msp

Geometry & Topology 22 (2018) 893–948

Long-time behavior of 3–dimensional Ricci flow C: 3–manifold topology and combinatorics of simplicial complexes in 3–manifolds

RICHARD H BAMLER

In the third part of this series of papers, we establish several topological results that will become important for studying the long-time behavior of Ricci flows with surgery. In the first part of this paper we recall some elementary observations in the topology of 3–manifolds. The main part is devoted to the construction of certain simplicial complexes in a given 3–manifold that exhibit useful intersection properties with embedded, incompressible solid tori.

This paper is purely topological in nature and Ricci flows will not be used.

57M50; 53C44, 57M15

1 Introduction

In this paper we establish several topological results that will be needed in the last part [D] of this series of papers.

In the first part, Section 2, we recall facts from the topology of 3–manifolds, which will be frequently used in this and the subsequent paper.

In the second part of this paper, Section 3, we prove a rather combinatorial–topological result (see Proposition 3.2). For clarity, we will first describe a much weaker version of this result.

Consider a closed 3–manifold M that does not contain any hyperbolic pieces in its geometric decomposition, eg a component of the thin part $\mathcal{M}_{\text{thin}}(t)$ from [A, Proposition 3.16]. We claim that then there is a finite, 2–dimensional simplicial complex V, as well as a continuous map $f_0\colon V \to M$ with the following property: if $\sigma \subset M$ is an embedded incompressible loop in M (meaning that the fundamental group of σ injects into the fundamental group of M), then σ intersects the image of every continuous map $f\colon V \to M$ that is homotopic to f_0.

Published: 22 January 2018 DOI: 10.2140/gt.2018.22.893

In the case in which M is a 3–torus, this statement follows in an elementary way: we can choose V to be the disjoint union of three 2–tori and $f_0\colon V \to M$ to be an embedding that maps those 2–tori to 2–tori that generate the second homology of M. The fact that every homotope of f_0 intersects every noncontractible loop of M can then be seen by a standard intersection number argument.

For the purposes of [D], we will however need a somewhat stronger statement, which as it turns out, is far more difficult to prove. This statement is also captured by Proposition 3.2. We will now describe this stronger statement in a somewhat restricted setting.

Assume again that M is a closed 3–manifold that cannot be covered by a 2–torus bundle over a circle. Then we claim that there is a finite, 2–dimensional simplicial complex V and a continuous map $f_0\colon V \to M$ such that the following holds: Consider an arbitrary solid torus $S \subset M$, $S \approx S^1 \times D^2$, that is incompressible in M (ie whose S^1–fibers are incompressible in M) and a map $f\colon V \to M$ that is homotopic to f_0. Moreover, consider an arbitrary Riemannian metric g on M. Then there is a "compressing domain" for S whose area, with respect to g, is bounded in terms of the area of f By this we mean that there is a smooth domain $\Sigma \subset \mathbb{R}^2$ and a smooth map $h\colon \Sigma \to M$ such that $h(\partial\Sigma) \subset \partial S$ and such that h restricted to the exterior circle of Σ is noncontractible in ∂S and h restricted to each interior circle of Σ is contractible in ∂S. This map h can be chosen such that its area with respect to g satisfies the bound

$$\operatorname{area}_g h < C \operatorname{area}_g f,$$

where C is a constant that only depends on the topology of M. Essentially, h will arise from the intersection $f(V) \cap S$, taking into account multiple overlaps and multiply counted faces. Note that this intersection can a priori be arbitrarily complex, without any bound on the number of edges, and it is a difficult task to extract a compressing domain out of it whose area is sufficiently controlled.

It is not known to the author whether such a statement remains true if M is covered by a 2–torus bundle over a circle. In this case, however, we can make use of the special topology of M and we will be able to prove a different statement, which will be sufficient for the arguments in [D]; more specifically, we will construct a sequence of continuous maps

$$f_1, f_2, \dots\colon V \to M$$

with the same simplicial complex V as a domain such that for every embedded, incompressible loop $\sigma \subset M$, every $n \geq 1$ and every map $f_n'\colon V \to M$ that is homotopic to $f_n\colon V \to M$, the image $f_n'(V)$ intersects σ at least n times. By this we mean that $f_n'^{-1}(\sigma)$ contains at least n points.

We refer to [0] for historical remarks and acknowledgements.

In the following we will assume that all manifolds are orientable and 3–dimensional, unless stated otherwise.

2 Preliminaries on 3–dimensional topology

In this section we present important topological facts, which we will frequently use in the course of this and the subsequent paper [D]. A more elaborate discussion of most of these results can be found in [3]. In the following, all manifolds are assumed to be connected and 3–dimensional.

Definition 2.1 (prime manifold) A manifold M is called *prime* if it cannot be represented as a connected sum $M = M_1 \# M_2$ of two manifolds M_1 and M_2 that are not diffeomorphic to spheres ($\approx S^3$).

Definition 2.2 (irreducible manifold) A manifold M is called *irreducible* if every smoothly embedded 2–sphere $S \subset M$ bounds a smoothly embedded 3–disk $D \subset M$, ie $\partial D = S$.

Recall that a manifold is prime if and only if it is either irreducible or diffeomorphic to $S^1 \times S^2$ (see [3, Proposition 1.4]). We furthermore have:

Proposition 2.3 *An orientable manifold M is irreducible if and only if $\pi_2(M) = 0$.*

Proof The backward direction follows from the sphere theorem (see [3, Theorem 3.8]). For the forward direction suppose that $\pi_2(M) = 0$ and let $S \subset M$ be a smoothly embedded sphere. So S is contractible and, by [3, Proposition 3.10], it bounds a compact contractible submanifold $N \subset M$. Following [3, Proposition 3.7], we conclude that if we attach a 3–disk to N, we obtain a closed, simply connected manifold M'. By the resolution of the Poincaré conjecture, $M' \approx S^3$ and hence N is a 3–disk. □

Definition 2.4 (incompressibility) Let X be a topological space and $Y \subset X$ a connected subspace. Then we call Y (*algebraically*) *incompressible in* X if the induced map $\pi_1(Y) \to \pi_1(X)$ is injective. Otherwise, we call Y (*algebraically*) *compressible*.

Proposition 2.5 (Dehn's lemma) *Let M be a manifold with boundary. If $C \subset \partial M$ is an embedded loop that is nullhomotopic in M, then C bounds an embedded disk D in M, ie $\partial D = C$ and $D \cap \partial M = C$.*

Proof See [3, Corollary 3.2]. □

Proposition 2.6 *Let M be a manifold (possibly with boundary) and $S \subset M$ a 2–sided embedded, connected surface. Then S is algebraically compressible if and only if there is an embedded loop $C \subset S$ that is homotopically nontrivial in S and that bounds an embedded* compressing disk *$D \subset M$ which meets S only in its boundary, ie $\partial D = C$ and $D \cap S = C$.*

In particular, the statement holds if $S = \partial M$.

Proof See [3, Corollary 3.3]. □

We can now define what we mean by a geometric decomposition.

Definition 2.7 (geometric decomposition of irreducible manifold) Let M be a compact, orientable and irreducible 3–manifold whose boundary consists of 2–tori. A *geometric decomposition* of M is a collection of pairwise disjoint, smoothly embedded 2–tori $T_1, \dots, T_m \subset M$ such that

(i) each torus T_i is incompressible in M (see Definition 2.4), and

(ii) each component of $M \setminus (T_1 \cup \dots \cup T_m)$ is either *hyperbolic* (ie it can be endowed with a complete metric of constant negative sectional curvature and finite volume) or it is *Seifert* (ie it carries a Seifert fibration that can be extended regularly to the closure of each of its ends).[1]

The decomposition is called *minimal* if no smaller subcollection of tori satisfies properties (i) and (ii).

If all components of $M \setminus (T_1 \cup \dots \cup T_m)$ are Seifert, then the manifold is called a *(prime) graph manifold* and the decomposition is called a *Seifert decomposition*.

[1] It may happen that the boundaries of two ends of such a component coincide. In this case it may not be possible to extend a (or any) Seifert fibration of this component to its closure, because the extensions of the fibration on the closure of each end may not be the same. Note also that a Seifert fibration on an orientable manifold can only have exceptional fibers of cone type.

For a minimal geometric decomposition, none of the components of $M \setminus (T_1 \cup \cdots \cup T_m)$ is diffeomorphic to $T^2 \times \mathbb{R}$ unless $m = 1$ and T_1 is nonseparating. Moreover, the following is true for a minimal geometric decomposition: Choose Seifert fibrations on each Seifert component of the decomposition, and consider a torus T_i that is adjacent to a Seifert component on both sides. Then the Seifert fibrations on either side of T_i extend to two nonisotopic fibrations on T_i.

Lastly, we mention that a minimal geometric decomposition is unique up to isotopy (see [3, Theorem 1.9]). So it is reasonable to speak of *the* (minimal) geometric decomposition of a manifold.

The statement of the geometrization conjecture is now the following:

Theorem 2.8 (geometrization conjecture) *Every closed, orientable, irreducible manifold admits a minimal geometric decomposition.*

Next, we will show that 3–manifolds that are not diffeomorphic to spherical space forms or $S^2 \times S^1$ have a sufficiently complex topology and hence cannot be covered by or decomposed into certain elementary pieces.

Lemma 2.9 *Let M be a closed, irreducible manifold and let $T \subset M$ be an embedded, 2–sided, compressible torus. Then T separates M into two components U and V (ie $M = U \cup V$ and $U \cap V = T$) and we can distinguish the following cases:*

(a) *Neither of the components U or V is diffeomorphic to a solid torus $S^1 \times D^2$. Then the compressing disks D for T either all lie in U or in V and for each such D a tubular neighborhood of $D \cup V$ or $D \cup U$ (depending on whether $D \subset U$ or $D \subset V$) is diffeomorphic to a 3–ball.*

(b) *Only one of the components U or V is diffeomorphic to a solid torus. Assume that this component is U. Then T has compressing disks in U. If it also has compressing disks in V, then U is contained in an embedded 3–ball in M and U is compressible in M (ie the map $\mathbb{Z} \cong \pi_1(U) \to \pi_1(M)$ is not injective).*

(c) *Both U and V are diffeomorphic to solid tori. Then M is diffeomorphic to a spherical space-form.*

Proof For the first part see [3, page 11]. Let D be a compressing disk for T and assume that $D \subset U$. Again by [3, page 11], we know that either U is a solid torus or a tubular neighborhood of $D \cup V$ is diffeomorphic to a 3–ball. So, if in case (a)

there are compressing disks for T in both U and V, then M is covered by two embedded 3–balls and we have $M \approx S^3$ by Lemma 2.10(a) (observe that the proof of Lemma 2.10(a) does not make use of this lemma). However, this contradicts the fact that an embedded 2–torus in S^3 bounds a solid torus on at least one side (see [3, page 11]). Case (b) follows similarly.

Consider now case (c). Let $K_1, K_2 \subset \pi_1(T) \cong \mathbb{Z}^2$ be the kernels of the projections $\pi_1(T) \to \pi_1(U)$ and $\pi_1(T) \to \pi_1(V)$. If $K_1 = K_2$, then $M \approx S^1 \times S^2$ contradicting the assumptions on M. So $K_1 \neq K_2$. Let $a_i \in K_i$ be generators. By an appropriate choice of coordinates, we can assume that $a_1 = (1,0) \in \mathbb{Z}^2$ and $a_2 = (p,q) \in \mathbb{Z}^2$ with $0 \leq p < q$. Then M is diffeomorphic to the lens space $L(p,q)$. □

Lemma 2.10 *Let M be a closed manifold and assume that $M = U \cup V$. Then:*

(a) *If U and V are diffeomorphic to a ball, then $M \approx S^3$.*

(b) *If U is diffeomorphic to a solid torus $S^1 \times D^2$ and V is diffeomorphic to a ball, then $M \approx S^3$.*

(c) *If U and V are diffeomorphic to a solid torus $\approx S^1 \times D^2$, then M is either not irreducible or it is diffeomorphic to a spherical space form.*

Proof In case (a), we can assume that U and V are the interiors of compact embedded 3–disks. So $\partial U \subset V$. By Alexander's theorem (see [3, Theorem 1.1]), ∂U bounds a 3–disk in V. So ∂U bounds a 3–disk on both sides and hence $M \approx S^3$.

Case (b) follows along the same lines; note that every embedded sphere in a solid torus bounds a ball.

For case (c) we can assume that M is irreducible. Moreover, by adding collar neighborhoods, we can assume that $\partial U \cap \partial V = \varnothing$. Let $T = \partial U$ and $V' = M \setminus \operatorname{Int} U$. Then T is compressible in V and, by Proposition 2.6, we find a spanning disk $D \subset \operatorname{Int} V$. If also $D \subset U$, then $U \setminus D$ is a 3–ball and $M = (U \setminus D) \cup V$ and we are done by case (b). So assume that $D \subset V'$. Then by Lemma 2.9(b), either V' is a solid torus or U is contained in an embedded 3–ball $B \subset M$. In the latter case, $M = B \cup V$ and we are again done by case (b). Finally, if V' is a solid torus, we are done by Lemma 2.9(c). □

Lemma 2.11 *Let M be a manifold and $T \subset M$ an embedded 2–torus that separates M into two connected components whose closures $U, V \subset M$ are diffeomorphic to $\mathrm{Klein}^2 \tilde{\times} I$ and $S^1 \times D^2$, respectively. Then M is either not irreducible or it is diffeomorphic to a spherical space form.*

Proof Consider the double cover $\widehat{U} \to U$ for which $\widehat{U} \approx T^2 \times I$. This cover extends to a double cover $\widehat{M} \to M$. Let $T' \subset \widehat{M}$ be the torus that projects to the zero section in $\mathrm{Klein}^2 \tilde{\times} I$. Then, as in the last part of the proof of Lemma 2.9(c), we can write $\widehat{M} = S_1 \#_{T'} S_2$, where S_1 and S_2 are solid tori. So $\widehat{M}$ is either diffeomorphic to $S^1 \times S^2$ or a lens space. In the first case, M is either diffeomorphic to $S^1 \times S^2$ or $\mathbb{R}P^3 \# \mathbb{R}P^3$ and in the second case, M is still spherical (see also [1]). □

The following lemma will be important in the proof of [D, Lemma 2.7].

Lemma 2.12 *Let M be compact, orientable, irreducible manifold (possibly with boundary) that is not diffeomorphic to a spherical space form. Consider a compact, connected 3–dimensional submanifold $N \subset M$ whose boundary components are tori and that carries a Seifert fibration. Assume that each boundary component $T \subset \partial N$ that is compressible in M either bounds a solid torus $\approx S^1 \times D^2$ on the other side or T separates M into two components and is incompressible in the component of $M \setminus T$ that does not contain N (if $T \subset \partial M$, then this component is empty).*

Then there are two cases: In the first case there is one boundary torus $T \subset \partial N$ that bounds a solid torus on the same side as N. In the second case every boundary component of N either bounds a solid torus on the side opposite to N or it is even incompressible in M. Moreover, in the second case, the generic Seifert fibers of N are incompressible in M.

Proof Some of the following arguments can also be found in [2; 5]. Denote the boundary tori of N by $T_1, \dots, T_m$. Assume that there is a component T_i that bounds a solid torus S_i on the side opposite to N such that the Seifert fibers in T_i are incompressible in S_i. Then we can extend the Seifert fibration of N to S_i. So assume in the following that for any T_i that bounds a solid torus S_i on the other side, the Seifert fibers of T_i are nullhomotopic in S_i. Denote by O the base orbifold of the Seifert fibration on N and call the projection $\pi\colon N \to O$. We remark that since M is orientable, the only singular points of O are cone points. Each T_i corresponds to a boundary circle $C_i = \pi(T_i) \subset \partial O$.

We first show that there is at most one T_i that bounds a solid torus S_i on the side opposite to N (we will call it from now on T_1): Assume, there were two such components T_1 and T_2 and denote the respective solid tori by S_1 and S_2. Let $\alpha \subset O$ be an embedded arc connecting C_1 and C_2 that does not meet any singular points. The preimage $Z_\alpha = \pi^{-1}(\alpha) \subset N$ is an annulus whose boundary components are each

nullhomotopic in S_1 or S_2, respectively. Let $D_1 \subset S_1$ and $D_2 \subset S_2$ be compressing disks for $Z_\alpha \cap \partial S_1$ and $Z_\alpha \cap \partial S_2$, respectively. Then $\Sigma_\alpha = D_1 \cup Z_\alpha \cup D_2$ is an embedded 2–sphere. Since D_1 and D_2 are nonseparating in S_1 and S_2, respectively, we conclude that Σ_α is nonseparating in M. This contradicts the assumption that M is irreducible.

Next, we show that if T_1 bounds a solid torus S_1 on the side opposite to N, then the topological surface underlying O is a planar domain: Assume not. Then there is an embedded, nonseparating arc $\alpha \subset O$ whose endpoints are distinct and lie in C_1. As before, this arc yields a nonseparating sphere $\Sigma_\alpha \subset M$, contradicting the irreducibility assumption of M.

Assume now for the rest of the proof that none of the tori T_i bound a solid torus on the same side as N. We will show in the following that then none of the tori T_i bounds a solid torus on either side, and that all T_i as well as the generic Seifert fibers on N are incompressible in M.

First assume that T_1 bounds a solid torus S_1 (on the side opposite to N). So the topological surface underlying O is a planar domain. We can find a collection of pairwise disjoint, embedded arcs $\alpha_1, \dots, \alpha_k \subset O$ with endpoints in C_1 that do not meet any singular points and that cut O into smaller pieces, each of which contain at most one singular point or one boundary component, and that are bounded by at most two of the arcs α_i and parts of C_1. The corresponding spheres $\Sigma_{\alpha_1}, \dots, \Sigma_{\alpha_k} \subset M$ bound closed 3–balls $B_1, \dots, B_k \subset M$. Any two such balls are either disjoint or one is contained in the other. Hence, either there is one B_i containing all other balls or there are two balls B_i, B_j such that any ball is contained in one of them. In the first case set $U = S_1 \cup B_i$ and in the second case set $U = S_1 \cup B_i \cup B_j$. From the position of the balls relatively to S_1 we conclude that U is diffeomorphic to a solid torus. Moreover, we can find a component $P \subset O \setminus (\alpha_1 \cup \dots \cup \alpha_k)$ whose boundary contains the arc α_i or the arcs α_i and α_j (depending on whether $U = S_1 \cup B_i$ or $U = S_1 \cup B_i \cup B_j$) such that $N = (B_i \cap N) \cup \pi^{-1}(P)$ or $N = ((B_i \cup B_j) \cap N) \cup \pi^{-1}(P)$, respectively. We can now distinguish the following cases:

- If P contains an orbifold singularity, then $\pi^{-1}(P)$ is diffeomorphic to a solid torus and $M = U \cup \pi^{-1}(P)$ and we obtain a contradiction using Lemma 2.10(c).
- If P contains a boundary component C_l of O, then we argue as follows: In this case P is diffeomorphic to an annulus. Let $\alpha' \subset P$ be an arc connecting C_l with C_1 and choose a compressing disk $D' \subset S_1$ for the arc $Z_{\alpha'} \cap \partial S_1$. Then

$Z_{\alpha'} \cup D'$ is a compressing disk for T_l. By our assumptions, T_l does not bound a solid torus. So, by Lemma 2.9(a), a tubular neighborhood of $T_l \cup Z_{\alpha'} \cup D'$ is diffeomorphic to a 3–ball. Since P is diffeomorphic to an annulus, this tubular neighborhood can be extended to a tubular neighborhood whose boundary is contained in U. This implies that M is covered by a solid torus and a ball, and Lemma 2.10(b) gives us a contradiction.

Hence, none of the T_i bound a solid torus on either side.

We argue that the generic Seifert fibers of N are incompressible in N: Using Lemma 2.10(c), we find that O cannot be a bad orbifold (ie the tear drop or the football) or a quotient of the 2–sphere. So, we can find a (possibly noncompact) cover $\widehat{O} \to O$ such that $\widehat{O}$ is smooth and corresponding to this a cover $\widehat{N} \to N$ such that we have an S^1–fibration $\widehat{N} \to \widehat{O}$. Observe that $\widehat{O}$ is not a 2–sphere, because otherwise, by Lemma 2.10(c), $\widehat{N} \approx S^3$, in contradiction to our assumptions. Using the long exact homotopy sequence and the fact that $\pi_2(\widehat{O}) = 0$, we conclude that a lift of any generic S^1–fiber $\gamma \subset N$ is incompressible in $\widehat{N}$ implying that γ is incompressible in N.

Next we show that any generic S^1–fiber γ of N is incompressible in M: Assume that there is a nullhomotopy $f\colon D^2 \to M$ for a nonzero multiple of γ. By a small perturbation, we can assume that f is transversal to the boundary tori $T_1, \dots, T_m$. So $f^{-1}(T_1 \cup \dots \cup T_m)$ consists of finitely many circles. Look at one of those circles $\gamma' \subset D^2$ that is innermost in D^2 and assume $f(\gamma') \subset T_i$. If $f|_{\gamma'}$ is homotopically trivial in T_i, then we can alter f such that γ' can be removed from the list. So assume that $f|_{\gamma'}$ is homotopically nontrivial in T_i. Let $D' \subset D^2$ be the disk that is bounded by γ'. Then by Proposition 2.6 and Lemma 2.9(b) we have $f(D') \subset N$. Since the generic Seifert fibers of N are incompressible in N, the restriction $f|_{\gamma'}$ cannot be homotopic to such a fiber, so it projects down to an arc that is homotopic to a nonzero multiple of the boundary circle C_i under π. Hence, a nonzero multiple of C_i is homotopically trivial in $\pi_{1,\mathrm{orbifold}}(O)$. We conclude that O can only be a disk with possibly one orbifold singularity. But this implies that N is diffeomorphic to a solid torus, in contradiction to our assumptions.

It remains to show that all tori T_i are incompressible in M. By Lemma 2.9(a), we conclude that if T_i is compressible in M, then T_i is contained in an embedded 3–ball. But this however contradicts the fact that the generic Seifert fibers of N are incompressible in M. □

3 Construction and analysis of simplicial complexes in M

3.1 Setup and statement of the results

In this section, we construct a simplicial complex V that will be used in [D] in combination with the area evolution result from [B]. We moreover analyze the intersections of images of V with solid tori in M. The results of this section are topological; however, we will need to make use of some combinatorial geometric arguments in the proofs.

We first recall the notion of simplicial complexes (compare also with [B, Definition 3.1]).

Definition 3.1 (simplicial complex) A (2–*dimensional*) *simplicial complex* V is a topological space that is the union of embedded, closed 2–simplices (triangles), 1–simplices (intervals) and 0–simplices (points) such that any two distinct simplices are either disjoint or their intersection is equal to another simplex whose dimension is strictly smaller than the maximal dimension of both simplices. V is called *finite* if the number of these simplices is finite.

In this paper, we assume V moreover to be *locally finite* and *pure*. The first property demands that every simplex of V is contained in only finitely many other simplices and the second property states that every 0– or 1–dimensional simplex is contained in a 2–simplex. We will also assume that all 2– and 1–simplices are equipped with differentiable parametrizations that are compatible with respect to restriction.

We will often refer to the 2–simplices of V as *faces*, the 1–simplices as *edges* and the 0–simplices as *vertices*. The 1–*skeleton* $V^{(1)}$ is the union of all edges and the 0–*skeleton* $V^{(0)}$ is the union of all vertices of V. The *valency* of an edge $E \subset V^{(1)}$ denotes the number of adjacent faces, ie the number of 2–simplices that contain E. The *boundary* ∂V is the union of all edges of valency 1.

Next let M be a closed, orientable, irreducible 3–manifold that is not a spherical space form. Consider a (not necessarily minimal) geometric decomposition of $T_1, \dots, T_m \subset M$ of M, ie the components of $M \setminus (T_1 \cup \dots \cup T_m)$ are either hyperbolic or Seifert (see Definition 2.7 for more details). We will assume from now on that the decomposition has been chosen so that no two hyperbolic components are adjacent to one another. This can always be achieved by adding a parallel torus next to a torus between two hyperbolic components and hence adding another Seifert piece $\approx T^2 \times (0,1)$. Let M_{hyp} be the union of the closures of all hyperbolic pieces of this decomposition and M_{Seif} the union of the closures of all Seifert pieces. Then

$M = M_{\mathrm{hyp}} \cup M_{\mathrm{Seif}}$ and $M_{\mathrm{hyp}} \cap M_{\mathrm{Seif}} = \partial M_{\mathrm{hyp}} = \partial M_{\mathrm{Seif}}$ is a disjoint union of embedded, incompressible 2–tori. Note that this construction parallels the "thick–thin decomposition" from [A, Proposition 3.16].

The goal of this section is to establish the following proposition. In this proposition, we need to distinguish the cases in which M is covered by a T^2–bundle over a circle (ie in which M is the quotient of a 3–torus, the Heisenberg manifold or the solvmanifold) and in which it is not. It is not known to the author whether part (a) of the proposition actually holds in both cases.

Proposition 3.2 *There is a finite simplicial complex V and a constant $C < \infty$ such that the following holds:*

(a) *In the case in which M is not covered by a T^2–bundle over a circle there is a map*

$$f_0\colon V \to M \quad \text{with } f_0(\partial V) \subset \partial M_{\mathrm{Seif}}$$

that is a smooth immersion on ∂V such that the following holds: Let $S \subset \operatorname{Int} M_{\mathrm{Seif}}$, $S \approx S^1 \times D^2$, be an embedded solid torus whose fundamental group injects into the fundamental group of M (ie S is incompressible in M). Let, moreover, $f\colon V \to M$ be a piecewise smooth map that is homotopic to f_0 relative ∂V and g a Riemannian metric on M. Then $f(V) \cap S \neq \varnothing$ and we can find a compact, smooth domain $\Sigma \subset \mathbb{R}^2$ and a smooth map $h\colon \Sigma \to S$ such that $h(\partial\Sigma) \subset \partial S$ and such that h restricted to the interior boundary circles of Σ of is contractible in ∂S and h restricted to the exterior boundary circle of Σ is noncontractible in ∂S and such that

$$\operatorname{area} h < C \operatorname{area} f.$$

(b) *In the case in which M is covered by a T^2–bundle over a circle the following holds: $\partial V = 0$ and there are continuous maps*

$$f_1, f_2, \dots\colon V \to M$$

such that for every $n \geq 1$, every map $f_n'\colon V \to M$ that is homotopic to f_n and every embedded loop $\sigma \subset M$, with the property that all nontrivial multiples of σ are noncontractible in M, the map f_n' intersects σ at least n times, ie $f_n'^{-1}(\sigma)$ contains at least n points.

We will first establish part (a) of the proposition in Sections 3.2–3.6 and then part (b) in Section 3.7.

3.2 Preliminary considerations for the case in which M is not covered by a T^2–bundle over a circle

Assume in this subsection that M is not covered by a T^2–bundle over a circle. In order to establish part (a) of Proposition 3.2, it suffices to construct a simplicial complex V and a map $f_0\colon V \to M$ with the desired properties for every component $M' \subset M_{\mathrm{Seif}}$, ie $f_0(\partial V) \subset \partial M'$, and check that the inequality involving the areas holds for every solid torus $S \subset M'$ and every homotope f of f_0. We will hence from now on fix a single component $M' \subset M_{\mathrm{Seif}}$.

The next lemma ensures that we can pass to a finite cover of M' and simplify the structure of M'. This simplification is not really needed in the following analysis, but it makes its presentation more comprehensible.

Lemma 3.3 *Under the assumptions of this subsection there is a finite cover $\widehat{\pi}'\colon \widehat{M}' \to M'$ such that the following holds: There is a Seifert decomposition $\widehat{T}_1, \dots, \widehat{T}_m \subset \widehat{M}'$ such that the components of $\operatorname{Int} \widehat{M}' \setminus (\widehat{T}_1 \cup \dots \cup \widehat{T}_m)$ are diffeomorphic to the interiors of manifolds $\widehat{M}_j = \Sigma_j \times S^1$ for $j = 1, \dots, k$, where each Σ_j is a compact orientable surface (possibly with boundary). The diffeomorphisms can be chosen in such a way that they can be smoothly extended to the boundary tori.*

Moreover, one of the following cases holds:

(A) *$\widehat{M}'$ is diffeomorphic to $T^2 \times I$ with $m = 0$ and $k = 1$.*

(B) *$\widehat{M}'$ is closed and diffeomorphic to an S^1–bundle over a closed, orientable surface Σ with $\chi(\Sigma) < 0$. In particular, we may assume that $m = k = 1$ and the surface Σ arises from Σ_1 by gluing together its two boundary circles.*

(C) *Σ_j has at least one boundary component and $\chi(\Sigma_j) < 0$ for all $j = 1, \dots, k$ and at each torus T_i the fibers coming from the S^1–fibration induced from either side are not homotopic to one another.*

Proof The arguments in this proof are similar to those in [4, Proposition 4.4].

Let $T_1, \dots, T_m \subset M'$ be a Seifert decomposition of M', that is, $T_1, \dots, T_m$ are pairwise disjoint, embedded, incompressible 2–tori such that the components of $\operatorname{Int} M' \setminus (T_1 \cup \dots \cup T_m)$ are diffeomorphic to the interiors of compact Seifert spaces $M'_1, \dots, M'_m$ of $M' \setminus (T_1 \cup \dots \cup T_m)$ whose quotient spaces are compact orbifolds $O_1, \dots, O_m$ (possibly with boundary) whose singularities are of cone type.

We first analyze the 2–orbifolds $O_1, \dots, O_m$. By Lemma 2.9(c) and the fact that M is aspherical we conclude that each O_j is good, ie its interior is diffeomorphic to an

isometric quotient of S^2, $\mathbb{R}^2$ or $\mathbb{H}^2$ (observe that otherwise we would be able to cover M by two solid tori). By the same argument and the fact that every orbifold covering of O_j induces a covering of M_j, it follows that O_j can also not be a quotient of S^2. So each O_j is an isometric quotient of $\mathbb{R}^2$ or $\mathbb{H}^2$.

If Int O_j is diffeomorphic to an isometric quotient of $\mathbb{R}^2$, then there is a finite covering $\widehat{O}_j \to O_j$ such that $\widehat{O}_j$ is diffeomorphic to a torus or an annulus. Let $\widehat{M}'_j \to M'_j$ be the induced covering. In the first case $m = j = 1$ and $\widehat{M} = \widehat{M}' = \widehat{M}'_j$ carries an S^1–fibration over T^2. Since T^2 fibers over a circle, this would however imply that $\widehat{M}$ fibers over a circle with T^2–fibers, in contradiction to our assumptions. So $\widehat{O}_j$ is diffeomorphic to an annulus and $\widehat{M}'_j \approx T^2 \times I$. We mention the following fact, which we will use later in the proof: For every natural number $N \geq 1$, the covering $\widehat{O}_j \to O_j$ can be chosen so that its restriction to every boundary component of $\widehat{O}_j$ is an N–fold covering over a circle. We can moreover pass to a covering $\widehat{M}_j \to \widehat{M}'_j$, $\widehat{M}_j \approx T^2 \times I$ such that the composition $\widehat{M}_j \to \widehat{M}'_j \to M'_j$ over each boundary torus of M'_j is an N^2–fold covering of $n_j := 1$ or $n_j := 2$ tori over a torus that is induced by the sublattice $N\mathbb{Z}^2 \subset \mathbb{Z}^2$.

If Int O_j is diffeomorphic to an isometric quotient of $\mathbb{H}^2$, then, by an argument from [4, Lemma 4.1], for every large enough $N \geq 2$ we can find a finite orbifold covering $\widehat{O}_j \to O_j$ such that $\widehat{O}_j$ is a manifold and such that the covering map restricted to each boundary component of $\widehat{O}_j$ is an N–fold covering of the circle. Consider the induced covering $\widehat{M}'_j \to M'_j$, where $\widehat{M}'_j$ is an S^1–bundle over Int $\widehat{O}_j$. If $\widehat{O}_j$ is closed, then we are in case (B) of the lemma, so assume in the following that none of the $\widehat{O}_j$ is closed. The S^1–fibration on each $\widehat{M}'_j$ can then be trivialized, ie $\widehat{M}'_j = \widehat{O}_j \times S^1$. We can hence pass to a further N–fold covering $\widehat{M}_j \to \widehat{M}'_j$ using an N–fold covering of the S^1–factor. Then for some $n_j \geq 1$ the composition $\widehat{M}_j \to \widehat{M}'_j \to M'_j$ over each boundary torus of M'_j is the disjoint union of n_j many N^2–fold coverings over the torus, induced by a sublattice $N\mathbb{Z}^2 \subset \mathbb{Z}^2$.

Now choose N large enough that the construction of the last two paragraphs can be carried out for every $j = 1, \dots, m$. Observe that the coverings over every T_i coming from the coverings over the two adjacent M'_j consist of equivalent pieces. Let $N_0 = n_1 \cdots n_k$ and consider N_0/n_j many disjoint copies of $\widehat{M}_j$ for each $j = 1, \dots, k$. Then these copies can be glued together along their boundary to obtain a covering $\widehat{M}' \to M'$. The Seifert decomposition on M' induces a Seifert decomposition $\widehat{T}'_1, \dots, \widehat{T}'_{m'}$ of $\widehat{M}'$ all of whose pieces are products.

We are now almost done. As a final step we successively remove tori $\widehat{T}_i'$ that are adjacent to Seifert components $\approx T^2 \times (0,1)$. Since $\widehat{M}$ cannot be a T^2–bundle over a circle, these Seifert components can never be adjacent to such a torus $\widehat{T}_i'$ from both sides. At the end of this process, we are either left with a single piece $\approx T^2 \times I$ and we are in case (A) of the lemma or none of the Seifert pieces are diffeomorphic to $T^2 \times I$. In the latter case we also remove tori $\widehat{T}_i'$ for which the S^1–fibers coming from either side are homotopic to one another. This will either result in two distinct Seifert components getting joined together or in identifying two boundary tori of a single Seifert component. If at any point in this process the new Seifert component is closed, then we undo the last step and we are in case (B). Otherwise, we are in case (C). □

We will now show that Proposition 3.2(a) is implied by the following proposition.

Proposition 3.4 *Let M_0 be an arbitrary 3–manifold with $\pi_2(M_0) = 0$ and $M \subset M_0$ be an embedded, connected, orientable, compact 3–manifold with incompressible toroidal boundary components such that the fundamental group of M injects into the fundamental group of M_0.*

Assume that M satisfies one of the following conditions:

(A) $M \approx T^2 \times I$.

(B) *M is the total space of an S^1–bundle over a closed, orientable surface Σ with $\chi(\Sigma) < 0$.*

(C) *M admits a Seifert decomposition $T_1, \dots, T_m \subset M$ such that the components of $\operatorname{Int} M \setminus (T_1 \cup \dots \cup T_m)$ are diffeomorphic to the interiors of $M_j = \Sigma_j \times S^1$ for $j = 1, \dots, k$, where each Σ_j is a compact orientable surface with at least one boundary component and $\chi(\Sigma_j) < 0$. The diffeomorphisms can be chosen in such a way that they can be smoothly extended to the boundary tori. Moreover, at each T_i the fibers of the S^1–fibrations induced from the manifold M_j on either side are not homotopic to one another.*

Then there is a constant $C < \infty$, a simplicial complex V and a continuous map

$$f_0 \colon V \to M \quad \text{with } f_0(\partial V) \subset \partial M$$

that is a smooth immersion on ∂V such that the following holds:

Let $S \subset \operatorname{Int} M$, $S \approx S^1 \times D^2$, be an embedded solid torus whose fundamental group injects into the fundamental group of M (ie S is incompressible in M). Let, moreover, $f \colon V \to M_0$ be a piecewise smooth map that is homotopic to f_0 relative to ∂V in M_0

and g a Riemannian metric on M_0. Then $f(V) \cap S \neq \varnothing$ and we can find a compact, smooth domain $\Sigma \subset \mathbb{R}^2$ and a smooth map $h\colon \Sigma \to S$ such that $h(\partial\Sigma) \subset \partial S$ and such that h restricted to the interior boundary circles of Σ of is contractible in ∂S and h restricted to the exterior boundary circle of Σ is noncontractible in ∂S, and such that

$$\operatorname{area} h < C \operatorname{area} f.$$

Proof that Proposition 3.4 implies Proposition 3.2(a) Let $M = M_{\mathrm{hyp}} \cup M_{\mathrm{Seif}}$ be a closed, orientable, irreducible manifold as defined in Section 3.1 and M' a component of M_{Seif}. By van Kampen's theorem the fundamental group of M' injects into that of M. Consider now the finite covering $\widehat{\pi}'\colon \widehat{M}' \to M'$ from Lemma 3.3. Choose $p \in \widehat{M}'$ and consider the push forward $\widehat{\pi}'_*(\pi_1(\widehat{M}', p))$ inside $\pi_1(M, \widehat{\pi}'_*(p))$. This subgroup induces a covering $\widehat{\pi}\colon \widehat{M} \to M$, which can be viewed as an extension of $\widehat{\pi}'\colon \widehat{M}' \to M'$. Still, the fundamental group of $\widehat{M}'$ injects into that of $\widehat{M}$.

The cases (A)–(C) of Lemma 3.3 for $\widehat{M}'$ correspond to the conditions (A)–(C) in Proposition 3.4. So we can apply Proposition 3.4 for $M \leftarrow \widehat{M}'$ and $M_0 \leftarrow \widehat{M}$, and obtain a simplicial complex V and a map $\hat{f}_0\colon V \to \widehat{M}'$ (observe that $\pi_2(\widehat{M}) = \pi_2(M) = 0$ by Proposition 2.3 and by the fact that M is irreducible). Set $f_0 = \widehat{\pi} \circ \hat{f}_0\colon V \to M$. Then we can lift any homotopy between f_0 and a map $f\colon V \to M$ to a homotopy between $\hat{f}_0$ and $\hat{f}\colon V \to \widehat{M}$ such that $f = \widehat{\pi} \circ \hat{f}$. Consider now an incompressible solid torus $S \subset M'$ and choose a component $\widehat{S} \subset \widehat{\pi}^{-1}(S) \cap \widehat{M}'$. Since $\widehat{\pi}'$ is a finite covering, we find that $\widehat{S}$ is a solid torus as well, which is incompressible in $\widehat{M}'$. So Proposition 3.4 provides a compact, smooth domain $\Sigma \subset \mathbb{R}^2$ and a map $\hat{h}\colon \Sigma \to \widehat{M}$ such that $\hat{h}$ restricted to the exterior boundary circle of Σ is noncontractible in $\partial\widehat{S}$, but $\hat{h}$ restricted to the other boundary circles is contractible in $\partial\widehat{S}$. Therefore, $h = \widehat{\pi} \circ \hat{h}$ has the desired topological properties and we have

$$\operatorname{area} h = \operatorname{area} \hat{h} < C \operatorname{area} \hat{f} = C \operatorname{area} f.$$

This finishes the proof. □

In the following four subsections, we will frequently refer to the conditions (A)–(C). We first finish off the case in which M satisfies condition (A).

Proposition 3.5 *Proposition 3.4 holds if M satisfies condition* (A).

Proof Observe that $M \approx T^2 \times I \approx S^1 \times S^1 \times I$. Denote by A_1 and A_2 the two embedded annuli of the form

$$\{\mathrm{pt}\} \times S^1 \times I, \quad S^1 \times \{\mathrm{pt}\} \times I \subset M.$$

Let V be their disjoint union and $f_0\colon V \to M$ be the map that restricts to the inclusion on each component of V. Moreover, fix a triangulation on V to give it the structure of a simplicial complex.

Note that every noncontractible loop $\sigma \subset \operatorname{Int} M$ has nonzero intersection number with one of the maps $f_0|_{A_1}$ or $f_0|_{A_2}$. Consider now the solid torus $S \subset M$ and let $\sigma \subset \operatorname{Int} S$ be a noncontractible loop inside S (and hence also inside M). Choose $i \in \{1,2\}$ such that $f_0|_{A_i}$ has nonzero intersection number with σ. Then so does $f|_{A_i}$. Let $f'\colon A_i \to M$ be a small perturbation of $f|_{A_i}$ that is transversal to ∂S and for which area $f' < 2$ area f. Still, f' has nonzero intersection number with σ.

Denote the components of $f'^{-1}(S)$ by $Q_1, \dots, Q_p \subset A_i \approx S^1 \times I$. The sum of the intersection numbers of $f'|_{Q_j}$ with σ is nonzero. Moreover, by the choice of i none of these components Q_j can contain a circle that is noncontractible in A_i. So each Q_j is contained in a closed disk $Q'_j \subset A_i$ with $\partial Q'_j \subset \partial Q_j$, which arises from filling in all its interior boundary circles. Note that any two such disks, Q'_{j_1}, Q'_{j_2} are either disjoint or one is contained in the other. By a maximality argument, we can choose $j \in \{1, \dots, p\}$ such that the intersection number of $f'|_{Q_j}$ with σ is nonzero, but such that Q'_j does not contain any other $Q_{j'}$ with the same property. Then f' has to have zero intersection number with σ on every component of $Q'_j \setminus Q_j$. Hence, f' restricted to every circle of $\partial Q_j \setminus \partial Q'_j$ is contractible in ∂S and f' restricted to $\partial Q'_j$ is noncontractible. So if we choose $\Sigma = Q_j \subset Q'_j \approx D^2 \subset \mathbb{R}^2$ and $h = f'|_{Q_j}$, then the desired properties are fulfilled and area $h <$ area $f' < 2$ area f. □

It remains to prove Proposition 3.4 in the cases in which M satisfies condition (B) or (C). Its proof in these two cases will be carried out in Section 3.6. The proof makes use of a simplicial complex V, which will be constructed and analyzed in the following subsection and relies on a certain combinatorial convexity estimate on V, which will be derived in Section 3.4 for case (C) and in Section 3.5 for case (B).

3.3 Combinatorial geometry of $\tilde{M}$ if M satisfies condition (B) or (C)

In this subsection we will set up the proof of Proposition 3.4. In particular, we will construct the simplicial complex V and introduce the tools that will be needed in the following two subsections.

Assume that M satisfies condition (B) or (C) in Proposition 3.4, ie M is a compact, connected, orientable 3–manifold with incompressible toroidal boundary components.

If M satisfies condition (C), we fix the Seifert decomposition $T_1, \dots, T_m$ of M as well as the identifications of the components of $\operatorname{Int} M \setminus (T_1 \cup \dots \cup T_m)$ with the interiors of the products $M_j \approx \Sigma_j \times S^1$ for $j = 1, \dots, k$. Here $\Sigma_1, \dots, \Sigma_m$ are compact surfaces with at least one boundary component and negative Euler characteristic. If M satisfies condition (B), then we set $m = k = 1$ and we can find a torus $T_1 \subset M$ such that $M \setminus T_1$ is diffeomorphic to the interior of the product $\Sigma_1 \times S^1$, where Σ_1 is a compact, orientable surface with two boundary circles, which can be obtained from Σ by cutting along a nonseparating, embedded loop. Moreover, $\chi(\Sigma_1) = \chi(\Sigma) < 0$. In either case, we assume that the diffeomorphisms that identify the interior of each M_j with the corresponding component of $\operatorname{Int} M \setminus (T_1 \cup \dots \cup T_m)$ can be continued smoothly up to the boundary tori. If M satisfies condition (C), then the fibrations coming from either side of each torus T_i are assumed to be nonhomotopic to one another and in case (B) we assume that the fibration on M_1 has been chosen such that both fibrations agree.

We will mainly be working in the universal covering $\widetilde{M}$ of M. Let $\pi\colon \widetilde{M} \to M$ be the covering projection.

Definition 3.6 (chambers) The closures $K \subset \widetilde{M}$ of components of the preimages of components of $M \setminus (T_1 \cup \dots \cup T_m)$ under π are called *chambers* and the set of chambers is denoted by $\mathcal{K}$.

Definition 3.7 (walls) The components W of $\partial \widetilde{M}$ and of the preimages $\pi^{-1}(T_i)$ for $i = 1, \dots, m$ are called *walls* and the set of walls is denoted by $\mathcal{W}$. We say that two distinct chambers $K_1, K_2 \in \mathcal{K}$ are *adjacent* if they share a common wall.

By van Kampen's theorem every chamber $K \in \mathcal{K}$ can be viewed as the universal cover of M_{j_K} for a unique $j_K \in \{1, \dots, k\}$. So $K \approx \widetilde{\Sigma}_{j_K} \times \mathbb{R}$. The boundary of K is a disjoint union of walls that cover exactly the tori T_i and the boundary tori of M that are adjacent to M_{j_K}, and these tori stand in one-to-one correspondence with the boundary circles of Σ_{j_K}. Moreover, every wall is diffeomorphic to $\mathbb{R}^2$. For later purposes, we will replace the j–index by K and write for example $M_K = M_{j_K}$ and $\Sigma_K = \Sigma_{j_K}$. Note that the interior of every chamber is disjoint from the union of all walls. So the complement of the union of all walls in $\widetilde{M}$ is equal to the union of the interiors of all chambers.

Lemma 3.8 *Every wall $W \in \mathcal{W}$ with $W \not\subset \partial \widetilde{M}$ separates $\widetilde{M}$ into two components. So every two distinct chambers $K_1, K_2 \in \mathcal{K}$ can only intersect in at most one wall $W = K_1 \cap K_2$ and the adjacency graph of $\mathcal{K}$ is a tree.*

Proof If $W \in \mathcal{W}$ did not separate $\widetilde{M}$, then we could find a loop $\gamma \subset \widetilde{M}$ that intersects W transversally and exactly once, ie its intersection number with W is 1. However $\gamma \subset \widetilde{M}$ must be contractible. □

On each torus T_i and boundary torus of M we fix an affine structure and a point $e_i \in T_i$ for the remainder of this subsection. These affine structures induce an affine structure on all walls $W \in \mathcal{W}$. We can assume that the product structures on each $M_j \approx \Sigma_j \times S^1$ are chosen so that the circle fibers on each boundary component of M_j coming from the S^1–factor and the boundary circle of the Σ_j are geodesic circles in the corresponding torus T_i.

Now, for each $j = 1, \dots, k$ we choose an embedded section $S_j \subset M_j \approx \Sigma_j \times S^1$ of the form $\Sigma_j \times \{\mathrm{pt}\}$. Next, we choose embedded and pairwise disjoint arcs inside each Σ_j, whose endpoints lie in the boundary of Σ_j and that cut the interior of Σ_j into a topological ball. Denote their union by $C_j^* \subset \Sigma_j$ and set $C_j = C_j^* \times S^1$. Now let

$$V = T_1 \cup \dots \cup T_m \cup S_1 \cup \dots \cup S_k \cup C_1 \cup \dots \cup C_k.$$

The complement $\operatorname{Int} M \setminus V$ is a disjoint union of k topological balls $\approx (\Sigma_1 \setminus C_1^*) \times (0,1)$, $\dots, (\Sigma_k \setminus C_k^*) \times (0,1)$.

By construction, V can be viewed as an embedded, finite and pure 2–dimensional polygonal complex with $\partial V \subset \partial M$. The notion of a "polygonal complex" generalizes the notion of a "simplicial complex" from Definition 3.1, by allowing the faces to be polygons instead of 2–simplices. The 1–skeleton $V^{(1)}$ of V, viewed as a polygonal complex, is the union of ∂S_j, ∂C_j and $C_j \cap S_j$ for all $j = 1, \dots, k$. And the set of vertices $V^{(0)}$ of V is contained in $T_1 \cup \dots \cup T_m \cup \partial M$. By subdividing the polygonal faces of V into triangles, we can give V the structure of a finite and pure simplicial complex. In doing so, we increase the number of edges and faces of V, but the topology of V still remains the same. In the following, we will fix this simplicial structure, and the structure of V as a polygonal complex will not be essential for us anymore.

Consider now the universal covering $\pi\colon \widetilde{M} \to M$ and set $\widetilde{V} = \pi^{-1}(V) \subset \widetilde{M}$. Then $\widetilde{V}$ inherits the structure of an infinite simplicial complex with $\partial \widetilde{V} \subset \partial \widetilde{M}$ and the components of $\operatorname{Int} \widetilde{M} \setminus \widetilde{V}$ are topological balls on which π is injective. Their boundaries are diffeomorphic to simplicial 2–spheres.

Definition 3.9 (cells) The closure Q of any component of $\widetilde{M} \setminus \widetilde{V}$ is called a *cell* and the set of cells is denoted by $\mathcal{Q}$. Two distinct cells are called *adjacent* if their intersection contains a point of $\widetilde{V} \setminus \widetilde{V}^{(1)}$. Here $\widetilde{V}^{(1)}$ denotes the 1–skeleton of $\widetilde{V}$,

viewed as a simplicial complex. (Note that the notion of adjacent cells would be the same if $\widetilde{V}^{(1)}$ denoted the 1–skeleton of $\widetilde{V}$ with respect to the inherited polygonal structure on $\widetilde{V}$ as opposed to the simplicial structure.)

So every chamber $K \in \mathcal{K}$ is equal to the union of cells $Q \subset K$. Identify K with $\widetilde{\Sigma}_K \times \mathbb{R}$ as before and set $C_K = C_{j_K}$ and $C_K^* = C_{j_K}^*$. The structure of $\widetilde{V}$ in K can then be understood as follows: Let $\widetilde{C}_K^*$ be the preimage of C_K^* under the universal covering map $\widetilde{\Sigma}_K \to \Sigma_K$. Then $\widetilde{V} \cap K$ is equal to the union of $\pi^{-1}(C_K) \cap K \approx \widetilde{C}_K^* \times \mathbb{R}$ with $\pi^{-1}(S_K) \cap K$ and ∂K. So the arrangement of the cells $Q \subset K$ is reflected by the identity

$$\bigcup_{Q \in \mathcal{Q},\, Q \subset K} \operatorname{Int} Q = (\operatorname{Int} \widetilde{\Sigma}_K \setminus \widetilde{C}_K^*) \times (\mathbb{R} \setminus \mathbb{Z}). \tag{3-1}$$

We will always refer to the first factor in this cartesian product as the *horizontal* direction and to the second factor as the *vertical* direction. In the next definition we group cells that share the same horizontal coordinates.

Definition 3.10 (columns) Consider a chamber $K \in \mathcal{K}$ and choose the identification $K \cong \widetilde{\Sigma}_K \times \mathbb{R}$ as in the last paragraph. Then the closure E of each component of $(\operatorname{Int} \widetilde{\Sigma}_K \setminus \widetilde{C}_K^*) \times \mathbb{R}$ is called a *column*. The set of columns of K is denoted by $\mathcal{E}_K$.

We say that two distinct columns $E_1, E_2 \in \mathcal{E}_K$ are *adjacent* if they intersect. An ordered tuple $(E_0, \dots, E_n)$ of columns for which E_i is adjacent to E_{i+1} is called a *chain between E_0 and E_n* and n is called its *length*. It is called *minimal* if its length is minimal amongst all chains between the same columns.

So each chamber $K \in \mathcal{K}$ is equal to the union of all its columns $E \in \mathcal{E}_K$ and every such column E consists of cells $Q \subset E$, which are arranged in a linear manner. Next, we define distance functions with respect to the horizontal and vertical direction in (3-1).

Definition 3.11 (horizontal and vertical distance within a chamber) Let $K \in \mathcal{K}$ be a chamber and $E_1, E_2 \in \mathcal{E}_K$ two columns. We define their *horizontal distance* $\operatorname{dist}_K^H(E_1, E_2)$ *(within K)* to be the minimal length of a chain between E_1 and E_2. For two cells $Q_1, Q_2 \subset K$ with $Q_1 \subset E_1$ and $Q_2 \subset E_2$ we define the *horizontal distance* $\operatorname{dist}_K^H(Q_1, Q_2) = \operatorname{dist}_K^H(E_1, E_2)$ (*within K*). We say that Q_1 and Q_2 are *vertically aligned* (*within K*) if $\operatorname{dist}_K^H(Q_1, Q_2) = 0$, ie if Q_1 and Q_2 lie in the same column.

For two cells $Q_1, Q_2 \subset K$ we define the *vertical distance* $\operatorname{dist}_K^V(Q_1, Q_2)$ (*within K*) by the minimal number of times that an arc $\gamma\colon [0,1] \to K$ with $\gamma(0) \in \operatorname{Int} Q_1$ and

$\gamma(1) \in \operatorname{Int} Q_2$ intersects $\pi^{-1}(S_K)$, ie the number of integers between the second coordinates of both cells in (3-1). We say that Q_1 and Q_2 are *horizontally aligned* (*within* K) if $\operatorname{dist}_K^V(Q_1, Q_2) = 0$.

Obviously, both distance functions satisfy the triangle inequality. Two cells $Q_1, Q_2 \subset K$ are adjacent if and only if $\operatorname{dist}_K^H(Q_1, Q_2) + \operatorname{dist}_K^V(Q_1, Q_2) = 1$. And they are disjoint if and only if this sum is ≥ 2 and not both summands are equal to 1.

Lemma 3.12 *Assume that M satisfies condition* (B) *or* (C). *Consider a chamber $K \in \mathcal{K}$. Then the set of columns $\mathcal{E}_K$ together with the adjacency relation describes a tree with constant valency ≥ 4. So, for every two columns $E_1, E_2 \in \mathcal{E}_K$, there is a unique minimal chain between E_1 and E_2, and a chain between E_1 and E_2 is the minimal one if and only if it contains each column not more than once. Moreover, for every three columns $E_1, E_2, E_3 \in \mathcal{E}_K$ there is a unique column $E^* \in \mathcal{E}_K$ that lies on all three minimizing chains between every pair of E_1, E_2 and E_3.*

Finally, for every two columns $E_1, E_2 \in \mathcal{E}_K$ with $\operatorname{dist}_K^H(E_1, E_2) \geq 2$, there is at most one wall $W \in \mathcal{W}$ that is adjacent to both E_1 and E_2.

Proof By an intersection number argument as in the proof of Lemma 3.8, we find that a loop in $\widetilde{\Sigma}_K$ cannot cross a component of $\widetilde{C}_K^* \subset \widetilde{\Sigma}_K$ exactly once. This establishes the tree property.

Now assume that there are two distinct boundary components $B_1, B_2 \subset \partial\widetilde{\Sigma}_K$ that are adjacent to two distinct components $U_1, U_2 \subset \widetilde{\Sigma}_K \setminus \widetilde{C}_K^*$ at the same time. Since $\widetilde{\Sigma}_K$ is simply connected, the closure of the set $B_1 \cup B_2 \cup U_1 \cup U_2$ separates $\widetilde{\Sigma}_K$ into two open components A_1 and A_2, one of which, say A_1, has compact closure. So A_1 only contains finitely many components of $\widetilde{\Sigma}_K \setminus \widetilde{C}_K^*$ and all these components are only adjacent to each other or to U_1 or U_2. This however contradicts the tree property. □

In the following we want to understand the adjacency structure of $\mathcal{Q}$ on $\widetilde{M}$. As a first step we analyze its structure near walls.

Lemma 3.13 *There is a constant $C_0 < \infty$ such that the following holds:*

Let $W \in \mathcal{W}$ with $W \not\subset \partial M$ be a wall and let $K, K' \in \mathcal{K}$ be the chambers that are adjacent to W from either side. Then the columns $E \in \mathcal{E}_K$ and $E' \in \mathcal{E}_{K'}$ intersect W in affine strips $E \cap W$ and $E' \cap W$ (ie domains bounded by two parallel straight lines). In the case in which M satisfies condition (B), *these strips are all parallel and if M*

satisfies condition (C), *each pair of strips coming from* K *and* K' *are not parallel to one another; so they intersect in a nonempty compact set.*

We furthermore have the following estimates between the horizontal and vertical distance functions in K *and* K'*:*

(a) *Assume that* M *satisfies condition* (B) *or* (C) *and let* $Q_1, Q_2 \subset K$ *be cells that are adjacent to a common cell* $Q' \subset K'$*. Then*

$$\mathrm{dist}_K^H(Q_1, Q_2),\ \mathrm{dist}_K^V(Q_1, Q_2) < C_0.$$

(b) *Assume that* M *satisfies condition* (C) *and let* $Q_1, Q_2 \subset K$ *and* $Q_1', Q_2' \subset K'$ *be cells such that* Q_1, Q_1' *and* Q_2, Q_2' *are adjacent pairs and such that* Q_1' *and* Q_2' *are vertically aligned. Then*

$$\mathrm{dist}_K^V(Q_1, Q_2),\ \mathrm{dist}_{K'}^V(Q_1', Q_2') < C_0\,\mathrm{dist}_K^H(Q_1, Q_2) + C_0.$$

(c) *Assume that* M *satisfies condition* (C) *and consider cells* $Q_1, Q_2, Q_3, Q_4 \subset K$*. Assume that* Q_1, Q_2 *and* Q_3, Q_4 *are vertically aligned pairs and assume that there are columns* $E_1', E_2' \in \mathcal{E}_{K'}$ *such that* Q_1 *and* Q_4 *are adjacent to some cells in* E_1' *and* Q_2 *and* Q_3 *are adjacent to some cells in* E_2'*. Then*

$$|\mathrm{dist}_K^V(Q_1, Q_2) - \mathrm{dist}_K^V(Q_3, Q_4)| < C_0.$$

(d) *Assume that* M *satisfies condition* (C) *and consider cells* $Q_1, Q_2 \subset K$ *and* $Q_1', Q_2' \subset K'$ *such that* Q_1, Q_1' *and* Q_2, Q_2' *are adjacent pairs and that* $\mathrm{dist}_K^H(Q_1, Q_2),\ \mathrm{dist}_{K'}^H(Q_1', Q_2') \le 3$*. Then*

$$\mathrm{dist}_K^V(Q_1, Q_2),\ \mathrm{dist}_{K'}^V(Q_1', Q_2') < C_0.$$

(e) *Assume that* M *satisfies condition* (B) *and consider cells* $Q_1, Q_2 \subset K$ *and* $Q_1', Q_2' \subset K'$ *such that* Q_1, Q_1' *and* Q_2, Q_2' *are adjacent pairs. Then*

$$\mathrm{dist}_K^V(Q_1, Q_2) < \mathrm{dist}_{K'}^V(Q_1', Q_2') + C_0\,\mathrm{dist}_K^H(Q_1, Q_2) + C_0.$$

(f) *Assume that* M *satisfies condition* (B) *or* (C) *and consider cells* $Q_1, Q_2 \subset K$ *and* $Q_1', Q_2' \subset K'$ *such that* Q_1, Q_1' *and* Q_2, Q_2' *are adjacent pairs. Then*

$$\begin{aligned}\mathrm{dist}_K^H(Q_1, Q_2),\ &\mathrm{dist}_K^V(Q_1, Q_2)\\ &< C_0\,\mathrm{dist}_{K'}^H(Q_1', Q_2') + C_0\,\mathrm{dist}_{K'}^V(Q_1', Q_2') + C_0.\end{aligned}$$

(g) *Assume that* M *satisfies condition* (B) *or* (C) *and consider cells* $Q_1, Q_2, Q_3, Q_4 \subset K$ *and* $Q_1', Q_2', Q_3', Q_4' \subset K'$ *such that* Q_i *and* Q_i' *are adjacent for all*

$i = 1, \dots, 4$. Assume moreover that $\mathrm{dist}_K^H(Q_1, Q_2) = \mathrm{dist}_K^H(Q_3, Q_4)$ *and* $\mathrm{dist}_K^V(Q_1, Q_2) = \mathrm{dist}_K^V(Q_3, Q_4)$ *in an oriented sense, ie the cells Q_1, Q_2, Q_3 and Q_4 form a "parallelogram" along W. Then*

$$|\mathrm{dist}_{K'}^H(Q_1', Q_2') - \mathrm{dist}_{K'}^H(Q_3', Q_4')|,\ |\mathrm{dist}_{K'}^V(Q_1', Q_2') - \mathrm{dist}_{K'}^V(Q_3', Q_4')| < C_0.$$

Proof Note first that in all cases (a)–(g) the cells Q_i and Q_i' are adjacent to W (meaning that Q_i and Q_i' intersect W). The cells of K and K' that are adjacent to W are arranged in a doubly periodic pattern along W. So we can introduce euclidean coordinates $(x_1, x_2)\colon W \to \mathbb{R}^2$ such that for any two cells $Q_1, Q_2 \subset K$ that are adjacent to W and any points $p_1 \in Q_1 \cap W$ and $p_2 \in Q_2 \cap W$ we have $\big|\mathrm{dist}_K^H(Q_1, Q_2) - |x_1(p_1) - x_1(p_2)|\big| < C$ and $\big|\mathrm{dist}_K^V(Q_1, Q_2) - |x_2(p_1) - x_2(p_2)|\big| < C$ for some uniform constant C. Similarly, we can find Euclidean coordinates $(x_1', x_2')\colon W \to \mathbb{R}^2$ with the analogous behavior for the cells of K' that are adjacent to W such that the origins of (x_1, x_2) and (x_1', x_2') agree. The transformation matrix $A \in \mathbb{R}^{2\times 2}$ with $A(x_1, x_2) = (x_1', x_2')$ is invertible. In case (C) we have $A_{12} \neq 0$ and in case (B) we have $A_{12} = 0$, $A_{11} \neq 0$ and $A_{22} = 1$. All assertions of the lemma now follow from the corresponding statements for these two coordinate systems. □

Next, consider a smooth arc $\gamma\colon [0, 1] \to \widetilde{M}$.

Definition 3.14 (general position) We say that γ is in *general position* if its endpoints $\gamma(0), \gamma(1) \notin \widetilde{V}$ and if γ intersects $\widetilde{V}$ transversally and only in $\widetilde{V} \setminus \widetilde{V}^{(1)}$. If $Q_1, Q_2 \in \mathcal{Q}$ are two cells with $\gamma(0) \in Q_1$ and $\gamma(1) \in Q_2$, then we say that γ *connects* Q_1 *with* Q_2.

Let η, $H > 0$ be constants whose value will be determined later, in Section 3.6. In the course of the following three subsections, we will need to assume that η is small enough and H is large enough to make certain arguments work out.

Definition 3.15 (length and distance) The (*combinatorial*) *length* $|\gamma|$ of an arc $\gamma\colon [0, 1] \to \widetilde{M}$ in general position is defined as

$$\begin{aligned}|\gamma| = \eta(&\text{number of intersections of } \gamma \text{ with } \pi^{-1}(S_1 \cup \dots \cup S_k))\\ &+ H(\text{number of intersections of } \gamma \text{ with } \pi^{-1}(T_1 \cup \dots \cup T_m))\\ &+ (\text{number of intersections of } \gamma \text{ with } \pi^{-1}(C_1 \cup \dots \cup C_k)).\end{aligned}$$

The (*combinatorial*) *distance* $\mathrm{dist}(Q_1, Q_2)$ between two cells $Q_1, Q_2 \in \mathcal{Q}$ is the minimal combinatorial length of all arcs in general position between Q_1 and Q_2. An

arc $\gamma\colon [0,1] \to \widetilde{M}$ in general position is said to be (*combinatorially*) *minimizing* if its length is equal to the combinatorial distance between the two cells that contain its endpoints.

Observe that $(\mathcal{Q}, \mathrm{dist})$ is a metric space. On a side note, it is an interesting "coincidence" that this metric space approximates the conjectured geometric behavior of the Ricci flow metric $t^{-1} g_t$ lifted to the universal cover $\widetilde{M}$.

Our main characterization of combinatorially minimizing arcs will be stated in Proposition 3.20 in case (C) and in Proposition 3.26 in case (B). We conclude this subsection by pointing out three basic properties of combinatorially minimizing arcs.

Lemma 3.16 *If $\gamma\colon [0,1] \to \widetilde{M}$ is combinatorially minimizing, then so is every orientation-preserving or -reversing reparametrization and every subsegment of γ whose endpoints don't lie in $\widetilde{V}$.*

Proof Obvious. □

Lemma 3.17 *For any cell $Q \in \mathcal{Q}$, the preimage $\gamma^{-1}(Q)$ under a combinatorially minimizing arc $\gamma\colon [0,1] \to \widetilde{M}$ is a closed interval, that is, γ does not reenter Q after exiting it.*

Proof Otherwise we could replace γ by a shorter arc. □

Lemma 3.18 *Assume that $\gamma\colon [0,1] \to \widetilde{M}$ is combinatorially minimizing and stays within some chamber $K \in \mathcal{K}$. Let $E_0, \dots, E_n \in \mathcal{E}_K$ be the columns that γ intersects in that order. Then $(E_0, \dots, E_n)$ is a minimal chain in $\mathcal{E}_K$.*

Moreover, γ intersects each component of $\pi^{-1}(S_1 \cup \dots \cup S_k)$ at most once. So, if the endpoints of γ lie in cells $Q_1, Q_2 \in \mathcal{Q}$, then

$$|\gamma| = \mathrm{dist}(Q_1, Q_2) = \mathrm{dist}_K^H(Q_1, Q_2) + \eta\, \mathrm{dist}_K^V(Q_1, Q_2).$$

Finally, for any two cells $Q_1, Q_2 \subset K$ we have

$$\mathrm{dist}(Q_1, Q_2) \le \mathrm{dist}_K^H(Q_1, Q_2) + \eta\, \mathrm{dist}_K^V(Q_1, Q_2).$$

Proof This follows from the cell structure of K (see also (3-1)). □

3.4 A combinatorial convexity estimate if M satisfies condition (C)

In this subsection we assume that M satisfies condition (C) in Proposition 3.4. We will analyze the combinatorial distance function on $\mathcal{Q}$ in this case. The main result in this section will be the combinatorial convexity estimate in Proposition 3.24.

Lemma 3.19 *There are constants $\eta^* > 0$ and $H^* < \infty$ such that if $\eta \le \eta^*$ and $H \ge H^*$, the following holds:*

Consider two chambers $K, K' \in \mathcal{K}$ that are adjacent to a common wall $W = K \cap K'$ from either side and assume that $\gamma\colon [0,1] \to K \cup K'$ is combinatorially minimizing. Then γ intersects W at most twice.

(a) *If γ intersects W exactly once, then there is a unique column $E \in \mathcal{E}_K$ in K that is both adjacent to W and that intersects γ. The same is true in K'.*

(b) *If γ intersects W exactly twice and $\gamma(0), \gamma(1) \in K$, then there is a unique column $E' \in \mathcal{E}_{K'}$ such that γ is contained in $K \cup E'$. Moreover there are exactly two columns $E_1, E_2 \in \mathcal{E}_K$ that are adjacent to W and that intersect γ. And we have $\operatorname{dist}_K^H(E_1, E_2) > H$.*

(c) *If γ does not intersect W, but intersects two columns $E_1, E_2 \in \mathcal{E}_K$ that are both adjacent to W, then $\operatorname{dist}_K^H(E_1, E_2) < 3H$.*

Proof We first establish assertion (a). Assume without loss of generality that $\gamma(0) \in K$ and $\gamma(1) \in K'$. Let $Q \subset K$ be the last cell that γ intersects inside K and $Q' \subset K'$ the first cell in K'. So Q, Q' are adjacent. Let $E \in \mathcal{E}_K$ be the column that contains Q and E' the column that contains Q'. Assume that contrary to the assertion there is another column $E_1 \ne E \in \mathcal{E}_K$ that is adjacent to W and intersects γ. Choose a cell $Q_1 \subset E_1$ that intersects γ. Then, by Lemma 3.18,

$$\operatorname{dist}(Q_1, Q') = \operatorname{dist}(Q_1, Q) + H = \operatorname{dist}_K^H(E_1, E) + \eta \operatorname{dist}_K^V(Q_1, Q) + H.$$

Let $Q_2 \subset E_1$ be the cell that is horizontally aligned with Q. Then, by the triangle inequality and Lemma 3.18,

$$\operatorname{dist}(Q_2, Q') \le \operatorname{dist}(Q_2, Q) + \operatorname{dist}(Q, Q') \le \operatorname{dist}_K^H(E_1, E) + H.$$

On the other hand, again by the triangle inequality and Lemma 3.18,

$$\begin{aligned}\operatorname{dist}(Q_2, Q') + \eta \operatorname{dist}_K^V(Q_1, Q) &\ge \operatorname{dist}(Q_2, Q') + \operatorname{dist}(Q_1, Q_2)\\ &\ge \operatorname{dist}(Q_1, Q')\\ &= \operatorname{dist}_K^H(E_1, E) + \eta \operatorname{dist}_K^V(Q_1, Q) + H.\end{aligned}$$

It follows that

$$\operatorname{dist}(Q_2, Q') = \operatorname{dist}_K^H(E_1, E) + H.$$

Since $E_1 \cap W$ and $E' \cap W$ are nonparallel strips in W, we can find cells $Q_3 \subset E_1$ and $Q_3' \subset E'$ that are adjacent to each other and, by Lemma 3.13(b), we can estimate

$$\operatorname{dist}_K^V(Q_2, Q_3) = \operatorname{dist}_K^V(Q, Q_3),\ \operatorname{dist}_{K'}^V(Q', Q_3') < C_0 \operatorname{dist}_K^H(E_1, E) + C_0.$$

We then conclude

$$\begin{aligned}\operatorname{dist}_K^H(E_1, E) + H = \operatorname{dist}(Q_2, Q') &\le \operatorname{dist}(Q_2, Q_3) + \operatorname{dist}(Q_3, Q_3') + \operatorname{dist}(Q_3', Q') \\ &< 2\eta(C_0 \operatorname{dist}_K^H(E_1, E) + C_0) + H.\end{aligned}$$

For $\eta < (4C_0)^{-1}$ this implies $\operatorname{dist}_K^H(E_1, E) < 1$ and hence $E_1 = E$.

Next, we show assertion (b). Let $E_1, E_2 \in \mathcal{E}_K$ be the columns that γ intersects right before intersecting W for the first time and right after intersecting W for the second time. Let $E_1', E_2' \in \mathcal{E}_{K'}$ be the first and last columns that γ intersects inside K'. Assertion (a) applied to the subsegments of γ between $\gamma(0)$ and E_2' and between E_1' and $\gamma(1)$ yields that $E' := E_1' = E_2'$. Since the subsegment of γ that is contained in K' has both of its endpoints in E', it has to be fully contained in it. Moreover, assertion (a) implies that there are no other columns than E_1 and E_2 in K that are adjacent to W and intersect γ.

It remains to show the inequality on the horizontal distance between E_1 and E_2. We will do this by comparing the intrinsic and extrinsic distance between these two columns. Choose $Q_1 \subset E_1$ and $Q_1' \subset E'$ such that γ crosses W between Q_1 and Q_1' for the first time and pick $Q_2 \subset E_2$ and $Q_2' \subset E'$ accordingly. So Q_1, Q_1' and Q_2, Q_2' are adjacent pairs. Lemma 3.13(b) provides the bound $\operatorname{dist}_K^V(Q_1, Q_2) < C_0 \operatorname{dist}_K^H(Q_1, Q_2) + C_0$. So

$$\begin{aligned}2H \le \operatorname{dist}(Q_1, Q_2) &\le \operatorname{dist}_K^H(Q_1, Q_2) + \eta \operatorname{dist}_K^V(Q_1, Q_2) \\ &< (1 + \eta C_0) \operatorname{dist}_K^H(Q_1, Q_2) + \eta C_0.\end{aligned}$$

The desired inequality follows for $\eta < (2C_0)^{-1}$ and $H > 2$. This finishes the proof of assertion (b).

We can now show that γ intersects W at most twice. Assume not. After passing to a subsegment and possibly reversing the orientation, we may assume that γ intersects W exactly three times and that $\gamma(0) \in K$, $\gamma(1) \in K'$. By assertion (b) applied to subsegments of γ which intersects W exactly twice, we find that there are columns $E_1, E_3 \in \mathcal{E}_K$ and $E_2', E_4' \in \mathcal{E}_{K'}$, which are all adjacent to W, such that γ crosses W first between E_1 and E_2', then between E_2' and E_3 and finally between E_3 and E_4'. Choose cells $Q_1 \subset E_1$, $Q_1', Q_2' \subset E_2'$, $Q_2, Q_3 \subset E_3$ and $Q_3' \subset E_4'$ such that γ

crosses W first between Q_1 and Q_1', then between Q_2' and Q_2 and finally between Q_3 and Q_3'. So

$$\begin{aligned}
&\operatorname{dist}(Q_1, Q_3') \\
&\quad = \operatorname{dist}(Q_1, Q_1') + \operatorname{dist}(Q_1', Q_2') + \operatorname{dist}(Q_2', Q_2) + \operatorname{dist}(Q_2, Q_3) + \operatorname{dist}(Q_3, Q_3') \\
&\quad = 3H + \eta \operatorname{dist}^V_{K'}(Q_1', Q_2') + \eta \operatorname{dist}^V_K(Q_2, Q_3).
\end{aligned}$$

Since $E_1 \cap W$ and $E_4' \cap W$ are nonparallel strips in W, we can find cells $Q^* \subset E_1$ and $Q^{*\prime} \subset E_4'$ that are adjacent to each other. By Lemma 3.13(c),

$$\operatorname{dist}^V_K(Q_1, Q^*) < \operatorname{dist}^V_K(Q_2, Q_3) + C_0, \quad \operatorname{dist}^V_{K'}(Q^{*\prime}, Q_3') < \operatorname{dist}^V_{K'}(Q_1', Q_2') + C_0.$$

So

$$\begin{aligned}
\operatorname{dist}(Q_1, Q_3') &\le \operatorname{dist}(Q_1, Q^*) + \operatorname{dist}(Q^*, Q^{*\prime}) + \operatorname{dist}(Q^{*\prime}, Q_3') \\
&< \eta \operatorname{dist}^V_K(Q_2, Q_3) + \eta C_0 + H + \eta \operatorname{dist}^V_{K'}(Q_1', Q_2') + \eta C_0 \\
&= \operatorname{dist}(Q_1, Q_3') - 2H + 2\eta C_0.
\end{aligned}$$

We obtain a contradiction for $\eta C_0 < H$.

Finally, we show assertion (c). Assume now that γ does not intersect W and choose cells $Q_1 \subset E_1$ and $Q_2 \subset E_2$ that intersect γ. Since γ stays within K we have

$$\operatorname{dist}(Q_1, Q_2) = \operatorname{dist}^H_K(E_1, E_2) + \eta \operatorname{dist}^V_K(Q_1, Q_2).$$

Let $Q_3 \subset E_2$ be the cell that is horizontally aligned with Q_1. By the triangle inequality and the previous equation,

$$\operatorname{dist}(Q_1, Q_3) = \operatorname{dist}^H_K(E_1, E_2).$$

Let $Q_1' \subset K'$ be a cell that is adjacent to Q_1 and let $E' \in \mathcal{E}_{K'}$ be the column that contains Q_1'. Since $E' \cap W$ and $E_2 \cap W$ are nonparallel strips, we can find cells $Q_2' \subset E'$ and $Q_2'' \subset E_2$ that are adjacent. By Lemma 3.13(b), we have $\operatorname{dist}^V_K(Q_2'', Q_3), \operatorname{dist}^V_{K'}(Q_1', Q_2') < C_0 \operatorname{dist}^H_K(E_1, E_2) + C_0$. So

$$\begin{aligned}
&\operatorname{dist}^H_K(E_1, E_2) \\
&\quad = \operatorname{dist}(Q_1, Q_3) \le \operatorname{dist}(Q_1, Q_1') + \operatorname{dist}(Q_1', Q_2') + \operatorname{dist}(Q_2', Q_2'') + \operatorname{dist}(Q_2'', Q_3) \\
&\quad < 2H + 2\eta C_0 \operatorname{dist}^H_K(E_1, E_2) + 2\eta C_0.
\end{aligned}$$

The desired inequality follows for $2\eta C_0 < \frac{1}{10}$ and $H > 1$. □

The next proposition provides an accurate characterization of the behavior of a minimizing arc.

Proposition 3.20 *Assume that M satisfies condition* (C). *There are constants $\eta^* > 0$ and $H^* < \infty$ such that if $\eta \le \eta^*$ and $H \ge H^*$, the following holds:*

Consider a combinatorially minimizing arc $\gamma\colon [0,1] \to \widetilde{M}$. Then:

(a) *For every chamber $K \in \mathcal{K}$ and every column $E \in \mathcal{E}_K$, the preimage $\gamma^{-1}(E)$ is a connected interval, ie γ does not exit and reenter E.*

(b) *γ intersects every wall $W \in \mathcal{W}$ at most twice. Assume that $K, K' \in \mathcal{K}$ are two chambers that are adjacent to a wall $W \in \mathcal{W}$ from either side. Then:*

(b1) *If γ intersects W exactly once, then there is a unique column $E \in \mathcal{E}_K$ that is both adjacent to W and intersects γ. Moreover, for every column $E^* \in \mathcal{E}_K$ that also intersects γ, the minimal chain between E and E^* intersects W in at most two columns.*

(b2) *If γ intersects W twice and its endpoints lie on the same side of W as K, then within both intersections it stays inside a column $E' \in \mathcal{E}_{K'}$ adjacent to W. Moreover, there are exactly two columns $E_1, E_2 \in \mathcal{E}_K$ that intersect γ in this order and that are adjacent to W and we have $\operatorname{dist}_K^H(E_1, E_2) > H$. The arc γ leaves K through W right after E_1 and reenters K through W right before E_2.*

(b3) *If γ does not intersect W, but intersects two columns $E_1, E_2 \in \mathcal{E}_K$ that are both adjacent to W, then $\operatorname{dist}_K^H(E_1, E_2) < 3H$.*

(c) *Consider a chamber $K \in \mathcal{K}$ and let $E_1, \dots, E_n \in \mathcal{E}_K$ be the columns of K that γ intersects in that order. Then there are columns $E_1^*, \dots, E_n^* \in \mathcal{E}_K$ such that:*

(c1) *$E_1^* = E_1$ and $E_n^* = E_n$.*

(c2) *$\operatorname{dist}_K^H(E_i^*, E_i) \le 1$ for all $i = 1, \dots, n$.*

(c3) *$E_1^*, \dots, E_n^*$ are pairwise distinct and lie on the minimal chain between E_1 and E_n in that order.*

(c4) *If $E_i^* \ne E_i$, then there are two walls $W, W' \subset \partial K$ that both intersect γ twice such that γ exits K through W' right after E_{i-1}, enters K through W' right before E_i, exits K through W right after E_i and enters K through W right before E_{i+1}. In particular E_i does not lie on the minimal chain between E_1, E_n and E_i is not adjacent to E_{i-1} or E_{i+1}.*

(c5) *If E_i and E_{i+1} are adjacent, then γ stays within $E_i \cup E_{i+1}$ between E_i and E_{i+1}.*

(c6) *If E_i and E_{i+1} are not adjacent, then there is a wall $W \subset \partial K$ such that γ exits K through W right after E_i and enters K through W right before E_{i+1}. The columns E_i, E_i^*, E_{i+1}^*, E_{i+1} lie in that order (some of these columns might be the same) on a minimal chain that runs along W.*

(c7) *If $i_1 < i_2$ and E_{i_1}, E_{i_2} are adjacent to a common wall $W \subset \partial K$, then either $(E_{i_1}, \dots, E_{i_2})$ form a minimal chain or $i_2 = i_1 + 1$ and γ intersects W right after E_{i_1} and right before E_{i_2}.*

Proof The proof uses induction on the combinatorial length $|\gamma|$ of γ. The case $|\gamma| = 0$ is obvious, so assume that $|\gamma| > 0$ and that all assertions of the proposition hold for all combinatorially shorter minimizing arcs.

Let $W \in \mathcal{W}$ be a wall and $K, K' \in \mathcal{K}$ the chambers that are adjacent to W from either side. We first check the first statement of assertion (b). Assume that γ intersects W three times or more. Then by assertion (b2) of the induction hypothesis applied to every subsegment of γ that intersects W exactly twice, we obtain that γ stays within $K \cup K'$ between its first and last intersection with W. This however contradicts Lemma 3.19.

Assertion (b2) follows similarly. Assume that γ intersects W twice and that both endpoints lie on the same side of W as K. Let $E \in \mathcal{E}_K$ be the first column that γ intersects in K and $E^* \in \mathcal{E}_K$ the last. By assertion (b1) of the induction hypothesis applied to the subsegment of γ from $\gamma(0)$ to E^*, we find that $E = E^*$. By assertion (a) of the induction hypothesis, γ remains within E between both intersections with W. So we obtain again that γ stays within $K \cup K'$ between its first and second intersection with W. The rest follows with Lemma 3.19(b).

For assertion (b1), observe that the complete assertion (c) holds for γ in the case in which γ crosses a wall exactly once, because in this case assertion (c) is only concerned with proper subsegments of γ. So consider the columns $E_1, \dots, E_n, E_1^*, \dots, E_n^* \in \mathcal{E}_K$. Without loss of generality, we may assume that $\gamma(1)$ lies on the same side of W as K. This implies that E_1 is adjacent to W. If E_i for some $2 \le i \le n$ was adjacent to W as well, then by assertions (c7) and (b2) the arc γ must be contained in K in between E_1 and E_i. This is however impossible by Lemma 3.19(a) applied to the subsegment of γ between the last column in K' and E_i. So the first part of (b1) holds. Consider now the minimal chain between E_1 and some E_i and assume that three of its columns are adjacent to W. Those columns need to be the first three columns in this chain. We can assume that $E_i = E_n$, because otherwise we could pass to a subsegment of γ. Since, by what we have already shown, γ cannot intersect any column that is adjacent to W

other than E_1, it cannot happen that E_2 is adjacent to E_1 (compare with assertions (c4)–(c5)). So, by assertion (c6), there is a wall $W' \subset \partial K$ with $W' \neq W$ that is intersected twice by γ in between E_1 and E_2. So, if H is sufficiently large, assertion (c6) implies that the first three columns on the minimal chain between E_1 and E_2^*, ie the first three columns on the minimal chain between E_1 and E_n are adjacent to both W' and W. This however contradicts Lemma 3.12 and finishes the proof of assertion (b1).

We now show assertions (c1)–(c7), (a) and (b3). Observe that by the induction hypothesis, it suffices to restrict our attention to the case in which $\gamma(0), \gamma(1) \in K$. Consider now the columns $E_1, \dots, E_n$ as defined in the proposition. If $n \leq 2$, we are done with the choice $E_1^* = E_1$ and $E_2^* = E_2$ by assertion (b2) and assuming $H > 2$. So assume that $n \geq 3$. Assertion (c5) and the first part of (c6) follows immediately by passing to the subsegment between E_i and E_{i+1} and using the induction hypothesis. We will now distinguish the cases of when E_{n-1} lies on the minimal chain between E_1 and E_n or not and establish assertions (c1)–(c4) and the second part of assertion (c6) in each case. Based on these assertions we will next conclude assertion (c7) in both cases.

Consider first the case in which E_{n-1} lies on the minimal chain between E_1 and E_n. Then we can apply the induction hypothesis to the subsegment of γ between E_1 and E_{n-1} and obtain the columns $E_1^*, \dots, E_{n-1}^*$ on the minimal chain between E_1 and E_{n-1}. Moreover, we set $E_n^* = E_n$. Assertions (c1)–(c6) follow immediately.

Next consider the case in which E_{n-1} does not lie on the minimal chain between E_1 and E_n. Define $E_1^*, \dots, E_{n-2}^*$ using the induction hypothesis applied to the subsegment of γ between E_1 and E_{n-1}.

Assume first that E_n and E_{n-1} are adjacent. Then E_n must lie on the minimal chain between E_1 and E_{n-1} (by our assumption and the tree property; see Lemma 3.12). So E_{n-2} cannot be adjacent to E_{n-1}, because that would imply by assertion (c4) of the induction hypothesis that $E_{n-2} = E_n$ and it is elementary that γ cannot reenter a column without exiting K. This means (by assertion (c6) of the induction hypothesis) that there is a wall $W \subset \partial K$ that intersects γ twice and that is adjacent to E_{n-2}, E_{n-2}^* and E_{n-1} and hence also E_n. This however contradicts assertion (b2).

So E_n and E_{n-1} are not adjacent and by assertion (b2) both columns are adjacent to a wall $W \subset \partial K$ such that γ intersects W right after E_{n-1} and right before E_n. By the tree property of $\mathcal{E}_K$ there is a column $E^* \in \mathcal{E}_K$ that lies on the three minimal chains between E_{n-1}, E_n and E_1, E_{n-1} and E_1, E_n. So E^* is adjacent to W

and horizontally lies between E_{n-1}, E_n. By our earlier assumption, $E^* \neq E_{n-1}$. Assertion (b1) applied to a subsegment of γ implies that $\mathrm{dist}_K^H(E^*, E_{n-1}) \leq 1$; so E^* is adjacent to E_{n-1}. If E_{n-2} was adjacent to E_{n-1}, then $E_{n-2} = E^*$ contradicting assertion (b2). So, by assertion (c6) of the induction hypothesis, E_{n-2}, E^*_{n-2} and E_{n-1} are adjacent to a wall $W' \subset \partial K$ such that γ intersects W' twice between E_{n-2} and E_{n-1}. This implies $W \neq W'$. Set $E^*_{n-1} = E^*$ and $E^*_n = E_n$. Assertions (c1)–(c3) follow immediately. Assertion (c4) and the second part of (c6) hold with the walls W and W' that we have just defined.

We now establish assertion (c7) in the general case (ie independently of whether E_{n-1} lies on the minimal chain between E_1 and E_n or not). Assume that $(E_{i_1}, \dots, E_{i_2})$ does not form a minimal chain. Then γ has to leave K in between E_{i_1} and E_{i_2}, ie by assertion (c5) there is a $j \in \{i_1, \dots, i_2 - 1\}$ such that E_j and E_{j+1} are not adjacent and hence, by assertion (c6), γ has to intersect a wall $W' \subset \partial K$ in between E_j and E_{j+1}. The columns on the minimal chain between E^*_j, E^*_{j+1} are adjacent to both W and W' and for $H > 10$ there are at least 3 such columns. So, by Lemma 3.12, $W = W'$ and, by assertion (b2), we must have $E_{i_1} = E^*_j$ and $E_{i_2} = E^*_{j+1}$.

Finally, assertion (a) is a direct consequence of assertion (c7) and assertion (b3) follows from assertion (c7) and Lemma 3.19(c). □

Next, we analyze the relative behavior of two combinatorially minimizing arcs.

Lemma 3.21 *There are constants $\eta^* > 0$ and $H^* < \infty$ such that if $\eta \leq \eta^*$ and $H \geq H^*$, then the following holds:*

Let $\gamma_1, \gamma_2 \colon [0,1] \to \widetilde{M}$ be two combinatorially minimizing arcs and consider a wall $W \in \mathcal{W}$ that is adjacent to two chambers $K, K' \in \mathcal{K}$ on either side. Assume that γ_1 and γ_2 intersect W exactly once and that $\gamma_1(0)$ and $\gamma_2(0)$ lie in a common chamber on the same side of W as K. If that chamber is K, we additionally require that the cells that contain these points are vertically aligned. Similarly, assume that $\gamma_1(1)$ and $\gamma_2(1)$ lie in a common chamber on the same side of W as K'. If that chamber is K', we also require that the cells that contain these points are vertically aligned.

Let $Q_1, Q_2 \subset K$ be the cells that γ_1 and γ_2 intersect right before crossing W and let $Q'_1, Q'_2 \subset K'$ be the cells that γ_1 and γ_2 intersect right after crossing W. Then every pair of the cells Q_1, Q_2, Q'_1 and Q'_2 has combinatorial distance bounded by 4 or $4 + H$ depending on whether they lie on the same side of W or not.

Proof Let $E_1, E_2 \in \mathcal{E}_K$ be the columns that contain Q_1 and Q_2. We first show that $\mathrm{dist}_K^H(E_1, E_2) \le 3$ (in fact, we can show that $\mathrm{dist}_K^H(E_1, E_2) \le 1$, but we don't need this result here).

Define $E_1^*, E_2^* \in \mathcal{E}_K$ to be the first columns in K that are intersected by γ_1 and γ_2. In the case $\gamma_1(0), \gamma_2(0) \in K$ we have $E_1^* = E_2^*$. So, in either case, we can find a wall $W^* \subset \partial K$ with $W^* \neq W$ that is adjacent to both E_1^* and E_2^*. Consider the minimal chain between E_1 and E_1^* and let E_1^{**} be the last column on that chain that is adjacent to W. Define E_2^{**} accordingly. By Proposition 3.20(b1), $\mathrm{dist}_K^H(E_1, E_1^{**}), \mathrm{dist}_K^H(E_2, E_2^{**}) \le 1$. We need to show that $\mathrm{dist}_K^H(E_1^{**}, E_2^{**}) \le 1$.

If W and W^* are adjacent to a common column, then both E_1^{**} and E_2^{**} have to be adjacent to W^* since in that case a minimal chain between E_1 and E_1^* first runs along W and then along W^*. Hence, in that case, $\mathrm{dist}_K^H(E_1^{**}, E_2^{**}) \le 1$ by Lemma 3.12. If W, W^* are not adjacent to a common column, we follow the minimal chain between E_1^{**} and E_1^*, then the minimal chain between E_1^*, E_2^* (along W^*) and finally the minimal chain between E_2^* and E_2^{**}, to obtain a chain that connects E_1^{**} with E_2^{**} and that intersects W only in its first and last column. By the tree property of $\mathcal{E}_K$ this chain covers the minimal chain between E_1^{**} and E_2^{**}, and hence it has to include all columns along W between E_1^{**} and E_2^{**}. So $\mathrm{dist}_K^H(E_1^{**}, E_2^{**}) \le 1$.

It follows that $\mathrm{dist}_K^H(Q_1, Q_2) = \mathrm{dist}_K^H(E_1, E_2) \le 3$. Analogously, $\mathrm{dist}_K^H(Q_1', Q_2') \le 3$. It now follows from Lemma 3.13(d) that $\mathrm{dist}_K^V(Q_1, Q_2), \mathrm{dist}_K^V(Q_1', Q_2') < C_0$. This establishes the claim for $\eta < C_0^{-1}$. □

Lemma 3.22 *There are constants $\eta^* > 0$ and $H^*, C_1 < \infty$ such that if $\eta \le \eta^*$ and $H \ge H^*$, the following holds:*

Let $K \in \mathcal{K}$ be a chamber and $Q_1, \overline{Q}_1, Q_2, \overline{Q}_2 \subset K$ be cells such that Q_1, $\overline{Q}_1$ and Q_2, $\overline{Q}_2$ are vertically aligned pairs in K. Assume that the vertical order of the pair Q_1, $\overline{Q}_1$ is opposite to the one of Q_2, $\overline{Q}_2$ (ie Q_1 is "above" $\overline{Q}_1$ and Q_2 is "below" $\overline{Q}_2$ or the other way round). Let $\gamma, \overline{\gamma}\colon [0,1] \to \widetilde{M}$ be minimizing arcs from Q_1 to Q_2 and from $\overline{Q}_1$ to $\overline{Q}_2$. Then we can find cells $Q', \overline{Q}' \subset K$ such that Q' intersects γ, $\overline{Q}'$ intersects $\overline{\gamma}$ and such that $\mathrm{dist}_K^H(Q', \overline{Q}') < 3H$, $\mathrm{dist}_K^V(Q', \overline{Q}') < C_1 H$ and $\mathrm{dist}(Q', \overline{Q}') < 4H$.

Proof Note that the last inequality follows from the first two inequalities if $\eta^* < C_1^{-1}$.

Let $E_0, E_\omega \in \mathcal{E}_K$ be the columns that contain Q_1, $\overline{Q}_1$ and Q_2, $\overline{Q}_2$. We first invoke Proposition 3.20(c) on γ to obtain columns $E_1, \dots, E_n$, $E_1^*, \dots, E_n^* \in \mathcal{E}_K$ with $E_1 = E_1^* = E_0$ and $E_n = E_n^* = E_\omega$. Then $E_1^*, \dots, E_n^*$ lie on the minimal chain

L between E_0 and E_ω. Let $S \subset L \cup E_1 \cup \cdots \cup E_n$ be the union of all cells in K that intersect γ and all cells in $L \cup E_1 \cup \cdots \cup E_n$ that are adjacent to those cells outside K which intersect γ. Then $Q_1, Q_2 \subset S$ and by the results of Proposition 3.20 these two cells lie in the same connected component of S. Based on the set S we construct another set $S' \subset L$ as follows: S' is the union of $S \cap L$ with all cells in each E_i^* that are horizontally aligned with a cell in $S \cap E_i$. Then again $Q_1, Q_2 \subset S'$ and both cells lie in the same connected component of S'. Similarly, we can invoke Proposition 3.20(c) on $\overline{\gamma}$, obtaining columns $\overline{E}_1, \dots, \overline{E}_{\overline{n}} \in \mathcal{E}_K$ and $\overline{E}_1^*, \dots, \overline{E}_{\overline{n}}^*$ on L and we can define $\overline{S}$ and $\overline{S}'$ in the same way. So $\overline{Q}_1, \overline{Q}_2 \subset \overline{S}'$ and both cells lie in the same connected component of $\overline{S}'$. Since the cells on L are arranged on a rectangular lattice and the cells Q_1, $\overline{Q}_1$ and Q_2, $\overline{Q}_2$ lie on opposite sides of L and have opposite vertical order, we conclude that the sets $\overline{S}$ and $\overline{S}'$ have to intersect. Let $Q^\circ \subset S' \cap \overline{S}'$ be a cell in the intersection and $E^\circ \subset L$ the column containing Q°. So we can find cells $Q, \overline{Q} \in \mathcal{Q}$ that intersect γ and $\overline{\gamma}$ in such a way that the following holds: Either $Q \subset L$ and $Q^\circ = Q$, or $Q \subset (E_1 \cup \cdots \cup E_n) \setminus \operatorname{Int} L$ and Q is adjacent and horizontally aligned with Q°, or $Q \not\subset K$ and Q is either adjacent to Q° or $Q^\circ \subset E_i^*$ for some $i \in \{1, \dots, n\}$ for which $E_i^* \neq E_i$ and Q is adjacent to a cell in E_i that is adjacent to Q° and horizontally aligned with it. In the first two cases we set $Q' := Q$. In the third case, we will define Q' later. So, if $Q \subset K$, then Q' intersects γ and $\operatorname{dist}_K^H(Q', Q^\circ) \leq 1$ and $\operatorname{dist}_K^V(Q', Q^\circ) = 0$. The analogous characterization holds for $\overline{Q}$ and we define $\overline{Q}'$ in the same way if $\overline{Q} \subset K$.

We now consider the case in which $\operatorname{dist}_K^H(E^\circ, E_i^*) \leq 1$ for some $i \in \{1, \dots, n\}$, and we establish the existence of a cell $Q' \subset K$ that intersects γ and that is within bounded distance from Q°. If $Q \subset K$, then we are done by the previous paragraph. So assume that $Q \not\subset K$. Let $K' \in \mathcal{K}$ be the chamber that contains Q and let $W = K \cap K' \in \mathcal{W}$ be the wall between K and K'. So γ intersects W twice and E° is adjacent to W. Choose $i' \in \{1, \dots, n-1\}$ such that γ intersects W between $E_{i'}$ and $E_{i'+1}$. If E° lies between $E_{i'}^*$ and $E_{i'+1}^*$, then $\operatorname{dist}_K^H(E^\circ, E_{i'}^*) \leq 1$ or $\operatorname{dist}_K^H(E^\circ, E_{i'+1}^*) \leq 1$, by our initial assumption. If E° lies on L not between $E_{i'}^*$ and $E_{i'+1}^*$, then we can conclude, by applying Proposition 3.20(b1) to subsegments of γ which intersect W exactly once, that we still have $\operatorname{dist}_K^H(E^\circ, E_{i'}^*) \leq 1$ or $\operatorname{dist}_K^H(E^\circ, E_{i'+1}^*) \leq 1$ (note that there is an $i'' \neq i', i'+1$ such that the minimal chain from $E_{i'}$ or $E_{i'+1}$ to $E_{i''}$ contains the columns $E_{i'}^*$ or $E_{i'+1}^*$ and E° in that order). So, in either case, $\operatorname{dist}_K^H(E^\circ, E_{i'}^*) \leq 1$ or $\operatorname{dist}_K^H(E^\circ, E_{i'+1}^*) \leq 1$ and, after possibly rechoosing i, we may assume that $i = i'$ or $i = i'+1$.

Now let $Q' \subset E_i$ be the cell that is intersected by γ right before or right after W, depending on whether $i = i'$ or $i = i'+1$. Then $\operatorname{dist}_K^H(Q^\circ, Q') \le 2$ and by Lemma 3.13(b) we get $\operatorname{dist}_K^V(Q^\circ, Q') < 3C_0 + 1$.

Combining the previous conclusion with the analogous conclusion for $\overline{\gamma}$ and the triangle inequality yields the desired result in the case in which there are indices $i \in \{1, \dots, n\}$ and $\bar{\imath} \in \{1, \dots, \overline{n}\}$ such that $\operatorname{dist}_K^H(E^\circ, E_i^*) \le 1$ and $\operatorname{dist}_K^H(E^\circ, \overline{E}_{\bar{\imath}}^*) \le 1$. So, after possibly interchanging the roles of γ and $\overline{\gamma}$, it remains to consider the case in which there is an index $i \in \{1, \dots, n-1\}$ such that E° lies strictly in between E_i^* and E_{i+1}^* and such that E° is not adjacent to either of these columns. We will henceforth always assume that. Let $W \subset \partial K$ be the wall that γ intersects between E_i and E_{i+1} and let $K' \in \mathcal{K}$ be the chamber on the other side. Then E_i, E_i^*, E°, E_{i+1}^*, E_{i+1} are arranged along W in that order and, by Lemma 3.12, we must have $Q \subset K'$. (Note that W is the only wall that both intersects γ twice and is adjacent to E°, because any such wall must be adjacent to at least 3 columns of the chain between E_i^* and E_{i+1}^*.) Let $E \in \mathcal{E}_{K'}$ be the column that contains Q. Finally, let $Q' \subset E_i$ be the cell that γ intersects right before W.

Consider the columns on L (not strictly) between $\overline{E}_{\bar{\imath}}^*$ and $\overline{E}_{\bar{\imath}+1}^*$ for each $\bar{\imath} = 1, \dots, \overline{n}-1$. If, for some $\bar{\imath}$, there are at least 3 such columns that are also (not strictly) between E_i^* and E_{i+1}^*, we must have $\operatorname{dist}_K^H(\overline{E}_{\bar{\imath}}^*, \overline{E}_{\bar{\imath}+1}^*) > 1$ and all columns between $\overline{E}_{\bar{\imath}}^*$ and $\overline{E}_{\bar{\imath}+1}^*$ have to be adjacent to W by Lemma 3.12. However, this situation can only occur for at most one index $\bar{\imath}$. So there are two different cases: In the first case there is no such $\bar{\imath}$ and hence all columns that are strictly between E_i^* and E_{i+1}^* are contained in $\overline{E}_1^* \cup \dots \cup \overline{E}_{\overline{n}}^*$. In this case all columns of this union that are strictly between E_i^* and E_{i+1}^* have to be adjacent to one another, and hence by Proposition 3.20(c4) $\overline{\gamma}$ has to intersect all these columns. In the second case there is exactly one such $\bar{\imath}$ and each column that is strictly between E_i^* and E_{i+1}^* either lies (not strictly) between $\overline{E}_{\bar{\imath}}^*$ and $\overline{E}_{\bar{\imath}+1}^*$ or in one of the unions $\overline{E}_0^* \cup \dots \cup \overline{E}_{\bar{\imath}-1}^*$ and $\overline{E}_{\bar{\imath}+1}^* \cup \dots \cup \overline{E}_{\bar{\imath}+1}^*$. Those two subsets can only cover the remaining columns, which are strictly between E_i^* and E_{i+1}^*, if the columns of $\overline{E}_0^*, \dots, \overline{E}_{\overline{n}}^*$ that lie strictly between E_i^* and E_{i+1}^* are adjacent to one another, which implies that either $\overline{E}_{\bar{\imath}}^*$ is adjacent to $\overline{E}_{\bar{\imath}-1}^*$ and/or $\overline{E}_{\bar{\imath}+1}^*$ is adjacent to $\overline{E}_{\bar{\imath}+1}^*$, depending on the side on which the remaining columns lie. So we conclude using Proposition 3.20(c4) that if not all columns that are strictly between E_i^* and E_{i+1}^* are also (not strictly) between $\overline{E}_{\bar{\imath}}^*$ and $\overline{E}_{\bar{\imath}+1}^*$, then $\overline{E}_{\bar{\imath}}^*$ and/or $\overline{E}_{\bar{\imath}+1}^*$ lies strictly between E_i^* and E_{i+1}^* and we have $\overline{E}_{\bar{\imath}}^* = \overline{E}_{\bar{\imath}}$ and/or $\overline{E}_{\bar{\imath}+1}^* = \overline{E}_{\bar{\imath}+1}$, respectively. Now by Proposition 3.20(b1) applied to the minimal chain between E_0 and $\overline{E}_{\bar{\imath}}$ and/or

the minimal chain between E_n and $\overline{E}_{\bar{\imath}+1}$, we conclude that $\overline{E}_{\bar{\imath}}$ has to be adjacent to E_i^* and/or $\overline{E}_{\bar{\imath}+1}$ has to be adjacent to $E_{\bar{\imath}+1}^*$, depending on which of these columns lie strictly between E_i^* and E_{i+1}^*. So, to summarize our findings: Either all columns that are strictly between E_i^* and E_{i+1}^* intersect $\overline{\gamma}$ or they lie (not strictly) between $\overline{E}_{\bar{\imath}}^*$ and $\overline{E}_{\bar{\imath}+1}^*$. In the second case, we can apply the same argument reversing the roles of γ and $\overline{\gamma}$ to conclude that there is no other index $i' \in \{1, \dots, n-1\}$ with $i' \neq i$ such that there are more than 2 columns that are between $E_{i'}^*$, $E_{i'+1}^*$ and $\overline{E}_{\bar{\imath}}^*$, $\overline{E}_{\bar{\imath}+1}^*$. This implies that $\operatorname{dist}_K^H(E_i^*, \overline{E}_{\bar{\imath}}^*), \operatorname{dist}_K^H(E_{i+1}^*, \overline{E}_{\bar{\imath}+1}^*) \leq 1$ in the second case.

In the first case, we use Proposition 3.20(b3) and (c7) to find that $\overline{\gamma}$ intersects fewer than $3H$ columns that are adjacent to W. Since $\overline{\gamma}$ intersects E° we have $\overline{Q}' := \overline{Q} = Q^\circ$. Hence $\operatorname{dist}_K^H(Q', \overline{Q}') < 3H$ and Lemma 3.13(b) yields $\operatorname{dist}_K^V(Q', \overline{Q}') < 3C_0 H + C_0$ and we are done.

In the second case, E° lies strictly between $\overline{E}_{\bar{\imath}}^*$ and $\overline{E}_{\bar{\imath}+1}^*$. So $\overline{Q} \not\subset K$ and $\overline{\gamma}$ intersects W between $\overline{E}_{\bar{\imath}}$ and $\overline{E}_{\bar{\imath}+1}$ (by Lemma 3.12). Let $\overline{K}' \in \mathcal{K}$ be the chamber that contains $\overline{Q}$ and $\overline{W} = K \cap \overline{K}' \in \mathcal{W}$ the wall between K and $\overline{K}'$. We will now show that $\overline{K}' = K'$ and $\overline{W} = W$. If not, then there must be an index $\bar{\imath}' \in \{1, \dots, n-1\}$ with $\bar{\imath}' \neq \bar{\imath}$ such that $\overline{W}$ is adjacent to $\overline{E}_{\bar{\imath}'}^*$, $\overline{E}_{\bar{\imath}'+1}^*$ and E°. If $\bar{\imath}' < \bar{\imath}$, then $\overline{W}$ is also adjacent to E_i^* (which is then between $\overline{E}_{\bar{\imath}'}^*$ and E°), contradicting Lemma 3.12. If $\bar{\imath}' > \bar{\imath}$, then $\overline{W}$ is also adjacent to E_{i+1}^*, contradicting Lemma 3.12 as well. So indeed $Q, \overline{Q} \subset \overline{K}' = K'$; let $\overline{E} \in \mathcal{E}_{K'}$ be the column that contains $\overline{Q}$. Now let $\overline{Q}' \subset \overline{E}_{\bar{\imath}}$ be the cell that $\overline{\gamma}$ intersects right before W.

Recall that $Q' \subset E_i$ and $\overline{Q}' \subset \overline{E}_{\bar{\imath}}$, that Q' is adjacent to E, that $\overline{Q}'$ is adjacent to $\overline{E}$ and that E and $\overline{E}$ are both adjacent to Q°. Moreover, by our previous conclusions, $\operatorname{dist}_K^H(Q', \overline{Q}') \leq 3$. Let $Q'' \subset E_i$ be a cell that is adjacent to $\overline{E}$. Then, by Lemma 3.13(b), $\operatorname{dist}_K^V(\overline{Q}', Q'') < 4C_0$ and, by Lemma 3.13(c), $\operatorname{dist}_K^V(Q'', Q') < C_0$. Hence $\operatorname{dist}_K^V(Q', \overline{Q}') < 5C_0$. This finishes the proof of the lemma. □

The next lemma is a preparation for the combinatorial convexity estimate stated in Proposition 3.24.

Lemma 3.23 *There are constants $\eta^* > 0$ and $H^* < \infty$ such that if $\eta \leq \eta^*$ and $H \geq H^*$, then the following holds:*

Let $K \in \mathcal{K}$ be a chamber and $Q_0, Q_1, Q_2 \subset K$ cells such that Q_1 and Q_2 are vertically aligned. Assume that $\operatorname{dist}(Q_0, Q_1), \operatorname{dist}(Q_0, Q_2) \leq R$ for some $R \geq 0$. Then, for any cell $Q^* \subset K$ between Q_1 and Q_2, we have $\operatorname{dist}(Q_0, Q^*) < R + 8H$.

Proof We prove this lemma by induction on R (observe that we are only interested in a discrete set of values of R) and then on $\mathrm{dist}_K^V(Q_1, Q_2)$. Consider the action $\varphi\colon \mathbb{Z} \curvearrowright \widetilde{M}$ by deck transformations of the universal covering $\widetilde{M} \to M$ that acts as a vertical shift on K, leaving $\widetilde{V}$ and hence the cell structure and combinatorial distance function invariant and choose $z \in \mathbb{Z}$ such that $Q^* = \varphi_z(Q_1)$. We may assume $z \neq 0$.

Without loss of generality, we can assume that Q^* lies between Q_1 and $Q^{**} = \varphi_{-z}(Q_2)$. Otherwise, we can interchange the roles of Q_1 and Q_2. Let γ_1 and γ_2 be minimizing arcs between Q_0 and Q_1, and Q_0 and Q_2, respectively. We can now apply Lemma 3.22 to $\varphi_z \circ \gamma_1$ and γ_2 to obtain cells $Q_1', Q_2' \subset K$ on $\varphi_z \circ \gamma_1$ and γ_2 with $\mathrm{dist}(Q_1', Q_2') < 4H$. Then $\varphi_{-z}(Q_1')$ lies on γ_1 and hence

$$\mathrm{dist}(Q_0, \varphi_{-z}(Q_1')) + \mathrm{dist}(\varphi_{-z}(Q_1'), Q_1) = \mathrm{dist}(Q_0, Q_1) \leq R. \tag{3-2}$$

We also have

$$\mathrm{dist}(Q_0, Q_2') + \mathrm{dist}(Q_2', Q_2) = \mathrm{dist}(Q_0, Q_2) \leq R. \tag{3-3}$$

If $\mathrm{dist}(Q_0, Q_2') + \mathrm{dist}(\varphi_{-z}(Q_1'), Q_1) \leq R + 4H$, then

$$\mathrm{dist}(Q_0, Q^*) \leq \mathrm{dist}(Q_0, Q_2') + \mathrm{dist}(Q_2', Q_1') + \mathrm{dist}(Q_1', \varphi_z(Q_1)) < R + 8H,$$

which proves the desired estimate. On the other hand, assume that

$$\mathrm{dist}(Q_0, Q_2') + \mathrm{dist}(\varphi_{-z}(Q_1'), Q_1) > R + 4H.$$

Then (3-2) and (3-3) give us

$$\mathrm{dist}(Q_0, \varphi_{-z}(Q_1')) + \mathrm{dist}(Q_2', Q_2) < R - 4H.$$

It follows that

$$\begin{aligned}\mathrm{dist}(\varphi_{-z}(Q_1'), Q^{**}) &= \mathrm{dist}(Q_1', Q_2) \leq \mathrm{dist}(Q_1', Q_2') + \mathrm{dist}(Q_2', Q_2) \\ &< 4H + R - 4H - \mathrm{dist}(Q_0, \varphi_{-z}(Q_1')) = R - \mathrm{dist}(Q_0, \varphi_{-z}(Q_1')).\end{aligned}$$

Also by (3-2),

$$\mathrm{dist}(\varphi_{-z}(Q_1'), Q_1) \leq R - \mathrm{dist}(Q_0, \varphi_{-z}(Q_1')).$$

So by the induction hypothesis, we find that

$$\mathrm{dist}(\varphi_{-z}(Q_1'), Q^*) < R - \mathrm{dist}(Q_0, \varphi_{-z}(Q_1')) + 8H.$$

This implies

$$\mathrm{dist}(Q_0, Q^*) \leq \mathrm{dist}(Q_0, \varphi_{-z}(Q_1')) + \mathrm{dist}(\varphi_{-z}(Q_1'), Q^*) < R + 8H. \qquad \square$$

Proposition 3.24 *Assume that M satisfies condition* (C)*. There are constants $\eta^* > 0$ and $H^* < \infty$ such that whenever $\eta \le \eta^*$ and $H \ge H^*$, the following holds:*

Consider a cell $Q_0 \in \mathcal{Q}$, a chamber $K \in \mathcal{K}$ (not necessarily containing Q_0) and cells $Q_1, Q_2 \subset K$ that are vertically aligned within K. Assume that $\operatorname{dist}(Q_0, Q_1)$, $\operatorname{dist}(Q_0, Q_2) \le R$ *for some $R \ge 0$. Then, for any cell $Q^* \subset K$ that is vertically aligned with Q_1 and Q_2 and vertically between Q_1 and Q_2, we have* $\operatorname{dist}(Q_0, Q^*) < R + 10H$.

Proof If $Q_0 \subset K$, we are done by the previous lemma. So assume that Q_0 lies outside of K and let γ_1 and γ_2 be minimizing arcs from Q_0 to Q_1 and Q_2, respectively.

Then there is a unique wall $W \subset \partial K$ through which both γ_1 and γ_2 enter K. Let $Q_1', Q_2' \subset K$ be the first cells in K that are intersected by γ_1 and γ_2. So both cells are adjacent to W. By Lemma 3.21 we know that $\operatorname{dist}(Q_1', Q_2') \le 4$.

So

$$\operatorname{dist}(Q_1', Q_1) \le R - \operatorname{dist}(Q_0, Q_1')$$

and

$$\begin{aligned}\operatorname{dist}(Q_1', Q_2) &\le \operatorname{dist}(Q_1', Q_2') + \operatorname{dist}(Q_2', Q_2) \le 4 + R - \operatorname{dist}(Q_0, Q_2') \\ &\le 4 + R - \operatorname{dist}(Q_0, Q_1') + \operatorname{dist}(Q_1', Q_2') \le R + 8 - \operatorname{dist}(Q_0, Q_1').\end{aligned}$$

We can now apply Proposition 3.24 to obtain

$$\operatorname{dist}(Q_1', Q^*) < R + 8 - \operatorname{dist}(Q_0, Q_1') + 8H.$$

So $\operatorname{dist}(Q_0, Q^*) < R + 10H$ for $H > 4$. □

3.5 A combinatorial convexity estimate if M satisfies condition (B)

Assume now that M satisfies condition (B) in Proposition 3.4, ie that M is the total space of an S^1–bundle over a closed, orientable surface of genus ≥ 2. In this setting we will establish the same combinatorial convexity estimate as in Proposition 3.24. It will be stated in Proposition 3.28. Its proof will resemble the proof in the previous subsection, except that most lemmas will be simpler.

We first let $\mathcal{E} = \bigcup_{K \in \mathcal{K}} \mathcal{E}_K$ be the set of all columns of $\widetilde{M}$. We say that two columns $E_1, E_2 \in \mathcal{E}$ are *adjacent* their intersection contains a point of $\widetilde{V} \setminus \widetilde{V}^{(1)}$. In other words, E_1 and E_2 are adjacent if and only if we can find cells $Q_1 \subset E_1$, $Q_2 \subset E_2$ such that Q_1 and Q_2 are adjacent. Observe that in the setting of condition (B) every column $E_1 \in \mathcal{E}$ is adjacent to only finitely many columns $E_2 \in \mathcal{E}$ and every wall $W \in \mathcal{W}$ intersects its adjacent columns $E \in \mathcal{E}$ from either side in parallel strips $E \cap W$ (see Lemma 3.13).

The first lemma is an analog of Lemma 3.19.

Lemma 3.25 *There are constants $\eta^* > 0$ and $H^* < \infty$ such that if $\eta \le \eta^*$ and $H \ge H^*$, the following holds:*

Consider two chambers $K, K' \in \mathcal{K}$ that are adjacent to one another across a wall $W = K \cap K' \in \mathcal{W}$. Assume that $\gamma\colon [0,1] \to \tilde{M}$ is combinatorially minimizing and that its image is contained in $K \cup K'$. Then the following holds:

(a) *The arc γ intersects W at most twice and γ does not reenter any column, ie $\gamma^{-1}(E)$ is an interval for all $E \in \mathcal{E}$.*

(b) *In the case in which γ intersects W exactly once, the following is true: the columns on γ that are adjacent to W form two minimal chains in K and K', moving in the same direction, which are adjacent to one another in a unique pair of columns $E \in \mathcal{E}_K$ and $E' \in \mathcal{E}_{K'}$.*

(c) *In the case in which γ intersects W exactly twice, the following is true: Assume that $\gamma(0), \gamma(1) \in K$. Let $E_1, E_2 \in \mathcal{E}_K$ be the columns that γ intersects right before and after W. Then γ does not intersect any column of K that is adjacent to W and that horizontally lies strictly between E_1 and E_2. Moreover, $\operatorname{dist}_K^H(E_1, E_2) > H$.*

Proof First note that every subsegment of γ that does not intersect W and whose endpoints lie in columns that are adjacent to W, stays within columns that are adjacent to W and does not reenter any column. So we can restrict our attention to the case in which γ intersects only columns that are adjacent to W.

Assume first that we are in the setting of assertion (b), meaning that γ intersects W exactly once and assume without loss of generality that $\gamma(0) \in K$. Let $E_1, \dots, E_n \in \mathcal{E}_K$ and $E'_1, \dots, E'_{n'} \in \mathcal{E}_{K'}$ be the columns that γ intersects in that order. Then both sequences of columns form minimal chains, which move along W, and E_n and E'_1 are adjacent across W. We now show that E_i can only be adjacent to $E'_{i'}$ if $i = n$ and $i' = 1$. This will also imply that the directions of both minimal chains agree. Assume that this was not the case and assume without loss of generality that E_i is adjacent to $E'_{i'}$ for some $i < n$ and $i' \ge 1$ (otherwise we reverse the orientation of γ). Let $Q_1 \subset E_1$ be the cell that contains $\gamma(0)$, $Q_2 \subset E_n$ and $Q_3 \subset E'_1$ the cells that γ intersects right before and after W, and $Q_4 \subset E_{i'}$ a cell that intersects γ. Choose, moreover, a cell $Q^* \subset E_i$ which is adjacent to Q_4. By Lemma 3.13(e),

$$\operatorname{dist}_K^V(Q_2, Q^*) < \operatorname{dist}_{K'}^V(Q_3, Q_4) + C_0 \operatorname{dist}_K^H(Q_2, Q^*) + C_0.$$

So

$$\begin{aligned}\operatorname{dist}_K^V(Q_1, Q^*) &\le \operatorname{dist}_K^V(Q_1, Q_2) + \operatorname{dist}_K^V(Q_2, Q^*)\\ &< \operatorname{dist}_K^V(Q_1, Q_2) + \operatorname{dist}_{K'}^V(Q_3, Q_4) + C_0 \operatorname{dist}_K^H(Q_2, Q^*) + C_0.\end{aligned}$$

Hence,

$$\begin{aligned}\operatorname{dist}(Q_1, Q_4) &\le \operatorname{dist}(Q_1, Q^*) + H\\ &< \operatorname{dist}_K^H(Q_1, Q^*) + \eta \operatorname{dist}_K^V(Q_1, Q_2) + \eta \operatorname{dist}_{K'}^V(Q_3, Q_4)\\ &\qquad + \eta C_0 \operatorname{dist}_K^H(Q_2, Q^*) + \eta C_0 + H.\end{aligned}$$

On the other hand, the minimizing property of γ yields

$$\begin{aligned}\operatorname{dist}(Q_1, Q_4) &= \operatorname{dist}(Q_1, Q_2) + H + \operatorname{dist}(Q_3, Q_4)\\ &\ge \operatorname{dist}_K^H(Q_1, Q_2) + \eta \operatorname{dist}_K^V(Q_1, Q_2) + \eta \operatorname{dist}_{K'}^V(Q_3, Q_4) + H.\end{aligned}$$

Combining both inequalities yields

$$\operatorname{dist}_K^H(Q_1, Q_2) < \operatorname{dist}_K^H(Q_1, Q^*) + \eta C_0 \operatorname{dist}_K^H(Q_2, Q^*) + \eta C_0.$$

Since $\operatorname{dist}_K^H(Q_1, Q_2) = n-1$, $\operatorname{dist}_K^H(Q_1, Q^*) = i-1$ and $\operatorname{dist}_K^H(Q_2, Q^*) = n-i \ge 1$, we obtain

$$n - 1 < i - 1 + \eta C_0(n-i) + \eta C_0.$$

This yields a contradiction if $\eta < (2C_0)^{-1}$. So assertion (b) holds.

Assume next that we are in the setting of assertion (c), meaning that γ intersects W exactly twice and that $\gamma(0), \gamma(1) \in K$. Define $E_1, E_2 \in \mathcal{E}_K$ as in the statement of the Lemma. We now establish the bound $\operatorname{dist}_K^H(E_1, E_2) > H$ (for sufficiently small η and large H). By assertion (b) applied to subsegments of γ, we find that γ cannot intersect any column of K that is adjacent to W and lies strictly between E_1 and E_2. Now let $Q_1 \subset E_1$ and $Q_2 \subset E_2$ be the cells that γ intersects right before and after W and let $Q_1', Q_2' \subset K'$ be the cells that γ intersects right after Q_1 and right before Q_2. Then

$$\operatorname{dist}(Q_1, Q_2) = 2H + \operatorname{dist}(Q_1', Q_2') \ge 2H + \eta \operatorname{dist}_{K'}^V(Q_1', Q_2').$$

By Lemma 3.13(e),

$$\operatorname{dist}_K^V(Q_1, Q_2) < \operatorname{dist}_{K'}^V(Q_1', Q_2') + C_0 \operatorname{dist}_K^H(Q_1, Q_2) + C_0.$$

So

$$\operatorname{dist}(Q_1, Q_2) < \operatorname{dist}_K^H(Q_1, Q_2) + \eta \operatorname{dist}_{K'}^V(Q_1', Q_2') + \eta C_0 \operatorname{dist}_K^H(Q_1, Q_2) + \eta C_0.$$

Hence,

$$2H < (1+\eta C_0)\,\mathrm{dist}_K^H(Q_1,Q_2)+\eta C_0.$$

Assertion (c) follows for $H>10$ and $\eta<(2C_0)^{-1}$.

It remains to show assertion (a). To do this, we first show that γ cannot intersect W more than twice. Assume it does. By passing to a subsegment and possibly interchanging the roles of K and K', we can assume that γ intersects W exactly three times and that $\gamma(0)\in K$. Let $Q_1,Q_2,Q_3\subset K$ be the cells in K that γ intersects before the first, after the second and before the third intersection with W and let $Q_1',Q_2',Q_3'\subset K'$ be the cells of K' that γ intersects after the first, before the second and after the third intersection with W. Then

$$\begin{aligned}\text{(3-4)}\quad \mathrm{dist}(Q_1,Q_3') = 3H+\mathrm{dist}_{K'}^H(Q_1',Q_2')+\eta\,\mathrm{dist}_{K'}^V(Q_1',Q_2')\\ +\mathrm{dist}_K^H(Q_2,Q_3)+\eta\,\mathrm{dist}_K^V(Q_2,Q_3).\end{aligned}$$

Let $Q^*\subset K$ be the cell that is adjacent to W and that is located relatively to Q_1, Q_2 and Q_3 such that Q_1, Q_2, Q_3, Q^* forms a "parallelogram", ie

$$\begin{aligned}&\mathrm{dist}_K^H(Q_1,Q^*)=\mathrm{dist}_K^H(Q_2,Q_3),\quad \mathrm{dist}_K^V(Q_1,Q^*)=\mathrm{dist}_K^V(Q_2,Q_3),\\ &\mathrm{dist}_K^H(Q^*,Q_3)=\mathrm{dist}_K^H(Q_1,Q_2),\quad \mathrm{dist}_K^V(Q^*,Q_3)=\mathrm{dist}_K^V(Q_1,Q_2)\end{aligned}$$

in an oriented sense. Let, moreover, $Q^{*\prime}\subset K'$ be a cell that is adjacent to Q^*. Then, by Lemma 3.13(g),

$$\mathrm{dist}_{K'}^H(Q^{*\prime},Q_3')<\mathrm{dist}_{K'}^H(Q_1',Q_2')+C_0,\quad \mathrm{dist}_{K'}^V(Q^{*\prime},Q_3')<\mathrm{dist}_{K'}^V(Q_1',Q_2')+C_0.$$

So

$$\begin{aligned}\mathrm{dist}(Q_1,Q_3') &\le \mathrm{dist}(Q_1,Q^*)+\mathrm{dist}(Q^*,Q^{*\prime})+\mathrm{dist}(Q^{*\prime},Q_3')\\ &\le \mathrm{dist}_K^H(Q_1,Q^*)+\eta\,\mathrm{dist}_K^V(Q_1,Q^*)+H+\mathrm{dist}_{K'}^H(Q^{*\prime},Q_3')\\ &\qquad+\eta\,\mathrm{dist}_{K'}^V(Q^{*\prime},Q_3')\\ &< H+\mathrm{dist}_K^H(Q_2,Q_3)+\eta\,\mathrm{dist}_K^V(Q_2,Q_3)+\mathrm{dist}_{K'}^H(Q_1',Q_2')+C_0\\ &\qquad+\eta\,\mathrm{dist}_{K'}^V(Q_1',Q_2')+\eta C_0.\end{aligned}$$

Together with (3-4) this yields

$$2H<C_0+\eta C_0$$

and hence a contradiction for $H>C_0$ and $\eta<1$.

Next, we show that γ does not reenter any column. If γ does not intersect W, then this fact is a consequence of Lemma 3.18. The same is true if γ intersects W exactly once, by passing to a subsegment. So it remains to consider the case in which γ intersects W exactly twice. Assume that the assertion was wrong. By passing to a subsequent, we can assume that there is a column $E \in \mathcal{E}_K$ such that $\gamma(0), \gamma(1) \in E$ and that γ intersects W exactly twice. Let $E_1, E_2 \in \mathcal{E}_K$ be the columns that γ intersects right before and after W, as in the last part of the lemma. Let, moreover, $E'_1 = E, \dots, E'_{n'} = E_1 \in \mathcal{E}_K$ and $E''_1 = E_2, \dots, E''_{n''} = E \in \mathcal{E}_K$ be the columns of K that γ intersects in that order. By Lemma 3.18 we know that $L_1 = (E'_1, \dots, E'_{n'})$ and $L_2 = (E''_{n''}, \dots, E''_1)$ form minimal chains between E, E_1 and E, E_2, respectively. Let L be the minimal chain between E_1 and E_2. By assertion (c) of this lemma, the pairs L_1, L and L_2, L only intersect in E_1 and E_2, respectively. So $L_1 \cup L$ is a minimal chain between E and E_2. By the tree property of $\mathcal{E}_K$, we must have $L_1 \cup L = L_2$, which is impossible, since L_2 intersects L only in E_2. This finishes the proof. □

The following proposition and its proof are similar to Proposition 3.20.

Proposition 3.26 *Assume that M satisfies condition* (B). *There are constants $\eta^* > 0$ and $H^* < \infty$ such that if $\eta \le \eta^*$ and $H \ge H^*$, the following holds:*

Consider a combinatorially minimizing arc $\gamma\colon [0,1] \to \tilde{M}$. Then:

(a) *For every column $E \in \mathcal{E}$, the preimage $\gamma^{-1}(E)$ is an interval.*

(b) *γ intersects every wall $W \in \mathcal{W}$ at most twice. Assume that $K, K' \in \mathcal{K}$ are two chambers that are adjacent to a wall $W \in \mathcal{W}$ from either side. Then:*

  (b1) *If γ intersects W exactly once then the following holds: Assume that $\gamma(0)$ lies on the same side of W as K. Let $E \in \mathcal{E}_K$ be the first column that is intersected by γ and that is adjacent to W. Then for every column $E^* \in \mathcal{E}_K$ that γ intersects before E, the minimal chain between E^* and E intersects W in at most two columns.*

  (b2) *If γ intersects W exactly twice and its endpoints lie on the same side of W as K, then γ stays within K' between both intersections and the columns $E_1, E_2 \in \mathcal{E}_K$ that γ intersects right before and after W satisfy $\mathrm{dist}_K^H(E_1, E_2) > H$. Moreover, γ does not intersect any column of K that is adjacent to W and that horizontally lies strictly between E_1 and E_2.*

  (b3) *If γ intersects two columns $E_1, E_2 \in \mathcal{E}$ that are both adjacent to W, then γ stays within $K \cup K'$ in between E_1 and E_2 and only intersects columns that are adjacent to W.*

(c) *Consider a chamber $K \in \mathcal{K}$ and let $E_1, \dots, E_n \in \mathcal{E}_K$ be the columns of K that γ intersects in that order. Then there are columns $E_1^*, \dots, E_n^* \in \mathcal{E}_K$ such that assertions* (c1)–(c6) *of Proposition 3.20 hold.*

Proof We use again induction on the combinatorial length $|\gamma|$ of γ. Assume that $|\gamma| > 0$, since for $|\gamma| = 0$ there is nothing to prove. The first part of assertion (b) follows as in the proof of Proposition 3.20.

We now establish assertion (b1). So assume that γ intersects W exactly once and that $\gamma(0)$ lies on the same side of W as K and consider the columns $E, E^* \in \mathcal{E}_K$. Note that it suffices to show that the second or third last element of the minimal chain between E^* and E is not adjacent to W. Apply assertions (b) and (c) of the induction hypothesis to the subsegment of γ between E^* and E. We obtain sequences $E_1, \dots, E_n$ and $E_1^*, \dots, E_n^*$ with $E_1 = E_1^* = E^*$ and $E_n = E_n^* = E$. If E_{n-1} and E_n are adjacent, then $E_{n-1} = E_{n-1}^*$ lies on the minimal chain between E^* and E and by assumption E_{n-1} cannot be adjacent to W; so we are done. If E_{n-1}, E_n are not adjacent, then γ intersects a wall $W' \subset \partial K$, $W' \neq W$ twice between E_{n-1}, E_n. All columns on the minimal chain between E_{n-1}^* and E_n are adjacent to W'. By Lemma 3.12 at most 2 of those columns can also be adjacent to W.

Assertion (b2) follows from assertion (b1) of the induction hypothesis and Lemma 3.25 by passing to subsegments of γ that intersect W exactly once and whose endpoints are contained in columns adjacent to W.

Next, we establish assertions (c) and (a). It suffices to consider the case in which $\gamma(0), \gamma(1) \in K$. Let $E_1, \dots, E_n \in \mathcal{E}_K$ be as defined in the proposition. If $n \leq 2$, then we are done using assertion (b2); so assume $n \geq 3$. If E_{n-1} lies on the minimal chain between E_1 and E_n, then we are done as in the proof of Proposition 3.20. So assume that E_{n-1} does not lie on the minimal chain between E_1 and E_n.

We show that E_{n-1} and E_n cannot be adjacent. Otherwise, as in the proof of Proposition 3.20, γ intersects a wall $W \subset \partial K$ twice between E_{n-2} and E_{n-1} and the columns E_{n-2}, E_{n-2}^*, E_n and E_{n-1} lie along W in that order. This contradicts assertion (b2).

So there is a wall $W \subset \partial K$ that is adjacent to both E_{n-1} and E_n, and γ crosses W twice between those two columns. We now proceed as in the proof of Proposition 3.20, but we have to be careful whenever we make use of assertion (b). As in that proof, we can find a column $E^* \in \mathcal{E}_K$ that lies on the three minimizing chains between E_{n-1}, E_n and E_1, E_{n-1} and E_1, E_n with $E^* \neq E_{n-1}$. We also know that E_{n-2}

cannot be adjacent to E_{n-1}, since otherwise it would lie on the minimal chain between E_{n-1} and E^* along W, in contradiction to assertion (b2). So, by assertion (c6) of the induction hypothesis, E_{n-2}, E^*_{n-2} and E_{n-1} are adjacent to a wall $W' \subset \partial K$ such that γ intersects W' twice between E_{n-2} and E_{n-1}. This implies $W' \neq W$ by assertion (b). Now both W and W' are adjacent to all columns on the minimal chain between E_{n-1} and E^* or between E_{n-1} and E^*_{n-2}, whichever is shorter. So, by Lemma 3.12, we must have $\mathrm{dist}^H_K(E^*, E_{n-1}) = 1$. Assertion (c1)–(c6) now follow as in the proof of Proposition 3.20.

Now, for assertion (a), we may assume that $\gamma(0), \gamma(1) \in E \in \mathcal{E}_K$ in view of the induction hypothesis. Then assertion (c) implies that γ is fully contained in E.

Finally, we establish assertion (b3). In view of the induction hypothesis it suffices to consider the case in which $E_1, E_2 \in \mathcal{E}_K$ and in which γ does not intersect W. Apply assertion (c) to obtain sequences $E'_1, \dots, E'_n$ and $E'^*_1, \dots, E'^*_n$ with $E'^*_1 = E'_1 = E_1$ and $E'^*_n = E'_n = E_2$. It follows that all columns $E^*_1, \dots, E^*_n$ are adjacent to W. If γ crossed a wall $W' \subset \partial K$ twice in between some E'_i and E'_{i+1}, then all columns between E'^*_i and E'^*_{i+1} would be adjacent to W' and W. This is impossible by Lemma 3.12. □

The next lemma is an analog of Lemma 3.22. Note that in the setting of condition (B), we don't need to work inside a single chamber. This fact will later compensate us for the lack of an analog for Lemma 3.21.

Lemma 3.27 *There are constants $\eta^* > 0$ and $H^* < \infty$ such that if $\eta \leq \eta^*$ and $H \geq H^*$, the following holds:*

Let $E^\circ_1, E^\circ_2 \in \mathcal{E}$ be two columns and $Q_1, \overline{Q}_1 \subset E^\circ_1$ and $Q_2, \overline{Q}_2 \subset E^\circ_2$ cells such that the vertical orders of Q_1, $\overline{Q}_1$ and Q_2, $\overline{Q}_2$ are opposite to each other. Let γ, $\overline{\gamma}\colon [0,1] \to \widetilde{M}$ be minimizing arcs from Q_1 to Q_2 and from $\overline{Q}_1$ to $\overline{Q}_2$. Then we can find cells $Q', \overline{Q}' \in \mathcal{Q}$ that intersect $\gamma, \overline{\gamma}$ and such that $\mathrm{dist}(Q', \overline{Q}') < 3H$.

Proof Consider first a wall $W \in \mathcal{W}$ that intersects γ (and hence also $\overline{\gamma}$) exactly once. Let $K, K' \in \mathcal{K}$ be the chambers that are adjacent to W from either side in such a way that $\gamma(0)$ and $\overline{\gamma}(0)$ lie on the same side of W as K. Let $E \in \mathcal{E}_K$ be the first column on γ that is adjacent to W and choose $\overline{E} \in \mathcal{E}_K$ analogously. We argue similarly to in the proof of Lemma 3.21 that $\mathrm{dist}^H_K(E, \overline{E}) \leq 3$. Let $E^* \in \mathcal{E}_K$ be the first column on γ inside K and define $\overline{E}^*$ analogously. Then either $E^* = \overline{E}^* = E^\circ_1$ or $E^\circ_1 \notin \mathcal{E}_K$. In both cases there is a wall $W^* \subset \partial K$ with $W^* \neq W$ that is adjacent to both E^* and $\overline{E}^*$.

Let $E^{**} \in \mathcal{E}_K$ be the last column on the minimal chain between E and E^* that is adjacent to W and define $\overline{E}^{**} \in \mathcal{E}_K$ analogously. By Proposition 3.26(b1) we have $\mathrm{dist}_K^H(E, E^{**})$, $\mathrm{dist}_K^H(\overline{E}, \overline{E}^{**}) \leq 1$. It now follows, as in the proof of Lemma 3.21, that $\mathrm{dist}_K^H(E^{**}, \overline{E}^{**}) \leq 1$ and hence $\mathrm{dist}_K^H(E, \overline{E}) \leq 3$ (observe that this part of the proof only makes use of the tree property of $\mathcal{E}_K$ from Lemma 3.12).

Now let $W_1, \dots, W_h$ be all the walls that γ intersects exactly once in this order. Then also $\overline{\gamma}$ intersects each of these walls exactly once in this order. For each $i = 1, \dots, h$, let $E_i' \in \mathcal{E}$ be the first and E_i'' the last column on γ that is adjacent to W_i. Define $\overline{E}_i'$ and $\overline{E}_i''$ accordingly. By the last paragraph, we obtain that E_i', $\overline{E}_i'$ and E_i'', $\overline{E}_i''$ have horizontal distance ≤ 3 in the chamber in which they are contained (the bound on the horizontal distance between E_i'' and $\overline{E}_i''$ can be obtained by reversing the orientation of γ and $\overline{\gamma}$). Choose cells $Q_i' \subset E_i'$, $Q_i'' \subset E_i''$ or $\overline{Q}_i' \subset \overline{E}_i'$, $\overline{Q}_i'' \subset \overline{E}_i''$ that intersect γ or $\overline{\gamma}$, respectively.

We first consider the case in which there is some $i \in \{1, \dots, h\}$ such that the vertical orders of the pairs Q_i', $\overline{Q}_i'$ and Q_i'', $\overline{Q}_i''$ are different. Observe that by Proposition 3.26(b3) the arc γ only intersects cells adjacent to W_i between Q_i' and Q_i''; the same is true for $\overline{\gamma}$. Let $S \subset \widetilde{M}$ be the union of all cells that γ intersects between Q_i' and Q_i'' and define $\overline{S}$ accordingly. Using Lemma 3.25 we find that either $S \cap W_i$ and $\overline{S} \cap W_i$ intersect or there is a cell $Q' \in \mathcal{Q}$ on γ with $\mathrm{dist}(Q', \overline{Q}_i') \leq 3 + H$ or $\mathrm{dist}(Q', \overline{Q}_i'') \leq 3 + H$ or there is a cell $\overline{Q}' \in \mathcal{Q}$ on $\overline{\gamma}$ with $\mathrm{dist}(\overline{Q}', Q_i') \leq 3 + H$ or $\mathrm{dist}(\overline{Q}', Q_i'') \leq 3 + H$. In all these cases we are done.

So assume from now on that the vertical orders of the pairs Q_i', $\overline{Q}_i'$ and Q_i'', $\overline{Q}_i''$ are the same for all $i = 1, \dots, h$. Choose $i \in \{1, \dots, h\}$ minimal such that the vertical order of Q_i', $\overline{Q}_i'$ differs from that of Q_1, $\overline{Q}_1$. If there is no such i, then the vertical orders of Q_h', $\overline{Q}_h'$ and Q_2, $\overline{Q}_2$ are opposite and we can get rid of this case by reversing the orientations of γ and $\overline{\gamma}$. Let $K \in \mathcal{K}$ be the chamber that contains Q_i' and $\overline{Q}_i'$. If $i > 1$, the choice of i implies that the vertical order Q_{i-1}'', $\overline{Q}_{i-1}''$ is different from that of Q_i', $\overline{Q}_i' \subset K$. If $i = 1$, then the vertical order of Q_1, $\overline{Q}_1 \subset K$ is different from that of Q_1', $\overline{Q}_1'$. Apply Proposition 3.26(c) to the subsegment of γ between Q_{i-1}' or Q_1 and Q_i'' to obtain columns $E_1, \dots, E_n$ and $E_1^*, \dots, E_n^* \in \mathcal{E}_K$. Similarly we obtain the columns $\overline{E}_1, \dots, \overline{E}_{\overline{n}}$ and $\overline{E}_1^*, \dots, \overline{E}_{\overline{n}}^* \in \mathcal{E}_K$ for the corresponding subsegment of $\overline{\gamma}$. Note that $\mathrm{dist}_K^H(E_1, \overline{E}_1), \mathrm{dist}_K^H(E_n, \overline{E}_{\overline{n}}) \leq 3$.

If $\mathrm{dist}_K^H(E_1, E_n), \mathrm{dist}_K^H(\overline{E}_1, \overline{E}_{\overline{n}}) \leq 6$, then by Proposition 3.26(c), all columns E_i and $\overline{E}_i$ have distance ≤ 17 from one another and hence we can just pick cells Q' and $\overline{Q}'$ that are horizontally aligned to show the lemma. So assume from now on that this

is not the case and let L and $\overline{L}$ be the minimal chains between the pairs E_1, E_n and $\overline{E}_1$, $\overline{E}_{\overline{n}}$. By the tree property, as explained in Lemma 3.12, L and $\overline{L}$ intersect in a minimal chain L° such that every column on $(L \cup \overline{L}) \setminus L^\circ$ has horizontal distance ≤ 3 from L°.

As in the proof of Lemma 3.22 define the sets $S \subset L \cup E_1 \cup \cdots \cup E_n$, $S' \subset L$ and $\overline{S} \subset \overline{L} \cup \overline{E}_1 \cup \cdots \cup \overline{E}_{\overline{n}}$, $\overline{S}' \subset \overline{L}$. Observe that S' and $\overline{S}'$ lie in different sets and might not intersect, as before. However, we can still find cells $Q^\circ \subset S'$ and $\overline{Q}^\circ \subset \overline{S}'$ such that

$$\operatorname{dist}_K^H(Q^\circ, \overline{Q}^\circ) \leq 6 \quad \text{and} \quad \operatorname{dist}_K^V(Q^\circ, \overline{Q}^\circ) = 0.$$

We will work with these cells now instead of Q° alone. By the definition of S' there is a cell $Q^{\circ\circ} \subset S$ that is either equal to Q° or adjacent to Q° and horizontally aligned with it, ie $\operatorname{dist}_K^H(Q^{\circ\circ}, Q^\circ) \leq 1$ and $\operatorname{dist}_K^V(Q^{\circ\circ}, Q^\circ) = 0$. Again, by the definition of S, there is a cell $Q' \in \mathcal{Q}$ on γ that is either equal to $Q^{\circ\circ}$ or adjacent to it across a wall, ie $\operatorname{dist}(Q', Q^{\circ\circ}) \leq H$. Altogether this implies that $\operatorname{dist}(Q', Q^\circ) \leq 1 + H$. By an analogous argument, we can find a cell $\overline{Q}'$ on $\overline{\gamma}$ with $\operatorname{dist}(\overline{Q}', \overline{Q}^\circ) \leq 1 + H$. Hence, $\operatorname{dist}(Q', \overline{Q}') \leq 5 + 2H < 3H$ for large enough H. □

Proposition 3.28 *Proposition 3.24 also holds if M satisfies condition* (B).

Proof We follow the proof of Lemma 3.23. Observe that since M satisfies condition (B), the action $\varphi\colon \mathbb{Z} \curvearrowright \widetilde{M}$ acts as a vertical shift on each column of $\widetilde{M}$. So we do not need to restrict to the case in which the cells Q_0, Q_1 and Q_2 lie in the same chamber. Instead of applying Lemma 3.22, we now make use of Lemma 3.27 to obtain cells $Q_1', Q_2' \subset K$ on $\varphi_z \circ \gamma_1$ and γ_2 with $\operatorname{dist}(Q_1', Q_2') < 3H < 4H$. The rest of the proof is exactly the same as that of Lemma 3.23. □

3.6 Proof of Proposition 3.4 if M satisfies condition (B) or (C)

We will now apply the combinatorial convexity estimates from Propositions 3.24 and 3.28 to construct large polyhedral balls in $\widetilde{M}$ which consist of cells. In the following we will always assume that M satisfies condition (B) or (C) and that η and H have been chosen smaller/larger than all constants η^* and H^*, respectively, which appeared in the lemmas and propositions of Sections 3.4 and 3.5.

Lemma 3.29 *Let $K \in \mathcal{K}$ be a chamber of $\widetilde{M}$ and consider a finite union of cells $S \subset K$ whose interior is connected. Assume that S has the property that for any two cells $Q_1, Q_2 \subset S$ that are vertically aligned, S also contains all cells that are vertically between Q_1 and Q_2. Then S is homeomorphic to a closed* 3*–disk and the intersection*

of S with every wall $W \subset \partial K$ has connected interior in W. More precisely, there is a continuous, injective map $b\colon D^3 \to \widetilde{M}$ with $b(D^3) = S$ that is an embedding on $B^3 \cup (S^2 \setminus b^{-1}(\widetilde{V}^{(1)}))$ and for all walls $W \subset \partial K$ the preimage $b^{-1}(W)$ is either empty or a (connected) topological disk that is the union of rectangles.

Proof The lemma is obviously true if S only consists of cells that are vertically aligned. Observe next that the columns of K are bounded by subsets of ∂K and components of $\pi^{-1}(C_K)$. Those components correspond to arcs of $\widetilde{C}_K^* \subset \widetilde{\Sigma}_K$, are diffeomorphic to $I \times \mathbb{R}$ and every two adjacent columns intersect in exactly one such component. Moreover, each such component separates K into two components.

Consider now such a component $X \subset \pi^{-1}(C_K)$ with the property that not all cells of S lie on one side of X. This is always possible if not all cells of S are vertically aligned. Let $S_1, S_2 \subset K$ be the closures of the two components of $S \setminus X$. Then $S_1 \cap S_2$ is a connected rectangle and so the interiors of S_1 and S_2 must be connected and hence S_1 and S_2 are homeomorphic to 3–disks. Since the interior of $S_1 \cap S_2$ in X is a (connected) disk, we find that $S = S_1 \cup S_2$ is a topological 3–disk as well. The fact that S is a topological disk follows from this argument by induction.

Next, let $W \subset \partial K$ be a wall and assume that two cells $Q, Q' \subset S$ are adjacent to W. Let $E, E' \in \mathcal{E}_K$ be the columns that contain Q and Q'. Since S is connected, we can find a chain $(E_0, \dots, E_n)$ between E and E' such that E_i contains a cell of S for all $i = 0, \dots, n$. We may assume that we have picked the chain so that n is minimal. Thus this chain cannot contain any column twice. Hence it is minimal and so all its columns are adjacent to W. Note that $E_i \cap S \cap W$ is a rectangle for each $i = 0, \dots, n$. By the previous paragraph, the rectangles $E_{i-1} \cap S \cap W$ and $E_i \cap S \cap W$ intersect in more than one point. It follows that $S \cap W$ is a topological disk.

It follows easily that we can connect Q with Q' through cells in K which are adjacent to W, and hence $S \cap W$ is connected. By the property of S, this intersection can only be a topological 2–disk.

The existence of the map b follows along the same lines as this proof. □

Let $Q_0 \in \mathcal{Q}$ be an arbitrary cell and $R > 0$ a positive number. Then we define

$$B_R(Q_0) = \bigcup \{Q \in \mathcal{Q} : \operatorname{dist}(Q, Q_0) < R\}.$$

Next, consider the distance function $\operatorname{dist}^{\mathcal{K}}\colon \mathcal{K} \times \mathcal{K} \to [0, \infty)$, which assigns to every pair of chambers K_1 and K_2 the length of the minimal chain between K_1 and K_2. This length is equal to the minimal number of intersections of an arc between K_1 and K_2

with the walls of $\widetilde{M}$. For two cells $Q_1 \subset K_1$ and $Q_2 \subset K_2$ we set $\operatorname{dist}^{\mathcal{K}}(Q_1, Q_2) = \operatorname{dist}^{\mathcal{K}}(K_1, K_2)$. Observe that

$$\operatorname{dist}(Q_1, Q_2) \geq H \operatorname{dist}^{\mathcal{K}}(K_1, K_2).$$

Let $J > 0$ be a large constant whose value we will determine later. We define a new distance function $\operatorname{dist}'(\cdot,\cdot)$ on $\mathcal{Q}$ as follows

$$\operatorname{dist}'(Q_1, Q_2) := \operatorname{dist}(Q_1, Q_2) + J \operatorname{dist}^{\mathcal{K}}(Q_1, Q_2).$$

Obviously, $(\mathcal{Q}, \operatorname{dist}')$ is a metric space. Set, moreover,

$$B'_R(Q_0) = \bigcup \{Q \in \mathcal{Q} : \operatorname{dist}'(Q, Q_0) < R\}.$$

Finally, we define

$$P_R(Q_0) = \bigcup \left\{ Q \in \mathcal{Q} \;\middle|\; \begin{array}{c} Q \subset K \in \mathcal{K} \text{ and there are cells } Q_1, Q_2 \subset B'_R(Q_0) \\ \text{in } K \text{ such that } Q_1, Q, Q_2 \text{ are vertically aligned} \\ \text{and } Q \text{ lies vertically between } Q_1 \text{ and } Q_2 \end{array} \right\}.$$

Proposition 3.30 *Assume that M satisfies condition* (B) *or* (C). *Then there are choices for η, H and J and a constant $C_2 < \infty$ such that the following holds:*

For all $Q_0 \in \mathcal{Q}$ and all $R > 0$ we have

$$B'_R(Q_0) \subset P_R(Q_0) \subset \operatorname{Int} B'_{R+C_2}(Q_0) \cup \partial \widetilde{M}.$$

Moreover, there is a continuous map $b_{R,Q_0}\colon D^3 \to \widetilde{M}$ such that $b_{R,Q_0}(D^3) = P_R(Q_0)$ and $b_{R,Q_0}(S^2) = \partial P_R(Q_0)$ and b_{R,Q_0} is an embedding on $B^3 \cup (S^2 \setminus b^{-1}_{R,Q_0}(\widetilde{V}^{(1)}))$.

Finally, let $K_0 \in \mathcal{K}$ be the chamber that contains Q_0. Then for all cells $Q \subset B'_R(Q_0) \cap K_0$ we have $\operatorname{dist}^H_{K_0}(Q, Q_0), \operatorname{dist}^V_{K_0}(Q, Q_0) < C_2 R$.

Proof We will see that the proposition holds for $J = 11H$.

We first show that

$$B'_R(Q_0) \subset P_R(Q_0) \subset B'_{R+10H}(Q_0). \tag{3-5}$$

The first inclusion property is trivial. For the second inclusion property, consider a cell $Q \subset P_R(Q_0)$. Let $K \in \mathcal{K}$ be the chamber that contains Q and choose cells $Q_1, Q_2 \subset B'_R(Q_0) \cap K$ such that Q_1, Q and Q_2 are vertically aligned and Q lies vertically in between Q_1 and Q_2. Then $\operatorname{dist}(Q_1, Q_0) = \operatorname{dist}'(Q_1, Q_0) - J \operatorname{dist}^{\mathcal{K}}(K, K_0) < R - J \operatorname{dist}^{\mathcal{K}}(K_0, K)$ and, similarly, $\operatorname{dist}(Q_2, Q_0) < R - J \operatorname{dist}^{\mathcal{K}}(K, K_0)$. It follows

from Proposition 3.24 in case (C) and Proposition 3.28 in case (B) that $\operatorname{dist}(Q, Q_0) < R + 10H - J \operatorname{dist}^{\mathcal{K}}(K, K_0)$. So $\operatorname{dist}'(Q, Q_0) < R + 10H$ and (3-5) follows. In order to establish the inclusion property of this proposition, it hence suffices to choose C_2 larger than $10H + J$ plus the maximal number of cells that can intersect in one point.

Next, choose a sequence $K_1, K_2, \ldots \in \mathcal{K}$ such that $\mathcal{K} = \{K_0, K_1, K_2, \ldots\}$ and such that $\operatorname{dist}^{\mathcal{K}}(K_n, K_0)$ is nondecreasing in n. We will first show that the interior of $B'_R(Q_0) \cap (K_0 \cup \cdots \cup K_n)$ is connected for each $n \geq 0$: Fix n, choose a cell $Q \subset B'_R(Q_0) \cap (K_0 \cup \cdots \cup K_n)$ with $Q \neq Q_0$, let K_i be the chamber that contains Q and consider a combinatorially minimizing arc $\gamma\colon [0,1] \to \widetilde{M}$ from Q_0 to Q. We show by induction on the number of cells that intersect γ that $\operatorname{Int} Q$ lies in the same connected component of $\operatorname{Int}(B'_R(Q_0) \cap (K_0 \cup \cdots \cup K_n))$ as $\operatorname{Int} Q_0$. Let $Q' \in \mathcal{Q}$ be the cell that γ intersects prior to Q. If $Q' \subset K_i$, then we are done by the induction hypothesis, since then $\operatorname{dist}'(Q', Q_0) < \operatorname{dist}'(Q, Q_0)$ and hence $Q' \subset B'_R(Q_0) \cap (K_0 \cup \cdots \cup K_n)$. Assume next that $Q' \subset K_j \in \mathcal{K}$ for $j \neq i$ and hence γ crosses a wall $W = K_i \cap K_j \in \mathcal{W}$ in between Q' and Q. Then $\operatorname{dist}(Q', Q_0) = \operatorname{dist}(Q, Q_0) - H$ and $\operatorname{dist}^{\mathcal{K}}(K_j, K_0) = \operatorname{dist}^{\mathcal{K}}(K_i, K_0) \pm 1$. It suffices to consider the case in which $\operatorname{dist}^{\mathcal{K}}(K_j, K_0) = \operatorname{dist}^{\mathcal{K}}(K_i, K_0) + 1$, since otherwise we are again done by the induction hypothesis. In this case, γ must cross W twice and there is a cell $Q'' \subset K_i$ that γ intersects right before intersecting W for the first time. By Proposition 3.20(b2) in case (C) or Proposition 3.26(b3) in case (B), the arc γ only intersects cells that lie in K_j and that are adjacent to W between Q'' and Q'. Consider now all cells $Q^* \subset K_i$ that are adjacent to a cell $Q^{**} \subset K_j$ which intersects γ. For each such Q^* we have

$$\operatorname{dist}(Q^*, Q_0) \leq H + \operatorname{dist}(Q^{**}, Q_0) \leq H + \operatorname{dist}(Q', Q_0) = \operatorname{dist}(Q, Q_0)$$

and thus $\operatorname{dist}'(Q^*, Q_0) \leq \operatorname{dist}'(Q, Q_0)$ and $Q^* \subset B'_R(Q_0) \cap K_i$. It follows that Q' and Q'' lie in the same connected component of $B'_R(Q_0) \cap K_i$. This finishes the induction argument.

So also the interior of $P_R(Q_0) \cap (K_0 \cup \cdots \cup K_n)$ is connected for all $n \geq 0$. We will now show by induction on n that there is a continuous map $b_n\colon D^3 \to \widetilde{M}$ whose image is equal to the closure of this interior and which is an embedding when restricted to $B^3 \cup (S^2 \setminus b_n^{-1}(\widetilde{V}^{(1)}))$. For $n = 0$ this statement follows immediately from Lemma 3.29 and the fact that the interior of $P_R(Q_0) \cap K_0$ is connected. Assume now that $n \geq 1$. There is a unique $i \in \{1, \ldots, n-1\}$ such that K_i is adjacent to K_n. So $\operatorname{dist}^{\mathcal{K}}(K_i, K_0) = \operatorname{dist}^{\mathcal{K}}(K_n, K_0) - 1$. Let $W = K_i \cap K_n \in \mathcal{W}$ be the wall between K_i

and K_n. Observe that for every cell $Q \subset P_R(Q_0) \cap K_n$ that is adjacent to W and every cell $Q' \subset K_i$ that is adjacent to Q we have $\operatorname{dist}(Q', Q_0) \leq \operatorname{dist}(Q, Q_0) + H$. So, by (3-5),

$$\operatorname{dist}'(Q', Q_0) \leq \operatorname{dist}'(Q, Q_0) + H - J < R + 11H - J = R.$$

Hence, $Q' \subset B'_R(Q_0) \subset P_R(Q_0)$. This implies

$$\overline{P_R(Q_0) \cap \operatorname{Int} K_n} \cap W \subset P_R(Q_0) \cap W = \overline{P_R(Q_0) \cap \operatorname{Int} K_i} \cap W = b_{n-1}(D^3) \cap W.$$

By Lemma 3.29, $\overline{P_R(Q_0) \cap \operatorname{Int} K_n}$ is the union of the images of maps $b'\colon D^3 \to \widetilde{M}$ with the appropriate regularity properties and any two such images intersect in at most an edge of $\widetilde{V}$. Moreover, the preimage of W under every such map b' is a (connected) topological disk that is contained in $b_{n-1}(D^3) \cap W$. So we can combine b_{n-1} with the maps b' to obtain a map whose image is equal to the closure of the interior of $P_R(Q_0) \cap (K_0 \cup \cdots \cup K_n)$. Smoothing this map in the interior of D^3 yields b_n. This finishes the induction and proves the second assertion of the proposition for large n.

Finally, we show the last statement. Let $Q \subset B'_R(Q_0) \cap K_0$. Then $\operatorname{dist}(Q, Q_0) < R$. Consider a minimizing arc $\gamma\colon [0,1] \to \widetilde{M}$ between Q_0 and Q. By Proposition 3.20 or Proposition 3.26 the arc γ stays within the union of K_0 with the chambers that are adjacent to K_0. Let $Q_0, Q_1, \dots, Q_n = Q' \subset K_0$ be the cells of K_0 that γ intersects in that order. Then for all $i = 0, \dots, n-1$ either $\operatorname{dist}^H_{K_0}(Q_i, Q_{i+1}) + \operatorname{dist}^V_{K_0}(Q_i, Q_{i+1}) = 1 \leq \eta^{-1} \operatorname{dist}(Q_i, Q_{i+1})$ or γ intersects a wall $W \subset \partial K$ right after Q_i and right before Q_{i+1}. In this case let $K' \in \mathcal{K}$ be the chamber on the other side of W and let $Q'_i, Q'_{i+1} \subset K'$ be the cells that γ intersects right after Q_i and right before Q_{i+1}. By Lemma 3.13(f) we have

$$\operatorname{dist}^H_{K_0}(Q_i, Q_{i+1}),\ \operatorname{dist}^V_{K_0}(Q_i, Q_{i+1}) < C_0 \eta^{-1} \operatorname{dist}(Q'_i, Q'_{i+1}) + C_0.$$

If $H > 1$, then the right-hand side is bounded by $C_0 \eta^{-1} \operatorname{dist}(Q_i, Q_{i+1})$. The rest follows from the triangle inequality for $\operatorname{dist}^H_{K_0}$ and $\operatorname{dist}^V_{K_0}$ with $C_2 > C_0 \eta^{-1}$. □

We can finally establish Proposition 3.4 and hence Proposition 3.2(a) (see Section 3.2).

Proof of Proposition 3.4 By Proposition 3.5, we may assume that M satisfies condition (B) or (C).

Observe first that the universal covering $\pi\colon \widetilde{M} \to M$ can be viewed as the restriction of the universal covering $\pi\colon \widetilde{M}_0 \to M_0$ to a component of $\pi^{-1}(M)$. Consider the simplicial complex $V \subset M$ as defined in Section 3.3 and let $f_0\colon V \to M$ be the

inclusion map. Recall that f_0 lifts to the inclusion map $\tilde{f}_0\colon \widetilde{V} \to \widetilde{M}$ in the universal covering $\pi\colon \widetilde{M} \to M$. Consider the Riemannian metric g on M_0 and the map $f\colon V \to M$ from the assumptions of the Proposition. Let $H\colon V \times [0,1] \to M_0$ be the homotopy between f_0 and f and let L be a strict upper bound on the length of the arcs $t \mapsto H(x,t)$ (note that V is compact). Since this homotopy leaves ∂V invariant and embedded in ∂M, we can extend H to a homotopy $H^*\colon (V \cup \partial M) \times [0,1] \to M_0$ between the inclusion map $f_0^*\colon V \cup \partial M \to M$ and the extension $f^*\colon V \cup \partial M \to M_0$ of f such that $H^*(\cdot, t)$ restricted to ∂M is the identity for all $t \in [0,1]$. Here we view $V \cup \partial M$ as a connected simplicial complex. The homotopy H^* can be lifted to a homotopy $\widetilde{H}^*\colon (\widetilde{V} \cup \partial \widetilde{M}) \times [0,1] \to \widetilde{M}_0$ between the inclusion map $\tilde{f}_0\colon \widetilde{V} \cup \widetilde{M} \to \widetilde{M}$ and a lift $\tilde{f}^*\colon \widetilde{V} \cup \partial \widetilde{M} \to \widetilde{M}_0$ of f^*, that is, $f^* \circ \pi|_{\widetilde{V} \cup \partial \widetilde{M}} = \pi \circ \tilde{f}^*$. Note that $\pi(\widetilde{H}^*(x,t)) = H^*(\pi(x),t)$ for all $(x,t) \in (\widetilde{V} \cup \partial \widetilde{M}) \times [0,1]$. Still, the lengths of the arcs $t \mapsto \widetilde{H}^*(x,t)$ are bounded by L.

Consider the solid torus $S \subset \operatorname{Int} M$ and pick a component $\widetilde{S} \subset \pi^{-1}(S) \subset \operatorname{Int} \widetilde{M}$. Fix a diffeomorphism $\Phi\colon S^1 \times D^2 \to S$ and an orientation on the S^1–factor and denote by $\sigma = \Phi(S^1 \times \{0\}) \subset S$ the core of S. By our assumptions, $\widetilde{S} \approx \mathbb{R} \times D^2$ and $\pi|_{\widetilde{S}}\colon \widetilde{S} \to S$ is a universal covering of S. So we can find a lifted diffeomorphism $\widetilde{\Phi}\colon \mathbb{R} \times D^2 \to \widetilde{S}$ such that $\pi \circ \widetilde{\Phi} = \Phi \circ \pi_{S^1 \times D^2}$, where $\pi_{S^1 \times D^2}\colon \mathbb{R} \times D^2 \to S^1 \times D^2$, $(u,x) \mapsto (e^{2\pi i u}, x)$ is the standard universal covering map. Let $F = \widetilde{\Phi}([0,1] \times D^2) \subset \widetilde{S}$. Then $\pi(F) = S$ and π restricted to the interior of F is injective. So F is a fundamental domain for the universal covering $\pi|_{\widetilde{S}}\colon \widetilde{S} \to S$. The central loop $\sigma \approx S^1 \times \{0\} \subset S \approx S^1 \times D^2$ induces a deck transformation $\varphi\colon \widetilde{M}_0 \to \widetilde{M}_0$, which is an isometry and $\widetilde{S}$ is covered by fundamental domains of the form $\varphi^{(n)}(F)$ where $n \in \mathbb{Z}$. Observe also that $\widetilde{\sigma} = \widetilde{\Phi}(\mathbb{R} \times \{0\}) = \pi^{-1}(\sigma) \cap \widetilde{S}$ is a properly embedded, infinite line, which is invariant under φ.

Choose a chamber $K_0 \in \mathcal{K}$ for which the displacement $\operatorname{dist}^{\mathcal{K}}(K_0, \varphi(K_0))$ is minimal. Next, if $\varphi(K_0) = K_0$ choose a column $E_0 \in \mathcal{E}_{K_0}$ for which the displacement $\operatorname{dist}^{H}_{K_0}(E_0, \varphi(E_0))$ is minimal. If $\varphi(K_0) \neq K_0$, the column $E_0 \in \mathcal{E}_{K_0}$ can be chosen arbitrarily. Finally, choose an arbitrary cell $Q_0 \subset E_0$. We will now show that there is a universal constant $c > 0$, which only depends on the structure of V (and not on S!) such that, for all $n \in \mathbb{Z}$,

$$\operatorname{dist}'(Q_0, \varphi^{(n)}(Q_0)) \geq c|n|. \tag{3-6}$$

If $\varphi(K_0) \neq K_0$, then we argue as follows. Consider the minimal chain between K_0 and $\varphi(K_0)$ in the adjacency graph of $\mathcal{K}$. The images of this minimal chain under the

deck transformations $\varphi^{(0)},\dots,\varphi^{(n-1)}$ are each minimal and can be concatenated along $\varphi^{(1)}(K_0),\dots,\varphi^{(n-1)}(K_0)$ to a chain between K_0 and $\varphi^{(n)}(K_0)$. We now claim that this chain is minimal. Otherwise, there are elements in this chain that occur at least twice. Since the adjacency graph of $\mathcal{K}$ is a tree (see Lemma 3.8), there must then be even two consecutive elements in this chain that are equal. These two elements can only come from two distinct images of the minimal chain between K_0 and $\varphi(K_0)$. So if $K_0', K_0''\in\mathcal{K}$ are the second and second-to-last elements on this minimal chain then we must have $\varphi^{(i+1)}(K_0')=\varphi^{(i)}(K_0'')$ for some $i\in\{0,\dots,n-1\}$. But this would imply that $\varphi(K_0')=K_0''$ and hence $\operatorname{dist}^{\mathcal{K}}(K_0',\varphi(K_0'))=\operatorname{dist}^{\mathcal{K}}(K_0',K_0'')=\operatorname{dist}^{\mathcal{K}}(K_0,\varphi(K_0))-2$, contradicting the minimal choice of K_0. So we conclude that the chain in question is minimal and hence $\operatorname{dist}^{\mathcal{K}}(K_0,\varphi^{(n)}(K_0))\ge|n|$ for all $n\in\mathbb{Z}$, which establishes (3-6) assuming $c<H+J$.

If $\varphi(K_0)=K_0$ but $\varphi(E_0)\neq E_0$, then we can draw the same conclusions for $\mathcal{E}_{K_0}$ instead of $\mathcal{K}$ and obtain $\operatorname{dist}^H_{K_0}(Q_0,\varphi^{(n)}(Q_0))\ge|n|$ for all $n\in\mathbb{Z}$. If $\varphi(E_0)=E_0$, then $\operatorname{dist}^V_{K_0}(Q_0,\varphi^{(n)}(Q_0))\ge|n|$ for all $n\in\mathbb{Z}$. So, in the latter two cases, (3-6) follows by the last assertion of Proposition 3.30.

Let $N\ge1$ be some large natural number whose value we will determine at the end of the proof. It will depend on g, f and L. The sets $\widetilde{S}_+=F\cup\varphi(F)\cup\cdots\varphi^{(N-1)}(F)$ and $\widetilde{S}_-=\varphi^{(-1)}(F)\cup\cdots\varphi^{(-N)}(F)$ are each diffeomorphic to a solid cylinder $\approx I\times D^2$ and are bounded by annuli inside $\partial\widetilde{S}$ as well as disks D_0, D_+ and D_0, D_-, where $D_+=\varphi^{(N)}(D_0)$ and $D_-=\varphi^{(-N)}(D_0)$. Let $\widetilde{\sigma}_+$ and $\widetilde{\sigma}_-$ be the subsegments of $\widetilde{\sigma}$ that connect D_0 with D_+ and D_0 with D_-, ie $\widetilde{\sigma}_+=\widetilde{\sigma}\cap\widetilde{S}_+$ and $\widetilde{\sigma}_-=\widetilde{\sigma}\cap\widetilde{S}_-$.

Choose $R_0>0$ large enough that $B'_{R_0}(Q_0)$ contains all points of $\widetilde{M}$ that have distance at most L from F. Then, for all $n\in\mathbb{Z}$, the set $B'_{R_0}(\varphi^{(n)}(Q_0))$ contains all points of $\widetilde{M}$ that have distance at most L from $\varphi^{(n)}(F)$. Consider for the moment some number R such that

$$R_0\le R\le cN-C_2-R_0.$$

(C_2 is the constant from Proposition 3.30.) Then we have

$$B'_{R_0}(Q_0)\subset B'_R(Q_0)\subset P_R(Q_0),$$

so every point of $\widetilde{M}$ that has distance at most L from $D_0\subset F$ is contained in $P_R(Q_0)$. We now claim that for all $n\in\mathbb{Z}$ with $|n|\ge N$ the set $B'_{R_0}(\varphi^{(n)}(Q_0))$ is disjoint from the interior of $B'_{R+C_2}(Q_0)$. Assume not and let $Q'\subset B'_{R_0}(\varphi^{(n)}(Q_0))\cap B'_{R+C_2}(Q_0)$ be a cell in the intersection. Then we obtain the following contraction using (3-6):

$$c|n|\le\operatorname{dist}'(Q_0,\varphi^{(n)}(Q_0))\le\operatorname{dist}'(Q_0,Q')+\operatorname{dist}'(Q',\varphi^{(n)}(Q_0))<R_0+R+C_2\le cN.$$

So $B'_{R_0}(\varphi^{(n)}(Q_0))$ is disjoint from $\operatorname{Int} B'_{R+C_2}(Q_0)$ and hence also from $P_R(Q_0) \subset \operatorname{Int} B'_{R+C_2}(Q_0) \cup \partial\widetilde{M}$. So $P_R(Q_0)$ does not contain any point of $\widetilde{M}$ that has distance at most L from $\widetilde{S} \setminus (\widetilde{S}_+ \cup \widetilde{S}_-)$ and thus also from D_+ or D_-. This implies in particular that the arcs $\widetilde{\sigma}_+$ and $\widetilde{\sigma}_-$ have intersection number 1 with the restriction $b_{R,Q_0}|_{S^2}\colon S^2 \to \widetilde{M}$, whose image is $\partial P_R(Q_0)$ (see Proposition 3.30).

Our conclusions imply that the homotopy $\widetilde{H}^*$ restricted to $\partial P_R(Q_0)$ does not intersect $D_0 \cup D_+ \cup D_-$ or, more generally, that it stays away from $\widetilde{S} \setminus (\widetilde{S}_- \cup \widetilde{S}_+)$. If we view $b_{R,Q_0}|_{S^2}$ as a map from S^2 to $\widetilde{V} \cup \partial\widetilde{M}$, then $(x,t) \mapsto \widetilde{H}^*(b_{R,Q_0}|_{S^2}(x),t)$ is a homotopy between

$$b_{R,Q_0}|_{S^2} = \tilde{f}_0^* \circ b_{R,Q_0}|_{S^2}\colon S^2 \to \widetilde{M} \quad \text{and} \quad s_{R,Q_0} = \tilde{f}^* \circ b_{R,Q_0}|_{S^2}\colon S^2 \to \widetilde{M}_0$$

whose image is disjoint from $D_0 \cup D_+ \cup D_-$ and $\widetilde{S} \setminus (\widetilde{S}_- \cup \widetilde{S}_+)$. So s_{R,Q_0} has intersection number 1 with $\widetilde{\sigma}_+$ and $\widetilde{\sigma}_-$. Choose a small perturbation $s'_{R,Q_0}\colon S^2 \to \widetilde{M}_0$ of s_{R,Q_0} that intersects $\partial\widetilde{S}$ transversally, that still stays away from D_0, D_+, D_- and $\widetilde{S} \setminus (\widetilde{S}_- \cup \widetilde{S}_+)$ and that satisfies

$$\text{area}\, s'_{R,Q_0}\big|_{s'^{-1}_{R,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)} < 2\,\text{area}\, s_{R,Q_0}\big|_{s^{-1}_{R,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)}. \tag{3-7}$$

(This can always be achieved by perturbing the composition of s_{R,Q_0} with a diffeomorphism of $\widetilde{M}_0$ which slightly expands $\widetilde{S}$.) Set

$$X = s'^{-1}_{R,Q_0}(\widetilde{S}), \quad X_+ = s'^{-1}_{R,Q_0}(\widetilde{S}_+), \quad X_- = s'^{-1}_{R,Q_0}(\widetilde{S}_-).$$

Then X, X_+ and X_- are compact smooth domains of S^2 and we have $X = X_+ \cup X_-$, $s'(\partial X) \subset \partial\widetilde{S}$ and that s'_{R,Q_0} restricted to X_+ and X_- has nonzero intersection number with $\widetilde{\sigma}$. Let $X'_+ \subset X_+$ be the union of all components of X_+ on which s'_{R,Q_0} has nonzero intersection number with $\widetilde{\sigma}$, define $X'_- \subset X_-$ analogously and set $X' = X'_+ \cup X'_-$. Then $X'_+, X'_- \neq \varnothing$ and $X'_+, X'_- \neq S^2$ and every component $Y \subset X'$ is bounded by at least one circle $Z \subset \partial Y$ such that $s'_{R,Q_0}|_Z\colon Z \to \partial\widetilde{S}$ is noncontractible in $\partial\widetilde{S}$. Each such circle Z bounds two disks $E_1, E_2 \subset S^2$ on either side (one of these disks contains Y and the other one doesn't). Consider now the set of all such disks $E \subset S^2$ coming from all components Y of X' and all boundary circles $Z \subset \partial Y$ for which $s'_{R,Q_0}|_Z\colon Z \to \partial\widetilde{S}$ is noncontractible in $\partial\widetilde{S}$. Any two such disks are either disjoint or one is contained in the other. We can hence choose a component $Y \subset X'$, a boundary circle $Z \subset \partial X'$ with the aforementioned property and a disk $E \subset S^2$ bounded by Z such that E is minimal with respect to inclusion. We argue that s'_{R,Q_0} restricted to every other boundary circle $Z' \subset \partial Y$ is contractible in $\widetilde{S}$: If this was not

the case, then Y must be disjoint from the interior of E, since otherwise $Z' \subset Y \subset E$ bounds a disk $E' \subsetneq E$. By the same argument, E cannot contain any other component Y' of X', because otherwise we would find a boundary circle $Z'' \subset \partial Y' \subset \operatorname{Int} E$ such that $s'_{R,Q_0}|_{Z''}$ is noncontractible in $\partial \widetilde{S}$. So $\operatorname{Int} E$ must be disjoint from X' and hence $s'_{R,Q_0}|_E$ describes a nullhomotopy of a noncontractible loop in $\partial \widetilde{S}$, which does not intersect $\widetilde{\sigma}$. Since $\pi_2(\widetilde{M}_0) = \pi_2(M_0) = 0$, this nullhomotopy can be homotoped relative boundary to a nullhomotopy that has nonzero intersection number with $\widetilde{\sigma}$. This is however impossible and we obtain a contradiction. So s'_{R,Q_0} restricted to all other boundary components of Y is noncontractible in $\widetilde{S}$ and hence we have shown that $\Sigma = Y$ and $h = \pi \circ s'_{R,Q_0}|_Y$ satisfy all the claims of the proposition except for the area bound.

In view of (3-7) it remains to choose R and N such that area $s_{R,Q_0}|_{s^{-1}_{R,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)}$ can be bounded by a uniform multiple of area f. To do this choose radii $R_i = R_0 + C_2 i$ where $i = 0, \dots, e$ with $e = \lfloor C_2^{-1}(cN - C_2 - 2R_0) \rfloor$. Then

$$R_0 < R_1 < \cdots < R_e \le cN - C_2 - R_0.$$

By Proposition 3.30 we know that $\partial P_{R_0}(Q_0) \setminus \partial \widetilde{M}, \dots, \partial P_{R_e}(Q_0) \setminus \partial \widetilde{M} \subset \widetilde{V} \subset \widetilde{M}$ are pairwise disjoint. So, since $b_{R_i,Q_0}(S^2) = \partial P_{R_i}(Q_0)$ and $s_{R_i,Q_0} = \tilde{f}^* \circ b_{R_i,Q_0}|_{S^2}$ and $\tilde{f}^*(\partial \widetilde{M}) = \partial \widetilde{M}$, we have

$$\operatorname{area} s_{R_0,Q_0}\big|_{s^{-1}_{R_0,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)} + \cdots + \operatorname{area} s_{R_e,Q_0}\big|_{s^{-1}_{R_e,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)} \le \operatorname{area} \tilde{f}^*|_{\tilde{f}_0^{*-1}(\widetilde{S}_+ \cup \widetilde{S}_-)}.$$

Since $\widetilde{S}_+ \cup \widetilde{S}_- = \varphi^{(-N)}(F') \cup \cdots \cup \varphi^{(N-1)}(F') \cup D_-$ for the half-open set $F' = F \setminus D_0$, we further have

$$\operatorname{area} \tilde{f}^*|_{\tilde{f}^{*-1}(\widetilde{S}_+ \cup \widetilde{S}_-)} = \operatorname{area} \tilde{f}^*|_{\tilde{f}^{*-1}(\varphi^{(-N)}(F'))} + \cdots + \operatorname{area} \tilde{f}^*|_{\tilde{f}^{*-1}(\varphi^{(N-1)}(F'))}.$$

Observe now that if $\tilde{f}^*(x) = \tilde{f}^*(y)$ and $\pi(x) = \pi(y)$ for $x, y \in \widetilde{V}$, then $x = y$, since the arcs $t \mapsto \widetilde{H}^*(x,t), \widetilde{H}^*(y,t)$ have the same endpoint and project to the same arc $t \mapsto H^*(\pi(x), t)$ under π. So, for all $n \in \mathbb{Z}$, the projection π restricted to $\tilde{f}^{*-1}(\varphi^{(n)}(F'))$ is injective. Since $\pi(\tilde{f}^{*-1}(\varphi^{(n)}(F'))) \subset f^{*-1}(S)$, we conclude $\operatorname{area} \tilde{f}^*|_{\tilde{f}^{*-1}(\varphi^{(n)}(F'))} \le \operatorname{area} f^*|_{f^{*-1}(S)} < \operatorname{area} f$. Putting all this together yields

$$\operatorname{area} s_{R_0,Q_0}\big|_{s^{-1}_{R_0,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)} + \cdots + \operatorname{area} s_{R_e,Q_0}\big|_{s^{-1}_{R_e,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)} < 2N \operatorname{area} f.$$

So we can find an index $i \in \{0, \dots, e\}$ such that

$$\operatorname{area} s_{R_i,Q_0}\big|_{s^{-1}_{R_i,Q_0}(\widetilde{S}_+ \cup \widetilde{S}_-)} < \frac{2N}{e+1} \operatorname{area} f \le \frac{2C_2 N}{cN - C_2 - 2R_0} \operatorname{area} f.$$

Choosing $N > 2c^{-1}(C_2 + 2R_0)$ yields

$$\text{area}\, s_{R_i,Q_0}\big|_{s^{-1}_{R_i,Q_0}(\tilde{S}_+\cup\tilde{S}_-)} < 4c^{-1}C_2 \text{ area } f.$$

This finishes the proof of Proposition 3.4. □

3.7 The case in which M is covered by a T^2–bundle over a circle

We finally present the proof of Proposition 3.2(b).

Lemma 3.31 *Let $A \in \mathrm{SL}(2,\mathbb{Z})$ be a 2×2–matrix with integral entries and determinant 1. Then for every $k \geq 1$ there is a number $1 \leq d \leq 6^k$ such that*

$$I + A + A^2 + \cdots + A^{d-1} \equiv 0 \mod 3^k.$$

(Here I is the identity matrix.)

Proof We first show the claim for $k = 1$. Since $\det A = 1$, the Cayley–Hamilton theorem yields that $I - (\operatorname{tr} A)A + A^2 = 0$. So we are done for $\operatorname{tr} A \equiv 2 \bmod 3$. If $\operatorname{tr} A \equiv 0 \bmod 3$, then $I + A + A^2 + A^3 = (I + A^2)(I + A) \equiv 0 \bmod 3$ and if $\operatorname{tr} A \equiv 1 \bmod 3$, then $I + \cdots + A^5 = (I - A + A^2)(I + 2A + 2A^2 + A^3) \equiv 0 \bmod 3$.

We now apply induction. Assume that the statement is true for all numbers up to $k \geq 1$. We will show that it then also holds for $k+1$. Choose $1 \leq d_1 \leq 6$ such that $I + A + \cdots + A^{d_1-1} \equiv 0 \bmod 3$. By the induction hypothesis applied to $A^{d_1} \in \mathrm{SL}(2,\mathbb{Z})$, there is a number $1 \leq d_2 \leq 6^k$ such that $I + A^{d_1} + A^{2d_1} + \cdots + A^{(d_2-1)d_1} \equiv 0 \bmod 3^k$. So

$$\begin{aligned} I + A + \cdots + A^{d_1 d_2 - 1} &= (I + A + \cdots + A^{d_1-1})(I + A^{d_1} + \cdots + A^{(d_2-1)d_1}) \\ &\equiv 0 \mod 3^{k+1}, \end{aligned}$$

and $1 \leq d_1 d_2 \leq 6^{k+1}$. □

Lemma 3.32 *Assume that M is the total space of a T^2–bundle over a circle. Then for every $n \geq 1$ there is a finite covering map $\pi_n\colon M \to M$ with the same domain and range such that for every embedded loop $\sigma \subset M$ the preimage $\pi_n^{-1}(\sigma)$ consists of at least n loops.*

Proof The manifold M is diffeomorphic to a mapping torus of an orientation-preserving diffeomorphism $\phi\colon T^2 \to T^2$. The diffeomorphism ϕ acts on $\pi_1(T^2) \cong \mathbb{Z}^2$ by an element in $A \in \mathrm{SL}(2,\mathbb{Z})$. The fundamental group $\pi_1(M)$ is isomorphic to a semidirect product of $\mathbb{Z}$ with $\mathbb{Z}^2$ coming from the action of A on $\mathbb{Z}^2$. So $\pi_1(M)$ can be identified with $\mathbb{Z}^2 \times \mathbb{Z}$ with the multiplication

$$\left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix}, z_1\right) \cdot \left(\begin{pmatrix} x_2 \\ y_2 \end{pmatrix}, z_2\right) = \left(\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + A^{z_1} \begin{pmatrix} x_2 \\ y_2 \end{pmatrix}, z_1 + z_2\right).$$

Since for any $m \geq 1$ the lattice $m\mathbb{Z}^2 \subset \mathbb{Z}^2$ is preserved by the action of A, the subset

$$U_m = \left\{\left(\begin{pmatrix} mx \\ my \end{pmatrix}, z\right) \mid x, y, z \in \mathbb{Z}\right\} \subset \pi_1(M)$$

is a subgroup of $\pi_1(M)$ of index m^2. Note that U_m is isomorphic to $\pi_1(M)$, so if we consider the corresponding m^2–fold covering $\pi'_m\colon \widehat{M}_m \to M$, then $\widehat{M}_m$ is diffeomorphic to M.

It remains to compute the number of components of ${\pi'}_m^{-1}(\sigma)$ and to show that this number can be made arbitrarily large for the right choice of m. Set $m = 3^k$ for some $k \geq 1$, which we will determine later. Let $\hat{\sigma} \subset \pi_m'^{-1}(\sigma)$ be an arbitrary loop in the preimage of σ. Then we can find an element $g = ((x, y), z) \in \pi_1(M)$ in the conjugacy class of $[\sigma]$ such that σ represents g in M and such that $\hat{\sigma}$ represents a multiple of g that is contained in $U_m = \pi_1(\widehat{M}_m)$ in $\widehat{M}_m$. Then the restriction $\pi'_m|_{\hat{\sigma}}\colon \hat{\sigma} \to \sigma$ is a covering of a circle and its degree is equal to the first exponent $d_0 \geq 1$ for which $g^{d_0} \in U_m$. We will show that $d_0 \leq 6^k$. To do this observe that, for all $i \geq 1$,

$$g^i = \left(\begin{pmatrix} x \\ y \end{pmatrix} + A^z \begin{pmatrix} x \\ y \end{pmatrix} + \cdots + A^{(i-1)z} \begin{pmatrix} x \\ y \end{pmatrix}, iz\right).$$

By Lemma 3.31, there is a number $1 \leq d \leq 6^k$ such that the first two entries of g^d are divisible by $m = 3^k$ and hence $g^d \in U_m$. This implies the desired bound.

Since the choice of $\hat{\sigma}$ was arbitrary, we conclude that every loop in ${\pi'}_m^{-1}(\sigma)$ covers σ at most 6^k times and hence the number of such loops is at least

$$\frac{m^2}{6^k} = \frac{3^{2k}}{6^k} = \left(\frac{3}{2}\right)^k.$$

So choosing k such that $\left(\frac{3}{2}\right)^k > n$ yields the desired result. □

Proof of Proposition 3.2(b) We only have to consider the case in which M is a T^2–bundle over a circle, since for any finite cover $\hat{\pi}\colon \widehat{M} \to M$ we can compose the maps $\hat{f}_1, \hat{f}_2, \ldots\colon V \to \widehat{M}$ obtained for $\widehat{M}$ with $\hat{\pi}$ to obtain the maps $f_1 = \hat{\pi} \circ \hat{f}_1$, $f_2 = \hat{\pi} \circ \hat{f}_2, \ldots$.

We first establish the assertion for the case $n = 1$. Fix a T^2–bundle projection $p\colon M \to S^1$ and let $T = p^{-1}(\{\mathrm{pt}\}) \subset M$ be a torus fiber. Then $M \setminus T$ is diffeomorphic to $T^2 \times (0, 1)$ and we can find a smooth local diffeomorphism $\Phi\colon T^2 \times [0, 1] = S^1 \times S^1 \times [0, 1] \to M$ such that $\Phi|_{T^2 \times (0,1)}$ is a diffeomorphism onto $M \setminus T$ and such that Φ restricted to $T^2 \times \{0\}$ and $T^2 \times \{1\}$ is a diffeomorphism onto T. Moreover,

we may assume that Φ is chosen in such a way that the map

$$(\Phi|_{T^2\times\{1\}})^{-1}\circ\Phi|_{T^2\times\{0\}}\colon T^2\times\{0\}\to T^2\times\{0\}$$

is affine with respect to the standard affine structure on $T^2=S^1\times S^1$. Now let $z\in S^1$ be an arbitrary point and set

$$V=T\cup\Phi(S^1\times\{z\}\times[0,1])\cup\Phi(\{z\}\times S^1\times[0,1]).$$

Then $M\setminus V$ is diffeomorphic to a 3–ball and V can be given the structure of a simplicial complex. Fix such a structure for the rest of the proof and let $f_1\colon V\to M$ be the inclusion map of V.

Consider the universal covering $\pi\colon\widetilde{M}\to M$ and choose a lift $F\subset\widetilde{M}$ of $M\setminus V$. Then F is a connected fundamental domain. Furthermore, denote by $\widetilde{V}=\widetilde{\pi}^{-1}(V)\subset\widetilde{M}$ the preimage of V and by $\tilde{f}_1\colon\widetilde{V}\to\widetilde{M}$ its inclusion map. The complement of $\widetilde{V}$ in $\widetilde{M}$ consists of open sets whose closures $Q\subset\widetilde{M}$ are finite polyhedra and which we call *cells*. Denote the set of cells again by $\mathcal{Q}$. Observe that every cell is the image of F under a deck transformation of $\pi\colon\widetilde{M}\to M$. We say that two cells $Q_1,Q_2\in\mathcal{Q}$ are adjacent if their intersection contains a point of $\widetilde{V}\setminus\widetilde{V}^{(1)}$. Choose a cell $Q_0\in\mathcal{Q}$ and consider for each $k\geq 0$ the union $B_k(Q_0)$ of all cells that have distance $\leq k$ in the adjacency graph of $\mathcal{Q}$. Then $S_k=\partial B_k(Q_0)\subset\widetilde{V}$ is the image of a continuous map $s_k\colon\Sigma_k\to\widetilde{V}$ where Σ_k is an orientable surface such that s_k is an embedding on $\Sigma_k\setminus s_k^{-1}(\widetilde{V}^{(1)})$.

Choose a component $\widetilde{\sigma}\subset\pi^{-1}(\sigma)$ that intersects Q_0. Then $\widetilde{\sigma}\subset\widetilde{M}$ is a noncompact, properly embedded line and there is a noncompact ray $\widetilde{\sigma}^+\subset\widetilde{\sigma}$ that starts in Q_0. This implies that $\widetilde{\sigma}^+$ has nonzero intersection number with the map $\tilde{f}_1\circ s_k\colon\Sigma_k\to\widetilde{M}$ for each $k\geq 1$.

Consider now the continuous map $f_1'\colon V\to M$ that is homotopic to $f_1\colon V\to M$ via a homotopy $H\colon V\times[0,1]\to M$ and use this homotopy to construct a lift $\tilde{f}_1'\colon\widetilde{V}\to\widetilde{M}$. Let N be a bound on the number of cells that each arc of the form $t\mapsto H(\,\cdot\,,t)$ intersects. Then H induces a homotopy from $\tilde{f}_1\circ s_N\colon\Sigma_N\to\widetilde{M}$ to $\tilde{f}_1'\circ s_N\colon\Sigma_N\to\widetilde{M}$, which is disjoint from Q_0. So both maps have the same, nonzero, intersection number with $\widetilde{\sigma}^+$. We conclude that $\tilde{f}_1'(s_N(\Sigma_N))\subset\tilde{f}_1'(\widetilde{V})$ intersects $\widetilde{\sigma}$. Hence, $f_1'(V)$ intersects σ.

We finally show the assertion for all remaining $n\geq 2$. Fix n, consider the covering map $\pi_n\colon M\to M$ from Lemma 3.32 and set

$$f_n=\pi_n\circ f_1\colon V\to M.$$

Moreover, the preimage $\sigma_n=\pi_n^{-1}(\sigma)$ is the union of at least n loops which all have the property that all its nontrivial multiples are noncontractible in M.

Consider a map $f_n'\colon V \to M$ and a homotopy between f_n and f_n'. This homotopy can be lifted via $\pi_n\colon M \to M$ to a homotopy between f_1 and a map $f_1'\colon V \to M$ such that $f_n' = \pi_n \circ f_1'$. We now have

$$f'^{-1}_n(\sigma) = f'^{-1}_1(\pi_n^{-1}(\sigma)) = f'^{-1}_1(\sigma_n) = \bigcup_{\sigma' \subset \sigma_n} f'^{-1}_n(\sigma'),$$

where the last union is to be understood as the union over all loops σ' of σ_n. By our previous conclusion, $f'^{-1}_1(\sigma') \neq \varnothing$ for all such σ' and all such sets are pairwise disjoint. This proves the desired result. □

References

[0] **R H Bamler**, *Long-time behavior of 3–dimensional Ricci flow: introduction*, Geom. Topol. 22 (2018) 757–774

[A] **R H Bamler**, *Long-time behavior of 3–dimensional Ricci flow, A: Generalizations of Perelman's long-time estimates*, Geom. Topol. 22 (2018) 775–844

[B] **R H Bamler**, *Long-time behavior of 3–dimensional Ricci flow, B: Evolution of the minimal area of simplicial complexes under Ricci flow*, Geom. Topol. 22 (2018) 845–892

[D] **R H Bamler**, *Long-time behavior of 3–dimensional Ricci flow, D: Proof of the main results*, Geom. Topol. 22 (2018) 949–1068

[1] **K Asano**, *Homeomorphisms of prism manifolds*, Yokohama Math. J. 26 (1978) 19–25 MR

[2] **D Faessler**, *On the topology of locally volume collapsed Riemannian 3–orbifolds*, preprint (2011) arXiv

[3] **A Hatcher**, *Notes on basic 3–manifold topology*, book project (2007) Available at `http://www.math.cornell.edu/~hatcher/3M/3Mdownloads.html`

[4] **J Luecke**, **Y-Q Wu**, *Relative Euler number and finite covers of graph manifolds*, from "Geometric topology" (W H Kazez, editor), AMS/IP Stud. Adv. Math. 2, Amer. Math. Soc., Providence, RI (1997) 80–103 MR

[5] **J Morgan**, **G Tian**, *Completion of the proof of the geometrization conjecture*, preprint (2008) arXiv

Department of Mathematics, University of California
Berkeley, CA, United States
`rbamler@math.berkeley.edu`
`https://math.berkeley.edu/~rbamler/`

Proposed: Tobias H. Colding
Seconded: Bruce Kleiner, Gang Tian

Received: 16 December 2014
Revised: 21 January 2016