A Theory of Spherical Diagrams

Giovanni Viglietta ✉ 🏠 🄳

Department of Computer Science and Engineering, University of Aizu, Japan

─── **Abstract** ───

We introduce an axiomatic theory of spherical diagrams as a tool to study certain combinatorial properties of polyhedra in $\mathbb{R}^3$, which are of central interest in the context of Art Gallery problems for polyhedra and other visibility-related problems in discrete and computational geometry.

1 Introduction

Geometric intuition

Consider a set $\mathcal{P}$ of internally disjoint opaque polygons in $\mathbb{R}^3$ and a viewpoint $v \in \mathbb{R}^3$ such that no vertex of any polygon in $\mathcal{P}$ is visible to v. An example is given by the set of six rectangles in Figure 1 (left) with respect to the point v located at the center of the arrangement.

Let S be a sphere centered at v that does not intersect any of the polygons in $\mathcal{P}$ (we may assume without loss of generality that S is the unit sphere), and let $S_{\mathcal{P}}$ be the *visibility map* of $\mathcal{P}$ with respect to v. That is, $S_{\mathcal{P}}$ is the set of radial projections onto S of the portions of edges of polygons in $\mathcal{P}$ that are visible to v (i.e., where polygons occlude projection rays). Figure 1 (right) shows an example of such a projection. The resulting structure $S_{\mathcal{P}}$ is a *Spherical Occlusion Diagram*; this name indicates that no vertices of $\mathcal{P}$ appear in the visibility map, because they are all occluded by polygons.

In this paper we set out to formalize an axiomatic theory of Spherical Occlusion Diagrams and study their combinatorial structure.
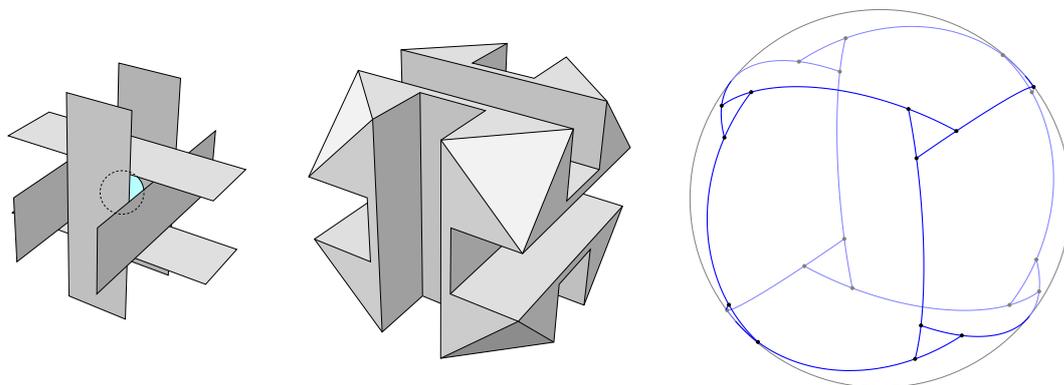
Figure 1 Construction of a Spherical Occlusion Diagram (right) as the visibility map of an arrangement of six rectangles (left) or the visibility map of a polyhedron with respect to the central point, which does not see any vertices (center).

Applications

Spherical Occlusion Diagrams naturally arise in visibility-related problems for arrangements of polygons in $\mathbb{R}^3$, and especially for polyhedra.

An example is found in [3], where an upper bound is given on the number of edge guards that solve the Art Gallery problem in a general polyhedron. That is, given a polyhedron $\mathcal{P}$, the problem is to find a (small) set of edges that collectively see the whole interior of $\mathcal{P}$. An edge e sees a point x if and only if there is a point $y \in e$ such that the open line segment xy does not intersect the boundary of $\mathcal{P}$.[1] (The reader can refer to [2, 14] for more results on this problem, as well as [9, 10] for surveys on the Art Gallery problem in 2-dimensional settings.)

The idea of [3] is to preliminarily select a (small) set E of edges that cover all vertices of $\mathcal{P}$. Note that E may be insufficient to guard the interior of $\mathcal{P}$, as some of its points may be invisible to all vertices; Figure 1 (center) shows an example. Thus, an additional (small) set of edges E' is selected, which collectively see all internal points of $\mathcal{P}$ that do not see any vertices. Clearly, $E \cup E'$ is a set of edges that see all internal points of $\mathcal{P}$.

The selection of the set of edges E' is carried out in [3] by means of an ad-hoc analysis of some properties of points that do not see any vertices of $\mathcal{P}$. Spherical Occlusion Diagrams offer a systematic and general tool to reason about points in a polyhedron that do not see any vertices.

Spherical Occlusion Diagrams have also provided a framework for proving the main result of [12, 13]: Any point that sees no vertex of a polyhedron must see at least 8 of its edges, and the bound is tight.

Summary

In Section 2 we give an axiomatic theory of Spherical Occlusion Diagrams and discuss their realizability as visibility maps. In Section 3 we prove some basic properties of Spherical Occlusion Diagrams, while in Section 4 we focus on an important pattern called *swirl*. Sections 5 and 6 are devoted to two interesting classes of Spherical Occlusion Diagrams: the *swirling* and the *uniform* ones, respectively. Section 7 contains some concluding remarks and directions for future research.

A preliminary version of this paper appeared at CCCG 2022 [15]. In the present version, most sections have been reworked and all missing proofs have been included. In addition, Proposition 13 and Theorems 28 and 29 are new contributions.

2   Axiomatic theory

In the following, we will abstract from a specific set of polygons $\mathcal{P}$ and a viewpoint v, and we will focus on the salient properties of the Spherical Occlusion Diagrams constructed in Section 1 in order to devise a small set of axioms that describe all of them.

---

[1] This definition of visibility is slightly more restrictive than the one most commonly adopted in the Art Gallery literature, which allows the segment xy to graze the boundary of $\mathcal{P}$ without properly crossing it. However, for the purposes of this paper, the two definitions are essentially equivalent and lead to the same theory of Spherical Occlusion Diagrams. The choice of the less common definition is due to the fact that it prevents visibility maps of polyhedra from having multiple coincident arcs, as well as other pathological configurations. In turn, this allows for a slightly more elegant and concise formulation of Proposition 3, which relates visibility maps and Spherical Occlusion Diagrams.

Definitions

Some terms will be useful. A *great circle* on a sphere S is a circle of maximum diameter within S. Equivalently, a great circle is the intersection between S and a plane through its center. Two points on a sphere are *antipodal* if the line through them contains the center of the sphere. A *great semicircle* is an arc of a great circle whose endpoints are antipodal.

The *geodesic arc* between two distinct non-antipodal points x and y on a sphere S is the unique shortest path within S having endpoints x and y. Equivalently, it is the arc of the unique great circle through x and y that has x and y as endpoints and is strictly shorter than a great semicircle. Two geodesic arcs are *collinear* if they lie on the same great circle.

▶ **Definition 1.** *Let a and b be two non-collinear geodesic arcs on a sphere. If an endpoint p of a lies in the relative interior of b, we say that a* hits *b at p (or* feeds into *b at p) and b* blocks *a at p.*
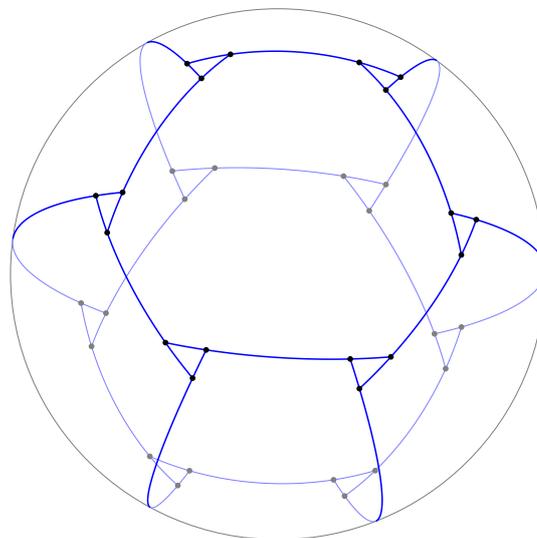
Axioms

We are now ready to formulate an abstract theory of Spherical Occlusion Diagrams.

▶ **Definition 2.** *A* Spherical Occlusion Diagram (SOD) *is a finite non-empty collection $\mathcal{D}$ of geodesic arcs on the unit sphere in $\mathbb{R}^3$ satisfying the following axioms.*
A1. Any two arcs in $\mathcal{D}$ are internally disjoint.
A2. Each arc in $\mathcal{D}$ is blocked by arcs of $\mathcal{D}$ at each endpoint.
A3. All arcs in $\mathcal{D}$ that hit the same arc of $\mathcal{D}$ reach it from the same side.

As an example, Figure 2 shows an SOD with 18 arcs.

▣ **Figure 2** Example of an SOD with 18 arcs.

Discussion

We remark that the SOD axioms introduced in [15] are slightly different. In particular, the axiom A1 therein states that, if two arcs $a, b \in \mathcal{D}$ have a non-empty intersection, then a hits b or b hits a. Also, in [15] it is not postulated that arcs are shorter than great semicircles, but this property is derived as a theorem.

The axioms in the present paper are slightly more inclusive, in that they allow multiple arcs to share an endpoint, as long as the common endpoint is an interior point of some other arc. Allowing such "degenerate configurations" in SODs allows us to prove a slightly more general Proposition 3, whereas the one in [15] required the extra assumption that the viewpoint v be in "general position" with respect to the polyhedron $\mathcal{P}$.

Moreover, if we removed the assumption that the arcs in an SOD are shorter than great semicircles, we would still be able to prove that arcs are *not longer* than great semicircles, but this would not exclude some degenerate cases in which some arcs are great semicircles.

Realizability

It is easy to recognize that the visibility maps $S_{\mathcal{P}}$ as constructed in Section 1 indeed provide a model for our theory, as they satisfy all its axioms, at least when $\mathcal{P}$ is a polyhedron.

▶ **Proposition 3.** *The visibility map $S_{\mathcal{P}}$ of any polyhedron $\mathcal{P}$ with respect to any viewpoint v that does not lie on the boundary of $\mathcal{P}$ and sees no vertices of $\mathcal{P}$ satisfies the axioms of Spherical Occlusion Diagrams.*

Proof. $S_{\mathcal{P}}$ is obviously finite. To prove that it is not empty, consider a shortest path from v to any vertex of $\mathcal{P}$; such a path must bend at the edges of $\mathcal{P}$, the first of which is visible to v.

For each arc $a \in S_{\mathcal{P}}$, let e_a be the edge of $\mathcal{P}$ whose radial projection on the sphere (partly occluded by faces of $\mathcal{P}$) contains a. Since e_a is a line segment of finite length, a must be an arc of a great circle that is shorter than a great semicircle, i.e., a geodesic arc.

Also, since e_a is an edge of a face $F \in \mathcal{P}$ that is partially visible to v, all arcs of $S_{\mathcal{P}}$ that touch the interior of a must reach it from the same side (axiom A3) and cannot continue past a (axiom A1), because such arcs correspond to edges partially hidden by F (hence, e_a must be a reflex edge of $\mathcal{P}$).

Finally, the fact that each vertex of $\mathcal{P}$ is occluded by some face translates into the property that each endpoint of each arc in $S_{\mathcal{P}}$ must lie in the interior of another arc of $S_{\mathcal{P}}$ (axiom A2). Note that here we used the fact that v does not lie on a face of $\mathcal{P}$. ◀

Note that Proposition 3 also holds (with the same proof) for any finite non-empty arrangement $\mathcal{P}$ of internally disjoint polygons, none of which is coplanar with the viewpoint.

It is known that the converse of Proposition 3 is not true, as not every SOD is the visibility map of a polyhedron [7]. There is also compelling evidence that a stronger statement holds, which we state next.

▶ **Definition 4.** *An SOD $\mathcal{D}$ is* irreducible *if no proper subset of $\mathcal{D}$ is an SOD.*

▶ **Conjecture 5.** *There is an* irreducible *Spherical Occlusion Diagram (satisfying the axioms in Definition 2) that is not the visibility map $S_{\mathcal{P}}$ of any polyhedron $\mathcal{P}$ with respect to any viewpoint v that sees no vertices of $\mathcal{P}$.*[2]

We remark that Conjecture 5 automatically extends to arbitrary arrangements $\mathcal{P}$ of disjoint polygons; actually, it is sufficient to prove Conjecture 5 for *simply connected* polyhedra. Indeed, a set of disjoint polygons $\mathcal{P}$ that gives rise to an SOD $\mathcal{D}$ with respect to a viewpoint v can easily be augmented by adding a mesh of polygons whose edges are either shared with $\mathcal{P}$ or concealed from v by polygons in $\mathcal{P}$. The resulting simply connected polyhedron gives rise to the same SOD $\mathcal{D}$ (a possible way of implementing this construction is found in [13]).

---

[2] The counterexample in [7] is not irreducible, because it is constructed by embedding a tiny non-realizable *planar* diagram into a larger SOD, for example within the eye of a swirl (refer to Section 4 for a definition of "swirl").

3   Elementary properties

We will prove some basic properties of SODs. In fact, all theorems in this section, with the only exception of Proposition 12, hold more generally for finite non-empty collections of geodesic arcs that satisfy axioms A1 and A2, but not necessarily A3.

It is immediate to prove a stronger form of axiom A2.

▶ **Proposition 6.** *Every arc in an SOD hits exactly two distinct arcs, one at each endpoint.*

Proof. Let p be an endpoint of an arc a. By axiom A2, a hits at least one arc b at p. If a hit a second arc b' at p, then p would be interior to both b and b', contradicting axiom A1.   ◀

The following statement implies that no two arcs in an SOD can hit each other.

▶ **Proposition 7.** *No two arcs in an SOD intersect in more than one point.*

Proof. Two internally disjoint arcs sharing two points cannot be shorter than great semicircles, contradicting the assumption that an SOD consists of geodesic arcs (see Figure 3).   ◀
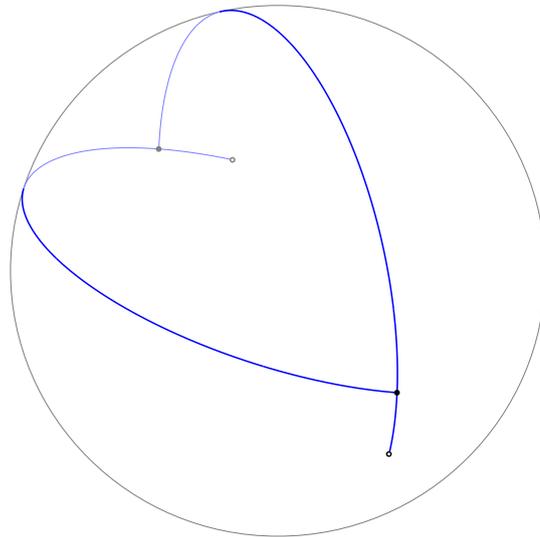
■ **Figure 3** Two non-collinear arcs can only intersect in two antipodal points.

For the next results, we describe a general construction called "monotonic walk"; refer to Figure 4. Let $\mathcal{D}$ be an SOD, and let p and p' be antipodal points on the unit sphere. A *clockwise monotonic walk around p* starting from a point $x_0 \notin \{p, p'\}$ on an arc $a_0 \in \mathcal{D}$ is a sequence of pairs (x_i, a_i), where $x_i \in a_i \in \mathcal{D}$ for all $i \geq 0$, defined inductively as follows. Assuming x_i is a point of $a_i \in \mathcal{D}$ distinct from p and p', we distinguish two cases. If the great circle containing a_i also contains p and p', then x_{i+1} is an endpoint of a_i such that the geodesic arc $x_i x_{i+1}$ contains neither p nor p' (such an endpoint must exist, or else a_i would not be shorter than a great semicircle; in the special case where x_i is an endpoint of a_i, we take $x_{i+1} = x_i$). Otherwise, the great circle containing a_i separates p from p'. Following this great circle in the clockwise direction with respect to p starting at x_i, we define x_{i+1} as the first endpoint of a_i encountered. In both cases, we let a_{i+1} be the (unique, due to Proposition 6) arc of $\mathcal{D}$ that blocks a_i at x_{i+1}.

A *counterclockwise monotonic walk* around a point is defined similarly.
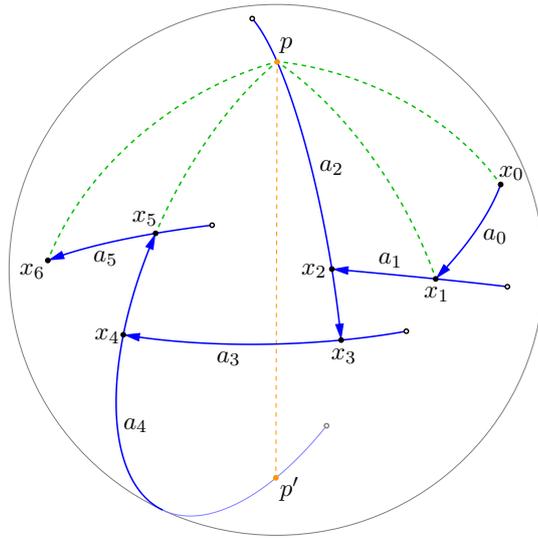
Figure 4 The initial steps of a clockwise monotonic walk around p starting from $x_0 \in a_0$.

▶ **Proposition 8.** *Given an SOD $\mathcal{D}$, the relative interior of any great semicircle on the unit sphere intersects at least one arc of $\mathcal{D}$.*

Proof. Let p and p' be the (antipodal) endpoints of a great semicircle c. Since $\mathcal{D}$ is not empty, there exists an arc $a_0 \in \mathcal{D}$. Recall that the endpoints of a geodesic arc are distinct; therefore, a_0 contains infinitely many points, and in particular it contains a point x_0 distinct from p and p'. Consider a clockwise monotonic walk $((x_i, a_i))_{i \geq 0}$ around p starting from x_0. Observe that any a_i that is an arc of a great circle through p and p' must be followed by an arc a_{i+1} that is not. Hence, the monotonic walk makes steady progress around p. Since $\mathcal{D}$ is finite, in a finite amount of steps the monotonic walk touches the interior of every great semicircle with endpoints p and p', including c.   ◀

▶ **Proposition 9.** *An SOD partitions the unit sphere into spherically convex regions.*

Proof. Let $\mathcal{D}$ be an SOD on the unit sphere S, let D be the union of the arcs in $\mathcal{D}$, and let p and q be two points in the same connected component of $S \setminus D$. We will prove that p and q are connected by a single geodesic arc that does not intersect D.

By assumption and by the finiteness of $\mathcal{D}$, there is a chain C consisting of k geodesic arcs (drawn in green in Figure 5) that connects p and q without intersecting D. Let us choose C so that k is minimum. If $k = 1$, there is nothing to prove. Thus, assume that $k \geq 2$.

Let pp' and $p'p''$ be the first two geodesic arcs of C. If p and p'' are antipodal, then pp' and $p'p''$ are collinear, and their union is a great semicircle. By Proposition 8, D intersects $pp' \cup p'p''$, which is a contradiction. Hence p and p'' are not antipodal, and there is a unique geodesic arc pp'' (drawn in orange).

Assume for a contradiction that D intersects pp'' in $x \notin \{p, p''\}$. Then, either a clockwise or a counterclockwise monotonic walk around p starting from x must intersect $pp' \cup p'p''$, which is again a contradiction (see Figure 5). Therefore, D does not intersect pp'', and we can replace pp' and $p'p''$ in C by the single geodesic arc pp''. Since this contradicts the minimality of k, we conclude that $k = 1$.   ◀

▶ **Definition 10.** *Each of the (spherically convex) regions into which the unit sphere is partitioned by an SOD is called a* tile.
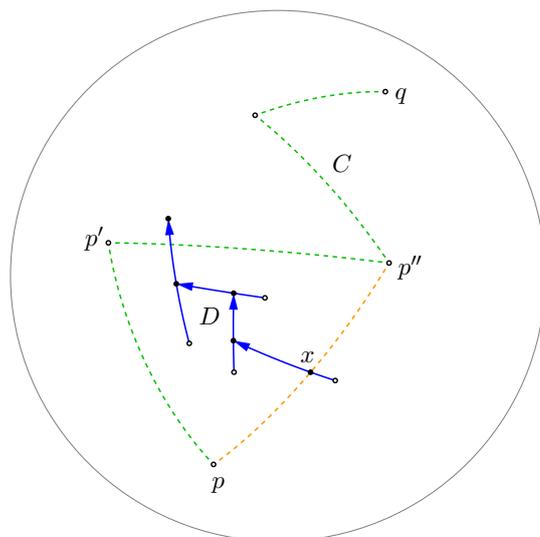
Figure 5 Proof of Proposition 9: The chain C can be simplified by connecting p and p''.

▶ **Corollary 11.** *In an SOD, no tile (including its boundary) contains two antipodal points.*

Proof. If two antipodal points p and p' were in a same tile T (or on its boundary), then by Proposition 9 there would be a great semicircle with endpoints p and p' whose interior is entirely contained in T. However, this would contradict Proposition 8. ◀

Since an SOD is a finite collection of arcs, tiles are *spherical polygons*, whose boundaries have finitely many vertices and edges.

▶ **Proposition 12.** *In an SOD, any arc coincides with an edge of a tile.*

Proof. Each arc a is incident to tiles on both sides. There are multiple tiles on the same side of a if and only if there are arcs hitting a from that side. Thus, axiom A3 implies that a cannot have multiple tiles on both sides; in other words, a side of a must have exactly one tile, and a is an edge of that tile. ◀

The *contact graph* of an SOD $\mathcal{D}$ is the undirected graph $(\mathcal{D}, \mathcal{E})$, where $\mathcal{E}$ is the set of pairs of arcs $\{a, b\} \subseteq \mathcal{D}$ such that a hits b. The following result implies that the contact graph of any SOD is 2-connected.

▶ **Proposition 13.** *Removing any one arc from an SOD and taking the union of the remaining arcs yields a connected subset of the unit sphere.*

Proof. Let $\mathcal{D}$ be an SOD, let $a \in \mathcal{D}$, and let D be the union of all arcs of $\mathcal{D}$ other than a. Our claim is equivalent to the statement that every connected component of $S \setminus D$ is simply connected, where S is the unit sphere.

Most of the connected components of $S \setminus D$ are interiors of tiles of $\mathcal{D}$, which are spherically convex (Proposition 9) and therefore simply connected. The only exceptions are the tiles incident to the interior of a, which constitute a single connected component C of $S \setminus D$. Precisely, C is the union of a and the interiors of all tiles of $\mathcal{D}$ that have an edge in a.

We will prove that C is contractible, and therefore simply connected. In fact, we will show that C deformation-retracts onto a, which is obviously contractible. Consider a tile T of $\mathcal{D}$ such that $e \subseteq a$ is an edge of T. Since both T and a are spherically convex, their

intersection is also spherically convex, hence a geodesic arc. It follows that $T \cap a = e$, and therefore $\mathring{T} \cup a$ deformation-retracts onto a, where $\mathring{T}$ is the interior of T. This induces a deformation retraction of C onto $C \setminus \mathring{T}$. By repeating the same process for all tiles with an edge in a, we conclude that C deformation-retracts onto a.   ◀

▶ **Corollary 14.** *The union of all the arcs in an SOD is a connected set.*

Proof. Immediate from Proposition 13.   ◀

It is easy to see that Proposition 13 cannot be improved, as there are SODs that are disconnected by the removal of just two arcs. Indeed, starting from any SOD, it is always possible to add a new arc a, connecting two arcs b and c, without violating the axioms of Definition 2. Now, removing b and c from this new SOD isolates a from the rest of the arcs. There are also irreducible SODs with this property, such as the one in Figure 6.
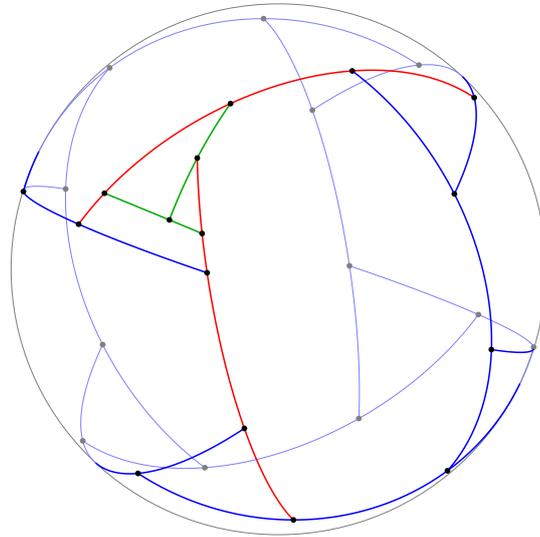
■ **Figure 6** An irreducible SOD whose contact graph is not 3-connected. Removing the two arcs in red disconnects the arrangement into a green and a blue part.

There is a very clean relationship between the number of arcs in an SOD and the number of tiles: There are two more tiles than there are arcs.

▶ **Proposition 15.** *An SOD with n arcs partitions the unit sphere into $n + 2$ tiles.*

Proof. Every endpoint of an arc of an SOD divides the arc it hits into two sub-arcs. The set of these sub-arcs induces a spherical drawing of a planar graph with k vertices and $n + k$ edges (if two arcs share an endpoint, this counts as a single vertex). Each face of this drawing coincides with a tile of the SOD. By Euler's formula, the number of faces is $(n + k) - k + 2 = n + 2$.   ◀

4   Swirls

There is a curious similarity between SODs and continuous vector fields on a sphere. According to the hairy ball theorem, "it is impossible to comb a hairy ball without creating cowlicks". Similarly, it is impossible to construct an SOD without creating "swirls", as we shall see in this section.

▶ **Definition 16.** *A* swirl *in an SOD is a cycle of arcs, each of which feeds into the next (and such that the last feeds into the first), going either all clockwise or all counterclockwise. The* degree *of a swirl is the number of arcs constituting it.*

Since no two arcs in an SOD can hit each other (Proposition 7), the minimum degree for a swirl is 3. Figure 2 shows an SOD with six clockwise swirls and six counterclockwise swirls, all of degree 3.

We will now prove some basic properties of swirls. Note that for the second time in this paper, after Proposition 12, we will be using axiom A3.

▶ **Proposition 17.** *In an SOD $\mathcal{D}$, let $\mathcal{S}$ be a swirl of degree k.*

■ *The union of the k arcs of $\mathcal{S}$ separates the unit sphere in two regions, exactly one of which is spherically convex; this region is a spherical k-gon called the* eye E *of $\mathcal{S}$.*

■ *The only points of intersection between pairs of arcs of $\mathcal{S}$ are the vertices of E.*

■ *The tiles of $\mathcal{D}$ adjacent to E are exactly k; any two such tiles are either disjoint or intersect only along a single arc of $\mathcal{S}$.*

Proof. Let $\mathcal{S} = (a_0, a_1, \ldots, a_{k-1})$, and let x_i be the unique endpoint of a_{i-1} lying in the interior of a_i, with $0 \leq i < k$ (indices are taken modulo k), as shown in Figure 7. By definition of swirl, each arc of $\mathcal{S}$ feeds into the next forming convex angles; moreover, turning in the same direction (either clockwise or counterclockwise) at every x_i yields a cycle of arcs. We conclude that following the a_i's in order traces out the boundary of a spherically convex k-gon with vertices x_0, x_1, $\ldots$, x_{k-1}. Let E be such a convex k-gon (drawn in yellow in Figure 7).

We already know that any two consecutive arcs of $\mathcal{S}$ intersect at a vertex of E and only there (Proposition 7). We will now prove that non-consecutive arcs a_i and a_j of $\mathcal{S}$ are disjoint. Since E is spherically convex, it completely lies in one of the two hemispheres bounded by a_i, and in one of the two hemispheres bounded by a_j. Thus, E lies in a spherical lune L determined by a and b, as shown in Figure 8. Clearly, a_i and a_j can intersect only at the vertices p and p' of L, which are antipodal points. Note that, because of the way the arcs of a swirl are oriented, a_i may be incident to a vertex of L, say p, but not to the other vertex p'. Likewise, a_j may be incident to p' but not to p, and thus a_i and a_j cannot intersect.

Since the arcs of $\mathcal{S}$ have no intersections away from of the boundary of E, the region of sphere external to E is not disconnected by $\mathcal{S}$, and is therefore a unique non-convex connected component.

It remains to prove that the tiles of $\mathcal{D}$ adjacent to E are exactly k and may only intersect each other along arcs of $\mathcal{S}$. Due to axiom A3, no arc of $\mathcal{D}$ lying outside of E can hit the boundary of E away from its vertices. Thus, there can be at most k tiles adjacent to E, and indeed these are the k tiles that have an a_i as an edge, for some $0 \leq i < k$ (recall that every arc coincides with an edge of a tile, due to Proposition 12).

Let us define x_i' as the point antipodal to x_i, with $0 \leq i < k$, and extend each a_i in the direction opposite to x_{i+1} until it reaches x_i', as shown in Figure 7. Our proof that non-adjacent swirl arcs are disjoint also implies that these arc extensions partition the sphere into exactly $k + 2$ regions: E, the polygon $x_0' x_1' \ldots x_{k-1}'$ (which is congruent to E and antipodal to it), and k spherical lunes, each of which entirely contains an arc a_i and has x_{i+1} and x_{i+1}' as vertices.

Note that the tile of $\mathcal{D}$ adjacent to E that has a_i as an edge must be contained in one of these k spherical lunes, namely the one whose vertices are x_{i+1} and x_{i+1}'. Thus, the pairs of tiles corresponding to non-consecutive a_i's must be disjoint, because their respective lunes

are disjoint. On the other hand, any two tiles corresponding to adjacent lunes intersect along an arc a_i. Since tiles are spherically convex, their intersection must be a sub-arc of a_i.   ◀
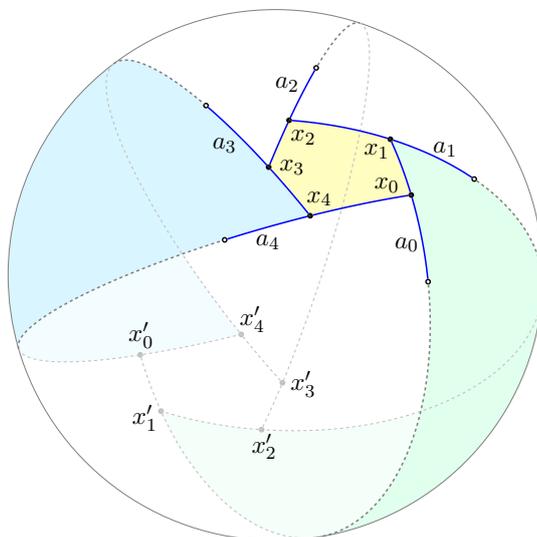
Figure 7 A counterclockwise swirl with its eye in yellow. The arcs constituting the swirl cannot intersect outside of the eye's vertices, and the two colored spherical lunes are disjoint.
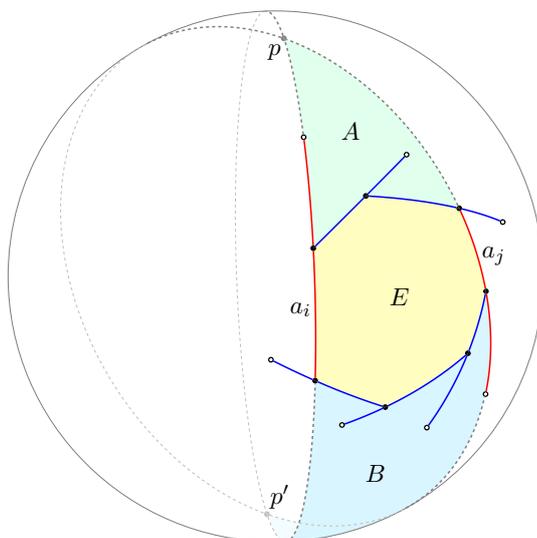
Figure 8 Any two non-consecutive arcs of a swirl are disjoint.

Observe that, in an irreducible SOD, the eye of each swirl coincides with a single tile; in general, the eye of a swirl is a union of tiles, as there may be inner arcs.

▶ **Proposition 18.** *In an SOD, if the eyes of two distinct swirls have intersecting interiors, then their boundaries are disjoint.*

Proof. Let $\mathcal{S} = (a_0, a_1, \ldots, a_{k-1})$ and $\mathcal{S}' = (b_0, b_1, \ldots, b_{k'-1})$ be two distinct swirls with eyes E and E', respectively. Note that, although the two swirls may share some arcs, not all

arcs of $\mathcal{S}$ can be arcs of $\mathcal{S}'$, and vice versa (this is a simple consequence of Proposition 7 and the definition of swirl).

Assume for a contradiction that E and E' have intersecting interiors and intersecting boundaries. Then, as we trace out the perimeter of E' following the arcs of $\mathcal{S}'$ in order, we eventually reach a b_i lying entirely in E that hits an a_j, as shown in Figure 9. Thus, $b_{i+1} = a_j$. This immediately implies that $\mathcal{S}$ and $\mathcal{S}'$ cannot be concordant, e.g., both clockwise. For otherwise, continuing to follow the arcs of $\mathcal{S}'$ clockwise we would reach $b_{i+2} = a_{j+1}$, $b_{i+3} = a_{j+2}$, etc., contradicting the fact that $\mathcal{S}$ and $\mathcal{S}'$ are distinct.

Hence, we may assume without loss of generality that $\mathcal{S}$ is a clockwise swirl and $\mathcal{S}'$ is a counterclockwise swirl. Now, continuing to trace out the perimeter of E' following the arcs of $\mathcal{S}'$ counterclockwise from $b_{i+1} = a_j$ immediately leads outside of E (see the green arrow starting at b_i in Figure 9).

We claim that once our walk around the perimeter of E' has exited E, it can no longer reach its interior, contradicting the fact that b_i is internal to E. Indeed, as soon as a counterclockwise walk reaches the boundary of E at a certain $a_{i'}$, whether coming from the side of $a_{i'-1}$ or from the opposite side,[3] it is immediately "pushed back" outside E, as illustrated by the green arrows in Figure 9. ◀
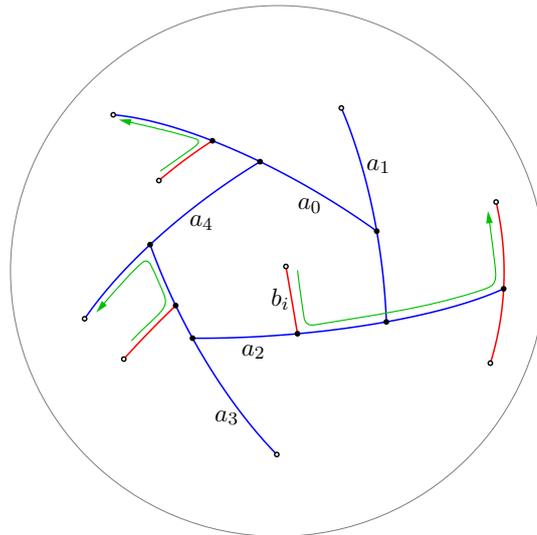

Figure 9 No two distinct swirls can have intersecting interiors and intersecting boundaries.

Proposition 18 implies that any two eyes in an SOD are either internally disjoint or one is strictly contained in the other. Observe that, if we exclude "degenerate" SODs in which some arcs may share an endpoint, the statement of Proposition 18 is strengthened to "distinct eyes have disjoint boundaries". The inclusion of degenerate SODs, however, implies that some eyes may be internally disjoint but have a common vertex.

We will now study the combinatorics of swirls within an SOD by analyzing a structure called "swirl graph".

---

[3] Note that the latter case cannot occur in an SOD, because it violates axiom A3. We chose to present a more general proof which does not use axiom A3, at the cost of slightly complicating the analysis.

▶ **Definition 19.** *The* swirl graph *of an SOD $\mathcal{D}$ is the undirected multigraph on the set of swirls of $\mathcal{D}$ having an edge between two swirls for every arc in $\mathcal{D}$ shared by the two swirls.*
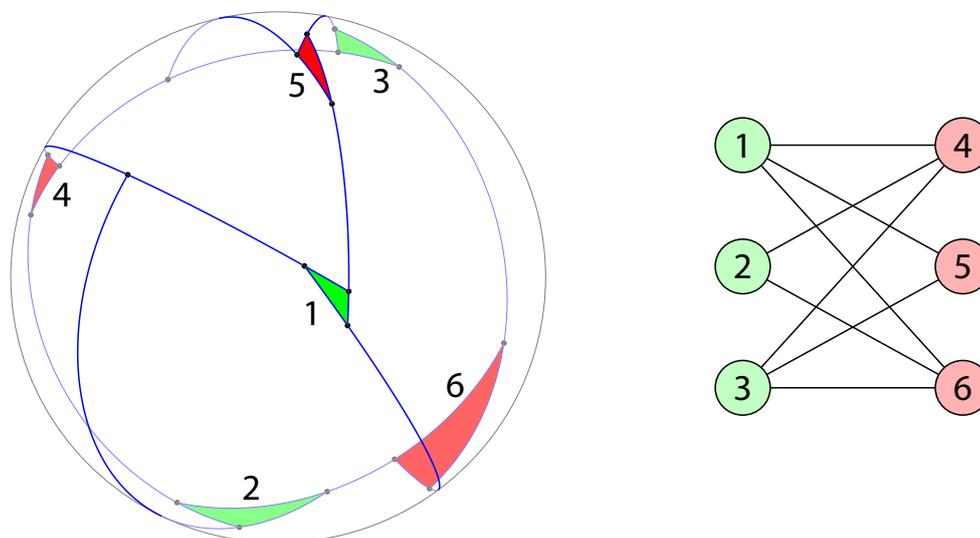
■ **Figure 10** An SOD and its swirl graph.

In Figure 10, the eyes of clockwise swirls are colored green, and the eyes of counterclockwise swirls are colored red. Observe that the swirl graph of this SOD is simple and bipartite; this is true in general, as the next theorem shows.

▶ **Theorem 20.** *The swirl graph of any SOD is a simple planar bipartite graph with non-empty partite sets.*

Proof. The swirl graph is spherical, hence planar. It is bipartite, where the partite sets correspond to clockwise and counterclockwise swirls, respectively. Indeed, if the same arc is shared by two concordant swirls (say, clockwise), then it is hit by arcs from both sides, violating axiom A3.

Figure 11 shows how to find a clockwise and a counterclockwise swirl in any SOD. For a clockwise swirl, start from any arc and follow it in any direction until it hits another arc. Then turn clockwise and follow this arc until it hits another arc, and so on. The sequence of arcs encountered is eventually periodic, and the period identifies a clockwise swirl. A counterclockwise swirl is found in a similar way.

To prove that the swirl graph is simple, assume for a contradiction that a swirl $\mathcal{S}$ shares two arcs a_i and a_j with another swirl $\mathcal{S}'$. Let E and E' be the eyes of $\mathcal{S}$ and $\mathcal{S}'$, respectively; observe that a_i and a_j are incident to both E and E'. Let α and β be the two great circles containing a_i and a_j, respectively. As proved above, $\mathcal{S}$ and $\mathcal{S}'$ must be discordant, i.e., one is a clockwise swirl and the other is a counterclockwise swirl. Thus, E and E' lie on the same side of a_i. Furthermore, since both eyes are spherically convex (Proposition 17), they lie in the same hemisphere bounded by α. For the same reason, E and E' lie in the same hemisphere bounded by β.

Therefore, there is a spherical lune L, bounded by α and β, such that both E and E' lie completely in L. Also, E separates L into two disjoint regions A and B, such that some internal points of a_i lie on the boundary of A (but not on B) and some internal points of a_j lie on the boundary of B (but not on A), as shown in Figure 8. Since the boundary of E'
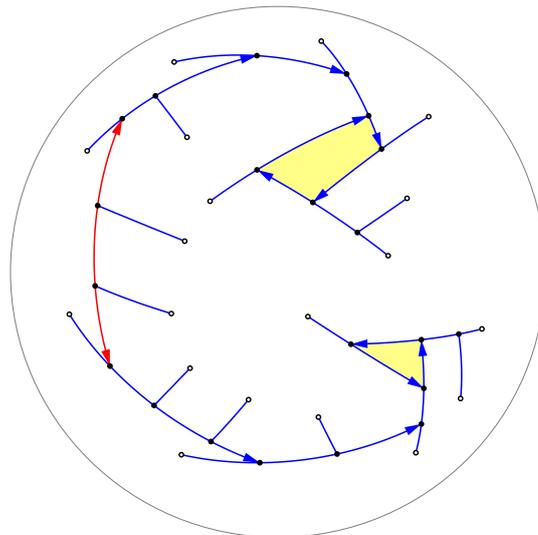
Figure 11 Finding swirls in an SOD.

connects a_i and a_j, it must intersect the boundary of E. Hence, by Proposition 18, E and E' must have disjoint interiors. Thus, as E' is bounded by a_i, it must lie entirely in A. But E' is also bounded by a_j, implying that it must lie in B, which is a contradiction. (In the special case where a_i and a_j are incident arcs, either A or B degenerates to the empty set, which of course cannot contain E'.) ◀

More is actually known about swirl graphs.

▶ **Theorem 21.** *Every SOD has at least four swirls.* ◀

This result was announced in [12], and a proof can be found in [13]. From Theorem 21, it easily follows that every SOD has at least eight arcs. On the other hand, Figure 12 shows an example of an SOD with exactly eight arcs and exactly four swirls, which is therefore minimal.

It is not yet clear whether there are SODs with only one clockwise swirl (or only one counterclockwise swirl), but there is overwhelming evidence that this is not the case.

▶ **Conjecture 22.** *Every SOD has at least two clockwise and two counterclockwise swirls.*

5 Swirling SODs

This section is devoted to a special type of SODs whose arcs always meet forming swirls. Patterns arising in these SODs are found in modular origami, globe knots, rattan balls, etc.

▶ **Definition 23.** *An SOD is* swirling *if each of its arcs is part of two swirls.*

We define the *degree* of a (swirling) SOD as the the maximum degree of its swirls. Obviously, the degree of an SOD is at least 3. An example of a swirling SOD of degree 3 is found in Figure 2; further examples of swirling SODs are in Figure 14. All of these SODs were obtained from convex polyhedra or, more precisely, from convex subdivisions of the sphere, by a general process that we call *swirlification*. As we will see in this section, there is a deep relationship between convex polyhedra and swirling SODs.
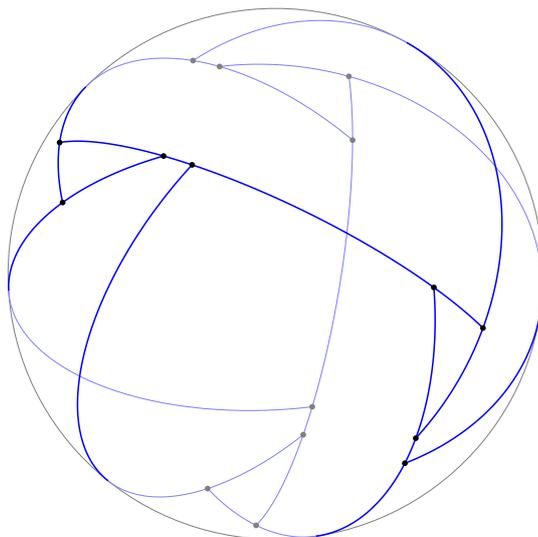
■ **Figure 12** An SOD with eight arcs and four swirls.

A *convex subdivision* of the unit sphere is a partition into spherically convex polygons called *faces* by a finite set of internally disjoint geodesic arcs called *edges*, such that no two incident edges are collinear. As usual, the graph induced by the edges is called the *1-skeleton* of the convex subdivision. Note that the visibility map of a convex polyhedron with respect to an internal point is the 1-skeleton of a convex subdivision of the unit sphere.

▶ **Definition 24.** *A* swirlable *subdivision of the unit sphere is a convex subdivision where each face has an even number of edges.*

▶ **Proposition 25.** *A convex subdivision of the unit sphere is swirlable if and only if its 1-skeleton is bipartite.*

Proof. The 1-skeleton is bipartite if and only if its has no odd cycles, which is true if and only if each face has an even number of edges.   ◀

Note that we can also obtain a swirlable subdivision of the sphere by taking the dual of a subdivision whose vertices have even degree, or by truncating it. More generally, we have the following.

▶ **Proposition 26.** *A convex subdivision of the unit sphere is swirlable if and only if its truncated dual is swirlable.*

Proof. Observe that the operations of truncation and taking the dual of a convex subdivision of the sphere preserve the convexity of the subdivision. To conclude the proof, we only have to show that a convex subdivision has even-sided faces if and only if its truncated dual does.

Figure 13 illustrates the procedure of truncating the dual of a convex subdivision $\mathcal{S}_1$ of the sphere. The edges of $\mathcal{S}_1$ are drawn in blue, and the ones in red are the edges of its dual $\mathcal{S}_2$. Note that $\mathcal{S}_2$ has a vertex in each face of $\mathcal{S}_1$, and vice versa. Moreover, the degree of each vertex of $\mathcal{S}_2$ is equal to the number of edges of the face of $\mathcal{S}_1$ containing it. Truncating $\mathcal{S}_2$ amounts to drawing the green edges and deleting the dashed portions of the red edges, as well as all red vertices; let $\mathcal{S}_3$ be the new subdivision. The last operation has two effects on the set of faces. On one hand, it doubles the edges of each face of $\mathcal{S}_2$; thus, these faces are

always even-sided. On the other hand, it adds a new face around each vertex of $\mathcal{S}_2$, with a number of edges equal to its degree. Hence, these faces are even-sided if and only if the faces of $\mathcal{S}_1$ are even-sided. ◀
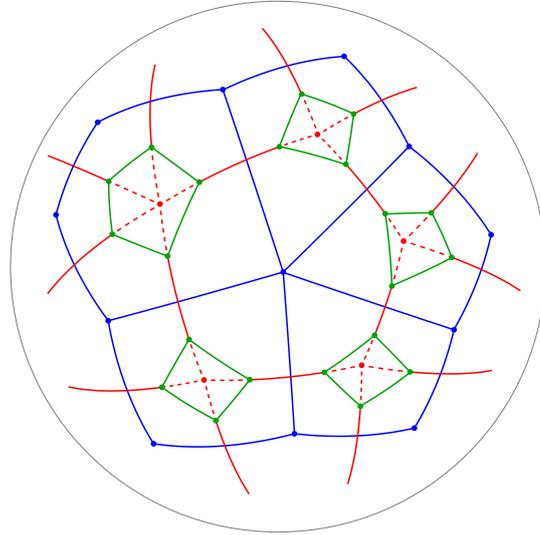
■ **Figure 13** Constructing the truncated dual of a convex subdivision of the sphere.

Given a swirlable subdivision of the sphere, the operation of turning each of its vertices into a swirl, going clockwise or counterclockwise according to the bipartition of the 1-skeleton, is called *swirlification*. As we will prove next, the swirlification operation allows us to turn any swirlable subdivision of the sphere into a swirling SOD. In fact, we will prove something more: In the case of swirling SODs of degree 3, the inverse operation is also possible, although not always trivial.

This will lead us to a characterization of the swirl graphs of swirling SODs as the 3-regular bipartite polyhedral graphs (a graph is *polyhedral* if can be realized as the 1-skeleton of a convex polyhedron). Note that, by a well-known theorem of Steinitz, these graphs can also be described as the 3-regular 3-vertex-connected planar bipartite graphs.[4] These are also known as *Barnette graphs* [6], because Barnette conjectured in 1969 that they are Hamiltonian.

We will start by illustrating the swirlification operation.

▶ **Proposition 27.** *The 1-skeleton of any swirlable subdivision of the unit sphere is isomorphic to the swirl graph of a swirling SOD.*

Proof. Let $\mathcal{S}$ be any swirlable subdivision of the unit sphere. Since each edge of $\mathcal{S}$ is a geodesic arc, we can perturb its endpoints within a small-enough ϵ-neighborhood without making them antipodal. This operation allows us to create a small swirl in lieu of each vertex of $\mathcal{S}$, while moving each edge by at most ϵ (this is possible in particular because the faces of $\mathcal{S}$ are convex, and hence do not have reflex angles). If the swirls go clockwise or counterclockwise according to the bipartition of the 1-skeleton of $\mathcal{S}$, and if ϵ is small enough, the resulting set of geodesic arcs is a swirling SOD. Since every arc in a swirling SOD is an edge of its swirl graph, the swirl graph must be G. ◀

---

[4] Steinitz's theorem states that the polyhedral graphs are precisely the 3-vertex-connected planar graphs (see [5, Chapter 13]).

Figure 14 shows two applications of Propositions 26 and 27. The two top SODs are swirlifications of a 6-trapezohedron and a truncated 6-antiprism; note that k-trapezohedron and k-antiprism are dual polyhedra. The two bottom SODs are swirlifications of a rhombic triacontahedron and a truncated icosidodecahedron; the rhombic triacontahedron and the icosidodecahedron are dual polyhedra. This process can be continued indefinitely: By repeatedly taking the truncated dual of a swirlable subdivision of the sphere and swirlifying it, one obtains an infinite family of swirling SODs.
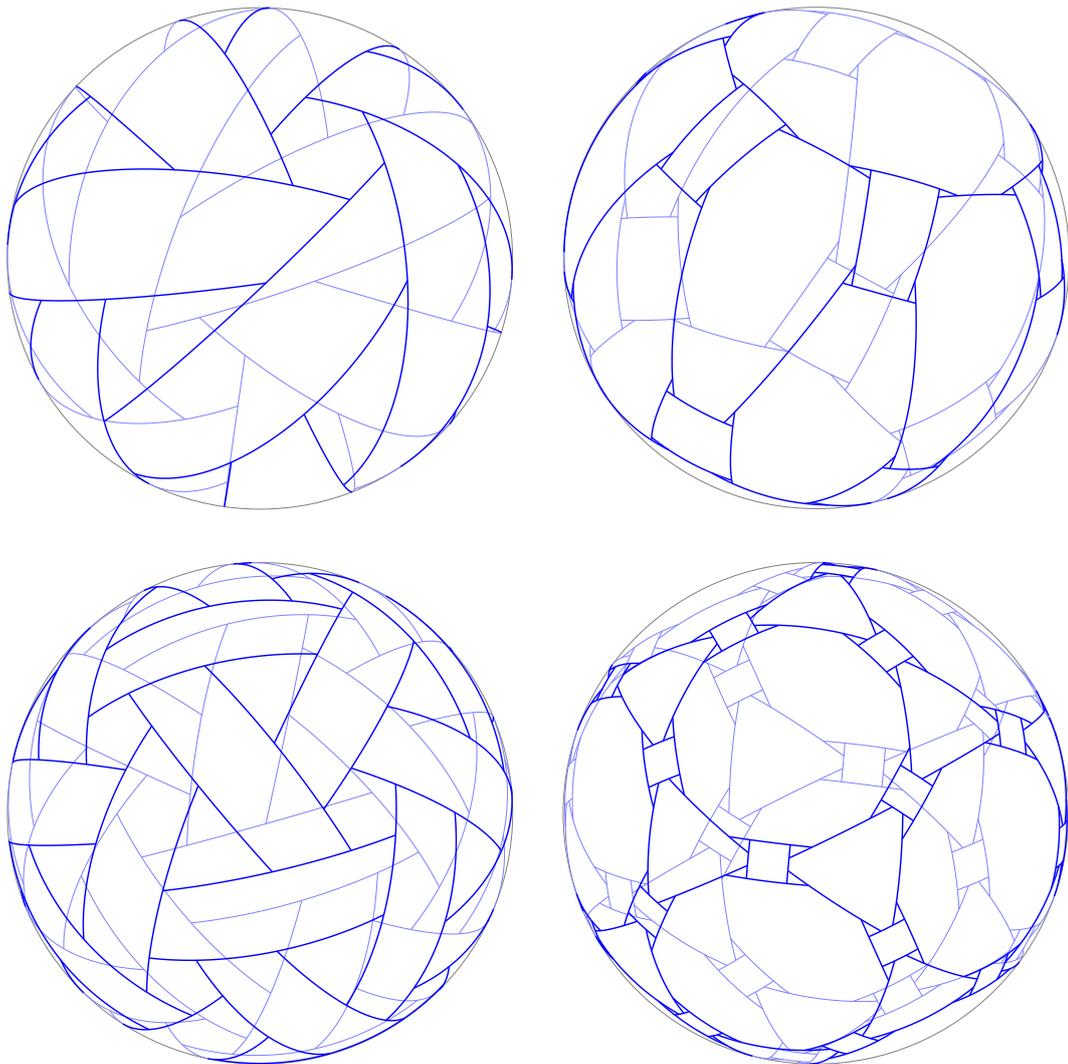
Figure 14 Examples of the swirlification method developed in Section 5 to produce swirling SODs from convex polyhedra with a bipartite 1-skeleton (or, equivalently, from swirlable subdivisions of the unit sphere). The pictures show swirling SODs resulting from a 6-trapezohedron, a truncated 6-antiprism, a rhombic triacontahedron, and a truncated icosidodecahedron, respectively.

We will now discuss the connectivity of the swirl graphs of swirling SODs. Observe that Proposition 13 immediately implies that the swirl graph of any swirling SOD is 2-edge-connected. The two following theorems strengthen this result.
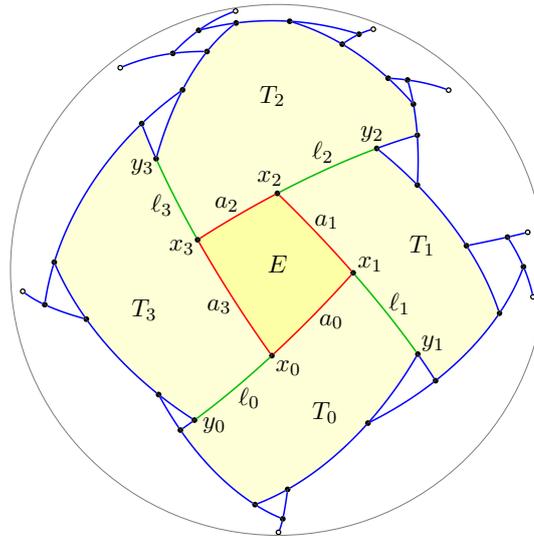
▶ **Theorem 28.** *The swirl graph of any swirling SOD is 2-vertex-connected.*

Proof. Let $\mathcal{S} = (a_0, a_1, \ldots, a_{k-1})$ be any swirl in a swirling SOD $\mathcal{D}$, and let x_i be the point where a_{i-1} feeds into a_i, with $0 \leq i < k$ (indices are taken modulo k). We define an *eyelash* ℓ_i of $\mathcal{S}$ as the maximal open geodesic arc $x_i y_i \subseteq a_i \setminus x_{i+1} x_i$ that contains no endpoints of arcs of $\mathcal{D}$, as show in Figure 15. By Proposition 17, no two arcs of $\mathcal{S}$ intersect away from x_0, $x_1, \ldots, x_{k-1}$; therefore, $\mathcal{S}$ has exactly k disjoint eyelashes $\ell_0, \ell_1, \ldots, \ell_{k-1}$.

Let D be the union of the arcs in $\mathcal{D}$. Our claim is equivalent to the statement that removing the perimeter of E and the k eyelashes of $\mathcal{S}$ from D does not disconnect D. In turn, since no arcs of $\mathcal{D}$ lie in the interior of E,[5] this is equivalent to the statement that every component of $R = (S \setminus D) \cup E \cup \ell_0 \cup \ell_1 \cup \cdots \cup \ell_{k-1}$ is simply connected, where S is the unit sphere.

Most of the connected components of R are interiors of tiles of $\mathcal{D}$, which are spherically convex, hence simply connected. The only exception is the component C (the yellow region in Figure 15) consisting of the union of the eye E, the k eyelashes of $\mathcal{S}$, and the interiors of the tiles $T_0, T_1, \ldots, T_{k-1}$ of $\mathcal{D}$ adjacent to E. By Proposition 17, any two tiles adjacent to E are either disjoint or intersect only along a single arc of $\mathcal{S}$. Thus, the tile T_i is only adjacent to T_{i-1} along (the closure of) ℓ_i and to T_{i+1} along (the closure of) ℓ_{i+1}. Also, T_i is adjacent to E along $x_i x_{i+1}$. Therefore, the common boundary between the open disk $\mathring{T}_i$ and $C \setminus \mathring{T}_i$ is $\ell_i \cup \ell_{i+1} \cup x_i x_{i+1}$, which is simply connected.

It follows that C deformation-retracts onto $C \setminus \mathring{T}_i$ for every $0 \leq i < k$. By composing these k deformation retractions, we obtain a deformation retraction of C onto $E \cup \ell_0 \cup \ell_1 \cup \cdots \cup \ell_{k-1}$, which is obviously contractible. We conclude that C is simply connected, as desired.   ◄

Figure 15 A swirl in a swirling SOD with its eye E, eyelashes ℓ_i, and neighboring tiles T_i.

▶ **Theorem 29.** *The swirl graph of any swirling SOD of degree 3 is 3-vertex-connected.*

Proof. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two swirls of a swirling SOD $\mathcal{D}$ of degree 3, and let E_1 and E_2 be their respective eyes. We define the "eyelashes" of the two swirls as in the proof of Theorem 28,

---

[5] It is intuitively true that no eye of a swirling SOD has inner arcs, and therefore E coincides with a single tile of $\mathcal{D}$. We will not formally prove it here because in Section 6 we will give a self-contained proof of a more general statement: All *uniform* SODs (hence all swirling SODs) are irreducible.

and we denote as P_1 the union of E_1 and the eyelashes of $\mathcal{S}_1$, while P_2 is the union of E_2 and the eyelashes of $\mathcal{S}_2$. Our claim is that every component of $R = (S \setminus D) \cup P_1 \cup P_2$ is simply connected, where S is the unit sphere and D is the union of the arcs of $\mathcal{D}$ (again, this relies on the fact that E_1 and E_2 have no inner arcs, as the SOD is swirling).

Most components of R are the interiors of tiles of $\mathcal{D}$, which are spherically convex, hence simply connected. Figure 16 sketches the three possible configurations for the remaining components of R; the boundaries of P_1 and P_2 are represented in red. In the first case, $\mathcal{S}_1$ and $\mathcal{S}_2$ do not share any arcs and no tile of $\mathcal{D}$ is adjacent to both E_1 and E_2. So, there is a component of R containing E_1 and a distinct component containing E_2. Reasoning as in Theorem 28, we conclude that both components are simply connected.

In the second case, $\mathcal{S}_1$ and $\mathcal{S}_2$ share an arc a with endpoints $p_1 \in E_1$ and $p_2 \in E_2$, as in Figure 16 (center). By Theorem 20, $\mathcal{S}_1$ and $\mathcal{S}_2$ cannot share more than one arc. Thus, there are two tiles A and B of $\mathcal{D}$ that are incident to the interior of a (on opposite sides of it) and adjacent to both E_1 and E_2. Note that A and B are distinct, due to Proposition 17. Since both swirls have degree 3, there is exactly one more tile C adjacent to E_1 and one more tile D adjacent to E_2. Observe that the (unique) component of R containing (the interiors of) all of these tiles is simply connected, provided that C and D are distinct tiles. So, assume for a contradiction that $C = D$, and thus both points p_1 and p_2 lie on the boundary of C. Since p_1 and p_2 are endpoints of the geodesic arc a, they are not antipodal, and therefore a is the unique geodesic arc connecting them. On the other hand, C is spherically convex, and hence the geodesic arc $p_1 p_2 = a$ must lie within C, contradicting the fact that the interior of a does not border C (cf. Proposition 17).

The only remaining case is shown in Figure 16 (right): $\mathcal{S}_1$ and $\mathcal{S}_2$ share no arcs, but there is a tile A adjacent to both E_1 and E_2. Let B_1 and B_2 be the two remaining tiles (other than A) adjacent to E_1, and let C_1 and C_2 be the two remaining tiles adjacent to E_2. There is a unique component of R containing (the interiors of) all of these tiles, and it is simply connected, provided that all the tiles involved are distinct. In particular, we only have to rule out the case that $B_i = C_j$ for $i, j \in \{1, 2\}$. Since both swirls have degree 3, there is an arc a of $\mathcal{S}_1$ incident to both A and B_i, and there is an arc a' of $\mathcal{S}_2$ incident to both A and C_j. Note that a and a' are distinct arcs, because $\mathcal{S}_1$ and $\mathcal{S}_2$ share no arcs. Assume for a contradiction that $B_i = C_j$. Then, the two tiles A and B_i, which are both adjacent to E_1, are also incident along the two arcs a and a', which contradicts Proposition 17. We conclude that all components of R are simply connected. ◀
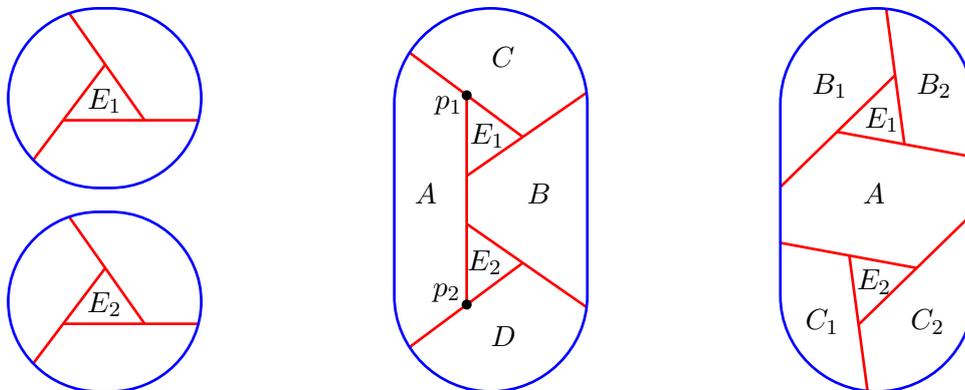
Figure 16 Sketch of the possible configurations of two swirls in a swirling SOD of degree 3.

We remark that Theorems 28 and 29 cannot be improved, as there are swirling SODs of degree 4 whose swirl graph is not 3-vertex-connected, such as the one in Figure 17. Similar counterexamples of any degree greater than 3 can be constructed, as well.
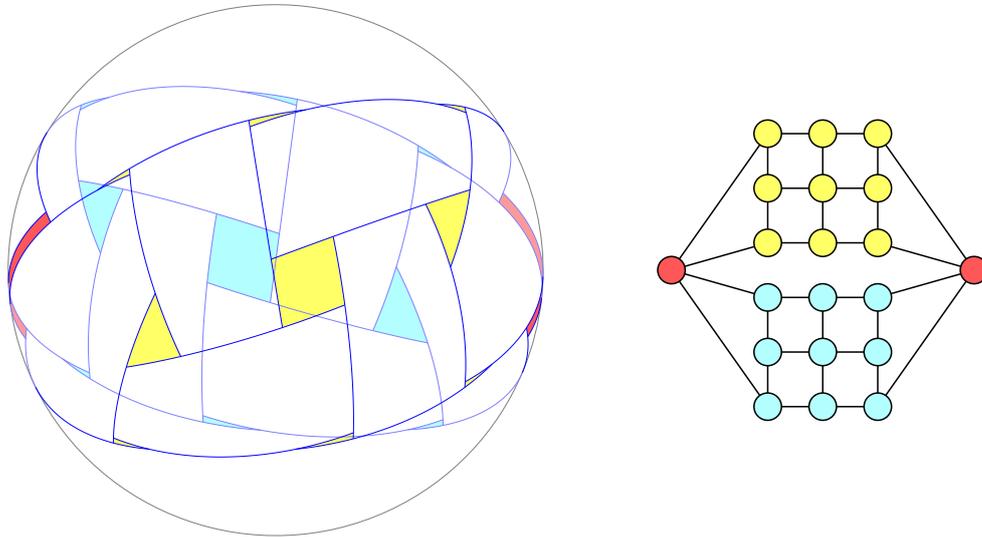
Figure 17 A swirling SOD of degree 4 (left) whose swirl graph is not 3-vertex-connected (right). Deleting the two red swirls disconnects the swirl graph into a yellow and a teal component, representing the swirls on the front side and on the back side of the sphere, respectively.

As a corollary to Theorem 29, we can now characterize the swirl graphs of swirling SODs of degree 3.

▶ **Corollary 30.** *The swirl graphs of swirling SODs of degree 3 are precisely the Barnette graphs.*

Proof. Let G be a Barnette graph, i.e., a 3-regular 3-connected planar bipartite graph. We will construct a swirling SOD of degree 3 whose swirl graph is G. By Steinitz's theorem, since G is 3-connected and planar, there is a convex polyhedron $\mathcal{P}$ whose 1-skeleton is isomorphic to G. Radially projecting the 1-skeleton of $\mathcal{P}$ onto a unit sphere whose center is internal to $\mathcal{P}$ yields a convex subdivision $\mathcal{S}$ of the sphere whose 1-skeleton is isomorphic to G. Since G is bipartite, Proposition 25 implies that $\mathcal{S}$ is swirlable. Finally, due to Proposition 27, the swirlification of $\mathcal{S}$ yields a swirling SOD whose swirl graph is G. Moreover, since G is 3-regular, all swirls of this SOD must have degree 3.

Conversely, let G be the swirl graph of a swirling SOD of degree 3. By Theorem 20, G is planar and bipartite; by Theorem 29, G is 3-connected. Also, since the SOD is swirling, each of its swirls has a number of neighbors in G equal to its degree. As the SOD has degree 3, and no swirl in an SOD can have degree lower than 3, it follows that all swirls in the SOD have degree 3, and the swirl graph G is 3-regular. We conclude that G is a Barnette graph. ◄

An easy consequence of Steinitz's theorem and Corollary 30 is that the swirl graphs of swirling SODs of degree 3 are precisely the 1-skeletons of swirlable subdivisions of the sphere whose vertices have degree 3. In other words, the converse of Proposition 27 is true, and the swirlification operation can be reversed, provided that all swirls have degree 3. However, the reversal process is not always trivial, and this explains why our proof is not direct but relies on Steinitz's theorem, which is a deep result in polyhedral combinatorics.

For instance, one might be tempted to reverse the swirlification operation by simply "contracting" the eye of each swirl of a swirling SOD of degree 3 into an interior point in order to obtain a swirlable subdivision of the sphere. Unfortunately, this is not always possible, and Figure 18 shows a counterexample consisting of a swirling SOD of degree 3 with a 3-fold rotational symmetry around the center of a large triangular eye E. Suppose we replace the eye of each swirl with an interior point and move the arcs of the SOD accordingly to form a subdivision of the sphere. No matter how an interior point $p \in E$ is chosen, the subdivision always has a non-convex face (the perimeter of such a face is drawn in red in Figure 18), hence it is not a swirlable subdivision.
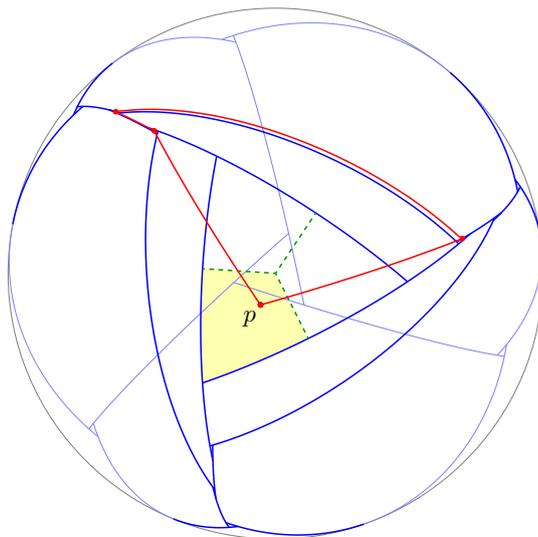
Figure 18 A swirling SOD of degree 3 where contracting the eye of each swirl into an interior point necessarily yields a non-convex subdivision of the sphere.

6   Uniform SODs

We now turn to a class of SODs that generalizes the swirling ones.

▶ **Definition 31.** *An SOD is* uniform *if every arc blocks the same number of arcs.*

▶ **Proposition 32.** *For an SOD $\mathcal{D}$, the following are equivalent.*
1. *$\mathcal{D}$ is uniform.*
2. *Every arc of $\mathcal{D}$ blocks exactly two arcs.*
3. *Every arc of $\mathcal{D}$ blocks at least two arcs.*
4. *Every arc of $\mathcal{D}$ blocks at most two arcs.*

Proof. By Proposition 6, each arc hits exactly two distinct arcs. Hence, each arc blocks two arcs on average. Thus, if all arcs blocks the same number of arcs, this number must be two. For the same reason, if every arc blocks at most two arcs (or at least two arcs), it must block exactly two arcs.   ◀

▶ **Proposition 33.** *Every uniform SOD is irreducible.*

Proof. Let $\mathcal{D}$ be a uniform SOD, and assume that there is a proper subset of arcs $\mathcal{D}' \subset \mathcal{D}$ that is itself an SOD. By Corollary 14, $\mathcal{D}$ is connected; thus, removing arcs from $\mathcal{D}$ causes

some arcs to block fewer than two arcs. Since $\mathcal{D}$ is uniform, it follows that the arcs of $\mathcal{D}'$ block fewer than two arcs on average, contradicting Proposition 6. ◄

▶ **Corollary 34.** *In a uniform SOD, the eye of each swirl coincides with a single tile.*

Proof. If the interior of the eye of a swirl contains some arcs, then such arcs can be removed without violating the SOD axioms. Hence, such an SOD is not irreducible, and by Proposition 33 it cannot be uniform. ◄

▶ **Theorem 35.** *Every swirling SOD is uniform.*

Proof. In a swirling SOD, each arc a is part of two distinct swirls. By Theorem 20, these two swirls share no arcs other than a, and hence a must block one arc from each of them. Therefore, every arc in a swirling SOD blocks at least two arcs, and by Proposition 32 the SOD is uniform. ◄

The converse of Theorem 35 is not true in general, as Figure 19 shows.
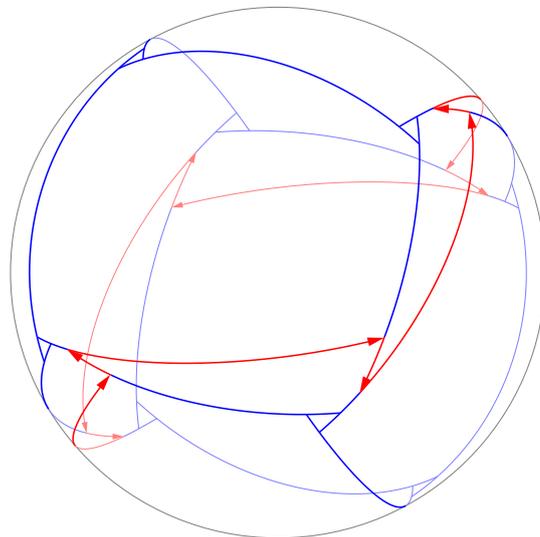
■ **Figure 19** A uniform SOD that is not swirling.

▶ **Definition 36.** *An endpoint of an arc of an SOD is called a* swirling vertex *if it is incident to the eye of a swirl. A walk on an SOD is* non-swirling *if it does not touch swirling vertices and, whenever it touches the interior of an arc, it follows it until it reaches one of its endpoints, without touching any other arc along the way. A cyclic non-swirling walk is called a* non-swirling cycle.

Observe that there is a non-swirling cycle that covers all the non-swirling vertices of the SOD in Figure 19 (drawn in red). As the next theorem shows, this is not a coincidence.

▶ **Theorem 37.** *In a uniform SOD where no two arcs share an endpoint, all non-swirling vertices are covered by disjoint non-swirling cycles.*

Proof. Suppose that no two arcs of a uniform SOD share an endpoint. Consider a non-swirling walk W terminating at a non-swirling vertex x_i, endpoint of an arc a_i, as Figure 20 shows. We will prove that W can be extended to a longer non-swirling walk in a unique way.
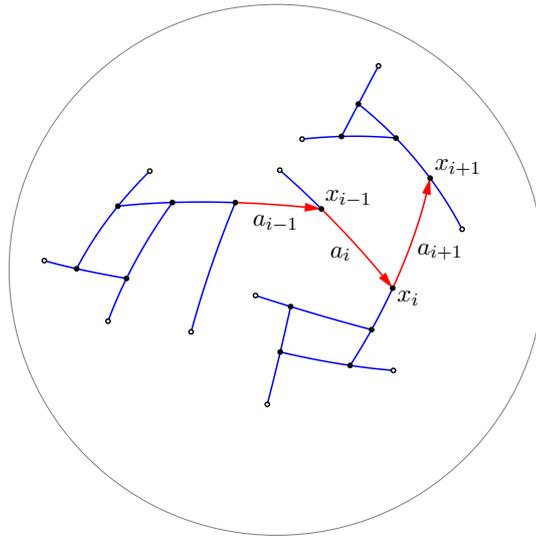
Figure 20 Proof of Theorem 37: A non-swirling walk can be extended indefinitely.

Let a_{i+1} be the arc that blocks a_i at x_i. Since exactly one arc other than a_i hits a_{i+1}, and it does so at an endpoint other than x_i, there is exactly one endpoint of a_{i+1}, say x_{i+1}, that can be reached from x_i without touching any arc other than a_{i+1}.

By definition of non-swirling walk, x_{i+1} can be used to extend W if and only if it is a non-swirling vertex. However, if x_{i+1} were incident to a swirl's eye E, then an arc of that swirl would either hit a_{i+1} between x_i and x_{i+1}, contradicting the fact that a_{i+1} blocks exactly two arcs, or it would hit a_{i+1} on the other side of x_i, implying that E contains the arc a_i in its interior, which contradicts Corollary 34.

Hence, W can be extended uniquely to a non-swirling walk. By a similar reasoning, we argue that W can also be uniquely extended backwards to a non-swirling walk. Thus, W is part of a unique non-swirling cycle. Now we conclude the proof by inductively repeating the same argument with any remaining non-swirling vertices.   ◀

It is straightforward to extend Theorem 37 to "degenerate" uniform SODs where arcs may share endpoints. However, in this case the portions of arcs that do not contribute to swirls are only partitioned into edge-disjoint cycles (as opposed to vertex-disjoint cycles).

We remark that we can construct uniform SODs with any number of unboundedly long non-swirling cycles. An example with two non-swirling cycles is shown in Figure 21.

7   Conclusions

We introduced the theory of Spherical Occlusion Diagrams and studied their basic properties, while also discussing some applications to visibility-related problems in discrete and computational geometry.

Although we strongly believe Conjecture 5 to be true, i.e., not all irreducible SODs can be realized as polyhedra, a related and more subtle question can be asked, inspired by previous work on weaving patterns [1, 11]. Namely, whether for every SOD $\mathcal{D}$ there is a *combinatorially equivalent* SOD $\mathcal{D}'$ and a polyhedron $\mathcal{P}$ such that $\mathcal{D}' = S_{\mathcal{P}}$. In other words, does every class of combinatorially equivalent SODs contain a realizable instance?
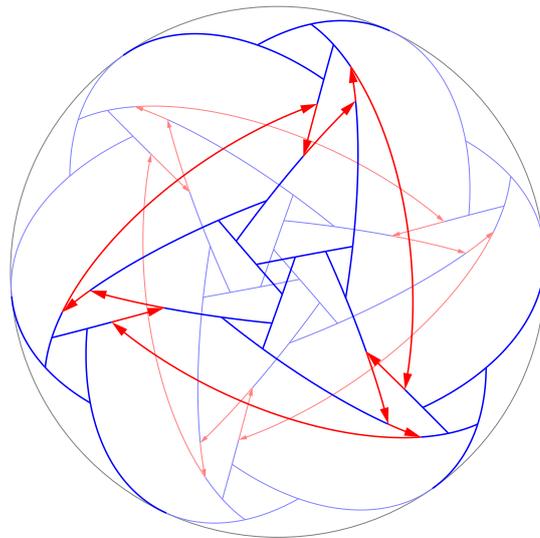
■ **Figure 21** A uniform SOD with two non-swirling cycles.

We have introduced three noteworthy families of SODs: irreducible, uniform, and swirling. We proved that all swirling SODs are uniform, and all uniform SODs are irreducible. In the case of swirling SODs of degree 3 (i.e., where all swirls consist of exactly three arcs), we were able to characterize their swirl graphs as the Barnette graphs, that is, the 3-regular 3-connected planar bipartite graphs (Corollary 30). It remains an open problem to characterize the swirl graphs of other families of SODs. Theorems 20, 28, and 37 are steps in this direction.

It is somewhat surprising that there is a wealth of swirling SODs whose swirl graphs are not bipartite polyhedral graphs (cf. Figure 17). Nonetheless, it is natural to wonder if the graphs in these two classes share any notable properties, such as the one established by Theorem 28 (i.e., being 2-vertex-connected).

Another feature shared by bipartite polyhedral graphs and the swirl graphs of swirling SODs is that all such graphs have vertices of degree 3.[6] In turn, this immediately implies that any swirling SOD has a swirl of degree 3. This contrasts with the fact that there are (non-swirling) SODs containing only swirls of arbitrarily large degree.

Theorem 37 reveals a structural connection between swirling and uniform SODs, which might be taken a step further. It seems to be possible to systematically transform any uniform SOD into a swirling SOD by "sliding" the endpoints of some arcs along other arcs and "merging" coincident arcs. Making this observation rigorous is left as a direction for future work.

More generally, we may wonder which SODs can be transformed into swirling ones by sequences of elementary operations on arcs (defining suitable "elementary operations" is in itself an open problem). The SOD in Figure 12 shows that the question is not trivial. Indeed, this is the unique configuration of any SOD with eight or fewer arcs; since the SOD itself is not swirling, it cannot be transformed into a swirling one by means of operations that only rearrange or merge arcs.

Tóth pointed out an interesting similarity between SODs and *one-sided rectangulations*,

---

[6] It is an elementary consequence of Euler's formula that any planar graph contains a triangle or a vertex of degree at most 3. Thus, a bipartite planar graph must have a vertex of degree at most 3, as it cannot contain triangles.

which are subdivisions of an axis-aligned rectangle into axis-aligned rectangles such that every maximal inner edge coincides with a side of one of the rectangles.[7] It is known that any one-sided rectangulation is *area-universal*, i.e., any assignment of areas to its rectangles can be realized by a combinatorially equivalent rectangular layout [4, 8]. We wonder if a similar property holds for SODs, and to what extent the theory of one-sided rectangulations can provide insights for the study of SODs.

References

1 Saugata Basu, Raghavan Dhandapani, and Richard Pollack. On the realizable weaving patterns of polynomial curves in $\mathbb{R}^3$. In *Proceedings of the 12th International Symposium on Graph Drawing (GD 2004)*, pages 36–42, 2004.

2 Nadia M. Benbernou, Erik D. Demaine, Martin L. Demaine, Anastasia Kurdia, Joseph O'Rourke, Godfried T. Toussaint, Jorge Urrutia, and Giovanni Viglietta. Edge-guarding orthogonal polyhedra. In *Proceedings of the 23rd Canadian Conference on Computational Geometry (CCCG 2011)*, pages 461–466, 2011.

3 Javier Cano, Csaba D. Tóth, Jorge Urrutia, and Giovanni Viglietta. Edge guards for polyhedra in three-space. *Computational Geometry: Theory and Applications*, 104:101859, 2022.

4 David Eppstein, Elena Mumford, Bettina Speckmann, and Kevin Verbeek. Area-universal and constrained rectangular layouts. *SIAM Journal on Computing*, 41(3):537–564, 2012.

5 Branko Grünbaum. *Convex polytopes*. Springer-Verlag New York, Inc., second edition, 2003.

6 Jochen Harant. A note on Barnette's conjecture. *Discussiones Mathematicae Graph Theory*, 33(1):133–137, 2013.

7 Kimberly Kokado and Csaba D. Tóth. Nonrealizable planar and spherical occlusion diagrams. In *Proceedings of the 24th Japan Conference on Discrete and Computational Geometry, Graphs, and Games (JCDCGGG 2022)*, pages 60–61, 2022.

8 Arturo Merino and Torsten Mütze. Efficient generation of rectangulations via permutation languages. In *37th International Symposium on Computational Geometry (SoCG 2021)*, volume 189, pages 54:1–54:18, 2021.

9 Joseph O'Rourke. *Art gallery theorems and algorithms*. Oxford University Press, 1987.

10 Joseph O'Rourke. Visibility. In Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, editors, *Handbook of discrete and computational geometry*, chapter 33, pages 875–896. CRC Press, 2017.

11 Jànos Pach, Richard Pollack, and Emo Welzl. Weaving patterns of lines and line segments in space. *Algorithmica*, 9(6):561–571, 1993.

12 Csaba D. Tóth, Jorge Urrutia, and Giovanni Viglietta. Minimizing visible edges in polyhedra. In *Proceedings of the 23rd Thailand-Japan Conference on Discrete and Computational Geometry, Graphs, and Games (TJCDCGGG 2020+1)*, pages 70–71, 2021.

13 Csaba D. Tóth, Jorge Urrutia, and Giovanni Viglietta. Minimizing visible edges in polyhedra. *arXiv:2208.09702 [cs.CG]*, pages 1–19, 2022.

14 Giovanni Viglietta. Optimally guarding 2-reflex orthogonal polyhedra by reflex edge guards. *Computational Geometry: Theory and Applications*, 86:101589, 2020.

15 Giovanni Viglietta. A theory of spherical diagrams. In *Proceedings of the 34th Canadian Conference on Computational Geometry (CCCG 2022)*, pages 306–313, 2022.

---

[7] This property is akin to our Proposition 12, and the term "one-sided" is reminiscent of our axiom A3.