On locally constructible spheres and balls

Bruno Benedetti[*]
Inst. Mathematics, MA 6-2
TU Berlin
D-10623 Berlin, Germany
benedetti@math.tu-berlin.de

Günter M. Ziegler[**]
Inst. Mathematics, MA 6-2
TU Berlin
D-10623 Berlin, Germany
ziegler@math.tu-berlin.de

February 2, 2009

Abstract

The unsolved question whether there are only exponentially-many combinatorial types of simplicial 3-spheres is crucial for the convergence of models for 3D quantum gravity. Working towards this question, Durhuus and Jonsson (1995) introduced the restriction to "locally constructible" (LC) 3-spheres, and showed that there are only exponentially-many LC 3-spheres.

We characterize the LC property for d-spheres ("the sphere minus a facet collapses to a $(d-2)$-complex") and for d-balls. Thus we link it to the classical notions of collapsibility, shellability and constructibility, and obtain hierarchies of such properties for simplicial balls and spheres. The main corollaries from this study are:

– Not all simplicial 3-spheres are locally constructible.
(This solves a problem by Durhuus and Jonsson.)
– There are only exponentially many shellable simplicial 3-spheres with given number of facets. (This answers a question by Kalai.)
– All simplicial constructible 3-balls are collapsible.
(This answers a question by Hachimori.)

1 Introduction

Ambjørn, Boulatov, Durhuus, Jonsson, and others have worked to develop a three-dimensional analogue of the simplicial quantum gravity theory, as provided for two dimensions by Regge [40]. (See [3] and [41] for surveys.) The discretized version of quantum gravity considers simplicial complexes instead of smooth manifolds; the metric properties are artificially introduced by assigning length a to any edge. (This approach is due to Weingarten [46] and known as "theory of dynamical triangulations".) A crucial path integral over metrics, the "partition function for gravity", is then defined via a weighted sum over all triangulated manifolds of fixed topology. In three dimensions, the whole model is convergent only if the number of triangulated

[*]Supported by DFG via the Berlin Mathematical School
[**]Partially supported by DFG

3-spheres with N facets grows not faster than C^N, for some constant C. Is this true? How many simplicial spheres are there with N facets, for N large?

This crucial question, which was put into the spotlight also by Gromov [18, pp. 156-157], still represents an open problem. Its 2D-analogue, however, was answered long time ago by Tutte [44] [45], who proved that there are asymptotically fewer than $\left(\frac{16}{3\sqrt{3}}\right)^N$ combinatorial types of triangulated 2-spheres. (By Steinitz' theorem, cf. [48, Lect. 4], this quantity equivalently counts the maximal planar maps on $n \geq 4$ vertices, which have $N = 2n - 4$ faces, and also the combinatorial types of simplicial 3-dimensional polytopes with N facets.)

In the following, the adjective "simplicial" will often be omitted when dealing with balls, spheres, or manifolds, as all the regular cell complexes and polyhedral complexes that we consider are simplicial.

Why are 2-spheres "not so many"? Every combinatorial type of triangulation of the 2-sphere can be generated as follows (Figure 1): First for some even $N \geq 4$ build a tree of N triangles (which combinatorially is the same thing as a triangulation of an $(N+2)$-gon), and then glue edges according to a complete matching of the boundary edges. A necessary condition in order to obtain a 2-sphere is that such a matching is *planar*. Planar matchings and triangulations of $(N+2)$-gons are both enumerated by a Catalan number C_{N+2}, and since the Catalan numbers satisfy a polynomial bound $C_N = \frac{1}{N+1}\binom{2N}{N} < 4^N$, we get an exponential upper bound for the number of triangulations.

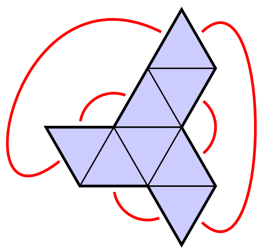

Figure 1: How to get an octahedron from a tree of 8 triangles (= a triangulated 10-gon).

Neither this simple argument nor Tutte's precise count can be easily extended to higher dimensions. Indeed, we have to deal with three different effects when trying to extend results or methods from dimension two to dimension three:

(i) Many combinatorial types of simplicial 3-spheres are not realizable as boundaries of convex 4-polytopes; thus, even though we observe below that there are only exponentially-many simplicial 4-polytopes with N facets, the 3-spheres could still be more numerous.

(ii) The counts of combinatorial types according to the number n of vertices and according to the number N of facets are not equivalent any more. We have $3n - 10 \leq N \leq \frac{1}{2}n(n-3)$ by the lower resp. upper bound theorem for simplicial 3-spheres. We know that there are more than $2^{n\sqrt[4]{n}}$ 3-spheres [29] [38], but less than $2^{20 n \log n}$ types of 4-polytopes with n vertices [16] [1], yet this does not answer the question for a count in terms of the number N of facets.

(iii) While it is still true that there are only exponentially-many "trees of N tetrahedra", the matchings that can be used to glue 3-spheres are not planar any more; thus, they could be more than exponentially-many. If, on the other hand, we restrict ourselves to "local gluings", we generate only a limited family of 3-spheres, as we will show below.

In the early nineties, new finiteness theorems by Cheeger [11] and Grove et al. [19] yielded a new approach, namely, to count d-manifolds of “fluctuating topology” (not necessarily spheres) but “bounded geometry” (curvature and diameter bounded from above, and volume bounded from below). This allowed Bartocci et al. [6] to bound for any d-manifold the number of triangulations with N or more facets, under the assumption that no vertex had degree higher than a fixed integer. However, for this it is crucial to restrict the topological type: Already for $d = 2$, there are more than exponentially many triangulated 2-manifolds of bounded vertex degree with N facets.

In 1995, the physicists Durhuus and Jonsson [13] introduced the class of “locally constructible” (LC) 3-spheres. An LC 3-sphere (with N facets) is a sphere obtainable from a tree of N tetrahedra, by identifying pairs of adjacent triangles in the boundary. “Adjacent” means here “sharing at least one edge”, and represents a dynamic requirement. Clearly, every 3-sphere is obtainable from a tree of N tetrahedra by matching the triangles in its boundary; according to the definition of LC, however, we are allowed to match only those triangles that *are* adjacent — or that have *become* adjacent by the time of the gluing.

Durhuus and Jonsson proved an exponential upper bound on the number of combinatorially distinct LC spheres with N facets. Based also on computer simulations ([4], see also [10] and [2]) they conjectured that all 3-spheres should be LC. A positive solution of this conjecture would have implied that spheres with N facets are at most C^N, for a constant C — which would have been the desired missing link to implement discrete quantum gravity in three dimensions.

In the present paper, we show that the conjecture of Durhuus and Jonsson has a negative answer: There are simplicial 3-spheres that are not LC. (With this, however, we do not resolve the question whether there are fewer than C^N simplicial 3-spheres on N facets, for some constant C.) Moreover, we establish the following theorem, which relates the “locally constructible” spheres defined by physicists to concepts that originally arose in topological combinatorics.

Main Theorem 1 (Theorem 2.1)**.** *A simplicial d-sphere, $d \geq 3$, is LC if and only if the sphere after removal of one facet can be collapsed down to a complex of dimension $d-2$. Furthermore, there are the following inclusion relations between families of simplicial d-spheres:*

$$\{\text{vertex decomposable}\} \subsetneq \{\text{shellable}\} \subseteq \{\text{constructible}\} \subsetneq \{\text{LC}\} \subsetneq \{\text{all } d\text{-spheres}\}.$$

The inclusions all hold with equality for $d = 2$: all 2-spheres are vertex-decomposable. We use the hierarchy in particular for $d = 3$, in conjunction with the following extension and sharpening of Durhuus and Jonsson’s theorem (who discussed only the case $d = 3$).

Main Theorem 2 (Theorem 4.4)**.** *For fixed $d \geq 2$, the number of combinatorially distinct LC d-spheres with N facets grows not faster than $2^{d^2 \cdot N}$.*

We will give a proof for this theorem in Section 4; the same type of upper bound, with the same type of proof, also holds for LC d-balls with N facets.

Already in 1988 Kalai [29] constructed for every $d \geq 4$ a family of more than exponentially many d-spheres on n vertices; Lee [33] later showed that all of Kalai’s spheres are shellable. Combining this with Theorem 4.4 and Theorem 2.1, we obtain the following asymptotic result:

Corollary. *For fixed $d \geq 4$, the number of shellable d-spheres grows more than exponentially with respect to the number n of vertices, but only exponentially with respect to the number N of facets.*

The hierarchy of Theorem 2.1 is not quite complete: It is still not known whether constructible, non-shellable 3-spheres exist (see [30] [14]). A shellable 3-sphere that is not vertex-decomposable was found by Lockeberg in his 1977 Ph.D. work (reported in [32, p. 742]; see also [23]). Again, the 2-dimensional case is much simpler and completely solved: All 2-spheres are vertex decomposable (see [39]).

In order to show that not all spheres are LC we study in detail simplicial spheres with a "knotted triangle"; these are obtained by adding a cone over the boundary of a ball with a knotted spanning edge (as in Furch's 1924 paper [15]; see also Bing [8]). Spheres with a knotted triangle cannot be boundaries of polytopes. Lickorish [34] had shown in 1991 that

a 3-*sphere with a knotted triangle is not shellable if the knot is at least* 3-*complicated.*

Here "at least 3-complicated" refers to the technical requirement that the fundamental group of the complement of the knot has no presentation with less than four generators. A concatenation of three or more trefoil knots satisfies this condition. In 2000, Hachimori and Ziegler [25] [21] demonstrated that Lickorish's technical requirement is not necessary for his result:

a 3-*sphere with* any *knotted triangle is not constructible.*

In the present work, we re-justify Lickorish's technical assumption, showing that this is exactly what we need if we want to reach a stronger conclusion, namely, a topological obstruction to local constructibility. Thus, the following result is established in order to prove that the last inclusion of the hierarchy in Theorem 2.1 is strict.

Main Theorem 3 (Theorem 2.14)**.** *A* 3-*sphere with a knotted triangle is not LC if the knot is at least* 3-*complicated.*

The knot complexity requirement is now necessary, as non-constructible spheres with a single trefoil knot can still be LC. (See Example 2.27.)

The combinatorial topology of 3-balls and that of 3-spheres are of course closely related — our study builds on the well-known connections and also adds new ones.

Main Theorem 4 (Theorem 3.1)**.** *A simplicial* 3-*ball is LC if and only if it collapses down to the boundary minus a facet. We have the following hierarchy:*

$$\{\text{vertex decomp.}\} \subsetneq \{\text{shellable}\} \subsetneq \{\text{constructible}\} \subseteq \{\text{LC}\} \subsetneq \{\text{collapsible}\} \subsetneq \{\text{all 3-balls}\}.$$

In particular, we settle a question of Hachimori (see e.g. [22, pp. 54, 66]) whether all constructible 3-balls are collapsible. The converse does not hold [35] [22, p. 54]. All the inclusions hold with equality for simplicial 2-balls. However, note that Main Theorem 4 does not easily generalize to $d \geq 3$, since our proofs for the inclusion $\{\text{LC}\} \subseteq \{\text{collapsible}\}$ (Corollary 3.12) and for its strictness (Corollary 3.20) are valid only for $d \leq 3$.

A result of Chillingworth can then be re-stated in our language as "if for any geometric simplicial complex Δ the support (union) $|\Delta|$ is a convex 3-dimensional polytope, then Δ is necessarily an LC 3-ball", see Theorem 3.21. From that we derive that any geometric subdivision of the 3-simplex is necessarily constructible, if all the vertices of the subdivision lie on the boundary of the simplex. The result is best possible, since Rudin's ball is a subdivided 3-simplex with no interior vertex, and Rudin's ball is constructible, but not shellable.

1.1 Definitions and Notations

1.1.1 Simplicial regular CW complexes

In the following, we present the notion of "local constructibility" (due to Durhuus and Jonsson). Although in the end we are interested in this notion as applied to finite simplicial complexes, the iterative definition of locally constructible complexes dictates that for intermediate steps we must allow for the greater generality of finite "simplicial regular CW complexes". A CW complex is *regular* if the attaching maps for the cells are injective on the boundary (see e.g. [9]). A regular CW-complex is *simplicial* if for every proper face F, the interval $[0,F]$ in the face poset of the complex is boolean. Every simplicial complex (and in particular, any triangulated manifold) is a simplicial regular CW-complex.

The k-dimensional cells of a regular CW complex C are called *k-faces*; the inclusion-maximal faces are called *facets*, and the inclusion-maximal proper subfaces of the facets are called *ridges*. The *dimension* of C is the largest dimension of a facet; *pure* complexes are complexes where all facets have the same dimension. All complexes that we consider in the following are finite, most of them are pure. A *d-complex* is a d-dimensional complex. Conventionally, the 0-faces are called *vertices*, and the 1-faces *edges*. (In the discrete quantum gravity literature, the ridges are sometimes called "hinges" or "bones", whereas the edges are sometimes referred to as "links".) If the union $|C|$ of all simplices of C is homeomorphic to a manifold M, then C is a *triangulation* of M; if C is a triangulation of a d-ball or of a d-sphere, we will call C simply a *d-ball* (resp. *d-sphere*).

1.1.2 Knots

All the knots we consider are *tame*, that is, realizable as 1-dimensional subcomplexes of some 3-sphere. A knot is *m-complicated* if the fundamental group of the complement of the knot in the 3-sphere has a presentation with $m+1$ generators, but no presentation with m generators. By "at least m-complicated" we mean "k-complicated for some $k \geq m$". There exist arbitrarily complicated knots: Goodrick [17] showed that the connected sum of m trefoil knots is at least m-complicated.

Another measure of how tangled a knot can be is the bridge index (see [42, pp. 114-117] or [31, p. 18] for the definition). If a knot has bridge index b, the fundamental group of the knot complement admits a presentation with b generators and $b-1$ relations [31, p. 82]. In other words, the bridge index of a t-complicated knot is at least $t+1$. As a matter of fact, the connected sum of t trefoil knots is t-complicated, and its bridge index is exactly $t+1$ [14].

1.1.3 The combinatorial topology hierarchy

In the following, we review the key properties from the inclusion

$$\{\text{shellable}\} \subsetneq \{\text{constructible}\}$$

valid for all simplicial complexes, and the inclusion

$$\{\text{shellable}\} \subsetneq \{\text{collapsible}\}$$

applicable only for *contractible* simplicial complexes, both known from combinatorial topology (see [9, Sect. 11] for details).

Shellability can be defined for pure simplicial complexes as follows:

- every simplex is shellable;
- a d-dimensional pure simplicial complex C which is not a simplex is shellable if and only if it can be written as $C = C_1 \cup C_2$, where C_1 is a shellable d-complex, C_2 is a d-simplex, and $C_1 \cap C_2$ is a shellable $(d-1)$-complex.

Constructibility is a weakening of shellability, defined by:

- every simplex is constructible;
- a d-dimensional pure simplicial complex C which is not a simplex is constructible if and only if it can be written as $C = C_1 \cup C_2$, where C_1 and C_2 are constructible d-complexes, and $C_1 \cap C_2$ is a constructible $(d-1)$-complex.

Let C be a d-dimensional simplicial complex. An *elementary collapse* is the simultaneous removal from C of a pair of faces (σ, Σ) with the following prerogatives:

- $\dim \Sigma = \dim \sigma + 1$;
- σ is a proper face of Σ;
- σ is not a proper face of any other face of C.

(The three conditions above are usually abbreviated in the expression "σ is a free face of Σ"; some complexes have no free face). If $C' := C - \Sigma - \sigma$, we say that the complex C *collapses onto* the complex C'. We also say that the complex C *collapses onto* the complex D, and write $C \searrow D$, if C can be reduced to D by a finite sequence of elementary collapses. Thus a *collapse* refers to a sequence of elementary collapses. A *collapsible* complex is a complex that can be collapsed onto a single vertex.

Since $C' := C - \Sigma - \sigma$ is a deformation retract of C, each collapse preserves the homotopy type. In particular, all collapsible complexes are contractible. The converse does not hold in general: For example, the so-called "dunce hat" is a contractible 2-complex without free edges, and thus with no elementary collapse to start with. However, the implication "contractible $\Rightarrow$ collapsible" holds for all 1-complexes, and also for shellable complexes of any dimension.

A connected 2-dimensional complex is collapsible if and only if it does *not* contain a 2-dimensional complex without a free edge. In particular, for 2-dimensional complexes, if $C \searrow D$ and D is not collapsible, then C is also not collapsible. This holds no more for complexes C of dimension larger than two [27].

1.1.4 LC pseudomanifolds

By a *d-pseudomanifold* [possibly with boundary] we mean a finite regular CW-complex P that is pure d-dimensional, simplicial, and such that each $(d-1)$-dimensional cell belongs to at most two d-cells. The *boundary* of the pseudomanifold P, denoted ∂P, is the smallest subcomplex of P containing all the $(d-1)$-cells of P that belong to exactly one d-cell of P.

According to our definition, a pseudomanifold needs not be a simplicial complex; it might be disconnected; and its boundary might not be a pseudomanifold.

Definition 1.1 (Locally constructible pseudomanifold)**.** Let C be a pure d-dimensional simplicial complex with N facets. A *local construction* for C is a sequence $T_1, T_2, \ldots, T_N, \ldots, T_k$ $(k \geq N)$ such that T_i is a d-pseudomanifold for each i and

(1) T_1 is a d-simplex;

(2) if $i \leq N-1$, then T_{i+1} is obtained from T_i by gluing a new d-simplex to T_i alongside one of the $(d-1)$-cells in ∂T_i;

(3) if $i \geq N$, then T_{i+1} is obtained from T_i by identifying a pair σ, τ of $(d-1)$-cells in the boundary ∂T_i that share a $(d-2)$-cell F;

(4) $T_k = C$.

We say that C is *locally constructible*, or *LC*, if a local construction for C exists. With a little abuse of notation, we will call each T_i an *LC pseudomanifold*. We also say that C is locally constructed *along* T, if T is the dual graph of T_N, and thus a spanning tree of the dual graph of C.

The identifications described in item (3) above are operations that are not closed with respect to the class of simplicial complexes. Local constructions where all steps are simplicial complexes produce only a very limited class of manifolds (see Corollary 3.17).

However, since by definition the local construction in the end must arrive at a pseudomanifold C that *is* a simplicial complex, each intermediate step T_i must satisfy severe restrictions: for each $t \leq d$,

- distinct t-simplices that are not in the boundary of T_i share at most one $(t-1)$-simplex;
- distinct t-simplices in the boundary of T_i that share more than one $(t-1)$-simplex will need to be identified by the time the construction of C is completed.

Moreover,

- if σ, τ are the two $(d-1)$-cells glued together in the step from T_i to T_{i+1}, σ and τ cannot belong to the same d-simplex of T_i; nor can they belong to two d-simplices that are already adjacent in T_i.

For example, in each step of the local construction of a 3-sphere, no two tetrahedra share more than one triangle. Moreover, any two distinct interior triangles either are disjoint, or they share a vertex, or they share an edge; but they cannot share two edges, nor three; and they also cannot share one edge and the opposite vertex. If we glued together two boundary triangles that belong to adjacent tetrahedra, no matter what we did afterwards, we would not end up with a simplicial complex any more. So,

> *a locally constructible* 3*-sphere is a triangulated* 3*-sphere obtained from a tree of tetrahedra* T_N *by repeatedly identifying two adjacent triangles in the boundary.*

As we mentioned, the boundary of a pseudomanifold need not be a pseudomanifold. However, if P is an LC d-pseudomanifold, then ∂P is automatically a $(d-1)$-pseudomanifold. Nevertheless, ∂P may be disconnected, and thus, in general, it is not LC. By a result of Durhuus and Jonsson [13, Theorem 2], the boundary of an LC 3-pseudomanifold is a finite disjoint union of "cacti" of 2-spheres; any two 2-spheres in such a cactus share at most one point. We will call the points shared by two or more spheres in the boundary of an LC 3-pseudomanifold *pinch points* (or *PPs*). The PPs are characterized by the property that their link inside ∂T_i is not a 1-sphere. (It is a disjoint union of 1-spheres).

Definition 1.2. [Steps of types (i)-(ix) in LC constructions] Any admissible step in a local construction of a 3-pseudomanifold falls into one of the following nine types:

(i) tree-wise gluing;
(ii) identifying two triangles that share exactly 1 edge;
(iii) identifying two triangles that share 1 edge and the opposite vertex;
(iv) identifying two triangles that share 2 edges that meet in a PP;
(v) identifying two triangles that share 2 edges that do not meet in a PP;

(vi) identifying two triangles that share 3 edges, all of whose vertices are PPs;
(vii) identifying two triangles that share 3 edges, two of whose vertices are PPs;
(viii) identifying two triangles that share 3 edges, one of whose vertices is a PP;
(ix) identifying two triangles that share 3 edges, none of whose vertices is a PP.

By definition, the first $N-1$ steps of an LC construction of a 3-pseudomanifold are the ones of type (i). The last step in the construction of an LC 3-sphere is of type (ix).

The following table summarizes the distinguished effects of the steps:

step type	*no. of interior vertices*	*no. of connected components of the boundary*
(i)	+ 0	+ 0
(ii)	+ 0	+ 0
(iii)	+ 0	+ 0 (*)
(iv)	+ 0	+ 1
(v)	+ 1	+ 0
(vi)	+ 0	+ 3
(vii)	+ 1	+ 2
(viii)	+ 2	+ 0
(ix)	+ 3	– 1

where the asterisk recalls that a type (iii) step *almost* disconnects the boundary, pinching it in a point.

2 On LC Spheres

In this section, we establish the following hierarchy announced in the introduction.

Theorem 2.1. *For all $d \geq 3$, we have the following inclusion relations between families of simplicial d-spheres:*

$$\{\text{vertex decomposable}\} \subsetneq \{\text{shellable}\} \subseteq \{\text{constructible}\} \subsetneq \{\text{LC}\} \subsetneq \{\text{all } d\text{-spheres}\}.$$

Proof. The first two inclusions are known; the third inclusion follows from Lemma 2.24 and is shown to be strict via Example 2.27 together with Lemma 2.25; finally, Corollary 2.23 establishes the strictness of the fourth inclusion for all $d \geq 3$. □

2.1 Some d-spheres are not LC

Let S be any simplicial d-sphere ($d \geq 2$), and T any spanning tree of the dual graph of S. We denote by K^T the subcomplex of S formed by all the $(d-1)$-faces of S that are not intersected by T.

Lemma 2.2. *Let S be any d-sphere. Then for every spanning tree T of the dual graph of S,*

- *K^T is a contractible pure $(d-1)$-dimensional simplicial complex with $\frac{dN-N+2}{2}$ facets;*
- *for any facet Δ of S, $S - \Delta \searrow K^T$.*

Any collapse of a d-sphere S minus a facet Δ to a complex of dimension at most $d-1$ proceeds along a dual spanning tree T. To see this, fix a collapsing sequence. We may assume that the collapse of $S - \Delta$ is ordered so that the pairs $((d-1)\text{-face}, d\text{-face})$ are removed first. Whenever both the following conditions are met:

1. σ is the $(d-1)$-dimensional intersection of the facets Σ and Σ' of S;
2. the pair (σ, Σ) is removed in the collapsing sequence of $S - \Delta$,

draw an oriented arrow from the center of Σ to the center of Σ'. This yields a directed spanning tree T of the dual graph of S, where Δ is the root. Indeed, T is *spanning* because all d-simplices of $S - \Delta$ are removed in the collapse; it is *acyclic*, because the center of each d-simplex of $S - \Delta$ is reached by exactly one arrow; it is *connected*, because the only free $(d-1)$-faces of $S - \Delta$, where the collapse can start at, are the proper $(d-1)$-faces of the "missing simplex" Δ. We will say that the collapsing sequence *acts along the tree T* (in its top-dimensional part). Thus the complex K^T appears as intermediate step of the collapse: It is the complex obtained after the N-th pair of faces has been removed from $S - \Delta$.

Definition 2.3. By a *facet-killing sequence* for a d-dimensional simplicial complex C we mean a sequence $C_0, C_1, \ldots, C_{t-1}, C_t$ of complexes such that $t = f_d(C)$, $C_0 = C$, and C_{i+1} is obtained by an elementary collapse that removes a free $(d-1)$-face σ of C_i, together with the unique facet Σ containing σ.

If C is a d-complex, and D is a lower-dimensional complex such that $C \searrow D$, there exists a facet-killing sequence C_0, …, C_t for C such that $C_t \searrow D$. In other words, the collapse of C onto D can be rearranged so that the pairs $((d-1)\text{-face}, d\text{-face})$ are removed first. In particular, for any d-complex C, the following are equivalent:

1. there exists a facet-killing sequence for C;
2. there exists a k-complex D with $k \leq d-1$ such that $C \searrow D$.

What we argued before can be rephrased as follows:

Proposition 2.4. *Let S be a d-sphere, and Δ a d-simplex of S. Let C be a k-dimensional simplicial complex, with $k \leq d-2$. Then,*

$$S - \Delta \searrow C \iff \exists\, T \text{ s.t. } K^T \searrow C.$$

The right-hand side in the equivalence of Proposition 2.4 does not depend on the Δ chosen. So, for any d-sphere Δ, either $S - \Delta$ is collapsible for every Δ, or $S - \Delta$ is not collapsible for any Δ.

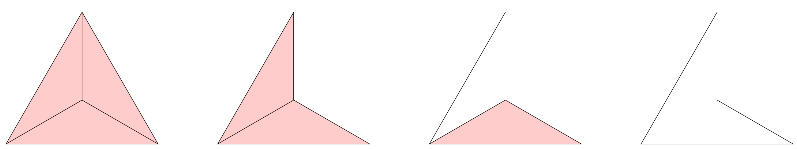

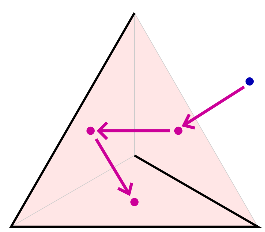

Figure 2: (Above): A facet-killing sequence of $S - \Delta$, where S is the boundary of a tetrahedron ($d = 2$), and Δ one of its facets. (Right): The 1-complex K^T [in black] onto which $S - \Delta$ collapses, and the directed spanning tree T [in purple] along which the collapse above acts.

One more convention: by a *natural labeling* of a rooted tree T on n vertices we mean a bijection $b : V(T) \longrightarrow \{1, \ldots, n\}$ such that if v is the root, $b(v) = 1$, and if v is not the root, there exists a vertex w adjacent to v such that $b(w) < b(v)$.

We are now ready to link the LC concept with collapsibility. Take a d-sphere S, a facet Δ of S, and a rooted spanning tree T of the dual graph of S, with root Δ. Since S is given, fixing T

is really the same as fixing the manifold T_N in the local construction of S; and at the same time, fixing T is the same as fixing K^T.

Once T, T_N, and K^T have been fixed, to describe the first part of a local construction of S (that is, $T_1, \ldots, T_N$) we just need to specify the order in which the tetrahedra of S have to be added, which is the same as to give a natural labeling of T. Besides, natural labelings of T are in bijection with collapses $S - \Delta \searrow K^T$ (the i-th facet to be collapsed is the node of T labelled i; see Proposition 2.4).

What if we do not fix T? Suppose S and Δ are fixed. Then the previous reasoning yields a bijection among the following sets:

1. the set of all facet-killing sequences of $S - \Delta$;
2. the set of "natural labelings" of spanning trees of S, rooted at Δ;
3. the set of the first parts $(T_1, \ldots, T_N)$ of local constructions for S, with $T_1 = \Delta$.

Can we understand also the second part of a local construction "combinatorially"? Let us start with a variant of the "facet-killing sequence" notion.

Definition 2.5. A *pure facet-massacre* of a pure d-dimensional simplicial complex P is a sequence $P_0, P_1, \ldots, P_{t-1}, P_t$ of (pure) complexes such that $t = f_d(P)$, $P_0 = P$, and P_{i+1} is obtained by P_i removing:

(a) a free $(d-1)$-face σ of P_i, together with the unique facet Σ containing σ, and

(b) all inclusion-maximal faces of dimension smaller than d that are left after the removal of type (a) or, recursively, after removals of type (b).

In other words, the (b) step removes lower-dimensional facets until one obtains a pure complex. Since $t = f_d(P)$, P_t has no facets of dimension d left, nor inclusion-maximal faces of smaller dimension; hence P_t is empty. The other P_i's are pure complexes of dimension d. Notice that the step $P_i \longrightarrow P_{i+1}$ is not a collapse, and does not preserve the homotopy type in general. Of course $P_i \longrightarrow P_{i+1}$ can be "factorized" in an elementary collapse followed by a removal of a finite number of k-faces, with $k < d$. However, this factorization is not unique, as the next example shows.

Example 2.6. Let P be the full triangle $\{1, 2, 3\}$. P admits admits three different facet-killing collapses (each edge can be chosen as free face), but it admits only one pure facet-massacre, namely $P, \emptyset$.

Lemma 2.7. *Let P be a pure d-dimensional simplicial complex. Every facet-killing sequence of P naturally induces a unique pure facet-massacre of P. All pure facet-massacres of P are induced by some (possibly more than one) facet-killing sequence.*

Proof. The map consists in taking a facet-killing sequence C_0, ..., C_t, and "cleaning up" the C_i by recursively killing the lower-dimensional inclusion-maximal faces. As the previous example shows, this map is not injective. It is surjective essentially because the removed lower-dimensional faces are of dimension "too small to be relevant". In fact, their dimension is at most $d-1$, hence their presence can interfere only with the freeness of faces of dimension at most $d-2$; so the list of all removals of the form $((d-1)\text{-face}, d\text{-face})$ in a facet-massacre yields a facet-killing sequence. □

Theorem 2.8. *Let S be a d-sphere; fix a spanning tree T of the dual graph of S. The second part of a local construction for S along T corresponds bijectively to a facet-massacre of K^T.*

Proof. Fix S and T; T_N and K^T are determined by this. Let us start with a local construction $(T_1, \ldots, T_{N-1},)\, T_N, \ldots, T_k$ for S along T. Topologically, $S = T_N/\sim$, where $\sim$ is the equivalence relation determined by the gluing (two distinct points of T_N are equivalent if and only if they will be identified in the gluing). Moreover, $K^T = \partial T_N/\sim$, by the definition of K_T.

Define $P_0 := K_T = \partial T_N/\sim$, and $P_j := \partial T_{N+j}/\sim$. We leave it to the reader to verify that $k - N$ and $f_d(K^T)$ are the same integer (see Lemma 2.2), which we called D; in particular $P_D = \partial T_k/\sim = \partial S/\sim = \emptyset$.

In the first LC step, $T_N \to T_{N+1}$, we remove from the boundary a free ridge r, together with the unique pair σ', σ'' of facets of ∂T_N sharing r. At the same time, r and the newly formed face σ are sunk into the interior. This step $\partial T_N \longrightarrow \partial T_{N+1}$ naturally induces an analogous step $\partial T_{N+j}/\sim \longrightarrow \partial T_{N+j+1}/\sim$, namely, the removal of r and of the (unique!) $(d-1)$-face σ containing it.

In the j-th LC step, $\partial T_{N+j} \longrightarrow \partial T_{N+j+1}$, we remove from the boundary a ridge r together with a pair σ', σ'' of facets sharing r; moreover, we sink into the interior a lower-dimensional face F if and only if we have just sunk into the interior all faces containing F. The induced step from $\partial T_{N+j}/\sim$ to $\partial T_{N+j+1}/\sim$ is precisely a "facet-massacre" step.

For the converse, we start with a "facet-massacre" $P_0, \ldots, P_D$ of K^T, and we have $P_0 = K_T = \partial T_N/\sim$. The unique $(d-1)$-face σ_j killed in passing from P_j to P_{j+1} corresponds to a unique pair of (adjacent!) $(d-1)$-faces σ_j', σ_j'' in ∂T_{N+j}. Gluing them together is the LC move that transforms T_{N+j} into T_{N+j+1}. □

Example 2.9. In a type (vii) step of a local construction for a 3-sphere S we remove from the boundary

– an edge, and the two adjacent triangles that share it;
– the other two edges shared by the previous triangles;
– only one of the three vertices: the one that is not a pinch point, or equivalently, the one that belonged to no other triangle in ∂T_{N+j}.

We do not remove the pinch point from the boundary, as they belong to other boundary triangles.

Remark 2.10. Summing up:

– The first part of a local construction along a tree T corresponds to a facet-killing collapse of $S - \Delta$ (that ends in K^T).
– The second part of a local construction along a tree T corresponds to a pure facet-massacre of K^T.
– A single facet-massacre of K^T corresponds to many facet-killing sequences of K^T.
– By Proposition 2.4, there exists a facet-killing sequence of K^T if and only if K^T collapses onto some $(d-2)$-dimensional complex C. This C is necessarily contractible, like K^T.

So S is locally constructible along T if and only if K^T collapses onto some $(d-2)$-dimensional contractible complex C, if and only id K^T has a facet-killing sequence. What if we do not fix T?

Theorem 2.11. *Let S be a d-sphere ($d \geq 3$). Then the following are equivalent:*

1. *S is LC;*
2. *for some spanning tree T of S, K^T is collapsible onto some $(d-2)$-dimensional (contractible) complex C;*
3. *there exists a $(d-2)$-dimensional (contractible) complex C such that for every facet Δ of S, $S - \Delta \searrow C$;*
4. *for some facet Δ of S, $S - \Delta$ is collapsible onto a $(d-2)$-dimensional contractible complex C.*

Proof. S is LC if and only if it is LC along some tree T; thus $(1) \Leftrightarrow (2)$ follows from Remark 2.10. Besides, $(2) \Rightarrow (3)$ follows from the fact that $S - \Delta \searrow K^T$ (Lemma 2.2), where K^T is independent of the choice of Δ. $(3) \Rightarrow (4)$ is trivial. To show $(4) \Rightarrow (2)$, take a collapse of $S - \Delta$ onto some $(d-2)$-complex C; by Lemma 2.4, there exists some tree T (along which the collapse acts) so that $S - \Delta \searrow K^T$ and $K^T \searrow C$. □

Corollary 2.12. *Let S be a* 3*-sphere. Then the following are equivalent:*
1. *S is LC;*
2. *K^T is collapsible, for some spanning tree T of the dual graph of S;*
3. *$S - \Delta$ is collapsible for every facet Δ of S;*
4. *$S - \Delta$ is collapsible for some facet Δ of S.*

Proof. This follows from the previous theorem, together with the fact that all contractible 1-complexes are collapsible. □

We are now in the position to exploit results by Lickorish about collapsibility.

Theorem 2.13 (Lickorish [34])**.** *Let $\mathfrak{L}$ be a knot on m edges in the* 1*-skeleton of a simplicial* 3*-sphere S. Suppose that $S - \Delta$ is collapsible, where Δ is some tetrahedron in $S - L$. Then $|S| - |\mathfrak{L}|$ is homotopy equivalent to a connected cell complex with one 0-cell and at most m* 1*-cells. In particular, the fundamental group of $|S| - |\mathfrak{L}|$ admits a presentation with m generators.*

Now assume that a certain sphere S containing a knot $\mathfrak{L}$ is LC. By Corollary 2.12, $S - \Delta$ is collapsible, for any tetrahedron Δ not in the knot $\mathfrak{L}$. Hence by Lickorish's criterion the fundamental group $\pi_1(|S| - |\mathfrak{L}|)$ admits a presentation with m generators.

Theorem 2.14. *Any* 3*-sphere with a* 3*-complicated triangular* 3*-edge knot is not LC. More generally, a* 3*-sphere with an m-gonal knot cannot be LC if the knot is at least m-complicated.*

Example 2.15. As for the "Furch–Bing ball" [15, p. 73] [8, p. 110] [49], drill a hole into a finely triangulated 3-ball along a triple pike dive of three consecutive trefoils; stop drilling one step before destroying the property of having a ball. (See Figure 3). Add a cone over the boundary. The resulting sphere has a three edge knot which is a connected sum of three trefoil knots. By Goodrick[17] the connected sum of m copies of the trefoil knot is at least m-complicated. So, this sphere has a knotted triangle, the fundamental group of whose complement has no presentation with 3 generators. Hence S cannot be LC.

From this we get a negative answer to the Durhuus–Jonsson conjecture:

Corollary 2.16. *Not all simplicial 3-spheres are LC.*

Lickorish proved also a higher-dimensional statement, basically by taking successive suspensions of the 3-sphere in Example 2.15.

Theorem 2.17 (Lickorish [34])**.** *For any fixed $d \geq 3$, there exists a PL d-sphere S such that $S - \Delta$ is not collapsible for any facet Δ of S.*

To exploit our Theorem 2.11 we need a sphere S such that $S - \Delta$ is not even collapsible to a $(d-2)$-complex. To establish that such a sphere exists, we strengthen Lickorish's result.

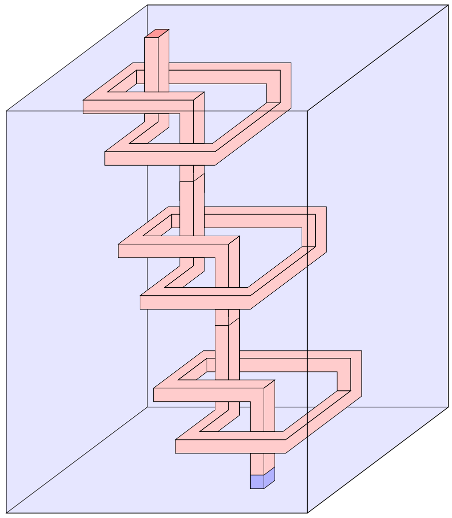

Figure 3: Furch–Bing ball with a (corked) tubular hole along a triple-trefoil knot. The cone over the boundary of this ball is a sphere that is *not* LC.

Definition 2.18. Let K be a d-manifold, A an r-simplex in K, and $\hat{A}$ the barycenter of A. Consider the barycentric subdivision $sd(K)$ of K. The *dual* A^* of A is the subcomplex of $sd(K)$ given by all flags

$$A \subset A_0 \subset A_1 \subset \cdots \subset A_r$$

where $r = \dim A$, and $\dim A_{i+1} = \dim A_i + 1$ for each i.

A^* is a cone with apex $\hat{A}$, and thus collapsible. We also have the following known result (see e.g. Munkres [37] or Hudson [28, pp. 29-30]).

Lemma 2.19. *Let K be a PL d-manifold (without boundary), and let A be a simplex in K of dimension r. Then*

- *A^* is a $(d-r)$-ball, and*
- *if A is a face of a $(r+1)$-simplex B, then B^* is a $(d-r-1)$-subcomplex of ∂A^*.*

We already know from Lemma 2.2 that for any d-sphere S, $S-\Delta$ is collapsible onto a $(d-1)$-complex: In other words, via collapses one can always get *one* dimension down. To get *two* dimensions down is not so easy: Our Theorem 2.11 states that $S-\Delta$ is collapsible onto a $(d-2)$-complex precisely when S is LC.

This "number of dimensions down you can get by collapsing" can be related to the minimal presentations of certain homotopy groups. The idea of the next theorem is that if one can get k dimensions down by collapsing, the $(k-1)$-th homotopy group of the complement of any $(d-k)$-subcomplex of the sphere cannot be too complicated to present.

Theorem 2.20. *Let t, d with $0 \le t \le d-2$, and let S be a PL d-sphere. Suppose that $S-\Delta$ collapses onto a t-complex, for some facet Δ of S. Then, for each t-dimensional subcomplex $\mathfrak{L}$ of S, the homotopy group*

$$\Pi_{d-t-1}(|S|-|\mathfrak{L}|)$$

has a presentation with exactly $f_t(\mathfrak{L})$ generators.

Proof. As usual, we assume that the collapse of $S-\Delta$ is ordered so that:
– first all pairs (d-face, $(d-1)$-face) are collapsed;
– then all pairs ($(d-1)$-face , $(d-2)$-face) are collapsed;
– $\vdots$
– finally, all pairs ($(t+1)$-face, t-face) are collapsed.
Let us put together all the faces that appear above, maintaining their order, to form a unique list of simplices

$$A_1, A_2, \ldots, A_{2M-1}, A_{2M}.$$

In such a list A_1 is a free face of A_2; A_3 is a free face of A_4 with respect to the complex $S-A_1-A_2$; and so on. In general, A_{2i-1} is a face of A_{2i} for each i, and in addition, if $j>2i$, A_{2i-1} is not a face of A_j.

We set $X_0 = A_0 := \hat{\Delta}$ and define a finite sequence $X_1, \ldots, X_M$ of subcomplexes of $sd(S)$ as follows:

$$X_j := \bigcup \{A_i^* \text{ s.t. } i \in \{0,\ldots,2j\} \text{ and } A_i \notin \mathfrak{L}\}, \qquad \text{for } j \in \{1,\ldots,M\}.$$

None of the A_{2i}'s can be in $\mathfrak{L}$, because $\mathfrak{L}$ is t-dimensional and $\dim A_{2i} \geq \dim A_{2M} = t+1$. However, exactly $f_t(\mathfrak{L})$ of the A_{2i-1}'s are in $\mathfrak{L}$. Consider how X_j differs from X_{j-1}. There are two cases:

- If A_{2j-1} is not in $\mathfrak{L}$,
$$X_j = X_{j-1} \cup A^*_{2j-1} \cup A^*_{2j}.$$
By Lemma 2.19, setting $r = \dim A_{2j-1}$, A^*_{2j-1} is a $(d-r)$-ball that contains in its boundary the $(d-r-1)$-ball A^*_{2j}. Thus $|X_j|$ is just $|X_{j-1}|$ with a $(d-r)$-cell attached via a cell in its boundary, and such an attachment does not change the homotopy type.
- If A_{2j-1} is in $\mathfrak{L}$, then
$$X_j = X_{j-1} \cup A^*_{2j}.$$
As this occurs only when $\dim A_{2j-1} = t$, we have that $\dim A_{2j} = t+1$ and $\dim A_{2j} = d-t-1$; hence $|X_j|$ is just $|X_{j-1}|$ with a $(d-t-1)$-cell attached via its whole boundary.

Only in the second case the homotopy type of $|X_j|$ changes at all, and this second case occurs exactly $f_t(\mathfrak{L})$ times. Since X_0 is one point, it follows that X_M is homotopy equivalent to a bouquet of $f_t(\mathfrak{L})$ many $(d-t-1)$-spheres.

Now let us list by (weakly) decreasing dimension the faces of S that do not appear in the previous list $A_1, A_2, \ldots, A_{2M-1}, A_{2M}$. We name the elements of this list

$$A_{2M+1}, A_{2M+2}, A_F$$

(where $\sum_{i=1}^{d} f_i(S) = F+1$ because all faces appear in $A_0, \ldots, A_F$).

Correspondingly, we recursively define a new sequence of subcomplexes of $sd(S)$ setting $Y_0 := X_M$ and

$$Y_h := \begin{cases} Y_{h-1} & \text{if } A_{2M+h} \in \mathfrak{L}, \\ Y_{h-1} \cup A^*_{2M+h} & \text{otherwise.} \end{cases}$$

Since $\dim A_{2M+h} \leq \dim A_{2M+1} = t$, we have that $|X_h|$ is just $|X_{h-1}|$ with possibly a cell of dimension at least $d-t$ attached via its whole boundary. Let us consider the homotopy groups of the Y_h 's : Recall that Y_0 was homotopy equivalent to a bouquet of $f_t(\mathfrak{L})$ $(d-t-1)$-spheres. Clearly, for all h,

$$\Pi_j(Y_h) = 0 \text{ for each } j \in \{1,\ldots,d-t-1\}.$$

Moreover, the higher-dimensional cell attached to $|Y_{h-1}|$ to get $|Y_h|$ corresponds to the addition of relators to a presentation of $\Pi_{d-t-1}(Y_{h-1})$ to get a presentation of $\Pi_{d-t-1}(Y_h)$. This means that for all h the group $\Pi_{d-t-1}(Y_h)$ is generated by (at most) $f_t(\mathfrak{L})$ elements.

The conclusion follows from the fact that, by construction, Y_{2M-F} is the subcomplex of $sd(S)$ consisting of all simplices of $sd(S)$ that have no face in $\mathfrak{L}$; and one can easily prove (see [34, Lemma 1]) that such a complex is a deformation retract of $|S|-|\mathfrak{L}|$. □

Corollary 2.21. *Let S be a d-sphere with a $(d-2)$-dimensional subcomplex $\mathfrak{L}$. If the fundamental group of $|S|-|\mathfrak{L}|$ has no presentation with $f_{d-2}(\mathfrak{L})$ generators, then S is not LC.*

Proof. Set $t = d-2$ in Theorem 2.20, and apply Theorem 2.11. □

Corollary 2.22. *Fix an integer $d \geq 3$. Let S be a 3-sphere with an m-gonal knot in its 1-skeleton, so that the knot is at least $(m \cdot 2^{d-3})$-complicated. Then the $(d-3)$-rd suspension of S is a PL d-sphere that is not LC.*

Proof. Let S' be the $(d-3)$-rd suspension of S, and let $\mathfrak{L}'$ be the subcomplex of S' obtained taking the $(d-3)$-rd suspension of the m-gonal knot $\mathfrak{L}$. Since $|S|-|\mathfrak{L}|$ is a deformation retract of $|S'|-|\mathfrak{L}'|$, they have the same homotopy groups. In particular, the fundamental group of $|S'|-|\mathfrak{L}'|$ has no presentation with $m \cdot 2^{d-3}$ generators. Now $\mathfrak{L}'$ is $(d-2)$-dimensional, and

$$f_{d-2}(\mathfrak{L}') = 2^{d-3} \cdot f_1(\mathfrak{L}) = m \cdot 2^{d-3},$$

whence we conclude via Corollary 2.21. □

Corollary 2.23. *For every $d \geq 3$, not all d-spheres are LC.*

Theorem 2.20 can be used in connection with the existence of knotted 2-spheres in $\mathbb{R}^4$ (see Kawauchi [31, p. 190]) to see that there are many non-LC 4-spheres beyond those that arise by suspension of 3-spheres. Thus, being "non-LC" is not simply induced by knots.

2.2 Many spheres are LC

Next we show that all constructible spheres are LC.

Lemma 2.24. *Let C be a d-pseudomanifold. Assume that $C = C_1 \cup C_2$, where C_1 and C_2 are d-pseudomanifolds. If C_1 and C_2 are both LC and $C_1 \cap C_2$ is a strongly connected $(d-1)$-pseudomanifold, then C is LC as well.*

Proof. Notice first that $C_1 \cap C_2 = \partial C_1 \cap \partial C_2$. In fact, every ridge of C belongs to two facets of C, hence every $(d-1)$-face σ of $C_1 \cap C_2$ is contained in exactly one d-face of C_1 and in exactly one d-face of C_2. In particular, there is an edge e_σ of the dual graph of C that punctures σ, and that connects a vertex of the dual graph of C_1 with a vertex of the dual graph of C_2.

Let C'_1 be a copy of C_1, and C'_2 a copy of C_2 (disjoint from C'_1). Each C'_i is LC; let us fix a local construction for each of them, and call T_i the tree along which C'_i is locally constructed. Furthermore, let C' be the pseudomanifold obtained merging C'_1 onto C'_2 by gluing together the two copies of σ. C' can be locally constructed along the tree $T_1 \cup T_2 \cup e_\sigma$ (just redo the same moves of the local constructions of the C_i's): so, C' is LC.

If $C_1 \cap C_2$ consists of one simplex only, then $C' \equiv C$ and we are already done. Otherwise, by the strongly connectedness assumption, the facets of $\partial C_1 \cap \partial C_2$ can be labeled $0, 1, \ldots, m$, so that:

- the facet labeled by 0 is σ;
- each facet labeled by $k \geq 1$ is adjacent to some facet labeled j with $j < k$.

Now for each $i \geq 1$, glue together the two copies of the facet i inside C'. All these gluings are *local* because of the labeling chosen, and we eventually obtain C. Thus, C is LC. $\square$

Since all constructible simplicial complexes are pure and strongly connected [9], we obtain for simplicial d-manifolds that

$$\{\text{constructible}\} \subseteq \{\text{LC}\}.$$

The previous containment is strict: Let C_1 be a $4 \times 4 \times 4$ pile of cubes, and let C_2 be the mirror image of C_1. Glue C_1 and C_2 together along the 2-dimensional annulus consisting of the external squares of one single face of C_1 (see Figure 4 below). Once triangulated properly,

Figure 4: Gluing the mirroring cubes along the green 2-dimensional region gives an LC, non-constructible 3-manifold.

$C_1 \cup C_2$ is a manifold which is LC (by Lemma 2.24) but not 2-connected (it retracts to a 2-sphere). So, LC d-manifolds are not necessarily $(d-1)$-connected. Since all constructible d-manifolds are $(d-1)$-connected [9, p. 1854], the previous argument produces many examples of d-manifolds with boundary that are LC, but not constructible. None of these examples, however, will be a sphere (or a ball). We do not know whether the containment $\{\text{constructible}\} \subseteq \{\text{LC}\}$ is strict for d-balls as well (see also Theorem 3.16).

We show now that for d-spheres, for every $d \geq 3$, the containment $\{\text{constructible}\} \subseteq \{\text{LC}\}$ is strict.

Lemma 2.25. *Suppose that we can find a* 3*-sphere* $\overline{S}$ *that is LC but not constructible. Then for all* $d \geq 3$*, the* $(d-3)$*-rd suspension of* $\overline{S}$ *is a* d*-sphere that is also LC but not constructible.*

Proof. Whenever S is an LC sphere, $v * S$ is an LC $(d+1)$-ball. (The proof is straightforward from the definition of “local construction”.) Now if S is LC, its suspension $(v * S) \cup (w * S)$ is also LC by Lemma 2.24. On the other hand, the suspension of a non-constructible sphere is a non-constructible sphere [25, Corollary 2]. $\square$

Of course, we should better show that the 3-sphere $\overline{S}$ in the assumption of Lemma 2.25 really exists. This will be established by Example 2.27. For this we make use of Corollary 2.12 as follows.

Lemma 2.26. *Let* B *be a* 3*-ball,* v *an external point, and* $B \cup v * \partial B$ *the* 3*-sphere obtained by adding to* B *a cone over its boundary. If* B *is collapsible, then* $B \cup v * \partial B$ *is LC.*

Proof. By Corollary 2.12, and since B is collapsible, all we need to prove is that $(B \cup v * \partial B) - (v * \sigma)$ collapses onto B, for some triangle σ in the boundary of B.

As all 2-balls are collapsible, and $\partial B - \sigma$ is a 2-ball, there is some vertex P in ∂B such that $\partial B - \sigma \searrow P$. This naturally induces a collapse of $v * \partial B \, - \, v * \sigma$ onto $\partial B \cup v * P$, according to the correspondence

$$\sigma \text{ is a free face of } \Sigma \iff v * \sigma \text{ is a free face of } v * \Sigma.$$

Collapsing the edge $v * P$ down to P, we get $v * \partial B \, - \, v * \sigma \searrow \partial B$.

It is crucial to notice that in the collapse given here, the pairs of faces removed are all of the form $(v * \sigma, v * \Sigma)$; thus, the facets of ∂B are always removed together with subfaces (and never with superfaces) in the collapse. This means that the freeness of the faces in ∂B is not needed; so when we glue back B the collapse $v * \partial B \, - \, v * \sigma \searrow \partial B$ can be read off as $B \cup v * \partial B \, - \, v * \sigma \searrow B$. $\square$

Example 2.27. In [35], Lickorish and Martin described a collapsible 3-ball B with a non-trivial knot in its 1-skeleton. An analogous result was obtained independently by Hamstrom and Jerrard [26]. The knot was an arbitrary 2-bridge index knot (for example, the trefoil knot). Merging B with the cone over its boundary, we obtain a knotted 3-sphere S which is LC (by Lemma 2.26; see also [34]) but not constructible (because it is knotted; see [25] [22, p. 54]).

Remark 2.28. In his 1991 paper [34, p. 530], Lickorish announces (without proof given) that "with a little ingenuity" one can get a sphere S with a 2-complicated triangular knot (the double trefoil), such that $S - \Delta$ is collapsible. Such a sphere is LC by Corollary 2.12.

Corollary 2.29. *For every fixed $d \geq 3$, not all LC d-spheres are constructible. In particular, a knotted 3-sphere can be LC if the knot is just 1-complicated or 2-complicated.*

The knot in the 1-skeleton of the ball B in Example 2.27 consists of a path on the boundary of B together with a "spanning edge", that is, an edge in the interior of B with both extremes on ∂B. This edge determines the knot, in the sense that any other path on ∂B between the two extremes of this edge closes it up into an equivalent knot. For these reasons such an edge is called a *knotted spanning edge*. More generally, a *knotted spanning arc* is a path of edges in the interior of a 3-ball, such that both extremes of the path lie on the boundary of the ball.[1]

The Example 2.27 can then be generalized by adopting the idea that Hamstrom and Jerrard used to prove their "Theorem B" [26, p. 331], as follows.

Theorem 2.30. *Let K be any 2-bridge knot (e.g. the trefoil knot). For any positive integer m, there exists a collapsible 3-ball B_m with a knotted spanning arc of m edges, such that the knot is the connected union of m copies of K.*

Proof. The case $m = 1$ is settled by Example 2.27; we will prove only the case $m = 2$, the general case being analogous.

Let B_1 be the 3-ball obtained when $m = 1$. In its 1-skeleton, B_1 has a knot consisting of a knotted spanning edge $\{A, Z\}$, together with some path $\mathfrak{L}$ from A to Z on the boundary of the ball. Choose a triangle $\{A, B, C\} \in \partial B_0$ such that the edge $\{A, B\}$ belongs to $\mathfrak{L}$. Let

[1] According to this definition, the relative interior of the arc is allowed to intersect the boundary of the 3-ball; this is the approach of Hachimori and Ehrenborg in [14].

$\Sigma = \{A,B,C,D\}$ be the unique tetrahedron of B_1 containing $\{A,B,C\}$, and fix a collapse of B_1. The idea is to glue B_1 to a "mirror copy" B'_1 of B_1, so that the triangle $\{A,B,C\}$ is identified with its mirror counterpart $\{A',B',C'\}$. Let B_2 be the resulting 3-ball. Clearly, the boundary path $\mathfrak{L}$ from A to Z in B_1, and its counterpart $\mathfrak{L}' \subset B'_1$, are merged together in a path on the boundary of B_2 that goes from Z to Z' avoiding A. At the same time, the spanning edge with its mirror image forms a knotted spanning arc of two edges.

It remains to show that B_2 is collapsible. This is not entirely trivial, because even if B_1 is collapsible, the "freeness" of the triangle $\{A,B,C\}$ was compromised in the very moment we glued B'_1 onto it. We should show that there is no loss of generality in assuming that $\{A,B,C\}$ was removed together with an *edge* in the collapse of B_1.

Suppose not; then $\{A,B,C\}$ was collapsed away together with some tetrahedron $\{A,B,C,D\}$. We modify B_1 a little bit, subdividing this tetrahedron into three tetrahedra via the insertion of a new vertex $\tilde{C}$ in the middle of the boundary triangle $\{A,B,C\}$. The pair $(\{A,B,C\}, \{A,B,C,D\})$ in the collapse of B_1 can then be replaced with the three pairs

$$(\{B,C,\tilde{C}\}, \{B,C,\tilde{C},D\}), \qquad (\{C,\tilde{C},D\}, \{A,C,\tilde{C},D\}), \qquad (\{A,\tilde{C},D\}, \{A,B,\tilde{C},D\});$$

in order to give a collapse of the subdivided ball, with the additional property that $\{A,B,\tilde{C}\}$ is removed together with an *edge*. Thus, up to replacing the ball B_1 with the subdivided one, and considering $\{A,B,\tilde{C}\}$ instead of $\{A,B,C\}$, we may assume that the merge of B_1 and B'_1 along $\{A,B,C\}$ collapses onto $\{A,B,C\}$. This means that B_2 is collapsible.

Since B_2 contains a knotted spanning arc of 2 consecutive edges ($\{Z,A\}$ and $\{A,Z'\}$), the cone over the boundary of B_2 is an LC 3-sphere with a knotted 4-gon – the knot being the sum of two copies of K. □

Corollary 2.31. *A* 3*-sphere with an m-complicated* $(m+2)$*-gonal knot can be LC.*

Proof. Let $S_m = B_m \cup v * B_m$, where B_m is the 3-ball constructed in the previous theorem. By Lemma 2.26, S_m is LC, and the spanning arc of m edges is closed up in v to form a $(m+2)$-gon. □

The spheres S_m are neither vertex decomposable, nor shellable, nor constructible, because of the following result about the bridge index.

Theorem 2.32 (Ehrenborg, Hachimori, Shimokawa [14] [24])**.** *Suppose that a* 3*-sphere (or a* 3*-ball) S contains a knot of m edges.*

- *If the bridge index of the knot exceeds* $\frac{m}{3}$*, then S is not vertex decomposable;*
- *If the bridge index of the knot exceeds* $\frac{m}{2}$*, then S is not constructible.*

The bridge index of a t-complicated knot is at least $t+1$, so if a knot is at least $\frac{m}{3}$-complicated, its bridge index automatically exceeds $\frac{m}{3}$. Thus, Ehrenborg–Hachimori–Shimokawa's theorem, the results of Hachimori and Ziegler in [25], the previous examples, and all our theorems and corollaries, blend into the following new hierarchy.

Theorem 2.33. *A* 3*-sphere with a non-trivial knot consisting of*

3 *edges,* 1*-complicated*	*is not constructible, but can be LC.*
3 *edges,* 2*-complicated*	*is not constructible, but can be LC.*
3 *edges,* 3*-complicated or more*	*is not LC.*
4 *edges,* 1*-complicated*	*is not vertex decomposable, but can be shellable.*
4 *edges,* 2*-complicated*	*is not constructible, but can be LC.*

4 edges, 3-complicated	*is not constructible.*
4 edges, 4-complicated or more	*is not LC.*
5 edges, 1-complicated	*is not vertex decomposable, but can be shellable.*
5 edges, 2- or 3-complicated	*is not constructible, but can be LC.*
5 edges, 4-complicated	*is not constructible.*
5 edges, 5-complicated or more	*is not LC.*
6 edges, 1-complicated	*can be vertex decomposable.*
6 edges, 2-, 3-, or 4-complicated	*is not constructible, but can be LC.*
6 edges, 5-complicated	*is not constructible*
6 edges, 6-complicated or more	*is not LC.*
$\vdots$	$\vdots$
m edges, k-complicated, with $k \geq \frac{m}{3}$	*is not vertex decomposable.*
m edges, k-complicated, with $k \geq \frac{m}{2}$	*is not constructible.*
m edges, k-complicated, with $k \leq m-2$	*can be LC.*
m edges, k-complicated, with $k \geq m$	*is not LC.*

We do not know whether for $m \geq 3$ a 3-sphere with an m-complicated $(m+1)$-gonal knot can be LC. However, an LC sphere with a 2-complicated triangular knot was found by Lickorish (see Remark 2.28).

One may also derive from Zeeman's theorem ("given any simplicial ball, there exists a positive integer r so that its r-th barycentric subdivision is collapsible" [47, Chapters I and III]) that any 3-sphere will become LC after sufficiently many barycentric subdivisions. On the other hand, there is no fixed number r of subdivisions that is sufficient to make *all* 3-spheres LC. (For this use sufficiently complicated knots, together with Theorem 2.14.)

3 On LC Balls

The combinatorial topology of d-balls and of d-spheres are intimately related: Removing any facet Δ from a PL d-sphere S we obtain a d-ball $S-\Delta$, and adding a cone over the boundary of a d-ball B we obtain a d-sphere S_B. Nevertheless, the combinatorics of triangulated balls seems to be more complicated (and/or less understood) than that of spheres. Thus the hierarchy for simplicial balls in this section — as collected in Theorem 3.1 — is given only for $d = 3$, while key questions remain open for $d > 3$. On the other hand, we do have a combinatorial characterization of LC d-balls, which we will reach in Theorem 3.10; it is a bit more complicated, but otherwise analogous to the characterization of LC d-spheres as given in Main Theorem 1.

Theorem 3.1. *For simplicial* 3*-balls,*

$$\{\text{vertex decomp.}\} \subsetneq \{\text{shellable}\} \subsetneq \{\text{constructible}\} \subseteq \{\text{LC}\} \subsetneq \{\text{collapsible}\} \subsetneq \{\text{all 3-balls}\}.$$

Proof. The first two inclusions are known (Ziegler's non-shellable ball from [49] is constructible by construction). We have already seen that all constructible complexes are LC (Lemma 2.24). Every LC 3-ball is collapsible by Corollary 3.12; the implication is strict due to Theorem 3.19. Finally, Bing and Goodrick showed that not every 3-ball is collapsible [8] [17]. □

3.1 Local constructions for d-balls

We begin with a relative version of the notions of "facet-killing sequence" and "facet massacre", which we introduced in the previous chapter.

Definition 3.2. Let P a pure d-complex. Let Q be a proper subcomplex of P, either pure d-dimensional or empty. A *facet-killing sequence of* (P,Q) is a sequence $P_0, P_1, \ldots, P_{t-1}, P_t$ of simplicial complexes such that $t = f_d(P) - f_d(Q)$, $P_0 = P$, and P_{i+1} is obtained by P_i removing a pair (σ, Σ) such that σ is a free $(d-1)$-face of Σ that does not lie in Q.

It is easy to see that P_t has the same $(d-1)$-faces as Q. The version of facet killing sequences given in Definition 2.3 is a special case of this one, namely the case when Q is empty.

Definition 3.3. Let P a pure d-dimensional simplicial complexes. Let Q be either the empty complex, or a pure d-dimensional proper subcomplex of P. A *pure facet-massacre of* (P,Q) is a sequence $P_0, P_1, \ldots, P_{t-1}, P_t$ of (pure) complexes such that $t = f_d(P) - f_d(Q)$, $P_0 = P$, and P_{i+1} is obtained by P_i removing:

(a) a pair (σ, Σ) such that σ is a free $(d-1)$-face of Σ that does not lie in Q, and

(b) all inclusion-maximal faces of dimension smaller than d that are left after the removal of type (a) or, recursively, after removals of type (b).

Necessarily $P_t = Q$ (and when $Q = \emptyset$ we recover the notion of facet-massacre of P, that we introduced in Definition 2.5). It is easy to see that a step $P_i \longrightarrow P_{i+1}$ can be factorized (not in an unique way) in an elementary collapse followed by a removal of k-faces, with $k < d$, that makes P_{i+1} a pure complex. Thus, a single pure facet-massacre of (P,Q) corresponds to many facet-killing sequences of (P,Q).

We will apply both definitions to the pair $(P,Q) = (K^T, \partial B)$, where K^T is defined for balls as follows.

Definition 3.4. If B be a d-ball with N facets, and T is a spanning tree of the dual graph of B, define K^T as the subcomplex of B formed by all $(d-1)$-faces of B that are not hit by T.

Lemma 3.5. *Under the previous notations,*

- K^T *is a pure* $(d-1)$*-dimensional simplicial complex, containing* ∂B *as a subcomplex;*
- K^T *has* $D + \frac{b}{2}$ *facets, where* b *is the number of facets in* ∂B*, and* $D := \frac{dN-N+2}{2}$*;*
- *for any* d*-simplex* Δ *of* B*,* $B - \Delta \searrow K^T$*;*
- K^T *is homotopy equivalent to a* $(d-1)$*-dimensional sphere.*

We introduce another convenient shortening.

Definition 3.6 (seepage)**.** Let B be a simplicial d-ball. A *seepage* is a $(d-1)$-dimensional subcomplex C of B whose $(d-1)$-faces are exactly given by the boundary of B.

Note that a seepage is not necessarily pure; actually there is only one pure seepage, namely ∂B itself. Since K^T contains ∂B, a collapse of K^T a seepage must remove all the $(d-1)$-faces of K^T that are not in ∂B: This is what we called a facet-killing sequence of $(K^T, \partial B)$.

Proposition 3.7. *Let* B *be a* d*-ball, and* Δ *a* d*-simplex of* B*. Let* C *be a seepage of* ∂B*. Then,*

$$B - \Delta \searrow C \iff \exists\, T \text{ s.t. } K^T \searrow C.$$

Proof. Analogous to the proof of Proposition 2.4. The crucial assumption is that no face of ∂B is removed in the collapse (since all boundary faces are still present in the final complex C). $\square$

If we fix a spanning tree T of the dual graph of B, we have then a 1-1 correspondence between the following sets:

1. the set of collapses $B - \Delta \searrow K^T$;
2. the set of "natural labelings" of T, where Δ is labelled by 1;
3. the set of the first parts $(T_1, \dots, T_N)$ of local constructions for B, with $T_1 = \Delta$.

Theorem 3.8. *Let B be a d-ball; fix a facet Δ, and a spanning tree T of the dual graph of B, rooted at Δ. The second part of a local construction for B along T corresponds bijectively to a facet-massacre of $(K^T, \partial B)$.*

Proof. Let us start with a local construction $[T_1, \dots, T_{N-1},]\, T_N, \dots, T_k$ for B along T. Topologically, $B = T_N/\!\sim$, where $\sim$ is the equivalence relation determined by the gluing, and $K^T = \partial T_N/\!\sim$.

K^T has $D + \frac{b}{2}$ facets (see Lemma 3.5), and all of them, except the b facets in the boundary, represent gluings. Thus we have to describe a sequence $P_0, \dots, P_t$ with $t = D - \frac{b}{2}$. But the local construction $(T_1, \dots, T_{N-1},)\, T_N, \dots, T_k$ produces B (which has b facets in the boundary) from T_N (which has $2D$ facets in the boundary, cf. Lemma 4.1) in $k - N$ steps, each removing a pair of facets from the boundary. So, $2D - 2(k-N) = b$, which implies $k - N = t$.

Define $P_0 := K_T = \partial T_N/\!\sim$, and $P_j := \partial T_{N+j}/\!\sim$. In the first LC step, $T_N \longrightarrow T_{N+1}$, we remove from the boundary a free ridge r, together with the unique pair σ', σ'' of facets of ∂T_N sharing r. At the same time, r and the newly formed face σ are sunk into the interior; so obviously neither σ nor r will appear in ∂B. This step $\partial T_N \longrightarrow \partial T_{N+1}$ naturally induces an analogous step $\partial T_{N+j}/\!\sim\ \longrightarrow \partial T_{N+j+1}/\!\sim$, namely, the removal of r and of the unique $(d-1)-$face σ containing it, with r not in ∂B.

The rest is analogous to the proof of Theorem 2.8. $\square$

Thus, B can be locally constructed along a tree T if and only if K^T collapses onto some seepage. What if we do not fix the tree T or the facet Δ?

Lemma 3.9. *Let B be a d-ball; let σ be a $(d-1)$-face in the boundary ∂B, and let Σ be the unique facet of B containing σ. Let C be a subcomplex of B. If C contains ∂B, the following are equivalent:*

1. *$B - \Sigma \searrow C$;*
2. *$B - \Sigma - \sigma \searrow C - \sigma$;*
3. *$B \searrow C - \sigma$.*

Theorem 3.10. *Let B be a d-ball. Then, the following are equivalent:*

1. *B is LC;*
2. *K^T collapses onto some seepage C, for some spanning tree T of the dual graph of B;*
3. *there exists a seepage C such that for every facet Δ of B one has $B - \Delta \searrow C$;*
4. *$B - \Delta \searrow C$, for some facet Δ of B, and for some seepage C;*
5. *there exists a seepage C such that for every facet σ of ∂B one has $B \searrow C - \sigma$;*
6. *$B \searrow C - \sigma$, for some facet σ of ∂B, and for some seepage C;*

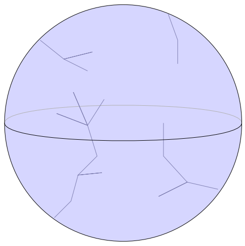

Figure 5: A seepage of a 3-ball.

Proof. The equivalences $1 \Leftrightarrow 2 \Leftrightarrow 3 \Leftrightarrow 4$ are established analogously to the proof of Theorem 2.11. Finally, Lemma 3.9 implies that $3 \Rightarrow 5 \Rightarrow 6 \Rightarrow 4$. □

Corollary 3.11. *Let B be a 3-ball. Then the following are equivalent:*

1. *B is LC;*
2. *$K^T \searrow \partial B$, for some spanning tree T of the dual graph of B;*
3. *$B - \Delta \searrow \partial B$, for every facet Δ of B;*
4. *$B - \Delta \searrow \partial B$, for some facet Δ of B;*
5. *$B \searrow \partial B - \sigma$, for every facet σ of ∂B;*
6. *$B \searrow \partial B - \sigma$, for some facet σ of ∂B.*

Proof. When B has dimension 3, any seepage C of ∂B is a 2-complex containing ∂B, plus some edges and vertices. If a complex homotopy equivalent to a S^2 collapses onto C, then C is also homotopy equivalent to S^2, thus C can only be ∂B with some trees attached (see Figure 5); which implies that $C \searrow \partial B$. □

Corollary 3.12. *For 3-balls,* LC $\Rightarrow$ *collapsible.*

Proof. If B is LC, it collapses to some 2-ball $\partial B - \sigma$, but all 2-balls are collapsible. □

Corollary 3.13. *All constructible 3-balls are collapsible.*

For example, Ziegler's ball, Grünbaum's ball, and Rudin's ball are collapsible (see [49]).

Remark 3.14. The locally constructible 3-balls with N facets are precisely the 3-balls that admit a "special collapse", namely such that after the first elementary collapse, in the next $N-1$ collapses, no triangle of ∂B is collapsed away. Such a collapse acts along a dual (directed) tree of the ball, whereas a generic collapse acts along an acyclic graph that might be disconnected.

One could argue that maybe "special collapses" are not that special: Perhaps every collapsible 3-ball has a collapse that removes only one boundary triangle in its top-dimensional phase? This does not hold, and we will produce a counterexample in the next paragraph (see Theorem 3.19).

3.2 3-Balls without interior vertices.

Here we show that balls with all vertices on the boundary are LC if and only if they are constructible. We use this fact to establish our hierarchy for 3-balls (Theorem 3.1). At the same time, we review, revise, and extend some of the results of Hachimori [21]. In particular we will show that a type (ii) LC step transforms a constructible 3-ball into a constructible 3-ball.

Proposition 3.15 (Hachimori)**.** *Let B be a constructible 3-ball. Let C be a 3-ball obtained from B by gluing together two triangles of ∂B. Then, C is constructible.*

The "obvious" thing to try is to recycle for C the constructible decomposition $B = B_1 \cup B_2$ of B: that is, to split C into two balls C_1 and C_2 with the same facets of B_1 and B_2, respectively. This does not always work in the way stated in [21, p. 227], since the intersection $C_1 \cap C_2$ might not be a 2-ball anymore. However, Hachimori's proof can be salvaged, and the result is valid; see Benedetti [7].

Theorem 3.16. *Let B be a* 3*-ball with no interior vertex. Then B is LC if and only if B is constructible.*

Proof. Assume B is LC. If B is a tree of tetrahedra there is nothing to prove; otherwise, we claim that B is obtained from a tree of tetrahedra, through a sequence of identifications only of type (ii).

In fact, steps of type (v), (viii) or (ix) sink respectively 1, 2 or 3 vertices into the interior of B; so, they cannot occur. Besides, any identification of type (vi), (vii), or (iv) increases the number of connected component in the boundary; hence, it must be followed by at least one step of type (ix), that destroys a connected component of the boundary. Since (ix) is forbidden, no identification of type (vi), (vii), or (iv) can occur. Finally, a pinching step like (iii), needs to be followed by one of the steps (vi), (vii), (viii) or (ix) in order to restore the ball topology – but such steps are forbidden, so also step (iii) cannot occur. Concluding this Sudoku, each identification step in a local construction for B must be of type (ii).

Each type (ii) step leaves the topology unchanged. Moreover, none of these steps leads out of the world of simplicial manifolds; for otherwise, to adjust things, a different step than (ii) would be needed. By Proposition 3.15, we conclude. $\square$

Hachimori's algorithm [21, Lemma 3, p. 227] [22, Chapter 4] to test *constructibility* of 3-balls without interior vertices is thus also an algorithm to decide whether such 3-balls are LC or not.

Furthermore, Theorem 3.16 shows that it is crucial to admit the generality of regular CW complexes in the definition of "LC". Recall that we have seen examples of LC simplicial complexes that have many interior vertices, and different topology than a ball (namely, LC spheres); some of these were not constructible, either. The situation is very different for LC simplicial complexes.

Corollary 3.17. *Let C be a simplicial complex. If every pseudomanifold in a local construction for C is a simplicial complex, then P is a constructible* 3*-ball with all its vertices on the boundary.*

Proof. All local gluings in a local construction of C are necessarily of type (ii). As these steps do not alter the topology and do not sink vertices into the interior, C is a 3-ball with all vertices on the boundary. Constructibility follows then from Theorem 3.16, or alternatively from Proposition 3.15. $\square$

We are going to exploit Theorem 3.16 to obtain examples of non-LC 3-balls. We already know that if a ball B is not collapsible, then B is not LC, by Corollary 3.12. Thus, a ball with a knotted spanning edge cannot be LC if the knot is the sum of two or more trefoil knots. (See also Bing [8] and Goodrick [17].)

What about knotted balls with a single trefoil knot inside? Via Theorem 3.16, we have an answer in the case that the ball has no interior vertices.

Corollary 3.18. *Knotted balls without interior vertices are not LC. Moreover, Hachimori's triangulations of Bing's house with two rooms and of Furch's ball are not LC.*

Proof. Knotted balls are not constructible [21], [25]. The triangulation in [20] of Furch's ball [15] has all the vertices on the boundary, thus it is not LC by Theorem 3.16. Hachimori also triangulated [21] a 3-dimensional thickening of Bing's house with two rooms with all the vertices on the boundary, proving its non-constructibility. Such a triangulated ball cannot be LC, either. □

We can now move on to complete the proof of our Theorem 3.1. The following result is inspired by Lickorish–Martin's Theorem ([35], quoted also in our Example 2.27): We realized a collapsible triangulation of a 3-ball with a knotted spanning edge and no interior vertices.

Theorem 3.19. *Not all collapsible* 3*-balls are LC.*

Proof. Start with a large $m \times m \times 1$ pile of cubes, triangulated in the standard way, and take away two distant cubes, leaving only their bottom squares (which will be called from now on the *free squares* X and Y). The 3-complex C obtained can be collapsed vertically onto its square basis; in particular, it is collapsible, and has no interior vertices.

Let C' be a 3-ball with two tubular holes drilled away, but where 1) each hole has been corked at a bottom with a 2-disk, and 2) the tubes are disjoint but intertwined, so that a closed path that passes through both holes and between these traverses the top resp. bottom face of C' yields a trefoil knot (see Figure 6).

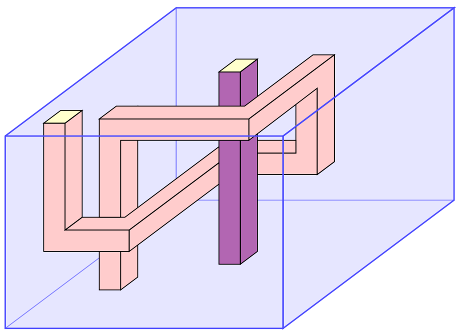

Figure 6: C' is obtained from a 3-ball drilling away two intertwined tubular holes, and then "corking" the holes on the bottom with 2-dimensional membranes.

C and C' are homeomorphic. Any homeomorphism induces on C' a collapsible triangulation with no interior vertices. The free squares of C correspond via the homeomorphism to the corking membranes of C', which we will call correspondingly X' and Y'. To get from C' to a ball with a knotted spanning edge we will carry out two more steps:

(i) create a single edge $\{x', y'\}$ that goes from X' to Y';

(ii) thicken the "bottom" of C' a bit, so that C' becomes a 3-ball and $\{x', y'\}$ becomes an interior edge (even if its extremes are still on the boundary).

We perform both steps by adding cones over 2-disks to the complex. Such steps preserve collapsibility, but in general they produce interior vertices; thus we choose "specific" disks with few interior vertices.

(i) Provided m is large enough, one finds a “nice” strip $F_1, F_{2,3}, \ldots, F_k$ of triangles on the bottom of C', such that $F_1 \cup F_2 \cup \cdots \cup F_k$ is a disk without interior vertices, F_1 has a single vertex x' in the boundary of X', while F_k has a single vertex y' in the boundary of Y', and the whole strip intersects $X' \cup Y'$ only in x' and y'. Then we add a cone to C', setting

$$C_1 := C' \cup \left(y' * (F_1 \cup F_2 \cup \cdots \cup F_{k-1})\right).$$

(An explicit construction of this type is carried out in [25, pp. 164-165].) Thus one obtains a collapsible 3-complex C_1 with no interior vertex, and with a direct edge from X' to Y'.

(ii) Let R be a 2-ball inside the boundary of C_1 that contains in its interior the 2-complex $X' \cup Y' \cup \{x', y'\}$, and such that every interior vertex of R lies either in X' or in Y'. Take a new point z' and define $C_2 := C_1 \cup (z' * R)$.

As $z' * R$ collapses onto R, it is easy to verify that C_2 is a collapsible 3-ball with a knotted spanning edge $\{x', y'\}$. By Corollary 3.18, C_2 is not LC. □

Corollary 3.20. *There exists a collapsible* 3*-ball B such that every collapse of B involves at least two pairs of the form* (boundary triangle, tetrahedron).

We conclude this chapter observing that Chillingworth’s theorem, “every geometric triangulation of a convex 3-dimensional polytope is collapsible”, can be strengthened as follows.

Theorem 3.21 (Chillingworth [12])**.** *Every* 3*-ball embeddable as a convex subset of the Euclidean* 3*-space* $\mathbb{R}^3$ *is LC.*

Proof. The argument of Chillingworth for collapsibility runs showing that $B \searrow \partial B - \sigma$, where σ is any triangle in the boundary of B. Now Theorem 3.11 ends the proof. □

Any subdivided 3-simplex is LC. If it has all vertices on the boundary, then it is constructible, by Theorem 3.16. Note that a subdivided 3-simplex with all vertices on the boundary might be non-shellable (e.g. Rudin’s ball).

4 Upper bounds on the number of LC d-spheres.

For fixed $d \geq 2$ and a suitable constant C that depends on d, there are less than C^N combinatorial types of d-spheres with N facets. Our proof for this fact is a d-dimensional version of the main theorem in [13], and allows us to determine an explicit constant C, for any d. It consists in two different phases:

1. we observe that there are less trees of d-simplices than planted plane d-ary trees, which are counted by order d Fuss–Catalan numbers;
2. we count the number of “LC matchings” according to ridges in the tree of simplices.

4.1 Counting the trees of d-simplices.

We will here establish that there are less than $C_d(N) := \frac{1}{(d-1)N+1}\binom{dN}{N}$ trees of N d-simplices.

Lemma 4.1. *Every tree of N d-simplices has $(d-1)N+2$ boundary facets of dimension $d-1$ and $N-1$ interior faces of dimension $d-1$.*
It has $\frac{d}{2}((d-1)N+2)$ faces of dimension $d-2$, all of them lying in the boundary.

By *rooted* tree of simplices we mean a tree of simplices B together with a distinguished facet δ of ∂B, whose vertices have been labelled $1, 2, \dots, d$. Rooted trees of d-simplices are in bijection with "planted plane d-ary trees", that is, plane rooted trees such that every non-leaf vertex has exactly d (left-to-right-ordered) sons; cf. [36].

Theorem 4.2. *There is a bijection between rooted trees of N d-simplices and planted plane d-ary trees with N non-leaf vertices, which in turn are counted by the Fuss–Catalan numbers $C_d(N) = \frac{1}{(d-1)N+1}\binom{dN}{N}$. Thus, the number of combinatorially-distinct trees of N d-simplices satisfies*

$$\frac{1}{(d-1)N+2}\,\frac{1}{d!}\,C_d(N) \;\leq\; \#\{\text{ trees of } N \;\; d\text{-simplices }\} \;\leq\; C_d(N).$$

Proof. Given a rooted tree of d-simplices with a distinguished facet δ in its boundary, there is a unique extension of the labeling of the vertices of δ to a labeling of all the vertices by labels $1, 2, \dots, d+1$, such that no two adjacent vertices get the same label. Thus each d-simplex receives all $d+1$ labels exactly once.

Now, label each $(d-1)$-face by the unique label that none of its vertices has. With this we get an edge-labeled rooted d-ary tree whose non-leaf vertices correspond to the N d-simplices; the root corresponds to the d-simplex that contains δ, and the labeled edges correspond to all the $(d-1)$-faces other than δ. We get a plane tree by ordering the down-edges at each non-leaf vertex left to right according to the label of the corresponding $(d-1)$-face.

The whole process is easily reversed, so that we can get a rooted tree of d-simplices from an arbitrary planted plane d-ary tree.

There are exactly $C_d(N) = \frac{1}{(d-1)N+1}\binom{dN}{N}$ planted plane d-ary trees with N interior vertices (see e.g. Aval [5]; the integers $C_2(N)$ are the "Catalan numbers", which appear in many combinatorial problems, see e.g. Stanley [43, Ex. 6.19]). Any tree of N d-simplices has $(d-1)N+2$ boundary facets, so it can be rooted in exactly $((d-1)N+2)\,d!$ ways, which however need not be inequivalent. This explains the first inequality claimed in the lemma. Finally, combinatorially-inequivalent trees of d-simplices also yield inequivalent rooted trees, whence the second inequality follows. $\square$

Corollary 4.3. *The number of trees of N d-simplices, for N large, is bounded by*

$$\binom{dN}{N} \sim \left(d \cdot \left(\tfrac{d}{d-1}\right)^{d-1}\right)^N < (de)^N.$$

4.2 Counting the matchings in the boundary.

We know from the previous section that there are exponentially many trees of N d-simplices. Our goal is to find an exponential upper bound for the LC spheres obtainable by a matching of adjacent facets in the boundary of one fixed tree of simplices.

Theorem 4.4. *Fix $d \geq 2$. The number of combinatorially distinct LC d-spheres with N facets, for N large, is not larger than*

$$\left(d \cdot \left(\tfrac{d}{d-1}\right)^{d-1} \cdot 2^{\frac{2d^2-d}{3}}\right)^N.$$

Proof. Let us fix a tree of N d-simplices B. We adopt the word "couple" to denote a pair of facets in the boundary of B that are glued to one another during the local construction of S.

Let us set $D := \frac{1}{2}(2+N(d-1))$, which is an integer. By Lemma 4.1, the boundary of the tree of N d-simplices contains $2D$ facets, so each perfect matching is just a set of D pairwise disjoint couples. We are going to partition every perfect matching into "rounds". The first round will contain couples that are adjacent in the boundary of the tree of simplices. Recursively, the $(i+1)$-th round will consist of all pairs of facets that *become* adjacent only after a pair of facets are glued together in the i-th round.

Selecting a pair of adjacent facets is the same as choosing the ridge between them; and by Lemma 4.1, the boundary contains dD ridges. Thus the first round of identifications consists in choosing n_1 ridges out of dD, where n_1 is some positive integer. After each identification, at most $d-1$ new ridges are created; so, after this first round of identifications, there are at most $(d-1)n_1$ new pairs of adjacent facets.

In the second round, we identify $2n_2$ of these newly adjacent facets: as before, it is a matter of choosing n_2 ridges, out of the at most $(d-1)n_1$ just created ones. Once this is done, at most $(d-1)n_2$ ridges are created. And so on.

We proceed this way until all the $2D$ facets in the boundary of B have been matched (after f steps, say). Clearly $n_1+\ldots+n_f = D$, and since the n_i's are positive integers, $f \le D$ must hold. This means there are at most

$$\sum_{f=1}^{D} \quad \sum_{\substack{n_1,\ldots,n_f \\ n_i \ge 1,\, \sum n_i = D \\ n_{i+1} \le (d-1)n_i}} \binom{dD}{n_1}\binom{(d-1)n_1}{n_2}\binom{(d-1)n_2}{n_3}\cdots\binom{(d-1)n_{f-1}}{n_f}$$

possible perfect matchings of $(d-1)$-simplices in the boundary of a tree of N d-simplices.

We sharpen this bound by observing that not all ridges may be chosen in the first round of identifications. For example, we should exclude those ridges that belong to just two d-simplices of B. An easy double-counting argument reveals that in a tree of d-simplices, the number of ridges belonging to at least 3 d-simplices is smaller or equal than $\frac{N}{3}\binom{d+1}{2}$. So in the upper bound above we may replace the first factor $\binom{dD}{n_1}$ with the smaller factor $\binom{\frac{N}{3}\binom{d+1}{2}}{n_1}$.

To bound the sum from above, we use $\binom{n}{k} \le 2^n$ and $n_1+\cdots+n_{f-1} < n_1+\cdots+n_f = D$, while ignoring the conditions $n_{i+1} \le (d-1)n_i$. Thus we obtain the upper bound

$$2^{\frac{N}{3}\binom{d+1}{2}+\frac{N}{2}(d-1)^2+(d-1)} \cdot \sum_{f=1}^{D}\binom{D-1}{f-1} \;=\; 2^{\frac{N}{3}(2d^2-d)+(d-1)}.$$

Thus the number of ways to fold a tree of N d-simplices into a sphere via a local construction sequence is smaller than than $2^{\frac{2d^2-d}{3}N}$. Combining this with Proposition 4.2, we conclude. $\square$

The upper bound above can be simplified in many ways. For example, when $d \ge 16$, it is smaller than $\sqrt[3]{4}^{d^2}$. Explicitly, from Theorem 4.4 we obtain the following upper bounds:

- There are less than 216^N LC 3-spheres with N facets,
- there are less than 6117^N LC 4-spheres with N facets,

and so on. We point out that these upper bounds are not sharp, as we overcounted both on the combinatorial side and on the algebraic side. When $d = 2$, Tutte's upper bound is asymptotically 3.08^N, whereas the one given by our formula is 16^N. When $d = 3$, however, our constant is smaller than what follows from Durhuus–Jonsson's original argument (216^N vs. 3405^N).

Corollary 4.5. *For any fixed $d \geq 2$, there are exponential lower and upper bounds for the number of LC d-spheres on N facets.*

Proof. We have just obtained an upper bound; however, we also obtain a lower bound from Proposition 4.2/Corollary 4.3, since the boundary of a tree of $(d+1)$-simplices is a stacked d-sphere, and for $d \geq 2$ the stacked d-sphere determines the tree of $(d+1)$-simplices uniquely. □

Acknowledgement. We are very grateful to Matthias Staudacher, Niko Witte, Raman Sanyal, Thilo Rörig, Frank Lutz, Gil Kalai, and Emo Welzl for useful discussions and references.

References

[1] N. ALON, *The number of polytopes, configurations and real matroids*, Mathematika, 33 (1986), pp. 62–71.

[2] J. AMBJØRN, D. V. BOULATOV, N. KAWAMOTO, AND Y. WATABIKI, *Recursive sampling simulations of 3D gravity coupled to scalar fermions*, Nucl. Phys. B, 480 (2000), pp. 319–330.

[3] J. AMBJØRN, B. DURHUUS, AND T. JONSSON, *Quantum Geometry*, Cambridge University Press, Cambridge, 1997.

[4] J. AMBJØRN AND S. VARSTED, *Three dimensional simplicial quantum gravity*, Nucl. Phys. B, 373 (1992), pp. 557–577.

[5] J.-C. AVAL, *Multivariate Fuss–Catalan numbers*, Discrete Math., 308 (2008), pp. 4660–4669.

[6] C. BARTOCCI, U. BRUZZO, M. CARFORA, AND A. MARZUOLI, *Entropy of random coverings and 4-D quantum gravity*, Journal of Geometry and Physics, 18 (1996), pp. 247–294.

[7] B. BENEDETTI, *Locally Constructible Balls and Spheres*, PhD thesis, TU Berlin, 2009. In preparation.

[8] R. H. BING, *Some aspects of the topology of 3-manifolds related to the Poincaré Conjecture*, in Lectures on Modern Mathematics, T. Saaty, ed., vol. II, Wiley, 1964, pp. 93–128.

[9] A. BJÖRNER, *Topological methods*, in Handbook of Combinatorics, R. Graham, M. Grötschel, and L. Lovász, eds., vol. II, Amsterdam, 1995, Elsevier, pp. 1819–1872.

[10] S. CATTERALL, J. KOGUT, AND R. RENKEN, *Is there an exponential bound in four dimensional simplicial gravity?*, Phys. Rev. Lett. 72, Issue 26 (1994), pp. 4062–4065.

[11] J. CHEEGER, *Critical points of distance functions and applications to geometry*, vol. 1504 of Springer Lect. Notes in Math., Montecatini Terme, 1991, pp. 1–38.

[12] D. R. J. CHILLINGWORTH, *Collapsing three-dimensional convex polyhedra*, Proc. Camb. Phil. Soc., 63 (1967), pp. 353–357. Erratum in Math. Proc. Camb. Phil. Soc. 88 (1980), pp. 307–310.

[13] B. DURHUUS AND T. JONSSON, *Remarks on the entropy of 3-manifolds*, Nuclear Physics B, 445 (1995), pp. 182–192.

[14] R. EHRENBORG AND M. HACHIMORI, *Non-constructible complexes and the bridge index*, Europ. J. Combinatorics, 22 (2001), pp. 475–491.

[15] R. FURCH, *Zur Grundlegung der kombinatorischen Topologie*, Abh. Math. Sem. Univ. Hamburg, 3 (1924), pp. 69–88.

[16] J. E. GOODMAN AND R. POLLACK, *There are asymptotically far fewer polytopes than we thought*, Bull. Amer. Math. Soc., 14 (1986), pp. 127–129.

[17] R. GOODRICK, *Non-simplicially collapsible triangulations of I^n*, Proc. Camb. Phil. Soc., 64 (1968), pp. 31–36.

[18] M. GROMOV, *Spaces and questions*, in GAFA 2000, N. Alon, J. Bourgain, A. Connes, M. Gromov, and V. Milman, eds., no. Special Volume, Part I, Birkhäuser, 2000, pp. 118–161.

[19] C. GROVE, P. PETERSEN, AND J.-Y. WU, *Geometric finiteness theorems via controlled topology*, Invent. Math., 99 (1990), pp. 205–213. Erratum in Invent. Math. 104 (1991), pp. 221–222.

[20] M. HACHIMORI, *Simplicial complex library*. http://infoshako.sk.tsukuba.ac.jp/~HACHI/math/library/index_eng.html.

[21] ——, *Nonconstructible simplicial balls and a way of testing constructibility*, Discrete Comput. Geom., 22 (1999), pp. 223–230.

[22] ——, *Combinatorics of constructible complexes*, PhD thesis, Tokyo University, 2000.

[23] ——, *Decompositions of two-dimensional simplicial complexes*, Discrete Mathematics, 308 (2008), pp. 2307–2312.

[24] M. HACHIMORI AND K. SHIMOKAWA, *Tangle sum and constructible spheres*, Journal of Knot Theory and Its Ramifications, 13 (2004), pp. 373–383.

[25] M. HACHIMORI AND G. M. ZIEGLER, *Decompositions of simplicial balls and spheres with knots consisting of few edges*, Math. Z., 235 (2000), pp. 159–171.

[26] M. E. HAMSTROM AND R. P. JERRARD, *Collapsing a triangulation of a "knotted" cell*, Proc. Amer. Math. Soc., 21 (1969), pp. 327–331.

[27] C. HOG-ANGELONI AND W. METZLER, *Geometric aspects of two-dimensional complexes*, in Two-dimensional homotopy and combinatorial group theory, vol. 197 of London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 1993, pp. 1–50.

[28] J. F. P. HUDSON, *Piecewise Linear Topology*, Benjamin, New York, 1969.

[29] G. KALAI, *Many triangulated spheres*, Discrete Comput. Geom., 3 (1988), pp. 1–14.

[30] S. KAMEI, *Cones over the boundaries of non-shellable but constructible 3-balls*, Osaka J. Math., 41 (2004), pp. 357–370.

[31] A. KAWAUCHI, *A survey of knot theory*, Birkhäuser, Basel, 1996.

[32] V. KLEE AND P. KLEINSCHMIDT, *The d-step conjecture and its relatives*, Math. Operations Research, 12 (1987), pp. 718–755.

[33] C. W. LEE, *Kalai's squeezed spheres are shellable*, Discrete Comput. Geometry, 24 (2000), pp. 391–396. The Branko Grünbaum birthday issue (G. Kalai, V. Klee, eds.).

[34] W. B. R. LICKORISH, *Unshellable triangulations of spheres*, Europ. J. Combinatorics, 12 (1991), pp. 527–530.

[35] W. B. R. LICKORISH AND J. M. MARTIN, *Triangulations of the 3-ball with knotted spanning 1-simplexes and collapsible r-th derived subdivisions*, Trans. Amer. Math. Soc., 170 (1972), pp. 451–458.

[36] J. MATOUŠEK AND J. NEŠETŘIL, *Invitation to Discrete Mathematics*, Clarendon Press, Oxford, 1998. (Second edition 2009).

[37] J. R. MUNKRES, *Elements of Algebraic Topology*, Addison-Wesley, Menlo Park, CA, 1984.

[38] J. PFEIFLE AND G. M. ZIEGLER, *Many triangulated 3-spheres*, Mathematische Annalen, 330 (2004), pp. 829–837.

[39] J. S. PROVAN AND L. J. BILLERA, *Decompositions of simplicial complexes related to diameters of convex polyhedra*, Math. Operations Research, 5 (1980), pp. 576–594.

[40] T. REGGE, *General relativity without coordinates*, Il Nuovo Cimento, 19 (1961), pp. 558–571.

[41] T. REGGE AND R. M. WILLIAMS, *Discrete structures in gravity*, J. Math. Phys., 41 (2000), pp. 3964–3983.

[42] D. ROLFSEN, *Knots and links*, Publish or Perish, University of British Columbia, 1976.

[43] R. STANLEY, *Enumerative Combinatorics Vol. 2*, Cambridge University Press, Cambridge, 1999. See also his *Catalan Addendum*, http://www-math.mit.edu/~rstan/ec/catadd.pdf.

[44] W. T. TUTTE, *A census of planar triangulations*, Canadian J. Math, 14 (1962), pp. 21–38.

[45] ——, *On the enumeration of convex polyhedra*, J. Combinatorial Theory, Ser. B, 28 (1980), pp. 105–126.

[46] D. WEINGARTEN, *Euclidean quantum gravity on a lattice*, Nucl. Phys. B, 210 (1982), pp. 229–245.

[47] E. C. ZEEMAN, *Seminar on Combinatorial Topology*, Institut des Hautes Études Scientifiques and University of Warwick, Paris – Coventry, 1966.

[48] G. M. ZIEGLER, *Lectures on Polytopes*, vol. 152 of Graduate Texts in Math., Springer Verlag, 1995. (Revised seventh printing 2007.)

[49] ——, *Shelling polyhedral* 3*-balls and* 4*-polytopes*, Discrete Comput. Geom., 19 (1998), pp. 159–174.