On graphs, homology bases, and triangulated homology spheres

Karim Adiprasito[*], Marc Lackenby[†], Juan Souto[‡], and Geva Yashfe[§]

Abstract

We describe a construction that takes as input a graph and a basis for its first homology, and returns a triangulation of a 3-dimensional homology sphere. This makes precise an idea of M. Gromov and A. Nabutovski described in [7]. The immediate application, essentially described by Gromov, is to translate problems about asymptotics of homology sphere triangulations to asymptotic counting problems for constant-degree graphs with "short" homology bases. We construct families of 3-sphere triangulations with dual graphs that are expanders, answering a relaxation of a question asked by G. Kalai in [9, 10]. Our results also imply that if the number of d-dimensional triangulated homology spheres with n facets is superexponential in n for some d then the same holds for $d = 3$.

1 Introduction

1.1 Motivation and main results

This paper is about constructing triangulations of the 3-sphere and of 3-dimensional homology spheres. This is interesting both from an enumerative and from a more structural point of view. In [7, p. 33], Mikhail Gromov posed the following problem:

Problem 1. Fix a smooth d-manifold X and let $t(X, N)$ be the number of combinatorial isomorphism types of triangulations of X with N d-simplices. Does t grow at most exponentially in N?

Remark 2. It is known and not difficult to prove that if $d \geq 2$ and X is a PL d-manifold then $t(X, N)$ is bounded below by an exponential function of N for all $d \geq 2$, and bounded above by a function of the form $c^{n \log n}$. This problem has a different character from the one in which we count triangulations by the number of vertices, for which there are faster-growing lower bounds (see [14])

With Alex Nabutovski, Gromov found a combinatorial analogue of Problem 1:

Problem 3. Evaluate the number $t_L(N)$ of connected 3-regular graphs G with N edges, such that cycles of length $\leq L$ normally generate $\pi_1(G)$ (a variation of the problem asks only that such cycles generate $H_1(G)$). Is $t_L(N)$ at most exponential in N?

[*]Sorbonne Université and Université Paris Cité, CNRS, IMJ-PRG, F-75005 Paris, France
[†]University of Oxford
[‡]CNRS, IRMAR - UMR 6625, Université de Rennes, Campus de Beaulieu, Rennes, 35042, France
[§]Hebrew University of Jerusalem and The University of Chicago

In [7, p. 33], Gromov states that the essential part of Problem 1 is to count triangulations of manifolds with fixed π_1 or fixed H_1, and that this is essentially equivalent to Problem 3. Our first main theorem formalizes part of this statement.

Theorem 4. *There exists $L_0 \in \mathbb{N}$ such that for all $L \geq L_0$ the following holds. Let $\mathbb{F}$ be a field and $d \geq 3$ an integer. Let $s_{d,\mathbb{F}}(N)$ be the number of isomorphism types of triangulated combinatorial d-manifolds M with $H_1(M;\mathbb{F}) = 0$ and with at most N d-simplices. Let $t_{L,\mathbb{F}}(N)$ be the number of connected 3-regular graphs G with at most N vertices such that cycles of length $\leq L$ generate $H_1(G;\mathbb{F})$. Then there exists $c > 0$ such that*

$$s_{d,\mathbb{F}}(\frac{1}{c}N) \leq t_{L,\mathbb{F}}(N) \leq s_{d,\mathbb{F}}(cN)$$

for all $N \in \mathbb{N}$.

This is deduced from a general result that constructs triangulated manifolds from bounded-degree graphs G with a given set of cycles. We can use the theorem to compare the behavior of the functions $s_{d,\mathbb{F}}$ in different dimensions d:

Corollary 5. *If $s_{3,\mathbb{F}}(N) = O(\exp(c \cdot n))$ for some constant $c > 0$ then $s_{d,\mathbb{F}}(N) = O(\exp(c' \cdot n))$ for some $c' > 0$ and for all $d \geq 4$. In particular, if $s_{d,\mathbb{F}}$ is superexponential for some d then so is $s_{3,\mathbb{F}}$. Further, if $s_{d,\mathbb{F}}(N) = \Omega(\exp(c \cdot n \log n))$ for some $c > 0$ then $s_{3,\mathbb{F}} = \Omega(\exp(c' \cdot n \log n))$ for some $c' > 0$.*

Our second main theorem is that, for each $d \geq 3$, there is an infinite family of d-sphere triangulations for which the dual graphs form an expander family.

Theorem 6. *For each $d \geq 3$ there is an infinite family of combinatorially distinct triangulated d-spheres such that their dual graphs form a family of $(d+1)$-regular expanders. Moreover, there is a uniform bound (depending only on the dimension d) on the number of simplices incident to each vertex of these triangulations.*

The next corollary is deduced by taking connected sums.

Corollary 7. *Let M be a triangulable d-manifold. Then there is an infinite family of combinatorially distinct triangulations of M in which the vertex degrees are uniformly bounded (with the bound depending on M) and for which the dual graphs form a family of $(d+1)$-regular expanders. If M is PL, there exists such a family in which the triangulations are combinatorial.*

An analogous result in the Riemannian setting was first proved in unpublished work of the second and third authors: there exist metrics on S^3 with arbitrarily large volume and Cheeger constant uniformly bounded from 0. The version here answers a topological version of a question about convex polytopes asked by Gil Kalai in [9, 10]. Examples of simplicial 4-polytopes with dual graphs that are close to expanders (but which are not quite expanders) were described by Lauri Loiskekoski and Günter Ziegler in [12].

1.2 Paper outline and organization

In Section 3 we define *short classes* of graphs and prove basic results about them. These are essentially bounded-degree classes of graphs with "short" cycle bases over a given field; in a sense they are the main object studied in the paper. In Section 4 we describe a topological construction that takes as input a 2-complex and outputs a triangulated manifold with boundary. Section 5 discusses the relations between the topology of the boundary of such a manifold and the topology of the input 2-complex. In Section 6 we explain how to injectively encode facet-colored triangulations of a manifold in slightly larger (uncolored) triangulations of the same manifold: this makes our constructions injective. In Section 8 we put these tools together and prove our main results. A main part of the construction of expanding triangulations of S^3 is deferred to Section 9.

Acknowledgments

K.A. is supported by the Centre National de Recherche Scientifique, and the Horizon Europe ERC Grant number: 101045750 / Project acronym: HodgeGeoComb.

M.L. was partially supported by the Engineering and Physical Sciences Research Council (grant number EP/Y004256/1).

G.Y. was partially supported by the Horizon Europe ERC Grant number: 101045750 / Project acronym: HodgeGeoComb. He also thanks Zlil Sela for support and helpful conversations.

For the purpose of open access, the authors have applied a CC BY public copyright licence to any author accepted manuscript arising from this submission.

2 Notation and preliminaries

A triangulation Δ of a d-manifold M is a simplicial complex with realization homeomorphic to M. The dual graph of such a triangulation Δ is the graph with vertices the set of d-simplices, and edges connecting each pair of d-simplices that meet in a $(d-1)$-simplex; it is always $(d+1)$-regular if M has no boundary.

In a simplicial complex Δ, $\Delta^{(i)}$ denotes the set of i-dimensional simplices and $\Delta^{\leq i}$ denotes the i-skeleton. We denote the barycentric subdivision of Δ by $b(\Delta)$.

The *length* $\ell(\gamma)$ of a cycle γ in a graph is its number of edges.

The d-disk is denoted D^d, and $I = [0,1]$. All topology is done in the PL category unless stated otherwise. We follow [15] for the theory of regular neighborhoods. Stellar operations are defined in [11].

2.1 Handle structures on manifolds

We recall and summarize some basic PL handle theory (details can be found in either of [8, 15]). For the following discussion, $I = [0,1]$, and the center point of a cube I^d is $\left(\frac{1}{2}, \dots \frac{1}{2}\right)$. A d-dimensional k-*handle* (or *handle of index* k) is a space $H \overset{\text{homeo.}}{\cong} I^k \times I^{d-k}$ which is attached to a d-manifold W along an embedding $f : \left(\partial I^k\right) \times I^{d-k} \to \partial W$. The map f is then the *attaching map* of H to $W \cup_f H$. The *core* of H is $\mathrm{core}(H) = I^k \times \left(\frac{1}{2}, \dots \frac{1}{2}\right)$. The manifold $W \cup_f H$ is thus determined by W together with the attaching map f. The image $\mathrm{im}(f)$ is (by definition) a regular neighborhood of $f\left(\partial \mathrm{core}(H)\right)$ in ∂W: the pair $(\mathrm{im}(f), f(\partial \mathrm{core}(H))$ is homeomorphic to $\left(\partial I^k \times I^{d-k}, \partial I^k \times \left(\frac{1}{2}, \dots \frac{1}{2}\right)\right)$, so the first space collapses onto the second (which is contained in the first space's interior).

A PL embedding $f : \partial I^k \simeq \partial \mathrm{core}(H) \to \partial W$ need not extend to an attaching map of a k-handle H to W. Identifying the image of such an f with ∂I^k and denoting a regular neighborhood of ∂I^k within ∂W by E, (any two such neighborhoods are equal up to an ambient isotopy fixing ∂I^k,) it is clear that an embedding f extends to an attaching map if and only if the pair $\left(E, \partial I^k\right)$ is homeomorphic to $\left(\partial I^k \times I^{d-k}, \partial I^k \times \left(\frac{1}{2}, \dots \frac{1}{2}\right)\right)$.

This condition always holds if $k = 0$ or $k = 1$: if $k = 0$ this is trivial, while if $k = 1$ it follows from Whitehead's theorem that a regular neighborhood of any point in a manifold without boundary is a disk that contains the point in its interior. For the construction below we need the case $k = 2$.

Fact 8. *If $k = 2$, $f : \partial I^2 \to \partial W$ extends to an attaching map of a 2-handle if and only if a regular neighborhood E of the image is orientable.*

This is well known.

Proof sketch. We prove that a regular neighborhood N of an embedded S^1 within a PL d-manifold M is either homeomorphic to $\partial I^2 \times I^{d-1}$ or nonorientable. First, choose a $(d-1)$-disk h intersecting the S^1 transversely at a point p. Cut M along h to obtain a manifold M' with a new boundary sphere S^{d-1} (consisting of two copies of h glued along their boundary). The embedded $S^1 \subset M$ is cut into a segment, which is properly embedded in M'; the regular neighborhood N is cut into a regular neighborhood N' of this segment. By Whitehead's theorem, N' is a d-disk D. Re-gluing N' to itself across the two copies of h recovers N from N', and topologically it has the effect of identifying two disjoint $(d-1)$-disks in $\partial N'$. We denote these disks by $(h \cap N)^+$ and $(h \cap N)^-$. The triple

$$\left(D, (h \cap N)^+, (h \cap N)^-\right) \quad \text{is homeomorphic to} \quad \left(I^d, \{1\} \times I^{d-1}, \{0\} \times I^{d-1}\right),$$

and the gluing N is therefore homeomorphic to a quotient space of I^d obtained by identifying $\{1\} \times I^{d-1}$ with $\{0\} \times I^{d-1}$ via some homeomorphism. The isotopy class of this homeomorphism determines the homeomorphism type of the quotient space, and hence of N. There are only two isotopy classes of self-homeomorphisms of a disk I^{d-1}, only one of which results in an orientable quotient manifold. $\square$

2.2 Quasi-isometries

We recall some definitions and basic results.

Definition 9. Let $L, c > 0$. An (L, c) quasi-isometry between metric spaces X and Y is a function $f : X \to Y$ (not necessarily continuous) such that:

1. For each $y \in Y$ there exists $x \in X$ such that $d(f(x), y) \leq c$.

2. If $x, x' \in X$ then $\frac{1}{L} d(f(x), f(x')) - c \leq d(x, x') \leq L d(f(x), f(x')) + c$.

The spaces X, Y are called quasi-isometric if there exists a quasi-isometry $X \to Y$.

Remark 10.
1. If the constants don't matter much, we can also discuss c quasi-isometries, by which we mean (c, c) quasi-isometries in the sense above.

2. Given an (L, c) quasi-isometry $f : X \to Y$, one can construct a quasi-isometry $g : Y \to X$ as follows: for each $y \in Y$, choose $x \in X$ such that $d(f(x), y) \leq c$, and set $g(y) = x$. This is a near-inverse to f in the sense that that $d(y, f \circ g(y)) < c$ for all $y \in Y$ and $d(x, g \circ f(x)) < c'$ for all $x \in X$, where $c' > 0$ is some constant depending only on (L, c).

3. A composition of (L, c) quasi-isometries is again a quasi-isometry, with constants depending only on (L, c) (but generally larger).

4. For quasi-isometry related purposes it is convenient to work with vertex sets of graphs rather than the graphs themselves (where the vertex sets are metrized by distance, measured by the lengths of paths). The difference is inessential: if the vertex sets of two graphs are (L, c) quasi-isometric, and we geometrize the graphs' realizations (as simplicial complexes, including edges) by setting all edge lengths to be 1, these new spaces are also (L', c') quasi-isometric where (L', c') depend only on (L, c).

The following definition and lemma are useful for producing quasi-isometries of graphs.

Definition 11. A metric space X is (L, c) *quasi-geodesic* if for any $x, x' \in X$ (with $x = x'$ allowed) there are $n+1$ points $x_0 = x, x_1, \dots, x_n = x' \in X$ satisfying $d(x_i, x_{i+1}) \leq c$ for all $i < n$ as well as $n \leq Ld(x, x') + 1$. (The $+1$ here is necessary in case $d(x, x') < \frac{1}{L}$).

Note that a connected graph is always $(1, 1)$ quasi-geodesic.

Lemma 12. *Let X and Y be (L, c) quasi-geodesic metric spaces. Suppose there exists a relation $R \subset X \times Y$ and an $M > 0$ such that:*

1. *For each $x \in X$ the set $\{y \in Y \mid xRy\}$ is nonempty, and similarly for each $y \in Y$ the set $\{x \in X \mid xRy\}$ is nonempty.*

2. *If $x, x' \in X$ (not necessarily distinct) satisfy $d(x, x') \leq c$ and $y, y' \in Y$ satisfy xRy and $x'Ry'$ then $d(y, y') \leq M$.*

3. *If $y, y' \in Y$ (not necessarily distinct) satisfy $d(y, y') \leq c$ and $x, x' \in X$ satisfy xRy and $x'Ry'$ then $d(x, x') \leq M$.*

Then X is quasi-isometric to Y, with quasi-isometry constants depending only on L, M.

Proof. We construct a quasi-isometry $f : X \to Y$ as follows. For each $x \in X$ choose $f(x)$ to be an arbitrary element of $\{y \in Y \mid xRy\}$. If $x, x' \in X$ then there exist $x_0 = x, x_1, \dots, x_n = x' \in X$ such that $d(x_i, x_{i+1}) \leq c$ for all $i < n$, and in addition $n \leq Ld(x, x') + 1$. By (2) we have

$$d(f(x), f(x')) \leq d(f(x_0), f(x_1)) + d(f(x_1), f(x_2)) + \ldots + d(f(x_{n-1}), f(x_n))$$

$$\leq Mn \leq MLd(x, x') + M.$$

Similarly, for $y = f(x)$ and $y' = f(x')$ there exist $y_0 = y, y_1, \dots, y_n = y' \in Y$ such that $d(y_i, y_{i+1}) \leq c$ for all $i < n$, and in addition $n \leq Ld(y, y') + 1$. For each $0 < i < n$ choose x_i such that $x_i R y_i$. By (3) we have

$$d(x, x') \leq d(x_0, x_1) + \ldots + d(x_{n-1}, x_n) \leq Mn \leq MLd(f(x), f(x')) + M.$$

Finally, note that if $y \in Y$ then there exists $x \in X$ with xRy, and $d(f(x), y) \leq M$ by (2), using xRy and $xRf(x)$. $\square$

Fact 13. *Let $\{G_n\}_{n\in\mathbb{N}}$ and $\{H_n\}_{n\in\mathbb{N}}$ be families of graphs. Assume that for each n there is an (L, c) quasi-isometry $f_n : H_n \to G_n$, and that there is a uniform bound k on the vertex degrees of the graphs in both families. If $\{G_n\}_{n\in\mathbb{N}}$ is an expander family then so is $\{H_n\}_{n\in\mathbb{N}}$.*

This is known: see for instance [3, Lemma 5.8 and Remark 5.12]. It seems difficult to find a proof in the literature, so we include one for completeness.

Proof. Assume otherwise for a contradiction. Then for any $\varepsilon > 0$ we may find some $n \in \mathbb{N}$ and a subset $A_n \subset V(H_n)$ containing at most half of the vertices of H_n such that

$$\frac{|E(A_n, A_n^c)|}{|A_n|} < \varepsilon.$$

In particular, this implies $|V(H_n)| \geq \frac{1}{\varepsilon}$ and hence we may also choose such an n such that H_n has at least m vertices, for any given m; and since the vertex degrees are bounded, we may also demand that H_n have diameter larger than any given integer. We now suppress n from the notation, so that for any given $\varepsilon > 0$ we have $f : H \to G$ a quasi-isometry and $A \subseteq V(H)$ as above, with H as large as desired and G an expander (with Cheeger constant bounded below by some fixed $\alpha > 0$.)

We use big-O notation as follows: for expressions X, Y defined in the proof (typically subsets of vertex or edge sets of G and H,) $X = O(Y)$ means that there exists a real $r > 0$, depending only on the parameters L, c, k, and α, (but not on ε,) such that $X \leq r \cdot Y$ holds whenever ε is small enough. Similarly, $Y = \Omega(X)$ means $X = O(Y)$ (or $Y \geq r \cdot X$ for some r as above.) We also denote $X = \Theta(Y)$ if both $X = O(Y)$ and $Y = O(X)$. Each of O and Ω is transitive, and Θ is an equivalence relation.

Note that $|V(H)| = \Omega(|V(G)|)$, because each vertex of G is at distance at most c from a vertex of $f(V(H))$, so that $|V(G)| \leq k^c |V(H)|$. Also, if $S \subseteq V(H)$ then $|f(S)| = \Omega(|S|)$. This holds because f maps at most k^{Lc+L+1} vertices of H to a vertex of G: if $v, v' \in V(H)$ have $d(v, v') \geq Lc + L$ then $d(f(v), f(v')) \geq \frac{1}{L} d(v, v') - c \geq 1$. Note that for any $d \in \mathbb{N}$ there are at most $\sum_{0 \leq i \leq d} k^i \leq k^{d+1}$ vertices at distance at most d from any $v \in V(H)$. In particular we have $|V(H)| = O(|V(G)|)$ and hence also $|V(H)| = \Theta(V(G))$.

Let N_A be a $2c+1$-neighborhood of $f(A)$. We now show that $|N_A^c| = \Omega(|V(G)|)$. Denote $s = \lceil (3c+1) L \rceil$. Then there are at most $k^s |E(A, A^c)|$ vertices in H which are not in A but have distance at most s from A, because each such vertex has a path from a nearest vertex in A; such a path begins with an edge $e \in E(A, A^c)$, and continues along one of at most k^s paths of length at most $s-1$ from the terminal vertex of e in A^c. In particular, if ε is small enough (so that $|A| + k^s |E(A, A^c)| \leq (1 + k^s \varepsilon) |A| \leq \frac{2}{3} |V(H)|$) then the set $B \subset V(H)$ of vertices having distance greater than s from A satisfies $|B| \geq \frac{1}{3} |V(H)|$. It follows that $|f(B)| = \Theta(|V(G)|)$. Observe that any vertex of $f(A)$ has distance strictly larger than $2c+1$ from any vertex of $f(B)$: for any $a \in A$ and $b \in B$ we have

$$d(f(a), f(b)) \geq \frac{1}{L} d(a, b) - c > \frac{1}{L} s - c \geq 2c + 1,$$

since $d(a, b) > s$. Hence $f(B) \subseteq N_A^c$, and also

$$|N_A^c| \stackrel{*}{=} \Theta(|V(G)|).$$

We have $|E(N_A, N_A^c)| = \Omega(|N_A|)$: if $|N_A| \leq \frac{1}{2} |V(G)|$ this holds by α-expansion of G. Otherwise, since $|N_A^c| \leq \frac{1}{2} |V(G)|$, we can apply the Cheeger inequality to N_A^c and obtain

$$|E(N_A, N_A^c)| \geq \alpha |N_A^c| \stackrel{\text{by } (*)}{=} \Theta(|V(G)|) = \Omega(|N_A|),$$

so the conclusion holds in either case.

Define $g : V(G) \to V(H)$ as follows: for each $v \in V(G)$ choose a nearest vertex $f(w) \in \operatorname{im}(f)$, and set $g(v) = w$. Then g is a quasi-isometry which is at most k^{c+1}-to-one, and satisfies that $d(v, g \circ f(v)) \le c'$ for all $v \in V(H)$, for some constant c' depending only on (L, c).

Let $C \subseteq N_A$ be the set of endpoints in N_A of edges in $E(N_A, N_A^c)$. Then

$$|C| \ge \frac{1}{k} |E(N_A, N_A^c)| = \Omega(|N_A|),$$

and we conclude

$$|g(C)| = \Omega(|A|)$$

because $N_A \supseteq f(|A|) = \Theta(|A|)$.

Observe that each vertex of $g(C)$ has bounded distance t (depending only on (L,c)) from A. To see this, note that $C \subseteq N_A$, so given $v \in C$ we may find $w \in A$ with $d(v, f(w)) \le 2c+1$. By definition of g we have $g(v) = w'$ for some $w' \in V(H)$ with $d(v, f(w')) \le c$, and hence $d(f(w'), f(w)) \le 3c+1$. This implies an upper bound t on $d(w', w) = d(g(v), w) \ge d(g(v), A)$. As above, there are at most $k^t |E(A, A^c)| \le \varepsilon |A|$ vertices in H which are not in A but have distance at most t from A. Hence we have both $|g(C)| = O(\varepsilon |A|)$ and $|g(C)| = \Omega(|A|)$. Taking ε small enough gives a contradiction. □

2.3 PL collapse

Our notion of collapse is the same as in [15, Ch. 3]. We recall the definition:

Definition 14 ([15])**.** Suppose $X \supset Y$ are polyhedra such that $X = Y \cup D^n$, where $Y \cap D^n$ is a face* D^{n-1} of D^n. Then we say there is an elementary collapse of X on Y and write $X \Downarrow Y$ (the collapse is *across* D^n and *onto* D^{n-1}, *from* the complementary face $\mathrm{cl}(\partial D^n \setminus D^{n-1})$). We say X collapses on Y and write $X \searrow Y$ if there is a sequence of elementary collapses $X = X_0 \Downarrow X_1 \Downarrow \ldots \Downarrow X_n = Y$.

Note that this notion is invariant under PL homeomorphism. It is also useful to work with specific cellulations - this is effective (in the sense of being computable, for example) but less flexible.

Definition 15. Let X be polyhedron with a cell decomposition satisfying that each cell is a subpolyhedron, the attaching maps are embeddings, and the boundary of each cell within X is a union of lower-dimensional cells. An $(n-1)$-face $\sigma \subset X$ is *free* if it is contained in a unique n-face τ (so there is an elementary collapse from σ across τ).

Given a cellulation of a polyhedron, we can perform sequences of one-cell elementary collapses if the required faces are free.

3 Graphs, cycle bases, and 2-complexes

We are interested in bounded-degree graphs G in which there is a basis for H_1 consisting of "short" cycles. It would be convenient to work with graphs having a cycle basis with all cycles of uniformly bounded length, but a slight generalization is useful.

*A face D^{n-1} of a disk D^n is a subdisk of the boundary $D^{n-1} \subset \partial D^n$ such that the pair (D^n, D^{n-1}) is PL-homeomorphic to the pair $(I^{n-1} \times I, I^{n-1} \times \{0\})$.

Definition 16. Let R be a ring. A class $\mathcal{C}$ of connected finite graphs is *short* (over R) if there exist $d, k, L \in \mathbb{N}$ such that the following holds: all vertex degrees of graphs in $\mathcal{C}$ are at most d, and for each $G \in \mathcal{C}$ there exists a collection $\{\gamma_i\}_{i\in I}$ of cycles in G satisfying:

1. $\{\gamma_i\}_{i\in I}$ spans $H_1(G; R)$,
2. Each edge of G participates in at most k of the cycles,
3. The average length $\frac{1}{|I|}\sum_{i\in I}\ell(\gamma_i)$ of the cycles is at most L.

Such a class $\mathcal{C}$ is *uniformly* short over R if (3) is strengthened to: Each γ_i satisfies $\ell(\gamma_i) \leq L$.

Remark 17. If the graphs in $\mathcal{C}$ have no vertices of degree 1 or 2 then condition (3) is redundant: it always holds with $L = 2k$. To see this, consider pairs (e, γ) of an edge $e \in E(G)$ and a cycle $\gamma \in \{\gamma_i\}_{i\in I}$ such that $e \in \gamma$. The number of pairs equals $\sum_{i\in I}\ell(\gamma_i)$, and by condition (2) we find

$$\sum_{i\in I}\ell(\gamma_i) \leq k \cdot |E(G)|\,.$$

By condition (1) we have $|I| \geq |E(G)| - |V(G)| \overset{\star}{\geq} \frac{1}{2}|E(G)|$ and dividing by $|I|$ yields the result ($\star$ holds because $2|E(G)| \geq 3|V(G)|$).

In the other direction, for uniformly short classes condition (2) follows automatically: since the vertex degrees are bounded, there is a uniformly bounded number of cycles of length $\leq L$ passing through any given edge.

Remark 18. There is a universal L_0 such that all short classes can be "encoded" in uniformly short 3-regular graphs with cycle length bound $L = L_0$: see Lemma 48.

Proposition 19. *Let $d \geq 2$. The class $\mathcal{C}$ of all dual graphs to combinatorially-triangulated d-dimensional homology spheres over a ring R is a short class over R.*

If R is a PID the same holds without restricting to combinatorial triangulations, i.e. for the class of all dual graphs to triangulated d-dimensional homology spheres over R.

Proof. Note that graphs in $\mathcal{C}$ are all $(d+1)$-regular. Let Δ be a combinatorial triangulation of a d-dimensional homology sphere over R. The set of d-simplices containing a given $(d-2)$-simplex of Δ forms a simple cycle in the dual graph G. Let $\{\gamma_i\}_{i\in I}$ be the collection of all such cycles, where $I = \Delta^{(d-2)}$. These are the boundaries of 2-faces of the dual cell complex X to Δ.

If Δ is a combinatorial triangulation then X is homeomorphic to Δ, and hence has no first homology over R. Otherwise, the cohomology of Δ over R in each degree i is isomorphic to the homology of X in degree $d-i$. Assuming that R is a PID and applying the universal coefficient theorem for cohomology gives the exact sequence

$$0 \to \operatorname{Ext}^1_R(H_{d-2}(\Delta; R), R) \to \underbrace{H^{d-1}(\Delta; R)}_{\simeq H_1(X;R)} \to \underbrace{\operatorname{Hom}_R(H_{d-1}(\Delta; R), R)}_{=\operatorname{Hom}_R(0,R)=0} \to 0$$

where if $d = 2$ then the Ext on the left is $\operatorname{Ext}^1_R(R, R) = 0$, and otherwise it is $\operatorname{Ext}^1_R(0, R) = 0$. This implies that the dual cell complex X has no first homology over R. Hence, in either case $\{\gamma_i\}_{i\in I}$ (boundaries of the 2-cells of X) span $H_1(G; R)$ (because $G \simeq X^{\leq 1}$ is the 1-skeleton of X). By Remark 17 it suffices to show that each edge of G participates in boundedly many of the cycles $\{\gamma_i\}_{i\in I}$.

A pair of edges $\{u,v\}$ and $\{v,w\}$ of G which meet at v is contained in at most one (in fact exactly one) of the cycles $\{\gamma_i\}_{i\in I}$: the corresponding triple $\sigma_u, \sigma_v, \sigma_w$ of d-simplices of Δ satisfy that each of $\sigma_u \cap \sigma_v$ and $\sigma_v \cap \sigma_w$ is a $(d-1)$-simplex, and these intersections are distinct (if they are equal we obtain a $(d-1)$-simplex contained in each of σ_u, σ_v,and σ_w, but this is impossible in a triangulated d-manifold). Hence an edge $\{u,v\}$ is contained in at most as many of the cycles as the number of edges containing v in G which are not $\{u,v\}$. Since G is $(d+1)$-regular, there are exactly d such edges. $\square$

A significant part of the rest of this paper is devoted to the details of a construction going in the other direction: from a short class of graphs to a class of triangulated d-dimensional homology spheres (where we are free to choose $d \geq 3$). That this construction can be made injective, together with the fact that it only increases the number of faces by a constant ratio, is the main idea behind Theorem 4. There is a natural intermediate step, from graphs to nulhomologous 2-complexes.

Lemma 20. *For each $m \geq 3$ there exists a triangulation of the closed 2-disk such that:*

- *The boundary is a cycle of length m, and is an induced subcomplex,*
- *Each vertex has degree at most 6,*
- *There are at most m interior vertices, and each is connected to a boundary vertex by an edge.*

Proof. Begin with the zigzag triangulations of [1, p.3]. For $m=8$ and $m=13$, they are:

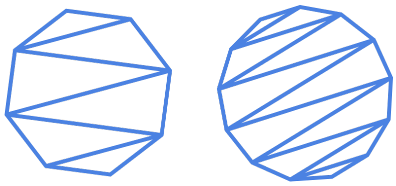

Then take a stellar subdivision at the center of each interior edge:

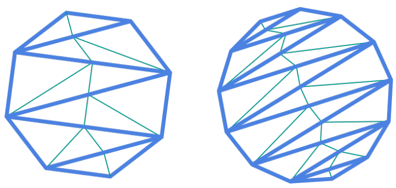

(The order in which these subdivisions are preformed is inessential, but here we have subdivided starting from the topmost interior edge, in vertical order.) $\square$

Proposition 21. *Let $\mathcal{C}$ be a class of short graphs over a field $\mathbb{F}$ with constants d,k,L as above. Then there is a class $\mathcal{D}$ of connected 2-dimensional simplicial complexes and an integer d' such that for all $\Delta \in \mathcal{D}$ we have $H_1(\Delta;\mathbb{F}) = H_2(\Delta;\mathbb{F}) = 0$, and all vertices of Δ have degree at most d'. Further, there exists a function $f : \mathcal{C} \to \mathcal{D}$, injective on isomorphism classes, such that if $f(G) = \Delta$ then:*

- *G is isomorphic to a subgraph of $\Delta^{\leq 1}$,*
- *Every vertex of Δ is connected by an edge to a vertex in the image of G. In particular, if G has n vertices then Δ has at most $d' \cdot n$ vertices.*

- *If $\mathcal{C}$ is uniformly short, then G is quasi-isometric to $\Delta^{\leq 1}$ with quasi-isometry constants depending only on d, k, L.*

Injectivity here is achieved by a somewhat ad-hoc process, but morally this is similar to the coloring arguments we use elsewhere.

Proof. Given $G \in \mathcal{C}$ with cycles $\{\gamma_i\}_{i\in I}$, construct $\Delta = f(G)$ as follows. First choose $J \subset I$ such that $\{\gamma_j\}_{j\in J}$ is a basis of $H_1(G;\mathbb{F})$. For each γ_j glue to G a triangulation of the 2-disk with boundary γ_j as given by the previous lemma. Denote the resulting complex by Δ'.

It is straightforward that $H_1(\Delta';\mathbb{F}) = 0$. We prove that Δ' has $H_2(\Delta';\mathbb{F}) = 0$: without loss of generality $J = \{1, 2, \dots, t\}$ and the corresponding 2-disks are $D_1, \dots, D_t$. We glue the disks one at a time starting from D_1 and prove the claim by induction on $G \cup \bigcup_{j\leq i} D_j$, where $G \cup \bigcup_{j\leq t} D_j = \Delta'$. For the gluing of D_i we have a Mayer–Vietoris sequence (with coefficients in $\mathbb{F}$, which we omit from the notation):

$$H_2(\gamma_i) = 0 \to H_2\left(G \cup \bigcup_{j<i} D_j\right) \oplus H_2(D_i) \to H_2\left(G \cup \bigcup_{j\leq i} D_j\right) \to$$

$$H_1(\gamma_i) \to H_1\left(G \cup \bigcup_{j<i} D_j\right) \oplus H_1(D_i) \to \dots.$$

Since the cycles $\{\gamma_j\}_{j\in J}$ form an $\mathbb{F}$-basis of $H_1(G)$ and gluing $D_1, \dots, D_{i-1}$ has the effect on H_1 of annihilating precisely the subspace spanned by $\gamma_1, \dots, \gamma_{i-1}$, the map $H_1(\gamma_i) \to H_1\left(G \cup \bigcup_{j<i} D_j\right)$ (induced from the inclusion) is injective. It follows that the boundary map $H_2\left(G \cup \bigcup_{j\leq i} D_j\right) \to H_1(\gamma_i)$ is 0. By induction $H_2\left(G \cup \bigcup_{j<i} D_j\right) = 0$, and hence also $H_2\left(G \cup \bigcup_{j<i} D_j\right) \oplus H_2(D_i) = 0$. So part of the sequence reads

$$0 \to H_2\left(G \cup \bigcup_{j\leq i} D_j\right) \to H_1(\gamma_i),$$

where both maps are 0, and therefore $H_2\left(G \cup \bigcup_{j\leq i} D_j\right) = 0$.

The vertices of Δ' are the vertices of G, together with the interior vertices of the triangulated disks produced by the previous lemma. Its vertex degrees can be bounded by $d(1+6k)$: each vertex not in G has degree at most 6 by construction. To bound degrees of vertices in G, observe that each edge of G is contained in at most k cycles of $\{\gamma_i\}_{i\in I}$, and hence a vertex v of G is contained in at most dk of these cycles. For each such cycle we have added at most 6 edges to v (here 6 is the degree bound on the disk triangulations of the previous lemma). In particular, each edge of Δ' participates in at most $d(1+6k)-1$ triangles: if $e \in \Delta'$ is an edge and v is one of its vertices, each triangle τ containing e must contain an additional edge $\{v, u\}$ ($u \notin e$) and this edge identifies τ uniquely.

We construct Δ from Δ' as follows: for each edge $e = \{v, w\}$ of $G \subset \Delta'$ add $d(1+6k)$ new vertices $u_1, \dots, u_{d(1+4k)}$, connect each of them to both v and w, and add the triangles $\{u_i, v, w\}$ for $i = 1, \dots, d(1+6k)$. Note that each of the new vertices is contained in precisely one triangle. Hence G can be identified within Δ as follows: it is the subgraph of the 1-skeleton containing all edges that participate in at least $d(1+6k)$ triangles, together with all these edges' vertices. Note further that $\tilde{H}_*(\Delta;\mathbb{F}) = 0$, because this holds for Δ' (which is a deformation retract).

Define $\mathcal{D}$ to be the class of all complexes Δ obtained from graphs $G \in \mathcal{C}$ as above, and define $f : \mathcal{C} \to \mathcal{D}$ by setting $f(G)$ to be the corresponding complex Δ.

It remains to prove that if $\mathcal{C}$ is uniformly short then each $G \in \mathcal{C}$ is quasi-isometric to $f(G)$, with constants depending only on $\mathcal{C}$. Let $G \in \mathcal{C}$ and let $\Delta \supset G$ be the corresponding complex. Then the embedding $\iota : G \to \Delta^{\leq 1}$ is a quasi-isometry as desired: each vertex of $\Delta^{\leq 1}$ is adjacent to a vertex of G. For any vertices v, v' of G we have $d(\iota(v), \iota(v')) \leq d(v, v')$. Further, if $w_0 = \iota(v), w_1, \ldots, w_n = \iota(v')$ is a path of minimal length in $\Delta^{\leq 1}$, and $w_{i_0} = \iota(v), w_{i_1}, \ldots, w_{i_t} = \iota(v')$ are its vertices in the image of G, (where $0 = i_0 \leq i_1 \leq \ldots \leq i_t = n$,) then the distance in G between each successive pair $v_{i_j}, v_{i_{j+1}}$ is at most L, because any such pair is contained within the same cycle γ_i. To see this, note that a connected component of the complement of G within $\Delta^{\leq 1}$ is either a set of the interior vertices in a disk filling a cycle γ_i, as produced by the previous lemma, or a single vertex in $\Delta^{(0)} \setminus \Delta'^{(0)}$. In the latter case a path through the vertex is never a geodesic, because such a vertex is connected to precisely two vertices of G, which are already adjacent in G. Hence $d(v, v') \leq L \cdot n \leq L \cdot d(\iota(v), \iota(v'))$. □

4 Handle structures with controlled triangulations

We describe a general "thickening" construction that takes 2-dimensional simplicial complexes Δ as input and outputs d-dimensional ($d \geq 4$) combinatorial manifolds $M(\Delta)$ with boundary. The results of this section can be summarized as follows:

Theorem 22. *Fix $d, k \in \mathbb{N}$ with $d \geq 4$. Then for every 2-dimensional simplicial complex Δ in which all vertex stars have at most k faces there exists a PL d-manifold $M(\Delta)$ with boundary such that:*

1. *$M(\Delta)$ has a handle structure consisting of 0-, 1-, and 2-handles in bijection with the faces of Δ:*

$$M(\Delta) = \bigcup_{v\in\Delta^{(0)}} H_{0,v} \cup \bigcup_{e\in\Delta^{(1)}} H_{1,e} \cup \bigcup_{\sigma\in\Delta^{(2)}} H_{2,\sigma}.$$

   Further, $H_{i,\sigma} \cap H_{j,\tau} \neq \emptyset$ if and only if $\sigma \cap \tau \neq \emptyset$ in Δ.

2. *There is a subpolyhedron $\Delta' \subset M(\Delta)$ such that $\Delta \simeq \Delta'$ are PL homeomorphic, and if Δ is finite then $M(\Delta) \searrow \Delta'$.*

3. *$M(\Delta)$ has a triangulation in which each handle $H_{i,\sigma}$ is a subcomplex. The total number of faces in each such subcomplex is uniformly bounded in terms of d, k (and in particular is independent of Δ).*

The construction of $M(\Delta)$ from Δ involves many arbitrary choices, but it can be made computable and explicit bounds for (3) can be found.

The general idea is to replace each i-simplex in Δ by a d-dimensional i-handle, and to glue these handles in the "same" manner as the faces of Δ (although this involves choices that affect the PL homeomorphism type nontrivially). We first carry out the construction in the PL category, then prove the wanted results on collapsibility, and finally find suitable triangulations.

In Section 5 we use part (2) to obtain conclusions on $H_*(\partial M)$. Part (3) implies that the dual graph of the induced triangulation of ∂M is quasi-isometric to Δ, with quasi-isometry constants depending only on d and k: we return to this point in the proof of Theorem 43.

4.1 Construction of $M(\Delta)$

It is useful for us to mark certain subpolyhedra on our 0- and 1-handles: all gluing will take place along these subpolyhedra. Having notation for them facilitates finding small triangulations, because it allows us to work with fixed triangulations of the handles (depending on the vertex degree bound k but not on Δ).

4.1.1 Constructing the 0- and 1-handles

Fix an oriented $(d+1)$-disk H_0 (a "reference 0-handle") together with k points $\{p_1, \ldots, p_k\}$ on ∂H_0. Take pairwise disjoint regular neighborhoods Σ_i of the points p_i $(1 \leq i \leq k)$ in ∂H_0, and fix k points in $\partial \Sigma_i$ for each i: call these $q_{i,1}, \ldots, q_{i,k}$.

Fix a collection of pairwise disjoint paths $\{p_t\}_t$ within ∂H_0 and disjoint from the interiors of $\Sigma_1, \ldots, \Sigma_k$, one path between each pair q_{i_1,j_1}, q_{i_2,j_2} where $i_1 \neq i_2$. Note that ∂H_0 is a d-sphere with $d \geq 3$, so such paths exist.

Now fix a d-disk Σ together with k marked points $\{u_1, \ldots, u_k\}$ in $\partial \Sigma$, and define $H_1 = \Sigma \times I$. Choose an orientation for H_1 (our "reference 1-handle").

4.1.2 Gluing 0- and 1-handles

Given an input complex Δ in which all vertex stars have at most k faces, let (V, E) be the 1-skeleton $\Delta^{\leq 1}$. Take a collection of copies of H_0 indexed by V and a collection of copies of H_1 indexed by E: denote these by $\{H_{0,v}\}_{v \in V}$ and $\{H_{1,e}\}_{e \in E}$ respectively.

For each $H_{1,e}$ we want to glue each of the two "ends" $\Sigma \times 0$ and $\Sigma \times 1$ to some Σ_i $(1 \leq i \leq k)$ in the appropriate $H_{0,v}$. These should be chosen such that no two 1-handles are glued to the same $\Sigma_i \subset H_{0,v}$ for any v. An arbitrary choice works (and is always possible): first, for each edge $e = \{v, w\}$ choose an arbitrary orientation (v, w). Then for each vertex v order the adjacent edges as $e_1, \ldots, e_s$ (where $s \leq k$). For each $1 \leq i \leq s$, if v is the initial vertex of the oriented edge e_i (i.e. $e_i = (v, w)$ for some w) we glue $\Sigma \times 0 \subset H_{1,e}$ to Σ_i; otherwise v is the terminal vertex of e_i, and we glue $\Sigma \times 1 \subset H_{1,e}$ to Σ_i. We arrange that $H_{1,e}$ and $H_{0,v}$ induce opposite orientations on the identified disk $\Sigma \times 0 \simeq \Sigma_i$ (or $\Sigma \times 1 \simeq \Sigma_i$). This ensures that the gluing is orientable. The gluing maps are chosen such that $\{u_1, \ldots, u_k\} \subset \Sigma$ is glued to $\{q_{i,1} \ldots, q_{i,k}\} \subset \Sigma_i$, and we relabel the points in $\Sigma_i \subset H_{0,v}$ so that the image of each u_j is $q_{i,j}$ (it is possible to arrange directly for u_j to be glued to $q_{i,j}$ without any relabeling, because the group of PL orientation-preserving homeomorphisms of the sphere of each dimension ≥ 2 is k-transitive for all k. But we will not do this in the triangulations we construct).

We denote the resulting manifold with boundary by M_1.

4.1.3 Gluing 2-handles

For each triangle $\sigma = \{u, v, w\} \in \Delta^{(2)}$ we construct a closed path γ_σ in ∂M_1 traversing the handles corresponding to $u, \{u, v\}, v, \{v, w\}, w, \{w, u\}$ in cyclic order. The part of γ_σ within $\partial H_{0,v}$ is one of the paths p_t constructed above (in the construction of subpolyhedra of H_0) and similarly for $H_{0,w}$ and $H_{0,u}$. For each edge e of the triangle, the part of γ_σ within $\partial H_{1,e}$ is of the form $u_j \times I$ for a certain $u_j \in \partial\Sigma$. These choices need to be made consistently, so that $(u_j, 0) \in \partial(\Sigma \times 0)$ and $(u_j, 1) \in \partial(\Sigma \times 1)$ are the endpoints of the appropriate paths p_t in the boundaries of the corresponding 0-handles, and also so that the same point u_j does not get chosen twice within the same 1-handle by different triangles. Like in the case of the 1-handles, this can always be done: first, for each edge e order the triangles σ containing e as $\sigma_1, \ldots, \sigma_s$ (where $s \leq k$ by assumption). Then assign each σ_i the interval $u_i \times I$ within $\partial H_{1,e}$ (and hence also within ∂M_1). At this point, for each triangle σ of Δ we have chosen the subpaths of γ_σ within the boundaries of the 1-handles corresponding to its edges. If v is a vertex of σ, two of these subpaths terminate at some $q_{i_1,j_1} \in \Sigma_{i_1}$ and $q_{i_2,j_2} \in \Sigma_{i_2}$ within $\partial H_{0,v}$ (by our gluing of 1-handles to 0-handles). Exactly one of the paths in $\{p_t\}_t$ connects these two points: this path p_t is the subpath of γ_σ within $H_{0,v}$.

The paths $\{\gamma_\sigma : \sigma \in \Delta^{(2)}\}$ are pairwise disjoint, and since ∂M_1 is orientable (as M_1 is) Fact 8 shows that their regular neighborhoods are PL-homeomorphic to $S^1 \times D^{d-1}$. Choose pairwise disjoint regular neighborhoods of $\{\gamma_\sigma : \sigma \in \Delta^{(2)}\}$ and glue a 2-handle to each of them.

The resulting manifold is M.

Remark 23. The gluing maps of the 2-handles involve choices that may change the homeomorphism type. This may be easiest to see when $d = 4$: in this case the effect on ∂M of gluing an additional 2-handle is a Dehn surgery. The choice of gluing map determines the slope, for which an arbitrary integer can be chosen.

4.2 The subpolyhedron Δ' of $M(\Delta)$

In order to identify Δ' within $M(\Delta)$ it is useful to have notation for certain additional subpolyhedra of our handles.

Identify the reference 0-handle H_0 with the cone over its boundary $c_0 * \partial H_0$, and similarly identify the $(d-1)$-disk Σ with the cone $c_1 * \partial\Sigma$, so that our reference 1-handle is $H_1 = \Sigma \times I = (c_1 * \partial\Sigma) \times I$. Also choose a "center" point d_i in the interior of each $\Sigma_i \subset \partial H_0$, so that $\Sigma_i = d_i * \partial\Sigma_i$. Without loss of generality we may assume that the gluing maps of 1-handles to 0-handles, which glue copies of $\Sigma \times i \subset \partial H_1$ (for $i \in \{0, 1\}$) to copies of $\Sigma_j \subset \partial H_0$ (for $j \in \{1, \ldots, k\}$) satisfy that $c_1 \in \Sigma \times i$ maps to $d_j \in \Sigma_j$. We further make the following "linearity" assumption: If $X \subset \partial\Sigma$ is a subpolyhedron, and X' is its image under one of the gluing maps of $\Sigma \times i \subset \partial H_{1,e}$ to $\Sigma_j \subset \partial H_{0,v}$, then $c_1 * X$ is glued to $d_i * X'$. This assumption is harmless: it can be omitted at the cost of introducing separate notation for the PL-homeomorph of $c_1 * X$ within Σ_j (but it is actually satisfied in our triangulations of M).

Each 2-handle $H_{2,\sigma}$ (for $\sigma \in \Delta^{(2)}$) in $M(\Delta)$ is a PL $I^2 \times I^{d-2}$, and it is attached along a regular neighborhood of $\partial I^2 \times \{(\frac{1}{2}, \ldots, \frac{1}{2})\}$ to a regular neighborhood of a closed path γ_σ. We may assume without loss of generality that the gluing map sends $\partial I^2 \times \{(\frac{1}{2}, \ldots, \frac{1}{2})\}$ to γ_σ.

Recall that in each $e \in \Delta^{(1)}$ the 2-faces of Δ containing e are ordered as $\sigma_1, \ldots, \sigma_s$

($s \leq k$) in the construction of M, and that 2-handles are glued to $\partial H_{1,e}$ along regular neighborhoods of $u_1 \times I, \ldots, u_s \times I$ respectively. We denote by $s(e)$ the number of 2-faces containing each $e \in \Delta^{(1)}$.

We define Δ' by giving its intersection with each handle $H_{i,\sigma}$. The idea is that these intersections contain the core of each handle $H_{i,\sigma}$ (a point for a 0-handle, a segment for a 1-handle, or a 2-disk for a 2-handle); but further, whenever σ contains a face τ, the intersection of Δ' with the handle corresponding to τ has a "flap" extending into it from the handle corresponding to σ. For example, the core of a 1-handle $H_{1,e}$ is $c_1 \times I$, and the core of an adjacent 0-handle $H_{0,v}$ is c_0. We extend $c_1 \times I$ into $H_{0,v}$ by a segment from c_0 to the corresponding endpoint of $c_1 \times I$.

- For each 2-handle $H_{2,\sigma} = I^2 \times I^{d-2}$ define $C(\sigma) = I^2 \times \left(\frac{1}{2}, \ldots, \frac{1}{2}\right)$ (this square is the core of the handle). We define $\Delta' \cap H_{2,\sigma} = C(\sigma)$.

- For each 1-handle $H_{1,e}$, the core of the handle is $C(e) = c_1 \times I \subset H_{1,e}$. Let $\sigma_1, \ldots, \sigma_s$ be the 2-faces containing e in the order above. Define the flaps
$$F_{\sigma_i,e} = (c_1 * u_i) \times I$$
(these are rectangles with one edge going vertically through the core of $H_{1,e}$) and set
$$\Delta' \cap H_{1,e} = C(e) \cup \bigcup_{i=1}^{s(e)} F_{\sigma_i,e}.$$
Observe that we also have $\Delta' \cap H_{1,e} = (c_1 * \{u_1, \ldots, u_{s(e)}\}) \times I$ (or just $c_1 \times I$ if $s(e) = 0$).

- For each $v \in \Delta^{(0)}$, the core of $H_{0,v}$ is $C(v) = \{c_0\} \subset H_{0,v}$. Given an edge e containing v, let i be the unique index $1 \leq i \leq k$ such that $H_{1,e}$ is glued to $\partial H_{0,v}$ along $\Sigma_i \subset \partial H_{0,v}$. Define
$$F_{e,v} = c_0 * d_i.$$
Given a 2-face σ containing v, denote the two edges of σ that contain v by e_1, e_2. Let $1 \leq i_1, i_2 \leq k$ be the indices such that H_{1,e_1} and H_{1,e_2} are glued to $\partial H_{0,v}$ along Σ_{i_1} and Σ_{i_2} respectively. Let j_1 be the index such that σ is the j_1-th face containing e_1 (in the order above), and similarly j_2 for e_2. In other words, j_1 is the index such that $F_{\sigma,e_1} = (c_1 * u_{j_1}) \times I$, and similarly for j_2. Among the paths $\{p_t\}_t \subset \partial H_0$, let p_t be the one running between q_{i_1,j_1} and q_{i_2,j_2}. Define
$$F_{\sigma,v} = [c_0 * p_t] \cup [F_{e_1,v} * q_{i_1,j_1}] \cup [F_{e_2,v} * q_{i_2,j_2}].$$
Note that $c_0 * p_t$ is a PL 2-disk with boundary $p_t \cup (c_0 * q_{i_1,j_1}) \cup (c_0 * q_{i_2,j_2})$. Each of the other two subpolyhedra is a triangle linearly embedded in $H_{0,v}$, with boundaries $\{c_0, q_{i_1,j_1}, d_{i_1}\}$ and $\{c_0, q_{i_2,j_2}, d_{i_2}\}$ respectively. The disk $c_0 * p_t$ meets the triangle $F_{e_1,v} * q_{i_1,j_1}$ along the segment $c_0 * q_{i_1,j_1}$ and similarly for e_2: the union is hence a PL 2-disk with boundary
$$\partial F_{\sigma,v} = \underbrace{(c_0 * d_{i_1})}_{F_{e_1,v}} \cup [(d_{i_1} * q_{i_1,j_1}) \cup p_t \cup (q_{i_2,j_2} * d_{i_2})] \cup \underbrace{(d_{i_2} * c_0)}_{F_{e_2,v}}.$$
We define
$$\Delta' \cap H_{0,v} = C(v) \cup \bigcup_{e \ni v} F_{e,v} \cup \bigcup_{\sigma \ni v} F_{\sigma,v}.$$

Proposition 24. $\Delta' \simeq \Delta$.

Proof. For each face $\sigma \in \Delta$ we define $D_\sigma = C(\sigma) \cup \bigcup_{\emptyset \neq \tau \subsetneq \sigma} F_{\sigma,\tau}$. Observe that if σ is a vertex then $D_\sigma = c_0 \in H_{0,v}$ is a point, and if $\sigma = \{v, w\}$ is an edge then D_σ is an arc with endpoints D_v, D_w. We prove that D_σ is a 2-disk for each $\sigma \in \Delta^{(2)}$, that $\Delta' = \bigcup_{\sigma \in \Delta} D_\sigma$, and that each point of Δ' is in the interior of precisely one of the disks $\{D_\sigma : \sigma \in \Delta\}$ (where the unique point of a 0-disk is considered "interior"). Therefore Δ' is homeomorphic to a cell complex with cells indexed by the faces of Δ. Finally, we prove that $\Delta' \simeq \Delta$ by showing that the cells $\{D_\sigma\}_{\sigma \in \Delta}$ are glued to each other in the same way as the faces of Δ are.

Claim. Let $\sigma = \{u, v, w\}$ be a 2-face of Δ. Then D_σ is a PL 2-disk with boundary $D_{\{u,v\}} \cup D_{\{v,w\}} \cup D_{\{w,u\}}$.

Proof. By definition,

$$D_\sigma = C(\sigma) \cup F_{\sigma,\{u,v\}} \cup F_{\sigma,\{v,w\}} \cup F_{\sigma,\{w,u\}} \cup F_{\sigma,u} \cup F_{\sigma,v} \cup F_{\sigma,w}$$

is a union of seven 2-disks. Their interiors are disjoint, because the interior of each is contained in the interior of a different handle of M. It follows from the construction that the disks are arranged as

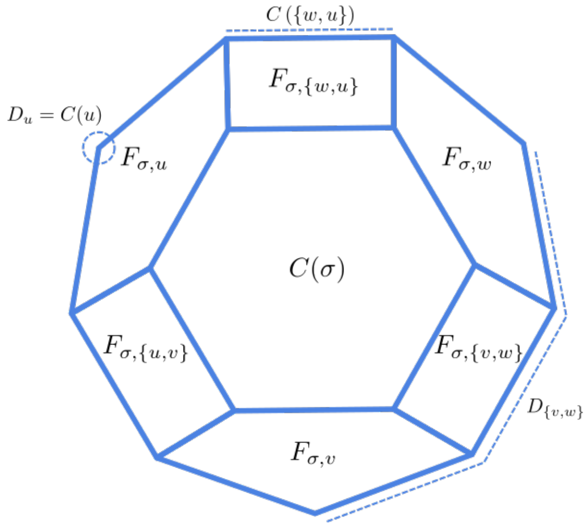

and hence form a disk as desired. The dashed line on the right is parallel to the boundary arc $D_{\{v,w\}}$. The rest of the boundary is similarly composed of $D_{\{w,u\}} \cup D_{\{u,v\}}$. □

Claim. $\Delta' = \bigcup_{\sigma \in \Delta} D_\sigma$.

Proof. This is direct from the definition: for each face $\tau \in \Delta$ of dimension i, the sub-polyhedron $\Delta' \cap H_{i,\tau}$ is of the form $C(\tau) \cup \bigcup_{\sigma \supset \tau} F_{\sigma,\tau}$. Hence

$$\Delta' = \bigcup_{\sigma \in \Delta} C(\sigma) \cup \bigcup_{\emptyset \neq \tau \subsetneq \sigma \in \Delta} F_{\sigma,\tau}.$$

Each D_σ is of the form $D_\sigma = C(\sigma) \cup \bigcup_{\emptyset \neq \tau \subsetneq \sigma} F_{\sigma,\tau}$, and hence $\Delta' = \bigcup_{\sigma \in \Delta} D_\sigma$. □

Claim. For each $\sigma \in \Delta$ the subpolyhedron $D_\sigma \subset \Delta'$ is a disk of dimension $\dim \sigma$. If $\sigma, \tau \in \Delta$ are distinct faces then the interiors of the disks D_σ and D_τ are disjoint.

Proof. For σ a vertex, $D_\sigma = C(\sigma)$ is a point. For $\sigma = \{v, w\}$ an edge, $D_\sigma = C(\sigma) \cup F_{\sigma,v} \cup F_{\sigma,w}$ is a union of three segments: for the point $c_0 \in H_{0,v}$ and the index $1 \leq i \leq k$ such that $H_{1,\sigma}$ is glued to $H_{0,v}$ along $\Sigma_i \subset \partial H_{0,v}$, we have $F_{\sigma,v} = c_0 * d_i$, where d_i is one of the two endpoints of $C(\sigma)$. Hence D_σ is a 1-disk (an arc) with endpoints $C(v)$ and $C(w)$. For σ a triangle the previous claim shows that D_σ is a 2-disk. This proves the first part of the claim.

Let $\sigma, \tau \in \Delta$ be distinct. Note that $D_\sigma = C(\sigma) \cup \bigcup_{\emptyset \neq \rho \subset \sigma} F_{\sigma,\rho}$, where $C(\sigma)$ is contained in $H_{\dim \sigma, \sigma}$ and each $F_{\sigma,\rho}$ is contained in $H_{\dim \rho, \rho}$. In particular, in order to compute $D_\sigma \cap D_\tau$ it suffices to consider the intersection within those handles of M that correspond to faces $\rho \subseteq \sigma \cap \tau$. For $\rho \subseteq \sigma \cap \tau$ of dimension i, we consider the possibilities for $D_\sigma \cap D_\tau \cap H_{i,\rho}$:

- The case $i = 2$ is impossible, because if $\sigma \cap \tau$ is a 2-face then $\tau = \sigma$.
- If $i = 1$ then $D_\sigma \cap H_{1,\rho}$ is either $C(\rho)$ (if $\rho = \sigma$) or of the form $F_{\sigma,\rho} = (c_1 * u_j) \times I$ for some $1 \leq j \leq k$ if $\rho \subsetneq \sigma$. Similarly $D_\tau \cap H_{1,\rho} = F_{\tau,\rho}$ is either $C(\rho)$ or $(c_1 * u_{j'}) \times I$ for some $1 \leq j' \leq k$, where $j \neq j'$ if both are defined. In all cases we have $F_{\sigma,\rho} \cap F_{\tau,\rho} = C(\rho)$.
- If $i = 0$ then $\rho = \{v\}$ is a vertex. Recall that $C(v) = \{c_0\} \subset H_{0,v}$. If σ is a proper superset of $\{v\}$ then $F_{\sigma,\rho}$ is a cone $c_0 * Z$ for some $Z \subset \partial H_{0,v}$: Z is either a single point d_j ($1 \leq j \leq k$) if σ is the edge such that $H_{1,\sigma}$ is glued onto $\Sigma_j \subset \partial H_{0,v}$, or (if σ is a 2-face) it is an arc of the general form $(d_{i_1} * q_{i_1,j_1}) \cup p_t \cup (q_{i_2,j_2} * d_{i_2})$ for some indices $1 \leq i_1, i_2, j_1, j_2 \leq k$ (where $i_1 \neq i_2$). Similarly if τ is a proper superset of ρ then $F_{\tau,\rho} = c_0 * Z'$ for Z' of the same general form as Z. If either $\sigma = \rho = \{v\}$ then $D_\sigma \cap D_\tau \cap H_{0,v} = D_\sigma \cap D_\tau = C(v)$, which is a boundary point of D_τ, and the interiors of the disks are disjoint. Otherwise ρ is a proper face of both σ and τ, and $F_{\sigma,\rho} \cap F_{\tau,\rho} = c_0 * (Z \cap Z')$. If $Z \cap Z' = \emptyset$ then $F_{\sigma,\rho} \cap F_{\tau,\rho} = \{c_0\}$ is a boundary point of D_σ and of D_τ and again we are done. Otherwise, $Z \cap Z' = \{d_{i_1}\}$ for some $1 \leq i_1 \leq k$. In this case at least one of σ or τ is a 2-face (else both are edges, but this would imply that $H_{1,\sigma}$ and $H_{1,\tau}$ have both been glued to $\Sigma_{i_1} \subset \partial H_{0,v}$, but this contradicts our choices above). Assuming without loss of generality that σ is a 2-face, $c_0 * d_{i_1} \subset H_{0,v}$ is an arc on the boundary of D_σ, and hence there is no interior point of D_σ in $D_\sigma \cap D_\tau \cap H_{0,v}$.

□

Observe that each point of $\Delta'^{\leq 1} := \bigcup_{\{\sigma \in \Delta : \dim \sigma \leq 1\}} D_\sigma$ is in the interior of precisely one of the disks D_σ: if it is not in the interior of some arc D_e (where $e \in \Delta^{(1)}$ is an edge) then it is the boundary point of such an arc, which is an interior point of the corresponding 0-disk. Since for each 2-face $\tau \in \Delta$ the 2-disk D_τ is glued onto $\Delta'^{\leq 1}$ along the entire boundary, each point of Δ' is either a point of $\Delta'^{\leq 1}$ or is in the interior of some disk in $\{D_\tau : \tau \in \Delta^{(2)}\}$. Hence Δ' is homeomorphic to a cell complex with cells $\{D_\sigma : \sigma \in \Delta\}$. (If Δ happens to be infinite, the topology induced on Δ' from M is the weak topology†.)

A PL homeomorphism $\Delta \simeq \Delta'$ can now be described by induction on skeleta, where $\Delta'^{\leq i} := \{D_\sigma : \sigma \in \Delta, \dim \sigma \leq i\}$. We have $\Delta^{\leq 0} \simeq \Delta'^{\leq 0}$, since both are 0-dimensional

†This follows from the fact that in our triangulation of M it is possible to realize Δ' as a subcomplex after a constant number of barycentric subdivisions.

and their vertex sets are in bijection $v \mapsto D_v$. Also $\Delta^{\leq 1} \simeq \Delta'^{\leq 1}$, because for each edge $e = \{u, v\}$ we have $\partial D_e = D_u \cup D_v$, and hence the map $\Delta^{\leq 0} \xrightarrow{\sim} \Delta'^{\leq 0}$ extends to 1-skeleta. It also extends to 2-skeleta, because for each $\sigma = \{u, v, w\} \in \Delta^{(2)}$ the boundary of D_σ is precisely the cycle $D_{\{u,v\}} \cup D_{\{v,w\}} \cup D_{\{w,u\}}$, which is the image of $\partial\sigma$ under the map $\Delta^{\leq 1} \xrightarrow{\sim} \Delta'^{\leq 1}$. □

4.3 Collapsing $M(\Delta)$ onto Δ'

We prove that for each i-face $\sigma \in \Delta$ it is possible to collapse $H_{i,\sigma}$ onto the union of $H_{i,\sigma} \cap \Delta'$ and the subpolyhedron along which $H_{i,\sigma}$ is glued to lower-index handles. We then show that these operations can be carried out within M in order of decreasing handle index, so that $M(\Delta) \searrow \Delta'$.

Our main tool is the next lemma.

Lemma 25. *Let $P \subset \partial I^{d-k}$ be a subpolyhedron, and identify the $(d-k)$-cube I^{d-k} with the cone $c * \partial I^{d-k}$ over its boundary. Then*

$$I^d = I^k \times I^{d-k} \searrow \left[I^k \times (c * P)\right] \cup \left[\partial I^k \times I^{d-k}\right].$$

Proof. Choose a triangulation X of ∂I^{d-k} such that P is a subcomplex, so that $c * X$ is a triangulation of I^{d-k}. Identify I^k with the k-simplex, which we denote Z (as Δ is taken). Consider the cellulation of $I^k \times I^{d-k}$ with cells given by $\{\sigma \times \tau \mid \sigma \in Z, \tau \in c * X\}$. Denote by $\delta \in Z$ the unique k-dimensional face.

Order $X \setminus P$ (i.e. the set of faces of X which are not faces of P) in order of decreasing dimension as $\tau_1, \tau_2, \ldots, \tau_k$. Observe that $\delta \times \tau_1$ is contained in no face other than $\delta \times (c * \tau_1)$, and hence we may collapse $\delta \times (c * \tau_1)$ from $\delta \times \tau_1$ onto $\delta \times (c * \partial\tau_1)$. We do this for every τ_i in turn for $2 \leq i \leq k$, collapsing $\delta \times (c * \tau_i)$ from $\delta \times \tau_i$. By the end of the process we have deleted precisely $\delta \times (c * X \setminus c * P)$, where δ is the interior of the closed disk Z. Hence we are left with

$$[\delta \times (c * P)] \cup [\partial Z \times (c * X)] = [Z \times (c * P)] \cup [\partial Z \times (c * X)]$$

(where the inclusion $\subset$ is clear, and $\supset$ follows from the fact $Z = \delta \cup \partial Z$, so that $Z \times (c * P)$ is contained in the union on the left).

By our identification of I^k and I^{d-k} with the appropriate triangulations we have obtained precisely $I^d = I^k \times I^{d-k} \searrow \left[I^k \times (c * P)\right] \cup \left[\partial I^k \times I^{d-k}\right]$. □

We apply this lemma for each handle $H_{i,\sigma}$, where $k = i = \dim\sigma$ and $H_{i,\sigma}$ is identified with $I^k \times I^{d-k}$:

- If $\sigma = \{v\}$ is a vertex, the lemma states that I^d collapses onto $c * P$ for any subpolyhedron $P \subset \partial I^d$. Since $\Delta' \cap H_{0,v}$ is of the form $c_0 * P$ for a certain subpolyhedron $P \subset \partial H_{0,v}$, we get the desired collapse $H_{0,v} \searrow \Delta' \cap H_{0,v}$.

- If $\sigma = e = \{u, v\}$ is an edge, $\Delta' \cap H_{1,e} = (c_1 * \{u_1, \ldots, u_{s(e)}\}) \times I$ for $c_1 \in \Sigma$ the chosen center point and $s(e)$ the number of 2-faces containing e. To bring this description to the form in the lemma, identify $H_{1,e} = \Sigma \times I$ with $I \times I^{d-1}$ (where the Σ factor maps to the I^{d-1} factor) and set $P \subset I^{d-1}$ the image of the subpolyhedron $\{u_1, \ldots, u_{s(e)}\}$ of Σ. In terms of $H_{1,e}$, we obtain

$$H_{1,e} \searrow \left[\Delta' \cap H_{1,e}\right] \cup [\Sigma \times \{0\}] \cup [\Sigma \times \{1\}].$$

Observe that this is the union of $\Delta' \cap H_{1,e}$ with the part of $\partial H_{1,e}$ that is glued onto the 0-handles of M.

- If σ is a 2-face, we apply the lemma with P empty, so that $c * P = \{c\}$. We obtain a collapse $I^2 \times I^{d-2} \searrow [I^2 \times c] \cup [\partial I^2 \times I^{d-2}]$. Identifying c with the interior point $\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)$ of I^{d-k}, we get

$$I^2 \times I^{d-2} \searrow \left[I^2 \times \left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right] \cup \left[\partial I^2 \times I^{d-2}\right].$$

  Here $I^2 \times \left(\frac{1}{2}, \ldots, \frac{1}{2}\right) = C(\sigma)$, and altogether $\left[I^2 \times \left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right] \cup \left[\partial I^2 \times I^{d-2}\right]$ is a union of $C(\sigma)$ and a regular neighborhood of $C(\sigma) \cap \partial I^d$, by the next lemma. In other words, we obtain a collapse of $H_{2,\sigma}$ onto the union of $\Delta' \cap H_{2,\sigma}$ with the part of $\partial H_{2,\sigma}$ that is glued onto the 0- and 1-handles of M.

Lemma 26. *$(\partial I^2) \times I^{d-2}$ is a regular neighborhood of $\left[I^2 \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}\right] \cap \partial I^d$ within ∂I^d.*

Proof. By [15, Cor. 3.30] it suffices to show that $(\partial I^2) \times I^{d-2} \subset \partial I^d$ is a compact manifold with boundary and that it collapses onto $\left[I^2 \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}\right] \cap \partial I^d$.

It is clear that $(\partial I^2) \times I^{d-2} \simeq S^1 \times D^{d-2}$ is a compact manifold with boundary.

For the statement on collapsing, first observe that $\left[I^2 \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}\right] \cap \partial I^d = \left(\partial I^2\right) \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}$. So we need to prove that $\left(\partial I^2\right) \times I^{d-2} \searrow \left(\partial I^2\right) \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}$. We prove more generally that $P \times I^n \searrow P \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}$ for any compact polyhedron P. It suffices to prove this for $n = 1$ by induction: if the claim is true for some n then

$$P \times I^{n+1} = (P \times I) \times I^n \searrow (P \times I) \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\},$$

and up to permuting coordinates the polyhedron on the right is just $\left(P \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}\right) \times I$, which we know collapses onto $P \times \left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\} \times \left\{\frac{1}{2}\right\} = P \times \underbrace{\left\{\left(\frac{1}{2}, \ldots, \frac{1}{2}\right)\right\}}_{\in I^{d+1}}$.

For $n = 1$, fix a triangulation of P and triangulate I as a union of two edges $\left[0, \frac{1}{2}\right] \cup \left[\frac{1}{2}, 1\right]$. Thus $P \times I$ has a cellulation in which the cells are $\{\sigma \times \tau \mid \sigma \in P, \tau \in I\}$. For each $\dim(P)$-dimensional face σ of P we have that $\sigma \times \{1\}$ is a free face of $P \times I$ (it is contained only in $\sigma \times \left[\frac{1}{2}, 1\right]$) and similarly $\sigma \times \{0\}$ is a free face of $P \times I$ (contained only in $\sigma \times \left[0, \frac{1}{2}\right]$). Collapsing $\sigma \times \left[\frac{1}{2}, 1\right]$ and $\sigma \times \left[0, \frac{1}{2}\right]$ along these free faces , we see

$$P \times I \searrow \left(P \times \left\{\frac{1}{2}\right\}\right) \cup \{\sigma \times \tau \mid \sigma \in P, \tau \in I, \dim \sigma \leq \dim P - 1\}.$$

We continue in this way with the faces of P of dimension $\dim(P) - 1, \ldots, 0$: at the i-th step, we have collapsed $P \times I$ onto

$$\left(P \times \left\{\frac{1}{2}\right\}\right) \cup \{\sigma \times \tau \mid \sigma \in P, \tau \in I, \dim \sigma \leq \dim P - i\},$$

and for each $(\dim P - i)$-dimensional face σ of P we have that $\sigma \times \{1\}$ is contained only in $\sigma \times \left[\frac{1}{2}, 1\right]$ and $\sigma \times \{0\}$ is contained only in $\sigma \times \left[0, \frac{1}{2}\right]$, so we may collapse further onto

$$\left(P \times \left\{\frac{1}{2}\right\}\right) \cup \{\sigma \times \tau \mid \sigma \in P, \tau \in I, \dim \sigma \leq \dim P - (i+1)\}.$$

Continuing this way we eventually collapse $P \times I$ onto $P \times \left\{\frac{1}{2}\right\}$ as desired. $\square$

Proposition 27. *If Δ is finite then $M(\Delta) \searrow \Delta'$.*

For the proof, it is useful to have a criterion for when a collapse $X \searrow Y$ extends to a collapse $X \cup Z \searrow Y \cup Z$.

Lemma 28. *Let $X \subset W$ be polyhedra and let $Z = \mathrm{cl}_W(W \setminus X)$. Let $Y \subset X$ be a subpolyhedron such that $X \searrow Y$ by a sequence of elementary collapses $X = X_0 \Downarrow X_1 \Downarrow \ldots \Downarrow X_t = Y$. For each $0 \leq i < t$ denote by $\sigma_i \subset \tau_i$ the disks such that $X_i \Downarrow X_{i+1}$ is a collapse across τ_i from σ_i. Assume that Z is disjoint from the interiors of σ_i and τ_i for all i. Then $X \cup Z \searrow Y \cup Z$.*

Proof. We first prove that $X_0 \cup Z \Downarrow X_1 \cup Z$. For brevity denote $\sigma = \sigma_0$ and $\tau = \tau_0$. We have $X_0 = X_1 \cup \tau$ where $X_1 \cap \tau$ is a face of τ and σ is the complementary face to $X_1 \cap \tau$, meaning $\sigma = \mathrm{cl}_\tau (\partial\tau \setminus X_1 \cap \tau)$. Hence $X_0 \cup Z = X_1 \cup Z \cup \tau$ and

$$(X_1 \cup Z) \cap \tau = (X_1 \cap \tau) \cup (Z \cap \tau)\,.$$

By assumption, Z is disjoint from the interiors of σ and τ. We have $Z \cap \tau \subseteq X_1 \cap \tau$, because $\tau \setminus X_1$ consists precisely of the union of the interiors of τ and σ (so $\tau \setminus X_1 \subseteq \tau \setminus Z$). Hence $(X_1 \cup Z) \cap \tau = X_1 \cap \tau$ is a face of τ complementary to σ, and we may collapse $X_1 \cup Z \cup \tau = X_0 \cup Z$ onto $X_1 \cup Z$.

Observe that $X_1 \subset X_1 \cup Z$ are polyhedra, and $Z = \mathrm{cl}_{X_1 \cup Z}([X_1 \cup Z] \setminus X_1)$: to see the equality, note that Z is a closed subset of $X_1 \cup Z$ (which is itself a closed subset of $X \cup Z$) and that $[X_1 \cup Z] \setminus X_1 = [X \cup Z] \setminus X$, because the interiors of σ and τ are in X. Hence, setting $W_1 = X_1 \cup Z$, we may apply the claim inductively to the shorter sequence $X_1 \Downarrow \ldots \Downarrow X_n = Y$, with $X_1 \subset W_1$ taking the place of $X \subset W$, and the collapse $X_1 \searrow Y$ taking the place of $X \searrow Y$. □

Proof of the proposition. As described above, we order the handles by decreasing index (2-handles first) as $H_1, \ldots, H_n$ (the ordering induced on the i-handles for each $0 \leq i \leq 2$ is arbitrary). Then, within M we collapse each handle H_i onto $H_i \cap \Delta' \cup \left[H_i \cap \bigcup_{j>i} H_j\right]$, in order. We use Lemma 28 to show that this is indeed possible.

Suppose the handles have been ordered as $H_1, \ldots, H_n$. After collapsing $H_1, \ldots, H_{i-1}$ onto their intersections with Δ' we are left with $M_i := \Delta' \cup \bigcup_{j\geq i} H_j$. Since $M_i = H_i \cup (M_i \setminus H_i)$, it suffices to show that $\mathrm{cl}_{M_i}(M_i \setminus H_i)$ is disjoint from the interiors of the disks removed in the collapse $H_i \searrow H_i \cap \Delta'$. We prove this with $\mathrm{cl}_{M_i}(M_i \setminus H_i)$ replaced by the larger set M_{i+1}. Observe that M_{i+1} intersects the interior of H_i only in Δ' (because distinct handles of M intersect only in their boundaries). Also, every elementary collapse in Lemma 25 is performed from a face σ in ∂H_i across a face τ whose interior is entirely contained in the interior of H_i. Therefore, for such an elementary collapse it suffices to check that M_{i+1} is disjoint from the interior of σ. Also, Δ' is disjoint from the interior of σ (since $\Delta' \cap H_i$ is contained in the subpolyhedron of H_i obtained by deleting the interiors of σ and of τ). Hence it suffices to check that for all $j > i$ the handle H_j is disjoint from the interior of σ. Note that any two k-handles ($k \in \{0, 1, 2\}$) within M have disjoint boundaries. Therefore it suffices to verify that the collapse of H described in Lemma 25 occurs from subdisks of ∂H that do not intersect any lower-index handle. But the collapse is precisely onto the union of $H \cap \Delta'$ with the part of ∂H glued onto the lower-index handles; so no subdisk in the intersection of ∂H with a lower-index handle is removed in the process. □

4.4 Triangulating $M(\Delta)$

We construct triangulations of the handles and their subpolyhedra as defined in Section 4.1. These triangulations depend on $d = \dim M$ and on k (the bound on vertex links of Δ) but not on Δ itself.

It is convenient to demand that k is even, by increasing it by 1 if necessary.

4.4.1 Triangulating the 0- and 1-handles

Choose a combinatorially triangulated $(d-1)$-sphere S_0 with an automorphism group that has an orbit of k vertices, which we denote $d_1, \ldots, d_k$. We additionally require that for each i there is an orientation-reversing automorphism that fixes d_i. An example of such a triangulation is a join $X * Y$, where X is a triangulation of S^1 with k vertices and Y is the boundary of a $(d-k-1)$-simplex.

Replace S_0 by its second barycentric subdivision: all automorphisms are still available because the subdivision respects all symmetries. Also, the closed stars of $d_1, \ldots, d_k$ are pairwise isomorphic and disjoint $(d-1)$-disks, which we denote $\Sigma_1, \ldots, \Sigma_k$. By assumption each Σ_i has an orientation-reversing automorphism that fixes the center point u_i. We choose such automorphisms $\varphi_1, \ldots, \varphi_k$ such that each pair φ_i, φ_j are conjugate via an automorphism of S_0 taking u_i to u_j.

Take further barycentric subdivisions until each of the disks Σ_i has k points in its boundary. These play the role of $q_{i,1}, \ldots, q_{i,k}$ in Section 4.1. We demand that $q_{i,1}, \ldots, q_{i,k}$ are a disjoint union of φ_i-orbits $\{q_{i,1}, q_{i,2}\}, \ldots, \{q_{i,k-1}, q_{i,k}\}$ for each i, and that choices are made symmetrically in the sense that for each $1 \leq i, j \leq k$ there exists an automorphism $\psi_{i,j}$ of S_0 taking this set of orbits to the set of φ_j-orbits $\{\{q_{j,1}, q_{j,2}\}, \ldots, \{q_{j,k-1}, q_{j,k}\}\}$. (This can be done by choosing the points within Σ_i and then taking their images under an appropriate isomorphism $\Sigma_1 \to \Sigma_j$ for each $j > 1$). Then, take further barycentric subdivisions until S_0 contains a system of pairwise disjoint paths $\{p_t\}_t$ connecting q_{i_1,j_1} with q_{i_2,j_2} for all pairs with $i_1 \neq i_2$ (where $1 \leq i_1, i_2, j_1, j_2 \leq k$), where the paths are disjoint from the interiors of the disks $\Sigma_1, \ldots, \Sigma_k$. The cone $c_0 * S_0$ is (almost) our reference handle H_0: S_0 will be replaced with a further subdivision during the construction of the 1-handles.

At this point, take a copy Σ of one of the (pairwise isomorphic) disks Σ_i. Denote the copies of the points $q_{i,1}, \ldots, q_{i,k}$ within Σ by $u_1, \ldots, u_k$. There is a cellulation of $\Sigma \times I$ in which the cells are all products of the form $\sigma \times \tau$ for $\sigma \in \Sigma$ and $\tau \in I$ (where I is triangulated by a single edge). Take a combinatorial triangulation refining this cellulation, in a way that the induced triangulation of $\Sigma \times \{0\}$ and $\Sigma \times \{1\}$ is precisely the original triangulation of Σ. One way of doing this is to stellate in the interior of each face which is not in $\Sigma \times \{0\}$ or $\Sigma \times \{1\}$, in order of decreasing dimension. In this triangulation of $\Sigma \times I$ the subpolyhedra $u_i \times I$ are pairwise disjoint subcomplexes. Further barycentrically subdivide this triangulation twice in order to obtain that the simplicial neighborhoods of the subcomplexes $u_i \times I$ are regular neighborhoods. This is our triangulation of H_1.

Finally, replace S_0 with its second barycentric subdivision, so that the copies of (subdivisions of) $\Sigma_1, \ldots, \Sigma_k$ within it are isomorphic to our triangulation of $\Sigma \times \{0\}$ and $\Sigma \times \{1\}$. Our triangulation of H_0 is $c_0 * S_0$. Note that the further subdivision ensures that simplicial neighborhoods of the paths $\{p_t\}_t$ are pairwise disjoint regular neighborhoods.

Proposition 29. *$M(\Delta)$ has a triangulation in which each handle $H_{i,\sigma}$ is a subcomplex.*

The total number of faces in each such subcomplex is uniformly bounded in terms of d, k (and in particular is independent of Δ).

Proof. It is possible to glue together M_1 from our triangulations of H_0 and H_1, where the gluing maps are simplicial isomorphisms gluing $\Sigma \times \{0\}$ or $\Sigma \times \{1\}$ in a copy of H_1 onto some Σ_i in a copy of H_0, because there are simplicial isomorphisms $\Sigma \times \{0\} \xrightarrow{\sim} \Sigma_i$ and $\Sigma \times \{1\} \xrightarrow{\sim} \Sigma_i$. In fact, there are at least two such simplicial isomorphisms, with opposite orientations - given one, a second can be obtained by composing with φ_i. So the requirements on orientation in M_1 can also be met: by construction of $q_{i,1}, \ldots, q_{i,k}$ as a φ_i-invariant set, both of these isomorphisms map $u_1, \ldots, u_k \in \partial (\Sigma \times \{0\})$ to $q_{i,1}, \ldots, q_{i,k} \in \partial\Sigma_i$.

Note that the triangulation of M_1 we have obtained satisfies the conditions of the proposition, and it remains to triangulate the 2-handles in some way. Each 2-handle $H_{2,\sigma}$ is glued onto a regular neighborhood of a path γ_σ within ∂M_1, which in our triangulation is a subcomplex. In fact, by construction our triangulation of M_1 is the second barycentric subdivision of a complex in which γ_σ is a subcomplex; in ∂M_1, its simplicial neighborhood is a regular neighborhood (and this regular neighborhood of γ_σ is disjoint from the regular neighborhood of γ_τ for any 2-face $\tau \neq \sigma$).

The simplicial neighborhood of γ_σ in ∂M_1 has one of only a finite number of isomorphism types. To see this, observe that if $\sigma = \{u, v, w\}$ then the simplicial neighborhood of γ_σ is a subcomplex of our triangulation of

$$H_{0,u} \cup H_{0,v} \cup H_{0,w} \cup H_{1,\{u,v\}} \cup H_{1,\{v,w\}} \cup H_{1,\{w,u\}},$$

which is a union of a bounded number of handles and hence has a uniformly bounded number of faces; there are only finitely many such isomorphism types of subcomplexes of simplicial complexes (and their number depends only on k and d). Let $\Gamma_1, \ldots, \Gamma_N$ be the total number of possible isomorphism types of simplicial neighborhoods of the paths $\gamma_\sigma \subset \partial M_1$. For each isomorphism type Γ_i fix a triangulation of a 2-handle $I^2 \times I^{d-2}$ satisfying that the induced triangulation of $\partial I^2 \times I^{d-2}$ is isomorphic to Γ_i: this is possible by [2]. Finiteness ensures that the number of faces in the triangulation of each 2-handle can be bounded in terms of d and k alone. □

5 Topology of the boundary

The goal of this section is to prove the following theorem.

Theorem 30. *Let M be a connected manifold.*

1. *If M is PL-collapsible then ∂M is a sphere.*
2. *If ∂M is connected and R is a PID such that $H_k(M; R) = 0$ for all $k > 0$ then ∂M is a triangulated R-homology sphere.*
3. *If M is contractible, $\dim M \geq 5$, and M has a handle structure consisting only of 0-, 1-, and 2-handles, then ∂M is a homotopy sphere.*

Remark 31. All three parts are routine or known, but we include the proof for completeness. The assumption that ∂M is connected in (2) is necessary: consider $\mathbb{RP}^2 \times I$ over $R = \mathbb{Q}$. One can replace it with the assumption that 2 is not invertible in R.

Proof. Denote $d = \dim M$.

1. If M is PL-collapsible then M is a PL disk by a theorem of Whitehead (see [15, Corollary 3.28]). Therefore ∂M is a sphere.

2. Consider the long exact sequence

$$\ldots \to H_k(M;R) \to H_k(M,\partial M;R) \to H_{k-1}(\partial M;R) \to H_{k-1}(M;R) \to \ldots$$

of the pair $(M, \partial M)$. The assumption that $H_k(M;R) = 0$ for all $k > 0$ implies that $H_k(M,\partial M;R) \cong H_{k-1}(\partial M;R)$ for all $k > 1$.

Observe that around $k = 1$ the sequence reads

$$0 = H_1(M;R) \to H_1(M,\partial M;R) \to H_0(\partial M;R) \to H_0(M;R) \to 0 = H_0(M,\partial M;R)$$

where $H_0(M;R) \simeq R$ and $H_0(\partial M;R) \simeq R$ by the assumption that ∂M is connected. Hence $H_1(M,\partial M;R) = 0$ so that the pair $(M,\partial M)$ is R-orientable, $H_d(M,\partial M;R) \simeq R$, and Lefschetz duality applies.

By Lefschetz duality $H_k(M,\partial M;R) \simeq H^{d-k}(M;R)$. The universal coefficient theorem for cohomology states that for each k there is an exact sequence of the form

$$0 \to \mathrm{Ext}^1_R(H_{k-1}(M;R), R) \to H^k(M;R) \to \mathrm{Hom}_R(H_k(M;R), R) \to 0.$$

If $k > 0$ we have $\mathrm{Hom}_R(H_k(M;R), R) = \mathrm{Hom}_R(0, R) = 0$. Further, if $k > 1$ we have

$$\mathrm{Ext}^1_R(H_{k-1}(M;R), R) = \mathrm{Ext}^1_R(0, R) = 0,$$

while if $k = 1$ we have

$$\mathrm{Ext}^1_R(H_{k-1}(M;R), R) = \mathrm{Ext}^1_R(R, R) = 0$$

since R is free. Therefore

$$0 = H^k(M;R) \simeq H_{d-k}(M,\partial M;R) \simeq H_{d-k-1}(\partial M;R) \text{ for all } 0 < k < d-1.$$

(For $k = d-1$ this fails because $0 = H_1(M,\partial M;R) \not\simeq H_0(\partial M;R) \simeq R$, as already observed above).

Finally, $H_{d-1}(\partial M;R) \simeq H_d(M,\partial M;R) \simeq R$, where the last isomorphism holds because $(M,\partial M)$ is R-orientable.

3. By part (2), ∂M is a $\mathbb{Z}$-homology sphere. A simply connected homology sphere is a homotopy sphere, so it suffices to prove ∂M is simply-connected.

By eliminating the 0-handles beyond the first from the given handle decomposition and reordering, we may assume that M is obtained from a disk D (the unique 0-handle) by attaching a collection of mutually disjoint 1-handles and then a collection of 2-handles which are disjoint from each other (but not from the 1-handles): there is a sequence of the form

$$D = M_0, M_1, \ldots, M_r, M_{r+1}, \ldots, M_{r+s} = M,$$

such that M_{i+1} is obtained from M_i by adding a 1-handle for each $i < r$, and M_{r+i+1} is obtained from M_{r+i} by adding a 2-handle for $i < s$. The boundary ∂M can be described by a sequence of surgeries on the boundary S^{d-1} of the disk D, one for each k-handle added to M. More explicitly, given an attaching map $f_j : S^{k-1} \times D^{d-k} \to \partial M_j$ of a k-handle to M_j, ∂M_{j+1} can be described as the result of the surgery step that removes $f_j\left(S^{k-1} \times \operatorname{int}\left(D^{d-k}\right)\right)$ and glues a copy of $D^k \times S^{d-k-1}$ onto $f_j\left(S^{k-1} \times S^{d-k-1}\right)$.

The inclusion map $\partial M_j \hookrightarrow M_j$ induces an isomorphism on fundamental groups for each j. Choose some $x \in \partial D$ such that x is not in the image of any f_j (and hence $x \in \partial M_j$ for each j). We prove the claim by induction on j: for $j = 0$ it is clear. If $j+1 \leq r$ then M_{j+1} is obtained by adding a 1-handle to M_j, and ∂M_{j+1} is obtained by removing two disjoint open disks from ∂M_j and gluing a cylinder $S^{d-2} \times I$ between their boundaries. Both steps add a free generator to the fundamental group, so that $\pi_1(M_r, x) \simeq F_r \simeq \pi_1(\partial M_r, x)$, and the inclusion $\partial M_r \hookrightarrow M_r$ induces such an isomorphism. If $r < j+1 \leq r+s$ then a 2-handle is added to M_j along a map $f_j : S^1 \times D^{d-2} \to \partial M_j$, which annihilates the conjugacy class of the restriction to $S^1 \times \{0\}$ in $\pi_1(M_j, x)$. Since $\dim \partial M_j \geq 4$, the corresponding surgery on ∂M_j has precisely the same effect of annihilating the conjugacy class of the corresponding element in $\pi_1(\partial M_j, x)$: see [13, Lemma 2].

$\square$

6 Encoding facet colorings in manifold triangulations

Let Δ be a triangulated d-dimensional homology manifold and $c : \Delta^{(d)} \to [r]$ a coloring of the facets with r colors. We wish to encode the pair (Δ, c) in a single triangulation of a homology manifold Δ', in such a way that Δ' is not much larger than Δ. In this section we provide such a construction and prove some of its basic properties.

Note that for many purposes (e.g. asymptotic enumeration) it would suffice to work with colored triangulations: so long as the number of colors r is uniformly bounded, the number of pairs (Δ, c) is larger than the number of triangulated manifolds Δ with N facets by only an exponential factor.

Definition 32. We call a pair (Δ, c) of a d-dimensional homology manifold Δ (without boundary) and a facet coloring $c : \Delta^{(d)} \to [r]$ an *r-colored triangulated homology d-manifold.* An isomorphism of r-colored triangulated homology d-manifolds $(\Delta_1, c_1) \xrightarrow{\sim} (\Delta_2, c_2)$ is a simplicial isomorphism $\phi : \Delta_1 \to \Delta_2$ such that for each facet σ of Δ_1 it holds that $c_2(\phi(\sigma)) = c_1(\sigma)$.

Theorem 33. *For each $d, r \in \mathbb{N}$ with $d \geq 2$ there is a constant $m \in \mathbb{N}$ and a function*

$\mathrm{Enc}_{d,r}$: {*r-colored triangulated homology d-manifolds*} $\to$ {*triangulated homology d-manifolds*}

such that if $\Delta' = \mathrm{Enc}_{d,r}(\Delta, c)$ then:

1. *Δ' is a subdivision of Δ in which each facet is replaced with a combinatorial d-disk on at most m facets. In particular, there is a PL homeomorphism $\Delta' \simeq \Delta$ and $\left|\Delta'^{(d)}\right| \leq m \left|\Delta^{(d)}\right|$.*

2. *If each vertex of Δ is incident to at most n edges, then each vertex of Δ' is incident to at most mn edges,*

3. In addition, $\mathrm{Enc}_{d,r}$ *is injective on isomorphism classes.*

Remark 34. The functions $\mathrm{Enc}_{d,r}$ we obtain are easy to compute, and so are their left inverses.

Summary. The idea of the proof is to subdivide Δ, and replace each of its facets σ with a PL-triangulated disk σ' such that the triangulation of σ' encodes the color $c(\sigma)$. It is automatic that condition (1) of the theorem holds for such a construction. Since we only need a finite "palette" of disk triangulations σ', condition (2) holds as well. There is a mild difficulty with condition (3): it is not clear that Δ can be reconstructed from the resulting complex. Our strategy is to first replace Δ with its barycentric subdivision $\mathrm{b}(\Delta)$, and the coloring c by the facet coloring $\mathrm{b}(c)$ that it induces (each facet of $\mathrm{b}(\Delta)$ is given the color of the unique facet of Δ that it refines). We then perform a further subdivision to encode the induced facet coloring.

The barycentric subdivision is useful because in $\mathrm{b}(\Delta)$ every vertex has high degree. Our further subdivision will introduce only vertices of lower degree, so that we will be able to identify the vertices of $b(\Delta)$ in the final subdivided complex Δ'. This will ensure that $\mathrm{b}(\Delta)$ and its coloring $\mathrm{b}(c)$ can both be reconstructed from the resulting complex. Except for certain degenerate cases, Δ can be reconstructed from $\mathrm{b}(\Delta)$ by [6]. We enrich the facet coloring c with more information in order to deal with these degenerate cases.

Lemma 35. *Let* Δ *be a triangulated homology d-manifold without boundary* $(d \geq 2)$*, and denote its barycentric subdivision by* $\mathrm{b}(\Delta)$*. Then each vertex of* $\mathrm{b}(\Delta)$ *is incident to at least* $d+2$ *edges.*

Proof. A vertex v of $\mathrm{b}(\Delta)$ is the barycenter of a nonempty face σ of Δ. Denote $k = \dim \sigma$. Each vertex w neighboring v in the 1-skeleton of $\mathrm{b}(\Delta)$ is either the barycenter of a face τ satisfying $\emptyset \subsetneq \tau \subsetneq \sigma$ (a total of $2^{k+1} - 2$ faces) or the barycenter of a face τ satisfying $\tau \supsetneq \sigma$. The number of faces that properly contain σ is the number of nonempty faces in $\mathrm{lk}(\sigma, \Delta)$, which is a homology sphere of dimension $d - k - 1$.

We split into cases depending on the value of $s = \dim \mathrm{lk}(\sigma, \Delta) = d - k - 1$. Note that $s \geq -1$.

Case 1. If $s = -1$ then $k = d$ and σ contains $2^{k+1} - 2 = 2^{d+1} - 2 \geq d + 2$ proper nonempty faces. Thus v neighbors at least $d + 2$ vertices.

Case 2. If $s = 0$ then $k = d-1$, and σ contains $2^{k+1}-2 = 2^d-2$ proper nonempty faces, and is contained in precisely 2 faces of Δ (since Δ is a homology manifold, and the link of σ is a 0-sphere). Again v neighbors at least

$$2^d \geq d + 2$$

vertices in the 1-skeleton of $\mathrm{b}(\Delta)$.

Case 3. Suppose $s \geq 1$. Any homology sphere Σ of dimension $s \geq 1$ has at least $s + 2$ vertices and $s + 2$ facets: it has at least one facet ρ of dimension s, which has $s + 1$ vertices, and at least one vertex not on this face. Each of the maximal proper faces of ρ is a ridge (or codimension-1 face) of Σ, and hence contained in a facet of Σ other than ρ. Hence there are at least $s + 1$ facets of Σ other than ρ itself. It therefore suffices to check that

$$\left(2^{k+1} - 2\right) + 2\,(d - k - 1 + 2) \geq d + 2.$$

Observe that

$$\left(2^{k+1}-2\right)+2\,(d-k+1)=2\left(d+2^{k}-k\right)$$

and that $2^k - k \geq 1$, which proves the claim.

□

The family of triangulated disks in the next definition will form the "color palette" for our constructions.

Definition 36. Let $d \in \mathbb{N}$. We define $\mathcal{S}_d$ to be the family of simplicial complexes that can be obtained from the d-simplex by iterated stellar subdivision at facets. In other words, a complex Δ is in $\mathcal{S}_d$ if there exists a finite sequence of complexes $\Delta_0, \Delta_1, \ldots, \Delta_n$ such that Δ_0 is the d-simplex, $\Delta_n = \Delta$, and Δ_{i+1} is obtained from Δ_i by stellating a facet.

Remark. Note that any $\Delta \in \mathcal{S}_d$ is a triangulated PL disk, and its boundary is isomorphic to the boundary of the simplex. Note also that Δ_1 in any sequence as above is uniquely defined: it is the cone over the boundary of the d-simplex.

Lemma 37. *Let (Δ, c) be an r-colored triangulated homology d-manifold ($d \geq 2$) in which all vertices have degree at least $d+2$. Choose any r non-isomorphic complexes $\Sigma_1, \ldots, \Sigma_r \in \mathcal{S}_d$. Construct a new complex Δ' as follows: delete the interior of each facet σ of Δ, and glue $\Sigma_{c(\sigma)}$ to the deleted complex via an (arbitrary) simplicial isomorphism $\partial\Sigma_{c(\sigma)} \to \partial\sigma$. Then the isomorphism class (Δ, c) can be reconstructed from Δ'.*

In particular, this construction is injective on the set of all isomorphism classes of r-colored triangulated homology d-manifolds ($d \geq 2$) with all vertex degrees at least $d+2$.

Claim 38. If $\Delta_0 = \Delta, \Delta_1, \ldots, \Delta_k$ is a sequence in which each Δ_{i+1} is constructed from Δ_i by stellar subdivision at a facet, and in Δ all vertices have degree at least $d+2$, then the process of greedily performing reverse stellations at facets while possible always terminates with Δ.

Proof. Note that reverse stellar subdivision at a facet is precisely the removal of the open star of a vertex of degree $d+1$ (whose link must be a simplicial sphere with $d+1$ vertices, hence is the boundary of a simplex) and the replacement of this open star by the interior of a single d-simplex. On the 1-skeleton this process has precisely the effect of deleting a vertex and its edges. Hence we can reverse-stellate and delete a given vertex if and only if its degree in the 1-skeleton is $d+1$. Note that this means that a vertex of Δ can never be deleted in the greedy process: if $v \in \Delta$ is the first vertex deleted this way, all its neighbors in Δ must have been present before the deletion, at which point its degree in the 1-skeleton was at least $d+2$ by assumption.

It is convenient to realize Δ in $\mathbb{R}^k$ (for some large k, e.g. $2d+1$) in order to use the following geometric notion of subdivision: a subdivision of Δ is a simplicial complex in $\mathbb{R}^k$ with the same support as Δ, and such that every face is contained in a unique face of Δ.

Call a vertex of a subdivision of Δ *young* if it is not a vertex of Δ. The neighbors of a young vertex are all contained in one facet of Δ, and any reverse stellation which deletes a given (necessarily young) vertex v preserves this property (since it preserves the neighbors' positions). It follows that a reverse stellation at a facet of a subdivision of Δ always results in a subdivision of Δ. The young vertices of Δ_k are ordered by age (each vertex is introduced at a different Δ_i, with the youngest being the last added,) and if $\overline{\Delta}$ is the result of any sequence of reverse facet stellations from Δ_k, a reverse facet stellation deleting a vertex v of $\overline{\Delta}$ may be performed whenever v is younger than all its neighbors in $\overline{\Delta}^{\leq 1}$: if v has no younger neighbor then its degree is at most $d+1$, but the degree of a vertex in a d-dimensional homology manifold is always at least $d+1$, so $\deg_{\overline{\Delta}}(v) = d+1$. Hence the process can never get stuck except at Δ (in which we know no reverse stellation is possible): if $\overline{\Delta} \neq \Delta$, we may always reverse-stellate further by deleting the youngest vertex left in $\overline{\Delta}$. $\square$

Proof of the lemma. The claim above shows that iteratively performing all possible reverse stellations at facets recovers Δ from Δ'. We can embed Δ in some $\mathbb{R}^k$ and re-perform the stellations to obtain Δ' as a geometric subdivision of Δ; the combinatorial isomorphism type of the simplicial disk refining a given facet σ of Δ contains precisely the information of the color $c(\sigma)$. $\square$

Proof of Theorem 33. The main theorem of [6] states that if K, L are connected, locally finite simplicial complexes such that $b(K)$ and $b(L)$ are isomorphic then K is isomorphic to L. Further, if $|K|$ is not a 1-manifold and K is neither a simplex nor the boundary of a simplex then the isomorphism between $b(K)$ and $b(L)$ induces an isomorphism of K and L.

In other words, given $b(K)$ where $\dim K \geq 2$ we can determine the isomorphism type of K. Unless K is the simplex or the boundary of a simplex, we can further determine those subcomplexes of $b(K)$ that refine faces of K. The "missing information" in the exceptional cases is the information assigning to each vertex v of $b(K)$ the dimension of the minimal face σ of K that contains it (equivalently, the dimension of the face σ of which v is the barycenter).

Fix $d, r \in \mathbb{N}$ and choose $r \cdot (d+1)$ nonisomorphic complexes $\Sigma_{(1,0)}, \ldots, \Sigma_{(r,d)}$ from $\mathcal{S}_d$. Let a triangulated homology d-manifold Δ be given together with a facet coloring $c : \Delta^{(d)} \to [r]$. We use the following observation: for any face $\sigma \in \Delta$ there exists a vertex $v_\sigma \in b(\Delta)$ at the barycenter of σ, and the stars of the vertices $\{v_\sigma\}_{\sigma \in \Delta}$ within $b(b(\Delta))$ partition the facets of $b(b(\Delta))$. Apply Lemma 37 to $b(b(\Delta))$ (in which all vertex degrees are at least $d+2$ by Lemma 35) as follows: given a facet τ in the star of v_σ within $b(b(\Delta))$, τ is contained in a unique facet τ' of Δ. Color τ by $(c(\tau'), \dim \sigma)$, or in other words, replace τ by $\Sigma_{(c(\tau'), \dim \sigma)}$. Perform this process simultaneously for all facets of $b(b(\Delta))$, and call the result $\mathrm{Enc}_{d,r}(\Delta, c)$. Note that the "color palette" $\Sigma_{(1,0)}, \ldots, \Sigma_{(r,d)}$ is part of $\mathrm{Enc}_{d,r}(\Delta, c)$, and is also needed in order to compute the left inverse.

Statements (1,2) of the theorem follow automatically from the construction. For statement (3), note that Lemma 37 guarantees we can reconstruct the isomorphism type of $b(b(\Delta))$ together with its facet coloring from the isomorphism type of $\mathrm{Enc}_{d,r}(\Delta, c)$. We can then reconstruct $b(\Delta)$ together with, for each vertex, the dimension of the face of Δ of which it is a barycenter; this allows us to reconstruct the isomorphism type of Δ. Further, the given facet coloring of $b(b(\Delta))$ by colors in $(1,0), \ldots, (r,d)$ induces the coloring c on the facets of Δ by taking just the first coordinate, because all facets of $b(b(\Delta))$ refining a given facet of Δ have colors with the same first coordinate. $\square$

7 Encoding manifold triangulations in colorings of the dual graph

Given a triangulated d-manifold M, the dual graph D does not always uniquely determine the simplicial isomorphism type of M: examples of this phenomenon exist in triangulations of the torus and the projective plane ([5]). We produce colorings of D with constantly many colors (depending on d) from which M can be reconstructed.

We begin by describing the data our colorings are designed to encode.

Definition 39. Let $<$ be a total order on $M^{(0)}$, and denote the vertices by $v_1 < \ldots < v_n$. Each facet σ of M is of the form $\left\{v_{i_1}, \ldots, v_{i_{d+1}}\right\}$ (where the indices are chosen such that $i_1 < i_2 < \ldots < i_{d+1}$). If τ is a facet of M that intersects σ at a $(d-1)$-face then $\sigma \setminus \tau = \left\{v_{i_j}\right\}$ is a singleton, and it has a well-defined index j within the order induced on the vertices of σ. In this situation, denote

$$\widetilde{\text{missing}}_{<}(\sigma, \tau) = j.$$

Thus $\widetilde{\text{missing}}_{<}$ is a function from ordered pairs of adjacent facets into $\{1, \ldots, d+1\}$. By identifying ordered pairs of adjacent facets in M with the corresponding oriented edges of D, we obtain a function

$$\text{missing}_{<} : \left\{(v, w) \in D^{(0)} \times D^{(0)} \mid \{v, w\} \text{ is an edge of } D\right\} \to \{1, \ldots, d+1\}.$$

This is the *missing vertex function* associated to $<$.

Claim 40. Let M be a triangulated homology d-manifold and let $<$ be any linear order on $M^{(0)}$. The simplicial isomorphism type of M can be reconstructed from its dual graph D together with the function $\text{missing}_{<}$.

Remark 41. More generally, the proof works whenever all maximal faces in M have dimension d, and the dual graph of the star of each vertex of M is connected.

Proof. Let X be the simplicial complex given by a disjoint union of d-simplices in bijection with $D^{(0)}$ (together with all their faces). Denote the vertices of a d-simplex σ corresponding to $v \in D^{(0)}$ by $u_{v,1}, \ldots, u_{v,d+1}$. Define an equivalence relation $\sim$ on the vertices of X: whenever $\{v, w\}$ is an edge of D with $i_v = \text{missing}_{<}(v, w)$ and $i_w = \text{missing}_{<}(w, v)$, define corresponding vertices in the tuples $(u_{v,1}, \ldots, \hat{u}_{v,i_v}, \ldots, u_{v,d+1})$ and $(u_{w,1}, \ldots, \hat{u}_{w,i_w}, \ldots, u_{w,d+1})$ to be equivalent, where $\hat{u}_{v,i_v}$ denotes that the element is excluded from the tuple. For a face $\tau \in X$, denote $[\tau]_{\sim} = \{[u]_{\sim} \mid u \in \tau\}$ (the set of equivalence classes of vertices of τ), and finally set

$$X/\sim := \{[\tau]_{\sim} \mid \tau \in X\}.$$

There is a simplicial map $\varphi : X \to M$ which is bijective on d-faces: if $\{u_{w,1}, \ldots, u_{w,d+1}\}$ is a facet of X and $\sigma \in M^{(d)}$ is the facet corresponding to $w \in D^{(0)}$, the vertices of σ are $v_{i_1} < \ldots < v_{i_{d+1}}$ (in the order induced by $<$). Define $\varphi(u_{w,j}) = v_{i_j}$ for each $1 \leq j \leq d+1$. Since each face of M is contained in a d-face, φ is surjective on i-faces for all i. It suffices to prove that φ factors through the quotient $X/\sim$ and that the induced simplicial map $X/\sim \to M$ is an isomorphism.

Observe that if $v \in M^{(0)}$ then $\varphi^{-1}(v)$ is a collection of vertices, one contained in each facet of X that corresponds (via φ) to a facet of $\text{st}_v(M)$. If σ, τ are two such facets of

X which are adjacent in the dual graph of $\mathrm{st}_v(M)$, then $\sim$ identifies the unique vertex in $\sigma \cap \varphi^{-1}(v)$ with the unique vertex in $\tau \cap \varphi^{-1}(v)$. Since the dual graph of $\mathrm{st}_v(M)$ is connected, $\varphi^{-1}(v)$ is precisely an equivalence class under $\sim$, so φ induces a bijection $(X/\sim)^{(0)} \to M^{(0)}$. It also follows that if $\rho \in X$ is any face then

$$\varphi([\rho]_\sim) = \{\varphi([u]_\sim) \mid u \in \rho\} = \{\varphi(u) \mid u \in \rho\} = \varphi(\rho),$$

so φ is well defined on the faces of $X/\sim$. We have $\varphi^{-1}(\varphi(\rho)) = \{[u]_\sim \mid u \in \rho\} = [\rho]_\sim$, and hence the induced map is injective on faces of $X/\sim$ (surjectivity follows from the surjectivity of φ). $\square$

Proposition 42. *Let D be the dual graph of a triangulated homology d-manifold M. There is a vertex coloring c of D with boundedly many colors (depending on d) such that (D, c) uniquely reconstructs the simplicial isomorphism type of M.*

Proof. Choose a linear ordering $<$ of the vertices of M and define the function $\mathrm{missing}_<$ as above. We first pick a coloring c_0 of D such that any two vertices at distance at most 2 have distinct colors. By Brooks' theorem this requires at most $k = (d+1)^2 + 1$ colors. We now define a coloring c_1 of D: for each $v \in D^{(0)}$, list its neighbors $w_1, \ldots, w_{d+1}$ (in some arbitrary ordering) and define $c_1(v)$ to be the following tuple of pairs:

$$c_1(v) = (c_0(w_1), \mathrm{missing}_<(v, w_1)), \ldots, (c_0(w_{d+1}), \mathrm{missing}_<(v, w_{d+1})) \in ([k] \times [d+1])^{d+1}.$$

Now define $c(v) = (c_0(v), c_1(v))$ for each $v \in D^{(0)}$. This is a coloring with at most $k \cdot (k(d+1))^{d+1}$ colors. It is clear that knowing c allows us to reconstruct $\mathrm{missing}_<$, and hence the simplicial isomorphism type of M. $\square$

8 Graphs and triangulated manifolds

We put together our previous results to prove the main results of the paper.

Theorem 43. *Let $\mathcal{C}$ be a class of simplicial 2-complexes in which all vertex degrees are bounded by some $k \in \mathbb{N}$. Fix a dimension $d \geq 3$. Then there is a function*

$$F : \mathcal{C} \to \{\textit{triangulated } d\textit{-manifolds}\}$$

and a constant $c > 0$ satisfying that for each $\Delta \in \mathcal{C}$:

1. *If Δ has n vertices then $F(\Delta)$ has at most cn facets.*
2. *If Δ is nulhomologous over a field $\mathbb{F}$ then $F(\Delta)$ is a homology sphere over $\mathbb{F}$; if Δ is collapsible then $F(\Delta)$ is a d-sphere; if Δ is contractible and $d \geq 4$ then $F(\Delta)$ is a homotopy sphere.*
3. *$F(\Delta)$ is quasi-isometric to the 1-skeleton of Δ with quasi-isometry constant c.*
4. *F is injective on isomorphism types: if $F(\Delta)$ is simplicially isomorphic to $F(\Delta')$ then Δ is simplicially isomorphic to Δ'.*

The idea is to thicken Δ into a triangulated manifold $M(\Delta)$ and pass to the boundary. The output of F is more or less $\partial M(\Delta)$: this satisfies all requirements except possibly the injectivity of the construction, which requires us to encode additional information in the triangulation. We do this using the results of Section 6.

Proof. From the complexes $\Delta \in \mathcal{C}$ we construct $(d+1)$-dimensional triangulated manifolds with boundary $M(\Delta)$ using Theorem 22 (the input degree bound is the degree bound of complexes in $\mathcal{C}$). For such a manifold $M(\Delta)$, we color the facets of $\partial M(\Delta)$ with 3 colors, as follows: if σ is such a facet then is contained in a unique handle H of $M(\Delta)$, and H is an i-handle corresponding to an i-face of Δ. Assign σ the color $c(\sigma) = i$.

Claim. The isomorphism type of Δ can be reconstructed from the colored triangulation $(\partial M(\Delta), c)$.

Proof. Observe that for each handle $H \subset M(\Delta)$ the boundary $\partial H \subset M(\Delta)$ is a triangulated d-sphere. The submanifold-with-boundary $\partial H \cap \partial M(\Delta)$ has path-connected interior: its complement within ∂H is exactly the intersection of H with the other handles of $M(\Delta)$, and these intersections are regular neighborhoods of points, arcs, and circles within ∂H. In particular, $\partial H \setminus (\partial H \cap \partial M(\Delta))$ is the regular neighborhood of a graph. The complement of a graph in a d-manifold (where $d \geq 3$) is path connected, and hence so is the complement of its closed regular neighborhood. In particular, $\partial H \cap \partial M(\Delta)$ is a triangulated d-manifold with boundary that has a connected dual graph.

Note that for handles $H, H' \subset M(\Delta)$ corresponding to faces ρ, ρ' of Δ we have $H \cap H' \neq \emptyset$ if and only if $\rho \subset \rho'$ or $\rho' \subset \rho$, and in this case also $\partial H \cap \partial H' \cap \partial M$ intersect in a $(d-1)$-dimensional subcomplex. Given $(\partial M(\Delta), c)$, let D be the dual graph of $\partial M(\Delta)$ (so that c induces a vertex coloring of D). The connected components in D of each color class $c(i)$ are in bijection with the i-faces of Δ, and a connected component K of a color class $c(i)$ has an edge to a connected component L of a color class $c(i+1)$ if and only if the face of Δ corresponding to K is contained in the face of Δ corresponding to L. Thus the isomorphism type of Δ can be reconstructed (D, c). $\square$

Apply Theorem 33 to the colored triangulations $(\partial M(\Delta), c)$ we obtain from $\mathcal{C}$. Composing these constructions gives a function

$$F : \mathcal{C} \to \{\text{triangulated } d\text{-manifolds}\}\,.$$

This function is injective on isomorphism classes by the claim above and by Theorem 33, which proves (4). The fact that the total number of faces in Δ is at most a constant times the number of vertices (because of the degree bound) together with part (1) of Theorem 33 and part (3) of Theorem 22 implies part (1). Part (2) follows from the fact that $\mathrm{Enc}_{d,3}\,(\partial M(\Delta), c)$ is PL-homeomorphic to $\partial M(\Delta)$, together with the fact (part (2) of Theorem 22) that $M(\Delta) \searrow \Delta$: this implies that if Δ is collapsible then so is $M(\Delta)$, and in general Δ is homotopy equivalent to $M(\Delta)$, and the topological conclusions follow directly from Theorem 30. It remains to prove (3).

Let $\Delta \in \mathcal{C}$ and let $X = F(\Delta) = \mathrm{Enc}_{d,3}(\partial M(\Delta), c)$ be the corresponding triangulated manifold. Denote the dual graph of X by D. Since X is a subdivision of $\partial M(\Delta)$ in which each facet is replaced by a combinatorial disk with (uniformly) boundedly many facets, for each handle $H \subset M(\Delta)$ we have a subdivision of $H \cap \partial M(\Delta)$ induced by X. Denote the set of vertices of D corresponding to the facets of $H \cap \partial M(\Delta)$ by D_H. Then $\{D_H \mid H \text{ is a handle of } M\}$ partitions the vertex set of D. By construction, each D_H contains boundedly many vertices, with the bound depending only on the degree bound of complexes in $\mathcal{C}$ and on the dimension d. By the proof of the claim above, each D_H is also nonempty and induces a connected subgraph of D. Fix M to be a uniform bound on the number of vertices in a subset D_H, and note that this is also a bound on the diameter of the induced subgraph of such a set. Define a relation

$$R \subset (\text{vertices of } D) \times (\text{vertices of } \Delta)$$

so that xRv if and only if $x \in D_H$ where H is a handle of M corresponding to a face $\sigma \in \Delta$ satisfying $v \in \sigma$. We prove that R satisfies the assumptions of Lemma 12:

- If x is a vertex of D then $x \in D_H$ for some H, and the face of Δ corresponding to H contains a vertex v for which xRv.
- If v is a vertex of Δ and H is the corresponding handle of $M(\Delta)$ then $D_H \neq \emptyset$, so there is some $x \in D_H$ with xRv.
- Let x, x' be vertices of D at distance at most 1, and assume xRv and $x'Rv'$. Then $x \in D_H$ and $x' \in D_{H'}$ for H, H' handles of M corresponding to either the same face $\sigma' = \sigma \in \Delta$ or (without loss of generality) to a pair $\sigma' \subset \sigma$ of faces of Δ. Further, $v \in \sigma$ and $v' \in \sigma'$ by definition of R. Therefore either $v = v'$ or v, v' are adjacent in the 1-skeleton (since both are vertices of the same face σ). Hence $d(v, v') \leq 1$.
- Let v, v' be vertices of Δ at distance at most 1, and assume xRv and $x'Rv'$. Then $x \in D_H$ and $x' \in D_{H'}$ for H, H' handles of M corresponding to faces σ, σ' containing v and v' respectively. If H_v and $H_{v'}$ are the handles corresponding to v and v', then there is an edge between each consecutive pair of subsets in the sequence $D_H, D_{H_v}, D_{H_{v'}}, D_{H'}$, and these subsets induce connected subgraphs of diameter at most M. Hence $d(x, x') \leq 4M + 3$.

By Lemma 12, D is quasi-isometric to $\Delta^{\leq 1}$ with constants depending only on the dimension d and the degree bound k of complexes in $\mathcal{C}$. □

Corollary 44 (Theorem 6)**.** *For each $d \geq 3$ there is an infinite family of combinatorially distinct triangulated d-spheres such that their dual graphs form a family of $(d+1)$-regular expanders. Moreover, there is a uniform bound (depending only on the dimension d) on the number of simplices incident to each vertex of these triangulations.*

Proof. It suffices to construct an infinite family of 2-complexes $\{\Delta_n\}_{n\in\mathbb{N}}$ such that each Δ_n is collapsible, the vertex degrees in the family are uniformly bounded by some k, and the 1-skeleta $\{\Delta_n^{\leq 1}\}_{n\in\mathbb{N}}$ form an expander family: applying Theorem 43 to such a family gives a corresponding family of triangulated d-spheres $\{\Sigma_n\}_{n\in\mathbb{N}}$ such that the dual graph of each Σ_n is quasi-isometric to $\Delta_n^{\leq 1}$, with uniformly bounded quasi-isometry constants. By Fact 13, this implies that the dual graphs to $\{\Sigma_n\}_{n\in\mathbb{N}}$ are an expander family. We construct a suitable family $\{\Delta_n\}_{n\in\mathbb{N}}$ in Theorem 54. □

Corollary 45 (Corollary 7)**.** *Let M be a triangulable d-manifold. Then there is an infinite family of combinatorially distinct triangulations of M in which the vertex degrees are uniformly bounded (with the bound depending on M) and for which the dual graphs form a family of $(d+1)$-regular expanders. If M is PL, there exists such a family in which the triangulations are combinatorial.*

Proof. Fix a triangulation X of M, taking X combinatorial if M is PL. Let $\{\Sigma_n\}_{n\in\mathbb{N}}$ be a family of triangulated d-spheres as in Theorem 6 (see the previous corollary). Fix a facet σ_0 of X, and for each n choose a facet σ_n of Σ_n, together with a simplicial isomorphism $f_n : \partial\sigma_0 \to \partial\sigma_n$. Delete the interiors of σ_0 and of σ_n and glue X to Σ_n along the boundaries of the two facets via f_n: this yields a simplicial complex Y_n which is PL homeomorphic to $X\#S^d \simeq X$. (Note that the gluing is indeed simplicial: $\partial\sigma_0$ is an induced subcomplex of $X \setminus \sigma_0$, and similarly for $\partial\sigma_n$ and $\Sigma_n \setminus \sigma_n$.)

We prove that the dual graphs of $\{Y_n\}_{n\in\mathbb{N}}$ are an expander family. By Fact 13, it suffices to show that they are quasi-isometric (with uniform constants) to the dual graphs of $\{\Sigma_n\}_{n\in\mathbb{N}}$. For each $n \in \mathbb{N}$ construct the following relation $R_n \subset Y_n^{(d)} \times \Sigma_n^{(d)}$:

$$R_n = \{(\rho, \sigma_0) \mid \rho \in X^{(d)} \setminus \{\sigma_0\}\} \cup \{(\rho, \rho) \mid \rho \in \Sigma_n^{(d)} \setminus \{\sigma_n\}\}.$$

That is, a pair of facets (ρ, σ) is in R_n if $\rho = \sigma$ is a facet of $\Sigma_n \setminus \sigma_n$ or if ρ is a facet of $X \setminus \sigma_0$ and $\sigma = \sigma_n$. Identifying the facets of X with the vertices of its dual graph $D(X)$ and similarly for Σ_n and Y_n, each such relation satisfies the conditions of Lemma 12 with $L = c = 1$ and $M = \mathrm{diam}(D(X)) + 1$, so we obtain the desired quasi-isometries. □

Lemma 46. *Let $\mathcal{C}$ be a class of bounded-degree simplicial 2-complexes which are nulhomologous over $\mathbb{F}$. Then the corresponding family of triangulated 3-dimensional homology manifolds output by Theorem 43 has dual graphs which are uniformly short over $\mathbb{F}$. The bound L on the length of a cycle in a homology basis for these graphs depends only on the degree bound for complexes in $\mathcal{C}$, and not on $\mathbb{F}$.*

Proof. A homology basis for the dual graph of a triangulated homology sphere Σ is given by the dual 2-cells (in the case of a 3-manifold, these are the cycles in the dual graph corresponding to the links of 1-simplices in Σ). If Σ is the output of running Theorem 43 on a 2-complex with all vertex degrees bounded by k then the star of any face $\sigma \in \Sigma$ is contained in the union of a uniformly bounded number m of handles (depending only on k) and each such handle contains a uniformly bounded number of faces N (again depending only on k). Hence the star of σ contains at most mN faces. Note that the construction of Σ does not depend on $\mathbb{F}$, so the bound $L = mN$ in the statement depends only on k. □

Lemma 47. *Let $\mathcal{C}$ be a class of 4-regular short graphs with parameters k, L. Then there is a class $\mathcal{D}$ of 3-regular short graphs with parameters k', L' depending only on k, L and a map $\mathcal{C} \to \mathcal{D}$ that is injective on isomorphism classes and increases the number of vertices by only a constant factor.*

Proof. Given a 4-regular graph, replace each vertex (with its four half-edges) with a cycle on 4-vertices with the four half-edges shown:

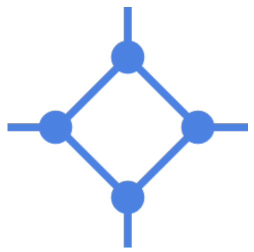

then connect each half-edge to a half-edge of a 4-cycle of a neighboring vertex. This construction is injective: contracting all 4-cycles in the resulting graph to vertices always yields the original graph (any cycle in the original graph is subdivided into a cycle of length at least 6). Further, the resulting graphs are short: each cycle of the input graph of given length m is subdivided into a cycle of length at most $3m$, and together with the

4-cycles these form a homology basis. It is straightforward that each edge of the new graph participates in a bounded number k' of cycles in this basis, depending on k. $\square$

Lemma 48. *There is an integer $L_0 \in \mathbb{N}$ such that the following holds. Let $\mathbb{F}$ be a field and let $\mathcal{D}$ be the class of uniformly short graphs of degree 3 over $\mathbb{F}$ in which there is a cycle basis consisting of cycles of length $\leq L_0$. For any class $\mathcal{C}$ of short graphs over $\mathbb{F}$ there is a constant c and a function $S : \mathcal{C} \to \mathcal{D}$, injective on isomorphism classes, such that every $G \in \mathcal{C}$ with n vertices is mapped to a graph $S(G)$ with at most cn vertices.*

The point is that while c depends on $\mathcal{C}$, L_0 doesn't.

Proof. Observe that if Σ is a triangulated 3-manifold then the dual graph of $b(b(\Sigma))$ uniquely determines the simplicial isomorphism type of Σ (alternatively one can avoid using this observation by an appropriate use of Proposition 42). We perform the following construction twice: given a short class $\mathcal{D}$ over $\mathbb{F}$, first use Proposition 21 to fill the cycles and obtain a class of 2-complexes. Then apply Theorem 43 to obtain triangulated 3-manifolds, and finally take the second barycentric subdivision of each of these and return the dual graphs. This construction is injective on $\mathcal{D}$ (because each step separately is injective) and it increases the number of vertices in the input graphs by only a constant factor, which depends on $\mathcal{D}$. The result is a short class $\mathcal{D}'$ of 4-regular graphs.

Apply this first to $\mathcal{D} = \mathcal{C}$, and then again to the resulting class $\mathcal{D}'$, to obtain some short class $\mathcal{D}''$. By Lemma 46, the output class $\mathcal{D}''$ is uniformly short, with cycle length bound L'' depending on the maximal vertex degrees in the simplicial 2-complexes that result from filling in the graphs of $\mathcal{D}'$. By the proof of Proposition 21, these vertex degrees depend only on the class $\mathcal{D}'$: specifically, if each $G \in \mathcal{D}'$ has a homology basis $\{\gamma_i\}_i$ such that each edge of G participates in at most k of the cycles, and the vertex degrees in each $G \in \mathcal{D}$ are bounded by d, then the vertex degrees of the 2-complexes are at most $d(1 + 4k)$. Since $\mathcal{D}'$ consists of dual graphs of triangulated 3-spheres, we can take $d = 4$ and $k = 3$ (each edge of the dual graph participates in precisely three dual 2-faces, see Proposition 19). Hence L'' is a universal constant.

Finally, the class $\mathcal{D}''$ is 4-regular. Applying Lemma 47 injects it into a uniformly short 3-regular class (with constant cycle bound the desired L_0) and finishes the proof. $\square$

Corollary 49 (Theorem 4)**.** *There exists $L_0 \in \mathbb{N}$ such that for all $L \geq L_0$ the following holds. Let $\mathbb{F}$ be a field and $d \geq 3$ an integer. Let $s_{d,\mathbb{F}}(N)$ be the number of isomorphism types of triangulated combinatorial d-manifolds M with $H_1(M;\mathbb{F}) = 0$ and with at most N d-simplices. Let $t_{L,\mathbb{F}}(N)$ be the number of connected 3-regular graphs G with at most N vertices such that cycles of length $\leq L$ generate $H_1(G;\mathbb{F})$. Then there exists $c, c' > 0$ such that*

$$s_{d,\mathbb{F}}(\frac{1}{c}N) \leq t_{L,\mathbb{F}}(N) \leq s_{d,\mathbb{F}}(cN)$$

for all $N \in \mathbb{N}$.

Proof. The inequality $t_{L,\mathbb{F}}(N) \leq s_{d,\mathbb{F}}(cN)$ follows from Proposition 21 and Theorem 43: the uniformly short class $\mathcal{C}$ consisting of 3-regular graphs with $\mathbb{F}$-homology bases having all cycles of length at most L can be injectively encoded by combinatorially triangulated $\mathbb{F}$-homology d-spheres in such a way that the encoding of a graph $G \in \mathcal{C}$ on n vertices has at most cn facets for some fixed $c > 0$.

For the inequality $s_{d,\mathbb{F}}(\frac{1}{c}N) \leq t_{L,\mathbb{F}}(N)$ we encode triangulated d-dimensional $\mathbb{F}$-homology spheres in 3-regular short graphs as follows. Given a triangulated homology sphere Σ, first pass to the dual graph, colored using Proposition 42 (so it can be used to reconstruct Σ). This is a $(d+1)$-regular colored graph (D, c), where c uses k colors, with k depending only on d. Encode it into a bounded-degree graph D' as follows: if $c(v) = t \in [k]$, add t isolated vertices to the graph and connect them to v. This construction yields bounded degree graphs (with degree at most $d+1+k$,) increases the number of vertices by a constant depending only on d, and is injective: D can be reconstructed from D' by erasing all degree-1 vertices, and the color $c(v)$ of a given vertex is the number of degree-1 vertices connected to v. Now by Proposition 19 and Lemma 48 we can encode these graphs injectively into 3-regular graphs with first homology generated by cycles of length at most L_0. $\square$

9 Sphere triangulations with expanding duals

We construct an infinite sequence of bounded-degree collapsible 2-complexes such that their 1-skeleta are an expander family. The idea is to construct the mapping telescopes of sequences of 2-covers obtained from [4].

9.1 Collapsibility and expansion of telescopes

Definition 50. Let $\{X_k\}_{k=1}^{n}$ be a sequence of cell complexes, and $\{f_k\}_{k=1}^{n-1}$ a sequence of maps $f_k : X_{k+1} \to X_k$. The *mapping telescope* of the sequence is

$$\frac{(X_1 \times \{1\}) \sqcup \bigsqcup_{k=2}^{n} X_k \times I}{\sim},$$

where the relation $\sim$ identifies $(x, 0)$ with $(f_k(x), 1)$ whenever $x \in X_{k+1}$.

In other words, the mapping telescope is given by gluing the mapping cylinders of the maps $\{f_k\}_k$ together, each on top of the previous one.

Definition 51. A *hierarchical sequence of graphs* $\mathcal{H}$ is a sequence of graphs $\{H_n\}_{n\in\mathbb{N}}$ together with functions $f_n : V(H_{n+1}) \to V(H_n)$ for each $n \in \mathbb{N}$, satisfying that for each pair $\{v, w\} \in E(H_{n+1})$, we have either $f_n(v) = f_n(w)$ or $\{f_n(v), f_n(w)\} \in E(H_n)$.

Given a hierarchical sequence $\mathcal{H} = \left(\{H_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}\right)$, we extend each f_n to a function $H_{n+1} \to H_n$ by mapping each edge $\{v, w\}$ of H_{n+1} linearly onto $\{f(v), f(w)\}$ (which may be a single vertex). The N-th mapping telescope $T_N(\mathcal{H})$ of the sequence is then the mapping telescope of $\{f_n\}_{n=1}^{N}$.

Observe that $T_N(\mathcal{H})$ is a 2-dimensional CW complex in which all 2-cells are squares or triangles.

Claim 52. Let $\mathcal{H} = \left(\{H_n\}_{n\in\mathbb{N}} \{f_n\}_{n\in\mathbb{N}}\right)$ be a hierarchical sequence of graphs. Then $T_N(\mathcal{H})$ collapses onto H_1 for all $N \in \mathbb{N}$. In particular, if H_1 is a single vertex, then each $T_N(\mathcal{H})$ is collapsible.

Proof. It suffices to show that $T_1(\mathcal{H})$ collapses onto H_1, and that $T_N(\mathcal{H})$ collapses onto $T_{N-1}(\mathcal{H})$ for each $N \geq 2$. If $N \geq 2$, note that

$$T_N(\mathcal{H}) \simeq \left[T_{N-1}(\mathcal{H}) \sqcup (H_{N+1} \times I)\right] / \sim,$$

where $\sim$ glues $H_{N+1} \times \{0\}$ onto $T_{N-1}(\mathcal{H})$. The 2-cells of $T_N(\mathcal{H})$ which are not contained in $T_{N-1}(\mathcal{H})$ have boundary $\{v, w, f_N(v), f_N(w)\}$, and in each of these $\{v, w\} \in E(H_{N+1})$

is a free face. Hence all such 2-cells can be collapsed. We are then left with the subcomplex

$$[T_{N-1}(\mathcal{H}) \sqcup (V(H_{N+1}) \times I)] / \sim,$$

and for each 1-cell of the form $v \times I$ (where $v \in V(H_{N+1})$) the vertex $v \times \{1\}$ is a free face. Hence these can be collapsed as well, leaving exactly $T_{N-1}(\mathcal{H})$.

In exactly the same way, $T_1(\mathcal{H}) = [(H_1 \times \{1\}) \sqcup (H_2 \times I)] / \sim$ collapses onto $H_1 \times \{1\}$ by first collapsing all 2-cells, and then all 1-cells of the form $v \times I$ for $v \in V(H_2)$ (from the free face $v \times \{1\}$). $\square$

Proposition 53. *Let $\mathcal{H} = \left(\{H_n\}_{n\in\mathbb{N}}, \{f_n\}_{n\in\mathbb{N}}\right)$ be a hierarchical sequence of graphs. Suppose $\{H_n\}_{n\in\mathbb{N}}$ forms a bounded degree expander family, and that there exists a constant c such that $2 \leq \left|f_n^{-1}(v) \cap V(H_{n+1})\right| \leq c$ for all $n \in \mathbb{N}$ and $v \in V(H_n)$. Then the sequence of graphs $\left\{T_N(\mathcal{H})^{\leq 1}\right\}_{N\in\mathbb{N}}$ is a bounded degree expander family.*

Proof. For convenience we denote $T_N \coloneqq T_N(\mathcal{H})^{\leq 1}$ for all N. Note that if the degrees of the graphs $\{H_n\}_{n\in\mathbb{N}}$ are bounded by some d then the degrees of the graphs $\{T_N\}_{N\in\mathbb{N}}$ are bounded by $d + c + 1$.

Suppose the graphs $\{H_n\}_{n\in\mathbb{N}}$ all have edge expansion ratio at least α. Let $N \in \mathbb{N}$ and let $W \subset V(T_N)$ have $|W| \leq \frac{1}{2}|V(T_N)|$.

Recall $V(T_N) = \bigsqcup_{n\leq N+1} V(H_n)$. Call $V(H_n)$ the n-th *layer* of $V(T_N)$ for each $n \leq N+1$, and denote the part of W in n-th layer by $L_n = W \cap V(H_n)$. Let

$$S = \bigcup \left\{ L_n \mid \frac{|L_n|}{|V(H_n)|} < \frac{9}{10} \right\},$$

$$D = \bigcup \left\{ L_n \mid \frac{|L_n|}{|V(H_n)|} \geq \frac{9}{10} \right\}.$$

Intuitively, S is the union of those layers L_n of W which are reasonably sparse within $V(H_n)$, while D is the union of those layers which are dense. It is clear that $\{S, D\}$ is a partition of W. We split into several cases, and show $|E(W, W^c)|$ is at least a constant multiple of $|W|$ in each case:

Case 1: $|S| \geq \frac{1}{2}|W|$. Let n be such that $L_n \subset S$, and observe that

$$\frac{1}{9}|L_n| \leq \min\{|L_n|, |V(H_n) \setminus L_n|\} \leq \frac{1}{2}|V(H_n)|.$$

Therefore

$$|E(L_n, V(H_n) \setminus L_n)| \geq \frac{1}{9}\alpha |L_n|.$$

Since $\bigcup_{n\leq N+1} E(L_n, V(H_n) \setminus L_n) \subset E(W, W^c)$, we have

$$E(W, W^c) \geq \sum_{\{n | L_n \subset S\}} \frac{1}{9}\alpha |L_n| = \frac{1}{9}\alpha \cdot |S| \geq \frac{\alpha}{18}|W|.$$

Case 2: $|D| \geq \frac{1}{2}|W|$. Let $k = \max\{n \leq N+1 \mid L_n \subset D\}$. We split this into two further subcases.

Case 2.1: $k < N+1$. Then $L_{k+1} \subset S$. We have

$$|E(W, W^c)| \overset{\star}{\geq} |E(L_k, V(H_{k+1}) \setminus L_{k+1})| + |E(L_{k+1}, V(H_{k+1}) \setminus L_{k+1})|.$$

Note that $|L_k| \geq \frac{1}{4}|W|$, because

$$|L_k| \geq \frac{9}{10}|V(H_k)| \geq \frac{9}{10} \cdot \frac{1}{2}\left|\bigcup_{j \leq k} V(H_j)\right|,$$

where $D \subseteq \bigcup_{j \leq k} V(H_j)$ contains at least half the vertices of $|W|$.

The idea is now that either L_{k+1} is relatively large compared to L_k, in which case an analysis similar to case 1 applies to the summand $|E(L_{k+1}, V(H_{k+1}) \setminus L_{k+1})|$, or that L_{k+1} is small compared to L_k, in which case $|E(L_k, V(H_{k+1}) \setminus L_{k+1})|$ is at least a constant times $|L_k|$.

In detail, if $|L_{k+1}| > |L_k|$ then (as in case 1) we have

$$|E(W, W^c)| \overset{\text{by } \star}{\geq} |E(L_{k+1}, V(H_{k+1}) \setminus L_{k+1})| \geq \frac{1}{9}\alpha |L_{k+1}| \geq \frac{1}{9}\alpha \cdot |L_k| \geq \frac{\alpha}{36}|W|$$

On the other hand, if $|L_{k+1}| \leq |L_k|$, then since each vertex of L_k is connected to at least two vertices of $V(H_{k+1}) \supseteq L_{k+1}$ (and each vertex of H_{k+1} is connected to at most one vertex of L_k,) we find

$$|E(W, W^c)| \overset{\text{by } \star}{\geq} |E(L_k, V(H_{k+1}) \setminus L_{k+1})| \geq |E(L_k, V(H_{k+1}))| - |E(L_k, L_{k+1})|$$

$$\geq 2|L_k| - |L_k| \geq |L_k| \geq \frac{1}{4}|W|.$$

Case 2.2: $k = N+1$. This case is very similar to case 2.1. First note that $L_N \not\subset D$, because otherwise

$$|W| \geq \frac{9}{10}|V(H_N) \cup V(H_{N+1})| \geq \frac{9}{10} \cdot \frac{3}{4}|V(T_N)| > \frac{1}{2}|V(T_N)|.$$

(This follows from the fact that $|V(H_{i+1})| \geq 2|V(H_i)|$ for all $i \leq N$.) Therefore

$$|E(W, W^c)| \geq |E(L_N, V(H_N) \setminus L_N)| + |E(L_{N+1}, V(H_N) \setminus L_N)|.$$

Observe that

$$|E(L_{N+1}, V(H_N) \setminus L_N)| = |E(L_{N+1}, V(H_N))| - |E(L_{N+1}, L_N)|$$

$$= |L_{N+1}| - \left|f^{-1}(L_N) \cap L_{N+1}\right| \geq |L_{N+1}| - c|L_N|.$$

Hence if $|L_N| \leq \frac{1}{2c}|L_{N+1}|$ we have $|E(W, W^c)| \geq \frac{1}{2}|L_{N+1}|$.

On the other hand, if $|L_N| > \frac{1}{2c}|L_{N+1}|$, we have $|E(L_N, V(H_N) \setminus L_N)| \geq \frac{1}{9}\alpha|L_N|$ by the considerations of case 1, and hence $|E(W, W^c)| \geq \frac{1}{9}\alpha|L_N| \geq \frac{\alpha}{18c}|L_{N+1}|$.

In either case, since

$$|L_{N+1}| \geq \frac{9}{10}|V(H_{N+1})| \geq \frac{9}{20}|V(T_N)| \geq \frac{9}{20}|W|,$$

we have a lower bound on $|E(W, W^c)|$ of the form constant $\cdot\, |W|$. □

9.2 An expanding family of collapsible 2-complexes

Theorem 54. *There exists an infinite family $\{X_n\}_{n\in\mathbb{N}}$ of finite 2-dimensional simplicial complexes satisfying:*

1. *Each of the complexes is collapsible to a point,*
2. *The 1-skeleta $\left\{X_n^{\leq 1}\right\}_n$ are a bounded-degree expander family.*

Proof. By the main result of [4], if $d \in \mathbb{N}$ is large enough then there exists an expander family $\{H_n\}_{n\in\mathbb{N}}$ of d-regular graphs such that each H_{n+1} is a 2-sheeted cover of H_n. Denote the covering maps by $f_n : H_{n+1} \to H_n$. By prefixing this sequence with a graph H_1 on a single vertex (incrementing all other indices for consistency, and letting $f_1 : H_2 \to H_1$ be the constant map) we obtain a hierarchical family $\mathcal{H}$ of bounded degree expanders. The sequence of mapping telescopes $\{T_N(\mathcal{H})\}_{N\in\mathbb{N}}$ is the required family of 2-complexes by Claim 52 and Proposition 53. □

References

[1] Oswin Aichholzer, Thomas Hackl, David Orden, Pedro Ramos, Günter Rote, André Schulz, and Bettina Speckmann. Flip graphs of bounded degree triangulations. *Graphs and Combinatorics*, 29:1577–1593, 2013. 9

[2] Mark Anthony Armstrong. Extending triangulations. *Proceedings of the American Mathematical Society*, 18(4):701–704, 1967. 21

[3] Goulnara Arzhantseva, Dawid Kielak, Tim de Laat, and Damian Sawicki. Spectral gap and origami expanders. *arXiv preprint arXiv:2112.11864*, 2021. 6

[4] Yonatan Bilu and Nathan Linial. Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica*, 26(5):495–519, 2006. 33, 36

[5] Cesar Ceballos and Joseph Doolittle. Subword complexes and Kalai's conjecture on reconstruction of spheres. *arXiv preprint arXiv:2206.15461*, 2022. 27

[6] Ross L Finney. The insufficiency of barycentric subdivision. *Michigan Mathematical Journal*, 12(3):263–272, 1965. 24, 26

[7] Misha Gromov. Spaces and questions. In *Visions in mathematics*, pages 118–161. Springer, 2010. 1, 2

[8] John FP Hudson. *Piecewise linear topology*. WA Benjamin, 1969. 3

[9] Gil Kalai. The diameter of graphs of convex polytopes and f-vector theory. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 387–412, 1991. 1, 2

[10] Gil Kalai. Polytope skeletons and paths. In *Handbook of Discrete and Computational Geometry*, pages 505–532. Chapman and Hall/CRC, 2017. 1, 2

[11] William Bernard Raymond Lickorish. Simplicial moves on complexes and manifolds. *Geometry and Topology Monographs*, 2(299-320):314, 1999. 3

[12] Lauri Loiskekoski and Günter M Ziegler. Simple polytopes without small separators, II: Thurston's bound. *Israel Journal of Mathematics*, 228(1):293–303, 2018. 2

[13] John Milnor. A procedure for killing homotopy groups of differentiable manifolds. In *Proceedings of Symposia in Pure Mathematics*, pages 39–55. American Mathematical Society, 1961. 23

[14] Eran Nevo, Francisco Santos, and Stedman Wilson. Many triangulated odd-dimensional spheres. *Mathematische Annalen*, 364(3):737–762, 2016. 1

[15] Colin P Rourke and Brian Joseph Sanderson. *Introduction to piecewise-linear topology.* Springer Science & Business Media, 2012. 3, 7, 18, 22