Minkowski Sums and Spherical Duals

John M. Sullivan
Technische Universität Berlin
Institut für Mathematik, MA 3–2
Straße des 17. Juni 136
DE–10623 Berlin
Email: jms@isama.org

Abstract

At Bridges 2001, Zongker and Hart [8] gave a construction for "blending" two polyhedra using an overlay of dual spherical nets. The resulting blend, they noted, is the Minkowski sum of the original polyhedra. They considered only a restricted class of polyhedra, with all edges tangent to some common sphere. This note defines spherical duals of general convex polyhedra and proves that the Zongker/Hart construction is always valid. It can be used visually, for instance, to "morph" from any polyhedron to any other.

Polyhedra and their spherical duals

The notion of dual convex polyhedra, like the cube and octahedron, or the dodecahedron and icosahedron, is familiar. The faces of the polyhedron P correspond to vertices of its dual and vice versa. The combinatorics are thus clear, but in general (moving away from the Platonic solids) it is not always clear what geometry to give the dual. Indeed, most useful for us will be a dual which is itself not a convex polyhedron, but instead a network drawn on the surface of a sphere. We still consider it a dual of the original polyhedron P because it does have the dual combinatorics: a node for each face of P, an arc for each edge of P, and a region on the sphere for each vertex.

The nodes of this spherical dual are easy to find: each face f of P has an outward unit normal vector ν_f. Since ν_f is a unit vector in space, it can be viewed as a point on the unit sphere. We can think of this correspondance as follows: put a flashlight down on P such that its base rests flat on f. Its beam will then shine outwards in the normal direction, and it will hit the celestial sphere in the point ν_f.

Now consider an edge e of P. It lies between some pair of faces f and f' of P. If our flashlight is on f, and we tip it slowly across the edge e towards f', the beam will trace out an arc in the celestial sphere. This will be the geodesic or great-circle arc η_e from ν_f to $\nu_{f'}$. Suppose d_e is the edge vector of e (the difference between the endpoints of e). Then the arc η_e lies exactly in the great circle perpendicular to d_e. (We can check this as follows: the face normals ν_f and $\nu_{f'}$, being the endpoints of η_e, both lie on this circle, but they are also both perpendicular to d_e.) The length of the arc η_e is exactly the exterior dihedral angle of P along e (the angle between ν_f and $\nu_{f'}$).

The nodes and arcs we have described form the network on the sphere that we call the *spherical dual* $\widetilde{P}$ of P. (By analogy to smooth surfaces, where the normal vector ν is given by the so-called Gauss map, this spherical dual is sometimes also called the Gauss image of P.) The network $\widetilde{P}$ cuts the sphere into regions corresponding to the vertices v of P. Indeed, the region ρ_v associated to a vertex v is the one bounded by the arcs η_e corresponding to the edges e incident to v.

Let us recall the definition of *supporting plane*. A supporting plane through a point p on P is a plane such that all of P lies to one side (or in the plane). At a point within a face f, the unique supporting plane is

the one with normal ν_f. At any point along an edge e, there is a one-parameter family of supporting planes; their normals are the points along the arc η_e. At a vertex v there is a two-parameter family of supporting planes: holding our flashlight at v, we can tip it to be perpendicular to any of these planes. The normal directions (in which the flashlight then shines) fill out the region ρ_v. The area of this region ρ_v of $\widetilde{P}$ is what one might call the *exterior solid angle* of P at v or the *Gauss curvature* of P at v.

Note that, by the combinatorial duality, an n-gon face f of P corresponds to a n-valent node ν_f in $\widetilde{P}$: indeed there are exactly n ways to tip our flashlight off f. Similarly, if k edges e_i meet at the vertex v of P, the the region ρ_v on the sphere has k sides, namely the k arcs η_{e_i}.

We can consider two ways in which a polyhedron can degenerate to be lower-dimensional. First, a planar n-gon in space, with its two sides thought of as two faces (with opposite normals), forms a *dihedron*; its spherical dual has two antipodal nodes $\pm\nu_f$, connected by n arcs. Second, a line segment in space has no faces but has an edge with vector, say, d_e; its spherical dual has no nodes, but consists of the entire great circle perpendicular to d_e.
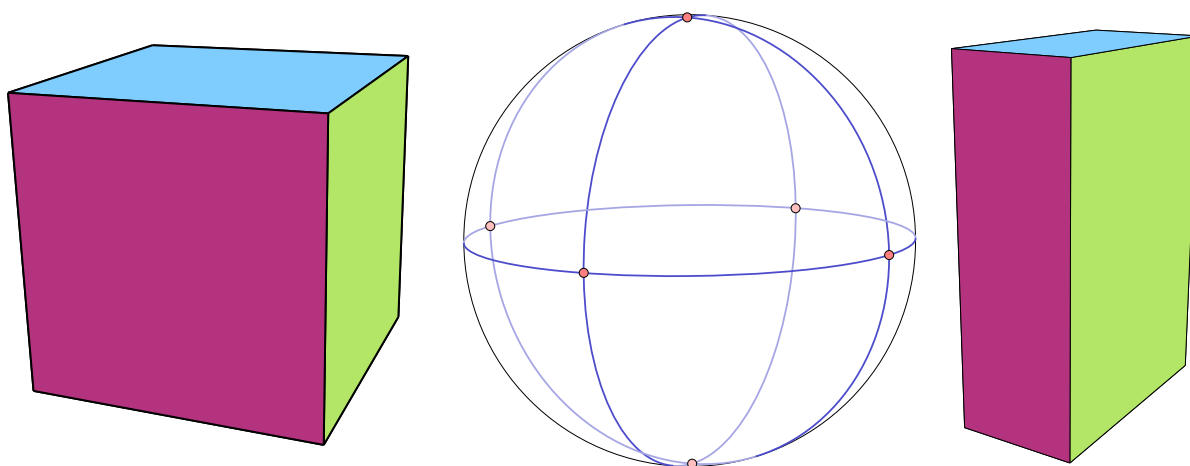

Figure 1: The spherical dual of a cube (left) is a network of three perpendicular great circles (center) which could be called a spherical octahedron. (Our spherical figures were drawn with the program "Spherical Easel" [1].) The six nodes of the network are the normals to the six cube faces; the twelve arcs correspond to the cube edges; these cut the sphere into eight congruent triangular regions corresponding to the cube vertices. The same network is also the dual of any rectangular parallelepiped (right). We can introduce labels on the arcs of the dual, recording the edge lengths of the original polyhedron, to distinguish these possibilities.

Recovering a polyhedron from its spherical dual

As a very simple example, the cube C and its dual $\widetilde{C}$ are shown in Figure 1. The dual consists of three great circles on the sphere, in the coordinate planes. These meet at six nodes (the north and south poles plus four along the equator) corresponding to the the cube's face normals. Just as the usual dual of a cube is an octahedron, this network $\widetilde{C}$ could be called as a spherical octahedron; indeed it is the radial projection of a regular octahedron to its circumscribed sphere.

Given this dual $\widetilde{C}$, can we reconstruct the original cube C? We know that if any polyhedron P has dual $\widetilde{C}$, then P must have six faces, with opposite pairs parallel to the coordinate planes. But any (axis-aligned) rectangular parallelepiped satisfies this conditions, and will indeed have this same dual network $\widetilde{C}$.

In general, given a spherical network N with convex regions, we can look for polyhedra P with dual N. Any such polyhedron will be the intersection of halfspaces H_ν, one for each node ν of N. We know that H_ν will be bounded by a plane p_ν with unit normal ν, but the location of this plane is not usually uniquely

determined. Given a polyhedron P, we can move its faces inwards or outwards a bit, keeping them parallel. As long as we don't move any face so far that the combinatorics of P changes, the result is new polyhedron with the same spherical dual $\widetilde{P}$.

Note, however, that there are also rigid examples, like the octahedron O of Figure 2. It is uniquely determined (up to homothety) by its spherical dual $\widetilde{O}$, since its combinatorics would change immediately, no matter how little we moved any face.
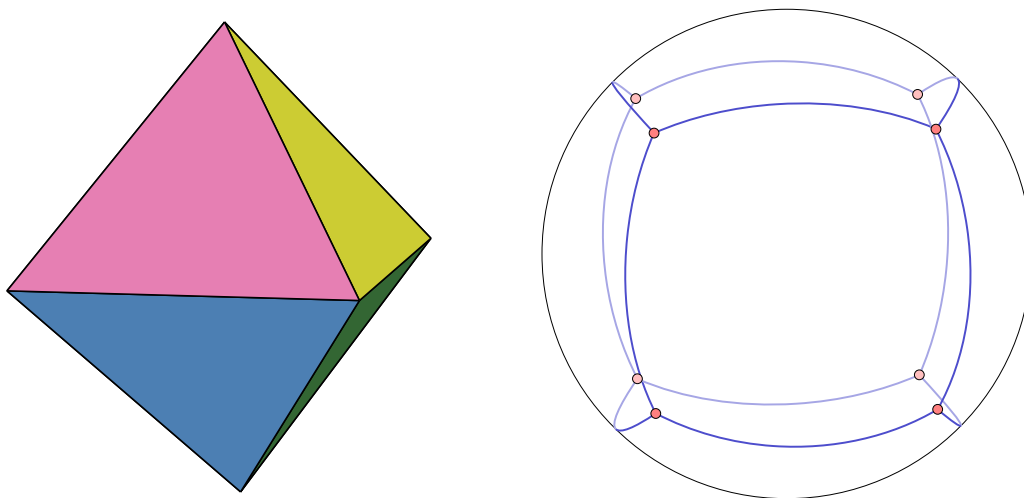

Figure 2: The octahedron (left) is uniquely determined by its spherical dual (right), since moving any face inwards or outwards would change the combinatorics. The dual $\widetilde{O}$, a spherical cube, has eight nodes (in the body-diagonal directions) and twelve arcs, dividing the sphere into six congruent square regions. Here, the only legal edge-labelings use equal labels on all twelve edges.

In general, though, to recover P from $\widetilde{P}$ we need some further information. One possibility would be to record, for each node ν_f, the distance from the origin to the plane of the face f. This would make the spherical dual carry exactly the information of the so-called *support function* of P. That is, for any direction we would know the distance from the origin to the supporting plane to P in that direction. It is well-known that any convex body is determined by its support function. (See for instance [7].) Given a spherical network, the various polyhedra with that dual correspond to the different labelings of the nodes. An n-faced polyhedron P is part of a large family with dual $\widetilde{P}$; any small adjustment of the n labels on $\widetilde{P}$ would lead to a particular member of this family.

Recent work of Fogel and Halperin [3] develops an exact algorithm for computing Minkowski sums based on spherical duals. (Actually, for computational efficiency, they project the dual radially onto a cube.) In order to be able to recover polyhedra from the dual networks, they store quite redundant information: namely, for each region of the dual, the three coordinates of the original vertex in space.

Zongker and Hart [8] suggest yet another way to encode the extra information needed to determine P from $\widetilde{P}$. They record, for each arc η_e, the length ℓ_e of the corresponding edge e. Since a polyhedron has more edges than it has faces, this leaves us with more labels than we need. That is, not every edge-labeling on $\widetilde{P}$ corresponds to a polyhedron; instead there are necessary conditions on these labels. However, we will show that certain constructions (in particular the overlay of two labeled dual nets) always do give correct labelings that correspond to polyhedra.

To understand the conditions on the labels, think first about a single arc η_e from ν_f to $\nu_{f'}$, with label ℓ_e. It must correspond to an edge e of length ℓ_e in the direction perpendicular to ν_f and $\nu_{f'}$. That is, the edge vector d_e is known to be the vector of length ℓ_e in the direction of the cross product $\nu_f \times \nu_{f'}$. Note that

this direction lies tangent to the sphere at ν_f, perpendicular to the direction in which the arc η_e leaves the node ν_f.

Of course, because the face f is a closed polygon, the edge vectors d_e around f must sum to 0. Equivalently, think of the weights ℓ_e as tensions in the arcs η_e. The closure condition around f becomes the following condition at ν_f: the tensions in all the arcs η_e sum to 0. (Since this is a vector equation in the tangent plane at ν_f, it really represents two linear conditions on the incident labels there.)

Minkowski sums

Given two polyhedra A and B, their Minkowski sum $A + B := \{a + b : a \in A, b \in B\}$ is again a polyhedron. It has faces parallel to the faces of the original polyhedra, and additional faces which are parallelograms generated by one edge from A and one from B.

As above, we will allow a segment in space to count as a degenerate polyhedron. Then the Minkowski sum of two segments is a parallelogram, the sum of three (linearly independent) segments is a parallelepiped, and in general the sum of k segments is a special kind of polyhedron called a zonohedron. Unless two of the segments are collinear, each edge of the zonohedron is parallel and equal in length to one of the original segments. Unless three of the segments are coplanar, the faces of the zonohedron are parallelograms. Figure 3 shows two zonohedra which have the same duals as certain Archimedean solids.
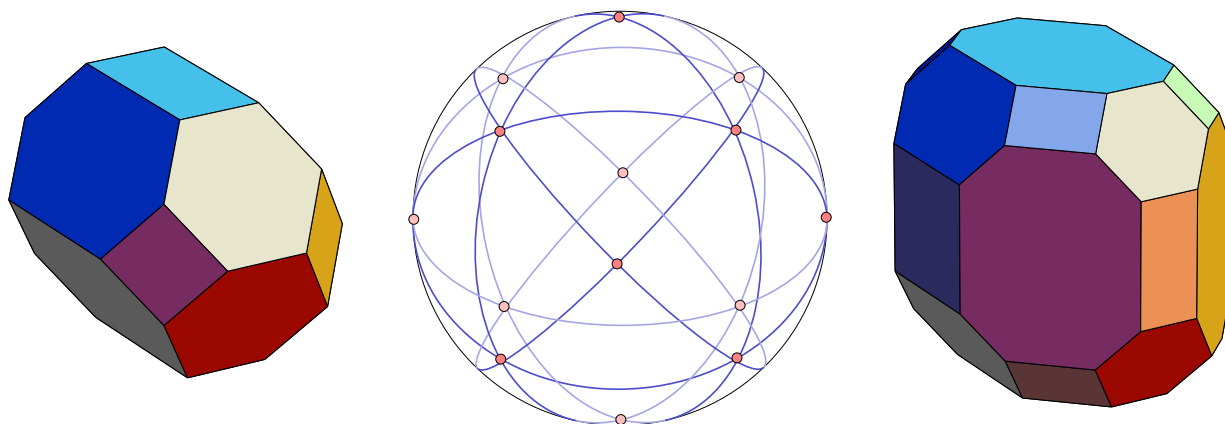

Figure 3: The zonohedron T (left) is a "stretched" truncated octahedron. Its spherical dual (center) is obtained by extending the arcs of the spherical cube $\widetilde{O}$ in Figure 2 until we have six full great circles. The truncated octahedron would have equal weights on all dual arcs; we obtain the stretched version T by varying the weights. The stretched truncated cuboctahedron (right), also a zonohedron, is the Minkowski sum of T with the box of Figure 1(right). Its dual (not shown) is obtained simply by overlaying their two duals.

We have noted that the spherical dual of a segment should be taken as the great circle perpendicular to this segment. If a zonohedron Z is the sum of k segments s_i, then its dual is the great-circle arrangement on the sphere consisting of the corresponding k great circles c_i. Each great circle in this arrangement is divided into a number of subarcs by its intersections with the other circles; these arcs correspond to a family of parallel edges (called a zone) on Z. To recover Z from the dual, we just need to know the length of the edges in each zone. But that is the same as the length ℓ_i of the original segment s_i. Thus, to properly label the dual of the Minkowski sum, we put the label ℓ_i on each segment of the great circle c_i.

The observation of Zongker and Hart [8] was that this overlay method works to produce Minkowski sums in general. Suppose A and B are two polyhedra, with labeled spherical duals $\widetilde{A}$ and $\widetilde{B}$. Then we construct a spherical network $\widetilde{P}$ by overlaying $\widetilde{A}$ and $\widetilde{B}$, simply drawing them both on the same sphere, inserting a new node wherever arcs cross. Each arc η of $\widetilde{P}$ is part of an arc from either $\widetilde{A}$ or $\widetilde{B}$, and inherits

the label of that arc. In excpetional cases, η will lie along parts of arcs from both $\widetilde{A}$ and $\widetilde{B}$, in which case it gets the sum of those two labels.

The analysis of this construction in [8] was limited to the special case where A and B were each mid-scribed around some sphere (that is, had all edges tangent to that sphere). Here we show that, in fact, it works in complete generality.

Generically, the nodes in the overlay network $\widetilde{P}$ either will be nodes from $\widetilde{A}$ or $\widetilde{B}$ or will arise where an edge of $\widetilde{A}$ crosses one of $\widetilde{B}$. In the first case, the node and its incident arcs and their labels are exactly the same as seen in $\widetilde{A}$ or $\widetilde{B}$, so the closure condition is satisfied. In the second case, the node ν has four incident arcs. They come in two opposite pairs, with equal labels ℓ_i on either pair. Clearly, this node is the dual of a parallelogram, and the closure condition is satisfied by symmetry.

In special cases, a node of $\widetilde{A}$ may lie exactly on an arc or node of $\widetilde{B}$. But then the closure condition in $\widetilde{P}$ is just the sum of the closure conditions from these two overlapping nodes. The case where an arc of $\widetilde{A}$ (partially) overlaps an arc of $\widetilde{B}$ also causes no problems, as long as we have used the sum of the original labels on the overlap, as specified above.

Since the closure conditions are satisfied everywhere, the labeled network $\widetilde{P}$ does correspond to a poly-hedron P, the Minkowski sum of A and B.

Connections and applications

This construction could be used to morph between any two convex polyhedra A and B. For time t ranging from 0 to 1, we would use the weighted sum $P_t := (1 - t)A + tB$. These intermediate polyhedra P_t all have the same spherical dual $\widetilde{P}$, obtained by overlaying the duals $\widetilde{A}$ and $\widetilde{B}$. All we need to do as t varies is to linearly interpolate the labels.

As an example, consider a symmetric pattern of great circles drawn on the sphere as in Figure 4, each tilted slightly from the equator. It is the dual to a so-called "polar zonohedron" (see [2, 6]). Suppose we
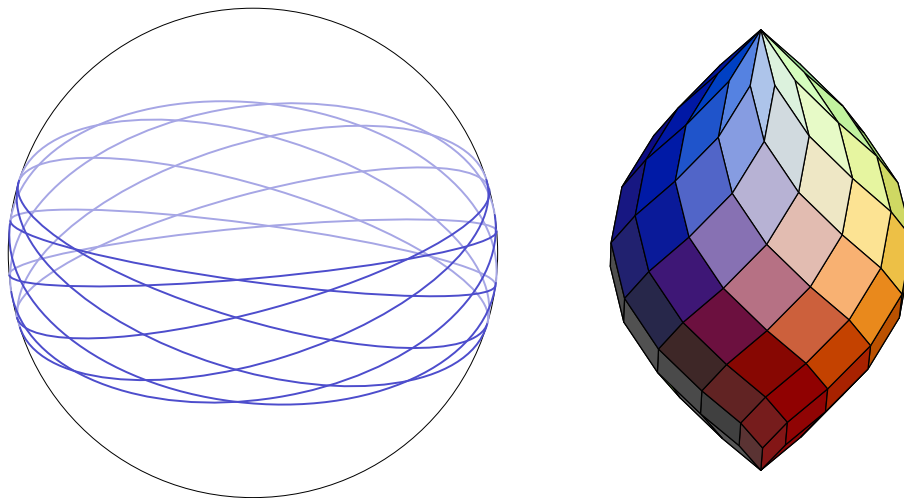

Figure 4: The symmetric pattern of great circles (left) is the spherical dual of a polar zonohedron like the one shown (right). Such a zonohedron is by definition generated by segments equally spaced around a cone.

place two different polar zonohedra in space with their axes perpendicular. Morphing between them the results in the sequence shown in Figure 5.

Aside from zonohedra, another interesting class to consider for this duality construction is deltahedra. A deltahedron is a polyhedron with equilateral-triangle faces. Its dual is then a network of great-circle arcs meeting in threes at equal ($120°$) angles. Such a network is reminiscent of a two-dimensional bubble cluster,
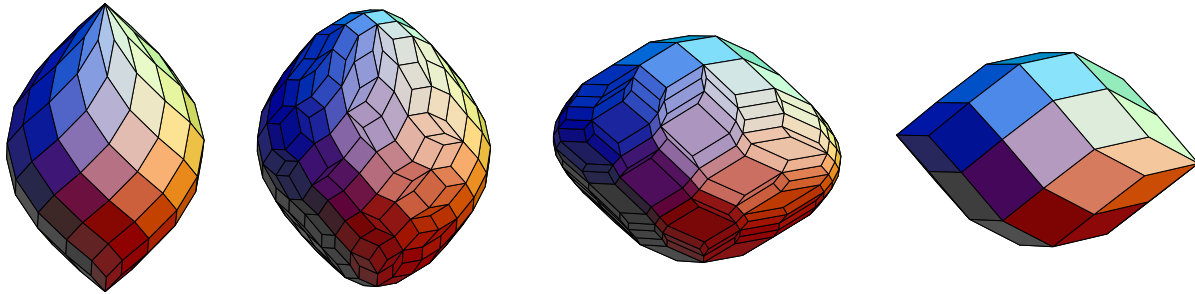
Figure 5: The polar zonohedra at left and right are generated by segments equally spaced around a cone. They have been placed with perpendicular axes. In the middle we see two intermediate stages of the morph between them.

and indeed the eight possible such networks describe the only candidate singularities for three-dimensional bubble clusters. These ideas generalize to arbitrary higher dimensions [5] where the classification of such soap-film singularities is not yet complete.

From a more abstract point of view, our spherical dual $\widetilde{P}$ with arcs labeled by edge lengths is simply the generalized mean curvature measure of the polyhedron P in the sense of Minkowski mixed volumes. (See for instance [4].) The overlay construction is then simply explained by the known fact that such measures are additive under Minkowski sum. It seems to be an open problem in general, however, to decide which measures on the sphere can arise as the generalized mean curvature of some convex body.

Acknowledgements

References

[1] David Austin and Will Dickinson. *Spherical Easel*. Software, 2002–2004. `http://merganser.math.gvsu.edu/easel/`

[2] David Eppstein. *Zonohedra and Zonotopes*. Mathematica in Education and Research **5**:4, pp. 15–21, 1996. `http://www.ics.uci.edu/~eppstein/junkyard/ukraine/ukraine.html`

[3] Efi Fogel and Dan Halperin. *Exact and Efficient Construction of Minkowski Sums of Convex Polyhedra with Applications*. Proc. ALENEX 2006, to appear. `http://www.math.tau.ac.il/~halperin/papers/exact_mink_3d.pdf`

[4] Jane R. Sangwine-Yager. *Mixed Volumes*. In *Handbook of Convex Geometry*, Gruber and Wills, eds., pp. 43–71, 1993.

[5] John M. Sullivan. *Convex Deltatopes in all Dimensions, and Polyhedral Soap Films*. Preprint, 1995. `http://torus.math.uiuc.edu/jms/Papers/`

[6] Russell Towle. *Graphics Gallery: Polar Zonohedra*. The Mathematica J. **6**:2, pp. 8–12, 1996.

[7] Günter M. Ziegler. *Lectures on Polytopes*. Springer GTM **152**, 1997.

[8] Douglas Zongker and George W. Hart. *Blending Polyhedra with Overlays*. Bridges 2001 Proceedings, Southwestern Coll., Kansas, pp. 167–174.