Flip graphs of stacked and flag triangulations of the 2-sphere

Benjamin A. Burton*

The University of Queensland
School of Mathematics and Physics
Brisbane QLD 4072, Australia

bab@maths.uq.edu.au

Basudeb Datta†

Department of Mathematics
Indian Institute of Science
Bangalore 560 012, India

dattab@iisc.ac.in

Jonathan Spreer‡

School of Mathematics and Statistics
The University of Sydney
NSW 2006, Australia

jonathan.spreer@sydney.edu.au

Abstract

It is well-known that the flip graph of n-vertex triangulated 2-spheres is connected, i.e., each pair of n-vertex triangulated 2-spheres can be turned into each other by a sequence of edge flips for each $n \geqslant 4$. In this article, we study various induced subgraphs of this graph. In particular, we prove that the subgraph of n-vertex *flag* 2-spheres distinct from the double cone is still connected. In contrast, we show that the subgraph of n-vertex *stacked* 2-spheres has at least as many connected components as there are trees on $\lfloor \frac{n-5}{3} \rfloor$ nodes with maximum node-degree at most four.

Keywords: Triangulated 2-sphere, planar triangulation, flip graph, Pachner graph, edge flip, flag 2-sphere, stacked 2-sphere.

Mathematics Subject Classifications: 57Q15; 57M20, 05C10

https://doi.org/10.37236/10292

1   Introduction

The *Pachner graph* of triangulated 2-spheres is the graph whose nodes are triangulated 2-spheres (also known as planar triangulations), and two nodes are connected by an arc if and only if their corresponding triangulations can be transformed into each other by a single *bistellar move*, i.e., an edge flip, a stellar subdivision of a triangle or its inverse, see Figure 2.

The Pachner graph of triangulated 2-spheres is connected. More precisely, starting from an arbitrary node representing an n-vertex 2-sphere, a path of length $O(n)$ can be found in the Pachner graph ending at the node representing the boundary of the tetrahedron. Conversely, it is not difficult to see that $\Omega(n)$ arcs are also necessary for the length of such a path.

The Pachner graph has a natural graded structure into induced subgraphs on the sets of nodes representing n-vertex triangulated 2-spheres with n fixed: The arcs within a level correspond to edge flips, the arcs corresponding to stellar subdivisions (and their inverses) connect different levels of the grading. It is well-known that each such level, sometimes called the *flip graph (of n-vertex triangulated 2-spheres)*, is connected [21]. Moreover, its diameter is bounded from above by $5n - 23$ due to work by Cardinal, Hoffmann, Kusters, Tóth and Wettstein [6] and bounded from below by $7n/3 - 34$ due to work by Frati [10]. These two results are the most recent additions to a series of papers aimed at reducing the gap between upper bounds and lower bounds for the diameter of the flip graph. One of the current open problems in this area is to find an upper bound and a lower bound which differ by a factor of two (the optimum achievable by bounding the diameter as twice the distance of a particular pair of triangulations). See [5] for a survey on previous attempts to bound the diameter of the flip graph of the 2-sphere.

Sulanke and Lutz [20] show that there are exactly 59 twelve-vertex triangulations of the orientable surface of genus six. Since they all must be neighbourly, none of them allows any edge flips. Thus, the flip graph of twelve-vertex triangulated orientable surfaces of genus six is the discrete graph on 59 nodes.

See various chapters of the book of De Loera, Rambau, and Santos [8] for further and closely related research concerning the flip graph and similar objects. See also the survery article of Bose and Hurtado [3], and Bose and Verdonshot [5] for survey on results on the flip graph of planar graphs and triangulations.

Structural results for, as well as bounds on flip distances in Pachner graphs (of spheres or, more generally, triangulated manifolds) which are as precise as the ones mentioned above, are unlikely to be provable in dimensions greater than two. For instance, the best upper bound for distances in the Pachner graph of *generalised triangulations of the 3-sphere* is given by $O(t^2 2^{ct^2})$ for the number of moves between a t-tetrahedron triangulation of S^3 and the boundary of the 4-simplex, see Mijatović [15]. Naturally, the corresponding upper bound in the simplicial setting must be at least as large. Moreover, the n-th level of the Pachner graph of simplicial triangulations of the 3-sphere is not even connected (in contrast to the setting of generalised triangulations, see [14]): Consider an n-vertex triangulation of the 3-sphere containing (i) no edge of degree three and (ii) the complete

graph with n vertices as edges. Such a triangulation only admits stellar subdivisions as bistellar moves and is thus isolated in the Pachner graph of n-vertex triangulated 3-spheres. See [9, 18] for a number of examples of such triangulated 3-spheres.

Even more, in dimensions greater than four, no such general upper bounds can exist at all due to the undecidability of the homeomorphism problem.

In this paper we focus on the connectedness of certain subgraphs of the flip graph of n-vertex triangulated 2-spheres. Namely, we consider what are called *stacked* and *flag 2-spheres* (see Sections 2.2 and 2.3 for details). In many ways, flag 2-spheres are the counterpart to stacked 2-spheres. While stacked 2-spheres contain the maximum number of induced 3-cycles, flag 2-spheres do not contain any such cycle. Moreover, every triangulated 2-sphere can be decomposed into a collection of flag 2-spheres and boundaries of the tetrahedron (called *standard 2-spheres*) by iteratively cutting along its induced 3-cycles and pasting the missing triangles. For a flag 2-sphere this decomposition is the 2-sphere itself. For stacked 2-spheres it yields the maximum number of connected components, each isomorphic to the standard 2-sphere.

Mori, Nakamoto, and Ota [16, Theorem 5] prove upper bounds for the number of edge flips connecting two flag 2-spheres within the class of Hamiltonian triangulations. Our main result states that such a sequence of edge flips exists even *within the class of flag 2-spheres* – as long as both triangulations are distinct from the double cone Γ_n over the $(n-2)$-gon (Figure 6(a)), see Theorem 6. Observe that excluding the n-vertex double cone Γ_n, $n \geqslant 6$, from Theorem 6 is necessary: Γ_n is a flag 2-sphere in which every edge contains a degree four vertex. Thus every edge flip on Γ_n produces a vertex of degree three (implying that there exist a missing triangle) and the resulting complex is not flag. In particular, Γ_n cannot be connected to any other flag 2-sphere by an edge flip.

This theorem complements a result by Lutz and Nevo [13] stating that for $d \geqslant 3$, every pair of piecewise linear homeomorphic d-dimensional flag complexes is connected by a sequence of edge subdivisions, and edge contractions.

In contrast, the subgraph of the flip graph of n-vertex stacked 2-spheres has much less uniform properties. In Section 4 we give a precise condition on when exactly an edge flip of a stacked 2-sphere produces another stacked 2-sphere (Theorem 14). Using this result, we prove that the flip graph of n-vertex stacked 2-spheres is not connected, and that there are at least as many connected components as there are trees on $\lfloor \frac{n-5}{3} \rfloor$ nodes and with degrees of nodes at most four. In particular, the number of connected components of the flip graph of n-vertex stacked 2-spheres is exponential in n (Corollary 20). Furthermore, we show that a pair of n-vertex stacked 2-spheres can be connected by a sequence of n-vertex stacked 2-spheres, each related to the previous one by an edge flip, if their associated stacked 3-balls have a dual graph without degree four vertices (Theorem 22). These results are complemented by additional experimental data for $n \leqslant 14$ vertices (Table 1).

Figure 1: Map of the flip graph $\mathcal{P}_n$ of n-vertex 2-spheres. For notations see Sections 2, 3 and 4. The containment of $\mathcal{S}_n^0 \subset \mathcal{H}_n$ is due to a simple but unpublished argument which is left as an exercise to the interested reader. A *Klee polytope* is a triangulated 2-sphere S' obtained from a triangulated 2-sphere S by stellarly subdividing each triangle of S once. We refer to Section 2.4 for the notation used in this figure.

Altogether, the results contained in this paper together with existing results on the flip graph discussed above allow us to draw a relatively precise map of the flip graph of n-vertex triangulated 2-spheres. Having more knowledge about the structural properties of the flip graph might be one key for challenging future endeavours such as sampling triangulated 2-spheres or even generating triangulated 2-spheres with certain properties under some conditions of randomness.

For a graphical summary of what is known about the flip graph at present see Figure 1.

2 Preliminaries

2.1 Triangulations of 2-spheres

A *triangulation of the 2-sphere*, sometimes also referred to as a *planar triangulation*, is an n-vertex graph embedded in the 2-sphere with $3n - 6$ edges for some $n \geqslant 4$. As a direct result, the embedding decomposes the 2-sphere into $2n - 4$ triangles. This graph together with the triangles is called a *triangulated 2-sphere*. The graph is also called the *edge graph* of the triangulated 2-sphere. The simplest example of a triangulated 2-sphere is the boundary of the tetrahedron, called the *standard 2-sphere*.

Every n-vertex triangulated 2-sphere can be identified with an *abstract simplicial complex*, that is, a set of subsets of a finite ground set V, called *faces*, closed under taking subsets. For this, label its vertices with the elements of $V = \{1, \dots, n\}$ and represent triangles, edges and vertices by subsets of V of cardinality three, two and one respectively. Note that, for the purpose of this article, we sometimes do not make the distinction between vertices of an abstract simplicial complex and elements of its ground set.

We say that two triangulated 2-spheres are *combinatorially isomorphic*, or just *isomorphic* for short, if their respective abstract simplicial complexes are equal possibly after relabelling the elements of the ground set. In this article, whenever we talk about triangulated 2-spheres we mean their corresponding isomorphism classes of abstract simplicial complexes. By a theorem of Steinitz [19], isomorphism types of triangulated 2-spheres are in one-to-one correspondence with isomorphism types of simplicial 3-polytopes. This fact does not generalise to higher dimensions [2, 11].

Given a triangulated 2-sphere S, we usually denote its set of vertices, edges and triangles by $V(S)$, $E(S)$ and $F(S)$ respectively. Analogous notation is used for arbitrary abstract simplicial complexes. For $v \in V(S)$, its *star* $\mathrm{st}_S(v)$ is the simplicial complex generated by all triangles in $F(S)$ containing v. The edges and vertices of $\mathrm{st}_S(v)$ not containing v (i.e., the boundary of $\mathrm{st}_S(v)$) constitute the *link* of v in S, denoted by $\mathrm{lk}_S(v)$. The star and the link of an arbitrary face of an arbitrary abstract simplicial complex are defined analogously. The number of edges containing v is called the *degree* of v, denoted by $\deg_S(v)$.

For a triangulated 2-sphere S on ground set V and $W \subseteq V$, the *subcomplex induced by W*, denoted $S[W]$, is the simplicial complex of all triangles, edges and vertices of S entirely contained in W. Induced subcomplexes on arbitrary abstract simplicial complexes are defined analogously. In the special case of a graph $G = (V, E)$ and one of its vertices

$v \in V$, the induced subgraph $G[V \setminus \{v\}]$ is referred to as the *vertex-deleted subgraph $G - v$*.

2.2 Flag and Hamiltonian 2-spheres

There are several special types of triangulated 2-spheres which are relevant for this article. The most important ones are introduced in this section and in Section 2.3.

Definition 1 (Flag **2**-sphere). A *flag 2-sphere* is a triangulated 2-sphere in which all minimal non-faces of the underlying simplicial complex are of size two. Equivalently, a flag 2-sphere is a triangulated 2-sphere distinct from the standard 2-sphere, in which every 3-cycle (i.e., cycle of three edges) bounds a triangle.

Every triangulated 2-sphere S can be decomposed into a collection of flag 2-spheres and standard 2-spheres: Simply cut along a 3-cycle not bounding a triangle, and fill in the missing triangle in both parts. Iterating this procedure results in a set of spheres called the primitive components of S. Identifying each one of them by a node, and the 3-cycles by arcs between nodes this defines a tree. If the tree is a single vertex, S is called *primitive*. A triangulated 2-sphere is called *4-connected* if its edge graph is 4-connected. A triangulated 2-sphere distinct from the standard 2-sphere is 4-connected if and only if it is primitive if and only if it is flag.

Definition 2 (Hamiltonian **2**-sphere). A *Hamiltonian 2-sphere* is a triangulated 2-sphere containing a Hamiltonian cycle in its edge graph.

Hamiltonian 2-spheres play an important role in the proofs of upper bounds for the diameter of the flip graph of n-vertex triangulated 2-spheres for a fixed n, see [5] for an overview. This is due to (i) the well-behaved structure of the flip graph of n-vertex Hamiltonian 2-spheres which admits relatively precise bounds on its diameter, see Theorem 5, and (ii) the fact that a flag 2-sphere is necessarily Hamiltonian [22]. The converse of (ii) is not true.

2.3 Stacked 3-balls and stacked 2-spheres

A *triangulated 3-ball* is a collection of tetrahedra (together with their faces) whose union is a topological 3-ball. If B is a triangulated 3-ball then its *boundary ∂B* is the complex generated by all triangles of B contained in only one tetrahedron of B. By the *standard 3-ball* we mean a single tetrahedron together with its faces. The boundary of the standard 3-ball is the standard 2-sphere.

A triangulated 3-ball B is called a *stacked 3-ball* if there is a sequence $B_1, \ldots, B_m$ of triangulated 3-balls such that B_1 is the standard 3-ball, $B_m = B$ and, for $2 \leqslant i \leqslant m$, B_i is constructed from B_{i-1} by gluing (or stacking) a standard 3-ball onto a single triangle of B_{i-1}. Note that, by construction, all edges and vertices of B are contained in ∂B.

Conversely, let B be a triangulated 3-ball with all of its edges and vertices in ∂B. If t is an interior triangle in B then the boundary of t is a 3-cycle in ∂B (i.e., an induced 3-cycle in ∂B). Since B is a union of tetrahedra, B is the union of two smaller 3-balls B_1 and B_2 glued together along t and all the edges and vertices of B_i are in ∂B_i for $i = 1, 2$.

Inductively, this shows that B is a stacked 3-ball. (See [7, Theorem 4.5] for a more general result with a rigorous proof.) A *stacked 2-sphere* is a triangulated 2-sphere isomorphic to the boundary of a stacked 3-ball. It follows from the definition of a stacked ball that an n-vertex stacked 2-sphere contains exactly $n - 4$ induced 3-cycles.

For an abstract simplicial complex C whose faces consist of tetrahedra and their subfaces, the graph whose nodes correspond to the tetrahedra of C and two nodes are connected by an arc if and only if their corresponding tetrahedra share a triangle is called the *dual graph* of C, denoted by $\Lambda(C)$. If B is a stacked 3-ball then $\Lambda(B)$ is a tree, and every node of $\Lambda(B)$ corresponds to a primitive component of the bounding stacked 2-sphere ∂B. It follows that a triangulated 2-sphere is stacked if and only if all of its primitive components are standard 2-spheres.

From [1, Lemma 4.6 and Remark 4.1] we know the following statement.

Lemma 3 (Bagchi, Datta [1]). *Let S be a stacked 2-sphere with edge graph G. Let $\overline{S}$ denote the simplicial complex whose faces are all the cliques of G. Then $\overline{S}$ is a stacked 3-ball and $S = \partial \overline{S}$. Moreover, up to isomorphism, $\overline{S}$ is the unique stacked 3-ball such that $S = \partial \overline{S}$.*

2.4   Bistellar moves

Bistellar moves are local combinatorial alterations of a simplicial complex which, in general, change the isomorphism type of the complex, but not the topology of the underlying space. For a triangulated 2-sphere S there are the following two bistellar moves to consider (see also Figure 2).

- Replace a triangle of S by three triangles joined around a new vertex. Such a stellar subdivision of a triangle is also called a *0-move* (because a 0-dimensional face is inserted) or *1-3-move* (because one triangle is replaced by three triangles). For its inverse operation, a so-called *2-move* (a 2-dimensional face is inserted) or *3-1-move* (three triangles are replaced by one), remove the vertex star of a vertex of degree three and replace it by a single triangle. This inverse operation is only possible if the new triangle is not already present in the triangulation. In particular, the standard 2-sphere does not allow any 2-moves.

- Replace two triangles of S which are joined along a common edge, say abx and aby, and replace them with triangles axy, bxy. This operation is possible if and only if xy is not an edge of S. This move is called a *1-move*, *2-2-move*, or, for obvious reasons, an *edge flip*. Throughout this article we denote it by $ab \mapsto xy$. The inverse of an edge flip is again an edge flip.

Definition 4. The *Pachner graph* $\mathcal{P}$ of triangulated 2-spheres is the graph whose nodes are triangulated 2-spheres up to combinatorial isomorphism, with arcs between all pairs of triangulated 2-spheres that can be transformed into isomorphic copies of each other by a single bistellar move.

Figure 2: The bistellar moves in dimension two.

The *flip graph* $\mathcal{P}_n$ is the induced subgraph of $\mathcal{P}$ whose nodes are triangulated 2-spheres with precisely n vertices.

Note that it is a fundamental and well-known fact that the Pachner graph $\mathcal{P}$ of triangulated 2-spheres is connected (see for example [17] for a much more general statement due to Pachner). Also note that all arcs in the flip graph $\mathcal{P}_n$ correspond to edge flips.

The flip graph of n-vertex flag 2-spheres is denoted by $\mathcal{F}_n$, the flip graph of n-vertex Hamiltonian 2-spheres by $\mathcal{H}_n$, and the flip graph of n-vertex stacked 2-spheres by $\mathcal{S}_n$. Note that, naturally, all of these graphs are induced subgraphs in the flip graph $\mathcal{P}_n$ of n-vertex 2-spheres. In particular, a priori it is not clear, whether or not any of them is connected. The following statement is due to work by Mori, Nakamoto and Ota.

Theorem 5 ([12, Theorem 5] and [16, Theorem 1]). *For $n \geqslant 5$, the flip graph $\mathcal{H}_n$ is connected and has diameter at least $2n - 15$ and at most $4n - 20$.*

In this article, we focus on structural properties of $\mathcal{F}_n$ and $\mathcal{S}_n$.

3 The flip graph $\mathcal{F}_n$ of n-vertex flag 2-spheres

In this section we prove that, for $n \geqslant 8$, the flip graph $\mathcal{F}_n$ of n-vertex flag 2-spheres contains exactly two components, one of them consisting of the double cone Γ_n, the other one containing all other n-vertex flag 2-spheres. If S and S' are two n-vertex flag 2-spheres, we write $S \sim S'$ to mean that there exists a sequence of edge flips connecting S and S' preserving the flagness property at each step. We prove the following statement.

Theorem 6. *If S and S' are two n-vertex flag 2-spheres distinct from Γ_n, then $S \sim S'$.*

See Figures 3 to 5 for illustrations of the flip graph $\mathcal{F}_n$ for $n \in \{8, 9, 10\}$.

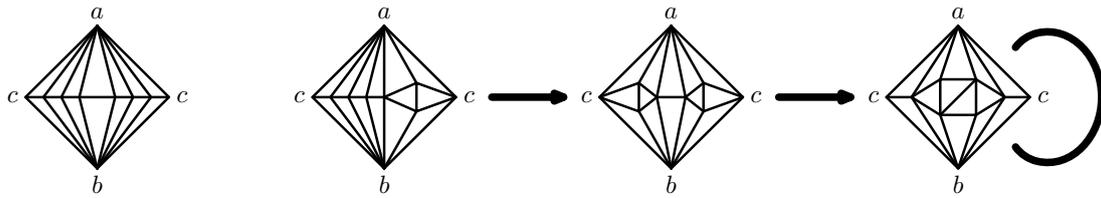Figure 3: The flip graph $\mathcal{F}_8$ of 8-vertex flag 2-spheres.
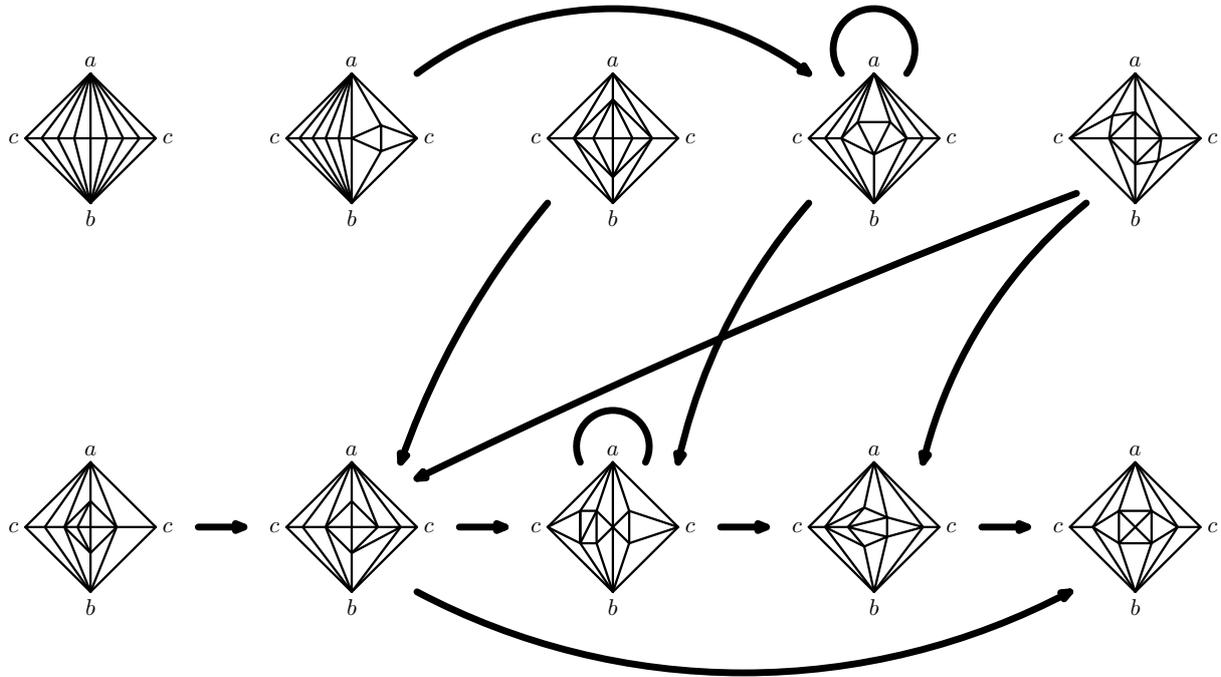
Figure 4: The flip graph $\mathcal{F}_9$.

Figure 5: The flip graph $\mathcal{F}_{10}$.

The proof of Theorem 6 relies on a number of lengthy and technical lemmas (Lemmas 9 to 13). We thus start by introducing all necessary terminology and a sketch of the proof, before proving all lemmas in detail.

Definition 7. Let S be a flag 2-sphere. A subcomplex Q of S is called a *quadrilateral* if it is a triangulated disc and its boundary is a 4-cycle. A quadrilateral Q in S with boundary *a-b-c-d-a* is called *proper*, if *a-b-c-d-a* is an induced cycle in S and $\deg_S(a), \deg_S(b), \deg_S(c), \deg_S(d) \geqslant 5$. Since the boundary is an induced cycle, a proper quadrilateral contains at least one interior vertex. A quadrilateral Q in S is called *ordered*, if it contains an interior vertex, and all of its interior vertices are of degree four. A path in an ordered quadrilateral Q joining two antipodal vertices on the boundary that are contained in only two triangles of Q, and going through all interior vertices but not any other vertex on the boundary, is called a *diagonal path*, or just a *diagonal of Q*. If Q is ordered and has more than one interior vertex, then its diagonal is unique.

Definition 8. For $n \geqslant 7$, let A_n in $\mathcal{F}_n$ be as in Figure 6(b). Note that $A_7 = \Gamma_7$, $A_n \neq \Gamma_n$ for $n \geqslant 8$ (see Figure 3), and that A_n is a vertex of degree one in $\mathcal{F}_n$ for $n \geqslant 9$ (see Figures 4 and 5). To see the latter claim, note that there are precisely four edges in A_n, $n \geqslant 8$, with endpoints of degree greater than four. All four edges are symmetric and hence flipping any such edge results in the same isomorphism type of flag 2-sphere triangulation.

For $k \geqslant 3$, let $\mathcal{Q}_k$ be the triangulated quadrilateral with k interior vertices shown in Figure 6(c). The path a_0-a_1-$\cdots$-a_k is said to be the *diagonal path* of $\mathcal{Q}_k$.
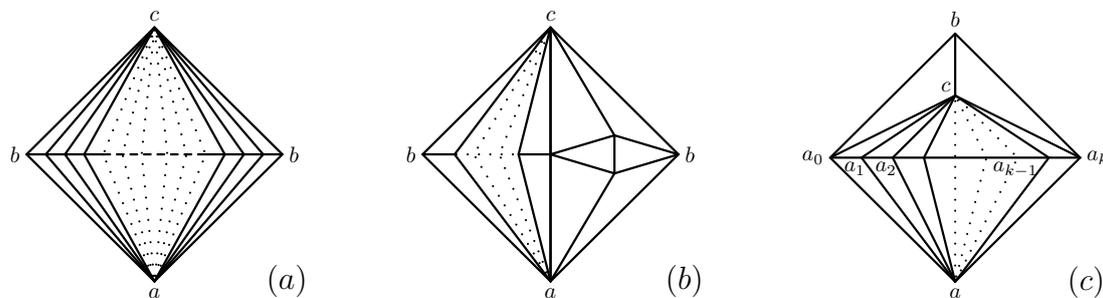

Figure 6: (a) Double cone Γ_n over the $(n-2)$-gon. (b) Target n-vertex flag 2-sphere A_n. (c) Quadrilateral $\mathcal{Q}_k$ with boundary vertices a_0, a, a_k, b and interior vertices $c, a_1, \ldots, a_{k-1}$.

We prove Theorem 6 by showing that $S \sim A_n$, for any n-vertex flag 2-sphere S distinct from the double cone. For this, we split S (or a slight variation thereof) along an induced 4-cycle into two triangulated quadrilaterals Q and R using Lemma 12. We then use Lemma 13 to turn both Q and R into ordered quadrilaterals. Finally, we use Lemma 9 to transport excess internal vertices from R to Q (or vice versa), until we obtain A_n.

The main difficulty in the above procedure is to prove Lemma 13. For this we need Lemma 11, which allows us to merge two smaller triangulated quadrilaterals, and Lemma 10, which allows us to resolve a pathological class of triangulations of the quadrilateral (triangulation $\mathcal{Q}_k$, shown in Figure 6(c)). In addition, all of Lemmas 10, 11 and 13 need Lemma 9 to transport internal vertices from one quadrilateral to another.

For a more precise but less descriptive outline, see the proof of Theorem 6 at the end of this section.

Lemma 9 (Transport Lemma). *Let S be a flag 2-sphere containing two ordered quadrilaterals α and β with disjoint interiors, but a common boundary edge vw. Furthermore, let $k \geqslant 2$ ($\ell \geqslant 1$) be the number of interior vertices of α (resp., β), and let v and w satisfy one of the following conditions:*

(1) $\deg_S(w) \geqslant 5$, *and the diagonal paths of α and β intersect in w;*

(2) $\deg_S(v) \geqslant 5$, $\deg_S(w) \geqslant 6$, *the diagonal path of α intersects v, and the diagonal path of β intersects w.*

Then there exists a flag 2-sphere S' such that (i) $S \sim S'$, (ii) S' contains two ordered quadrilaterals α' and β', (iii) $S' = (S \setminus \{\alpha, \beta\}) \cup \{\alpha', \beta'\}$, (iv) vw is a common edge of α' and β' in S', and (v) the number of interior vertices of α' is $k-1$, and the number of interior vertices of β' is $\ell + 1$.

Lemma 9 gives precise conditions on when exactly we can "transport" an interior vertex of an ordered quadrilateral of S into an adjacent ordered quadrilateral without changing anything else in S. Both Condition *(1)* and *(2)* for Lemma 9 are satisfied as soon as α and β only share one edge. If α and β share two edges, the situation is different: In Condition *(1)* we can then have $\deg_S(w) = 4$ if both the common edges of α and β are through w, in Condition *(2)* and for $k = 2$ and $\ell = 1$ we can have both $\deg_S(v) = 4$ and $\deg_S(w) = 5$.

Proof. Each ordered quadrilateral of S must be subdivided by a diagonal path containing all of its interior vertices all of which are of degree four. Hence, up to exchanging the roles of v and w, there are two possible initial configurations to consider: The diagonal paths of α and β either meet, or one ends in v and the other in w. The former corresponds to Condition *(1)* of the Lemma, the latter one to Condition *(2)*.

Condition *(1)* The diagonal paths of α and β meet in w. In this case, the sequence of flips transforming S to S' is shown in Figure 7(a) (top to bottom). The dotted edge denotes the edge to be flipped next, the dashed line denotes the newly inserted edge. The integer next to a vertex indicates the change of the respective vertex degree with respect to the initial vertex degree.

Throughout this edge flip sequence the degrees of w, v and the upper left vertex of α are, at some point, decreased to the initial degree minus one. The degrees of all other boundary vertices are never decreased below the initial degree. Since all three vertices of the former group are initially of degree at least five (w by assumption and the other two by the flagness of S), the flagness condition is preserved in each step. The preconditions of the lemma ensure that no 3-cycle is introduced in the first flip, the edges introduced by flip two and three end in the interior of $\alpha \cup \beta$ and hence cannot introduce a new 3-cycle, and the last flip re-introduces the edge removed by the first flip.

Condition *(2)* α and β have diagonal paths ending in v and w respectively. To comply with the labelling of the statement of the lemma, let the diagonal of α contain v and the diagonal of β contain w. The sequence of edge flips transforming S to S' in this case is shown in Figure 7(b) (top to bottom). The meaning of dotted and dashed lines as well as integers next to vertices is the same as in Condition *(1)*.

Note that, in this procedure, only the degree of w is, at one stage, decreased to the initial degree minus two. In addition, v and the lower left vertex of α are, at some point, decreased to the initial degree minus one. The degrees of all other boundary vertices are never decreased below the initial degree. By assumption, w is of initial degree at least six and v is of initial degree five. Again, the other vertex of α not containing the diagonal must be of initial degree at least five by the flagness of S. It follows that the flagness condition is preserved in each step. Again, no 3-cycle is introduced by the flip sequence for reasons analogous to the ones described in the previous case. $\square$
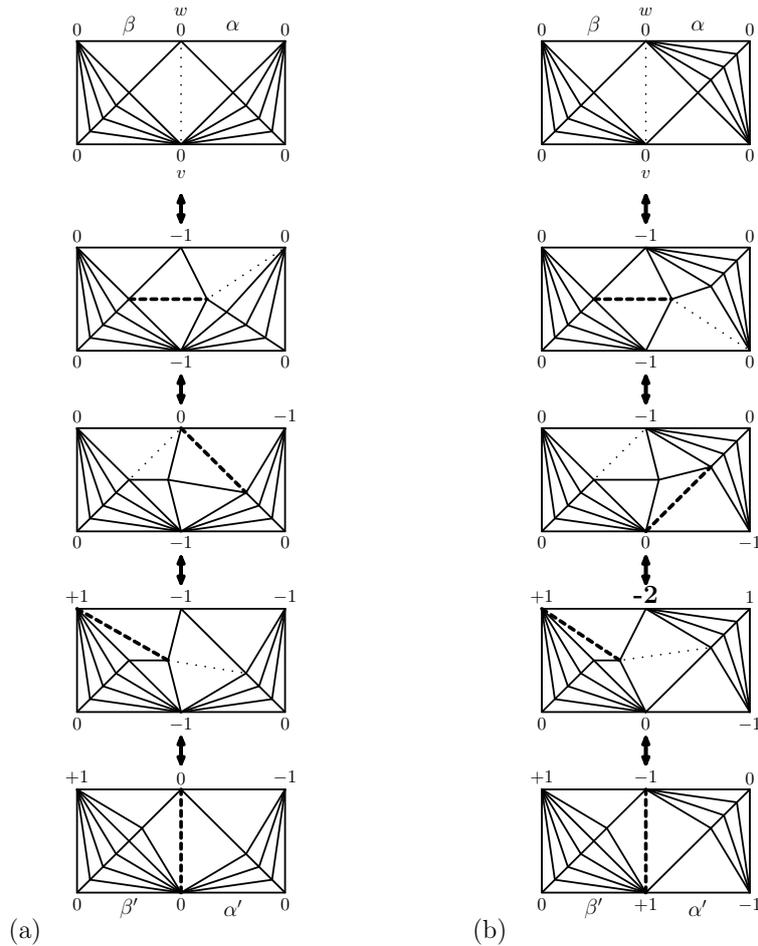
Figure 7: Transport Lemma. (a) sequence of edge flips for intersecting diagonal paths (Cond. *(1)*). (b) sequence of edge flips for diagonal paths ending in v and w (Cond. *(2)*).

Lemma 10. *Let S be an n-vertex flag 2-sphere, $n \geqslant 8$, with induced 4-cycle a-a_0-b-a_k-a bounding $\mathcal{Q}_k$. Then either $S = \Gamma_n$, $S \sim A_n$, or there exists an n-vertex flag 2-sphere S' with $S \sim S'$, such that (i) a-a_0-b-a_k-a is an induced 4-cycle in S' bounding an ordered quadrilateral Q, and (ii) $S \setminus \mathcal{Q}_k = S' \setminus Q$.*

Proof. We use the notation for $\mathcal{Q}_k$ as introduced in Figure 6(c) and in accordance with the vertex labels of the induced 4-cycle a-a_0-b-a_k-a bounding $\mathcal{Q}_k$.

Case $k = 3$: Refer to Figure 8(a). Consider the two triangles $a_0ax_1, a_3ax_1' \in F(S)$ outside but adjacent to $\mathcal{Q}_3$. If $x_1 = x_1'$ (i.e., $\deg_S(a) = 5$) consider triangles $a_0x_ix_{i+1}, a_3x_ix_{i+1}' \in F(S)$, $i \geqslant 1$, until either $x_{\ell+1} \neq x_{\ell+1}'$, that is, $\deg_S(x_\ell) \geqslant 5$, or $x_\ell' = x_\ell = b$.

The case $x_1' = x_1 = b$ is not possible because a-a_0-b-a_k-a is induced (and because $n \geqslant 8$). If $x_\ell' = x_\ell = b$, $\ell \geqslant 2$, S must be isomorphic to $A_{\ell+6}$ and we are done. Otherwise, consider the two triangles $a_0x_\ell x_{\ell+1}$ and $a_3x_\ell x_{\ell+1}'$, $x_{\ell+1}' \neq x_{\ell+1}$. Neither $a_0x_{\ell+1}'$ nor $a_3x_{\ell+1}$

can be edges of S since otherwise there are induced 3-cycles a_0-$x'_{\ell+1}$-x_ℓ-a_0 or a_3-$x_{\ell+1}$-x_ℓ-a_3.

Keeping these observations in mind, we perform edge flip $a_0 x_\ell \mapsto x_{\ell+1} x_{\ell-1}$ (see Figure 8(b)), followed by edge flips $a_0 x_{\ell-1} \mapsto x_{\ell+1} x_{\ell-2}$, etc. all the way down to $a_0 a \mapsto x_{\ell+1} a_1$ (see Figure 8(c)). For each of them we have that, since $a_3 x_{\ell+1}$ is not an edge, a_3-x_i-$x_{\ell+1}$-a_3 is not a 3-cycle of S.

It follows that we can perform flips $a_1 c \mapsto a_0 a_2$ (Figure 8(d)) and $a a_2 \mapsto a_1 a_3$ (Figure 8(e)), followed by the initial sequence of edge flips in reverse, i.e., $x_{\ell+1} a_1 \mapsto a_0 a$, $x_{\ell+1} a \mapsto a_0 x_1$, $x_{\ell+1} x_1 \mapsto a_0 x_2$, all the way up to $x_{\ell+1} x_{\ell-1} \mapsto a_0 x_\ell$ (Figure 8(f)). Observe that now all vertices inside $\mathcal{Q}_3$ are of degree four and outside $\mathcal{Q}_3$ the triangulation is unchanged. This proves the result for $k = 3$.
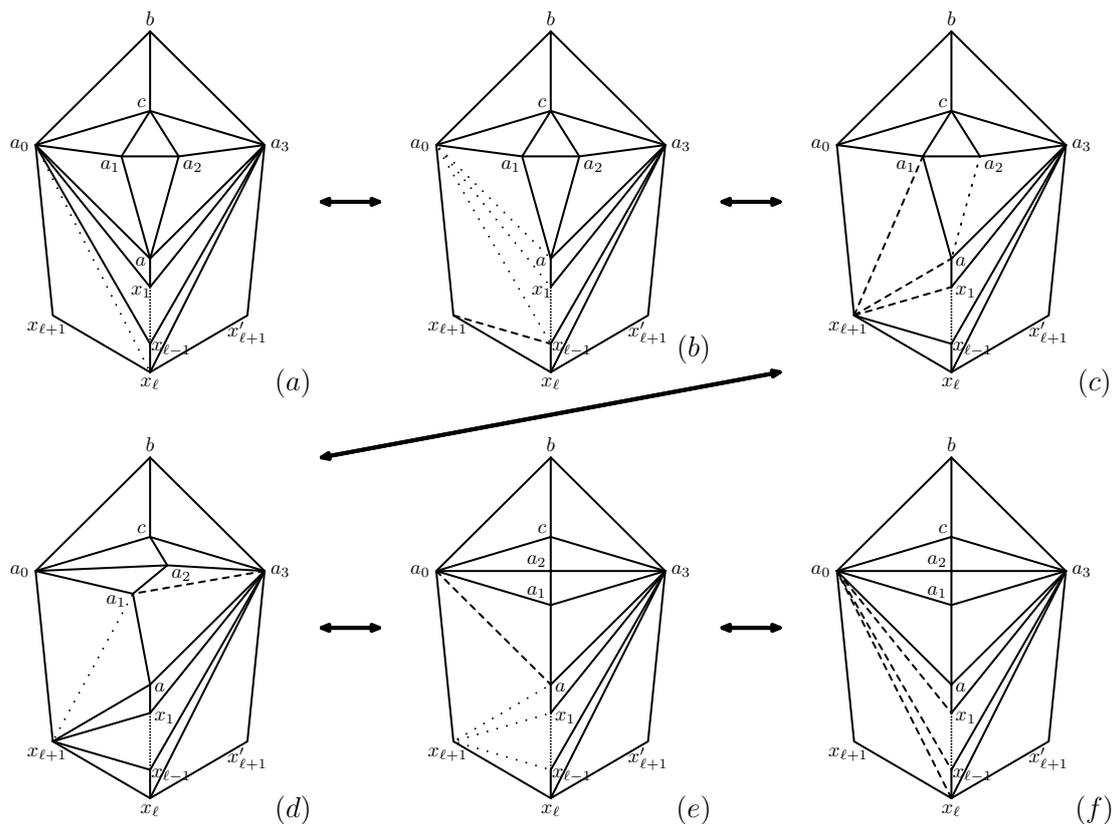

Figure 8: Resolving $\mathcal{Q}_3$ into a quadrilateral with three interior vertices of degree four.

Case $k = 4$: Refer to Figure 9(a). The case $k = 4$ is very similar to the case $k = 3$. Again, the case $x'_1 = x_1 = b$ is not possible because a-a_0-b-a_k-a is induced. If $x'_\ell = x_\ell = b$ for $\ell \geqslant 2$, S decomposes into two ordered proper quadrilaterals along induced 4-cycle a_0-a-a_4-c-a_0 to which we can apply Lemma 9: The ordered proper quadrilateral contained in $\mathcal{Q}_4$, the rest of S, a and a_0 take the roles of α, β, w and v. The diagonals are disjoint, $\deg_S(a) = 6$ and $\deg_S(a_0) \geqslant 5$. In particular, Condition (2) is satisfied with $k = 3$ and $\ell \geqslant 1$ and we can transport a_1 or a_3 away from its quadrilateral to conclude that $S \sim A_n$, $n \geqslant 8$.
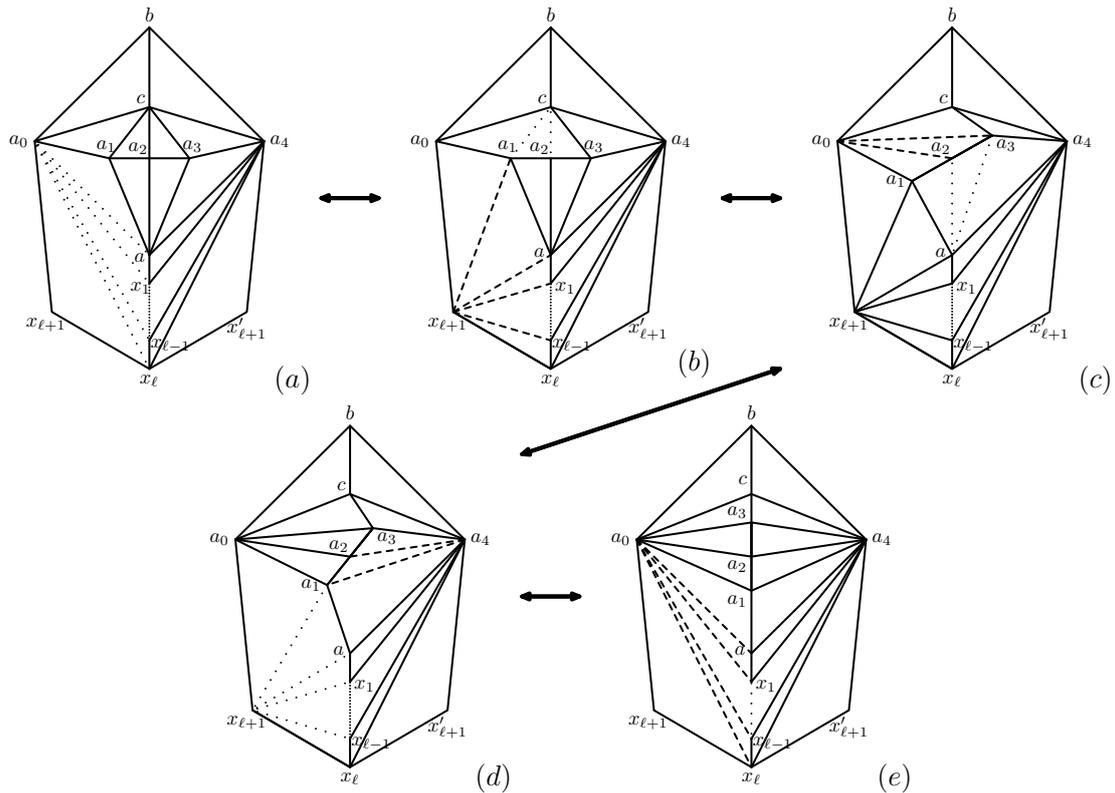
Figure 9: Resolving $\mathcal{Q}_4$ into a quadrilateral with four interior vertices of degree four.

If $x_{\ell+1} \neq x'_{\ell+1}$ for some $\ell \geqslant 0$ we perform a sequence of edge flips similar to the one in the case $k = 3$ above. More precisely, the initial set of flips (Figure 9(b)) and the final set of flips (Figure 9(e)) are identical with the initial two and the final two steps of case $k = 3$. Once flip $a_0 a \mapsto x_{\ell+1} a_1$ is performed, we can perform $a_1 c \mapsto a_0 a_2$ and $a_2 c \mapsto a_0 a_3$ (Figure 9(c)), followed by $a_3 a \mapsto a_2 a_4$ and $a_2 a \mapsto a_1 a_4$ (Figure 9(d)).

Case $k > 4$: Refer to Figure 10(a). From $k > 4$ it follows that $n \geqslant 10$. Moreover, $a_0 \neq a_k$, and $a_0 a_k$ is a non-edge of S since a_0-a-a_k-b-a_0 is an induced 4-cycle. We start by performing flips $a_0 c \mapsto b a_1$, $a_1 c \mapsto b a_2$, all the way to $a_{k-4} c \mapsto b a_{k-3}$ (see Figure 10(b)). The resulting quadrilateral splits into two parts. One with only degree four interior vertices (at least one), the other one being isomorphic to $\mathcal{Q}_3$ with diagonal path going from a_{k-3} to a_k (see Figure 10(c) for a re-arranged version of the top centre quadrilateral emphasising this fact).

Use the case $k = 3$ to turn $\mathcal{Q}_3$ into a quadrilateral containing only interior vertices of degree four with the diagonal path running from a to b (see Figure 10(d)). Since $k > 4$, the overall quadrilateral again splits into two parts, one with only degree four interior vertices (possibly none), the other one being isomorphic to $\mathcal{Q}_4$ with diagonal path going from a to b (see Figure 10(e) for a re-arranged version of the bottom left quadrilateral emphasising this fact). Use the case $k = 4$ to either conclude that $S \sim A_n$, or to turn $\mathcal{Q}_4$ into a quadrilateral containing only degree four interior vertices and diagonal running
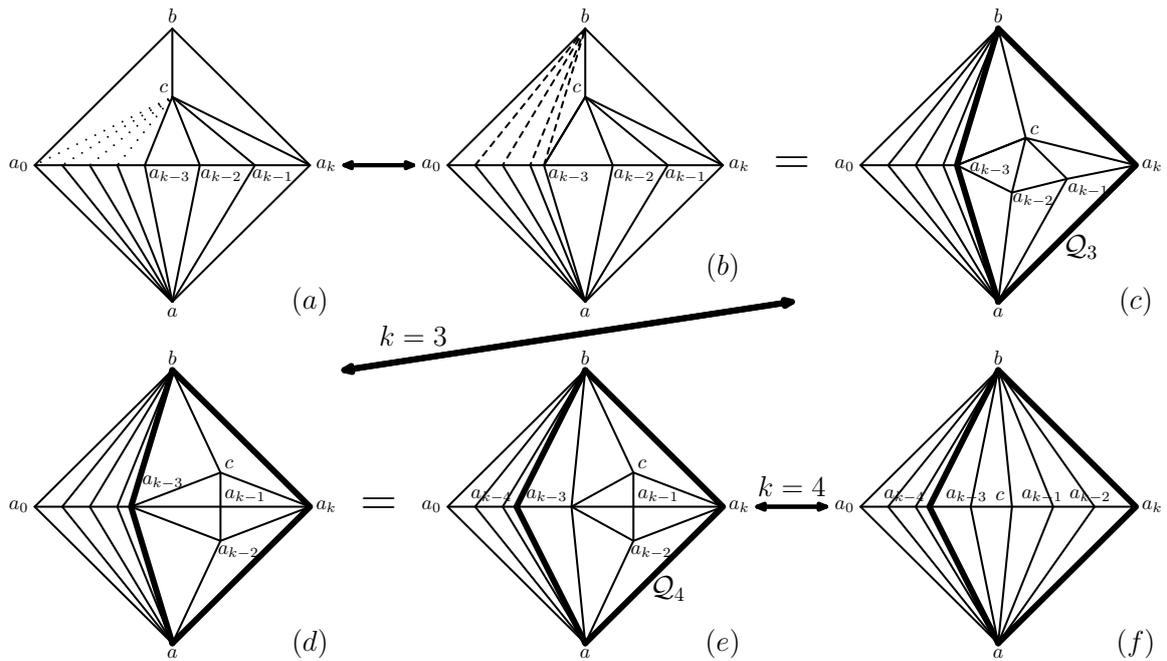
Figure 10: Resolving $\mathcal{Q}_k$, $k > 4$, into a quadrilateral with k interior vertices of degree four.

from a_{k-4} to a_k. In the latter case the overall quadrilateral now only has interior vertices of degree four which proves the lemma (see Figure 10(f)). □

Lemma 11 (Merge Lemma). *Let S be an n-vertex flag 2-sphere containing two ordered quadrilaterals α and β with disjoint interiors, but common outer edges uv and uw. Then either $S = \Gamma_n$, $S \sim A_n$, or $S \sim S'$ where S' has an ordered quadrilateral γ with boundary $\partial(\alpha \cup \beta)$ and $S' = (S \setminus \{\alpha, \beta\}) \cup \{\gamma\}$.*

Proof. We have four cases for the initial configuration of α and β emerging from the different possible relative orientations of the diagonal paths of α and β, see Figure 11. Throughout this proof, vertices a and b denote the only boundary vertices of α and β that, possibly, are not contained in both α and β.

Case 1: If $a = b$ then $\alpha \cup \beta = S$. In this case, $S = \Gamma_n$ with cone apices v and w and we are done.
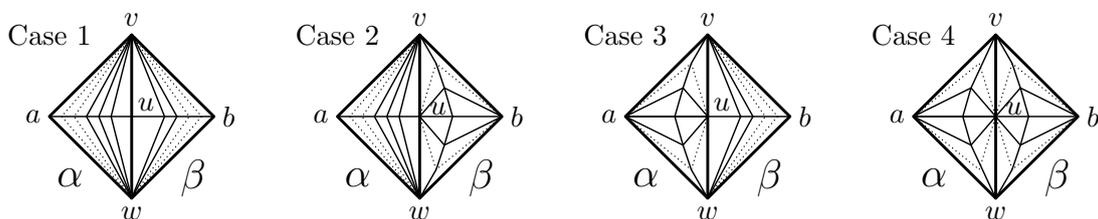

Figure 11: The four initial configurations for α and β in the proof of Lemma 11.

If $a \neq b$ we can merge α and β into one larger ordered quadrilateral with boundary $\partial(\alpha \cup \beta)$.
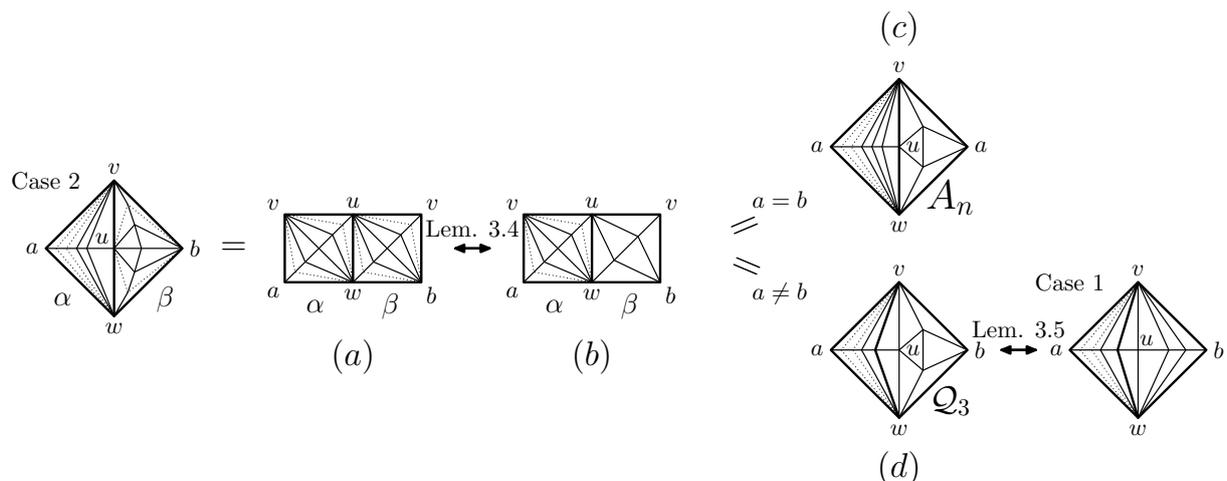

Figure 12: Transporting vertices in Case 2: (a) Case 2 redrawn after cutting along edge uv. (b) After transporting interior vertices away from β (Lemma 9). (c) Case $a = b$ yields A_n. (d) Case $a \neq b$ yields $\mathcal{Q}_3$. In the latter case apply Lemma 10 to fall back to Case 1.

Case 2: Refer to Figure 12. As before, if $a = b$ then $\alpha \cup \beta = S$. If, in this case, β contains only one interior vertex, then we have $S = \Gamma_n$ with cone apices v and w and we are done. If α contains only one interior vertex, both v and w are of degree four, and we have $S = \Gamma_n$ with cone apices u and $a = b$.

Thus, we can assume both α and β have at least two interior vertices. In this case, we iteratively apply Lemma 9 to transport interior vertices from β to α across edge uw until u is of degree five (Figure 12(b)) and we obtain A_n (Figure 12(c)).

If $a \neq b$, we, again, apply Lemma 9 to transport interior vertices from β to α across edge uw until u is of degree five (Figure 12(b)). The quadrilateral β together with the two rightmost triangles of α now form a quadrilateral isomorphic to $\mathcal{Q}_3$ with diagonal path from v to w (see Figure 12(d)). This can be resolved into a quadrilateral with interior vertices all of degree four and diagonal intersecting b (note that a is of degree greater than four and thus (i) the preconditions of Lemma 10 are satisfied and (ii) we can always resolve $\mathcal{Q}_3$ in this case) and we are back to Case 1.

Case 3: This is completely analogous to Case 2.

Case 4: Again, if $a = b$ then $\alpha \cup \beta = S$, and S is equal to Γ_n with cone apices u and $a = b$.

Hence, let $a \neq b$. If α contains only a single interior vertex we fall back to Case 2, if β contains only a single vertex we fall back to Case 3. Thus we can assume both α and β have at least two interior vertices. In this case, $\deg_S(u) \geqslant 6$, $\deg_S(v), \deg_S(w) \geqslant 5$, and we apply Lemma 9 to transport vertices from α to β until α contains only a single interior vertex. Then we proceed with Case 2. □

Lemma 12. *For $n \geqslant 8$, let $S \in \mathcal{F}_n \setminus \{\Gamma_n\}$. Then there exists $S' \in \mathcal{F}_n \setminus \{\Gamma_n\}$ with $S \sim S'$, and $a, b, c, d \in V(S')$ such that (i) a-b-c-d-a is an induced 4-cycle in S', and (ii) $\deg_{S'}(a)$, $\deg_{S'}(b)$, $\deg_{S'}(c)$, $\deg_{S'}(d) \geqslant 5$. In particular, S' splits into two proper quadrilaterals both bounded by a-b-c-d-a.*

Proof. If S contains a vertex v of degree four, then, by the flagness of S, the link of v is an induced 4-cycle, say a-b-c-d-a. If any of these vertices, say a, is of degree four, then, since $n \geqslant 8$, the boundary of the union of the stars v and a is an induced 4-cycle. Moreover, b and d are of degree at least five. Iterating this process either yields an induced 4-cycle x-b-c-d-x, for some vertex x of S of degree at least five, or $x = c$, and S is isomorphic to Γ_n, a contradiction. Hence, assume $\deg_S(x) \geqslant 5$, and thus $x \neq c$. If the degree of c is 4, consider the union of the quadrilateral containing v and bounded by x-b-c-d-x and the star of vertex c. As before, iterate this procedure until we obtain an induced 4-cycle x-b-y-d-x in S (possibly $y = c$) with x and y necessarily distinct and both of degree at least five (note that $x = y$ implies S isomorphic to Γ_n and thus $\deg_S(x) = 4$, a contradiction).

Since S is flag, it cannot contain a vertex of degree three. If, in addition, S does not contain a vertex of degree four, then S must contain a vertex w of degree five (this is a consequence of Euler's formula which implies that the average vertex degree of a triangulated 2-sphere must be less than six). Let auw and buw be two adjacent triangles in the star of w. If a and b have a common neighbour x distinct from w and u, then x-a-w-b-x is an induced 4-cycle, and we are done since S has no vertex of degree four. Otherwise the flip $uw \mapsto ab$ yields a flag 2-sphere in which w has degree four. Now the link of w is an induced 4-cycle with all four vertices being of degree at least five. $\square$

Lemma 13. *Let S be an n-vertex flag 2-sphere which splits into two proper quadrilaterals Q and R along an induced 4-cycle a-c-b-d-a. Then there exists an n-vertex flag 2-sphere S' with $S \sim S'$, such that $S' = Q' \cup R$, and the interior of Q' contains only degree four vertices.*

Note that, in S', neither Q' nor R need to be proper quadrilaterals. However, both Q' and R contain interior vertices. In particular, each of a, b, c, and d is contained in at least two triangles of both Q' and R. We deal with this issue separately whenever we need to, namely in the proof of Theorem 6.

Proof. We prove this statement by induction on the number k of interior vertices in Q. First note that $k > 0$, and that the statement is true for $k \leqslant 2$.

Let a-c-b-d-a be the boundary of a quadrilateral Q in S with $k \geqslant 3$ interior vertices, such that $\deg_S(a), \deg_S(b), \deg_S(c), \deg_S(d) \geqslant 5$. Since a-c-b-d-a is induced, ab and cd cannot be edges of S.

Claim: There exist a triangulation S' with $S \sim S'$, such that $S' = Q' \cup R$, and in the interior of Q' either a and b or c and d have at least one common neighbour.

We first complete the proof of the lemma assuming the claim is true. This is then followed by a proof of the claim. We can thus assume that we have an n-vertex flag 2-sphere S', $S \sim S'$, such that either a and b or c and d have at least one common neighbour in Q'.

Assume that there exist at least one common neighbour of a and b (the case that c and d have at least one common neighbour is completely analogous). If all such neighbours are of degree four, all interior vertices must be neighbours of a and b of degree four and we are done. Otherwise, choose a common neighbour e of degree at least five, and split Q' into two smaller quadrilaterals Q_1 and Q_2 with boundaries e-a-c-b-e and e-a-d-b-e respectively. Without loss of generality, let Q_2 be the quadrilateral with at least three triangles containing e.

If Q_1 has interior vertices, use the induction hypothesis to obtain a 2-sphere S'', $S' \sim S''$, in which Q_1 is transformed into a quadrilateral Q_1' with boundary e-a-c-b-e, $S' \setminus Q_1 = S'' \setminus Q_1'$, and in which all interior vertices of Q_1' have degree four. In S'' vertex d is still of degree at least five, vertices a and b must be of degree at least six, and vertex e must be of degree at least five since at least three triangles containing e are outside Q_1'. In particular, Q_2 is proper and we can apply the induction hypothesis to Q_2 to obtain a triangulated 2-sphere S''' with two ordered quadrilaterals Q_1' and Q_2' joined along two adjacent edges. Use Lemma 11 to merge both quadrilaterals, or conclude that $S \sim A_n$. We have that $S \not\sim \Gamma_n$, since Γ_n does not split into to proper quadrilaterals, as required by the statement of Lemma 13.

Hence, without loss of generality let Q_1 be without interior vertices. Use the induction hypothesis to transform Q_2 into Q_2' with only degree four vertices inside. Now either e is of degree four, all interior vertices of $Q' = Q_1 \cup Q_2'$ are of degree four, and we are done. Or Q' is isomorphic to $\mathcal{Q}_k$ and, by Lemma 10, can be transformed into a quadrilateral containing only degree four vertices (or $S \sim A_n$), and again we are done.

Proof of the claim: Refer to Figure 13. In the following procedure we always denote the flag 2-sphere by S and the quadrilateral enclosed by a-c-b-d-a by Q, although both objects are altered in the process.

1. Denote all neighbours of a in Q from left to right by $c = a_0, a_1, \ldots, a_m = d$.

2. If a_0 and a_m have a common neighbour in Q other than a and b we are done.

3. If no such neighbour exists, let $1 \leqslant j \leqslant m-1$ be the largest index for which a_0 and a_j have common neighbours outside the star of a.

   There exist an outermost neighbour x_1 in Q, bounding a quadrilateral x_1-a_0-a-a_j-x_1 that contains all other common neighbours of a_0 and a_j. Note that, in this case, a_j must be of degree at least five. If $x_1 = b$, a and b have a common neighbour and we are done. If $x_1 \neq b$, then there is at least one triangle inside Q containing a_0 but not contained in the quadrilateral inside Q and bounded by x_1-a_0-a-a_j-x_1. In particular, a_0 is of degree at least five in S (although S might have changed during this proof).

4. If the quadrilateral inside Q and bounded by x_1-a_0-a-a_j-x_1 does not contain interior vertices, we must have $j = 1$ and the quadrilateral consists of the two triangles $a_0 a_1 a$ and $a_0 a_1 x_1$. Note that $x_1 \neq a_i$ by the flagness of the triangulation, and $x_1 a_i$ is a non-edge for $2 \leqslant i \leqslant m$ by construction of the procedure.

As explained in detail above, both a_0 and a_1 are of degree at least five, and a and x_1 do not have common neighbours other than a_0 and a_1. Hence, we can perform flip $a_0 a_1 \mapsto a x_1$ which strictly increases the degree of a inside Q. We then start over at step 1 with $a_0' = a_0, a_1' = x_1, a_2' = a_1, \ldots a_{m+1}' = a_m$.

5. If the quadrilateral inside Q and bounded by x_1-a_0-a-a_j-x_1, say Q_1, contains interior vertices, we have $\deg_S(x_1) \geqslant 5$. Moreover, as explained above $\deg_S(a_0), \deg_S(a_j) \geqslant 5$, and $\deg_S(a) \geqslant 5$ by assumption. In particular, Q_1 is a proper quadrilateral with fewer interior vertices than Q. We can thus use the induction hypothesis to rearrange the interior of Q_1 to contain only interior vertices of degree four. Note that, in the new triangulation, all of x_1, a_0, a and a_j still have degree at least five (i.e., the rearranged quadrilateral is an ordered proper quadrilateral). This is important later on in the proof.

6. After rearranging Q_1, bounded by x_1-a_0-a-a_j-x_1, into an ordered proper quadrilateral, repeat steps 3 to 5 by looking for the largest index $j < \ell \leqslant m - 1$ for which a_j and a_ℓ have common neighbours outside the star of a. Note that, whenever we flip an edge in step 4 we start over at step 1 with a strictly larger degree of vertex a in Q.

This process either yields the desired result, or it terminates with Q having a sequence of smaller ordered quadrilaterals $Q_1, \ldots, Q_p$ around vertex a, $p > 1$, see Figure 13.

Call the "peaks" of the quadrilaterals $x_1, \ldots, x_p$, and the "valleys" between quadrilaterals $a_0 = y_0, \ldots, y_p = a_m$ (cf. Figure 13). By construction, all x_i, $1 \leqslant i \leqslant p$, and y_j, $0 \leqslant j \leqslant p$ are of degree at least five (see step 5 above). That is, the quadrilaterals Q_i, $1 \leqslant i \leqslant p$, are ordered and proper.

Recall that all quadrilaterals Q_i, $1 \leqslant i \leqslant p$, contain only degree four interior vertices. We want all of the diagonal paths of Q_i, $1 \leqslant i \leqslant p$, to run from y_{i-1} to y_i. If Q_i only has one interior vertex, this is automatically the case. Thus, assume that there exist a pair of quadrilaterals Q_i and Q_{i+1}, $1 \leqslant i \leqslant p - 1$, sharing common edge $a y_i$, and, without loss of generality, assume that Q_i has a diagonal path from a to x_i of length at least two.

Observe that in this particular situation, both a and y_i must be of degree at least six. Hence we can apply Lemma 9 to "transport" all but one interior vertices of Q_i to the diagonal path of Q_{i+1}, and declare the diagonal path in Q_i to run from y_{i-1} to y_i. If the diagonal path of Q_{i+1} connects y_i with y_{i+1} we are done. If not, note that, again, both a and y_i must be of degree at least six. We proceed by transporting all but one interior vertices of Q_{i+1} onto the new diagonal path from y_{i-1} to y_i of Q_i, and declare the diagonal path in Q_{i+1} to run from y_i to y_{i+1}. Repeating this with all pairs of quadrilaterals containing at least one diagonal intersecting a yields the desired result. Note that this procedure terminates with the degree of a being at least as large as it was before starting the process at step 1 (that is, the degree of a in Q is at least $m + 1$).

In Figure 13, denote the vertices in the upper link of y_j by

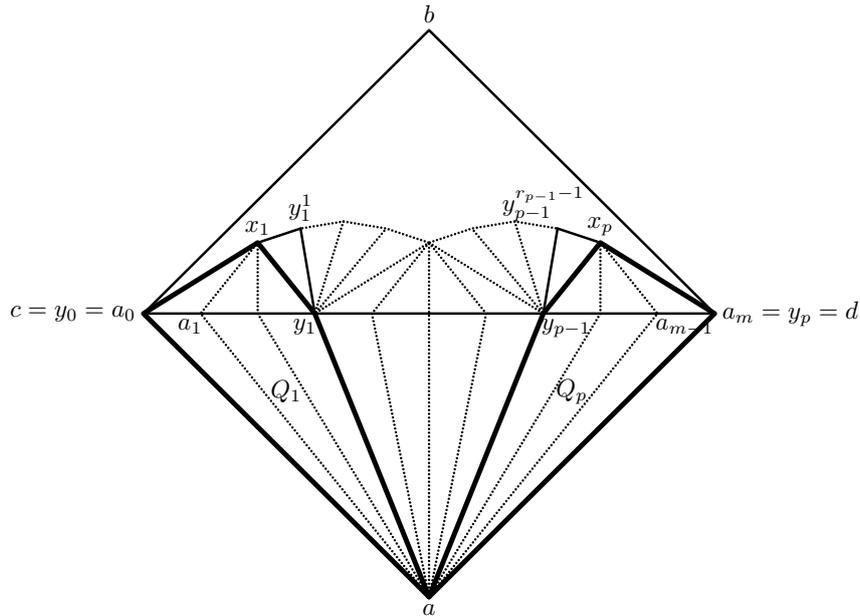$$x_j = y_j^0, y_j^1, y_j^2, \ldots, y_j^{r_j} = x_{j+1}.$$

Figure 13: The quadrilateral Q after performing steps 1-6, and after reorganising the interior vertices of quadrilaterals Q_i, $1 \leqslant i \leqslant p$.

By construction we have $r_j > 0$ for all j.

Refer to Figure 14(a). Since $p > 1$, x_1, $y_1 = a_j$, and x_2 are in the interior of Q. Moreover, both $a_j = y_1$, $j > 1$, and x_2 are of degree at least five and, by design of the procedure, $y_0 y_1^\ell$, $1 \leqslant \ell \leqslant r_1$, is a non-edge (otherwise y_1^ℓ is a better choice for x_1). It follows that we can perform the flips $x_1 a_j \mapsto a_{j-1} y_1^1$, $x_1 a_{j-1} \mapsto a_{j-2} y_1^1$, etc., all the way down to $x_1 a_2 \mapsto a_1 y_1^1$ (see Figure 14(b)). Note that y_1 and x_1 are now both of degree at least four, the degree of y_1^1 is larger than before, a_1 is of degree five, and all other degrees have not changed. Since $x_1 a_i$, $2 \leqslant i \leqslant m$, must be non-edges, a and x_1 do not have common neighbours. We can thus perform the flip $a_0 a_1 \mapsto a x_1$, see Figure 14(c).
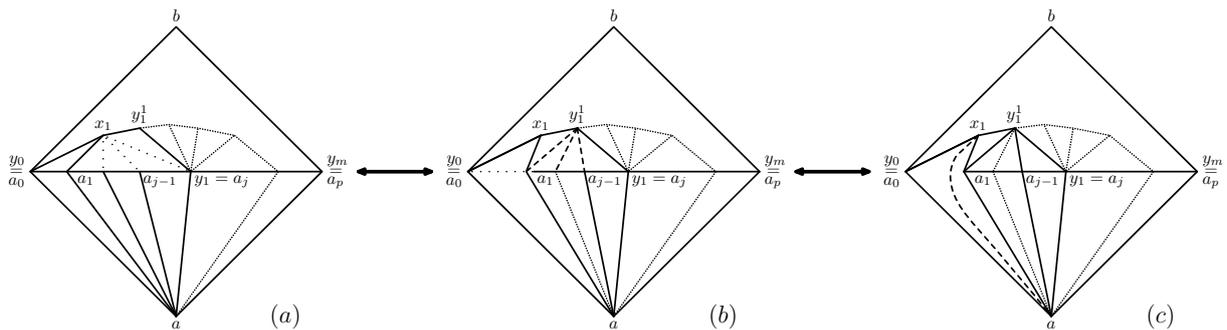

Figure 14: Increasing the size of the link of a.

This strictly increases the degree of a. We now start over with our procedure at step 1.

Since there are only finitely many vertices inside Q, this procedure must terminate with Q containing a common neighbour of a and b. This proves the claim and completes the proof of the lemma. $\qquad\square$

Proof of Theorem 6. To prove the theorem it suffices to show that $S \sim A_n$ for all $S \in \mathcal{F}_n \setminus \{\Gamma_n\}$.

Apply Lemma 12 to split S into two proper quadrilaterals $S = Q \cup R$. This is always possible since $S \neq \Gamma_n$. Use Lemma 13 to turn all interior vertices of both Q and R into vertices of degree four.

If, after the first or second application of Lemma 13, any of the boundary vertices of Q (or R) are of degree four, we grow Q (or R) such that eventually it is bounded by vertices of degree at least five, or $S \sim \Gamma_n$. However, since all edge flips on Γ_n produce a non-flag 2-sphere triangulation, the latter case implies $S = \Gamma_n$, a contradiction.

Thus, S can be transformed into a triangulation S'' of the 2-sphere which splits into two ordered proper quadrilaterals. This corresponds to the cases $a = b$ in the proof of Lemma 9. In particular, either $S'' = \Gamma_n$, which is impossible, $S'' = A_n$, or the degrees of all vertices of the separating induced 4-cycle satisfy the preconditions of Lemma 9, and we can conclude that $S \sim A_n$. $\qquad\square$

4 The flip graph $\mathcal{S}_n$ of n-vertex stacked 2-spheres

Every pair of n-vertex stacked 2-spheres is, by definition, connected in the Pachner graph of stacked 2-spheres by a sequence of $(n-4)$ 2-moves, followed by a sequence of $(n-4)$ 0-moves. However, if we look at the flip graph $\mathcal{S}_n$ of n-vertex stacked 2-spheres, the situation is different.

In this section we show that the structure of $\mathcal{S}_n$ is very special. More precisely, we prove that $\mathcal{S}_n$ is not connected for $n \geqslant 7$ (Corollary 19), and that the number of connected components rapidly increases with the number of vertices (Corollary 20). More precisely, for n fixed, the number of connected components is at least as large as the number of isomorphism classes of trees of maximum degree at most four on $\lfloor \frac{n-5}{3} \rfloor$ vertices. See Table 1 for the number and cardinalities of connected components of $\mathcal{S}_n$ for $n \leqslant 14$.

For a stacked 2-sphere S, let $\overline{S}$ be the unique stacked 3-ball whose boundary is S, see Lemma 3. If α is a triangle of S then α is a face of a unique tetrahedron of $\overline{S}$ (i.e., a clique of size four in the edge graph of S). We denote this unique tetrahedron by $\overline{\alpha}$. Naturally, $\overline{\alpha}$ is a node in the dual graph $\Lambda(\overline{S})$.

Theorem 14. *Let S be a stacked 2-sphere other than the standard sphere. Let $\alpha = abc$, $\beta = abd$ be two triangles of S. Let $\overline{\alpha}$ (resp., $\overline{\beta}$) be the unique tetrahedron in $\overline{S}$ containing α (resp., β). Then cd is not an edge of S and the 2-sphere T obtained from S by the edge flip $ab \mapsto cd$ is stacked if and only if the nodes $\overline{\alpha}$ and $\overline{\beta}$ of $\Lambda(\overline{S})$ are adjacent in $\Lambda(\overline{S})$.*

Proof. Suppose $\overline{\alpha}$ and $\overline{\beta}$ are adjacent in the dual graph $\Lambda(\overline{S})$, $\overline{\alpha} \neq \overline{\beta}$. Then there exists a vertex e of S such that $\overline{\alpha} = abce$ and $\overline{\beta} = abde$ ($e \notin \{d, c\}$ since $\overline{\alpha} \neq \overline{\beta}$). If cd is an edge

n	$\#(\mathcal{S}_n)$	# cc	size of connected components
4	1	1	1
5	1	1	1
6	1	1	1
7	3	1	3
8	7	2	$1, 6$
9	24	2	$1, 23$
10	93	3	$3, 4, 86$
11	434	5	$1, 7, 10, 19, 397$
12	2110	8	$1, 2, 6, 43, 46, 57, 82, 1873$
13	11002	15	$1, 2, 2, 3, 4, 6, 6, 7, 57, 222$ $223, 246, 326, 394, 9503$
14	58713	33	$1, 1, 3, 4, 4, 4, 5, 6, 6, 6, 6, 7, 7, 9, 9, 9, 12,$ $15, 19, 27, 28, 36, 36, 246, 304, 339, 757,$ $1165, 1182, 1571, 1944, 1987, 48958$

Table 1: Number and cardinalities of the connected components of $\mathcal{S}_n$ for $n \leqslant 14$.

of S then $\{a, b, c, d, e\}$ is a clique in the edge graph of S and hence, by Lemma 3, $abcde$ is a simplex of $\overline{S}$. This is not possible since $\overline{S}$ is 3-dimensional.

Let $B = \overline{S} \cup abcd$. Since $\overline{S} \cap abcd$ is a 2-disk, B is a triangulated 3-ball. The link $\mathrm{lk}_B(ab)$ is the induced 3-cycle c-d-e-c in B. Let D be obtained from B by the 3-dimensional bistellar 2-move that replaces the three tetrahedra $abcd$, $abce$ and $abde$ around edge ab with the two tetrahedra $\overline{\gamma} = acde$ and $\overline{\delta} = bcde$ sharing triangle cde, denoted by $ab \mapsto cde$ in short. By construction we have (i) $\partial D = T$, where T is the 2-sphere obtained from S by the edge flip $ab \mapsto cd$ and (ii) all edges of D are boundary edges (ab is the only edge of B not in the boundary which is removed by the bistellar move $ab \mapsto cde$) and thus T is stacked (cf. Section 2.3).

Conversely, suppose cd is not an edge of S and the triangulated 2-sphere T obtained from the stacked 2-sphere S by the edge flip $ab \mapsto cd$ is a stacked 2-sphere. Observe that both $\gamma = acd$ and $\delta = bcd$ are triangles of T.

Since $ab, abc, abd \in S = \partial\overline{S}$, $\mathrm{lk}_{\overline{S}}(ab)$ is a path in $E(\overline{S})$ from c to d. Let $\mathrm{lk}_{\overline{S}}(ab) = e_0\text{-}e_1\text{-}\cdots\text{-}e_k\text{-}e_{k+1}$ for some $k \geqslant 1$, where $e_0 = c$ and $e_{k+1} = d$. We have that $abce_1 = abe_0e_1, abe_1e_2, \ldots, abe_{k-1}e_k, abe_ke_{k+1} = abde_k$ are tetrahedra in $\overline{S}$. Thus, $abe_1, \ldots, abe_k$ are interior triangles of $\overline{S}$. Since abe_1 is an interior triangle with all of its edges in $\partial\overline{S}$, it follows that $\overline{S}[V(S) \setminus \{a, b, e_1\}]$ has two components, one containing e_0 and the other containing e_2. Thus, the common neighbours of e_0 and e_2 in $E(\overline{S}) = E(S)$ are a, b and e_1. Similarly, the set of common neighbours of e_{i-1} and e_{i+1} is $\{a, b, e_i\}$ for $1 \leqslant i \leqslant k$. This implies that the set of common neighbours of $c = e_0$ and $d = e_{k+1}$ in $E(T)$ is $\{a, b, e_1\} \cap \{a, b, e_k\}$ (note that $E(S)$ differs from $E(T) = E(\overline{T})$ only in edges ab and cd).

On the other hand the triangles $\gamma = acd = ae_0e_{k+1}$ and $\delta = bcd = be_0e_{k+1}$ are contained in unique tetrahedra $\overline{\gamma} = acdx$ and $\overline{\delta} = bcdy$ of $\overline{T}$ and hence a, b, x and y are common neighbours of c and d. By the above this is only possible if $e := x = y = e_1 = e_k$. In particular, $\mathrm{lk}_{\overline{S}}(ab)$ is a path from c to d of length two, $\overline{\alpha} = abce$, $\overline{\beta} = abde$, and in particular $\overline{\alpha}$ and $\overline{\beta}$ are adjacent in $\Lambda(\overline{S})$. $\qquad\square$

Remark 15. For an edge flip $ab \mapsto cd$ on a stacked 2-sphere S to be valid, we must have $\alpha = abc, \beta = abd \in F(S)$ and $cd \notin E(S)$. We have seen that an n-vertex 2-sphere T can be obtained from a stacked 2-sphere S by an edge flip $ab \mapsto cd$ (that is, the edge flip is valid) and T is stacked if and only if the nodes corresponding to tetrahedra $\overline{\alpha}$ and $\overline{\beta}$ of $\overline{S}$ are adjacent in $\Lambda(\overline{S})$.

Note that we can replace this latter condition in Theorem 14 by any of the following equivalent conditions:

- The path in the link of ab in $\overline{S}$ from c to d is of length exactly two.

- Edge ab is contained in exactly two tetrahedra of $\overline{S}$.

- The vertices a and b have exactly three common neighbours in S.

- There exists a unique vertex $e \notin \{c, d\}$ such that ae and be are edges of S.

While some of these conditions are easier to grasp, others are more efficient for implementations. It is thus useful to keep all of them in mind.

Remark 16. Let T be obtained from S by the edge flip $ab \mapsto cd$ and e, $\alpha = abc$, $\beta = abd$, $\gamma = acd$, $\delta = bcd$ as in the proof of Theorem 14. Then $\overline{\alpha} = abce$, $\overline{\beta} = abde \in \overline{S}$ and $\overline{\gamma} = acde$, $\overline{\delta} = bcde \in \overline{T}$. Moreover, let the (up to) two nodes adjacent to $\overline{\alpha}$ in $\Lambda(\overline{S})$ be $acex$ and $bcey$, and let the (up to) two nodes adjacent to $\overline{\beta}$ in $\Lambda(\overline{S})$ be $adez$ and $bdew$.
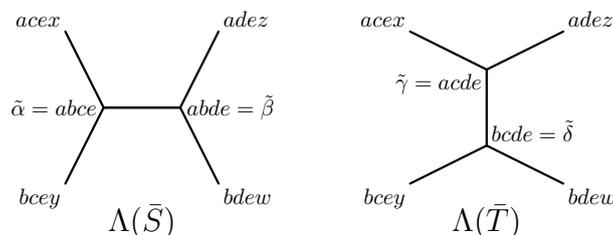

Figure 15: Transformation of dual graph by edge flip $ab \mapsto cd$ in the proof of Theorem 14.

Then the dual graph $\Lambda(\overline{T})$ is the tree built from $\Lambda(\overline{S})$, with set of nodes $U = \{\sigma \in \overline{S} \mid \sigma$ is a tetrahedron $\} \setminus \{\overline{\alpha}, \overline{\beta}\}) \cup \{\overline{\gamma}, \overline{\delta}\}$ with all arcs in $\Lambda(\overline{S})$ adjacent to $\overline{\alpha}$ and $\overline{\beta}$ removed, and arcs added between $\overline{\gamma}$ and $\overline{\delta}$ (corresponding to triangle cde), $\overline{\gamma}$ and $acex$ (corresponding to ace), $\overline{\delta}$ and $bcey$ (bce), $\overline{\gamma}$ and $adez$ (ade), and $\overline{\delta}$ and $bdew$ (bde), see Figure 15.

Corollary 17. *Let S be a stacked 2-sphere, $\alpha = abc$, $\beta = abd$ two triangles of S, $\overline{\alpha}$ (resp., $\overline{\beta}$) the unique tetrahedron of $\overline{S}$ containing α (resp., β), $\sigma \in \overline{S}$ correspond to a degree four node in $\Lambda(\overline{S})$, and let G_1, G_2, G_3, G_4 be the connected components of $\Lambda(\overline{S}) - \sigma$. If the 2-sphere T obtained from S by the edge flip $ab \mapsto cd$ is also a stacked 2-sphere then*

(i) *σ is a tetrahedron of $\overline{T}$,*

(ii) *σ is a degree four node in $\Lambda(\overline{T})$,*

(iii) *both $\overline{\alpha}$ and $\overline{\beta}$ are in one component of $\Lambda(\overline{S}) - \sigma$, say in G_4, and*

(iv) *the components of $\Lambda(\overline{T}) - \sigma$ are G_1, G_2, G_3, G'_4 for some tree G'_4.*

Proof. It follows from Theorem 14 that $\mathrm{lk}_{\overline{S}}(ab)$ is a path of the form c-e-d and $\overline{\alpha} = abce$, $\overline{\beta} = abde$ for some vertex e. In particular, $\overline{\alpha}$ and $\overline{\beta}$ are the only two tetrahedra in $\overline{S}$ containing ab. Since all the 2-dimensional faces of σ are interior triangles, we have $\sigma \notin \{\overline{\alpha}, \overline{\beta}\}$. Thus, σ cannot contain the edge ab. Since σ forms a clique in $E(S)$, this implies that σ forms a clique in $E(T)$. Hence $\sigma \in \overline{T}$. This proves part (i).

Observe that $\{a, c, d, e\}$ and $\{b, c, d, e\}$ span cliques in $E(T)$. Therefore, $\overline{\gamma} := acde$, $\overline{\delta} := bcde \in \overline{T}$. Let τ be a 2-dimensional face of σ. Then τ is an interior face in $\overline{S}$. Let $\tau = \sigma \cap \mu$ for some tetrahedron $\mu \in \overline{S}$. If $ab \not\subset \mu$ then μ forms a clique in $E(T)$ and hence $\mu \in \overline{T}$. Then $\tau = \sigma \cap \mu$ is an interior triangle of $\overline{T}$. If $ab \subset \mu$ then μ is $\overline{\alpha}$ or $\overline{\beta}$. Assume, without loss of generality, that $\mu = \overline{\alpha} = abce$. Since $ab \not\subset \sigma$ and $\mu \cap \sigma$ is a face of μ, $\tau = \mu \cap \sigma = ace$ or bce. Assume, without loss of generality, that $\tau = ace$. Then $\sigma = acex$ for some vertex x and $\tau = \sigma \cap \overline{\gamma}$. Thus, τ is an interior triangle of $\overline{T}$. Thus, each 2-dimensional face of σ is an interior triangle of $\overline{T}$. Part (ii) follows from this.

Part (iii) follows from the fact that $\overline{\alpha}$ and $\overline{\beta}$ share a triangle in $\overline{S}$ which (necessarily) is not a face of σ.

The four 2-dimensional faces of $\overline{\gamma} = acde$ are acd, ace, ade and cde. Since cd is a non-edge in $\overline{S}$, we have that acd, cde are not in $\overline{S}$ and $ace = \overline{\gamma} \cap \overline{\alpha}$, $ade = \overline{\gamma} \cap \overline{\beta}$. Thus, by part (iii), $\overline{\gamma}$ is not adjacent to any nodes of $G_1 \cup G_2 \cup G_3$. Similarly, $\overline{\delta}$ is not adjacent to any nodes of $G_1 \cup G_2 \cup G_3$. Part (iv) now follows since the set of nodes of $\Lambda(\overline{T})$ is $(\{\tau : \tau$ is a tetrahedron in $\overline{S}\} \setminus \{\overline{\alpha}, \overline{\beta}\}) \cup \{\overline{\gamma}, \overline{\delta}\}$. $\qquad\square$

Corollary 18. *Let S be a stacked 2-sphere, T a stacked 2-sphere obtained from S by an edge flip, and let V_S (resp., V_T) be the set of degree four nodes in $\Lambda(\overline{S})$ (resp., in $\Lambda(\overline{T})$). Then the induced subgraphs $\Lambda(\overline{S})[V_S]$ and $\Lambda(\overline{T})[V_T]$ are isomorphic.*

Proof. By Corollary 17, $V_S = V_T$. For $\sigma_1, \sigma_2 \in V_S = V_T$, σ_1 and σ_2 are adjacent in $\Lambda(\overline{S})[V_S]$ if and only if $\sigma_1 \cap \sigma_2$ is an interior triangle of $\overline{S}$ if and only if $\sigma_1 \cap \sigma_2$ contains three vertices if and only if $\sigma_1 \cap \sigma_2$ is an interior triangle of $\overline{T}$ if and only if σ_1 and σ_2 are adjacent in $\Lambda(\overline{T})[V_T]$. The corollary follows from this observation. $\qquad\square$

Corollary 19. *The flip graph $\mathcal{S}_n$ of n-vertex stacked 2-spheres is disconnected for $n \geqslant 8$.*

Proof. The stacked 3-ball associated to an n-vertex stacked 2-sphere, $n \geqslant 8$, has a dual graph with $m = n - 3 \geqslant 5$ nodes, and every m-node tree (with degrees of nodes $\leqslant 4$) is the dual graph of at least one stacked 3-ball. Hence there exist a stacked 3-ball B_1 with dual graph having one node of degree four and $m - 1$ nodes of degree at most three, and there exist a stacked 3-ball B_2 with dual graph with all m nodes of degree at most two. Then, by Corollary 18, the n-vertex stacked 2-spheres ∂B_1 and ∂B_2 are in different connected components of $\mathcal{S}_n$. $\qquad\square$

Corollary 20. *For $m \in \mathbb{Z}^+$, let $t(m)$ be the number of non-isomorphic m-node trees with degrees of nodes at most four. Moreover, let $n = 3m + 5$. Then the flip graph $\mathcal{S}_n$ of n-vertex stacked 2-spheres has $t(m)$ components each containing a single stacked 2-sphere.*

Proof. Let H be an m-node tree in which the degrees of all the nodes are at most four. Consider a new graph G by connecting each node of H of degree i to $(4 - i)$ new nodes. Then G is a connected acyclic graph and hence a tree. By construction, the number of new nodes in G equals the number of new arcs in G which is $\sum_{v \in V(H)} (4 - \deg_H(v)) = 4m - \sum_{v \in V(H)} \deg_H(v) = 4m - 2(m - 1) = 2m + 2$. Therefore, G has $(m - 1) + (2m + 2) = 3m + 1$ arcs, and thus $3m + 2$ nodes. It follows that G has m nodes of degree four and $2m + 2$ nodes of degree one, and each degree one node of G is adjacent to a degree four node.

Let B be a stacked 3-ball whose dual graph $\Lambda(B)$ is G. It follows from the definition that we can always construct such a stacked 3-ball. Let $S = \partial B$. Since S is stacked it must have $3m + 5$ vertices. Let $\alpha = abc$, $\beta = abd$ be two triangles of S, and let $\overline{\alpha}$ (resp., $\overline{\beta}$) be the unique tetrahedron of B containing α (resp., β). Then $\deg_{\Lambda(B)}(\overline{\alpha}), \deg_{\Lambda(B)}(\overline{\beta}) < 4$ and hence $\deg_{\Lambda(B)}(\overline{\alpha}) = 1 = \deg_{\Lambda(B)}(\overline{\beta})$. If $\overline{\alpha} = \overline{\beta}$, then cd is an edge and hence we cannot perform the edge flip $ab \mapsto cd$. If $\overline{\alpha} \neq \overline{\beta}$, then $\overline{\alpha}$ and $\overline{\beta}$ are not adjacent in $\Lambda(B)$ (degree one nodes are only adjacent to degree four nodes in $\Lambda(B)$) and hence, by Theorem 14, the 2-sphere T obtained from S by the edge flip $ab \mapsto cd$ is not stacked. Thus S is isolated in $\mathcal{S}_n$.

If H_1 and H_2 are non-isomorphic trees on m nodes, then the above construction carried out for both H_1 and H_2 leads to two non-isomorphic trees G_1 and G_2, leading to two non-isomorphic stacked 3-balls B_1 and B_2 with, by Lemma 3, non-isomorphic boundaries S_1 and S_2. Since there exist at least $t(m)$ non-isomorphic m-node trees with degree of nodes at most four, we have at least $t(m)$ singleton components in $\mathcal{S}_n$. $\qquad\square$

Corollary 21. *The number of connected components in $\mathcal{S}_n$ is bounded from below by C^n, for some real number $C > 1$.*

Proof. Let $m = \lfloor \frac{n-5}{3} \rfloor$. Let $t(m)$ be the number of non-isomorphic m-node trees with degree of nodes at most four as in Corollary 20.

Claim: The number of components in $\mathcal{S}_n$ is at least $t(m)$.

Let $\mathcal{T}$ be the set of all m-node trees with node-degrees at most four. For each $H \in \mathcal{T}$, we can construct a $(3m + 2)$-nodes tree G whose degree four nodes are the nodes of H and all others are of degree one (as in the proof of Corollary 20). By randomly attaching

an additional $n - 3m - 5 \leqslant 2$ extra nodes to existing degree one nodes of G we obtain a new tree G' having the same set of degree four nodes as in G. Let B be a stacked 3-ball whose dual graph is G' and let $S = \partial B$. By construction, S is a stacked 2-sphere with exactly n vertices. Let V_S be as in Corollary 18. Then $G'[V_S] = G[V_S] = H$. Therefore, by Corollary 18, the n-vertex stacked 2-spheres obtained in this process corresponding to different graphs in $\mathcal{T}$ are in different components of $\mathcal{S}_n$. This proves the claim.

Since $t(m)$ is exponential in m, the result follows from the claim. $\qquad\square$

Following arguments along the lines of Corollary 17 we can observe that, apart from a large number of isolated singleton components in $\mathcal{S}_n$, there are also larger connected components corresponding to dual graphs with no, or very few nodes of degree four. For instance, the largest connected component in $\mathcal{S}_n$, $n \leqslant 14$, shown in Table 1, corresponds to boundaries S of stacked balls $\overline{S}$ with dual graphs without nodes of degree four (i.e., $V_S = \emptyset$). Let $\mathcal{S}_n^0$ denote the flip graph consisting of this class of stacked 2-spheres. We have the following result.

Theorem 22. *The flip graph $\mathcal{S}_n^0$ is connected.*

We split the proof of Theorem 22 into two lemmas.

Lemma 23. *Each stacked 2-sphere $S \in \mathcal{S}_n^0$ is connected to a stacked 2-sphere T in the flip graph $\mathcal{S}_n^0$, where the dual graph $\Lambda(\overline{T})$ of $\overline{T}$ is a path.*

Proof. The idea of the proof is to show that, for every $S \in \mathcal{S}_n^0$ with $\Lambda(\overline{S})$ not a path, S is connected in $\mathcal{S}_n^0$ to a stacked 2-sphere $T \in \mathcal{S}_n^0$ with the number of nodes of degree three in $\Lambda(\overline{T})$ less than that in $\Lambda(\overline{S})$.

For $S \in \mathcal{S}_n^0$ and α, β nodes in $\Lambda(\overline{S})$, let $d_{\overline{S}}(\alpha, \beta)$ be the length of the unique path from α to β in the tree $\Lambda(\overline{S})$. Moreover, if $\overline{S}$ has a degree three node in $\Lambda(\overline{S})$, let $\ell(S) = \min\{d_{\overline{S}}(\alpha, \beta) \mid \alpha \text{ leaf}, \beta \text{ degree three in } \Lambda(\overline{S})\}$.

Claim 1: Let $S \in \mathcal{S}_n^0$ be a stacked 2-sphere such that $\Lambda(\overline{S})$ is not a path. If $\ell(S) \geqslant 2$ then there exists a stacked 2-sphere $T \in \mathcal{S}_n^0$ such that (i) S is connected to T in $\mathcal{S}_n^0$, (ii) the number of degree three nodes in $\Lambda(\overline{T})$ is the same as in $\Lambda(\overline{S})$ and (iii) $\ell(T) = \ell(S) - 1$.

Let $\ell = \ell(S) = d_{\overline{S}}(\gamma, \delta)$, where γ is a degree three node and δ is a leaf in $\Lambda(\overline{S})$. Let γ_0-γ_1-$\cdots$-γ_ℓ be the path in $\Lambda(\overline{S})$ from $\gamma = \gamma_0$ to $\delta = \gamma_\ell$. Then $\deg_{\Lambda(\overline{S})}(\gamma_0) = 3$, $\deg_{\Lambda(\overline{S})}(\gamma_i) = 2$ for $1 \leqslant i \leqslant \ell - 1$, and $\deg_{\Lambda(\overline{S})}(\gamma_\ell) = 1$. Let the other nodes adjacent to γ be α and β. Assume, without loss of generality, that $\gamma = 1234$, $\alpha = 124a$, $\beta = 134b$ and $\gamma_1 = 123x_1$. Then, the link of 23 in $\overline{S}$ is of the form 4-1-x_1-$\cdots$-x_k for some $k \leqslant \ell$.

Case 1. Let $k = 1$. It follows that $23x_1$ is a face of $S = \partial \overline{S}$. By Theorem 14, the triangulated 2-sphere T obtained from S by the edge flip $23 \mapsto 4x_1$ is stacked and hence, by Corollary 17, is in $\mathcal{S}_n^0$. By Lemma 3, $\gamma' := 134x_1$ and $\gamma_1' := 124x_1$ are tetrahedra in $\overline{T}$.

Following the transformation of the dual graph of a stacked ball under an edge flip, as shown in Figure 15, the dual graph $\Lambda(\overline{T})$ is obtained from $\Lambda(\overline{S})$ by replacing the three edges adjacent to γ with the path β-γ'-γ_1'-α, and attaching the path γ_2-$\cdots$-γ_ℓ to either γ' or γ_1'. In either case, the path from the new degree three node to γ_ℓ is of length $\ell - 1$,

and since the remaining part of $\Lambda(\overline{T})$ is equal to the remaining part of $\Lambda(\overline{S})$, we have $\ell(T) = \ell(S) - 1$ and Claim 1 is true in this case.

Case 2. Let $k \geqslant 2$. In this case we can assume that $\gamma_i = 23x_{i-1}x_i$ for $2 \leqslant i \leqslant k$, and that the triangles $21x_1, 2x_1x_2, \ldots, 2x_{k-2}x_{k-1}, 31x_1, 3x_1x_2, \ldots, 3x_{k-2}x_{k-1} \in S$ (i.e., are in the boundary of $\overline{S}$). Since $\deg_{\Lambda(\overline{S})}(\gamma_k) \leqslant 2$ ($= 1$ if $k = \ell$ and $= 2$ if $k < \ell$), at least two 2-dimensional faces of γ_k are triangles of S. This implies that at least one of the triangles $2x_{k-1}x_k$ and $3x_{k-1}x_k$ is a triangle of S.
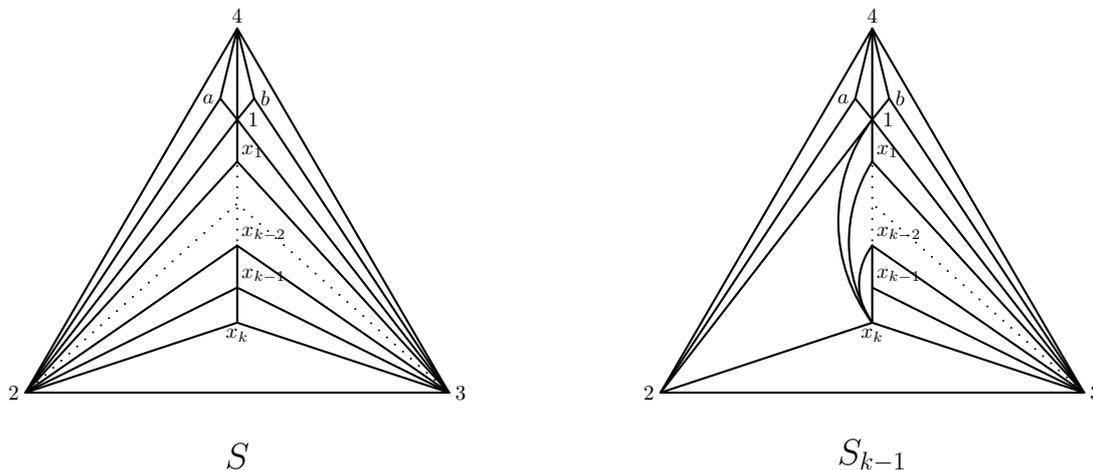

$$S \qquad\qquad\qquad\qquad\qquad S_{k-1}$$

Figure 16: Sequence of edge flips as performed in the proof of Lemma 23, Claim 1, Case 2.

Assume, without loss of generality, that $2x_{k-1}x_k \in S$. (In that case, γ_{k+1} is of the form $3x_{k-1}x_kx_{k+1}$ for some $x_{k+1} \in V(S)$ when $k < \ell$.) Let S_1 be obtained from $S = S_0$ by the edge flip $2x_{k-1} \mapsto x_kx_{k-2}$. Since $\mathrm{lk}_{\overline{S}}(2x_{k-1}) = x_{k-2}\text{-}3\text{-}x_k$, by Theorem 14, S_1 is stacked. Observe that the path $\gamma_{k-2}\text{-}\gamma_{k-1}\text{-}\gamma_k\text{-}\gamma_{k+1}$ in $\Lambda(\overline{S})$ is replaced by $\gamma_{k-2}\text{-}(23x_kx_{k-2})\text{-}(3x_kx_{k-1}x_{k-2})\text{-}\gamma_{k+1}$ in $\Lambda(\overline{S_1})$ when $k < \ell$, and $\gamma_{k-2}\text{-}\gamma_{k-1}\text{-}\gamma_k$ is replaced by $\gamma_{k-2}\text{-}(23x_kx_{k-2})\text{-}(3x_kx_{k-1}x_{k-2})$ when $k = \ell$. Thus, $\Lambda(\overline{S_1})$ is isomorphic to $\Lambda(\overline{S})$.

Inductively, for $1 \leqslant i \leqslant k - 1$, $\mathrm{lk}_{\overline{S}}(2x_{k-i}) = x_{k-i-1}\text{-}3\text{-}x_k$ and hence the sphere S_i obtained from S_{i-1} by the edge flip $2x_{k-i} \mapsto x_kx_{k-i-1}$ is stacked. Then $\Lambda(\overline{S_i})$ is isomorphic to $\Lambda(\overline{S_{i-1}})$, see Figure 16. (Note that S_{k-1} is obtained by the sequence of edge flips $2x_{k-1} \mapsto x_kx_{k-2}, 2x_{k-2} \mapsto x_kx_{k-3}, \ldots, 2x_2 \mapsto x_kx_1, 2x_1 \mapsto x_k1$.)

It follows that S_{k-1} is stacked, S can be joined to S_{k-1} in $\mathcal{S}_n^0$, $\Lambda(\overline{S_{k-1}})$ is isomorphic to $\Lambda(\overline{S})$, and $\mathrm{lk}_{\overline{S_{k-1}}}(23) = 4\text{-}1\text{-}x_k$. In particular, S_{k-1} satisfies the hypothesis of Case 1, $\ell(S_{k-1}) = \ell(S)$ and the number of degree three nodes in $\Lambda(\overline{S_{k-1}})$ is the same as that in $\Lambda(\overline{S})$. Consequently, by Case 1, S_{k-1} is connected to some T in $\mathcal{S}_n^0$, such that the number of degree three nodes in $\Lambda(\overline{T})$ is the same as that in $\Lambda(\overline{S_{k-1}})$ (which is the same as that in $\Lambda(\overline{S})$) and $\ell(T) = \ell(S_{k-1}) - 1 = \ell(S) - 1$. This completes the proof of Claim 1.

Claim 2: For $S \in \mathcal{S}_n^0$, if $\Lambda(\overline{S})$ has a leaf which is adjacent to a degree three node in $\Lambda(\overline{S})$ (i.e., $\ell(S) = 1$) then there exists $T \in \mathcal{S}_n^0$ which can be obtained from S by an edge flip and the number of nodes of degree three in $\Lambda(\overline{T})$ is one less than that in $\Lambda(\overline{S})$.

Let $\delta = 123d$ be a leaf node which is adjacent to a degree three node $\gamma = 1234$. Assume, as above, that the adjacent nodes of γ are $\alpha = 124a$ and $\beta = 134b$. Then edge 23 is in two tetrahedra and, by Theorem 14, the 2-sphere T obtained from S by the edge flip $23 \mapsto 4d$ is stacked and hence in $\mathcal{S}_n^0$ by Corollary 17. Moreover, by Lemma 3, $\gamma' := 124d$ and $\delta' := 134d$ are in $\overline{T}$. Again, by following the transformation shown in Figure 15, $\Lambda(\overline{T})$ contains the path α-γ'-δ'-β instead of the three edges adjacent to γ in $\Lambda(\overline{S})$. Since the remaining parts of $\Lambda(\overline{S})$ and $\Lambda(\overline{T})$ coincide, Claim 2 follows.

The result follows inductively using Claims 1 and 2. $\qquad\square$

Lemma 24. *Let $\partial\Delta_n$ be as shown in Figure 17 and let $S \in \mathcal{S}_n^0$. If $\Lambda(\overline{S})$ is a path then S is connected to $\partial\Delta_n$ in $\mathcal{S}_n^0$.*
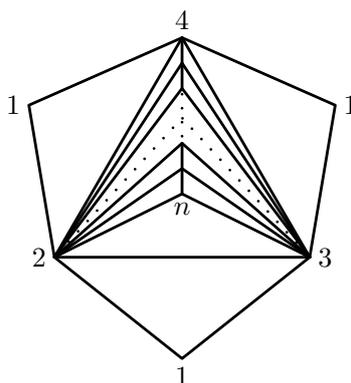

Figure 17: The stacked 3-ball Δ_n. Note that this complex is also used as a target in [4] to prove upper bounds on the diameter of the flip graph $\mathcal{P}_n$ of n-vertex 2-spheres.

Proof. Let $\Lambda(\overline{S}) = \gamma_1$-$\gamma_2$-$\cdots$-$\gamma_{n-3}$.

Since the stacked 3-ball Δ_n is characterised by the property of being the star of an edge (all of its tetrahedra share an edge), we prove this lemma by verifying the following claim.

Claim: If $\gamma_1, \ldots, \gamma_k$ have a common edge and $\gamma_1, \ldots, \gamma_{k+1}$ have no common edge, $k \leqslant n - 4$, then S can be joined to $T \in \mathcal{S}_n^0$, where $\Lambda(\overline{T})$ is a path of the form α_1-α_2-$\cdots$-α_{n-3} such that $\alpha_1, \ldots, \alpha_{k+1}$ have a common edge.

Since $\gamma_1, \ldots, \gamma_{k+1}$ have no common edge, we can assume that $k \geqslant 3$. Let $\gamma_i = abx_ix_{i+1}$ for $1 \leqslant i \leqslant k$. Assume without loss of generality that $\gamma_{k+1} = bx_kx_{k+1}x_{k+2}$. Then $\text{lk}_{\overline{S}}(ax_k) = x_{k-1}$-$b$-$x_{k+1}$. Thus, by Theorem 14, the 2-sphere S_1 obtained from S by the edge flip $ax_k \mapsto x_{k+1}x_{k-1}$ is stacked. Similarly, the 2-sphere S_2 obtained from S_1 by the edge flip $ax_{k-1} \mapsto x_{k+1}x_{k-2}$ is stacked. Continuing this way, we obtain a stacked sphere $T = S_{k-1}$ from S_{k-2} by the edge flip $ax_2 \mapsto x_{k+1}x_1$, see Figure 18. Hence S can be joined to T in $\mathcal{S}_n^0$ and $\Lambda(\overline{T}) = \alpha_1$-$\alpha_2$-$\cdots$-$\alpha_{k+1}$-$\gamma_{k+2}$-$\cdots$-$\gamma_{n-3}$, where $\alpha_1 = bx_{k+1}ax_1$, $\alpha_i = bx_{k+1}x_{i-1}x_i$, $2 \leqslant i \leqslant k$, and $\alpha_{k+1} = bx_{k+1}x_kx_{k+2}$. This proves the claim.

The lemma follows by induction using the claim. $\qquad\square$

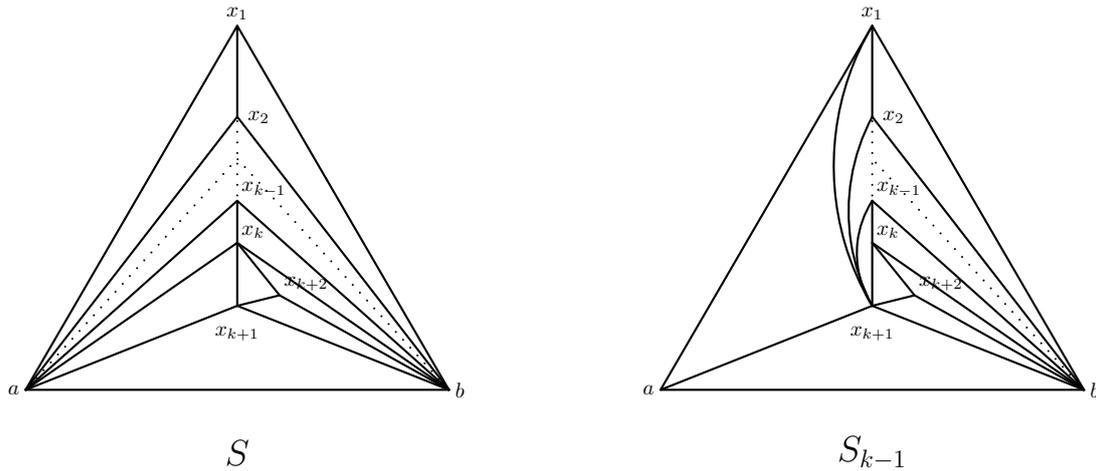Figure 18: Sequence of edge flips as performed in the proof of Lemma 24.

Proof of Theorem 22. The result follows from Lemmas 23 and 24. ☐

References

[1] B. Bagchi and B. Datta. Lower bound theorem for normal pseudomanifolds. *Expo. Math.*, 26(4):327–351, 2008.

[2] D. Barnette. The triangulations of the 3-sphere with up to 8 vertices. *J. Combin. Theory Ser. A*, 14:37–52, 1973.

[3] P. Bose and F. Hurtado. Flips in planar graphs. *Comput. Geom.*, 42(1):60–80, 2009.

[4] P. Bose, D. Jansens, A. van Renssen, M. Saumell, and S. Verdonschot. Making triangulations 4-connected using flips. *Comput. Geom.*, 47(2, part A):187–197, 2014.

[5] P. Bose and S. Verdonschot. A history of flips in combinatorial triangulations. In *Computational geometry*, volume 7579 of *Lecture Notes in Comput. Sci.*, pages 29–44. Springer, Cham, 2011.

[6] J. Cardinal, M. Hoffmann, V. Kusters, C. D. Tóth, and M. Wettstein. Arc Diagrams, Flip Distances, and Hamiltonian Triangulations. *Comput. Geom.*, 68:206–225, 2018.

[7] B. Datta and S. Murai. On stacked triangulated manifolds. *Electron. J. Comb.*, 24(4): #P4.12, 2017.

[8] J. A. De Loera, J. Rambau, and F. Santos. *Triangulations*, volume 25 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2010. Structures for algorithms and applications.

[9] R. Dougherty, V. Faber, and M. Murphy. Unflippable tetrahedral complexes. *Discrete Comput. Geom.*, 32(3):309–315, 2004.

[10] F. Frati. A Lower Bound on the Diameter of the Flip Graph. *Electron. J. Comb.*, 24(1): #P1.43, 2017.

[11] B. Grünbaum and V. P. Sreedharan. An enumeration of simplicial 4-polytopes with 8 vertices. *J. Combin. Theory*, pages 437–465, 1967.

[12] H. Komuro. The diagonal flips of triangulations on the sphere. *Yokohama Math. J.*, 44(2):115–122, 1997.

[13] F. H. Lutz and E. Nevo. Stellar theory for flag complexes. *Math. Scand.*, 118(1):70–82, 2016.

[14] S. Matveev. *Algorithmic topology and classification of 3-manifolds*, volume 9 of *Algorithms and Computation in Mathematics*. Springer, Berlin, second edition, 2007.

[15] A. Mijatović. Simplifying triangulations of S^3. *Pacific J. Math.*, 208(2):291–324, 2003.

[16] R. Mori, A. Nakamoto, and K. Ota. Diagonal flips in hamiltonian triangulations on the sphere. *Graphs and Combinatorics*, 198:413–418, 2003.

[17] U. Pachner. Konstruktionsmethoden und das kombinatorische Homöomorphieproblem für Triangulierungen kompakter semilinearer Mannigfaltigkeiten. *Abh. Math. Sem. Uni. Hamburg*, 57:69–86, 1987.

[18] J. Spreer. Combinatorial 3-manifolds with transitive cyclic symmetry. *Discrete Comput. Geom.*, 51(2):394–426, 2014.

[19] E. Steinitz. Polyeder und Raumeinteilungen. In *Encyclopädie der mathematischen Wissenschaften*, volume III.1.2, pages 1–139. B. G. Teubner, Leipzig, 1922.

[20] T. Sulanke and F. H. Lutz. Isomorphism-free lexicographic enumeration of triangulated surfaces and 3-manifolds. *European J. Combin.*, 30(8):1965–1979, 2009.

[21] K. Wagner. Bemerkungen zum vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 46:26–32, 1936.

[22] H. Whitney. A theorem on graphs. *Ann. of Math.*, 32:378–390, 1931.