the map $H_2(E_{L'}) \rightarrow H_2(\pi / \Gamma_n) / K_j(\pi / \Gamma_n)$ is surjective, and therefore so too is the map $H_2(\pi') \rightarrow H_2(\pi / \Gamma_n) / K_j(\pi / \Gamma_n)$. By the Stallings–Dwyer theorem with $A = \Gamma'$ and $B = \Gamma / \Gamma_n$, g induces an isomorphism $\Gamma' / \Gamma'_n \xrightarrow{\cong} \Gamma / \Gamma_n$, hence an isomorphism $\pi' / \Gamma'_n \xrightarrow{\cong} \pi / \Gamma_n$ by the five lemma. This proves Claim 3.

To complete the proof of the theorem, we note that the isomorphism $\phi : \pi' / \Gamma'_n \xrightarrow{\cong} \pi / \Gamma_n$ we obtain from Claim 3 along with the isomorphism $\phi_\partial : H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L)$ constitutes an n -basing for L' relative to L . Thus, $h_n(L', \phi)$ is defined, and

$$h_n(L', \phi) = g = f \in [M, X_n(L)]_0. \quad \square$$

By the above proof, we observe:

Corollary 6.2 *Let $L \subset M$ be an m -component link, and let $f \in [M, X_n(L)]_0$. Suppose that $\pi_2(M) = 0$ or, more generally, that the following composition is zero:*

$$H_2(E_L; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(M; \mathbb{Z}[\pi_1(M)]) \rightarrow H_2(X_n(L); \mathbb{Z}[\pi_1(M)]).$$

If f satisfies conditions (1) and (2) of Theorem C, then f also satisfies condition (3).

This fact is particularly useful for applications of Theorem C to open questions like the Almost-concordance conjecture, as it shows there are many instances where we do not need to check the somewhat technical condition (3).

Proof All homology groups in this proof have coefficients in $\mathbb{Z}[\pi_1(M)]$. Suppose that $\pi_2(M) = 0$. Then $H_2(M) = \pi_2(M) = 0$. Beginning with the isomorphism $\Gamma' / \Gamma'_2 \xrightarrow{\cong} \Gamma / \Gamma_2$ from the proof of Theorem C which did not depend on condition (3), we inductively show $\Gamma' / \Gamma'_k \xrightarrow{\cong} \Gamma / \Gamma_k$ for all $k \leq n$.

Suppose we have an isomorphism for some $2 \leq k < n$. By the final Mayer–Vietoris diagram appearing in the proof of Theorem C, the homomorphisms $H_2(\partial E_{L'}) \rightarrow H_2(E_{L'})$ and $H_2(\partial E_L) \rightarrow H_2(E_L)$ are surjective. This implies that the images of $H_2(\pi')$ and $H_2(\pi)$ in $H_2(\pi' / \Gamma'_k)$ and $H_2(\pi / \Gamma_k)$, respectively, are equal to the images of $H_2(\partial E_{L'})$ and $H_2(\partial E_L)$ and are therefore isomorphic. By the Stallings exact sequence, we obtain an isomorphism $\Gamma'_k / \Gamma'_{k+1} \xrightarrow{\cong} \Gamma_k / \Gamma_{k+1}$, hence an isomorphism $\Gamma' / \Gamma'_{k+1} \xrightarrow{\cong} \Gamma / \Gamma_{k+1}$. This completes the induction.

By the proof of the Stallings–Dwyer theorem, the map $H_2(\pi') \rightarrow H_2(\pi / \Gamma_n) / K_j(\pi / \Gamma_n)$ is surjective for all $j < n$. From diagram (*) in the proof of Theorem C, $H_2(E_{L'})$ surjects onto the image of $H_2(\pi / \Gamma_n) / K_j(\pi / \Gamma_n)$ in $H_2(X_n(L)) / \text{im}(K_j(\pi / \Gamma_n))$. This image is 0 because the homomorphism from $H_2(E_{L'})$ factors through $H_2(M) = 0$. Thus, condition (3) of Theorem C holds, regardless of the image of $\text{proj} \circ (\cap f_*[M])$. (In fact, the image of the map $\text{proj} \circ (\cap f_*[M])$ is also 0 because $H^1(X_n(L)) \cong H^1(M) \cong H_2(M) = 0$.)

The proof in the general case where the composition $H_2(E_L) \rightarrow H_2(M) \rightarrow H_2(X_n(L))$ is zero is nearly identical and leverages the same Mayer–Vietoris diagram. $\square$

7 Characterization and realization for θ_n and $\bar{\mu}_n$

In this section, we prove Theorems D₀, D and E concerning the homology concordance invariants θ_n and $\bar{\mu}_n$. Recall that Theorem D states that the $\mathbb{Z}[G]$ -homology concordance invariants $\bar{\mu}_n$ determine

whether an $(n+1)$ -basing exists, assuming the existence of an n -basing. In other words, the G -Milnor invariants $\bar{\mu}_n$ determine the G -lower central quotients one step at a time. Theorem D will follow quickly from Theorem D₀, which characterizes when the invariants θ_n vanish in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$. When G is the trivial group, Theorem E gives precise conditions under which elements of $H_3(X_n(L))$ are realized as (integral) homology concordance invariants of links in 3-manifolds. In this section, links are not assumed to live in the same 3-manifold. Throughout this section, given a properly embedded surface Σ in some 4-dimensional cobordism W of 3-manifolds M and M' , we will use the notation $\partial\nu\Sigma$ to denote the intersection $\partial E_\Sigma \cap \nu\Sigma$, not the manifold boundary of $\nu\Sigma$ which also includes $(M \sqcup M') \cap \nu\Sigma$. Some of the proofs in this section are similar to arguments seen in [Cha and Orr 2024].

7.1 Proof of Theorem D

We prove the following theorem toward eventually proving Theorem D₀ and its generalization, Theorem D. Recall that Theorem D and its siblings, Theorem D₀ and Corollaries D₁ and D₂, say that an n -basing lifts to an $(n+1)$ -basing if and only if the corresponding $\bar{\mu}$ -invariant (or θ -invariant, depending on context) vanishes in an appropriate sense. Furthermore, these invariants are defined inductively and characterize when the h -invariants satisfy a lifting property.

Theorem 7.1 Fix an m -component link $L \subset M$, and let $n \geq 2$. Suppose $L' \subset M'$ and $L'' \subset M''$ are m -component links equipped with n -basings (ϕ', ϕ'_∂) and $(\phi'', \phi''_\partial)$ over G relative to L , so that $(\psi, \psi_\partial) = ((\phi')^{-1} \circ \phi'', (\phi'_\partial)^{-1} \circ \phi''_\partial)$ is an n -basing over G for L'' relative to L' . If

$$\theta_n^G(L', \phi') = \theta_n^G(L'', \phi'') \in H_3(X_n^G(L)),$$

then the n -basing (ψ, ψ_∂) lifts to an $(n+1)$ -basing over G for L'' relative to L' .

Proof We suppress the group G from our notation. Suppose $\theta_n(L', \phi') = \theta_n(L'', \phi'')$. Recall that $H_3(X_n(L))$ is isomorphic to the oriented bordism group $\Omega_3^{SO}(X_n(L))$. By definition, $\theta_n(L', \phi') = h_n(L', \phi')_*[M']$ and $\theta_n(L'', \phi'') = h_n(L'', \phi'')_*[M'']$. Since $\theta_n(L', \phi') = \theta_n(L'', \phi'')$, there is an oriented cobordism W from M' to M'' over the space $X_n(L)$ as in the diagram below:

$$\begin{array}{ccc} M'' & & \\ \downarrow & \searrow h_n(L'', \phi'') & \\ W & \xrightarrow{f} & X_n(L) \\ \uparrow & \nearrow h_n(L', \phi') & \\ M' & & \end{array}$$

Taking the transverse preimage of $L \subset X_n(L)$ under the map f , we obtain a surface $\Sigma \subset W$ cobounded by the links L' and L'' . This surface may have more than m components. We use an argument similar to the one seen in Claim 1 of (1) $\implies$ (2) in the proof of Theorem B₀ to reduce the number of components of Σ .

Suppose $f^{-1}(L_i)$ has more than one component for some i . Using the same method as the proof of Theorem B₀, we attach 5-dimensional 1-handles to $W \times \{1\} \subset W \times [0, 1]$ in which we connect the

components of $f^{-1}(L_i)$. It is not necessary in this case, as it was in the proof of Theorem B₀, to attach canceling 2-handles. Nevertheless, we can attach canceling 5-dimensional 2-handles and extend the map to $X_n(L)$ over these 2-handles. This produces a homotopy rel boundary $W \times [0, 1] \rightarrow X_n(L)$. The preimage of L under the map $W \times \{1\} \rightarrow X_n(L)$ is an m -component surface, which we continue to call Σ , cobounded by L and L' . We also refer to this new map $W \rightarrow X_n(L)$ as f .

The pair $(E_\Sigma, \partial v\Sigma)$ is a cobordism of pairs between $(E_{L'}, \partial E_{L'})$ and $(E_{L''}, \partial E_{L''})$. The following diagram of pairs is obtained by removing vL from $X_n(L)$ and its corresponding preimages under the maps $h_n(L', \phi')$, $h_n(L'', \phi'')$ and f :

$$\begin{array}{ccc} (E_{L''}, \partial E_{L''}) & & \\ \downarrow & \searrow & \\ (E_\Sigma, \partial v\Sigma) & \longrightarrow & (M_{p_n}^\times, \partial E_L) \\ \uparrow & \nearrow & \\ (E_{L'}, \partial E_{L'}) & & \end{array}$$

We postcompose with the homotopy inverse of the homotopy equivalence $h_{\phi'} : (M_{p_n}^\times, \partial E_{L'}) \xrightarrow{\cong} (M_{p_n}^\times, \partial E_L)$ provided by Proposition 4.2 induced by the n -basing ϕ' to obtain

$$\begin{array}{ccccc} (E_{L''}, \partial E_{L''}) & & & & \\ \downarrow & \searrow & h_\psi \circ p_n'' & & \\ (E_\Sigma, \partial v\Sigma) & \longrightarrow & (M_{p_n}^\times, \partial E_L) & \xrightarrow{\cong} & (M_{p_n}^\times, \partial E_{L'}) \\ \uparrow & \nearrow & p_n' & & \\ (E_{L'}, \partial E_{L'}) & & & & \end{array}$$

The n -basing (ψ, ψ_∂) of L'' relative to L' is now encoded in this diagram via the corresponding homotopy equivalence $h_\psi : (M_{p_n'}^\times, \partial E_{L''}) \xrightarrow{\cong} (M_{p_n}^\times, \partial E_{L'})$. In order to lift ψ to an $(n+1)$ -basing, we will obtain a lift

$$\begin{array}{ccc} & (M_{p_{n+1}'}^\times, \partial E_{L'}) & \\ & \nearrow \text{---} & \downarrow p_{n+1,n}' \\ (E_{L''}, \partial E_{L''}) & \xrightarrow{h_\psi \circ p_n''} & (M_{p_n}^\times, \partial E_{L'}) \end{array}$$

The lift of $\partial E_{L''}$ is already determined by the isomorphism ψ_∂ . We thus have a relative lifting problem

$$\begin{array}{ccc} \partial E_{L''} & \xrightarrow{\psi_\partial} & M_{p_{n+1}'}^\times \\ \downarrow & \nearrow \text{---} \text{?} & \downarrow p_{n+1,n}' \\ E_{L''} & \xrightarrow{h_\psi \circ p_n''} & M_{p_n}^\times \end{array}$$

Since $p'_{n+1,n}$ is not necessarily a fibration, we replace $M_{p'_{n+1}}^\times$ with the mapping path space P of $p'_{n+1,n}$. It suffices to find a relative lift of $h_\psi \circ p''_n$ to P . We may lift the 1-skeleton with no obstruction. Since P is a $K(\pi'/\Gamma'_{n+1}, 1)$, if we can lift the 2-skeleton of $E_{L''}$, then we can also lift the 3-skeleton (i.e., all of $E_{L''}$). The map $P \rightarrow M_{p'_n}^\times$ is a fibration from a $K(\pi'/\Gamma'_{n+1}, 1)$ to a $K(\pi'/\Gamma'_n, 1)$. Its fiber is therefore a $K(\Gamma'_n/\Gamma'_{n+1}, 1)$. This fibration defines a system of local coefficients given by $A = \Gamma'_n/\Gamma'_{n+1}$ as a module over $\mathbb{Z}[\pi'/\Gamma'_n]$. This system corresponds to a homomorphism $\pi'/\Gamma'_n \rightarrow \text{Aut}(A)$. Since A is central in Γ'/Γ'_{n+1} , the restriction of this homomorphism to Γ'/Γ'_n is trivial, inducing a homomorphism $G \rightarrow \text{Aut}(A)$ and giving A the structure of a $\mathbb{Z}[G]$ -module:

$$\begin{array}{ccc}
 \Gamma'/\Gamma'_n & & \\
 \downarrow & \searrow 0 & \\
 \pi'/\Gamma'_n & \longrightarrow & \text{Aut}(A) \\
 \downarrow & \nearrow \exists & \\
 G & &
 \end{array}$$

The obstruction to lifting the 2-skeleton of $E_{L''}$ is an element $o_{L''} \in H^2(E_{L''}, \partial E_{L''}; A)$, where we understand coefficients in A to mean the system of local coefficients from the fibration pulled back over the appropriate spaces. This obstruction vanishes if and only if we have a relative lift of $E_{L''}$ to P and, in turn, to $M_{p'_{n+1}}^\times$. Analogously, we have obstruction classes $o_{L'}$ and o_Σ corresponding to lifting $E_{L'}$ and E_Σ . The naturality of obstruction classes gives us

$$H^2(E_{L'}, \partial E_{L'}; A) \leftarrow H^2(E_\Sigma, \partial v\Sigma; A) \rightarrow H^2(E_{L''}, \partial E_{L''}; A), \quad o_{L'} \leftarrow o_\Sigma \mapsto o_{L''}.$$

The relative lift of $E_{L'}$ exists, as we have the inclusion map $E_{L'} \hookrightarrow M_{p'_{n+1}}^\times$. Thus, $o_{L'} = 0$. It therefore suffices to prove that

$$\ker(H^2(E_\Sigma, \partial v\Sigma; A) \rightarrow H^2(E_{L'}, \partial E_{L'}; A)) = \ker(H^2(E_\Sigma, \partial v\Sigma; A) \rightarrow H^2(E_{L''}, \partial E_{L''}; A)).$$

We prove this equality through a series of reductions.

First, consider the following diagram, where the horizontal arrows are Poincaré duality and each lowest vertical arrow is projection onto the first factor:

$$\begin{array}{ccc}
 H^2(E_\Sigma, \partial v\Sigma; A) & \xrightarrow{\cong} & H_2(E_\Sigma, E_{L'} \sqcup E_{L''}; A) \\
 \text{inc}^* \downarrow & & \downarrow \Delta \\
 H^2(E_{L'} \sqcup E_{L''}, \partial E_{L'} \sqcup \partial E_{L''}; A) & \xrightarrow{\cong} & H_1(E_{L'} \sqcup E_{L''}; A) \\
 \parallel & & \parallel \\
 H^2(E_{L'}, \partial E_{L'}; A) \oplus H^2(E_{L''}, \partial E_{L''}; A) & \xrightarrow{\cong} & H_1(E_{L'}; A) \oplus H_1(E_{L''}; A) \\
 \downarrow & & \downarrow p \\
 H^2(E_{L'}, \partial E_{L'}; A) & \xrightarrow{\cong} & H_1(E_{L'}; A)
 \end{array}$$

As in the proof of Claim 3 of Theorem C, the top square commutes by the relationship between cap product and the connecting homomorphism in a Mayer–Vietoris sequence; see [Spanier 1995]. By the diagram, $\ker(H^2(E_\Sigma, \partial v \Sigma; A) \rightarrow H^2(E_{L'}, \partial E_{L'}; A)) = \text{PD}^{-1}(\ker(p \circ \Delta))$, where PD is the Poincaré duality map. If $p|_{\text{im} \Delta}$ injects, then this equals $\text{PD}^{-1}(\ker \Delta)$. On the other hand,

$$\ker(H^2(E_\Sigma, \partial v \Sigma; A) \rightarrow H^2(E_{L''}, \partial E_{L''}; A)) = \text{PD}^{-1}(\ker(q \circ \Delta)) \supseteq \text{PD}^{-1}(\ker \Delta),$$

where $q : H_1(E_{L'}; A) \oplus H_1(E_{L''}; A) \twoheadrightarrow H_1(E_{L''}; A)$ is projection onto the second factor. Thus, the equality we wish to establish follows if both $p|_{\text{im} \Delta}$ and $q|_{\text{im} \Delta}$ inject. But

$$\begin{aligned} \text{im} \Delta &= \ker(H_1(E_{L'} \sqcup E_{L''}; A) \rightarrow H_1(E_\Sigma; A)) \\ &= \{(x, y) \in H_1(E_{L'}; A) \oplus H_1(E_{L''}; A) \mid i'_*(x) + i''_*(y) = 0\}, \end{aligned}$$

where $i'_* : H_1(E_{L'}; A) \rightarrow H_1(E_\Sigma; A)$ and $i''_* : H_1(E_{L''}; A) \rightarrow H_1(E_\Sigma; A)$ are induced by inclusion. Suppose $(x, y) \in \text{im} \Delta$ with $p(x, y) = x = 0$. Then $i''_*(y) = 0$. Similarly, if $q(x, y) = y = 0$, then $i'_*(x) = 0$. It therefore suffices to prove that i'_* and i''_* inject.

Similar to the proof of Theorem C, we invoke the (homology) universal coefficient spectral sequence, a homology spectral sequence with $E_{p,q}^2 = \text{Tor}_p^{\mathbb{Z}[G]}(H_q(-; \mathbb{Z}[G]), A)$ which converges to $H_*(-; A)$; see [Levine 1977, Theorem 2.3] or [Rotman 2009, Theorem 11.34]. From this spectral sequence, we obtain the following natural short exact sequence akin to the universal coefficient theorem for homology:

$$0 \rightarrow \frac{H_1(-; \mathbb{Z}[G]) \otimes_{\mathbb{Z}[G]} A}{\text{im}(\text{Tor}_2(H_0(-; \mathbb{Z}[G]), A))} \rightarrow H_1(-; A) \rightarrow \text{Tor}_1(H_0(-; \mathbb{Z}[G]), A) \rightarrow 0.$$

The composition $E_{L'} \rightarrow M_{p_n}^\times \xrightarrow{\cong} M_{p_n}^\times$ induces maps of spectral sequences, and therefore a diagram of short exact sequences of the type seen above. The isomorphisms

$$\begin{aligned} H_0(E_{L'}; \mathbb{Z}[G]) &\xrightarrow{\cong} H_0(\pi' / \Gamma'_n; \mathbb{Z}[G]) \xrightarrow{\cong} H_0(\pi / \Gamma_n; \mathbb{Z}[G]), \\ H_1(E_{L'}; \mathbb{Z}[G]) &\xrightarrow{\cong} H_1(\pi' / \Gamma'_n; \mathbb{Z}[G]) \xrightarrow{\cong} H_1(\pi / \Gamma_n; \mathbb{Z}[G]), \end{aligned}$$

imply that $H_1(E_{L'}; A) \xrightarrow{\cong} H_1(\pi / \Gamma_n; A)$ is an isomorphism. Note that we use $n \geq 2$ to obtain the isomorphism $H_1(E_{L'}; \mathbb{Z}[G]) \xrightarrow{\cong} H_1(\pi / \Gamma_n; \mathbb{Z}[G])$. The isomorphism $H_1(E_{L'}; A) \xrightarrow{\cong} H_1(\pi / \Gamma_n; A)$ factors as

$$\begin{array}{ccc} H_1(E_{L'}; A) & \xrightarrow{\cong} & H_1(\pi / \Gamma_n; A) \\ \downarrow i'_* & \nearrow & \\ H_1(E_\Sigma; A) & & \end{array}$$

Thus, i'_* injects. To show i''_* injects, we repeat the argument and apply the spectral sequence to $E_{L''}$.

The injectivity of i'_* and i''_* implies

$$\ker(H^2(E_\Sigma, \partial v \Sigma; A) \rightarrow H^2(E_{L'}, \partial E_{L'}; A)) = \ker(H^2(E_\Sigma, \partial v \Sigma; A) \rightarrow H^2(E_{L''}, \partial E_{L''}; A)).$$

Since the obstruction class $o_{L'}$ vanishes, the obstruction class $o_{L''}$ vanishes. We obtain a lift of $h_\psi \circ p_n'' : E_{L''} \rightarrow M_{p_n}^\times$ to $M_{p_{n+1}}^\times$ rel ψ_∂ , that is, a map of pairs $(E_{L''}, \partial E_{L''}) \rightarrow (M_{p_{n+1}}^\times, \partial E_{L''})$.

This map factors as

$$\begin{array}{ccc} (E_{L''}, \partial E_{L''}) & \longrightarrow & (M_{p'_{n+1}}^\times, \partial E_{L'}) \\ \downarrow & \nearrow \tilde{h} & \\ (M_{p'_{n+1}}^\times, \partial E_{L''}) & & \end{array}$$

since the obstruction to such a factorization vanishes: The homomorphism $\pi'' \twoheadrightarrow \pi' / \Gamma'_{n+1}$ factors through π'' / Γ''_{n+1} and yields a homomorphism $\tilde{\psi} : \pi'' / \Gamma''_{n+1} \rightarrow \pi' / \Gamma'_{n+1}$ lifting ψ . We wish to show that $(\tilde{\psi}, \psi_\partial)$ is an $(n+1)$ -basis for L'' relative to L' which is a lift of (ψ, ψ_∂) . By Proposition 4.2, it suffices to show that $\tilde{\psi} : \pi'' / \Gamma''_{n+1} \rightarrow \pi' / \Gamma'_{n+1}$ is an isomorphism, so that $\tilde{h}$ is a homotopy equivalence of pairs.

To do so, we reverse the roles of L' and L'' and instead consider the diagram

$$\begin{array}{ccccc} (E_{L''}, \partial E_{L''}) & & & & \\ \downarrow & \searrow p''_n & & & \\ (E_\Sigma, \partial v \Sigma) & \longrightarrow & (M_{p_n}^\times, \partial E_L) & \xrightarrow{\cong} & (M_{p''_n}, \partial E_{L''}) \\ \uparrow & \nearrow h_{\psi^{-1} \circ p'_n} & & & \\ (E_{L'}, \partial E_{L'}) & & & & \end{array}$$

We solve the relative lifting problem

$$\begin{array}{ccc} \partial E_{L'} & \xrightarrow{\psi_\partial^{-1}} & M_{p'_{n+1}}^\times \\ \downarrow & \nearrow ?\exists & \downarrow p''_{n+1,n} \\ E_{L'} & \xrightarrow{h_{\psi^{-1} \circ p'_n}} & M_{p'_n}^\times \end{array}$$

and obtain a homomorphism $\tilde{\psi}' : \pi' / \Gamma'_{n+1} \rightarrow \pi'' / \Gamma''_{n+1}$ lifting ψ^{-1} .

We now show that $\tilde{\psi}$ and $\tilde{\psi}'$ are mutually inverse. Consider the following diagram of short exact sequences, noting that $\tilde{\psi} \circ \tilde{\psi}'$ is a lift of $\psi \circ \psi^{-1} = \text{id}_{\pi' / \Gamma'_n}$:

$$\begin{array}{ccccccc} 1 & \longrightarrow & \Gamma'_n / \Gamma'_{n+1} & \hookrightarrow & \pi' / \Gamma'_{n+1} & \twoheadrightarrow & \pi' / \Gamma'_n \longrightarrow 1 \\ & & \downarrow & & \downarrow \tilde{\psi} \circ \tilde{\psi}' & & \parallel \\ 1 & \longrightarrow & \Gamma'_n / \Gamma'_{n+1} & \hookrightarrow & \pi' / \Gamma'_{n+1} & \twoheadrightarrow & \pi' / \Gamma'_n \longrightarrow 1 \end{array}$$

To show $\tilde{\psi} \circ \tilde{\psi}'$ is the identity, it suffices to show $\tilde{\psi} \circ \tilde{\psi}'$ is the identity on $\Gamma'_n / \Gamma'_{n+1}$. Since $\tilde{\psi} \circ \tilde{\psi}'$ is a lift of $\text{id}_{\pi' / \Gamma'_n}$, it is a homomorphism over G and therefore sends Γ' / Γ'_{n+1} to itself. We show that $\tilde{\psi} \circ \tilde{\psi}'$ is, in fact, the identity on $\Gamma'_2 / \Gamma'_{n+1}$.

Let $g \in \Gamma'_2 / \Gamma'_{n+1}$. Then g may be written as $g = \prod_i [a_i, b_i]$, where $a_i, b_i \in \Gamma' / \Gamma'_{n+1}$. Since $\tilde{\psi} \circ \tilde{\psi}'$ is a lift of $\text{id}_{\pi' / \Gamma'_n}$, we have $\tilde{\psi} \circ \tilde{\psi}'(a_i) = a_i u_i$ and $\tilde{\psi} \circ \tilde{\psi}'(b_i) = b_i v_i$ for some $u_i, v_i \in \Gamma'_n / \Gamma'_{n+1}$.

But $\Gamma'_n / \Gamma'_{n+1}$ is central in Γ' / Γ'_{n+1} , so $[a_i u_i, b_i v_i] = [a_i, b_i]$ for all i . Thus,

$$(\tilde{\psi} \circ \tilde{\psi}')(g) = \prod_i [(\tilde{\psi} \circ \tilde{\psi}')(a_i), (\tilde{\psi} \circ \tilde{\psi}')(b_i)] = \prod_i [a_i u_i, b_i v_i] = \prod_i [a_i, b_i] = g.$$

It follows that $\tilde{\psi} \circ \tilde{\psi}' = \text{id}_{\pi'/\Gamma'_{n+1}}$.

To complete the proof, reverse the roles of $\tilde{\psi}$ and $\tilde{\psi}'$ to see that $\tilde{\psi}' = \tilde{\psi}^{-1}$. Thus, $(\tilde{\psi}, \psi_\partial)$ constitutes an $(n+1)$ -basing over G for L'' relative to L' which is a lift of the n -basing (ψ, ψ_∂) . $\square$

We now apply Theorem 7.1 to characterize the notion of “vanishing in the cokernel”. Recall that Theorem D₀ asserts that the invariant $\theta_n(L')$ relative to L vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial)$ if and only if L' admits some $(n+1)$ -basing relative to L .

Proof of Theorem D₀ We again suppress the fixed group G . Fix $L \subset M$, and suppose $L' \subset M'$ admits an n -basing (ϕ, ϕ_∂) relative to L . The proof of this theorem follows quickly from Theorem 7.1.

Suppose L' admits an $(n+1)$ -basing $\tilde{\phi}$ over G relative to L which is a lift of ϕ . The invariant $\theta_{n+1}(L', \tilde{\phi})$ is sent to $\theta_n(L', \phi)$ under the map $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$. Thus, by definition, $\theta_n(L', \phi)$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$.

Conversely, suppose $\theta_n(L', \phi)$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$. Then there exists a 3-manifold (M'', φ'') with $\varphi'' : \pi_1(M'') \twoheadrightarrow G$ and a link $L'' \subset M''$ equipped with an $(n+1)$ -basing $\tilde{\phi}'$ over G relative to L such that $\theta_{n+1}(L'', \tilde{\phi}')$ is sent to $\theta_n(L', \phi)$ under the map $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$. Let ϕ' be the n -basing for L'' induced by the $(n+1)$ -basing $\tilde{\phi}'$. Then $\theta_n(L'', \phi') = \theta_n(L', \phi)$. By Theorem 7.1, the n -basing $\psi = (\phi')^{-1} \circ \phi$ of L' relative to L'' lifts to an $(n+1)$ -basing $\tilde{\psi}$. We obtain an $(n+1)$ -basing $\tilde{\phi}$ for L' relative to L which is a lift of ϕ by composing with the $(n+1)$ -basing $\tilde{\psi}$. Define $\tilde{\phi}$ to be the $(n+1)$ -basing $\tilde{\phi}' \circ \tilde{\psi}$. This proves the equivalence of (1) and (2).

We now prove the equivalence of the analogous basing-independent statements (1') and (2'). Suppose L' admits an $(n+1)$ -basing $\tilde{\phi}'$ over G relative to L . Let ϕ' be the n -basing induced by $\tilde{\phi}'$. Note that $\psi := \phi \circ (\phi')^{-1}$ is a self- n -basing of L . Then $\theta_{n+1}(L', \tilde{\phi}')$ maps to $\theta_n(L', \phi')$ under $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$, and $\theta_n(L', \phi) = \psi \cdot \theta_n(L', \phi')$, that is, $\theta_n(L')$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial)$.

Conversely, suppose $\theta_n(L')$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)/\text{Aut}(\pi/\Gamma_n, \partial)$. Then there exists a self- n -basing ψ for L such that $\psi \cdot \theta_n(L', \phi) = \theta_n(L', \psi \circ \phi)$ is in the image of $\mathcal{R}_{n+1}(L)$. By statement (2) of this theorem, L' admits an $(n+1)$ -basing $\tilde{\phi}'$ relative to L which is a lift of $\psi \circ \phi$. $\square$

Before proceeding to the proof of Theorem D, we recall the definition of the equivalence relation $\sim$ defined on $\mathcal{R}_n(L)$. For $\theta \in \mathcal{R}_n(L)$, there exists a link L' in some 3-manifold M' with $\pi_1(M') \twoheadrightarrow G$ which admits an n -basing ϕ over G relative to L such that $\theta_n(L', \phi) = \theta \in H_3(X_n(L))$. The set I_θ is the image of the composition $\mathcal{R}_{n+1}(L') \rightarrow \mathcal{R}_n(L') \xrightarrow[\phi_*]{\cong} \mathcal{R}_n(L)$. We say $\theta \sim \theta'$ if $\theta' \in I_\theta$.

Lemma 7.2 *The set I_θ is well-defined and $I_\theta = I_{\theta'}$ if and only if $\theta' \in I_\theta$, that is, the sets I_θ form a partition of $\mathcal{R}_n(L)$ and $\sim$ is a well-defined equivalence relation.*

Proof Let $\theta \in \mathcal{R}_n(L)$ with $\theta = \theta_n(L', \phi)$ for some link L' in some 3-manifold M' which admits the n -basing ϕ relative to L . Suppose $\theta' \in I_\theta$. In particular this means $\theta' \in \mathcal{R}_n(L)$, so that $\theta' = \theta_n(L'', \phi')$

for some link $L'' \subset M''$. Then $\phi^{-1} \circ \phi'$ is an n -basing for L'' relative to L' . Since $\theta' \in I_\theta$, we have that $\theta_n(L'', \phi^{-1} \circ \phi') \in \text{im}(\mathcal{R}_{n+1}(L') \rightarrow \mathcal{R}_n(L'))$. By Theorem D₀, the n -basing $\phi^{-1} \circ \phi'$ lifts to an $(n+1)$ -basing ψ for L'' relative to L' .

Consider the following diagram, where the horizontal arrows are bijections of sets:

$$\begin{array}{ccc} \mathcal{R}_{n+1}(L'') & \xrightarrow[\psi_*]{\cong} & \mathcal{R}_{n+1}(L') \\ \downarrow & & \downarrow \\ \mathcal{R}_n(L'') & \xrightarrow[\phi'_*]{\cong} \mathcal{R}_n(L) \xleftarrow[\phi_*]{\cong} & \mathcal{R}_n(L') \end{array}$$

To prove $I_{\theta'} = I_\theta$, we need to see that this diagram commutes. Suppose $\theta'' \in \mathcal{R}_{n+1}(L'')$ with $\theta'' = \theta_{n+1}(L''', \tilde{\phi}'')$, where $\tilde{\phi}''$ is an $(n+1)$ -basing for L''' relative to L'' inducing the n -basing ϕ'' . Considering the images of $\theta_{n+1}(L''', \tilde{\phi}'')$ under the maps in the diagram above, we have

$$\begin{array}{ccc} \theta_{n+1}(L''', \tilde{\phi}'') & \xrightarrow{\psi_*} & \theta_{n+1}(L''', \psi \circ \tilde{\phi}'') \\ \downarrow & & \downarrow \\ \theta_n(L''', \phi'') & \xrightarrow{\phi'_*} \theta_n(L''', \phi' \circ \phi'') \xleftarrow{\phi_*} & \theta_n(L''', \phi^{-1} \circ \phi' \circ \phi'') \end{array}$$

Thus, $I_{\theta'} = I_\theta$. Taking $\theta' = \theta$ shows that the set I_θ is well-defined, independent of our choice of (L', ϕ) . $\square$

Lemma 7.2 allows us to prove a generalization of Theorem D₀ involving the equivalence relation $\sim$. Instead of comparing one link L' to a fixed link L , we can again compare links L' and L'' over the fixed link L .

Corollary 7.3 Fix an m -component link $L \subset M$, and let $n \geq 2$. Suppose $L' \subset M'$ and $L'' \subset M''$ are m -component links equipped with n -basings (ϕ', ϕ'_∂) and $(\phi'', \phi''_\partial)$ over G relative to L , so that $(\psi, \psi_\partial) = ((\phi')^{-1} \circ \phi'', (\phi'_\partial)^{-1} \circ \phi''_\partial)$ is an n -basing over G for L'' relative to L' . Then the n -basing (ψ, ψ_∂) lifts to an $(n+1)$ -basing for L'' over G relative to L' if and only if $\theta_n(L', \phi') \sim \theta_n(L'', \phi'')$ in $\mathcal{R}_n(L)$.

This is a generalization of Theorem D₀ because $\theta \sim \theta_n(L, \text{id})$ if and only if θ vanishes in the cokernel of $\mathcal{R}_{n+1}(L) \rightarrow \mathcal{R}_n(L)$.

Proof By definition, $\theta_n(L'', \phi'') \in I_{\theta_n(L', \phi')}$ if and only if $\theta_n(L'', \psi) = \theta_n(L'', (\phi')^{-1} \circ \phi'')$ vanishes in the cokernel of $\mathcal{R}_{n+1}(L') \rightarrow \mathcal{R}_n(L')$. By Theorem D₀, this happens if and only if ψ lifts to an $(n+1)$ -basing of L'' relative to L' . $\square$

Recall that the Milnor invariant of length n of the pair (L', ϕ) , where ϕ is an n -basing over G relative to L , is the equivalence class $\bar{\mu}_n(L', \phi) = [\theta_n(L', \phi)]$ of the θ -invariant under the equivalence relation $\sim$. The Milnor invariant of length n of the link L' is the equivalence class $\bar{\mu}_n(L') = [\theta_n(L', \phi)]$ under the equivalence relation $\approx$, where we further quotient by the action of $\text{Aut}(\pi / \Gamma_n, \partial)$.

We now prove Theorem D, which generalizes Theorem D₀. We leverage this theorem in Section 8 to show that the vanishing of the lower central homotopy invariants, taken in appropriate contexts, implies the vanishing of all previous versions of Milnor's invariants in dimension 3.

Proof of Theorem D The equivalence of statements (1)–(3) and the equivalence of statements (1')–(3') are the conclusions of Theorem D₀. The equivalence of statements (4) and (5), which compare links $L' \subset M'$ and $L'' \subset M''$ relative to the fixed link $L \subset M$, is a rewriting of the conclusion of Corollary 7.3. Finally, we prove that the analogous basing-independent statements (4') and (5') are equivalent.

Let $L \subset M$, $L' \subset M'$, $L'' \subset M''$, ϕ' and ϕ'' be as in Corollary 7.3. Suppose there exists an $(n+1)$ -basing $\tilde{\psi}$ of L'' relative to L' . Let ψ be the induced n -basing, so that $\psi = (\phi')^{-1} \circ (\phi' \circ \psi)$ lifts. By Corollary 7.3, $\theta_n(L', \phi') \sim \theta_n(L'', \phi' \circ \psi)$. Equivalently, $\theta_n(L', \phi') \sim ((\phi' \circ \psi) \circ (\phi'')^{-1}) \cdot \theta_n(L'', \phi'')$, that is, $\bar{\mu}_n(L') = \bar{\mu}_n(L'')$. Conversely, suppose $\bar{\mu}_n(L') = \bar{\mu}_n(L'')$. Then there exists a self- n -basing ψ of L such that $\psi \cdot \theta_n(L'', \phi'') \sim \theta_n(L', \phi')$, that is, $\theta_n(L'', \psi \circ \phi'') \sim \theta_n(L', \phi')$. We obtain the desired $(n+1)$ -basing for L'' relative to L' from Corollary 7.3. $\square$

7.2 Proof of Theorem E

For the remainder of this section, let G be the trivial group. As G will indicate the trivial group throughout, we will write $\mathcal{R}_n(L)$, $X_n(L)$, $h_n(L', \phi)$ and $\theta_n(L', \phi)$ when we mean $\mathcal{R}_n^G(L)$, $X_n^G(L)$, $h_n^G(L', \phi)$ and $\theta_n^G(L', \phi)$.

Recall that Theorem E characterizes which elements of $H_3(X_n(L))$ can be realized as $\theta_n(L', \phi)$ for some link L' in some 3-manifold M' admitting an n -basing over G relative to L (i.e., in the case of integral homology concordance). See Section 3.7 for the full statement. Theorem E relies heavily on a 3-dimensional homology surgery theorem of Turaev [1984]. In the case where $G \neq \{1\}$, a realization theorem for θ -invariants should follow from similar arguments, but one would first need to prove a $\mathbb{Z}[G]$ -coefficient version of Turaev's homology surgery result. This is a topic of future consideration. We copy Turaev's result below as Lemma 7.4.

Lemma 7.4 [Turaev 1984, Lemma 2.2] *Suppose $g : N \rightarrow X$ is a map of a closed oriented 3-manifold N to a CW-complex X with finitely generated $\pi_1(X)$ such that the cap product*

$$\cap g_*[N] : tH^2(X) \rightarrow tH_1(X)$$

is an isomorphism. Then (N, g) is oriented bordant over X to a pair (M, f) of a closed oriented 3-manifold M and a map $f : M \rightarrow X$ which induces an isomorphism $f_* : H_1(M) \xrightarrow{\cong} H_1(X)$.

Before proving Theorem E, we briefly remark that conditions (2) and (3) in the statement of the theorem can be viewed as Poincaré duality restrictions imposed on the homology class $\theta \in H_3(X_n(L))$. The utility of condition (2) in the following proof is to apply Turaev's theorem, Lemma 7.4. A homology class $\theta \in H_3(X_n(L))$ yields a bordism class of 3-manifolds over $X_n(L)$. Turaev's theorem ensures this class has a representative with the desired first homology. See Section 3.7 for a comparison of the statement of Theorem E with the statement of Theorem C, the realization theorem for h -invariants.

Proof of Theorem E ($\Rightarrow$): Let $\theta \in \mathcal{R}_n(L)$. Then there exists a 3-manifold M' and an m -component link $L' \subset M'$, where L' admits an n -basing (ϕ, ϕ_∂) relative to L such that $\theta_n(L', \phi) = \theta$. In other words, $\theta = h_n(L', \phi)_*[M']$. For simplicity of notation, let f denote the map $h_n(L', \phi)$. First, since f is a map of pairs $(M', \nu L') \rightarrow (X_n(L), \nu L)$ restricting to an admissible homeomorphism $\partial E_{L'} \xrightarrow{\cong} \partial E_L$, $\Delta(\theta) = \Delta f_*[M'] = \sum_i [T_i]$, where Δ is the connecting homomorphism in the Mayer–Vietoris sequence corresponding to the decomposition $X_n(L) = M_{p_n}^\times \cup_{\partial E_L} \nu L$, and where $\partial E_L = \bigsqcup_i T_i$. Thus, condition (1) holds.

Next, consider the commutative square

$$\begin{array}{ccc} tH^2(M') & \xleftarrow{f^*} & tH^2(X_n(L)) \\ \cap[M'] \downarrow & & \downarrow \cap \theta \\ tH_1(M') & \xrightarrow{f_*} & tH_1(X_n(L)) \end{array}$$

The cap product $\cap[M']$ is an isomorphism by Poincaré duality, and

$$\pi_1(X_n(L)) \cong \pi_1(M)/\pi_1(M)_n \cong \pi_1(M')/\pi_1(M')_n$$

(see Proposition A.2). Thus, $f_* : H_1(M') \rightarrow H_1(X_n(L))$ and its restriction $tH_1(M') \rightarrow tH_1(X_n(L))$ are isomorphisms. But $tH^2(-) \cong \text{Ext}(H_1(-), \mathbb{Z})$, so $f^* : tH^2(X_n(L)) \rightarrow tH^2(M')$ is also an isomorphism. By the diagram, $\cap \theta$ is an isomorphism, too. This proves condition (2) holds.

To prove condition (3) holds, we consider diagram (**) below, where all coefficients are in $\mathbb{Z}$:

$$(**) \quad \begin{array}{ccccccc} & & H^1(M') & \xleftarrow[\cong]{f^*} & H^1(X_n(L)) & & \\ & & \cap[M'] \downarrow \cong & & \downarrow \cap \theta & \searrow \text{proj} \circ (\cap \theta) & \\ H_2(E_{L'}) & \xrightarrow{j'_*} & H_2(M') & \xrightarrow{f_*} & H_2(X_n(L)) & \xrightarrow{\text{proj}} \twoheadrightarrow & H_2(X_n(L))/\text{im}(K_j(\pi/\pi_n)) \\ \downarrow \Downarrow & & & & \uparrow & & \uparrow \\ H_2(\pi') & \longrightarrow & H_2(\pi'/\pi'_n) & \xrightarrow[\cong]{\phi_*} & H_2(\pi/\pi_n) & \xrightarrow{\text{proj}} \twoheadrightarrow & H_2(\pi/\pi_n)/K_j(\pi/\pi_n) \end{array}$$

Diagram (**) is similar to diagram (*) seen in the proof of Theorem C.

The map $f^* : H^1(X_n(L)) \rightarrow H^1(M')$ is an isomorphism, since $H^1(-) \cong \text{Hom}(H_1(-), \mathbb{Z})$, and $f_* : H_1(M') \xrightarrow{\cong} H_1(X_n(L))$ is an isomorphism. By the Stallings–Dwyer theorem, the map $H_2(\pi') \rightarrow H_2(\pi/\pi_n)/K_j(\pi/\pi_n)$ is surjective since $\phi : \pi'/\pi'_n \xrightarrow{\cong} \pi/\pi_n$ is an isomorphism. Just as in Theorem C, this implies the image of $\text{proj} \circ (\cap \theta)$ contains the image of $H_2(\pi/\pi_n)/K_j(\pi/\pi_n)$ for all $j < n$. This proves condition (3) holds.

($\Leftarrow$): The converse combines Turaev's result, Lemma 7.4, with ingredients from the proof of Theorem C. When possible, instead of repeating arguments from the proof of Theorem C, we will reference that proof and sketch analogues of the arguments in this setting.

Suppose conditions (1)–(3) hold for some class $\theta \in H_3(X_n(L))$. Since $H_3(-) \cong \Omega_3^{\text{SO}}(-)$, there exists a 3-manifold M'' and a map $g : M'' \rightarrow X_n(L)$ such that $g_*[M''] = \theta$. We first invoke the homology surgery result of Turaev to alter (M'', g) . We then use the transversality techniques and arguments seen in the proof of Theorem C to produce the desired link realizing the class θ .

By Turaev's result and condition (2), (M'', g) is oriented bordant over $X_n(L)$ to some (M', f) , where M' is a closed orientable 3-manifold and $f : M' \rightarrow X_n(L)$ induces an isomorphism on $H_1(-)$. Since (M', f) is oriented bordant to (M'', g) , we have $f_*[M'] = g_*[M''] = \theta$. As in the proof of Theorem C, we may assume that the image of the map f contains L . Define $L' \subset M'$ to be the transverse preimage $f^{-1}(L)$ of L .

Claim 1 *The restriction of f to $E_{L'}$ induces a surjective homomorphism $\pi'/\pi'_n \rightarrow \pi/\pi_n$.*

We consider a Mayer–Vietoris diagram similar to one seen in the proof of Theorem C:

$$\begin{array}{ccccccc} H_1(\partial E_{L'}) & \longrightarrow & H_1(E_{L'}) \oplus H_1(vL') & \longrightarrow & H_1(M') & \longrightarrow & 0 \\ \downarrow & & \downarrow & & \downarrow \cong & & \\ H_1(\partial E_L) & \longrightarrow & H_1(\pi/\pi_n) \oplus H_1(vL) & \longrightarrow & H_1(X_n(L)) & \longrightarrow & 0 \end{array}$$

The middle vertical arrow splits as a direct sum of homomorphisms. A diagram chase proves that $\pi'/\pi'_2 = H_1(E_{L'}) \rightarrow H_1(\pi/\pi_n) = \pi/\pi_2$ is surjective. As in the proof of Theorem C, Claim 1 now follows from a technique used in the proof of the Stallings–Dwyer theorem; see [Cha and Orr 2024].

Claim 2 *We may reduce the number of components of L' without altering the homotopy class of the map f so that the resulting link, which we continue to call L' , has m components and maps homeomorphically to L .*

The proof of Claim 2 is entirely similar to that of Claim 2 in the proof of Theorem C and requires condition (1); see also Claim 1 of (1) $\Rightarrow$ (2) from the proof of Theorem B₀.

Claim 3 *The map f induces an isomorphism $\pi'/\pi'_n \xrightarrow{\cong} \pi/\pi_n$.*

The proof of Claim 3 proceeds in the same way as the proof of Claim 3 of Theorem C. We briefly summarize the argument.

We first show $\pi'/\pi'_2 \rightarrow \pi/\pi_2$ is an isomorphism. Consider the diagram of Mayer–Vietoris sequences

$$\begin{array}{ccccccccc} H_2(M') & \longrightarrow & H_1(\partial E_{L'}) & \longrightarrow & H_1(E_{L'}) \oplus H_1(vL') & \longrightarrow & H_1(M') & \longrightarrow & 0 \\ \downarrow f_* & & \downarrow \cong & & \downarrow & & \downarrow \cong & & \\ H_2(X_n(L)) & \longrightarrow & H_1(\partial E_L) & \longrightarrow & H_1(\pi/\pi_n) \oplus H_1(vL) & \longrightarrow & H_1(X_n(L)) & \longrightarrow & 0 \\ \uparrow \iota_n(L)_* & & \parallel & & \uparrow \cong & & \uparrow \cong & & \\ H_2(M) & \longrightarrow & H_1(\partial E_L) & \longrightarrow & H_1(E_L) \oplus H_1(vL) & \longrightarrow & H_1(M) & \longrightarrow & 0 \end{array}$$

Using this diagram, it suffices to prove that the composition $H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L)$ of the isomorphisms in the second column from the left, induced by restrictions of f and $\iota_n(L)$, isomorphically

sends $\text{im}(H_2(M') \rightarrow H_1(\partial E_{L'}))$ to $\text{im}(H_2(M) \rightarrow H_1(\partial E_L))$. We accomplish this by appealing to the commutative diagram

$$\begin{array}{ccccc}
 H^1(M') & \xrightarrow[\cong]{\cap[M']} & H_2(M') & \xrightarrow{\Delta} & H_1(\partial E_{L'}) \\
 f^* \uparrow \cong & & & & \downarrow \cong \\
 H^1(X_n(L)) & & & & H_1(\partial E_L) \\
 \iota_n(L)^* \downarrow \cong & & & & \parallel \\
 H^1(M) & \xrightarrow[\cong]{\cap[M]} & H_2(M) & \xrightarrow{\Delta} & H_1(\partial E_L)
 \end{array}$$

The desired isomorphism $\pi'/\pi'_2 \rightarrow \pi/\pi_2$ then follows from a diagram chase.

To complete the proof of Claim 3, we again consider diagram (**). By assumption (3) of Theorem E, the image of the composition $\text{proj} \circ (\cap \theta) = \text{proj} \circ (\cap f_*[M'])$ contains the image of $H_2(\pi/\pi_n)/K_j(\pi/\pi_n)$ in $H_2(X_n(L))/\text{im}(K_j(\pi/\pi_n))$. The surjectivity of $H_2(\pi') \rightarrow H_2(\pi/\pi_n)/K_j(\pi/\pi_n)$ then follows from diagram (**) and a different part of the above diagram of Mayer–Vietoris sequences, namely

$$\begin{array}{ccccccc}
 H_2(\partial E_{L'}) & \longrightarrow & H_2(E_{L'}) & \longrightarrow & H_2(M') & \longrightarrow & H_1(\partial E_{L'}) \\
 \downarrow \cong & & \downarrow & & \downarrow f_* & & \downarrow \cong \\
 H_2(\partial E_L) & \longrightarrow & H_2(\pi/\pi_n) & \longrightarrow & H_2(X_n(L)) & \longrightarrow & H_1(\partial E_L) \\
 \parallel & & \uparrow & & \uparrow \iota_n(L)_* & & \parallel \\
 H_2(\partial E_L) & \longrightarrow & H_2(E_L) & \longrightarrow & H_2(M) & \longrightarrow & H_1(\partial E_L)
 \end{array}$$

After a diagram chase similar to the proof of Theorem C, Claim 3 then follows from the Stallings–Dwyer theorem.

To complete the proof of the theorem, we note that the isomorphism $\pi'/\pi'_n \xrightarrow{\cong} \pi/\pi_n$ we obtain from Claim 3, along with the isomorphism $H_1(\partial E_{L'}) \xrightarrow{\cong} H_1(\partial E_L)$, constitutes an n -basing over the trivial group for L' relative to L . Thus, the homotopy class $h_n(L', \phi)$ is defined, and $h_n(L', \phi)_*[M'] = f_*[M'] = \theta \in H_3(X_n(L))$. $\square$

8 Relationships to previous work

In this section, we investigate relationships between the lower central homotopy invariants and previous versions of Milnor's invariants. We apply Theorem D in Section 8.1 to show that for links in S^3 and $n \geq 2$ the invariant $\bar{\mu}_n$ relative to the unlink is equivalent to Milnor's $\bar{\mu}$ -invariant [1957] of length $n + 1$. The connection between the stronger invariants h_n and θ_n and Milnor's $\bar{\mu}$ -invariants passes through invariants of Orr [1989]. In Section 8.2, we show that the invariants θ_n and $\bar{\mu}_n$ for empty links are precisely the Cha–Orr “Milnor invariants of 3-manifolds” [2024]. Finally, in Section 8.3, we again leverage Theorem D to prove that the vanishing of the lower central homotopy invariants, taken in specific contexts, implies the vanishing of other extensions of Milnor's invariants. We consider the $\bar{\tau}$ -invariants for knots in Seifert-fibered 3-manifolds of D Miller [1995], and the Dwyer number of M Kuzbary [2024].

8.1 Invariants of links in S^3

We begin by understanding the relationship between the new invariants defined in this paper and Milnor's original $\bar{\mu}$ -invariants. Orr's homotopy-theoretic version [1989] of Milnor's $\bar{\mu}$ -invariants [1957] is essential for understanding this connection. His approach significantly improved the understanding of Milnor's invariants and inspired much of the work in this paper. Orr [1989] proves a realization theorem for his invariants, shows his invariants are equivalent to Milnor's, computes the number of linearly independent invariants for any fixed length, and establishes a connection to n -cobordism. We first apply Theorem D to establish the equivalence for $n \geq 2$ of the invariant $\bar{\mu}_n$ relative to the unlink and Milnor's $\bar{\mu}$ -invariant of length $n + 1$. Then, we show that the lower central homotopy invariants for links in S^3 relative to the m -component unlink are equivalent to Orr's invariants. Finally, we leverage work of Igusa and Orr [2001] to establish a relationship between the h_n and Milnor's $\bar{\mu}$ -invariants. For all of Section 8.1, we assume all links live in S^3 , we always choose the unlink as our fixed link, and we take the word "longitude" to mean "preferred longitude".

We briefly summarize Orr's construction. Suppose $L \subset S^3$ is an m -component link with vanishing Milnor $\bar{\mu}$ -invariants of lengths $\leq n$ for some $n \geq 2$. Equivalently, the longitudes of L are in π_n , so their images under the map $E_L \rightarrow K(\pi/\pi_n, 1)$ are nullhomotopic. Choosing meridians for L (i.e., a basing for L) defines an isomorphism $\tau : F/F_n \xrightarrow{\cong} \pi/\pi_n$, where $F = F(m)$ is the m -generator free group. We may view its inverse $\phi : \pi/\pi_n \xrightarrow{\cong} F/F_n$ as an isomorphism to the n -th lower central quotient of the link group of the m -component unlink U . This yields a map $E_L \rightarrow K(F/F_n, 1)$, which we may assume sends the components of ∂E_L by projection onto the images of the corresponding meridians. The map $E_L \rightarrow K(F/F_n, 1)$ therefore extends to a map $S^3 \rightarrow K_n$, where $K_n = K(F/F_n, 1) \cup (\bigsqcup_{m_i} (\bigsqcup D^2))$ is obtained by gluing m disks D^2 to the Eilenberg–Mac Lane space along the images of the meridians. The map $S^3 \rightarrow K_n$ sends the regular neighborhood of each component of L by projection onto the corresponding D^2 .

Orr denotes the resulting homotopy class of maps $S^3 \rightarrow K_n$ by $\theta_n(L, \tau) \in \pi_3(K_n)$. To avoid confusion with our notation and to emphasize the relationship with the invariants from Section 3.1, we will instead denote this class by $h_n^O(L, \phi) \in \pi_3(K_n)$. Similarly to the invariant h_n , one could remove the dependency on ϕ and define $h_n^O(L)$. Note that $h_n^O(U, \text{id})$ is nullhomotopic because the map $E_U \rightarrow K(F/F_n, 1)$ factors through the wedge of circles to which the disks are attached. Thus, $h_n^O(L, \phi) = h_n^O(U, \text{id})$ if and only if $h_n^O(L, \phi)$ is nullhomotopic. In other words, $h_n^O(L, \phi)$ vanishes if and only if it is nullhomotopic. We will construct a map $r_n : X_n(U) \rightarrow K_n$ relating the lower central homotopy invariant h_n to Orr's invariant h_n^O , so that $r_n \circ h_n(L, \phi) = h_n^O(L, \phi)$.

Lemma 8.1 *An n -basing for L relative to U exists if and only if Milnor's $\bar{\mu}$ -invariants for L of lengths $\leq n$ vanish.*

Proof ($\implies$): An n -basing (ϕ, ϕ_∂) for L relative to U includes an isomorphism $\phi : \pi/\pi_n \xrightarrow{\cong} F/F_n$ which sends (based) longitudes of L to (based) longitudes of U . This implies the longitudes of L are trivial in π/π_n or, equivalently, Milnor's $\bar{\mu}$ -invariants of lengths $\leq n$ vanish.

($\Leftarrow$): Now suppose Milnor's $\bar{\mu}$ -invariants of lengths $\leq n$ vanish. The longitudes of L are therefore in π_n , so we may define an isomorphism $\phi : \pi/\pi_n \xrightarrow{\cong} F/F_n$ which preserves meridians and longitudes. In other words, there is an admissible isomorphism $\phi_\partial : H_1(\partial E_L) \xrightarrow{\cong} H_1(\partial E_U)$ compatible with an isomorphism $\phi : \pi/\pi_n \xrightarrow{\cong} F/F_n$. Thus, we have produced an n -basing (ϕ, ϕ_∂) for L relative to U . $\square$

Corollary 8.2 *For $n \geq 2$, the invariant $\bar{\mu}_n(L)$ vanishes if and only if Milnor's $\bar{\mu}$ -invariants for L of lengths $\leq n+1$ vanish.*

Proof Let $n \geq 2$. By Theorem D, $\bar{\mu}_n(L)$ vanishes if and only if there exists an $(n+1)$ -basing for L relative to U . By Lemma 8.1, this is true if and only if Milnor's $\bar{\mu}$ -invariants for L of lengths $\leq n+1$ vanish. $\square$

Corollary 8.3 *For $n \geq 2$, the invariant $h_n(L)$ is defined if and only if Orr's invariant $h_n^O(L)$ is defined.*

Relative versions of Lemma 8.1 and Corollary 8.2 hold:

Lemma 8.4 *Suppose L and L' are links with vanishing Milnor $\bar{\mu}$ -invariants of lengths $\leq n$. An $(n+1)$ -basing for L' relative to L exists if and only if Milnor's $\bar{\mu}$ -invariants for L and L' of length $n+1$ agree.*

Corollary 8.5 *Suppose L and L' are links with vanishing Milnor $\bar{\mu}$ -invariants of lengths $\leq n$ for some $n \geq 2$. Then $\bar{\mu}_n(L) = \bar{\mu}_n(L')$ if and only if Milnor's $\bar{\mu}$ -invariants for L and L' of length $n+1$ agree.*

Thus, for $n \geq 2$, the invariant $\bar{\mu}_n$ is equivalent to Milnor's (total) $\bar{\mu}$ -invariant of length $n+1$. We provide a brief sketch of the proof. Details involving the additivity of the first nonvanishing Milnor invariants appear in [Cochran 1990] and [Orr 1989].

Sketch of the proofs of Lemma 8.4 and Corollary 8.5 For links L and L' with vanishing Milnor $\bar{\mu}$ -invariants of lengths $\leq n$, Milnor's $\bar{\mu}$ -invariants of length $n+1$ for L and L' agree if and only if $L \# (-L')$ (with any chosen bands for the connected sum) has vanishing Milnor invariants of length $n+1$ or, equivalently, the differences of the longitudes of L and L' lie in the $(n+1)$ -st lower central subgroup. Thus, Milnor's $\bar{\mu}$ -invariants of length $n+1$ for L and L' agree if and only if we can define an isomorphism $\pi'/\pi'_{n+1} \xrightarrow{\cong} \pi/\pi_{n+1}$ preserving (based) meridians and longitudes, that is, an $(n+1)$ -basing for L' relative to L . Corollary 8.5 then follows from Theorem D. $\square$

We now explore the relationship between the space $X_n(U)$ and Orr's space K_n . Recall that $X_n(U)$ is formed by gluing νU to $M_{p_n^U}^{\times U}$, the reduced mapping cylinder of $p_n^U : E_U \rightarrow K(F/F_n, 1)$. We view this gluing in two stages: First, attach a disk D^2 to each meridian of U . Then, fill in the remainder of each component of ∂E_U with a 3-ball. Since $M_{p_n^U}^{\times U}$ is also a $K(F/F_n, 1)$, the inclusion map $E_U \hookrightarrow M_{p_n^U}^{\times U}$ is homotopic to a map which projects each component of ∂E_U to the corresponding meridian in $M_{p_n^U}^{\times U}$. This factors through a homotopy self-equivalence $M_{p_n^U}^{\times U} \xrightarrow{\cong} M_{p_n^U}^{\times U}$ which behaves this way on ∂E_U .

After the first stage of gluing, we see that $M_{p_n^U}^{\times U} \cup (\bigsqcup m_i) (\bigsqcup D^2)$ is a model for the space K_n . We may extend the homotopy self-equivalence $M_{p_n^U}^{\times U} \xrightarrow{\cong} M_{p_n^U}^{\times U}$ by the identity on each D^2 , giving a homotopy equivalence $M_{p_n^U}^{\times U} \cup (\bigsqcup m_i) (\bigsqcup D^2) \xrightarrow{\cong} K_n$. This map extends over each 3-ball we attach to form $X_n(U)$

by viewing $D^3 = D^1 \times D^2$ and projecting to $D^2 \subset K_n$. We thus obtain a map $r_n : X_n(U) \rightarrow K_n$. By construction, $r_n \circ h_n(U, \text{id}) = r_n \circ \iota_n(U) = h_n^O(U, \text{id})$. We have the diagram

$$\begin{array}{ccccc}
 E_U & \xrightarrow{\quad} & M_{p_n^U}^\times & & \\
 \downarrow & \searrow & \nearrow \simeq & & \downarrow \\
 & & M_{p_n^U}^\times & & \\
 \downarrow & & \downarrow h_n^O(U, \text{id}) & & \downarrow \\
 S^3 & \xrightarrow{\quad} & X_n(U) & \xrightarrow{r_n} & K_n.
 \end{array}$$

$\iota_n(U)$

The image in K_n of the attaching map of each 3-ball we glue to $M_{p_n^U}^\times \cup (\bigsqcup m_i) (\bigsqcup D^2)$ to form $X_n(U)$ is nullhomotopic. Thus, $X_n(U) \simeq K_n \vee (\bigvee^m S^3)$, and the map r_n may be viewed as collapsing each of these S^3 factors to the basepoint.

Now suppose L admits an n -basing (ϕ, ϕ_∂) relative to U . By Corollary 8.2, this is the case if and only if Milnor's $\bar{\mu}$ -invariants of lengths $\leq n$ vanish for L . Equivalently, the longitudes of L are in π_n . Define a homotopy self-equivalence $M_{p_n}^\times \xrightarrow{\simeq} M_{p_n}^\times$, similar to the self-equivalence of $M_{p_n^U}^\times$ above, which projects each component of ∂E_L to the corresponding meridian. The rightmost square in the diagram above fits into the diagram

$$\begin{array}{ccccccc}
 E_L & \xrightarrow{\quad} & M_{p_n}^\times & \xrightarrow{\quad \simeq \quad} & M_{p_n^U}^\times & & \\
 \downarrow & \searrow & \nearrow \simeq & & \nearrow \simeq & & \downarrow \\
 & & M_{p_n}^\times & \xrightarrow[h_\phi]{\simeq} & M_{p_n^U}^\times & & \\
 \downarrow & & \downarrow & & \downarrow & & \downarrow \\
 S^3 & \xrightarrow{\quad} & X_n(L) & \xrightarrow[h_\phi]{\simeq} & X_n(U) & \xrightarrow{r_n} & K_n
 \end{array}$$

$\iota_n(L)$ $h_n^O(L, \phi)$ $h_n(L, \phi)$

The maps $M_{p_n}^\times \xrightarrow{\simeq} M_{p_n^U}^\times$ on the top square of the diagram are both homotopy equivalences of pairs; in particular, the map on the front face sends ∂E_L to ∂E_U by an admissible homeomorphism, and the map on the back face sends the meridians of L to the meridians of U . The extensions of the maps on the top of the diagram to the maps on the bottom therefore commute as claimed. Commutativity of the diagram proves the following proposition.

Proposition 8.6 *The map $r_n : X_n(U) \rightarrow K_n$ satisfies $r_n \circ h_n(L, \phi) = h_n^O(L, \phi)$.*

Recall from Section 3.1 that we say $h_n(L, \phi)$ *vanishes* if $h_n(L, \phi) = h_n(U, \text{id})$. By Proposition 8.6, the vanishing of $h_n(L, \phi)$ implies the vanishing of $h_n^O(L, \phi)$: If $h_n(L, \phi) = h_n(U, \text{id})$, then $h_n^O(L, \phi) = r_n \circ h_n(L, \phi) = r_n \circ h_n(U, \text{id}) = h_n^O(U, \text{id})$. The converse also holds.

Proposition 8.7 *The invariants h_n and h_n^O are equivalent, i.e., $h_n(L) = h_n(L')$ if and only if $h_n^O(L) = h_n^O(L')$. In particular, $h_n(L)$ vanishes if and only if $h_n^O(L)$ is nullhomotopic.*

Proof We could prove this directly, but a faster proof leverages the relationships between the invariants in question and n -cobordism. By Theorem B, the invariants $h_n(L)$ and $h_n(L')$ agree if and only if L and L' are n -cobordant. By Theorem 7 of [Orr 1989], this is the case if and only if $h_n^O(L) = h_n^O(L')$. $\square$

To prove Proposition 8.7 directly, one views $X_n(U)$ as $K_n \vee (\bigvee^m S^3)$ and observes that all invariants differ only on K_n .

Proposition 8.7 implies that the sets of realizable classes in $\pi_3(X_n(U))$ and $\pi_3(K_n)$ are in bijective correspondence. By Theorem 4 of [Orr 1989], all classes in $\pi_3(K_n)$ are realizable. Since $\pi_2(S^3) = 0$, condition (3) of the realization theorem for h -invariants, Theorem C, follows from conditions (1) and (2) by Corollary 6.2. Condition (2) of Theorem C is vacuously satisfied when $M = S^3$, so every homotopy class in $\pi_3(X_n(U))$ satisfying condition (1) of Theorem C is realizable. Thus, homotopy classes in $\pi_3(X_n(U))$ satisfying condition (1) of Theorem C correspond bijectively to elements of $\pi_3(K_n)$.

To conclude this section, we consider other versions of Orr's invariant. Consider the homology class $\theta_n^O(L, \phi) = h_n^O(L, \phi)_*[S^3] \in H_3(K_n)$, and let $\bar{\mu}_n^O(L, \phi)$ denote the image of this class in $\text{coker}(H_3(K_{n+1}) \rightarrow H_3(K_n))$, which is isomorphic to $\text{coker}(\pi_3(K_{n+1}) \rightarrow \pi_3(K_n))$ as seen in [Orr 1989; Igusa and Orr 2001]. Again, one could remove the dependency on ϕ and define $\theta_n^O(L)$ and $\bar{\mu}_n^O(L)$; in fact, Orr shows that $\bar{\mu}_n^O(L, \phi)$ is independent of the choice of ϕ . Theorem 12 of [Orr 1989] states that $\bar{\mu}_n^O(L)$ vanishes if and only if L has vanishing Milnor $\bar{\mu}$ -invariants of lengths $\leq n + 1$. Combining this discussion with Corollary 8.2, we have:

Corollary 8.8 *For $n \geq 2$, the invariant $\bar{\mu}_n(L)$ vanishes if and only if the invariant $\bar{\mu}_n^O(L)$ vanishes.*

Work of Igusa and Orr [2001] establishes the exact relationship between Milnor's $\bar{\mu}$ -invariants and the stronger invariants $\theta_n^O(L)$ and $h_n^O(L)$ for $n \geq 2$. We summarize their results in Table 1.

The final row of this table is known as the k -slice theorem due to the relationship between the invariant h_k^O and k -cobordism. A link which is k -cobordant to the unlink is called k -slice. By Proposition 8.7, the k -slice theorem still holds if we replace Orr's invariant h_k^O with the lower central homotopy invariant h_k . One may ask whether analogous relationships among the new invariants h_n , θ_n and $\bar{\mu}_n$ from Sections 3.1 and 3.5 hold for links in other 3-manifolds. This is a topic of future exploration; see Section 9.1.

Orr's invariant	equivalent Milnor $\bar{\mu}$ -invariant(s)
$\bar{\mu}_n^O$	$\bar{\mu}$ of length $n + 1$
θ_n^O	$\bar{\mu}$ of lengths $n + 1 \leq k \leq 2n - 1$
h_n^O	$\bar{\mu}$ of lengths $n + 1 \leq k \leq 2n$

Table 1

8.2 Milnor invariants of 3-manifolds

One of the most recent versions of Milnor's invariants is a collection of homology cobordism invariants of 3-manifolds due to Cha and Orr [2024]. In addition to Orr's work [Orr 1989], the paper [Cha and Orr 2024] also inspired many ideas in this paper. Cha and Orr refer to their invariants as "Milnor invariants of 3-manifolds" for good reason; the invariants exhibit behavior similar to Milnor's link invariants. In fact, Cha and Orr show one can recover Milnor's $\bar{\mu}$ -invariants of a link $L \subset S^3$ as invariants of the 3-manifold which is 0-framed surgery on every component of L .

In this section, we recover the Cha–Orr invariants from the lower central homology invariants θ_n and Milnor invariants $\bar{\mu}_n$ for empty links. Comparing empty links is equivalent to comparing 3-manifolds. Theorem A' proves that θ_n^G and $\bar{\mu}_n^G$ are invariants of $\mathbb{Z}[G]$ -homology cobordism, as a homology concordance of empty links is the same as a homology cobordism of 3-manifolds. For this section, we use the notation from Section 3.5. In particular, any 3-manifold M we consider will be equipped with a fixed surjective homomorphism $\varphi : \pi_1(M) \twoheadrightarrow G$ onto a fixed group G .

Consider such a pair (M, φ) . The empty link $\emptyset_M$ in M has exterior $E_{\emptyset_M} = M$ and link group $\pi_1(M)$. The appropriate $\mathbb{Z}[G]$ -lower central quotients are the $\mathbb{Z}[G]$ -lower central quotients of $\pi_1(M)$. Following our standard notation, we set $\pi = \pi_1(M)$ and $\Gamma = \ker \varphi$. The homotopy pushout diagram $\mathcal{D}_n^G(L)$ from Section 3.4 becomes

$$\begin{array}{ccc} M & \longrightarrow & K(\pi/\Gamma_n, 1) \\ \parallel & & \\ M & & \end{array}$$

Thus, the space $X_n(\emptyset_M)$ is just the reduced mapping cylinder of $M \rightarrow K(\pi/\Gamma_n, 1)$, which is itself a $K(\pi/\Gamma_n, 1)$, and the map $\iota_n(\emptyset_M)$ is a map $M \rightarrow K(\pi/\Gamma_n, 1)$ induced by the canonical projection $\pi \twoheadrightarrow \pi/\Gamma_n$.

Now fix the empty link $\emptyset_M$, and suppose we wish to compare another empty link $\emptyset_{M'}$ to $\emptyset_M$. To define the invariants from Section 3.5, we require an n -basing over G for $\emptyset_{M'}$ relative to $\emptyset_M$. Because the empty link exteriors have no boundary, an n -basing over G for $\emptyset_{M'}$ relative to $\emptyset_M$ is just an isomorphism $\phi : \pi'/\Gamma'_n \xrightarrow{\cong} \pi/\Gamma_n$ over G . Given such an isomorphism ϕ , the invariant $\theta_n(\emptyset_{M'}, \phi)$ is the image of $[M']$ under

$$H_3(M') \rightarrow H_3(\pi'/\Gamma'_n) \xrightarrow[\phi_*]{\cong} H_3(\pi/\Gamma_n).$$

This produces a $\mathbb{Z}[G]$ -version of the Cha–Orr invariant from [2024].

Taking G to be the trivial group (so that we obtain invariants of $\mathbb{Z}$ -homology cobordism) yields exactly the Cha–Orr invariant $\theta_n(M', \phi)$. In this setting, the set of realizable classes $\mathcal{R}_n(\emptyset_M) \subseteq H_3(\pi/\pi_n)$ is the set of realizable classes of 3-manifold invariants $\mathcal{R}_n(\pi)$ from [Cha and Orr 2024]. The definitions of the equivalence relations on these sets of realizable classes coincide, so the invariant $\bar{\mu}_n(\emptyset_{M'})$ agrees with the Cha–Orr Milnor invariant $\bar{\mu}_n(M')$.

Observe that the invariants $\theta_n(L')$ and $\bar{\mu}_n(L')$ for a link $L' \subset M'$ relative to a fixed link $L \subset M$ determine the 3-manifold invariants $\theta_n(M')$ and $\bar{\mu}_n(M')$ relative to M . An n -basing of links induces an

n -basing of corresponding sublinks, and the invariant for a link determines the invariants of its sublinks, including the empty link. Thus, if one wishes to compute the “number” of Milnor invariants relative to a fixed link L , one should restrict to counting invariants $\bar{\mu}_n(L')$ whose associated 3-manifold invariants $\bar{\mu}_n(M')$ vanish.

Because the invariants from Section 3.5 reduce to the Cha–Orr invariants, some of the main results from [Cha and Orr 2024] are special cases of results in this paper. When restricted to the natural numbers, Theorems B, C and E and Corollary D from [Cha and Orr 2024] follow from our Theorem D. The realization result Theorem G from [Cha and Orr 2024] restricted to natural numbers follows from our Theorem E. Cha and Orr extend their invariants to the transfinite setting, and their results hold in this more general context. We discuss such an extension of our work in Section 9.2.

8.3 Other extensions of Milnor's invariants

Theorem D states that the Milnor invariants $\bar{\mu}_n^G$ determine the $\mathbb{Z}[G]$ -lower central quotients associated to a link $L \subset M$ one step at a time. The implications of Theorem D include the equivalence for $n \geq 2$ of $\bar{\mu}_n$ for links in S^3 relative to the unlink with Milnor's $\bar{\mu}$ -invariant of length $n + 1$ as seen in Section 8.1. In this section, we again leverage Theorem D to show that the vanishing of $\bar{\mu}_n$, defined in particular settings, implies the vanishing of extensions of Milnor's invariants due to Miller [1995] and Kuzbary [2024].

8.3.1 Miller's $\bar{\tau}$ -invariants Miller [1995] considers knots in aspherical Seifert-fibered 3-manifolds in the free homotopy class of the Seifert fiber K . Miller fixes the Seifert fiber K and compares other knots in the homotopy class of the Seifert fiber relative to K . We remind the reader that Seifert-fibered 3-manifolds are precisely those M for which $\pi_1(M)$ contains a normal subgroup isomorphic to $\mathbb{Z}$. This normal subgroup is generated by the class of the Seifert fiber, and $\pi_1(M)$ may be described as the extension

$$1 \rightarrow \mathbb{Z} \hookrightarrow \pi_1(M) \twoheadrightarrow S \rightarrow 1,$$

where S is the orbifold fundamental group of the base surface of the fibration.

In our notation, Miller takes $G = S$, so that $\Gamma = \ker(\pi \twoheadrightarrow \pi_1(M) \twoheadrightarrow S)$, where $\pi = \pi_1(E_K)$ and $\Gamma/\Gamma_n \cong \mathbb{Z} \times F/F_n$, where F is a free group. These quotients are analogous to the lower central quotients of the unlink $U \subset S^3$. Given a knot K' freely homotopic to K , Miller finds a presentation for Γ'/Γ'_n analogous to Milnor's for $L \subset S^3$. He then proceeds in a manner analogous to Milnor [1957], using the Magnus expansion to define integer-valued invariants. These integers, which Miller calls the $\bar{\tau}$ -invariants of K' , are the coefficients of the images of distinguished longitudes of the S -cover of K' in a power series ring. Note that the image in S of the homotopy class of the Seifert fiber is trivial, so the preimage of any knot in that class under the S -cover of M is a link with infinitely many components, indexed by elements of S .

Miller proves that all $\bar{\tau}$ -invariants vanish for the fixed knot K . If $\Gamma'/\Gamma'_n \cong \Gamma/\Gamma_n$ preserving meridians and distinguished longitudes, then the $\bar{\tau}$ -invariants for K' of lengths $\leq n$ vanish. Miller's work implies a theorem analogous to Milnor's theorem, although it is not directly proven in [Miller 1995]: We have

an isomorphism $\Gamma'/\Gamma'_{n+1} \cong \Gamma/\Gamma_{n+1}$ or, equivalently, an isomorphism $\Gamma'/\Gamma'_n \cong \Gamma/\Gamma_n$ preserving meridians and distinguished longitudes *if and only if* the $\bar{\tau}$ -invariants for K' of lengths $\leq n$ vanish. Miller concludes by proving a constructive realization theorem for the $\bar{\tau}$ -invariants.

Proposition 8.9 *Let M be an aspherical Seifert-fibered 3-manifold, K the Seifert fiber of M , and K' some knot freely homotopic to K . Let S be the quotient of $\pi_1(M)$ by the infinite cyclic normal subgroup corresponding to the Seifert fiber. Suppose for some $n \geq 2$ that the n -th S -Milnor invariant $\bar{\mu}_n^S(K')$ vanishes, that is, $\bar{\mu}_n^S(K')$ is defined and equals $\bar{\mu}_n^S(K)$. Then Miller's $\bar{\tau}$ -invariants for K' of lengths $\leq n+1$ vanish.*

Proof By Theorem D, since $\bar{\mu}_n^S(K) = \bar{\mu}_n^S(K')$, K' admits an $(n+1)$ -basing over S relative to K . This consists of an isomorphism $\phi: \pi'/\pi'_{n+1} \xrightarrow{\cong} \pi/\pi_{n+1}$ which preserves meridians and distinguished longitudes. This yields an isomorphism $\Gamma'/\Gamma'_{n+1} \xrightarrow{\cong} \Gamma/\Gamma_{n+1}$ which preserves meridians and distinguished longitudes. By Miller's work, the $\bar{\tau}$ -invariants of lengths $\leq n+1$ vanish. $\square$

8.3.2 Kuzbary's Dwyer number Kuzbary [2024] develops the *Dwyer number*, an invariant for knots in closed oriented 3-manifolds that behaves like the first nonvanishing Milnor invariant. This invariant compares the lower central quotients π'/π'_n associated to a knot $K' \subset M$ to the quotients π/π_n for a fixed knot $K \subset M$. Using the geometrically defined Dwyer subgroups $\Phi_n(-) \leq H_2(-; \mathbb{Z})$ first formulated in [Freedman and Teichner 1995] (see also [Cochran and Harvey 2008]), Kuzbary defines the *Dwyer number of K' relative to K* as

$$D(K', K) = \max \left\{ n \mid \frac{H_2(E_{K'})}{\Phi_n(E_{K'})} \cong \frac{H_2(E_K)}{\Phi_n(E_K)} \right\} \in \mathbb{N} \cup \{\infty\}.$$

Kuzbary shows that $H_2(X)/\Phi_n(X) \cong H_2(\pi_1(X))/K_n(\pi_1(X))$. From the statements of Theorems C and E, recall that for a group H , we define an abelian group $K_n(H) = \ker(H_2(H) \rightarrow H_2(H/H_n))$. The Stallings–Dwyer theorem implies $D(K', K) \geq n$ if and only if there is an isomorphism $\pi'/\pi'_{n+1} \cong \pi/\pi_{n+1}$. Kuzbary leverages the Dwyer number to analyze concordance of nullhomologous knots in $\#^k(S^1 \times S^2)$ and establish a relationship with Massey products.

Proposition 8.10 *Fix a knot K in a closed orientable 3-manifold M , and let $K' \subset M$ be another knot freely homotopic to K . Let G be the trivial group. Suppose for some $n \geq 2$ that the n -th Milnor invariant $\bar{\mu}_n^G(K')$ vanishes, that is, $\bar{\mu}_n^G(K')$ is defined and equals $\bar{\mu}_n^G(K)$. Then $D(K', K) \geq n$.*

Proof By Theorem D, since $\bar{\mu}_n^G(K') = \bar{\mu}_n^G(K)$, K' admits an $(n+1)$ -basing over G relative to K . In particular, this includes an isomorphism $\phi: \pi'/\pi'_{n+1} \xrightarrow{\cong} \pi/\pi_{n+1}$. Thus, $D(K', K) \geq n$. $\square$

9 Future directions

We conclude this paper with a discussion of several current and future directions, including a number of open questions.

9.1 Further investigation of new invariants

The fact that the invariants from Sections 3.1 and 3.5 reduce to Milnor's $\bar{\mu}$ -invariants in the case where we compare links in S^3 to the unlink motivates us to ask whether known results for Milnor's $\bar{\mu}$ -invariants extend to the more general setting of this work.

Algebraic information captured by h_n and θ_n Igusa and Orr [2001] answer a longstanding open question known as the k -slice conjecture affirmatively: A link $L \subset S^3$ is k -cobordant to the unlink if and only if Milnor's $\bar{\mu}$ -invariants of lengths $2 \leq i \leq 2k$ vanish. They prove this using Igusa pictures [1979] and Orr's invariants [1989] which characterize k -cobordism and are analogous to the h -invariants from Section 3.1. Recall from Section 8.1 that the lower central homotopy invariant h_k is equivalent to Orr's invariant of length k . As discussed in Section 8.1, Igusa and Orr show that Orr's length k invariant captures Milnor's $\bar{\mu}$ -invariants of lengths $k + 1 \leq i \leq 2k$. In the more general setting of this work, it is probable that we also lose significant information passing from h_k or θ_k to the invariants $\bar{\mu}_k$. Understanding how much more information these stronger invariants capture is equivalent to understanding to what extent a relative version of the k -slice conjecture holds for links in closed orientable 3-manifolds.

Question 9.1 Does h_k vanish if and only if $\bar{\mu}_i$ vanishes for $k \leq i \leq 2k - 1$?

Whitney towers and surface towers Interpretations of Milnor's $\bar{\mu}$ -invariants in terms of twisted Whitney towers and surface towers (also called gropes in the literature) are given in [Freedman and Teichner 1995; Conant et al. 2011; Conant et al. 2014; Cha 2018]. Cha and Orr [2024] give a similar interpretation for “Milnor invariants of 3-manifolds” and extend it to the transfinite setting.

Goal 9.2 Provide an interpretation of the invariants from this paper in terms of Whitney towers and surface towers.

Finite type invariants Bar-Natan [1995] and X Lin [1991] showed that Milnor's invariants for string links are of finite type. Habegger and Masbaum [2000] determine a formula for computing Milnor's $\bar{\mu}$ -invariants for string links from the Kontsevich integral.

Goal 9.3 Recast the invariants from this paper as finite type invariants.

M Powell has suggested that the theory of claspers due to Gusarov [2000] and Habiro [2000] may seed a constructive realization machine for these invariants.

Invariants comparing links in S^3 to a nontrivial link In the case where we compare links in S^3 to the unlink, the invariants h_n from Section 3.1 reduce to invariants of Orr [1989] which capture Milnor's $\bar{\mu}$ -invariants.

Question 9.4 Do we gain anything by fixing a nontrivial link $L \subset S^3$ instead of the unlink?

Despite the additivity of Milnor's invariants, due to Cochran [1990] and Orr [1989] in the case of the first nonvanishing invariant and more generally due to Krushkal [1998], we expect the answer is yes.

9.2 Extensions of this work

In addition to the potential new interpretations of the invariants in this paper described in Section 9.1, we may consider whether certain known extensions of Milnor’s $\bar{\mu}$ -invariants carry over to the setting of this work.

Extending to the transfinite setting The existence of nontrivial transfinite Milnor invariants for links has been an open question since Milnor’s work [1957]. Milnor originally asked about invariants detecting phenomena in the transfinite lower central quotients of link groups for links in S^3 . Using J Levine’s algebraic closure of groups [1989b; 1989a], equivalent to Vogel’s homology localization of groups [1978], Cha and Orr [2024] extend their homology cobordism invariants to the transfinite setting and answer a similar question. They exhibit 3-manifolds with vanishing finite-length Milnor invariants and nontrivial transfinite invariants. One expects analogous extensions of the invariants h_n , θ_n and $\bar{\mu}_n$ should exist, in particular because the θ - and $\bar{\mu}$ -invariants contain the Cha–Orr invariants as special cases. These extensions would involve a tower of spaces $X_\omega(L)$ with maps from M , where ω is a countable ordinal.

Goal 9.5 Leverage the examples from [Cha and Orr 2024] to find the first instances of nontrivial transfinite Milnor invariants for links.

J Levine’s group closure One expects an extension of the invariants h_n , θ_n and $\bar{\mu}_n$ even beyond the previous idea to invariants that detect phenomena at the level of J Levine’s algebraic closure [1989b; 1989a]. These invariants would probe this algebraic closure without passing to (finite or transfinite) lower central quotients. They would be “universal π_1 invariants” in the sense that they would detect information about the $\mathbb{Z}[G]$ -homology of link exteriors coming from link groups at the deepest possible level. In other words, these invariants would detect information about the fundamental group of the $\mathbb{Z}[G]$ -Vogel homology localization of the link exterior invisible to any lower central quotient; see [Vogel 1978]. Such invariants would be analogous to J Levine’s invariant [1989a] for links in S^3 . Whether this invariant is nontrivial is an open question.

Just as in the transfinite setting, Cha and Orr [2024] extend their homology cobordism invariants to this universal π_1 setting and prove the existence of 3-manifolds with vanishing finite- and transfinite-length Milnor invariants but with nontrivial universal invariants. Again, one expects analogous extensions of the invariants h_n , θ_n and $\bar{\mu}_n$ since the θ - and $\bar{\mu}$ -invariants contain the Cha–Orr invariants as special cases. These invariants would involve spaces $X(L, 1)$ equipped with maps from M . The various maps from M would fit into a diagram as follows:

$$\begin{array}{ccccccc}
 & & M & & & & \\
 & \swarrow & & \searrow & \searrow & \searrow & \\
 X(L, 1) & \longrightarrow & \cdots & \longrightarrow & X_\omega(L) & \longrightarrow & \cdots \longrightarrow X_n(L) \longrightarrow \cdots \longrightarrow X_1(L)
 \end{array}$$

Goal 9.6 Leverage the work from [Cha and Orr 2024] to find links with nontrivial universal Milnor invariants.

Milnor invariants for knotted surfaces The study of surface links (including single-component knotted surfaces) in 4-manifolds is currently a very active area in 4-dimensional topology. Concordance in this setting has garnered a lot of recent attention [Sunukjian 2015; Schwartz 2019; Audoux et al. 2021; Klug and Miller 2021; Miller 2021; Klug and Miller 2022; Schneiderman and Teichner 2022], but little is understood overall. Despite recent work, there is a general shortage of concordance invariants for links of surfaces, particularly when the surfaces have positive genus.

Audoux, Meilhan and Yasuhara [Audoux et al. 2021] recently explored an analogue of Milnor's invariants for concordance of knotted surfaces in 4-space. They leverage the combinatorial method of cut-diagrams to apply tools similar to Milnor's work [1957] and develop concordance invariants. Analogous versions of Orr's invariants [1989] in this setting are likely amenable to computing the number of such invariants.

Goal 9.7 Extend the invariants of [Audoux et al. 2021] to links of positive genus surfaces in closed orientable 4-manifolds. More broadly, extend the invariants from this work to concordance invariants of embeddings of n -manifolds in $(n+2)$ -manifolds.

Progress toward Goal 9.7 will appear in forthcoming joint work with M Miller.

Variations on the definition of concordance We list a few variations of this work adapted to studying links in slightly different contexts:

- One could study concordance of unordered links, unoriented links, or both unordered and unoriented links by slightly changing the definitions of an n -basing and the group $\text{Aut}(\pi / \Gamma_n, \partial)$ of self- n -basings given in Section 4. Versions of the main results of this paper still follow.
- Instead of studying $\mathbb{Z}[G]$ -homology concordance of links in 3-manifolds M with fixed homomorphisms $\varphi : \pi_1(M) \twoheadrightarrow G$ to a fixed group G , as in Section 3.5, one could remove such a fixed homomorphism φ as part of the structure of M and only ask that $\pi_1(M)$ admits some homomorphism onto G . Again, this change would slightly alter the definitions of an n -basing and the group of self- n -basings.

9.3 Applications of invariants

We list several possible applications of the new invariants defined in this work to existing open questions.

Computing collections of realizable classes One can leverage Theorems C and E to compute the collection of Milnor invariants relative to a fixed link. One might even hope for a constructive algorithm to realize invariants by specific links in M . Such computations for specific knots $K \subset M$ can partially answer the Almost-concordance conjecture of [Celoria 2018; Friedl et al. 2019]. Partial progress toward this conjecture has been made in [Miller 1995; Celoria 2018; Yildiz 2018; Friedl et al. 2019; Nagel et al. 2019; Kuzbary 2024]. We conjecture that the invariants h_n are sufficient for proving the remaining cases of the Almost-concordance conjecture.

Applications to homology cobordism Recall that the invariants θ_n and $\bar{\mu}_n$ are homology concordance invariants and reduce to the homology cobordism invariants of Cha and Orr [2024] when specialized to empty links.

Question 9.8 Are there applications of the invariants in this work to the study of homology cobordism of 3-manifolds which do not pass through the Cha–Orr invariants?

The invariants θ_n and $\bar{\mu}_n$ may also be useful in answering the following question:

Question 9.9 Fix a 3-manifold M . Does there exist a 3-manifold M' , homology cobordant to M , and a knot $K \subset M$ such that K is not topologically homology concordant to any knot in M' ?

A Levine [2016] provides an example of a knot in a smoothly homology nullbordant homology sphere which is not smoothly concordant to any knot in S^3 . The existence of such an example in the topological category is still open; Davis [2020] shows that such an example cannot be detected by the powerful Cochran–Orr–Teichner filtration [Cochran et al. 2003] on the topological concordance group. The invariants defined in this work may be able to provide analogous examples in the topological setting in 3-manifolds which are not homology spheres. They may also interact in interesting ways with recent work on homology concordance such as [Zhou 2021] and [Hom et al. 2022].

Appendix Properties of $X_n(L)$ and $\iota_n(L)$

In this brief appendix, we collect a number of properties of the spaces $X_n(L)$ and maps $\iota_n(L) : M \hookrightarrow X_n(L)$ defined in Section 3.1. These include a number of well-definedness and naturality properties; see [Steas 2023] for more details.

A.1 Well-definedness and naturality properties

Observe the following well-definedness and naturality properties of $X_n(L)$ and $\iota_n(L)$. These follow from the naturality of the constructions of Section 3.1 and algebraic-topological properties of homotopy pushouts.

- The space $X_n(L)$ is well-defined up to homotopy equivalence, independent of the choices of model for $K(\pi/\Gamma_n, 1)$ and based map $p_n : E_L \rightarrow K(\pi/\Gamma_n, 1)$ induced by the canonical projection $\pi \twoheadrightarrow \pi/\Gamma_n$.
- The model obtained for $X_n(L)$ depends naturally on the choices of model for $K(\pi/\Gamma_n, 1)$ and based map p_n . Two sets of choices yield a canonical homotopy class of homotopy equivalences ψ between the two resulting models.
- The homotopy class $\iota_n(L) \in [M, X_n(L)]_0$ is well-defined.
- The inclusions $\iota_n(L) : M \hookrightarrow X_n(L)$ into two models for $X_n(L)$ are related, up to homotopy, via the canonical homotopy equivalence ψ above.

Recall the maps $\psi_{m,n} : X_m(L) \rightarrow X_n(L)$ for $m \geq n$ defined in Section 3.1. The $\psi_{m,n}$ satisfy the following:

- The homotopy classes $\psi_{m,n} \in [X_m(L), X_n(L)]_0$ are well-defined.
- The maps $\psi_{m,n}$ between two different models of each of $X_m(L)$ and $X_n(L)$ are related, up to based homotopy, via canonical homotopy equivalences.
- The collection $\{\psi_{i,j}\}$ satisfies $\psi_{m,p} \simeq \psi_{n,p} \psi_{m,n}$.
- The collections $\{\psi_{i,j}\}$ and $\{\iota_k(L)\}$ satisfy $\psi_{i,j} \iota_i(L) \simeq \iota_j(L)$.

The spaces $X_n^G(L)$ and maps $\iota_n^G(L)$ and $\psi_{m,n}^G$ defined in Section 3.4 satisfy analogous well-definedness and naturality properties.

A.2 Algebraic properties

Proposition A.1 *Let the fixed group G be $G = \pi_1(M)$, and let $\varphi : \pi_1(M) \twoheadrightarrow G$ be the identity homomorphism.*

- (1) *The map $\iota_n(L) : M \hookrightarrow X_n(L)$ induces an isomorphism $\iota_n(L)_* : \pi_1(M) \xrightarrow{\cong} \pi_1(X_n(L))$.*
- (2) *The abelian group $H_3(X_n(L))$ contains a summand canonically isomorphic to $H_3(M) \cong \mathbb{Z}$.*

Proof Since $X_n(L)$ may be obtained from the reduced mapping cylinder $M_{p_n}^\times$ by gluing ν_L to $\partial E_L \times \{1\}$, the homomorphism $\pi / \Gamma_n \rightarrow \pi_1(X_n(L))$ is surjective and its kernel is normally generated by the images of the meridians of L in π / Γ_n . This is precisely the subgroup Γ / Γ_n , since $\Gamma \trianglelefteq \pi$ is the subgroup of π normally generated by the meridians of L . Thus,

$$\pi_1(X_n(L)) \cong (\pi / \Gamma_n) / (\Gamma / \Gamma_n) \cong \pi / \Gamma \cong \pi_1(M).$$

To see that this isomorphism is induced by the map $\iota_n(L)$, we instead view $X_n(L)$ as obtained from M by attaching cells of dimension 2 and higher. The attaching maps of the 2-cells represent elements of Γ_n , which are already nullhomotopic in M , and the higher-dimensional cells do not affect the fundamental group. Therefore, $\iota_n(L)_*$ induces an isomorphism. This proves (1).

To prove statement (2), we consider the Mayer–Vietoris sequence corresponding to the decomposition $X_n(L) = M_{p_n}^\times \cup_{\partial E_L} \nu L$. The connecting homomorphism $H_3(X_n(L)) \rightarrow H_2(\partial E_L) \cong \mathbb{Z}^m$ sends the inclusion-induced image of $[M]$ to $(1, 1, \dots, 1)$ (i.e., the image of $[M]$ is the sum of the fundamental classes of the tori which form ∂E_L). Projecting $H_2(\partial E_L)$ onto $\mathbb{Z}$ generated by $(1, 1, \dots, 1)$ gives a surjective homomorphism $H_3(X_n(L)) \twoheadrightarrow \mathbb{Z}$ which sends the image of $[M]$ to 1. This yields a splitting of $H_3(X_n(L))$ as $H_3(M) \oplus A$ for some abelian group A . As the inclusion map $\iota_n(L) : M \hookrightarrow X_n(L)$ is canonical, this splitting is canonical. This proves (2). $\square$

The spaces $X_n^G(L)$ from Section 3.4 have analogous algebraic properties:

Proposition A.2 *Let G be some group with $\varphi : \pi_1(M) \twoheadrightarrow G$.*

- (1) *Let $\Gamma^G = \ker \varphi$. Then $\pi_1(X_n^G(L)) \cong \pi_1(M) / \Gamma_n^G$, the n -th G -lower central quotient of $\pi_1(M)$, and the map $\iota_n^G(L) : M \hookrightarrow X_n^G(L)$ induces the canonical projection $\pi_1(M) \twoheadrightarrow \pi_1(M) / \Gamma_n^G$.*
- (2) *The abelian group $H_3(X_n^G(L))$ contains a summand canonically isomorphic to $H_3(M) \cong \mathbb{Z}$.*

Acknowledgements

The author is grateful to Kent Orr for his unwavering support of this project, his many insightful comments and his guidance. The author also thanks Jae Choon Cha, whose work with Kent Orr in [Cha and Orr 2024] inspired several ideas in this paper. Additionally, the author thanks Anthony Conway, Matthew Hedden, Alexandra Kjuchukova and Mark Powell for valuable conversations. Finally, the author is grateful to the referee for their careful reading of this work and for many helpful suggestions.

References

- [Audoux et al. 2021] **B Audoux, J-B Meilhan, A Yasuhara**, *Milnor-type invariants for surface-links and cut-diagrams*, preprint (2021) arXiv 2109.14578
- [Bar-Natan 1995] **D Bar-Natan**, *Vassiliev homotopy string link invariants*, *J. Knot Theory Ramifications* 4:1 (1995) 13–32 MR
- [Brown 1982] **K S Brown**, *Cohomology of groups*, Graduate Texts in Mathematics 87, Springer (1982) MR
- [Cappell and Shaneson 1974] **S E Cappell, J L Shaneson**, *The codimension two placement problem and homology equivalent manifolds*, *Ann. of Math.* (2) 99 (1974) 277–348 MR
- [Casson 1975] **A J Casson**, *Link cobordism and Milnor’s invariant*, *Bull. London Math. Soc.* 7 (1975) 39–40 MR
- [Celoria 2018] **D Celoria**, *On concordances in 3-manifolds*, *J. Topol.* 11:1 (2018) 180–200 MR
- [Cha 2018] **J C Cha**, *Rational Whitney tower filtration of links*, *Math. Ann.* 370:3-4 (2018) 963–992 MR
- [Cha and Orr 2024] **J C Cha, K E Orr**, *Transfinite Milnor invariants for 3-manifolds*, *J. Eur. Math. Soc.* 26:8 (2024) 2971–3046 MR
- [Cochran 1990] **T D Cochran**, *Derivatives of links: Milnor’s concordance invariants and Massey’s products*, *Mem. Amer. Math. Soc.* 427, Amer. Math. Soc., Providence, RI (1990) MR
- [Cochran and Harvey 2008] **T D Cochran, S Harvey**, *Homology and derived series of groups, II: Dwyer’s theorem*, *Geom. Topol.* 12:1 (2008) 199–232 MR
- [Cochran et al. 2003] **T D Cochran, K E Orr, P Teichner**, *Knot concordance, Whitney towers and L^2 -signatures*, *Ann. of Math.* (2) 157:2 (2003) 433–519 MR
- [Conant et al. 2011] **J Conant, R Schneiderman, P Teichner**, *Higher-order intersections in low-dimensional topology*, *Proc. Natl. Acad. Sci. USA* 108:20 (2011) 8131–8138 MR
- [Conant et al. 2014] **J Conant, R Schneiderman, P Teichner**, *Milnor invariants and twisted Whitney towers*, *J. Topol.* 7:1 (2014) 187–224 MR
- [Davis 2020] **C W Davis**, *Topological concordance of knots in homology spheres and the solvable filtration*, *J. Topol.* 13:1 (2020) 343–355 MR
- [Davis et al. 2018] **C W Davis, M Nagel, J Park, A Ray**, *Concordance of knots in $S^1 \times S^2$* , *J. Lond. Math. Soc.* (2) 98:1 (2018) 59–84 MR
- [Dwyer 1975] **W G Dwyer**, *Homology, Massey products and maps between groups*, *J. Pure Appl. Algebra* 6:2 (1975) 177–190 MR
- [Freedman and Teichner 1995] **M H Freedman, P Teichner**, *4-manifold topology, II: Dwyer’s filtration and surgery kernels*, *Invent. Math.* 122:3 (1995) 531–557 MR
- [Friedl et al. 2019] **S Friedl, M Nagel, P Orson, M Powell**, *Satellites and concordance of knots in 3-manifolds*, *Trans. Amer. Math. Soc.* 371:4 (2019) 2279–2306 MR
- [Gauss 1877] **C F Gauss**, *Zur mathematischen Theorie der electrodynamischen Wirkungen*, from “Carl Friedrich Gauss: Werke”, volume V, Königlichen Gesellschaft der Wissenschaften, Göttingen (1877) 601–629
- [Goldsmith 1978] **D L Goldsmith**, *A linking invariant of classical link concordance*, from “Knot theory” (Plans-sur-Bex, 1977) (J-C Hausmann, editor), *Lecture Notes in Math.* 685, Springer (1978) 135–170 MR

- [Gusarov 2000] **M N Gusarov**, *Variations of knotted graphs: the geometric technique of n -equivalence*, *Algebra i Analiz* 12:4 (2000) 79–125 MR In Russian; translated in *St. Petersburg Math. J.* 12 (2001) 569–604
- [Habegger and Masbaum 2000] **N Habegger, G Masbaum**, *The Kontsevich integral and Milnor's invariants*, *Topology* 39:6 (2000) 1253–1289 MR
- [Habiro 2000] **K Habiro**, *Claspers and finite type invariants of links*, *Geom. Topol.* 4 (2000) 1–83 MR
- [Heck 2011] **P Heck**, *Homotopy properties of knots in prime manifolds*, preprint (2011) arXiv 1110.6903
- [Hom et al. 2022] **J Hom, A S Levine, T Lidman**, *Knot concordance in homology cobordisms*, *Duke Math. J.* 171:15 (2022) 3089–3131 MR
- [Igusa 1979] **K Igusa**, *The generalized grassman invariant $K_3(\mathbb{Z}[\pi]) \rightarrow H_0(\pi; \mathbb{Z}_2[\pi])$* , unpublished manuscript (1979) Available at <https://people.brandeis.edu/~igusa/Papers/GrassInvCompr.pdf>
- [Igusa and Orr 2001] **K Igusa, K E Orr**, *Links, pictures and the homology of nilpotent groups*, *Topology* 40:6 (2001) 1125–1166 MR
- [Klug and Miller 2021] **M R Klug, M Miller**, *Concordance of surfaces in 4-manifolds and the Freedman–Quinn invariant*, *J. Topol.* 14:2 (2021) 560–586 MR
- [Klug and Miller 2022] **M Klug, M Miller**, *Concordance of spheres in 4-manifolds with an immersed dual sphere*, preprint (2022) arXiv 2211.07177
- [Krushkal 1998] **V S Krushkal**, *Additivity properties of Milnor's $\bar{\mu}$ -invariants*, *J. Knot Theory Ramifications* 7:5 (1998) 625–637 MR
- [Kuzbary 2024] **M Kuzbary**, *An analogue of Milnor's invariants for knots in 3-manifolds*, *Algebr. Geom. Topol.* 24:6 (2024) 3043–3067 MR
- [Levine 1977] **J Levine**, *Knot modules, I*, *Trans. Amer. Math. Soc.* 229 (1977) 1–50 MR
- [Levine 1989a] **J P Levine**, *Link concordance and algebraic closure, II*, *Invent. Math.* 96:3 (1989) 571–592 MR
- [Levine 1989b] **J P Levine**, *Link concordance and algebraic closure of groups*, *Comment. Math. Helv.* 64:2 (1989) 236–255 MR
- [Levine 2016] **A S Levine**, *Nonsurjective satellite operators and piecewise-linear concordance*, *Forum Math. Sigma* 4 (2016) art. id.e34 MR
- [Lin 1991] **X-S Lin**, *Null k -cobordant links in S^3* , *Comment. Math. Helv.* 66:3 (1991) 333–339 MR
- [Miller 1995] **D Miller**, *An extension of Milnor's $\bar{\mu}$ -invariants*, *Topology Appl.* 65:1 (1995) 69–82 MR
- [Miller 2021] **M Miller**, *A concordance analogue of the 4D light bulb theorem*, *Int. Math. Res. Not.* 2021:4 (2021) 2565–2587 MR
- [Milnor 1957] **J Milnor**, *Isotopy of links*, from “Algebraic geometry and topology: a symposium in honor of S. Lefschetz” (R H Fox, D C Spencer, A W Tucker, editors), *Princeton Math. Series* 12, Princeton Univ. Press (1957) 280–306 MR
- [Nagel et al. 2019] **M Nagel, P Orson, J Park, M Powell**, *Smooth and topological almost concordance*, *Int. Math. Res. Not.* 2019:23 (2019) 7324–7355 MR
- [Orr 1989] **K E Orr**, *Homotopy invariants of links*, *Invent. Math.* 95:2 (1989) 379–394 MR
- [Porter 1980] **R Porter**, *Milnor's $\bar{\mu}$ -invariants and Massey products*, *Trans. Amer. Math. Soc.* 257:1 (1980) 39–71 MR
- [Raoux 2020] **K Raoux**, *τ -invariants for knots in rational homology spheres*, *Algebr. Geom. Topol.* 20:4 (2020) 1601–1640 MR
- [Rotman 2009] **J J Rotman**, *An introduction to homological algebra*, 2nd edition, Springer (2009) MR
- [Schneiderman 2003] **R Schneiderman**, *Algebraic linking numbers of knots in 3-manifolds*, *Algebr. Geom. Topol.* 3 (2003) 921–968 MR
- [Schneiderman and Teichner 2022] **R Schneiderman, P Teichner**, *Homotopy versus isotopy: spheres with duals in 4-manifolds*, *Duke Math. J.* 171:2 (2022) 273–325 MR
- [Schwartz 2019] **H R Schwartz**, *Equivalent non-isotopic spheres in 4-manifolds*, *J. Topol.* 12:4 (2019) 1396–1412 MR
- [Spanier 1995] **E H Spanier**, *Algebraic topology*, corrected reprint of 1st edition, Springer (1995) MR

- [Stallings 1965] **J Stallings**, *Homology and central series of groups*, J. Algebra 2 (1965) 170–181 MR
- [Steas 2023] **R Steas**, *Milnor’s invariants for knots and links in closed orientable 3-manifolds*, PhD thesis, Indiana University (2023) Available at <https://www.proquest.com/docview/2829577776>
- [Steas 2025] **R Steas**, *Almost-concordance of knots in aspherical 3-manifolds*, preprint (2025) arXiv 2508.14638
- [Sunukjian 2015] **NS Sunukjian**, *Surfaces in 4-manifolds: concordance, isotopy, and surgery*, Int. Math. Res. Not. 2015:17 (2015) 7950–7978 MR
- [Turaev 1976] **V G Turaev**, *The Milnor invariants and Massey products*, Zap. Nauch. Sem. Leningrad. Otdel. Mat. Inst. Steklov. 66 (1976) 189–203, 209–210 MR In Russian; translated in J. Soviet Math. 12 (1979) 128–137
- [Turaev 1984] **V G Turaev**, *Nilpotent homotopy types of closed 3-manifolds*, from “Topology” (Leningrad, 1982) (L D Faddeev, A A Mal’cev, editors), Lecture Notes in Math. 1060, Springer (1984) 355–366 MR
- [Vogel 1978] **P Vogel**, *Localization of spaces with respect to a class of maps*, preprint, Univ. de Nantes (1978) Available at <https://webhomes.maths.ed.ac.uk/~v1ranick/papers/vogelnantes.pdf>
- [Yildiz 2018] **E Z Yildiz**, *A note on knot concordance*, Algebr. Geom. Topol. 18:5 (2018) 3119–3128 MR
- [Zhou 2021] **H Zhou**, *Homology concordance and an infinite rank free subgroup*, J. Topol. 14:4 (2021) 1369–1395 MR

RYAN STEES rs2sf@virginia.edu

Department of Mathematics, University of Virginia, Charlottesville, VA, United States

Proposed: Cameron Gordon

Received: March 12, 2024

Seconded: Robion Kirby, David Gabai

Revised: February 21, 2025

Chromatic defect, Wood’s theorem and higher real K-theories

CHRISTIAN CARRICK

Using Ravenel’s Thom spectrum $X(n)$, we introduce the concept of chromatic defect, which measures how far a spectrum is from being complex-orientable. We compute the chromatic defect of various examples of interest, such as finite spectra, the Real Johnson–Wilson spectra $ER(n)$, fixed points of Morava E -theories (with respect to finite subgroups of the Morava stabilizer group), and the connective image of J spectrum. Moreover, an obstruction theory is developed for determining chromatic defect. Having finite chromatic defect is closely related to the existence of analogues of the classical Wood equivalence. We show that such equivalences exist in a wide generality and use them to construct $\mathbb{Z}$ -indexed Adams–Novikov towers.

1. Introduction	1051
2. Orientability and Ravenel’s $X(n)$	1060
3. The $T(n)$ and $\mathbb{F}_p$ -homology	1068
4. Chromatic defect and Wood equivalences for fp spectra	1079
5. Hopkins’ stacks and the $X(n)$	1083
6. Real-oriented theories and higher real K -theories	1092
7. $\mathbb{Z}$ -indexed Adams–Novikov spectral sequences	1097
Acknowledgements	1105
References	1105

1 Introduction

The chromatic approach to stable homotopy uses the complex cobordism spectrum MU to detect phenomena in stable homotopy. The MU -homology MU_*E of a spectrum E is often much more computable than its stable homotopy π_*E . However, the former is the input to the Adams–Novikov spectral sequence, which can be used to recover information about the latter. This is an extremely powerful tool that detects nilpotence and periodicity phenomena, as conjectured by Ravenel [52] and proven by Devinatz, Hopkins and Smith [19]. When E is complex-orientable, however, there is essentially no difference between π_*E and MU_*E ; the former is recovered as the comodule primitives in the latter, and thus E has no complexity from the point of view of chromatic homotopy theory.

There are many useful spectra E that are not complex-orientable, but are in some sense only finitely many steps away from being complex-orientable. Their complexity is thus small from the point of view of chromatic homotopy, which is reflected in the computability of their Adams–Novikov spectral sequence.

MSC2020: 55N22, 55P42, 55P91, 55P92, 55T15.

© 2026 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

The clearest example of this is ko , connective real K -theory. The classical Wood's theorem [56] states that there is an equivalence

$$ko \otimes C(\eta) \simeq ku,$$

where $C(\eta)$ is the cofiber of the Hopf map $\eta : \mathbb{S}^1 \rightarrow \mathbb{S}^0$, and ku is connective complex K -theory. The spectrum ku is complex-orientable and ko is not. However, $C(\eta)$ is a 2-cell complex with chromatic type zero, in the sense of the thick subcategory theorem; see [52].

Wood's theorem is useful for calculations involving ko , and it also implies that ko is in the thick tensor ideal generated by MU , so that for example the Adams–Novikov spectral sequence of ko has a horizontal vanishing line on a finite page and ko satisfies the conditions of the telescope conjecture at all heights. We study analogues of Wood's theorem and say that a spectrum E is *Wood-type* if there is a finite BP-projective F such that $E \otimes F$ is complex-orientable (Definition 2.30).

Chromatic defect Using Bott periodicity, Ravenel [51] introduced a filtration

$$\mathbb{S} = X(1) \rightarrow \cdots \rightarrow X(n) \rightarrow \cdots \rightarrow X(\infty) = MU$$

of MU , where $X(n)$ is a certain Thom spectrum over $\Omega\mathrm{SU}(n)$. The $X(n)$ provide a convenient interpolation between stable homotopy and MU -homology, and quite a lot is known about this filtration. In fact, the Devinatz–Hopkins–Smith proof [19] of the nilpotence theorem proceeds by induction downward along this filtration: they show that each of the maps $X(n) \rightarrow X(n+1)$ detects nilpotence, then by compactness conclude that $\mathbb{S} \rightarrow MU$ detects nilpotence.

Recovering stable homotopy from MU -homology entails infinite descent along the Ravenel filtration. In general, calculating $X(n)$ -homology is much more difficult than calculating MU -homology. Typically, calculating $X(n)$ -homology is as difficult as calculating stable homotopy as, in some sense, the $X(n)$ are too close to the sphere spectrum.

We study theories E with the property that $E \otimes X(n)$ is complex-orientable for some $n < \infty$, and we say in this case that E has *chromatic defect* $\leq n$; see Definition 2.15. When E has chromatic defect $\leq n$, the role of MU -homology may be replaced by that of $X(n)$ -homology, and thus the stable homotopy of E may be recovered by a *finite* descent along Ravenel's filtration. As E has chromatic defect 1 if and only if E is itself complex-orientable, chromatic defect measures the failure of E to be complex-orientable in a precise sense. The condition of E having finite chromatic defect is closely related to that of E being Wood-type; for example ko has chromatic defect 2. In fact, finite chromatic defect is a necessary condition for being Wood-type; see Corollary 2.35.

Ravenel [51, Section 7.1] has studied extensively the process of recovering stable homotopy by descent along his filtration from a computational point of view. In fact, in an unpublished work, Hill, Hopkins and Ravenel [24] use this filtration to compute the homotopy groups of the fixed points of Morava E -theories with respect to a natural C_p -action at chromatic height $n = 2(p-1)$. We compute chromatic defect in several cases of interest, which determines precisely the stage of Ravenel's filtration at which these descent computations must begin, and we focus on more general fixed-point spectra of Morava E -theories with respect to finite subgroups of the Morava stabilizer group.

$K(n)$ -local homotopy and the Morava stabilizer group Let $K(n)$ be the n -th Morava K -theory. By a theorem of Devinatz and Hopkins [18], the $K(n)$ -local sphere can be described as the homotopy fixed-point spectrum

$$L_{K(n)}\mathbb{S} \simeq E(k, \Gamma)^{h\mathbb{G}_n},$$

where $E(k, \Gamma)$ is a height n Morava E -theory, for Γ a height n formal group over a perfect field k of characteristic p , and $\mathbb{G}_n = \text{Aut}(\Gamma) \rtimes \text{Gal}(k/\mathbb{F}_p)$ the corresponding Morava stabilizer group. It was an observation of Ravenel [50] and Hopkins and Miller [53] that $L_{K(n)}\mathbb{S}$ is approximated by the fixed points of $E(k, \Gamma)$ at *finite* subgroups G of $\mathbb{G}_n$, and that these theories are more computable. These theories

$$(1-1) \quad \text{EO}_n(G) := E(k, \Gamma)^{hG}$$

are known as the *Hopkins–Miller higher real K -theories*.

We compute the chromatic defect of higher real K -theories using a construction of Hopkins, which associates a stack $\mathcal{M}_E$ to any homotopy commutative ring spectrum E . The stack $\mathcal{M}_E$ is the one corresponding to the Hopf algebroid

$$(\text{MU}_{2*}E, \text{MU}_{2*}(\text{MU} \otimes E)).$$

Via the unit map $\mathbb{S} \rightarrow E$, the stack $\mathcal{M}_E$ comes equipped with a canonical affine morphism $p_E : \mathcal{M}_E \rightarrow \mathcal{M}_{\text{FG}}(1)$ to the moduli stack of formal groups and strict automorphisms. This stack allows one to describe the E_2 -page of the Adams–Novikov spectral sequence of E from an algebrogeometric point of view, and it is the most essential tool in our analysis of chromatic defect. Indeed, applying Hopkins' stack construction to Ravenel's filtration produces the tower

$$\text{Spec}(L) \rightarrow \cdots \rightarrow \mathcal{M}_{\text{FG}}(n) \rightarrow \cdots \rightarrow \mathcal{M}_{\text{FG}}(1),$$

where L is the Lazard ring, and $\mathcal{M}_{\text{FG}}(n)$ is the moduli stack of formal group laws equipped with a coordinate through degree n . This often allows one to turn the question of whether E has finite chromatic defect into a rigidity question about $\mathcal{M}_E$: namely, whether the automorphism groups of objects in $\mathcal{M}_E$ become trivial after fixing coordinates on the associated formal groups through some finite degree.

1.1 Main results and outline of the paper

Section 2 In the classical Wood equivalence $\text{ko} \otimes C(\eta) \simeq ku$, a priori the left-hand side has no ring structure, as $C(\eta)$ does not admit a unital multiplication. Our definition of Wood-types thus requires a more flexible definition of complex-orientability for a spectrum E that does not require that E be a ring spectrum. Therefore, we make the following definition.

Definition 1.1 (Definition 2.3) A spectrum E is complex-orientable if the map

$$E \otimes \sigma_k : E \otimes \mathbb{S}^{2k+1} \rightarrow E \otimes \mathbb{C}\mathbb{P}^k$$

is nullhomotopic for all $k \geq 1$, where σ_k is the attaching map for the top cell in $\mathbb{C}\mathbb{P}^{k+1}$.

Remark 1.2 In the literature, complex orientations are defined only for ring spectra. We say a ring spectrum E is complex-orientable if the tautological line bundle over $\mathbb{C}\mathbb{P}^\infty$ is orientable with respect to E . In Section 2, we verify that this definition of complex-orientability is equivalent to Definition 1.1 for any homotopy associative ring spectrum.

Quillen [49] showed that a homotopy commutative ring spectrum E is complex-orientable if and only if there exists a homotopy ring map $\mathrm{MU} \rightarrow E$. We extend this result to complex-orientable spectra (in the sense of Definition 1.1) in the following way.

Theorem 1.3 (Theorem 2.9) *A spectrum E is complex-orientable if and only if it is a weak MU -module; that is, the unit map $E \rightarrow \mathrm{MU} \otimes E$ admits a retraction.*

This flexible definition of complex-orientability results in a straightforward definition of Wood-types and chromatic defect that makes sense in complete generality. In what follows, we work p -locally at a fixed prime p . We say a finite spectrum F is a finite BP-projective if BP_*F is a projective BP_* -module.

Definition 1.4 (Definitions 2.30 and 2.15) We say a spectrum E is *Wood-type* if there exists a finite BP-projective F such that $E \otimes F$ is complex-orientable.

Definition 1.5 (Definition 2.15) We say a spectrum E has *chromatic defect* $\leq n$ if $E \otimes X(n)$ is complex-orientable. In this case, we use the notation $\Phi(E) \leq n$.

We tie the notion of Wood-types to that of chromatic defect with the following result.

Theorem 1.6 (Corollary 2.35) *If E is Wood-type, then it has finite chromatic defect.*

It can be difficult to say whether a certain spectrum E is Wood-type since this requires producing a complex F such that $E \otimes F$ is complex-orientable. We will see, however, that computing chromatic defect is tractable in many cases, and when it is infinite for example, this rules out the possibility of being Wood-type; see Corollaries 3.15 and 5.19, for example.

Section 3 The mod p homology $H_*(X(n); \mathbb{F}_p)$ is determined as an $\mathcal{A}_*$ -comodule algebra via the Thom isomorphism. We use this to describe the Adams E_2 -page of $X(n)$ and give a May spectral sequence converging to it. This recovers known results about the homotopy groups of $X(n)$, such as that the bottom odd-dimensional homotopy group of $X(n)_{(p)}$ is

$$\pi_{2p^{m+1}-3}X(n)_{(p)} = \mathbb{Z}/p\{\chi_{m+1}\},$$

where $m = \lfloor \log_p(n) \rfloor$; see also [7] for this result. The classes χ_{m+1} are shown to give a sequence of obstructions for a homotopy associative ring spectrum E to have chromatic defect $\leq n$.

Theorem 1.7 (Proposition 3.7) *Let E be a p -local homotopy associative ring spectrum. Then $\Phi(E) \leq n$ if and only if*

$$\chi_{m+1} = 0 \in \pi_{2p^{m+1}-3}E \otimes X(p^m) \quad \text{for all } m \geq \lfloor \log_p(n) \rfloor.$$

We calculate the relative dual Steenrod algebra $\mathcal{A}^{X(n)*} := \pi_*(\mathbb{F}_p \otimes_{X(n)} \mathbb{F}_p)$, and this allows us to show that finite spectra are too close to the sphere to have finite chromatic defect.

Corollary 1.8 (Corollary 3.15) *If F is a finite spectrum, then $\Phi(F) = \infty$.*

Set $m := \lfloor \log_p(n) \rfloor$. We show that, when the $\mathcal{A}_*$ -comodule $H_*(X(n); \mathbb{F}_p)$ is restricted to $\mathcal{A}(m)_*$, it splits as a sum of finite even comodules, all of which are isomorphic to $\mathcal{P}(m-1)_*$; see Section 3 for these Hopf algebras. With an eye toward Wood equivalences, we thus construct finite complexes F that are free over $\mathcal{P}(m-1)$. This uses the idempotent construction of Jeff Smith [52], used to construct finite $\mathcal{A}(m)$ -free type $m+1$ complexes, and our proof is essentially the same.

Section 4 The spectra ko and tmf are our prototypical examples of Wood-types, the former via $C(\eta)$ in the classical Wood equivalence, and the latter (at the prime 2) via the complex $D\mathcal{A}(1)$; see [39]. These are also prototypical examples of fp spectra in the sense of Mahowald and Rezk [37], that is, a p -complete bounded below spectrum E that admits an isomorphism

$$H_*(E; \mathbb{F}_p) \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$$

for some $n < \infty$ and M a finite $\mathcal{A}(n)_*$ -comodule.

We show that the Wood equivalences for ko and tmf actually follow from certain evenness conditions that hold for a much broader class of fp spectra. This is obtained by using the Adams spectral sequence to sharpen the obstruction statement of Theorem 1.7 to an algebraic context. We use the $\mathcal{P}(n)$ -free complexes of Section 3 to obtain the following. Here, we use the following standard notation for certain exterior quotient Hopf algebras of the dual Steenrod algebra:

$$\mathcal{E}(n)_* := \begin{cases} E(\xi_1, \dots, \xi_{n+1}) & \text{if } p = 2, \\ E(\tau_0, \dots, \tau_n) & \text{if } p > 2. \end{cases}$$

Theorem 1.9 (Corollary 4.10) *Let E be a homotopy associative ring spectrum that is an fp spectrum, so that $H_*E \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$, and suppose that*

$$\mathrm{Ext}_{\mathcal{E}(n)_*}^{s,t}(\mathbb{F}_p, M)$$

is concentrated in even stems $t - s$. Then $\Phi(E) \leq p^n$ and E is Wood-type.

Theorem 1.9 is useful because computing Ext over the exterior Hopf algebra $\mathcal{E}(n)_*$ is much more tractable than over $\mathcal{A}(n)_*$. We actually prove more specific statements by tracking the degrees of the obstructions χ_{m+1} . When E as above is assumed to have finite chromatic defect, we can prove a partial converse to Theorem 1.6 that does not necessarily require evenness conditions.

Notation 1.10 For a spectrum X , we let $\mathrm{ASS}(X)$ and $\mathrm{ANSS}(X)$ denote the Adams and Adams–Novikov spectral sequences of X , respectively.

Theorem 1.11 (Theorem 4.9) *Let E be a homotopy associative ring spectrum that is an fp spectrum, and suppose E has finite chromatic defect. If $\mathrm{ASS}(E \otimes \mathrm{BP})$ collapses on the E_2 -page, then E is Wood-type.*

Giving ourselves the mod 2 homology of ko and tmf , this gives a new proof of the following result of Hopkins [20].

Corollary 1.12 (Example 4.11) *The spectra ko and tmf are Wood-type. The chromatic defect of ko is 2, and the chromatic defect of tmf is 4.*

Remark 1.13 The spectra ko and tmf are good connective models of certain Hopkins–Miller higher real K -theories; see [15, Section 1.1] for a discussion of this. In [9], Beaudry–Hill–Shi–Zeng construct good connective models of the spectra $\mathrm{EO}_{2^{n-1}m}(C_{2^n})$ known as the $\mathrm{BP}^{\langle G \rangle}(m)$. In work with Mike Hill [14], we show that these are fp spectra. The determination of their chromatic defect can thus be approached via the methods of Section 4, and we intend to return to this in future work.

Section 5 When E has finite chromatic defect, the E_2 -page of its Adams–Novikov spectral sequence simplifies substantially. For example, the E_2 -page in general is determined by cohomology over the group scheme $\mathrm{Spec}(\mathbb{Z}[b_1, b_2, \dots])$ corepresenting power series $x + b_1x^2 + \dots$. However, when E has chromatic defect $\leq n$, the E_2 -page is determined by cohomology over the subgroup scheme $\mathrm{Spec}(\mathbb{Z}[b_1, b_2, \dots, b_{n-1}])$, which has finite Krull dimension. Hopkins’ stacks framework allows for these sorts of change-of-rings isomorphisms to take a more conceptual form. We use this to verify our “finite subgroups” philosophy on chromatic defect in the following way.

Proposition 1.14 (Proposition 5.9) *Let E be a homotopy commutative ring spectrum with finite chromatic defect. Then, for any algebraically closed field k and any $x \in \mathcal{M}_E(k)$ such that $p_E(x) \in \mathcal{M}_{\mathrm{FG}}(1)(k)$ has finite height, the image of the homomorphism*

$$\psi : \mathrm{Aut}_{\mathcal{M}_E(k)}(x) \rightarrow \mathrm{Aut}_{\mathcal{M}_{\mathrm{FG}}(1)(k)}(p_E(x))$$

is a finite subgroup.

In other words, at finite-height geometric points, the morphism $p_E : \mathcal{M}_E \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ lands in finite subgroups. We provide various converses to this when $\mathcal{M}_E$ has certain strong finiteness conditions. These conditions hold for example for the $\mathrm{EO}_n(G)$ and also for ko and tmf . In the following, we suppose $E \otimes X(n)$ is MU nilpotent-complete for all n sufficiently large, for example if E is connective or MU-nilpotent.

Proposition 1.15 (Proposition 5.13) *Let E be a homotopy commutative ring spectrum, and suppose there is a faithfully flat, finite morphism $\mathrm{Spec}(R) \rightarrow \mathcal{M}_E$, for some Noetherian commutative ring R . Then E has finite chromatic defect.*

Finally, we move our stacks analysis to the $K(n)$ -local category to compute the chromatic defect of the $K(n)$ -local sphere.

Theorem 1.16 (Theorem 5.16) *The spectrum $L_{K(n)}\mathbb{S}$ does not have finite chromatic defect.*

Remark 1.17 Let j be the connective image of J spectrum. Since

$$L_{K(1)}j \simeq L_{K(1)}\mathbb{S},$$

it follows also that j has infinite chromatic defect; see Theorem 5.18. This is an important example

because j is an fp spectrum, and in fact MU-nilpotent, but it does not have finite chromatic defect and thus is not Wood-type.

Section 6 We study chromatic defect in two examples where a ring spectrum E admits an equivalence $\mathcal{M}_E \simeq \mathrm{Spec}(R)/G$ to a quotient stack by a finite group. Working at the prime 2, let $E(n)$ denote the height- n Johnson–Wilson theory. The spectrum $E(n)$ admits a natural lift to a C_2 -spectrum called $E_{\mathbb{R}}(n)$, constructed by Hu and Kriz [30], where C_2 acts by complex-conjugation. The fixed-point spectra $\mathrm{ER}(n) := E_{\mathbb{R}}(n)^{hC_2}$ have been studied extensively by Kitchloo and Wilson [34; 33]; they used these theories to prove new nonimmersion results for real projective spaces.

Theorem 1.18 (Theorem 6.5) *The chromatic defect of $\mathrm{ER}(n)$ is 2^n .*

This generalizes the result of Hopkins at height 1 [20], which states that $\mathrm{KO} \otimes X(2)$ is complex-orientable, and it is a result of Atiyah [4] that $\mathrm{ER}(1) \simeq \mathrm{KO}$. At higher heights, some care is needed since $\mathrm{ER}(n)$ is not known to admit a ring structure; see [32]. To compute the chromatic defect of $\mathrm{ER}(n)$, we thus prove the following, which may be of independent interest.

Theorem 1.19 (Theorem 6.5) *For all $m \geq 2^n$, the spectrum $\mathrm{ER}(n) \otimes X(m)$ has a (Landweber exact) homotopy commutative ring structure with respect to which the restriction map $\mathrm{ER}(n) \otimes X(m) \rightarrow E(n) \otimes X(m)$ is one of ring spectra.*

We turn to the Hopkins–Miller theories. Fixing as before a height n formal group Γ over a perfect field k of characteristic p , we let G be a finite subgroup of $\mathbb{G}_n = \mathrm{Gal}(k/\mathbb{F}_p) \rtimes \mathrm{Aut}(\Gamma)$ the corresponding Morava stabilizer group. By the Goerss–Hopkins–Miller theorem [22], G acts on $E(k, \Gamma)$ by $\mathbb{E}_{\infty}$ -ring maps. The chromatic defect of the fixed points $\mathrm{EO}_n(G)$ will be expressed in terms of the standard valuation on the endomorphism ring $\mathrm{End}(\Gamma)$, which we recall here. Every nonzero endomorphism $f(x)$ of Γ may be expressed uniquely as $f(x) = g(x^{p^k})$, where $g(x) \in \mathrm{End}(\Gamma)^{\times}$ is invertible. The valuation v is then defined by $v(f) = k/n$, normalized by the height n , so that $p \in \mathrm{End}(\Gamma)$ is a uniformizer. In the statement below, we let $\pi : \mathbb{G}_n \rightarrow \mathrm{Aut}(\Gamma)$ denote the projection map, where we identify the underlying set of the semidirect product as a cartesian product.

Theorem 1.20 (Theorem 6.7) *Let $N(G) := n \cdot \max\{v(\pi(g) - 1) : e \neq \pi(g)\}_{g \in G}$, where e is the identity element of G . The chromatic defect of $\mathrm{EO}_n(G)$ is $p^{N(G)}$.*

When the height h is of the form $p^{n-1}(p-1)m$, there is a C_{p^n} -subgroup of $\mathbb{G}_h$ (see Section 6 for more details), and we have the following in this case.

Example 1.21 (Corollaries 6.8 and 6.9) *The chromatic defect of $\mathrm{EO}_{p^{n-1}(p-1)m}(C_{p^n})$ is $p^{p^{n-1}m}$.*

To apply Ravenel's method of descent to compute $\pi_* \mathrm{EO}_n(G)$ as in [24], for example, the above theorem implies that the computation begins by tensoring with $X(p^{N(G)})$.

Remark 1.22 In fact, our methods also imply that the MU-homology of the $\mathrm{ER}(n)$ and the $\mathrm{EO}_n(G)$ is even and torsion-free. This implies these theories are *quasi-syntomic* in the sense of Hahn, Raksit and Wilson [23].

Section 7 One can conceive of variations on our definition of Wood-types E , for example by asking the complex F to be just a type zero instead of a BP-projective, or simply the condition of E being BP-nilpotent. Neither of these imply finite chromatic defect, but they do guarantee strong conditions on the Adams–Novikov spectral sequence of E , such as a horizontal vanishing line on a finite page. The Adams–Novikov spectral sequence of a Wood-type, however, is further restricted by way of a spectral sequence we introduce called the $\mathbb{Z}$ -indexed Adams–Novikov spectral sequence of a Wood-type. The relationship between the ANSS and the $\mathbb{Z}$ -ANSS of a Wood-type is very similar to that of the HFPSS and the Tate SS for a G -spectrum E with vanishing Tate spectrum E^{tG} .

Mahowald–Rezk introduced a $\mathbb{Z}$ -indexed Adams spectral sequence for an fp spectrum E , which extends its Adams spectral sequence to a full plane spectral sequence converging to $\pi_* L_n^f E$ when E is fp type n . Our definition of Wood-types is chosen in part to make their construction work instead with the Adams–Novikov spectral sequence. Rather than $L_n^f E$, the $\mathbb{Z}$ -ANSS of a Wood-type always converges to zero.

Theorem 1.23 (Theorem 7.9) *The $\mathbb{Z}$ -ANSS of a Wood-type E has the following properties:*

- (1) *The $\mathbb{Z}$ -ANSS is independent of the choice of finite BP-projective F from the E_2 -page on.*
- (2) *The natural map*

$$\text{ANSS}(E) \rightarrow \mathbb{Z}\text{-ANSS}(E)$$

is an isomorphism on E_2 -pages in positive filtrations and an epi in filtration zero.

- (3) *There is a one-to-one correspondence along the map*

$$\text{ANSS}(E) \rightarrow \mathbb{Z}\text{-ANSS}(E)$$

of differentials whose source is in nonnegative filtrations.

- (4) *The $\mathbb{Z}$ -ANSS converges to zero.*

This can be very useful for example for determining differentials and vanishing lines in the ANSS of a Wood-type; cf. [21]. We run the ANSS and $\mathbb{Z}$ -ANSS in detail for ko and use this to deduce the famous d_3 therein. We also give various descriptions of the E_2 -page of the $\mathbb{Z}$ -ANSS of a Wood-type in terms of Tate cohomology in BP_*BP -comodules.

1.2 Some numerology and questions

For a p -local homotopy commutative ring spectrum E , there are many numerical measures of E coming from chromatic homotopy that are often related. One has, for instance:

- (1) The chromatic height of E :

$$\text{ht}(E) = \min\{n \geq 0 : E \simeq L_n E\}.$$

- (2) The BP-nilpotence exponent of E ; see [41].

- (3) The chromatic defect $\Phi(E)$ of E .
- (4) For a Wood-type E , the minimum $n \geq 1$ such that there exists a finite $X(n)$ -projective F such that $E \otimes F$ is complex-orientable.
- (5) The orientation order of Bhattacharya and Chatham [11]:

$$\Theta(E) = \min\{n \geq 1 : \xi^{\oplus n} \text{ is } E\text{-orientable}\},$$

where ξ is the tautological line bundle on $\mathbb{CP}^\infty$.

At first glance, chromatic height does not fit so well in this list: if E is complex-orientable, then all the quantities (2)–(5) are equal to 1 and thus are not sensitive to the chromatic height of E . However, when E is the fixed points of a complex-oriented theory R , the quantities (2)–(5) will often be functions of the chromatic height of R , as in our Theorems 6.7 and 6.5 for example. In fact, our results suggest a heuristic that, in many cases, the chromatic defect of R^{hG} should be $\leq p^{\text{ht}(R)/(p-1)}$.

For ko and tmf at the prime 2, it can be shown that all the quantities (3)–(5) coincide and they are a lower bound for (2). In general, when E is a Wood-type, (4) is an upper bound for (3) by the proof of Corollary 2.35.

Question 1.24 Suppose E has finite chromatic defect. Is $\Phi(E)$ a lower bound for the BP-nilpotence exponent of E ?

Question 1.25 Does there exist a spectrum with finite chromatic defect that is not Wood type?

When E does not have finite chromatic defect, and is therefore not Wood-type, the quantities above don't seem to be closely related. For instance, $\Phi(j) = \infty$, but j is BP-nilpotent and therefore has finite BP-nilpotence exponent.

There are of course many other numerical quantities one may assign to E . One may speculate that various chromatic height-shifting phenomena also result in shifts of chromatic defect in certain cases.

Question 1.26 How does chromatic defect interact with chromatic redshift?

For instance, the computations of Angelini-Knoll, Ausoni and Rognes in [2] suggest that while $\Phi(\text{ko}) = 2$, it may be the case that $\Phi(K(\text{ko})) = 4$. Strong forms of the Ausoni–Rognes chromatic redshift conjecture state that the algebraic K-theory of an $\mathbb{E}_1$ -ring spectrum of fp type n is fp of type $n + 1$; the methods in Section 4 may then be useful for this question.

We finish by remarking that the quantities (3) and (5) coincide for $\text{ER}(n)$, by our Theorem 6.5 and work of Kitchloo and Wilson [34]. Tying these quantities together for EO_n -theories would allow for our Theorem 6.7 to shed light on conjectures of Bhattacharya and Chatham on the orientation orders of EO_n -theories; see for example [11, Conjecture 1.13]. The author investigates this in work in progress with Prasit Bhattacharya and Yang Hu.

2 Orientability and Ravenel's $X(n)$

In this section, we revisit the notion of a complex-orientable spectrum and extend the usual definition for ring spectra to all spectra. The resulting class of complex-orientable spectra is in particular closed under tensoring with an arbitrary spectrum and under taking retracts. This gives us flexible notions of Wood-types and chromatic defect.

2.1 Complex-orientability

Definition 2.1 Let σ_k be the map

$$S^{2k+1} \rightarrow S^{2k+1}/S^1 \cong \mathbb{CP}^k,$$

where $S^{2k+1} \subset \mathbb{C}^{k+1}$ is given its usual $S^1 \subset \mathbb{C}^\times$ -action.

Remark 2.2 The cofiber of σ_k is $\mathbb{CP}^{k+1}$. In other words, σ_k is the attaching map for the top cell in $\mathbb{CP}^{k+1}$. For simplicity, since we work stably, we will not distinguish between σ_k and its double desuspension $\Sigma^{-2}\sigma_k : S^{2k-1} \rightarrow \Sigma^{-2}\mathbb{CP}^k$.

Definition 2.3 We say a spectrum E is complex-orientable if the map

$$E \otimes \sigma_k : E \otimes S^{2k+1} \rightarrow E \otimes \mathbb{CP}^k$$

is nullhomotopic for all $k \geq 1$.

As we will see, when E is a homotopy associative ring spectrum, this is equivalent to the classical definition of complex orientability, as in [51]. The above definition is much more flexible, however, as it does not require E to be a ring spectrum. Moreover, it is closed under taking retracts and tensoring with an arbitrary spectrum. To be in line with this flexibility, we will work with weak forms of modules in the homotopy category, so we take care now to define our terms.

Definition 2.4 We say E is a *homotopy associative/commutative ring spectrum* if it is an associative/commutative monoid in the symmetric monoidal 1-category $\mathrm{Ho}(\mathrm{Sp})$. A (left/right) *homotopy E -module* is a (left/right) E -module in $\mathrm{Ho}(\mathrm{Sp})$.

Definition 2.5 An $\mathbb{E}_0$ -algebra in Sp is the data of a map of spectra $\eta_E : S \rightarrow E$. For an $\mathbb{E}_0$ -algebra E , a *weak E -module* is a spectrum M such that there exists $m : E \otimes M \rightarrow M$ making the diagram

$$\begin{array}{ccc} M & \xrightarrow{\eta_E} & E \otimes M \\ & \searrow & \downarrow m \\ & & M \end{array}$$

commute up to homotopy.

Remark 2.6 M is a weak left E -module if and only if it is a weak right E -module, hence we will just say weak E -module.

Proposition 2.7 *If E is an $\mathbb{E}_0$ -algebra that is complex-orientable in the sense of Definition 2.3, then the unit map $\eta_E : \mathbb{S} \rightarrow E$ extends to a map $\Sigma^{-2}\mathbb{CP}^\infty \rightarrow E$. Conversely, if E is a homotopy associative ring spectrum, and η_E extends over $\Sigma^{-2}\mathbb{CP}^\infty$, then E is complex-orientable.*

Proof If $E \otimes \sigma_k = 0$ for all $k \geq 1$, then there is a splitting

$$E \otimes \Sigma^{-2}\mathbb{CP}^\infty \simeq \bigoplus_{k \geq 0} \Sigma^{2k} E$$

indexed by the cells $\{b_0, b_1, \dots\}$ of $\Sigma^{-2}\mathbb{CP}^\infty$. The composite

$$\Sigma^{-2}\mathbb{CP}^\infty \xrightarrow{\eta_E} E \otimes \Sigma^{-2}\mathbb{CP}^\infty \simeq \bigoplus_{k \geq 0} \Sigma^{2k} E \rightarrow E$$

extends η_E , where the last map is projection onto b_0 .

Conversely, if the unit map of E extends over $\Sigma^{-2}\mathbb{CP}^\infty$, and E is a homotopy associative ring spectrum, then the corresponding class $x \in \pi_{-2}F(\mathbb{CP}_+^k, E)$ defines a map

$$\bigoplus_{0 \leq n \leq k} \Sigma^{-2n} E \rightarrow F(\mathbb{CP}_+^k, E)$$

which, on the n -th component, is given by

$$\Sigma^{-2n} E \xrightarrow{x^n} E \otimes F(\mathbb{CP}_+^k, E) \rightarrow F(\mathbb{CP}_+^k, E),$$

using the multiplication on E to define both maps in the sequence. Filtering the left-hand side in n and the right-hand side via the cellular filtration of $\mathbb{CP}^k$, it suffices to show the given map induces an equivalence upon taking associated graded. This follows from the fact that $H^*(\mathbb{CP}^k; \mathbb{Z}) \cong \mathbb{Z}[x]/x^{k+1}$; see for example [36, Lecture 4, Proposition 7]. This implies in particular that $E \otimes \mathbb{CP}^{k-1} \rightarrow E \otimes \mathbb{CP}^k$ admits a splitting, so that $E \otimes \sigma_{k-1}$ is null. $\square$

It is a classical result of Quillen that a homotopy commutative ring spectrum E is complex-orientable if and only if there is a homotopy ring map $\text{MU} \rightarrow E$. The proof breaks down if E is not homotopy commutative as the rings $E^*(\text{BU}(n))$ will not necessarily be commutative. This prevents one from constructing a system of Thom classes $u_n \in E^{2n}(\text{MU}(n))$ for the universal bundles that is compatible with the tensor product of bundles. We have the following replacement in the associative case.

Proposition 2.8 *Let E be a complex-orientable homotopy associative ring spectrum. Then E is a weak MU-module.*

Proof The differentials in the Atiyah–Hirzebruch spectral sequence (AHSS)

$$E_2 = H_*(\Sigma^{-2}\mathbb{CP}^\infty; E_*) \Rightarrow E_*(\Sigma^{-2}\mathbb{CP}^\infty)$$

come from the boundary maps in the cellular filtration for $\mathbb{CP}^\infty$. Definition 2.3 assumes that these become zero after tensoring with E , so this AHSS collapses on E_2 , and the universal coefficients theorem computes the E_2 -page

$$H_*(\Sigma^{-2}\mathbb{CP}^\infty; E_*) \cong E_*\{b_0, b_1, \dots\}$$

as a left E_* -module.

The map $\Sigma^{-2}\mathbb{CP}^\infty \rightarrow \text{MU}$ induces a map of AHSS's, which takes the form

$$E_* \otimes H_*(\Sigma^{-2}\mathbb{CP}^\infty; \mathbb{Z}) \rightarrow E_* \otimes H_*(\text{MU}; \mathbb{Z})$$

on the E_2 page, as maps of left E_* -modules. The map $H_*(\Sigma^{-2}\mathbb{CP}^\infty; \mathbb{Z}) \rightarrow H_*(\text{MU}; \mathbb{Z})$ exhibits the latter as

$$H_*(\text{MU}; \mathbb{Z}) = \text{Sym}(H_*(\Sigma^{-2}\mathbb{CP}^\infty; \mathbb{Z})) / (b_0 - 1) = \mathbb{Z}[b_1, b_2, \dots].$$

Since E is a homotopy associative ring spectrum, the E -based AHSS for MU is a multiplicative spectral sequence. It is also a spectral sequence of E_* -modules, so by use of the Leibniz rule, we see it must collapse on E_2 as the b_i are permanent cycles.

Let M be a monomial basis for $\mathbb{Z}[b_1, b_2, \dots]$ and consider the map of homotopy left E -modules

$$E\{M\} \rightarrow E \otimes \text{MU}$$

given by lifting the b_i and using the multiplication in $E \otimes \text{MU}$. Giving elements of M their corresponding AHSS filtration, this map becomes a map of filtered left homotopy E -modules. Since the AHSS for $E \otimes \text{MU}$ collapses, the map on associated graded is an equivalence, so the map is an equivalence since the filtrations are bounded below and exhaustive. Now the map

$$m : E \otimes \text{MU} \simeq E\{M\} \rightarrow E$$

given by projection onto the summand indexed by $1 \in M$ exhibits E as a weak MU-module. $\square$

Theorem 2.9 *A spectrum E is complex-orientable in the sense of Definition 2.3 if and only if E is a weak MU-module.*

Proof If E is complex-orientable, then so is $\text{End}(E) = F(E, E)$ since

$$F(E, E) \otimes \sigma_k \simeq F(E, E \otimes \sigma_k),$$

using that the domain and codomain of σ_k are dualizable. By Proposition 2.8, $\text{End}(E)$ is a weak homotopy MU-module. E is naturally an $\text{End}(E)$ -module, and in particular a weak $\text{End}(E)$ -module.

The “only if” direction of the claim then follows from the following observation. If $\eta_R : \mathbb{S} \rightarrow R$ and $\eta_T : \mathbb{S} \rightarrow T$ are $\mathbb{E}_0$ -algebras, and M is a weak R -module and R is a weak T -module, then M is a weak T -module. In the diagram

$$\begin{array}{ccc} M & \xrightarrow{\eta_T \otimes 1} & T \otimes M \\ \downarrow \eta_R \otimes 1 & & \downarrow 1 \otimes \eta_R \otimes 1 \\ R \otimes M & \xrightarrow{\eta_T \otimes 1 \otimes 1} & T \otimes R \otimes M \\ & \searrow & \downarrow m_R \otimes 1 \\ & & R \otimes M \\ & & \downarrow m_M \\ & & M \end{array}$$

the clockwise composite exhibits M as a weak T -module, where m_R is a weak T -module structure on R and m_M is a weak R -module structure on M .

Conversely, if E is a weak MU-module, then the map $E \otimes \sigma_k$ is a retract of the map $\text{MU} \otimes E \otimes \sigma_k$, which is null as MU is complex-orientable. $\square$

Corollary 2.10 *If E is complex-orientable, the Adams–Novikov spectral sequence of E collapses on the zero-line, i.e., it is concentrated in filtration zero from the E_2 -page on.*

Proof The ANSS of E comes from the coaugmented cosimplicial spectrum

$$E \rightarrow \text{MU} \otimes E \Rightarrow \text{MU} \otimes \text{MU} \otimes E \Rightarrow \text{MU} \otimes \text{MU} \otimes \text{MU} \otimes E \cdots$$

If E is a weak MU-module, this admits a (-1) -st codegeneracy, so that it is equivalent to the constant cosimplicial object at E . $\square$

Remark 2.11 If E is a weak MU-module, it is in particular MU-nilpotent, i.e., E is in the thick tensor ideal of Sp generated by MU. In fact, E is a weak MU-module if and only if E is MU-nilpotent of exponent 1, in the sense of [41, Part 1, Section 4]. As we will discuss in Section 7, this is also equivalent to asking E to be an MU-injective in the sense of MU-based Adams resolutions, which gives another proof of the preceding corollary.

2.2 $X(n)$ -orientations and chromatic defect

Theorem 2.9 tells us that the condition of being a weak homotopy MU-module is determined by the attaching maps for $\mathbb{CP}^\infty$. We may ask if there is a spectrum playing a similar role for $\mathbb{CP}^n$, i.e., for only the first $n - 1$ attaching maps of $\mathbb{CP}^\infty$. This is exactly Ravenel's Thom spectrum $X(n)$, which we now introduce. For a more thorough introduction to the $X(n)$, we strongly recommend Hopkins' thesis [27].

Definition 2.12 For $n \geq 1$, the spectrum $X(n)$ is the Thom spectrum

$$\text{Thom}(\Omega\text{SU}(n) \rightarrow \Omega\text{SU} \simeq \text{BU}),$$

where the equivalence $\Omega\text{SU} \simeq \text{BU}$ is Bott periodicity.

Remark 2.13 The equivalence $\Omega\text{SU} \simeq \text{BU}$ may be chosen to be a double loop map, so that $X(n)$ acquires the structure of an $\mathbb{E}_2$ -algebra in Sp . It is known that $X(n)$ is not $\mathbb{E}_3$ for $n > 1$ [35, Example 5.31].

The space $\Omega\text{SU}(n)$ admits a cell structure with even cells and $\mathbb{CP}^{n-1}$ as a subcomplex. Moreover, a Serre spectral sequence computation shows that

$$H_*(\Omega\text{SU}(n); \mathbb{Z}) \cong \text{Sym}(\tilde{H}_*(\mathbb{CP}^{n-1}; \mathbb{Z})) = \mathbb{Z}[b_1, \dots, b_{n-1}]$$

as a ring, with its product from the double loop space structure on $\Omega\text{SU}(n)$. The composite $\mathbb{CP}^{n-1} \rightarrow \Omega\text{SU}(n) \rightarrow \text{BU}$ classifies $L - 1$, where L is the tautological line bundle on $\mathbb{CP}^{n-1}$, and the Thom spectrum of $L - 1$ is $\Sigma^{-2}\mathbb{CP}^n$. Applying the Thom isomorphism, one finds that

$$H_*(X(n); \mathbb{Z}) \cong \text{Sym}(H_*(\Sigma^{-2}\mathbb{CP}^n; \mathbb{Z})) / (b_0 - 1) = \mathbb{Z}[b_1, \dots, b_{n-1}]$$

as a ring. With these facts, the proofs of the previous subsection now may be repeated to show the following.

Proposition 2.14 *A spectrum E is a weak $X(n)$ -module if and only if $E \otimes \sigma_k = 0$ for $1 \leq k \leq n-1$.*

The primary interest of this paper is to investigate spectra that are—in some sense—only finitely many steps away from being complex-orientable. Phrased as in Definition 2.3, we are interested in spectra E that become complex-orientable after forcing finitely many of the maps $E \otimes \sigma_k$ to be null, say for $k < n$. Proposition 2.14 tells us in some sense that, up to taking retracts, this passes through forming the tensor product $E \otimes X(n)$. We thus make the following definition.

Definition 2.15 A spectrum E has *chromatic defect* $\Phi(X) \leq n$ if $E \otimes X(n)$ is complex-orientable. We say $\Phi(X) = n$ if n is the smallest positive integer with this property.

Remark 2.16 The inequality appearing in the above definition is justified in the following sense. If $E \otimes X(n)$ is complex-orientable, then so is $E \otimes X(m)$ for any $m \geq n$. This follows from the definition of complex-orientability, as $X(m)$ is an $X(n)$ -module. In fact, the same argument shows that if $R \rightarrow S$ is a map of homotopy associative ring spectra, then $\Phi(R) \geq \Phi(S)$.

Example 2.17 We have $X(1) = \mathbb{S}$, so E has chromatic defect $\Phi(E) = 1$ if and only if E is complex-orientable. We give two nontrivial examples due to Hopkins [20, Chapter 9]:

- $\Phi(\mathrm{ko}) = \Phi(\mathrm{KO}) = 2$,
- $\Phi(\mathrm{tmf}) = \Phi(\mathrm{Tmf}) = \Phi(\mathrm{TMF}) = 4$.

We will return frequently to these examples. The spectra ko and tmf are prototypical examples of fp spectra, in the sense of Mahowald and Rezk [37], and we will investigate in Section 4 to what extent fp spectra admit finite chromatic defect.

If one localizes at a prime p , the spectrum MU splits

$$\mathrm{MU}_{(p)} \simeq \mathrm{BP}[x_i : i \neq p^k - 1]$$

as a sum of shifts of the p -primary Brown–Peterson spectrum BP , with $|x_i| = 2i$. Just as with MU , maps of homotopy ring spectra $X(m) \rightarrow X(m)$ are in bijection with the set of polynomials

$$x + b_1 x^2 + \cdots + b_{n-1} x^n$$

where $b_i \in \pi_{2i} X(m)$. By the Hurewicz theorem, the map $X(m) \rightarrow \mathrm{MU}$ induces an isomorphism in π_i for $i \leq 2m-2$. It follows that the Quillen idempotent ϵ on $\mathrm{MU}_{(p)}$ that defines BP restricts to an idempotent homotopy ring map $\epsilon : X(m)_{(p)} \rightarrow X(m)_{(p)}$.

Definition 2.18 We let $T(n)$ denote the spectrum $X(p^n)_{(p)}[\epsilon^{-1}]$.

Remark 2.19 $T(n)$ inherits the structure of a homotopy commutative ring spectrum from $X(p^n)$. Beardsley and Lawson [8] showed in fact that $T(n)$ is an $\mathbb{E}_1$ -summand of $X(p^n)$, but it is not known if it admits more structure; see [3].

Just as with BP, $X(m)$ splits p -locally. To describe this splitting, we use the following notation.

Definition 2.20 Let E be a homotopy associative ring spectrum, and let $\{x_1, x_2, \dots\}$ be a graded set. We let $E[x_1, x_2, \dots]$ denote the free E -module on the graded set given by the standard monomial basis of the polynomial ring $\mathbb{Z}[x_1, x_2, \dots]$.

Proposition 2.21 For $p^n \leq m < p^{n+1}$, there is a splitting of $T(n)$ -modules

$$X(m)_{(p)} \simeq T(n)[x_i : i \neq p^k - 1, i < m].$$

Proof See [51, 6.5.1]. □

For the same reasons that BP is often easier to work with than MU, it is often more convenient to work one prime at a time and phrase chromatic defect in terms of the $T(n)$.

Definition 2.22 For E a p -local spectrum, we let

$$\Phi_p(E) = \min\{n \geq 0 : E \otimes T(n) \text{ is complex-orientable}\}.$$

In particular,

$$\Phi_p(E) = \lfloor \log_p \Phi(E) \rfloor,$$

which follows from the previous proposition.

The filtration of BP by the $T(n)$ is especially nice for making inductive arguments due to the existence of a nice cell structure on $T(n+1)$ as a module over the $\mathbb{E}_1$ -ring $T(n)$. This was an essential ingredient in the Devinatz–Hopkins–Smith proof [19, Proposition 1.5] of the nilpotence theorem. We will also see a more structured description of this filtration due to Beardsley in Theorem 3.12.

Construction 2.23 One may use the Thom isomorphism to compute $H_*(X(n); \mathbb{Z})$, using that $X(n)$ is the Thom spectrum of a complex bundle. Applying the Quillen idempotent, one may compute that $H_*(T(n); \mathbb{Z}) = \mathbb{Z}[t_1, \dots, t_n]$ with $|t_i| = 2(p^i - 1)$; see [27, Corollary 1.3.8].

This allows one to form the diagram of $T(n)$ -modules

$$\begin{array}{ccccc} T(n) & \Sigma^{|t_{n+1}|} T(n) & \Sigma^{2|t_{n+1}|} T(n) & & \\ \downarrow t_{n+1}^0 & \downarrow t_{n+1}^1 & \downarrow t_{n+1}^2 & & \\ T(n+1) & \longrightarrow X_{(1)} & \longrightarrow X_{(2)} & \longrightarrow \cdots \end{array}$$

where $X_{(k)}$ is the cofiber of t_{n+1}^{k-1} , and t_{n+1}^k is the map of $T(n)$ -modules adjoint to the map $\mathbb{S}^{k|t_{n+1}|} \rightarrow X_{(k)}$ provided by the Hurewicz theorem, using that

$$H_*(X_{(k)}; \mathbb{Z}_{(p)}) = H_*(T(n); \mathbb{Z}_{(p)})\{t_{n+1}^m : m \geq k\}.$$

Using the Hurewicz theorem, we see that $X_{(\infty)} = 0$, hence defining

$$X^{(k)} = \text{fib}(T(n+1) \rightarrow X_{(k+1)})$$

we have an exhaustive cell structure

$$X^{(0)} \rightarrow X^{(1)} \rightarrow \dots \rightarrow X^{(\infty)} = T(n+1)$$

on the $T(n)$ -module $T(n+1)$ whose associated graded is the free module $T(n)[t_{n+1}]$.

Proposition 2.24 *Let E be any spectrum with $\Phi_p(E) \leq n \leq m \leq \infty$. There is an equivalence of $T(n)$ -modules*

$$E \otimes T(m) \simeq E \otimes T(n)[t_{n+1}, t_{n+2}, \dots, t_m].$$

Proof The decomposition of the $T(n)$ -module follows inductively from the cell structure of Construction 2.23 along with the fact that each of the cofiber sequences therein split after tensoring with BP. Indeed, this latter fact follows from the evenness of $\text{BP} \otimes T(n)$. $\square$

We finish this section with the following conjecture.

Conjecture 2.25 *Any spectrum with finite chromatic defect satisfies the condition of the telescope conjecture at all heights and primes.*

Remark 2.26 In unpublished work, Robert Burklund proves that for any k and n , there is a compact $\text{Tel}(k)$ -local spectrum $F_{n,k}$ that is a module over $L_{\text{Tel}(k)}X(n)$, where $\text{Tel}(k)$ is the v_k -telescope on a finite type k complex. This result would give a proof for the above conjecture. Indeed, if E has finite chromatic defect, then E satisfies the telescope conjecture at height k if and only if $L_{\text{Tel}(k)}(E \otimes F_{n,k})$ satisfies the telescope conjecture at height k , by use of the thick subcategory theorem. Since $F_{n,k}$ is a $L_{\text{Tel}(k)}X(n)$ -module, it follows that $L_{\text{Tel}(k)}(E \otimes F_{n,k})$ is a retract of $L_{\text{Tel}(k)}(E \otimes X(n) \otimes F_{n,k})$. The property of satisfying the telescope conjecture at height k is closed under retracts, and $L_{\text{Tel}(k)}(E \otimes X(n) \otimes F_{n,k})$ satisfies this property since it is complex-orientable by assumption.

Note that having infinite chromatic defect does not prevent a spectrum satisfying the condition of the telescope conjecture at all heights. For example, the connective image of J spectrum j satisfies the condition of the telescope conjecture at all heights since it is MU-nilpotent, but it has infinite chromatic defect by Theorem 5.18.

2.3 Wood-types

In this section, we work p -locally. First we recall the following standard definition.

Definition 2.27 A *finite BP-projective* is a finite spectrum F such that BP_*F is a projective BP_* -module.

In fact, we may as well work just with finite BP-frees, due to the following lemma.

Lemma 2.28 *Every finite BP-projective F is finite BP-free. That is, $\text{BP} \otimes F$ is a finite free BP-module.*

Proof By finite typeness, it suffices to prove the claim after p -completion. Since BP_*F is projective, one uses Brown representability to exhibit $\text{BP} \otimes F$ as a retract of $\text{BP}\{T\}$, for some finite graded indexing set T .

Since the $\mathbb{F}_p$ -based ASS of BP collapses on E_2 , it follows that that of $\text{BP} \otimes F$ does as well. It follows from the indecomposability of the $\mathcal{A}_*$ -comodule $H_*(\text{BP}; \mathbb{F}_p)$ that the inclusion $H_*(\text{BP} \otimes F; \mathbb{F}_p) \rightarrow H_*(\text{BP}\{T\}; \mathbb{F}_p)$ is of the form $H_*(\text{BP}; \mathbb{F}_p)\{T'\} \rightarrow H_*(\text{BP}; \mathbb{F}_p)\{T\}$ for some $T' \subset T$. The corresponding statement thus holds on E_∞ -pages and on homotopy groups. $\square$

Remark 2.29 Lemma 2.28 may also be proven inductively using a cell structure on F . In fact, all of the examples we will discuss concern finite complexes F with only even-dimensional cells, which are automatically BP-free since BP_* is concentrated in even degrees.

Definition 2.30 A p -local spectrum E is said to be *Wood-type* if there exists a finite BP-projective F such that $E \otimes F$ is complex-orientable.

Example 2.31 Fix a height- n formal group Γ over a perfect field k of characteristic p and a finite subgroup $G \subset \mathbb{G}_n$ of the corresponding Morava stabilizer group, and we let $\text{EO}_n(G)$ denote the fixed points of $E(k, \Gamma)$ with respect to G . It is a theorem of Meier, Naumann and Noel that $\text{EO}_n(G)$ is Wood-type; see [16, Appendix B].

Example 2.32 The above example gives higher height generalizations of Wood's theorem for the periodic theory KO, which becomes $\text{EO}_1(C_2)$ after 2-completion. Wood equivalences for connective theories like ko are harder to come by, and we revisit this in Section 4 in the context of fp spectra.

Some known examples beyond ko include the equivalence

$$\text{tmf} \otimes D\mathcal{A}(1) \simeq \text{tmf}_1(3)$$

at the prime 2, and the equivalence

$$\text{tmf} \otimes X_2 \simeq \text{tmf}_1(2)$$

at the prime 3, where X_2 is the 8-skeleton of $T(1)$; see [39]. The spectra $\text{tmf}_1(3)$ and $\text{tmf}_1(2)$ are complex-orientable, so these equivalences imply that tmf is Wood-type.

As before, the thick subcategory theorem implies the following.

Proposition 2.33 *If E is Wood-type, then E is BP-nilpotent.*

The notion of Wood-types is closely tied to that of chromatic defect via the following useful fact.

Proposition 2.34 *Every finite BP-free F is a finite $T(n)$ -free for some $0 \leq n < \infty$.*

Proof If F is a BP-free, there is an equivalence of BP-modules $\text{BP} \otimes F \simeq \text{BP}\{T\}$ for some finite graded indexing set T . For each $\alpha \in T$, the corresponding map $\mathbb{S}^{|\alpha|} \rightarrow \text{BP} \otimes F$ factors through $T(n) \otimes F$ for some n , since $T(n) \rightarrow \text{BP}$ is a $(2p^{n+1}-4)$ -equivalence. Since T is finite, we may choose n such that such a factorization exists for all $\alpha \in T$, and this defines a map

$$T(n)\{T\} \rightarrow T(n) \otimes F.$$

Applying $\text{BP} \otimes_{T(n)} -$, this map becomes an equivalence, thus it is an equivalence as BP is a free $T(n)$ -module by Proposition 2.24. $\square$

This gives a necessary condition for E to be Wood-type.

Corollary 2.35 *If E is Wood-type, then E has finite chromatic defect.*

Proof By the proposition, there is some finite $T(n)$ -free F such that $E \otimes F$ is complex-orientable. It follows that $E \otimes T(n)$ is a retract of the complex-orientable spectrum $E \otimes F \otimes T(n)$. $\square$

Corollary 2.36 *For any $E(k, \Gamma)$ and $G \subset \mathbb{G}_n$ as in Example 2.31, the spectrum $\mathrm{EO}_n(G)$ has finite chromatic defect.*

In Section 6 we will return to this example and compute the chromatic defect $\Phi(\mathrm{EO}_n(G))$ precisely in terms of the valuation on the ring $\mathrm{End}(\Gamma)$.

3 The $T(n)$ and $\mathbb{F}_p$ -homology

In this section, we use the Thom isomorphism to analyze the $T(n)$ and chromatic defect from the point of view of the Adams spectral sequence.

3.1 The Adams spectral sequence of $T(n)$

We fix a prime p and let $T(n)$ be the summand of $X(p^n)_{(p)}$ as in Definition 2.18. The Thom isomorphism implies that the map $T(n) \rightarrow \mathbb{F}_p$ induces an injection of $\mathcal{A}_*$ -comodule algebras

$$H_*(T(n); \mathbb{F}_p) \rightarrow H_*(\mathbb{F}_p; \mathbb{F}_p) = \mathcal{A}_*,$$

with image $\mathbb{F}_2[\zeta_1^2, \dots, \zeta_n^2]$ for $p = 2$ and $\mathbb{F}_p[\zeta_1, \dots, \zeta_n]$ for p odd. This can be described as a coinduced comodule.

Proposition 3.1 *There is an isomorphism of $\mathcal{A}_*$ -comodule algebras*

$$H_*(T(n); \mathbb{F}_p) \cong \begin{cases} \mathcal{A}_* \square_{\mathcal{A}_*/(\xi_1^2, \dots, \xi_n^2)} \mathbb{F}_2 & \text{if } p = 2, \\ \mathcal{A}_* \square_{\mathcal{A}_*/(\xi_1, \dots, \xi_n)} \mathbb{F}_p & \text{if } p > 2. \end{cases}$$

The quotient Hopf algebra $\mathcal{A}_*/(\xi_1^2, \dots, \xi_n^2)$ at $p = 2$ is given by

$$E(\xi_1, \dots, \xi_n) \otimes P(\xi_{n+1}, \xi_{n+2}, \dots)$$

and, at odd primes, $\mathcal{A}_*/(\xi_1, \dots, \xi_n)$ is given by

$$E(\tau_1, \tau_2, \dots) \otimes P(\xi_{n+1}, \xi_{n+2}, \dots).$$

A change-of-rings isomorphism then gives the following.

Corollary 3.2 *The E_2 -page of the Adams spectral sequence for $T(n)$ is given by*

$$\mathrm{Ext}_{\mathcal{A}_*}(\mathbb{F}_p, H_*T(n)) \cong \begin{cases} \mathrm{Ext}_{E(\xi_1, \dots, \xi_n) \otimes P(\xi_{n+1}, \xi_{n+2}, \dots)}(\mathbb{F}_2, \mathbb{F}_2) & \text{if } p = 2, \\ \mathrm{Ext}_{E(\tau_1, \tau_2, \dots) \otimes P(\xi_{n+1}, \xi_{n+2}, \dots)}(\mathbb{F}_p, \mathbb{F}_p) & \text{if } p > 2. \end{cases}$$

We now define an increasing filtration on these quotient Hopf algebras to obtain a corresponding May spectral sequence. This is the usual May filtration on $\mathcal{A}_*$ (see [51, Section 3.2]) projected to the quotient. For $p = 2$, this is obtained by giving $\xi_i^{2^j}$ filtration $2i - 1$, and for $p > 2$ by giving $\xi_i^{2^j}$ and τ_{i-1} filtration $2i - 1$, and extending multiplicatively. At $p = 2$, the associated graded of this filtration is a tensor product of primitively generated exterior algebras. At odd primes, it is a tensor product of primitively generated exterior algebras and primitively generated truncated polynomial algebras of height p . As in the case for the sphere, we deduce the following.

Proposition 3.3 *There is a May spectral sequence converging to the E_2 -page of the Adams spectral sequence for $T(n)$.*

- When $p = 2$, the May SS has signature

$$E_1^{s,t,w} = \mathbb{F}_2[h_{i,j} : j = 0 \text{ if } i \leq n] \Rightarrow \text{Ext}_{\mathcal{A}_*}^{s,t}(\mathbb{F}_2, H_*T(n)),$$

where $h_{i,j} = [\xi_i^{2^j}]$. In (s, t, w) tridegrees, we have $|h_{i,j}| = (1, 2^j(2^i - 1), 2i - 1)$.

- When $p > 2$, the May SS has signature

$$E_1^{s,t,w} = E(h_{i,j} : i > n) \otimes \mathbb{F}_p[b_{i,j} : i > n] \otimes \mathbb{F}_p[a_i] \Rightarrow \text{Ext}_{\mathcal{A}_*}^{s,t}(\mathbb{F}_p, H_*T(n)),$$

where $h_{i,j} = [\xi_i^{p^j}]$, $a_i = [\tau_i]$ and $b_{i,j}$ is the p -fold Massey power of $h_{i,j}$, which is represented explicitly by the cocycle

$$\sum_{k=1}^p \frac{1}{p} \binom{p}{k} [\xi_i^{p^{jk}} | \xi_i^{p^{j(p-k)}}].$$

The (s, t, w) tridegrees of the generators are $|h_{i,j}| = (1, 2p^j(p^i - 1), 2i - 1)$, $|a_i| = (1, 2p^i - 1, 2i - 1)$ and $|b_{i,j}| = (2, 2p^j(p^i - 1), 4i - 2)$.

This spectral sequence may be used to compute the E_2 -page of the Adams spectral sequence for $T(n)$ through a range. We plot in Figure 1 the E_1 -page for $T(1)$ at $p = 2$. When working with the May spectral sequence, it can be convenient to ignore the May filtration degree, and plot the spectral sequence in Adams bigrading $(t - s, s)$ for a class in $\text{Ext}^{s,t}$, where s is the cohomological degree and t is the internal degree. In this grading, all May differentials have the signature of an Adams d_1 . By direct use of the coproduct formula and Nakamura's lemma [47], one can compute some May differentials for $T(1)$ in low stems.

Proposition 3.4 *In the May SS for $T(1)$ at $p = 2$, we have the differentials*

$$d_1(h_{3,0}) = h_{1,0}h_{2,1}, \quad d_1(h_{4,0}) = h_{1,0}h_{3,1} + h_{2,0}h_{2,2}, \quad d_2(h_{3,0}^2) = h_{1,0}^2h_{2,2}.$$

The May SS for $T(n)$ in general exhibits a similar pattern. The following comes by a degree check on the generators of the E_1 -page, and the differential follows from the coproduct formula.




Figure 1: The E_1 -page of the May SS for $T(1)$ at $p = 2$.

Proposition 3.5 *The inclusion*

$$\mathbb{F}_2[h_{i,0}, h_{n+1,1} : i \leq n+2] \rightarrow E_1\text{-May}(T(n))$$

is an isomorphism in stems $t-s \leq 2^{n+2}-2 = \text{stem}(h_{n+2,0})$ at $p = 2$. There is a differential $d_1(h_{n+2,0}) = h_{1,0}h_{n+1,1}$.

The inclusion

$$E(h_{n+1,0}) \otimes \mathbb{F}_p[a_i : i \leq n+2] \rightarrow E_1\text{-May}(T(n))$$

is an isomorphism in stems $t-s \leq 2p^{n+1}-2 = \text{stem}(a_{n+1})$ for $p > 2$. There is a differential $d_1(a_{n+1}) = h_{1,0}h_{n+1,0}$.

Proof We give the proof for $p = 2$, the odd prime case being analogous. The inclusion in question is the canonical inclusion


$$\mathbb{F}_2[h_{i,0}, h_{n+1,1} : i \leq n+2] \rightarrow \mathbb{F}_2[h_{i,j} : j = 0 \text{ if } i \leq n].$$

In particular, the polynomial generators not in the image of this inclusion consist of $h_{n+1,j}$ where $j > 1$, $h_{n+2,j}$ where $j > 0$, and $h_{n+k,j}$ for $k > 2$ and for all j . The $t-s$ degree of each of these generators exceeds that of $h_{n+2,0}$, so these generators do not contribute to the range in question.

The May differential $d_1(h_{n+2,0}) = h_{1,0}h_{n+1,1}$ follows from the coproduct formula

$$\Delta(\xi_{n+2}) = \xi_{n+2} \otimes 1 + 1 \otimes \xi_{n+2} + \xi_{n+1}^2 \otimes \xi_1 + \sum_{j=2}^{n+1} \xi_{n+2-j}^2 \otimes \xi_j$$

so that, in the cobar complex, $d([\xi_{n+2}]) = [\xi_{n+1}^2 | \xi_1]$ modulo higher May filtration terms. $\square$

Figure 2: The E_2 -page of the Adams SS for $T(1)$ at $p = 2$.

Corollary 3.6 *The first nonzero odd homotopy group of $T(n)$ is*

$$\pi_{2p^{n+1}-3}T(n) = \mathbb{Z}/p.$$

A generator is detected by $h_{n+1,1}$ at $p = 2$, and $h_{n+1,0}$ at $p > 2$. The same is true for $X(m)_{(p)}$ whenever $p^n \leq m < p^{n+1}$.

Proof By Proposition 3.5, the $2p^{n+1} - 3$ stem of the E_2 -page of the May SS for $T(n)$ has a single nonzero bidegree $(2p^{n+1} - 3, 1)$ generated by $h_{n+1,1}$ when $p = 2$ and by $h_{n+1,0}$ when $p > 2$. The $2p^{n+1} - 4$ stem is $h_{1,0}$ -torsion free and $h_{1,0}h_{n+1,1} = 0$ (resp. $h_{1,0}h_{n+1,0} = 0$ at $p > 2$), so the generator of the $2p^{n+1} - 3$ stem cannot support a differential. The May SS thus collapses on E_2 in this stem, and the same argument shows the generator cannot support an Adams differential.

The statement about $X(m)$ then follows immediately from the splitting of Proposition 2.21. $\square$

In fact, the attaching maps σ_k discussed in the previous section identify for us a canonical generator of $\pi_{2p^{n+1}-3}T(n)$. Since $T(n)$ is a weak $X(p^{n+1}-1)$ -module, one has a splitting

$$T(n) \otimes \Sigma^{-2}\mathbb{C}\mathbb{P}^{p^{n+1}-1} \simeq T(n)\{b_0, \dots, b_{p^{n+1}-1}\}.$$

Since the odd homotopy groups of $T(n)$ vanish in degrees less than $2p^{n+1} - 3$, the attaching map

$$\sigma_{p^{n+1}-1} : \mathbb{S}^{2p^{n+1}-3} \rightarrow T(n) \otimes \Sigma^{-2}\mathbb{C}\mathbb{P}^{p^{n+1}-1}$$

has nonzero component at b_i if and only if $i = 0$. This component defines an element of $\pi_{2p^{n+1}-3}T(n)$ which we call χ_{n+1} .

Proposition 3.7 *The element χ_{n+1} is a generator of $\pi_{2p^{n+1}-3}T(n) = \mathbb{Z}/p$.*

Proof The above discussion shows that $\chi_{n+1} = 0$ if and only if $T(n) \otimes \sigma_{p^{n+1}-1}$ is null. If this is the case, then $T(n)$ is a weak $X(p^{n+1})$ -module by Proposition 2.14, and therefore a weak $T(n+1)$ -module. However, the map $T(n+1) \rightarrow \text{BP}$ is $(2p^{n+2}-3)$ -connected, hence the same is true for the map

$$T(n) \otimes T(n+1) \rightarrow T(n) \otimes \text{BP}.$$

The homotopy groups of $T(n) \otimes \text{BP}$ are concentrated in even degrees, thus $\pi_{2p^{n+1}-3}(T(n) \otimes T(n+1)) = 0$, contradicting that $\pi_{2p^{n+1}-3}T(n)$ is a retract. $\square$

Remark 3.8 The source $h_{n+2,0}$ (resp. a_{n+1} at $p > 2$) of the May d_1 of Proposition 3.5 detects the class

$$v_{n+1} \in \pi_{2(p^{n+1}-1)}T(n+1) \cong \pi_{2(p^{n+1}-1)}\text{BP}.$$

This implies in fact that $v_{n+1} \in \pi_{2(p^{n+1}-1)}(T(n+1)/p)$ lifts to a class in

$$\pi_{2(p^{n+1}-1)}(T(n)/p),$$

whose Bockstein is χ_{n+1} . This gives a Toda bracket description in $T(n+1)$

$$v_{n+1} \in \langle p, \chi_{n+1}, 1 \rangle,$$

as one can show using, for example, [42, Lemma 4.6.1].

The element χ_{n+1} gives a convenient criterion to compute the chromatic defect of a homotopy associative ring spectrum E , namely an explicit sequence of obstructions in $T(n)$ homology to $\Phi_p(E)$ being $\leq n$. We continue to work in the p -local setting.

Proposition 3.9 *Let E be a p -local homotopy associative ring spectrum. Then $\Phi_p(E) \leq n$ if and only if*

$$(3-1) \quad \chi_{m+1} = 0 \in \pi_{2p^{m+1}-3}(E \otimes T(m))$$

for all $m \geq n$.

Proof By Proposition 2.14 and the discussion preceding Proposition 3.7, $E \otimes T(m)$ is a weak $T(m+1)$ -module if and only if

$$E \otimes T(m) \otimes \sigma_{p^{m+1}-1} = 0,$$

whose only potentially nonzero component is $E \otimes \Sigma^{2p^{m+1}-3}T(m) \xrightarrow{E \otimes \chi_{m+1}} E \otimes T(m)$. Since $E \otimes T(m)$ is a homotopy associative ring spectrum, this is null if and only if (3-1) holds. By induction, if (3-1) holds for all $m \geq n$, then $E \otimes T(m)$ is a weak $T(m)$ -module for all $m \geq n$, hence complex-orientable. $\square$

3.2 The chromatic defect of finite spectra

We begin this section with a toy example, on which we will elaborate to show that the chromatic defect of any finite spectrum is infinite. In this example, we will show that the finite spectrum $C(\eta)$ is not complex-orientable; i.e., that $\Phi(C(\eta)) > 1$. The fact that $C(\eta)$ has infinite chromatic defect is the claim that $\Phi(C(\eta)) > n$ for all n , which follows by a completely analogous argument in the category of $X(n)$ -modules; see Corollary 3.15.

We work 2-locally throughout this example. We have $X(1) = \mathbb{S}$, and $T(0) = X(1)_{(p)} = \mathbb{S}_{(p)}$. The attaching map $\sigma_1 : \mathbb{S}^1 \rightarrow \mathbb{S}$ is by definition the Hopf map η , which corresponds to the element χ_1 when $p = 2$. The quotient $\mathbb{S}/\eta = C(\eta)$ does not admit a unital multiplication. Indeed, if it did, the endomorphism

$$\eta : \Sigma C(\eta) \rightarrow C(\eta)$$

must be null. Taking the cofiber, this would give an equivalence

$$C(\eta) \otimes C(\eta) \simeq C(\eta) \oplus \Sigma^2 C(\eta),$$

but the mod 2 cohomology of the left-hand side admits a nontrivial Sq^4 -action, while the right-hand side does not.

We may instead embed the quotient $C(\eta)$ into a ring spectrum by forming the $\mathbb{E}_1$ -quotient of $\mathbb{S}$ by η as the pushout

$$\begin{array}{ccc} \text{Free}_{\mathbb{E}_1}(\mathbb{S}^1) & \xrightarrow{\bar{\eta}} & \mathbb{S} \\ \bar{0} \downarrow & & \downarrow \\ \mathbb{S} & \longrightarrow & \mathbb{S} // \eta \end{array}$$

in the category of $\mathbb{E}_1$ -rings, where $\bar{0}$ and $\bar{\eta}$ denote the adjoint maps with respect to the free-forget adjunction. This quotient is in fact $T(1)$.

Proposition 3.10 *There is an equivalence of $\mathbb{E}_1$ -rings*

$$\mathbb{S} // \eta \rightarrow T(1)$$

2-locally, and an equivalence of $\mathbb{E}_1$ -rings

$$\mathbb{S} // \alpha \rightarrow T(1)$$

p -locally for odd primes p .

Proof We give the proof for $p = 2$, the odd primary case being analogous. Choosing a nullhomotopy of η in $\pi_* T(1)$ determines a map $\mathbb{S} // \eta \rightarrow T(1)$ by universal property, and since both sides are finite type, we may check that is an isomorphism in mod 2 cohomology. Since η is zero in mod 2 homology, when we tensor this pushout diagram up to $\mathbb{F}_2$, we have a pushout

$$\begin{array}{ccc} \text{Free}_{\mathbb{E}_1 - \mathbb{F}_2}(\mathbb{S}^1) & \xrightarrow{\text{Free}_{\mathbb{E}_1 - \mathbb{F}_2}(0)} & \text{Free}_{\mathbb{E}_1 - \mathbb{F}_2}(*) \\ \text{Free}_{\mathbb{E}_1 - \mathbb{F}_2}(0) \downarrow & & \downarrow \\ \text{Free}_{\mathbb{E}_1 - \mathbb{F}_2}(*) & \longrightarrow & \mathbb{S} // \eta \otimes \mathbb{F}_2 \end{array}$$

Since the Free functor preserves pushouts, one has

$$\pi_*(\mathbb{S} // \eta \otimes \mathbb{F}_2) = \pi_*(\text{Free}_{\mathbb{E}_1 - \mathbb{F}_2}(\mathbb{S}^2)) = \mathbb{F}_2[x_2],$$

where $|x_2| = 2$.

We claim now that in mod 2 cohomology, Sq^2 of the nonzero class in degree 0 is equal to the nonzero class in degree 2 for both $H^*(\mathbb{S}/\eta)$ and $H^*(T(1))$. With this in place, by linearity over the Steenrod algebra, the map $\mathbb{S}/\eta \rightarrow T(1)$ must induce an isomorphism on $H^{\leq 2}(-; \mathbb{F}_2)$, and therefore also on $H_{\leq 2}(-; \mathbb{F}_2)$. Since the map induced on homology is a ring map, it is therefore an isomorphism.

To establish the claim about Sq^2 , note that for a ring spectrum R , there is an extension of the unit map $\mathbb{S} \rightarrow R$ over $C(\eta)$. If R is connective and $H^0(R; \mathbb{F}_2) = \mathbb{F}_2$, then the map $C(\eta) \rightarrow R$ must induce an isomorphism on $H^0(-; \mathbb{F}_2)$ since it extends the unit. But then since Sq^2 is nonzero on $H^0(C(\eta); \mathbb{F}_2)$ it must be nonzero on $H^0(R; \mathbb{F}_2)$. Note now that $\mathbb{S}/\eta$ and $T(1)$ satisfy the conditions on R . $\square$

Remark 3.11 It follows that if E is any p -local spectrum, then E admits the structure of a $T(1)$ -module if and only if $E \otimes \eta$ is null for $p = 2$ and if and only if $E \otimes \alpha_1$ is null if $p > 2$. Indeed, if $E \otimes \eta$ (resp. $E \otimes \alpha_1$ for $p > 2$) is null, then the $\mathbb{E}_1$ -ring map $\mathbb{S} \rightarrow \mathrm{End}(E)$ factors canonically through an $\mathbb{E}_1$ -ring map $T(1) \rightarrow \mathrm{End}(E)$ by universal property, giving E a $T(1)$ -module structure. Conversely, the map η (resp. α_1 for $p > 2$) acts by zero on any $T(1)$ -module.

We finish the toy example by noting the following: if $C(\eta)$ were complex-orientable, then the map η must act by zero on $C(\eta)$. The above proposition would then provide a ring map

$$f : T(1) \rightarrow \mathrm{End}(C(\eta)).$$

This map f would send $1 \mapsto 1$ in mod 2 homology. However, in $H_*T(1) = \mathbb{F}_2[x_2]$, there is a coaction formula

$$\psi(x_2^{2^n}) = \zeta_1^{2^{n+1}} \otimes 1 + 1 \otimes x_2^{2^n}.$$

Since $H_*(f)$ is a map of comodule algebras, this implies that $x_2^{2^n}$ must be sent to a nonzero element for all n , contradicting the boundedness of $H_*\mathrm{End}(C(\eta))$.

We adapt this argument now to a general finite spectrum and to the more general notion of finite chromatic defect. This hinges on a theorem of Beardsley [7], which may be shown by adapting Proposition 3.10 to the category of $X(p^n)_{(p)}$ -modules, with χ_{n+1} in place of η .

Theorem 3.12 (Beardsley) *Let $p^n \leq m < p^{n+1}$. If $m < p^{n+1} - 1$, then*

$$X(m+1)_{(p)} \simeq X(m)_{(p)}[b_m]$$

is the free $\mathbb{E}_1$ - $X(m)_{(p)}$ -algebra on a class in degree $2m$.

If $m = p^{n+1} - 1$, $X(m+1)_{(p)}$ is the free $\mathbb{E}_1$ - $X(m)_{(p)}$ -algebra with a nullhomotopy of χ_{n+1} . That is, there is a pushout

$$\begin{array}{ccc} \mathrm{Free}_{\mathbb{E}_1 - X(p^{n+1}-1)_{(p)}}(\mathbb{S}^{2p^{n+1}-3}) & \xrightarrow{\overline{\chi_{n+1}}} & X(p^{n+1}-1)_{(p)} \\ \bar{0} \downarrow & & \downarrow \\ X(p^{n+1}-1)_{(p)} & \longrightarrow & X(p^{n+1}) \end{array}$$

in $\mathbb{E}_1$ - $X(p^{n+1}-1)_{(p)}$ -algebras.

Essentially, Beardsley's theorem says if one attaches an $\mathbb{E}_1$ - $T(n)$ -cell to $T(n)$ to kill χ_{n+1} , one gets $T(n+1)$. However, since $T(n)$ is not known to be an $\mathbb{E}_2$ -ring, it doesn't make sense to speak of $\mathbb{E}_1$ - $T(n)$ -algebras, so we must state things in terms of the $\mathbb{E}_2$ -rings $X(m)_{(p)}$.

The discussion following Proposition 3.10 may be repeated in the category of $X(p^n)_{(p)}$ -modules, using the relative dual Steenrod algebra. We will only need the following piece of information.

Lemma 3.13 *Let $p^n \leq m < p^{n+1}$. The map*

$$\mathcal{A}_* = \pi_*(\mathbb{F}_p \otimes \mathbb{F}_p) \rightarrow \pi_*(\mathbb{F}_p \otimes_{X(m)} \mathbb{F}_p) =: \mathcal{A}_*^{X(m)}$$

sends $\zeta_{n+1}^{p^k}$ to a nonzero coalgebra primitive for all k when $p > 2$. When $p = 2$, the map sends $\zeta_{n+1}^{2^k}$ to a nonzero coalgebra primitive for all $k > 0$.

Proof We have an isomorphism

$$\pi_*(\mathbb{F}_p \otimes_{X(m)} \mathbb{F}_p) \cong \pi_*((\mathbb{F}_p \otimes \mathbb{F}_p) \otimes_{\mathbb{F}_p \otimes_{X(m)} \mathbb{F}_p} \mathbb{F}_p) \cong \mathcal{A}_* \otimes_{H_*X(m)} \mathbb{F}_p.$$

Indeed, the map $H_*T(n) \rightarrow \mathcal{A}_*$ is identified with the flat inclusion

$$\mathbb{F}_2[\zeta_1^2, \dots, \zeta_n^2] \rightarrow \mathbb{F}_2[\zeta_1, \zeta_2, \dots].$$

Moreover, $H_*T(n)$ is flat over $H_*X(m)$ as it is a retract of a free module. This implies $\mathcal{A}_*$ is flat over $H_*X(m)$, and the above isomorphism follows from collapse of the corresponding Künneth spectral sequence.

The splitting $H_*X(m) \cong H_*T(n)[x_i : i \neq p^k - 1, i < m]$ implies that

$$\mathcal{A}_* \otimes_{H_*X(m)} \mathbb{F}_p \cong (\mathcal{A}_* \otimes_{H_*T(n)} \mathbb{F}_p)[x_i : i \neq p^k - 1, i < m].$$

It follows from Proposition 3.1 that

$$\mathcal{A}_* \otimes_{H_*T(n)} \mathbb{F}_p \cong \begin{cases} E(\xi_1, \dots, \xi_n) \otimes P(\xi_{n+1}, \xi_{n+2}, \dots) & \text{when } p = 2, \\ E(\tau_1, \tau_2, \dots) \otimes P(\xi_{n+1}, \xi_{n+2}, \dots) & \text{when } p > 2, \end{cases}$$

from which the claim now follows. $\square$

Proposition 3.14 *Suppose $p^n \leq m < p^{n+1}$, and $p^{n+1} \leq m+k \leq \infty$. There are no nontrivial compact $X(m)_{(p)}$ -modules that are weak $X(m+k)_{(p)}$ -modules.*

Proof In what follows, we implicitly work p -locally and drop this from the notation. If F is a compact $X(m)$ -module that is a weak $X(m+k)$ -module, then by Theorem 3.12 there is an $\mathbb{E}_1$ - $X(m)$ -algebra map

$$f : X(p^{n+1}) \rightarrow \text{End}_{X(m)}(F)$$

as $0 = \chi_{n+1} \in \pi_*X(m+k)$.

The relative homology

$$H_*^{X(m)}(\text{End}_{X(m)}(F)) = \pi_*(\mathbb{F}_p \otimes_{X(m)} \text{End}_{X(m)}(F))$$

is a comodule over the relative dual Steenrod algebra $\mathcal{A}_*^{X(m)}$. These relative homology groups are bounded above as $\text{End}_{X(m)}(F)$ is compact, and $H_*^{X(m)}(X(m))$ is bounded above.

Via the map of Hopf algebras $\mathcal{A}_* \rightarrow \mathcal{A}_*^{X(m)}$, any $\mathcal{A}_*$ -comodule may be regarded as an $\mathcal{A}_*^{X(m)}$ -comodule, and the composite

$$H_*T(n+1) \rightarrow H_*(X(p^{n+1})) \rightarrow H_*^{X(m)}(X(p^{n+1})) \xrightarrow{f_*} H_*^{X(m)}(\text{End}_{X(m)}(F))$$

is a map of $\mathcal{A}_*^{X(m)}$ -comodules sending $1 \mapsto 1$. By Lemma 3.13, one has the coaction formula

$$\psi(\zeta_{n+1}^{p^k}) = 1 \otimes \zeta_{n+1}^{p^k} + \zeta_{n+1}^{p^k} \otimes 1$$

in the $\mathcal{A}_*^{X(m)}$ -comodule $H_*T(n+1)$ at odd primes, and similarly at $p = 2$ with ζ_{n+1}^2 in place of ζ_{n+1} .

The above composite must therefore send $\zeta_{n+1}^{p^k}$ (resp. $\zeta_{n+1}^{2^{k+1}}$ at $p = 2$) to a nonzero element for all $k \geq 0$. This gives a contradiction if $1 \neq 0 \in \pi_0(\text{End}_{X(m)}(F))$ by boundedness. $\square$

Corollary 3.15 *If F is a nontrivial finite spectrum, then $\Phi(F) = \infty$. Working p -locally, if F is a nontrivial finite p -local spectrum, then $\Phi_p(F) = \infty$.*

Proof If F is a nontrivial finite spectrum, then $F_{(p)}$ is nontrivial for some prime p , so the latter claim implies the former. The $X(m)_{(p)}$ -module $F \otimes X(m)_{(p)}$ is compact, so if $F \otimes X(m)_{(p)}$ is complex-orientable, it is a weak MU-module, contradicting Proposition 3.14 if $m < \infty$. $\square$

3.3 $\mathcal{P}(n)$ -free complexes

An important fact in stable homotopy is the existence of type $n+1$ finite complexes. To construct a type $n+1$ finite complex, it suffices to construct a finite complex F such that $H^*(F; \mathbb{F}_p)$ is free over the subalgebra $\mathcal{A}(n)$ of the Steenrod algebra $\mathcal{A}$. The Steenrod algebra has an even variant $\mathcal{P}$ and corresponding subalgebras $\mathcal{P}(n)$. We mimic the $\mathcal{A}(n)$ case to construct even cell complexes whose cohomologies are free over $\mathcal{P}(n)$. We discuss also the structure of $H^*T(n)$ as a module over $\mathcal{A}(n)$; this will allow us to more closely relate chromatic defect with Wood-types.

Definition 3.16 We define the following Hopf algebras.

- Let $\mathcal{A}(n)$ be the subalgebra of the Steenrod algebra $\mathcal{A}$ generated by $\{\text{Sq}^{2^i} : 0 \leq i \leq n\}$ for $p = 2$ and $\{\beta, P^{p^i} : 0 \leq i \leq n-1\}$ for $p > 2$.
- Let $\mathcal{P} = \mathcal{A}/\langle \text{Sq}^1 \rangle$ where $\langle \text{Sq}^1 \rangle$ is the 2-sided ideal generated by Sq^1 for $p = 2$, and let $\mathcal{P}$ be the subalgebra of $\mathcal{A}$ generated by the P^i for $p > 2$.
- Let $\mathcal{P}(n)$ be the subalgebra of $\mathcal{P}$ generated by $\{\text{Sq}^{2^i} : 1 \leq i \leq n+1\}$ for $p = 2$ and $\{P^{p^i} : 0 \leq i \leq n\}$ for $p > 2$.

The algebras $\mathcal{A}$ and $\mathcal{P}$ are both Hopf algebras and $\mathcal{A}(n)$ and $\mathcal{P}(n)$ are subHopf algebras. The description of their duals is standard, and we recall this below. As is standard with the Steenrod algebra and its dual, we use the notation $\mathcal{A}(n)_*$ and $\mathcal{P}(n)_*$ to denote the dual of $\mathcal{A}(n)$ and $\mathcal{P}(n)$, respectively.

Proposition 3.17 *The inclusion $\mathcal{A}(n) \rightarrow \mathcal{A}$ is dual to the quotient*

$$\begin{cases} \mathbb{F}_2[\xi_1, \xi_2, \dots] \rightarrow \mathbb{F}_2[\xi_1, \dots, \xi_{n+1}]/(\xi_1^{2^{n+1}}, \xi_2^{2^n}, \dots, \xi_{n+1}^{2^2}) & \text{when } p = 2, \\ \mathbb{F}_p[\xi_1, \xi_2, \dots] \otimes E(\tau_0, \tau_1, \dots) \rightarrow \mathbb{F}_p[\xi_1, \dots, \xi_n]/(\xi_1^{p^n}, \xi_2^{p^{n-1}}, \dots, \xi_n^p) \otimes E(\tau_0, \dots, \tau_n) & \text{when } p > 2. \end{cases}$$

The inclusion $\mathcal{P}(n) \rightarrow \mathcal{P}$ is dual to the quotient

$$\begin{cases} \mathbb{F}_2[\xi_1^2, \xi_2^2, \dots] \rightarrow \mathbb{F}_2[\xi_1^2, \dots, \xi_{n+1}^2]/((\xi_1^2)^{2^{n+1}}, (\xi_2^2)^{2^n}, \dots, (\xi_{n+1}^2)^2) & \text{when } p = 2, \\ \mathbb{F}_p[\xi_1, \xi_2, \dots] \rightarrow \mathbb{F}_p[\xi_1, \dots, \xi_{n+1}]/(\xi_1^{p^{n+1}}, \xi_2^{p^n}, \dots, \xi_{n+1}^p) & \text{when } p > 2. \end{cases}$$

Remark 3.18 When $p = 2$, the Hopf algebra $\mathcal{P}(n)$ is sometimes referred to as $D\mathcal{A}(n)$ — or the “double” of $\mathcal{A}(n)$ — as there is a degree-doubling isomorphism of Hopf algebras $\mathcal{A} \cong \mathcal{P}$ that restricts to a degree-doubling isomorphism of Hopf algebras $\mathcal{A}(n) \cong \mathcal{P}(n)$.

The key fact that allows one to construct $\mathcal{A}(n)$ - and $\mathcal{P}(n)$ -free complexes is the following theorem of Adams and Margolis [1] at $p = 2$ and Miller and Wilkerson [45] at $p > 2$; see also [52, Section 6.2]. We let P_t^s be the element of $\mathcal{A}$ dual to $\xi_t^{p^s} \in \mathcal{A}_*$ with respect to the monomial basis, and for $p > 2$ we let Q_t denote the element dual to τ_t .

Theorem 3.19 *Let $\mathcal{B}$ be a Hopf subalgebra of the Steenrod algebra $\mathcal{A}$. At $p = 2$, a $\mathcal{B}$ -module M is free if and only if the Margolis homologies $H_*(M; P_t^s) = 0$ for all $P_t^s \in \mathcal{B}$, and for $p > 2$, M is free if and only if $H_*(M; P_t^s) = H_*(M; Q_t) = 0$ for all $P_t^s, Q_t \in \mathcal{B}$.*

The theorem applies verbatim to $\mathcal{B} = \mathcal{A}(n)$ for any prime p and to $\mathcal{B} = \mathcal{P}(n)$ for $p > 2$. For $p = 2$, the theorem also applies to $\mathcal{B} = \mathcal{P}(n) \subset \mathcal{P}$ because the degree-doubling isomorphism of Hopf algebras $\mathcal{A} \cong \mathcal{P}$ sends P_t^s to P_t^{s+1} . As outlined by Ravenel [52], Jeff Smith used idempotents in the group algebra of the symmetric group to construct finite $\mathcal{A}(n)$ -free complexes by way of Theorem 3.19, and we adapt this argument to $\mathcal{P}(n)$.

Proposition 3.20 *Let M be an $\mathcal{A}$ -module on which P_t^0 acts nontrivially for $1 \leq t \leq n$. Then there is an $N \gg 0$ and an idempotent $e \in \mathbb{Z}_{(p)}[\Sigma_N]$ such that*

$$H_*(eM^{\otimes N}; P_t^s) = 0$$

for $s + t \leq n + 1$ and $(s, t) \neq (0, n + 1)$ when $p = 2$, and for $s + t \leq n$ when $p > 2$.

Proof By [52, Appendix C.3.1], each P_t^s and Q_t generates a subalgebra of $\mathcal{A}$ of the form $E(x) = \mathbb{F}_p[x]/x^2$ or $T(x) = \mathbb{F}_p[x]/x^{p^n}$. If M is a free module over $E(x)$ or $T(x)$, the corresponding Margolis homology $H_*(M; x)$ vanishes. The claim now follows immediately as in the proof of [52, Appendix C.3.2], replacing $H^*(X)$ throughout with M . $\square$

Taking $M = H^*\mathbb{C}\mathbb{P}^{p^{n+1}}$, it follows from the Künneth formula that a retract of a smash power of $\mathbb{C}\mathbb{P}^{p^{n+1}}$ is free over $\mathcal{P}(n)$.

Theorem 3.21 *There is a finite complex F such that F is a retract of $(\mathbb{C}\mathbb{P}^{p^{n+1}})^{\otimes N}$ for some N , and $H^*(F; \mathbb{F}_p)$ is a free $\mathcal{P}(n)$ -module.*

Proof For $p > 2$, P_t^0 acts nontrivially on $H^*(\mathbb{CP}^{p^{n+1}}; \mathbb{F}_p)$ for $t \leq n + 1$, hence by combining Proposition 3.20 and Theorem 3.19, there exists a retract F of $(\mathbb{CP}^{p^{n+1}})^{\otimes N}$ for some N such that $H^*(F; \mathbb{F}_p)$ is free over $\mathcal{P}(n)$.

For $p = 2$, P_t^0 acts nontrivially on $H^*(\mathbb{RP}^{2^{n+1}}; \mathbb{F}_2)$ for $t \leq n + 1$, and hence there exists a retract F of $(\mathbb{RP}^{2^{n+1}})^{\otimes N}$ for some N such that $H^*(F; \mathbb{F}_2)$ is free over $\mathcal{A}(n)$. Now, there is a degree-doubling symmetric-monoidal equivalence of categories $\text{Mod}(\mathcal{A}) \simeq \text{Mod}(\mathcal{P})$ which sends $H^*(\mathbb{RP}^{2^{n+1}}; \mathbb{F}_2)$ to $H^*(\mathbb{CP}^{2^{n+1}}; \mathbb{F}_2)$. Since the degree-doubling isomorphism $\mathcal{A} \cong \mathcal{P}$ sends $\mathcal{A}(n)$ to $\mathcal{P}(n)$, it follows that using the same idempotent to split $(\mathbb{CP}^{2^{n+1}})^{\otimes N}$ gives a $\mathcal{P}(n)$ -free summand. $\square$

Remark 3.22 This theorem constructs *some* complex F whose cohomology is free over $\mathcal{P}(n)$, but it can be difficult to work with F in a direct manner. However, in some cases one can do better via a direct construction. The 2-cell complex $\mathbb{S}/\eta$ at $p = 2$ and $\mathbb{S}/\alpha_1$ at $p > 2$ has cohomology free of rank 1 over $\mathcal{P}(0)$. At $p = 2$, one may construct a complex called $D\mathcal{A}(1)$ which is free of rank 1 over $\mathcal{P}(1)$; see [39].

These complexes may be used to construct Wood equivalences for certain fp spectra in a wide generality, which we will study in Section 4. The main fact we use is that, as an $\mathcal{A}(n)_*$ -comodule, $\mathcal{P}(n-1)_*$ is coinduced from an exterior algebra. In the following, recall that $\mathcal{E}(n)_*$ denotes the quotient Hopf algebra of the dual Steenrod algebra given by $E(\xi_1, \dots, \xi_{n+1})$ when $p = 2$, and $E(\tau_0, \dots, \tau_n)$ when $p > 2$.

Proposition 3.23 Suppose that M is an $\mathcal{A}_*$ -comodule concentrated in even degrees and that M is cofree as a $\mathcal{P}(n-1)_*$ -comodule, i.e., that there is an isomorphism of $\mathcal{P}(n-1)_*$ -comodules $M \cong \mathcal{P}(n-1)_* \otimes V$ for some $\mathbb{F}_p$ -vector space V . Then there is an isomorphism of $\mathcal{A}(n)_*$ -comodules

$$M \cong \mathcal{A}(n)_* \square_{\mathcal{E}(n)_*} V.$$

Proof Note that an $\mathcal{A}(n)_*$ -comodule concentrated in even degrees is the same data as a $\mathcal{P}(n-1)_*$ -comodule concentrated in even degrees. The inclusion of $\mathcal{A}(n)_*$ -comodules $\mathcal{P}(n-1)_* \rightarrow \mathcal{A}(n)_*$ factors through an isomorphism

$$\mathcal{P}(n-1)_* \rightarrow \mathcal{A}(n)_* \square_{\mathcal{E}(n)_*} \mathbb{F}_p$$

of $\mathcal{A}(n)_*$ -comodules, as follows from Proposition 3.17. One now uses the isomorphism

$$\left(\mathcal{A}(n)_* \square_{\mathcal{E}(n)_*} \mathbb{F}_p \right) \otimes V \cong \mathcal{A}(n)_* \square_{\mathcal{E}(n)_*} V. \quad \square$$

Note that the $\mathcal{P}(n-1)$ -free complexes constructed in Theorem 3.21 are in fact $T(n)$ -projectives, since $\mathbb{CP}^{p^n}$ is $T(n)$ -free. This relationship is reflected on homology as follows. When $H_*T(n)$ is restricted to $\mathcal{A}(n)_*$, since it is concentrated in even degrees, it is equivalently a $\mathcal{P}(n-1)_*$ -comodule. Just like our $\mathcal{P}(n-1)$ -free complexes, $H_*T(n)$ is cofree as a $\mathcal{P}(n-1)_*$ -comodule.

Proposition 3.24 There is an isomorphism of $\mathcal{A}(n)_*$ -comodules

$$H_*T(n) \cong \mathcal{P}(n-1)_* \otimes \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2] \cong \mathcal{A}(n)_* \square_{\mathcal{E}(n)_*} \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2],$$

where both appearances of $\mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]$ have trivial coactions.

Proof As before, we regard all $\mathcal{A}(n)_*$ -comodules concentrated in even degrees equivalently as $\mathcal{P}(n-1)_*$ -comodules. The map of $\mathcal{A}(n)_*$ -comodule algebras

$$H_*T(n) \rightarrow H_*\mathbb{F}_p = \mathcal{A}_* \rightarrow \mathcal{A}(n)_*$$

has image $\mathcal{P}(n-1)_*$, and thus factors through a surjection of $\mathcal{P}(n-1)_*$ -comodules $H_*T(n) \rightarrow \mathcal{P}(n-1)_*$. By the Milnor–Moore theorem (see for instance [51, A1.1.20]), there is an isomorphism of $\mathcal{P}(n-1)_*$ -comodules $H_*T(n) \cong \mathcal{P}(n-1)_* \otimes (H_*T(n) \square_{\mathcal{P}(n-1)_*} \mathbb{F}_p)$, where the latter has the extended comodule structure (i.e., $H_*T(n) \square_{\mathcal{P}(n-1)_*} \mathbb{F}_p$ is simply regarded as a vector space).

One may now use Proposition 3.17 to compute that

$$H_*T(n) \square_{\mathcal{P}(n-1)_*} \mathbb{F}_p \cong \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]$$

to obtain the first claimed isomorphism, and directly apply Proposition 3.23 to obtain the second. $\square$

Remark 3.25 The splitting of the proposition cannot be made into an isomorphism of $\mathcal{A}_*$ -comodules because $H^*T(n)$ supports arbitrarily many Sq^{2^i} for $p = 2$ and P^{p^i} for $p > 2$.

4 Chromatic defect and Wood equivalences for fp spectra

In this section, we will discuss the behavior of chromatic defect on a special class of spectra called fp spectra, which we now recall.

Definition 4.1 [37, Proposition 3.2] A p -complete, bounded below spectrum E is said to be an fp spectrum if it satisfies one of the following equivalent conditions:

- There exists a finite p -local spectrum F such that $E \otimes F$ is π -finite. That is, $|\pi_*(E \otimes F)| < \infty$.
- There exists a finite p -local spectrum F such that $E \otimes F$ is a finite sum of shifts of $\mathbb{F}_p$.
- There exists a finite $\mathcal{A}(n)_*$ -comodule M and an isomorphism of $\mathcal{A}_*$ -comodules

$$H_*(E; \mathbb{F}_p) \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M.$$

- $H^*(E; \mathbb{F}_p)$ is a finitely presented $\mathcal{A}$ -module.

One often considers fp spectra, such as ko and tmf for example, because they are connective spectra with strong finiteness properties that make them amenable to computation. The isomorphisms in Definition 4.1 give rise to change-of-rings isomorphisms for the Adams spectral sequence (ASS) that give one considerable computational control. The resulting computations are interesting in their own right (see [12] for example) and allow for naturality arguments to deduce behavior in the Adams spectral sequence for more complicated theories, such as the sphere. Moreover, these theories often provide connective models for key finite-height chromatic spectra of interest, such as the $\mathrm{EO}_n(G)$, which are otherwise inaccessible from the point of view of the Adams spectral sequence because their mod p homology vanishes.

It is difficult, however, to say things in general about the behavior of the Adams–Novikov spectral sequence of an fp spectrum, such as the existence of vanishing lines or change-of-rings isomorphisms. We discuss some results of this form in this section in certain nice cases of fp spectra, by giving obstructions in the Adams spectral sequence to an fp spectrum having finite chromatic defect, and by using the $\mathcal{P}(n)$ -free complexes of Theorem 3.21 to give sufficient conditions for an fp spectrum to be Wood-type.

4.1 Chromatic defect for fp spectra

Proposition 3.24 allows us to use a change-of-rings isomorphism to identify the homology of $E \otimes T(n)$ when E is an fp spectrum.

Proposition 4.2 *Suppose that E is an fp spectrum such that $H_*(E; \mathbb{F}_p) \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$. There is an isomorphism of $\mathcal{A}_*$ -comodules*

$$H_*(E \otimes T(n)) \cong \left(\mathcal{A}_* \square_{\mathcal{E}(n)_*} M \right) \otimes \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2],$$

where $\mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]$ has trivial coaction. In particular, the ASS of $E \otimes T(n)$ has signature

$$E_2 = \text{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, M) \otimes \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2] \Rightarrow \pi_*(E \otimes T(n)).$$

Proof The first claimed isomorphism follows directly from the Künneth isomorphism and Proposition 3.24. A change-of-rings isomorphism (as in [51, A1.3.13]) for the Hopf algebra quotient $\mathcal{A}_* \rightarrow \mathcal{E}(n)_*$ now implies an isomorphism

$$\text{Ext}_{\mathcal{A}_*}(\mathbb{F}_p, \left(\mathcal{A}_* \square_{\mathcal{E}(n)_*} M \right) \otimes \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]) \cong \text{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, M \otimes \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]).$$

Finally, since $\mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]$ has trivial coaction, it follows directly from the cobar complex that

$$\text{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, M \otimes \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]) \cong \text{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, M) \otimes \mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]. \quad \square$$

The upshot of this change-of-rings isomorphism is that computing Ext over the exterior algebra $\mathcal{E}(n)_*$ is much easier than over $\mathcal{A}(n)_*$. If E is a homotopy associative ring spectrum, we can fit the obstructions of Proposition 3.9 into this picture and obtain the following.

Corollary 4.3 *If E is an fp homotopy associative ring spectrum such that $H_*(E; \mathbb{F}_p) \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$, and*

$$\text{Ext}_{\mathcal{E}(n)_*}^{s, 2p^{m+1}-3-2(p-1)*+s}(\mathbb{F}_p, M) = 0$$

for $s \geq 2$ and $m \geq n$, then $\Phi_p(E) \leq n$.

Proof The obstructions to having $\Phi_p(E) \leq n$ lie in the $2p^{m+1}-3$ stems for $m \geq n$ by Proposition 3.9, and the graded vector space $\mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2]$ is concentrated in degrees $2(p-1)*$, so by Proposition 4.2, it remains only to explain why one may take $s \geq 2$. This is because the obstruction χ_{m+1} is detected by $h_{m+1,1}$ when $p = 2$, and $h_{m+1,0}$ when $p > 2$ in the cobar complex for $H_*T(m)$. The map

$$\text{Ext}_{\mathcal{A}_*}(\mathbb{F}_p, H_*T(m)) \rightarrow \text{Ext}_{\mathcal{A}(n)_*}(\mathbb{F}_p, H_*T(m))$$

sends this class to zero when $m \geq n$. Indeed, this follows from the fact that the quotient map $\mathcal{A}_* \rightarrow \mathcal{A}(n)_*$ sends $\xi_{m+1}^2 \mapsto 0$ when $p = 2$ and $\xi_{m+1} \mapsto 0$ when $p > 2$ for $m \geq n$, by Proposition 3.17. Since χ_{m+1} has Adams filtration 1 in $\pi_* T(m)$, it must therefore be detected in filtration > 1 in $\pi_* E \otimes T(m)$. $\square$

Remark 4.4 In the preceding proposition and corollary, and throughout this section, we often fix an isomorphism $H_*(E; \mathbb{F}_p) \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$ for an fp spectrum E and study the condition that $\Phi_p(E) \leq n$. Asking that these two numbers agree does not result in much loss of generality: indeed, if there is an isomorphism $H_*(E; \mathbb{F}_p) \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$, then for any $m \geq n$, there is an isomorphism $H_*(E; \mathbb{F}_p) \cong \mathcal{A}_* \square_{\mathcal{A}(m)_*} M'$. Similarly, of course, if $\Phi_p(E) \leq n$, then $\Phi_p(E) \leq m$ as in Remark 2.16.

Example 4.5 Suppose that E is an fp homotopy associative ring spectrum and that M is concentrated in even degrees. It follows that M has trivial coaction as an $\mathcal{E}(n)_*$ -comodule, since the generators of $\mathcal{E}(n)_*$ are in odd degrees. Therefore

$$\mathrm{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, M) \cong \mathrm{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, \mathbb{F}_p) \otimes M,$$

and since M is concentrated in even degrees t and $s = 0$, and $\mathrm{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, \mathbb{F}_p)$ is concentrated in bidegrees (s, t) where $t - s$ is even, it follows that $\mathrm{Ext}_{\mathcal{E}(n)_*}^{s,t}(\mathbb{F}_p, M) = 0$ bidegrees (s, t) where $t - s$ is even. By Corollary 4.3, we then have that $\Phi_p(E) \leq n$.

We can in this way immediately recover Hopkins' examples from Example 2.17 using that

$$H_* \mathrm{ko} = \mathcal{A}_* \square_{\mathcal{A}(1)_*} \mathbb{F}_2 \quad \text{and} \quad H_* \mathrm{tmf} = \mathcal{A}_* \square_{\mathcal{A}(2)_*} \mathbb{F}_2.$$

Example 4.6 One may take the spectrum $\mathrm{BP}_{\mathbb{R}} \langle 2 \rangle^{C_2}$ — where $\mathrm{BP}_{\mathbb{R}} \langle 2 \rangle$ is any form of the second truncated Real Brown–Peterson spectrum (e.g., $\mathrm{tmf}_0(3) = \mathrm{tmf}_1(3)^{C_2}$ by [26]) — as an example where M is not itself even but $\mathrm{Ext}_{\mathcal{E}(n)_*}(\mathbb{F}_p, M)$ is concentrated in even stems. Indeed, by [15, Theorem 1.9], we have a decomposition

$$H_*(\mathrm{BP}_{\mathbb{R}} \langle 2 \rangle^{C_2}) \cong \mathcal{A}_* \square_{\mathcal{A}(2)_*} M,$$

where M is a certain 10-dimensional comodule described in [loc. cit.] It can be checked by hand or with an Ext resolver that $\mathrm{Ext}_{\mathcal{E}(2)_*}(\mathbb{F}_p, M)$ is even.

4.2 Wood-type fp spectra

The description of $H_*(E \otimes T(n))$ of Proposition 4.2 may be used to give an algebraic Wood equivalence for any fp spectrum E with finite chromatic defect. In this section, we discuss these algebraic equivalences and a sufficient condition to lift them to the spectrum level. The algebraic equivalence is a simple consequence of Proposition 3.23. In the following statement, we fix an $\mathcal{A}_*$ -comodule structure extending the canonical $\mathcal{A}(n)_*$ -comodule structure on $\mathcal{P}(n-1)_*$. Such structures exist by [46, Theorem A].

Proposition 4.7 *Let E be an fp spectrum so that $H_* E \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$. Then there is an isomorphism of $\mathcal{A}_*$ -comodules*

$$H_* E \otimes \mathcal{P}(n-1)_* \cong \mathcal{A}_* \square_{\mathcal{E}(n)_*} M.$$

In particular, by Proposition 3.24, $H_*E \otimes \mathcal{P}(n-1)_*$ is a retract of $H_*(E \otimes T(n))$ via the projection $\mathbb{F}_p[t_1^{2^n}, t_2^{2^{n-1}}, \dots, t_n^2] \rightarrow \mathbb{F}_p\{1\}$. Hence if E has finite chromatic defect, H_*E is Wood-type in the sense that $\mathcal{P}(n-1)_*$ is a finite comodule concentrated in even degrees, and $H_*(E) \otimes \mathcal{P}(n-1)_*$ is a weak $H_*\text{MU}$ -module.

Remark 4.8 This algebraic Wood equivalence can be given a proper home in Hovey's category $\text{Stable}(\mathcal{A})$, equivalently, by working mod τ in $\text{Syn}_{\mathbb{F}_p}$, Pstrągowski's category [48]. However, we will not make use of this, and therefore comment on it only briefly.

Indeed one may define complex-orientability in terms of the attaching maps $\sigma_k : \mathbb{S}^{2k+1,1} \rightarrow \nu\mathbb{C}\mathbb{P}^k$ as in Definition 2.3, and one has an immediate analog of Theorem 2.9 in both $\text{Syn}_{\mathbb{F}_p}$ and $\text{Stable}(\mathcal{A})$ with νMU . According to the analogous definition of Wood-types using compact νBP -projectives, if E is an fp spectrum with finite chromatic defect, the above proposition then implies that $\nu E/\tau$ is Wood-type in $\text{Stable}(\mathcal{A})$.

To lift these algebraic Wood equivalences to the spectrum level, we will need two main ingredients. First we must lift the comodule $\mathcal{P}(n-1)_*$ to a finite BP-projective; this is accomplished, up to taking shifts and direct sums, by the complexes of Theorem 3.21. Then we must lift the classes in $\mathbb{F}_p[t_1^{2^n}, \dots, t_n^2]$. Such lifts are guaranteed by the collapse of $\text{ASS}(E \otimes \text{BP})$.

Theorem 4.9 *Let E be an fp homotopy associative ring spectrum with finite chromatic defect. If $\text{ASS}(E \otimes \text{BP})$ collapses on E_2 , then E is Wood-type.*

Proof Since E is an fp spectrum, we have an isomorphism $H_*E \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$, and we have assumed $\Phi_p(E) \leq m$ for some m . As in Remark 4.4, we can take $m = n$. By Proposition 2.24, we have a splitting $E \otimes \text{BP} \simeq E \otimes T(n)[t_{n+1}, t_{n+2}, \dots]$ so that $\text{ASS}(E \otimes \text{BP})$ collapses on E_2 if and only if $\text{ASS}(E \otimes T(n))$ does.

Now we fix a finite complex F with the property that F is a retract of $(\mathbb{C}\mathbb{P}^{p^n})^{\otimes N}$ for some N and that $H^*(F; \mathbb{F}_p)$ is a free $\mathcal{P}(n-1)$ -module, using Theorem 3.21. As a retract of $(\mathbb{C}\mathbb{P}^{p^n})^{\otimes N}$, F is a finite BP-projective, and the AHSS computing $[F, T(n)]$ collapses on E_2 , as it is a retract of the corresponding spectral sequence for $(\mathbb{C}\mathbb{P}^{p^n})^{\otimes N}$. Fixing a basis $\{b_1, \dots, b_k\}$ of the free $\mathcal{P}(n-1)$ -module $H^*(F; \mathbb{F}_p)$ with $|b_i| = n_i$, we thereby fix a map $\iota : F \rightarrow T(n)\{b_1, \dots, b_k\}$ lifting the element in

$$\bigoplus_j H^j(F; T(n)\{b_1, \dots, b_k\}^{-j}) \cong \bigoplus_j \text{Hom}_{\mathbb{Z}}(H_j(F), T(n)\{b_1, \dots, b_k\}_j)$$

that sends $b_i \mapsto b_i$, and is zero on all other basis elements, using that $H_j(F)$ is a finitely generated free abelian group.

One may determine the effect of the map ι in homology as follows. The composition

$$\Sigma^{n_i} \mathcal{P}(n-1)_* \xrightarrow{b_i} H_*F \xrightarrow{\iota_*} H_*T(n)\{b_1, \dots, b_k\} \xrightarrow{b_j} \Sigma^{n_j} H_*T(n)$$

is either zero or an isomorphism in the bottom dimension of the target. Over $\mathcal{A}(n)_*$, the target is cofree on $\mathbb{F}_p[t_1^{2^n}, \dots, t_n^2]$, so when the above map is nonzero, its projection onto the cocyclic summand indexed

by 1 must agree with that of the canonical inclusion $\Sigma^{n_i} \mathcal{P}(n-1)_* \rightarrow \Sigma^{n_i} H_* T(n)$ of Proposition 3.24. Using this map along with the ring structure of $H_* T(n)$, we have a map $\Sigma^{n_i} \mathcal{P}(n-1)_*[t_1^{2^n}, \dots, t_n^{2^n}] \rightarrow H_* T(n)\{b_i\}$, which is therefore surjective by Proposition 3.24 and hence an isomorphism for dimension reasons.

Finally, since $\text{ASS}(E \otimes T(n))$ collapses on E_2 , there are classes $t_1^{2^n}, \dots, t_n^{2^n} \in \pi_*(E \otimes T(n))$ lifting the corresponding classes in homology provided by Proposition 4.2. Using these classes, the map ι , and the multiplication on $E \otimes T(n)$, we may define a map

$$(E \otimes F)[t_1^{2^n}, \dots, t_n^{2^n}] \rightarrow E \otimes T(n)\{b_1, \dots, b_k\}$$

that induces an isomorphism in homology by the above argument, and it is thus an equivalence. It follows that $E \otimes F$ is a retract of $E \otimes T(n)\{b_1, \dots, b_k\}$ and therefore complex-orientable. $\square$

Corollary 4.10 *Let E be an fp homotopy associative ring spectrum so that $H_* E \cong \mathcal{A}_* \square_{\mathcal{A}(n)_*} M$ and suppose that*

$$\text{Ext}_{\mathcal{E}(n)_*}^{s,t}(\mathbb{F}_p, M) = 0$$

whenever $t - s$ is odd. Then $\Phi_p(E) \leq n$ and E is Wood-type.

Proof The evenness conditions guarantee that $\Phi_p(E) \leq n$ by Corollary 4.3, and therefore $E \otimes \text{BP}$ is a sum of shifts of $E \otimes T(n)$ by Proposition 2.24. The latter implies that if $\text{ASS}(E \otimes T(n))$ collapses, so does $\text{ASS}(E \otimes \text{BP})$. To see that $\text{ASS}(E \otimes T(n))$ collapses, by Proposition 4.2 $\text{Ext}_{\mathcal{A}_*}^{s,t}(\mathbb{F}_p, H_*(E \otimes T(n))) = 0$ unless $t - s$ is even, so there is no room for differentials. $\square$

Example 4.11 As explained in the previous section, ko , tmf and $\text{BP}_{\mathbb{R}}\langle 2 \rangle^{C_2}$ all satisfy the conditions of Corollary 4.10, and thus are seen to be Wood-type. As a nonexample, the connective image of J spectrum j is an fp spectrum, but is not Wood-type as it does not have finite chromatic defect; we will show this in Section 5.

Remark 4.12 The conditions of the above Theorem 4.9, however, are not necessary. For example, the spectrum $ku/4$ is a Wood-type (complex-orientable, in fact) fp spectrum with the property that $\text{ASS}(E \otimes \text{BP})$ has nonzero differentials. It seems plausible that the conditions of the theorem may be relaxed to asking for collapse on a finite page.

5 Hopkins' stacks and the $X(n)$

An fp spectrum E admits a change-of-rings isomorphism

$$\text{Ext}_{\mathcal{A}_*}(\mathbb{F}_p, H_* X) \cong \text{Ext}_{\mathcal{A}(n)_*}(\mathbb{F}_p, M)$$

on the E_2 -page of the classical Adams spectral sequence of E , for some n and M . We will see in this section that spectra with finite chromatic defect play a similar role with respect to the Adams–Novikov spectral sequence. In this setting, the relevant change-of-rings isomorphisms take a more conceptual form, via the language of stacks.

5.1 Chromatic defect and the Adams–Novikov spectral sequence

Hopkins observed that many of the known change-of-rings isomorphisms in the Adams–Novikov spectral sequence could be reinterpreted and more conceptually derived in the language of stacks [20, Chapter 9]. To any homotopy commutative ring spectrum E , we may associate a stack $\mathcal{M}_E$ equipped with a canonical $\mathbb{G}_m$ -action and a $\mathbb{G}_m$ -equivariant affine morphism $p_E : \mathcal{M}_E \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$, where $\mathcal{M}_{\mathrm{FG}}(1)$ is the moduli stack of formal group laws and strict isomorphisms, viewed as a $\mathbb{G}_m$ -torsor over $\mathcal{M}_{\mathrm{FG}}$, the moduli stack of formal group laws and *all* isomorphisms.

Definition 5.1 Let E be a homotopy commutative ring spectrum. We let the $\mathbb{G}_m$ -stack $\mathcal{M}_E$ be the one associated to the graded Hopf algebroid

$$(\mathrm{MU}_{2*}E, \mathrm{MU}_{2*}(\mathrm{MU} \otimes E)).$$

Remark 5.2 It is often useful to mod out by the $\mathbb{G}_m$ -action and get an affine morphism

$$\mathcal{M}_E / \mathbb{G}_m \rightarrow \mathcal{M}_{\mathrm{FG}}.$$

We will instead primarily work $\mathbb{G}_m$ -equivariantly so that when E is complex-orientable, $\mathcal{M}_E$ is an affine scheme, as opposed to a stack of the form $\mathrm{Spec}(E_{2*})/\mathbb{G}_m$. This is convenient as we wish to capture more general, non-even-periodic cases, such as those arising from Johnson–Wilson theories $E(n)$.

By Quillen’s theorem [49] on MU_* , one has a $\mathbb{G}_m$ -equivariant equivalence $\mathcal{M}_{\mathbb{S}} \simeq \mathcal{M}_{\mathrm{FG}}(1)$, whereby the map p_E is induced by the unit map $\mathbb{S} \rightarrow E$ for E a homotopy commutative ring spectrum. More generally, *any* spectrum E gives rise to a pair of $\mathbb{G}_m$ -equivariant quasicoherent sheaves $\mathcal{F}_0(E)$ and $\mathcal{F}_1(E)$ on $\mathcal{M}_{\mathrm{FG}}(1)$, corresponding to the comodules $\mathrm{MU}_{2*}E$ and $\mathrm{MU}_{2*+1}E$. When E is a homotopy-commutative ring spectrum, the stack $\mathcal{M}_E$ is realized as a relative Spec construction on the sheaf of algebras $\mathcal{F}_0(E)$, and the sheaves $\mathcal{F}_0(E)$ and $\mathcal{F}_1(E)$ are pushed forward from sheaves on $\mathcal{M}_E$. This gives our change-of-rings isomorphisms.

Proposition 5.3 *There is a graded isomorphism of the E_2 -page of $\mathrm{ANSS}(E)$*

$$E_2^{*,*} \cong H^*(\mathcal{M}_{\mathrm{FG}}(1); \mathcal{F}_0(E) \oplus \mathcal{F}_1(E)) \cong H^*(\mathcal{M}_E; \mathbb{O}_{\mathcal{M}_E} \oplus \mathbb{O}_1),$$

where $\mathbb{O}_1$ is the $\mathbb{G}_m$ -equivariant quasicoherent sheaf on $\mathcal{M}_E$ defined by the comodule $\mathrm{MU}_{2*+1}E$.

Proof This follows by directly analyzing the cobar complex, whose cohomology is the E_2 -page of the ANSS of E and arises from the cosimplicial object

$$\mathrm{MU}_*E \rightrightarrows \mathrm{MU}_*\mathrm{MU} \otimes_{\mathrm{MU}_*} \mathrm{MU}_*E \rightrightarrows \cdots.$$

Since the Hopf algebroid $(\mathrm{MU}_*, \mathrm{MU}_*\mathrm{MU})$ is concentrated in even degrees, this cosimplicial object splits into an even piece

$$\mathrm{MU}_{2*}E \rightrightarrows \mathrm{MU}_*\mathrm{MU} \otimes_{\mathrm{MU}_*} \mathrm{MU}_{2*}E \rightrightarrows \cdots$$

and an odd piece

$$\mathrm{MU}_{2*+1}E \rightrightarrows \mathrm{MU}_*\mathrm{MU} \otimes_{\mathrm{MU}_*} \mathrm{MU}_{2*+1}E \rightrightarrows \cdots.$$

The identification $\mathcal{M}_{\mathrm{FG}}(1) \simeq \mathcal{M}_{(\mathrm{MU}_*, \mathrm{MU}_* \mathrm{MU})}$ gives the first isomorphism.

For the second identification, the cohomology of the even piece is also $H^*(\mathcal{M}_E; \mathbb{C}_{\mathcal{M}_E})$ by definition of $\mathcal{M}_E$, and the odd piece is identified with the cohomology of a sheaf on $\mathcal{M}_E$ whose underlying $\mathrm{MU}_{2*}E$ -module is $\mathrm{MU}_{2*+1}E$ via the isomorphism

$$\mathrm{MU}_* \mathrm{MU} \otimes_{\mathrm{MU}_*} \mathrm{MU}_{2*+1}E \cong (\mathrm{MU}_* \mathrm{MU} \otimes_{\mathrm{MU}_*} \mathrm{MU}_{2*}E) \otimes_{\mathrm{MU}_{2*}E} \mathrm{MU}_{2*+1}E \quad \square$$

Example 5.4 When $E = \mathrm{ko}$, the stack $\mathcal{M}_{\mathrm{ko}}$ is given by the moduli stack of quadratic equations $\mathcal{M}_{\mathrm{quad}}$ with strict coordinate transformations. When $E = \mathrm{tmf}$, the stack $\mathcal{M}_{\mathrm{tmf}}$ is given by the moduli stack of cubic equations $\mathcal{M}_{\mathrm{cub}}$ with strict coordinate transformations; see [20, Chapter 9] or [39]. These spectra have the property that $\mathrm{MU}_{2*+1}E = 0$, so that for instance $\mathrm{ANSS}(\mathrm{tmf})$ takes the form

$$H^*(\mathcal{M}_{\mathrm{cub}}; \mathbb{C}_{\mathcal{M}_{\mathrm{cub}}}) \Rightarrow \pi_* \mathrm{tmf},$$

which, in this form, is sometimes called the elliptic spectral sequence.

The construction of Definition 5.1 has two features that are especially useful. It has the property that it relates vanishing lines on the E_2 -page of the ANSS to cohomological vanishing of sheaves over $\mathcal{M}_E$, and it can often be used to determine the E_2 -page of the ANSS of a tensor product of spectra in terms of the corresponding stacky pullback over $\mathcal{M}_{\mathrm{FG}}(1)$.

Proposition 5.5 *Let E and E' be homotopy commutative ring spectra.*

- (1) *If E is complex-orientable, $\mathcal{M}_E \simeq \mathrm{Spec}(E_{2*})$, and p_E is the map classifying the formal group over E_* .*
- (2) *Suppose that $\mathrm{MU}_{2*}E'$ is a flat MU_* -module. Then one has a pullback square*

$$\begin{array}{ccc} \mathcal{M}_{E \otimes E'} & \longrightarrow & \mathcal{M}_{E'} \\ \downarrow & & \downarrow p_{E'} \\ \mathcal{M}_E & \xrightarrow{p_E} & \mathcal{M}_{\mathrm{FG}}(1) \end{array}$$

Proof The first claim follows as in Corollary 2.10 from the (-1) -st codegeneracy furnished by the MU -module structure in the coaugmented cosimplicial object

$$E \rightarrow \mathrm{MU} \otimes E \Rightarrow \mathrm{MU} \otimes \mathrm{MU} \otimes E \Rightarrow \mathrm{MU} \otimes \mathrm{MU} \otimes \mathrm{MU} \otimes E \cdots$$

The second claim follows by a Hopf algebroid manipulation using that $p_{E'}$ is an affine morphism; see [13, Proposition 6.1.6] or [10, 1.6.6]. $\square$

Claim (1) in Proposition 5.5 implies in particular that the stackiness of $\mathcal{M}_E$ is a measure of the failure of E to be complex-orientable, and indeed our chromatic defect of Definition 2.15 gives a way of quantifying this. In fact, the definition of chromatic defect is chosen precisely to reflect a specific algebrogeometric measure of the stackiness of $\mathcal{M}_E$ over $\mathcal{M}_{\mathrm{FG}}(1)$, and the connection comes via an identification of $\mathcal{M}_{X(n)}$.

Definition 5.6 We let $\mathcal{M}_{\text{FG}}(n)$ denote the moduli stack of formal groups together with an n -jet. That is,

$$\mathcal{M}_{\text{FG}}(n)(R) = \begin{cases} \text{objects:} & \text{formal group laws over } R, \\ \text{morphisms:} & f : F \rightarrow G \text{ such that } f(x) \equiv x \pmod{x^{n+1}}. \end{cases}$$

Proposition 5.7 *There is a $\mathbb{G}_m$ -equivariant equivalence of stacks $\mathcal{M}_{\text{FG}}(n) \simeq \mathcal{M}_{X(n)}$ over $\mathcal{M}_{\text{FG}}(1)$. In particular, the map $\mathcal{M}_{\text{FG}}(n) \rightarrow \mathcal{M}_{\text{FG}}(1)$ is a faithfully flat affine morphism with the property that*

$$\begin{array}{ccc} \mathcal{M}_{E \otimes X(n)} & \longrightarrow & \mathcal{M}_{\text{FG}}(n) \\ \downarrow & & \downarrow \\ \mathcal{M}_E & \longrightarrow & \mathcal{M}_{\text{FG}}(1) \end{array}$$

is a pullback, for E a homotopy commutative ring spectrum.

Proof By the Thom isomorphism, one has an isomorphism of Hopf algebroids

$$(\text{MU}_* X(n), \text{MU}_*(\text{MU} \otimes X(n))) \cong (\text{MU}_*[b_1, \dots, b_{n-1}], \text{MU}_* \text{MU}[b_1, \dots, b_{n-1}])$$

The map of Hopf algebroids

$$(\text{MU}_* X(n), \text{MU}_*(\text{MU} \otimes X(n))) \rightarrow (\text{MU}_* \text{MU}, \text{MU}_*(\text{MU} \otimes \text{MU})).$$

is an inclusion, so we may compute $\eta_R(b_i)$ in the latter. In $\text{MU}_* \text{MU}$, the b_i are by definition the coefficients of the canonical strict isomorphism $\eta_L^* F \rightarrow \eta_R^* F$, where F is the universal formal group law over MU_* . Note that a map of rings

$$\phi : \text{MU}_* \text{MU} \otimes_{\text{MU}_*} \text{MU}_* \text{MU} \rightarrow R$$

corresponds to a sequence

$$F_1 \xrightarrow{g} F_2 \xrightarrow{f} F_3$$

of isomorphisms of formal group laws over R , where F_i is the pushforward of F_{MU} along the map $\text{MU} \rightarrow \text{MU} \otimes \text{MU} \otimes \text{MU}$ including MU as the i -th tensor factor. Since η_R in the Hopf algebroid $(\text{MU}_* \text{MU}, \text{MU}_*(\text{MU} \otimes \text{MU}))$ is induced by the map

$$\text{MU} \otimes \text{MU} \simeq \text{MU} \otimes \mathbb{S} \otimes \text{MU} \rightarrow \text{MU} \otimes \text{MU} \otimes \text{MU},$$

we have that $\phi(\eta_L(b_i))$ are the coefficients of g , and $\phi(\eta_R(b_i))$ are the coefficients of $f \circ g$. Therefore, if $\eta_L(b_i) = \eta_R(b_i)$ for $i \leq m-1$, then $f(x) \equiv x \pmod{x^{n+1}}$. It follows that the map

$$\mathcal{M}_{X(n)} \rightarrow \mathcal{M}_{\text{FG}}(1)$$

factors through an equivalence onto the substack $\mathcal{M}_{\text{FG}}(n)$. □

Corollary 5.8 *Let E be a homotopy commutative ring spectrum. If $\Phi(E) \leq n$, then the pullback*

$$\mathcal{M}_E \times_{\mathcal{M}_{\text{FG}}(1)} \mathcal{M}_{\text{FG}}(n)$$

is an affine scheme. Conversely, if $E \otimes X(n)$ is MU nilpotent complete, and the above pullback is an affine scheme, then $\Phi(E) \leq n$.

Proof The first claim follows from Proposition 5.5(1) and Proposition 5.7. For the second, since the above pullback is an affine scheme, we see that $\text{ANSS}(\mathbb{CP}^k \otimes E \otimes X(n))$ collapses on the zero line for all k by Proposition 5.3. Since $\text{MU} \otimes \sigma_k = 0$ for all k , the class σ_k has Adams–Novikov filtration > 0 , and hence must be zero in $\pi_* \mathbb{CP}^k \otimes E \otimes X(n)$ since $\text{ANSS}(\mathbb{CP}^k \otimes E \otimes X(n))$ converges, as $E \otimes X(n)$ is MU nilpotent complete. This implies that $E \otimes X(n)$ is complex-orientable by definition since $E \otimes X(n)$ is a ring spectrum. $\square$

5.2 Chromatic defect and quotient stacks

We would like sufficient algebrogeometric conditions on the morphism $\mathcal{M}_E \rightarrow \mathcal{M}_{\text{FG}}(1)$ that guarantee finite chromatic defect for E , by way of Corollary 5.8. One can give various finiteness conditions to guarantee this, which apply in particular when $\mathcal{M}_E$ is a quotient stack by a finite group action and also in various connective cases such as $\mathcal{M}_{\text{ko}}$ and $\mathcal{M}_{\text{tmf}}$ from Example 5.4. We begin with a necessary condition.

Proposition 5.9 *Let E be a homotopy commutative ring spectrum with finite chromatic defect. Then, for any algebraically closed field k and any $x \in \mathcal{M}_E(k)$ such that $p_E(x) \in \mathcal{M}_{\text{FG}}(1)(k)$ has finite height, the image of the homomorphism*

$$\psi : \text{Aut}_{\mathcal{M}_E(k)}(x) \rightarrow \text{Aut}_{\mathcal{M}_{\text{FG}}(1)(k)}(p_E(x))$$

is a finite subgroup.

Proof First suppose k has characteristic zero; then $p_E(x)$ is isomorphic to $\widehat{\mathbb{G}_a}$, which has a trivial and therefore finite strict automorphism group. If k has characteristic $p > 0$, then since $k = \bar{k}$, Lazard's theorem implies that $p_E(x)$ is isomorphic to the Honda formal group law F_n for some height $n < \infty$, all of whose automorphisms are defined over the finite field $\mathbb{F}_{p^n}$; see [36, Lecture 19].

Since E has finite chromatic defect, it follows from Corollary 5.8 that for some $N > 0$, the stack

$$\mathcal{M}_E \times_{\mathcal{M}_{\text{FG}}(1)} \mathcal{M}_{\text{FG}}(N)$$

is an affine scheme, and therefore a discrete stack. This implies in particular that if $f(x)$ is in the image of ψ , and $f(x) \equiv x \pmod{x^{N+1}}$, then $f(x) = x$. Suppose for the sake of contradiction that the image of ψ is infinite. Since all the automorphisms of F_n are defined over the finite field $\mathbb{F}_{p^n}$, it follows that for any m , there are two *distinct* automorphisms $f_1(x) \neq f_2(x)$ in the image of ψ such that $f_1(x) \equiv f_2(x) \pmod{x^{m+1}}$. Indeed, were this not true, then each automorphism in the image of ψ would be uniquely determined by its truncation mod x^{m+1} , which contradicts that the image of ψ is infinite, as there are only finitely many such polynomials over a finite field.

Letting $m = N$, and choosing f_1, f_2 as above, the automorphism $g := f_1 \circ f_2^{-1}$ has the property that $g(x) \equiv x \pmod{x^{N+1}}$, but $g(x) \neq x$, contradicting discreteness of the above pullback. $\square$

To obtain a converse, one must ask for conditions on the morphism $p_E : \mathcal{M}_E \rightarrow \mathcal{M}_{\text{FG}}(1)$ guaranteeing that

$$\mathcal{P}(n) := \mathcal{M}_E \times_{\mathcal{M}_{\text{FG}}(1)} \mathcal{M}_{\text{FG}}(n)$$

is an affine scheme for some n . If the image of the map ψ appearing in Proposition 5.9 is always finite, $\mathcal{P}(n)$ still may not be discrete. For instance, if $\mathcal{M}_E$ has height ∞ points, or if the size of the image of ψ is not uniformly bounded above, the stack $\mathcal{P}(n)$ may not be discrete. One strong condition that guarantees affineness, however, is when $\mathcal{M}_E$ is the quotient of a Noetherian ring by a finite group.

Proposition 5.10 *Let E be a homotopy commutative ring spectrum such that $E \otimes X(n)$ is MU -nilpotent complete for all n sufficiently large. Suppose there is an equivalence $\mathcal{M}_E \simeq \mathrm{Spec}(R)/G$ for G a finite group and R a Noetherian commutative ring. Then E has finite chromatic defect.*

Proof Each $g \in G$ acts on the formal group over R via a power series

$$[g](x) = x + \sum_{n \geq 0} b_n(g)x^{n+1},$$

and the image of any automorphism of $\mathrm{Spec}(R)/G$ in $\mathcal{M}_{\mathrm{FG}}(1)$ is obtained from these power series by base change. For each g and n , consider the ideal in R

$$J_n(g) := (b_i(g) : i < n).$$

Since R is Noetherian, the $J_n(g)$ must stabilize for each g . It follows from the fact that G is finite that, by taking n large enough, if after base change to any ring S one has $[g](x) \equiv x \pmod{x^{n+1}}$, then $[g](x) = x$ over S . This implies the stack

$$\mathcal{M}_E \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathcal{M}_{\mathrm{FG}}(n)$$

is discrete. However, if $\mathcal{M}$ is any discrete stack with an affine morphism to $\mathrm{Spec}(R)/G$, then $\mathcal{M}$ is affine. This follows for example from [31, Theorem A7.1.1]; see also [13, Lemma 6.2.6]. The claim then follows from Corollary 5.8. $\square$

An identification $\mathcal{M}_E \simeq \mathrm{Spec}(R)/G$ suitable to apply Proposition 5.10 often corresponds to an identification $E = R^{hG}$, for R a homotopy commutative ring spectrum with G -action. In fact, the former identification follows from the latter under the following assumptions.

Proposition 5.11 *Let $R \in \mathrm{Sp}^{BG}$ be a homotopy commutative ring spectrum with G -action so that R is complex-orientable as a spectrum, and set $E = R^{hG}$. Suppose the map p in the diagram*

$$\begin{array}{ccc} \mathrm{Spec}(R_{2*}) & \longrightarrow & \mathcal{M}_{\mathrm{FG}}(1) \\ \downarrow & \nearrow p & \\ \mathrm{Spec}(R_{2*})/G & & \end{array}$$

is affine. Then if $\mathrm{MU}^{\otimes s} \otimes E \simeq (\mathrm{MU}^{\otimes s} \otimes R)^{hG}$ for $s = 1, 2$, there is an equivalence of stacks $\mathcal{M}_E \simeq \mathrm{Spec}(R_{2*})/G$.

Proof Set $\mathcal{N} := \mathrm{Spec}(R_{2*})/G$. By affineness of p , we have a diagram of pullback squares

$$\begin{array}{ccc}
 \mathrm{Spec}(\Gamma(\mathcal{N} \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathrm{Spec}(L) \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathrm{Spec}(L); \mathbb{O})) & \longrightarrow & \mathrm{Spec}(\Gamma(\mathcal{N} \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathrm{Spec}(L); \mathbb{O})) \\
 \downarrow & & \downarrow \\
 \mathrm{Spec}(\Gamma(\mathcal{N} \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathrm{Spec}(L); \mathbb{O})) & \longrightarrow & \mathcal{N} \\
 \downarrow & & \downarrow p \\
 \mathrm{Spec}(L) & \longrightarrow & \mathcal{M}_{\mathrm{FG}}(1)
 \end{array}$$

and

$$\Gamma(\mathcal{N} \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathrm{Spec}(L); \mathbb{O}) \cong \Gamma(\mathcal{N}; p^* \mathcal{F}_{\mathrm{MU}}) \cong H^0(G; \mathrm{MU}_{2*} R) = \mathrm{MU}_{2*} E.$$

The last identification follows from the HFPSS

$$H^*(G; \mathrm{MU}_* R) \Rightarrow \mathrm{MU}_* E$$

using the equivalence $\mathrm{MU} \otimes E \simeq (\mathrm{MU} \otimes R)^{hG}$. Indeed, the higher cohomology on this E_2 -page vanishes by a change-of-rings isomorphism using affineness of the pullback $\mathcal{N} \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathrm{Spec}(L)$. Similar considerations apply to the pullback in the upper left corner.

The top right stack is thereby identified with $\mathcal{M}_{\mathrm{MU} \otimes E}$, and the top left stack with $\mathcal{M}_{\mathrm{MU} \otimes \mathrm{MU} \otimes E}$. This gives the identification

$$\mathcal{N} \simeq \mathcal{M}_{((\mathrm{MU} \otimes E)_{2*}, (\mathrm{MU} \otimes \mathrm{MU} \otimes E)_{2*})} = \mathcal{M}_E. \quad \square$$

As remarked in [40, Section 6.2], if R is a Landweber exact $\mathbb{E}_\infty$ -ring, then the map p in Proposition 5.11 is affine precisely when, for every field valued point x of $\mathrm{Spec}(R_{2*})$, the stabilizer of x in G acts faithfully on $p(x)$. Moreover, the map $E \rightarrow R$ must then be a faithful Galois extension so that the condition $\mathrm{MU}^{\otimes s} \otimes E \simeq (\mathrm{MU}^{\otimes s} \otimes R)^{hG}$ holds automatically by [40, Theorem 5.10], and we have the following.

Corollary 5.12 *Let R be an even-periodic, Landweber exact $\mathbb{E}_\infty$ -ring with G -action for a finite group G and R_0 Noetherian. Suppose for every field valued point $x : \mathrm{Spec}(k) \rightarrow \mathrm{Spec}(R_0)$, the stabilizer of x in G acts faithfully on the corresponding formal group over k . Then $E = R^{hG}$ has finite chromatic defect.*

For example, we recover immediately Corollary 2.36. We finish by giving a more general converse to Proposition 5.9, for a class of stacks $\mathcal{M}_E$ that behave in many ways like a quotient by a finite group.

Proposition 5.13 *Let E be a homotopy commutative ring spectrum such that $E \otimes X(n)$ is MU -nilpotent complete for all n sufficiently large. Suppose there is a faithfully flat, finite morphism $q : \mathrm{Spec}(R) \rightarrow \mathcal{M}_E$ for some Noetherian commutative ring R . Then E has finite chromatic defect.*

Proof The argument for Proposition 5.10 needs only slight modification. An automorphism in $\mathcal{M}_E$ is given by a point of the affine scheme

$$\mathrm{Spec}(T) = \mathrm{Spec}(R) \times_{\mathcal{M}_E} \mathrm{Spec}(R),$$

which is finite over $\mathrm{Spec}(R)$, and therefore Noetherian. The morphism $\mathcal{M}_E \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ determines a morphism

$$\mathrm{Spec}(T) \rightarrow \mathrm{Spec}(L[b_1, b_2, \dots]) = \mathrm{Spec}(L) \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathrm{Spec}(L),$$

and we may consider the chain of ideals

$$(b_1) \subset (b_1, b_2) \subset (b_1, b_2, b_3) \subset \dots$$

in T , which must stabilize. As before we conclude that

$$\mathcal{P}(n) := \mathcal{M}_E \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathcal{M}_{\mathrm{FG}}(n)$$

is discrete for some n . Under the conditions above, it follows that $\mathcal{P}(n)$ is an Artin stack, and pulling back the cover q to $\mathcal{P}(n)$, we may apply [55, Lemma 69.17.3] to see that $\mathcal{P}(n)$ is affine. $\square$

Example 5.14 The stacks associated to ko and tmf as in Example 5.4 satisfy the conditions of Proposition 5.13. Indeed, one has the finite flat covers discussed by Hopkins in [20, Chapter 9]. For ko this cover is the morphism $\mathcal{M}_{ku} \rightarrow \mathcal{M}_{\mathrm{ko}}$, and for tmf this is given 2-locally by the morphism $\mathcal{M}_{\mathrm{tmf}_1(3)} \rightarrow \mathcal{M}_{\mathrm{tmf}}$, 3-locally by $\mathcal{M}_{\mathrm{tmf}_1(2)} \rightarrow \mathcal{M}_{\mathrm{tmf}}$, and $\mathcal{M}_{\mathrm{tmf}}$ is itself affine p -locally for $p > 3$. These finite covers predict the Wood equivalences of Example 2.32.

5.3 $K(n)$ -local stacks and the image of J spectrum

Yet another very useful aspect of Hopkins' stack construction of Definition 5.1 is that it allows one to access information about various chromatic localizations of a homotopy commutative ring spectrum E in terms of the stack $\mathcal{M}_E$. We refer the reader to [10, Section 1.6] for an excellent discussion of this aspect. Behrens' framework in [loc. cit.] can be used to compute chromatic defect $K(n)$ - or $E(n)$ -locally, and we use this to verify that $\Phi(j) = \infty$.

Behrens extends Hopkins' stack definition to include more general descent spectral sequences into this framework, such as that of Tmf , which do not coincide with the corresponding ANSS. We will not need this, however, and all our stacks are in the ordinary sense of Definition 5.1. Behrens gives the following pullback formula following Proposition 5.5(2) for $K(n)$ -localizations, where as before we work instead with the $\mathbb{G}_m$ -torsors over his stacks.

Proposition 5.15 *Let $(\mathcal{M}_{\mathrm{FG}})_{(p)}^{[n]}(1)$ be the formal neighborhood of the locus of formal group laws in characteristic p of exact height n in $\mathcal{M}_{\mathrm{FG}}(1)$. Let E be a homotopy commutative ring spectrum such that MU_*E is flat over MU_* . Then there is a pullback of (formal) stacks*

$$\begin{array}{ccc} \mathcal{M}_{L_{K(n)}E} & \longrightarrow & (\mathcal{M}_{\mathrm{FG}})_{(p)}^{[n]}(1) \\ \downarrow & & \downarrow \\ \mathcal{M}_E & \longrightarrow & \mathcal{M}_{\mathrm{FG}}(1) \end{array}$$

In particular, if $E = X(n)$, one may place $\mathcal{M}_{\mathrm{FG}}(n)$ in the bottom left corner, and if $L_{K(n)}E$ is complex-orientable, the pullback will be affine. This yields the following.

Theorem 5.16 *The $K(n)$ -local sphere $L_{K(n)}\mathbb{S}$ has infinite chromatic defect.*

Proof Suppose $L_{K(n)}\mathbb{S}$ has finite chromatic defect. Then there exists a homotopy ring map $\mathrm{MU} \rightarrow (L_{K(n)}\mathbb{S}) \otimes X(m)$ for some $m < \infty$, and by composing, a homotopy ring map

$$\mathrm{MU} \rightarrow (L_{K(n)}\mathbb{S}) \otimes X(m) \rightarrow L_{K(n)}X(m).$$

By Proposition 5.15, it follows that the pullback

$$\mathcal{M}_{\mathrm{FG}}(n) \times_{\mathcal{M}_{\mathrm{FG}}(1)} (\mathcal{M}_{\mathrm{FG}})_{(p)}^{[n]}(1)$$

is affine and therefore discrete. As in the proof of Proposition 5.9, this implies that the automorphism group of the Honda formal group law F_n must be finite, but the Morava stabilizer group $\mathbb{G}_n$ is infinite. $\square$

We finish the section with a connective version of this result at height 1, regarding Mahowald's connective image of J spectrum j . Working p -locally, one has Adams operations ψ^{p+1} acting on ko when $p = 2$, and on the Adams summand ℓ when $p > 2$. When $p = 2$, the map $\psi^3 - 1 : \mathrm{ko} \rightarrow \mathrm{ko}$ factors through the connective cover $\tau_{\geq 4}\mathrm{ko}$, and when $p > 2$, the map $\psi^{p+1} - 1 : \ell \rightarrow \ell$ factors through $\tau_{\geq 2p-2}\ell$; see for example [17].

Definition 5.17 When $p = 2$, we define

$$j = \mathrm{fib}(\psi^3 - 1 : \mathrm{ko} \rightarrow \tau_{\geq 4}\mathrm{ko}),$$

and when $p > 2$,

$$j = \mathrm{fib}(\psi^{p+1} - 1 : \ell \rightarrow \tau_{\geq 2p-2}\ell).$$

The spectrum j is an fp spectrum in the sense of Definition 4.1 and in fact admits an $\mathbb{E}_\infty$ -structure. One of its claims to fame is that it is a connective model of the $K(1)$ -local sphere; that is, $L_{K(1)}j = L_{K(1)}\mathbb{S}$ as $\mathbb{E}_\infty$ -rings at all primes; see for example [29, Lemma 2.2].

Theorem 5.18 *The connective image of J spectrum j does not have finite chromatic defect, at any prime p .*

Proof If $j \otimes X(n)$ were complex-orientable for some n , then so would be $L_{K(1)}j = L_{K(1)}\mathbb{S}$, contradicting Theorem 5.16. $\square$

Immediately from Corollary 2.35, we have the following.

Corollary 5.19 *The connective image of J spectrum j is not Wood-type. That is, there is no finite BP-projective F such that $j \otimes F$ is complex-orientable.*

6 Real-oriented theories and higher real K -theories

We compute the chromatic defect of the Real Johnson–Wilson theories $\mathrm{ER}(n)$ and Goerss–Hopkins–Miller theories $\mathrm{EO}_n(G)$, using Corollary 5.12. As in Proposition 5.11, this comes down to understanding the quotient stacks $\mathrm{Spec}(E(n)_*)/C_2$ and $\mathrm{Spec}((E_n)_*)/G$, and in particular showing that the corresponding morphisms to $\mathcal{M}_{\mathrm{FG}}(1)$ are affine.

6.1 The Real Johnson–Wilson theories

The fixed points $\mathrm{ER}(n) = (E(n))^{hC_2}$ with respect to the complex conjugation action give higher-height versions of periodic real K -theory $\mathrm{KO} = \mathrm{ER}(1)$. At higher heights, however, the C_2 -action on $E(n)$ is not known to give an equivariant ring spectrum, and hence the fixed-point spectrum $\mathrm{ER}(n)$ is not known to be a ring spectrum. Kitchloo, Lorman and Wilson [32] show that $\mathrm{ER}(n)$ has a homotopy commutative ring structure, up to phantom maps, but no more than this is known at the moment. We will show that $\mathrm{ER}(n) \otimes X(2^n)$ has homotopy groups concentrated in even degrees; its complex-orientability follows from Proposition 3.7 if it is a homotopy associative ring spectrum, so we will instead verify a ring structure on $\mathrm{ER}(n) \otimes X(2^n)$.

The lack of ring structure also makes it difficult to apply the methods of Mathew and Meier [40], say, to obtain equivalences such as

$$\mathrm{MU}^{\otimes s} \otimes \mathrm{ER}(n) \simeq (\mathrm{MU}^{\otimes s} \otimes E(n))^{hC_2},$$

as in Proposition 5.11. We can use a trick from genuine equivariant homotopy to get around this.

Lemma 6.1 *Let E be a Borel-complete genuine C_p -spectrum, X a spectrum, and let $i_* : \mathrm{Sp} \rightarrow \mathrm{Sp}^{C_p}$ denote the unique symmetric monoidal colimit preserving functor. Then the natural map*

$$E^{hC_p} \otimes X \rightarrow (E \otimes i_* X)^{hC_p}$$

is an equivalence if and only if $E \otimes i_* X$ is also Borel complete. This is always true if E is a module over a C_p -ring spectrum R such that $\Phi^{C_p}(R) \simeq *$.

Proof This is immediate from the fact that $E^{C_p} \simeq E^{hC_p}$ and $(-)^{C_p}$ commutes with colimits. For the second claim, note that $\Phi^{C_p}(E \otimes i_* X)$ and $(E \otimes i_* X)^{tC_p}$ are both modules over $\Phi^{C_p}(R)$. $\square$

We now state the crucial algebrogeometric input necessary to compute $\Phi(\mathrm{ER}(n))$.

Proposition 6.2 *Suppose that the pullback stack*

$$\mathrm{Spec}(E(n)_*)/C_2 \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathcal{M}_{\mathrm{FG}}(m)$$

is an affine scheme. Then $\mathrm{ER}(n) \otimes X(m)'$ admits a Landweber exact homotopy commutative ring structure for all $m' \geq m$, and in particular, $\Phi(\mathrm{ER}(n)) \leq m$.

Proof Let $\mathrm{Spec}(R_*)$ be the affine scheme given by the above pullback; then the map $\mathrm{Spec}(R_*) \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ is flat. Indeed this follows from the fact that $\mathcal{M}_{\mathrm{FG}}(n) \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ is flat and $\mathrm{Spec}(E(n)_*)/C_2 \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$

is flat. The former is Proposition 5.7 and the latter follows from the facts that $\mathrm{Spec}(E(n)_*) \rightarrow \mathrm{Spec}(E(n)_*)/C_2$ is faithfully flat and $\mathrm{Spec}(E(n)_*) \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ is flat as $E(n)$ is Landweber exact.

There is therefore a Landweber exact homotopy commutative ring spectrum R whose formal group induces the above map $\mathrm{Spec}(R_*) \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$. We can identify the spectra R and $\mathrm{ER}(n) \otimes X(m)$ by producing an isomorphism of homology theories

$$\mathrm{MU}_*(Z) \otimes_{\mathrm{MU}_*} R_* \rightarrow (\mathrm{ER}(n) \otimes X(m))_*(Z).$$

Since

$$\mathrm{Spec}(E(n)_*)/C_2 \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathcal{M}_{\mathrm{FG}}(m)$$

is an affine scheme, the homotopy fixed-point spectral sequence

$$H^*(C_2; (E(n) \otimes X(m))_* Z) \Rightarrow (\mathrm{ER}(n) \otimes X(m))_*(Z)$$

is concentrated on the zero-line, where we have used Lemma 6.1 to identify the target of the spectral sequence. The natural map

$$(\mathrm{ER}(n) \otimes X(m))_*(Z) \rightarrow ((E(n) \otimes X(m))_*(Z))^{C_2}$$

is then an isomorphism. Note now that the map

$$\mathrm{Spec}((E(n) \otimes X(m))_*) \rightarrow \mathrm{Spec}(R_*)$$

is a C_2 -torsor, as it is the base change of the C_2 -torsor $\mathrm{Spec}(E(n)_*) \rightarrow \mathrm{Spec}(E(n)_*)/C_2$ along the map $\mathcal{M}_{\mathrm{FG}}(m) \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$. This gives a natural identification

$$((E(n) \otimes X(m))_*(Z))^{C_2} \cong (\mathrm{MU}_* Z \otimes_{\mathrm{MU}_*} (E(n) \otimes X(m))_*)^{C_2} \cong \mathrm{MU}_* Z \otimes_{\mathrm{MU}_*} R_*,$$

using in the first identification that $E(n) \otimes X(m)$ is Landweber exact.

Since $\mathcal{M}_{\mathrm{FG}}(m') \rightarrow \mathcal{M}_{\mathrm{FG}}(m)$ is an affine morphism for $m' \geq m$, if the pullback in the statement of the proposition is an affine scheme, then so is

$$\mathrm{Spec}(E(n)_*)/C_2 \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathcal{M}_{\mathrm{FG}}(m'),$$

and the same argument now applies to $\mathrm{ER}(n) \otimes X(m)'$. □

Under the assumptions of Proposition 6.2, we can let $m = \infty$, and we deduce that $\mathrm{ER}(n) \otimes \mathrm{MU}$ is a Landweber exact homotopy commutative ring spectrum, and hence we can associate a stack to $\mathrm{ER}(n)$ exactly as in Definition 5.1. We set

$$\mathcal{M}_{\mathrm{ER}(n)} := \mathcal{M}_{(\mathrm{MU}_{2*}\mathrm{ER}(n), \mathrm{MU}_{2*}\mathrm{MU} \otimes_{\mathrm{MU}_{2*}} \mathrm{MU}_{2*}\mathrm{ER}(n))}.$$

All the usual properties for $\mathcal{M}_E$ hold as well for $\mathcal{M}_{\mathrm{ER}(n)}$, such as Proposition 5.5, since these properties only ever make use of the ring structure on $\mathrm{MU}_{2*}E$.

Proposition 6.3 Assuming that, for some m ,

$$\mathrm{Spec}(E(n)_*)/C_2 \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathcal{M}_{\mathrm{FG}}(m)$$

is an affine scheme, there is an equivalence of $\mathbb{G}_m$ -stacks

$$\mathcal{M}_{\mathrm{ER}(n)} \simeq \mathrm{Spec}(E(n)_*)/C_2$$

over $\mathcal{M}_{\mathrm{FG}}(1)$.

Proof Since the morphism $\mathcal{M}_{\mathrm{FG}}(m) \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ is faithfully flat, the assumptions imply that the map $p : \mathrm{Spec}(E(n)_*)/C_2 \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ is affine. This, along with Lemma 6.1, imply that the conditions of Proposition 5.11 are satisfied. $\square$

Remark 6.4 Hopkins gives a description of $\mathcal{M}_{\mathrm{KO}}$ as the moduli stack of nonsingular quadratic equations and strict coordinate transformations, as in Example 5.4. In fact, this description falls immediately out of Proposition 6.3 using that $\mathrm{KO} \simeq \mathrm{ER}(1)$. It is possible to make similar identifications to give a modular description of $\mathcal{M}_{\mathrm{ER}(n)}$ for all heights n , which we give in [13] and intend to return to in future work.

We turn now to proving that the pullback in Proposition 6.2 is indeed affine.

Theorem 6.5 For all $m \geq 2^n$, $\mathrm{ER}(n) \otimes X(m)$ is a Landweber exact homotopy commutative ring spectrum. Moreover, $\Phi(\mathrm{ER}(n)) = 2^n$.

Proof We begin by showing that $\Phi(\mathrm{ER}(n)) \leq 2^n$; we need to show that the pullback

$$\mathcal{P} := \mathrm{Spec}(E(n)_*)/C_2 \times_{\mathcal{M}_{\mathrm{FG}}(1)} \mathcal{M}_{\mathrm{FG}}(2^n)$$

is an affine scheme. As in the proof of Proposition 5.10 following [31, Theorem A7.1.1], since $\mathrm{Spec}(E(n)_*)/C_2$ is a quotient stack, it suffices to show $\mathcal{P}$ is discrete. An object $P \in \mathcal{P}(R)$ has a nontrivial automorphism if and only if the nontrivial element $\gamma \in C_2$ acts on the formal group F_P over R by a power series $f(x)$ with the property that $f(x) \equiv x \pmod{x^{2^n+1}}$.

The element γ , however, acts by formal inversion on F_P . More specifically, the complex-conjugation action map $\gamma : \mathrm{MU} \rightarrow \mathrm{MU}$ induces the map on homotopy groups classifying the conjugate of the universal formal group law F_{univ} by $-[-1]_{F_{\mathrm{univ}}}(x)$; see [25, Example 11.19]. Thus if P has a nontrivial automorphism, then

$$[-1]_{F_P}(x) \equiv x \pmod{x^{2^n+1}}.$$

Now since v_n is a unit in R , and $\gamma^* F_P = F_P$, it follows from the fact that $\gamma(v_n) = -v_n$ that $2v_n = 0 \in R$, so R is an $\mathbb{F}_2$ -algebra, and in particular $-[-1]_{F_P}(x) = [-1]_{F_P}(x)$. This implies that

$$0 = F_P(x, [-1]_{F_P}(x)) \equiv F_P(x, x) \pmod{x^{2^n+1}},$$

but the right-hand side is $[2]_{F_P}(x)$, and F_P has height $\leq n$, a contradiction for a nonzero ring R .

To see that $\Phi(\mathrm{ER}(n)) \geq 2^n$, suppose to the contrary that $\mathrm{ER}(n) \otimes X(2^n - 1)$ is complex orientable. By Proposition 6.3 it follows that there is a pullback square

$$\begin{array}{ccc} \mathcal{M}_{\mathrm{ER}(n) \otimes X(2^n - 1)} & \longrightarrow & \mathcal{M}_{\mathrm{FG}}(2^n - 1) \\ \downarrow & & \downarrow \\ \mathrm{Spec}(E(n)_*)/C_2 & \longrightarrow & \mathcal{M}_{\mathrm{FG}}(1) \end{array}$$

If the pullback $\mathcal{M}_{\mathrm{ER}(n) \otimes X(2^n - 1)}$ were affine, it would be discrete. This is a contradiction because we can consider the point of the pullback at the ring $\mathbb{F}_2[v_n^\pm]$ determined by the ring map $E(n)_* \rightarrow \mathbb{F}_2[v_n^\pm]$ which kills $(2, v_1, \dots, v_{n-1})$. This has a nontrivial automorphism because

$$[-1]_{F_{\mathrm{univ}}}(x) \equiv x \pmod{(2, v_1, \dots, v_{n-1}, v_n x^{2^n});}$$

see [9, Proposition 3.5]. □

6.2 Higher real K -theories

We turn now to the determination of the chromatic defect of Goerss–Hopkins–Miller higher real K -theories. As before we fix a height n formal group Γ over a perfect field k of characteristic p and a finite subgroup $G \subset \mathbb{G}_n$ of the corresponding Morava stabilizer group, and we let $\mathrm{EO}_n(G)$ denote the fixed points of $E(k, \Gamma)$ with respect to G . This follows a similar path to the previous subsection: we identify $\mathcal{M}_{\mathrm{EO}_n(G)}$ with a quotient stack and then use a hands-on stacks argument to determine the minimum m such that the stack becomes affine after pulling back to $\mathcal{M}_{\mathrm{FG}}(m)$. We have already made the first observation in the previous section, which we recall here.

Proposition 6.6 *There is a $\mathbb{G}_m$ -equivariant equivalence of stacks*

$$\mathcal{M}_{\mathrm{EO}_n(G)} \simeq \mathrm{Spec}((E_n)_*)/G$$

over $\mathcal{M}_{\mathrm{FG}}(1)$.

Proof As in [40, Theorem 5.10], the stabilizer of any point $x \in \mathrm{Spec}((E_n)_*)/G$ acts faithfully on the formal group since G is a subgroup of the Morava stabilizer group, thus the morphism $\mathrm{Spec}((E_n)_*)/G \rightarrow \mathcal{M}_{\mathrm{FG}}(1)$ is affine, and $\mathrm{EO}_n(G) \rightarrow E_n$ is a G -Galois extension. The latter implies that $(\mathrm{MU}^{\otimes s} \otimes E_n)^{hG} \simeq \mathrm{MU}^{\otimes s} \otimes \mathrm{EO}_n(G)$, and Proposition 5.11 applies. □

We will use the valuation on $\mathrm{End}(\Gamma)$ to calculate the chromatic defect of $\mathrm{EO}_n(G)$ in the following way. Fix a p -typical universal deformation $\tilde{\Gamma}$ of Γ , and an automorphism $g \in \mathrm{End}(\Gamma)^\times$ of Γ . By universal property, there is an induced isomorphism

$$[g] : \tilde{\Gamma} \rightarrow [g]^* \tilde{\Gamma}$$

of formal group laws over $W(k)[[u_1, \dots, u_{n-1}]]$. Since we have chosen a p -typical coordinate, we may write

$$[g](x) = x + \sum_{i \geq 1} g^* \tilde{\Gamma} t_i(g) x^{p^i},$$

and thus

$$[g-1](x) = \sum_{i \geq 1} g^* \tilde{\Gamma} t_i(g) x^{p^i}.$$

Working mod the maximal ideal $\mathfrak{m}$ in $W(k)[[u_1, \dots, u_{n-1}]]$ we have an expression

$$[g-1](x) = \sum_{i \geq 1} \Gamma t_i(g) x^{p^i},$$

so that $n \cdot v(g-1)$ is the minimum i such that $t_i(g) \not\equiv 0 \pmod{\mathfrak{m}}$. Since $W(k)[[u_1, \dots, u_{n-1}]]$ is a complete local ring, this is equivalent to asking that $t_i(g)$ be a unit. In the statement below, we let $\pi : \mathbb{G}_n \rightarrow \text{Aut}(\Gamma)$ denote the projection map, where we identify the underlying set of the semidirect product as a cartesian product.

Theorem 6.7 *Let $N(G) := n \cdot \max\{v(\pi(g) - 1) : e \neq \pi(g)\}_{g \in G}$, where e is the identity element of G . Then $\Phi(\text{EO}_n(G)) = p^{N(G)}$.*

Proof Our proof closely follows that of Theorem 6.5, namely to see that $\Phi(\text{EO}_n(G)) \leq p^{N(G)}$, we need to show that the stack

$$\mathcal{P} := \text{Spec}((E_n)_*) / G \times_{\mathcal{M}_{\text{FG}}(1)} \mathcal{M}_{\text{FG}}(p^{N(G)})$$

is an affine scheme, for which it suffices to show in this case that it is discrete. An object $P \in \mathcal{P}(R)$ has a nontrivial automorphism if and only if a nontrivial element $g \in G$ acts on the formal group Γ_P on R by a power series $[\pi(g)](x)$ with the property that $[\pi(g)](x) \equiv x \pmod{x^{p^{N(G)}+1}}$. However, by definition of $N(G)$, we have that

$$[\pi(g)](x) = x + g^* \tilde{\Gamma} \sum_{i \geq 1} g^* \tilde{\Gamma} t_i(g) x^{p^i}$$

and $t_i(g) \not\equiv 0 \pmod{\mathfrak{m}}$ for some $i \leq N(G)$, so that $t_i(G)$ is a unit in $(E_n)_*$ and therefore in R , which gives a contradiction for a nonzero ring R .

As before, to see that $\Phi(\text{EO}_n(G)) \geq p^{N(G)}$ it suffices to show that

$$\text{Spec}((E_n)_*) / G \times_{\mathcal{M}_{\text{FG}}(1)} \mathcal{M}_{\text{FG}}(p^{N(G)} - 1)$$

is not discrete. A nontrivial automorphism in this stack is found over $k[u^{\pm}]$ at Γ by taking g so that $N(G) = nv(\pi(g) - 1)$. $\square$

Computing this valuation is a purely number-theoretic problem, and can be done with ease in many cases using general facts about valuations on division algebras, for which we refer the reader to [54]. At a given height n we let F_n be the Honda formal group law over $\overline{\mathbb{F}}_p$. When $n = k(p-1)$, there is a tower of division algebras

$$\mathbb{Q}_p \subset \mathbb{Q}_p(\zeta_p) \subset \text{End}(F_n) \left[\frac{1}{p} \right].$$

The element $\zeta_p \in \mathbb{Q}_p(\zeta_p)^\times$ has order p and has positive valuation, therefore giving a copy of C_p in $\text{Aut}(F_n)$ and a C_p -action on E_n .

Corollary 6.8 $\Phi(\mathrm{EO}_{k(p-1)}(C_p)) = p^k$.

Proof For all $0 < k < p$, $\zeta^k - 1$ is a uniformizer of $\mathbb{C}_{\mathbb{Q}_p}(\zeta_p)$, and since $\mathbb{Q}_p(\zeta_p)/\mathbb{Q}_p$ is totally ramified, we therefore have $v(\zeta^k - 1) = 1/(p-1)$, so that $N(C_p) = k$. $\square$

More generally, $D = \mathrm{End}(F_h)[1/p]$ is the central $\mathbb{Q}_p$ -division algebra with Hasse invariant $1/h$, and Serre [54] showed that a field extension $K/\mathbb{Q}_p$ is contained in D if and only if $[K:\mathbb{Q}_p]$ divides h . Letting $h = p^{n-1}(p-1)m$, we have that

$$[\mathbb{Q}_p(\zeta_{p^n}) : \mathbb{Q}_p] = \phi(p^n) = p^{n-1}(p-1),$$

so a choice of embedding of $\mathbb{Q}_p(\zeta_{p^n})$ gives as before a C_{p^n} -action on E_n .

Corollary 6.9 $\Phi(\mathrm{EO}_{p^{n-1}(p-1)m}(C_{p^n})) = p^{p^{n-1}m}$.

Proof Again $\mathbb{Q}_p(\zeta_{p^n})/\mathbb{Q}_p$ is a totally ramified extension. One thus computes the valuations

$$v(\zeta_{p^n} - 1) = \frac{1}{p^{n-1}(p-1)}, \quad \dots, \quad v(\zeta_{p^n}^{p^{n-1}} - 1) = \frac{1}{p-1},$$

so that $N(C_{p^n}) = p^{n-1}m$. $\square$

Remark 6.10 When $p = 2$, the $\mathrm{ER}(n)$'s $K(n)$ -localize to $\mathrm{EO}_n(C_2)$'s, and thus the agreement of the numbers in Theorem 6.5 and Corollary 6.8 does not come as a surprise. In fact, Beaudry, Hill, Shi and Zeng [9] construct genuine C_{2^n} -spectra known as the $D^{-1}\mathrm{BP}^{(G)}\langle m \rangle$, whose fixed points $K(n)$ -localize to the $\mathrm{EO}_{2^{n-1}m}(C_{2^n})$ of Corollary 6.9. In this way the fixed-point spectra $(D^{-1}\mathrm{BP}^{(G)}\langle m \rangle)^G$ give generalizations of the $\mathrm{ER}(n)$'s that capture larger cyclic 2-groups in the Morava stabilizer group. It can be shown that their chromatic defect coincides with the number given in Corollary 6.8; we would like to return to these and their associated stacks in a future work.

7 $\mathbb{Z}$ -indexed Adams–Novikov spectral sequences

In their study of fp spectra, Mahowald and Rezk observed that the ASS of an fp spectrum E of type n could be extended to a full plane SS converging to the homotopy groups of the chromatic localization $L_n^f E$. The essential idea is that, choosing a type $n+1$ complex F so that $E \otimes F$ is an $\mathbb{F}_p$ -module as in Definition 4.1, one can form an Adams tower for E using F in place of $\mathbb{F}_p$. Since F is finite, the entire tower may be dualized, thus forming a $\mathbb{Z}$ -indexed Adams tower, resulting in a full plane spectral sequence.

Our definition of a Wood-type E gives E precisely the properties necessary to emulate this construction with the Adams–Novikov spectral sequence. Unlike the Adams case, the $\mathbb{Z}$ -indexed ANSS will instead converge to 0. We develop the basics of these $\mathbb{Z}$ -ANSSs in this section, relate their E_2 -pages to various forms of Tate cohomology, and then plug in our motivating example, ko.

There is an obvious generalization of the notion of Wood-types and $\mathbb{Z}$ -indexed Adams spectral sequences to a general ring spectrum E in place of BP or $\mathbb{F}_p$, producing full plane spectral sequences which converge to $L_n^f E$, whenever E is an E -Wood-type via a finite type $n+1$ complex F , but we do not pursue this here.

7.1 The ANSS for a Wood-type

We recall a general framework for Adams towers and their associated spectral sequences developed by Haynes Miller; see [44] as well as [28, Section 4]. For an $\mathbb{E}_0$ -algebra R , Miller defines an R -injective exactly as we define a weak R -module in Definition 2.5. Using this, he defines R -injective resolutions and relates them to R -based ASSs, from which the following may be deduced.

Proposition 7.1 *Let R be an $\mathbb{E}_0$ -algebra and E a spectrum, and suppose we are given a tower*

$$\begin{array}{ccccc} \cdots & \longrightarrow & E_2 & \longrightarrow & E_1 & \longrightarrow & E \\ & & \downarrow i_2 & & \downarrow i_1 & & \downarrow i_0 \\ & & C_2 & & C_1 & & C_0 \end{array}$$

where $E_j = \text{fib}(i_{j-1})$, $R \otimes i_j$ is a split monomorphism, and C_j is a weak R -module for all j . Then the spectral sequence obtained by applying $\pi_*(-)$ to this tower coincides with the R -based ASS of E from the E_2 -page on.

Suppose then that F is an $\mathbb{E}_0$ -algebra, and E a spectrum, and form the F -based Adams tower of E :

$$(7-1) \quad \begin{array}{ccccc} \cdots & \longrightarrow & E \otimes \bar{F} \otimes \bar{F} & \longrightarrow & E \otimes \bar{F} & \longrightarrow & E \\ & & \downarrow i_2 & & \downarrow i_1 & & \downarrow i_0 \\ & & E \otimes \bar{F} \otimes \bar{F} \otimes F & & E \otimes \bar{F} \otimes F & & E \otimes F \end{array}$$

Suppose now that R is a homotopy associative ring spectrum. If $E \otimes F$ is a weak R -module, and F is an R -projective, it follows from the proposition that the associated spectral sequence recovers the R -based ASS of E from the E_2 -page on. Using that $F = T(n)$ is a $R = \text{BP}$ -free, we have the following.

Corollary 7.2 *Suppose $\Phi_p(E) \leq n$, so that $E \otimes T(n)$ is complex-orientable. Then the $T(n)$ -based ASS of E coincides with the ANSS of E from the E_2 -page on.*

We have the following further simplification when E is Wood-type, setting now F to be a finite BP-projective and $R = \text{BP}$. Here we shift F so that its bottom cells are in degree zero and fix a cell $\mathbb{S} \rightarrow F$, thereby equipping F with an $\mathbb{E}_0$ -algebra structure.

Corollary 7.3 *Suppose that E is Wood-type, and choose a finite BP-projective F so that $E \otimes F$ is complex-orientable. Then the F -based ASS of E coincides with the ANSS of E from the E_2 -page on.*

Example 7.4 Let $E = \text{ko}$ and $F = C(\eta)$, i.e., the classical Wood equivalence. Corollary 7.3 gives an identification of ANSS(ko) with the η -Bockstein spectral sequence for ko.

Similar statements apply to the cases $E = \text{tmf}$ with $F = D\mathcal{A}(1)$ at the prime 2 and $F = X_2$ at the prime 3 from Example 2.32. The identification of the E_2 -page of the ANSS this provides corresponds to the stack-theoretic identifications of the descent SS for tmf as in [6], using the 2- and 3-local covers of $\mathcal{M}_{ell}$ by $\mathcal{M}_1(3)$ and $\mathcal{M}_1(2)$, respectively.

Remark 7.5 It is often the case, as in the example above, that for a Wood-type $\mathbb{E}_\infty$ -ring spectrum E , the corresponding complex-oriented spectrum $T := E \otimes F$ has the structure of an E -algebra, as in with $ko \rightarrow ku \simeq ko \otimes C(\eta)$. In this case, $\text{ANSS}(E)$ may be identified with relative Adams spectral sequence for the ring map $E \rightarrow T$ of [5]. This has the advantage of endowing the spectral sequence of (7-1) with the structure of a spectral sequence of algebras from the E_1 -page on.

7.2 $\mathbb{Z}$ -indexed ANSS and Tate cohomology

As above we fix a bottom cell $S \rightarrow F$, and we let $\bar{F} := \text{fib}(S \rightarrow F)$. The tower of (7-1) extends to the right whenever F is dualizable:

$$(7-2) \quad \begin{array}{ccccccc} \cdots & \longrightarrow & E \otimes \bar{F} & \longrightarrow & E & \longrightarrow & E \otimes \mathbb{D}\bar{F} \longrightarrow \cdots \\ & & \downarrow & & \downarrow & & \downarrow \\ & & E \otimes \bar{F} \otimes F & & E \otimes F & & E \otimes \Sigma \mathbb{D}F \end{array}$$

Definition 7.6 If E is Wood-type, fix a finite BP-projective F with $E \otimes F$ complex-orientable. The $\mathbb{Z}$ -indexed Adams–Novikov spectral sequence of E ($\mathbb{Z}$ -ANSS(E)) is the spectral sequence associated to the tower of (7-2).

We will show now that the $\mathbb{Z}$ -ANSS of a Wood-type E almost completely determines the ANSS of E . For this, we will need some terminology introduced by Meier, Shi and Zeng [43, Definition 2.5].

Definition 7.7 Let $f : \mathcal{E} \rightarrow \mathcal{E}'$ be a morphism of spectral sequences. We say f is an *isomorphism of spectral sequences in nonnegative filtrations* if the following statements hold:

- (1) On the E_2 -page, f induces an isomorphism in positive filtrations and an epimorphism in filtration zero.
- (2) For all $r \geq 2$, every nonzero d_r -differential in $\mathcal{E}$ whose source is in nonnegative filtration is mapped via f to a nonzero d_r -differential in $\mathcal{E}'$.
- (3) For all $r \geq 2$, every nonzero d_r -differential in $\mathcal{E}'$ whose source is in nonnegative filtration is the image along f of a nonzero d_r -differential in $\mathcal{E}$.

In fact, conditions (2) and (3) of Definition 7.7 are redundant; they follow from condition (1). Indeed, this may be shown using an induction argument given in [43, Proof of Theorem 2.6], which we now briefly recall; see also [21, Proof of Theorem 3.3].

Proposition 7.8 Let $f : \mathcal{E} \rightarrow \mathcal{E}'$ be a morphism of spectral sequences such that, on the E_2 -page, f induces an isomorphism in positive filtrations and an epimorphism in filtration zero. Then f is an **isomorphism of spectral sequences in nonnegative filtrations**.

Proof Conditions (2) and (3) of Definition 7.7 follow for $r = 2$ immediately from condition (1) of Definition 7.7, i.e., from the assumption of the proposition. We then let $r > 2$ and suppose that conditions (2) and (3) of Definition 7.7 hold for all $r' < r$.

Suppose $d_r(x) = y$ is a nonzero differential in $\mathcal{E}$. We claim that $d_r(f(x)) = f(y)$ is a nonzero differential in $\mathcal{E}'$. If not then $f(y) = 0$ on the E_r -page. However, since $f(y) \neq 0$ on the E_2 -page, there must be a differential $d_{r'}(z) = f(y)$ in $\mathcal{E}'$ for $r' < r$. By induction using condition (3), this differential must lift to $\mathcal{E}$, contradicting that y is nonzero on E_r . This establishes condition (2).

Similarly, suppose $d_r(x) = y$ is a nonzero differential in $\mathcal{E}'$. Since x admits a lift $\tilde{x}$ to $\mathcal{E}$ on the E_2 -page, if the differential does not lift on the E_r -page, then there must be a nonzero differential $d_{r'}(\tilde{x}) = z$ for $r' < r$. By induction using condition (2), it follows that x supports a nonzero $d_{r'}$ -differential in $\mathcal{E}'$, contradicting that x survives to the E_r -page. $\square$

Theorem 7.9 *The $\mathbb{Z}$ -ANSS of a Wood-type E has the following properties:*

- (1) *The $\mathbb{Z}$ -ANSS is independent of the choice of finite BP-projective F from the E_2 -page on.*
- (2) *The natural map*

$$\text{ANSS}(E) \rightarrow \mathbb{Z}\text{-ANSS}(E)$$

is an isomorphism of spectral sequences in nonnegative filtrations in the sense of Definition 7.7.

- (3) *The $\mathbb{Z}$ -ANSS converges to zero.*

Proof For (1), as in [28, Lemma 4.8], the BP-injective resolutions

$$E \rightarrow E \otimes F_i \rightarrow \Sigma E \otimes \bar{F}_i \otimes F_i \rightarrow \dots$$

corresponding to the towers (7-2) for any pair of finite BP-projectives F_1, F_2 with the property that $E \otimes F_i$ is BP-injective are chain homotopy equivalent, since they are Adams resolutions. One may dualize the chain maps and chain homotopies to see that the resolutions of E

$$\dots \rightarrow \Sigma^{-1} E \otimes \mathbb{D} \bar{F}_i \otimes \mathbb{D} F_i \rightarrow E \otimes \mathbb{D} F_i \rightarrow E$$

are chain homotopy equivalent. The two spectral sequences arise by splicing these resolutions together and thus must be isomorphic from the E_2 -page on.

For (2), note that condition (1) of Definition 7.7 follows by construction of the tower. The claim then follows immediately from Proposition 7.8.

For (3), we first show that the colimit of the tower of (7-2) becomes zero after tensoring with F . Each of the downward maps in (7-2) admits a section after tensoring with F . Indeed, since $E \otimes F$ is complex-orientable, the map $E \otimes F \rightarrow E \otimes F \otimes F$ is a retract of the map $E \otimes F \otimes \text{BP} \rightarrow E \otimes F \otimes F \otimes \text{BP}$. Since F is a BP-projective, the map $F \otimes \text{BP} \rightarrow F \otimes F \otimes \text{BP}$ admits a section. It follows that all of the horizontal maps in (7-2) become zero after tensoring with F , as desired.

Since F is a type-zero complex, a spectrum X is zero if and only if $X \otimes F = 0$, as follows from the thick subcategory theorem. Indeed, the set of finite spectra K such that $X \otimes K = 0$ is a thick subcategory of finite spectra, and if it contains a type zero complex, it contains $\mathbb{S}$. It follows that the colimit of the tower of (7-2) is zero, so that the $\mathbb{Z}$ -ANSS converges to zero. $\square$

Remark 7.10 Readers familiar with the homotopy fixed-point spectral sequence (HFPSS) and the Tate spectral sequence (TateSS) will notice similarities between these and the ANSS and $\mathbb{Z}$ -ANSS of a Wood-type. Indeed, for a G -spectrum E , the canonical map $\text{HFPSS}(E) \rightarrow \text{TateSS}(E)$ is given on the E_2 -page by the map

$$H^*(G; \pi_* E) \rightarrow \widehat{H}^*(G; \pi_* E)$$

from group cohomology to Tate cohomology. This map is an isomorphism in positive filtrations and a surjection in filtration zero, and hence the map $\text{HFPSS}(E) \rightarrow \text{TateSS}(E)$ is an isomorphism of spectral sequences in nonnegative filtrations in the sense of Definition 7.7.

When $E^{tG} = 0$, the Tate SS of E converges to zero, as in condition (3) of Theorem 7.9. This vanishing is often used to deduce vanishing lines and differentials in the HFPSS; see [21] for example.

We turn now to identifying the E_2 page of $\mathbb{Z}$ -ANSS(E). In the following statements we will assume that E is a Wood-type and choose a corresponding complex F equipped with an $\mathbb{E}_1$ -multiplication. Since the $\mathbb{Z}$ -ANSS(E) does not depend on the choice of F , this results in no loss of generality since we may replace F with its endomorphism ring $F \otimes \mathbb{D}F$. The E_2 -page is computed via a certain localization in Hovey's category $\text{Stable}(\text{BP}_*\text{BP})$ -comodules, equivalently mod τ in Pstragowski's category Syn_{BP} [48]. We will need to recall some terminology from [41, Section 3.1].

Definition 7.11 Let $\mathcal{C}$ be a presentably symmetric-monoidal stable ∞ -category, and $A \in \text{Alg}(\mathcal{C})$ a dualizable algebra.

- Let $\mathcal{C}_{A\text{-tors}}$ — the category of A -torsion objects — be the localizing subcategory of $\mathcal{C}$ generated by objects of the form $A \otimes X$, where X is dualizable.
- Let $\mathcal{C}_{A^{-1}\text{-local}}$ — the category of A^{-1} -local objects — be the subcategory of $\mathcal{C}$ consisting of objects Y with the property that $\text{Map}_{\mathcal{C}}(Z, Y) \simeq *$ for all $Z \in \mathcal{C}_{A\text{-tors}}$.

Any object $E \in \mathcal{C}$ fits into a natural cofiber sequence

$$Z_{A^{-1}} E \rightarrow E \rightarrow L_{A^{-1}} E,$$

where $Z_{A^{-1}} E \in \mathcal{C}_{A\text{-tors}}$ and $L_{A^{-1}} E \in \mathcal{C}_{A^{-1}\text{-local}}$. Mathew, Naumann and Noel [41, Proposition 3.5] give an explicit description of the localization functor $L_{A^{-1}}(-)$ as follows.

Proposition 7.12 Let $\text{CB}(A)$ denote the cosimplicial object

$$A \rightrightarrows A \otimes A \rightrightarrows A \otimes A \otimes A \cdots,$$

i.e., the cobar complex of A . Then for any $X \in \mathcal{C}$, the map $Z_{A^{-1}} E \rightarrow E$ is equivalent to the map

$$|\mathbb{D}(\text{CB}(A)) \otimes E| \rightarrow E$$

given by dualizing the cobar complex.

As above we fix an $\mathbb{E}_1$ -ring F such that $E \otimes F$ is complex-orientable, and we let

$$\nu(-)/\tau : \mathrm{Sp} \rightarrow \mathrm{Stable}(\mathrm{BP}_*\mathrm{BP})$$

denote the functor sending a spectrum E to the $\mathrm{BP}_*\mathrm{BP}$ -comodule BP_*E regarded as a discrete object in $\mathrm{Stable}(\mathrm{BP}_*\mathrm{BP})$.

Proposition 7.13 *The E_2 -page of the $\mathbb{Z}$ -ANSS of a Wood-type E is given by*

$$E_2^{*,*} = \pi_{*,*} L_{(\nu F/\tau)^{-1}}(\nu E/\tau)$$

in $\mathrm{Stable}(\mathrm{BP}_*\mathrm{BP})$.

Proof Mathew [38, Theorem 2.4] gives a similar categorical description of Tate cohomology, and the proof carries over to our setting without much change. The key point is to show that $\pi_{*,0} L_{(\nu F/\tau)^{-1}}(\nu E/\tau)$ coincides with the zero line on our E_2 -page, which is the cohomology of the complex

$$\pi_*(E \otimes \mathbb{D}F) \rightarrow \pi_*(E \otimes F) \rightarrow \pi_*(\Sigma E \otimes \bar{F} \otimes F).$$

Since $E \otimes F$ and therefore also $E \otimes \mathbb{D}F$ are weak BP-modules, this is isomorphic to

$$\mathrm{Ext}_{\mathrm{BP}_*\mathrm{BP}}^{0,*}(\mathrm{BP}_*(E \otimes \mathbb{D}F)) \rightarrow \mathrm{Ext}_{\mathrm{BP}_*\mathrm{BP}}^{0,*}(\mathrm{BP}_*(E \otimes F)) \rightarrow \mathrm{Ext}_{\mathrm{BP}_*\mathrm{BP}}^{0,*}(\mathrm{BP}_*(\Sigma E \otimes \bar{F} \otimes F)),$$

which is isomorphic to

$$\pi_{*,0}(\nu E/\tau \otimes \mathbb{D}\nu F/\tau) \rightarrow \pi_{*,0}(\nu E/\tau \otimes \nu F/\tau) \rightarrow \pi_{*,0}(\Sigma \nu E/\tau \otimes \overline{\nu F/\tau} \otimes \nu F/\tau).$$

The kernel of the second map, by Corollary 7.3, is isomorphic to $\mathrm{Ext}_{\mathrm{BP}_*\mathrm{BP}}^{0,*}(\mathrm{BP}_*(E)) \cong \pi_{*,0}(\nu E/\tau)$. Since $\nu E/\tau$ and $\nu F/\tau$ are connective objects in $\mathrm{Stable}(\mathrm{BP}_*\mathrm{BP})$, it follows that

$$\pi_{*,0}(|\mathbb{D}(\mathrm{CB}(\nu F/\tau)) \otimes \nu E/\tau|) \cong \pi_{*,0}(\nu E/\tau \otimes \mathbb{D}\nu F/\tau),$$

and now the result follows from Proposition 7.12. $\square$

The description of the E_2 -page resembles Tate cohomology, as the latter is given by $\pi_* L_{k[G]^{-1}}(-)$ in the stable module category associated to a finite group G . However, our description rings a bit hollow at this level of generality, as one would hope for a description in terms of $L_{\mathrm{BP}_*\mathrm{BP}^{-1}}(-)$. The analogous description does hold in the Mahowald–Rezk context replacing BP with $\mathbb{F}_p$ (see [37, Section 2]), where one may use the fact that $\mathcal{A}_*$ is projective in the category of finitely presented $\mathcal{A}_*$ -comodules, and we see no direct analog of this in our setting.

In practice, however, we often find ourselves in the situation of Remark 7.5, namely when E is an $\mathbb{E}_\infty$ ring spectrum and $E \otimes F$ is a finite Adams-flat E -algebra. In this setting, the E_2 -page of the $\mathbb{Z}$ -ANSS(E) can be viewed more aptly as an instance of Tate cohomology. The proof of the following is exactly as in the previous proposition.

Proposition 7.14 *Let E be an E_∞ -ring and $T \in \text{Alg}(E)$ with the property that $\Gamma := \pi_*(T \otimes_E T)$ is flat over $A = \pi_*T$. For any $X \in \text{Mod}(E)$, the spectral sequence associated to the $\mathbb{Z}$ -indexed tower*

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X \otimes_E \bar{T} & \longrightarrow & X & \longrightarrow & X \otimes_E D\bar{T} \longrightarrow \cdots \\ & & \downarrow & & \downarrow & & \downarrow \\ & & X \otimes_E \bar{T} \otimes_E T & & X \otimes_E T & & X \otimes_E \Sigma DT \end{array}$$

has E_2 -page given by $\pi_{*,*}L_{\Gamma-1}(\nu X/\tau)$ in $\text{Stable}(A, \Gamma)$.

Setting $T = E \otimes F$ and $X = E$ as in Remark 7.5, one recovers a description of the E_2 -page of $\mathbb{Z}$ -ANSS(E) as $\pi_{*,*}L_{\Gamma-1}\mathbb{1}$.

7.3 The $\mathbb{Z}$ -ANSS for ko

We begin by calculating ANSS(ko), which by Corollary 7.3 can be identified with the η -Bockstein SS from the E_2 -page on, equivalently, with the relative ASS for the map of commutative ring spectra $\text{ko} \rightarrow ku$, in the sense of [5]. From this latter perspective, one sees that the spectral sequence has the structure of a spectral sequence of algebras, with associative E_1 -page and commutative E_r -page for $r > 1$. Since $ku_* = \mathbb{Z}[u]$ with $|u| = 2$, one has an additive isomorphism

$$E_1 = \pi_*(\text{ko}/\eta)[\eta] = \mathbb{Z}[u, \eta].$$

The class η is a permanent cycle, and the differential d_1 is thus determined by the following.

Proposition 7.15 *In ANSS(ko), the d_1 -differential satisfies*

$$d_1(u^n) = \begin{cases} 2\eta u^{2k} & \text{when } n = 2k + 1, \\ 0 & \text{when } n = 2k. \end{cases}$$

Proof By the Leibniz rule, it suffices to establish $d_1(u) = 2\eta$ and $d_1(u^2) = 0$. These both follow from the Wood sequence: we have an exact sequence

$$\pi_2 ku \xrightarrow{\partial} \pi_0 ko \xrightarrow{\eta} \pi_1 ko \rightarrow 0.$$

Since $\pi_0 ko = \mathbb{Z}$ is generated by trivial bundles and $\pi_1 ko = \mathbb{Z}/2$ via the Mobius bundle, we see that $\partial(u) = 2$, which implies that $d_1(u) = 2\eta$. Since ∂ is thus a monomorphism on $\pi_2 ku$, the exact sequence

$$0 \rightarrow \pi_1 ko \xrightarrow{\eta} \pi_2 ko \rightarrow \pi_2 ku \xrightarrow{\partial} \pi_0 ko$$

implies $\pi_2 ko$ is generated by classes divisible by η , which implies that $d_1(u^2) = 0$, since d_1 is given by the composite

$$\pi_4 ku \xrightarrow{\partial} \pi_2 ko \rightarrow \pi_2 ku,$$

and the latter map has kernel those elements divisible by η . □




Figure 3: The ANSS for ko.

As a consequence of the Leibniz rule and the fact that $d_1(u^2) = 0$, we see that in fact $u\eta = -\eta u$ on E_1 . In any case it follows that $E_2 = \mathbb{Z}[u^2, \eta]/(2\eta)$, as a ring. It follows from the Leibniz rule that either this spectral sequence collapses on E_2 or there is a differential $d_3(u^2) = \eta^3$. We will use the $\mathbb{Z}$ -ANSS to show that this differential must happen. By construction $\mathbb{Z}$ -ANSS(ko) satisfies

$$E_1 = \mathbb{Z}[u, \eta^\pm]$$

as a module over ANSS(ko). It follows that d_1 is determined by the same formulae as in Proposition 7.15, and we have the following.

Proposition 7.16 *The E_2 -page of the $\mathbb{Z}$ -ANSS(ko) is given as a module over that of ANSS(ko) by*

$$E_2 = \mathbb{F}_2[u^2, \eta^\pm].$$


Corollary 7.17 *There is a differential*

$$d_3(u^2) = \eta^3$$

in both ANSS(ko) and $\mathbb{Z}$ -ANSS(ko).

Proof If the claimed differential does not happen in ANSS(ko), then the spectral sequence collapses on E_2 . Indeed if $d_3(u^2) = 0$, then u^2 is a permanent cycle for degree reasons, as is clear from Figure 3. Since ANSS(ko) is a multiplicative spectral sequence with E_2 generated as an algebra by u^2 and η , if u^2 is a permanent cycle, then there are no nonzero differentials in the spectral sequence, since η is also a permanent cycle.

By η -linearity it follows that $\mathbb{Z}$ -ANSS(ko) collapses on E_2 . This contradicts Theorem 7.9, namely that the latter spectral sequence must converge to zero. $\square$

Figure 4: The $\mathbb{Z}$ -ANSS for ko .

Acknowledgements

We would like to thank the referee for their thorough reading and many comments and suggestions, which substantially improved this paper. We would also like to acknowledge helpful conversations with, and comments from, Mike Hill, Lennart Meier, Jack Davies, Shaul Barkan, Ishan Levy and Robert Burklund. We thank Guy Boyde for suggesting the terminology “chromatic defect”. This material is based upon work supported by the National Science Foundation under grant DMS-1811189.

References

- [1] J F Adams, H R Margolis, *Modules over the Steenrod algebra*, Topology 10 (1971) 271–282 MR
- [2] G Angelini-Knoll, C Ausoni, J Rognes, *Algebraic K -theory of real topological K -theory*, preprint (2023) arXiv 2309.11463
- [3] G Angelini-Knoll, J D Quigley, *The Segal conjecture for topological Hochschild homology of Ravenel spectra*, J. Homotopy Relat. Struct. 16:1 (2021) 41–60 MR

- [4] **M F Atiyah**, *K-theory and reality*, Quart. J. Math. Oxford Ser. (2) 17 (1966) 367–386 MR
- [5] **A Baker, A Lazarev**, *On the Adams spectral sequence for R -modules*, Algebr. Geom. Topol. 1 (2001) 173–199 MR
- [6] **T Bauer**, *Computation of the homotopy of the spectrum tmf* , from “Groups, homotopy and configuration spaces” (N Iwase, T Kohno, R Levi, D Tamaki, J Wu, editors), Geom. Topol. Monogr. 13, Geom. Topol. Publ., Coventry (2008) 11–40 MR
- [7] **J Beardsley**, *A theorem on multiplicative cell attachments with an application to Ravenel’s $X(n)$ spectra*, J. Homotopy Relat. Struct. 14:3 (2019) 611–624 MR
- [8] **J Beardsley, T Lawson**, *Skeleta and categories of algebras*, Adv. Math. 457 (2024) art. id. 109944 MR
- [9] **A Beaudry, M A Hill, X D Shi, M Zeng**, *Models of Lubin–Tate spectra via real bordism theory*, Adv. Math. 392 (2021) art. id. 108020 MR
- [10] **M Behrens**, *Topological modular and automorphic forms*, from “Handbook of homotopy theory” (H Miller, editor), CRC, Boca Raton, FL (2020) 221–261 MR
- [11] **P Bhattacharya, H Chatham**, *On the EO-orientability of vector bundles*, J. Topol. 15:4 (2022) 2017–2044 MR
- [12] **R R Bruner, J Rognes**, *The Adams spectral sequence for topological modular forms*, Mathematical Surveys and Monographs 253, Amer. Math. Soc., Providence, RI (2021) MR
- [13] **C D Carrick**, *Stacks and real-oriented homotopy theory*, PhD thesis, University of California, Los Angeles (2022) Available at <https://www.proquest.com/docview/2679729095>
- [14] **C Carrick, M A Hill**, *On higher real K -theories and finite spectra*, preprint (2025) arXiv 2507.07051
- [15] **C Carrick, M A Hill, D C Ravenel**, *The homological slice spectral sequence in motivic and real bordism*, Adv. Math. 458 (2024) art. id. 109955 MR
- [16] **D Clausen, A Mathew, N Naumann, J Noel**, *Descent in algebraic K -theory and a conjecture of Ausoni–Rognes*, J. Eur. Math. Soc. 22:4 (2020) 1149–1200 MR
- [17] **D M Davis**, *The cohomology of the spectrum bJ* , Bol. Soc. Mat. Mexicana (2) 20:1 (1975) 6–11 MR
- [18] **E S Devinatz, M J Hopkins**, *The action of the Morava stabilizer group on the Lubin–Tate moduli space of lifts*, Amer. J. Math. 117:3 (1995) 669–710 MR
- [19] **E S Devinatz, M J Hopkins, J H Smith**, *Nilpotence and stable homotopy theory, I*, Ann. of Math. (2) 128:2 (1988) 207–241 MR
- [20] **C L Douglas, J Francis, A G Henriques, M A Hill** (editors), *Topological modular forms*, Mathematical Surveys and Monographs 201, Amer. Math. Soc., Providence, RI (2014) MR
- [21] **Z Duan, G Li, X D Shi**, *Vanishing lines in chromatic homotopy theory*, Geom. Topol. 29:2 (2025) 903–930 MR
- [22] **P G Goerss, M J Hopkins**, *Moduli spaces of commutative ring spectra*, from “Structured ring spectra” (A Baker, B Richter, editors), London Math. Soc. Lecture Note Ser. 315, Cambridge Univ. Press (2004) 151–200 MR
- [23] **J Hahn, A Raksit, D Wilson**, *A motivic filtration on the topological cyclic homology of commutative ring spectra*, preprint (2022) arXiv 2206.11208 To appear in Ann. Math.
- [24] **M A Hill, M J Hopkins, D C Ravenel**, *The homotopy of $EO_{2(p-1)}$* , unpublished note (2008) Available at <https://www.math.ucla.edu/~mikehill/Research/2pminus1.pdf>
- [25] **M A Hill, M J Hopkins, D C Ravenel**, *On the nonexistence of elements of Kervaire invariant one*, Ann. of Math. (2) 184:1 (2016) 1–262 MR
- [26] **M A Hill, L Meier**, *The C_2 -spectrum $\mathrm{Tmf}_1(3)$ and its invertible modules*, Algebr. Geom. Topol. 17:4 (2017) 1953–2011 MR
- [27] **M J Hopkins**, *Stable decompositions of certain loop spaces*, PhD thesis, Northwestern University (1984) Available at <https://www.proquest.com/docview/303306354>
- [28] **M J Hopkins**, *Complex-oriented cohomology theories and the language of stacks*, course notes (1999) Available at <https://www.sas.rochester.edu/mth/sites/doug-ravenel/otherpapers/coctalos.pdf>
- [29] **M J Hopkins, M Mahowald, H Sadofsky**, *Constructions of elements in Picard groups*, from “Topology and representation theory” (Evanston, IL, 1992) (E M Friedlander, M E Mahowald, editors), Contemp. Math. 158, Amer. Math. Soc., Providence, RI (1994) 89–126 MR

- [30] **P Hu, I Kriz**, *Real-oriented homotopy theory and an analogue of the Adams–Novikov spectral sequence*, *Topology* 40:2 (2001) 317–399 MR
- [31] **N M Katz, B Mazur**, *Arithmetic moduli of elliptic curves*, *Annals of Mathematics Studies* 108, Princeton Univ. Press (1985) MR
- [32] **N Kitchloo, V Lorman, W S Wilson**, *Multiplicative structure on real Johnson–Wilson theory*, from “New directions in homotopy theory”, *Contemp. Math.* 707, Amer. Math. Soc., Providence, RI (2018) 31–44 MR
- [33] **N Kitchloo, W S Wilson**, *The second real Johnson–Wilson theory and nonimmersions of RP^n* , *Homology Homotopy Appl.* 10:3 (2008) 223–268 MR
- [34] **N Kitchloo, W S Wilson**, *The second real Johnson–Wilson theory and nonimmersions of RP^n , II*, *Homology Homotopy Appl.* 10:3 (2008) 269–290 MR
- [35] **T Lawson**, *E_n -spectra and Dyer–Lashof operations*, from “Handbook of homotopy theory” (H Miller, editor), CRC, Boca Raton, FL (2020) 793–849 MR
- [36] **J Lurie**, *Chromatic homotopy theory*, lecture notes (2010) Available at <https://people.math.harvard.edu/~lurie/252x.html>
- [37] **M Mahowald, C Rezk**, *Brown–Comenetz duality and the Adams spectral sequence*, *Amer. J. Math.* 121:6 (1999) 1153–1177 MR
- [38] **A Mathew**, *Torus actions on stable module categories, Picard groups, and localizing subcategories*, preprint (2015) arXiv 1512.01716
- [39] **A Mathew**, *The homology of tmf* , *Homology Homotopy Appl.* 18:2 (2016) 1–29 MR
- [40] **A Mathew, L Meier**, *Affineness and chromatic homotopy theory*, *J. Topol.* 8:2 (2015) 476–528 MR
- [41] **A Mathew, N Naumann, J Noel**, *Nilpotence and descent in equivariant stable homotopy theory*, *Adv. Math.* 305 (2017) 994–1084 MR
- [42] **L Meier**, *United elliptic homology*, PhD thesis, Rheinische Friedrich-Wilhelms-Universität Bonn (2012) Available at <https://hdl.handle.net/20.500.11811/5378>
- [43] **L Meier, XD Shi, M Zeng**, *A stratification of the equivariant slice filtration*, preprint (2023) arXiv 2310.12105
- [44] **HR Miller**, *On relations between Adams spectral sequences, with an application to the stable homotopy of a Moore space*, *J. Pure Appl. Algebra* 20:3 (1981) 287–312 MR
- [45] **H Miller, C Wilkerson**, *Vanishing lines for modules over the Steenrod algebra*, *J. Pure Appl. Algebra* 22:3 (1981) 293–307 MR
- [46] **SA Mitchell**, *Finite complexes with $A(n)$ -free cohomology*, *Topology* 24:2 (1985) 227–246 MR
- [47] **O Nakamura**, *On the squaring operations in the May spectral sequence*, *Mem. Fac. Sci. Kyushu Univ. Ser. A* 26:2 (1972) 293–308 MR
- [48] **P Pstragowski**, *Synthetic spectra and the cellular motivic category*, *Invent. Math.* 232:2 (2023) 553–681 MR
- [49] **D Quillen**, *On the formal group laws of unoriented and complex cobordism theory*, *Bull. Amer. Math. Soc.* 75 (1969) 1293–1298 MR
- [50] **DC Ravenel**, *The non-existence of odd primary Arf invariant elements in stable homotopy*, *Math. Proc. Cambridge Philos. Soc.* 83:3 (1978) 429–443 MR
- [51] **DC Ravenel**, *Complex cobordism and stable homotopy groups of spheres*, *Pure and Applied Mathematics* 121, Academic, Orlando, FL (1986) MR
- [52] **DC Ravenel**, *Nilpotence and periodicity in stable homotopy theory*, *Annals of Mathematics Studies* 128, Princeton Univ. Press (1992) MR
- [53] **C Rezk**, *Notes on the Hopkins–Miller theorem*, from “Homotopy theory via algebraic geometry and group representations” (Evanston, IL, 1997) (M Mahowald, S Priddy, editors), *Contemp. Math.* 220, Amer. Math. Soc., Providence, RI (1998) 313–366 MR
- [54] **J-P Serre**, *Local fields*, *Graduate Texts in Mathematics* 67, Springer (1979) MR
- [55] *The Stacks project*, electronic reference (2005–) Available at <http://stacks.math.columbia.edu>

[56] **R Wood**, *Banach algebras and Bott periodicity*, Topology 4 (1965/66) 371–389 MR

CHRISTIAN CARRICK carrick@math.uni-bonn.de

Mathematisches Institut, Universität Bonn, Bonn, Germany

Proposed: Haynes R Miller

Received: September 17, 2024

Seconded: Stefan Schwede, Mark Behrens

Revised: June 10, 2025

Compact holonomy G_2 manifolds need not be formal

LUCÍA MARTÍN-MERCHÁN

We construct a compact simply connected manifold with holonomy G_2 that is nonformal. This manifold is the resolution of a flat G_2 orbifold with $\mathbb{Z}_2$ isotropy, which can be resolved using both the generalized Kummer construction by D. D. Joyce and the method developed by Joyce and S. Karigiannis. A non-vanishing triple Massey product is obtained by arranging the singular locus in a particular configuration.

1 Introduction

The link between special holonomy and formality was discovered by P. Deligne, P. Griffiths, J. Morgan, and D. Sullivan in [10]. They proved that compact complex manifolds satisfying the dd^c -lemma are formal. In particular, Kähler, Calabi–Yau, and hyper-Kähler compact manifolds are formal. Later, M. Amann and V. Kapovitch showed in [1] that compact quaternionic-Kähler manifolds with positive scalar curvature are also formal. For that, they used that their twistor space admits a Kähler metric, along with results from rational homotopy.

The question of formality for compact manifolds with exceptional holonomy has remained open since they were first constructed by D. D. Joyce in [13; 14; 15; 16]. In fact, these were conjectured to be formal in [1, Conjecture 2]. Various approaches have been explored to better understand the G_2 case. Analytical methods were employed by K. F. Chan, S. Karigiannis and C. C. Tsang [7] and M. Verbitsky [27], but these seem to be inconclusive on their own. More precisely, in [27] Verbitsky extended the Kähler identities to Riemannian manifolds equipped with a parallel form, and built a cohomology algebra with special properties using them. From this, he provided an alternative proof of formality for compact Kähler manifolds. This idea was later applied in [7] to torsion-free G_2 manifolds, where it is shown that these are almost formal — meaning that all triple Massey products vanish except possibly those involving three degree-2 cohomology classes. This condition is satisfied by all 7-manifolds with $b_1 = 0$.

Another strategy is contained in [9], where D. Crowley and J. Nordström studied the rational homotopy type of simply connected 7-manifolds, among others. In particular, they built the Bianchi–Massey tensor as a tool to detect nonformality of 7-manifolds, and they concluded that if there was a nonformal simply connected manifold with holonomy G_2 , it would have $b_2 \geq 4$. Prior to [9], G. R. Cavalcanti proved a weaker version of this result in [6], and also obtained a compact nonformal 7-dimensional manifold that satisfies the known topological obstructions for holonomy G_2 manifolds.

The relative scarcity of holonomy G_2 manifolds makes it hard to determine what phenomena might influence whether they are formal. Formality of one of Joyce’s examples was proved by Amann and

Taimanov [2]. The examples of Nordström in [22] are formal because they have $b_2 \leq 1$. The author is unaware of any attempts to address the formality of the examples with $b_2 \geq 4$ provided by A. Kovalev [19], Kovalev and N. H. Lee [20], or A. Corti, M. Haskins, Nordström and T. Pacini [8]. We provide a negative answer to the question of formality for compact manifolds with holonomy G_2 , proving the following result:

Theorem 1 *There exists a compact simply connected manifold with holonomy G_2 that is nonformal.*

The constructed manifold $(\tilde{M}, \tilde{\varphi})$ is obtained by resolving an orbifold (X, φ) which itself is the quotient of a flat 7-manifold under the action of an involution with fixed points. The singular locus of X consists of 10 disjoint associative manifolds, $N_1, \dots, N_{10}$, each of which is 2:1 covered by a torus. In Theorem 3, we obtain $(\tilde{M}, \tilde{\varphi})$ using the desingularization method of Joyce and Karigiannis [18], and in Remark 4, we argue that a resolution diffeomorphic to $\tilde{M}$ is obtained from the generalized Kummer construction of Joyce [15]. As explained in Remark 4, the choice of the first was made due to the author's familiarity with the topological properties of resolutions of orbifolds fitting into the setup of [18].

It was known that the cohomology groups of $\tilde{M}$ are determined by those of X and $N_1, \dots, N_{10}$ (see (17)), and the product structure depends on the embedding of the singular locus of X . More precisely, the square of the Thom class x_j of a connected component E_j in the exceptional divisor lies in $H^*(X)$, and it depends on the Thom class of the corresponding connected component N_j from which it arises; see (20). Here the orientation of N_j is determined by φ .

The novelty of this example is that we obtain a nonvanishing triple Massey product, even when the orbifold and the singular locus are formal, by positioning four of the connected components in a special configuration. Aside from the cohomology product formulas, the following facts ensure that the triple Massey product $\langle x_1 + x_2, x_7 + x_3, x_7 - x_3 \rangle$ is both well-defined and nonvanishing:

- (i) There are oriented cobordisms C and $\tilde{C}$ with $\partial C = N_3 \sqcup -N_7$ and $\partial \tilde{C} = N_1 \sqcup -N_2$. Here the minus sign indicates that the orientation is reversed.
- (ii) The cobordism C intersects N_1 negatively and N_2 positively.

Alternatively, this can be expressed by saying that $N_1 \sqcup -N_2$ and $N_3 \sqcup -N_7$ are nullcobordant and their linking number is nonzero.

This contrasts with the example examined by Amann and Taimanov [2], where both the orbifold and the singular locus are formal, but the cobordisms between singular components can be chosen to avoid transverse intersections with other components. This also differs from the examples of compact manifolds with $b_1 = 1$ that are equipped with a closed G_2 -structure and obtained by orbifold resolution techniques. The formal manifold in work of Fernández, Fino, Kovalev, and Muñoz [12] satisfies the same properties as the example analyzed in [2]. In work of Martín-Merchán [21], the orbifold is nonformal and the resolution carries a nonvanishing triple Massey product obtained by pulling back a nontrivial one from it. For topological reasons, the examples in [12; 21] do not admit metrics with holonomy contained in G_2 .

This paper is organized as follows: In Section 2 we construct $\tilde{M}$ and we prove that it is simply connected and it admits a metric with holonomy G_2 . In Section 3 we obtain its cohomology algebra and

show it is nonformal. The author has chosen to present most of the computations explicitly to ensure that the discussion is self-contained.

2 A simply connected compact manifold with holonomy G_2

Consider the flat 6-torus $T^6 = (\mathbb{C}/\Gamma)^3$, where $\Gamma = \mathbb{Z}\langle 1, i \rangle$, endowed with its standard flat $SU(3)$ -structure $(\omega, \Theta) = (\frac{1}{2}i \sum_k dz_k \bar{z}_k, dz_{123})$. The isometry

$$(1) \quad f: T^6 \rightarrow T^6, \quad f[z_1, z_2, z_3] = [iz_1, iz_2, -z_3]$$

preserves the $SU(3)$ -structure. The manifold

$$(2) \quad M = \mathbb{R} \times T^6 / (t, [z]) \sim (t + k, f^k[z]) \quad \text{for } k \in \mathbb{Z},$$

equipped with the flat metric has a parallel G_2 -structure $\varphi = dt \wedge \omega + \text{Re}(\Theta)$. Since $f^4 = \text{Id}$, the 7-torus $\mathbb{T} = \mathbb{R}/4\mathbb{Z} \times T^6$ covers M . The deck transformation group of the covering is $\mathbb{Z}_4$, generated by $F[t, [z]] = [t + 1, f[z]]$. We quotient M by the involutions

$$(3) \quad \iota_1, \iota_2: M \rightarrow M, \quad \iota_1[t, [z]] = [t, [-z_1, -z_2, z_3 + \tfrac{1}{2}]], \quad \iota_2[t, [z]] = [t, [-z_1, -z_2, z_3 + \tfrac{1}{2}i]].$$

Neither of these involutions, nor their compositions, have fixed points. Since they preserve φ , the quotient space $M' = M / \langle \iota_1, \iota_2 \rangle$ is a smooth manifold with a parallel G_2 -structure and a flat metric. Observe that there is a fibration $T^6 / \mathbb{Z}_2^2 \hookrightarrow M' \rightarrow S^1$. In addition, the involution

$$(4) \quad \kappa: M \rightarrow M, \quad \kappa[t, [z]] = [1 - t, [\bar{z}_1, \bar{z}_2, \bar{z}_3]],$$

is well-defined because $\kappa[t + 1, f[z]] = [-t, [-i\bar{z}_1, -i\bar{z}_2, -\bar{z}_3]] = [1 - t, [\bar{z}_1, \bar{z}_2, \bar{z}_3]] = \kappa[t, [z]]$. It determines a map $\kappa: M' \rightarrow M'$ as it commutes with both ι_1 and ι_2 . Moreover, κ has fixed points and $\kappa^*(\varphi) = \varphi$.

The orbifold $X = M' / \kappa$ can be alternatively described as the quotient of the torus $\mathbb{T} = \mathbb{R}/4\mathbb{Z} \times T^6$ by the group $D_4 \times \mathbb{Z}_2^2$. Here D_4 is the dihedral group spanned by κ and F (note that $\kappa \circ F = F^{-1} \circ \kappa$), whereas $\mathbb{Z}_2^2$ is generated by ι_1 and ι_2 .

Remark 2 The orbifold X is related to the examples provided by Joyce in [15]. More precisely, a change of variables transforms $X_1 = \mathbb{T}^7 / \langle F, \kappa \rangle$ into Joyce's Example 9, and $X_2 = \mathbb{T}^7 / \langle F, \kappa, \iota_1 \circ \iota_2 \rangle$ into Example 10. In particular, there is a branched covering $X_2 \rightarrow X$, because ι_1 has fixed points when it acts on X_2 (a fixed point of ι_1 on X_2 is the projection of $p = [0, [0, 0, \frac{1}{4}]]$ as $\iota_1(\kappa(p)) = [1, [0, 0, \frac{3}{4}]] = p$).

2.1 Singular locus

The singular locus of X is the projection of the fixed locus of $\kappa, \kappa \circ \iota_1, \kappa \circ \iota_2$, and $\kappa \circ \iota_1 \circ \iota_2$ on M . We observe that if γ is one of these maps, and we denote it by $\gamma[t, [z]] = [1 - t, \gamma'[z]]$, then a point $[t, [z]]$ with $t \in [0, 1)$ is fixed by γ if and only if $t = \frac{1}{2}$ and $[z] \in \text{Fix}(\gamma')$, or $t = 0$ and $[z] \in \text{Fix}(f^{-1} \circ \gamma')$. We now compute these and analyze their projections to X .

We begin with the level set $t = 0$. From $f^{-1}[z] = [-iz_1, -iz_2, -z_3]$, we obtain

$$\begin{aligned} f^{-1} \circ \kappa'[z] &= [-i\bar{z}_1, -i\bar{z}_2, -\bar{z}_3], & f^{-1} \circ (\kappa \circ \iota_1)'[z] &= [i\bar{z}_1, i\bar{z}_2, -\bar{z}_3 + \tfrac{1}{2}], \\ f^{-1} \circ (\kappa \circ \iota_2)'[z] &= [i\bar{z}_1, i\bar{z}_2, -\bar{z}_3 + \tfrac{1}{2}i], & f^{-1} \circ (\kappa \circ \iota_1 \circ \iota_2)'[z] &= [-i\bar{z}_1, -i\bar{z}_2, -\bar{z}_3 + \tfrac{1}{2}(1+i)]. \end{aligned}$$

Hence $\text{Fix}(f^{-1} \circ (\kappa \circ \iota_2)') = \text{Fix}(f^{-1} \circ (\kappa \circ \iota_1 \circ \iota_2)') = \emptyset$ (these act on y_3 as $y_3 \mapsto y_3 + \frac{1}{2}$), and

$$(5) \quad \text{Fix}(f^{-1} \circ \kappa') = \{[x_1 - ix_1, x_2 - ix_2, \varepsilon_3 + iy_3] \mid \varepsilon_3 \in \{0, \tfrac{1}{2}\}, x_1, x_2, y_3 \in [-\tfrac{1}{2}, \tfrac{1}{2}]\},$$

$$(6) \quad \text{Fix}(f^{-1} \circ (\kappa \circ \iota_1)') = \{[x_1 + ix_1, x_2 + ix_2, \delta_3 + iy_3] \mid \delta_3 \in \{-\tfrac{1}{4}, \tfrac{1}{4}\}, x_1, x_2, y_3 \in [-\tfrac{1}{2}, \tfrac{1}{2}]\}.$$

We denote by $N_1^1, N_1^2 \subset M$ the connected components of $\text{Fix}(\kappa)$ at $t = 0$ determined by $\varepsilon_3 = 0$ and $\varepsilon_3 = \frac{1}{2}$, respectively. Similarly, we denote by $N_2^1, N_2^2 \subset M$ the components of $\text{Fix}(\kappa \circ \iota_1)$ at $t = 0$ determined by $\delta_3 = -\frac{1}{4}$ and $\delta_3 = \frac{1}{4}$. Each of these is diffeomorphic to a 3-torus. The map ι_1 interchanges N_j^1 with N_j^2 , while ι_2 preserves each connected component. In addition, $N_j^k \rightarrow N_j^k/\iota_2$ is a 2:1 covering, and $H^1(N_j^k/\iota_2) = \langle [dy_3] \rangle$. The form dy_3 is harmonic because N_j^k is equipped with the flat metric. In fact, the map $N_1^1/\iota_2 \rightarrow \mathbb{R}/\mathbb{Z}$, $[x_1 - ix_1, x_2 - ix_2, iy_3] \mapsto [2y_3]$ is a submersion with a T^2 fiber. A similar statement holds for the remaining connected components. Therefore the fixed locus of κ on M' at $t = 0$ has two connected components, N'_1 and N'_2 . Both admit the nowhere-vanishing harmonic 1-form dy_3 . We denote by N_1 and N_2 their projections to X .

To analyze the singularities at $t = \frac{1}{2}$, we compute

$$\begin{aligned} \kappa'[z] &= [\bar{z}_1, \bar{z}_2, \bar{z}_3], & (\kappa \circ \iota_1)'[z] &= [-\bar{z}_1, -\bar{z}_2, \bar{z}_3 + \tfrac{1}{2}], \\ (\kappa \circ \iota_2)'[z] &= [-\bar{z}_1, -\bar{z}_2, \bar{z}_3 + \tfrac{1}{2}i], & (\kappa \circ \iota_1 \circ \iota_2)'[z] &= [\bar{z}_1, \bar{z}_2, \bar{z}_3 + \tfrac{1}{2}(1+i)]. \end{aligned}$$

Hence $\text{Fix}((\kappa \circ \iota_1)') = \text{Fix}((\kappa \circ \iota_1 \circ \iota_2)') = \emptyset$, and

$$(7) \quad \text{Fix}(\kappa') = \{[x_1 + i\varepsilon_1, x_2 + i\varepsilon_2, x_3 + i\varepsilon_3] \mid \varepsilon_1, \varepsilon_2, \varepsilon_3 \in \{0, \tfrac{1}{2}\}, x_1, x_2, x_3 \in [-\tfrac{1}{2}, \tfrac{1}{2}]\},$$

$$(8) \quad \text{Fix}((\kappa \circ \iota_2)') = \{[\varepsilon_1 + iy_1, \varepsilon_2 + iy_2, x_3 + i\delta_3] \mid \varepsilon_1, \varepsilon_2, \delta_3 - \tfrac{1}{4} \in \{0, \tfrac{1}{2}\}, y_1, y_2, x_3 \in [-\tfrac{1}{2}, \tfrac{1}{2}]\}.$$

We denote by $N_3^1, N_3^2 \subset M$ the connected components of $\text{Fix}(\kappa)$ at $t = \frac{1}{2}$ determined by $\varepsilon_1 = \varepsilon_2 = 0$ and $\varepsilon_3 = 0$ or $\varepsilon_3 = \frac{1}{2}$, respectively. We use lexicographic order for the pairs $(\varepsilon_1, \varepsilon_2)$ and label the remaining connected components of $\text{Fix}(\kappa)$ at $t = \frac{1}{2}$ by $N_4^1, N_4^2, N_5^1, N_5^2, N_6^1$, and N_6^2 , with the upper index 1 corresponding to $\varepsilon_3 = 0$. Similarly, we denote the connected components of $\text{Fix}(\kappa \circ \iota_2)$ by N_j^k , where $j = 7, \dots, 10$ and $k = 1, 2$, and the upper index $k = 1$ now indicates $\delta_3 = -\frac{1}{4}$. For instance,

$$N_7^1 = \{[\tfrac{1}{2}, [iy_1, iy_2, x_3 - \tfrac{1}{4}i]] \mid y_1, y_2, x_3 \in [-\tfrac{1}{2}, \tfrac{1}{2}]\}.$$

All these components are diffeomorphic to a 3-torus. We observe that ι_1 preserves each connected component, while ι_2 swaps N_j^1 with N_j^2 . Also, $N_j^k \rightarrow N_j^k/\iota_1$ is a 2:1 covering, and $H^1(N_j^k/\iota_1) = \langle [dx_3] \rangle$. Again, dx_3 is indeed harmonic on N_j^k and there are submersions $N_j^k/\iota_1 \rightarrow \mathbb{R}/\mathbb{Z}$ with a T^2 fiber. The fixed locus of κ on M' at $t = \frac{1}{2}$ has eight connected components, $N'_3, \dots, N'_{10}$. These admit the nowhere-vanishing harmonic 1-form dx_3 . We denote by $N_3, \dots, N_{10}$ their projections to X .

The submanifolds N_j^k and N_j' are associative submanifolds of (M, φ) and (M', φ) , respectively; see [18, Proposition 2.13]. We always assume that N_j^k (resp. N_j' , N_j) are oriented by $\varphi|_{N_j^k} = \operatorname{Re}(dz_{123})|_{N_j^k}$ (resp. $\varphi|_{N_j'}$, $\varphi|_{N_j}$). Using the coordinates at $t = 0$ and $t = \frac{1}{2}$ provided by (5)–(8), these forms are

$$(9) \quad \varphi|_{N_1^k} = 2 dx_1 \wedge dx_2 \wedge dy_3 \quad \text{for } k = 1, 2,$$

$$(10) \quad \varphi|_{N_2^k} = -2 dx_1 \wedge dx_2 \wedge dy_3 \quad \text{for } k = 1, 2,$$

$$(11) \quad \varphi|_{N_j^k} = dx_1 \wedge dx_2 \wedge dx_3 \quad \text{for } j = 3, 4, 5, 6 \text{ and } k = 1, 2,$$

$$(12) \quad \varphi|_{N_j^k} = -dy_1 \wedge dy_2 \wedge dx_3 \quad \text{for } j = 7, 8, 9, 10 \text{ and } k = 1, 2.$$

Note that if $j = 1$ (or $j = 2$), the projections of N_j^k to the first and second factors of $(\mathbb{C}/\Gamma)^3$ are the loops $[x - ix]$ (or $[x + ix]$) for $x \in [-\frac{1}{2}, \frac{1}{2}]$. These have length $\sqrt{2}$. The projection to the third factor has, of course, length 1. Therefore the volume of these tori is 2, which is why the factor of 2 appears in (9) and (10).

2.2 Resolution

The resolution method developed in [18] allows us to desingularize the orbifold X and to obtain a metric with holonomy G_2 on the resolution.

Theorem 3 *The resolution $\tilde{M}$ of X is a compact simply connected manifold that admits a metric with holonomy G_2 .*

Proof Since each connected component of the singular locus of X admits a nowhere-vanishing harmonic 1-form, the resolution method [18, Theorem 1.1] ensures that there is a resolution $\rho: \tilde{M} \rightarrow X$ and a torsion-free G_2 -structure $\tilde{\varphi}$ on $\tilde{M}$. It satisfies $\pi_1(\tilde{M}) = \pi_1(X)$. According to [17, Proposition 10.2.2], the holonomy of $(\tilde{M}, \tilde{\varphi})$ equals G_2 if and only if $\pi_1(\tilde{M})$ is finite. To finish the proof, it suffices to show that X is simply connected.

We first compute $\pi_1(M, p_0)$ with $p_0 = [\frac{1}{2}, [0, 0, 0]]$. From the long exact sequence of the fibration $T^6 \hookrightarrow M \rightarrow \mathbb{R}/\mathbb{Z}$, $[t, [z]] \mapsto [t]$, we obtain the short exact sequence

$$\{1\} = \pi_2(S^1) \rightarrow \pi_1(T^6) \rightarrow \pi_1(M) \rightarrow \pi_1(S^1) \rightarrow \{1\}.$$

Of course, the loop $\gamma_0(s) = [\frac{1}{2} + s, [0, 0, 0]]$ in M maps to the generator of $\pi_1(S^1)$, and provides a splitting of the short exact sequence. Therefore $\pi_1(M) = \pi_1(S^1) \ltimes \pi_1(T^6)$. We now analyze the product of $[\gamma_0]$ with elements in the image of $\pi_1(T^6) \hookrightarrow \pi_1(M)$. Given a loop $\gamma(s) = [\frac{1}{2}, \hat{\gamma}(s)]$ on M with $\hat{\gamma}(0) = [0, 0, 0]$, the loop $\gamma_0^{-1} \cdot \gamma \cdot \gamma_0$ is homotopic to $s \mapsto [\frac{3}{2}, \hat{\gamma}(s)] = [\frac{1}{2}, f^{-1} \circ \hat{\gamma}(s)]$ on M . (As usual, we write $\gamma^{-1}(s) = \gamma(1 - s)$.) The homotopy $H_r(s)$ is a sequence of three paths. First, move along the t -direction from $[\frac{3}{2}, [0, 0, 0]]$ to $[\frac{3}{2} - r, [0, 0, 0]]$. Then travel on the level set $\frac{3}{2} - r$ via the map $s \mapsto [\frac{3}{2} - r, \hat{\gamma}(s)]$, and finally return to $[\frac{3}{2}, [0, 0, 0]]$ by reversing the initial path.

We denote by $\gamma_k(s) = [\frac{1}{2}, \hat{\gamma}_k(s)]$ for $k = 1, \dots, 6$ the canonical generators of $\pi_1(\{\frac{1}{2}\} \times T^6, p_0)$. Namely, $\hat{\gamma}_1(s) = [s, 0, 0], \dots, \hat{\gamma}_6(s) = [0, 0, is]$. For $k = 1, 2$ we have $f^{-1} \circ \hat{\gamma}_{2k-1} = \hat{\gamma}_{2k}^{-1}$ and $f^{-1} \circ \hat{\gamma}_{2k} = \hat{\gamma}_{2k-1}$.

In addition, if $k = 5, 6$ then $f^{-1} \circ \hat{\gamma}_k = \hat{\gamma}_k^{-1}$. This implies the relation

$$(13) \quad [\gamma_{2k}]^{-1} = [\gamma_0]^{-1}[\gamma_{2k-1}][\gamma_0], \quad [\gamma_{2k-1}] = [\gamma_0]^{-1}[\gamma_{2k}][\gamma_0], \quad \text{for } k = 1, 2,$$

$$(14) \quad [\gamma_5]^{-1} = [\gamma_0]^{-1}[\gamma_5][\gamma_0], \quad [\gamma_6]^{-1} = [\gamma_0]^{-1}[\gamma_6][\gamma_0].$$

Since $\langle \iota_1, \iota_2 \rangle$ acts freely on M , the projection $q: M \rightarrow M'$ is a 4:1 regular cover with deck transformation group $\langle \iota_1, \iota_2 \rangle = \mathbb{Z}_2^2$; see [5, Chapter 3, Proposition 7.2]. Indeed, there is a short exact sequence

$$\{1\} \rightarrow \pi_1(M) \rightarrow \pi_1(M') \rightarrow \langle \iota_1, \iota_2 \rangle \rightarrow \{1\},$$

where the last map is defined by lifting a representative of an element in $\pi_1(M')$ to M with basepoint p_0 and then considering the map of the deck group that transforms p_0 to the endpoint of the lift. Consider the paths

$$\tilde{\gamma}_5, \tilde{\gamma}_6: [0, 1] \rightarrow M, \quad \tilde{\gamma}_5(s) = \left[\frac{1}{2}, [0, 0, \frac{1}{2}s]\right], \quad \tilde{\gamma}_6(s) = \left[\frac{1}{2}, [0, 0, \frac{1}{2}is]\right].$$

Then $\gamma'_5 = q \circ \tilde{\gamma}_5$ and $\gamma'_6 = q \circ \tilde{\gamma}_6$ are loops on M' . In the exact sequence, their homotopy classes project onto ι_1 and ι_2 , respectively. If $0 \leq k \leq 4$ we write $\gamma'_k = q \circ \gamma_k$. In addition, define the paths

$$\tilde{\gamma}_0, \gamma_7: [0, 1] \rightarrow M, \quad \tilde{\gamma}_0(s) = \left[\frac{1}{2} + s, [0, 0, \frac{1}{2}]\right], \quad \gamma_7(s) = \left[\frac{3}{2}, [0, 0, \frac{1}{2}s]\right] = \left[\frac{1}{2}, [0, 0, -\frac{1}{2}s]\right].$$

There is a homotopy with fixed endpoints on M from $\gamma_0^{-1} \cdot \tilde{\gamma}_5 \cdot \tilde{\gamma}_0$ to γ_7 . The way it is constructed is the same as before. Since $\gamma'_0 = q \circ \tilde{\gamma}_0$ and $q \circ \gamma_7(s) = q\left[\frac{1}{2}, [0, 0, \frac{1}{2} - \frac{1}{2}s]\right] = (\gamma'_5)^{-1}(s)$ we obtain $[\gamma'_0]^{-1}[\gamma'_5][\gamma'_0] = [\gamma'_5]^{-1}$ on M' ; a similar argument shows $[\gamma'_0]^{-1}[\gamma'_6][\gamma'_0] = [\gamma'_6]^{-1}$.

The involution $\kappa: M' \rightarrow M'$ has fixed points. Therefore the map $q'_*: \pi_1(M') \rightarrow \pi_1(X)$ induced by the projection $q': M' \rightarrow X$ is surjective; see [4, Chapter 2, Corollary 6.3]. Observe that κ reverses the direction of $\gamma'_0, \gamma'_2, \gamma'_4$, and γ'_6 . That is, if γ is one of these loops, then $\kappa \circ \gamma(s) = \gamma^{-1}(s)$. This implies that $q'_*[\gamma] = 1$ because $q' \circ \gamma = q' \circ \gamma''(s)$, where

$$\gamma''(s) = \begin{cases} \gamma(s) & \text{if } t \leq \frac{1}{2}, \\ \gamma^{-1}(s) & \text{if } t \geq \frac{1}{2}, \end{cases}$$

and γ'' represents the unit element in $\pi_1(M')$. Hence $q'_*[\gamma'_0] = q'_*[\gamma'_2] = q'_*[\gamma'_4] = q'_*[\gamma'_6] = 1$. From the relations (13), and $[\gamma'_5]^{-1} = [\gamma'_0]^{-1}[\gamma'_5][\gamma'_0]$, we obtain $q'_*[\gamma'_1] = q'_*[\gamma'_3] = q'_*[\gamma'_5]^2 = 1$. We finally show that, indeed, $q'_*[\gamma'_5] = 1$. We first consider the paths

$$\gamma_8, \gamma_9, \gamma_{10}: [0, 1] \rightarrow M, \quad \gamma_8(s) = \left[\frac{1}{2} + \frac{1}{2}s, [0, 0, 0]\right], \quad \gamma_9(s) = \left[1, [0, 0, \frac{1}{2}s]\right], \quad \gamma_{10}(s) = \left[\frac{1}{2} + \frac{1}{2}s, [0, 0, \frac{1}{2}]\right].$$

There is a basepoint preserving homotopy on M from $\gamma_8 \cdot \gamma_9 \cdot \gamma_{10}^{-1}$ to $\tilde{\gamma}_5$. Since $q \circ \gamma_8 = q \circ \gamma_{10}$, this implies that $(q \circ \gamma_8) \cdot (q \circ \gamma_9) \cdot (q \circ \gamma_8)^{-1}$ is homotopic to γ'_5 . We now observe that

$$\kappa \circ (q \circ \gamma_9)(s) = q[0, [0, 0, \frac{1}{2}s]] = q[1, [0, 0, -\frac{1}{2}s]] = q[1, [0, 0, \frac{1}{2} - \frac{1}{2}s]] = (q \circ \gamma_9)^{-1}(s).$$

The argument above ensures that the loop $q' \circ q \circ \gamma_9$ is trivial on X . Hence

$$q' \circ \gamma'_5 \sim (q' \circ q \circ \gamma_8) \cdot (q' \circ q \circ \gamma_9) \cdot (q' \circ q \circ \gamma_8)^{-1} \sim (q' \circ q \circ \gamma_8) \cdot (q' \circ q \circ \gamma_8)^{-1} \sim 1.$$

Therefore $\pi_1(\tilde{M}) = \pi_1(X) = \{1\}$. □

We denote by $\rho: \tilde{M} \rightarrow X$ the projection map, and $E_j = \rho^{-1}(N_j)$. The resolution described in [18] occurs on the fibers of the normal bundle $\nu(N_j)$ of $N_j \subset X$, which are of the form $\mathbb{C}^2/\mathbb{Z}_2$. These singular fibers are replaced by their algebraic resolution,

$$Y = \{(w, \ell) \in \mathbb{C}^2 \times \mathbb{CP}^1 \mid w \in \ell\} / (w, \ell) \sim (-w, \ell).$$

We describe this procedure in our setup. First, $N'_j = N_j^1/\iota$, where $\iota = \iota_2$ if $j = 1, 2$ or $\iota = \iota_1$ if $3 \leq j \leq 10$. Indeed, there is a $2:1$ cover between the normal bundles $\nu(N_j^1) \rightarrow \nu(N'_j)$ induced by the projection map $M \rightarrow M'$. The deck transformation group of the cover is generated by the involution ι_* . Both bundles have a complex structure I_j induced by the cross product with the vector field χ_j determined by the harmonic form; see [18, Remark 4.1]. That is, $\chi_j = \partial_{y_3}$ if $j = 1, 2$ or $\chi_j = \partial_{x_3}$ if $3 \leq j \leq 10$, and $I_j(X)$ is determined by the expression

$$g(I_j(X), Y) = \varphi(\chi_j, X, Y).$$

The action of ι_* on $\nu(N_j^1)$ preserves I_j because ι_* fixes φ , χ_j and g , and the projection $\nu(N_j^1) \rightarrow \nu(N'_j)$ is a complex homomorphism. In addition, the bundle $\nu(N_j^1)$ is trivial. Since $I_j(\partial_t) \in \langle \partial_{x_3}, \partial_{y_3} \rangle$, we pick a trivialization of the form $(\partial_t, I_j(\partial_t), X, I_j(X))$ where $X \in \langle \partial_{x_1}, \partial_{y_1}, \partial_{x_2}, \partial_{y_2} \rangle \cap \nu(N_j^1)$ is a unit-length vector field. For instance, if $j = 1$ this is $(\partial_t, \partial_{x_3}, (1/\sqrt{2})(\partial_{x_1} + \partial_{y_1}), (-1/\sqrt{2})(\partial_{x_2} + \partial_{y_2}))$, and if $j = 3$ this is $(\partial_t, -\partial_{y_3}, \partial_{y_1}, -\partial_{y_2})$. The orientation of $\nu(N_j^1)$ induced from N_j and M is determined by the form $(\chi_j \lrcorner \varphi|_{\nu(N_j^1)})^2$, and the frame above is positively oriented. Hence the trivialization yields an isomorphism of oriented vector bundles $\nu(N_j^1) \cong N_j^1 \times \mathbb{C}^2$, where $\mathbb{C}^2$ has the standard orientation. For any j , under the isomorphism, the generator of the deck group of $\nu(N_j^1) \rightarrow \nu(N'_j)$ is $\iota_*(x, w_1, w_2) = (\iota(x), w_1, -w_2)$.

The involution κ induces the $\mathbb{Z}_2$ -action on the fibers of $\nu(N'_j)$ given by $v_x \rightarrow \kappa_*(v_x) = -v_x$; see the proof of [21, Lemma 2.8]. This can be lifted to a map $\hat{\kappa}: N_j^1 \times \mathbb{C}^2 \rightarrow N_j^1 \times \mathbb{C}^2$ by means of the expression $\hat{\kappa}(x, w_1, w_2) = (x, -w_1, -w_2)$, which is induced by the involution on M fixing the points of that connected component. This is κ_* if $j = 1, 3, 4, 5, 6$, or $(\kappa \circ \iota_1)_*$ if $j = 2$, or $(\kappa \circ \iota_2)_*$ if $j = 7, 8, 9, 10$. This action commutes with ι_* . Thus $\nu(N_j) = \nu(N'_j)/\kappa_* \cong (N_j^1 \times \mathbb{C}^2)/(\iota_*, \hat{\kappa}) \cong (N_j^1 \times \mathbb{C}^2/\mathbb{Z}_2)/\iota_*$, and there is a $2:1$ cover $N_j^1 \times (\mathbb{C}^2/\mathbb{Z}_2) \rightarrow (N_j^1 \times \mathbb{C}^2/\mathbb{Z}_2)/\iota_*$.

The resolution of $N_j^1 \times \mathbb{C}^2/\mathbb{Z}_2$ is $N_j^1 \times Y$. The deck transformation ι_* lifts to $N_j^1 \times Y$ by means of the expression $\iota_*(x, [(w_1, w_2), [W_1 : W_2]]) = (\iota(x), [(w_1, -w_2), [W_1 : -W_2]])$. The resolution $P(N_j)$ of $\nu(N_j)$ is the quotient $(N_j^1 \times Y)/\iota_*$. Its exceptional divisor is $Q(N_j) = (N_j^1 \times \mathbb{CP}^1)/\iota_* \subset (N_j^1 \times Y)/\iota_*$, which inherits the product orientation from $N_j^1 \times \mathbb{CP}^1$, where $\mathbb{CP}^1$ is oriented by the Fubini–Study form. It turns out that $Q(N_j)$ is isomorphic to the trivial bundle because it is homotopic to it via

$$(15) \quad E_t = N_j^1 \times \mathbb{CP}^1 / (x, [W_1 : W_2]) \sim (\iota(x), [W_1 : e^{\pi i t} W_2]) \quad \text{for } t \in [0, 1].$$

To obtain $\tilde{M}$, the authors of [18] replace a neighborhood of N_j in X with a neighborhood of $Q(N_j)$ in $P(N_j)$. Therefore $E_j = Q(N_j)$. The identification between these pieces is given by the composition of the natural projection $P(N_j) \rightarrow \nu(N_j)$, a rescaling of the fibers of $\nu(N_j)$, and a modification of the exponential map obtained by precomposing it with the flow of a vector field of $\nu(N_j)$. The latter coincides

with the exponential map over the zero section, preserves the orientation, and is compatible with the involution; see [18, Definition 3.2, Lemma 3.5, and (6.1)].

Remark 4 The resolution method of Joyce [15, Theorem 2.2.3] also provides a resolution of the orbifold X and a metric with holonomy G_2 on the resolution. The orbifold X fits into this setup as $X = \mathbb{T}/(D_4 \times \mathbb{Z}_2^2)$, and a neighborhood of N_j is isometric to a neighborhood of the singular set of $(N_j^1 \times \mathbb{C}^2/\mathbb{Z}_2)/\iota_*$ where N_j^1 is a flat 3-torus and ι_* is an isometry of $N_j^1 \times \mathbb{C}^2/\mathbb{Z}_2$ acting freely on N_j^1 . The resolution is obtained by replacing a neighborhood of N_j in X with a neighborhood of $Q(N_j)$ in $P(N_j)$; the identification is the composition of the natural projection and the exponential map. Hence the resolution is diffeomorphic to $\tilde{M}$. We used the method in [18], even if it involves harder analytic tools that are not necessary here, to take advantage of the author's previous knowledge of the cohomology algebra of the resolution; see [21].

For later purposes, we also consider $q_0 = [1 : 0] \in \mathbb{CP}^1$ and observe that

$$(16) \quad L_j := (N_j^1 \times \{q_0\})/\iota_* = N_j \times \{q_0\}$$

is a submanifold of E_j that is diffeomorphic to N_j . Its orientation is determined by that on N_j .

3 Nonformality of the constructed manifold

3.1 Cohomology groups

We introduce some notation. Fix small tubular neighborhoods $O'_j \subset O''_j$ of N'_j on M' , and denote by $\pi_j: O''_j \rightarrow N'_j$ the nearest-point projection. Define $O_j = q'(O'_j)$ and $O_j^2 = q(O''_j)$, where $q': M' \rightarrow X$ is the quotient projection. To ease notation, let $\pi_j: O_j^2 \rightarrow N_j$ denote the nearest-point projection. In addition, we let $\text{pr}_j = \pi_j \circ \rho: \rho^{-1}(O_j^2) \rightarrow N_j$. From the discussion in Section 2.2 we deduce that it is a smooth map.

According to [18, Proposition 6.1], there are group isomorphisms

$$(17) \quad H^k(\tilde{M}) \cong H^k(X) \oplus \bigoplus_{j=1}^{10} H^{k-2}(N_j) \otimes \langle x_j \rangle.$$

Recall that Satake defines in [24] the de Rham algebra for orbifolds and proves that its cohomology algebra is isomorphic to the Čech cohomology with real coefficients. In [12, Remark 5], Fernández, Fino, Kovalev, and Muñoz deal with orbifolds obtained as the orbit space of a finite group action on a compact manifold. In this case, the de Rham algebra of the orbifold coincides with the subalgebra of invariant forms on the manifold. By averaging cohomology classes with respect to the group action, they observe that the cohomology algebra coincides with the subalgebra of invariant cohomology classes. Hence $H^*(X)$ is computed from the complex $(\Omega(M')^k, d)$, and $H^*(X) = H^*(M')^k$.

The inclusion $H^k(X) \subset H^k(\tilde{M})$ is determined by ρ^* (see Lemma 5). Each x_j has degree 2, and is represented by the Thom class $[\tau_j]$ of the submanifold $E_j \subset \tilde{M}$, oriented as in Section 2.2. Recall that the Thom class of E_j is the generator of the compactly supported cohomology group $H_c^2(\rho^{-1}(O_j))$ that integrates to 1 over each fiber of the normal bundle of E_j inside $\tilde{M}$; see [3, Proposition 6.18]. In addition, a closed form in $H^*(\tilde{M})$ that maps to $[\alpha] \otimes x_j \in H^*(N_j) \otimes x_j$ is $\text{pr}_j^*(\alpha) \wedge \tau_j$.

Given $\alpha \in \Omega^k(X)$, the restriction $\rho^*(\alpha)|_{\tilde{M} - \bigcup_j E_j}$ is a smooth form since $\rho: \tilde{M} - \bigcup_j E_j \rightarrow X - \bigcup_j N_j$ is a diffeomorphism. The next lemma describes how to find a representative α_s of a cohomology class $[\alpha] \in H^*(X)$ whose pullback to $\tilde{M}$ is smooth.

Lemma 5 *Let $W'_j \subset W''_j \subset O''_j$ be tubular neighborhoods of N'_j on M' and define $W_j = q(W'_j)$ and $W_j^2 = q(W''_j)$. For any $1 \leq j \leq 10$, if $\alpha \in \Omega^k(X)$ is closed on W_j^2 , then there is $\alpha_j \in \Omega^k(X)$ such that:*

- (i) *The forms α_j and α coincide on $X - W_j^2$. In addition, $\alpha_j - \alpha = d\beta_j$ where $\beta_j \in \Omega^{k-1}(X)$ is supported on W_j^2 .*
- (ii) *On W_j we have $\alpha_j = \pi_j^*(\alpha_j|_{N_j})$. In particular, $\rho^*(\alpha_j)$ is smooth on $\rho^{-1}(W_j)$ because $\rho^*(\alpha_j) = \text{pr}_j^*(\alpha_j|_{N_j})$ there.*

Therefore, given $[\alpha] \in H^k(X)$, there is a representative $\alpha_s \in \Omega^k(X)$ such that $\rho^*(\alpha_s)$ is smooth and $\rho^*(\alpha_s)|_{\rho^{-1}(O_j)} = \text{pr}_j^*(\alpha_s|_{N_j})$ for every $1 \leq j \leq 10$.

Proof The form $\beta = \pi_j^*(\alpha|_{N'_j}) \in \Omega^2(W''_j)$ is κ -invariant because $\pi_j \circ \kappa = \pi_j$. In addition, it is closed and satisfies $\beta|_{N'_j} = \alpha|_{N'_j}$. The Poincaré lemma allows us to find a form θ_j on W''_j such that $d\theta_j = (\beta - \alpha)|_{W''_j}$. Indeed, since β and α are κ -invariant, we can assume that θ_j is. Let h_j be a κ -invariant bump function which equals 0 on $M' - W''_j$ and 1 on W'_j . Then $h_j\theta_j \in \Omega^{k-1}(X)$ is supported on W_j^2 and $\alpha_j = \alpha + d(h_j\theta_j)$ satisfies the required properties. $\square$

Remark 6 In the proof of [18, Proposition 6.1], and following our notation, they use that the bundles $Q(N_j) \rightarrow N_j$ are trivial. Their justification is that the normal bundle $\nu(N_j)$ is trivial as a real bundle (see their Remark 2.14). There could be a gap between these statements, because the chosen complex structure might not be constant in the trivialization suggested by the authors. The space of complex lines could then differ from fiber to fiber. This is why we showed in (15) that $Q(N_j)$ is trivial.

Proposition 7 *The Betti numbers of $\tilde{M}$ are $(b_1, b_2, b_3) = (0, 11, 16)$. The cohomology groups are*

$$\begin{aligned} H^2(\tilde{M}) &= \langle [dz_{1\bar{2}} + dz_{\bar{1}2}] \rangle \oplus \langle \mathbf{x}_1, \dots, \mathbf{x}_{10} \rangle, \\ H^3(\tilde{M}) &= \langle [i dt \wedge dz_{1\bar{1}}], [i dt \wedge dz_{2\bar{2}}], [i dt \wedge dz_{3\bar{3}}], [i dt \wedge (dz_{1\bar{2}} - dz_{\bar{1}2})] \rangle \\ &\quad \oplus \langle [dz_{123} + dz_{\bar{1}\bar{2}\bar{3}}], [dz_{12\bar{3}} + dz_{\bar{1}\bar{2}3}] \rangle \oplus \langle [dy_3] \otimes \mathbf{x}_1, [dy_3] \otimes \mathbf{x}_2, [dx_3] \otimes \mathbf{x}_3, \dots, [dx_3] \otimes \mathbf{x}_{10} \rangle. \end{aligned}$$

Proof We first compute the cohomology groups of M . We observe,

$$H^k(M) = H^k(\mathbb{T})^F = [dt] \wedge H^{k-1}(T^6)^f \oplus H^k(T^6)^f.$$

It is then clear that $H^1(M) = \langle [dt] \rangle$; in addition, $\kappa^*(dt) = -dt$. For the remaining degrees, we analyze the map $f^*: H^k(T^6, \mathbb{C}) \rightarrow H^k(T^6, \mathbb{C})$ and then take real parts to find $H^k(T^6)^f$. That is, since f is a holomorphic map, it preserves each of the subspaces $H^{p,q}(T^6)$, and given a (p, q) -form α with $p > q$, the map f^* fixes $[\alpha + \bar{\alpha}]$ if and only if it fixes $[\alpha]$. Then it suffices to study the action of f^* on $H^{p,q}(T^6)$ with $p \geq q$.

The action of f^* in terms of the basis $([dz_{13}], [dz_{23}], [dz_{12}])$ of $H^{2,0}(T^6, \mathbb{C})$ is $\text{diag}(-i, -i, -1)$. For the basis of $H^{1,1}(T^6, \mathbb{C})$ given by $([dz_{1\bar{1}}], [dz_{2\bar{2}}], [dz_{3\bar{3}}], [dz_{1\bar{2}}], [dz_{\bar{1}2}], [dz_{1\bar{3}}], [dz_{2\bar{3}}], [dz_{\bar{1}3}], [dz_{\bar{2}3}])$,

the action of f^* is $\text{diag}(1, 1, 1, 1, 1, -i, -i, i, i)$. Hence

$$H^2(T^6)^f = \langle [i dz_{1\bar{1}}], [i dz_{2\bar{2}}], [i dz_{3\bar{3}}], [dz_{1\bar{2}} + dz_{\bar{1}2}], [i(dz_{1\bar{2}} - dz_{\bar{1}2})] \rangle.$$

All these classes are invariant under ι_1 and ι_2 . However, κ acts as $\text{diag}(-1, -1, -1, 1, -1)$.

The map f^* acts trivially on $H^{3,0}(T^6)$. For $H^{2,1}(T^6, \mathbb{C})$, consider the basis $([dz_{12\bar{3}}], [dz_{13\bar{2}}], [dz_{23\bar{1}}])$. Then f^* acts as $\text{diag}(1, -1, -1)$. Thus

$$H^3(T^6)^f = \langle [dz_{123} + dz_{\bar{1}\bar{2}\bar{3}}], [i(dz_{123} - dz_{\bar{1}\bar{2}\bar{3}})], [dz_{12\bar{3}} + dz_{\bar{1}\bar{2}3}], [i(dz_{12\bar{3}} - dz_{\bar{1}\bar{2}3})] \rangle.$$

These forms are fixed by ι_1 and ι_2 . In addition, the action of κ is $\text{diag}(1, -1, 1, -1)$. Finally, one deduces the result from our previous calculations, the equality $H^k(X) = H^k(M)^{\langle \kappa, \iota_1, \iota_2 \rangle}$, and (17). $\square$

Remark 8 As [17, Figure 12.3] shows, there is no holonomy G_2 manifold with $(b_2, b_3) = (11, 16)$ within the examples provided in Section 12 of that book. This pair does not appear in [23], where the Joyce–Karigiannis construction is used to resolve orbifolds of the form $(K3 \times T^3)/\mathbb{Z}_2^2$. The examples in [19, section 8] have $b_2 \leq 9$, and those in [20, section 6] with $b_2 = 11$ have $b_3 \geq 66$. The author could not find a twisted connected sum in [8] with this pair of Betti numbers.

3.2 Algebra structure

We now focus on the product under the isomorphism provided by (17). This discussion is based on [21, Proposition 6.4] and uses the notation introduced in Section 2.2. Similar computations have been done in [12; 25].

We first claim that the Thom form of the normal bundle of $N_j^1 \times \mathbb{CP}^1 \subset N_j^1 \times Y$ is the extension of

$$(18) \quad \tau_j = \frac{1}{\pi i} d(b(r^2) \partial \log(r^2)) = \frac{2}{\pi i} b'(r^2) r dr \wedge \partial \log(r^2) + \frac{b(r^2)}{\pi i} \bar{\partial} \partial \log r^2,$$

where r is the radial coordinate on $\mathbb{C}^2$, and $b(s)$ is an increasing function that equals -1 on $s \leq \varepsilon^2$ and 0 on $s \geq 4\varepsilon^2$. Of course, we choose ε small enough that τ_j is supported in $\rho^{-1}(O_j)$. Observe that τ_j is smooth because the second term is proportional to the pullback of the Fubini–Study form on $\mathbb{CP}^1$.

To verify this claim, consider a unit-length vector $(u_1, u_2) \in \mathbb{C}^2$ and the holomorphic embedding to a fiber of the tautological line bundle $F = \{(w, [u_1 : u_2]) \mid w \in [u_1 : u_2]\}$, given by $e: \mathbb{C} \rightarrow F$, $e(\lambda) = (\lambda u_1, \lambda u_2, [u_1 : u_2])$. This induces an orientation-preserving embedding onto the fiber $F_p = \{p\} \times F/\mathbb{Z}_2$ of the bundle $N_j^1 \times Y \rightarrow N_j^1 \times \mathbb{CP}^1$ that we denote by $e_p: \mathbb{C}/\mathbb{Z}_2 \rightarrow F_p$.

We observe that $e^*(\partial \log(r^2)) = \partial(\log |\lambda|^2) = (d\lambda)/\lambda$, so that $e^*\tau_j = 1/(\pi i) b'(|\lambda|^2) d\bar{\lambda} \wedge d\lambda$. Since $e^*\tau_j$ vanishes when $|\lambda| \notin [\varepsilon, 2\varepsilon]$ we have

$$\begin{aligned} \int_{\mathbb{C}/\mathbb{Z}_2} e_p^*(\tau_j) &= \frac{1}{2} \int_{\mathbb{C}} e^*(\tau_j) = \frac{1}{2\pi i} \int_{|\lambda| \in [\varepsilon, 2\varepsilon]} d\left(\frac{b(|\lambda|^2) d\lambda}{\lambda}\right) \\ &= \frac{1}{2\pi i} \int_{|\lambda|=2\varepsilon} \frac{b(|\lambda|^2) d\lambda}{\lambda} - \frac{1}{2\pi i} \int_{|\lambda|=\varepsilon} \frac{b(|\lambda|^2) d\lambda}{\lambda} = \frac{1}{2\pi i} \int_{|\lambda|=\varepsilon} \frac{d\lambda}{\lambda} = 1. \end{aligned}$$

This proves that τ_j integrates to 1 over the fibers of the normal bundle, and therefore it is a representative of the Thom class of that bundle. Note in addition that the Thom class of the tautological line bundle over $\mathbb{CP}^1$ is $\frac{1}{2}\tau_j$.

The form τ_j is ι_* invariant, and hence the Thom form of $Q(N_j) \subset P(N_j)$ is the pushforward of τ_j . We also denote it by τ_j . Since the second term in (18) is proportional to the pullback of the Fubini–Study form of $\mathbb{CP}^1$, the form τ_j^2 vanishes in a neighborhood of $Q(N_j)$. Hence it induces by pushforward a form $\rho_*(\tau_j^2)$ on X with compact support on O_j . Let V_j be a tubular neighborhood of N_j with $O_j \subset V_j$ and let $\text{Th}[N_j] \in H_c^4(V_j)$ be the Thom class of $N_j \subset X$ oriented by $\varphi|_{N_j}$. Then

$$(19) \quad [\rho_*(\tau_j^2)] = -2 \text{Th}[N_j] \in H_c^4(V_j) \subset H^4(X),$$

because $\int_{\mathbb{C}^2/\mathbb{Z}_2} \rho_*(\tau_j^2) = \int_Y \tau_j^2 = \int_{\mathbb{CP}^1} \tau_j = 1/(\pi i) \int_{\mathbb{CP}^1} \partial \bar{\partial} \log(r^2) = -2$ (or alternatively, $\int_Y \tau_j^2 = [\mathbb{CP}^1][\mathbb{CP}^1] = -2$). In addition, when $j \neq k$ the supports of τ_j and τ_k are disjoint; hence $\tau_j \wedge \tau_k = 0$. Finally, letting $[\alpha] \in H^k(X)$, by Lemma 5 we can assume that it satisfies $\alpha = \pi_j^*(\alpha|_{N_j})$ on O_j ; then $\rho^*(\alpha) \wedge \tau_j = \text{pr}_j^*(\alpha|_{N_j}) \wedge \tau_j$. Under the isomorphism in (17), our discussion yields the following product rules:

$$(20) \quad \mathbf{x}_j^2 = -2 \text{Th}[N_j], \quad \mathbf{x}_j \cdot \mathbf{x}_k = 0 \quad \text{if } j \neq k,$$

$$(21) \quad [\alpha] \cdot \mathbf{x}_j = [\alpha|_{N_j}] \otimes \mathbf{x}_j \quad \text{if } [\alpha] \in H^*(X).$$

The Thom class of N_j (viewed in $H^*(X)$) coincides with its Poincaré dual. Recall (see [3, p. 51]) that given a compact oriented n -dimensional manifold M , the Poincaré dual of a k -dimensional compact oriented submanifold $N \subset M$ is the unique cohomology class $\text{PD}[N] \in H^{n-k}(M)$ such that for any representative $v \in \Omega^{n-k}(M)$,

$$(22) \quad \int_N \alpha = \int_M \alpha \wedge v \quad \text{for all } [\alpha] \in H^k(M).$$

This is a consequence of Poincaré duality, namely $(H^k(M))^* = H^{n-k}(M)$, which also holds for orbifolds; see [24, Theorem 3]. Hence a similar definition applies to orbifolds. In Lemma 9 we find $\text{PD}[N_j] \in H^4(X)$ and in Lemma 10 we deduce from (19) that the self-intersection of $[E_j]$ is $-2[L_j]$, where L_j is defined by (16).

Lemma 9 Write $\beta_1 = i \, dt \wedge (dz_{123} - dz_{\bar{1}\bar{2}\bar{3}})$, and $\beta_2 = i \, dt \wedge (dz_{12\bar{3}} - dz_{\bar{1}\bar{2}3})$. Then

$$\text{PD}[N_1] = \text{PD}[N_2] = [\beta_1 + \beta_2], \quad \text{PD}[N_3] = \text{PD}[N_j] = \frac{1}{2}[\beta_1 - \beta_2] \quad \text{for } 4 \leq j \leq 10.$$

Proof Following the proof of Proposition 7, we obtain that $[\beta_1], [\beta_2] \in H^4(X)$. Let $[\alpha] \in H^3(X)$; then we can choose $\alpha = dt \wedge \beta + \lambda_1(dz_{123} + dz_{\bar{1}\bar{2}\bar{3}}) + \lambda_2(dz_{12\bar{3}} + dz_{\bar{1}\bar{2}3})$, where $f^*[\beta] = [\beta]$ and $\kappa^*[\beta] = -[\beta]$. Let $\mu_1, \mu_2 \in \mathbb{R}$. Since $M - \bigcup_{j,k} N_j^k \rightarrow X - \bigcup_j N_j$ is an 8 : 1 cover,

$$\int_X \alpha \wedge (\mu_1 \beta_1 + \mu_2 \beta_2) = \frac{1}{8} \int_M -2i(\mu_1 \lambda_1 - \mu_2 \lambda_2) \, dt \wedge dz_{1\bar{1}2\bar{2}3\bar{3}} = 2(\mu_1 \lambda_1 - \mu_2 \lambda_2).$$

Using the coordinates described in (5), we compute $\alpha|_{N_1^k} = 4(\lambda_1 - \lambda_2) dx_{12} \wedge dy_3$. Similarly, from (6) we obtain $\alpha|_{N_2^k} = -4(\lambda_1 - \lambda_2) dx_{12} \wedge dy_3$. Taking into account the induced orientations (9) and (10) of N_j^k , we deduce $\int_{N_j} \alpha = \frac{1}{2} \int_{N_j^1} \alpha = 2(\lambda_1 - \lambda_2)$ for $j = 1, 2$. In addition, if $3 \leq j \leq 6$ then $\alpha|_{N_j^k} = 2(\lambda_1 + \lambda_2) dx_{123}$ and if $7 \leq j \leq 10$ then $\alpha|_{N_j^k} = -2(\lambda_1 + \lambda_2) dy_{12} \wedge dx_3$. Of course, we used the coordinates in (7) and (8). Equations (11) and (12) ensure $\int_{N_j} \alpha = \frac{1}{2} \int_{N_j^1} \alpha = (\lambda_1 + \lambda_2)$ for $j \geq 3$. The statement follows from these calculations. For instance, if $3 \leq j \leq 7$,

$$\int_{N_j} \alpha = \lambda_1 + \lambda_2 = 2\left(\frac{1}{2}\lambda_1 - \left(-\frac{1}{2}\lambda_2\right)\right) = \int_X \alpha \wedge \left(\frac{1}{2}\beta_1 - \frac{1}{2}\beta_2\right),$$

and (22) implies $\text{PD}[N_j] = \frac{1}{2}[\beta_1 - \beta_2]$. $\square$

Lemma 10 *The following equalities hold on $H^4(\tilde{M})$:*

$$[\tau_1]^2 = [\tau_2]^2 = -2 \text{PD}[L_1] = -2 \text{PD}[L_2], \quad [\tau_3]^2 = [\tau_j]^2 = -2 \text{PD}[L_3] = -2 \text{PD}[L_j] \quad \text{for } 4 \leq j \leq 10.$$

Proof Fix $1 \leq j \leq 10$. We first prove that $[\tau_j]^2 = -2 \text{PD}[L_j]$. Let $[\alpha] \in H^3(\tilde{M})$; according to (17) and Lemma 5 we can write $\alpha = \rho^*(\beta) + \sum_{k=1}^{10} \text{pr}_k^*(\lambda_k) \wedge \tau_k$, where $[\lambda_k] \in H^1(N_k)$ and $[\beta] \in H^3(X)$ satisfies $\beta|_{O_k} = \pi_k^*(\beta|_{N_k})$ for every $1 \leq k \leq 10$.

On the one hand, $\alpha \wedge \tau_j^2 = \rho^*(\beta) \wedge \tau_j^2$ because $\tau_j \wedge \tau_k = 0$ if $j \neq k$ and $\tau_j^3 = 0$. For the last equality note that Y is 4-dimensional and that for every $[p, q] \in (N_j^1 \times Y)/\iota_*$ we have $\tau_j|_{[p, q]} \in \Lambda^2 T_q Y$ by definition (18). Equation (19) yields $\int_{\tilde{M}} \alpha \wedge \tau_j^2 = \int_{\tilde{M}} \rho^*(\beta) \wedge \tau_j^2 = \int_X \beta \wedge \rho_*(\tau_j^2) = -2 \int_{N_j} \beta$.

On the other hand, $\alpha|_{L_j} = \rho^*(\beta)|_{L_j}$ because $\tau_k|_{L_j} = 0$ for any $1 \leq k \leq 10$. Observe that if $j = k$ this follows from $\tau_j|_{[p, q_0]} \in \Lambda^2 T_{q_0} Y$ when $(p, q_0) \in L_j$. Hence $\int_{L_j} \alpha = \int_{L_j} \rho^*(\beta) = \int_{N_j \times \{q_0\}} \text{pr}_j^*(\beta|_{N_j}) = \int_{N_j} \beta = -\frac{1}{2} \int_{\tilde{M}} \alpha \wedge \tau_j^2$. This implies $[\tau_j]^2 = -2 \text{PD}[L_j]$; the result now follows from Lemma 9. $\square$

3.3 A nonvanishing triple Massey product

We first recall the definition of a triple Massey product; see [11, Definition 2.89] or [26, Definition 6.1]. For cohomology classes $[\alpha_1], [\alpha_2], [\alpha_3] \in H^*$, we say that their triple Massey product $\langle [\alpha_1], [\alpha_2], [\alpha_3] \rangle$ is well-defined if

$$[\alpha_1][\alpha_2] = 0 \quad \text{and} \quad [\alpha_2][\alpha_3] = 0.$$

If these conditions are satisfied, we pick forms ξ_{12} and ξ_{23} with $d\xi_{12} = \alpha_1 \wedge \alpha_2$ and $d\xi_{23} = \alpha_2 \wedge \alpha_3$. Then

$$(23) \quad y = \xi_{12} \wedge \alpha_3 - (-1)^{\deg(\alpha_1)} \alpha_1 \wedge \xi_{23}$$

is a closed form. Let $\mathcal{I}$ be the ideal in H^* generated by $[\alpha_1]$ and $[\alpha_3]$, and denote the projection map by $\pi: H^* \rightarrow H^*/\mathcal{I}$. The triple Massey product is defined as $\langle [\alpha_1], [\alpha_2], [\alpha_3] \rangle = \pi[y]$. We say that it vanishes if $\pi[y] = 0$, that is, if $[y] \in \mathcal{I}$.

It is well-known that all triple Massey products vanish on a formal manifold; see, for instance, [11, Proposition 2.90]. The goal of this section is to prove that the triple Massey product on $H^*(\tilde{M})$, $\langle [\tau_1 + \tau_2], [\tau_7 + \tau_3], [\tau_7 - \tau_3] \rangle$, does not vanish. As shown in the proof of Theorem 17 and Remarks 13, 15, and 18, this occurs because the oriented submanifolds $N'_1 \sqcup -N'_2$ and $N'_3 \sqcup -N'_7$ of M' are nullcobordant,

and their linking number is nonzero. Here the minus sign indicates that the orientation of the submanifold is reversed. The linking number (see [3, pp. 229–234]) of two disjoint nullhomologous oriented 3-dimensional submanifolds A and B of M' is

$$(24) \quad \int_{M'} \xi_A \wedge \tau_B = \int_{M'} \tau_A \wedge \xi_B, \quad d\xi_A = \tau_A, \quad d\xi_B = \tau_B, \quad [\tau_A] = \text{Th}[A], \quad [\tau_B] = \text{Th}[B].$$

If $A = \partial C_A$ and C_A intersects B transversally, the linking number is the oriented intersection number between C_A and B .

Note that $\langle [\tau_1 + \tau_2], [\tau_7 + \tau_3], [\tau_7 - \tau_3] \rangle$ is well-defined because $\tau_j \wedge \tau_k = 0$ when $j \neq k$ and $[\tau_7]^2 - [\tau_3]^2 = 0$ by Lemma 10. To find a primitive of $\tau_7^2 - \tau_3^2$, we first construct a form $\beta \in \Omega^3(M)$ with $d\beta = v_3^1 - v_7^1$, where v_j^1 is a representative of the Thom class of N_j^1 for $j = 3, 7$. The construction uses a cobordism between N_3^1 and N_7^1 .

Proposition 11 *There are small tubular neighborhoods V_3^1 and V_7^1 of N_3^1 and N_7^1 and forms $\beta \in \Omega^3(M)$ and $v_3^1, v_7^1 \in \Omega^4(M)$ such that:*

- (i) *The forms v_3^1 and v_7^1 are closed and supported on V_3^1 and V_7^1 , respectively. The cohomology class $[v_j^1] \in H_c^4(V_j^1)$ is the Thom class of N_j^1 .*
- (ii) *The form β vanishes in a neighborhood of the level set $t = \frac{1}{2}$ and $d\beta = v_3^1 - v_7^1$. In particular, it is closed on $M - (V_3^1 \cup V_7^1)$.*
- (iii) *We have $[\beta|_{N_1^k}] = -\frac{1}{2}[\varphi|_{N_1^k}]$, $[\beta|_{N_2^k}] = \frac{1}{2}[\varphi|_{N_2^k}]$.*

Proof For convenience, we use $t \in [\frac{1}{2}, \frac{3}{2}]$ instead of $t \in [0, 1]$. We first find a cobordism on M between N_3^1 and N_7^1 . We view points in N_3^1 at the level set $t = \frac{1}{2}$ and we consider the coordinates in (7) and the orientation in (11). Instead, we view N_7^1 inside the level set $t = \frac{3}{2}$. That is, since $[\frac{1}{2}, [iy_1, iy_2, x_3 - \frac{1}{4}i]] = [\frac{3}{2}, [-y_1, -y_2, -x_3 + \frac{1}{4}i]]$, we describe

$$(25) \quad N_7^1 = \{[\frac{3}{2}, [x_1, x_2, x_3 + \frac{1}{4}i]] \mid x_1, x_2, x_3 \in [-\frac{1}{2}, \frac{1}{2}]\}.$$

Its orientation is given by $\varphi|_{N_7^1} = dx_{123}$. Let $f: [\frac{1}{2}, \frac{3}{2}] \rightarrow [0, \frac{1}{4}]$ be a nondecreasing function that equals 0 if $t \leq \frac{9}{8}$, and $\frac{1}{4}$ if $t \geq \frac{5}{4}$. Define the embedding $\Psi: [\frac{1}{2}, \frac{3}{2}] \times N_3^1 \rightarrow M$ by

$$(26) \quad \Psi(t, [\frac{1}{2}, [x_1, x_2, x_3]]) = [t, [x_1, x_2, x_3 + if(t)]].$$

Then $C = \text{Im}(\Psi)$ is a manifold with boundary, $N_3^1 \cup N_7^1$. We trivialize and orient the bundle TC using the frame $(-\partial_t - f'(t)\partial_{y_3}, \partial_{x_1}, \partial_{x_2}, \partial_{x_3})$. The normal bundle $\nu(C)$ is then oriented and trivialized by $(\partial_{y_1}, \partial_{y_2}, \partial_{y_3} - f'(t)\partial_t)$. The basis $(\partial_{x_1}, \partial_{x_2}, \partial_{x_3})$ orients N_j^1 , and the outward-pointing vector is $V_o = -\partial_t$ on N_3^1 and $V_o = \partial_t$ on N_7^1 . Since a basis (e_1, e_2, e_3) of ∂C is positive if (V_o, e_1, e_2, e_3) is positive on C , we have that $\partial C = N_3^1 \sqcup -N_7^1$ as oriented manifolds.

We now find β using the cobordism; the Thom forms v_3^1 and v_7^1 will later be constructed from $d\beta$. We observe that if $j = 3, 7$, the bundle $\nu(N_j^1)$ is oriented by $(\partial_t, \partial_{y_1}, \partial_{y_2}, \partial_{y_3})$. Hence

$$\nu(N_j^1) \cong \langle \partial_t|_{N_j^1} \rangle \oplus \nu(C)|_{N_j^1}$$

as oriented bundles. The Thom form of $\nu(N_j^1)$ is then the product of the Thom forms of the bundles $\langle \partial_t|_{N_j^1} \rangle$ and $\nu(C)|_{N_j^1}$; see [3, Proposition 6.19]. This is the idea that makes our strategy possible.

Let $\zeta: [-\frac{1}{2}, \frac{1}{2}]^3 \rightarrow \mathbb{R}$ be a positive radial function whose integral is 1 and whose support is contained in $B_{2\delta}(0) - B_\delta(0)$, for a small δ . Consider a periodic extension of ζ to $\mathbb{R}^3/\mathbb{Z}^3$. Set $h: [\frac{1}{2}, \frac{3}{2}] \rightarrow \mathbb{R}$ a bump function that equals 0 on $[\frac{1}{2}, \frac{1}{2} + \varepsilon] \cup [\frac{3}{2} - \varepsilon, \frac{3}{2}]$, and 1 on $[\frac{1}{2} + 2\varepsilon, \frac{3}{2} - 2\varepsilon]$. Then

$$(27) \quad \beta|_{[t, [x_1 + iy_1, x_2 + iy_2, x_3 + iy_3]]} = h(t)\zeta(y_1, y_2, y_3 - f(t)) dy_1 \wedge dy_2 \wedge (dy_3 - f'(t) dt) \quad \text{for } t \in [\frac{1}{2}, \frac{3}{2}]$$

induces a well-defined form on M that vanishes around $t = \frac{1}{2}$. Indeed, around a small neighborhood of the interior of C , the expression $\zeta(y_1, y_2, y_3 - f(t)) dy_1 \wedge dy_2 \wedge (dy_3 - f'(t) dt)$ defines a closed form whose pullback to $\nu(\text{Int}(C))$ represents the Thom class of that oriented bundle. The support of the form

$$d\beta = h'(t)\zeta(y_1, y_2, y_3 - f(t)) dt \wedge dy_1 \wedge dy_2 \wedge dy_3 \quad \text{for } t \in [\frac{1}{2}, \frac{3}{2}]$$

is contained on a neighborhood of $N_3^1 \cup N_7^1$, and we can write $d\beta = v_3^1 - v_7^1$, where the support of v_j^1 is contained in the following neighborhood V_j^1 of N_j^1 :

$$\begin{aligned} V_3^1 &= \{[t, [x_1 + iy_1, x_2 + iy_2, x_3 + iy_3]] \mid |t - \frac{1}{2}| < 2\varepsilon, \|(y_1, y_2, y_3)\| < 2\delta\}, \\ V_7^1 &= \{[t, [x_1 + iy_1, x_2 + iy_2, x_3 + iy_3]] \mid |t - \frac{3}{2}| < 2\varepsilon, \|(y_1, y_2, y_3 - \frac{1}{4})\| < 2\delta\}. \end{aligned}$$

Moreover, if $[t, [z]] \in V_3^1$ satisfies $\frac{1}{2} - 2\varepsilon < t < \frac{1}{2}$ and $\|(y_1, y_2, y_3)\| < 2\delta$, from $f(t+1) = \frac{1}{4}$ we obtain $d\beta|_{[t, [z]]} = d\beta|_{[t+1, [iz_1, iz_2, -z_3]]} = h'(t+1)\zeta(x_1, x_2, -y_3 - \frac{1}{4}) dt \wedge dx_1 \wedge dx_2 \wedge d(-y_3)$. This vanishes because $|-y_3 - \frac{1}{4}| \geq \frac{1}{4} - 2\delta$ and ζ is supported around 0. Since h is constant on $[\frac{1}{2}, \frac{1}{2} + \varepsilon]$, we also have $d\beta = 0$ at points $[t, [z]] \in V_3^1$ with $\frac{1}{2} \leq t < \frac{1}{2} + \varepsilon$. In the same way, $d\beta = 0$ at points $[t, [z]] \in V_7^1$ with $\frac{3}{2} - \varepsilon < t < \frac{3}{2} + 2\varepsilon$.

We now check that v_3^1 represents the Thom class of N_3^1 . First, v_3^1 is closed because $d\beta$ is and the support of v_3^1 is disjoint from that of v_7^1 . In addition, let $p \in N_3^1$ and let F_p be the fiber of V_3^1 at p , oriented by $(\partial_t, \partial_{y_1}, \partial_{y_2}, \partial_{y_3})$. Since v_3^1 is supported on the level sets $\frac{1}{2} + \varepsilon < t < \frac{1}{2} + 2\varepsilon$ where $f = 0$,

$$\int_{F_p} v_3^1 = \int_{t=\frac{1}{2}+\varepsilon}^{\frac{1}{2}+2\varepsilon} h'(t) dt \int_{\|(y_1, y_2, y_3)\| < 2\delta} \zeta(y_1, y_2, y_3) dy_1 \wedge dy_2 \wedge dy_3 = h(\frac{1}{2} + 2\varepsilon) - h(\frac{1}{2} + \varepsilon) = 1.$$

Hence $[v_3^1] \in H_c^4(V_3^1)$ is the Thom class of N_3^1 . A similar computation shows that v_7^1 represents the Thom class of N_7^1 .

Of course, β is closed on $M - (V_3^1 \cup V_7^1)$. In particular, if $j, k = 1, 2$, then $\beta|_{N_j^k}$ determines an element in $H^3(N_j^k)$. To prove the third claim, it is sufficient to show $\int_{N_j^k} \beta|_{N_j^k} = \sigma_j$ where $\sigma_1 = -1$ and $\sigma_2 = 1$ because $\int_{N_j^k} \varphi = 2$. For convenience, we rewrite the formulas for N_j^k at the level set $t = 1$:

$$(28) \quad N_1^k = \{[1, [x_1 + ix_1, x_2 + ix_2, -\varepsilon_k - iy_3]] \mid x_1, x_2, y_3 \in [-\frac{1}{2}, \frac{1}{2}]\}, \quad \varepsilon_1 = 0, \quad \varepsilon_2 = \frac{1}{2},$$

$$(29) \quad N_2^k = \{[1, [-x_1 + ix_1, -x_2 + ix_2, -\delta_k - iy_3]] \mid x_1, x_2, y_3 \in [-\frac{1}{2}, \frac{1}{2}]\}, \quad \delta_1 = -\frac{1}{4}, \quad \delta_2 = \frac{1}{4}.$$

These are oriented by $\varphi|_{N_1^k} = 2 dx_1 \wedge dx_2 \wedge dy_3$ and $\varphi|_{N_2^k} = -2 dx_1 \wedge dx_2 \wedge dy_3$. Since $h(1) = 1$ and $f(1) = 0$, we have

$$\begin{aligned} \int_{N_1^k} \beta &= \int_{\|(x_1, x_2, -y_3)\| < 2\delta} \zeta(x_1, x_2, -y_3) dx_1 \wedge dx_2 \wedge d(-y_3) = -1, \\ \int_{N_2^k} \beta &= \int_{\|(x_1, x_2, -y_3)\| < 2\delta} \zeta(x_1, x_2, -y_3) dx_1 \wedge dx_2 \wedge d(-y_3) = 1. \end{aligned}$$

Here we used that $(x_1, x_2, y_3) \rightarrow (x_1, x_2, -y_3)$ reverses the orientation, and the expressions for $\varphi|_{N_j^k}$. $\square$

Remark 12 Following the notation of the proof, we provide a geometric justification of the equality $\int_{N_j^k} \beta|_{N_j^k} = \sigma_j$ for $j, k = 1, 2$. Observe that $N_j^k \cap C = \{p_j^k\}$, where $p_1^k = [1, [0, 0, -\varepsilon_k]]$ and $p_2^k = [1, [0, 0, -\delta_k]]$. Pick the basis $b_0 = (-\partial_t, \partial_{x_1}, \partial_{x_2}, \partial_{x_3})$ of $TC|_{N_j^k}$, and consider the positively oriented basis of N_j^k viewed at the level set $t = 1$ given by $b_j = (-\sigma_j \partial_{x_1} + \partial_{y_1}, -\sigma_j \partial_{x_2} + \partial_{y_2}, \sigma_j \partial_{y_3})$ —note that $\varphi(-\sigma_j \partial_{x_1} + \partial_{y_1}, -\sigma_j \partial_{x_2} + \partial_{y_2}, \sigma_j \partial_{y_3}) = 2$. Then (b_0, b_1) is a negative basis for $T_{p_1^k} M$ and (b_0, b_2) is a positive basis for $T_{p_2^k} M$ because

$$\begin{aligned} (-\partial_t) \wedge \partial_{x_1} \wedge \partial_{x_2} \wedge \partial_{x_3} \wedge (-\sigma_j \partial_{x_1} + \partial_{y_1}) \wedge (-\sigma_j \partial_{x_2} + \partial_{y_2}) \wedge (\sigma_j \partial_{y_3}) \\ = -\sigma_j \partial_t \wedge \partial_{x_1} \wedge \partial_{x_2} \wedge \partial_{x_3} \wedge \partial_{y_1} \wedge \partial_{y_2} \wedge \partial_{y_3} \\ = \sigma_j \partial_t \wedge \partial_{x_1} \wedge \partial_{y_1} \wedge \partial_{x_2} \wedge \partial_{y_2} \wedge \partial_{x_3} \wedge \partial_{y_3}. \end{aligned}$$

Around C and on the level sets where $h = 1$, namely $\frac{1}{2} + 2\varepsilon < t < \frac{3}{2} - 2\varepsilon$, the form

$$\beta = \zeta(y_1, y_2, y_3 - f(t)) dy_1 \wedge dy_2 \wedge (dy_3 - f'(t) dt)$$

is a representative of the Thom class of $\nu(C)$. Therefore $[\beta|_{N_j^k}]$ is the Thom class of p_j^k in N_j^k , oriented negatively if $j = 1$ or positively if $j = 2$. Hence the integral of β on N_1^k (resp. N_2^k) equals -1 (resp. 1).

Remark 13 The submanifold $N_1^1 \sqcup -N_2^1$ is nullcobordant. Similar to (26) we set

$$\tilde{\Psi}: [0, 1] \times N_1^1 \rightarrow M, \quad \tilde{\Psi}(t, [0, [x_1 - ix_1, x_2 - ix_2, iy_3]]) = [t, [x_1 - ix_1, x_2 - ix_2, \tilde{f}(t) + iy_3]],$$

where $\tilde{f}: [0, 1] \rightarrow [0, \frac{1}{4}]$ is a nondecreasing bump function with $\tilde{f}|_{[0, 3/4]} = 0$ and $\tilde{f}|_{[7/8, 1]} = \frac{1}{4}$. Then $\tilde{C} = \text{Im}(\tilde{\Psi})$ oriented by $(-\partial_t - \tilde{f}'(t)\partial_{x_3}, \partial_{x_1} - \partial_{y_1}, \partial_{x_2} - \partial_{y_2}, \partial_{y_3})$ satisfies $\partial\tilde{C} = N_1^1 \sqcup -N_2^1$. The linking number on M of the submanifolds $N_3^1 \sqcup -N_7^1$ and $N_1^1 \sqcup -N_2^1$ is then well-defined and coincides with the oriented intersection number of C and $N_1^1 \sqcup -N_2^1$. As a consequence of Remark 12, this equals -2 . In fact, set $\tilde{C}_t = \tilde{\Psi}(\{t\} \times N_1^1)$ oriented by $(\partial_{x_1} - \partial_{y_1}, \partial_{x_2} - \partial_{y_2}, \partial_{y_3})$. Then the oriented intersection number of C and $\tilde{C}_t$ is -1 if $t \in [0, \frac{1}{2})$, undefined if $t = \frac{1}{2}$, and 1 if $t \in (\frac{1}{2}, 1]$.

The projection to M' of the cobordism C obtained in the proof of Proposition 11 determines a cobordism between N'_3 and N'_7 that intersects N'_1 negatively and N'_2 positively. The average of the form α provides a primitive for the difference of the Thom forms of N'_3 and N'_7 . We make this precise in the following result:

Lemma 14 *There are small tubular neighborhoods V'_3 and V'_7 of N'_3 and N'_7 on M' and κ -invariant forms $\alpha \in \Omega^3(M')$ and $\nu_3, \nu_7 \in \Omega^4(M')$ such that:*

- (i) The forms v_3 and v_7 are closed and supported on V'_3 and V'_7 , respectively. The cohomology class $[v_j] \in H_c^4(V'_j)$ is the Thom class of N'_j .
- (ii) The pushforwards of $2v_3$ and $2v_7$ to X represent the Thom classes of N_3 and N_7 on X , respectively.
- (iii) The form α vanishes in a neighborhood of the level set $t = \frac{1}{2}$ and $d\alpha = v_3 - v_7$. In particular, α is closed on $M' - (V'_3 \cup V'_7)$.
- (iv) There are tubular neighborhoods V'_1 and V'_2 of N'_1 and N'_2 such that $\alpha|_{V'_1} = -\pi_1^*(\varphi|_{N'_1})$, and $\alpha|_{V'_2} = \pi_2^*(\varphi|_{N'_2})$, where $\pi_j: V'_j \rightarrow N'_j$ denotes the nearest-point projection for $j = 1, 2$.

Proof Let $K = \langle \iota_1, \iota_2, \kappa \rangle = \mathbb{Z}_2^3$. For $j = 3, 7$ we write $V_j^2 = \iota_2(V_j^1) \subset M$, and set $K_3^1 = \{\text{Id}, \iota_1, \kappa, \kappa \circ \iota_1\}$ and $K_7^1 = \{\text{Id}, \iota_1, \kappa \circ \iota_2, \kappa \circ \iota_1 \circ \iota_2\}$. We first observe that every element of K_j^1 maps N_j^1 onto itself, and since the elements in K_j^1 are isometries, they preserve V_j^1 (and its fiber bundle structure). In addition, the elements of $K_j^2 = K - K_j^1$ swap N_j^1 with N_j^2 , and therefore V_j^1 with V_j^2 . Of course, if $\gamma \in K$ maps N_j^1 to N_j^k with $k = 1$ or $k = 2$, then the restrictions $\gamma: N_j^1 \rightarrow N_j^k$ and $\gamma_*: \nu(N_j^1) \rightarrow \nu(N_j^k)$ preserve the fixed orientations because the submanifolds N_j^k are oriented by $\varphi|_{N_j^k}$ and all the elements in K fix φ and preserve the orientation of M . Hence, if $j = 3, 7$, the form $v'_j = \frac{1}{4} \sum_{\gamma \in K_j^1} \gamma^* v_j^1$ is another representative of the Thom class of N_j^1 in M , and $v''_j = \frac{1}{4} \sum_{\gamma \in K_j^2} \gamma^* v_j^1$ represents the Thom class of N_j^2 .

The pushforward v_j of $v'_j + v''_j = \frac{1}{4} \sum_{\gamma \in K} \gamma^* v_j^1$ to M' is a κ -invariant form that represents the Thom class of $N'_j \subset M'$. It is supported in the projection of $V_j^1 \cup V_j^2$ to M' ; we denote it by V'_j . We observe that (the pushforward of) $2v_j$ induces a representative of the Thom class of N_j on X . The reason is that the fiber of the normal bundle at a point $p \in N'_j$ is $F_p \cong \mathbb{C}^2$, and on X this is $F_p/\mathbb{Z}_2$; hence, $\int_{F_p/\mathbb{Z}_2} 2v_j = \frac{1}{2} \int_{F_p} 2v_j = 1$. Consider the K -invariant form $\alpha' = \frac{1}{4} \sum_{\gamma \in K} \gamma^* \beta$. On M we have

$$d\alpha' = \frac{1}{4} \sum_{\gamma \in K} \gamma^* v_3^1 - \frac{1}{4} \sum_{\gamma \in K} \gamma^* v_7^1 = (v'_3 + v''_3) - (v'_7 + v''_7).$$

Hence the pushforward of α' to M' (which we keep denoting by α') satisfies $d\alpha' = v_3 - v_7$. Since β vanishes in a neighborhood of the level set $t = \frac{1}{2}$ and the elements of K preserve that level set, α' vanishes in a neighborhood of $t = \frac{1}{2}$.

Similarly to Lemma 5, we now modify α' around the level set $t = 0$ to obtain α satisfying all these conditions. We again set $\sigma_1 = -1$ and $\sigma_2 = 1$. For $j = 1, 2$, if $\gamma \in K$ then γ maps N_j^1 onto itself, or it swaps N_j^1 and N_j^2 . Since γ preserves the orientations in both cases, $\int_{N_j^k} \gamma^* \beta = \int_{\gamma(N_j^k)} \beta = \sigma_j = \frac{1}{2} \sigma_j \int_{N_j^k} \varphi$. Hence $[(\gamma^* \beta)|_{N_j^k}] = \frac{1}{2} \sigma_j [\varphi|_{N_j^k}]$ and therefore $[\alpha'|_{N_j^k}] = \sigma_j [\varphi|_{N_j^k}]$. Thus, $[\alpha'|_{N'_j}] = \sigma_j [\varphi|_{N'_j}]$ on M' .

By the Poincaré lemma, given a tubular neighborhood W'_j of N'_j , there is $\theta_j \in \Omega^2(W'_j)$ such that $d\theta_j = \sigma_j \pi_j^*(\varphi|_{N'_j}) - \alpha'|_{W'_j}$. Indeed, we can assume that θ_j is κ -invariant as both $\pi_j^*(\varphi|_{N'_j})$ and α' are. Consider V'_1 and V'_2 smaller tubular neighborhoods with $\bar{V}'_j \subset W'_j$, and κ -invariant bump functions h_j that equal 1 on V'_j and 0 on $M' - W'_j$. The form

$$\alpha = \alpha' + d(h_1 \theta_1) + d(h_2 \theta_2).$$

satisfies all the stated conditions. □

Remark 15 Similar to Remark 13, the linking number of $N'_3 \sqcup -N'_7$ and $N'_1 \sqcup -N'_2$ on M' is -2 .

The equality (19), and Lemma 14 enable us to find a primitive of $\tau_7^2 - \tau_3^2$ in Lemma 16. To state it, we introduce some notation. For $j = 1, 2, 3, 7$, we define V_j as the projection of $V'_j \subset M'$ onto X and we let $U_j = \rho^{-1}(V_j)$. Recall that at the beginning of Section 3.1 we fixed tubular neighborhoods $O_j \subset O_j^2$ of $N_j \subset X$ and we assumed that τ_j is supported on $\rho^{-1}(O_j)$. Observe that it is not restrictive to assume that $O_j^2 \subset V_j$; we do so.

Lemma 16 *There is a form $\tilde{\alpha} \in \Omega^3(\tilde{M})$ such that $d(\tilde{\alpha}) = \tau_7^2 - \tau_3^2$. In particular, it is closed on $U_1 \cup U_2$. In addition,*

$$\tilde{\alpha} \wedge \tau_1 = -4 \operatorname{pr}_1^*(\varphi|_{N_1}) \wedge \tau_1, \quad \tilde{\alpha} \wedge \tau_2 = 4 \operatorname{pr}_2^*(\varphi|_{N_2}) \wedge \tau_2,$$

where $\operatorname{pr}_j = \pi_j \circ \rho: \rho^{-1}(V'_j) \rightarrow N_j$.

Proof Let $j = 3, 7$. We first observe that $\rho^*(v_j)$ is smooth because it vanishes on a small neighborhood of N_j as α vanishes on a small neighborhood of $t = \frac{1}{2}$ and $d\alpha = v_3 - v_7$. Since τ_j^2 vanishes on a neighborhood of E_j , the pushforward $\rho_*(\tau_j^2)$ represents an element in $H_c^4(V_j)$. This group is generated by the Thom class of N_j , which is represented by $2v_j$ according to Lemma 14. In fact, (19) shows $[\rho_*(\tau_j^2)] = -2[2v_j] \in H_c^4(V_j)$. Hence there are forms $\eta_3, \eta_7 \in \Omega(X)$ with compact support on V_j such that $d\eta_j = \rho_*(\tau_j^2) + 4v_j$.

Lemma 5 allows us to modify η_j around N_j and find a form η'_j supported on V_j whose pullback $\rho^*(\eta'_j)$ is smooth on $\tilde{M}$ and $d\eta'_j = \rho_*(\tau_j^2) + 4v_j$. This is possible because on a small tubular neighborhood W_j^2 of N_j both $\rho_*(\tau_j^2)$ and v_j vanish, so $d\eta_j|_{W_j^2} = (\rho_*(\tau_j^2) + 4v_j)|_{W_j^2} = 0$. We now argue that $\rho^*(\alpha)$ is smooth. Since $\alpha = 0$ on a neighborhood of the level set $t = \frac{1}{2}$, it is smooth on $\tilde{M} - (E_1 \cup E_2)$. In addition, if $j = 1, 2$, then $\alpha|_{V'_j} = \sigma_j \pi_j^*(\varphi|_{N'_j})$ on M' , where $\sigma_1 = -1$ and $\sigma_2 = 1$. The same identity holds for their pushforwards to X , so that $\rho^*(\alpha)|_{U_j} = \rho^*(\alpha|_{V'_j}) = \sigma_j \operatorname{pr}_j^*(\varphi|_{N_j})$. Hence it is smooth.

Finally, we define $\tilde{\alpha} = \rho^*(\eta'_7 + 4\alpha - \eta'_3)$. Then $d\tilde{\alpha} = \tau_7^2 + 4\rho^*v_7 + 4\rho^*(v_3 - v_7) - (\tau_3^2 + 4\rho^*(v_3)) = \tau_7^2 - \tau_3^2$. Given $j = 1, 2$, the support of τ_j is contained in $\rho^{-1}(O_j) \subset U_j$ and both $\rho^*(\eta'_3)$ and $\rho^*(\eta'_7)$ vanish there. Hence $\tilde{\alpha} \wedge \tau_j = 4\rho^*(\alpha) \wedge \tau_j = \sigma_j 4 \operatorname{pr}_j^*(\varphi|_{N_j}) \wedge \tau_j$. $\square$

We finish by showing that $\{[\tau_1 + \tau_2], [\tau_7 + \tau_3], [\tau_7 - \tau_3]\}$ does not vanish.

Theorem 17 *The triple Massey product $\langle [\tau_1 + \tau_2], [\tau_7 + \tau_3], [\tau_7 - \tau_3] \rangle$ is not trivial. Therefore $\tilde{M}$ is nonformal.*

Proof Note that $(\tau_1 + \tau_2) \wedge (\tau_7 + \tau_3) = 0$ and $(\tau_7 + \tau_3) \wedge (\tau_7 - \tau_3) = \tau_7^2 - \tau_3^2 = d\tilde{\alpha}$. The cohomology class determined by (23) is

$$[y] = -[\tilde{\alpha} \wedge (\tau_1 + \tau_2)],$$

and the triple Massey product vanishes if and only if $[y] \in \mathcal{J}$, where $\mathcal{J}$ is the ideal generated by $[\tau_1 + \tau_2]$ and $[\tau_7 - \tau_3]$. Observe that for any closed 5-form x , if $\int_{\tilde{M}} x \wedge (\tau_1 - \tau_2) \neq 0$ then $[x] \notin \mathcal{J}$ because if $[x] = [c_1(\tau_1 + \tau_2) + c_2(\tau_7 - \tau_3)] \in \mathcal{J}$ with $[c_1], [c_2] \in H^3(\tilde{M})$, then $[x \wedge (\tau_1 - \tau_2)] = [c_1](\tau_1^2 - \tau_2^2) = 0$ by Lemma 10. So we focus on proving $\int_{\tilde{M}} y \wedge (\tau_1 - \tau_2) \neq 0$.

By Lemma 16, $[y] = 4[\text{pr}_1^*(\varphi|_{N_1}) \wedge \tau_1] - 4[\text{pr}_2^*(\varphi|_{N_2}) \wedge \tau_2]$. Therefore

$$[y \wedge (\tau_1 - \tau_2)] = 4[\text{pr}_1^*(\varphi|_{N_1}) \wedge \tau_1^2] + 4[\text{pr}_2^*(\varphi|_{N_2}) \wedge \tau_2^2] = 4[\rho^*(\beta)] \wedge ([\tau_1^2] + [\tau_2^2]),$$

where $\beta \in \Omega^3(X)$ is a closed form such that $[\beta] = [\varphi]$ and $\rho^*(\beta)|_{\rho^{-1}(O_j)} = \text{pr}_j^*(\varphi|_{N_j})$ for every $1 \leq j \leq 10$, as in Lemma 5. From Lemma 10 we deduce

$$\int_{\tilde{M}} y \wedge (\tau_1 - \tau_2) = -16 \int_{L_1} \rho^*(\beta) = -16 \int_{L_1} \text{pr}_1^*(\varphi|_{N_1}) = -16.$$

So $[y] \notin \mathcal{J}$. □

Remark 18 Following the notation introduced in the proof of Theorem 17, we claim that the value of $\int_{\tilde{M}} y \wedge (\tau_1 - \tau_2)$ is related to the linking number of $N'_3 \sqcup -N'_7$ and $N'_1 \sqcup -N'_2$ on M' .

In the proof of Lemma 16 we obtained $\tilde{\alpha} = \rho^*(\eta'_7 + 4\alpha - \eta'_3)$, where α is defined in Lemma 14 and η'_j is supported around N'_j for $j = 3, 7$. Since $y \wedge (\tau_1 - \tau_2) = -\tilde{\alpha} \wedge (\tau_1^2 - \tau_2^2)$ we obtain, $\int_{\tilde{M}} y \wedge (\tau_1 - \tau_2) = -4 \int_X \alpha \wedge (\rho_*(\tau_1^2) - \rho_*(\tau_2^2)) = -2 \int_{M'} \alpha \wedge (\rho_*(\tau_1^2) - \rho_*(\tau_2^2))$.

Equation (19) implies that the pullback of $\rho_*(\tau_j^2)$ to M' (which we denote by the same name) equals $-4\text{Th}[N'_j] \in H_c^4(V'_j)$ because the fiber of the normal bundle of N'_j is $\mathbb{C}^2$ and that of N_j is $\mathbb{C}^2/\mathbb{Z}_2$. A representative of the Thom class of $N'_1 \sqcup -N'_2$ is then the pullback of $-\frac{1}{4}(\rho_*(\tau_1^2) - \rho_*(\tau_2^2))$ to M' . Lemma 14 ensures that $d\alpha = v_3 - v_7$, where $[v_j] = \text{Th}[N'_j]$ for $j = 3, 7$. Equation (24) shows that the linking number of $N'_3 \sqcup -N'_7$ and $N'_1 \sqcup -N'_2$ on M' is

$$-\frac{1}{4} \int_{M'} \alpha \wedge (\rho_*(\tau_1^2) - \rho_*(\tau_2^2)) = \frac{1}{8} \int_{\tilde{M}} y \wedge (\tau_1 - \tau_2).$$

Note that in the proof of Theorem 17 we obtained $\int_{\tilde{M}} y \wedge (\tau_1 - \tau_2) = -16$, so the linking number is $-\frac{16}{8} = -2$. This is consistent with Remark 15.

Acknowledgements

I would like to thank Spiro Karigiannis for carefully proofreading this paper and for providing suggestions that improved its exposition. I am also grateful to Dominic Joyce and Johannes Nordström for useful remarks, and to Xuemiao Chen, Aleksandar Milivojević, Vicente Muñoz, and Thomas Walpuski for their interest and opinions. I also acknowledge the referee for their valuable comments and suggestions.

The research and preparation of this manuscript, up to its submission, were carried out while the author was at the University of Waterloo. The revisions were completed at Humboldt University of Berlin, where the author's position was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy — the Berlin Mathematics Research Center MATH+ (EXC-2046/1, EXC-2046/2, project ID 390685689).

References

- [1] **M Amann, V Kapovitch**, *On fibrations with formal elliptic fibers*, Adv. Math. 231:3-4 (2012) 2048–2068 MR
- [2] **M Amann, I A Taimanov**, *On the formality of nearly Kähler manifolds and of Joyce's examples in G_2 -holonomy*, Mosc. Math. J. 24:4 (2024) 495–512 MR

- [3] **R Bott, L W Tu**, *Differential forms in algebraic topology*, Graduate Texts in Mathematics 82, Springer (1982) MR
- [4] **G E Bredon**, *Introduction to compact transformation groups*, Pure and Applied Mathematics 46, Academic, New York (1972) MR
- [5] **G E Bredon**, *Topology and geometry*, Graduate Texts in Mathematics 139, Springer (1993) MR
- [6] **G R Cavalcanti**, *Formality of k -connected spaces in $4k + 3$ and $4k + 4$ dimensions*, Math. Proc. Cambridge Philos. Soc. 141:1 (2006) 101–112 MR
- [7] **K F Chan, S Karigiannis, C C Tsang**, *The $\mathcal{L}_B$ -cohomology on compact torsion-free G_2 manifolds and an application to ‘almost’ formality*, Ann. Global Anal. Geom. 55:2 (2019) 325–369 MR
- [8] **A Corti, M Haskins, J Nordström, T Pacini**, *G_2 -manifolds and associative submanifolds via semi-Fano 3-folds*, Duke Math. J. 164:10 (2015) 1971–2092 MR
- [9] **D Crowley, J Nordström**, *The rational homotopy type of $(n-1)$ -connected manifolds of dimension up to $5n - 3$* , J. Topol. 13:2 (2020) 539–575 MR
- [10] **P Deligne, P Griffiths, J Morgan, D Sullivan**, *Real homotopy theory of Kähler manifolds*, Invent. Math. 29:3 (1975) 245–274 MR
- [11] **Y Félix, J Oprea, D Tanré**, *Algebraic models in geometry*, Oxford Graduate Texts in Mathematics 17, Oxford Univ. Press (2008) MR
- [12] **M Fernández, A Fino, A Kovalev, V Muñoz**, *A compact G_2 -calibrated manifold with first Betti number $b_1 = 1$* , Adv. Math. 381 (2021) art.id. 107623 MR
- [13] **DD Joyce**, *Compact 8-manifolds with holonomy $\text{Spin}(7)$* , Invent. Math. 123:3 (1996) 507–552 MR
- [14] **DD Joyce**, *Compact Riemannian 7-manifolds with holonomy G_2 , I*, J. Differential Geom. 43:2 (1996) 291–328 MR
- [15] **DD Joyce**, *Compact Riemannian 7-manifolds with holonomy G_2 , II*, J. Differential Geom. 43:2 (1996) 329–375 MR
- [16] **D Joyce**, *A new construction of compact 8-manifolds with holonomy $\text{Spin}(7)$* , J. Differential Geom. 53:1 (1999) 89–130 MR
- [17] **DD Joyce**, *Compact manifolds with special holonomy*, Oxford Univ. Press (2000) MR
- [18] **D Joyce, S Karigiannis**, *A new construction of compact torsion-free G_2 -manifolds by gluing families of Eguchi–Hanson spaces*, J. Differential Geom. 117:2 (2021) 255–343 MR
- [19] **A Kovalev**, *Twisted connected sums and special Riemannian holonomy*, J. Reine Angew. Math. 565 (2003) 125–160 MR
- [20] **A Kovalev, N-H Lee**, *K3 surfaces with non-symplectic involution and compact irreducible G_2 -manifolds*, Math. Proc. Cambridge Philos. Soc. 151:2 (2011) 193–218 MR
- [21] **L Martín-Merchán**, *A compact non-formal closed G_2 manifold with $b_1 = 1$* , Math. Nachr. 295:8 (2022) 1562–1590 MR
- [22] **J Nordström**, *Extra-twisted connected sum G_2 -manifolds*, Ann. Global Anal. Geom. 64:1 (2023) art.id. 2 MR
- [23] **F Reidegeld**, *A construction of G_2 -manifolds from K3 surfaces with a $\mathbb{Z}_2^2$ -action*, Differential Geom. Appl. 88 (2023) art.id. 101998 MR
- [24] **I Satake**, *On a generalization of the notion of manifold*, Proc. Nat. Acad. Sci. U.S.A. 42 (1956) 359–363 MR
- [25] **IA Taimanov**, *The generalized Kummer construction and cohomology rings of G_2 -manifolds*, Mat. Sb. 209:12 (2018) 139–148 MR In Russian; translated in Sb. Math. 209 (2018) 1803–1811
- [26] **A Tralle, J Oprea**, *Symplectic manifolds with no Kähler structure*, Lecture Notes in Mathematics 1661, Springer (1997) MR
- [27] **M Verbitsky**, *Manifolds with parallel differential forms and Kähler identities for G_2 -manifolds*, J. Geom. Phys. 61:6 (2011) 1001–1016 MR

LUCÍA MARTÍN-MERCHÁN lucia.martin.merchan@hu-berlin.de
 Institut für Mathematik, Humboldt Universität zu Berlin, Berlin, Germany

Proposed: Simon Donaldson
 Seconded: Nathalie Wahl, Aaron Naber

Received: November 4, 2024
 Revised: April 6, 2025

K3 surfaces associated with varieties of generalized Kummer type

SALVATORE FLOCCARI

Any hyper-Kähler variety K of generalized Kummer type is associated via Hodge theory with a K3 surface S_K . We show how they are related geometrically through a moduli space of sheaves on S_K . As a consequence, building fundamentally on the works of O’Grady, Markman, Voisin and Varesco, we establish the Hodge conjecture for all powers of any of these K3 surfaces as well as for all abelian fourfolds of Weil type with discriminant 1 and their powers, strengthening a result of Markman.

1 Introduction

A manifold of generalized Kummer type is, by definition, a compact Kähler manifold which is deformation equivalent to the generalized Kummer variety $K^n(A)$ on an abelian surface A constructed by Beauville [1983]. Manifolds of generalized Kummer type give examples of hyper-Kähler manifolds [Beauville 1983; Huybrechts 1999] in each even dimension $2n$, commonly referred to as manifolds of Kum ^{n} -type; by a variety of Kum ^{n} -type we shall indicate one such manifold which is projective. For any $n > 1$, these projective varieties form countably many 4-dimensional families. Besides Beauville’s construction of generalized Kummer varieties, other examples can be obtained by considering moduli of stable sheaves on abelian surfaces [Yoshioka 2001]; however, it is known that the general projective variety of generalized Kummer type cannot be realized from such a construction.

In this article we consider some K3 surfaces naturally attached via Hodge theory to varieties of generalized Kummer type. Let K be a variety of Kum ^{n} -type. By the Torelli theorem there exists a unique K3 surface S_K such that there exists a Hodge isometry $H_{\text{tr}}^2(S_K, \mathbb{Z}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Z})(2)$ of transcendental lattices, where on the right-hand side the Beauville–Bogomolov form is multiplied by a factor 2. If $K = K^n(A)$ is the generalized Kummer variety on an abelian surface A , then S_K is the Kummer K3 surface $\text{Km}(A)$, which motivates our definition.

We will explain how the K3 surface S_K is related geometrically to K , for any variety K of Kum ^{n} -type. Indeed, there exists a natural subvariety W_K of K which is a hyper-Kähler variety of K3^[m]-type where either $2m = n$ or $2m = n + 1$, with the key property of deforming everywhere together with K ; i.e., W_K is a trianalytic submanifold of K in the sense of Verbitsky [1995]. It is obtained as the component of maximal dimension of the fixed locus of an involution of K which acts trivially on the second cohomology group; such automorphisms were studied already in [Hassett and Tschinkel 2013; Boissière et al. 2011; Oguiso 2020; Kamenova et al. 2022; Floccari 2024] and used to produce trianalytic cycles in [Kaledin and Verbitsky 1998]. Our contribution lies in the determination of the period of W_K :

MSC2020: 14C15, 14C25, 14C30, 14J28, 14J42.

© 2026 The Author, under license to MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

we show that the pullback along $\iota : W_K \hookrightarrow K$ induces a primitive embedding of lattices and Hodge structures $\iota^* : H_{\text{tr}}^2(K, \mathbb{Z})(2) \xrightarrow{\sim} H_{\text{tr}}^2(W_K, \mathbb{Z})$, where on the left-hand side the form is multiplied by 2. It follows that $W_K \subset K$ is a smooth and projective moduli space of stable sheaves on the K3 surface S_K .

As a consequence of the above construction, we obtain the following result.

Theorem 1.1 *There exists an algebraic cycle on $S_K \times K$ inducing a Hodge isometry*

$$H_{\text{tr}}^2(S_K, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Q})(2).$$

This algebraic cycle gives in fact an isomorphism between the transcendental parts of the homological motives of S_K and K in degree 2; see Theorem 5.9.

Building on results of O’Grady [2021], Markman [2023], Voisin [2022] and Varesco [2025], we use Theorem 1.1 to prove some cases of the Hodge conjecture.

Theorem 1.2 *Let K be any projective variety of Kumⁿ-type, and let S_K be the K3 surface with transcendental lattice $H_{\text{tr}}^2(S_K, \mathbb{Z}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Z})(2)$. Then:*

- (i) *The Kuga–Satake correspondence is algebraic for S_K .*
- (ii) *The Hodge conjecture holds for any power of S_K .*

The theorem gives countably many new families of K3 surfaces of general Picard rank 16 for which it is possible to prove (i) and (ii), generalizing previous results from [Paranjape 1988; Ingalls et al. 2022; Floccari 2024]. Using that isogenies of K3 surfaces are algebraic by [Buskin 2019; Huybrechts 2019], we show that the conclusions of Theorem 1.2 hold for any K3 surface S such that there exists an isometric embedding of rational quadratic spaces $H_{\text{tr}}^2(S, \mathbb{Q}) \hookrightarrow (U_{\mathbb{Q}}^{\oplus 3} \oplus \langle -m \rangle_{\mathbb{Q}})$ for some positive integer m , where U denotes a hyperbolic plane; see Theorem 5.11.

Let us summarize the main points of the proof. Recall that the Kuga–Satake construction [1967] (see also [Deligne 1972]) produces an abelian variety $\text{KS}(X)$ from a polarized hyper-Kähler variety X , with an embedding of Hodge structures $H^2(X, \mathbb{Q})_{\text{prim}} \hookrightarrow H^2(\text{KS}(S)^2, \mathbb{Q})$ called the Kuga–Satake correspondence; according to the Hodge conjecture, this embedding should be induced by an algebraic cycle.

It was shown by O’Grady [2021] and Markman [2023] that the Kuga–Satake variety of a Kumⁿ-variety is isogenous to a power of its third intermediate Jacobian $J^3(K)$, which is an abelian fourfold of Weil type with discriminant 1. Moreover, varying K in polarized families, one obtains all complete up to isogeny families of such abelian fourfolds; see also [van Geemen 2023]. Abelian fourfolds of Weil type are characterized by the presence of exceptional Hodge classes: there is a 2-dimensional space of Hodge classes in their middle cohomology which, for the very general such fourfold, cannot be represented as intersection of divisors; see [van Geemen 2000; Moonen and Zarhin 1995]. Markman [2023] has proven that these Hodge–Weil classes are algebraic for any fourfold of Weil type of discriminant 1; his results further imply that the Kuga–Satake correspondence is algebraic for any variety of Kumⁿ-type, as shown by Voisin [2022].

Theorem 1.2 is then obtained as follows. By definition, the Kuga–Satake variety of S_K is isogenous to a power of the Kuga–Satake variety of K , and, hence, to a power of a Weil fourfold with discriminant 1. We deduce Theorem 1.2(i) from Theorem 1.1 and the algebraicity of the Kuga–Satake correspondence for the Kum^n -variety K . Statement (ii) is a consequence of (i) thanks to the work of Varesco [2025].

The Hodge conjecture for a very general fourfold A of Weil type with discriminant 1 is proven by Markman [2023], which implies the Hodge conjecture also for powers of such A by [Varesco 2025; Milne 2021]. However, the results of [Markman 2023] do not directly imply the Hodge conjecture for an arbitrary Weil fourfold with discriminant 1. We complete the proof of the Hodge conjecture for all Weil fourfolds with discriminant 1, and their powers.

Theorem 1.3 *Let A be an abelian fourfold of Weil type with discriminant 1. Then the Hodge conjecture holds for A and all of its powers.*

To prove this statement, we show in Theorem 3.5 that if S is any K3 surface such that the Kuga–Satake correspondence is algebraic and the Hodge conjecture holds for any power of S , then the Hodge conjecture holds for any power of the Kuga–Satake variety of S as well. We then obtain Theorem 1.3 from Theorem 1.2 using that any Weil fourfold with discriminant 1 appears as isogeny factor of the Kuga–Satake variety of some K3 surface S_K as in Theorem 1.2.

2 Motives

2.1 A neutral Tannakian category $\mathcal{C}$ is a $\mathbb{Q}$ -linear abelian rigid tensor category in which $\text{End}(1) = \mathbb{Q}$ and which admits an exact faithful tensor functor $\omega : \mathcal{C} \rightarrow \text{Vect}_{\mathbb{Q}}$. We refer the reader to [Deligne and Milne 1982] for more details. The group of tensor automorphisms of ω is a (pro-)algebraic group $\mathcal{G}(\mathcal{C})$ with the property that $\mathcal{C}$ is $\otimes$ -equivalent to the category of finite-dimensional representations of $\mathcal{G}(\mathcal{C})$. Thus, neutral Tannakian subcategories of $\mathcal{C}$ correspond bijectively to quotients of $\mathcal{G}(\mathcal{C})$. Given $M \in \mathcal{C}$, we denote by $\langle M \rangle_{\mathcal{C}}$ the neutral Tannakian subcategory of $\mathcal{C}$ generated by M ; more directly, it is defined as the smallest thick and full subcategory containing M and closed under direct sums, duals, tensor products and subobjects. Then $\langle M \rangle_{\mathcal{C}}$ is equivalent to the category of finite-dimensional representations of a linear algebraic group $G_{\mathcal{C}}(M) \subset \text{GL}(\omega(M))$, called the Tannakian fundamental group of M .

2.2 The category HS of polarizable pure $\mathbb{Q}$ -Hodge structures is an abelian and semisimple rigid tensor category, which is neutral Tannakian via the forgetful functor to $\mathbb{Q}$ -vector spaces. The Tannakian fundamental group of a polarizable Hodge structure V is a reductive and connected algebraic group over $\mathbb{Q}$ ([Moonen 1999]) denoted by $\text{MT}(V) \subset \text{GL}(V)$ and called the Mumford–Tate group of V . It is equivalently defined as the subgroup of $\text{GL}(V)$ which fixes all Hodge classes in any tensor construction

$$V^{\otimes i} \otimes V^{\vee, \otimes j} \otimes \mathbb{Q}(k),$$

where $\mathbb{Q}(k)$ is the Tate Hodge structure of weight $-2k$. Provided that V is not entirely of weight 0, there is a central torus $\mathbb{G}_{m, \mathbb{Q}} \subset \text{MT}(V)$ whose action induces the weight decomposition of the objects in $\langle V \rangle_{\text{HS}}$.

Remark 2.1 Let $V \in \text{HS}$ be a Hodge structure of odd weight. Denote by $\langle V \rangle_{\text{HS}}^{\text{ev}}$ the full subcategory of $\langle V \rangle_{\text{HS}}$ of Hodge structures of even weight, which is a neutral Tannakian subcategory of HS . If $w : \mathbb{G}_{m, \mathbb{Q}} \rightarrow \text{MT}(V)$ denotes the inclusion of the weight torus, the objects of $\langle V \rangle_{\text{HS}}^{\text{ev}}$ are precisely the Hodge structures in $\langle V \rangle_{\text{HS}}$ on which $w(-1)$ acts trivially. Hence, we have $\text{MT}(\langle V \rangle_{\text{HS}}^{\text{ev}}) = \text{MT}(V)/w(-1)$. It follows that if W is a Hodge structure of even weight in $\langle V \rangle_{\text{HS}}$ such that the induced surjective morphism $\phi : \text{MT}(V) \rightarrow \text{MT}(W)$ of Mumford–Tate groups is an isogeny of degree 2, we then have the equality $\langle V \rangle_{\text{HS}}^{\text{ev}} = \langle W \rangle_{\text{HS}}$ of subcategories of HS . Indeed, since W has even weight, it belongs to $\langle V \rangle_{\text{HS}}^{\text{ev}}$, and, therefore, ϕ factors as the composition of the projection $\psi : \text{MT}(V) \rightarrow \text{MT}(\langle V \rangle_{\text{HS}}^{\text{ev}})$ followed by a surjective morphism $\phi' : \text{MT}(\langle V \rangle_{\text{HS}}^{\text{ev}}) \rightarrow \text{MT}(W)$ of algebraic groups. But both ϕ and ψ are isogenies of degree 2, and, hence, ϕ' is an isomorphism; it follows that $\langle W \rangle_{\text{HS}} = \langle V \rangle_{\text{HS}}^{\text{ev}}$.

2.3 An algebraic class on a smooth and projective complex variety X is a cohomology class α in $H^{2k}(X, \mathbb{Q})$ which is $\mathbb{Q}$ -linear combination of fundamental classes of subvarieties of X . Any algebraic class is a Hodge class, and the Hodge conjecture predicts that, conversely, all Hodge classes are algebraic. A correspondence from X to Y of degree k is by definition an algebraic class in $H^{2k+2\dim X}(X \times Y, \mathbb{Q})$; correspondences may be composed [Fulton 1998]. Grothendieck’s theory of motives is very useful in the study of algebraic cycles; in brief, the main insight is to use correspondences as morphisms. We denote by Mot the category of homological motives over $\mathbb{C}$ with rational coefficients; see [André 2004]. Its objects are triples (X, p, n) where $X \in \text{SmProj}_{\mathbb{C}}$, $p \in H^{2\dim X}(X \times X, \mathbb{Q})$ is an idempotent correspondence, and n is an integer. Morphisms from (X, p, n) to (Y, q, m) are defined to be the correspondences $f \in H^{2m-2n+2\dim X}(X \times Y, \mathbb{Q})$ of degree $m - n$ such that $f \circ p = q \circ f$. There is a realization functor $R : \text{Mot} \rightarrow \text{HS}$ which sends (X, p, n) to $p_*(H^\bullet(X, \mathbb{Q})(n))$. Moreover, there is a contravariant functor $h : \text{SmProj}_{\mathbb{C}} \rightarrow \text{Mot}$ which attaches to the smooth and projective variety X its motive $h(X) := (X, \Delta_X, 0)$, where $\Delta_X \in H^{2\dim X}(X \times X, \mathbb{Q})$ is the class of the diagonal. The composition $R \circ h$ is the functor $H^\bullet : \text{SmProj}_{\mathbb{C}} \rightarrow \text{HS}$ which attaches to the smooth and projective complex variety X its rational Hodge structure $H^\bullet(X, \mathbb{Q})$.

2.4 The category Mot of homological motives is a pseudoabelian tensor category via the product of varieties. However, it is not known to be abelian nor Tannakian: these properties are conditional to the validity of Grothendieck’s standard conjectures [1969]. By the work of Jannsen [1992] and André [1996b], we have the following. Given $m \in \text{Mot}$, we let $\langle m \rangle_{\text{Mot}}$ be the smallest thick and full subcategory of Mot containing m and closed under direct sums, duals, tensor products and subquotients.

Theorem 2.2 [Arapura 2006, Theorem 4.1] *Let X be a smooth and projective variety. The standard conjectures hold for X if and only if $\langle h(X) \rangle_{\text{Mot}}$ is a semisimple abelian neutral Tannakian category.*

Thus, provided that the standard conjectures hold for X , we can attach to it a reductive algebraic group $G_{\text{Mot}}(h(X))$ over $\mathbb{Q}$, whose representation category is equivalent to $\langle h(X) \rangle_{\text{Mot}}$. The algebraic group $G_{\text{Mot}}(h(X))$ is characterized as the subgroup of $\prod_j \text{GL}(H^j(X, \mathbb{Q}))$ fixing exactly the algebraic classes in $H^{2j}(X^k, \mathbb{Q}(j))$ for any j and k . There is a canonical inclusion $i_X : \text{MT}(H^\bullet(X, \mathbb{Q})) \hookrightarrow G_{\text{Mot}}(h(X))$,

which is an isomorphism if and only if the Hodge conjecture holds for X and all of its powers. Moreover, any object of $\langle h(X) \rangle_{\text{Mot}}$ admits a weight decomposition induced by the action of the weight torus $\mathbb{G}_m \subset \text{MT}(H^\bullet(X, \mathbb{Q})) \subset G_{\text{Mot}}(X)$. In particular, we have the Künneth decomposition $h(X) = \bigoplus_i h^i(X)$ in Mot .

Remark 2.3 Let X be a smooth and projective variety and assume that the Künneth component $h^2(X)$ of $h(X)$ is well-defined. Thanks to the Lefschetz theorem on divisors, this motive splits as $h^2(X) = h_{\text{tr}}^2(X) \oplus h_{\text{alg}}^2(X)$: the realization of $h_{\text{alg}}^2(X)$ is the Néron-Severi group $\text{NS}(X)_{\mathbb{Q}}$ while that of $h_{\text{tr}}^2(X)$ is the transcendental part $H_{\text{tr}}^2(X, \mathbb{Q})$ of $H^2(X, \mathbb{Q})$, the smallest sub-Hodge structure whose complexification contains $H^{2,0}(X)$.

Remark 2.4 In light of Theorem 2.2 we shall say that the standard conjectures hold for a motive $m \in \text{Mot}$ if the category $\langle m \rangle_{\text{Mot}}$ is an abelian and semisimple neutral Tannakian category, thus equivalent to the category of representations of a reductive algebraic group $G_{\text{Mot}}(m)$. When this happens, the realization functor $R : \langle m \rangle_{\text{Mot}} \rightarrow \text{HS}$ is conservative [André 2004, Corollaire 5.1.3.3], i.e., a morphism in $\langle m \rangle_{\text{Mot}}$ is an isomorphism if and only if its realization is an isomorphism of $\mathbb{Q}$ -Hodge structures.

Remark 2.5 Let $m \in \text{Mot}$ be a motive for which the standard conjectures hold. The torus $w : \mathbb{G}_{m, \mathbb{Q}} \hookrightarrow G_{\text{Mot}}(m)$ inducing the weight decomposition for the objects in $\langle m \rangle_{\text{Mot}}$ coincides with the weight torus $\mathbb{G}_{m, \mathbb{Q}} \subset \text{MT}(R(m))$ from Hodge theory. Assume that m is a motive of odd weight and let $\langle m \rangle_{\text{Mot}}^{\text{ev}}$ be the full subcategory of $\langle m \rangle_{\text{Mot}}$ of objects of even weight, which is again abelian, semisimple and neutral Tannakian. Analogously to Remark 2.1, the motives in $\langle m \rangle_{\text{Mot}}^{\text{ev}}$ correspond to the representations of $G_{\text{Mot}}(m)$ on which $w(-1)$ acts trivially. Therefore, the natural morphism of algebraic groups $G_{\text{Mot}}(m) \rightarrow G_{\text{Mot}}(\langle m \rangle_{\text{Mot}}^{\text{ev}})$ is an isogeny of degree 2. Moreover, by the same argument given in Remark 2.1, if $n \in \langle m \rangle_{\text{Mot}}$ is a motive of even weight such that the corresponding morphism $G_{\text{Mot}}(m) \rightarrow G_{\text{Mot}}(n)$ is an isogeny of degree 2, we then have the equality $\langle n \rangle_{\text{Mot}} = \langle m \rangle_{\text{Mot}}^{\text{ev}}$ of subcategories of Mot .

2.5 The standard conjectures are known to hold for curves, surfaces and abelian varieties; they hold for varieties X and Y if and only if they hold for $X \times Y$; see [Kleiman 1968]. The standard conjectures are not known to hold in general for hyper-Kähler varieties. They hold for the varieties of $\text{K3}^{[n]}$, OG10 and OG6-type built via moduli spaces of sheaves on K3 or abelian surfaces by [Bülles 2020; Floccari et al. 2021; Floccari 2023b], and for some hyper-Kähler varieties admitting a Lagrangian fibration [Ancona et al. 2025]. The standard conjectures have been fully proven for varieties of $\text{K3}^{[n]}$ -type by Charles and Markman [2013]. Recently, Foster [2024] proved that they hold in certain degrees for varieties of Kum^n -type; for later reference we state the following consequence of his results; see also [Floccari and Varesco 2025, Remark 3.2].

Theorem 2.6 [Foster 2024] *Let K be a variety of Kum^n -type. The Künneth projector $H^\bullet(K, \mathbb{Q}) \rightarrow H^2(K, \mathbb{Q})$ is algebraic. Moreover, the standard conjectures hold for the Künneth component $h^2(K) \in \text{Mot}$ of the motive of K .*

3 The Kuga–Satake construction

3.1 Let V be a $\mathbb{Q}$ -vector space equipped with a nondegenerate symmetric bilinear form q . The Clifford algebra

$$C(V) := \bigoplus_{j \geq 0} V^{\otimes j} / \langle v \cdot v - q(v, v) \rangle_{v \in V}$$

is naturally $\mathbb{Z}/2\mathbb{Z}$ -graded: $C(V) = C^+(V) \oplus C^-(V)$. The Clifford algebra acts on itself via left and right multiplication. We fix a vector $v_0 \in V$ such that $q(v_0, v_0) \neq 0$. Setting $\mu(v)(a) := v \cdot a \cdot v_0$ we obtain an embedding

$$(1) \quad \mu : V \hookrightarrow \text{End}(C^+(v)).$$

Remark 3.1 We denote by $\text{End}_{C^+}(C^+(V))$ the algebra of linear endomorphisms of $C^+(V)$ commuting with the action of $C^+(V)$ on itself by right multiplication. There is a canonical isomorphism $C^+(V) = \text{End}_{C^+}(C^+(V))$, i.e., any $f \in \text{End}_{C^+}(C^+(V))$ is induced by left multiplication by some $\alpha_f \in C^+(V)$; see [Deligne 1972, §3] or [André 1996a, §4.2].

3.2 Assume now that (V, q) is a polarized $\mathbb{Q}$ -Hodge structure of K3-type, by which we mean that V is an effective Hodge structure of weight 2 with Hodge numbers $(1, k, 1)$. Following [Kuga and Satake 1967] and [Deligne 1972], the Hodge structure on V determines an effective polarizable Hodge structure of weight 1 on $C^+(V)$. This thus defines an abelian variety $\text{KS}(V)$ up to isogeny, such that $H^1(\text{KS}(V), \mathbb{Q}) \cong C^+(V)$. With a suitable Tate twist to make V of weight 0, the embedding $\mu : V(1) \rightarrow \text{End}(H^1(\text{KS}(V), \mathbb{Q}))$ of (1) is a morphism of Hodge structures (for any choice of the vector $v_0 \in V$).

Remark 3.2 It is explained by Deligne [1972, §4] how the Hodge structure on $C^+(V)$ is induced through the action of the Clifford group $\text{CSpin}(V)$, which is a subgroup of $C^+(V)^*$ acting on $C^+(V)$ via left multiplication. This means that the Mumford–Tate group of $H^1(\text{KS}(V), \mathbb{Q})$ is contained in $\text{CSpin}(V) \subset \text{GL}(C^+(V))$. Moreover the morphism $\text{MT}(H^1(\text{KS}(V), \mathbb{Q})) \rightarrow \text{MT}(V)$ induced by the embedding μ is an isogeny of degree 2, being the restriction of the double cover $\text{CSpin}(V) \rightarrow \text{GO}(V)$, where $\text{GO}(V)$ is the group of linear automorphisms of V preserving the form up to scalar.

3.3 Let S be a polarized K3 surface or hyper-Kähler variety, and consider the polarized Hodge structure of K3-type $V = H^2(S, \mathbb{Q})_{\text{prim}}$, the orthogonal to the given polarization, equipped with the restriction q of the form. We may also take for q the restriction of the Beauville–Bogomolov form, a natural pairing on the second cohomology of a hyper-Kähler manifold [Beauville 1983]. We define the Kuga–Satake variety $\text{KS}(S)$ of S as the abelian variety obtained from (V, q) via the Kuga–Satake construction. As above, we have an embedding of Hodge structures

$$\mu : H^2(S, \mathbb{Q})_{\text{prim}}(1) \hookrightarrow \text{End}(H^1(\text{KS}(S), \mathbb{Q}));$$

identifying $H^1(\mathrm{KS}(S), \mathbb{Q})^\vee$ with $H^1(\mathrm{KS}(S), \mathbb{Q})(1)$ by means of a polarization, we get the embedding of Hodge structures

$$(2) \quad \mu' : H^2(S, \mathbb{Q})_{\mathrm{prim}} \hookrightarrow H^1(\mathrm{KS}(S), \mathbb{Q})^{\otimes 2} \subset H^2(\mathrm{KS}(S)^2, \mathbb{Q}).$$

Conjecture 3.3 *The morphism μ' is induced by an algebraic cycle on $S \times \mathrm{KS}(S)^2$.*

This is the Kuga–Satake Hodge conjecture, which is of course a special case of the Hodge conjecture. It has been proven for many K3 surfaces of Picard rank at least 17 (see [Morrison 1985]), but it is widely open otherwise. In [Paranjape 1988] and [Ingalls et al. 2022] the conjecture is proven for two 4-dimensional families of K3 surfaces of general Picard rank 16; countably many more such families of K3 surfaces for which the conjecture holds are found in [Flocari 2024]. The paper [Bolognesi and Laterveer 2024] proves Conjecture 3.3 for a 9-dimensional family of K3 surfaces related to certain cubic fourfolds. The Kuga–Satake Hodge conjecture has been proven for all hyper-Kähler varieties of generalized Kummer type by Markman [2023] and Voisin [2022].

Remark 3.4 As a variant, we may define $\mathrm{KS}'(S)$ as the Kuga–Satake variety built from the smaller Hodge structure of K3-type $H_{\mathrm{tr}}^2(S, \mathbb{Q})$; then the Kuga–Satake variety $\mathrm{KS}(S)$ built from $H^2(S, \mathbb{Q})_{\mathrm{prim}}$ is isogenous to a power of $\mathrm{KS}'(S)$; see [Huybrechts 2016, Chapter 4, Example 2.4]. Moreover, Conjecture 3.3 is equivalent to the statement that the embedding $\mu' : H_{\mathrm{tr}}^2(S, \mathbb{Q}) \hookrightarrow H^1(\mathrm{KS}'(S), \mathbb{Q})^{\otimes 2} \subset H^2(\mathrm{KS}'(S)^2, \mathbb{Q})$ analogous to (2) is induced by an algebraic cycle.

3.4 We are now ready to prove the main result of this section.

Theorem 3.5 *Let S be a projective K3 surface. Assume that the Kuga–Satake Hodge conjecture holds for S . Then the Hodge conjecture holds for all powers of S if and only if it holds for all powers of its Kuga–Satake variety $\mathrm{KS}(S)$.*

We will deduce this result from the following proposition. Recall (Section 2.5) that the standard conjectures hold for S as well as for $\mathrm{KS}(S)$.

Proposition 3.6 *Let S be a projective K3 surface. Assume that the Kuga–Satake Hodge conjecture holds for S . Then, the motive $h^2(S)$ belongs to the subcategory $\langle h^1(\mathrm{KS}(S)) \rangle_{\mathrm{Mot}}$ of the category Mot of homological motives and the induced morphism of algebraic groups $G_{\mathrm{Mot}}(h^1(\mathrm{KS}(S))) \rightarrow G_{\mathrm{Mot}}(h^2(S))$ is an isogeny of degree 2.*

Proof By Theorem 2.2, the subcategory $\langle h(\mathrm{KS}(S) \times S) \rangle_{\mathrm{Mot}}$ of Mot is an abelian and semisimple neutral Tannakian category. Hence, if Conjecture 3.3 holds for S , then the motive $h^2(S)$ is a direct summand of $h^2(\mathrm{KS}(S)^2)$, and so it belongs to $\langle h^1(\mathrm{KS}(S)) \rangle_{\mathrm{Mot}}$. We obtain the commutative diagram

$$\begin{array}{ccc} G_{\mathrm{Mot}}(h^1(\mathrm{KS}(S))) & \xrightarrow{\psi} & G_{\mathrm{Mot}}(h^2(S)) \\ i_{\mathrm{KS}} \uparrow & & \uparrow i_S \\ \mathrm{MT}(H^1(\mathrm{KS}(S), \mathbb{Q})) & \xrightarrow{\phi} & \mathrm{MT}(H^2(S, \mathbb{Q})) \end{array}$$

We know from Remark 3.2 that the morphism ϕ is an isogeny of degree 2; we will show that $\ker(\psi) = \iota_{\text{KS}}(\ker(\phi))$, which implies that ψ is also an isogeny of degree 2.

To ease notation, we let $A := \text{KS}(S)$ and $V := H^2(S, \mathbb{Q})_{\text{prim}}$; we denote by q the polarization form on V . By Remark 3.2, the action of $C^+(V)$ on $H^1(A, \mathbb{Q})$ via right multiplication commutes with $\text{MT}(H^1(A, \mathbb{Q}))$, i.e., $C^+(V)$ acts via Hodge endomorphisms. By Lefschetz (1,1)-theorem, any such endomorphism is algebraic. Therefore, $G_{\text{Mot}}(h^1(A)) \subset \text{GL}(H^1(A, \mathbb{Q}))$ commutes with the $C^+(V)$ -action given by right multiplication. By Remark 3.1, we can thus identify $G_{\text{Mot}}(h^1(A))$ with a subgroup of $C^+(V)^*$, the group of units in the Clifford algebra, acting on $C^+(V)$ via left multiplication.

Let $g \in \ker(\psi : G_{\text{Mot}}(h^1(A)) \rightarrow G_{\text{Mot}}(h^2(S)))$; then $g : H^1(A, \mathbb{Q}) \rightarrow H^1(A, \mathbb{Q})$ is left multiplication by some $\alpha_g \in C^+(V)^*$. We claim that α_g is a central element of $C^+(V)$.

By assumption, the embedding $\mu : V(1) \hookrightarrow \text{End}(H^1(A, \mathbb{Q}))$ is induced by an algebraic cycle. As $g \in \ker(\psi)$, it follows that g commutes with the image of μ . Recall (Section 3.1) that the definition of μ depends on a vector $v_0 \in V$ with $q(v_0, v_0) \neq 0$ such that $\mu(v)(\beta) := v \cdot \beta \cdot v_0$. In the Clifford algebra we have $v_0^{2k} = q(v_0, v_0)^k \cdot e$, where e is the identity element, and any $\beta \in C^+(V)$ is a linear combination of products $v_1 \cdots v_{2k}$ of an even number of vectors in V . As g commutes with $\mu(v)$ for any $v \in V$, we find

$$\begin{aligned} \alpha_g(v_1 v_2 \cdots v_{2k}) &= \frac{1}{q(v_0, v_0)^k} \cdot \alpha_g \cdot (v_1 v_2 \cdots v_{2k}) \cdot v_0^{2k} \\ &= \frac{1}{q(v_0, v_0)^k} \cdot (g \circ \mu(v_1) \circ \mu(v_2) \circ \cdots \circ \mu(v_{2k}))(e) \\ &= \frac{1}{q(v_0, v_0)^k} \cdot (\mu(v_1) \circ \mu(v_2) \circ \cdots \circ \mu(v_{2k}) \circ g)(e) = (v_1 v_2 \cdots v_{2k}) \alpha_g. \end{aligned}$$

Therefore, $\ker(\psi)$ is contained in the center of $C^+(V)$. But $V = H^2(S, \mathbb{Q})_{\text{prim}}$ has odd dimension 21, and in such case $C^+(V)$ is a central simple algebra; see [Deligne 1972, page 216]. The units in the center of $C^+(V)$ thus form a torus $\mathbb{G}_{m, \mathbb{Q}}$ acting on $C^+(V)$ via scalar multiplication; this torus induces the weight filtration on objects of $\langle H^1(A, \mathbb{Q}) \rangle_{\text{HS}}$ and so it is contained in $\text{MT}(H^1(A, \mathbb{Q}))$. Hence, $\ker(\psi)$ is contained in $i_{\text{KS}}(\ker(\phi))$. $\square$

Proof of Theorem 3.5 One direction is immediate: as Conjecture 3.3 holds for S by assumption, we have $h^2(S) \in \langle h^1(\text{KS}(S)) \rangle_{\text{Mot}}$, and the Hodge conjecture for all powers of $\text{KS}(S)$ clearly implies the Hodge conjecture for all powers of S .

Let us prove the converse implication. The Hodge conjecture for all powers of $\text{KS}(S)$ predicts that the realization functor $R : \text{Mot} \rightarrow \text{HS}$ induces a surjection

$$(3) \quad \text{Hom}_{\text{Mot}}(\mathbb{Q}(-j), h^{2j}(\text{KS}(S)^k)) \rightarrow \text{Hom}_{\text{HS}}(\mathbb{Q}(-j), H^{2j}(\text{KS}(S)^k, \mathbb{Q}))$$

for any k and j . Only motives and Hodge structures of even weight are involved in (3), and so these Hom-spaces are actually taken in the categories $\langle h^1(\text{KS}(S)) \rangle_{\text{Mot}}^{\text{ev}}$ and $\langle H^1(\text{KS}(S), \mathbb{Q}) \rangle_{\text{HS}}^{\text{ev}}$, respectively. Proposition 3.6 and Remark 2.5 yield the equality $\langle h^1(\text{KS}(S)) \rangle_{\text{Mot}}^{\text{ev}} = \langle h^2(S) \rangle_{\text{Mot}}$ of subcategories of Mot ;

in particular, $h^{2j}(\mathrm{KS}(S)^k)$ belongs to $\langle h^2(S) \rangle_{\mathrm{Mot}}$ for any j and k . Therefore, the map in (3) is the map

$$\mathrm{Hom}_{\langle h^2(S) \rangle_{\mathrm{Mot}}}(\mathbb{Q}(-j), h^{2j}(\mathrm{KS}(S)^k)) \rightarrow \mathrm{Hom}_{\langle H^2(S, \mathbb{Q}) \rangle_{\mathrm{HS}}}(\mathbb{Q}(-j), H^{2j}(\mathrm{KS}(S)^k, \mathbb{Q}))$$

given by the realization functor R . If the Hodge conjecture holds for all powers of S , then $R: \langle h^2(S) \rangle_{\mathrm{Mot}} \rightarrow \langle H^2(S, \mathbb{Q}) \rangle_{\mathrm{HS}}$ is full, and the maps in (3) are surjective. $\square$

3.5 For later use, we recall the following functoriality property of the Kuga–Satake construction, following Varesco [2023]. Let X and Y be hyper-Kähler varieties, not necessarily deformation equivalent and possibly of different dimensions. Assume that $\phi: H_{\mathrm{tr}}^2(X, \mathbb{Q}) \rightarrow H_{\mathrm{tr}}^2(Y, \mathbb{Q})$ is a rational Hodge similitude, i.e., ϕ is an isomorphism of $\mathbb{Q}$ -Hodge structures which multiplies the Beauville–Bogomolov form on the left-hand side by some nonzero $k \in \mathbb{Q}$. As in Remark 3.4, let $\mathrm{KS}'(X)$ and $\mathrm{KS}'(Y)$ be the Kuga–Satake varieties constructed from $H_{\mathrm{tr}}^2(X, \mathbb{Q})$ and $H_{\mathrm{tr}}^2(Y, \mathbb{Q})$, respectively. Varesco [2023, Proposition 3.1] then proved that there exists an isogeny $\Psi: \mathrm{KS}'(X) \rightarrow \mathrm{KS}'(Y)$ of abelian varieties such that the diagram

$$\begin{array}{ccc} H_{\mathrm{tr}}^2(X, \mathbb{Q}) & \xrightarrow{\phi} & H_{\mathrm{tr}}^2(Y, \mathbb{Q}) \\ \downarrow \mu'_X & & \downarrow \mu'_Y \\ H^1(\mathrm{KS}'(X), \mathbb{Q})^{\otimes 2} & \xrightarrow{\Psi_*} & H^1(\mathrm{KS}'(Y), \mathbb{Q})^{\otimes 2} \end{array}$$

commutes, where μ'_X and μ'_Y are the respective Kuga–Satake correspondences and the bottom arrow is the isomorphism induced by the isogeny Ψ .

Remark 3.7 With notation as above, assume that $\phi: H_{\mathrm{tr}}^2(X, \mathbb{Q}) \rightarrow H_{\mathrm{tr}}^2(Y, \mathbb{Q})$ is induced by an algebraic cycle on $X \times Y$. Then, Conjecture 3.3 for Y implies Conjecture 3.3 for X . Indeed, the above diagram yields $\mu'_X = (\Psi_*)^{-1} \circ \mu'_Y \circ \phi$; note that $(\Psi_*)^{-1}: H^1(\mathrm{KS}'(Y), \mathbb{Q})^{\otimes 2} \rightarrow H^1(\mathrm{KS}'(X), \mathbb{Q})^{\otimes 2}$ is induced by an algebraic cycle, so that the Kuga–Satake correspondence μ'_X is algebraic if μ'_Y is so.

4 Some $\mathrm{K3}^{[m]}$ -type submanifolds of generalized Kummer varieties

4.1 Let A be an abelian surface. We denote by $K^n(A)$ the generalized Kummer variety on A of dimension $2n$ introduced in [Beauville 1983]. It is a hyper-Kähler variety constructed as follows: consider the Hilbert scheme $A^{[n+1]}$ of 0-dimensional subschemes of A of length $n+1$, which is a crepant resolution $\nu: A^{[n+1]} \rightarrow A^{(n+1)}$ of the symmetric power. Denoting by $\Sigma: A^{(n+1)} \rightarrow A$ the summation map $\Sigma(a_1, \dots, a_{n+1}) = \sum_{i=1}^{n+1} a_i$, the variety $K^n(A)$ is defined as the fiber of $\Sigma \circ \nu$ over $0 \in A$. We let $A_0^{(n+1)} \subset A^{(n+1)}$ be the fiber of Σ over 0 ; the Hilbert–Chow morphism ν restricts to a crepant resolution $\nu: K^n(A) \rightarrow A_0^{(n+1)}$. For $n=1$, the construction yields nothing but the Kummer K3 surface $\mathrm{Km}(A)$ associated to A , that is, the minimal resolution of the quotient $A/\pm 1$.

We will study certain natural subvarieties of $K^n(A)$, $n \geq 2$. Consider the morphisms

$$(4) \quad \begin{aligned} f: A^m &\rightarrow A_0^{(2m)}, & (a_1, \dots, a_m) &\mapsto (a_1, -a_1, \dots, a_m, -a_m), \\ f': A^m &\rightarrow A_0^{(2m+1)}, & (a_1, \dots, a_m) &\mapsto (a_1, -a_1, \dots, a_m, -a_m, 0). \end{aligned}$$

Definition 4.1 (i) If n is odd, $n = 2m - 1$, we let $\iota : W \hookrightarrow K^n(A)$ be the strict transform of the image of $f : A^m \rightarrow A_0^{(2m)}$ under $\nu : K^n(A) \rightarrow A_0^{(2m)}$.
(ii) If n is even, $n = 2m$, we let $\iota : W \hookrightarrow K^n(A)$ be the strict transform of the image of $f' : A^m \rightarrow A_0^{(2m+1)}$ under $\nu : K^n(A) \rightarrow A_0^{(2m+1)}$.

Notice that the automorphism -1 of A naturally acts on $K^n(A)$ for all n , as the natural action of -1 on $A^{[n+1]}$ stabilizes $K^n(A)$. It is known that -1 acts trivially on $H^2(K^n(A), \mathbb{Z})$; see [Boissière et al. 2011]. The map f (resp. f') factors through an embedding of $(A/\pm 1)^{(m)}$ into $A_0^{(2m)}$ (resp. into $A_0^{(2m+1)}$); a natural crepant resolution of $(A/\pm 1)^{(m)}$ is given by $\text{Km}(A)^{[m]}$, the Hilbert scheme of points on the Kummer K3 surface $\text{Km}(A)$.

Lemma 4.2 For any $n \geq 2$, the subvariety $W \subset K^n(A)$ is the unique component of maximal dimension of the fixed locus of -1 acting on $K^n(A)$. Moreover, if $n = 2m$ (resp. $n = 2m - 1$), then f' (resp. f) induces an embedding $\iota : \text{Km}(A)^{[m]} \hookrightarrow K^{2m}(A)$ (resp. $\iota : \text{Km}(A)^{[m]} \hookrightarrow K^{2m-1}(A)$) with image W .

Proof The lemma is proven by Kamenova, Mongardi and Oblomkov [Kamenova et al. 2022]. They show in [loc. cit., Theorem 1.3] that the fixed locus of -1 on $K^n(A)$ is the union of components which are varieties of $\text{K3}^{[k]}$ -type or points, and that there exists a unique component of maximal dimension $2m$ for $n = 2m - 1$ or $n = 2m$. As our $W \subset K^n(A)$ is evidently fixed by -1 , it must be the component of maximal dimension of the fixed locus. The local study done in [loc. cit., Appendix B] shows that the birational map $\text{Km}(A)^{[m]} \dashrightarrow W$ induced by f' if $n = 2m$ or f if $n = 2m - 1$ extends to an isomorphism. $\square$

The main result of this section is the following.

Proposition 4.3 Let $\iota : W \hookrightarrow K^n(A)$ denote the embedding of Definition 4.1. Then the pullback along ι induces a primitive embedding

$$\iota^* : H^2(K^n(A), \mathbb{Z})(2) \hookrightarrow H^2(W, \mathbb{Z})$$

of lattices and Hodge structures, where on the left-hand side the form is multiplied by a factor 2.

The cases $n = 2$ and $n = 3$ may be deduced from [Hassett and Tschinkel 2013] and [Floccari 2023a], respectively.

4.2 Let us fix some notation for the cohomology groups of the varieties we shall consider. By [Beauville 1983], for $n \geq 2$ there is a canonical primitive embedding of $H^2(A, \mathbb{Z})$ into the second cohomology of $K^n(A)$. We have

$$(5) \quad H^2(K^n(A), \mathbb{Z}) = H^2(A, \mathbb{Z}) \oplus \mathbb{Z} \cdot \xi,$$

where ξ is half the class of the divisor $E \subset K^n(A)$ parametrizing nonreduced subschemes, which is the exceptional divisor of the Hilbert–Chow resolution. The class ξ has square $-2(n+1)$ with respect to the Beauville–Bogomolov form; as a lattice, $H^2(K^n(A), \mathbb{Z})$ is thus isometric to $U^{\oplus 3} \oplus \langle -2(n+1) \rangle$.

The Kummer surface $\text{Km}(A)$ associated to A contains 16 exceptional curves C_τ , parametrized by the 2-torsion points $\tau \in A_2$. The pushforward along the rational map $\pi : A \dashrightarrow \text{Km}(A)$ induces a primitive

embedding of $H^2(A, \mathbb{Z})(2)$ into $H^2(\mathrm{Km}(A), \mathbb{Z})$ whose orthogonal is the Kummer lattice (see [Nikulin 1975]), which is the saturation of the sublattice generated by the 16 pairwise orthogonal (-2) -classes $[C_\tau]$. Hence, $H^2(\mathrm{Km}(A), \mathbb{Z})$ contains with finite index the sublattice $H^2(A, \mathbb{Z})(2) \oplus \langle [C_\tau] \rangle_{\tau \in A_2}$.

For the cohomology of the Hilbert scheme $\mathrm{Km}(A)^{[m]}$, $m \geq 2$, we have by [Beauville 1983] a canonical primitive embedding of $H^2(\mathrm{Km}(A), \mathbb{Z})$ into $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$, and

$$(6) \quad H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z}) = H^2(\mathrm{Km}(A), \mathbb{Z}) \oplus \mathbb{Z} \cdot \delta,$$

where δ is half the class of the divisor $D \subset \mathrm{Km}(A)^{[m]}$ of nonreduced subschemes; the Beauville–Bogomolov square of δ is $-2(m-1)$. Under the embedding of $H^2(\mathrm{Km}(A), \mathbb{Z})$ into $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$, the class $[C_\tau]$ of the exceptional curve over the node given by $\tau \in A_2$ corresponds to the class $[R_\tau]$ of the divisor R_τ of subschemes whose support intersects C_τ . Therefore, $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$ contains with finite index a sublattice $H^2(A, \mathbb{Z})(2) \oplus \langle [R_\tau] \rangle_{\tau \in A_2} \oplus \langle \delta \rangle$ and the saturation of $\langle [R_\tau] \rangle_{\tau \in A_2}$ is the Kummer lattice.

4.3 The next lemma is the first step towards Proposition 4.3.

Lemma 4.4 Consider the embedding $\iota : W \hookrightarrow K^n(A)$ for $n \geq 2$.

(i) For $n = 2$, the pullback along $\iota : \mathrm{Km}(A) \hookrightarrow K^2(A)$ yields

$$\iota^*(\xi) = [C_0] + \frac{1}{2} \sum_{\tau \in A_2} [C_\tau] \quad \text{in } H^2(\mathrm{Km}(A), \mathbb{Z}).$$

(ii) For $n = 2m - 1$, the pullback along $\iota : \mathrm{Km}(A)^{[m]} \hookrightarrow K^{2m-1}(A)$ yields

$$\iota^*(\xi) = 2\delta + \frac{1}{2} \sum_{\tau \in A_2} [R_\tau] \quad \text{in } H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z}).$$

(iii) For $n = 2m \geq 4$, the pullback along $\iota : \mathrm{Km}(A)^{[m]} \hookrightarrow K^{2m}(A)$ yields

$$\iota^*(\xi) = 2\delta + [R_0] + \frac{1}{2} \sum_{\tau \in A_2} [R_\tau] \quad \text{in } H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z}).$$

Notice that $\frac{1}{2} \sum_{\tau} [C_\tau]$ is indeed an integral class; see [Huybrechts 2016, Chapter 14, §3.3]. The proof of the lemma will be reduced to some local computations.

Example 4.5 Consider the involution $-1 : \mathbb{C}^2 \rightarrow \mathbb{C}^2$. The quotient $\mathbb{C}^2/\pm 1$ has a nodal singularity at the image $\bar{0}$ of the origin, which is resolved by the blowup $\mathrm{Bl}_{\bar{0}}(\mathbb{C}^2/\pm 1) \rightarrow \mathbb{C}^2$. We denote by $C \subset \mathrm{Bl}_{\bar{0}}(\mathbb{C}^2/\pm 1)$ the exceptional curve.

(a) The graph map $\hat{\phi} : \mathbb{C}^2 \rightarrow (\mathbb{C}^2)^2$ given by $y \mapsto (y, -y)$ yields a Cartesian diagram

$$(7) \quad \begin{array}{ccc} \mathrm{Bl}_{\bar{0}}(\mathbb{C}^2/\pm 1) & \xhookrightarrow{\phi} & \mathbb{C}^{[2]} \\ \downarrow \nu & & \downarrow \nu \\ (\mathbb{C}^2/\pm 1) & \xhookrightarrow{\bar{\phi}} & (\mathbb{C}^2)^{(2)} \end{array}$$

By [Kamenova et al. 2022, Lemma 3.13], the image of ϕ is the fixed locus of -1 acting on $\mathbb{C}^{[2]}$. Let $F \subset \mathbb{C}^{[2]}$ be the divisor of nonreduced subschemes. Then, as divisors, $\phi^*(F) = C$; indeed, the image of ϕ and F intersect transversely in the curve $\phi(C)$.

(b) Consider now the map $\hat{\phi}: \mathbb{C}^2 \rightarrow (\mathbb{C}^2)^3$ given by $y \mapsto (y, -y, 0)$. It induces the Cartesian commutative diagram

$$(8) \quad \begin{array}{ccc} \mathrm{Bl}_0(\mathbb{C}^2/\pm 1) & \xhookrightarrow{\phi} & \mathbb{C}^{[3]} \\ \downarrow \nu & & \downarrow \nu \\ (\mathbb{C}^2/\pm 1) & \xhookrightarrow{\bar{\phi}} & (\mathbb{C}^2)^{(3)} \end{array}$$

By [Kamenova et al. 2022, Lemma 3.13], the fixed locus of -1 acting on $\mathbb{C}^{[3]}$ consists of the image of ϕ and an isolated fixed point. Let $F \subset \mathbb{C}^{[3]}$ be the divisor of nonreduced subschemes. Then we claim that, as divisors, we have $\phi^*(F) = 3C$.

To see this, let $B \subset (\mathbb{C}^2)^{[3]}$ be the Briançon variety at 0, that is, the subvariety $B := \nu^{-1}((0, 0, 0))$ that parametrizes subschemes fully supported at $0 \in \mathbb{C}^2$. Let $j: B \hookrightarrow \mathbb{C}^{[3]}$ denote the embedding. The set-theoretic intersection of F and the image of ϕ is $\phi(C)$, which is contained in B . In our case, by [Briançon 1977, IV.2], the Briançon variety B is isomorphic to the weighted projective plane $\mathbb{P}(1, 1, 3)$; following [Hassett and Tschinkel 2013, §4], there exists a closed embedding $\mathbb{P}(1, 1, 3) \hookrightarrow \mathbb{P}^4$ realizing B as the cone over the twisted cubic curve and the restriction $j^*(\mathcal{O}_{(\mathbb{C}^2)^{[3]}}(F))$ equals $\mathcal{O}_B(-2H)$, where we denote by $\mathcal{O}_B(H)$ the line bundle which induces the embedding $B \hookrightarrow \mathbb{P}^4$. Moreover, $-1: \mathbb{C}^{[3]} \rightarrow \mathbb{C}^{[3]}$ acts on B with fixed locus the vertex of the cone and the twisted cubic curve at the base, which is thus identified with $\phi(C)$. The embedding $\phi|_C: C \hookrightarrow (\mathbb{C}^2)^{[3]}$ factors as the composition of embeddings $C \xrightarrow{\phi'} B \xrightarrow{j} (\mathbb{C}^2)^{[3]}$. Therefore,

$$(9) \quad \deg(\mathcal{O}_{(\mathbb{C}^2)^{[3]}}(F)|_{\phi(C)}) = \deg(\mathcal{O}_B(-2H)|_{\phi'(C)}) = -6,$$

as H is a hyperplane divisor in $\mathbb{P}^4$ and the twisted cubic curve has degree 3. If k denotes the multiplicity of $\phi^*(F)$ at C , we have $-2k = \deg(\mathcal{O}_{(\mathbb{C}^2)^{[3]}}(F)|_{\phi(C)}) = -6$, and we conclude that $k = 3$.

To justify the last claim, consider a smooth and projective surface Y with an involution θ with isolated fixed points. Each fixed point gives a node on the quotient Y/θ ; the blowup of these nodes is a smooth and projective surface X . We choose a fixed point $p \in Y$, and let $C_p \subset X$ be the exceptional curve over the corresponding node. We then have the Cartesian diagram of projective varieties

$$\begin{array}{ccc} X & \xhookrightarrow{\phi} & Y^{[3]} \\ \downarrow \nu & & \downarrow \nu \\ Y/\theta & \xhookrightarrow{\bar{\phi}} & Y^{(3)} \end{array}$$

induced by the map $Y \rightarrow Y^3$ sending y to $(y, \theta(y), p)$. We can find an analytic neighborhood $U \subset Y$ of p stable under θ which is identified with a neighborhood V of the origin in $\mathbb{C}^2$ in such a way that the

action of θ on U corresponds to the action of -1 on V . Thus, denoting by F the exceptional divisor of $Y^{[3]} \rightarrow Y$ and by k the multiplicity of $\phi^*(F)$ at C_p , we have to show that $-2k = \deg(\mathcal{O}_{Y^{[3]}}(F)|_{\phi(C_p)})$. But $\phi^*(F)$ is supported over the exceptional curves C_q of $X \rightarrow Y/\theta$, which are pairwise orthogonal (-2) -curves. Therefore, by the projection formula, we have

$$-2k = \int_X [C_p] \cdot [\phi^*(F)] = \int_{Y^{[3]}} [\phi(C_p)] \cdot [F] = \deg(\mathcal{O}_{Y^{[3]}}(F)|_{\phi(C_p)}).$$

Proof of Lemma 4.4 The exceptional divisor $E \subset K^n(A)$ is the restriction to $K^n(A) \subset A^{[n+1]}$ of the exceptional divisor of the Hilbert–Chow resolution $A^{[n+1]} \rightarrow A^{(n+1)}$, which we still denote by E . We consider the Cartesian diagram

$$(10) \quad \begin{array}{ccc} \mathrm{Km}(A)^{[m]} & \xhookrightarrow{\iota} & A^{[n+1]} \\ \downarrow & & \downarrow \nu \\ (A/\pm 1)^{(m)} & \xhookrightarrow{\bar{\iota}} & A^{(n+1)} \end{array}$$

where $n = 2m$ or $n = 2m - 1$: in the first case, $\bar{\iota}(\bar{a}_1, \dots, \bar{a}_m) = (a_1, -a_1, \dots, a_m, -a_m, 0)$, while in the second $\bar{\iota}(\bar{a}_1, \dots, \bar{a}_m) = (a_1, -a_1, \dots, a_m, -a_m)$. It is immediate to check that, if $m \geq 2$, the restriction $\iota^*(E)$ is supported over the union of the Hilbert–Chow divisor $D \subset \mathrm{Km}(A)^{[m]}$ and the 16 divisors $R_\tau \subset \mathrm{Km}(A)^{[m]}$, for $\tau \in A_2$, while if $m = 1$ then $\iota^*(E)$ is supported over the 16 exceptional curves $C_\tau \subset \mathrm{Km}(A)$, $\tau \in A_2$, which we will also denote by R_τ to ease notation. We now have the following cases to consider.

Case 1: Let $\tau \in A_2$ be a 2-torsion point; if n is even, assume further that $\tau \neq 0$. We claim that then $\iota^*(E)$ has multiplicity 1 along the divisor R_τ .

Indeed, the image in $(A/\pm 1)^{(m)}$ of a general subscheme $\zeta \in R_\tau$ is a point of the form $(\bar{\tau}, \bar{a}_2, \dots, \bar{a}_m)$ with the $\bar{a}_i$ pairwise distinct smooth points of $A/\pm 1$, while $\iota(\zeta)$ is a subscheme in $A^{[n+1]}$ which is supported over the point $(\tau, \tau, a_2, -a_2, \dots, a_m, -a_m)$ of $A^{(n+1)}$ if $n = 2m - 1$ is odd, or over $(\tau, \tau, a_2, -a_2, \dots, a_m, -a_m, 0)$ if $n = 2m$ is even. Let $\mathcal{U}$ be an analytic neighborhood of $\iota(\zeta)$ in $A^{[n+1]}$ which is stable under -1 , and let $\mathcal{V}$ be the neighborhood $\iota^{-1}(\mathcal{U})$ of ζ in $\mathrm{Km}(A)^{[m]}$. Assume first that $n = 2m - 1$, for some $m \geq 2$. Choosing $\mathcal{U}$ small enough, diagram (10) restricts to

$$\begin{array}{ccc} \mathcal{V} \cong (\mathrm{Bl}_{\bar{\tau}}(U_\tau/\pm 1) \times \prod_{j=2}^m U_j) & \xhookrightarrow{\iota} & \mathcal{U} \cong (U_\tau^{[2]} \times \prod_{j=2}^m (U_j \times U_j^-)) \\ \downarrow & & \downarrow \nu \\ (U_\tau/\pm 1) \times \prod_{j=2}^m U_j & \xhookrightarrow{\bar{\iota}} & U_\tau^{(2)} \times \prod_{j=2}^m (U_j \times U_j^-) \end{array}$$

where: $U_\tau \subset A$ is a neighborhood of τ which is stable under -1 ; $U_j \subset A$ is a neighborhood of a_j and $U_j^- := -1(U_j)$ for $j = 2, \dots, m$; the morphism ν is the product of the Hilbert–Chow resolution $U_\tau^{[2]} \rightarrow U_\tau^{(2)}$ with the identity of each factor $(U_j \times U_j^-)$; the morphism $\bar{\iota}$ is given by $\bar{\iota}(\bar{\alpha}, \alpha_2, \dots, \alpha_m) = ((\alpha, -\alpha), \alpha_2, -\alpha_2, \dots, \alpha_m, -\alpha_m)$. If instead $n = 2m$ is even, choosing $\mathcal{U}$ small enough, diagram (10)

restricts to

$$\begin{array}{ccc} \mathcal{V} \cong (\mathrm{Bl}_{\bar{\tau}}(U_{\tau}/\pm 1) \times \prod_{j=2}^m U_j) & \xhookrightarrow{\iota} & \mathcal{U} \cong (U_{\tau}^{[2]} \times \prod_{j=2}^m (U_j \times U_j^{-}) \times U_0) \\ \downarrow & & \downarrow \nu \\ (U_{\tau}/\pm 1) \times \prod_{j=2}^m U_j & \xhookrightarrow{\bar{\iota}} & U_{\tau}^{(2)} \times \prod_{j=2}^m (U_j \times U_j^{-}) \times U_0 \end{array}$$

where: $U_{\tau} \subset A$ is a neighborhood of τ stable under -1 ; $U_j \subset A$ is a neighborhood of a_j and $U_j^{-} := -1(U_j)$ for $j = 2, \dots, m$; $U_0 \subset A$ is a neighborhood of 0 ; the morphism ν is the product of the Hilbert–Chow resolution $U_{\tau}^{[2]} \rightarrow U_{\tau}^{(2)}$ with the identity of each factor $(U_j \times U_j^{-})$ and of U_0 ; we have $\bar{\iota}(\bar{\alpha}, \alpha_2, \dots, \alpha_m) = ((\alpha, -\alpha), \alpha_2, -\alpha_2, \dots, \alpha_m, -\alpha_m, 0)$.

In both cases, we have $E|_{\mathcal{U}} = \mathrm{pr}_1^*(F)$, where $\mathrm{pr}_1 : \mathcal{U} \rightarrow U_{\tau}^{[2]}$ is the projection onto the first factor and $F \subset U_{\tau}^{[2]}$ is the divisor of nonreduced subschemes. Moreover, the divisor $R_{\tau|_{\mathcal{V}}}$ equals $\widetilde{\mathrm{pr}}_1^*(C)$, where $\widetilde{\mathrm{pr}}_1 : \mathcal{V} \rightarrow \mathrm{Bl}_{\bar{\tau}}(U_{\tau}/\pm 1)$ is the projection and $C \subset \mathrm{Bl}_{\bar{\tau}}(U_{\tau}/\pm 1)$ is the exceptional curve. We thus have the commutative diagram

$$\begin{array}{ccc} \mathcal{V} & \xhookrightarrow{\iota} & \mathcal{U} \\ \downarrow \widetilde{\mathrm{pr}}_1 & & \downarrow \mathrm{pr}_1 \\ \mathrm{Bl}_{\bar{\tau}}(U_{\tau}/\pm 1) & \xhookrightarrow{\phi} & U_{\tau}^{[2]} \end{array}$$

where ϕ coincides with the restriction of the map $\mathrm{Bl}_0(\mathbb{C}^2/\pm 1) \rightarrow (\mathbb{C}^2)^{[2]}$ of Example 4.5(a) in a neighborhood of the exceptional curve $C \subset \mathrm{Bl}_0(\mathbb{C}^2 \pm 1)$. We deduce that, as divisors, we have

$$\iota^*(E|_{\mathcal{U}}) = (\iota \circ \mathrm{pr}_1)^*(F) = (\widetilde{\mathrm{pr}}_1^*(\phi^*(F))) = \widetilde{\mathrm{pr}}_1^*(C) = R_{\tau|_{\mathcal{V}}}.$$

Hence, $\iota^*(E)$ has multiplicity 1 along the divisor $R_{\tau} \subset \mathrm{Km}(A)^{[m]}$.

Case 2: Let $n = 2m$ be even and consider the divisor R_0 of $\mathrm{Km}(A)^{[m]}$. We claim that $\iota^*(E)$ has multiplicity 3 along the divisor R_0 .

Indeed, the image in $(A/\pm 1)^{(m)}$ of a general subscheme $\zeta \in R_0$ is a point of the form $(\bar{0}, \bar{a}_2, \dots, \bar{a}_m)$ with the $\bar{a}_i$ pairwise distinct smooth points of $A/\pm 1$, while $\iota(\zeta)$ is a subscheme in $A^{[n+1]}$ supported over the point $(0, 0, a_2, -a_2, \dots, a_m, -a_m, 0)$ of $A^{(n+1)}$. Let $\mathcal{U}$ be an analytic neighborhood of $\iota(\zeta)$ in $A^{[n+1]}$ such that $-1(\mathcal{U}) = \mathcal{U}$, and let $\mathcal{V}$ be the analytic neighborhood $\iota^{-1}(\mathcal{U})$ of ζ in $\mathrm{Km}(A)^{[m]}$. Choosing $\mathcal{U}$ small enough, diagram (10) restricts to

$$\begin{array}{ccc} \mathcal{V} \cong (\mathrm{Bl}_0(U_0/\pm 1) \times \prod_{j=2}^m U_j) & \xhookrightarrow{\iota} & \mathcal{U} \cong (U_0^{[3]} \times \prod_{j=2}^m (U_j \times U_j^{-})) \\ \downarrow & & \downarrow \nu \\ (U_0/\pm 1) \times \prod_{j=2}^m U_j & \xhookrightarrow{\bar{\iota}} & U_0^{(3)} \times \prod_{j=2}^m (U_j \times U_j^{-}) \end{array}$$

where: $U_0 \subset A$ is a neighborhood of 0 stable under -1 ; $U_j \subset A$ is a neighborhood of a_j and $U_j^{-} := -1(U_j)$, for $j = 2, \dots, m$; the morphism ν is the product of the Hilbert–Chow resolution $U_0^{[3]} \rightarrow U_0^{(3)}$ with the identity of each factor $(U_j \times U_j^{-})$; we have $\bar{\iota}(\bar{\alpha}, \alpha_2, \dots, \alpha_m) = ((\alpha, -\alpha, 0), \alpha_2, -\alpha_2, \dots, \alpha_m, -\alpha_m)$.

Then $E|_{\mathcal{U}} = \text{pr}_1^*(F)$, where $\text{pr}_1 : \mathcal{U} \rightarrow U_0^{[3]}$ is the projection and $F \subset U_0^{[3]}$ is the divisor of nonreduced subschemes. Moreover, $R_0|_{\mathcal{V}} = \widetilde{\text{pr}}_1^*(C)$, where $\widetilde{\text{pr}}_1 : \mathcal{V} \rightarrow \text{Bl}_0(U_0/\pm 1)$ is the projection and $C \subset \text{Bl}_0(U_0/\pm 1)$ is the exceptional curve. We have the commutative diagram

$$\begin{array}{ccc} \mathcal{V} & \xrightarrow{\iota} & \mathcal{U} \\ \downarrow \widetilde{\text{pr}}_1 & & \downarrow \text{pr}_1 \\ \text{Bl}_0(U_0/\pm 1) & \xrightarrow{\phi} & U_0^{[3]} \end{array}$$

where ϕ coincides with the restriction of the map $\text{Bl}_0(\mathbb{C}^2/\pm 1) \rightarrow (\mathbb{C}^2)^{[3]}$ of Example 4.5(b) in a neighborhood of $C \subset \text{Bl}_0(\mathbb{C}^2/\pm 1)$, and we deduce that

$$\iota^*(E|_{\mathcal{U}}) = (\iota \circ \text{pr}_1)^*(F) = (\widetilde{\text{pr}}_1^*(\phi^*(F))) = \widetilde{\text{pr}}_1^*(3C) = 3R_0|_{\mathcal{V}}$$

as divisors. Hence, $\iota^*(E)$ has multiplicity 3 along the divisor $R_0 \subset \text{Km}(A)^{[m]}$.

Case 3: Assume now that $m \geq 2$ and consider the Hilbert–Chow divisor $D \subset \text{Km}(A)^{[m]}$. We claim that $\iota^*(E)$ has multiplicity 2 along D .

In fact, the image in $(A/\pm 1)^{(m)}$ of a general subscheme $\zeta \in D$ is a point $(\bar{x}, \bar{x}, \bar{a}_3, \dots, \bar{a}_m)$, for $m-1$ distinct smooth points $\bar{x}, \bar{a}_3, \dots, \bar{a}_m$ of $(A/\pm 1)$. The subscheme $\iota(\zeta)$ in $A^{[n+1]}$ is supported over $(x, x, -x, -x, a_3, -a_3, \dots, a_m, -a_m)$ if $n = 2m-1$ is odd, or $(x, x, -x, -x, a_3, -a_3, \dots, a_m, -a_m, 0)$ if $n = 2m$ is even. Let $\mathcal{U} \subset A^{[n+1]}$ be an analytic neighborhood of $\iota(\zeta)$, and let $\mathcal{V} \subset \text{Km}(A)^{[m]}$ be its preimage $\iota^{-1}(\mathcal{U})$. If $n = 2m-1$ is odd, shrinking $\mathcal{U}$ if necessary, there exist neighborhoods U_x of x and U_j of a_j in A for $j = 3, \dots, m$, such that diagram (10) restricts to

$$\begin{array}{ccc} \mathcal{V} \cong (U_x^{[2]} \times \prod_{j=3}^m U_j) & \xrightarrow{\iota} & \mathcal{U} \cong (U_x^{[2]} \times (U_x^-)^{[2]} \times \prod_{j=3}^m (U_j \times U_j^-)) \\ \downarrow & & \downarrow \nu \\ U_x^{(2)} \times \prod_{j=3}^m U_j & \xrightarrow{\bar{\iota}} & U_x^{(2)} \times (U_x^-)^{(2)} \times \prod_{j=3}^m (U_j \times U_j^-) \end{array}$$

where: U_x^- and U_j^- denote the image via -1 of U_x and U_j respectively, the map ν is the product of the Hilbert–Chow morphisms $U_x^{[2]} \rightarrow U_x^{(2)}$ and $(U_x^-)^{[2]} \rightarrow (U_x^-)^{(2)}$ with the identity of each factor $(U_j \times U_j^-)$ for $j = 3, \dots, m$, and $\bar{\iota}$ is given by

$$\bar{\iota}((x_1, x_2), \alpha_3, \dots, \alpha_m) = ((x_1, x_2), (-x_1, -x_2), \alpha_3, -\alpha_3, \dots, \alpha_m, -\alpha_m).$$

If instead $n = 2m$ is even, there exist neighborhoods $U_x \subset A$ of x and $U_j \subset A$ of a_j for $j = 3, \dots, m$, and a neighborhood $U_0 \subset A$ of 0 such that diagram (10) restricts to

$$\begin{array}{ccc} \mathcal{V} \cong (U_x^{[2]} \times \prod_{j=3}^m U_j) & \xrightarrow{\iota} & \mathcal{U} \cong (U_x^{[2]} \times (U_x^-)^{[2]} \times \prod_{j=3}^m (U_j \times U_j^-) \times U_0) \\ \downarrow & & \downarrow \nu \\ U_x^{(2)} \times \prod_{j=3}^m U_j & \xrightarrow{\bar{\iota}} & U_x^{(2)} \times (U_x^-)^{(2)} \times \prod_{j=3}^m (U_j \times U_j^-) \times U_0 \end{array}$$

where: U_x^- and U_j^- denote the image via -1 of U_x and U_j respectively, the map ν is the product of the Hilbert–Chow morphisms $U_x^{[2]} \rightarrow U_x^{(2)}$ and $(U_x^-)^{[2]} \rightarrow (U_x^-)^{(2)}$ with the identity of U_0 and of each

factor $(U_j \times U_j^-)$ for $j = 3, \dots, m$, and we have

$$\overline{\iota}((x_1, x_2), \alpha_3, \dots, \alpha_m) = ((x_1, x_2), (-x_1, -x_2), \alpha_3, -\alpha_3, \dots, \alpha_m, -\alpha_m, 0).$$

In both cases, the divisor $E|_{\mathcal{U}}$ breaks into two components E_1 and E_2 such that $E_1 = \text{pr}_1^*(F)$ and $E_2 = \text{pr}_2^*(F^-)$, where $\text{pr}_1 : \mathcal{U} \rightarrow U_x^{[2]}$ and $\text{pr}_2 : \mathcal{U} \rightarrow (U_x^-)^{[2]}$ denote the projections and $F \subset U_x^{[2]}$ (resp. $F^- \subset (U_x^-)^{[2]}$) is the divisor of nonreduced subschemes. Moreover, denoting the projection by $\widetilde{\text{pr}}_1 : \mathcal{V} \rightarrow U_x^{[2]}$, we have $D|_{\mathcal{V}} = \widetilde{\text{pr}}_1^*(F)$. Notice that the element $-1 \in \text{Aut}(A)$ induces an isomorphism $-1 : U_x^{[2]} \rightarrow (U_x^-)^{[2]}$ which identifies F with F^- . From the commutative diagram

$$\begin{array}{ccc} \mathcal{V} & \xrightarrow{\iota} & \mathcal{U} \\ \downarrow \widetilde{\text{pr}}_1 & & \downarrow \text{pr}_1 \times \text{pr}_2 \\ U_x^{[2]} & \xrightarrow{\text{id} \times (-1)} & U_x^{[2]} \times (U_x^-)^{[2]} \end{array}$$

we deduce that $\iota^*(E_1) = D|_{\mathcal{V}} = \iota^*(E_2)$, and thus $\iota^*(E)$ has multiplicity 2 at the divisor $D \subset \text{Km}(A)^{[m]}$.

Conclusion: Recalling that we have $[E] = 2\xi$ in $H^2(K^n(A), \mathbb{Z})$ and $[D] = 2\delta$ in $H^2(\text{Km}(A)^{[m]}, \mathbb{Z})$, the three cases discussed above yield (i), (ii) and (iii). $\square$

4.4 The key property of the embedding $\iota : \text{Km}(A)^{[m]} \rightarrow K^n(A)$ of Definition 4.1 is that it deforms together with $K^n(A)$. More precisely, we have the following.

Proposition 4.6 *Let $\mathcal{K} \rightarrow B$ be a smooth proper family of complex manifolds of Kumⁿ-type over a connected manifold B , such that for some $0 \in B$ we have $\mathcal{K}_0 = K^n(A)$ for some abelian surface A . Let $\iota_0 : \text{Km}(A)^{[m]} \rightarrow K^n(A)$ be the embedding discussed above, where $n = 2m - 1$ or $n = 2m$. Then, up to a finite étale base-change, there exist a smooth and proper family $\mathcal{W} \rightarrow B$ of manifolds of K3^[m]-type with $\mathcal{W}_0 = \text{Km}(A)^{[m]}$, and a closed embedding $\iota : \mathcal{W} \hookrightarrow \mathcal{K}$ over B extending ι_0 .*

Proof By [Hassett and Tschinkel 2013, Theorem 2.1], the group $\text{Aut}_0(X)$ of automorphisms inducing the identity on $H^2(X, \mathbb{Z})$ is deformation invariant for any hyper-Kähler manifold X . It was computed in [Boissière et al. 2011] that for $K^n(A)$ we have $\text{Aut}_0(K^n(A)) = A_{n+1} \rtimes \langle -1 \rangle$; thus, any K of Kumⁿ-type admits an action of $\text{Aut}_0(K) \cong (\mathbb{Z}/(n+1)\mathbb{Z})^4 \rtimes \mathbb{Z}/2\mathbb{Z}$. In particular, if $\mathcal{K} \rightarrow B$ is a family as in the statement, the action of $-1 \in \text{Aut}(A)$ on the fiber $\mathcal{K}_0$ extends to a fiberwise automorphism of $\mathcal{K} \rightarrow B$, up to a finite étale base-change possibly needed in order to trivialize the monodromy action on $\text{Aut}_0(\mathcal{K}_0)$. For each $b \in B$, the action of -1 on $\mathcal{K}_b$ is a symplectic involution. By [Camere 2012], each component $\mathcal{Y}$ of the fixed locus $\mathcal{K}^{(-1)}$ is a smooth and proper family $\mathcal{Y} \rightarrow B$ of symplectic manifolds, or just the image of a section of $\mathcal{K} \rightarrow B$. By Lemma 4.2, for $n = 2m$ or $n = 2m - 1$ there exists a unique such component $\mathcal{W}$ of maximal relative dimension $2m$, such that $\mathcal{W}_0 = \text{Km}(A)^{[m]}$ is embedded into $\mathcal{K}_0$ via $\iota_0 : \text{Km}(A)^{[m]} \hookrightarrow K^n(A)$. We thus have obtained the desired subfamily $\mathcal{W} \rightarrow B$ of manifolds of K3^[m]-type. $\square$

As a consequence, the cohomology class $[\iota(\text{Km}(A)^{[m]})]$ remains algebraic, and in particular a Hodge class, on any deformation of $K^n(A)$.

Definition 4.7 A canonical Hodge class on a hyper-Kähler manifold X is a Hodge class $\eta \in H^{2k}(X, \mathbb{Q})$ which remains Hodge on any deformation of X .

The most obvious examples are given by the Chern classes of X . Any canonical Hodge class $\alpha \in H^{4k}(X, \mathbb{Q})$ on X satisfies the generalized Fujiki relation [1987]: there exists a constant $c(\alpha) \in \mathbb{Q}$ such that for any $\gamma \in H^2(X, \mathbb{Q})$ we have

$$(11) \quad \int_X \alpha \cdot (\gamma)^{2n-2k} = \frac{(2n-2k)!}{(n-k)! 2^{n-k}} \cdot c(\alpha) \cdot q_X(\gamma, \gamma)^{n-k},$$

where $\dim X = 2n$ and q_X is the Beauville–Bogomolov form on $H^2(X, \mathbb{Q})$. The constant $c(1)$ is called the reduced Fujiki constant of X .

Let now $\mathcal{K} \rightarrow B$ be a family of manifolds of Kum^n -type as in Proposition 4.6. Up to a finite étale base-change, that proposition yields a subfamily $\mathcal{W} \rightarrow B$ of manifolds of $\text{K3}^{[m]}$ -type, where $n = 2m$ or $n = 2m - 1$. We let $\iota : \mathcal{W} \rightarrow \mathcal{K}$ denote the embedding.

Lemma 4.8 For any $b \in B$, the pullback $\iota_b^* : H^2(\mathcal{K}_b, \mathbb{Z}) \rightarrow H^2(\mathcal{W}_b, \mathbb{Z})$ is injective and multiplies the form by a factor 2.

Proof Extending our family $\mathcal{K} \rightarrow B$ if necessary, we may assume that the very general fiber has Picard rank 0, so that $H^2(\mathcal{K}_b, \mathbb{Z})$ is an irreducible Hodge structure for very general b . A symplectic form on $\mathcal{K}_b$ restricts to a symplectic form on $\mathcal{W}_b$ (see [Camere 2012, Proposition 3]), and hence ι_b^* is not zero. But then $\iota_b^* : H^2(\mathcal{K}_b, \mathbb{Q}) \rightarrow H^2(\mathcal{W}_b, \mathbb{Q})$ must be an embedding of rational Hodge structures for very general $b \in B$. Since $H^2(\mathcal{K}_b, \mathbb{Z})$ is torsion-free, we conclude that $\iota_b^* : H^2(\mathcal{K}_b, \mathbb{Z}) \rightarrow H^2(\mathcal{W}_b, \mathbb{Z})$ is injective for all $b \in B$.

For the second assertion, it suffices to consider $\iota : \text{Km}(A)^{[m]} \rightarrow K^n(A)$. Assume first that $n = 2$ and $m = 1$. The cohomology class $[\iota(\text{Km}(A))] \in H^4(K^2(A), \mathbb{Z})$ is a canonical Hodge class and, hence, there exists a constant c such that

$$q_{\text{Km}(A)}(\iota^*(\gamma), \iota^*(\gamma)) = \int_{\text{Km}(A)} \iota^*(\gamma)^2 = \int_{K^2(A)} [\text{Km}(A)] \cdot \gamma^2 = c \cdot q_{K^2(A)}(\gamma, \gamma)$$

for any $\gamma \in H^2(K^2(A), \mathbb{Z})$. For the class $\xi \in H^2(K^2(A), \mathbb{Z})$, we have $q_{K^2(A)}(\xi, \xi) = -6$ while $q_{\text{Km}(A)}(\iota^*(\xi), \iota^*(\xi)) = -12$ by Lemma 4.4(i), and we get $c = 2$.

Consider next the case when $n = 2m - 1$ is odd, with $m \geq 2$. The cohomology class $[\iota(\text{Km}(A)^{[m]})]$ in $H^{4m-4}(K^{2m-1}(A), \mathbb{Z})$ is a canonical Hodge class. Using the generalized Fujiki relation (11) for $[\iota(\text{Km}(A)^{[m]})]$ and Fujiki relation for $\text{Km}(A)^{[m]}$, there exists a constant c such that

$$(12) \quad \frac{(2m)!}{m! 2^m} \cdot c_{\text{K3}^{[m]}} \cdot q_{\text{Km}(A)^{[m]}}(\iota^*(\gamma), \iota^*(\gamma))^m = \int_{\text{Km}(A)^{[m]}} \iota^*(\gamma)^{2m} = \frac{(2m)!}{m! 2^m} \cdot c \cdot q_{K^{2m-1}(A)}(\gamma, \gamma)^m$$

for any $\gamma \in H^2(K^{2m-1}(A), \mathbb{Z})$, where $c_{\text{K3}^{[m]}} = 1$ is the reduced Fujiki constant of manifolds of $\text{K3}^{[m]}$ -type (computed in [Beauville 1983]). For the class $\xi \in H^2(K^{2m-1}(A), \mathbb{Z})$ which is half the class of the

divisor of nonreduced subschemes, we have $q_{K^{2m-1}(A)}(\xi, \xi) = -4m$ and, by Lemma 4.4(ii), we have $q_{\mathrm{Km}(A)^{[m]}}(\iota^*(\xi), \iota^*(\xi)) = -8m$. We thus obtain

$$(-8m)^m = c \cdot (-4m)^m,$$

which gives $c = 2^m$. Therefore, for any $\gamma \in H^2(K^{2m-1}(A), \mathbb{Z})$ equation (12) yields

$$q_{\mathrm{Km}(A)^{[m]}}(\iota^*(\gamma), \iota^*(\gamma))^m = 2^m \cdot q_{K^{2m-1}(A)}(\gamma, \gamma)^m,$$

and hence ι^* multiplies the form by 2 (for m even, the sign is determined as ι^* sends Kähler classes to Kähler classes).

Finally, assume $n = 2m \geq 4$ is even. As above, there exists a constant c such that

$$(13) \quad c_{K3[m]} \cdot q_{\mathrm{Km}(A)^{[m]}}(\iota^*(\gamma), \iota^*(\gamma))^m = \int_{\mathrm{Km}(A)^{[m]}} \iota^*(\gamma)^{2m} = c \cdot q_{K^{2m}(A)}(\gamma, \gamma)^m$$

for all $\gamma \in H^2(K^{2m}(A), \mathbb{Z})$. For the class $\xi \in H^2(K^{2m}(A), \mathbb{Z})$ which is half the class of nonreduced subschemes, we have $q_{K^{2m}(A)}(\xi, \xi) = -4m - 2$, while we have $q_{\mathrm{Km}(A)^{[m]}}(\iota^*(\xi), \iota^*(\xi)) = -8m - 4$ by Lemma 4.4(iii). As $c_{K3[m]} = 1$, we obtain

$$(-8m - 4)^m = c \cdot (-4m - 2)^m,$$

and hence $c = 2^m$; by (13), we have $q_{\mathrm{Km}(A)^{[m]}}(\iota^*(\gamma), \iota^*(\gamma))^m = 2^m \cdot q_{K^{2m}(A)}(\gamma, \gamma)^m$ for any γ in $H^2(K^{2m}(A), \mathbb{Z})$, and again we deduce that ι^* multiplies the form by 2. $\square$

4.5 We can finally conclude our study of the pullback ι^* .

Proof of Proposition 4.3 Thanks to Lemma 4.8, it remains to show that the image of

$$\iota^* : H^2(K^n(A), \mathbb{Z}) \hookrightarrow H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$$

is a saturated sublattice. As in (5), we have $H^2(K^n(A), \mathbb{Z}) = H^2(A, \mathbb{Z}) \oplus \langle \xi \rangle$, and the Hodge structure on $H^2(K^n(A), \mathbb{Z})$ is determined by that on $H^2(A, \mathbb{Z})$ via this equality. Moreover, the Hodge structure on $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$ is determined by that on its primitive sublattice $H^2(A, \mathbb{Z})(2)$ (Section 4.2). Deforming to a very general complex torus A such that $H^2(A, \mathbb{Z})$ is an irreducible Hodge structure, we see that the restriction ι^* necessarily maps the summand $H^2(A, \mathbb{Z})$ of $H^2(K^n(A), \mathbb{Z})$ to the sublattice $H^2(A, \mathbb{Z})(2)$ of $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$. In fact, the restriction of ι^* to $H^2(A, \mathbb{Z}) \subset H^2(K^n(A), \mathbb{Z})$ is surjective onto $H^2(A, \mathbb{Z})(2) \subset H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$, as follows from the discriminant-index formula (see [Huybrechts 2016, Chapter 14]) since ι^* is injective and multiplies the form by 2 by Lemma 4.8.

Assume first that $n = 2$ and $m = 1$. Then, by the above and Lemma 4.4, the image of ι^* is the sublattice of $H^2(\mathrm{Km}(A), \mathbb{Z})$ generated by $H^2(A, \mathbb{Z})(2)$ and the class $[C_0] + \frac{1}{2} \sum_{\tau \in A_2} [C_\tau]$. This sublattice is saturated: if $w := \alpha u + \beta \cdot ([C_0] + \frac{1}{2} \sum_{\tau \in A_2} [C_\tau])$ is an integral class in $H^2(\mathrm{Km}(A), \mathbb{Z})$ for some primitive $u \in H^2(A, \mathbb{Z})(2)$ and some rational numbers α and β , then from $q_{\mathrm{Km}(A)}(w, [C_\tau]) = -\beta$ for any $\tau \neq 0$ we deduce that α and β are integers, and therefore w belongs to $\mathrm{im}(\iota^*)$.

For $m \geq 2$, the lattice $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$ contains with finite index the sublattice

$$H^2(A, \mathbb{Z})(2) \oplus \langle [R_\tau] \rangle_{\tau \in A_2} \oplus \langle \delta \rangle.$$

By Lemma 4.4, if $n = 2m - 1$, the image of ι^* is the sublattice of $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$ generated by $H^2(A, \mathbb{Z})(2)$ and the class $2\delta + \frac{1}{2} \sum_{\tau \in A_2} [R_\tau]$, while if $n = 2m$ the image of ι^* is the sublattice of $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$ generated by $H^2(A, \mathbb{Z})(2)$ and the class $2\delta + [R_0] + \frac{1}{2} \sum_{\tau \in A_2} [R_\tau]$. In both cases, the same argument used above shows that the image of ι^* is a saturated sublattice of $H^2(\mathrm{Km}(A)^{[m]}, \mathbb{Z})$. $\square$

Remark 4.9 A manifold of Kum^n -type K admits at least $(n + 1)^4$ involutions acting trivially on the second cohomology, whose fixed loci are union of $\mathrm{K3}^{[j]}$ -manifolds by [Kamenova et al. 2022]. Our results adapt immediately to all components of maximal dimension. For a lower-dimensional component $\iota : Z \hookrightarrow K$, the argument of Lemma 4.8 gives that $\iota^* : H^2(K, \mathbb{Z}) \rightarrow H^2(Z, \mathbb{Z})$ is injective and multiplies the form by some positive constant k ; however, determining this constant and the saturation of the image of ι^* becomes more difficult.

5 K3 surfaces associated with varieties of generalized Kummer type

5.1 In this section we shall conclude the proof of our main results.

Definition 5.1 Let K be a variety of Kum^n -type. The associated K3 surface S_K is the K3 surface with transcendental lattice Hodge isometric to $H_{\mathrm{tr}}^2(K, \mathbb{Z})(2)$.

In the special case of the generalized Kummer variety $K^n(A)$ on an abelian surface A , we have $H_{\mathrm{tr}}^2(K^n(A), \mathbb{Z}) = H_{\mathrm{tr}}^2(A, \mathbb{Z})$, and the K3 surface $S_{K^n(A)}$ is isomorphic to the Kummer K3 surface $\mathrm{Km}(A)$. This motivates our Definition 5.1. Note that for K very general S_K is projective of Picard rank 16, and thus it is not a Kummer K3 surface.

Lemma 5.2 *The K3 surface S_K exists and it is unique up to isomorphism.*

Proof The lattice $H_{\mathrm{tr}}^2(K, \mathbb{Z})(2)$ has rank at most 6 and signature $(2, k)$. Any such lattice can be embedded primitively in the K3 lattice and, by the Torelli theorem, there exists a K3 surface S_K whose transcendental lattice $H_{\mathrm{tr}}^2(S_K, \mathbb{Z})$ is Hodge isometric to $H_{\mathrm{tr}}^2(K, \mathbb{Z})(2)$; see [Morrison 1984, Corollary 2.10]. The K3 surface S_K is necessarily projective, of Picard rank at least 16, and it is determined up to isomorphism because projective K3 surfaces with Hodge isometric transcendental lattices and Picard rank at least 12 are isomorphic; see [Huybrechts 2016, Chapter 16, Corollary 3.8]. $\square$

The K3 surfaces associated with varieties of generalized Kummer type can be easily characterized via the Torelli theorem.

Lemma 5.3 *Let S be a projective K3 surface. Then S is isomorphic to the K3 surface S_K associated to some variety K of Kum^n -type if and only if there exists a primitive embedding of lattices*

$$H_{\mathrm{tr}}^2(S, \mathbb{Z}) \hookrightarrow \mathrm{U}(2)^{\oplus 3} \oplus \langle -4n - 4 \rangle.$$

Proof The lattice $U(2)^{\oplus 3} \oplus \langle -4n-4 \rangle$ is the Beauville–Bogomolov lattice $\Lambda_{\text{Kum}^n}(2)$ of Kum^n -manifolds with the form multiplied by 2. Thus, by definition, for the K3 surface S_K associated to a variety K of Kum^n -type we have a primitive embedding of $H_{\text{tr}}^2(S_K, \mathbb{Z})$ into the lattice $U(2)^{\oplus 3} \oplus \langle -4n-4 \rangle$. Conversely, assume that S is a projective K3 surface and $l : H_{\text{tr}}^2(S, \mathbb{Z}) \hookrightarrow \Lambda_{\text{Kum}^n}(2)$ is a primitive embedding. Let $T \subset \Lambda_{\text{Kum}^n}$ be the primitive sublattice such that $\text{im}(l) = T(2)$, and equip T with the Hodge structure induced by that on $H_{\text{tr}}^2(S, \mathbb{Z})$. By the surjectivity of the period map for hyper-Kähler manifolds [Huybrechts 1999, Theorem 8.1], there exists a manifold K of Kum^n -type with $H_{\text{tr}}^2(K, \mathbb{Z})$ Hodge isometric to T . By construction, we have a Hodge isometry

$$H_{\text{tr}}^2(S, \mathbb{Z}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Z})(2);$$

K is projective by [Huybrechts 1999, Theorem 3.11], since its transcendental lattice has signature $(2, k)$. Moreover, the K3 surface S is isomorphic to S_K , as they have Hodge isometric transcendental lattices and Picard rank at least 16. $\square$

Passing to rational coefficients, a similar argument yields the following.

Lemma 5.4 *Let S be a projective K3 surface. The following are equivalent:*

- (i) *There exists a rational Hodge isometry $H_{\text{tr}}^2(S, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(S_K, \mathbb{Q})$, where S_K is the K3 surface associated to a variety K of Kum^n -type for some $n \geq 2$.*
- (ii) *There exists an isometric embedding of $H_{\text{tr}}^2(S, \mathbb{Q})$ into $(U^{\oplus 3} \oplus \langle -m \rangle) \otimes_{\mathbb{Z}} \mathbb{Q}$ for some positive integer m .*

Proof For any nonzero rational number k , we have $U(k) \otimes_{\mathbb{Z}} \mathbb{Q} \cong U \otimes_{\mathbb{Z}} \mathbb{Q}$. If (i) holds, then $H_{\text{tr}}^2(S, \mathbb{Q})$ is Hodge isometric to $H_{\text{tr}}^2(K, \mathbb{Q})(2)$, and, hence, it embeds isometrically into $(U(2)^{\oplus 3} \oplus \langle -4n-4 \rangle) \otimes_{\mathbb{Z}} \mathbb{Q} \cong (U^{\oplus 3} \oplus \langle -n-1 \rangle) \otimes_{\mathbb{Z}} \mathbb{Q}$. Conversely, if S is a projective K3 surface with an embedding $H_{\text{tr}}^2(S, \mathbb{Q}) \hookrightarrow (U^{\oplus 3} \oplus \langle -m \rangle) \otimes_{\mathbb{Z}} \mathbb{Q}$ for some positive $m \in \mathbb{Z}$, we find as well an isometric embedding

$$j : H_{\text{tr}}^2(S, \mathbb{Q}) \hookrightarrow (U(2)^{\oplus 3} \oplus \langle -4n-4 \rangle) \otimes_{\mathbb{Z}} \mathbb{Q}$$

for some $n \geq 2$. The quadratic space on the right-hand side is $\Lambda_{\text{Kum}^n}(2) \otimes_{\mathbb{Z}} \mathbb{Q}$. We choose a primitive sublattice $T \subset \Lambda_{\text{Kum}^n}$ such that $\text{im}(j)$ coincides with $T(2) \otimes_{\mathbb{Z}} \mathbb{Q}$; under this identification, the Hodge structure on $H_{\text{tr}}^2(S, \mathbb{Q})$ induces a Hodge structure of K3-type on T . By the surjectivity of the period map, there exists a projective variety K of Kum^n -type such that $H_{\text{tr}}^2(K, \mathbb{Z})$ is Hodge isometric to T , equipped with this Hodge structure. Hence, by construction, $H_{\text{tr}}^2(S, \mathbb{Q})$ is Hodge isometric to $H_{\text{tr}}^2(S_K, \mathbb{Q})$. $\square$

5.2 A priori, the K3 surface S_K and the variety K of Kum^n -type are only related at the level of lattices and Hodge structures. But in fact they are related geometrically, as follows. Choose a family $\mathcal{H} \rightarrow B$ of manifolds of Kum^n -type, with one fiber $\mathcal{H}_{b_0} = K$ and another fiber $\mathcal{H}_{b_1} = K^n(A)$ isomorphic to the generalized Kummer variety on an abelian surface A . By Proposition 4.6, up to a finite étale base-change, we obtain a family $\mathcal{W} \rightarrow B$ of $\text{K3}^{[m]}$ -manifolds with an embedding $\iota : \mathcal{W} \rightarrow \mathcal{H}$ over B , where $n = 2m - 1$

or $n = 2m$, which extends the embedding $\iota_{b_1} : \mathrm{Km}(A)^{[m]} \hookrightarrow K^n(A)$ of Lemma 4.2. Taking the fiber over b_0 , we get a submanifold $\iota : W_K \hookrightarrow K$ of $\mathrm{K3}^{[m]}$ -type.

Lemma 5.5 *The pullback along $\iota : W_K \hookrightarrow K$ induces a primitive embedding of lattices and Hodge structures*

$$\iota^* : H^2(K, \mathbb{Z})(2) \hookrightarrow H^2(W_K, \mathbb{Z}).$$

In particular, ι^* induces a Hodge isometry $H_{\mathrm{tr}}^2(K, \mathbb{Z})(2) \xrightarrow{\sim} H_{\mathrm{tr}}^2(W_K, \mathbb{Z})$ of transcendental lattices.

Proof The lemma follows immediately from Proposition 4.3 as $\iota : W_K \hookrightarrow K$ deforms to the embedding $\mathrm{Km}(A)^{[m]} \hookrightarrow K^n(A)$ there considered. $\square$

We refer to [Huybrechts and Lehn 2010; Bridgeland 2008; Bayer and Macrì 2014b] for background on moduli spaces of sheaves on K3 surfaces.

Proposition 5.6 *Let S_K be the K3 surface associated to the Kum^n -variety K . Then:*

- (i) *If $n = 2$, W_K is isomorphic to the K3 surface S_K .*
- (ii) *If $n \geq 3$, $W_K \subset K$ is birational to a smooth and projective moduli space $M_{S_K, H}(v)$ for some primitive Mukai vector v and a v -generic polarization H .*

Proof A Hodge isometry $H_{\mathrm{tr}}^2(W_K, \mathbb{Z}) \xrightarrow{\sim} H_{\mathrm{tr}}^2(S_K, \mathbb{Z})$ of transcendental lattices exists by Lemma 5.5 and the definition of S_K . If $n = 2$, W_K and S_K are isomorphic as they are K3 surfaces with Hodge isometric transcendental lattice and large Picard rank. If $n \geq 3$, the proposition follows applying [Addington 2016, Proposition 4]. $\square$

Remark 5.7 In fact, for $n \geq 3$, W_K is isomorphic to a moduli space of Bridgeland stable objects on S_K , by [Bayer and Macrì 2014a, Theorem 1.2(c)]; see also [Mongardi and Wandel 2017, Proposition 2.3].

5.3 Let S be a projective K3 surface, and let $M = M_{S, H}(v)$ be a smooth and projective moduli space of stable sheaves of dimension ≥ 4 , for a primitive Mukai vector v and a v -generic polarization H . Then, following [Mukai 1987; O'Grady 1997], there exists a quasiuniversal sheaf $\mathcal{U}$ over $S \times M$ of similitude ρ , and the map

$$\alpha \mapsto \frac{1}{\rho} \cdot \mathrm{pr}_{M,*}[\mathrm{ch}(\mathcal{U})^\vee \cdot \sqrt{\mathrm{td}(S)} \cdot \mathrm{pr}_S^*(\alpha)]_3$$

induces a Hodge isometry $\theta : H_{\mathrm{tr}}^2(S, \mathbb{Z}) \xrightarrow{\sim} H_{\mathrm{tr}}^2(M, \mathbb{Z})$ of transcendental lattices. Here, the notation $[-]_3$ indicates that we take the Künneth component in $H^6(S \times M, \mathbb{Z})$ and $\mathrm{pr}_M, \mathrm{pr}_S$ denote the projections from $S \times M$ onto the two factors. The standard conjectures hold for both the surface S and the $\mathrm{K3}^{[n]}$ -variety M , as we recalled in Section 2.5. Hence, they hold for $S \times M$, and it follows that θ is the correspondence induced by the algebraic cycle $[(1/\rho)\mathrm{ch}(\mathcal{U})^\vee \cdot \sqrt{\mathrm{td}(S)}]_2$ on $S \times M$.

Remark 5.8 In fact, the Künneth components $h^2(S)$ and $h^2(M)$ as well as their transcendental parts $h_{\mathrm{tr}}^2(S)$ and $h_{\mathrm{tr}}^2(M)$ are well-defined homological motives; the Hodge isometry θ is the realization of an isomorphism of motives $\Theta : h_{\mathrm{tr}}^2(S) \rightarrow h_{\mathrm{tr}}^2(M)$ in Mot . This again uses the standard conjectures for

$S \times M$ and in particular the conservativity of the realization functor when restricted to $\langle h(S \times M) \rangle_{\text{Mot}}$ (Remark 2.4).

5.4 Let K be a variety of Kum^{*n*}-type. By the results of Foster [2024] recalled in Theorem 2.6, the Künneth component $h^2(K)$ of the motive of K is well-defined in Mot and $\langle h^2(K) \rangle_{\text{Mot}}$ is a semisimple and abelian neutral Tannakian category; the transcendental part $h_{\text{tr}}^2(K)$ of $h^2(K)$ is also a well-defined homological motive. The next result implies Theorem 1.1 from the introduction.

Theorem 5.9 *Let K be a variety of Kum^{*n*}-type and let S_K be the K3 surface associated with K . There exists an isomorphism*

$$h_{\text{tr}}^2(S_K) \xrightarrow{\sim} h_{\text{tr}}^2(K)$$

of motives, whose realization is a Hodge isometry $H_{\text{tr}}^2(S_K, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Q})(2)$.

Proof For $n = 2m$ or $n = 2m - 1$, we find as in Section 5.2 a K3^[*m*]-variety W_K which is isomorphic to S_K for $n = 2$ and birational to a moduli space of stable sheaves on S_K if $n \geq 3$, and that admits an embedding $\iota : W_K \hookrightarrow K$ such that the restriction induces a Hodge isometry $\iota^* : H_{\text{tr}}^2(K, \mathbb{Z})(2) \xrightarrow{\sim} H_{\text{tr}}^2(W_K, \mathbb{Z})$.

As ι^* is clearly induced by an algebraic cycle, it is the realization of a morphism $t : h_{\text{tr}}^2(K) \rightarrow h_{\text{tr}}^2(W_K)$ of motives. But in fact t is a morphism in the subcategory $\langle h_{\text{tr}}^2(K) \oplus h_{\text{tr}}^2(W_K) \rangle_{\text{Mot}}$, which is an abelian and semisimple neutral Tannakian category as the standard conjectures hold for W_K and $h^2(K)$. Since the realization of t is an isomorphism of rational Hodge structures, it follows that t is an isomorphism of motives; the realization of its inverse t^{-1} is a Hodge isometry $H_{\text{tr}}^2(W_K, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Q})(2)$. This concludes the proof for $n = 2$, since S_K is isomorphic to W_K in this case.

If $n \geq 3$, let $\phi : M_{S_K, H}(v) \dashrightarrow W_K$ be a birational map, where $M_{S_K, H}(v)$ is a smooth and projective moduli space of stable sheaves on S_K . Then ϕ induces a morphism $\Phi_* : h^2(M_{S_K, H}(v)) \rightarrow h^2(W_K)$ of motives whose realization is a Hodge isometry $\phi_* : H^2(M_{S_K, H}(v), \mathbb{Z}) \xrightarrow{\sim} H^2(W_K, \mathbb{Z})$, by [Huybrechts 1999, Corollary 5.2]. As the standard conjectures hold for K3^[*m*]-varieties, Φ_* is an isomorphism of motives. From Remark 5.8 we obtain an isomorphism of motives $\Theta : h_{\text{tr}}^2(S_K) \xrightarrow{\sim} h_{\text{tr}}^2(M_{S_K, H}(v))$, whose realization is a Hodge isometry $H_{\text{tr}}^2(S_K, \mathbb{Z}) \xrightarrow{\sim} H_{\text{tr}}^2(M_{S_K, H}(v), \mathbb{Z})$. We conclude that the composition $t^{-1} \circ \Phi_* \circ \Theta : h_{\text{tr}}^2(S_K) \rightarrow h_{\text{tr}}^2(K)$ is an isomorphism of motives, whose realization is the desired Hodge isometry $H_{\text{tr}}^2(S_K, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Q})(2)$. $\square$

Thanks to the work of O’Grady, Markman, Voisin and Varesco, recalled in the introduction, Theorem 5.9 yields the following result, which was stated as Theorem 1.2.

Theorem 5.10 *Let K be a variety of Kum^{*n*}-type with associated K3 surface S_K . Then:*

- (i) *The Kuga–Satake correspondence for S_K is algebraic.*
- (ii) *The Hodge conjecture holds for any power of S_K .*

Proof By Theorem 5.9 there exists an algebraic cycle inducing a Hodge isometry

$$H_{\text{tr}}^2(S_K, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Q})(2).$$

The Kuga–Satake Hodge conjecture (Conjecture 3.3) holds for K by [Voisin 2022, Theorem 0.5]. By Remark 3.7, the same conclusion holds for S_K , and we obtain (i).

The K3 surfaces S_K come in countably many 4-dimensional families of general Picard rank 16, corresponding to the polarized families of varieties of generalized Kummer type. By construction (see Section 3.5) the Kuga–Satake variety of S_K is isogenous to a power of $\text{KS}(K)$, and, hence, to a power of a Weil fourfold with discriminant 1 by [Markman 2023; O’Grady 2021]. As Conjecture 3.3 holds for any of the K3 surfaces S_K , we can apply [Varesco 2025, Theorem 0.2] to deduce that the Hodge conjecture holds for all powers of the K3 surfaces S_K . $\square$

5.5 We finally complete the proof of our results.

Theorem 5.11 *Let S be a projective K3 surface such that there exists an isometric embedding*

$$H_{\text{tr}}^2(S, \mathbb{Q}) \hookrightarrow (U^{\oplus 3} \oplus \langle -m \rangle) \otimes_{\mathbb{Z}} \mathbb{Q}$$

for some positive integer m . Then:

- (i) *The Kuga–Satake correspondence for S is algebraic.*
- (ii) *The Hodge conjecture holds for any power of S .*

Proof Let S be as in the statement. By Lemma 5.4, for some $n \geq 2$ there exist a variety K of Kum^n -type and a rational Hodge isometry $\phi : H_{\text{tr}}^2(S, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(S_K, \mathbb{Q})$. By [Buskin 2019; Huybrechts 2019], ϕ is algebraic and the motives of S and S_K are isomorphic. Therefore, by Remark 3.7 and Theorem 5.10(i), the Kuga–Satake correspondence for S is algebraic, and, by Theorem 5.10(ii), the Hodge conjecture holds for any power of S . $\square$

We next prove Theorem 1.3, which we state again below.

Theorem 5.12 *Let A be an abelian fourfold of Weil type with discriminant 1. Then the Hodge conjecture holds for A and any of its powers.*

Proof By [O’Grady 2021] and [Markman 2023], there exists a Kum^n -variety K such that A is isogenous to the intermediate Jacobian $J^3(K)$; moreover, the Kuga–Satake variety of K is isogenous to a power of A . The Kuga–Satake variety of the K3 surface S_K associated to K is thus also isogenous to a power of A . By Theorem 5.10, the Kuga–Satake Hodge conjecture holds for S_K and the Hodge conjecture holds for any power of S_K . Therefore, Theorem 3.5 implies that the Hodge conjecture holds for A and all of its powers. $\square$

We conclude with the following corollary, which follows from a result of Varesco [2023].

Corollary 5.13 *Let S be a projective K3 surface and let K be a projective variety of Kum^n -type. Assume that $\phi : H_{\text{tr}}^2(S_K, \mathbb{Q}) \xrightarrow{\sim} H_{\text{tr}}^2(K, \mathbb{Q})(k)$ is a Hodge isometry, where on the right-hand side the form is multiplied by some positive $k \in \mathbb{Q}$. Then ϕ is induced by an algebraic cycle on $S \times K$.*

Proof As $\Lambda_{\text{Kum}}^n = U^{\oplus 3} \oplus \langle -2n - 2 \rangle$ and $U(k) \otimes_{\mathbb{Z}} \mathbb{Q} \cong U \otimes_{\mathbb{Z}} \mathbb{Q}$ for any $0 \neq k \in \mathbb{Q}$, there exists by hypothesis an isometric embedding

$$H_{\text{tr}}^2(S_K, \mathbb{Q}) \hookrightarrow U_{\mathbb{Q}}^{\oplus 3} \oplus \langle -m \rangle_{\mathbb{Q}}$$

for some positive integer m . Hence, by Theorem 1.2, the Kuga–Satake Hodge conjecture holds for S . Moreover, this conjecture holds for K by [Voisin 2022, Theorem 0.5], and the Lefschetz standard conjecture in degree 2 for K is proven in [Foster 2024, Theorem 1.1]. We can therefore apply [Varesco 2023, Corollary 4.6] to obtain that ϕ is algebraic. $\square$

Acknowledgements

I am grateful to Lie Fu, Bert van Geemen, Ben Moonen, Ángel Ríos Ortiz and Charles Vial for useful conversations and to the referee for helpful comments. I thank Stefan Schreieder for his support. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Project-ID 491392403 - TRR 358.

References

- [Addington 2016] **N Addington**, *On two rationality conjectures for cubic fourfolds*, Math. Res. Lett. 23:1 (2016) 1–13 MR
- [Ancona et al. 2025] **G Ancona, M Cavicchi, R Laterveer, G Saccà**, *Relative and absolute Lefschetz standard conjectures for some Lagrangian fibrations*, J. Lond. Math. Soc. (2) 111:4 (2025) art. id. e70133 MR
- [André 1996a] **Y André**, *On the Shafarevich and Tate conjectures for hyper-Kähler varieties*, Math. Ann. 305:2 (1996) 205–248 MR
- [André 1996b] **Y André**, *Pour une théorie inconditionnelle des motifs*, Inst. Hautes Études Sci. Publ. Math. :83 (1996) 5–49 MR
- [André 2004] **Y André**, *Une introduction aux motifs (motifs purs, motifs mixtes, périodes)*, Panoramas et Synthèses 17, Soc. Math. France, Paris (2004) MR
- [Arapura 2006] **D Arapura**, *Motivation for Hodge cycles*, Adv. Math. 207:2 (2006) 762–781 MR
- [Bayer and Macrì 2014a] **A Bayer, E Macrì**, *MMP for moduli of sheaves on K3s via wall-crossing: nef and movable cones, Lagrangian fibrations*, Invent. Math. 198:3 (2014) 505–590 MR
- [Bayer and Macrì 2014b] **A Bayer, E Macrì**, *Projectivity and birational geometry of Bridgeland moduli spaces*, J. Amer. Math. Soc. 27:3 (2014) 707–752 MR
- [Beauville 1983] **A Beauville**, *Variétés Kähleriennes dont la première classe de Chern est nulle*, J. Differential Geom. 18:4 (1983) 755–782 MR
- [Boissière et al. 2011] **S Boissière, M Nieper-Wißkirchen, A Sarti**, *Higher dimensional Enriques varieties and automorphisms of generalized Kummer varieties*, J. Math. Pures Appl. (9) 95:5 (2011) 553–563 MR
- [Bolognesi and Laterveer 2024] **M Bolognesi, R Laterveer**, *A 9-dimensional family of K3 surfaces with finite-dimensional motive*, Rend. Circ. Mat. Palermo (2) 73:7 (2024) 2313–2331 MR
- [Brianc¸on 1977] **J Brianc¸on**, *Description de $\text{Hilb}^n C\{x, y\}$* , Invent. Math. 41:1 (1977) 45–89 MR
- [Bridgeland 2008] **T Bridgeland**, *Stability conditions on K3 surfaces*, Duke Math. J. 141:2 (2008) 241–291 MR
- [Bülles 2020] **T-H Bülles**, *Motives of moduli spaces on K3 surfaces and of special cubic fourfolds*, Manuscripta Math. 161:1-2 (2020) 109–124 MR
- [Buskin 2019] **N Buskin**, *Every rational Hodge isometry between two K3 surfaces is algebraic*, J. Reine Angew. Math. 755 (2019) 127–150 MR

- [Camere 2012] **C Camere**, *Symplectic involutions of holomorphic symplectic four-folds*, Bull. Lond. Math. Soc. 44:4 (2012) 687–702 MR
- [Charles and Markman 2013] **F Charles, E Markman**, *The standard conjectures for holomorphic symplectic varieties deformation equivalent to Hilbert schemes of K3 surfaces*, Compos. Math. 149:3 (2013) 481–494 MR
- [Deligne 1972] **P Deligne**, *La conjecture de Weil pour les surfaces K3*, Invent. Math. 15 (1972) 206–226 MR
- [Deligne and Milne 1982] **P Deligne, JS Milne**, *Tannakian categories*, from “Hodge cycles, motives, and Shimura varieties” (A Dold, B Eckmann, editors), Lecture Notes in Mathematics 900, Springer (1982) 101–228 MR
- [Floccari 2023a] **S Floccari**, *The Hodge and Tate conjectures for hyper-Kähler sixfolds of generalized Kummer type*, preprint (2023) arXiv 2308.02267
- [Floccari 2023b] **S Floccari**, *On the motive of O’Grady’s six-dimensional hyper-Kähler varieties*, Épijournal Géom. Algébrique 7 (2023) art. id. 4 MR
- [Floccari 2024] **S Floccari**, *Sixfolds of generalized Kummer type and K3 surfaces*, Compos. Math. 160:2 (2024) 388–410 MR
- [Floccari and Varesco 2025] **S Floccari, M Varesco**, *Algebraic cycles on hyper-Kähler varieties of generalized Kummer type*, Math. Ann. 391:3 (2025) 4443–4453 MR
- [Floccari et al. 2021] **S Floccari, L Fu, Z Zhang**, *On the motive of O’Grady’s ten-dimensional hyper-Kähler varieties*, Commun. Contemp. Math. 23:4 (2021) art. id. 2050034 MR
- [Foster 2024] **J Foster**, *The Lefschetz standard conjectures for IHSMs of generalized Kummer deformation type in certain degrees*, Eur. J. Math. 10:2 (2024) art. id. 34 MR
- [Fujiki 1987] **A Fujiki**, *On the de Rham cohomology group of a compact Kähler symplectic manifold*, from “Algebraic geometry” (Sendai, 1985) (T Oda, editor), Adv. Stud. Pure Math. 10, North-Holland, Amsterdam (1987) 105–165 MR
- [Fulton 1998] **W Fulton**, *Intersection theory*, 2nd edition, Ergebnisse der Math. (3) 2, Springer (1998) MR
- [van Geemen 2000] **B van Geemen**, *Kuga–Satake varieties and the Hodge conjecture*, from “The arithmetic and geometry of algebraic cycles” (Banff, 1998) (B B Gordon, JD Lewis, S Müller-Stach, S Saito, N Yui, editors), NATO Sci. Ser. C Math. Phys. Sci. 548, Kluwer, Dordrecht (2000) 51–82 MR
- [van Geemen 2023] **B van Geemen**, *Fourfolds of Weil type and the spinor map*, Expo. Math. 41:2 (2023) 418–447 MR
- [Grothendieck 1969] **A Grothendieck**, *Standard conjectures on algebraic cycles*, from “Algebraic geometry” (Bombay, 1968) (S S Abhyankar, editor), Tata Inst. Fundam. Res. Stud. Math. 4, Tata Inst. Fund. Res., Bombay (1969) 193–199 MR
- [Hassett and Tschinkel 2013] **B Hassett, Y Tschinkel**, *Hodge theory and Lagrangian planes on generalized Kummer fourfolds*, Mosc. Math. J. 13:1 (2013) 33–56, 189 MR
- [Huybrechts 1999] **D Huybrechts**, *Compact hyper-Kähler manifolds: basic results*, Invent. Math. 135:1 (1999) 63–113 MR
- [Huybrechts 2016] **D Huybrechts**, *Lectures on K3 surfaces*, Cambridge Studies in Advanced Mathematics 158, Cambridge Univ. Press (2016) MR
- [Huybrechts 2019] **D Huybrechts**, *Motives of isogenous K3 surfaces*, Comment. Math. Helv. 94:3 (2019) 445–458 MR
- [Huybrechts and Lehn 2010] **D Huybrechts, M Lehn**, *The geometry of moduli spaces of sheaves*, 2nd edition, Cambridge Univ. Press (2010) MR
- [Ingalls et al. 2022] **C Ingalls, A Logan, O Patashnick**, *Explicit coverings of families of elliptic surfaces by squares of curves*, Math. Z. 302:2 (2022) 1191–1238 MR
- [Jannsen 1992] **U Jannsen**, *Motives, numerical equivalence, and semi-simplicity*, Invent. Math. 107:3 (1992) 447–452 MR
- [Kaledin and Verbitsky 1998] **D Kaledin, M Verbitsky**, *Partial resolutions of Hilbert type, Dynkin diagrams, and generalized Kummer varieties*, preprint (1998) arXiv math/9812078
- [Kamenova et al. 2022] **L Kamenova, G Mongardi, A Oblomkov**, *Symplectic involutions of $K3^{[n]}$ type and Kummer n type manifolds*, Bull. Lond. Math. Soc. 54:3 (2022) 894–909 MR
- [Kleiman 1968] **SL Kleiman**, *Algebraic cycles and the Weil conjectures*, from “Dix exposés sur la cohomologie des schémas” (A Grothendieck, NH Kuiper, editors), Adv. Stud. Pure Math. 3, North-Holland, Amsterdam (1968) 359–386 MR
- [Kuga and Satake 1967] **M Kuga, I Satake**, *Abelian varieties attached to polarized K_3 -surfaces*, Math. Ann. 169 (1967) 239–242 MR

- [Markman 2023] **E Markman**, *The monodromy of generalized Kummer varieties and algebraic cycles on their intermediate Jacobians*, J. Eur. Math. Soc. (JEMS) 25:1 (2023) 231–321 MR
- [Milne 2021] **J S Milne**, *The Tate and standard conjectures for certain abelian varieties*, preprint (2021) arXiv 2112.12815
- [Mongardi and Wandel 2017] **G Mongardi, M Wandel**, *Automorphisms of O’Grady’s manifolds acting trivially on cohomology*, Algebr. Geom. 4:1 (2017) 104–119 MR
- [Moonen 1999] **B Moonen**, *Notes on Mumford–Tate groups*, unpublished lecture notes, Centre Emile Borel (1999) Available at <https://www.math.ru.nl/~bmoonen/Lecturenotes/CEBnotesMT.pdf>
- [Moonen and Zarhin 1995] **B J J Moonen, Y G Zarhin**, *Hodge classes and Tate classes on simple abelian fourfolds*, Duke Math. J. 77:3 (1995) 553–581 MR
- [Morrison 1984] **D R Morrison**, *On K3 surfaces with large Picard number*, Invent. Math. 75:1 (1984) 105–121 MR
- [Morrison 1985] **D R Morrison**, *The Kuga–Satake variety of an abelian surface*, J. Algebra 92:2 (1985) 454–476 MR
- [Mukai 1987] **S Mukai**, *On the moduli space of bundles on K3 surfaces, I*, from “Vector bundles on algebraic varieties” (Bombay, 1984) (MF Atiyah, editor), Tata Inst. Fund. Res. Stud. Math. 11, Tata Inst. Fund. Res., Bombay (1987) 341–413 MR
- [Nikulin 1975] **V V Nikulin**, *Kummer surfaces*, Izv. Akad. Nauk SSSR Ser. Mat. 39:2 (1975) 278–293, 471 MR In Russian; translated in Math. USSR-Izv. 9:2 (1976), 261–275
- [O’Grady 1997] **K G O’Grady**, *The weight-two Hodge structure of moduli spaces of sheaves on a K3 surface*, J. Algebraic Geom. 6:4 (1997) 599–644 MR
- [O’Grady 2021] **K G O’Grady**, *Compact tori associated to hyperkähler manifolds of Kummer type*, Int. Math. Res. Not. 2021:16 (2021) 12356–12419 MR
- [Oguiso 2020] **K Oguiso**, *No cohomologically trivial nontrivial automorphism of generalized Kummer manifolds*, Nagoya Math. J. 239 (2020) 110–122 MR
- [Paranjape 1988] **K Paranjape**, *Abelian varieties associated to certain K3 surfaces*, Compositio Math. 68:1 (1988) 11–22 MR
- [Varesco 2023] **M Varesco**, *Hodge similarities, algebraic classes, and Kuga–Satake varieties*, Math. Z. 305:4 (2023) art. id. 69 MR
- [Varesco 2025] **M Varesco**, *The Hodge conjecture for powers of K3 surfaces of Picard number 16*, Michigan Math. J. 75:5 (2025) 1085–1116 MR
- [Verbitsky 1995] **M Verbitsky**, *Tri-analytic subvarieties of hyper-Kähler manifolds*, Geom. Funct. Anal. 5:1 (1995) 92–104 MR
- [Voisin 2022] **C Voisin**, *Footnotes to papers of O’Grady and Markman*, Math. Z. 300:4 (2022) 3405–3416 MR
- [Yoshioka 2001] **K Yoshioka**, *Moduli spaces of stable sheaves on abelian surfaces*, Math. Ann. 321:4 (2001) 817–884 MR

SALVATORE FLOCCARI salvatore.floccari@hu-berlin.de

Fakultät für Mathematik, Universität Bielefeld, Bielefeld, Germany

Current address: *Institut für Mathematik, Humboldt-Universität zu Berlin, Berlin, Germany*

Proposed: Sándor J Kovács

Received: January 18, 2025

Seconded: Marc Levine, Dan Abramovich

Revised: July 12, 2025

Ulrich sheaves, the arithmetic writhe and algebraic isotopies of space curves

DANIELE AGOSTINI AND MARIO KUMMER

We establish a connection between the theory of Ulrich sheaves and $\mathbb{A}^1$ -homotopy theory. For instance, we prove that the $\mathbb{A}^1$ -degree of a morphism between projective varieties, that is relatively oriented by an Ulrich sheaf, is constant on the target even when it is not $\mathbb{A}^1$ -chain connected or $\mathbb{A}^1$ -connected. Further if an embedded projective variety is the support of a symmetric Ulrich sheaf of rank one, the $\mathbb{A}^1$ -degree of all its linear projections can be read off in an explicit way from the free resolution of the Ulrich sheaf. Finally, we construct an Ulrich sheaf on the secant variety of a curve and use this to define an arithmetic version of Viro’s encomplexed writhe for curves in $\mathbb{P}^3$. This can be considered to be an arithmetic analogue of a knot invariant. Namely, we define a notion of algebraic isotopy under which the arithmetic writhe is invariant. For rational curves of degree at most four in $\mathbb{P}^3$ we obtain a complete classification up to algebraic isotopies.

1. Introduction	1155
2. Notation and conventions	1160
3. Ulrich sheaves	1160
4. Ulrich sheaves and the $\mathbb{A}^1$ -degree	1161
5. Secant varieties of curves	1174
6. The arithmetic writhe	1184
7. Algebraic isotopies	1189
8. Isotopic embeddings of curves in 3-space	1195
Acknowledgements	1199
References	1199

1 Introduction

In the emerging field of $\mathbb{A}^1$ -enumerative geometry, building upon the $\mathbb{A}^1$ -homotopy theory developed by Morel and Voevodsky [1999], solutions to enumerative problems over a field K are counted in the Grothendieck–Witt group $\mathrm{GW}(K)$ of this field in a way that the result does not depend on the chosen instance of the enumerative problem. Typical examples of enumerative problems whose solutions can be “arithmetically enriched” in this way include among others the 27 lines on a cubic surface [Kass and Wickelgren 2021], Bézout’s theorem [McKean 2021] or Gromov–Witten invariants [Kass et al. 2023]. Many results of this type are achieved by studying arithmetic versions of topological invariants like the Euler characteristic [Levine 2019] or the Brouwer degree [Morel 2012; Kass and Wickelgren 2019]. Here we focus on the latter, namely the $\mathbb{A}^1$ -degree of a morphism of smooth schemes as considered

MSC2020: 14F42, 14H50, 14J60, 14N07.

© 2026 MSP (Mathematical Sciences Publishers). Distributed under the Creative Commons Attribution License 4.0 (CC BY). Open Access made possible by subscribing institutions via [Subscribe to Open](#).

in [Kass et al. 2023; Pauli and Wickelgren 2021]. This is defined for a finite and surjective morphism $f : X \rightarrow Y$ of smooth K -varieties starting from a *relative orientation* of f , this is a line bundle L on X , with an isomorphism

$$\psi : L \otimes L \rightarrow \mathcal{H}om(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y).$$

With this, we can compute the $\mathbb{A}^1$ -degree at a closed point $y \in Y$ as a sum of local degrees on the fiber, as in the classical case; see Section 4 for more details. When this is independent of the point $y \in Y$, one says that the $\mathbb{A}^1$ -degree of the map f is *well-defined*. This is not always the case. One important case, when this happens, is if Y is $\mathbb{A}^1$ -chain connected. Our first result is a different condition for the well-definedness of the $\mathbb{A}^1$ -degree in terms of *Ulrich sheaves*: a sheaf $\mathcal{F}$ on X is called *f -Ulrich* if its pushforward along f is trivial, i.e., $f_* \mathcal{F} \cong \mathcal{O}_Y^{\oplus N}$ for $N \in \mathbb{N}$.

Theorem A *If the relative orientation L is f -Ulrich and if $H^0(Y, \mathcal{O}_Y) = K$, then the $\mathbb{A}^1$ -degree of f is well-defined.*

We refer to Theorem 4.17 for a precise statement of the result and for the proof. In particular, we consider the case of standard Ulrich sheaves on an embedded irreducible projective variety $X \subset \mathbb{P}^n$ of dimension $\dim X = k$: these are sheaves $\mathcal{F}$ on X which are π -Ulrich for any finite linear projection $\pi : X \rightarrow \mathbb{P}^k$. Ulrich sheaves were introduced by Eisenbud and Schreyer [2003] and they satisfy a plethora of nice properties. In particular, there has been ample interest in the question of whether every closed subvariety of $\mathbb{P}^n$ carries an Ulrich sheaf [Eisenbud and Schreyer 2003; Beauville 2018].

One of the main features of Ulrich sheaves is to give a determinantal representation of the *Chow form* of X . Recall that a finite linear projection $\pi : X \rightarrow \mathbb{P}^k$ has the form $\pi = [s_0, \dots, s_k]$, where $s_0, \dots, s_k \in H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ are linearly independent, and such that the linear space $\Pi = \{s_0 = \dots = s_k = 0\} \subset \mathbb{P}^n$ does not intersect X . The $s_0, \dots, s_k$ satisfy these conditions if and only if $\mathcal{C}_X(s_0 \wedge \dots \wedge s_k) \neq 0$, for a homogeneous form on $\wedge^{k+1} H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ called the *Chow form* of the embedded variety $X \subset \mathbb{P}^n$. One of the main results of [Eisenbud and Schreyer 2003] is that if X carries an Ulrich sheaf of rank one, then there is a matrix Λ of linear forms on $\wedge^{k+1} H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ whose determinant gives the Chow form. Furthermore, there is a natural notion of *symmetry* for the Ulrich sheaf which guarantees that Λ can be taken to be symmetric. This matrix encodes the $\mathbb{A}^1$ -degrees of all finite linear projections of X :

Theorem B *Let $\mathcal{F}$ be a symmetric Ulrich sheaf of rank one on $X \subset \mathbb{P}^n$. Then this sheaf induces a relative orientation on each finite linear projection*

$$[s_0, \dots, s_k] : X \rightarrow \mathbb{P}^k$$

such that the resulting $\mathbb{A}^1$ -degree is well-defined. Furthermore, there is a symmetric matrix Λ of linear forms on $\wedge^{k+1} H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ whose determinant is the Chow form of X , and such that the class

$$[\Lambda(s_0 \wedge \dots \wedge s_k)] \in \mathrm{GW}(K)$$

is precisely the $\mathbb{A}^1$ -degree of the linear projection $[s_0, \dots, s_k] : X \rightarrow \mathbb{P}^k$ with respect to the above orientation.




Figure 1: A knot diagram of the trefoil knot.

We refer to Theorem 4.32 for a precise statement. We also note that the symmetric matrix Λ is explicitly computable from a resolution of the sheaf $\mathcal{F}$, for example via a computer algebra system such as Macaulay2 [1992] or OSCAR [2024], as already demonstrated in [Eisenbud and Schreyer 2003].

With these results in hand, we pass to the second main theme of the paper, an $\mathbb{A}^1$ -analogue of knot theory. Classical knots can be studied via *knot diagrams*. These are planar projections of the knot, which, for each point where two arcs cross, keeps track of the over- and under-passing arc; see, e.g., Figure 1. The *local writhe* of a crossing in a knot diagram is a number in $\{-1, +1\}$ and is defined as in Figure 2. Note that one has to choose an orientation of the knot but the local writhe numbers are independent of this choice. The *writhe* of a knot diagram is then the sum of the local writhe numbers of all its crossings. For example, the writhe of the knot diagram in Figure 1 is equal to $+3$.

While the writhe is unaffected by Reidemeister moves of type II and III, a Reidemeister move of type I changes the writhe by ± 1 . In particular, the writhe is only an invariant of the knot diagram but not of the knot itself. When the knot is the real part of an algebraic curve $C \subset \mathbb{P}^3$, this issue corresponds to the fact that some plane projections might have isolated real nodes; see Figure 3. Viro [2001] showed that the problem can be circumvented by also assigning a local writhe to *isolated* real nodes of a planar projection of the algebraic curve: he defined the *encomplexed writhe number* as the sum over all local writhe numbers, including those at isolated nodes, and proved that this is indeed independent of the projection center, hence it is an invariant of the real curve $C \subset \mathbb{P}^3$. This invariant played a major role in the recent breakthrough on real algebraic links by Mikhalkin and Orevkov [2019].

We generalize Viro's encomplexed writhe to any field K via Ulrich bundles on secant varieties. This is natural, since the nodes of a linear projection $\pi : C \rightarrow \mathbb{P}^2$ corresponds to the secant lines to C passing through the center of the projection. Let us be more precise here. If $\psi : C \hookrightarrow \mathbb{P}^3$ is a nondegenerate embedding and $L = \psi^* \mathcal{O}_C(1)$ the corresponding line bundle, then the embedding ψ is obtained from composing the embedding of the complete linear system $\varphi_L : C \hookrightarrow \mathbb{P}(H^0(C, L)^\vee) = \mathbb{P}^n$




Figure 2: Local writhe numbers.



Figure 3: When applying a Reidemeister move of type I to a real algebraic curve, a crossing of two arcs becomes an isolated node.

with a linear projection $[s_0, s_1, s_2, s_3] : \varphi_L(C) \rightarrow \mathbb{P}^3$. In order for ψ to be an embedding, the sections $s_0, \dots, s_3 \in H^0(C, L)$ must be linearly independent, and the space $\Pi = \{s_0 = \dots = s_3\}$ should not intersect the secant variety $\Sigma = \Sigma(C, L)$ of $\varphi_L(C)$ in $\mathbb{P}^n$. In this setting, the induced projection $[s_0, \dots, s_3] : \Sigma \rightarrow \mathbb{P}^3$ is a finite and surjective map. Assume now that the secant variety Σ is *identifiable* in the sense that every point in $\Sigma \setminus \varphi_L(C)$ is contained in a unique secant line. Then, by construction, the fiber of $[s_0, \dots, s_3] : \Sigma \rightarrow \mathbb{P}^3$ over a point $q \in \mathbb{P}^3$ correspond to secant lines to $\psi(C)$ passing through q . Thus, if we can find a suitable orientation, the $\mathbb{A}^1$ -degree of this map gives an arithmetic count of secant lines through q , or, equivalently, an arithmetic count of nodes of the projection of $\psi(C)$ from q . We can actually find a suitable orientation via an Ulrich sheaf:

Theorem C *Let C be a smooth K -curve and L be a line bundle such that the secant variety Σ of $\varphi_L : C \hookrightarrow \mathbb{P}^n$ is identifiable.¹ Let also α be a noneffective theta characteristic on C , i.e., a line bundle such that $\alpha \otimes \alpha \cong \omega_C$ and $h^0(C, \alpha) = 0$. Then there is a symmetric Ulrich sheaf of rank one $\mathcal{F}_\alpha$ on Σ , depending on α .*

With such a symmetric Ulrich sheaf we can then use Theorem B, and give the following definition:

Definition (arithmetic writhe) In the above setting, the *arithmetic writhe*

$$w(\psi(C), \alpha) \in \text{GW}(K)$$

of the embedded curve $\psi : C \hookrightarrow \mathbb{P}^3$ and of the theta characteristic α , is the $\mathbb{A}^1$ -degree of the map

$$[s_0, \dots, s_3] : \Sigma \rightarrow \mathbb{P}^3$$

with respect to the relative orientation induced by $\mathcal{F}_\alpha$. Here $[s_0, \dots, s_3]$ is such that $\psi = [s_0, \dots, s_3] \circ \varphi_L$.

By the previous discussion, the arithmetic writhe is an arithmetically enriched count of the nodes of plane projections $\psi(C) \rightarrow \mathbb{P}^2$. The fact that the arithmetic writhe is well-defined as an $\mathbb{A}^1$ -degree means precisely that this does not depend on the center of the plane projection. We also provide an explicit local description of the arithmetic writhe which for real curves agrees with the local description of Viro's encomplexed writhe. Hence, for real curves, the arithmetic writhe is equal to the encomplexed writhe.

Observe that the appearance of a noneffective theta characteristic α on C is not surprising: indeed these correspond on the one hand to symmetric rank one Ulrich sheaves on the curve itself [Eisenbud and Schreyer 2003, Theorem 4.3] and, on the other hand, if the curve is real, to classical orientations;

¹In the upcoming article [Agostini et al. ≥ 2026] it is shown that one can in fact replace the condition of Σ being identifiable by the weaker condition that φ_L is an embedding.

see Example 4.5. Over an arbitrary field, every smooth rational curve has a noneffective theta characteristic. We also remark that the condition of identifiability on the secant variety of $\varphi_L : C \hookrightarrow \mathbb{P}^n$ is rather natural and general. For example, Riemann–Roch shows that this happens as soon as the degree $\deg(C)$ is at least $2g(C) + 3$, where $g(C)$ is the genus of the curve. In particular, this is true for any smooth rational curve of degree at least 3. Another important case are canonical curves of gonality at least 5, such as any general real or complex curve in the moduli space $\mathcal{M}_g$, $g \geq 7$. We note that, in a similar spirit, Lemarié-Rieusset [2026] has defined an arithmetic linking degree for two embeddings of $\mathbb{A}^2 \setminus \{0\}$ to $\mathbb{A}^4 \setminus \{0\}$ using motivic Seifert classes. On the other hand, to our knowledge, the arithmetic writhe is the first *knot* invariant defined in the $\mathbb{A}^1$ -context.

In the last part of the paper, we define and study *algebraic isotopies* between embeddings of a smooth projective curve C in 3-space in analogy with topological isotopies: they are morphisms

$$I : \mathbb{A}^1 \times C \rightarrow \mathbb{P}^3, \quad (t, x) \mapsto I_t(x),$$

such that $I_t : C_{\kappa(t)} \hookrightarrow \mathbb{P}^3_{\kappa(t)}$ is an embedding for all $t \in \mathbb{A}^1$. Two embeddings $\psi, \psi' : C \hookrightarrow \mathbb{P}^3$ are algebraically isotopic if they can be connected by a chain of algebraic isotopies. We extend our definition of the arithmetic writhe $w(\psi(C), \alpha)$ for an embedded curve $\psi : C \hookrightarrow \mathbb{P}^3$ to a definition of the *arithmetic writhe* $w(\psi, \alpha)$ for the embedding ψ itself, and we show that this is invariant under algebraic isotopies, see Remark 8.3. We use it to study algebraic isotopy classes of rational curves of low degree in $\mathbb{P}^3$.

Theorem D *Let K be a field of $\text{char}(K) \neq 2$.*

- (1) *If $d \in \{1, 2\}$, then any two embeddings $\psi, \psi' : \mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ over K of degree d are algebraically isotopic.*
- (2) *The set of algebraic isotopy classes of embeddings $\mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ over K of degree 3 is in bijection to $K^\times / K^{\times 4}$.*
- (3) *Two embeddings $\psi, \psi' : \mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ over K of degree 4 are algebraically isotopic if and only if they have the same arithmetic writhe.*

In the case $K = \mathbb{R}$, Theorem D strengthens a result by Björklund [2011] who proved that the encomplexed writhe characterizes real rational curves of degree at most four up to *rigid isotopy*, a (potentially strictly) coarser equivalence relation than algebraic isotopy. While Björklund’s result holds true in degree five as well, we do not know whether this is the case for our Theorem D. It was also shown in [Björklund 2011] that the encomplexed writhe is not enough to distinguish rigid isotopy classes of embeddings $\mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ of degree at least six, hence the corresponding statement also fails for the arithmetic writhe and algebraic isotopies. However, Ulrich sheaves other than those produced in Theorem C can be used to give further isotopy invariants. In Theorem 5.18, we construct a higher-rank Ulrich sheaf on the secant variety of a rational normal curve. From this we obtain a new invariant for algebraic isotopy classes of rational curves of degree six, and we construct an example of two embeddings of the same arithmetic writhe for which the new invariant is different; see Example 8.8. We do not know whether this is enough to distinguish all isotopy classes, nor do we have a local description as for the arithmetic writhe.

The paper is organized as follows. In Section 3, we collect some facts on Ulrich sheaves, and in Section 4 we show how they are connected to the $\mathbb{A}^1$ -degree, proving Theorem A in Theorem 4.17 and Theorem B in both Theorem 4.28 and Theorem 4.32. In Section 5, we turn our attention to secants of curves and on Ulrich sheaves on them: we prove Theorem C in Theorem 5.16 and we also construct, in Theorem 5.18, higher-rank Ulrich sheaves on secants of rational normal curves. In Section 6, we define the arithmetic writhe for an embedded curve, and we also show how to compute it via a sum of explicit local writhes; see Remark 6.2. Finally, in Section 7 we collect some general results about isotopies for projective embeddings, which imply the first two parts of Theorem D. In Section 8 we focus on the case of curves in 3-space. We define the arithmetic writhe of an embedding in $\mathbb{P}^3$ and we prove that it is invariant under algebraic isotopies; see Remark 8.3. We then prove Theorem D in Theorem 8.4 and Corollary 8.5, and we conclude with some remarks on rational curves of higher degree.

2 Notation and conventions

Let K always denote a field. By a K -variety we mean an integral, separated scheme of finite type over K . We will use the word curve to denote a projective, nonsingular K -variety of dimension 1. If X is a scheme and $x \in X$, then we denote by $\kappa(x)$ the residue field of X at x . If A, B are two divisors on a smooth variety X , we write $A \sim B$ to denote that A is linearly equivalent to B .

3 Ulrich sheaves

We first recall the basic properties of Ulrich sheaves as introduced in [Eisenbud and Schreyer 2003]. Consider the projective space $\mathbb{P}^n$ and the corresponding homogeneous coordinate ring $S = K[x_0, \dots, x_n]$. Let $\mathcal{F}$ be a coherent sheaf on $\mathbb{P}^n$ with scheme-theoretic support $\iota : X \hookrightarrow \mathbb{P}^n$ of pure dimension $k > 0$ and codimension $c = n - k$. Let also $\Gamma_*(\mathcal{F}) = \bigoplus H^0(\mathcal{F}(q))$ be the module of twisted global sections, seen as a graded S -module. The following important equivalence was proven in [Eisenbud and Schreyer 2003, Theorem 2.1]:

Definition 3.1 (Ulrich sheaf) The sheaf $\mathcal{F}$ is called an *Ulrich sheaf* if it satisfies one of the following equivalent conditions:

- (i) $\Gamma_*(\mathcal{F})$ has a linear minimal free resolution as an S -module:

$$F : \quad 0 \rightarrow F_c \xrightarrow{\varphi_c} F_{c-1} \xrightarrow{\varphi_{c-1}} \dots \xrightarrow{\varphi_2} F_1 \xrightarrow{\varphi_1} F_0 \rightarrow \mathcal{F} \rightarrow 0,$$

where F_i is a direct sum of copies of $S(-i)$.

- (ii) $H^i(X, \mathcal{F}(-i)) = 0$ for $i > 0$ and $H^i(X, \mathcal{F}(-i-1)) = 0$ for $i < k$.
- (iii) If $\pi : X \rightarrow \mathbb{P}^k$ is a finite surjective linear projection, then $\pi_* \iota^* \mathcal{F} \cong \mathcal{O}_{\mathbb{P}^k}^{\oplus t}$ for a certain $t > 0$.

Remark 3.2 It follows from the Auslander–Buchsbaum formula and part (i) of Definition 3.1 that an Ulrich sheaf is Cohen–Macaulay.

The resolution of $\mathcal{F}$ can be used to compute a *Chow form* of X . Recall that this is a polynomial in the Plücker coordinates of the Grassmannian $\mathbb{G}(c-1, n)$ that cuts out the locus of $(c-1)$ -planes that intersect X . In the coordinate ring of $\mathbb{G}(c-1, n)$ it is unique up to a scalar factor. A power of the Chow form can be written as the determinant of a matrix with entries linear forms in the Plücker coordinates [Eisenbud and Schreyer 2003, §3]. This is constructed from the resolution $\mathbf{F}$ as follows: after choosing a basis of each F_i the maps φ_i are given by matrices A_i whose entries are linear forms in the variables $x_0, \dots, x_n$. We consider these linear forms as degree one elements of the tensor algebra $T(K^{n+1})^\vee$ and define $\gamma(\mathbf{F})$ to be the product $A_1 \cdots A_c$ over $T(K^{n+1})^\vee$. The entries of $\gamma(\mathbf{F})$ are multilinear forms on K^{n+1} and since $\varphi_i \circ \varphi_{i+1} = 0$ these multilinear forms are alternating. Therefore, the entries of $\gamma(\mathbf{F})$ are elements of $\wedge^c(K^{n+1})^\vee$ and thus linear in the Plücker coordinates of $\mathbb{G}(c-1, n)$. Evaluating $\gamma(\mathbf{F})$ at a $(c-1)$ -plane $V \subset \mathbb{P}^n$ gives a singular matrix if and only if V intersects X . This implies that the determinant of $\gamma(\mathbf{F})$ is the $\text{rank}(\mathcal{F})$ -th power of the Chow form of X . It follows from the definition of the matrix $\gamma(\mathbf{F})$ that it actually depends only on the choice of bases of F_0 and F_c . Thus $\gamma(\mathbf{F})$ is uniquely determined by $\mathcal{F}$ up to multiplication from left and right by invertible matrices over K . Finally, if there exists a symmetric isomorphism $\mathcal{F} \rightarrow \mathcal{E} \text{xt}^c(\mathcal{F}, \mathcal{O}_{\mathbb{P}^n})(-c)$, then by choosing a suitable basis the determinantal representation can be made symmetric [Eisenbud and Schreyer 2003, §3.1]. We will come back to the symmetric case later.

Example 3.3 We consider the usual twisted cubic curve $\iota: \mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ given by $\iota(s:t) = (s^3 : s^2t : st^2 : t^3)$. By Definition 3.1 and a direct computation of cohomology the sheaf $\mathcal{F} = \iota_*(\mathcal{O}_{\mathbb{P}^1}(2))$ is Ulrich. A quick computation using the computer algebra system [Macaulay2] gives us the minimal resolution

$$\mathbf{F}: 0 \rightarrow S(-2)^3 \xrightarrow{\varphi_2} S(-1)^6 \xrightarrow{\varphi_1} S^3 \rightarrow \Gamma_*(\mathcal{F}) \rightarrow 0,$$

where the maps φ_1 and φ_2 are given by the matrices

$$A_1 = \begin{pmatrix} -x_1 & -x_2 & -x_2 & -x_3 & -x_3 & 0 \\ x_0 & 0 & x_1 & 0 & x_2 & -x_3 \\ 0 & x_0 & 0 & x_1 & 0 & x_2 \end{pmatrix} \quad \text{and} \quad A_2 = \begin{pmatrix} -x_2 & x_1 & 0 & -x_0 & x_0 & 0 \\ -x_3 & x_2 & -x_2 & 0 & x_1 & -x_0 \\ 0 & 0 & x_3 & -x_2 & 0 & x_1 \end{pmatrix}^t,$$

and one calculates that

$$\gamma(\mathbf{F}) = \begin{pmatrix} -x_{12} & -x_{13} & x_{23} \\ x_{02} & x_{12} + x_{03} & -x_{13} \\ -x_{01} & -x_{02} & x_{12} \end{pmatrix},$$

where the x_{ij} are the usual Plücker coordinates on $\mathbb{G}(1, 3)$.

4 Ulrich sheaves and the $\mathbb{A}^1$ -degree

In this section we discuss a connection of Ulrich sheaves and a recently developed notion of degree in the context of $\mathbb{A}^1$ -enumerative geometry.

4.1 Orientations and $\mathbb{A}^1$ -degree

In this subsection we recall some preliminaries from $\mathbb{A}^1$ -enumerative geometry, mostly following the exposition of [Pauli and Wickelgren 2021, §8]. Let K always denote a field.

Definition 4.1 (algebraic orientation) Let X be a nonsingular K -variety. An *algebraic orientation* of X is an isomorphism $L \otimes L \rightarrow \omega_X$, where L is a line bundle on X and ω_X the canonical sheaf. We also say that X is *oriented* by $L \otimes L \rightarrow \omega_X$. A *theta characteristic* of X is a line bundle L such that $L \otimes L$ is isomorphic to ω_X .

Example 4.2 If $X = \mathbb{P}^n$ with $n = 2h - 1$ odd, then $\mathcal{O}_{\mathbb{P}^n}(-h)$ is a theta characteristic of X . If n is even, then $\mathbb{P}^n$ does not have a theta characteristic.

Definition 4.3 Let X be a nonsingular K -variety and L a theta characteristic. Two algebraic orientations $\psi_1, \psi_2 : L \otimes L \rightarrow \omega_X$ are *equivalent* if $\psi_1 \circ \psi_2^{-1}$ is multiplication by a square of a global section of $\mathcal{O}_X$.

Remark 4.4 Let $K = \mathbb{R}$ and X be a nonsingular $\mathbb{R}$ -variety of dimension n . Let $\psi : L \otimes L \rightarrow \omega_X$ be an algebraic orientation. Let f some rational section of L and $\omega = \psi(f \otimes f)$. All zeros and poles of the rational differential n -form ω on X are of even order. It can be shown that this implies the existence of a rational function g on X which is nonnegative wherever it is defined on $X(\mathbb{R})$ and such that $g \cdot \omega$ does not have any zeros or poles on $X(\mathbb{R})$. Thus $g \cdot \omega$ defines an orientation on $X(\mathbb{R})$ in the classical topological sense. It is straightforward to check that this does not depend on the choices we made and that two algebraic orientations that are equivalent in the sense of Definition 4.3 induce the same classical orientation on $X(\mathbb{R})$.

Example 4.5 Let X be a curve of genus g over $\mathbb{R}$. Then every topological orientation of $X(\mathbb{R})$ is induced by an algebraic orientation $L \otimes L \rightarrow \omega_X$. This follows for instance from [Geyer 1977, Satz 2.4.c] or by the same argument as in [Kummer 2019, Corollary 2.4]. However, this is not true in general. For instance, let X be a real K3 surface: since the Picard group of X is torsion-free, there is, up to a scalar multiple, only one algebraic orientation on X . Thus only two topological orientations on $X(\mathbb{R})$ arise in this way. However, if $X(\mathbb{R})$ is not connected, then there exist at least four different topological orientations of $X(\mathbb{R})$.

There is a relative version of algebraic orientations for morphisms [Pauli and Wickelgren 2021, Definition 7]:

Definition 4.6 (relative orientation) Let $f : X \rightarrow Y$ be a finite surjective morphism of nonsingular K -varieties. A *relative orientation* of f is an isomorphism

$$\psi : L \otimes L \rightarrow \mathcal{H}\text{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y),$$

where L is a line bundle on X , and $\mathcal{T}_X$ and $\mathcal{T}_Y$ are the tangent bundles on X and Y , respectively. Two relative orientations ψ_1, ψ_2 as above are *equivalent* if $\psi_1 \circ \psi_2^{-1}$ is multiplication by a square of a global section of $\mathcal{O}_X$.

Remark 4.7 If $f : X \rightarrow Y$ is a finite surjective morphism of nonsingular K -varieties with algebraic orientations $\psi_1 : L_1 \otimes L_1 \rightarrow \omega_X$ and $\psi_2 : L_2 \otimes L_2 \rightarrow \omega_Y$, then a relative orientation of f is given by the induced isomorphism

$$\psi : (L_1 \otimes f^* L_2^\vee)^{\otimes 2} \rightarrow \omega_X \otimes f^* \omega_Y^\vee.$$

Algebraic orientations ψ'_1 and ψ'_2 on X and Y that are equivalent to ψ_1 and ψ_2 induce a relative orientation ψ' for f that is equivalent to ψ .

Recall the set of equivalence classes of nondegenerate symmetric bilinear forms on finite-dimensional vector spaces over K form a monoid via the orthogonal sum. The Grothendieck–Witt group $\mathrm{GW}(K)$ of K is the Grothendieck group of this monoid. It is generated by all equivalence classes $\langle a \rangle$ of one-dimensional bilinear forms

$$K \times K \rightarrow K, \quad (x, y) \mapsto axy,$$

for $a \in K^\times$. Note that $\langle a \rangle = \langle ab^2 \rangle$ for all $b \in K^\times$.

For a finite morphism $f : X \rightarrow Y$ of nonsingular oriented K -varieties of the same dimension one can define a notion of degree that takes its values in $\mathrm{GW}(K)$ under appropriate hypotheses. Actually it suffices to make the weaker assumption of requiring f to have a relative orientation $\psi : L \otimes L \rightarrow \mathcal{H}\mathrm{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y)$. In this case, let $x \in X$ be a closed point where f is not ramified and set $y = f(x) \in Y$. Note that this in particular requires the field extension $\kappa(x)/\kappa(y)$ to be separable. The differential of f defines a morphism $T_x f : \mathcal{T}_X \rightarrow f^* \mathcal{T}_Y$ and thus $\kappa(x)$ -linear maps

$$T_x f : T_x X \rightarrow T_y Y \otimes_{\kappa(y)} \kappa(x) \quad \text{and} \quad J_x f = \det T_x f : \det T_x X \rightarrow \det T_y Y \otimes_{\kappa(y)} \kappa(x).$$

Two bases of $T_x X$ and $T_y Y$ are called compatible with respect to the relative orientation if the determinant of the linear map that sends one base to the other is an element in the fiber of $\mathcal{H}\mathrm{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y)$ at x which is the image of a square under ψ [Pauli and Wickelgren 2021, Definition 8]. After choosing such compatible bases, we can identify the determinant $J_x f$ with an element in $\kappa(x)$. Thanks to the requirement of compatibility, it is straightforward to check that the class $\langle J_x f \rangle \in \mathrm{GW}(\kappa(x))$ does not depend on the bases. We can also define $\langle J_x f \rangle$ without mentioning bases as follows: taking the fiber of the orientation at x we see that $J_x f = g \cdot \psi(t \otimes t)$ for a certain $t \in L_x$ and $g \in \kappa(x)$. Then $\langle J_x f \rangle = \langle g \rangle$.

Then the *local $\mathbb{A}^1$ -degree* $\deg_x^{\mathbb{A}^1}(f)$ of f at x as the class in $\mathrm{GW}(\kappa(y))$ of the symmetric K -bilinear form defined by

$$(4-1) \quad \mathrm{Tr}_{\kappa(x)/\kappa(y)} \langle J_x f \rangle : \kappa(x) \times \kappa(x) \rightarrow \kappa(y), \quad (a, b) \mapsto \mathrm{Tr}_{\kappa(x)/\kappa(y)}(J_x ab).$$

Finally, for a closed point $y \in Y$ outside the branch locus of f , one defines the *$\mathbb{A}^1$ -degree* $\deg_y^{\mathbb{A}^1}(f)$ of f at y as

$$\deg_y^{\mathbb{A}^1}(f) = \sum_{x \in f^{-1}(y)} \deg_x^{\mathbb{A}^1}(f) \in \mathrm{GW}(\kappa(y)).$$

Remark 4.8 One could argue that equation (4-1) is rather a *formula* for the local $\mathbb{A}^1$ degree [Morel 2006; Kass and Wickelgren 2021] than a definition.

Remark 4.9 The (local) $\mathbb{A}^1$ -degree only depends on the equivalence class of the relative orientation.

Remark 4.10 In the case $K = \mathbb{R}$, the signature of the $\mathbb{A}^1$ -degree of $f : X \rightarrow Y$ is the topological degree of the restriction of f to the real parts of X and Y . This is clear from the description of the local degree in equation (4-1) together with the observation that every scaled trace bilinear form $\mathbb{C} \times \mathbb{C} \rightarrow \mathbb{R}$ has signature zero.

Definition 4.11 Let $f : X \rightarrow Y$ be a relatively oriented finite surjective morphism of nonsingular K -varieties. If there is an element $d \in \mathrm{GW}(K)$ such that for every closed point $y \in Y$ outside the branch locus we have that

$$\deg_y^{\mathbb{A}^1}(f) = d \otimes_K \kappa(y) \in \mathrm{GW}(\kappa(y)),$$

then we say that $\deg^{\mathbb{A}^1}(f)$ is *well-defined* and we write $\deg^{\mathbb{A}^1}(f) = d$.

Remark 4.12 In the situation of Definition 4.11, writing $\deg^{\mathbb{A}^1}(f) = d$ is a slight abuse of notation since there might be several different $d \in \mathrm{GW}(K)$ satisfying the requirements from Definition 4.11 — for instance, in the case that Y has no K -rational points.

The degree of a morphism is not always well-defined. An important case where this happens is when the base of the morphism is $\mathbb{A}^1$ chain-connected:

Example 4.13 Let $f : X \rightarrow Y$ be as in Definition 4.11.

- (1) If Y is $\mathbb{A}^1$ -chain connected in the sense of [Pauli and Wickelgren 2021, Definition 9], then $\deg^{\mathbb{A}^1}(f)$ is well-defined by [Pauli and Wickelgren 2021, Theorem 8]. This remains true under the potentially weaker hypothesis that Y is $\mathbb{A}^1$ -connected [Kass et al. 2023, §2.5].
- (2) If $K = \mathbb{R}$, then it suffices that $Y(\mathbb{R})$ is connected in the classical topology because the (signature of the) $\mathbb{A}^1$ -degree is the Brouwer degree of the map $X(\mathbb{R}) \rightarrow Y(\mathbb{R})$.
- (3) In general $\deg^{\mathbb{A}^1}(f)$ is not necessarily well-defined; see for example [Pauli and Wickelgren 2021, Example 16].

4.2 Relative Ulrich sheaves and the $\mathbb{A}^1$ -degree

We now derive a further case when $\deg^{\mathbb{A}^1}(f)$ is well-defined, which is a condition on the relative orientation rather than on the target. Motivated by Definition 3.1(iii), one makes the following definition.

Definition 4.14 Let $f : X \rightarrow Y$ a morphism of schemes. A coherent sheaf $\mathcal{F}$ on X is called *f-Ulrich* if there exists an integer $t > 0$ such that $f_*\mathcal{F} \cong \mathcal{O}_Y^{\oplus t}$.

We further recall the following construction of a trace map for differential forms.

Construction 4.15 [Hartshorne 1977, Exercise III.7.2] Let $f : X \rightarrow Y$ be a finite surjective morphism of nonsingular varieties over K . We recall the construction of a natural trace map $t : f_*\omega_X \rightarrow \omega_Y$. First let $X = \mathrm{Spec}(B)$ and $Y = \mathrm{Spec}(A)$ be affine, with $\dim(X) = \dim(Y) = n$, and let $E = \mathrm{Quot}(A)$ and $F = \mathrm{Quot}(B)$ be the fields of fractions. Choose any nonzero $\omega \in \bigwedge^n \Omega_{A/K}$. Then, for any $\omega' \in \bigwedge^n \Omega_{B/K}$,

there is a $b \in F$ such that $\omega' = b \cdot \omega$. The trace map sends ω' to $t(\omega') = \text{tr}_{F/E}(b) \cdot \omega$. Note that while b is not necessarily in B , we always have that $t(\omega') \in \bigwedge^n \Omega_{A/K}$. In the general, not necessarily affine case, the map t is obtained by gluing.

Now let $f : X \rightarrow Y$ be a finite surjective morphism of nonsingular K -varieties, relatively oriented by the isomorphism

$$\psi : L \otimes L \rightarrow \mathcal{H}\text{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y) = \omega_X \otimes f^* \omega_Y^{-1}.$$

By the projection formula the pushforward of ψ induces a map

$$f_* L \otimes f_* L \rightarrow f_* \omega_X \otimes \omega_Y^{-1}.$$

Tensoring the trace morphism from Construction 4.15 by ω_Y^{-1} , we obtain a morphism $f_* \omega_X \otimes \omega_Y^{-1} \rightarrow \mathcal{O}_Y$, which we precompose with the map above to obtain a symmetric $\mathcal{O}_Y$ -bilinear form $\tilde{\psi} : f_* L \otimes f_* L \rightarrow \mathcal{O}_Y$. The following simple lemma can also be seen as a consequence of [Kass et al. 2023, Corollary 3.10], but we give here a self-contained proof.

Lemma 4.16 *Let $f : X \rightarrow Y$ be a finite surjective morphism of nonsingular K -varieties, relatively oriented by the isomorphism*

$$\psi : L \otimes L \rightarrow \mathcal{H}\text{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y).$$

Let $y \in Y$ be a closed point outside the branch locus. The class of the fiber of $\tilde{\psi}$ at y in $\text{GW}(\kappa(y))$ is $\deg_y^{\mathbb{A}^1}(f)$.

Proof Let $U \subset Y$ an open affine neighborhood of y and $V = f^{-1}(U)$. Then we have a finite ring extension $A \subset B$, where $A = \mathcal{O}_Y(U)$ and $B = \mathcal{O}_X(V)$. If we choose U sufficiently small, then $\Omega_{A/K}$ is a free A -module and $\Omega_{B/K}$ and $M = L(V)$ are free B -modules. Since y is not in the branch locus of f we can, after further shrinking U if necessary, assume moreover that $\Omega_{B/K} = \Omega_{A/K} \otimes_A B$. Then there are $a_1, \dots, a_n \in A$ such that $da_1, \dots, da_n$ is an A -basis of $\Omega_{A/K}$ and $da_1 \otimes 1, \dots, da_n \otimes 1$ is a B -basis of $\Omega_{B/K}$. Let $v_1, \dots, v_n \in \Omega_{A/K}^\vee$ and $w_1, \dots, w_n \in \Omega_{B/K}^\vee$ be the dual bases so that $T_x f(w_i) = v_i$ for all $x \in V$, and let $t \in M$ be a generator of M . Then

$$\psi(t \otimes t) = g \cdot ((da_1 \otimes 1) \wedge \dots \wedge (da_n \otimes 1)) \otimes (v_1 \wedge \dots \wedge v_n)$$

for some $g \in B^\times$. Then it follows that $\langle J_{x_i} f \rangle = \langle g(x_i) \rangle$. For $s_1, s_2 \in M$ we can write $s_i = h_i \cdot t$ for $i = 1, 2$ and $h_i \in B$. By the definition of $\tilde{\psi}$ we then have

$$(4-2) \quad \tilde{\psi}(s_1, s_2) = \text{Tr}_{B/A}(g \cdot h_1 \cdot h_2).$$

Since $B \otimes_A \kappa(y) = \kappa(x_1) \oplus \dots \oplus \kappa(x_r)$, equation (4-2) evaluated at y equals

$$\sum_{i=1}^r \text{Tr}_{\kappa(x_i)/\kappa(y)}(g(x_i) \cdot h_1(x_i) \cdot h_2(x_i)).$$

Therefore, the class of the fiber of $\tilde{\psi}$ at y equals

$$\sum_{i=1}^r \mathrm{Tr}_{\kappa(x_i)/\kappa(y)} \langle g(x_i) \rangle = \sum_{i=1}^r \mathrm{Tr}_{\kappa(x_i)/\kappa(y)} \langle J_{x_i} f \rangle = \sum_{i=1}^r \deg_{x_i}^{\mathbb{A}^1}(f) = \deg_y^{\mathbb{A}^1}(f),$$

where we use that $\langle g(x_i) \rangle = \langle J_{x_i} f \rangle$. $\square$

Passing from $\tilde{\psi} : f_* L \otimes f_* L \rightarrow \mathcal{O}_Y$ to global sections, we obtain a symmetric bilinear form

$$\bar{\psi} : H^0(Y, f_* L) \times H^0(Y, f_* L) \rightarrow H^0(Y, \mathcal{O}_Y).$$

In the case that $H^0(Y, \mathcal{O}_Y) = K$, this is a symmetric K -valued bilinear form.

Theorem 4.17 *Let $f : X \rightarrow Y$ be a finite surjective morphism of nonsingular K -varieties, relatively oriented by the isomorphism*

$$\psi : L \otimes L \rightarrow \mathcal{H}\mathrm{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y).$$

If L is f -Ulrich and $H^0(Y, \mathcal{O}_Y) = K$, then the K -bilinear form

$$\bar{\psi} : H^0(Y, f_* L) \times H^0(Y, f_* L) \rightarrow K$$

is nondegenerate. In this case, $\deg^{\mathbb{A}^1}(f)$ is well-defined and is equal to the class of $\bar{\psi}$.

Proof At every $y \in Y$, the fiber of a K -basis of $H^0(Y, f_* L)$ is a $\kappa(y)$ basis of the fiber of $f_* L$ at y because L is f -Ulrich. Therefore, the fiber of $\tilde{\psi}$ at y is isometric to $\bar{\psi} \otimes_K \kappa(y)$. Now the claim follows from Lemma 4.16. $\square$

Example 4.18 Let X be a smooth projective curve and $f : X \rightarrow \mathbb{P}^1$ a finite surjective morphism. Every relative orientation of f is given by a line bundle of the form $L = M(1)$, where M is a theta characteristic on X . Then L is f -Ulrich if and only if $h^0(M) = 0$. In particular, for $X = \mathbb{P}^1$ every relative orientation is given by an Ulrich line bundle.

One has the following converse of Lemma 4.16.

Lemma 4.19 *Let $f : X \rightarrow Y$ be a finite surjective morphism of nonsingular complete K -varieties, relatively oriented by the isomorphism*

$$\psi : L \otimes L \rightarrow \mathcal{H}\mathrm{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y),$$

and assume that Y is geometrically irreducible. The following are equivalent:

- (1) L is f -Ulrich.
- (2) The bilinear form $\bar{\psi}$ is nondegenerate and, for every closed point $y \in Y$, one has

$$\deg_y^{\mathbb{A}^1}(f) = \bar{\psi} \otimes_K \kappa(y) \in \mathrm{GW}(\kappa(y)).$$

Proof Theorem 4.17 proves that (1) implies (2). On the other hand, condition (2) implies that $\bar{\psi}$ is nondegenerate and for $y \in Y$,

$$\dim(H^0(X, L)) = \text{rank}(\bar{\psi}) = \text{rank}(\deg_y^{\mathbb{A}^1}(f)) = \deg(f).$$

Therefore, by [Hanselka and Kummer 2024, Theorem 4.8], the line bundle L is f -Ulrich. $\square$

We conclude this section with two examples showing that the two sufficient conditions for the $\mathbb{A}^1$ -degree being well-defined that we have seen here do not imply each other.

Example 4.20 Consider a smooth plane cubic curve $X \subset \mathbb{P}^2$ and let $f : X \rightarrow \mathbb{P}^1$ be the linear projection from a point not on X . Then f is relatively oriented by

$$\psi : \mathcal{O}_X(1) \otimes \mathcal{O}_X(1) \rightarrow \mathcal{O}_X(2) \cong \mathcal{H}\text{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_{\mathbb{P}^1}).$$

However, the line bundle $\mathcal{O}_X(1)$ is not f -Ulrich because it violates part (ii) of Definition 3.1. Since $\dim(H^0(X, \mathcal{O}_X(1))) = 3 = \deg(f)$, this implies in particular that $\bar{\psi}$ cannot be a nondegenerate bilinear form. On the other hand, the target $\mathbb{P}^1$ is clearly $\mathbb{A}^1$ -chain connected.

Example 4.21 Assume here that the characteristic of K is not 2 or 3, and consider the smooth sextic curve X in $\mathbb{P}^3$ defined over $\mathbb{Q}$ by the two equations

$$Q = x_1^2 + x_2^2 - x_3^2 \quad \text{and} \quad T = x_0^3 + x_1^3 + x_2^3,$$

and the smooth cubic curve Y in $\mathbb{P}^2$ defined by T . The linear projection from the point $[0, 0, 0, 1]$ defines a finite surjective morphism $f : X \rightarrow Y$ of degree two. The ramification divisor of f is the zero divisor of x_3 on X . This shows that

$$\mathcal{H}\text{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y) \cong \mathcal{O}_X(1).$$

Furthermore, the zero divisor of $x_2 - x_3$ on X is of the form $2D$, where D is an effective divisor of degree three, because this hyperplane intersects the singular quadric in $\mathbb{P}^3$ defined by Q in a line with multiplicity two. Thus the corresponding line bundle L satisfies $L \otimes L \cong \mathcal{O}_X(1)$. In particular, we obtain a relative orientation

$$L \otimes L \rightarrow \mathcal{H}\text{om}(\det \mathcal{T}_X, f^* \det \mathcal{T}_Y)$$

of f . A calculation with the computer algebra system [Macaulay2] further shows that $f_* L = \mathcal{O}_Y^{\oplus 2}$. Thus $\deg^{\mathbb{A}^1}(f)$ is well-defined by Theorem 4.17. On the other hand, by Lüroth's theorem there is no nonconstant morphism $\mathbb{A}^1 \rightarrow Y$ and therefore Y is not $\mathbb{A}^1$ -chain connected (actually not even $\mathbb{A}^1$ -connected [Asok and Morel 2011, Corollary 2.4.4]). Let us compute the degree of this map using the method of Theorem 4.17. The quadratic form $\bar{\psi}$ is given by the composition

$$H^0(X, \mathcal{O}_X(D)) \otimes H^0(X, \mathcal{O}_X(D)) \rightarrow H^0(X, \mathcal{O}_X(2D)) \xrightarrow{\sim} H^0(X, \mathcal{O}_X(R)) \rightarrow H^0(Y, \mathcal{O}_Y),$$

where the first map is the multiplication map, the second map is given by the linear equivalence between $2D$ and the ramification divisor $R = \{x_3 = 0\}$ and the last map is the trace. A basis of $H^0(X, \mathcal{O}_X(D))$

is given by the rational functions $1, x_1/(x_2 - x_3)$ and multiplying them together we obtain the rational functions $1, x_1/(x_2 - x_3), (x_1/(x_2 - x_3))^2$. These get mapped to the elements $(x_2 - x_3)/x_3, x_1/x_3, x_1^2/((x_2 - x_3)x_3)$ in $H^0(X, \mathcal{O}_X(R))$, and finally the traces of these are

$$\begin{aligned}\mathrm{Tr}\left(\frac{x_2 - x_3}{x_3}\right) &= \frac{x_2 - x_3}{x_3} - \frac{x_2 + x_3}{x_3} = \frac{-2x_3}{x_3} = -2, \\ \mathrm{Tr}\left(\frac{x_1}{x_3}\right) &= \frac{x_1}{x_3} - \frac{x_1}{x_3} = 0, \\ \mathrm{Tr}\left(\frac{x_1^2}{(x_2 - x_3)x_3}\right) &= \frac{x_1^2}{(x_2 - x_3)x_3} - \frac{x_1^2}{(x_2 + x_3)x_3} = \frac{2x_1^2}{x_2^2 - x_3^2} = -2,\end{aligned}$$

where the last equality comes from the identity $x_1^2 = x_3^2 - x_2^2$ on X . In conclusion, the quadratic form is given by $(\lambda, \mu) \mapsto -2\lambda^2 - 2\mu^2$ so that the $\mathbb{A}^1$ -degree is $\langle -2 \rangle + \langle -2 \rangle$.

4.3 Symmetric Ulrich sheaves

Let us first set up some notation. We let $S = \Gamma_*(\mathcal{O}_{\mathbb{P}^n}) = \bigoplus_{j \in \mathbb{Z}} \Gamma(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(j))$. If M is a graded S -module, we denote by $\tilde{M}$ the sheaf associated to M on $\mathbb{P}^n$. If $\varphi : M \rightarrow N$ is a homomorphism of graded S -modules, then we denote by $\tilde{\varphi} : \tilde{M} \rightarrow \tilde{N}$ the associated morphism of quasicoherent sheaves on $\mathbb{P}^n$. One has

$$\mathcal{F} = \widetilde{\Gamma_*(\mathcal{F})}$$

for any quasicoherent sheaf $\mathcal{F}$ on $\mathbb{P}^n$; see [Hartshorne 1977, Proposition II.5.15]. Conversely, if M is a finitely generated graded S -module of $\mathrm{depth}(M) \geq 2$, one also has $M = \Gamma_*(\tilde{M})$; see [Eisenbud 2005, Corollary A.1.13]. This applies in particular when M is Cohen–Macaulay of dimension at least two.

Now let $\mathcal{F}$ be an Ulrich sheaf on $\mathbb{P}^n$ whose support is a closed subvariety $\iota : X \hookrightarrow \mathbb{P}^n$ of dimension at least one. Let us denote the rank of $\iota^*\mathcal{F}$ by r and by $M = \Gamma_*(\mathcal{F})$ the module of twisted global sections over S . As explained in Section 3, we can construct from $\mathcal{F}$ a matrix $\gamma(\mathbf{F})$ obtained from the free resolution $\mathbf{F}$ of M whose entries are linear forms in Plücker coordinates and whose determinant is the r -th power of the Chow form of X . We now recall from [Eisenbud and Schreyer 2003, §3.1] a condition for this matrix to be symmetric. Consider the contravariant functor

$$(4-3) \quad D : \mathcal{G} \mapsto \mathcal{E}xt^c(\mathcal{G}, \mathcal{O}_{\mathbb{P}^n}(-c)),$$

where c is the codimension of X . If $\mathcal{F}$ is Ulrich supported on X , the sheaf $D(\mathcal{F})$ is again an Ulrich sheaf with support X and there is a canonical isomorphism $\beta : \mathcal{F} \rightarrow DD\mathcal{F}$. A morphism $\sigma : \mathcal{F} \rightarrow D\mathcal{F}$ is *symmetric* if $\sigma = D(\sigma) \circ \beta$. If such a morphism exists, then a suitable choice of bases makes the matrix $\gamma(\mathbf{F})$ symmetric; see [Eisenbud and Schreyer 2003, §3.1].

We will now recall a more explicit description of a minimal free resolution of M . To this end recall the following definition.

Definition 4.22 Let $A = (A_1, \dots, A_s)$ be a tuple of pairwise commuting $m \times m$ matrices over a ring R . We can define on R^m the structure of an $R[t_1, \dots, t_s]$ -module P by letting t_i act on R^m via multiplication with the matrix A_i from the left. Consider the complex $P \otimes \mathbf{K}(t)$, where $\mathbf{K}(t)$ is the Koszul complex of the sequence $t = (t_1, \dots, t_s)$. We can view this complex as a complex of R -modules instead of $R[t_1, \dots, t_s]$ -modules. The resulting complex of free R -modules is called *the Koszul complex associated to the matrices $A_1, \dots, A_s$* and we denote it by $\mathbf{K}(A)$. The maps of $\mathbf{K}(A)$ are obtained from the maps of the Koszul complex $\mathbf{K}(t)$ by replacing every occurrence of t_i by A_i for $i = 1, \dots, s$.

Let $k = n - c$ be the dimension of X and $s_0, \dots, s_k \in H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ be sections that do not have a common zero on X . Then the morphism

$$\pi : X \rightarrow \mathbb{P}^k, \quad x \mapsto [s_0(x), \dots, s_k(x)],$$

defines a finite surjective morphism. Let $s_0, \dots, s_n$ a basis of $H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ and consider the inclusion

$$R := K[s_0, \dots, s_k] \subset K[s_0, \dots, s_n] = S$$

of graded polynomial rings. Because $\mathcal{F}$ is an Ulrich sheaf, it holds that $\pi_* \mathcal{F} \cong \mathcal{O}_{\mathbb{P}^k}^{\oplus t}$, so that there is an isomorphism

$$\psi : M \rightarrow R^m$$

of R -modules, where $m = t \cdot \deg(X)$. Under this identification, multiplication by s_i can be represented by an $m \times m$ matrix B_i with entries from R_1 for $i = 0, \dots, n$. Let $A_i = s_i \cdot I_m - B_i$, where I_m is the $m \times m$ identity matrix.

Theorem 4.23 *The Koszul complex $\mathbf{K}(A)$ associated to the matrices $A_{k+1}, \dots, A_n$ is a minimal free resolution of M . In particular, we can describe M as the cokernel of the matrix $(A_{k+1} | \dots | A_n)$ and $\text{Ext}_S^c(M, S)(-c)$ as the cokernel of $(A_{k+1}^t | \dots | A_n^t)$.*

Proof See [Kummer 2016, Proof of Theorem 6.2.5] or [Kummer and Shamovich 2020, Remark 4.8]. $\square$

Remark 4.24 Since the Koszul complex is self-dual, Theorem 4.23 also implies that a minimal free resolution of $\text{Ext}_S^c(M, S)(-c)$ is given by the Koszul complex $\mathbf{K}(A^t)$ associated to the matrices $A_{k+1}^t, \dots, A_n^t$. This is one way to see that $\text{Ext}_S^c(M, S)(-c)$ is an Ulrich module.

Corollary 4.25 *There is a natural isomorphism*

$$\text{Ext}_S^c(M, S)(-c) \cong \text{Hom}_R(M, R)$$

of S -modules. Here the S -module structure on $\text{Hom}_R(M, R)$ is defined by $(s\varphi)(x) := \varphi(sx)$ for $s \in S$ and $x \in M$.

Proof Under the identification

$$\psi^\vee : R^m = \text{Hom}_R(R^m, R) \rightarrow \text{Hom}_R(M, R),$$

multiplication by s_i on $\mathrm{Hom}_R(M, S)$ is represented by the $m \times m$ matrix B_i^t . Hence the S -module $\mathrm{Hom}_R(M, R)$ is isomorphic to the cokernel of $(A_{k+1}^t \cdots A_n^t)$, which is $\mathrm{Ext}_S^c(M, S)(-c)$ by Theorem 4.23. It is straightforward to check that the resulting isomorphism $\mathrm{Ext}_S^c(M, S)(-c) \rightarrow \mathrm{Hom}_R(M, R)$ does not depend on the choice of the R -module isomorphism $\psi : M \rightarrow R^m$. $\square$

The next lemma shows that we can work with $\mathrm{Ext}_S^c(M, S)$ and $\mathcal{E}\mathrm{xt}^c(\mathcal{F}, \mathcal{O}_{\mathbb{P}^n})$ interchangeably.

Lemma 4.26 *Let $N = \mathrm{Ext}_S^c(M, S)$ and $\mathcal{G} = D(\mathcal{F})$.*

- (1) *There is a natural isomorphism of coherent sheaves $\tilde{N} \cong \mathcal{G}$.*
- (2) *There is a natural isomorphism of graded S -modules $N \cong \Gamma_*(\mathcal{G})$.*

Proof A minimal free resolution of M gives rise to a free resolution of $\mathcal{F}$ of length c . Thus by [Hartshorne 1977, Proposition III.6.5] we can compute $\mathcal{E}\mathrm{xt}^c(\mathcal{F}, \mathcal{O}_{\mathbb{P}^n}(-c))$ as the cokernel of the dual of the last map of this resolution. This map is induced by the dual of the last map of the resolution of M whose cokernel is $\mathrm{Ext}_S^c(M, S)$. This implies part (1). Since N is a twist of an Ulrich module and therefore Cohen–Macaulay, we have $N = \Gamma_*(\tilde{N})$, which implies part (2). $\square$

Recall that for a finite surjective morphism $f : X \rightarrow Y$ of Noetherian schemes the right adjoint functor of f_* (considered as a functor from quasicoherent sheaves on X to quasicoherent sheaves on Y) can be described as follows. For a quasicoherent $\mathcal{O}_Y$ -module $\mathcal{G}$ the sheaf $\mathcal{H}\mathrm{om}_Y(f_*\mathcal{O}_X, \mathcal{G})$ is a quasicoherent $f_*\mathcal{O}_X$ -module. The corresponding quasicoherent $\mathcal{O}_X$ -module is then $f^!\mathcal{G}$. See for example [Hartshorne 1977, Exercise III.6.10].

Lemma 4.27 *Let $f : X \rightarrow Y$ be a finite surjective morphism of smooth varieties over K . Then $f^!\mathcal{O}_Y$ is a line bundle, which is naturally isomorphic to $\mathcal{H}\mathrm{om}(\det \mathcal{T}_X, f^*\det \mathcal{T}_Y)$.*

Proof This is, for example, shown in [Kummer 2016, Remark 2.2.19]. $\square$

Theorem 4.28 *Let $\mathcal{F}$ be an Ulrich sheaf on $\mathbb{P}^n$ and $\sigma : \mathcal{F} \rightarrow D\mathcal{F}$ an isomorphism. Assume that the support of $\mathcal{F}$ is a closed subvariety $\iota : X \hookrightarrow \mathbb{P}^n$ of dimension $k = \dim(X) > 0$, and that $\iota^*\mathcal{F}$ has rank one. Let $\pi : X \rightarrow \mathbb{P}^k$ be a finite surjective linear projection and $U \subset \mathbb{P}^k$ an open subset such that $V = \pi^{-1}(U) \subset X$ is smooth. Then there is a natural relative orientation*

$$\sigma(\pi|_V) : L \otimes L \rightarrow \mathcal{H}\mathrm{om}(\det \mathcal{T}_V, (\pi|_V)^*\det \mathcal{T}_U)$$

of $\pi|_V$, where $L = (\iota^*\mathcal{F})|_V$.

Proof Corollary 4.25 and Lemma 4.26 imply that there is a natural isomorphism

$$\iota^*(D\mathcal{F}) \rightarrow \mathcal{H}\mathrm{om}_X(\iota^*\mathcal{F}, \pi^!\mathcal{O}_{\mathbb{P}^k}).$$

Precomposing this with the pullback of σ under ι gives an isomorphism

$$\iota^*(\mathcal{F}) \rightarrow \mathcal{H}\mathrm{om}_X(\iota^*\mathcal{F}, \pi^!\mathcal{O}_{\mathbb{P}^k}).$$

Because every Ulrich sheaf is Cohen–Macaulay, the restriction $(\iota^*\mathcal{F})|_V$ is a line bundle on V . Indeed, in this situation being Cohen–Macaulay implies being locally free by [EGA IV₂ 1965, Proposition 6.1.5]. Thus restricting to V and tensoring with $\iota^*\mathcal{F}$ gives the isomorphism

$$(4-4) \quad (\iota^*\mathcal{F})|_V \otimes (\iota^*\mathcal{F})|_V \rightarrow (\pi^!\mathcal{O}_{\mathbb{P}^k})|_V = (\pi|_V)^!\mathcal{O}_U.$$

Finally, by Lemma 4.27 we obtain on V the isomorphism

$$(\iota^*\mathcal{F})|_V \otimes (\iota^*\mathcal{F})|_V \rightarrow \mathcal{H}om(\det \mathcal{T}_V, (\pi|_V)^* \det \mathcal{T}_U). \quad \square$$

Remark 4.29 One can prove in a similar way that if $\mathcal{F}$ is a rank one sheaf on X with an isomorphism $\sigma : \mathcal{F} \xrightarrow{\sim} D\mathcal{F}$, and such that the restriction $\mathcal{F}|_V$ is a line bundle L , then L is a relative orientation for the map $\pi|_V : V \rightarrow U$. This works also when $\mathcal{F}$ is not Ulrich. However, when $\mathcal{F}$ is Ulrich, the proof of Theorem 4.28 shows how to compute the $\mathbb{A}^1$ -degree of the map $\pi|_V$ via Theorem 4.17. We will make this clear in Theorem 4.32.

Definition 4.30 In the situation of Theorem 4.28 we call $\sigma(\pi|_V)$ the *relative orientation induced by σ* .

Lemma 4.31 *If the rank of $\iota^*\mathcal{F}$ is one, then every isomorphism $\sigma : \mathcal{F} \rightarrow D\mathcal{F}$ is symmetric.*

Proof Since $\mathcal{F}$ is supported on X , it suffices to prove the pullback of the equality $\sigma = D(\sigma) \circ \beta$ under ι . Let V as in Theorem 4.28. Since V is dense in X , it further suffices to prove the equality on V . As in equation (4-4) in the proof of Theorem 4.28 the isomorphism σ induces an isomorphism

$$\rho : (\iota^*\mathcal{F})|_V \otimes (\iota^*\mathcal{F})|_V \rightarrow (\pi|_V)^!\mathcal{O}_U$$

and the condition on σ being symmetric translates to the condition that $\rho_x(a \otimes b) = \rho_x(b \otimes a)$ for all $x \in V$ and $a, b \in (\iota^*\mathcal{F})|_x$. Since $(\iota^*\mathcal{F})|_V$ is locally free of rank one, there is $t \in (\iota^*\mathcal{F})|_x$ and $f, g \in \mathcal{O}_{X,x}$ such that $a = f \cdot t$ and $b = g \cdot t$. Then

$$\rho_x(a \otimes b) = fg \cdot \rho_x(t \otimes t) = \rho_x(b \otimes a). \quad \square$$

Now we are ready to state the main result of this section.

Theorem 4.32 *Let $\mathcal{F}$ be an Ulrich sheaf on $\mathbb{P}^n$ whose support is a closed subvariety $\iota : X \hookrightarrow \mathbb{P}^n$ of dimension $k > 0$ and degree d . Assume that the rank of $\iota^*\mathcal{F}$ is one and let $\sigma : \mathcal{F} \rightarrow D\mathcal{F}$ be an isomorphism. There is a symmetric $d \times d$ matrix Λ whose entries are linear forms on $\bigwedge^{k+1} H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ such that for all $s_0, \dots, s_k \in H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ it holds that:*

- (1) *The entry-wise evaluation $\Lambda(s_0 \wedge \dots \wedge s_k)$ is a singular matrix if and only if the s_j have a common zero on X .*
- (2) *If the s_j do not have a common zero on X , let $\pi : X \rightarrow \mathbb{P}^k$ be the morphism defined by $s_0, \dots, s_k$ and $U \subset \mathbb{P}^k$ an open subset such that $V = \pi^{-1}(U) \subset X$ is smooth. The class of $\Lambda(s_0 \wedge \dots \wedge s_k)$ in $\text{GW}(K)$ equals $\deg^{\mathbb{A}^1}(\pi|_V)$, where $\pi|_V$ is relatively oriented by $\sigma(\pi|_V)$.*

Proof We first note that by Lemma 4.31 the isomorphism $\sigma : \mathcal{F} \rightarrow D\mathcal{F}$ is symmetric. Let $M = \Gamma_*(\mathcal{F})$ and let

$$F : 0 \rightarrow F_c \xrightarrow{\varphi_c} F_{c-1} \xrightarrow{\varphi_{c-1}} \dots \xrightarrow{\varphi_2} F_1 \xrightarrow{\varphi_1} F_0 \rightarrow M \rightarrow 0$$

be a minimal free resolution of M . We choose any S -basis $\mathcal{B}_0$ of F_0 . The image of $\mathcal{B}_0$ under the map $F_0 \rightarrow M$ is sent by σ to a generating set of $\text{Ext}^c(M, S)(-c)$. A preimage of this generating set in the degree-zero part of $\text{Hom}_S(F_c, S)(-c)$ under the natural map $\text{Hom}_S(F_c, S)(-c) \rightarrow \text{Ext}^c(M, S)(-c)$ is a basis of $\text{Hom}_S(F_c, S)(-c)$, and thus of $\text{Hom}_S(F_c, S)$. We denote by $\mathcal{B}_c$ the dual of this basis, which is a basis of F_c . We claim that with this choice of bases of F_0 and F_c (and any choice of basis of F_i for $0 < i < c$) the matrix $\gamma(F)$ has the desired properties. We already know that $\gamma(F)$ is a $d \times d$ matrix whose entries are linear forms on $\bigwedge^{k+1} H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ which satisfies (1). It thus remains to show that $\gamma(F)$ is symmetric and satisfies (2). To this end let $s_0, \dots, s_k \in H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ be sections that do not have a common zero on X . Then, as in the discussion before Theorem 4.23, the morphism $\pi : X \rightarrow \mathbb{P}^k$ defined by $s_0, \dots, s_k$ corresponds to an inclusion of graded polynomial rings

$$R := K[s_0, \dots, s_k] \subset K[s_0, \dots, s_n] = S.$$

We consider the isomorphism

$$\psi : M \rightarrow R^d$$

of R -modules which sends the image of $\mathcal{B}_0$ under the map $F_0 \rightarrow M$ to the standard basis of R^d . We denote by B_i the representing matrix of multiplication by s_i with respect to this basis and let $A_i = s_i \cdot I_d - B_i$, where I_d is the $d \times d$ identity matrix. Further let $\mathbf{K}(A)$ the Koszul complex associated to the matrices $A_{k+1}, \dots, A_n$. We then have

$$(4-5) \quad \gamma(F) \cdot Q = \gamma(\mathbf{K}(A)),$$

where Q is the representing matrix of the R -bilinear form on M obtained by composing $\sigma : M \rightarrow \text{Ext}^c(M, S)(-c)$ with the natural isomorphism from Corollary 4.25 with respect to our chosen basis. Note that Q is a symmetric matrix because σ is symmetric which further satisfies

$$(4-6) \quad QA_i = A_i^t Q$$

for all $i = k+1, \dots, n$ because σ is a homomorphism of S -modules. By Theorem 4.17 and our choice of relative orientation the class of Q in $\text{GW}(K)$ equals to $\deg^{\mathbb{A}^1}(\pi|_V)$. Since Q and Q^{-1} represent the same class in $\text{GW}(K)$ and by equation (4-5), it now suffices to show that $\gamma(\mathbf{K}(A)) \cdot Q^{-1}$ is symmetric and that $\gamma(\mathbf{K}(A))(s_0 \wedge \dots \wedge s_k) = I_d$. To this end we regard $\gamma(\mathbf{K}(A))$ as an alternating multilinear map

$$(H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1)))^{\vee c} \rightarrow \text{Mat}_d(K),$$

and use the explicit description

$$(4-7) \quad \gamma(\mathbf{K}(A))(v_1, \dots, v_c) = \sum_{\tau \in S_c} \text{sgn}(\tau) \cdot A_{k+\tau(1)}(v_1) \cdots A_{k+\tau(c)}(v_c),$$

where $v_1, \dots, v_c \in H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))^\vee$; see [Kummer 2016, Example 6.1.4]. Equation (4-7) together with equation (4-6) now imply

$$\begin{aligned} \gamma(\mathbf{K}(A))(v_1, \dots, v_c) \cdot Q^{-1} &= \sum_{\tau \in S_c} \operatorname{sgn}(\tau) \cdot A_{k+\tau(1)}(v_1) \cdots A_{k+\tau(c)}(v_c) \cdot Q^{-1} \\ &= \sum_{\tau \in S_c} \operatorname{sgn}(\tau) Q^{-1} \cdot A_{k+\tau(1)}^t(v_1) \cdots A_{k+\tau(c)}^t(v_c) \\ &= Q^{-1} \gamma(\mathbf{K}(A))^t(v_1, \dots, v_c). \end{aligned}$$

This shows that $\gamma(\mathbf{K}(A)) \cdot Q^{-1}$ is symmetric. Letting $s_0^\vee, \dots, s_n^\vee$ be the dual basis of $s_0, \dots, s_n$ we further have

$$\begin{aligned} \gamma(\mathbf{K}(A))(s_0 \wedge \cdots \wedge s_k) &= \gamma(\mathbf{K}(A))(s_{k+1}^\vee, \dots, s_n^\vee) \\ &= \sum_{\tau \in S_c} \operatorname{sgn}(\tau) \cdot A_{k+\tau(1)}(s_{k+1}^\vee) \cdots A_{k+\tau(c)}(s_{k+c}^\vee) = I_d, \end{aligned}$$

since $A_{k+i}(s_{k+j}^\vee) = \delta_{ij} \cdot I_d$. This concludes the proof. $\square$

Example 4.33 Consider the rational normal curve $\iota: \mathbb{P}^1 \hookrightarrow \mathbb{P}^n$ of degree n and the symmetric Ulrich sheaf $\mathcal{F} = \iota_*(\mathcal{O}_{\mathbb{P}^1}(n-1))$. Theorem 4.32 gives an $n \times n$ matrix Λ whose entries are linear forms on $\wedge^2 K[x_0, x_1]_n$ such that for all binary forms $p, q \in K[x_0, x_1]_n$ we have:

- (1) $\Lambda(p \wedge q)$ has full rank if and only if p and q are coprime.
- (2) If p and q are coprime, then the $\mathbb{A}^1$ -degree of the map

$$\psi: \mathbb{P}^1 \rightarrow \mathbb{P}^1, \quad (x_0 : x_1) \mapsto (p(x_0, x_1) : q(x_0, x_1)),$$

relatively oriented by an isomorphism

$$\mathcal{O}_{\mathbb{P}^1}(n-1) \otimes \mathcal{O}_{\mathbb{P}^1}(n-1) = \mathcal{O}_{\mathbb{P}^1}(2(n-1)) \cong \omega_{\mathbb{P}^1} \otimes \psi^* \omega_{\mathbb{P}^1}^{-1}$$

is given by $\Lambda(p \wedge q)$.

This resulting matrix $\Lambda(p \wedge q)$ is the so-called *Bézout matrix* of p and q . Its connection to $\mathbb{A}^1$ -homotopy theory has already been observed in [Cazanave 2012].

Example 4.34 Consider the elliptic curve $(E, \mathcal{O})$ with Weierstrass equation

$$y^2 = x^3 - x$$

over a field K with $\operatorname{char}(K) = 0$. Letting $P_0 = (0, 0)$ and $P_{-1} = (-1, 0)$, the line bundle $\mathcal{L}$ associated with the divisor $P_{-1} - P_0$ is 2-torsion. We consider the embedding of $\iota: E \rightarrow \mathbb{P}^3$ via the linear system W spanned by $(s_0, s_1, s_2, s_3) = (1, x, y, x^2)$. The sheaf $\mathcal{F} = \mathcal{O}_{\mathbb{P}^3}(1) \otimes \iota_* \mathcal{L}$ is a symmetric Ulrich sheaf.

From its free resolution we obtain the symmetric matrix

$$\Lambda = \begin{pmatrix} -x_{23} + x_{12} + x_{02} & x_{23} - x_{12} - x_{02} & x_{13} + x_{03} & x_{03} + x_{01} \\ x_{23} - x_{12} - x_{02} & x_{12} + x_{02} & -x_{03} & -x_{01} \\ x_{13} + x_{03} & -x_{03} & x_{23} + x_{12} & x_{02} \\ x_{03} + x_{01} & -x_{01} & x_{02} & x_{02} \end{pmatrix},$$

from which we can read off the $\mathbb{A}^1$ -degree of maps $\psi : E \rightarrow \mathbb{P}^1$ given by two elements of W and relatively oriented by an isomorphism

$$\mathcal{L}(1) \otimes \mathcal{L}(1) \cong \omega_E \otimes \psi^* \omega_{\mathbb{P}^1}^{-1}.$$

Here x_{ij} denotes the linear form on $\wedge^2 W$ that evaluates on $s_k \wedge s_l$ to 1 if $(i, j) = (k, l)$, and to 0 if $\{i, j\} \neq \{k, l\}$. For instance, the map defined by

$$E \rightarrow \mathbb{P}^1, \quad (x, y) \mapsto (1 : x^2),$$

has $x_{03} = 1$ and all other Plücker coordinates zero. Thus its $\mathbb{A}^1$ -degree equals

$$\begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & -1 & 0 \\ 1 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \cong 2 \cdot (\langle 1 \rangle \quad \langle -1 \rangle).$$

On the other hand, the map defined by

$$E \rightarrow \mathbb{P}^1, \quad (x, y) \mapsto (y : (2x + 1)(x - 2)),$$

has $x_{02} = 2$, $x_{12} = 3$, $x_{23} = 2$ and all other Plücker coordinates zero. Thus its $\mathbb{A}^1$ -degree equals

$$\begin{pmatrix} 3 & -3 & 0 & 0 \\ -3 & 5 & 0 & 0 \\ 0 & 0 & 5 & 2 \\ 0 & 0 & 2 & 2 \end{pmatrix} \cong 2 \cdot (\langle 3 \rangle + \langle 2 \rangle).$$

5 Secant varieties of curves

In this section we prove some cohomological statements on the secant variety that we will need later to construct Ulrich bundles. Furthermore, we define Viro frames on (the desingularization of) the secant variety.

5.1 Symmetric products of curves

We recall some preliminaries on symmetric products of curves, their tautological bundles and the connections to secant varieties. Unless we give another reference, the statements made in this section can be found in [Ein et al. 2020, §3]. If C is a curve, we will denote by C_n its n -th symmetric product. It can

be considered as the quotient $\sigma : C^n \rightarrow C_n$ with respect to the natural $\mathfrak{S}_n$ -action. This is a nonsingular projective variety of dimension n that parametrizes effective divisors of degree n on C . The addition map

$$\sigma : C \times C_{n-1} \rightarrow C_n, \quad (x, D) \mapsto x + D,$$

makes $C \times C_{n-1}$ into the universal family over C_n : the fiber over $D \in C_n$ is naturally isomorphic to the subscheme $D \subset C$. If L is a line bundle on C we can pull it back to the universal family $C \times C_{n-1}$ and then pushforward to define the tautological bundle $E_{n,L} = \sigma_* \operatorname{pr}_C^* L$. The sheaf $E_{n,L}$ is a vector bundle of rank n on C_n whose fiber at D can be identified with $H^0(C, L \otimes \mathcal{O}_D)$. Here $\mathcal{O}_D$ denotes the structure sheaf of D considered as subscheme of C .

Definition 5.1 A line bundle L on C is called *k-very ample* if

$$h^0(C, L \otimes \mathcal{O}_C(-D)) = h^0(C, L) - (k + 1)$$

for all effective divisors D of degree $k + 1$.

Remark 5.2 A line bundle L is 0-very ample if and only if it is globally generated, and it is 1-very ample if and only if it is very ample. If $k > 1$, then L is k -very ample if and only if it induces an embedding $\varphi_L : C \hookrightarrow \mathbb{P}^n$ such that any, possibly coincident, $k + 1$ points on C span a k -plane in $\mathbb{P}^n$.

Remark 5.3 A line bundle L is 3-very ample if and only if its complete linear system induces an embedding $\varphi_L : C \hookrightarrow \mathbb{P}^n$ such that the secant variety $\Sigma = \Sigma(C, L)$ of $\varphi_L(C)$ is identifiable. This means that any point in $\Sigma \setminus \varphi_L(C)$ is contained in a unique secant line. Indeed, two secant lines $\ell(p_1, p_2)$ and $\ell(p_3, p_4)$, with $p_i \in C$ meet, if and only if the points $p_1, \dots, p_4$ span a 2-plane.

We have $H^0(C, L) = H^0(C_n, E_{n,L})$. Further, if the line bundle L is $(n-1)$ -very ample, then $E_{n,L}$ is globally generated. Thus in this case there is a surjective map

$$\operatorname{ev}_L : H^0(C, L) \otimes \mathcal{O}_{C_n} \rightarrow E_L$$

of vector bundles on C_n which is an isomorphism on global sections.

5.2 Line bundles on the symmetric product

We now define some line bundles on C_n and describe their relations to each other. More precisely, to any line bundle L on C we can associate two line bundles on C_n . Firstly, the line bundle $N_{n,L}$ on C_n is defined as the determinant of E_L . We write $\mathcal{O}_{C_n}(\delta) = N_{\mathcal{O}_C}^\vee$, and δ is the corresponding divisor class. Further, on the direct product C^n we have the line bundle $L^{\boxtimes n} = \operatorname{pr}_1^* L \otimes \cdots \otimes \operatorname{pr}_n^* L$. Since $L^{\boxtimes n}$ is invariant under permuting the components, it descends to a line bundle $S_{n,L}$ on C_n satisfying $\sigma^* S_L \cong L^{\boxtimes n}$ and the induced map

$$\operatorname{Pic}(C) \rightarrow \operatorname{Pic}(C_n), \quad L \mapsto S_{n,L},$$

is a group homomorphism. For a closed point $P \in C$ we have the divisor $S_P = P + C_{n-1} \subset C_n$ and the associated line bundle is precisely $S_{\mathcal{O}_C(P)}$. Extending this linearly, we can define the divisor S_D on C_n

for every divisor D on C_n . The associated line bundle of S_D is $S_{\mathcal{O}_C(D)}$ and we will use both notations interchangeably. Finally, a distinguished divisor on C_n is the locus $\Delta \subset C_n$ of all nonreduced divisors, and for any line bundle L on C we set $A_{n,L} = S_{n,L}(-\Delta)$ as in [Ein et al. 2021]. The following lemma summarizes the relations of these line bundles and divisors to each other.

Lemma 5.4 *Letting $\mathcal{O}_{C_n}(\delta) = \det E_{n,\mathcal{O}_C}^\vee$, we have:*

- (a) $\mathcal{O}_{C_n}(2\delta) \cong \mathcal{O}_{C_n}(\Delta)$.
- (b) $N_{n,L} \otimes \mathcal{O}_{C_n}(\delta) \cong S_{n,L}$ for every line bundle L on C .
- (c) The cotangent vector bundle is $\Omega_{C_n}^1 \cong E_{n,\omega_n}$ and the canonical line bundle on C_n is $\omega_{C_n} \cong N_{n,\omega_C}$.

The cohomology of these line bundles is known:

Lemma 5.5 *Let L and M be line bundles on C . Then we have isomorphisms*

$$\begin{aligned} H^i(C_n, N_{n,L}) &\cong \wedge^{n-i} H^0(C, L) \otimes \operatorname{Sym}^i H^1(C, L), \\ H^i(C_n, S_{n,L}) &\cong \operatorname{Sym}^{n-i} H^0(C, L) \otimes \wedge^i H^1(C, L), \\ H^i(C_n, E_{n,M} \otimes S_{n,L}) &\cong H^0(C, L \otimes M) \otimes \operatorname{Sym}^{n-1-i} H^0(C, L) \otimes \wedge^i H^1(C, L) \\ &\quad \oplus H^1(C, L \otimes M) \otimes \operatorname{Sym}^{n-i} H^0(C, L) \otimes \wedge^{i-1} H^1(C, L). \end{aligned}$$

Proof The first two follow from the discussion after [Krug 2018, Proposition 6.3]. For the last one, we can use the same strategy as in [Agostini 2024, Lemma 4.4] and use the universal family $\sigma : C_{n-1} \times C \rightarrow C_n$. By definition of $E_{n,L}$ and by the projection formula $E_{n,M} \otimes S_{n,L} \cong \sigma_*(\operatorname{pr}_C^* M \otimes \sigma^* S_{n,L})$. Since the map σ is finite, it also follows that $H^i(C_n, E_{n,M} \otimes S_{n,L}) \cong H^i(C \times C_{n-1}, \operatorname{pr}_C^* L \otimes \sigma^* S_{n,M})$. To conclude, it is enough to observe that $\sigma^* S_{n,M} \cong \operatorname{pr}_C^*(M) \otimes \operatorname{pr}_{C_{n-1}}^* S_{n-1,M}$ and use the formulas for $H^i(C_{n-1}, S_{n-1,M})$. $\square$

5.3 Secants and Viro frames

Let L be a 3-very ample line bundle on the curve C and $\iota : C \hookrightarrow \mathbb{P}(H^0(C, L)^\vee) = \mathbb{P}^r$ the corresponding embedding. The fact that L is 3-very ample means precisely that no four (possibly coincident) points on C are contained in a plane in $\mathbb{P}^r$. In this section we recall some observations on the projective space bundle $B = \mathbb{P}(E_{2,L})$, where E_L is the tautological bundle on C_2 defined in Section 5. Since E_L is globally generated, we have a natural identification

$$B = \{(p, x + y) \in \mathbb{P}^r \times C_2 \mid p \in \ell(x + y)\},$$

where $\ell(x + y) \subset \mathbb{P}^r$ denotes the line in $\mathbb{P}^r$ spanned by the subscheme $x + y$ of C . We consider the projections

$$p : B \rightarrow C_2 \quad \text{and} \quad \varepsilon : B \rightarrow \Sigma,$$

where Σ is the secant variety of $\iota(C) \subset \mathbb{P}^r$. If L is 3-very ample, then the map ε is an isomorphism outside of $C \subsetneq \Sigma$ and we denote the preimage of C by E . In this case we can identify E with $C \times C$ via

$$C \times C \rightarrow B, \quad (x, y) \mapsto (\iota(x), x + y).$$

In particular, the fiber of ε over $x \in C$ is naturally identified with $\{x\} \times C$. Now we would like to compute the class of the divisor E , but instead of working with classes, we will introduce some explicit rational functions, that will be useful later on.

Let z_0, z_1 be two linearly independent global sections of $\mathcal{O}_{\mathbb{P}^r}(1)$. Let $H_\infty \subset \mathbb{P}^r$ be the hyperplane $H_\infty = \{z_0 = 0\}$. This pulls back to the effective divisors $D = \iota^*(H_\infty)$ on C and $H = \varepsilon^*(H_\infty)$ on B . Note that $L \cong \mathcal{O}_C(D)$ and $\mathcal{O}_B(1) \cong \mathcal{O}_B(H)$. Consider the rational functions $q = z_1/z_0$ on $\mathbb{P}^r$, $f = \iota^*q$ on C and $h = \varepsilon^*q$ on B . Write $f_1 = f \otimes 1$ and $f_2 = 1 \otimes f$, which are rational functions on $C \times C$. Further, we let $g_1 = f_1 + f_2$ and $g_2 = f_1 \cdot f_2$. These are rational functions on C_2 , since they are $\mathfrak{S}_2$ -invariant. Finally, we denote by $\tilde{\Delta}$ the closure in C_2 of the set

$$\{x + y \mid x, y \in (C \setminus D), x \neq y, f(x) = f(y)\}.$$

Lemma 5.6 Consider the rational function $G = p^*g_2 - p^*g_1 \cdot h + h^2$ on B . Its divisor is

$$\operatorname{div}(G) = E + p^*(2\tilde{\Delta} - S_D) - 2H.$$

Proof A point on B has the form $(z, x + y)$, where z is a point on the line spanned by $\iota(x)$ and $\iota(y)$. Thus G evaluated at this point is

$$\begin{aligned} (5-1) \quad G &= f_1(x) \cdot f_1(y) - (f_1(x) + f_2(y)) \cdot h(z) + h(z)^2 \\ &= (f_1(x) - h(z))(f_2(y) - h(z)) \\ &= (q(\iota(x)) - q(z))(q(\iota(y)) - q(z)). \end{aligned}$$

Since z is on the line spanned by $\iota(x)$ and $\iota(y)$, we find that $f(x) = f(y)$ implies $q(z) = q(\iota(x)) = q(\iota(y))$. This means that G vanishes of order 2 along $p^*(\tilde{\Delta})$. If instead $f(x) \neq f(y)$, then G vanishes on $(z, x + y)$ if and only if $z = \iota(x)$ or $z = \iota(y)$, which means that $(z, x + y)$ lies on E . This shows that the zero divisor of G is $E + p^*(2\tilde{\Delta})$. Finally, G has a pole of order 2 whenever z lies on H_∞ , and an ordinary pole whenever $\iota(x)$ or $\iota(y)$ lies on H_∞ . This shows that the pole divisor of G is $p^*(S_D) + 2H$. $\square$

Lemma 5.7 The divisor of the rational differential 2-form $dg_1 \wedge dg_2$ on C_2 is $S_{\operatorname{div}(df)} - S_D + \tilde{\Delta}$.

Proof One computes that

$$(5-2) \quad dg_1 \wedge dg_2 = (f_1 - f_2) \cdot (df_1 \wedge df_2).$$

The divisor of this rational differential 2-form on $C \times C$ is $\Delta_C + \sigma^*(S_{\operatorname{div}(df)} - S_D + \tilde{\Delta})$, where $\Delta_C \subset C \times C$ is the diagonal. Now the claim follows from the fact that Δ_C is the ramification divisor of σ . $\square$

Now, fix some rational differential ω on C . Then there is a rational function t on C such that $\omega = t \cdot df$. Let $t_1 = t \otimes 1$ and $t_2 = 1 \otimes t$, and consider the rational function $F = t_1 t_2$ on C_2 . Define at this point the rational 3-form on B

$$(5-3) \quad \Omega(q, \omega) := p^* \left(\frac{F}{G} \cdot (dg_2 \wedge dg_1) \right) \wedge dh.$$

Lemma 5.8 *The divisor of the form $\Omega(q, \omega)$ is*

$$(5-4) \quad \operatorname{div} \Omega(q, \omega) = p^*(S_{\operatorname{div}(\omega)}) - E.$$

Proof By Lemma 5.7 the divisor of the rational differential 2-form $F \cdot (dg_1 \wedge dg_2)$ on C_2 is $S_{\operatorname{div}(\omega)} - S_D + \tilde{\Delta}$. If the differential dh is defined and nonzero at a point $(z, x + y)$ on B , then it is in the span of dg_1 and dg_2 at this point if and only if $f(x) = f(y)$. Thus the divisor of the differential 3-form $F \cdot (dh \wedge dg_1 \wedge dg_2)$ is

$$p^*(\operatorname{div}(F \cdot (dg_1 \wedge dg_2))) + p^*(\tilde{\Delta}) - 2H = p^*(S_{\operatorname{div}(\omega)} - S_D + 2\tilde{\Delta}) - 2H.$$

Combining this with Lemma 5.6, we conclude. $\square$

Remark 5.9 In particular, this computes one canonical divisor of B as $K_B = p^*S_{K_C} - E$. Using the formula $K_B \sim p^*(K_{C_2}) + (p^*N_L) - 2H$ for the canonical class of the projective bundle $p : B \rightarrow C_2$, we can also compute that $E \sim 2H - p^*(S_D - \Delta)$, where D is any divisor such that $\mathcal{O}_C(D) \cong L$.

The differential 3-form $\Omega(q, \omega)$ depends a priori on the choice of the rational function $q = z_1/z_0$ on $\mathbb{P}^r$ and on the differential ω on C . Letting $\omega' = t' \cdot \omega$ be another differential, we get that

$$(5-5) \quad \Omega(q, \omega') = p^*(t' \otimes t') \cdot \Omega(q, \omega),$$

where we see $t' \otimes t'$ as a rational function on C_2 . Equation (5-5) shows, in particular:

Corollary 5.10 *If ω' differs from ω only by a scalar or by the square of a rational function, then $\Omega(q, \omega')$ differs from $\Omega(q, \omega)$ only by a square.*

Next we will show that $\Omega(q, \omega)$ does not depend on q . Equation (5-4) shows that the divisor of $\Omega(q, \omega)$ does not depend on q . Therefore, it suffices to evaluate $\Omega(q, \omega)$ at one basis of the tangent space of B at one point, and show that the result does not depend on the choice of q . To this end let a, b be two distinct points of C where ω does not have a zero or a pole, and let $z \in \ell(a, b) \setminus \{a, b\}$ be a point on the line spanned by them. Then set $Q = (p, a + b)$. Since $a \neq b$, the differential

$$T_{(a,b)}(\sigma) : T_a C \times T_b C \rightarrow T_{a+b} C_2$$

is an isomorphism. Further, let $F = \ell(a + b)$ be the fiber of $p : B \rightarrow C_2$ over $a + b$. There is a unique isomorphism $\psi : \mathbb{P}^1 \rightarrow F$ that satisfies $\psi(1 : 0) = \iota(a)$, $\psi(1 : 1) = z$ and $\psi(0 : 1) = \iota(b)$. Via the differentials $T_{(a,b)}(\sigma)$ and $T_{(1:1)}(\psi)$ we obtain the short exact sequence of tangent spaces

$$0 \rightarrow T_{(1:1)} \mathbb{P}^1 \rightarrow T_Q B \rightarrow T_a C \times T_b C \rightarrow 0,$$

whose determinant gives the isomorphism

$$\Psi : T_a C \otimes T_{(1:1)} \mathbb{P}^1 \otimes T_b C \rightarrow \wedge^3 T_Q B.$$

Let $v \in T_a C$ and $w \in T_b C$ vectors on which ω evaluates to one. Let x_0, x_1 be the homogeneous coordinates on $\mathbb{P}^1$, let $t = x_1/x_0$ and $u \in T_{(1:1)}$ be the vector on which dt evaluates to one. Now consider

$$\Psi(v \otimes u \otimes w) \in \wedge^3 T_Q B.$$

Besides the choice of the differential ω , the construction depends a priori on the choice of the preimage (a, b) of $a + b$ under σ . However, choosing the other preimage, we obtain

$$\Psi(w \otimes (-u) \otimes v) = \Psi(v \otimes u \otimes w).$$

Thus $\Psi(v \otimes u \otimes w)$ is independent of this choice.

Definition 5.11 We call

$$V(\omega, Q) := \Psi(v \otimes u \otimes w) \in \wedge^3 T_Q B$$

the ω -Viro frame at Q .

Now we evaluate $\Omega(q, \omega)$ at $V(\omega, Q)$. We can write

$$s := q \circ \psi = \frac{\alpha t + \beta}{\gamma t + \delta}$$

for some scalars α, β, γ and δ . Using equations (5-1) and (5-2) we calculate that the evaluation of $\Omega(q, \omega)$ at $V(\omega, Q)$ equals the evaluation of

$$\frac{q(\iota(a)) - q(\iota(b))}{(q(\iota(a)) - q(p))(q(\iota(b)) - q(p))} ds = \frac{s(1:0) - s(0:1)}{(s(1:0) - s(1:1))(s(0:1) - s(1:1))} ds$$

at u . On the one hand we have

$$\frac{s(1:0) - s(0:1)}{(s(1:0) - s(1:1))(s(0:1) - s(1:1))} = \frac{(\gamma + \delta)^2}{\alpha\delta - \beta\gamma},$$

and on the other hand one calculates

$$ds = \frac{\alpha(\gamma t + \delta) - \gamma(\alpha t + \beta)}{(\gamma t + \delta)^2} dt,$$

which evaluates at u to $(\alpha\delta - \beta\gamma)/(\gamma + \delta)^2$. Thus we have proven the following.

Theorem 5.12 The differential 3-form $\Omega(\omega) := \Omega(q, \omega)$ does not depend on the choice of q . At every ω -Viro frame it evaluates to 1.

5.4 Some cohomology

Now we prove some technical cohomological vanishing statements that we will use to construct an Ulrich sheaf on the secant variety.

Suppose again that $\iota : C \hookrightarrow \mathbb{P}(H^0(C, L)^\vee) = \mathbb{P}^r$ is an embedding by a 3-very ample line bundle. Let Σ be the secant variety and

$$C_2 \xleftarrow{p} B \xrightarrow{\varepsilon} \Sigma$$

as before. We take a vector bundle F on C_2 and we consider the induced sheaf $\varepsilon_* p^* F$ on Σ . Our lemma allows to transfer the cohomology of this sheaf from Σ to B :

Lemma 5.13 *Assume that*

$$H^1(\{x\} \times C, \sigma^* F|_{\{x\} \times C}) = 0 \quad \text{for all } x \in C,$$

or equivalently, that $H^1(x + C, F|_{x+C}) = 0$ for all $x \in C$. Then

$$H^j(\Sigma, \varepsilon_* p^* F \otimes \mathcal{O}_{\mathbb{P}^n}(i)) \cong H^j(B, p^* F \otimes \mathcal{O}_B(i)) \quad \text{for all } i, j.$$

Proof Suppose $R^1 \varepsilon_* p^* F = 0$. Since $\varepsilon : B \rightarrow \Sigma$ has fibers of dimension ≤ 1 , the Leray spectral sequence together with the projection formula shows that $H^i(B, p^* F \otimes \mathcal{O}_B(s)) \cong H^i(\Sigma, \varepsilon_* p^* F \otimes \mathcal{O}_\Sigma(q))$ for all $i \geq 0$. To prove that $R^1 \varepsilon_* p^* F = 0$ we follow an argument of [Ullery 2016] and [Ein et al. 2020] that we repeat here. Consider the fiber $\varepsilon^{-1}(x)$ over $x \in C$, and let $\mathcal{I} = \mathcal{I}_{\varepsilon^{-1}(x)/B}$ be its sheaf of ideals on B . Then the Theorem of Formal Functions proves that the vanishing $R^1 \varepsilon_* p^* F$ follows from $H^1(B, p^* F \otimes \mathcal{O}_B/\mathcal{I}^k) = 0$ for all $k \geq 1$. Looking at the sequences

$$0 \rightarrow \mathcal{I}^k/\mathcal{I}^{k+1} \rightarrow \mathcal{O}_B/\mathcal{I}^{k+1} \rightarrow \mathcal{O}_B/\mathcal{I}^k \rightarrow 0,$$

it is enough to show that $H^1(B, p^* F \otimes \mathcal{I}^k/\mathcal{I}^{k+1}) = 0$ for all $k \geq 0$. Now, we know that in our case $\varepsilon^{-1}(x)$ is a smooth subvariety of the smooth variety B , hence it is a locally complete intersection, so that the sheaf $\mathcal{I}^k/\mathcal{I}^{k+1}$ is naturally isomorphic to the symmetric power of the conormal bundle $\text{Sym}^k(\mathcal{I}/\mathcal{I}^2) \cong \mathcal{I}^k/\mathcal{I}^{k+1}$. In summary, we are left to prove the vanishings

$$H^1(\varepsilon^{-1}(x), (p^* F|_{\varepsilon^{-1}(x)} \otimes \text{Sym}^k(\mathcal{I}/\mathcal{I}^2))) = 0 \quad \text{for all } k \geq 0.$$

Now we identify the pieces in our situation: we know from before that the fiber is identified with $\{x\} \times C \subset C \times C$, and [Ullery 2016, Lemma 2.3] shows that $\mathcal{I}/\mathcal{I}^2 \cong \mathcal{O}_C \oplus L(-2x)$, hence $\text{Sym}^k(\mathcal{I}/\mathcal{I}^2) \cong \bigoplus_{h=0}^k L(-2x)^{\otimes h}$. Since we are assuming that L is 3-very ample, it must be that $h^0(C, L) \geq 3$, so that $L(-2x)$ is effective. Thus, all the vanishings that we need are implied by the single vanishing

$$H^1(\varepsilon^{-1}(x), (p^* F|_{\varepsilon^{-1}(x)})) = 0.$$

With the identification of $\varepsilon^{-1}(x)$ with $\{x\} \times C$, this is exactly the vanishing in our assumption. $\square$

Now, we want to look for an Ulrich sheaf of the form

$$\mathcal{F} = \varepsilon_* p^* F \otimes \mathcal{O}_{\mathbb{P}^n}(t)$$

for a certain $t \in \mathbb{Z}$.

Lemma 5.14 *Let F be a vector bundle on C_2 and let $t \in \mathbb{Z}$ be such that*

- (1) $H^1(x + C, F|_{x+C}) = 0$ for all $x \in C$,
- (2) $H^i(C_2, F \otimes S^{t-i} E_{2,L}) = H^{i-1}(C_2, F \otimes S^{t-i} E_{2,L}) = 0$ for $1 \leq i \leq 3$ and $i \leq t$,
- (3) $H^i(C_2, F^\vee \otimes S^{2-i-t} E_{2,L} \otimes A_{2,\omega_C} \otimes L) = 0$ and $H^{i-1}(C_2, F^\vee \otimes S^{2-i-t} E_{2,L} \otimes A_{2,\omega_C} \otimes L) = 0$ for $1 \leq i \leq 3$ and $i \leq 2 - t$.

Then $\mathcal{F} = \varepsilon_* p^* F \otimes \mathcal{O}_{\mathbb{P}^n}(t)$ is an Ulrich sheaf on Σ of rank equal to the rank of F .

Proof Condition (1) and Lemma 5.13 show that the conditions for being Ulrich are equivalent to

$$H^i(B, p^* F \otimes \mathcal{O}_B(t-i)) = H^{i-1}(B, p^* F \otimes \mathcal{O}_B(t-i)) = 0 \quad \text{for } 1 \leq i \leq 3.$$

We compute these via the Leray spectral sequence for $p : B \rightarrow C_2$. The projection formula gives that $R^a p_*(p^* F \otimes \mathcal{O}_B(t-i)) = F \otimes R^a p_* \mathcal{O}_B(t-i)$ for all a , and [Hartshorne 1977, Exercise III.8.4] shows that

$$R^a p_* \mathcal{O}_B(t-i) \cong \begin{cases} S^{t-i} E_{2,L} & \text{if } a = 0 \text{ and } i \leq t, \\ S^{i-t-2} E_{2,L}^\vee \otimes N_{2,L}^\vee & \text{if } a = 1 \text{ and } i \geq t+2, \\ 0 & \text{otherwise,} \end{cases}$$

and then the Leray spectral sequence shows that the conditions for being Ulrich are equivalent to the following vanishings for $1 \leq i \leq 3$:

$$\begin{aligned} H^i(C_2, F \otimes S^{t-i} E_{2,L}) &= H^{i-1}(C_2, F \otimes S^{t-i} E_{2,L}) = 0 \quad \text{when } i \leq t, \\ H^{i-1}(C_2, F \otimes S^{i-t-2} E_{2,L}^\vee \otimes N_{2,L}^\vee) &= H^{i-2}(C_2, F \otimes S^{i-t-2} E_{2,L}^\vee \otimes N_{2,L}^\vee) = 0 \quad \text{when } i \geq t+2. \end{aligned}$$

To get the conditions that we are looking for, we can apply Serre duality, together with the observation that $\omega_{C_2} \cong N_{2,\omega_C}$. $\square$

We can also control the effect of duality:

Lemma 5.15 *Let F be a vector bundle on C_2 such that $H^1(C+x, F) = H^1(C+x, F^\vee \otimes A_{2,L \otimes \omega_C}) = 0$ for all $x \in C$, and let $\mathcal{F} \cong \varepsilon_* p^* F \otimes \mathcal{O}_{\mathbb{P}^r}(t)$ for one $t \in \mathbb{Z}$. Then*

$$D(\varepsilon_* p^* F \otimes \mathcal{O}_{\mathbb{P}^r}(t)) \cong \varepsilon_* p^*(F^\vee \otimes A_{2,L \otimes \omega_C}) \otimes \mathcal{O}_{\mathbb{P}^r}(2-t).$$

Proof The hypothesis implies that $R^1 \varepsilon_*(p^* F) = R^1 \varepsilon_*(p^* F^\vee \otimes A_{2,\omega_C} \otimes L) = 0$ as in the proof of Lemma 5.13. Now we use Grothendieck duality: define $f : B \rightarrow \mathbb{P}^r$ as the composition $B \rightarrow \Sigma \hookrightarrow \mathbb{P}^r$. If $\mathcal{G}$ is a vector bundle on B , then

$$D(\varepsilon_* \mathcal{G}) = \mathcal{E}xt_{\mathbb{P}^r}^{-3}(f_* \mathcal{G}, \mathcal{O}_{\mathbb{P}^r}(3-r))$$

by definition. In particular, assume that $R^1 f_* \mathcal{G} = 0$, so that $f_* \mathcal{G} \cong Rf_* \mathcal{G}$ in the derived category $D^b(\mathbb{P}^r)$.

Then Grothendieck duality [Huybrechts 2006, Theorem 3.34] gives an isomorphism in $D^b(\mathbb{P}^r)$:

$$\begin{aligned} R\mathcal{H}om_{\mathbb{P}^r}(f_*\mathcal{G}, \mathcal{O}_{\mathbb{P}^r}(3-r))[r-3] &\cong Rf_*R\mathcal{H}om_B(\mathcal{G}, f^*\mathcal{O}_{\mathbb{P}^r}(3-r) \otimes \omega_B \otimes f^*\omega_{\mathbb{P}^r}^\vee) \\ &\cong Rf_*R\mathcal{H}om_B(\mathcal{G}, \omega_B \otimes \mathcal{O}_B(4)) \\ &\cong Rf_*R\mathcal{H}om_B(\mathcal{G}, p^*(S_{L \otimes \omega_C}(-\Delta)) \otimes \mathcal{O}_B(2)) \\ &\cong Rf_*(\mathcal{G}^\vee \otimes p^*(S_{L \otimes \omega_C}(-\Delta)) \otimes \mathcal{O}_B(2)), \end{aligned}$$

where we used the fact in Remark 5.9 that $\omega_B \cong p^*(S_{L \otimes \omega_C}(-\Delta)) \otimes \mathcal{O}_B(2)$. Since $R^1\varepsilon_*p^*F = 0$, we can apply this reasoning to $\mathcal{G} = p^*F$, and we obtain

$$\begin{aligned} R\mathcal{H}om(f_*p^*F, \mathcal{O}_{\mathbb{P}^r}(3-r))[r-3] &\cong Rf_*(p^*(F^\vee \otimes A_{2,L \otimes \omega_C})) \otimes \mathcal{O}_{\mathbb{P}^r}(2) \\ &\cong \varepsilon_*(p^*(F^\vee \otimes A_{2,L \otimes \omega_C})) \otimes \mathcal{O}_{\mathbb{P}^r}(2), \end{aligned}$$

where the last isomorphism comes from the fact that $R^1\varepsilon_*(p^*(F^\vee \otimes A_{2,L \otimes \omega_C})) = 0$. If we tensor both sides by $\mathcal{O}_{\mathbb{P}^r}(-t)$, we get what we want. $\square$

Now we can finally construct Ulrich sheaves on the secant variety:

Theorem 5.16 *Let α be a line bundle on C such that $H^0(C, \alpha) = H^1(C, \alpha) = 0$. Then the coherent sheaf*

$$\mathcal{F} = \varepsilon_*p^*A_{2,L \otimes \alpha}$$

is an Ulrich sheaf of rank one on Σ . If, furthermore, α is a theta characteristic, meaning that $\alpha \otimes \alpha \cong \omega_C$, then this is a symmetric Ulrich sheaf of rank one on Σ .

Proof To prove that the sheaf is Ulrich, we check the conditions of Lemma 5.14 for $F = A_{2,L \otimes \alpha}$. The first condition, $H^1(x + C, A_{2,L \otimes \alpha}|_{x+C}) = 0$ for each $x \in C$, becomes $H^1(C, \alpha \otimes L(-2x)) = 0$ for all $x \in C$. Since $h^0(C, L) \geq 3$ we see that $L(-2x)$ is effective, and since $H^1(C, \alpha) = 0$, it follows that $H^1(C, \alpha \otimes L(-2x)) = 0$ as well. Observe that in the notation of Lemma 5.14, we have $t = 0$, so that the second condition is empty. For the third condition we have to check some cohomology vanishings, which in our case are

$$\begin{aligned} H^0(C_2, S_{2, \omega_C \otimes \alpha^\vee} \otimes E_{2,L}) &= H^1(C_2, S_{2, \omega_C \otimes \alpha^\vee} \otimes E_{2,L}) = 0, \\ H^1(C_2, S_{2, \omega_C \otimes \alpha^\vee}) &= H^2(C_2, S_{2, \omega_C \otimes \alpha^\vee}) = 0. \end{aligned}$$

By Serre duality, we know that $H^0(C, \omega_C \otimes \alpha^\vee) = H^1(C, \omega_C \otimes \alpha^\vee) = 0$ and then the above vanishings follow from the formulas in Lemma 5.5. This proves that $\mathcal{F}$ is Ulrich. We can also compute $D(\mathcal{F})$ using Lemma 5.15: the other condition that we need to check is that $H^1(x + C, A_{2,L \otimes (\omega_C \otimes \alpha^\vee)}) = 0$ for all $x \in C$, but this can be proven as in the beginning of this proof. Then Lemma 5.15 shows that

$$D(\mathcal{F}) \cong \varepsilon_*p^*S_{2, \omega_C \otimes \alpha^\vee} \otimes \mathcal{O}_{\mathbb{P}^r}(2).$$

To conclude, we will show that this is isomorphic to $\varepsilon_* p^*(A_{2,\omega_C \otimes \alpha^\vee})$, so that if α is a theta characteristic, we are done thanks to Lemma 4.31. We see that

$$\varepsilon_* p^* S_{2,\omega_C \otimes \alpha^\vee} \otimes \mathcal{O}_{\mathbb{P}^r}(2) \cong \varepsilon_*(p^*(S_{\omega_C \otimes \alpha^\vee}) \otimes \mathcal{O}_B(2)) \cong \varepsilon_*(p^* A_{2,L \otimes \omega_C \otimes \alpha^\vee} \otimes \mathcal{O}_B(E)),$$

where the last isomorphism comes from Remark 5.9. Now consider the exact sequence on B

$$(5-6) \quad 0 \rightarrow p^* A_{2,L \otimes \omega_C \otimes \alpha^\vee} \rightarrow p^* A_{L \otimes \omega_C \otimes \alpha^\vee} \otimes \mathcal{O}_B(E) \rightarrow p^* A_{L \otimes \omega_C \otimes \alpha^\vee} \otimes \mathcal{O}_E(E) \rightarrow 0,$$

and observe that under the isomorphism $C \times C \rightarrow E$ given by $(x, y) \mapsto (\iota(x), x + y)$ it holds that

$$p^* A_{L \otimes \omega_C \otimes \alpha^\vee} \otimes \mathcal{O}_E(E) \cong p^*(S_{\omega_C \otimes \alpha^\vee}) \otimes \mathcal{O}_E(2) \cong \text{pr}_1^*(L^{\otimes 2} \otimes \omega_C \otimes \alpha^\vee) \otimes \text{pr}_2^*(\omega_C \otimes \alpha^\vee),$$

and since the composition $C \times C \rightarrow E \xrightarrow{f} \mathbb{P}^r$ is identified with the first projection $C \times C \rightarrow C$, we see that

$$\begin{aligned} Rf_*(p^* A_{L \otimes \omega_C \otimes \alpha^\vee} \otimes \mathcal{O}_E(E)) &\cong R\text{pr}_{1*}(\text{pr}_1^*(L^{\otimes 2} \otimes \omega_C \otimes \alpha^\vee) \otimes \text{pr}_2^*(\omega_C \otimes \alpha^\vee)) \\ &\cong (L^{\otimes 2} \otimes \omega_C \otimes \alpha^\vee) \otimes R\text{pr}_{1*} \text{pr}_2^*(\omega_C \otimes \alpha^\vee) = 0, \end{aligned}$$

where the last equality comes from Grauert's theorem, together with the fact that $\omega_C \otimes \alpha^\vee$ has no cohomology. At this point equation (5-6) shows that $\varepsilon_*(p^* A_{L \otimes \alpha'} \otimes \mathcal{O}(E)) \cong \varepsilon_*(p^* A_{L \otimes \alpha'})$. $\square$

Remark 5.17 Looking at the proofs in this section, we see that we do not need the full assumption of 3-very ampleness. What is actually needed is that $R^1 \varepsilon_* p^*(A_{L \otimes \alpha}) = 0$ along the map $\varepsilon : B \rightarrow \Sigma$. For example, the proof of Lemma 5.13 shows that this holds whenever the map $\varepsilon : B \rightarrow \Sigma$ has only zero-dimensional fibers over $\Sigma \setminus C$. Geometrically, this means that there are no infinitely secant lines passing through a point $x \in \mathbb{P}^r \setminus C$.

5.5 Rational normal curves

In the case of rational normal curves, we can also find Ulrich sheaves of higher rank. Thus, we now let $C = \mathbb{P}^1$ and $L = \mathcal{O}_{\mathbb{P}^1}(n)$, with $n \geq 3$. Recall that in this case there is a natural identification of the second symmetric product $(\mathbb{P}^1)_2$ with a projective space $\mathbb{P}^2$, so that the divisors $x + C \subset C_2$ are lines in $\mathbb{P}^2$ and the locus of nonreduced divisors $\Delta \subset (\mathbb{P}^1)_2$ is a smooth conic. In particular, we see that $S_{2,\mathcal{O}_{\mathbb{P}^1}(d)} \cong \mathcal{O}_{\mathbb{P}^2}(d)$ and $\mathcal{O}_{(\mathbb{P}^1)_2}(\delta) \cong \mathcal{O}_{\mathbb{P}^2}(1)$. We also recall that the quotient bundle Q on $\mathbb{P}^2$ is defined by the Euler exact sequence

$$(5-7) \quad 0 \rightarrow \mathcal{O}_{\mathbb{P}^2}(-1) \rightarrow \mathcal{O}_{\mathbb{P}^2}^{\oplus 3} \rightarrow Q \rightarrow 0,$$

and furthermore, $Q \cong T_{\mathbb{P}^2}(-1)$.

Theorem 5.18 *With the previous notation, the coherent sheaf*

$$\mathcal{F}_n = \varepsilon_* p^*(S^{n-2} Q \otimes \mathcal{O}_{\mathbb{P}^2}(-1)) \otimes \mathcal{O}_{\mathbb{P}^n}(1)$$

is an Ulrich sheaf of rank $n - 1$ on the secant Σ of the rational normal curve in $\mathbb{P}^n$. Furthermore, $D(\mathcal{F}_n) \cong \mathcal{F}_n$.

Proof To prove the sheaf is Ulrich, we check the conditions of Lemma 5.14 for $F = S^{n-2}Q \otimes \mathcal{O}_{\mathbb{P}^2}(-1)$. We first observe that in this case

$$(5-8) \quad F^\vee \otimes A_{2,\omega_C} \otimes L \cong S^{n-2}Q^\vee \otimes \mathcal{O}_{\mathbb{P}^2}(1) \otimes \mathcal{O}_{\mathbb{P}^2}(n-4) \cong S^{n-2}Q \otimes \mathcal{O}_{\mathbb{P}^2}(-1) \cong F,$$

where the last isomorphism comes from the fact that Q is a rank-two bundle of determinant $\mathcal{O}_{\mathbb{P}^2}(1)$, so that $Q^\vee \cong Q \otimes \mathcal{O}_{\mathbb{P}^2}(-1)$. Now we check the conditions of Lemma 5.14: for the first condition, we show that $H^1(\ell, F|_\ell) = 0$ for every line in $\mathbb{P}^2$, and this follows from the fact that $Q|_\ell \cong \mathcal{O}_\ell \oplus \mathcal{O}_\ell(1)$ for every line, so that $F|_\ell \cong \mathcal{O}_\ell(-1) \oplus \mathcal{O}_\ell \oplus \cdots \oplus \mathcal{O}_\ell(n-3)$.

For the other conditions, we observe that in the notation of Lemma 5.14, we have $t = 1$, so that, thanks to equation (5-8), we are reduced to the two vanishings

$$H^0(\mathbb{P}^2, S^{n-2}Q(-1)) = H^1(\mathbb{P}^2, S^{n-2}Q(-1)) = 0,$$

which hold because the bundle $S^{n-2}Q(-1)$ has no cohomology. This last statement follows from Borel–Weil–Bott in characteristic zero, but there is an easier proof valid in all characteristics, for which we thank Claudiu Raicu: taking symmetric powers in the exact sequence (5-7), we obtain another exact sequence

$$0 \rightarrow S^{n-3}(\mathcal{O}_{\mathbb{P}^2}^{\oplus 3}) \otimes \mathcal{O}_{\mathbb{P}^2}(-2) \rightarrow S^{n-2}(\mathcal{O}_{\mathbb{P}^2}^{\oplus 3}) \otimes \mathcal{O}_{\mathbb{P}^2}(-1) \rightarrow S^{n-2}Q \otimes \mathcal{O}_{\mathbb{P}^2}(-1) \rightarrow 0,$$

and since the line bundles $\mathcal{O}_{\mathbb{P}^2}(-1)$ and $\mathcal{O}_{\mathbb{P}^2}(-2)$ have no cohomology, the statement follows. This proves that $\mathcal{F}_n$ is Ulrich. The fact that $D(\mathcal{F}_n) \cong \mathcal{F}_n$ follows from Lemma 5.15 and equation (5-8). $\square$

6 The arithmetic writhe

6.1 An algebraic orientation on $B \setminus E$

We use the notation from Section 5.3. Namely, let L be a 3-very ample line bundle on the curve C over the field K and $\iota : C \hookrightarrow \mathbb{P}(H^0(C, L)^\vee) = \mathbb{P}^n$ be the corresponding embedding. We have the tautological projective bundle $p : B \rightarrow C_2$, i.e., $B = \mathbb{P}(E_{2,L})$ where $E_{2,L}$ is the tautological bundle (see Section 5) on the second symmetric power C_2 of C , and the projection $\varepsilon : B \rightarrow \Sigma$ to the secant variety Σ of C . The latter is an isomorphism outside the preimage E of C . Let α be a theta characteristic on C such that $H^0(C, \alpha) = 0$. We have seen in Theorem 5.16 that this defines a symmetric Ulrich sheaf of rank one on Σ , and consequently a relative orientation as in Theorem 4.28. We will now make this concrete and explain how α defines an algebraic orientation on $B \setminus E$, even when α has sections. To that end recall from Remark 5.9 that the canonical bundle on B is given by

$$p^*(S_{2,\omega_C}) \otimes \mathcal{O}_B(-E) \cong p^*(S_{2,\alpha}) \otimes p^*(S_{2,\alpha}) \otimes \mathcal{O}_B(-E).$$

This implies that $\omega_{B \setminus E}$ is isomorphic to the restriction of $p^*(S_{2,\alpha}) \otimes p^*(S_{2,\alpha})$. Thus $p^*(S_{2,\alpha})|_{B \setminus E}$ is a theta characteristic on $B \setminus E$. In order to even define an algebraic orientation on $B \setminus E$, choose an isomorphism $\sigma : \alpha \otimes \alpha \rightarrow \omega_C$ for now. Let τ be a rational section of α and let $\omega = \sigma(\tau \otimes \tau)$. Let ρ be a

rational section of $p^*(S_{2,\alpha})$ with divisor $S_{\text{div}(\tau)}$. Then by equation (5-4) an isomorphism

$$\psi : p^*(S_{2,\alpha})|_{B \setminus E} \otimes p^*(S_{2,\alpha})|_{B \setminus E} \rightarrow \omega_{B \setminus E}$$

can be defined by mapping $\rho \otimes \rho$ to the differential 3-form $\Omega(\omega)$ from Theorem 5.12. Note that by Corollary 5.10 another choice of σ , τ and ρ would lead to an equivalent algebraic orientation. Therefore, we can call ψ the *algebraic orientation induced by α* .

6.2 An arithmetic count of secant lines

Let $x_0, \dots, x_3$ be the homogeneous coordinates on $\mathbb{P}^3$ and $x = x_1/x_0$, $y = x_2/x_0$ and $z = x_3/x_0$. We consider on $\mathbb{P}^3$ the algebraic orientation defined by the rational differential 3-form $dx \wedge dy \wedge dz$ (meaning that we choose an isomorphism for which this 3-form is the image of a square). Now let $\iota : C \hookrightarrow \mathbb{P}^3$ be an embedded curve over K such that $L = \iota^* \mathcal{O}_{\mathbb{P}^3}(1)$ is 3-very ample. Geometrically, this means that C is the isomorphic image via a linear projection of a curve $C \subset \mathbb{P}^r$ such that no four points on C lie on a plane in $\mathbb{P}^r$.

Let α be a theta characteristic on C and let $\Sigma \subset \mathbb{P}^r$ be the secant variety to C . Composing the map $p : B \rightarrow \Sigma$ with a suitable projection and restricting to $B \setminus E$, we obtain the finite surjective morphism $p' : B \setminus E \rightarrow \mathbb{P}^3 \setminus C$. The algebraic orientation on $B \setminus E$ induced by α together with our fixed algebraic orientation on $\mathbb{P}^3$ define a relative orientation of p' and since $\mathbb{P}^3$ is $\mathbb{A}^1$ -chain connected we have that $\deg^{\mathbb{A}^1}(p')$ is well-defined; see [Pauli and Wickelgren 2021, Theorem 9] applied to the proper map $B \rightarrow \mathbb{P}^3$.

Definition 6.1 We define the *arithmetic writhe* of the 3-very ample curve C semi-oriented by α as $w(C, \alpha) := \deg^{\mathbb{A}^1}(p')$.

Remark 6.2 The arithmetic writhe is the result of an arithmetic count of secant lines S to C passing through a given point $q \in \mathbb{P}^3 \setminus C$. Indeed, such secant lines are in bijection to points q_S in the preimage of q under p' . We define the *local writhe* $w_{q,S}(C, \alpha)$ at such a secant to be the local $\mathbb{A}^1$ -degree $\deg_{q_S}^{\mathbb{A}^1}(p')$. The sum of the local writhe over all secants that contain a point q is independent of q and equals $w(C, \alpha)$. Note that $w(C, \alpha)$ *does*, however, depend on the embedded curve C (and α), as for example Proposition 6.6 shows.

We now describe how to compute the local writhe explicitly. By Theorem 5.12 $w_{q,S}(C, \alpha)$ can be computed by evaluating $dx \wedge dy \wedge dz$ (or any other differential 3-form that differs from this by the square of a rational function) at the ω -Viro frame at q_S , where ω is any differential on C whose divisor is of the form $2D$, where D is a divisor whose line bundle is α . More precisely, we choose tangent vectors $u \in T_q S$, $v \in T_a C$ and $w \in T_b C$ as in Definition 5.11. Namely, $v \in T_a C$ and $w \in T_b C$ are chosen in a way that ω evaluates to one in v and w . After identifying the secant S with $\mathbb{P}^1$ sending a to $(1 : 0)$, b to $(0 : 1)$ and q to $(1 : 1)$, the tangent vector u is chosen in a way that dt evaluates to one in u . Then write u , v and w as vectors $\tilde{u}$, $\tilde{v}$ and $\tilde{w}$ with respect to the basis given by dx , dy and dz . The local writhe then

is equal to

$$w_{q,S}(C, \alpha) = \text{Tr}_{F/K} \left\langle \det \begin{pmatrix} | & | & | \\ \tilde{v} & \tilde{u} & \tilde{w} \\ | & | & | \end{pmatrix} \right\rangle,$$

where F is the field of definition of S .

Remark 6.3 The above description of the arithmetic writhe as a sum of local writhe numbers shows that in the case $K = \mathbb{R}$ it agrees with the encomplexed writhe number introduced by Viro [2001]. Realizing it as the degree of a morphism gives another proof that the encomplexed writhe number does not depend on the choice of the center of projection.

Example 6.4 Let K a field with $\text{char}(K) \neq 2$ and consider the rational curve $C \subset \mathbb{P}^3$ of degree four defined as the image of

$$\varphi : \mathbb{P}^1 \rightarrow \mathbb{P}^3, \quad (r : s) \mapsto (r^4 : r^3 s : r s^3 : s^4).$$

In the following, every (local) writhe number will be computed with respect to the unique theta characteristic on $C \cong \mathbb{P}^1$ given by $\mathcal{O}_{\mathbb{P}^1}(-1)$. In order to compute the local writhe of secants to C , we choose the algebraic orientation σ on C given by $dt = t^2 d\bar{t}$, where $t = s/r$ and $\bar{t} = -r/s$. The genus–degree formula implies that the projection from a general point has three nodes. We will now compute the local writhe of all three secants that contain the point $P = (1 : 0 : 0 : 1)$. For the secant L_1 spanned by the points $a = \varphi(1 : 1)$ and $b = \varphi(1 : -1)$ we can work on the affine chart $x_0 \neq 0$ and the coordinates x, y, z . In this chart the curve is parametrized by

$$t \mapsto (t, t^3, t^4).$$

Expressed in these coordinates, our points a and b correspond to $\tilde{a} = (1, 1, 1)$ (when $t = 1$) and $\tilde{b} = (-1, -1, 1)$ (when $t = -1$). Tangent vectors at $\tilde{a}$ and $\tilde{b}$ in direction of t are given by $\tilde{v} = (1, 3, 4)$ and $\tilde{w} = (1, 3, -4)$. For computing the vector $\tilde{u}$, we have to consider the parametrization

$$\lambda \mapsto \left(\frac{1-\lambda}{1+\lambda}, \frac{1-\lambda}{1+\lambda}, 1 \right).$$

The tangent vector at $\lambda = 1$ equals $\tilde{u} = (-\frac{1}{2}, -\frac{1}{2}, 0)$. We can thus compute the local writhe as

$$w_{P,L_1}(C) = \left\langle \det \begin{pmatrix} 1 & -\frac{1}{2} & 1 \\ 3 & -\frac{1}{2} & 3 \\ 4 & 0 & -4 \end{pmatrix} \right\rangle = \langle -8 \rangle = \langle -2 \rangle.$$

Let $i \in \bar{K}$ be a square root of -1 . For the secant L_2 spanned by the points $a = \varphi(1 : i)$ and $b = \varphi(1 : -i)$ we can work on the same affine chart. Note that although a and b might not be K -rational points, the line L_2 is defined over K . We have $\tilde{a} = (i, -i, 1)$ (when $t = i$) and $\tilde{v} = (1, -3, -4i)$ as well as




Figure 4: The quartic curve considered in Example 6.4 projected from the point P . The nodes of this planar curve correspond to secant lines containing P and we denoted their local writhe in the picture.

$\tilde{b} = (-i, i, 1)$ (when $t = -i$) and $\tilde{w} = (1, -3, 4i)$. In this chart the line L_2 is parametrized by

$$\lambda \mapsto \left(i \cdot \frac{1-\lambda}{1+\lambda}, -i \cdot \frac{1-\lambda}{1+\lambda}, 1 \right).$$

Thus we compute $\tilde{u} = (-i/2, i/2, 0)$. Now we can compute the local writhe as

$$w_{p,L_2}(C) = \left\langle \det \begin{pmatrix} 1 & -\frac{1}{2}i & 1 \\ -3 & \frac{1}{2}i & -3 \\ -4i & 0 & 4i \end{pmatrix} \right\rangle = \langle 8 \rangle = \langle 2 \rangle.$$

For the third secant L_3 spanned by $\varphi(1:0)$ and $\varphi(0:1)$ we work on the affine chart $x_0 + x_3 \neq 0$ and the coordinates $x_1/(x_0 + x_3)$, $x_2/(x_0 + x_3)$, $x_3/(x_0 + x_3)$. A short computation verifies that these are compatible with our chosen algebraic orientation. In this chart, the curve is parametrized by

$$t \mapsto \left(\frac{t}{1+t^4}, \frac{t^3}{1+t^4}, \frac{t^4}{1+t^4} \right) \quad \text{and} \quad \bar{t} \mapsto \left(\frac{-\bar{t}^3}{1+\bar{t}^4}, \frac{-\bar{t}}{1+\bar{t}^4}, \frac{1}{1+\bar{t}^4} \right).$$

We have $\tilde{a} = (0, 0, 0)$ (when $t = 0$) and $\tilde{b} = (0, 0, 1)$ (when $\bar{t} = 0$). Tangent vectors at $\tilde{a}$ and $\tilde{b}$ in the direction of t and $\bar{t}$ are given by $\tilde{v} = (1, 0, 0)$ and $\tilde{w} = (0, -1, 0)$, respectively. Moreover, in this chart the line L_3 is parametrized by

$$\lambda \mapsto \left(0, 0, \frac{\lambda}{1+\lambda} \right),$$

which gives $\tilde{u} = (0, 0, \frac{1}{4})$. We can thus compute the local writhe as

$$w_{p,L_3}(C) = \left\langle \det \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & \frac{1}{4} & 0 \end{pmatrix} \right\rangle = \left\langle \frac{1}{4} \right\rangle = \langle 1 \rangle.$$

Summing up the local writhe numbers, we find that the arithmetic writhe of C is

$$w(C) = \langle -2 \rangle + \langle 2 \rangle + \langle 1 \rangle = 2 \cdot \langle 1 \rangle + \langle -1 \rangle.$$

If the theta characteristic α does not have global sections, then we can compute the arithmetic writhe of C directly from the Ulrich bundle constructed in Theorem 5.16.

Theorem 6.5 *Let L be a 3-very ample line bundle on the curve C and α a theta characteristic of C with $h^0(C, \alpha) = 0$. There exists a symmetric matrix Λ whose entries are linear forms on $\wedge^4 H^0(C, L)$ such that for all $s_0, \dots, s_3 \in H^0(C, L)$ it holds that:*

- (1) $\Lambda(s_0 \wedge \dots \wedge s_3)$ is nonsingular if and only if the linear system spanned by the $s_0, \dots, s_3$ is very ample.
- (2) If the linear system spanned by the $s_0, \dots, s_3$ is very ample, then the writhe $w(\psi(C), \alpha)$ is the class of $\Lambda(s_0 \wedge \dots \wedge s_3)$, where $\psi : C \rightarrow \mathbb{P}^3$ is the embedding defined by $s_0, \dots, s_3$.

Proof Let Σ be the secant variety of C embedded to $\mathbb{P}(H^0(C, L)^\vee)$. Then the linear system spanned by $s_0, \dots, s_3 \in H^0(C, L)$ is very ample if and only if $s_0, \dots, s_3$ do not have a common zero on Σ . If so, then $w(\psi(C), \alpha)$ is the $\mathbb{A}^1$ -degree of the linear projection $\Sigma \rightarrow \mathbb{P}^3$ given by $s_0, \dots, s_3$ restricted to the preimage of $\mathbb{P}^3 \setminus C$ relatively oriented by the rank-one symmetric Ulrich sheaf from Theorem 5.16. Hence the claim follows from Theorem 4.32. $\square$

We end this section with some thoughts on which classes $w \in \text{GW}(K)$ can be realized as the writhe of a spatial curve.

Proposition 6.6 *For every $a, b, c \in K^\times$ there is a rational curve $X \subset \mathbb{P}^3$ of degree four with arithmetic writhe $w(X) = \langle a \rangle + \langle b \rangle + \langle c \rangle$.*

Proof Let $C = \mathbb{P}^1$, $L = \mathcal{O}_{\mathbb{P}^1}(4)$ and $\alpha = \mathcal{O}_{\mathbb{P}^1}(-1)$. The matrix Λ from Theorem 6.5 is then given by the Hankel matrix

$$(6-1) \quad \Lambda = \begin{pmatrix} x_0 & x_1 & x_2 \\ x_1 & x_2 & x_3 \\ x_2 & x_3 & x_4 \end{pmatrix},$$

where x_i denotes the linear form

$$f \mapsto f \wedge (r^{4-i}s^i) \in \wedge^5 K[r, s]_4 = K$$

on $\wedge^4 K[r, s]_4$. On the other hand, for arbitrary $a, b, c \in K^\times$ we have

$$\begin{pmatrix} a + (b^2/c) & 0 & b \\ 0 & b & 0 \\ b & 0 & c \end{pmatrix} \cong \langle a \rangle + \langle b \rangle + \langle c \rangle.$$

This proves the claim. $\square$

Example 6.7 We have for instance

$$\Lambda(r^4 \wedge r^3 s \wedge r s^3 \wedge s^4) = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \cong 2 \cdot \langle 1 \rangle + \langle -1 \rangle,$$

which confirms our calculations from Example 6.4.

Question 1 Let $d \in \mathbb{N}$, $N = \frac{1}{2}(d-1)(d-2)$ and $a_1, \dots, a_N \in K^\times$. Is there a rational curve $X \subset \mathbb{P}^3$ of degree d with $w(X) = \sum_{i=1}^N \langle a_i \rangle$? More generally, for any $N \in \mathbb{N}$ and $a_1, \dots, a_N \in K^\times$, does there exist a curve $X \subset \mathbb{P}^3$ and a theta characteristic α on X with $w(X, \alpha) = \sum_{i=1}^N \langle a_i \rangle$?

7 Algebraic isotopies

Let C be a curve over a field K of characteristic $\text{char}(K) \neq 2$. In analogy to classical knot and link theory, we want to study embeddings of C into the three-dimensional space $\mathbb{P}^3$ up to algebraic isotopies. We first make some general definitions and observations.

Definition 7.1 Let X be a projective variety over K . An *algebraic isotopy* of two closed embeddings $\psi_i : X \rightarrow \mathbb{P}^r$, $i = 0, 1$, is a morphism

$$I : \mathbb{A}^1 \times X \rightarrow \mathbb{P}^r, \quad (t, x) \mapsto I_t(x),$$

such that for all $t \in \mathbb{A}^1$ the induced map $I_t : X_{\kappa(t)} \rightarrow Y_{\kappa(t)}$ is an embedding and $I_t = \psi_i$ for $i = 0, 1$. Two embeddings ψ, ψ' are *algebraically isotopic* if they are connected by a chain of algebraic isotopies. We denote by $[X, \mathbb{P}^r]^1$ the quotient of the set of embeddings $\text{Emb}(X, \mathbb{P}^r)(K)$ by the equivalence relation generated by algebraic isotopies. We write $[\psi]$ for the class of ψ in $[X, \mathbb{P}^r]^1$ and call it the *isotopy class* of ψ .

A first observation is that two isotopic embeddings have isomorphic underlying line bundles:

Lemma 7.2 Suppose that $\psi, \psi' : X \hookrightarrow \mathbb{P}^r$ are two algebraically isotopic embeddings. Then

$$\psi^* \mathcal{O}_{\mathbb{P}^r}(1) \cong \psi'^* \mathcal{O}_{\mathbb{P}^r}(1).$$

Proof Any algebraic isotopy $I : \mathbb{A}^1 \times X \rightarrow \mathbb{P}^r$ induces a morphism $I : \mathbb{A}^1 \rightarrow \text{Pic}(X)$, $t \mapsto I_t^* \mathcal{O}_{\mathbb{P}^r}(1)$, which is constant, since any morphism from a rational curve to an abelian variety is constant. Alternatively, the map

$$\text{pr}_X^* : \text{Pic}(X) \rightarrow \text{Pic}(\mathbb{A}^1 \times X), \quad L \mapsto \text{pr}_X^* L,$$

is an isomorphism [Hartshorne 1977, Proposition II.6.6], so that the isomorphism class of the line bundle $I^* \mathcal{O}_{\mathbb{P}^r}(1)$ is constant on the fibers $\{t\} \times X$. $\square$

Definition 7.3 Let X be a projective variety over K and L a line bundle on X . We denote the set of embeddings $\psi : X \hookrightarrow \mathbb{P}^r$ with $\psi^* \mathcal{O}_{\mathbb{P}^r}(1) \cong L$ by $\text{Emb}_L(X, \mathbb{P}^r)(K)$, and its quotient by chains of algebraic isotopies by $[X, \mathbb{P}^r]_L^1$.

We now recall how to parametrize the set $\text{Emb}_L(X, \mathbb{P}^r)(K)$. First of all, if this is nonempty, then L is very ample. In this case, consider the embedding given by the complete linear system $V = H^0(X, L)$,

$$\varphi_L : X \hookrightarrow \mathbb{P}^n = \mathbb{P}(V^\vee),$$

and the secant variety $\Sigma = \Sigma(X, L)$ of X in $\mathbb{P}(V^\vee)$. Any embedding ψ from $\text{Emb}_L(X, \mathbb{P}^r)$ is obtained

by composing φ_L with a linear projection

$$[s_0, \dots, s_r] : \varphi_L(X) \rightarrow \mathbb{P}^r, \quad \text{where } s_i \in V,$$

and, in order for $\psi = [s_0, \dots, s_r] \circ \varphi_L$ to be an embedding, the condition is that the linear space $\Pi = \{s_0 = \dots = s_r = 0\}$ does not intersect the secant variety Σ . Consider now the set

$$(7-1) \quad \mathcal{U}_{L,r} = \{(s_0, \dots, s_r) \in V^{\oplus(r+1)} \mid \{s_0 = \dots = s_r = 0\} \cap \Sigma = \emptyset\}.$$

This is open: its complement in $V^{\oplus(r+1)}$ is the image of the incidence correspondence

$$\mathcal{I} = \{((s_0, \dots, s_r), p) \in V^{\oplus(r+1)} \times \Sigma \mid s_0(p) = \dots = s_r(p) = 0\}$$

along the proper map $\text{pr}_{V^{\oplus(r+1)}} : V^{\oplus(r+1)} \times \Sigma \rightarrow V^{\oplus(r+1)}$. The previous discussion shows that $\text{Emb}_L(X, \mathbb{P}^r)(K)$ is the set of K -rational points of a (possibly empty) open subset of a projective space

$$(7-2) \quad \text{Emb}_L(X, \mathbb{P}^r) = \mathcal{U}_{L,r}/K^\times = \{[s_0, \dots, s_r] \in \mathbb{P}(V^{\oplus(r+1)}) \mid (s_0, \dots, s_r) \in \mathcal{U}_{L,r}\}.$$

This comes equipped with a universal embedding

$$\text{Emb}_L(X, \mathbb{P}^r) \times X \rightarrow \mathbb{P}^r, \quad ([s_0, \dots, s_r], x) \mapsto [s_0(x), \dots, s_r(x)],$$

so that an algebraic isotopy between $\psi, \psi' \in \text{Emb}_L(X, \mathbb{P}^r)(K)$ is simply a map $\mathbb{A}^1 \rightarrow \text{Emb}_L(X, \mathbb{P}^r)$ whose image contains these two points. This leads naturally to the notion of naive connected components in $\mathbb{A}^1$ -geometry, as for instance in [Cazanave 2012].

Definition 7.4 Let X be a scheme over the field K . The set $\pi_0^N X(K)$ of *naive connected components* of $X(K)$ is the quotient of $X(K)$ by the finest equivalence relation under which any two points $x_0, x_1 \in X(K)$ for which there exists a morphism $f : \mathbb{A}^1 \rightarrow X$ with $f(i) = x_i$ for $i = 0, 1$ are equivalent.

Remark 7.5 The previous discussion shows that the set of isotopy classes of embeddings with line bundle L can be considered as the set of naive connected components

$$[X, \mathbb{P}^r]_L^1 = \pi_0^N \text{Emb}_L(X, \mathbb{P}^r)(K).$$

We could also replace $\mathbb{P}^r$ by any other projective variety Y by using Grothendieck's Hom-scheme, although we will not consider this here.

Remark 7.6 Recall that a K -scheme X is $\mathbb{A}^1$ -connected if and only if for any extension K'/K , any two K' -rational points in X can be connected by a chain of maps from $\mathbb{A}_{K'}^1$. This means precisely that $\pi_0^N \text{Emb}_L(X_{K'}, \mathbb{P}^r)(K')$ has at most one element for any extension K'/K .

Example 7.7 The group scheme SL_{r+1} is $\mathbb{A}^1$ -connected: indeed, for any field K , any matrix in $\text{SL}_{r+1}(K)$ can be written as a product of upper- or lower-triangular elementary matrices with all ones on the diagonal. For such a matrix A , the map

$$\mathbb{A}^1 \rightarrow \text{SL}_{r+1}, \quad t \mapsto I_{r+1} + t(A - I_{r+1}),$$

connects A at $t = 1$ with the identity I_{r+1} at $t = 0$.

Corollary 7.8 Let $\psi : X \hookrightarrow \mathbb{P}^r$ be an embedding and let $A \in \mathrm{SL}_{r+1}(K)$. The two embeddings ψ and $A \circ \psi$ are algebraically isotopic.

Proof Since SL_{r+1} is $\mathbb{A}^1$ -connected, we can find a chain of algebraic isotopies connecting A with the identity I_{r+1} . $\square$

Lemma 7.9 Let $\psi : X \hookrightarrow \mathbb{P}^r$ be an embedding and let $A \in K^{(r+1) \times (r+1)}$ be of nonmaximal rank, such that the composition of ψ with the rational map $A : \mathbb{P}^r \dashrightarrow \mathbb{P}^r$ is still an embedding. Then ψ and $A \circ \psi$ are algebraically isotopic.

Proof Let $\mathrm{rk}(A) = m + 1$ with $m < r$. Then there is an $M \in \mathrm{GL}_{r+1}(K)$ such that MA is in reduced row-echelon form:

$$M \cdot A = \begin{pmatrix} A' \\ 0 \end{pmatrix},$$

where A' is an $(m+1) \times (r+1)$ matrix. If we let $M' \in \mathrm{GL}_{r-m}(K)$ be any matrix with determinant $\det(M') = \det(M)^{-1}$, we see that

$$N = \begin{pmatrix} I_{m+1} & 0 \\ 0 & M' \end{pmatrix} \quad \text{satisfies} \quad NM \in \mathrm{SL}_{r+1}(K) \quad \text{and} \quad NM \cdot A = \begin{pmatrix} A' \\ 0 \end{pmatrix}.$$

Thus, thanks to Corollary 7.8, we can assume that A has reduced row-echelon form. For simplicity, let's assume that the pivots of A are in the first $m + 1$ columns, i.e.,

$$A = \begin{pmatrix} I_{m+1} & B \\ 0 & 0 \end{pmatrix}$$

for another $(m+1) \times (r-m)$ matrix B , with the general case being similar. There is an algebraic isotopy

$$\mathbb{A}^1 \times X \rightarrow \mathbb{P}^r, \quad (t, x) \mapsto \begin{pmatrix} I_{m+1} & B \\ 0 & tI_{r-m} \end{pmatrix} \cdot \psi(x),$$

between $A \circ \psi$ at $t = 0$ and another embedding ψ' at $t = 1$. To conclude, we need to show that ψ' is algebraically isotopic to ψ . It is enough to observe that

$$\begin{pmatrix} I_{m+1} & -B \\ 0 & I_{r-m} \end{pmatrix} \cdot \begin{pmatrix} I_{m+1} & B \\ 0 & I_{r-m} \end{pmatrix} \cdot \psi(x) = \begin{pmatrix} I_{m+1} & 0 \\ 0 & I_{r-m} \end{pmatrix} \cdot \psi(x) = \psi(x),$$

so that we conclude by invoking Corollary 7.8 again. $\square$

The following corollary implies the first part of Theorem D.

Corollary 7.10 Let L be a line bundle on X with $h^0(X, L) \leq r$. Any two embeddings $\psi, \psi' : X \hookrightarrow \mathbb{P}^r$ in $\mathrm{Emb}_L(X, \mathbb{P}^r)$ are algebraically isotopic. Equivalently, $[X, \mathbb{P}^r]_L^1$ has at most one element.

Proof If L is not very ample, there is nothing to prove. If L is very ample, consider the embedding $\varphi_V : X \hookrightarrow \mathbb{P}(V^\vee)$. Since $\dim \mathbb{P}(V^\vee) < r$, we can compose this with a linear embedding $j : \mathbb{P}(V^\vee) \hookrightarrow \mathbb{P}^r$ to get another embedding $\psi'' = j \circ \varphi_V$. By construction, there is an $A \in K^{(r+1) \times (r+1)}$ not of maximal rank, such that $\psi = A \circ \psi''$. Then Lemma 7.9 shows that the resulting embedding ψ'' and ψ are algebraically isotopic. The same reasoning holds for ψ' . $\square$

Corollary 7.11 Consider an embedding $\psi : X \hookrightarrow \mathbb{P}^r$ in $\text{Emb}_L(X, \mathbb{P}^r)$ and let $\langle \psi(X) \rangle$ be the linear span. Consider also the secant variety Σ of $\varphi_L(X)$ in $\mathbb{P}(V^\vee)$.

- (1) It holds that $\dim \langle \psi(X) \rangle \geq \dim \Sigma$.
- (2) If K is infinite, ψ is algebraically isotopic to an embedding $\psi' : X \hookrightarrow \mathbb{P}^r$ such that $\dim \langle \psi'(X) \rangle = \dim \Sigma$.

Proof By the description in (7-2), we know that $\psi = [s_0, \dots, s_r]$, with $\Pi = \{s_0 = \dots = s_r = 0\} \subset \mathbb{P}(V^\vee)$ disjoint from Σ . Furthermore, the dimension of the linear span $\langle \psi(X) \rangle$ is

$$\dim \langle \psi(X) \rangle = \dim \langle s_0, \dots, s_r \rangle - 1 = \text{codim}((\Pi, \mathbb{P}(V^\vee))) - 1,$$

where $\langle s_0, \dots, s_r \rangle \subset V$ is the vector subspace generated by the s_i .

- (1) Since $\Pi \cap \Sigma = \emptyset$, it must be that $\text{codim}(\Pi, \mathbb{P}(V^\vee)) \geq \dim \Sigma + 1$. Then $\dim \langle \psi(X) \rangle \geq \dim \Sigma$.
- (2) Assume that $\dim \langle \psi(X) \rangle > \dim \Sigma$. Since the image of the map

$$[s_0, \dots, s_r] : \Sigma \rightarrow \mathbb{P}^r$$

is the secant variety Σ' to $\psi(X)$, we see that $\dim \Sigma' < \dim \langle \psi(X) \rangle$. Since K is infinite, we can find a K -point in $\langle \psi(X) \rangle \setminus \Sigma'$, and projecting from this point onto a hyperplane $H \subset \langle \psi(X) \rangle$ we get an embedding $\psi' : X \hookrightarrow H$, which is algebraically isotopic to ψ , thanks to Lemma 7.9. By repeating this procedure, we conclude. $\square$

In the following, we fix a very ample line bundle L on X , the space of global sections $V = H^0(X, L)$ and we consider the secant variety $\Sigma(X, L)$ of X inside $\mathbb{P}(V^\vee)$. We set

$$b = \dim \Sigma(X, L).$$

Remark 7.12 One always has

$$b \leq \min\{\dim \mathbb{P}(V^\vee), 2 \cdot \dim X + 1\},$$

and when the equality holds, the secant variety has expected dimension. This happens for example if X is a curve [Lange 1984] or if X is arbitrary and the line bundle L is 3-very ample.

Now we want to strengthen Corollary 7.10 and show that, if K is infinite, the isotopy classes of embeddings $X \hookrightarrow \mathbb{P}^r$ with $r > b$ are trivial. This is analogous to the fact that any knot can be unraveled in dimension higher than three.

Proposition 7.13 Assume that K is infinite and let $r > b$. Then any two embeddings $\psi, \psi' : X \hookrightarrow \mathbb{P}^r$ in $\text{Emb}_L(X, \mathbb{P}^r)$ are algebraically isotopic. Hence, if $r \neq b$, the set $[X, \mathbb{P}^r]_L^1$ has at most one element. In particular, if X is a curve, then $[X, \mathbb{P}^r]_L^1$ has at most one element whenever $r \geq 4$.

Proof Since K is infinite, we can use Corollary 7.11 and assume that both ψ and ψ' have images contained in a linear subspace of dimension b . Because the group $\text{SL}_{r+1}(K)$ acts transitively on these

subspaces, we can use Corollary 7.8 and assume that this subspace is the one where all coordinates are zero, apart from the first $b + 1$. This means

$$\psi = [s_0, \dots, s_b, 0, \dots, 0] \quad \text{and} \quad \psi' = [s'_0, \dots, s'_b, 0, \dots, 0]$$

for some $s_i, s'_i \in V$. Observe that in the above expressions there is a positive number of zeroes, since we are assuming $b < r$. We claim that there is an $s''_0 \in V$ such that

$$[s''_0, s_1, \dots, s_b, 0, \dots, 0], [s''_0, s'_1, \dots, s'_b, 0, \dots, 0]: X \hookrightarrow \mathbb{P}^r$$

are both embeddings. Assume for a moment that this claim is true: then Lemma 7.9 shows that both ψ and $[s''_0, s_1, \dots, s_b, 0, \dots, 0]$ are algebraically isotopic to the embedding given by

$$[s_0, s_1, \dots, s_b, s''_0, 0, \dots, 0]: X \hookrightarrow \mathbb{P}^r.$$

Notice that this is the step where we use that $b < r$. An analogous statement holds for ψ' , so that we can assume that ψ and ψ' have the form

$$\psi = [s''_0, s_1, \dots, s_b, 0, \dots, 0] \quad \text{and} \quad \psi' = [s''_0, s'_1, \dots, s'_b, 0, \dots, 0],$$

and by repeating this reasoning with other sections $s''_1, \dots, s''_b$, we conclude. We are left with proving the above claim. Consider the set

$$\begin{aligned} U &= \{s''_0 \in V \mid [s''_0, s_1, \dots, s_b, 0, \dots, 0] \text{ is an embedding}\} \\ &= \{s''_0 \in V \mid \{s''_0 = s_1 = \dots = s_b = 0\} \cap \Sigma = \emptyset\}. \end{aligned}$$

One can show as for the set $\mathcal{U}_{L,r}$ of (7-1) that U is open in V , and since $s_0 \in U$ and V is irreducible, U is a dense open subset. The same reasoning shows that

$$U' = \{s''_0 \in V \mid [s''_0, s'_1, \dots, s'_b, 0, \dots, 0] \text{ is an embedding}\}$$

is a dense open subset, and since K is infinite, there must be a K -rational point in $U \cap U'$, which is the section s''_0 that we are looking for.

The last statement follows from what we have just proved for $r > b$ and from Corollary 7.11 for $r < b$. $\square$

In light of Proposition 7.13, we will only consider the classes of embeddings

$$X \hookrightarrow \mathbb{P}^b, \quad \text{where } b = \dim \Sigma.$$

By Corollary 7.11, any such embedding is nondegenerate, meaning that if $\psi = [s_0, \dots, s_b]$, the sections s_i must be linearly independent. Now, consider the morphism of affine spaces

$$w: \mathbb{A}(V^{\oplus(b+1)}) \rightarrow \mathbb{A}(\wedge^{b+1} V), \quad (s_0, \dots, s_b) \mapsto s_0 \wedge \dots \wedge s_b,$$

and let $\mathcal{C}_\Sigma$ be the Chow form of the secant variety: this means that $\Pi = \{s_0 = \dots = s_b = 0\}$ intersects Σ if and only if $\mathcal{C}_\Sigma(s_0 \wedge \dots \wedge s_b) = 0$. Define $\mathcal{W}_{L,b} \subset \mathbb{A}(\wedge^{b+1} V)$ to be the intersection of the principal open

affine set $\{\mathcal{C}_\Sigma \neq 0\}$ with the locus of decomposable multivectors: we have that $\mathcal{U}_{L,b} = w^{-1}(\mathcal{W}_{L,b})$, where $\mathcal{U}_{L,b}$ is the open subset in (7-1). We have seen in (7-2) and Remark 7.5, that $\text{Emb}_L(X, \mathbb{P}^b) = \mathcal{U}_{L,b}/K^\times$ and that $[X, \mathbb{P}^b]_L^1 = \pi_0^N \text{Emb}_L(X, \mathbb{P}^b)(K)$. Thus, we have a surjective map

$$(7-3) \quad \pi_0^N \mathcal{U}_{L,b}(K) \rightarrow [X, \mathbb{P}^b]_L^1, \quad (\sigma_0, \dots, \sigma_b) \mapsto \psi = [\sigma_0, \dots, \sigma_b].$$

Our next goal is to show that it factors through the map $\pi_0^N \mathcal{U}_{L,b}(K) \rightarrow \pi_0^N \mathcal{W}_{L,b}(K)$ induced by w .

Lemma 7.14 *Let $0 \leq d < h^0(X, L)$ and $f : \mathbb{A}^1 \rightarrow \mathbb{A}(\wedge^{d+1} V)$ be a morphism whose image is contained in the locus of decomposable multivectors. There exists a morphism $g : \mathbb{A}^1 \rightarrow \mathbb{A}(V^{\oplus d+1})$ such that $f = w \circ g$:*

$$\begin{array}{ccc} & \mathbb{A}(V^{\oplus d+1}) & \\ \nearrow \exists g & \downarrow w & \\ \mathbb{A}^1 & \xrightarrow{f} & \mathbb{A}(\wedge^{d+1} V) \end{array}$$

Proof Set $\ell = h^0(X, L)$. Choosing a basis, we can identify $V = K^\ell$. We let M be the subsets of $\{1, \dots, \ell\}$ with $d+1$ elements. For $J = \{j_0, \dots, j_d\}$ with $1 \leq j_0 < \dots < j_d \leq \ell$ we write $e_J = e_{j_0} \wedge \dots \wedge e_{j_d}$, where $e_i \in K^\ell$ is the i -th unit vector. The morphism f is given by

$$f(t) = \sum_{J \in M} P_J(t) \cdot e_J$$

for some polynomials $P_J(t) \in K[t]$. Let $G(t) \in K[t]$ be the greatest common divisor of the $P_J(t)$. Since the image is contained in the locus of decomposable multivectors, we can write

$$f(t) = v_0(t) \wedge \dots \wedge v_d(t)$$

for some $v_j(t) \in K(t)^\ell$. We need to show that one can even choose $v_j \in K[t]^\ell$. Let $A(t)$ be the $\ell \times (d+1)$ matrix whose columns are the $v_j(t)$. Let $Q(t) \in K[t]$ be the smallest common denominator of the entries of A and let $A'(t) = Q(t) \cdot A(t)$. Because $K[t]$ is a principal ideal domain, there exist matrices $S(t) \in \text{SL}_\ell(K[t])$ and $T(t) \in \text{SL}_{d+1}(K[t])$ such that

$$S(t) \cdot A'(t) \cdot T(t) = \begin{pmatrix} D(t) \\ 0 \end{pmatrix},$$

where $D(t)$ is a $(d+1) \times (d+1)$ diagonal matrix with diagonal entries $p_0, \dots, p_d \in K[t]$. Suppose that $w_0(t), \dots, w_d(t)$ are the first $d+1$ columns of the matrix $S(t)^{-1}$. Then

$$Q(t) \cdot f(t) = p_0(t) \cdots p_d(t) \cdot (w_0(t) \wedge \dots \wedge w_d(t)).$$

Since $S(t)^{-1}$ is invertible over $K[t]$, the greatest common divisor of the Plücker coordinates of the right-hand side is $p_0(t) \cdots p_d(t)$. By construction, the greatest common divisor of the Plücker coordinates of the left-hand side is $G(t) \cdot Q(t)$. Thus, there is some $\lambda \in K^\times$ such that

$$f(t) = \lambda \cdot G(t) \cdot (w_0(t) \wedge \dots \wedge w_d(t)).$$

This proves the claim. □

Corollary 7.15 *The map $\pi_0^N \mathcal{U}_{L,b}(K) \rightarrow \pi_0^N \mathcal{W}_{L,b}(K)$ induced by w is a bijection.*

Proof It is surjective because w is surjective on the level of K -points. For proving injectivity let $x, y \in \mathcal{U}_{L,b}(K)$ be such that $w(x)$ and $w(y)$ belong to the same equivalence class in $\pi_0^N \mathcal{W}_{L,b}(K)$. By Lemma 7.14 there exists y' in the equivalence class of y such that $w(y') = w(x)$. The K -points of any fiber of w over a nonzero point form an $\mathrm{SL}_{b+1}(K)$ -orbit. Therefore, by Example 7.7, the two points x and y' , and thus x and y , belong to the same equivalence class in $\pi_0^N \mathcal{U}_{L,b}(K)$. $\square$

Corollary 7.16 *Let $(s_0, \dots, s_n), (s'_0, \dots, s'_n) \in \mathcal{U}_{L,b}(K)$. Then the embeddings*

$$\psi, \psi' : X \hookrightarrow \mathbb{P}^b,$$

defined respectively by the s_i and the s'_i , are algebraically isotopic if and only if there exists $\lambda \in K^\times$ such that $s_0 \wedge \dots \wedge s_b$ and $\lambda^{b+1} \cdot (s'_0 \wedge \dots \wedge s'_{b+1})$ are in the same naive connected component of $\mathcal{W}_{L,b}(K)$. In other words, there is a natural bijection

$$\pi_0 \mathcal{W}_{L,b}(K) / (K^{\times b+1}) \xrightarrow{\sim} [X, \mathbb{P}^b]_L^1.$$

Proof Since $\mathrm{Emb}_L(X, \mathbb{P}^b)(K) = \mathcal{U}_{L,b}(K) / K^\times$, we see that ψ and ψ' are algebraically isotopic if and only if and only if there exists $\lambda \in K^\times$ such that $(s_0, \dots, s_b)$ and $(\lambda s'_0, \dots, \lambda s'_b)$ are in the same naive connected component of $\mathcal{U}_{L,b}(K)$. By Corollary 7.15 this is the case if and only if $s_0 \wedge \dots \wedge s_b$ and $\lambda^{b+1} \cdot (s_0 \wedge \dots \wedge s_{b+1})$ are in the same naive connected component of $\mathcal{W}_{L,b}(K)$. $\square$

The following corollary implies the second part of Theorem D.

Corollary 7.17 *Assume that $h^0(X, L) = b + 1$. There is a bijection*

$$K^\times / K^{\times b+1} \rightarrow [X, \mathbb{P}^b]_L^1.$$

Proof Since $\dim(V) = b + 1$, we have that $\mathcal{W}_{L,b} \cong \mathbb{A}^1 \setminus \{0\}$. This shows that every point of $\mathcal{W}_{L,b}(K)$ is its own naive connected component, and $\mathcal{W}_{L,b}(K)$ can be identified with $K^\times$. The statement follows then from Corollary 7.16. $\square$

8 Isotopic embeddings of curves in 3-space

In the following, C will be a smooth and connected projective curve over K . It follows for example from Corollary 7.11 that the secant variety of C has the expected dimension three if and only if C is not planar. Thus, in the following we will consider embeddings

$$\psi : C \hookrightarrow \mathbb{P}^3$$

such that $L = \psi^* \mathcal{O}_{\mathbb{P}^3}(1)$ has $h^0(C, L) \geq 4$ and the isotopy classes in $[C, \mathbb{P}^3]_L^1$, so that the results of the previous section apply, in particular Corollary 7.16.

In the case of curves, we will prove that the arithmetic writhe is invariant under algebraic isotopies. We will make use of the following version of Harder's theorem.

Theorem 8.1 Let $n \in \mathbb{N}$ and $\mathcal{S}_n$ the scheme of nonsingular symmetric $n \times n$ matrices over K . We consider the fiber product $\mathrm{GW}(K) \times_{K^\times/K^{\times 2}} K^\times$ with respect to the canonical map $K^\times \rightarrow K^\times/K^{\times 2}$ and the discriminant map $\mathrm{GW}(K) \rightarrow K^\times/K^{\times 2}$. There is a well-defined injection

$$(8-1) \quad \pi_0^N \mathcal{S}_n(K) \rightarrow \mathrm{GW}(K) \times_{K^\times/K^{\times 2}} K^\times$$

that takes the naive connected component of a symmetric matrix to the pair of its class in $\mathrm{GW}(K)$ and its determinant.

Proof See for example [Cazanave 2012, Proposition 3.9]. $\square$

Recall that the writhe is constructed as follows: let α be a theta characteristic of C with $h^0(C, \alpha) = 0$ and let $\psi : C \hookrightarrow \mathbb{P}^3$ be an embedding such that $L = \psi^* \mathcal{O}_{\mathbb{P}^3}(1)$ is 3-very ample. The corresponding Ulrich sheaf $\mathcal{F}_\alpha$ of Theorem 6.5 provides a symmetric matrix Λ of linear forms on $\wedge^4 V$. This defines a morphism $\mathcal{W} \rightarrow \mathcal{S}_N$, where N is the degree of the secant variety. Taking the naive connected component, and using equation (8-1), this gives us a map

$$(8-2) \quad w(-, \alpha) : \pi_0^N \mathcal{W}_{L,3}(K) \rightarrow \mathrm{GW}(K) \times_{K^\times/K^{\times 2}} (K^\times).$$

By Corollary 7.16 this induces a map

$$(8-3) \quad w(-, \alpha) : [C, \mathbb{P}^3]_L^I \rightarrow \mathrm{GW}(K) \times_{K^\times/K^{\times 2}} (K^\times/K^{\times 4N}).$$

Definition 8.2 In the above setting, we define the *arithmetic writhe* of the embedding ψ to be the image $w([\psi], \alpha)$ of the isotopy class of ψ under the map from equation (8-3).

Remark 8.3 The arithmetic writhe of the embedding is automatically invariant under algebraic isotopies. Furthermore, the first component of $w([\psi], \alpha)$ is the arithmetic writhe of the embedded curve $\psi(C)$. In particular, this shows that if $\psi, \psi' : C \hookrightarrow \mathbb{P}^3$ are two algebraically isotopic embeddings, then the writhes of $\psi(C)$ and $\psi'(C)$ are the same.

We also notice that, if K is an algebraically or real closed field, then we have $K^{\times 4N} = K^{\times 2}$, which implies that the second factor of $w([\psi], \alpha)$ is superfluous: the arithmetic writhe of the embedding ψ coincides with the arithmetic writhe of the embedded curve $\psi(C)$.

8.1 Rational curves of degree four

In this subsection we consider the case $C = \mathbb{P}^1$. As $\alpha = \mathcal{O}_{\mathbb{P}^1}(-1)$ is the only theta characteristic of $\mathbb{P}^1$, we write $w(-) = w(-, \alpha)$. For $n \geq 3$ we write $[\mathbb{P}^1, \mathbb{P}^3]_n^I = [\mathbb{P}^1, \mathbb{P}^3]_{\mathcal{O}_{\mathbb{P}^1}(n)}^I$. Our goal is to prove that

$$w : [\mathbb{P}^1, \mathbb{P}^3]_4^I \rightarrow \mathrm{GW}(K) \times_{K^\times/K^{\times 2}} (K^\times/K^{\times 12})$$

is injective in this case. In other words, the isotopy type of two embeddings of $\mathbb{P}^1$ to $\mathbb{P}^3$ of degree four is completely determined by the arithmetic writhe. To this end, we will prove that $\pi_0^N \mathcal{W}_{\mathcal{O}_{\mathbb{P}^1}(4),3}(K)$ is in bijection to the set of pointed naive homotopy classes of morphisms $\mathbb{P}^1 \rightarrow \mathbb{P}^1$ of degree three. These were characterized in [Cazanave 2012]. We first have to introduce some notation. Following [Cazanave

2012, Definition 2.1] we consider the scheme $\mathcal{T}_n$, $n \in \mathbb{N}$, of pointed rational functions on $\mathbb{P}^1$, realized as the open subscheme of the affine space

$$\mathbb{A}^{2n} = \operatorname{Spec} K[a_0, \dots, a_{n-1}, b_0, \dots, b_{n-1}],$$

where the resultant of the two polynomials $t^n + a_{n-1}t^{n-1} + \dots + a_0$ and $b_{n-1}t^{n-1} + \dots + b_0$ does not vanish. Hence $\mathcal{T}_n(K)$ is the set of all pairs (f, g) of coprime polynomials $f, g \in K[t]$ where f is monic of degree n and g has degree at most $n-1$. For such a pair one can write

$$\frac{g}{f} = h_0(f, g) \cdot t^{-1} + h_1(f, g) \cdot t^{-2} + h_2(f, g) \cdot t^{-3} + \dots$$

for suitable $h_i(f, g) \in K$, and we define the *Hankel matrix* of f and g as

$$H_n(f, g) = (h_{i+j}(f, g))_{0 \leq i, j \leq n-1}.$$

A related matrix is the *Bézout matrix*, which is constructed as follows. The polynomial $f(x)g(y) - f(y)g(x)$ is divisible by $x - y$, so we can write

$$\frac{f(x)g(y) - f(y)g(x)}{x - y} = \sum_{1 \leq i, j \leq n} c_{ij} x^{i-1} y^{j-1}$$

for $c_{ij} \in K$. The Bézout matrix is defined as

$$B_n(f, g) = (c_{ij})_{1 \leq i, j \leq n}.$$

It is classical knowledge that $H_n(f, g)$ is equivalent as a symmetric bilinear form to the Bézout matrix $B_n(f, g)$ and these two matrices have the same determinant; see for example [Knebusch and Scheiderer 2022, Proposition 1.9.7] for a proof that, albeit formulated over real closed fields only, works over every field. In particular, the determinant of $H_n(f, g)$ is nonzero if and only if f and g are coprime.

Theorem 8.4 *The morphism $h : \mathcal{T}_3 \rightarrow \mathcal{W}_{\mathcal{O}_{\mathbb{P}^1}(4), 3}$ that sends a pair (f, g) to*

$$\varphi(f, g) = h_4(f, g) \cdot (t \wedge t^2 \wedge t^3 \wedge t^4) + \dots + h_0(f, g) \cdot (1 \wedge t \wedge t^2 \wedge t^3)$$

is well-defined and surjective on K -points. It satisfies $H_3 = \Lambda \circ h$, where Λ is the matrix from Theorem 6.5.

Proof The property $H_3 = \Lambda \circ h$ follows from the explicit expression of Λ in equation (6-1). This also implies that φ is well-defined because the determinant of H_3 does not vanish on $\mathcal{T}_3$. For the surjectivity statement it suffices to show that every $h_0, h_1, h_2, h_3 \in K$ whose Hankel matrix has rank three can be completed to the sequence of coefficients of the power series expansion of a suitable rational function. This is an easy linear algebra exercise. $\square$

Corollary 8.5 *The arithmetic writhe*

$$w : [\mathbb{P}^1, \mathbb{P}^3]_4^1 \rightarrow \operatorname{GW}(K) \times_{K^\times/K^{\times 2}} (K^\times/K^{\times 12})$$

is injective.

Proof We consider the composition of maps

$$\pi_0^N \mathcal{T}_3(K) \rightarrow \pi_0^N \mathcal{W}_{\mathcal{O}_{\mathbb{P}^1}(4),3}(K) \rightarrow \mathrm{GW}(K) \times_{K^\times/K^{\times 2}} K^\times.$$

This is the map which was proven in [Cazanave 2012, §3.3] to be injective. Indeed, this follows from $H_3 = \Lambda \circ h$ and from the fact that $H_3(f, g)$ and $B_3(f, g)$ have the same determinant and the same class in $\mathrm{GW}(K)$. By Theorem 8.4 the first map is surjective, which implies that the second map is injective. Now the claim follows from Corollary 7.16. $\square$

Remark 8.6 By the same argument as in Proposition 6.6 we can deduce that the image of the map of Corollary 8.5 consists of all pairs whose first component can be represented as $\sum_{i=1}^3 \langle a_i \rangle$ for $a_i \in K^\times$.

Remark 8.7 In the case $K = \mathbb{R}$, Remark 8.3 shows that the arithmetic writhe of the embedding coincides with the arithmetic writhe of the embedded curve. Thus, we can see it as a map

$$w : [\mathbb{P}^1, \mathbb{P}^3]_n^I \rightarrow \mathrm{GW}(\mathbb{R}).$$

As the rank of $w([\psi])$ is determined by n , we do not lose information when postcomposing with the signature. We can thus view w as a map

$$w : [\mathbb{P}^1, \mathbb{P}^3]_n^I \rightarrow \mathbb{Z},$$

which is the encomplexed writhe introduced by Viro [2001], who showed that w is an invariant under the possibly coarser equivalence relation of *rigid isotopy*. In other words, the map w factors through the set $[\mathbb{P}^1, \mathbb{P}^3]_n^{\mathrm{rig}}$ of rigid isotopy classes:

$$[\mathbb{P}^1, \mathbb{P}^3]_n^I \rightarrow [\mathbb{P}^1, \mathbb{P}^3]_n^{\mathrm{rig}} \rightarrow \mathbb{Z}.$$

It was shown in [Björklund 2011, Remark 3.21] that the map $[\mathbb{P}^1, \mathbb{P}^3]_6^{\mathrm{rig}} \rightarrow \mathbb{Z}$ is not injective. This implies that $w : [\mathbb{P}^1, \mathbb{P}^3]_6^I \rightarrow \mathrm{GW}(\mathbb{R})$ is also not injective. On the other hand, by [Björklund 2011, Theorem 1.2] the map $[\mathbb{P}^1, \mathbb{P}^3]_5^{\mathrm{rig}} \rightarrow \mathbb{Z}$ is injective. We do not know whether this generalizes to our setup, i.e., whether $w : [\mathbb{P}^1, \mathbb{P}^3]_5^I \rightarrow \mathrm{GW}(\mathbb{R})$ is injective or maybe even $w : [\mathbb{P}^1, \mathbb{P}^3]_5^I \rightarrow \mathrm{GW}(K) \times_{K^\times/K^{\times 2}} (K^\times/K^{\times 4N})$ is injective for every K .

Example 8.8 The Macaulay2 code “invariants.m2” (available as an ancillary file to the arXiv submission with identifier 2307.07543) computes two Ulrich sheaves on the secant variety $\Sigma \subset \mathbb{P}^6$ of a rational normal curve of degree 6: the first one is the Ulrich sheaf $\mathcal{F}$ of rank one coming from the theta characteristic $\alpha = \mathcal{O}_{\mathbb{P}^1}(-1)$ as in Theorem 5.16, and the second is the sheaf $\mathcal{F}_6$ of rank five coming from Theorem 5.16. Then it computes for each of those a symmetric matrix of linear forms as in Theorem 4.17. Note that, although we have not shown that the sheaf $\mathcal{F}_6$ is a *symmetric* Ulrich sheaf, it turns out that the resulting matrix is symmetric. We denote these matrices by Λ and Λ' , respectively. The signature of Λ evaluated at a linear subspace E of dimension two in $\mathbb{P}^6$ that is disjoint from the secant variety is the writhe of the projected curve in $\mathbb{P}^3$ by Theorem 6.5. The signature of Λ' at such a space is a new invariant of the curve,

which also does not change along algebraic or rigid isotopies. Our code computes these two invariants for three different linear spaces E . Here we record the results:

- (1) The linear space E is the row-span of the matrix

$$\begin{pmatrix} 2 & 9 & 3 & 1 & 8 & 7 & 3 \\ 1 & 6 & 5 & 1 & 9 & 6 & 8 \\ 7 & 8 & 1 & 1 & 2 & 6 & 9 \end{pmatrix}.$$

Both invariants of the projected curve in $\mathbb{P}^3$ are equal to 0.

- (2) The linear space E is the row-span of the matrix

$$\begin{pmatrix} 2 & 1 & 7 & 2 & 4 & 2 & 0 \\ 7 & 6 & 9 & 0 & 3 & 6 & 8 \\ 9 & 6 & 8 & 3 & 5 & 1 & 7 \end{pmatrix}.$$

The writhe of the projected curve in $\mathbb{P}^3$ is 0 while the new invariant is 4.

- (3) The linear space E is the row-span of the matrix

$$\begin{pmatrix} 0! & 1! & 2! & 3! & 4! & 5! & 6! \\ 1! & 2! & 3! & 4! & 5! & 6! & 7! \\ 2! & 3! & 4! & 5! & 6! & 7! & 8! \end{pmatrix}.$$

The writhe of the projected curve in $\mathbb{P}^3$ is 10. Note that this is predicted by [Kummer and Sinn 2022, Corollary 4.15]. The new invariant is -2 .

In particular, these computations show that the new invariant can distinguish nonisotopic curves which cannot be distinguished by the writhe. It would be interesting to have a local description for this new invariant similar to the one of the writhe.

Acknowledgements

We thank the organizers of the Thematic Einstein Semester on “*Algebraic geometry: varieties, polyhedra, computation*” in Berlin during which this project was initiated. We further thank Claudiu Raicu, Stephen McKean, Jinhyung Park and Kirsten Wickelgren for their helpful comments.

Daniele Agostini was supported by the Deutsche Forschungsgemeinschaft under grant 530132094, and by the SFB-TRR 195. Mario Kummer was supported by the Deutsche Forschungsgemeinschaft under grant 502861109.

References

- [Agostini 2024] **D Agostini**, *The Martens–Mumford theorem and the Green–Lazarsfeld secant conjecture*, J. Algebraic Geom. 33:4 (2024) 629–654 MR
- [Agostini et al. $\geq$ 2026] **D Agostini, M Kummer, J Park**, *Ulrich sheaves on secant varieties of curves*, in preparation
- [Asok and Morel 2011] **A Asok, F Morel**, *Smooth varieties up to $\mathbb{A}^1$ -homotopy and algebraic h -cobordisms*, Adv. Math. 227:5 (2011) 1990–2058 MR

- [Beauville 2018] **A Beauville**, *An introduction to Ulrich bundles*, Eur. J. Math. 4:1 (2018) 26–36 MR
- [Björklund 2011] **J Björklund**, *Real algebraic knots of low degree*, J. Knot Theory Ramifications 20:9 (2011) 1285–1309 MR
- [Cazanave 2012] **C Cazanave**, *Algebraic homotopy classes of rational functions*, Ann. Sci. Éc. Norm. Supér. (4) 45:4 (2012) 511–534 MR
- [EGA IV₂ 1965] **A Grothendieck**, *Éléments de géométrie algébrique, IV: Étude locale des schémas et des morphismes de schémas, II*, Inst. Hautes Études Sci. Publ. Math. 24 (1965) 5–231 MR
- [Ein et al. 2020] **L Ein, W Niu, J Park**, *Singularities and syzygies of secant varieties of nonsingular projective curves*, Invent. Math. 222:2 (2020) 615–665 MR
- [Ein et al. 2021] **L Ein, W Niu, J Park**, *On blowup of secant varieties of curves*, Electron. Res. Arch. 29:6 (2021) 3649–3654 MR
- [Eisenbud 2005] **D Eisenbud**, *The geometry of syzygies: a second course in commutative algebra and algebraic geometry*, Graduate Texts in Mathematics 229, Springer (2005) MR
- [Eisenbud and Schreyer 2003] **D Eisenbud, F-O Schreyer**, *Resultants and Chow forms via exterior syzygies*, J. Amer. Math. Soc. 16:3 (2003) 537–579 MR
- [Geyer 1977] **W-D Geyer**, *Reelle algebraische Funktionen mit vorgegebenen Null- und Polstellen*, Manuscripta Math. 22:1 (1977) 87–103 MR
- [Hanselka and Kummer 2024] **C Hanselka, M Kummer**, *Positive Ulrich sheaves*, Canad. J. Math. 76:3 (2024) 881–914 MR
- [Hartshorne 1977] **R Hartshorne**, *Algebraic geometry*, Graduate Texts in Mathematics 52, Springer (1977) MR
- [Huybrechts 2006] **D Huybrechts**, *Fourier–Mukai transforms in algebraic geometry*, Oxford Univ. Press (2006) MR
- [Kass and Wickelgren 2019] **J L Kass, K Wickelgren**, *The class of Eisenbud–Khimshiashvili–Levine is the local $\mathbb{A}^1$ -Brouwer degree*, Duke Math. J. 168:3 (2019) 429–469 MR
- [Kass and Wickelgren 2021] **J L Kass, K Wickelgren**, *An arithmetic count of the lines on a smooth cubic surface*, Compos. Math. 157:4 (2021) 677–709 MR
- [Kass et al. 2023] **J L Kass, M Levine, J P Solomon, K Wickelgren**, *A quadratically enriched count of rational curves*, preprint (2023) arXiv 2307.01936
- [Knebusch and Scheiderer 2022] **M Knebusch, C Scheiderer**, *Real algebra: a first course*, Springer (2022) MR
- [Krug 2018] **A Krug**, *Remarks on the derived McKay correspondence for Hilbert schemes of points and tautological bundles*, Math. Ann. 371:1-2 (2018) 461–486 MR
- [Kummer 2016] **M Kummer**, *From hyperbolic polynomials to real fibered morphisms*, PhD thesis, Universität Konstanz (2016) Available at <https://d-nb.info/1114894567/34>
- [Kummer 2019] **M Kummer**, *Totally real theta characteristics*, Ann. Mat. Pura Appl. (4) 198:6 (2019) 2141–2150 MR
- [Kummer and Shamovich 2020] **M Kummer, E Shamovich**, *Real fibered morphisms and Ulrich sheaves*, J. Algebraic Geom. 29:1 (2020) 167–198 MR
- [Kummer and Sinn 2022] **M Kummer, R Sinn**, *Hyperbolic secant varieties of M -curves*, J. Reine Angew. Math. 787 (2022) 125–162 MR
- [Lange 1984] **H Lange**, *Higher secant varieties of curves and the theorem of Nagata on ruled surfaces*, Manuscripta Math. 47:1-3 (1984) 263–269 MR
- [Lemarié-Rieusset 2026] **C Lemarié-Rieusset**, *The quadratic linking degree*, Ann. Inst. Fourier (Grenoble) (online publication March 2026)
- [Levine 2019] **M Levine**, *Motivic Euler characteristics and Witt-valued characteristic classes*, Nagoya Math. J. 236 (2019) 251–310 MR
- [Macaulay2 1992] **D R Grayson, M E Stillman**, *Macaulay2, a software system for research in algebraic geometry* (1992) Available at <https://macaulay2.com/>
- [McKean 2021] **S McKean**, *An arithmetic enrichment of Bézout’s theorem*, Math. Ann. 379:1-2 (2021) 633–660 MR
- [Mikhalkin and Orevkov 2019] **G Mikhalkin, S Orevkov**, *Maximally writhed real algebraic links*, Invent. Math. 216:1 (2019) 125–152 MR

- [Morel 2006] **F Morel**, $\mathbb{A}^1$ -algebraic topology, from “International Congress of Mathematicians” (Madrid, 2006) (M Sanz-Solé, J Soria, J L Varona, J Verdera, editors), volume 2, Eur. Math. Soc., Zürich (2006) 1035–1059 MR
- [Morel 2012] **F Morel**, $\mathbb{A}^1$ -algebraic topology over a field, Lecture Notes in Mathematics 2052, Springer (2012) MR
- [Morel and Voevodsky 1999] **F Morel, V Voevodsky**, $\mathbb{A}^1$ -homotopy theory of schemes, Inst. Hautes Études Sci. Publ. Math. :90 (1999) 45–143 MR
- [OSCAR 2024] *Oscar: open source computer algebra research system, v1.0.0* (2024) Available at <https://www.oscar-system.org/>
- [Pauli and Wickelgren 2021] **S Pauli, K Wickelgren**, Applications to $\mathbb{A}^1$ -enumerative geometry of the $\mathbb{A}^1$ -degree, Res. Math. Sci. 8:2 (2021) art. id. 24 MR
- [Ullery 2016] **B Ullery**, On the normality of secant varieties, Adv. Math. 288 (2016) 631–647 MR
- [Viro 2001] **O Viro**, Encomplexing the writhe, from “Topology, ergodic theory, real algebraic geometry” (V Turaev, A Vershik, editors), Amer. Math. Soc. Transl. Ser. 2 202, Amer. Math. Soc., Providence, RI (2001) 241–256 MR

DANIELE AGOSTINI daniele.agostini@uni-tuebingen.de
 Eberhard Karls Universität Tübingen, Tübingen, Germany

MARIO KUMMER mario.kummer@tu-dresden.de
 Technische Universität, Dresden, Germany

Proposed: Marc Levine
 Seconded: Kirsten Wickelgren, Arend Bayer

Received: January 28, 2025
 Revised: September 11, 2025

Guidelines for Authors

Submitting a paper to Geometry & Topology

Papers must be submitted using the upload page at the GT website. You will need to choose a suitable editor from the list of editors' interests and to supply MSC codes.

The normal language used by the journal is English. Articles written in other languages are acceptable, provided your chosen editor is comfortable with the language and you supply an additional English version of the abstract.

Preparing your article for Geometry & Topology

At the time of submission you need only supply a PDF file. Once accepted for publication, the paper must be supplied in $\text{\LaTeX}$. More information on preparing articles in $\text{\LaTeX}$ for publication in GT is available on the GT website.

arXiv papers

If your paper has previously been deposited on arXiv, we will need its arXiv number at acceptance time. This allows us to deposit the DOI of the published version on the paper's arXiv page.

References

Bibliographical references should be listed alphabetically at the end of the paper. All references in the bibliography should be cited at least once in the text. Use of Bib $\text{\TeX}$ is preferred but not required. Any bibliographical citation style may be used, but will be converted to the house style (see a current issue for examples).

Figures

Figures, whether prepared electronically or hand-drawn, must be of publication quality. Fuzzy or sloppily drawn figures will not be accepted. For labeling figure elements consider the pinlabel $\text{\LaTeX}$ package, but other methods are fine if the result is editable. If you're not sure whether your figures are acceptable, check with production by sending an email to graphics@msp.org.

Proofs

Page proofs will be made available to authors (or to the designated corresponding author) in PDF format. Failure to acknowledge the receipt of proofs or to return corrections within the requested deadline may cause publication to be postponed.

GEOMETRY & TOPOLOGY

Volume 30 Issue 3 (pages 835–1201) 2026

The level structure in quantum K -theory and mock theta functions	835
YONGBIN RUAN, MING ZHANG and YAOXIONG WEN	
Springer theory for symplectic Galois groups	883
KEVIN MCGERTY and THOMAS NEVINS	
Unipotent morphisms	929
DANIEL BRAGG, JACK HALL and SIDDHARTH MATHUR	
On the Donaldson–Scaduto conjecture	959
SAMAN HABIBI ESFAHANI and YANG LI	
Milnor’s invariants for knots and links in closed orientable 3-manifolds	983
RYAN STEES	
Chromatic defect, Wood’s theorem and higher real K -theories	1051
CHRISTIAN CARRICK	
Compact holonomy G_2 manifolds need not be formal	1109
LUCÍA MARTÍN-MERCHÁN	
$K3$ surfaces associated with varieties of generalized Kummer type	1129
SALVATORE FLOCCARI	
Ulrich sheaves, the arithmetic writhe and algebraic isotopies of space curves	1155
DANIELE AGOSTINI and MARIO KUMMER	