

AG
T

*Algebraic & Geometric
Topology*

Volume 24 (2024)

Comparison of period coordinates and Teichmüller distances

IAN FRANKEL



Comparison of period coordinates and Teichmüller distances

IAN FRANKEL

We show that when two unit-area quadratic differentials are ϵ -close with respect to good systems of period coordinates and lie over a compact subset K of the moduli space of Riemann surfaces $\mathcal{M}_{g,n}$, then their underlying Riemann surfaces are $C\epsilon^\alpha$ -close in the Teichmüller metric. Here, α depends only on the genus g and the number of marked points, while C depends on K .

30F60

1. Introduction and statement of main result	2451
2. Sketch of proof	2455
3. Teichmüller spaces and quadratic differentials	2456
4. Quadratic differentials on the sphere	2464
5. Perturbing quadratic differentials on the sphere	2481
6. Building the quasiconformal map	2492
Appendix A. Local finiteness of period coordinate systems	2498
Appendix B. δ -clusters and the Euclidean metric	2501
References	2508

1 Introduction and statement of main result

1.1 Preliminaries

The moduli space of compact Riemann surfaces of genus g with n unlabeled marked points (or n deleted points), $\mathcal{M}_{g,n}$, is a complex variety and orbifold of dimension $3g - 3 + n$, with each point of $\mathcal{M}_{g,n}$ representing a biholomorphism class a of compact genus Riemann surface with n points deleted (or with n marked points). The topology of $\mathcal{M}_{g,n}$ is induced by the Kobayashi metric. (The Kobayashi pseudometric on a complex analytic X space is the largest pseudometric for which all holomorphic maps of the hyperbolic disk into X are nonexpanding, and it is a nondegenerate metric in all cases we will consider.) Teichmüller constructed a metric whereby a K -quasiconformal homeomorphism (see Section 3) between X and Y exists if and only if the distance between $(X; p_1, \dots, p_n)$ and $(Y; q_1, \dots, q_n)$ is at most K . Royden [1971] showed that this is the same as the Kobayashi metric. The orbifold universal cover

of $\mathcal{M}_{g,n}$, the Teichmüller space $\mathcal{T}_{g,n}$, is a contractible complex manifold of dimension $3g - 3 + n$ when $3g - 3 + n > 0$. In the same paper, Royden showed that the biholomorphism group of $\mathcal{T}_{g,n}$ is naturally isomorphic to the Teichmüller modular group, or mapping class group $\text{Mod}(S_{g,n})$, of the surface $S_{g,n}$, and the quotient by this action is $\mathcal{M}_{g,n}$.

Even though the Teichmüller–Kobayashi metric on $\mathcal{T}_{g,n}$ is not Riemannian, it has a geodesic flow, which is usually described as a dynamical system on the unit *cotangent bundle*. For $X \in \mathcal{T}_{g,n}$, the cotangent space to $\mathcal{T}_{g,n}$ at X consists of those meromorphic sections of the tensor square of the cotangent bundle to the complete Riemann surface corresponding to X for which the only poles are simple and occur at the marked points of X . We refer to such sections as *quadratic differentials*. We will write $\text{QD}(\mathcal{T}_{g,n})$ to denote the space of nonzero quadratic differentials on surfaces of genus g with n marked points.

Many analogies have been made between the geodesic flow for this metric and the geodesic flow on a closed negatively curved Riemannian manifold. Euclidean geometry is the main tool in the study of the Teichmüller geodesic flow; we describe briefly the connection here. Given a nonzero quadratic differential q on X , there is an associated flat (Gaussian curvature = 0) metric on X with a finite number of cone-type singularities, which we will call the q -metric. The metric can be defined locally by taking the integral of the holomorphic 1-form \sqrt{q} to give an isometric chart to \mathbb{C} , ie

$$z \mapsto \int_{z_0}^z \sqrt{q}$$

gives an isometry between a neighborhood of $z_0 \in X$ with the q -metric and an open set in $\mathbb{C} = \mathbb{R}^2$ with the standard Euclidean metric, for each point z_0 where q is holomorphic and nonvanishing.

Singularities of the q -metric occur at points where such charts cannot be defined — at zeros and poles of q . At these points, we have cone-type singularities with cone angle $(n + 2)\pi$ at each zero of order $n \geq -1$. (If we allowed q to have poles of higher order, they would be an infinite distance away, so the metric cannot extend.) It should be noted that, since \sqrt{q} is only defined up to sign, these charts to \mathbb{C} are unique up to the group $\{z \mapsto C \pm z : C \in \mathbb{C}\}$. We will refer to a surface equipped with such a metric as a *half-translation surface*. We note that it is also possible to take a surface with charts and transition maps lying in $\{z \mapsto C \pm z : C \in \mathbb{C}\}$, and give it a metric for which those charts are isometric and recover a complex structure and a quadratic differential which is $dz \otimes dz$ in local coordinates for the system of charts. If, in addition, the resulting metric space can be completed by adding only finitely many points with cone-type singularities, the complex structure extends uniquely to the completion and the quadratic differential extends meromorphically. We also establish the following convention:

Convention 1.1 Given $(X; p_1, \dots, p_n; q)$ a Riemann surface with n marked points p_1, \dots, p_n and a nonzero meromorphic quadratic differential q on X whose only poles are simple and occur at a subset of $\{p_1, \dots, p_n\}$, the set of *singularities* of $(X; p_1, \dots, p_n; q)$ are the zeros of q on X together with the collection of all marked points, regardless of whether or not q has poles at those points.

1.2 Period coordinates and the main theorem

We can integrate \sqrt{q} along q -geodesic arcs $\gamma_i : [0, 1] \rightarrow X$ chosen so that $\gamma_i((0, 1))$ contains no singularities and γ_i is injective on $(0, 1)$, but $\gamma(0)$ and $\gamma(1)$ are (not necessarily distinct) singularities. Such arcs are called *saddle connections*. A collection of such integrals forms a local holomorphic coordinate chart for $\text{QD}(\mathcal{T}_{g,n})$ in a neighborhood of (X, q) , provided that q has $4g - 4 + n$ simple zeros and n simple poles. We call such charts *period coordinate charts*. A period coordinate chart can be extended to the boundary of any open precompact subset of $\text{QD}(\mathcal{T}_{g,n})$ on which it is well defined, and embeds into $\mathbb{C}^{6g-6+2n}$ as a convex set. $\text{QD}(\mathcal{T}_{g,n})$ is locally finitely covered by the closures of such sets by Corollary A.7.

If a set $S \subset \text{QD}(\mathcal{T}_{g,n})$ is homeomorphic to a compact convex set with nonempty interior $K \subset \mathbb{C}^{6g-6+2n}$ via a homeomorphism $f : S \rightarrow K$ whose restriction to the interior of S is a period coordinate embedding, we say S is a *compact convex period coordinate patch*. In Appendix A we show that $\text{QD}(\mathcal{T}_{g,n})$ is locally finitely covered by compact convex period coordinate patches.

We will define a $\text{Mod}(S_{g,n})$ -invariant path metric $d_{\text{Euclidean}}$ (see Definition 3.2) that has the property that every compact convex period coordinate patch is locally bi-Lipschitz to the corresponding subset of $\mathbb{C}^{6g-6+2n}$.

We now state our main theorem:

Theorem 1.2 *Let g and n be nonnegative integers such that $3g - 3 + n > 0$. Let $a_n = 1$ if $n = 0$ and 2 if $n > 0$. Let K be a compact subset of $\mathcal{T}_{g,n}$. Then for any unit-area quadratic differentials (X_1, q_1) and (X_2, q_2) with $X_1, X_2 \in K$, there is a constant C_K such that*

$$d_{\text{Teich}}(X_1, X_2) < C_K d_{\text{Euclidean}}(q_1, q_2)^{2/[2+a_n(4g-4+n)]}.$$

Remark Theorem 1.2 is true whether we restrict $d_{\text{Euclidean}}$ to unit-area quadratic differentials intrinsically or extrinsically (whether or not we consider paths that leave the space of unit-area differentials). The reason for this is that $\|q\|$, the area of the q -metric, is locally Lipschitz as a function of (X, q) with respect to the metric $d_{\text{Euclidean}}$. Therefore if (X_1, q_1) and (X_2, q_2) are unit area and sufficiently close, the shortest path between them stays near the set of unit-area quadratic differentials. A path in $\text{QD}(\mathcal{T}_{g,n})$ that starts and ends in K can be projected onto the unit-area subspace, and if the path stays sufficiently close to K this projection will only increase its length by a bounded factor.

The real content of the theorem is that it remains valid even when zeros of the quadratic differential are allowed to coincide with each other, or with marked points. If we required K to be a compact subset of the *principal stratum* of quadratic differentials, ie the space of quadratic differentials with n simple poles and $4g - 4 + n$ simple zeros, then we could replace the Hölder exponent $2/[2 + a_n(4g - 4 + n)]$ in Theorem 1.2 with 1, by finding a common triangulation of (X_1, q_1) and (X_2, q_2) by saddle connections and mapping triangles to triangles.

We further remark that no inequality in the other direction is possible, since two quadratic differentials $q_1 \neq q_2$ may have the same underlying Riemann surface $X = X_1 = X_2$.

The *strongly stable leaves* for the Teichmüller geodesic flow are the collections of unit-area quadratic differentials that, on each chart given by period coordinates, can be obtained from each other by changing only the imaginary parts of period coordinates and changing systems of period coordinates. One might hope for a converse inequality if we make the additional assumption that (X_1, q_1) and (X_2, q_2) are on the same strongly stable leaf or strongly unstable leaf of the Teichmüller geodesic flow, for in such cases X_1 and X_2 are known to be different points in $T_{g,n}$, by the main result of [Hubbard and Masur 1979].

Before we proceed toward the proof, we would like to make a remark about the optimal Hölder exponent. We do not make any claim that the exponent we produce is sharp, but we do claim that, asymptotically, it can only be improved by a factor of 2, at least when $n \geq 1$.

The reason for this is the following example:

Let (X_t, q_t) be a family of quadratic differentials that are all isometric to some (X, q) , but on which one marked point moves according to a holomorphic parameter t . Moreover, assume that the quadratic differential q has the form $z^k dz^2$ in some local coordinate z on the underlying Riemann surface. The q -metric is then given by $ds^2 = |z|^k |dz|^2$, and we take the parameter t to be the location of the marked point in the coordinate z , so t varies over a neighborhood of zero.

For t near 0, we have that $C^{-1}|t| \leq d_{\text{Teich}}(X_t, X_0) < C|t|$ since t is a coordinate in a holomorphic coordinate system for Teichmüller space, and the Teichmüller metric is Finsler with respect to this complex structure.

On the other hand, the Euclidean distance between q_0 and q_t is comparable to the q -distance between the points $z = 0$ and $z = t$, which equals $\int_0^{|t|} x^{k/2} dx$.

Thus $C^{-1}|t|^{(2+k)/2} < d_{\text{Euclidean}}(q_0, q_t) < C|t|^{(2+k)/2}$.

(The actual behavior of the Euclidean metric near q_0 is treated more carefully in the appendices, but the intuition is that one of our period coordinates comes from a saddle connection joining 0 to t , and this coordinate accounts for the Euclidean distance from q_0 to q_t .)

The largest possible k is $4g - 4 + n$, so the best possible Hölder exponent is then $2/[4g - 2 + n]$.

Acknowledgments

This research was conducted for the author's PhD thesis at the University of Chicago. The author would like to thank Howard Masur for suggesting this problem. The author would also like to thank Howard Masur, Kasra Rafi, and especially Alex Eskin for helpful conversations, and the referee(s) whose input helped to clarify the exposition.

2 Sketch of proof

The proof is by construction of a quasiconformal map between nearby surfaces. The easiest way to build a quasiconformal map between (X_1, q_1) and (X_2, q_2) is by finding a common triangulation by saddle connections and applying a piecewise linear (PL) map. Then one can simply estimate the dilatation (the best possible constant K such that the map is K -quasiconformal) on each triangle, and in so doing obtain a bound for the dilatation of the entire map.

Unfortunately, this strategy yields poor estimates unless all of the edge lengths and angles of the triangles can be bounded away from zero, and usually can't be used to build maps between surfaces in different strata. Our solution to this problem is to find a system of disks, which we call *nearly regular right polygons* or NRRPs, that isolate clusters of singularities. We can triangulate the remainder of each of the surfaces, and build a PL map in the complement of these disks. As for the disks themselves, we need to find a quasiconformal extension of the boundary map to the NRRP, and estimate its dilatation. For this we use a Beurling–Ahlfors extension, which requires as input an estimate of the quasisymmetry of the boundary map. This is carried out Section 6.

The boundary maps between disks are fixed since the map is PL outside of the disks. However, we need to understand the uniformization of these disks to the upper half-plane in order to estimate its dilatation. Our disks will be chosen so that, when doubled along the boundary, they have the isometry type of half-translation surfaces coming from conjugation-invariant quadratic differentials on \hat{C} . Almost all of the work consists of bounding the changes in the locations of the singularities as a function of the changes in period coordinates. Since the quadratic differentials are actually of the form $f(z) dz^2$ for some rational function $f \in \mathbb{R}(z)$, it is actually easier to give a lower bound for the changes in period coordinates in terms of the changes of locations of the zeros and poles, for a sequence of perturbations tending to 0. This is most of the work, and it is carried out in Section 5.

The key to finding a lower bound is finding a scale on which perturbations do not have canceling effects on period coordinates. For this we use a partial compactification of strata similar to the compactification used in [Bainbridge et al. 2019] to handle collisions of singularities. For us, the most important organizing information is which collections of singularities are colliding faster than others. The limit of a differential in this compactification is a stable curve with one component for each such cluster, and a meromorphic quadratic differential on each component. A colliding cluster of singularities will correspond to a component with a quadratic differential that has a higher-order pole on ∞ . The partial compactification is described in Definition 4.35, and a convergent sequence is defined in Notation 5.3.

When a cluster of colliding singularities moves mostly in the same direction, we show that it essentially behaves as a point, and shadows a perturbation in a lower-dimensional stratum. This is the content of Proposition 5.4. When a cluster of singularities move in what would appear to be canceling directions according to Proposition 5.4, we find a leading-order approximation for how the perturbation affects saddle connections in this cluster. If there are subclusters the leading term may again be zero, but we

can break the perturbation into components, at least one of which must be detectable on some cluster. Appendix B is dedicated to showing that these kinds of perturbations do not cancel out. The end result of this argument is the key estimate, Lemma 5.11, which bounds the change in the location of the singularities as a function of the change in the periods. (More precisely, we prove the contrapositive.) We finish Section 5 by converting Lemma 5.11 to the form we need to apply to the specific disks, Lemma 5.12.

3 Teichmüller spaces and quadratic differentials

In this section, we assemble basic facts about Teichmüller spaces and coordinate systems on strata.

3.1 Flat metrics and period coordinates

On a half-translation surface, the slope of a tangent vector at any nonsingular point is well defined; in particular, vertical and horizontal are well-defined notions. Length and area are also well defined in the q -metric, since the change of charts preserves them. It is also evident that if X has the q -metric, then geodesics are polygonal arcs and only change direction at cone points, and when they do they turn by an angle of at least π measured either way around the cone point. If q is given in local coordinates by $(dx + i dy)^{\otimes 2}$, then the area of X with the q metric, given by $|q| = \int_X dx \wedge dy$, is the norm of q in the Teichmüller cometric. (Note that this area form does not depend on which square root of q is picked.)

The collection of all quadratic differentials whose associated metrics have finite area is a vector bundle over $\mathcal{T}_{g,n}$ (if we include the zero section), and the set of nonzero quadratic differentials admits a holomorphic stratification whereby each stratum consists of differentials whose metrics have the same number of singularities of each type; two singularities have the same type if they are both marked points and have the same cone angle, or they are both unmarked points and have the same cone angle. Two different strata can have the same underlying Riemann surfaces and quadratic differentials, but differ in which points are marked. The *principal stratum* consists of those quadratic differentials with a simple pole at each of the n marked points and $4g - 4 + n$ simple zeros. The complement of the principal stratum is an analytic subvariety, so the principal stratum sits inside $\text{QD}(\mathcal{T}_{g,n})$ as a dense set of full measure (with respect to Lebesgue measure class).

Definition 3.1 The *universal half-translation surface* is a surface bundle over $T^*\mathcal{T}_{g,n}$ with fiber S_g , and where the fiber over (X, q) is a copy of \bar{X} equipped with the q -metric.

We will give a more explicit description of the topology of this bundle later in this section.

With our convention, the set of all singularities of q -metrics is an *incidence subvariety* of the universal half-translation surface. Locally they define a collection of sections of the universal half-translation surface, at least in a neighborhood of each point in the principal stratum. If p_1, \dots, p_r are the singularities at a point (X, q) in the principal stratum, any nearby (X', q') will have singularities p'_1, \dots, p'_r , and each p'_i varies holomorphically with (X', q') in a bundle over $T^*\mathcal{M}_{g,n}$ whose fibers are the Riemann surfaces

represented. It is thus meaningful to talk about points being the same singularity on different surfaces. We may pick homotopy classes of arcs whose endpoints are the singularities, and the integrals $\int_{p'_i}^{p'_j} \sqrt{q'}$ over such arcs will vary holomorphically. When the arcs are represented by saddle connections, we call the integrals $\int_{p'_i}^{p'_j} \sqrt{q'}$ *period coordinates*. In a neighborhood of any point in the principal stratum, some such collection of period coordinates associated to saddle connections will form a holomorphic coordinate chart to $\mathbb{C}^{6g-6+2n}$. On a dense subset of such coordinate charts, we can paste together Euclidean metrics to form a path metric on the moduli space of quadratic differentials, described in the next section.

3.2 The Euclidean metric on $\text{QD}(\mathcal{T}_{g,n})$

A choice of norm on a vector space is equivalent to a choice of the closed unit ball, ie a compact convex set with nonempty interior which is symmetric about the origin. Given a finite collection of norms we can simply take the convex hull of the union of their unit balls to be the unit ball of the pasted metric. The resulting norm is of course equivalent (bi-Lipschitz) to any of the original Euclidean norms at each point. Call this the *union convex hull operation*. It produces the largest norm that is less than or equal to a given set of norms. If we have an arc in a stratum, at each point of which some system of period coordinates has locally constant derivative, we describe a norm which is the speed of the arc at almost every point. For each real $L > 0$, there are locally only finitely many coordinate systems represented by saddle connections of length $\leq L$ (see Appendix A) in the following sense: for any point $(X, q) \in \text{QD}(\mathcal{T}_{g,n})$ with $q \neq 0$, there is a neighborhood of (X, q) in which only finitely many such systems exist, *even if* (X, q) is not in the principal stratum.

Recall the definition of compact convex period coordinate patch in the discussion immediately before Theorem 1.2. At each point (X, q) , we consider all good embeddings of sets containing (X, q) in which (X, q) maps to $\{z : |z| \leq \Theta(X, q)\}^n$ for some continuous proper function Θ which is large enough to ensure that the set of such good embeddings is not empty for any (X, q) . From Appendix A it is clear such a function exists. For example, 4 times the q -diameter of (X, q) is such a function.

Definition 3.2 Fix a continuous function $\Theta : \text{QD}(\mathcal{M}_{g,n}) \rightarrow (0, \infty)$ as above. We define the *Euclidean metric* $d_{\text{Euclidean}}$ on the space of half-translation surfaces to be the path metric obtained by applying the union convex hull operation to the norms coming from compact convex period coordinate patches containing (X, q) , whose defining saddle connections have length less than $\Theta((X, q))$.

It should be clear from the definition that $d_{\text{Euclidean}}$ is the largest path metric such that, for all (X, q) , all compact convex period coordinate patches containing (X, q) that take it into $\{v : |v|_{\infty} < \Theta(X, q)\}^{6g-6+2n}$ are noncontracting in a neighborhood of (X, q) .

Remark The local Lipschitz class of the metric $d_{\text{Euclidean}}$ depends only on the fact that it comes from a locally finite collection of systems of period coordinates. It does not depend on the specific coordinate systems chosen. Theorem 1.2 and any other statements that depend only on the local Lipschitz class remain true if we define $d_{\text{Euclidean}}$ differently, for instance by applying the union convex hull operation on

a different locally finite collection of period coordinate systems that is rich enough to cover quadratic differential space.

We note that the Euclidean metric extends to a path metric on the whole of $\text{QD}(\mathcal{T}_{g,n})$, but in the absence of the lower bound for the area, arbitrarily short paths with the Euclidean metric move arbitrarily far in Teichmüller space. For this reason, and for the reason that the Teichmüller flow is usually studied on unit-area differentials, we find it prudent to restrict to the unit-area locus.

Also, when we are referring to a particular stratum or stratum closure, the Euclidean metric is understood to be the *intrinsic* metric on the stratum or stratum closure, rather than the Euclidean metric in the ambient moduli space.

3.3 Teichmüller's metric

The Teichmüller space of Riemann surfaces of genus g with n marked points is the set of marked complex structures on a Riemann surface of genus g with n points deleted. One way to define this more precisely is, given a smooth oriented surface S diffeomorphic to a surface of genus g with n points deleted, to look at all orientation-preserving homeomorphisms of complete finite volume hyperbolic Riemann surfaces into S , subject to the following equivalence relation: for $f_1: M_1 \rightarrow S$ and $f_2: M_2 \rightarrow S$, we say (M_1, f_1) and (M_2, f_2) are equivalent if $f_1^{-1} \circ f_2$ is homotopic to a biholomorphism. The set of equivalence classes of pairs (M, f) is the *Teichmüller* space of genus g surfaces with n punctures, which we will call $\mathcal{T}_{g,n}$. When it is convenient, we will sometimes view the punctures as marked points on a compact surface instead of deleted points.

$\mathcal{T}_{g,n}$ carries several metrics; the one we will be concerned with is the *Teichmüller* metric. We will briefly summarize the geometric properties of this metric by characterizing its geodesics. Proofs of these can be found in, for instance, [Hubbard 2006] or [Farb and Margalit 2012].

- Start with any nonzero holomorphic quadratic differential α , ie a nonzero holomorphic section of the tensor square of the cotangent bundle of a Riemann surface. If z is a local coordinate on a Riemann surface, then locally $\alpha = f(z) dz^2$, with f holomorphic in z ; at each point at which f does not vanish we can find a holomorphic coordinate z for which $\alpha = dz^2$.
- In any simply connected region on such a chart, the level sets of the real and imaginary parts of z give a pair of transverse *measured foliations*, that is to say, foliations equipped with a transverse measure on the local leaf space, and the transverse measure is invariant under transition maps of a system of charts defining the foliation. (See [Fathi et al. 2012, Section 1.3] for a precise construction of the space of measured foliations.) Given any local coordinates $z = x + iy$, the integrals $\int_C |dx|$ and $\int_C |dy|$ are well defined for any smooth contour of integration C . There is also an invariant volume form, and in the case that our surface has punctures we will only consider those quadratic differentials for which the area form $dx \otimes dy$ assigns finite measure. This is equivalent to assuming the quadratic differential extends meromorphically to a compact Riemann surface with at most simple poles at each of finitely many new points.

- Now, fix a Riemann surface M_0 and a pair of transverse measured foliations associated to a holomorphic quadratic differential α with finite area. Then for each $\lambda \in \mathbb{R}$ there is a Riemann surface M_λ with a homeomorphism $f : M_0 \rightarrow M_\lambda$, smooth away from the zeros of α , with holomorphic charts such that the pushforwards of $|dx|$ and $|dy|$ are $|e^\lambda dx|$ and $|e^{-\lambda} dy|$, and $\lambda \mapsto M_\lambda$ is a unit-speed geodesic. The map f is called a *Teichmüller map*.
- All geodesics in $\mathcal{T}_{g,n}$ arise in this way, and the length of a geodesic between two points in $\mathcal{T}_{g,n}$ is equal to the distance between those points in the Teichmüller metric. Any two distinct points in $\mathcal{T}_{g,n}$ lie on a unique geodesic.

Suppose $f : U \rightarrow V$ is an almost-everywhere differentiable orientation-preserving homeomorphism of bounded domains in \mathbb{C} , with first-order distributional derivatives in $L^2(U)$. Then the partial derivatives

$$f_z := \frac{1}{2}(f_x - if_y) \quad \text{and} \quad f_{\bar{z}} := \frac{1}{2}(f_x + if_y)$$

are defined for almost all z and $|f_{\bar{z}}| < |f_z|$ for almost all $z \in U$.

Definition 3.3 If

$$K_{z_0}(f) := \left| \frac{|f_z(z_0)| + |f_{\bar{z}}(z_0)|}{|f_z(z_0)| - |f_{\bar{z}}(z_0)|} \right|$$

satisfies $K_{z_0}(f) \leq K$ for almost all z_0 , with respect to the Lebesgue measure class, then we say that f is K -quasiconformal.

Since the quantity $K_{z_0}(f)$ is invariant under holomorphic changes of coordinates on U and V , we say that a homeomorphism between Riemann surfaces is K -quasiconformal if it is K -quasiconformal with respect to a choice of holomorphic coordinate charts; by the above discussion the choice of charts does not matter. It is a theorem that 1-quasiconformal maps are actually conformal.

The Teichmüller metric also has the following characterization:

Theorem 3.4 (Teichmüller) *If (Y_1, f_1) and (Y_2, f_2) represent two points in $\mathcal{T}_{g,n}$, then their distance is at most $\frac{1}{2} \log(K)$ if and only if there is a K -quasiconformal homeomorphism $g : Y_1 \rightarrow Y_2$ such that $f_2 \circ g \circ f_1^{-1}$ is homotopic to the identity on S . Teichmüller maps are the unique dilatation-minimizing maps in their homotopy classes (that is, they are K -quasiconformal for the smallest possible K).*

With this metric, $\mathcal{T}_{g,n}$ is homeomorphic to $\mathbb{R}^{6g-6+2n}$. This manifold admits a complex structure, and $\text{QD}(\mathcal{T}_{g,n})$ is a holomorphic vector bundle over $\mathcal{T}_{g,n}$ with the zero section deleted (as we have defined it). However we will not need to use this complex structure.

In addition to quadratic differentials that we used to characterize the Teichmüller metric, we will sometimes consider quadratic differentials on \mathbb{C} of the form $f(z) dz^2$, where f is a polynomial or a rational function with at most one pole in \mathbb{C} , which is simple. They extend to $\hat{\mathbb{C}}$ with a higher-order pole at ∞ . The induced flat metrics on \mathbb{C} have infinite area, and the distance to ∞ is infinite in such metrics.

We can use Teichmüller maps to create various bundles over Teichmüller space. We will pick a basepoint $X \in \mathcal{T}_{g,n}$. If q is a quadratic differential over the Riemann surface X , we can associate to q the Teichmüller map ϕ_q that moves X along the geodesic described by q by a distance equal to the area of (X, q) . This identifies Teichmüller space with the space of quadratic differentials over X , and we can use coordinates (x, q) to refer to the point $\phi_q(x)$.

We can also trivialize the bundles of quadratic differentials and half-translation surface structures using Teichmüller maps. To do so, we need the following definition from [Fathi et al. 2012] and theorem from [Hubbard and Masur 1979]. In the definition below, our base surface $S_{g,n}$ and all Riemann surfaces are understood to have n points *deleted*.

Definition 3.5 Let q_1 and q_2 be vertical foliations of quadratic differentials on two Riemann surfaces $X_1, X_2 \in \mathcal{T}_{g,n}$, with vertical measured foliations $\mathcal{F}_v(q_1)$ and $\mathcal{F}_v(q_2)$, respectively. For a simple closed curve C on the base surface $S_{g,n}$ with the n points *deleted*, let $i(\mathcal{F}_v(q_j), C)$ denote the infimum of the transverse measure evaluated on curves homotopic to C in the q_j -metric. If $i(\mathcal{F}_v(q_1), C) = i(\mathcal{F}_v(q_2), C)$ for all C , we say that $\mathcal{F}_v(q_1)$ and $\mathcal{F}_v(q_2)$ are equivalent and we write $\mathcal{F}_v(q_1) = \mathcal{F}_v(q_2)$.

Definition 3.6 The set of all equivalence classes of measured foliations is the space \mathcal{MF} . Its quotient by the action of $\mathbb{R}^{>0}$ by scalars is *PMF*.

The values of the numbers $i(\mathcal{F}, C)$, where C ranges over all (homotopy classes of) essential simple closed curves, give a weak- $*$ topology on the set of measured foliations, which gives \mathcal{MF} the homeomorphism type of the product of $\mathbb{R}^{6g-6+2n} \setminus \{0\}$; see exposé 6 of [Fathi et al. 2012]. Moreover, if we quotient by the obvious $\mathbb{R}^{>0}$ action by scalars, the result is homeomorphic to a sphere of dimension $6g - 7 + 2n$.

Theorem 3.7 [Hubbard and Masur 1979] *The map from $\text{QD}(\mathcal{T}_{g,n})$ to $\mathcal{T}_{g,n} \times \mathcal{MF}$ defined by*

$$(X, q) \mapsto (X, \mathcal{F}_v(q))$$

is a homeomorphism.

Theorem 3.8 [Hubbard and Masur 1979] *The map from the set of unit-area quadratic differentials in $\text{QD}(\mathcal{T}_{g,n})$ to $\mathcal{T}_{g,n} \times \mathcal{PMF}$ defined by*

$$(X, q) \mapsto (X, \mathbb{R}^{>0} \mathcal{F}_v(q))$$

is a homeomorphism.

Corollary 3.9 [Hubbard and Masur 1979] *Pick $(X, q) \in \text{QD}(\mathcal{T}_{g,n})$. For each Y in $\mathcal{T}_{g,n}$, there is a unique $q' \in T_Y^* \mathcal{T}_{g,n}$ such that $\mathcal{F}_v(q) = \mathcal{F}_v(q')$. Moreover, given any basepoint $Y \in \mathcal{T}_{g,n}$, there is a trivialization of the vector bundle $T^* \text{QD}(\mathcal{T}_{g,n})$ over $\text{QD}(\mathcal{T}_{g,n})$ that sends each (X, q) to (X, q') , where q' is the unique quadratic differential on Y whose vertical foliation is equivalent to that of q .*

We note that the foliation of $\text{QD}(\mathcal{T}_{g,n})$ whose leaves are the sets of quadratic differentials with equivalent vertical foliations have the following property: if we fix a system of period coordinates, then moving along a leaf of this foliation only changes the imaginary parts of the period coordinates.

Now, we describe a trivialization of another bundle. Let q be a quadratic differential over the basepoint surface X , let q' be a nonzero quadratic differential over X , and let $x \in \bar{X}$, the completion of X . We can use coordinates (x, q, q') to refer to the point $\phi_q(x)$ on the surface $\phi_q(\bar{X})$ equipped with the metric given by the unit-area quadratic differential on $\phi_q(\bar{X})$ whose vertical foliation is equivalent to $\mathcal{F}_v(q')$. The union of all such points (x, q, q') is the *universal half-translation surface of type $S_{g,n}$* . Similarly, we can take the universal cover of this bundle, which results in replacing each fiber with its universal cover in our coordinates (Teichmüller maps lift to universal covers). The coordinates are not particularly important, but the induced manifold topologies will be used for various compactness statements — for instance, we may refer to a sequence of saddle connections on quadratic differentials (X_i, q_i) converging to a saddle connection on (X, q) . This means that the convergence is in the Hausdorff metric on the space of compact subsets of the bundle. In the case when the surfaces in question have genus 0 and are explicitly uniformized to $\hat{\mathbb{C}}$, we have an alternative but topologically equivalent trivialization of the same bundles, so the notion of Hausdorff convergence is the same.

Definition 3.10 Given a quadratic differential (X, q) , we may form a branched double cover as follows: Let (\tilde{X}, \tilde{q}) be the metric completion of a metric double cover of the set of nonsingular points of X given by $\{(X, v) : v \text{ is a vertical unit vector with respect to } q\}$. We call this the *orienting double cover* of (X, q) , and it admits a degree-2 map to X that is branched over all points where the cone angle is of the form $(2n - 1)\pi$ for $n \in \mathbb{N}$.

The orienting double cover is the square of a holomorphic one-form on \tilde{X} , and \tilde{X} is connected only if q is not the square of a holomorphic one-form. We can see that the orienting double cover comes with a flat metric and a unit vertical vector field at all nonsingular points. Moreover, the map $i : X \rightarrow X$ defined by $(x, v) \mapsto (x, -v)$ is an involution fixing the ramification points. If Σ is the set of singular points of X , then $\tilde{\Sigma}$ is the preimage of the set of singular points in the double cover. Now, on \tilde{X} we can find a holomorphic 1-form α such that $f^*q = \alpha^{\otimes 2}$, and α is unique up to sign.

3.4 Triangulations and degenerations

Every nonzero quadratic differential admits a cellular decomposition whose 1-cells are saddle connections and whose open 2-cells are isometric to the interiors of triangles in Euclidean space. We shall refer to such decompositions as *triangulations*, even though not every edge has distinct vertices.

There is one triangulation, due to [Masur and Smillie 1991], which can be constructed as follows: If we delete all the poles of a nonzero meromorphic quadratic differential (X, α) , we can take the metric universal cover of $X \setminus \{x_1, \dots, x_n\}$ where the x_i are the marked points; let Γ be the deck group of this covering map. This makes sense because on $X \setminus \{x_1, \dots, x_n\}$ the flat metric associated to α is still a path metric. Then complete the resulting space \tilde{X} . (The new points added will be points of cone angle ∞ .) Then, for each point in \tilde{X} , we consider the set of singularities and cone points of angle ∞ that are closest to X . If for some point p there are three or more such points distance R_p from p , then there is an

embedded Euclidean disk of radius $R > 0$ around p . Since the length of the shortest saddle connection is nonzero, there are only finitely many singularities on the boundary of this disk, and they have a cyclic ordering. For each such p draw the chords that are the boundary of the convex hull of these points in the disk, and label them as 1–cells. We now have a Γ –invariant cell decomposition. Now quotient by Γ . The result is a cell decomposition in which the 0–cells are singularities, the 1–cells are saddle connections, and the 2–cells are convex polygons, which can then be triangulated by diagonals.

Definition 3.11 We refer to the above construction as the *Delaunay triangulation*.

The Delaunay triangulation has a number of useful properties. It is invariant under scaling the quadratic differential by nonzero complex numbers, and for a typical surface it does not involve a choice — if it does, only finitely many choices are possible. (By a typical surface, we mean a full-measure set with respect to a Lebesgue measure, described below.) One of them is that the lengths of the saddle connections are bounded by twice the diameter of the flat metric, and another is that all of the angles of the triangles are bounded away from zero given an upper bound on the diameter of the surface and a lower bound on the length of the shortest saddle connection.

A triangulation by saddle connections, together with the collection of cone angles, imposes relations between the periods of the saddle connections. Given a triangulation, there is some number of periods that determines all of the others.

In [Masur 1982] the following complex manifold structure on strata in $\text{QD}(\mathcal{T}_{g,n})$ was described, using the homology of the orienting double cover:

Definition 3.12 Let (X, q) be a quadratic differential with double cover (\tilde{X}, \tilde{q}) , let Σ be the set of singularities of X with inverse image $\tilde{\Sigma}$, and let ι be the involution of \tilde{X} . Then

$$H_1^{\text{odd}}(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) := \{\varphi \in H_1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) : \iota_*(\varphi) = -\varphi\},$$

$$H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) := \{\alpha \in H^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) : \iota^*(\alpha) = -\alpha\}.$$

Similarly, define H_1^{even} and H_{even}^1 to be the eigenspaces of ι_* and ι^* with eigenvalue 1.

Since ι^2 , ι_*^2 and $(\iota^*)^2$ are the identity map on their domains, we have

$$H_1^{\text{odd}}(\tilde{X}, \tilde{\Sigma}, \mathbb{C}) \oplus H_1^{\text{even}}(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) = H_1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$$

and

$$H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) \oplus H_{\text{even}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) = H^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}).$$

We also have the usual duality between homology and cohomology groups since we are using field coefficients.

Fix a subset S of a stratum of $\text{QD}(\mathcal{T}_{g,n})$ sharing a triangulation T by saddle connections in fixed homotopy classes relative to the singularities; each saddle connection γ on any $(X, q) \in S$ has two lifts γ' and γ'' to \tilde{X} . Fix T as an oriented graph, and assume the edges have names. On the orienting double cover

of a (X, q) , there is a triangulation \tilde{T} that maps to T as an oriented graph. For each γ , fix a choice of γ' and γ'' . Then \tilde{T} comes equipped with this choice of names as well. In other words, if we say two surfaces have the same triangulation, then an isomorphism of the two as simplicial complexes determines a name-preserving isomorphism of their orienting double covers as simplicial complexes, which is unique up to composition with the involution. We require that one of the two possible isomorphisms of the double cover preserves the edge names.

Definition 3.13 A triangulation T by saddle connections *degenerates* to a cell decomposition T' if there is a sequence of unit-area half-translation surfaces $\{X_i\}_{i=1}^\infty$, all with the same marked triangulation T , that converges in the Hausdorff topology in the universal curve, and each cell of T in the sequence of surfaces X_i converges in the Hausdorff topology to a cell or union of cells of a surface X with the cell decomposition T' . If v is a vertex of T in each X_i , we say that v *limits to* v' . If a collection of vertices all limit to the same $v \in v'$ in X , we say that those vertices *collide*.

Definition 3.14 Let T be a triangulation of S_g , together with a prescribed cone angle for each vertex of T . Assume T has $4g - 4 + n$ vertices with prescribed cone angle 3π (the zeros of T) and n vertices with cone angle π (the poles of T). We say a simply connected subset A of $\text{QD}(\mathcal{T}_{g,n})$ is T -convex if each $(X, q) \in A$ has a triangulation by saddle connections isotopic to T or a degeneration T' of T such that:

- The order of vanishing of q at each vertex v of T' is the sum of the orders of vanishing of all vertices v_k of a sequence (X_i, q_i) triangulated by T and degenerating to v in (X, q) . (Given a maximal collection D of vertices joined by degenerate edges of T whose union is connected and contains only nullhomotopic closed curves, the order of vanishing of q at the degenerate singularity is equal to the number of zeros minus the number of poles in D .)
- For each oriented edge γ of T there is a half-plane $H_\gamma \subset \mathbb{C}$ not containing 0 in its interior such that if we vary (X, q) continuously, we always have

$$\int_{\gamma'} \sqrt{q} \in H_\gamma.$$

- For (X, q) let $v(X, q)$ be the vector of periods of saddle connections of (X, q) that belong to T , such that for each edge γ the coordinate is in H_γ . Then for any $t \in (0, 1)$ and $(X_1, q_1), (X_2, q_2) \in A$, there is some $(X_3, q_3) \in A$ such that

$$v(X_3, q_3) = tv(X_1, q_1) + (1 - t)v(X_2, q_2).$$

We will show in Appendix A that there is a locally finite collection of triangulations T_i on $\text{QD}(\mathcal{T}_{g,n})$ such that $\text{QD}(\mathcal{T}_{g,n})$ is a locally finite union of T_i -convex sets, and the sets are invariant under scaling by $\mathbb{R}^{>0}$. The purpose of this definition is to allow us to analyze collisions of singularities while letting them retain their individual identities.

A T -convex subset of a stratum admits an embedding into $H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$ by integrating a choice of square root of \tilde{q} along relative cycles in $H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$. We always choose the square root so that

$$\int_{\gamma'} \sqrt{\tilde{q}} \in H_{\gamma}.$$

The choice of triangulation gives a local trivialization of the vector space $H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$, so this gives us a system of charts into \mathbb{C}^n , where n is the dimension of $H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$; in fact they form an atlas of charts, and the dimension of the connected component of a stratum is the dimension of $H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$.

There is a canonical identification

$$H_{\text{odd}}^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) = \text{Hom}(H_1^{\text{odd}}(\tilde{X}, \tilde{\Sigma}; \mathbb{Z}), \mathbb{C})$$

which gives rise to a natural Lebesgue measure on each stratum, by picking a basis for the free \mathbb{Z} -module $H_1^{\text{odd}}(\tilde{X}, \tilde{\Sigma}; \mathbb{Z})$; we can write any quadratic differential in terms of this basis as an element of \mathbb{C}^d for some d by taking the integral over this basis, and the natural Lebesgue measure on $\mathbb{C}^d = \mathbb{R}^{2d}$ is invariant under the adjoint action by $\text{GL}_n(\mathbb{Z})$. We refer to this measure on any stratum of $\text{QD}(\mathcal{T}_{g,n})$ as the *Masur–Veech* measure. We can obtain a measure on the space of unit-area differentials by a standard cone construction; the coned measure of a set of unit-area differentials is defined to be the measure of the union of their multiples by scalars in $(0, 1)$.

It is well known that these measures are finite on the whole stratum, and invariant under the Teichmüller geodesic flow. They were constructed by Masur [1982] and Veech [1986] in order to apply ergodic theory to the study of the Teichmüller flow. It is of use to us that the condition of having a vertical or horizontal saddle connection is Masur–Veech measure 0. If μ and ν are strata and $\nu \subset \bar{\mu}$, then ν is nowhere dense and measure 0 with respect to the Masur–Veech measure on $\bar{\mu}$, in the sense that there are neighborhoods of ν whose intersections with $\bar{\mu}$ have arbitrarily small measure.

4 Quadratic differentials on the sphere

In this section, we collect basic facts about collisions of singularities in quadratic differentials. We show that every cluster of singularities is biholomorphic to a cluster of singularities on a meromorphic quadratic differential on $\hat{\mathbb{C}}$, and describe a partial compactification of strata in terms of this uniformization.

4.1 Singularities and δ -clusters

Notation 4.1 Let A and B be two positive real quantities that depend on a common variable. We write $A \prec B$ or $B \succ A$ if there is a positive constant c such that $A \leq cB$. We write $A \asymp B$ if $A \prec B \prec A$.

We will use the following version of Mumford’s compactness criterion:

Theorem 4.2 *Let K be a closed set in the moduli space of unit-area quadratic differentials $\text{QD}^1(\mathcal{M}_{g,n})$. For $(X, q) \in K$, let $S(q)$ denote the infimum of the q -lengths of simple closed curves in X which are not nullhomotopic or homotopic to a loop around a puncture. Then $S(q) \succ 1$ as (X, q) ranges over K if and only if K is compact.*

This is really a combination of Mumford's original theorem [1971], which was about lengths of the curves in hyperbolic metrics, with any of several other works that can be used to compare q -metrics to hyperbolic metrics. See for instance [Rafi 2007] or [Maskit 1985].

There are two concrete equivalent descriptions of complex manifold structures on strata of $\text{QD}(\mathcal{T}_{0,n})$. Given a quadratic differential, we may assume three of its poles are at 0 , 1 , and ∞ , since it necessarily has four more poles than zeros (counted with multiplicity) and all of the poles are simple. Once this is fixed, it follows that each stratum can be represented locally as a space of differentials of the form $P(z)/(z(z-1)Q(z)) dz^2$, where the roots of P have prescribed multiplicities, the roots of Q do not repeat and are distinct from 0 and 1 , and the number of roots of P of each multiplicity that are also roots of Q is fixed. The locations of the roots of P and Q , together with the ratio of leading coefficients, give local coordinate charts.

It follows easily that, once we restrict to any fixed triangulation, the periods are defined by holomorphic functions of the coefficients on each stratum, or equivalently, by the locations of the singularities in $\mathbb{C} \setminus \{0, 1\}$. Periods vary holomorphically with respect to these coordinates. If the locations of singularities vary holomorphically in the plane, we can assume that the endpoints of any saddle connection remain fixed at a and b , and one other singularity remains fixed at ∞ by applying a holomorphic choice of Möbius transformation. Along a saddle connection, \sqrt{q} varies in a 1-parameter family h_t . We can compute the derivative as

$$\frac{d}{dt} \int_a^b h_t(z) dz,$$

which we can simply differentiate under the integral over a fixed contour of integration. In fact, even if the endpoints do not remain fixed, differentiation under the integral remains valid when the endpoints are zeros of the differential. If we never allow our singularities to come close to colliding, this is an effective way to estimate the order of magnitude of $d_{\text{Euclidean}}$ as a function of the change of the locations of the singularities. When collections of singularities collide or come close, we can make estimates from two perspectives: one, that they never really collided (we zoom in) or two, that they stay together (we zoom out). To aid us in analyzing these clusters of singularities we make the following definition:

Definition 4.3 Given $0 < \delta < 1$, a metric space X , and a discrete set $S \subset X$, we say that $D \subset S$ is a δ -cluster in S if D contains at least two points, and the distance between any two points in D is at most δ times the distance from any point in D to any point in $S \setminus D$.

We will first show that when a subset of the singularities forms a δ -cluster D , and the diameter of D is small compared to the lengths of all nontrivial simple closed curves, there are enough saddle connections connecting singularities in D .

Definition 4.4 Given a graph G and a subset W of the vertices, the *induced* subgraph with respect to W is the maximal subgraph whose vertex set is W .

The Delaunay triangulation has the following property:

Lemma 4.5 Suppose $(X, (x_1, \dots, x_n), q) \in \text{QD}(\mathcal{T}_{g,n})$ or $\text{QD}(\mathcal{M}_{g,n})$. Let $S(q)$ denote the infimum length of simple closed curves on $X \setminus (x_1, \dots, x_n)$ that are not homotopic to punctures or constant maps, where lengths are taken in the q -metric, and let $\Delta(q)$ be the diameter of X in the q -metric. Then for each $\epsilon > 0$ there is a number $\delta_0 > 0$ such that whenever $S(q)/\Delta(q) > \epsilon$ and $\delta < \delta_0$, for every δ -cluster D in the singularity set of q , the induced subgraph on the vertex set D from the Delaunay triangulation is a connected graph. Moreover, D contains at most one marked point and all cycles in this induced subgraph are nullhomotopic in X .

Proof We may scale the metric by a real number so that $\Delta(q) = 1$, and we may pick δ to be less than $\frac{1}{8}S(q)$. We may also assume $\delta < \frac{1}{8}$. We note that the distance between two distinct poles is at least half of $S(q)$, so no δ -cluster contains two distinct marked points.

Now suppose that T_1 and T_2 are disjoint subsets of T with no edges between them in the Delaunay triangulation, and $T_1 \cup T_2 = T$. We may consider the completion of the metric universal cover of $X \setminus \{x_1, \dots, x_n\}$. Let v be a saddle connection connecting points that project into T_1 and T_2 whose length is minimal. Indeed, the shortest path from T_1 to T_2 must be a single saddle connection, since if it were a sequence of more than one saddle connection, any intermediate singularity would be closer to the start and end than the start and end are to each other, and must therefore belong to any δ -cluster containing the start and end. Pick a lift of this saddle connection and, starting at its midpoint, move along the locus of points equidistant from its ends until a point is reached that is equidistant from a third singularity. It may be the case that we do not have to move at all. We eventually meet a point p which is equidistant from our two original singularities plus some more, which is the center of a 2-cell in the decomposition into convex polygons, a triangulation of which forms the Delaunay triangulation. Moving around the boundary of the circle in one direction or the other, it must be the case that all singularities we hit are at least as close to one of the ends of the saddle connection we started on the midpoint of as to each other, and therefore they are edges in the Delaunay triangulation (no matter how it is chosen). Thus they are in the original δ -cluster, and the sequence of chords connecting them projects down to sequence of edges in the Delaunay triangulation that connect T_1 and T_2 , all of which have both ends in T . \square

Notation 4.6 Let (X, q) be a quadratic differential, and let γ be an (oriented) arc such that the endpoints of γ are the only singularities of the q -metric contained in γ . Let $\tilde{\gamma}^1$ and $\tilde{\gamma}^2$ be the two oriented lifts of γ in the orienting double cover (\tilde{X}, \tilde{q}) .

Lemma 4.7 *Given a quadratic differential $(X, q) \in \text{QD}(\mathcal{T}_{0,n})$ with a simple pole at ∞ and a collection of saddle connections forming a tree B whose vertex set is the set of singularities of (X, q) that are not ∞ , let $\gamma_1, \dots, \gamma_r$ be oriented saddle connections in X and let γ_i'' and γ_i' be the lifts of γ_i in the orienting double cover of (X, q) . Then $\{[\gamma_i''] - [\gamma_i'] : 1 \leq i \leq r\}$ is a basis for $H_1^{\text{odd}}(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$.*

Proof In \tilde{X} , the inverse image of B is a graph \tilde{B} , and $\tilde{X} \setminus \tilde{B}$ is a ramified double cover of $\hat{\mathbb{C}} \setminus B$ whose only branching is of order 2 over ∞ ; call the ramification point $\tilde{\infty}$. This is biholomorphic to a disk. Thus \tilde{B} is the 1–skeleton of a CW decomposition of \tilde{X} with a unique 2–cell, which we may call U . The boundary of U is zero (which can be seen since $H^2(X, \mathbb{Z}) \neq 0$ and there is only one 2–cell, or alternatively because its boundary wraps twice around the graph, on opposite sheets, and hence ∂U traverses each 1–cell twice in each direction). So $\tilde{B} \hookrightarrow X$ induces isomorphisms $H_i(\tilde{B}; \mathbb{C}) \rightarrow H_i(\tilde{X}; \mathbb{C})$ for $i \leq 1$. Since \tilde{B} is invariant with respect to the involution, the isomorphism on the first homology respects the even–odd decomposition. So all absolute cycles in $H_1^{\text{odd}}(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$ are homologous to cycles in $H_1^{\text{odd}}(\tilde{B}, \tilde{\Sigma} \setminus \{\tilde{\infty}\}; \mathbb{C})$. It is clear that all of the absolute cycles are in the odd part, for both \tilde{B} and \tilde{X} , since $\hat{\mathbb{C}}$ and B have no homology in dimension 1.

From the long exact sequences of the pairs $(\tilde{X}, \tilde{\Sigma})$ and $(\tilde{B}, \tilde{\Sigma} \setminus \{\tilde{\infty}\})$ we get natural inclusions

$$H_1(\tilde{X}; \mathbb{C}) \hookrightarrow H_1(\tilde{X}, \tilde{\Sigma}; \mathbb{C}) \quad \text{and} \quad H_1(\tilde{B}; \mathbb{C}) \hookrightarrow H_1(\tilde{B}, \tilde{\Sigma} \setminus \{\tilde{\infty}\}; \mathbb{C}).$$

An element of either quotient group is exactly determined by the image of any cycle representing it under the connecting map into $H_0(\tilde{\Sigma})$ or $H_0(\tilde{\Sigma} \setminus \{\tilde{\infty}\})$. If such an element is odd, it must be the case that its boundary is a linear combination terms of the form $\{p_i - \iota(p_i)\}$ for $p_i \in \Sigma$. Since $\tilde{\infty}$ is invariant under ι , the exact same relative cycle classes are realizable as well. □

Corollary 4.8 *The periods of B form a holomorphic coordinate chart.* □

A similar argument shows that, when $\tilde{\cdot}$ denotes the preimage of \cdot in the orienting double cover:

Corollary 4.9 *If γ is a simple closed curve containing no singularities which bounds a disk $\mathbb{D} \subset X$, and the singularities in D form a set S , then the periods of any tree with vertex set S and all edges saddle connections contained in \mathbb{D} form a basis for $H^1(\tilde{\mathbb{D}}, \tilde{S}; \mathbb{C})$ as a vector space over \mathbb{C} .*

Lemma 4.10 *Let $(X, q) \in \text{QD}(\mathcal{T}_{0,n})$. If Σ is its singularity set and p is a pole in Σ , then any Delaunay triangulation has a connected induced subgraph with respect to $\Sigma \setminus \{p\}$.*

Remark This property is shared with the L^∞ Delaunay triangulations considered in Appendix A.

Proof There is only one ray of each slope emanating from p , so there cannot be a saddle connection that starts and ends at p . The edges emanating from p are cyclically ordered counterclockwise, and there is an edge connecting each consecutive pair of them, so for any pair of neighbors of p there is a path in the graph avoiding p . Therefore p is not a cut vertex of the 1–skeleton. □

On compact sets, any two holomorphic coordinate systems are bi-Lipschitz, so the main difficulty will be in dealing with degenerations to lower-dimensional strata. Our main object will be to examine periods of clusters of singularities close to 0.

In the sequel, we will often have to deal with a meromorphic quadratic differential q defined on \mathbb{C} , and we will often have to speak of distances between points, diameters of sets, etc in two different metrics: the usual Euclidean metric on $\mathbb{C} = \mathbb{R}^2$ (which does not depend on the choice of q) and the singular metric that depends on q .

Notation 4.11 We use $d_{\mathbb{C}}$ and $\text{diam}_{\mathbb{C}}$ to denote the distance and diameter in the usual metric on \mathbb{C} , and d_q and diam_q to denote the distance and diameter in the singular metric defined by q . We use $\text{perim}_{\mathbb{C}}$ and perim_q to denote perimeters of regions, and $\text{rds}_{\mathbb{C}}(A)$ and $\text{rds}_q(A)$ will denote the maximum radius of a ball contained entirely in A with respect to $d_{\mathbb{C}}$ and d_q .

We now discuss a certain class of quadratic differential on \mathbb{C} which we use to model clusters of singularities. These differentials will extend meromorphically to $\hat{\mathbb{C}}$, but will have higher-order poles at ∞ .

Definition 4.12 We say q is a *cluster differential* if is a quadratic differential on $X = \mathbb{C}$ which takes one of the following forms: $q = p(z) dz^2$, where p is a monic polynomial of degree at least 2 whose roots sum to 0, or $q = p(z)/z dz^2$. If p is of the second type, we say that 0 is a *marked point*.

We consider two cluster differentials to be distinct if one has a marked point and the other does not, even though they may have the same underlying quadratic differential on \mathbb{C} . That is to say, 0 is *allowed* to be a root of p .

The sets of cluster differentials with a fixed degree and number of marked points are complex manifolds, with the coefficients of the polynomial as coordinates. The complex dimension of the space of cluster differentials is $\deg(p) - 1$ if the differentials are of the form $p(z)$ and there is no marked point, and $\deg(p)$ if there is a marked point.

Spaces of cluster differentials are also stratified by the numbers and types of singularities, and the periods of any spanning tree of saddle connections forms a local holomorphic coordinate system.

In fact, the periods of the saddle connections in the tree and the cone angles determine the q -metric entirely, since the metric has constant curvature 0 away from the tree. This determines (\mathbb{C}, q) as a metric space, from which the conformal structure can be recovered; this determines the differential and the locations of its singularities up to a Möbius transformation fixing ∞ . Since we know that p is monic and we know which direction is vertical, we can recover p up to multiplying all of its singularities by a root of unity. If the singularities have names, then locally there is a unique correct choice. We may write q as $q(z_1, \dots, z_n, w)$ where n is the number of zeros of q and w is the marked point. We are of course constrained to the hyperplane $\sum(z_i) = 0$ if there is no marked point.

If a cluster differential has no poles, then its infinitesimal metric is locally nonpositively curved in the sense of Alexandrov. If a quadratic differential induces a complete Alexandrov nonpositively curved metric on a simply connected Riemann surface, then there is a unique length-minimizing geodesic between every pair of points. In particular, a cluster differential has only finitely many saddle connections if it has no poles. We also see that if we take a branched double cover of a cluster differential branched only over the pole and pull back the differential, we get a cluster differential with no poles and a marked point. All saddle connections on our original cluster differential lift to one of finitely many saddle connections upstairs so there were only finitely many saddle connections on the original cluster differential, and the number of saddle connections is bounded by the number of pairs of singularities in the double cover.

There is an analogous definition of T -convex for cluster differentials: we may say that a set A is Γ -convex if Γ is the graph whose vertices are singularities and whose edges are saddle connections, the periods of each saddle connection that is an edge of Γ vary continuously in a half-plane, and we can take convex combinations in period coordinates without leaving A .

Proposition 4.13 *The effect of multiplying the zeros and poles of a cluster differential q by t is to scale all periods by $t^{(m+2)/2}$, where m is the rational function degree of q , and angles between saddle connections meeting at a singularity are preserved.*

Proof We give the proof for $q = p(z) dz^2$ without the marked point since the proof with a marked point is identical. The first statement is a straightforward change of variables:

$$\int_{ta}^{tb} [t^m p(z/t)]^{1/2} dz = \int_a^b t^{m/2} [p(z)]^{1/2} t dz.$$

The two sides are the values of the two periods in question, which clearly differ by $t^{(m+2)/2}$. The only way that the cone angles can differ is by integral multiples of π , but they don't differ at all since we can vary t continuously and cone angles vary continuously with t . □

A similar change of variables tells us the following:

Proposition 4.14 *Fix a stratum \mathcal{Q} of quadratic differential with n singularities in \mathbb{C} , and let m be the total number of zeros minus the number of poles (counted with multiplicity). For an n -tuple p of distinct points in \mathbb{C} which can be the singularity set of an element in \mathcal{Q} , let q_p be the cluster differential with those singularities (it is assumed that for each coordinate the singularity has a prescribed type). Then for any $v \in \mathbb{C}^n$ and any period P_i of any saddle connection, we have*

$$\left. \frac{d}{dt} P_i(q_{sp+tv}) \right|_{t=0} = s^{m/2} \left. \frac{d}{dt} P_i(q_{p+tv}) \right|_{t=0}.$$

Corollary 4.15 *Let q be a cluster differential of rational function degree m . Let Σ be the set of singularities of q . Then*

$$\text{diam}_q(\Sigma) \asymp \text{diam}_{\mathbb{C}}(\Sigma)^{(2+m)/2}.$$

Proof By Proposition 4.13 this reduces to the following claim: if $\text{diam}_{\mathbb{C}}(\Sigma) = 1$, then $\text{diam}_q(\Sigma) \asymp 1$.

We will only deal with the case in which there is a marked point since that is more difficult. What we will do is prove an upper bound on the q -diameter of the $d_{\mathbb{C}}$ ball of radius 1 about the marked point and a lower bound on the q -length of an arc whose endpoints belong to Σ and are distance 1 apart with respect to $d_{\mathbb{C}}$.

Indeed, suppose that the marked point is 0. We will give an upper bound for the q -length of any radius of the unit circle $|z| = 1$, since any two points in Σ can be joined by a path contained in a pair of two such radii. By rotating our coordinate system, we may assume that the radius is the interval $[0, 1] \subset \mathbb{R}$.

If $q = f(z)/z dz^2$ and f has degree $m + 1$, then for all points $x \in [0, 1]$ we have $|f(x)/x| \leq (x + 1)^n/x$. Therefore, the q -length of the line segment $[0, 1]$ is bounded by

$$d_q(0, 1) \leq \int_0^1 x^{-1/2} dx.$$

This gives us the upper bound.

For a lower bound, note that a rectifiable arc γ from a to b with $d_{\mathbb{C}}(a, b) = 1$ must travel Euclidean distance at least $\frac{1}{2}$ outside of the following union of $m + 2$ disks:

$$\bigcup_{z_i \in \Sigma} \left\{ z : |z - z_i| < \frac{1}{4(m+2)} \right\}.$$

As before, assume that $q = [f(z)/z] dz^2$ and z_1, \dots, z_{m+1} are the roots of f . Then

$$\left| \frac{f(z)}{z} \right| \geq \left| \frac{z - z_{m+1}}{z} \right| \prod_{j=1}^m |z - z_j| \geq \frac{|z - z_{m+1}|}{1 + |z - z_{m+1}|} \prod_{j=1}^m |z - z_j| \geq \frac{1/(4m + 8)}{1 + 1/(4m + 8)} \frac{1}{(4m + 8)^m}.$$

The length of γ is therefore bounded below by

$$\int_{\gamma} |\sqrt{q}| \geq \frac{1}{2} \left| \frac{1/(4m + 8)}{1 + 1/(4m + 8)} \frac{1}{(4m + 8)^m} \right|^{1/2}. \quad \square$$

Proposition 4.16 *Let f be a monic polynomial of degree d with $f(z) \neq 0$ whenever $|z| < 1$, and assume $r < 1$. Then there is a constant $C_{r,d}$, depending on r and d but not on f , such that on the disk $\{|z| < r\}$ we have $C_{r,d}^{-1} \leq |f(z)/f(0)| \leq C_{r,d}$. Moreover, if r is small enough we may take C arbitrarily close to 1.*

Proof For a monic polynomial $z - w$, this follows from the fact that whenever $|w| \geq 1$ and $|z_1|, |z_2| < r$, we have

$$\left| \frac{z_1 - w}{z_2 - w} - 1 \right| \leq \frac{|z_1| + |z_2|}{|w| - r} \leq \frac{2r}{1 - r}. \quad \square$$

Notation 4.17 For the purpose of Corollary 4.18, let B_s denote the disk $\{z : z < s\}$.

Corollary 4.18 *Let $r < \frac{1}{4}$. Let $p(z) dz^2$ be a cluster differential with a pole of order $m + 4$ at ∞ and no singularities in \mathbb{C} outside B_r . Let f and g be polynomials with fixed degrees and no zeros in B_1 , and let $q = (f(z)/g(z))p(z) dz^2$.*

Then

$$d_q(0, \partial B_r) \asymp \left| \frac{f(0)}{g(0)} \right|^{1/2} r^{(2+m)/2}, \quad \text{rds}_q(B_r) \asymp \left| \frac{f(0)}{g(0)} \right|^{1/2} r^{(2+m)/2},$$

$$\text{diam}_q(B_r) \asymp \left| \frac{f(0)}{g(0)} \right|^{1/2} r^{(2+m)/2}, \quad \text{perim}_q(B_r) \asymp \left| \frac{f(0)}{g(0)} \right|^{1/2} r^{(2+m)/2}.$$

The implied constants depend on m and the degrees of f and g .

Proof For the base case where f and g are both constant, we can apply Corollary 4.15. For other cases, we may simply apply Proposition 4.16. \square

Corollary 4.19 *Let q be as in Corollary 4.18. Then for each $C \in (0, 1)$ there is a constant $c > 0$ such that if all zeros and poles of p are contained in B_{cR} , any length-minimizing path in the d_q -metric between points in B_{cR} is contained in B_{CR} .*

Proof By the estimates in Corollary 4.18, if c/C is small enough, then the diameter of B_{cR} is less than the d_q distance from ∂B_{cR} to ∂B_{CR} . \square

The following lemma says that while a faraway singularity may change the size of a δ -cluster, it does little to change the shape:

Lemma 4.20 *Let f, g , and B_s be as in Corollary 4.18. Let $q = p(z) dz^2$ be a cluster differential with all singularities in B_r , and let $q' = (f(z)/g(z))p(z) dz^2$. Then for each $\epsilon > 0$ there is some $r_0 > 0$ such that whenever $r < r_0$ and $a, b, c, d \in B_r$,*

$$1 - \epsilon < \frac{d_q(a, b)d_{q'}(c, d)}{d_q(c, d)d_{q'}(a, b)} < 1 + \epsilon.$$

Moreover, if Γ is a tree whose vertices are the singularities of p and whose edges are saddle connections whose lengths with respect to d_q are as small as possible, then there is a tree Γ' whose endpoints are saddle connections of q' and whose edges are homotopic (rel endpoints) to those of Γ whenever $r < r_0$ is sufficiently small, and angles between corresponding edges differ by less than ϵ .

Proof The first claim is immediate from the pointwise estimates (Proposition 4.16), since the metric is scaled pointwise by a near constant.

For the second claim, we need to show that if we start with our embedding of the graph $\Gamma \subset \mathbb{C}$ and modify it by making its edges geodesic with respect to $d_{q'}$, the edges remain saddle connections. Γ is obtained by a greedy algorithm: pick saddle connections one at a time, repeatedly picking the shortest saddle connection that does not join two vertices belonging to the same connected component of the graph with edges already picked. (This is because for any saddle connection e not used, the fundamental cycle of $\Gamma \cup e$ with e must consist of edges no longer than e .) Therefore, whenever an edge e is drawn, it must be the case that there is no singularity of distance less than the length of e away from both ends

of e , and there are two equilateral triangles sharing the edge e with no singularities in the interior of either triangle. Also there are no other singularities on e . It follows that the angle formed by e and any previously drawn edge is at least $\frac{1}{3}\pi$.

The estimate of Proposition 4.16 then implies that the lengths and angles of all tangent vectors are scaled by nearly the same constant, so for any points a and b joined by a q -geodesic with no singularities, we have

$$\frac{\int_a^b \sqrt{q'}}{\int_a^b \sqrt{q}} = \sqrt{f(0)/g(0)}[1 + o(1)] \quad \text{as } r_0 \rightarrow 0.$$

We may scale the entire differential by a constant without changing the geodesics, so we may assume $\sqrt{f(0)/g(0)} = 1$. Now, if we have a map from a region bounded by the union of two equilateral triangles into \mathbb{R}^2 sharing an edge, and its derivative is close enough to $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, then the geodesic in \mathbb{R}^2 joining the endpoints is contained in the image of the region. □

Proposition 4.21 *Let f and g be polynomials such that g has no repeated roots, and $\deg(f) \geq \deg(g) - 3$. Let $q = (f(z)/g(z)) dz^2$ be a meromorphic quadratic differential on $\hat{\mathbb{C}}$ whose poles are all simple, except perhaps for a higher-order pole at ∞ . Assume that the roots of f and g are all contained in B_{r_0} for some $r_0 > 0$. Then for each $r > 0$ there is a number R depending only on r, r_0 , and the degrees of f and g , such that the shortest path between any points in B_r with respect to the metric d_q is contained in B_R .*

Proof We will assume f and g are monic, since scaling the metric by a constant doesn't change whether or not a path is length-minimizing. Let $m = \deg(f) - \deg(g)$. As $|z| \rightarrow \infty$, we have $|f(z)/g(z)|^{1/2}/|z^{m/2}| \rightarrow 1$ uniformly for all q satisfying our hypotheses. Integrating this pointwise bound gives us the following:

The distance from the circle of radius r to the circle of radius r^2 has asymptotics

$$\int_r^{r^2} |\sqrt{q}| \sim \int_r^{r^2} |z^{m/2} dz| \sim \begin{cases} 2r^{-1/2} & \text{if } m = -3, \\ \log(r) & \text{if } m = -2, \\ 2r^{m+2}/(m+2) & \text{if } m \geq -1. \end{cases}$$

As $r \rightarrow \infty$, the ratio between the actual q -distance between ∂B_r and ∂B_{r^2} and the asymptotic value converges to 1 at a rate that depends only on r_0 , and the degrees of f and g .

The q -lengths of semicircles on ∂B_r are asymptotic to $\pi r \cdot r^{m/2}$. Thus, the distance between two points on ∂B_r is at most $[1 + o(1)]\frac{\pi}{4}$ times the distance from ∂B_r to ∂B_{r^2} and back when $m = -3$, and $o(1)$ times the distance from ∂B_r to ∂B_{r^2} and back for $m \geq -3$. □

We could have used B_{cr} in place of B_{r^2} and gotten the same conclusion in the case $m = -3$, for any sufficiently large C . By performing the change of coordinates $z \mapsto 1/z$ and $c = 1/C$, we conclude the following:

Proposition 4.22 *If a quadratic differential q on $\widehat{\mathbb{C}}$ has only simple poles, one occurring at $z = 0$, and no other poles or zeros in the ball B_r , then there is some $c \in (0, 1)$, depending only the number of zeros and poles of q of each multiplicity, such that no length-minimizing path from two points outside of B_r passes through B_{cr} .* \square

Proposition 4.23 *Fix integers $e_1, \dots, e_k \geq -1$ with $\sum_{i=1}^k e_i \geq -3$. Let $q(\alpha, z_1, \dots, z_k)$ denote the quadratic differential $[\prod_{i=1}^k (z - z_i)^{e_i}] dz^2$ on \mathbb{C} . (It will be meromorphic on the Riemann sphere with a pole at ∞ which is simple if and only if $\sum e_i = -3$.) Assume the set of tuples $(\alpha, z_1, \dots, z_k)$ is restricted to the set where $z_i \neq z_j$ if $e_i = e_j = -1$. Then for all $i, j \in \{1, \dots, k\}$, we have that $d_{q(\alpha, z_1, \dots, z_k)}(z_i, z_j)$ varies continuously with respect to $(\alpha, z_1, \dots, z_k)$.*

Proof Let $\epsilon > 0$. By the estimates of Corollary 4.18, the q -diameter, radius, perimeter, and distance from the boundary to the point z_i of

$$B_{r_i}(z_i) := \{z : |z - z_i| < r_i\}$$

are all $O(r_i^{1/2})$ as $r \rightarrow 0$. We pick around all singularities a ball small enough that each has diameter less than ϵ , and also a ball B_∞ about ∞ such that no length-minimizing path between z_i and z_j can enter. We can pick the balls much smaller than the distance between any two points z_i and z_j unless $z_i = z_j$. In this case we pick $r_i = r_j$ in a neighborhood of $(\alpha, z_1, \dots, z_k)$. We can fix r_1, \dots, r_k and α so that each of these balls has diameter, radius, and perimeter less than ϵ in the q -metric in a neighborhood of $(\alpha, z_1, \dots, z_k)$. Then, as (β, y_1, \dots, y_k) converges to $(\alpha, z_1, \dots, z_k)$, the sizes of the balls $B_i(z_i)$ still satisfy

$$\text{diam}_{q(\beta, y_1, \dots, y_k)}(B_{r_i}(z_i)) \prec \epsilon.$$

The distances $d_{q(\beta, y_1, \dots, y_k)}(B_{r_i}(z_i), B_{r_j}(z_j))$ vary continuously since the metric varies smoothly on $\widehat{\mathbb{C}} \setminus (B_\infty \cup \bigcup_{i=1}^k B_{r_i}(z_i))$. Since length-minimizing paths from $B_{r_i}(z_i)$ to $B_{r_j}(z_j)$ stay in this region, it follows that $d_{q(\beta, y_1, \dots, y_k)}(y_i, y_j)$, viewed as a function of (β, y_1, \dots, y_k) , can be written as a sum of a continuous function and a function taking values in $[-\epsilon, \epsilon]$. Since ϵ was arbitrary, the proposition follows. \square

Definition 4.24 *A Möbius normalization of $\mathcal{T}_{0,n}$ or $\text{QD}(\mathcal{T}_{0,n})$ is a collection of uniformization maps of the underlying Riemann surfaces to $\widehat{\mathbb{C}}$ such that the points mapped to 0, 1, and ∞ define three continuous sections of the universal curve or universal half-translation surface, and all each of these three sections intersects each fiber at a marked point. A Möbius normalized collection A of quadratic differentials on $\mathcal{T}_{0,n}$ is the set of quadratic differentials on $\widehat{\mathbb{C}}$ that pull back to elements of A across the uniformization maps.*

Proposition 4.25 *The distance $d_{q(\alpha, z_1, \dots, z_k)}(a, b)$ is jointly continuous in $(\alpha, z_1, \dots, z_k, a, b)$.*

Proof We can just pretend a and b are singularities z_{k+1} and z_{k+2} with $e_{k+1} = e_{k+2} = 0$ and apply Proposition 4.23. \square

Corollary 4.26 Fix a Möbius normalization of $\text{QD}(\mathcal{T}_{0,n})$ and a compact set $K \subset \text{QD}(\mathcal{T}_{0,n})$ for which ∞ is always a pole. Then for any compact subset $V \subset \mathbb{C}$ and any $\epsilon > 0$, there is a number $\delta > 0$ such that for all $q \in K$, $d_q(x, y) < \epsilon$ whenever $x, y \in V$ and $d_{\mathbb{C}}(x, y) < \delta$. The same is true if we allow K to vary over a compact set of cluster differentials.

Proof The set $\{(x, y, q) : x, y \in V \text{ and } d_q(x, y) \geq \epsilon\}$ is compact so $d_{\mathbb{C}}(x, y)$ attains a minimum there, which we can take to be δ . \square

Corollary 4.27 Suppose $\{t_m\} \rightarrow t_\infty$ and $\{u_m\} \rightarrow u_\infty$ are convergent sequences in \mathbb{C} and $\{q_m\} \rightarrow q_\infty$ is a Möbius-normalized convergent sequence in $\text{QD}(\mathcal{T}_{0,n})$ or a space of cluster differentials. Then any sequence $\gamma_m : [0, 1] \rightarrow \mathbb{C}$ of d_{q_m} -geodesics from t_m to u_m which are length minimizing has a subsequence that converges uniformly (as \mathbb{C} -valued functions) to a constant-speed length-minimizing geodesic in q_m along a subsequence.

Proof By Proposition 4.25 it is clear that any subsequential limit of constant-speed length-minimizing geodesics must be a constant-speed length-minimizing geodesic. By a diagonalization argument, we can pass to a subsequence such that $\gamma_m(t)$ converges for all $t \in [0, 1] \cap \mathbb{Q}$. Uniform convergence follows from Corollary 4.26. \square

Corollary 4.28 Suppose that $q_n \rightarrow q_\infty$ in $\text{QD}(\mathcal{T}_{g,n})$ or a space of cluster differentials. Let $\gamma_n : [0, 1] \rightarrow \mathbb{C}$ be a geodesic arc for the q_n -metric, parametrized to be constant speed, such that the γ_n converge to a geodesic arc γ_∞ in the q_∞ -metric that does not pass through any singularities. Then the sequence of maps γ_n converges to γ_∞ in the C^1 topology on $C^1([0, 1])$.

Proof We already know that the lengths converge. It is easy to see that the directions converge as well, since the arcs formed by concatenating the q_∞ geodesic from $\gamma_\infty(0)$ to $\gamma_n(0)$, the q_n -geodesic from $\gamma_n(a)$ to $\gamma_n(b)$, and the q_∞ -geodesic from $\gamma_n(1)$ to $\gamma_\infty(1)$ are homotopic rel endpoints, via homotopies that pass through no singularities, to γ_∞ for all sufficiently large n . The integrands of $\sqrt{q_n}$ converge to $\sqrt{q_\infty}$ on an open set containing all of these arcs. It thus follows that, if we take the correct branch of the square root,

$$\int_{\gamma_n(0)}^{\gamma_n(1)} \sqrt{q_n} \rightarrow \int_{\gamma_\infty(0)}^{\gamma_\infty(1)} \sqrt{q_\infty}.$$

This implies that the complex lengths ℓ_n of the segments $\gamma_n([0, 1])$ converge to some ℓ_∞ . By Corollary 4.27 and the fact that the tangent vector that maps \sqrt{q} to 1 varies continuously in (q, z) away from singularities, we conclude that ℓ_i / \sqrt{q} , the lengths and directions of the tangent vectors $\{\gamma_n'(t) : t \in [0, 1]\}$ converge uniformly in $T\mathbb{C}$. \square

Fix a space of cluster differentials, or else fix $c > 0$ and a compact subset K of $\text{QD}(\mathcal{T}_{0,n})$, and assume that for each $(X, q) \in K$, we pick a marking such that ∞ is a pole and such that for all singularities p , we have $d_q(p, \infty) > c\Delta(q)$. (Recall that $\Delta(q)$ is the diameter of the q -metric.) This choice is invariant under the group of Möbius transformations fixing ∞ . Then we have the following:

Proposition 4.29 *The notions of δ -cluster with respect to \mathbb{C} -metric and q -metric are comparable on K in the following sense:*

- *For each sufficiently small $\delta > 0$ there exists $\delta' > 0$, depending only on δ and K , such that whenever S is a δ -cluster of the singularity set of $q \in K$ with respect to the \mathbb{C} -metric, S is a δ' -cluster of singularities with respect to the q -metric.*
- *Conversely, for each sufficiently small $\delta > 0$ we can find $\delta' > 0$, depending only on δ and K , such that whenever S is a δ -cluster of the singularity set of $q \in K$ with respect to the q -metric, S is a δ' -cluster of singularities with respect to the \mathbb{C} -metric.*
- *The two statements above are true for strata of cluster differentials in place of differentials in $\text{QD}(\mathcal{T}_{0,n})$.*

In other words, we can tell when a collection of singularities is collapsing strictly faster than any proper superset using either metric.

Proof By scale invariance of the ratios of distances in both metrics, we will assume all quadratic differentials in K are of the form $(P(z)/Q(z)) dz^2$ with P and Q monic of fixed degrees. P is allowed to have multiple roots and Q is not, and P and Q may have roots in common. By applying Möbius translations fixing ∞ , we can also assume that the set Σ of singularities of each q is contained in a ball of radius R about 0 but not in a ball of radius r , for some $R > r > 0$. A closed subset of the space of differentials with these restrictions is compact if and only if there is a lower bound on the distance between any two roots of Q ; conversely, for any K we can find some R and r constraining the singularities of q for all $q \in K$, after a Möbius transformation.

Let $\{q_n\}_{n=1}^\infty$ be a sequence in K , uniformized to $\widehat{\mathbb{C}}$ so as to satisfy the constraints in the above paragraph, and let S_n be a collection of singularities in the q_n -metrics. Then by Corollary 4.18, $\text{diam}_{\mathbb{C}}(S_n) \rightarrow 0$ if and only if $\text{diam}_{q_n}(S_n) \rightarrow 0$. We may rephrase this as $\text{diam}_{q_n}(S_n)/\Delta(q_n) \rightarrow 0$ if for each $\delta > 0$, S_n is eventually contained in a δ -cluster with respect to $d_{\mathbb{C}}$, if and only if for each $\delta > 0$, S_n is eventually contained in a δ -cluster with respect to d_{q_n} .

Now let $S_n \subsetneq S'_n$ be sets of singularities such that S'_n is a δ_n cluster in q_n , and $\text{diam}_{q_n}(S_n) \rightarrow 0$. By Lemma 4.20 we can ignore all singularities outside of S'_n and use the above argument to deduce that $\text{diam}_{\mathbb{C}}(S_n)/\text{diam}_{\mathbb{C}}(S'_n) \rightarrow 0$ if and only if $\text{diam}_{q_n}(S_n)/\text{diam}_{q_n}(S'_n) \rightarrow 0$.

The proposition then follows by induction on the total number of singularities. □

Definition 4.30 Let $\delta < \frac{1}{10}$. We say that a saddle connection is *internal* to a δ -cluster D in the q -metric if it stays in the $2 \text{diam}_q(D)$ -neighborhood of D (in the q -metric). (In particular, this includes any saddle connection that is the shortest path from one endpoint to the other.)

Definition 4.31 Given a meromorphic quadratic differential $(P(z)/Q(z)) dz^2$ on \mathbb{C} whose poles are all simple, we say that a δ -cluster of singularities (where δ -cluster is with respect to the flat metric) is *shrunk*

if it contains more zeros than poles (counted with multiplicity). Given a shrunk δ -cluster of singularities, we say that the *center* of the δ -cluster is the weighted average of the singularities, where each point is counted with weight equal to the order of vanishing of $P(z)/Q(z)$ at that point (simple poles count as -1 .)

These definitions are motivated by the fact that the diameter in the singular q -metric shrinks much faster than the diameter in the nonsingular metric as a cluster converges to a point and all other singularities remain fixed, and all other periods in a good coordinate system are well approximated by the periods of the quadratic differential that is obtained when we replace the δ -cluster with a single singularity at its center, which is a zero of order equal to the sum of the orders of vanishing of P/Q at the singularities being replaced.

The equivalence of complex coordinate systems then follows because a holomorphic local homeomorphism of complex manifolds is a local biholomorphism — that is, the locations of the singularities in \mathbb{C} are holomorphic functions of the periods.

4.2 Modeling collisions of singularities by cluster differentials

We claim that cluster differentials form a model for a neighborhood of any disk without poles:

Proposition 4.32 *If (X, q) is a simply connected Riemann surface equipped with a holomorphic quadratic differential with finitely many zeros and the metric d_q is complete, then (X, q) is determined up to isometry by the isometry type of the geodesic convex hull of its zeros and the exterior angles at the vertices of the convex hull.*

Proof Assume that a ball of radius r about a singularity p contains all singularities. By Alexandrov nonpositive curvature, there is a unique geodesic between any pair of points. Moreover, geodesic rays from a fixed point p have a circular order from their initial direction, and the directions they exit singularities if they enter the same singularity. We can therefore construct all geodesic rays from p , and reconstruct the complement of the geodesic convex hull of the zeros as a union of sectors of disks of infinite radii. \square

Proposition 4.33 *If (X, q) is a (noncompact) simply connected Riemann surface with a complete metric d_q coming from a holomorphic quadratic differential q , and q has finitely many zeros and no poles, then (X, q) is a cluster differential.*

Proof Given (X, q) we will construct a sequence of quadratic meromorphic differentials $\{q_n\}$ on $\widehat{\mathbb{C}}$ that converge to a holomorphic quadratic differential q_∞ on every compact subset of \mathbb{C} . Consider the balls $B_R(p)$ in (X, q) for some fixed singularity $p \in X$. Now, the circle $c_R(p)$ of radius R about p has geodesic curvature $1/(R - s)$ at z if the geodesic from p to z last changed direction at a cone point distance s from z . In particular, the geodesic curvature of $c_R(p)$ is between $1/R$ and $1/(R - r)$ if $B_r(p)$ contains all singularities of d_q , except at finitely many points where the curvature is undefined, and never changes sign where the curvature is discontinuous.

Step 1 (construction of $\{q_n\}$) There are geodesic tangent lines to $c_R(p)$, and a finite number of these lines are vertical or horizontal. For large R , the number of such lines remains fixed, their points of tangency vary continuously in R , and they alternate vertical and horizontal. If $V(R)$ and $H(R)$ are consecutive vertical and horizontal tangent lines with respect to the circular order, then V intersects H . The convex hull of all such intersection points is a piecewise geodesic arc with right-angles at its singularities, and it bounds a disk D_R . We will take (X_n, q_n) , a sequence of quadratic differentials, to be the union of disks D_{R_n} with their Schwarz reflections about the boundary (as metric spaces). Note that this Schwarz reflection gives us a canonical way to extend the vertical and horizontal foliations, and creates cone points of angle π at the corners of ∂D_R . However, we have not yet described the sequence of uniformization maps on X_n that makes q_n convergent on compact sets.

Step 2 (picking coordinates for $\{X_n\}$) We note that if we rescale q_n to have unit area, then the d_q distances between poles remain bounded below. (The area of D_R is a quadratic polynomial in R for large R , and the distances between poles grow linearly.) Thus we can apply Mumford’s compactness criterion. By the equivalence of δ -clusters in d_{q_n} and $d_{\mathbb{C}}$, for any sufficiently small positive δ there is an $N < \infty$ such that set of zeros on D_{R_n} forms a δ -cluster in \mathbb{C} for all $n > N$ in both $d_{\mathbb{C}}$ and d_q . However, it is also true that for any sufficiently small $\delta > 0$ no proper subset of the zeros in D_{R_n} forms a δ -cluster, for any sufficiently small δ and large n . Fix any pair of vertical and horizontal tangent lines V and H to have intersection equal to ∞ under choice of uniformizations $X_n \rightarrow \hat{\mathbb{C}}$. Then, there is a sequence of Möbius transformations $T_n(z) = a_n z + b_n$ such that the collection of zeros converges in the configuration space $\text{Conf}_n(\mathbb{C})$, along a subsequence, since ratios of \mathbb{C} -distances between zeros in D_{R_n} remain bounded above and below. Moreover, the locations of all other singularities go to ∞ .

Step 3 (proving uniform convergence on compact sets) It follows that on \mathbb{C} , each q_n takes the form $p_n(z)h_n(z) dz^2$, where we may assume (passing to a subsequence if necessary) that $p_n(z)$ has convergent coefficients, and h_n is a rational function whose zeros and poles tend to ∞ uniformly in n . Therefore, there are numbers C_n such that $h_n(z) = C_n(1 + o(1))$ on any ball as $n \rightarrow \infty$. Moreover, since the periods of saddle connections converge, h_n must converge uniformly to a *fixed constant* on each compact subset of \mathbb{C} . Moreover, judicious choices of a_n and b_n allow us to assume that p_n is monic and its roots sum to 0. \square

Corollary 4.34 *Let B be the closed ball of d_q -radius 1 in a quadratic differential (X, q) about a point $p \in X$. If B has compact closure in X and is homeomorphic to a closed disk, and q has no poles except possibly at p , then the interior of B is isometric to a ball in a cluster differential.*

Proof The case of no poles is obvious, and the case of one pole follows from taking a double cover branched only over the pole. \square

Definition 4.35 Let μ be a set whose elements correspond to singularities of a quadratic differential in $\mathcal{T}_{g,n}$ or a space of cluster differentials. To each element of μ we associate an integer no less than -1 , corresponding to order of vanishing, and a boolean variable which is 1 or 0 for marked or unmarked points, respectively. Let $\mathcal{Q}(\mu)$ be the corresponding stratum in the Teichmüller space of half-translation surfaces.

For each stratum $\mathcal{Q}(\mu)$ of quadratic differential or cluster differential with at least one zero and at least two singularities, there is a partial compactification $\bar{\mathcal{Q}}(\mu)$, each of whose elements consist of the following data, up to an equivalence relation:

- S_ℓ is a collection of subsets of μ of size at least 2, such that each S_ℓ contains at most one singularity corresponding to a marked point. We require that any two subsets in this collection are either nested or disjoint. We also require $S_0 = \mu$ to be one of the subsets. Let m_ℓ be the sum of the orders of vanishing of singularities in S_ℓ (poles count as -1).
- T is a directed tree whose vertices consist of the sets S_ℓ . Let the edges be such that there is a directed path from v_{ℓ_1} to v_{ℓ_2} if and only if $S_{\ell_1} \subset S_{\ell_2}$. For each vertex $S_\ell \neq S_0$, let $\phi(\ell)$ be the unique subset of μ such that there is a directed edge from S_ℓ to $S_{\phi(\ell)}$.
- (X_ℓ, q_ℓ) is a quadratic differential with one singularity for each element p of μ in S_ℓ not in any $S_j \subsetneq S_\ell$, and one singularity for each $S_j \subsetneq S_\ell$. Each p not in any S_ℓ corresponds to the same type of singularity as p in μ , and for each S_j there is a singularity whose order of vanishing is the sum m_j of orders of vanishing of the singularities in S_j , and which is a marked point if and only if S_j contains a marked point.
- We require that (X_0, q_0) have genus g with n marked points. For $\ell \neq 0$ we require that (X_ℓ, q_ℓ) be a cluster differential with a pole of order $m_\ell + 4$ at ∞ , with a marked point if and only if S_ℓ has a marked point. We also require that the $d_{\mathbb{C}}$ -diameter of the set of singularities of S_ℓ is 1.
- To each oriented saddle connection γ_i of the quadratic differential corresponding to vertex $\phi(S_\ell)$ that starts at the singularity corresponding to S_ℓ , we associate an angle $\theta(\gamma_i) \in \mathbb{R}/2\pi\mathbb{Z}$ such that the direction of γ_i is counterclockwise from the direction of any other such saddle connection γ_k by a (cone) angle of $[\theta(\gamma_i) - \theta(\gamma_k)] \cdot \frac{1}{2}(m_\ell + 2)$.

The equivalence relation is generated by the condition that if $q_\ell = f(z) dz^2$ with the collection of angles $\theta(\gamma_k)$, it is equivalent to $Af(e^{i\alpha}z)[d(e^{i\alpha}z)]^2$ and the collection of angles $\theta(\gamma_k) - \text{Re } \alpha$ for $A \in \mathbb{R}$ and $\alpha \in \mathbb{C}$.

This equivalence relation is simply because we can't distinguish isometric cluster differentials under coordinate changes on \mathbb{C} or scaling by real constant.

This is similar to the compactification of strata in [Bainbridge et al. 2019], except that we do not consider quadratic differentials that escape to ∞ in the moduli space of quadratic differentials, and we consider the limiting objects to be different if their components are rotated relative to each other but are less specific with regard to the relative sizes of different components.

We will not need to construct the topology of $\bar{\mathcal{Q}}(\mu)$, though it can be constructed real analytically by repeatedly taking blowups (over \mathbb{R}) of the closure of $\mathcal{Q}(\mu)$ in $\mathcal{QD}(\mathcal{T}_{g,n})$ along the loci where singularities collide and taking finite branched covers of the blown up subspace. Such a construction would go beyond our scope here. However, for our purposes it is enough describe when a sequence in $\mathcal{Q}(\mu)$ converges to an element of $\bar{\mathcal{Q}}(\mu)$.

Definition 4.36 Let $\{(Y_N, q_N)\}_{N=1}^\infty$ be a sequence of quadratic differentials in $\mathcal{Q}(\mu)$. We say the sequence converges to a given element of $\overline{\mathcal{Q}}(\mu)$ if for every $\delta \in (0, 1)$ there exists some $M < \infty$ such that:

- $d_{\text{Euclidean}}((Y_N, q_N), (X_0, q_0)) < \delta$ whenever $N > M$.
- δ -clusters in the set of singularities of (Y_m, q_m) correspond exactly to the sets S_ℓ whenever $M > N$. That is, they have the same number of singularities of each type, same nesting, same graph of inclusions, etc. Call these clusters $S_\ell(N)$. For each ℓ we will refer to the sequence $S_\ell(N)$ as a *vanishing cluster*.
- For each $S_\ell, \ell \neq 0$, there is a disk $\mathbb{D}_\ell(N) \subset Y_N$ containing only the singularities in $S_\ell(N)$. For some $\lambda_N \in \mathbb{C}^\times$, $(\mathbb{D}_\ell(N), d_{\lambda_N q_N})$ is locally isometric to the disk $\{z : |z| < 1/\delta\}$ in the q'_ℓ -metric on \mathbb{C} , for some cluster differential (\mathbb{C}, q'_ℓ) , where $d_{\text{Euclidean}}(X'_\ell, q'_\ell) < \delta$.
- For each oriented saddle connection γ_k , if we pick a sequence of saddle connections γ_k^N converging to γ_k in the universal half-translation surface containing $q_{\phi(\ell)}$, and the choice of disks and isometries above, the $d_{q_i^N}$ ray in \mathbb{C} that has the same initial location and direction as the image of γ_k , when parametrized by arc length, is of the form $\mathbb{R}(t)e^{i\alpha(t)}$ for $t \in [0, \infty)$ and $|\theta(\gamma_k) - \lim_{t \rightarrow \infty} \alpha(t)| < \epsilon \pmod{2\pi/\mathbb{Z}}$.

Proposition 4.37 Every convergent sequence in $\text{QD}(\mathcal{T}_{g,n})$ consisting of elements of $\mathcal{Q}(\mu)$ has a subsequence that converges to an element of $\overline{\mathcal{Q}}(\mu)$.

Proof This is just a simple matter of extracting subsequences, since the space of cluster differentials with $\text{diam}_{\mathbb{C}} = 1$ is compact for each stratum of cluster differentials, as well as the fact that the space of directions θ is compact. □

We remark that there is an equivalent characterization of the convergence to an element of $\overline{\mathcal{Q}}(\mu)$ when the underlying Riemann surface is $\widehat{\mathbb{C}}$ and the differentials are Möbius normalized. In this case, there is a sequence of Möbius transformations $T_i(N)$, all fixing ∞ , such that the images of singularities in $S_\ell(N)$ under $T_i(N)$ converge to those of a suitable cluster differential. By Propositions 4.13 and 4.16, and Lemma 4.20, it follows that this sequence of cluster differentials can be rescaled to converge in the Euclidean metric so that the limit is isometric to (\mathbb{C}, d_{q_i}) . It should also be noted that if q_i has no poles, then for some numbers $\lambda_N \in \mathbb{C}$, $\{\lambda_N (T_N^*)^{-1}(q_N) dz^{-2}\}_{N=1}^\infty$ is a uniformly convergent sequence of holomorphic functions on each compact subset of \mathbb{C} . For each vanishing cluster, let $z_{j,N}$ be a sequence of elements corresponding to a singularity in $S_\ell(N)$. After possibly permuting singularities of the same type, we conclude that there are sequences of numbers $a_N \in \mathbb{R}$ and $b_N \in \mathbb{C}$ such that $T_N(z) = a_N(z) + b_N$, with $a_N \rightarrow \infty$, such that the singularities $T_N(z_{j,N})$ all converge to the singularities of the differential q_i . In summary, we have the following:

Proposition 4.38 Let $\{q_N\}_{N=1}^\infty$ be a Möbius normalized sequence of quadratic differentials on $\widehat{\mathbb{C}}$ in a stratum of cluster differentials or in $\text{QD}(\mathcal{T}_{0,n})$, converging to an element of $\overline{\mathcal{Q}}(\mu)$ with a pole at ∞ . If S_ℓ is a vanishing cluster, then after permuting singularities $z_{j,N}$ of the same type in $S_\ell(N)$, there exist Möbius transformations $T_N(z) = a_N z + b_N$ for $a_N \in \mathbb{R}^+$ and $b_N \in \mathbb{C}$ with $a_N \rightarrow \infty$, such that the singularities $T_N(z_{j,N})$ converge to the singularities of the differential q_i . There are scalars $\lambda_N \in \mathbb{C}$ such

that $\lambda_N((T_N)^*)^{-1}$ converges uniformly on compact subsets of \mathbb{C} if S_ℓ does not contain a pole, and on compact subsets of $\mathbb{C} \setminus \{0\}$ if S_ℓ contains a pole. \square

Because of the convergence of the rescaled quadratic differentials, we may also conclude the following version of Corollaries 4.27 and 4.28:

Corollary 4.39 *Let λ_N , T_n , q_N , and $S_\ell(N)$ be as in Proposition 4.38. If $\gamma_N : [0, 1] \rightarrow \mathbb{C}$ is a sequence of constant-speed saddle connections internal to $S_\ell(N)$, then $T_N(\gamma_N)$ has a Hausdorff convergent subsequence. If the limit is a saddle connection in q_ℓ then the convergence is C^1 on compact subsets of $(0, 1)$. \square*

We conclude this section by collecting several simple facts about cluster differentials, whose proofs we defer to Appendix B. These results are needed for the proof of Lemma 5.11.

Definition 4.40 Let $\delta \in (0, 1)$. A possibly collapsed δ -cluster is a δ -cluster or a single point.

Proposition 4.41 *Let K be a compact subset of $\text{QD}^1(\mathcal{T}_{g,n})$ and let $(X, q) \in K$ have singularity set Σ . Let $\gamma : [0, 1] \rightarrow \text{QD}(\mathcal{T}_{g,n})$ be a rectifiable path with respect to the Euclidean metric on $\text{QD}(\mathcal{T}_{g,n})$ starting at $\gamma(0) = (X, q) \in K$. We write $\gamma(t) = (X(t), q(t))$, and we let $\Sigma(t)$ be the singularity set of $\gamma(t)$.*

Let $\Sigma' \subset \Sigma$, let $\Sigma'(t)$ be a finite subset of $X(t)$ that varies continuously in the Hausdorff topology along $\gamma(t)$, and assume that $\Sigma'(0)$ is a possibly collapsed δ -cluster for some $\delta > 1$.

Then for any $\delta_1 \in (0, 1)$ there are positive real numbers $C, \delta_2 > 0$, depending only on K and δ_1 , such that the following holds: if $\Sigma'(0)$ is a δ_2 -cluster of Euclidean length at most $C \text{diam}_q(\Sigma)/\delta_2$ in $\text{QD}(\mathcal{T}_{g,n})$, there is a unique way to choose $\Sigma'(t)$ for all t such that the number of marked points in $\Sigma(t)$ and the number of zeros minus poles (counted with multiplicity) remain constant along γ , even if the cardinality of $\Sigma'(t)$ does not remain constant.

Proposition 4.42 *Let $K \subset \text{QD}(\mathcal{T}_{g,n})$ be compact. Let U be a contractible open subset of $\text{QD}(\mathcal{T}_{g,n})$ with closure in K . Suppose that there is a continuous choice of possibly collapsed δ -cluster in U , that is to say, a continuous function h from U to the space of compact sets in the universal half-translation surface of type g, n (endowed with the topology of the Hausdorff metric) that maps each element $(X, q) \in U$ to a finite collection of singularities of (X, q) , such that $h(X, q)$ is a possibly collapsed δ -cluster in the singularity set of (X, q) , and moreover the number of marked points and sum of the orders of vanishing of quadratic differentials at points in the image of $h(X, q)$ are constant. (By Proposition 4.41, if δ is small enough and the diameter of U is small enough, then any such function h is uniquely determined by its value at any point in U .)*

Let M be the moduli space of cluster differentials with the same number of marked points and total order of vanishing as $h(X, q)$. Then there is a number δ_K such that if $\delta < \delta_K$ and U and M are given Euclidean metrics in the sense of Definition 3.2, there is a Lipschitz map $F : U \rightarrow M$ with the following property: a disk containing all of the singularities of $F(X, q)$ maps isometrically into (X, q) so that singularities of $F(X, q)$ map to $h(X, q)$ and the vertical foliation is preserved. The Lipschitz constant depends only on K .

Proof This is immediate from local finiteness of period coordinate systems coming from bounded length saddle connections, ie from Proposition A.1. \square

Lemma 4.43 *Let $\{q_N\}$ be a sequence in $\text{QD}(\mathcal{T}_{0,n})$ that converges to an element of $\bar{\mathcal{Q}}$, and fix a vanishing cluster $\{S_\ell(N)\}$ for the sequence q_N . Assume q_N is given the normalization such that one pole, which is not part of a vanishing cluster, is at ∞ , the center of $S_\ell(N)$ is always 0, and $q(z) dz^{-2}$ is a quotient of monic polynomials. Assume the sum of the orders of vanishing of q_N on $S_\ell(N)$ is m .*

Let $z_{j,N}$ range over the noninfinite singularities of q_N , with q_N vanishing at $z_{j,N}$ to order $e_{j,N}$, and let

$$t_N = \prod_{z_{j,N} \notin S_\ell(N)} (-z_{j,N})^{e_{j,N}} \quad \text{and} \quad \alpha_N = \prod_{z_{j,N} \in S_\ell(N)} (z - t_N^{1/(m+2)} z_{j,N})^{e_j} dz^2.$$

Let the map F_N be the locally defined map F for the cluster $S_\ell(N)$ to the space of cluster differentials associated to $S_\ell(N)$, as defined in Proposition 4.42. Then, for an appropriate choice of the $(m+2)^{\text{nd}}$ root of t_N , $d_{\text{Euclidean}}(\alpha_N, F_N(q_N)) = o(\text{diam}_{q_N}(S_\ell(N)))$.

The following proposition says that if we perturb a quadratic differential but preserve the isometry type of all δ -clusters for some sufficiently small δ , then for a reasonable period coordinate chart the change in the periods is comparable to the Euclidean distance.

Proposition 4.44 *Let K be a compact subset of $\text{QD}(\mathcal{T}_{g,n})$ or a space of cluster differentials, let $L > 0$, and assume that on K the systems of saddle connections associated to the Euclidean metric on K consist of saddle connections of length less than L . Then there is some $\delta_0 > 0$ such that for all $\delta < \delta_0$, the following holds: As X varies in K with singularity set Σ and orienting double cover $(\tilde{X}, \tilde{\Sigma})$, let $H_\delta^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$ denote the subspace of $H^1(\tilde{X}, \Sigma, \mathbb{C})$ that vanishes on all saddle connections internal to δ -clusters of singularities.*

If $\{X_m\}_{m=1}^\infty$ is a convergent sequence in K with singularities $\{\Sigma_m\}$, Y_m is another such sequence lying on a common convex T -convex period coordinate chart $U \subset \mathbb{C}^n$ with X_m for some T , U is generated by periods of saddle connections of length less than L , X_m is connected to Y_m by a line segment of length $\epsilon_m \rightarrow 0$ in U , and X_m and Y_m are such that, in local period coordinates, $X_m - Y_m \in H_\delta^1(\tilde{X}_m, \tilde{\Sigma}_m; \mathbb{C}) \cap H_{\text{odd}}^1(\tilde{X}_m, \Sigma_m, \mathbb{C})$, then $d_{\text{Euclidean}}(X_m, Y_m) \asymp \epsilon_m$.

5 Perturbing quadratic differentials on the sphere

We now establish a setup to discuss saddle connections in convergent sequences of quadratic differentials on $\hat{\mathbb{C}}$. We will consider all strata of quadratic differentials with a pole fixed at ∞ , and such that all poles in \mathbb{C} are simple. We simply want to know that our quadratic differentials have enough saddle connections for us to apply Corollary 4.9. The construction in Definition 4.35 also applies to such differentials.

5.1 Derivatives of period coordinates

Notation 5.1 Fix integers $e_1, \dots, e_r \geq -1$. Assume $\sum_{m=0}^r e_m \geq -3$. For any r -tuple of distinct complex numbers (z_1, \dots, z_r) , let $q(z_1, \dots, z_r) = \prod_{j=1}^r (z - z_j)^{e_j} dz^2$.

Given a finite system of saddle connections γ_i for $q(w_1, \dots, w_r)$ with γ_i having endpoints w_{i_1} and w_{i_2} , there is a neighborhood U of (w_1, \dots, w_r) such that for all $(z_1, \dots, z_r) \in U$ the saddle connection in $(\mathbb{C}, q(z_1, \dots, z_r))$ with endpoints z_{i_1} and z_{i_2} can be chosen Hausdorff continuously for all $\ell \in I$, such that γ_i is the collection of saddle connections at (w_1, \dots, w_r) . The periods vary holomorphically with respect to (z_1, \dots, z_r) ; call these periods P_i , and their partial derivatives $\partial P_i / \partial z_j$.

Differentiation under the integral gives us the following basic formula:

Proposition 5.2 If γ_i is a saddle connection of the quadratic differential $q(z) = f(z) dz^2$ with period P_i , and either $e_j \geq 1$ or v_j is not an endpoint of γ_i , then

$$\frac{\partial P_i}{\partial z_j} = \int_{\gamma_i} \frac{e_j \sqrt{q(z)}}{2(z_j - z)}.$$

In all other cases we can recover the partial derivative $\partial P_i / \partial z_j$ by performing a 1-parameter family of Möbius transformations so that the endpoints of γ_i remain fixed, pushing forward the family of differentials, and differentiating under the integral.

Proof The only cases in which the formula for $\partial P_i / \partial z_j$ is not the immediate result of differentiation under the integral are those in which one endpoint of γ_i is z_j and $e_j \geq 1$. Suppose a and b are the endpoints of γ_i and $b = z_j$. Assume $q_h(z) = q(z)((z - b - h)/(z - b))^{e_j}$. Fix a contour of integration from a to b to be the saddle connection γ_i for the q -metric. Then

$$\frac{1}{h} \left[\int_a^{b+h} \sqrt{q_h(z)} - \int_a^b \sqrt{q(z)} \right] = \int_a^b \frac{1}{h} [\sqrt{q_h(z)} - \sqrt{q}] + \int_b^{b+h} \frac{1}{h} q_h(z).$$

We claim that the last integral on the right-hand side is $O(h^{e_j/2})$ because we can take our contour of integration from b to $b + h$ to have length h , and the integrand is $O(h^{e_j})$ along the entire contour. The remaining term limits to the desired formula.

In other cases, the integrals need not converge. To prove our remaining claim, we simply note that any holomorphic 1-parameter family of quadratic differentials $\prod (z - z_j(t))^{e_i} dz^2$ admits a holomorphic 1-parameter family of Möbius transformations sending $z_{j_1}(t)$, $z_{j_2}(t)$ and ∞ to $z_{j_1}(0)$ and $z_{j_2}(0)$. The pushforwards of $q(t)$ will be of the form $(z - z_{j_1})$. \square

5.2 The limit of the matrix $\partial P_i / \partial z_j$

Notation 5.3 Fix a sequence of quadratic differentials $q_N = q(z_{1,N}, \dots, z_{r,N})$ that converge to an element of the partial compactification of Definition 4.35. Assume that the length-minimizing spanning trees of saddle connections are isomorphic to some fixed Γ as ribbon graphs, where each $z_{j,N}$ corresponds

to a fixed vertex v_j of Γ for all N . Also assume that for each vanishing cluster S_ℓ with $z_{j,N} \in S_\ell(N)$, there is some sequence T_n as in Proposition 4.38 such that $T_N(z_{j,N})$ converges for all j , where $z_{j,N} \in S_\ell(N)$, to the corresponding singularity of the cluster differential q_ℓ corresponding to S_ℓ .

For this sequence, let γ_i vary over the edges of Γ and have period P_i . Let $\Gamma_{i,N}$ be the saddle connection corresponding to γ in (\mathbb{C}, q_N) . Assume we have picked $\sqrt{q_N}$ so that, for some sequence of positive real numbers t_N , the sequence $\int_{\gamma_{i,N}} t_N \sqrt{q_N}$ converges to a nonzero value, and let P_i be the holomorphic function associated to γ_i .

Finally, let M_N be the matrix $(M_N)_{ij} = \partial P_i / \partial z_j |_{(z_1, \dots, z_r) = (z_{1,N}, \dots, z_{r,N})}$. Finally, let M'_N be obtained from M_N by the following operation: for each maximal vanishing cluster S_ℓ , delete the rows corresponding to the saddle connections internal to vanishing clusters, and replace the columns corresponding to singularities in vanishing clusters by a single column that is the sum of the columns corresponding to singularities in $S_\ell(N)$.

Proposition 5.4 *Let $\{q_N\}_{N=1}^\infty = \{q(z_{1,N}, \dots, z_{r,N})\}_{N=1}^\infty$ be a sequence satisfying the hypotheses of Notation 5.3, and suppose that the root of the associated tree is q_0 . Let m be such that q has a pole of order $m + 4$ at ∞ . Assume no vanishing cluster contains a pole, and there is at most one singularity of cone angle 2π in each vanishing cluster. Then we have the following:*

- (1) For each i , $P_i(z_{1,N}, \dots, z_{r,N})$ converges as $N \rightarrow \infty$.
- (2) M_N converges to a matrix M_∞ .
- (3) If $(z_j)_\infty = (z_k)_\infty$, then $e_j(M_\infty)_{ik} = e_k(M_\infty)_{ij}$.
- (4) M'_N converges to a matrix M'_∞ .
- (5) The kernel of M'_∞ has dimension 1 and it is spanned by the all-1 vector.
- (6) M'_∞ is the matrix of partial derivatives for the periods of the saddle connections and vertices of the degeneration of T that occurs as $q_n \rightarrow q_\infty$.
- (7) If γ_i corresponds to a sequence of saddle connections internal to a vanishing cluster, then $\partial P_i / \partial z_j \rightarrow 0$ for all j .

More concisely, if \cdot' denotes the object that \cdot degenerates to as $N \rightarrow \infty$, then

$$\frac{\partial P_i}{\partial z_j} \rightarrow \frac{e_j}{e'_j} \frac{\partial P'_i}{\partial z'_j}.$$

If $e_j = e'_j = 0$ then $e_j/e'_j = 1$. (In particular, the period of a saddle connection that degenerates to a point has partial derivatives tending to 0.)

Proof The first claim follows easily from Corollaries 4.27, 4.28, and 4.18.

Items (2)–(7) determine the value of M , up to the value of the derivative of the saddle connection of a nonshrunk vanishing cluster.

Suppose there are no vanishing clusters. Then the period coordinates P_i are a holomorphic coordinate system for our stratum of cluster differentials in a neighborhood of q_0 . Indeed, the periods of the $r - 1$ edges of Γ form a local coordinate system for the stratum of the differentials $q(z_{1,N}, \dots, z_{r,N})$. It is also true that the values of $r - 1$ elements of $\{z_1, \dots, z_r\}$ determine a local coordinate system if the remaining z_j are fixed, since $q(z_1, \dots, z_r)$ is biholomorphically equivalent to $q(z_1 + \alpha, \dots, z_r + \alpha)$ for any $\alpha \in \mathbb{C}$ (each is the pullback of the other via translation). Since any injective holomorphic map from a domain in \mathbb{C}^n to a domain in \mathbb{C}^n is a biholomorphism onto its image, it follows that the kernel of M consists only of perturbations of (z_0, \dots, z_r) that preserve the isomorphism type of the quadratic differential q_0 . This subspace is precisely the span of the all-1 vector. In this case, the dimension of the stratum of the limiting dimension equals the claimed rank of M . It follows that all claims hold in this case.

Now, recall that the sequences of saddle connections $\gamma_{i,N} : [0, 1] \rightarrow \mathbb{C}$ each have the property that for some sequence of Möbius transformations $T_{i,N}(z) = t_{i,N}(z) + u_{i,N}$ fixing ∞ as in Proposition 4.38, $T_{i,N} \circ \gamma_{i,N}$ converges uniformly and c_1 on compact subsets of $(0, 1)$ to a saddle connection in some cluster differential q_ℓ .

Assume that $T_{i,N}(z) = t_{i,N}z + u_{i,N}$ is chosen so that $T_{i,N}(z)\gamma_i(q_N)$ converges to the saddle connection corresponding to γ_i on one of the quadratic differentials q_ℓ associated to the limit of the sequence $\{q_N\}$.

Now, we would like to apply Proposition 5.2 with the change of coordinates $T_{i,N}(z) = \zeta$, so let

$$\alpha_{i,N}(\zeta) = t_{i,N}^{(m+2)} \prod_{j=1}^r (\zeta - \zeta_j)^{e_j} d\zeta^2 = (T_{i,N})_*(q_N).$$

If $\gamma_{i,N}$ has endpoints $a_{i,N}$ and $b_{i,N}$ for the differential q_N , we then have

$$\frac{\partial Q_i}{\partial z_j} \Big|_{(z_1, \dots, z_r) = (z_{1,N}, \dots, z_{r,N})} = \int_{a_{i,N}}^{b_{i,N}} \frac{1}{2} e_j \frac{\sqrt{q_N(z)}}{z_{j,N} - z} = \int_{T_{i,N}(a_{i,N})}^{T_{i,N}(b_{i,N})} t_{i,N}^{-m/2} \cdot \frac{1}{2} e_j \frac{\sqrt{\alpha_{i,N}(\zeta)}}{\zeta_{j,N} - \zeta}.$$

We would like to apply the Lebesgue dominated convergence theorem to the right-hand side, parametrizing the integrals to be constant-speed $d_{(T_{i,N})_*(q_N)}$ -geodesic maps from $[0, 1]$ to \mathbb{C} . To do this, we need the following:

Claim 5.5 *A dominating function for the sequence $\int_{T_{i,N}(a_{i,N})}^{T_{i,N}(b_{i,N})} t_{i,N}^{-m/2} \cdot \frac{1}{2} e_j \sqrt{\alpha_{i,N}(\zeta)} / (\zeta_{j,N} - \zeta)$, where the contours are parametrized as constant-speed geodesics from $[0, 1]$ to \mathbb{C} with respect to $d_{T_*q_N}$, exists whenever $z_{j,N}$ is a zero of q_N or not an endpoint of $\gamma_{i,N}$.*

Proof Suppose that $\{\gamma_{i,N}\}$ is not internal to a vanishing cluster.

Since we are assuming that our contours of integration belong to length-minimizing spanning trees, for each z on our contour of integration, we have

$$d_{q_N}(z, (z_{j,N})) \geq \min(d_{q_N}(z, (a_{i,N})), d_{q_N}(z, (b_{i,N}))).$$

If this were not the case, we would have

$$d_{q_N}(a_{i,N}, b_{i,N}) > \min(d_{q_N}(a_{i,N}, z_{j,N}), d_{q_N}(b_{i,N}, z_{j,N})),$$

and a saddle connection from $a_{i,N}$ to $b_{i,N}$ could not belong to a length-minimizing spanning tree.

All singularities of q_N remain bounded, the poles are all bounded away from $z_{j,N}$, and $z_{j,N}$ is a zero of q_N , so

$$d_{q_N}(z, z_{j,N}) \prec \int_0^{|z-z_{j,N}|} t^{1/2} dt = \frac{2}{3}|z - z_{j,N}|^{3/2}.$$

We thus have, along our contour of integration,

$$|z - z_{j,N}|^{-1} \prec d_{q_N}(z, z_{j,N})^{-2/3} \leq \min(d_{q_N}(z, T_{i,N}(a_{i,N})), d_{q_N}(z, T_{i,N}(b_{i,N})))^{-2/3}.$$

We have parametrized our arcs to be d_{q_N} geodesics instead of d_{α_n} geodesics, but applying Lemma 4.20 to the proofs of Corollaries 4.27 and 4.28 tells us that the contours of integration converge (Hausdorff, in total length, and C^1 on compact subsets of $(0, 1)$) to geodesics with respect to d_{α_∞} , where $\alpha_\infty = \lim_{n \rightarrow \infty} \alpha_n$. Therefore, for some $C > 0$, our sequence of integrals is dominated by $\int_0^1 C[x(1-x)]^{-2/3} dx$. This completes the proof of the claim in this case. \square

5.3 The derivative of period coordinates, rescaled

We may do the same for the case of a saddle connection internal to a vanishing cluster. By translating our original sequence q_N so that one endpoint of $\gamma_{i,N}$ is 0 and not equal to z_j , we can assume that $\zeta = T_{i,N}(z) = t_{i,N}z$ with $t_{i,N} \rightarrow \infty$. If $\{S_\ell(N)\}_{N=1}^\infty$ is the minimal vanishing cluster to which $\gamma_{i,N}$ is internal, write $q_N = f_{i,N}(z)g_{i,N}(z) dz^2$ with $g(z) = \prod_{j \in S_\ell(N)} (z - z_{j,N}^{e_j})$. Then the sequence of integrals becomes

$$\begin{aligned} \frac{\partial P_i}{\partial z_j} \Big|_{(z_1, \dots, z_r) = (z_{1,N}, \dots, z_{r,N})} &= \int_{a_{i,N}}^{b_{i,N}} \sqrt{f(z)g(z)} \frac{e_j dz}{2(z - z_{j,N})} \\ &= \int_{\alpha_{i,N}}^{\beta_{i,N}} [f_N(\zeta/t_{i,N})]^{1/2} [g_N(\zeta/t_{i,N})]^{1/2} \frac{e_j d\zeta}{2(\zeta - \zeta_{j,N})}. \end{aligned}$$

Now, the contour of integration converges, and $f_N(\zeta/t_{i,N}) = f_N(z)$ converges to a constant, possibly zero. If m_ℓ is the degree of $g(z)$ then $[t_{i,N}^{-m_\ell} h_N(\zeta)]$ converges as a function of ζ . By Lemma 4.20 and the argument for the saddle connections not internal to vanishing clusters, for $x \in [0, 1]$ we have

$$\left| \frac{h_N(\zeta)}{\zeta - \zeta_{j,N}} \circ T_{i,N} \gamma_{i,N}(x) \right| \leq C[x(1-x)]^{-2/3}.$$

We therefore see that not only does a dominating function exist, but we can take the dominating function to be $t_{i,N}^{-m/2} C[x(1-x)]^{-2/3}$ for $m > 0$, if we are willing to start the sequence at some large N . This proves item (7). The rest of the claims follow from Proposition 5.2 and the dominated convergence theorem. \square

Corollary 5.6 Let $q_N = \prod_{j=1}^s (z - z_{j,N})^{e_j} dz^2$ be a sequence of cluster differentials with no poles satisfying the hypotheses of Proposition 5.4, and assume that at least one of the sequences $\{z_{j,N}\}$ does not converge to 0. Fix integers e_{s+1}, \dots, e_r and let $(z_{s+1,N}, \dots, z_{r,N})$ be a sequence of $(s-r)$ -tuples of complex numbers converging to $(\infty, \infty, \dots, \infty)$ in $\widehat{\mathbb{C}}^{r-s}$. Assume $\{\gamma_{i,N}\}$ is a sequence of saddle connections of the differential q_N that converges (Hausdorff), and its period P_i is locally a holomorphic function of (z_1, \dots, z_s) . Let $Q_{i,N}$ be the periods of a sequence saddle connections in the metric associated to the quadratic differential $\prod_{j=1}^r (z - z_{j,N})^{e_j}$ which converge (Hausdorff) to same limit as $\gamma_{i,N}$, which are locally given by the holomorphic function $Q_i(z_1, \dots, z_r)$.

If $\gamma_{i,N}$ does not converge to a point, for $1 \leq j \leq s$ the following limits exist and are equal:

$$\lim_{N \rightarrow \infty} \frac{\partial Q_{i,N}}{\partial z_j} = \frac{\partial \log Q_{i,N}}{\partial z_j} = \lim_{N \rightarrow \infty} \frac{\partial \log P_i}{\partial z_j}.$$

If $\gamma_{i,N}$ converges to a point, then let $\gamma_{i',N}$ be another convergent sequence of saddle connections that does not converge to a point, and let them have periods $P_{i',N}$ and $Q_{i',N}$. Then

$$0 = \lim_{N \rightarrow \infty} \frac{\partial Q_{i,N}}{\partial z_j} = \frac{\partial Q_{i,N}}{\partial z_j} \frac{1}{Q_{i',N}} = \lim_{N \rightarrow \infty} \frac{\partial P_i}{\partial z_j} \frac{1}{P_{i',N}}.$$

For $s + 1 \leq j \leq r$,

$$\lim_{N \rightarrow \infty} \frac{\partial Q_{i,N}}{\partial z_j} = \frac{\partial \log Q_{i,N}}{\partial z_j} = 0.$$

Proof We use an identical dominated convergence argument. We can treat the faraway singularities (those outside $S_\ell N$) as scalars in the limit, and by Lemma 4.20 we can use the same system of saddle connections γ_i we would if there were no finite singularities outside of $S_\ell(N)$. □

Corollary 5.7 Given a sequence $\{q_N\}$ converging as in Proposition 5.4, let $M_{\ell,N}$ be the submatrix of M_N corresponding to saddle connections internal to the vanishing cluster $\{S_\ell(N)\}$. Then we have the following as $N \rightarrow \infty$:

- $\partial P_i / \partial z_j |_{(z_1, \dots, z_r) = (z_{1,N}, \dots, z_{r,N})} = o(|M_{\ell,N}|)$ if $j \notin S_\ell$.
- If $\gamma_{i,N}$ corresponds to a row of $M_{\ell,N}$ then $\sum_{j \in S_\ell} \partial P_i / \partial z_j |_{(z_1, \dots, z_r) = (z_{1,N}, \dots, z_{r,N})} = o(|M_{\ell,N}|)$.
- If $S_l \subsetneq S_\ell$ is a vanishing cluster and $\sum_{j \in S_l} a_j e_j = 0$, then

$$\sum_{j \in S_l} a_j \frac{\partial P_i}{\partial z_j} \Big|_{(z_1, \dots, z_r) = (z_{1,N}, \dots, z_{r,N})} = o(|M_{\ell,N}|).$$

Proof This is immediate from Corollary 5.6 and Proposition 4.14. □

Corollary 5.8 Let $\{q_N\}$ and $M_{\ell,N}$ be as above, and assume the center of $S_{\ell,N}$ is 0 for all N . Letting $t_N = \prod_{z_{j,N} \notin S_\ell(N)} z_j (-z_{j,N})^{e_{j,N}}$, we have that $|M_{\ell,N}| \asymp |t_N|^{1/2} \text{diam}_{\mathbb{C}}(S_\ell(N))^{m/2}$, where m is the number of zeros in $S_\ell(N)$, counted with multiplicity. In fact, the sequence of matrices $t_N^{-1/2} \text{diam}_{\mathbb{C}}(S_\ell(N))^{m/2} M_{\ell,N}$ converges to a matrix whose cokernel corresponds to the space spanned by periods of saddle connections internal to proper vanishing subclusters of $\{S_\ell(N)\}$.

Proof We are free to consider the same system of saddle connections we would use in the absence of faraway singularities by Lemma 4.20.

Assume z_j is a zero of order e_j that degenerates to $z_{j'}$, a zero of order $e_{j'}$ times the corresponding column of the matrix of partial derivatives for the limiting differential q_ℓ . If we change coordinates on \mathbb{C} by the positive real scalar $\text{diam}_{\mathbb{C}}(S_\ell(N))$ so that $S_\ell(N)$ has diameter 1, then Propositions 4.16, 4.14 and 5.4 imply that column j of the matrix $|t_N|^{-1/2} \text{diam}_{\mathbb{C}}(S_\ell(N))^{-m/2} M_{\ell,N}$ converges to $e_j/e_{j'}$, together with additional zeros for the saddle connections internal to proper subclusters of q_ℓ . Since the matrix of partial derivatives for q_ℓ has a nonzero limit, so does $|t_N|^{-1/2} \text{diam}_{\mathbb{C}}(S_\ell(N))^{-m/2} M_{\ell,N}$, and the result follows. \square

We would have liked to extend Proposition 5.4 and its corollaries to clusters that include a single pole. Unfortunately, the partial derivatives of a period need not remain bounded when a zero collides with a pole. Consider the family of quadratic differentials $(z(z-1) dz^2)/(z+t)$ where t varies. If t is a negative real number, then the segment $[0, 1]$ of the real line is a vertical saddle connection; call its period $P(t)$. Then, by Proposition 5.2, we have

$$P'(t) = \frac{1}{2} \int_0^1 \sqrt{\frac{x(x-1)}{(x+t)^3}} dx.$$

The integrand is purely imaginary and does not change sign. Thus, by the monotone convergence theorem,

$$\lim_{t \rightarrow 0^-} |P'(t)| = \int_0^1 \frac{(1-x)^{1/2}}{x} dx \geq \int_0^{1/2} \frac{dx}{2x} = +\infty.$$

However, for the estimates we need, it is possible to take a double cover branched over the pole and pull back the quadratic differential; this converts the pole into a singularity with cone angle 2π at the cost of duplicating other singularities.

5.4 The main estimate

Notation 5.9 Fix a stratum of quadratic differentials in $\text{QD}(\mathcal{M}_{0,n})$ with no zeros of order 2 or higher. Assume all quadratic differentials are normalized to have simple poles at ∞ and two other fixed points A and B . Then each has the form $\lambda_q[(z-A)(z-B)]^{-1} \prod_{j=1}^s (z-z_j)^{e_j} dz^2$. Fix integers $\{e_j\}_{j=1}^s$ with $-1 \leq e_j \leq 1$ for all j and $e_1 + \dots + e_s = -1$, and denote such a differential by $q(\lambda, z_1, \dots, z_r)$. Let

$$d_{\text{Sym}}(q_1, q_2) = \inf_{\phi} \sum_w |\phi(w) - w| + |\lambda_{q_1} - \lambda_{q_2}|,$$

where w ranges over all noninfinite singularities of q_1 , and ϕ ranges over all bijections between the singularities of q_1 and q_2 that preserve the type of singularity.

Proposition 5.10 Suppose that $\{q_N\}$ and $\{r_N\}$ are sequences converging to possibly distinct elements of the partial compactification $\bar{\mathcal{Q}}$, but such that $d_{\text{Sym}}(q_N, r_N) \rightarrow 0$ with respect to a common Möbius normalization. Let ϕ_N be the type-preserving bijection between the singularities of q_N and the singularities of

r_N that minimizes $\sum |\phi(z_{j,N}) - z_{j,N}|$ for this Möbius normalization. Let $\{S_\ell(N)\}$ be a vanishing cluster for the sequence $\{q_N\}$. Suppose that as $N \rightarrow \infty$, $d_{\text{Sym}}(q_N, r_N)/\text{diam}_{\mathbb{C}}(S_m(N)) \rightarrow 0$ for every vanishing cluster $\{S_m(N)\}$ in q_N with $S_m(N) \supseteq S_\ell(N)$ for all N . Then $\phi_N(S_\ell(N))$ is a vanishing cluster for r_N .

Proof It suffices to consider the minimal S_m that properly contains S_ℓ . It is clear from the definition of vanishing cluster that for any sequence of sets of singularities $S(N)$ with $S_m(N) \supseteq S(N) \supseteq S_\ell(N)$, we must have $\text{diam}_{\mathbb{C}}(T(N)) \asymp \text{diam}_{\mathbb{C}}(S_m(N))$, since otherwise we could pass to a subsequence along which $T(N)$ belonged to a vanishing cluster with $S_m(N) \supseteq S(N) \supseteq S_\ell(N)$. By minimality of S_m this can't happen, since the nesting of the vanishing clusters is determined by the limit in \bar{Q} . \square

Lemma 5.11 Suppose $\{q_N\}_{N=1}^\infty = \{q((\lambda_N, z_{1,N}, \dots, z_{r,N}))\}_{N=1}^\infty$ is a sequence in $\text{QD}(\mathcal{T}_{0,n})$, normalized in the sense of Notation 5.9 that converges in the sense of Definition 4.35, such that no vanishing cluster contains a pole and every vanishing cluster contains a zero. Suppose that in addition, no vanishing cluster for the sequence $\{q_N\}$ has more than k zeros (counted with multiplicity) for $k \geq 0$. Let $\{r_N\}$ be another sequence with the same Möbius normalization, converging to a point in the partial compactification of Definition 4.35, and assume that for each singularity type m and each $p \in \mathbb{C}$, the number of singularities of type m that limit to p along the sequences $\{q_N\}$ and $\{r_N\}$ is the same. (In particular, $\{q_N\}$ and $\{r_N\}$ have the same limit in $\text{QD}(\mathcal{T}_{0,n})$.) Then

$$d_{\text{Euclidean}}(q_N, r_N) \asymp d_{\text{Sym}}(q_N, r_N)^{2/(2+k)} \quad \text{as } N \rightarrow \infty.$$

Proof The proof is divided into two steps. In the first we fix a subsequence of potential counterexamples with specified combinatorics of the singularities belonging to vanishing clusters, and isolate some piece of the difference between q_N and r_N . In the second we show that this piece causes a nontrivial change in period coordinates that is not canceled by any other components.

Step 1 Each maximal vanishing cluster of $\{q_N\}$ is also a maximal vanishing cluster for $\{r_N\}$, and we will only consider bijections that preserve maximal vanishing clusters since d_{Sym} is only realized by such a bijection for large N . In fact, we can take this further. For each vanishing cluster $\{S_\ell(N)\}$ of $\{q_N\}$ we can pass to a subsequence whereby for each vanishing cluster $S_\ell(N)$,

$$d_{\text{Sym}}(q_N, r_N) \star_\ell \text{diam}_{\mathbb{C}}(S_\ell(N)),$$

where \star_ℓ depends the choice of ℓ but is one of $\prec, \succ,$ or $\dot{\asymp}$.

Assume

$$q_N = \frac{\lambda_N}{(z-A)(z-B)} \prod_{j=1}^r (z - z_{j,N})^{e_j} dz^2 \quad \text{and} \quad r_N = \frac{\mu_N}{(z-A)(z-B)} \prod_{j=1}^r (z - w_{j,N})^{e_j} dz^2,$$

and let ϕ_N be the bijection realizing $d_{\text{Sym}}(q_N, r_N)$ that takes z_j to w_j for all j . We will describe how to break the perturbation realized by ϕ_N into components, which we will call $v_{\ell,N}$ and define below.

Let $\{S_\ell(N)\}$ be either a vanishing cluster of the sequence $\{q_N\}$ or the entire set of singularities (besides ∞). Let the sequences $\{v_N\}$ and $\{v_{\ell,N}\}$ satisfy the following:

- $v_{\ell,N}$ is an $(r+1)$ -tuple of complex numbers, corresponding to the coordinates $(\lambda, z_1, \dots, z_r)$ and (μ, w_1, \dots, w_r) .
- If $\{S_\ell(N)\}$ is a vanishing cluster of $\{q_N\}$, then the support of $v_{\ell,N}$ is in the entries corresponding to singularities in S_ℓ .
- If $S_k \subsetneq S_\ell$ is a vanishing cluster of $\{q_N\}$, then $v_{\ell,N}$ is constant on the entries corresponding to singularities in D' . (All points in each proper subcluster move the same distance and in the same direction.)
- For each N , $\sum_\ell v_{\ell,N} = (\mu - \lambda, w_1 - z_1, \dots, w_s - z_s) =: v_N$.
- If S_ℓ is a vanishing cluster and $v_{\ell,N} = (0, a_1, \dots, a_s)$ then $\sum_{j=1}^s a_j e_j = 0$. (Together with other conditions, this means that $v_{\ell,N}$ does not move the center of $S_{\ell,N}$.)

The purpose of this notation, and the idea of the rest of the proof, is that for some ℓ , the component $v_{\ell,N}$ can be detected either by the periods of saddle connections internal to $S_\ell(N)$, or by periods internal to the largest $S_k(N) \supset S_\ell(N)$ with $\text{diam}_\mathbb{C}(S_k(N)) \dot{\asymp} \|v_{\ell,N}\|$. Then we show that this does not get canceled by other components of the perturbation.

We may pass to a subsequence with the following property: for any S_k and S_ℓ , not necessarily distinct, the following converge (possibly to 0 or ∞):

$$\frac{\text{diam}_\mathbb{C}(S_\ell(N))}{\text{diam}_\mathbb{C}(S_k(N))}, \quad \frac{\|v_{\ell,N}\|}{\text{diam}_\mathbb{C} S_{k,N}}, \quad \frac{v_{\ell,N}}{\|v_N\|}.$$

In other words, the relative sizes of the various clusters and components of v_N converge, and if $v_{\ell,N}$ is not vanishingly small relative to any $v_{k,N}$, the direction of $v_{\ell,N}$ converges. Moreover, the sizes of perturbations of vanishing clusters converge relative to the sizes of all vanishing clusters.

Now pick maximal S_ℓ , subject to the following condition: either $|v_{\ell,N}| \dot{\asymp} |v_N|$, or $\sum_{S_l \subseteq S_\ell} |v_{l,N}| \dot{\asymp} |v_N|$ and $\sum_{S_l \subseteq S_\ell} |v_{l,N}| \dot{\asymp} \text{diam}_\mathbb{C}(S_\ell(N))$.

In particular, this means that either $S_\ell(N)$ is always the entire set of singularities, or $S_\ell(N)$ and $\phi_N(S_\ell(N))$ are both vanishing clusters of the sequence $\{r_N\}$, by Proposition 5.10.

In both cases, let $T_N(z) = a_N z + b_N$ be a sequence of Möbius transformations fixing ∞ with the property that $T_N(S_\ell(N))$ converges (Hausdorff) to the limiting differential associated to S_ℓ ; we may assume $a_N \in \mathbb{R}$. Then $T_N(\phi_N(S_\ell(N)))$ also converges to the singularity set of a cluster differential along a subsequence.

Step 2 (estimating the perturbation in period coordinates) Now we have two cases. In both, we can detect a change in Euclidean distance from the periods internal to $S_\ell(N)$ by Proposition 4.42.

Case 1 ($|v_N|/\text{diam}_\mathbb{C}(q_N) \dot{\asymp} 1$) Here, for all large N we can define the projection map F associated to the clusters $S_\ell(N)$ and $\phi_N(S_\ell(N))$ in Proposition 4.42 on a connected open set containing both, by Proposition 5.10. By Lemma 4.43 and Corollary 4.15 it follows that

$$d_{\text{Euclidean}}(q_N, r_N) \dot{\asymp} d_{\text{Euclidean}}(F_N(q_N), F_N(r_N)) \dot{\asymp} \max\{\text{diam}_{q_N}(S_\ell(N)), \text{diam}_{r_N} \phi_N(S_\ell(N))\}.$$

The last inequality can be justified as follows: First, the ratios between the factors t_N in Lemma 4.43 for q_N and r_N differ by a ratio of $1 + o(1)$. Now, if two cluster differentials q and r have $d_{\mathbb{C}}$ -diameter at most 1, one has diameter 1, and their distance with respect to $d_{\mathbb{C}}$ is bounded below, then $d_{\text{Euclidean}}(q, r)$ is also bounded below by compactness. Our last inequality then comes from Lemma 4.43 and the fact that scalars induce similarities on spaces of cluster differentials.

Now by Corollary 4.15 and Lemma 4.43, if $\max\{\text{diam}_{\mathbb{C}}(q_N), \text{diam}_{\mathbb{C}}(r_N)\} = \epsilon_N$ then

$$\begin{aligned} \max\{\text{diam}_{q_N}(S_{\ell}(N)), \text{diam}_{r_N}(\phi_N(S_{\ell}(N)))\} &\asymp \epsilon_N^{(m+2)/2}, \\ \max\{\text{diam}_{q_N}(S_{\ell}(N)), \text{diam}_{r_N}(\phi_N(S_{\ell}(N)))\} &\asymp \epsilon_N^{(m+2)/2}, \end{aligned}$$

only if q_N contains all m singularities.

Case 2 ($|v_N|/\text{diam}_{\mathbb{C}}(q_N) \rightarrow 0$) Here $|v_{\ell,N}| \asymp |v_N|$. Again we choose F_N to be the Lipschitz projection associated to $S_{\ell}(N)$. Assume the center of $|v_{\ell,N}|$ is always zero and let $t_N = \prod_{z_{j,N} \notin S_{\ell}(M)} (-z_{j,N})^{e_{j,N}}$. Assume m is the number of zeros in $S_{\ell}N$ (each zero has multiplicity 1, so we don't need to specify that we are counting with multiplicity).

If $M_{\ell,N}$ is as in Corollary 5.6, then the sequence of matrices

$$\{t_N^{-1/2} \text{diam}_{\mathbb{C}}(S_{\ell,N})^{-m/2} M_{\ell,N}\}_{N=1}^{\infty}$$

converges (by Corollary 5.7 and Proposition 4.14), and its kernel is the span of those vectors that are allowed to be $v_{k,N}$ with $k \neq \ell$. What we mean by this is that when we defined $v_{k,N}$, we required that it belong to a certain subspace of \mathbb{C}^{r+1} , and all vectors in this subspace are in the kernel of the limiting matrix. The periods of saddle connections internal to proper subclusters of $S_{\ell}(N)$ span the cokernel of the limiting matrix. The limiting matrix

$$\lim_{N \rightarrow \infty} t_N^{-1/2} \text{diam}_{\mathbb{C}}(S_{\ell,N})^{-m/2} M_{\ell,N}$$

has full rank when restricted to the space of vectors that are allowed to be $v_{\ell,N}$.

We may move from q_N to r_N along a piecewise-smooth path β_N , so that the motion of each of the singularities $z_{j,N}$ to $\phi(z_{j,N})$ is a straight line segment with constant speed along each of two segments of β_N . On the first piece of β_N the total displacement of the vectors $(\lambda_N, z_{1,N}, \dots, z_{r,N})$ will be $v_{\ell,N}$, and on the second piece the total displacement will be the remainder.

The singularities starting at $S_{\ell}(N)$ will be in the stratum of $S_{\ell}(N)$ for all but finitely many points along each segment. For any such sequence of paths, if we pick s_N to be one point in this stratum along each β_N and construct the matrix $\widehat{M}_{\ell,N}$ analogously to the matrix $M_{\ell,N}$ at s_N , we can pass to a subsequence that converges in the sense of Notation 5.3. We also note that $\partial P_i / \partial \lambda = o(M_{\ell,N})$, since $\partial \log P_i / \partial z_j$ is not bounded for all j but $\partial P_i / \partial \lambda$ is. The limit of the sequence of matrices must be the same, since it only depended on the limiting differential for S_{ℓ} , which is the same for any subsequence of the sequence $\{s_N\}$. We therefore conclude that the matrices $M_{\ell,N}$ converge *uniformly after rescaling by constants c_N along β_N* in the sense that the matrix $\widehat{M}_{\ell,N}$ differs from $M_{\ell,N}$ at q_N by

$o(|M_\ell, N|)$. In particular, this does not depend on which choices of saddle connections that degenerate to the proper edges of the limiting differential for S_ℓ we picked at each s_ℓ . Moreover, the displacement of periods of S_ℓ is $M_{\ell,N}(v_{\ell,N}) + o(|M_{\ell,N}(v_N)|)$ along the first segment of β_N and the displacement is $o(|M_{\ell,N}(v_N - v_{\ell,N})|)$ along the second segment of β_N .

This, combined with Proposition 4.44 tells us that $d_{\text{Euclidean}}(F_N(q_N), F_N(r_N))$ is comparable to the size of the leading-order approximation $M_{\ell,N}v_{\ell,N}$. By Corollary 5.8 we have

$$|M_{\ell,N}v_N| \asymp |t_N|^{1/2}|s_N|^{m/2}|v_{\ell,N}| \asymp \text{diam}_{\mathbb{C}}(S_{\ell,N})^{(k-m)/2}[\text{diam}_{\mathbb{C}}(S_\ell(N))]^{m/2}|v_N| \cdot |v_N|^{(2+k)/2}.$$

We explain the sources of the inequalities: The size of t_N is at least comparable to the product of the zeros outside of $S_\ell(N)$, since all poles remain bounded away from 0. Since we only care about a multiplicative constant, we may ignore all zeros and poles that do not belong to a vanishing cluster containing $\{S_\ell(N)\}$, and there are at most $k - m$ additional zeros inside any such vanishing cluster, each of which is larger than $\text{diam}_{\mathbb{C}}(S_\ell)$. Hence $|t_N| \asymp |v_N|^{(k-m)/2}$. The remaining inequalities are clear from $\text{diam}_{\mathbb{C}}(S_\ell(N)) \asymp |v_{\ell,N}| \asymp |v_N|$.

Then by Proposition 4.42 we get

$$d_{\text{Euclidean}}(q_N, r_N) \asymp d_{\text{Euclidean}}(F_N(q_N), F_N(r_N)) \asymp |v_N|^{(k+2)/2}. \quad \square$$

Lemma 5.12 Fix a compact subset H in the open upper half-plane \mathbb{H} . Fix a real number $C > 1$ and an integer $m \geq 0$. Consider the space K of quadratic differentials in $\mathcal{T}_{0,2m}$ with the following restrictions:

- They are of the form $(F(z)/G(z)) dz^2$ with $F, G \in \mathbb{R}[z]$, $\deg(F) = 2m - 4$, and $\deg(G) = 2m - 1$.
- G may have simple roots at i and $-i$, but all other roots of G are real. (Either way, G has real coefficients.) The roots of G and ∞ are the marked points of the underlying Riemann surface.
- G has no repeated roots, and the distance between any two consecutive real roots of G is between $1/C$ and C .
- G has a root in $[-C, C]$ if i and $-i$ are roots of G ; otherwise, 0 and 1 are roots of G .
- No root of $F(z)$ is real, and one root from each conjugate pair belongs to H .
- The ratio of the leading coefficients of F and G is between $1/C$ and C .

Assume $q_1 = (F_1(z)/G_1(z)) dz^2$ and $q_2 = (F_2(z)/G_2(z)) dz^2$ are elements of K . Let ϕ range over all conjugation-invariant bijections of the roots of F_1 with the roots of F_2 that fix i if i is a root of G , and let ψ be the bijection of the roots of G_1 with the roots of G_2 that preserves the order of the real roots and fixes i and $-i$ if they are roots of G_1 and G_2 . Let

$$\epsilon = \min_{\phi} \max\{\phi(w) - w : F_1(w) = 0\} \cup \{\psi(w) - w : G_1(w) = 0\}.$$

Then there are constants $M, \epsilon_0 > 0$ depending only on K such that the following hold:

- If $\pm i$ are not roots of G , then $d_{\text{Euclidean}}(q_1, q_2) \asymp \epsilon^{(m+2)/2}$.
- If $\pm i$ are roots of G , then $d_{\text{Euclidean}}(q_1, q_2) \asymp \epsilon^{m+1}$.

Proof First we will explain how both conclusions follow from Lemma 5.11. If $\pm i$ are roots of G_1 and G_2 , then we can take a double cover of each that preserves the real line and is branched at $\pm i$ via $z = h(w) = (w^2 - 1)/(2w)$. Then $\pm i$ are marked points, but not poles on the pullbacks $h^*(q_1)$ and $h^*(q_2)$. Then $\iota^*(q)$ is a conjugation-invariant quadratic differential on $\widehat{\mathbb{C}}$ whose poles are all simple and occur on the projective real line. There are uniform upper and lower bounds on the distance between real poles, and the zeros with positive real part are in a fixed compact part of the upper half-plane.

Now, suppose that z_1 is a root of F_1 , z_2 is a root of F_2 , and $|z_1 - z_2| = \epsilon$. Then if $h(w_1) = z_1$ and $h(w_2) = z_2$, we have

$$|w_1 - w_2| = |h(w_1) - h(w_2)| \cdot \left| \frac{2w_1w_2}{1 + w_1w_2} \right| \dot{>} \epsilon.$$

We have a similar estimate for the inverse images of the roots of G_1 .

We can apply Lemma 5.11 to both. □

6 Building the quasiconformal map

In this section, we build a quasiconformal map with the desired properties to prove Theorem 1.2. Note that for us it will suffice to consider the case when our two quadratic differentials lie in the principal stratum, since the principal stratum is dense in $\mathcal{QD}^1(\mathcal{T}_{g,n})$.

6.1 Cutting into triangles and nearly regular right polygons

Definition 6.1 A *nearly regular right polygon*, or *NRRP*, is a subset P of a half-translation surface (X, q) with the following properties:

- P is homeomorphic to a closed disk.
- P has at least one singularity in its interior, and at most one pole, but has no singularities on its boundary.
- ∂P is piecewise geodesic, and all pieces are vertical or horizontal segments.
- Every interior angle of ∂P is $\frac{\pi}{2}$.
- If P is doubled along its boundary to form a sphere $P \cup \bar{P}$, where \bar{P} is the Schwarz reflection of P , then there is a choice of holomorphic coordinate z on $P \cup \bar{P}$ for which ∂P is the real projective line, and the quadratic differential on $P \cup \partial P$ is of the form $(F(z)/G(z)) dz^2$ where F and G satisfy the conditions of Lemma 5.12. (That is, they have real coefficients, and G is normalized to have two roots at i and $-i$, or to have roots at 0 and 1, etc.)

Let R be the minimum q -distance from ∂P to a singularity in $P \setminus \partial P$. We call R the *radius* of the NRRP.

On a quadratic differential it is always possible to find an NRRP around each singularity with all sides of the same length. If we change the differential by a small enough amount, and in so doing break up such a singularity into multiple singularities, we can still find an NRRP that stays Hausdorff close to our original NRRP. The motivation for this definition is that if one wants to build a quasiconformal map

between two half-translation surfaces, piecewise-affine maps which are affine on some triangulation will tend to have large dilatation even when two Riemann surfaces are nearly conformal, if the triangulation is nearly degenerate. Given an edgewise-linear map between boundaries of similarly shaped NRRPs in two quadratic differentials that are close in $\text{QD}(\mathcal{T}_{g,n})$, we can extend to a quasiconformal map of NRRPs, and explicitly estimate its quasiconformal dilatation.

Proposition 6.2 *On any compact subset K of $\text{QD}(\mathcal{T}_{g,n})$ there are constants $C = C(K) > 0$ and $\delta = \delta(K)$ such that for every surface X in K we can partition the set of singularities of X into δ -clusters and singletons (for some $\delta \in (0, 1)$), and associate to each δ -cluster or singleton an NRRP containing it such that the following hold:*

- *Each NRRP contains only the singularities associated to it.*
- *Each NRRP has radius and all boundary side lengths at least $C(K)$.*
- *If we consider all lifts of the NRRPs in the universal cover of the underlying compact Riemann surface, the distance between any two NRRPs is more than twice as large as the side length of any NRRP.*

Proof It is clearly possible to choose a constant $C(X)$ at each $X \in K$ such that the properties hold for $C(X)$ by picking a small NRRP around each singularity. (NRRPs of a given radius persist under sufficiently small perturbation.) For each X , we can find a neighborhood of X where the conditions all hold with the constant $\frac{1}{2}C(X)$ instead of $C(X)$, and K is finitely covered by such neighborhoods. It follows that we can take $C(K)$ to be half the minimum value of $C(X)$ used in a finite subcover. \square

Let Y be the closure of the complement of the union of a system of NRRPs satisfying the hypotheses of Proposition 6.2 in a half-translation surface X , and assume every singularity of X is contained in some NRRP. If we double Y along ∂Y , the foliations by vertical and horizontal segments extend by reflection, and Y acquires the structure of a quadratic differential, with singularities of cone angle 3π at the vertices of ∂Y .

Moreover, if we take the Delaunay triangulation of the resulting surface, the midpoint of any boundary edge is strictly closer to the endpoints of that edge than to any other singularity, so all edges of ∂Y belong to the Delaunay triangulation of the double cover.

Therefore, *it makes sense simply to speak of the Delaunay triangulation of Y* . That is to say, if Σ is the set of vertices of Y (which are all on the boundary of Y), then we have proved:

Proposition 6.3 *Y has a triangulation with the property that the circumcenter of each triangle is the boundary of a standard Euclidean disk of radius equal to the distance to any vertex of the triangle. The boundary edges of Y are all edges.* \square

Now, we describe how to build a quasiconformal map between two nearby quadratic differentials, provided they are within some distance $c(K)$ that depends only on K .

Let X_1 and X_2 be two such Riemann surfaces, and let Y_1 be the complement of a system of NRRPs for X_1 satisfying the hypotheses of Proposition 6.2.

The circumradii of Delaunay triangles on any half-translation surface are bounded above by twice the diameter of the surface, which is bounded on K . The side lengths of Y_1 are bounded below by $C(K)$. Thus by the law of sines, the angles of Delaunay triangles are bounded below. So we could have picked $c(K)$ small enough that when we change all of the period coordinates by at most $c(K)$, all angles of the Delaunay triangulation remain bounded away from zero. So the Delaunay triangulation for Y is uniformly bounded away from degeneration, ie all of the side lengths and angles are uniformly bounded away from 0. If $d_{\text{Euclidean}}(X_1, X_2) < c(K)$, then along a path from X_1 to X_2 we can make a choice of NRRP decomposition so that the radii, edges of NRRPs, and edges of Delaunay triangles that start as edges of Y_1 vary in a Lipschitz manner with respect to arc length, and the Lipschitz constant depends only on K . Let Y_2 be the complement of the NRRPs in X_2 , as chosen along the path.

By Proposition B.2 the triangulation of Y_1 persists to a triangulation Y_2 , and this determines an affine map on triangles of Y_1 to a triangulation of Y_2 . Since no edge gets too short and no angle gets too close to 0, it follows that if the length of the path is ϵ , then as $\epsilon \rightarrow 0$ this part of the map is $(1 + O(\epsilon))$ -quasiconformal on these triangles, and the implied constants in the $O(\epsilon)$ depend only on K .

6.2 The Beurling–Ahlfors extension

The final step is to show how to build a quasiconformal map between NRRPs, and to estimate its dilatation. We will do this by means of a Beurling–Ahlfors extension. Given an orientation-preserving self-homeomorphism h of \mathbb{R}^+ , we have the following:

Definition 6.4 The *Beurling–Ahlfors extension* of h to the upper half-plane with parameter $r > 0$ is the function f_r given by

$$f_r(x + iy) = \frac{1}{2} \int_0^1 [h(x + yt) + h(x - yt)] dt + \frac{ir}{2} \int_0^1 [h(x + yt) - h(x - yt)] dt.$$

Definition 6.5 Let $\rho \in [1, \infty)$. We say that an orientation-preserving homeomorphism $h: \mathbb{R} \rightarrow \mathbb{R}$ is ρ -*quasisymmetric* on \mathbb{R} if for all $x \in \mathbb{R}, t > 0$ we have

$$\frac{1}{\rho} \leq \frac{h(x + t) - h(x)}{h(x) - h(x - t)} \leq \rho.$$

Theorem 6.6 [Beurling and Ahlfors 1956] *If h is ρ -quasisymmetric, then for some explicit choice of r the Beurling–Ahlfors extension f_r is ρ^2 -quasiconformal. Moreover, $r = 2 + O(\rho - 1)$ as $\rho \rightarrow 1^+$.*

Note that $r = 2$ extends a Möbius transformation of \mathbb{RP}^1 to a Möbius transformation of \mathbb{CP}^1 . Any choice of r yields a quasiconformal extension, but possibly with worse dilatation.

In the event that an NRRP does not contain a marked point, the Beurling–Ahlfors extension, together with the piecewise-affine maps with compatible boundary conditions, completes our construction of a quasiconformal map. In the case with marked points, we will have to compose with a quasiconformal map on each NRRP which is trivial on the boundary and moves the marked point to the correct location. For this we will uniformize our map to the upper half-plane and apply the unique \mathbb{R} -linear map on the complex plane fixing the real line pointwise and taking our marked point to its desired location.

We have established uniform control over the NRRPs. If the only interior singularity is a pole, then our map of NRRPs is in fact already an isometry in the singular metric, and hence conformal. In all other cases we have at least four poles on the boundary, and we can assume three of them are 0, 1, and ∞ . Subject to this normalization, if we double each NRRP along its boundary, the singularities in the interior of each NRRP can be chosen to belong to some fixed compact subset of the upper half-plane.

This motivates the following:

Lemma 6.7 *Suppose that P_1 and P_2 are NRRPs. Let $h_q : \partial P_1 \rightarrow \partial P_2$ be a piecewise-affine identification of sides with respect to the singular metrics on P_1 and P_2 , and affine on each pair of corresponding sides. Assume that when we uniformize each of P_1 and P_2 to the upper half-plane, three pairs of corresponding corners are sent to 0, 1, and ∞ , the distance between any two finite singularities on the boundary is between $1/B$ and B for $B > 1$, and all interior singularities belong to some fixed compact subset K of the open upper half-plane. For all sufficiently small ϵ , if the singularities of P_1 and P_2 can be put in bijection so that corresponding singularities are distance at most ϵ apart in the plane when uniformized, then the quasisymmetry constant ρ of the boundary map $h : \mathbb{R} \rightarrow \mathbb{R}$ induced by h_q and our chosen uniformizations is at most $1 + C\epsilon$, for some C depending only on K , the number of sides of P_1 and P_2 , and B .*

Proof The boundary map h is continuous, and it is differentiable except possibly at points corresponding to vertices of the NRRP. It is enough to show that the derivative of the boundary map satisfies $h'(x) - 1 = O(\epsilon)$ uniformly on the complement of these points.

Suppose that one pair of corresponding sides is $h([a_1, b_1]) = [a_2, b_2]$, where a_1, a_2, b_1 , and b_2 are all finite. (We will deal with the intervals with an endpoint at ∞ later.) Assume the quadratic differentials are $(p_i(z)/((z - a_i)(z - b_i)q_i(z))) dz^2$ on $[a_i, b_i]$ where p_i and q_i are polynomials. Then for $a_1 < c < b_1$, our boundary map h is defined so that

$$\frac{\int_{a_1}^c \left(\frac{p_1(z)}{(z-a_1)(z-b_1)q_1(z)}\right)^{1/2} dz}{\int_{a_1}^{b_1} \left(\frac{p_1(z)}{(z-a_1)(z-b_1)q_1(z)}\right)^{1/2} dz} = \frac{\int_{a_2}^{h(c)} \left(\frac{p_2(z)}{(z-a_2)(z-b_2)q_2(z)}\right)^{1/2} dz}{\int_{a_2}^{b_2} \left(\frac{p_2(z)}{(z-a_2)(z-b_2)q_2(z)}\right)^{1/2} dz}.$$

We differentiate both sides of the above equation in c and solve for $h'(c)$:

$$h'(c) = \frac{\int_{a_2}^{b_2} \left(\frac{p_2(z)}{(z-a_2)(z-b_2)q_2(z)}\right)^{1/2} dz}{\int_{a_1}^{b_1} \left(\frac{p_1(z)}{(z-a_1)(z-b_1)q_1(z)}\right)^{1/2} dz} \left(\frac{p_1(c)q_2(h(c))(h(c) - a_2)(h(c) - b_2)}{p_2(h(c))q_1(c)(c - a_1)(c - b_1)} \right)^{1/2}.$$

By applying changes of coordinates consisting of translations by real numbers to each NRRP, we can make a few simplifying assumptions: at the cost of increasing ϵ to 2ϵ we can do a translation to assume $a_1 = a_2 = 0$. After this simplifying assumption, our argument can be rephrased. Let

$$f_i(z) = \frac{\left(\frac{p_i(z)}{z(z-b_i)q_i(z)}\right)^{1/2}}{\int_0^{b_i} \left(\frac{p_i(t)}{t(t-b_i)q_i(t)}\right)^{1/2} dt}.$$

Then $h'(c) = f_1(c)/f_2(h(c))$. So we need to show $\log[f_1(c)/f_2(h(c))] = O(\epsilon)$.

We will actually do this under the assumption that $0 < c \leq \frac{2}{3}b_i$. A basically identical argument will cover the cases $\frac{1}{3}b_i < c < b$, and a final argument will extend this to cover the two pairs of sides of P_1 and P_2 that have ∞ as an endpoint after we uniformize.

In what follows, C will be used to denote various positive constants that depend only on choices we have already made — its meaning may vary from one line to the next. The following will be immediate from our compactness assumptions. To simplify matters, we first prove:

Claim 6.8 $|\log(h(c)/c)| \leq C\epsilon.$

The idea is that the f_i are equal to $z^{-1/2}$ times two very similar functions. Recall that h is defined so that $\int_0^c f_1(z) dz = \int_0^{h(c)} f_2(z) dz$.

Define $g_i(z) = z^{1/2} f_i(z)$.

We know the following:

- The g_i are either both purely imaginary or both purely real, and never change sign.
- $-C \leq \log |g_i(z)| \leq C$ and $|g'_i(z)| \leq C$ for $z \in [0, \frac{3}{4}b_i]$.
- $|\log(f_1(z)/f_2(z))| = |\log(g_1(z)/g_2(z))| < C\epsilon$ for $z \in (0, \frac{3}{4}b_i]$.
- $-C \leq \log\left[z^{1/2} \int_0^z f_i(t) dt\right] \leq C$.
- $\frac{d}{dz} \log\left|\int_0^z t^{-1/2} g_i(t) dt\right| \geq z^{-1/2}/(Cz^{1/2}) = 1/(Cz)$ for $z \in (0, \frac{3}{4}b_i]$.

From the last observation we see that the amount that we need to perturb z to change the value of $\log \int_0^z t^{-1/2} g_i(t) dt$ by ϵ is at most $Cz\epsilon$, for all sufficiently small ϵ . In particular, for $i = 2$ we can start with $\int_0^c f_2(t) dt$, and by moving the upper limit of integration most $C\epsilon c$ away from c we can adjust the value of the integral of f_2 by a factor of at least $1 + C\epsilon$. The intermediate value theorem implies that the adjustment of the upper limit of integration needed to get from c to $h(c)$ is at most $C\epsilon c$. This establishes the claim.

Now that we have the claim, the main case of the lemma follows quite easily: the logarithmic derivatives of all factors of f_i are bounded by some constant C on $[c, h(c)]$ or $[h(c), c]$, and $\log(f_1/f_2)$ is uniformly close to $(0, \frac{3}{4}b)$. So the claim, plus boundedness of logarithmic derivatives of the remaining factors, gives us $g_1(c)/g_2(h(c)) = 1 + O(\epsilon)$, and by the claim, $c/h(c) = 1 + O(\epsilon)$. This proves the main case of the lemma. Interchanging the roles of a_1 and b_1 handles the case in which c is in the right half of the interval $[a_1, b_1]$.

Now we must handle the infinite intervals. For $1 < c < \infty$, we can conjugate $h(z)$ by the map $z \rightarrow 1/z$ to a function $H(c)$ satisfying an estimate of the type we originally had. The lemma and its proof hold for H ; in particular $H(1/c)/(1/c) = 1 + O(\epsilon)$. Finally, we have

$$h'(c) = \frac{d}{dc} \left(\frac{1}{H(1/c)} \right) = H'(1/c) \frac{1}{c^2 H(1/c)^2} = 1 + O(\epsilon). \quad \square$$

Proof of Theorem 1.2 We describe how to build a quasiconformal map between two nearby quadratic differentials, provided they are within some distance $c(K)$ depending only on a compact set $K \subset \text{QD}(\mathcal{T}_{g,n})$.

Let X_1 and X_2 be two such Riemann surfaces, and let Y_1 be the complement of a system of NRRPs for X_1 satisfying the hypotheses of Proposition 6.2.

The circumradii of Delaunay triangles on any half-translation surface are bounded above by twice the diameter of the surface, which is bounded on K . The side lengths of Y_1 are bounded below by $C(K)$. Thus by the law of sines, the angles of Delaunay triangles are bounded below. So we could have picked $c(K)$ small enough that when we change all of the period coordinates by at most $c(K)$, all angles of the Delaunay triangulation remain bounded away from zero. So the Delaunay triangulation for Y is uniformly bounded away from degeneration, ie all of the side lengths and angles are uniformly bounded away from 0. If $d_{\text{Euclidean}}(X_1, X_2) < c(K)$, then along a path from X_1 to X_2 we can make a choice of NRRP decomposition so that the radii, edges of NRRPs, and edges of Delaunay triangles that start as edges of Y_1 vary in a Lipschitz manner with respect to arc length, and the Lipschitz constant depends only on K . Let Y_2 be the complement of the NRRPs in X_2 , as chosen along the path.

Since the triangulation persists, this determines an affine map on triangles of Y_1 to a triangulation of Y_2 . Since no edge gets too short, and no angle gets too close to 0, it follows that if the length of the path is ϵ , then as $\epsilon \rightarrow 0$ this part of the map is $(1 + O(\epsilon))$ -quasiconformal on these triangles, and the implied constants in the $O(\epsilon)$ depend only on K . We need only extend this map to the interiors of the disks, which are pairs of NRRPs in Y_1 and Y_2 .

The NRRPs themselves are represented by quadratic differentials on disks, which can be uniformized to the closure of the upper half-plane in $\widehat{\mathbb{C}}$, and the metrics induce quadratic differentials on these close disks.

We may assume three pairs of corresponding vertices map to 0, 1, and ∞ under this uniformization if the NRRPs do not contain marked points; we may assume that one vertex is ∞ and the marked point is sent to i if there is a marked point. The quadratic differentials on corresponding pairs of NRRPs are Euclidean distance $O(\epsilon)$ in the moduli spaces of quadratic differentials in Lemma 5.12.

Thus, the locations of the zeros and poles in uniformized coordinates differ by $O(\epsilon^{\alpha_{g,n}})$, where $\alpha_{g,n} = 2/[2 + a_n(4g - 4 + n)]$, by Lemma 5.11. This allows us to use Lemma 6.7 to estimate the dilatation of the map on the NRRPs as $1 + O(\epsilon^{\alpha_{g,n}})$.

We have now produced a quasiconformal map f between the compact Riemann surfaces that are the completions of X_1 and X_2 with dilatation $1 + O(\epsilon^{\alpha_{g,n}})$. Unfortunately, it might not send marked points on X_1 to marked points on X_2 , and this problem occurs precisely when the marked points are in NRRPs.

However, if we can find another quasiconformal map $\tau: X_2 \rightarrow X_2$ with dilatation $1 + O(\epsilon^{\alpha_{g,n}})$ that fixes the complement of each NRRP in X_2 , but such that $\tau(f(p_1)) = p_2$ whenever p_1 and p_2 are corresponding pairs of marked points in X_1 and X_2 , respectively, then the log of the dilatation of $\tau \circ f$ is $O(\epsilon^{\alpha_{g,n}})$, ie the Teichmüller distance between X_1 and X_2 is $O(\epsilon^{\alpha_{g,n}})$.

In uniformized coordinates, the map from the boundary of each NRRP in X_1 to the corresponding NRRP of X_2 is nearly the identity on any compact subset of the real line. More precisely, if S is a compact subset of the real line, then for all $x \in S$ we have

$$f(x) - x = O(\epsilon^{\alpha_{g,n}}).$$

It follows that for any compact $K \subset \mathbb{H}$ the Beurling–Ahlfors extension moves points in K by $O(\epsilon^{\alpha_{g,n}})$ as well.

We therefore pick τ to be the \mathbb{R} –linear map that is the identity on $\partial\mathbb{H}$ and moves the marked point in K from $f(i)$ to i , which is a distance of $O(\epsilon^{\alpha_{g,n}})$. Then on the interior of each NRRP, τ also has constant dilatation that is $1 + O(\epsilon^{\alpha_{g,n}})$.

We have constructed a homeomorphism that satisfies the dilatation bound at almost every point, as its first partial derivatives are defined off of a measure 0 set, and these partial derivatives are clearly bounded away from every point except possibly the singularities and vertices of the NRRPs since there is a conformal metric in which the partial derivatives converge piecewise. This shows that if we delete a finite set, the first partial derivatives are in L^2_{loc} . To finish, we simply recall that a homeomorphism which is M –quasiconformal on the complement of a finite set is M –quasiconformal; see eg [Hubbard 2006]. \square

Appendix A Local finiteness of period coordinate systems

In this appendix we show the following, leaving the proof to the end of the section:

Proposition A.1 *There are only finitely many systems of period coordinates represented by saddle connections of length at most D in any compact subset K of $\text{QD}^1(\mathcal{T}_{g,n})$.*

To this end, we will show that every such period coordinate system is related to a basis of the edges of the Delaunay triangulation by one of finitely many transition matrices, and there are only finitely many choices of Delaunay triangulation on K . We will always assume that our triangulations are *labeled* and *marked*, that is, the vertices and edges are distinguishable, we know which marked points (if any) on the base surface correspond to which singularities, and we know which homotopy classes of curves on the surface are represented by which homotopy classes of curves on the graph. Finally, we will also include in the data of the Delaunay triangulation the sign of the slope of each saddle connection.

Proposition A.2 *Up to the action of the mapping class group, there are only finitely many triangulations that can occur as Delaunay triangulations in $\text{QD}^1(\mathcal{T}_{g,n})$.*

Proof The number of triangles is a function only of the stratum, since the sum of the angles of a triangle is π and the sum of the angles of all triangles is the sum of the cone angles of the singularities. (There is one triangle for each vertical separatrix emanating from a singularity.) \square

Proposition A.3 *Only finitely many Delaunay triangulations occur in any compact $K \subset \text{QD}^1(\mathcal{T}_{g,n})$.*

Proof If not, then infinitely many such triangulations occur in a single orbit of the $\text{Mod}(S_{g,n})$. Given a fixed Delaunay triangulation T , the space of complex lengths we can assign to the edges and still have a unit-area quadratic differential with systole $> \epsilon$ and such that the triangulation remains Delaunay (but possibly becomes degenerate) is compact. (If there are no short simple closed curves, Mumford's compactness criterion implies the diameter of the surface, and hence the diameter of each Delaunay triangle, is bounded. The inequalities that guarantee that the triangulation is a Delaunay triangulation are closed conditions.) This contradicts the proper discontinuity of the action of $\text{Mod}(S_{g,n})$ on $\text{QD}^1(\mathcal{T}_{g,n})$. \square

Proposition A.4 *There is a uniform bound on the number of geodesic concatenations of saddle connections of length at most D on K . In particular, there is a uniform bound on the number of saddle connections of length at most D on K .*

Proof We first reduce to the case in which $(X, q) \in K$ has no poles. If q has marked points, we simply find q' belonging to a compact K' in a moduli space of higher-genus surfaces by taking a pole-free degree-2 branched cover that is branched over all poles such that any geodesic concatenation of saddle connections on (X, q) can be lifted to a saddle connection or geodesic concatenation of saddle connections on (X', q') ; we then bound the number of geodesic segments joining cone points and marked points on (X', q') .

To form (X', q') in a way that guarantees that X' stays in a compact K' , the branching will be over all marked points, and, if the number of marked points is odd, one additional point. We pick this extra branch point to be as far away as possible from the marked points of X . To this double cover we give the metric associated to the quadratic differential that is $1/\sqrt{2}$ times the pullback of q (this is so that the area is 1). Because we have a lower bound on the distance between branch points in (X, q) in K , the space K' of double covers has a lower bound on the injectivity radius, ie it satisfies Mumford's compactness criterion. Thus we have reduced to the case of dealing with a compact subset of a moduli space without poles.

To deal with surfaces without poles, we fix a word metric on $\pi_1(S_g)$. We first claim that there exist real numbers $k, c > 0$ such that whenever $(X, q) \in K$, the fundamental group of S_g in this word metric is equivariantly (k, c) -quasi-isometric to the universal cover (\tilde{X}, \tilde{q}) of $(X, q) \in K$ via a (k, c) quasi-isometry $\phi_{(X,q)}$ that sends $h \in \pi_1(K)$ to $h(x, q)$, where $x \in X$ is chosen in a continuous section of the universal disk bundle over $\text{QD}^1(\mathcal{T}_g)$, ie the universal cover of the surface bundle over $\text{QD}(\mathcal{T}_{g,n})$. The uniformity of quasi-isometries then follows from finiteness of $\{h \subset \pi_1(S_g) : d_{\tilde{q}}(x, hx) < R \text{ for some } x \in K\}$.

Now, for each Delaunay triangulation T that appears in K , we make the quasi-isometries between the universal covers of surfaces with triangulation T more explicit by picking the basepoint to be a vertex p . By virtue of the quasi-isometry constants between the fundamental group of X and the universal cover

of X , there is a constant M depending on K and D such a ball of radius D in the universal cover of (X, q) can contain at most M lifts of any point in X distance less than D from any fixed lift of the basepoint.

If this were not the case, then there would be no upper bound on the number of closed geodesics in X with basepoint p of length at most $D + \max_{(X,q) \in K} \text{diam}_q(X)$ as X varies over K ; this would contradict the existence of the uniform quasi-isometry constants between the fundamental group and the universal cover of (X, q) .

In particular, we conclude that for any Delaunay vertices v_1 and v_2 , and a fixed lift w of v_1 , there are at most M lifts of geodesic paths of length at most D that start at w and end at lifts of v_2 , and therefore at most M geodesics from v_1 to v_2 on X of length at most D .

Since we may have chosen T to be any Delaunay triangulation and p to be any vertex, and the number of vertices is obviously bounded in K , this means that there is an upper bound on the number of geodesics of length at most D whose endpoints are singularities of q -metrics for all (X, q) . \square

Proposition A.5 *If $q_n \in K$ all have the same Delaunay triangulation and $q_n \rightarrow q$, then the common Delaunay triangulation of the q_n is a possibly degenerate Delaunay triangulation of q .*

Proof The conditions for a triangulation to be Delaunay are nonstrict inequalities in the absolute values of period coordinates, which persist under taking limits. \square

Corollary A.6 *For any $D > 0$ and any compact $K \subset \text{QD}^1(\mathcal{T}_{g,n})$ there are only finitely many period coordinate systems consisting of saddle connections of length less than D .*

Proof There are only finitely many period coordinate systems that can be Delaunay triangulations in a neighborhood of each point in K because there are only finitely many Delaunay triangulations at each point, and only finitely many Delaunay triangulations that can degenerate to a triangulation T without changing the topology. Thus each point in K has a neighborhood with a uniformly bounded number of possible Delaunay triangulations, and a simple compactness argument tells us that there are only finitely many Delaunay triangulations of surfaces in K . For each subset of K sharing a common Delaunay triangulation we can fix a basis of period coordinates coming from edges of the triangulation, and we can write the period of every homotopy class of geodesic arc whose endpoints are singularities in terms of this basis.

Finally, we claim that the saddle connections in a fixed homotopy class can only have finitely many representations in terms of any fixed basis for $H_1^{\text{odd}}(\tilde{X}, \tilde{\Sigma})$ consisting of saddle connections that are edges of the Delaunay triangulation, if we force (X, q) to remain in K . Indeed, by uniform quasi-isometry of the fundamental group with \tilde{X} (endowed with the pulled back q metric), such a saddle connection can only pass through a fixed finite list of Delaunay triangles, and it cannot leave and then enter a triangle since triangles are geodesically convex in the q -metrics. Thus there is a finite list of sequences of triangles a saddle connection can pass through, and by developing this chain of triangles in the plane, we can write its period as one of finitely many sums of sides of periods of the edges of the triangles. \square

Corollary A.7 For each compact $K \subset \mathcal{T}_{g,n}$ there is a finite set of triangulations $\{T_i\}_i \in I$ such that all quadratic differentials lying over K belong to a T_i -convex set for some $i \in I$.

Proof the q -metric on (X, q) is obtained by pulling back the Euclidean metric on \mathbb{R}^2 from a collection of charts, but one could just as easily have pulled back the L^∞ -metric $d(a+bi, c+di) = \max(|a-c|, |b-d|)$ to form an L^∞ q -metric on X . We may consider Delaunay triangulations with respect to the L^∞ flat metric, following [Guéritaud 2016]. Every triangle in the L^∞ Delaunay triangulation has its vertices on the boundary of an open square that is maximal with respect to the conditions that its lifts to the universal cover are embedded and that its edges be vertical and horizontal. If a triangulation is the L^∞ triangulation for a set of surfaces in $\text{QD}(\mathcal{T}_{g,n})$ that has nonzero Masur–Veech measure, then no edge of the triangulation is always horizontal or always vertical. So we may further restrict each saddle connection in the triangulation to have nonpositive slope or nonnegative slope, and include that in the combinatorial data of the triangulation. We may also disregard any triangulations that arise only on sets of measure zero.

We may now claim the conditions that a triangulation be L^∞ Delaunay are given by (nonstrict) linear inequalities in the period coordinates. Moreover, no triangle can have three edges of positive slope or three edges of negative slope. Hence, if we orient all edges so that the imaginary parts of periods are nonnegative, then for each L^∞ triangulation T the space of surfaces for which T is an L^∞ triangulation is T -convex.

Moreover, for each triangle there is a point p that is equidistant from the three vertices in the L^∞ -metric, and those three vertices are the nearest three vertices to p in the L^∞ metric (possibly along with some others, if the L^∞ triangulation is not unique). If this distance is d then they do not all have positive slope or all negative slope, since they touch at least three sides of a $2d \times 2d$ square with vertical and horizontal sides centered at p . This imposes a uniform bound on the lengths of saddle connections in the L^∞ Delaunay triangulation of $2\sqrt{2}$ times the diameter of the surface. Since it is obviously enough to consider unit-area surfaces lying over K , which have bounded diameter, this imposes a bound on the number of possible L^∞ Delaunay triangulations. \square

Proof of Proposition A.1 There are locally finitely many L^∞ -Delaunay triangulations; for each such triangulation we can locally fix names of the singularities. The collection of all singularities in the universal cover of (X, q) is quasi-isometric to the fundamental group of the compact surface \bar{X} , uniformly for all $(X, q) \in K$. In particular, there are finitely many ways to choose the names of the endpoints of a saddle connection of length at most L in the metric universal cover of (\bar{X}, q) , up to the action of $\pi_1(\bar{X})$, and therefore only finitely many systems of saddle connections of length at most L for each L^∞ Delaunay triangulation that occurs in K . \square

Appendix B δ -clusters and the Euclidean metric

First, we would like to establish a basic fact about systems of period coordinates persisting under perturbation. This will be useful throughout:

Proposition B.1 *Let $\mathcal{Q}(\mu)$ be a stratum of cluster differentials. Then there is a constant $\epsilon > 0$, depending only on μ , such that whenever Γ is a length-minimizing system of saddle connections on the vertex set of singularities of $q \in \mathcal{Q}(\mu)$ there is a holomorphic coordinate system $U \subset \mathbb{C}^r$ given by the logs of the periods $(\log(P_1), \dots, \log(P_r))$ of saddle connections in Γ , and moreover U contains an ϵ -ball about q in which the coordinates are logarithms of periods of saddle connections.*

Proof First we reduce to the case in which the stratum has no poles, except for a higher-order pole at ∞ . In the case of a cluster differential there is a length-minimizing tree of saddle connections Γ , and we claim that the length-minimizing tree $\tilde{\Gamma}$ for the double cover (\tilde{X}, \tilde{q}) of (X, q) branched only over the pole contains the lifts of all the edges of the base graph Γ . Therefore a small perturbation of the logs of the periods of Γ arises as the quotient of a perturbation of the logs of the periods of $\tilde{\Gamma}$ which descends. The proof is by following the greedy algorithm for length-minimizing spanning trees: repeatedly pick the shortest edge that does not form a cycle with previous edges. Every saddle connection of (\tilde{X}, \tilde{q}) is a lift of a saddle connection of (X, q) . Following the greedy algorithm to build $\tilde{\Gamma}$ will thus start with a lift of the shortest edge of Γ and then the other lift; this does not form a cycle because the double cover of the graph Γ branched over the pole is still a tree. Now, continuing inductively, assume that the first $2k$ edges of $\tilde{\Gamma}$ are the lifts of the first k edges of Γ . The shortest saddle connection of Γ that has not been picked has lifts which do not form cycles with the previously picked edges of $\tilde{\Gamma}$, because the double cover of Γ is a tree. However, any shorter edge would project down to an edge e , which, when added to the first k edges of Γ , creates a cycle. Adding the two lifts of e to Γ and Γ' would therefore produce a graph with first homology group of rank 2, since it would be obtained by taking two disjoint copies of a graph with first integral homology group \mathbb{Z} and identifying a pair of vertices. Since deleting one edge of a graph can only decrease the rank of the first homology group by 1, it follows that adding just one lift of e would create a cycle with the first $2k$ edges of Γ . Therefore the lift of the saddle connection e is not available to pick as the next edge of $\tilde{\Gamma}$. Hence the two lifts of the $(k+1)^{\text{st}}$ edge of Γ can be chosen by the greedy algorithm, and it is possible to choose one and then the other as the next two edges of $\tilde{\Gamma}$ while following the greedy algorithm. By induction, our claim follows, as does the reduction to the case with no poles.

Now suppose the proposition is false. We can find a sequence of counterexamples $\{q_m\}$ that would converge in the sense of Definition 4.35, and the graph Γ , as well as which collections of vertices correspond to vanishing clusters, are the same for all m . There would be perturbation vectors h_m of the logs of periods with $\|h_m\| \rightarrow 0$ such that $\{(\log P_1(q_m), \dots, \log P_r(q_m)) + th_m : t \in [0, 1)\}$ represents log-period coordinates of saddle connections with the graph Γ , but $(\log P_1(q_m), \dots, \log P_r(q_m)) + h_m$ does not. If the sequence $\{q_m\}$ has no vanishing clusters, this follows easily from the fact that $\log P_i$ is a system of period coordinates. So we induct on the number of nested clusters.

First, we note that the edges of Γ contain a spanning tree for each vanishing cluster of $\{q_m\}$, ie there are $n - 1$ edges that are internal to each vanishing cluster with n singularities. Otherwise there would be a trivial improvement of Γ by adding an edge joining two singularities in the same vanishing cluster that were not connected by a path in Γ that does not go outside the cluster, and deleting some other edge. In fact,

the edge set of Γ is naturally in bijection with the disjoint union of the edge sets of length-minimizing trees of saddle connections on the quadratic differentials corresponding the various cluster differentials (X_ℓ, q_ℓ) associated to the limit of the sequence $\{q_m\}$ in the compactification of Definition 4.35. To obtain the saddle connections corresponding to a length-minimizing tree in (X_ℓ, q_ℓ) , simply delete all singularities outside S_ℓ and contract all edges belonging to each $S_j \subsetneq S_\ell$; the remaining graph will have edges corresponding to S_ℓ .

We will show that for large enough m , each singularity is bounded away from each edge of Γ of which it is not a vertex along $\gamma_m([0, 1]) := \{(\log P_1(q_m), \dots, \log P_r(q_m)) + th_m : t \in [0, 1]\}$, contradicting the fact that the coordinate system does not extend the point where $t = 1$.

Step 1 For any two singularities a and b , $d_{q_m}(a, b)$ and the distance between a and b along the metric graph Γ is uniformly comparable as $m \rightarrow \infty$, ie $d_{q_m}(a, b) \asymp d_{q_m}|_\Gamma(a, b)$. The implied constants in \asymp depend on the sequence but not on m . Moreover, both of these distances are comparable to (\asymp) the q_m -diameter of the smallest vanishing cluster $S_\ell(a, b)$ containing a and b , or the entire singularity set if no such vanishing cluster exists, and comparable to the length of any saddle connection that is internal to $S_\ell(a, b)$ but not to a proper vanishing subcluster.

Proof This is clear from the definitions and Proposition 4.29, since any saddle connection whose limit is less than every positive multiple of $\text{diam}_q(S_\ell(a, b))$ would belong to a proper subcluster. \square

Step 2 Let $d_{q_m+th_m}$ denote the metric at time t along γ_m . For any two singularities a and b , $(d/dt) \log d_{q_m+th_m}(a, b)$ is defined for almost every t , and converges to 0 uniformly along points of definition as $m \rightarrow \infty$. Moreover, $\log d_{q_m+th_m}(a, b)$ is locally absolutely continuous, so the integral of this derivative represents the change in length.

Proof Since the angle formed by two saddle connections varies real-analytically, the set of times at which the configuration of saddle connections on the geodesic from a to b changes is discrete, consisting only of isolated times at which two of the saddle connections form an angle of π . This establishes the final claim, so we must now bound the logarithmic derivatives of distances. This is done by induction on the nesting of the cluster containing a and b .

First, we note that the distance from a to b is the sum of the lengths of finitely many saddle connections, and a saddle connection from c to d has period equal to the sum of the periods of saddle connections of the path from c to d in Γ with appropriate signs. There are therefore only finitely many ways to express the periods of saddle connections comprising the geodesic from a to b as linear combinations of periods from Γ , and since the geodesic does not revisit an edge, there is a bound on how many saddle connections are in the geodesic. So the derivatives (but not necessarily logarithmic derivatives) of distances are all uniformly bounded and tend to 0. In particular, this means that for each $\delta > 0$, every maximal vanishing cluster remains a δ -cluster along γ_m for almost every m .

Now, by induction, assume that some vanishing cluster S_ℓ is, for each $\delta > 0$, a δ -cluster in the singularity set of each point in γ_m for all sufficiently large m . Moreover assume that for all proper subclusters

$S_k \supsetneq S_\ell$, the conclusion holds for any singularities that belong to S_k but not to a common subcluster of S_k . Then along all but finitely many γ_m , the unique geodesic joining each pair of singularities in S_k consists entirely of saddle connections that are internal to S_k . We can then rescale each q_m to assume S_k has q_m -diameter 1, and ignore all singularities outside of S_k . By the same arguments, the conclusion holds for all pairs of singularities in S_k that are not in any proper vanishing subcluster. By induction, the conclusion holds for all vertices a and b of Γ . \square

Step 3 By Step 2, no singularities collide along γ_m as $t \rightarrow 1$ and γ_m may be extended to all $[0, 1]$, but with the possibility that some edges of Γ may degenerate to concatenations of saddle connections rather than saddle connections at $\gamma(1)$. Moreover, for each pair of singularities a and b , the direction of the geodesic from a to b varies almost everywhere differentiably, locally absolutely continuously, and with derivative uniformly tending to 0 on points of definition as $m \rightarrow 0$. In particular, for each saddle connection that exists along γ , the angle changes absolutely continuously, and differentiably with uniformly bounded derivative.

Proof It is sufficient to show that the imaginary part of the logarithmic derivative of the period of every saddle connection goes to 0 as $\gamma_m \rightarrow \infty$. If a and b are joined by a saddle connection at some point along γ_m and m is sufficiently large, let $S_\ell(a, b)$ be the minimal vanishing cluster containing a and b . For all points $q_m + th_m$ on γ we have $d_{q_m+th_m}(a, b) \asymp \text{diam}_{q_m}(S_\ell(a, b))$. The period of the saddle connection joining a and b is a bounded linear combination of periods of saddle connections of Γ internal to $S_\ell(a, b)$, all of which are at most a bounded multiple of $d_{q_m+th_m}(a, b)$, and each of these periods has logarithmic derivative going to 0 as $m \rightarrow \infty$. It follows that the logarithmic derivative of the saddle connection joining a to b goes to 0. \square

Step 4 By Step 2 no singularities collide along γ_m as $t \rightarrow 1$, so it must be the case that the distance from some singularity a to an edge of Γ joining two vertices $b, c \neq a$ goes to 0 along γ_m as $t \rightarrow 1$. The only way this can happen is if $d_{q_m}(a, b)/d_{q_m}(b, c) \rightarrow 0$ or $d_{q_m}(a, c)/d_{q_m}(b, c) \rightarrow 0$.

Proof By the length-minimizing property, each saddle connection of Γ is a side of two equilateral singularity-free triangles in q_m . Then $d(a, b) + d(a, c) \rightarrow d(b, c)$ along γ_m as $t \rightarrow 1$, but since distances are preserved up to a small multiplicative error, $\lim_{m \rightarrow \infty} (d_{q_m}(a, b) + d_{q_m}(a, c))/d_{q_m}(b, c) \rightarrow 1$. Now, if b and c belong to a vanishing cluster not containing a this is impossible, and if the only vanishing clusters containing a and one of b or c contains all three, then a is not on the geodesic joining b and c in the cluster differential associated to this vanishing cluster. \square

Step 5 Assume that as $t \rightarrow 1$ along infinitely many γ_m , the singularities a and b belong to a vanishing cluster S_ℓ , and b is joined to a vertex $c \notin S_\ell$ by an edge of Γ . Then the angle formed at b by the geodesic from b to c and the geodesic from b to a remains bounded away from 0 along γ_m .

Proof This is obvious from Step 3, since the geodesic joining a to b is a concatenation of saddle connections. This completes the proof, since as $t \rightarrow 1$, the distance between a and any saddle connection in Γ not containing a does not tend to 0 along γ_m as $t \rightarrow 1$. \square

Proposition B.2 *Let K be a compact subset of $\mathcal{M}_{g,n}$ or a moduli space of cluster differentials. Then there is some $\epsilon(K) > 0$ such that for all half-translation surfaces (X, q, Σ) , length-minimizing systems of saddle connections whose pairs of lifts added with opposite signs form a basis for $H^1_{\text{odd}}(\tilde{X}, \tilde{\Sigma})$ persist when the logs of their periods all change by at most $\epsilon(K)$.*

Proof The first thing we observe is that given a sequence that converges in the sense of Definition 4.35, the collection of saddle connections that belong to saddle connections internal to vanishing clusters is either linearly independent or has exactly one dependence relation, which we will describe. Cutting \tilde{X} along this system of saddle connections produces either one connected component or two, and if produces two components the relation is that the sum of the lifts of the edges that divide the length-minimizing tree into two components, each of which contains an odd number of cone points of cone angle an odd multiple of π , taken with appropriate signs, is the shared boundary of two surfaces. Therefore all but the longest such edge will be included in any length-minimizing system of saddle connections. By Proposition B.2 all of the saddle connections in the length-minimizing system internal to vanishing clusters will persist, since a perturbation that changes the log-periods of all k of the remaining saddle connections by at most ϵ changes the log-period of the longest one by $O(k\epsilon)$. The proof that saddle connections that are not internal to vanishing clusters persist is essentially the same as in the cluster differential case, since (\tilde{X}, \tilde{q}) admits a cover branched only over the poles, and the universal cover of this is complete and nonpositively curved, so there is a unique geodesic joining each pair of points. \square

Proof of Lemma 4.43 This is immediate from Lemma 4.20, and Propositions 4.13 and B.2. \square

Definition B.3 A length-minimizing period-coordinate system for (X, q) is *persistent* for (X', q') if for each coordinate, the change in the log-period is less than the quantity $\epsilon(K)$ for some compact set K containing (X, q) in its interior.

Proof of Proposition 4.41 First, we note that there is an upper bound on the systole of (X, q) and a lower bound on the diameter of (X, q) coming from K . We can fix a compact set $K' \subset \text{QD}^1(\mathcal{M}_{g,n})$ which contains the projection of K to the moduli space $\mathcal{M}_{g,n}$, defined to be all surfaces (X, q) with some fixed lower bound on the d_q -systole and some fixed upper bound on the d_q -diameter. For all period coordinate systems on K' with upper bound L on the length of the saddle connections used, as (X, q) varies, the q -distances between pairs of singularities, the systole, and the diameter of (X, q) are uniformly Lipschitz with respect to the path metric $d_{\text{Euclidean}}$, since there are locally finitely many period coordinate systems by Proposition A.1. Since there is a neighborhood of K which projects to K' , and the path γ can be assumed to lie entirely in K' , the proposition follows easily from the Lipschitz property. \square

Proof of Proposition 4.44 Clearly the only way a sequence of counterexamples is possible is if

$$\frac{d_{\text{Euclidean}}(X_m, Y_m)}{\epsilon_m} \rightarrow 0.$$

Obviously a sequence of counterexamples contains a subsequence in which $\{X_m\}$ converges in the sense of Notation 5.3 by passing to a subsequence. Moreover, we can assume that for every vanishing cluster D of $\{X_m\}$, the ratio $\text{diam}_{q_m}(D)/\epsilon_m$ converges in $[0, \infty]$.

Moreover, we can pick a system of period coordinates that consists of a maximal set of saddle connections internal to vanishing clusters whose periods are linearly independent in $H^1(\tilde{X}, \tilde{\Sigma}; \mathbb{C})$, and a length-minimizing set of complementary geodesic arcs in X_m whose endpoints are singularities, that complete the system of period coordinates. Note that the saddle connections defining coordinate system are internal to vanishing clusters if and only if their lengths go to 0. Then, in these systems of period coordinates, we make the following claim:

Claim B.4 *The complementary geodesic arcs can be chosen (Hausdorff) continuously on the straight line-segment path from X_m to Y_m , and the change in their periods is $\asymp \epsilon_m$ as $m \rightarrow \infty$.*

Proof The only periods that change are those whose lengths are bounded away from zero, and the perturbations are close to zero. Thus the length-minimizing coordinate system is persistent. The claim is thus immediate from Proposition B.2 and the fact that distance in each of the finite permissible coordinate systems is uniformly comparable to distance in the length-minimizing system.

If a sequence of counterexamples exists then we can of course choose them to lie in the principal stratum, since the principal stratum is dense on each coordinate chart.

Now, we assume that we have a counterexample sequence $\{(X_m, Y_m)\}_{m=1}^\infty$ with

$$d_{\text{Euclidean}}(X_m, Y_m) < \epsilon'_m = o(\epsilon_m),$$

ie $\lim_{m \rightarrow \infty} \epsilon'_m/\epsilon_m = 0$. Then we will show that pairs of corresponding complementary saddle connections of X_m and Y_m in the persistent coordinate chart have periods differing by $o(\epsilon_m)$.

Consider a path $\gamma_m(t)$ from X_m to Y_m that is rectifiable, parametrized by arc length, and of length ϵ'_m with respect to $d_{\text{Euclidean}}$. For $1 \leq j \leq 4g - 4 + 2n$, make Hausdorff continuous choices of singularities $a_m^j(t)$ with the property that for each t , the collection of cone points $a_m^j(t)$, counted with multiplicity, are the singularities of $\gamma_m(t)$, where $a_m^j(t)$ for $1 \leq j \leq n$ correspond to the marked points and $a_m^j(t)$ for $n + 1 \leq j \leq 4g - 4 + n$ correspond to the zeros with multiplicity (or multiplicity -1 if they collide with marked points). These choices are not unique, but they are unique up to a permutation of the singularities that preserves maximal vanishing clusters, since all nonuniqueness is caused by collisions of singularities, and this can only happen to two singularities in the same vanishing cluster. For a proof that it is possible to continuously choose a root along the path γ , see eg [Armstrong 1968, Lemma 1].

Then t runs from 0 to ϵ'_m , and let $\ell_m(t)$ be a continuous choice of a simple (non-self-intersecting) geodesic joining $a_m^j(t)$ and $a_m^k(t)$ such that $\ell_m(0)$ is one of the complementary saddle connections that is part of the length-minimizing system for X_m . We note that this simple geodesic is not quite determined by its

endpoints and homotopy class because there may be poles, but there are at most two choices for each pair of endpoints and they can be chosen continuously by passing to an appropriate finite branched cover branched only over marked points and branched over all marked points, if such exist. (In particular, this operation is defined over the base surface, so this finite cover varies in a Lipschitz manner over $\text{QD}(\mathcal{T}_{g,n})$.) In fact, we may assume that we picked the functions $a_m^r(t)$ to be projections of continuous choices of singularities in this branched double cover; each such choice determines a unique choice of $\ell_m(t)$. For the remainder of the proof we may assume $\ell_m(t)$ is chosen continuously in the universal cover of some branched cover.

In particular, the length of $\ell_m(t)$ is bounded above and bounded away from 0. Then, by the definition of the Euclidean path metric, the following are uniformly Lipschitz in t , ie Lipschitz with a constant that holds for all sufficiently large m :

- (1) the length of $\ell_m(t)$,
- (2) the distance between any two choices of $\alpha_m^r(t)$ and $\beta_m^r(t)$, if these are two possible ways of choosing $a_m^r(t)$, for each r ,
- (3) the diameter of the maximal vanishing cluster containing a_m^r along $\gamma_m(t)$, for each r ,
- (4) the distance from any singularity $a_m^r(t)$ to $\ell_m(t)$, and
- (5) the slope of the longest segment of $\ell_m(t)$, measured as an angle in $\mathbb{R}/\pi\mathbb{Z}$.

Let b_m^j be the singularity that corresponds to $a_m^j(0)$ in the persistent coordinate system for X_m that extends to Y_m . In order for us to get a contradiction it must be the case that for some choice of a_j and a_k the singularity $a_m^j(\epsilon_m)$ in Y_m is not equal to b_m^j , since otherwise the length and angle of $\ell_m(\epsilon_m)$, and hence also the period of the corresponding saddle connection, would each differ by $O(\epsilon_m)$ from $\ell_m(0)$. In fact, it must be the case that for some j , $\epsilon'_m/d_Y(b_m, a_m^j(\epsilon'_m)) \rightarrow 0$. Otherwise, we could move the endpoints of the geodesic joining $a_m^j(t)$ and $a_m^k(t)$ to b_m^j and b_m^k , and so the length of the geodesic, and the slope of the longest segment, would vary in a manner that is Lipschitz in the change of the endpoints. The resulting geodesic joining b_m^j and b_m^k would be a single saddle connection, whose length and angle are $o(\epsilon_m)$ away from the period of the saddle connection joining $a_m(0)$ and $b_m(0)$.

Finally, we show that $a_m^j(\epsilon'_m)$ and $a_m^k(\epsilon'_m)$ are $O(\epsilon'_m)$ away from the singularities b_m^j and b_m^k in Y_m , which makes this contradiction impossible. If we vary the endpoints in Y_m along rectifiable curves, the distance and angle of the longest segment again vary in a uniformly Lipschitz manner with respect to the endpoints. So we can therefore move the endpoints to the singularities corresponding to a_m and b_m , and conclude that their lengths and angles have changed by $o(\epsilon_m)$. However, this is also impossible, since a singularity that is not an endpoint of a saddle connection of length M in a length-minimizing system and is distance d away from both ends of the saddle connection is distance at least $\min(\frac{\sqrt{3}}{2}d, M)$ away from the saddle connection. (This is due to the fact that the original saddle connection is a side of two singularity-free equilateral triangles). This gives a contradiction to item (4) above. □

References

- [Armstrong 1968] **MA Armstrong**, *The fundamental group of the orbit space of a discontinuous group*, Proc. Cambridge Philos. Soc. 64 (1968) 299–301 MR Zbl
- [Bainbridge et al. 2019] **M Bainbridge, D Chen, Q Gendron, S Grushevsky, M Möller**, *Strata of k -differentials*, Algebr. Geom. 6 (2019) 196–233 MR Zbl
- [Beurling and Ahlfors 1956] **A Beurling, L Ahlfors**, *The boundary correspondence under quasiconformal mappings*, Acta Math. 96 (1956) 125–142 MR Zbl
- [Farb and Margalit 2012] **B Farb, D Margalit**, *A primer on mapping class groups*, Princeton Math. Ser. 49, Princeton Univ. Press (2012) MR Zbl
- [Fathi et al. 2012] **A Fathi, F Laudenbach, V Poénaru**, *Thurston’s work on surfaces*, Math. Notes 48, Princeton Univ. Press (2012) MR Zbl
- [Guéritaud 2016] **F Guéritaud**, *Veering triangulations and Cannon–Thurston maps*, J. Topol. 9 (2016) 957–983 MR Zbl
- [Hubbard 2006] **JH Hubbard**, *Teichmüller theory and applications to geometry, topology, and dynamics, I: Teichmüller theory*, Matrix Editions, Ithaca, NY (2006) MR Zbl
- [Hubbard and Masur 1979] **J Hubbard, H Masur**, *Quadratic differentials and foliations*, Acta Math. 142 (1979) 221–274 MR Zbl
- [Maskit 1985] **B Maskit**, *Comparison of hyperbolic and extremal lengths*, Ann. Acad. Sci. Fenn. Math. 10 (1985) 381–386 MR Zbl
- [Masur 1982] **H Masur**, *Interval exchange transformations and measured foliations*, Ann. of Math. 115 (1982) 169–200 MR Zbl
- [Masur and Smillie 1991] **H Masur, J Smillie**, *Hausdorff dimension of sets of nonergodic measured foliations*, Ann. of Math. 134 (1991) 455–543 MR Zbl
- [Mumford 1971] **D Mumford**, *A remark on Mahler’s compactness theorem*, Proc. Amer. Math. Soc. 28 (1971) 289–294 MR Zbl
- [Rafi 2007] **K Rafi**, *Thick-thin decomposition for quadratic differentials*, Math. Res. Lett. 14 (2007) 333–341 MR Zbl
- [Royden 1971] **HL Royden**, *Automorphisms and isometries of Teichmüller space*, from “Advances in the theory of Riemann surfaces”, Ann. of Math. Stud. 66, Princeton Univ. Press (1971) 369–383 MR Zbl
- [Veech 1986] **WA Veech**, *The Teichmüller geodesic flow*, Ann. of Math. 124 (1986) 441–530 MR Zbl

Monroe Township, NJ, United States

ian.frankel@gmail.com

Received: 16 April 2021 Revised: 15 April 2023

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Robert Lipshitz	University of Oregon lipshitz@uoregon.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Thomas Nikolaus	University of Münster nikolaus@uni-muenster.de
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Robert Oliver	Université Paris 13 bobol@math.univ-paris13.fr
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futер	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
Ian Hambleton	McMaster University ian@math.mcmaster.ca	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu	Chris Wendl	Humboldt-Universität zu Berlin wendl@math.hu-berlin.de
Thomas Koberda	University of Virginia thomas.koberda@virginia.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr		


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

The subscription price for 2024 is US \$705/year for the electronic version, and \$1040/year (+\$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<https://msp.org/>

© 2024 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 24 Issue 5 (pages 2389–2970) 2024

Formal contact categories	2389
BENJAMIN COOPER	
Comparison of period coordinates and Teichmüller distances	2451
IAN FRANKEL	
Topological Hochschild homology of truncated Brown–Peterson spectra, I	2509
GABRIEL ANGELINI-KNOLL, DOMINIC LEON CULVER and EVA HÖNING	
Points of quantum SL_n coming from quantum snakes	2537
DANIEL C DOUGLAS	
Algebraic generators of the skein algebra of a surface	2571
RAMANUJAN SANTHAROUBANE	
Bundle transfer of L -homology orientation classes for singular spaces	2579
MARKUS BANAGL	
A reduction of the string bracket to the loop product	2619
KATSUHIKO KURIBAYASHI, TAKAHITO NAITO, SHUN WAKATSUKI and TOSHIHIRO YAMAGUCHI	
Asymptotic dimensions of the arc graphs and disk graphs	2655
KOJI FUJIWARA and SAUL SCHLEIMER	
Representation stability for homotopy automorphisms	2673
ERIK LINDELL and BASHAR SALEH	
The strong Haken theorem via sphere complexes	2707
SEBASTIAN HENSEL and JENNIFER SCHULTENS	
What are GT-shadows?	2721
VASILY A DOLGUSHEV, KHANH Q LE and AIDAN A LORENZ	
A simple proof of the Crowell–Murasugi theorem	2779
THOMAS KINDRED	
The Burau representation and shapes of polyhedra	2787
ETHAN DLUGIE	
Turning vector bundles	2807
DIARMUID CROWLEY, CSABA NAGY, BLAKE SIMS and HUIJUN YANG	
Rigidification of cubical quasicategories	2851
PIERRE-LOUIS CURIEN, MURIEL LIVERNET and GABRIEL SAADIA	
Tautological characteristic classes, I	2889
JAN DYMARA and TADEUSZ JANUSZKIEWICZ	
Homotopy types of suspended 4-manifolds	2933
PENGCHENG LI	
The braid indices of the reverse parallel links of alternating knots	2957
YUANAN DIAO and HUGH MORTON	