# Blending Polyhedra with Overlays

Douglas Zongker
University of Washington
Box 352350 Seattle, WA 98195
dougz@cs.washington.edu

George W. Hart
http://www.georgehart.com/
george@georgehart.com

### Abstract

We introduce a novel method for constructing polyhedra by blending together two or more existing polyhedra, using a dualization method. The method results in a Minkowski sum, analogous to that of overlay tilings [5], which uses a topological dual operation to create tilings from a network (a graph embedded in a plane). This paper extends that method to networks on the surface of the unit sphere. We begin with polyhedra in canonical form, dualize them to form networks, overlay the networks, and dualize the result to obtain a new polyhedron that blends together the faces of the original polyhedra.
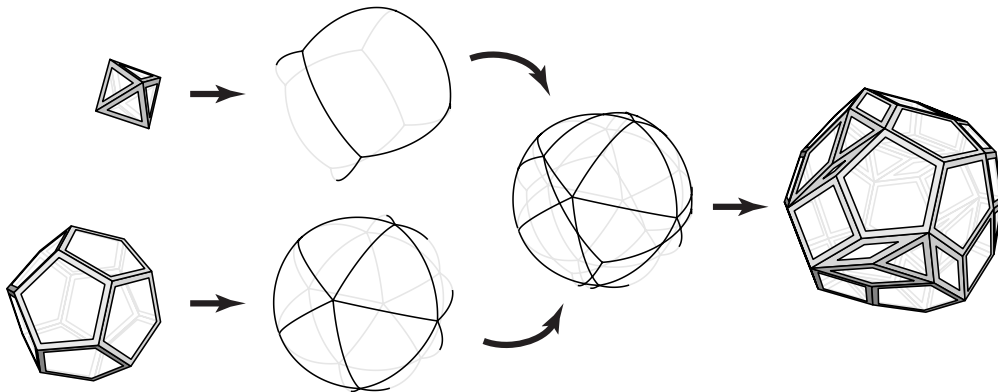
**Figure 1** An example of blending polyhedra.

## 1   Introduction

We introduce a method for blending together polyhedra of certain classes. In the example above, we have blended a regular octahedron with a regular dodecahedron. This is accomplished via dualizing each source polyhedron into a network on the surface of the sphere (shown in the second column). The source networks are overlaid, introducing new vertices wherever edges cross (third column). Finally this overlay network is dualized to form a polyhedron. Typically this polyhedron consists of the faces of the original polyhedra, with normals preserved, with a number of rhombi introduced to fill in the gaps. By exploiting symmetries in selecting and orienting the source polyhedra, a variety of interesting and attractive blend polyhedra can be created.

The technique lends itself well to animation. Rotating one polyhedron with respect to the other produces a series of blend polyhedra that shows both continuous and discrete changes. The blended polyhedra can

be given weights which scale their edge lengths in the result, and smoothly varying the weights produces a second type of animation, in which one polyhedron can morph into another.

## 2 Polyhedra and networks

The first step, representing each input polyhedron as its "dual" spherical network, is simple. Note that we use a nonstandard notion of "dual" in which the arcs of the spherical network are each labeled with real numbers to encode edge lengths. For each polyhedral face with outward unit normal $\vec{n}$, a vertex is placed on the unit sphere at $\vec{n}$. Vertices corresponding to neighboring faces are connected with arcs that are sections of great circles, choosing the *shorter* of the two possible arcs connecting any pair of vertices. Each of these arcs will thus be less than a semicircle (in fact, the angle spanned by the arc is $\pi$ minus the polyhedron's corresponding dihedral angle). Each arc is labeled with the length of the polyhedron edge that generated it.

Two types of degenerate polyhedra also give rise to networks which prove useful in blending. Imagine an *n*-ary prism of zero thickness, that is, two *n*-gons glued together edge to edge. The network dual of this "polyhedron" is a pair of antipodal vertices, joined by *n* evenly spaced semicircular arcs, slicing the sphere like a melon. Another is the even more degenerate polyhedron consisting of a single line segment, whose dual is a network with no vertices but a single edge — an entire great circle of the sphere.

The overlay procedure is straightforward. Consider two or more networks on the unit sphere. Wherever arcs cross they are split, introducing a new vertex at the intersection point. Coincident vertices are merged, and vertices falling exactly on arcs cause the arcs to be split. Whenever an arc is split, its label is propagated to the newly created subarcs.

Once the overlay network is constructed, we perform another dualization to obtain the corresponding blended polyhedron. We will explain this procedure here in high-level terms, and offer arguments that the result is actually a polyhedron in Appendix A.

To dualize a network of arcs back to a polyhedron, we start by constructing a *dual edge* for each arc of the source network. The dual edge is a straight line segment of length given by the arc label, normal to the plane in which the source arc lies. The two endpoints of the dual edge are associated with the two spherical faces which meet along the source arc.

We now attempt to collect all the dual edges together to form a polyhedron. We choose one edge and fix its position. Each endpoint of that edge corresponds to a face of the network. For each of those faces, we gather together all of the dual edges sharing that face and translate them so the corresponding endpoints coincide, forming a single complex of edges containing the fixed edge. This is illustrated in Figure 2. This process is repeated for each face at the frontier of this complex, translating each dual edge and attaching it to the complex. When the process is complete, if we have started with the right kind of network, the result is that the dual edges form a polyhedron.

## 3 Which polyhedra can be blended?

Blendable polyhedra are those in which all edges are tangent to a unit sphere. As a consequence each face can be circumscribed about a circle. (The circle is the intersection of the face's plane with the sphere.) Many familiar classes of polyhedra are already in this form, including the Platonic solids, the Archimedean solids, and the Archimedean duals, the Catalan polyhedra.
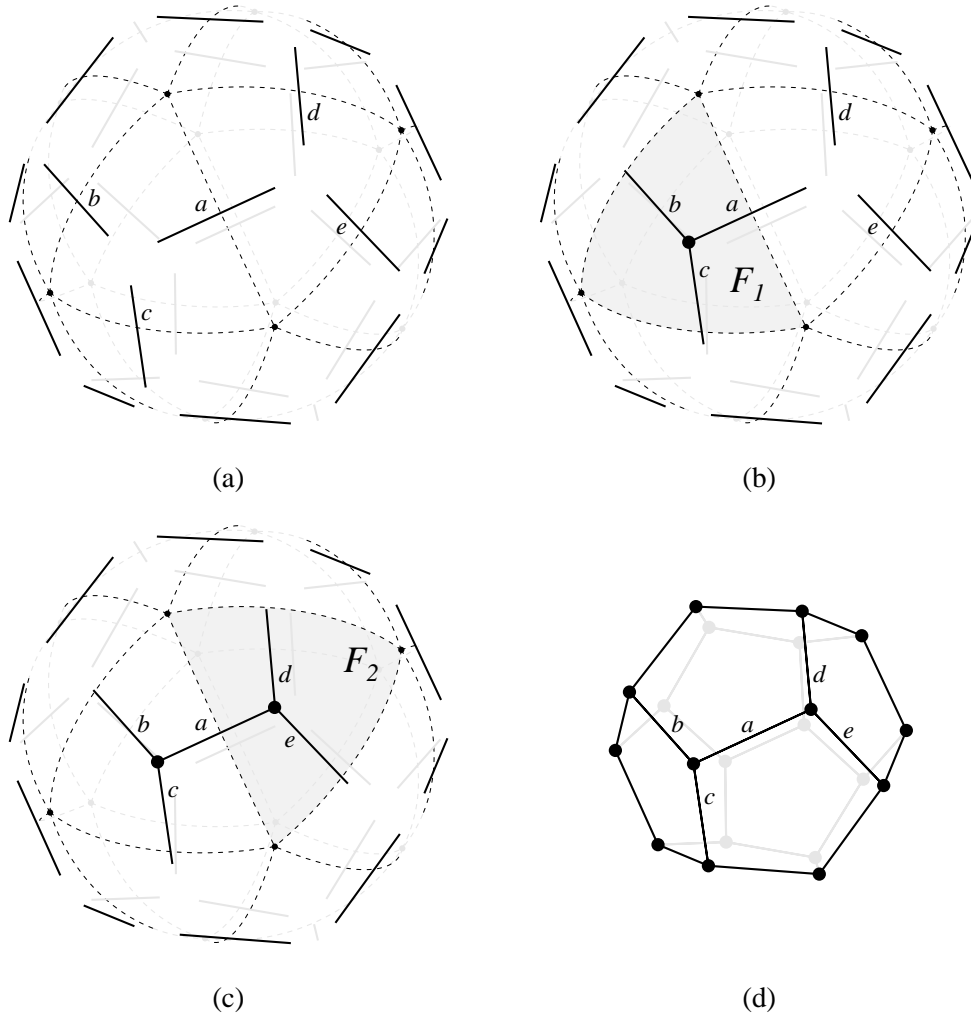
**Figure 2** Part (a) shows a spherical network (dashed lines), along with the initial set of dual edges (solid lines). Edges on the back side of the sphere are shown in light gray. Edge $a$ is fixed. Part (b) shows the result of the first gather operation, on the face $F_1$. Edges $b$ and $c$ are translated to join to $a$. The second gather, for face $F_2$, joins edges $d$ and $e$ to the complex as shown in image (c). When all the edges have been agglomerated into a single component, the result (for this network) is a polyhedron, shown in image (d) (the network has been omitted).

There is a well studied class of polyhedra which are in *canonical form*. These have the property that all edges are tangent to a sphere, plus an additional property which requires an even distribution of the edges around the sphere. For our purposes, this second property is not necessary. So canonical polyhedra are a standard class of polyhedra which can be blended, but are not the only blendable polyhedra. An introduction to canonical polyhedra can be found in Ziegler [4]. An iterative numerical method for finding the canonical form of a given topological polyhedron can be found in Hart [1].

Our polyhedra are represented by $(x, y, z)$ vertex coordinates, so are actually "oriented polyhedra" in the sense that their orientation affects the blend. We can scale the edge weights of any blendable polyhedron $P$ by $\alpha$ to get $\alpha P$, and form blended linear combinations $\alpha_1 P_1 + \alpha_2 P_2 + \ldots$, where "+" represents the blend operations. Thus we generate a large space of polyhedra, each of which we can further operate upon.

# 4   Results

The blend procedure we have described has a number of useful properties. It is commutative and associative, so one can talk unambiguously about the blend of a set of oriented polyhedra without needing to specify a sequence of binary blend operations. Blending polyhedra requires no new edge directions to be added — each edge of the blend is parallel to some edge of one of the input polyhedra. Similarly, each edge length of the blend is an edge length of one of the input polyhedrons, or possibly a sum of edge lengths if the set of inputs contains parallel edges. Incidentally, because of these properties, if one can make a Zometool[3] model of two polyhedra, then one can make a Zometool model of their blend—at least for those relative orientations of the two polyhedra which are compatible with the angles available in the Zometool system.

Figures 3 and 4 show some polyhedra created with the blending technique. For each, both the overlaid network and the blend polyhedron are shown, with the source polyhedra listed underneath. In most cases these are created with multiple copies of a single polyhedron, rotated to take advantage of some underlying symmetry. The relative orientation of the polyhedra being blended affects the result. We have chosen some particularly symmetrical arrangements.

Figure 5 shows an "animation" of intersecting the network corresponding to a snub dodecahedron with a rotating great circle. Intersection of a network with a set of great circles is equivalent to the technique for creating "zonish polyhedra." [2] Figure 6 shows another animation, based on three rotating copies of the same enneahedron.
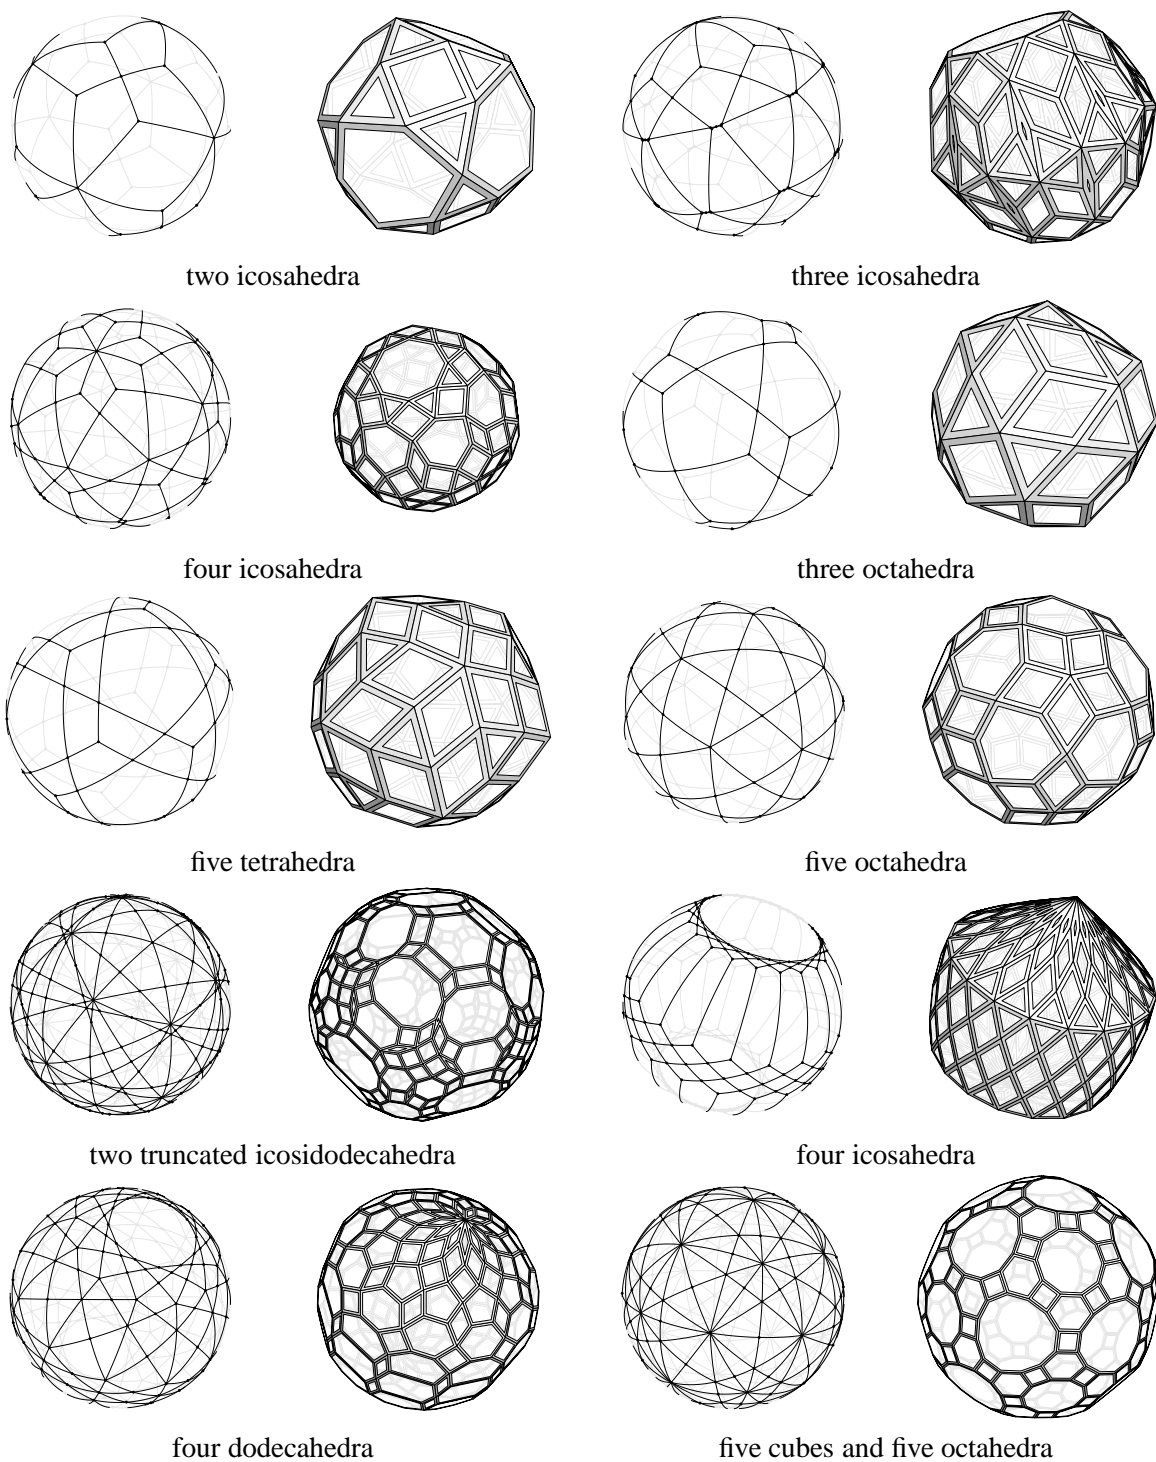
Animations like this one can have a pleasing effect, displaying both continuous motion — the angles of rhombs change continuously as the angles of intersection in the network change — as well as discrete jumps, when edges move through vertices. Another kind of animation is shown in Figure 7, where we adjust the weights on the two networks being blended to produce a kind of linear interpolation between the two polyhedra. The shape starts as one of the input polyhedra, then its faces become separated by those of the other polyhedron, which grow until they overtake those of the first shape.

# 5   Conclusions

We have introduced a method for creating attractive polyhedra through a technique that blends existing shapes. It does not apply to all polyhedra, but it does work with a wide range of well-known polyhedra. The technique lends itself well to animation, producing a changing sequence of result polyhedra that is both visually interesting and temporally coherent.

# References

[1] George W. Hart. Calculating canonical polyhedra. *Mathematica in Education and Research*, 6(3):5–10, Summer 1997.

[2] George W. Hart. Zonish polyhedra. In *Proceedings of Mathematics and Design '98*, page 653, 1998.

[3] George W. Hart and Henri Picciotto. *Zome Geometry*. Key Curriculum Press, 2001.

[4] Günter M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, 1994.

[5] Douglas Zongker. Creation of overlay tilings through dualization of regular networks. In Nathaniel Friedman and Javier Barrallo, editors, *Proceedings of ISAMA 99*, pages 495–502, June 1999.

two icosahedra

three icosahedra

four icosahedra

three octahedra

five tetrahedra

five octahedra

two truncated icosidodecahedra

four icosahedra

four dodecahedra

five cubes and five octahedra

**Figure 3** Polyhedral blends. Each pair shows a blended result (right) and the overlayed spherical network from which it is derived (left).

networks from snub dodecahedra of opposite handedness      overlay      dual polyhedron

**Figure 4** Another polyhedron created with the dualization technique.
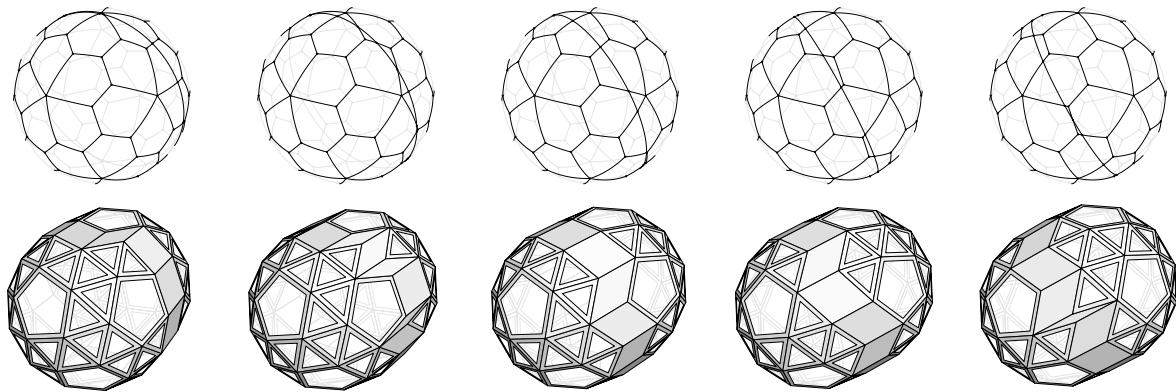


**Figure 5** Snub dodecahedron, overlaid with a rotating great circle. Faces produced by intersection with the great circle are shown as solid polygons.
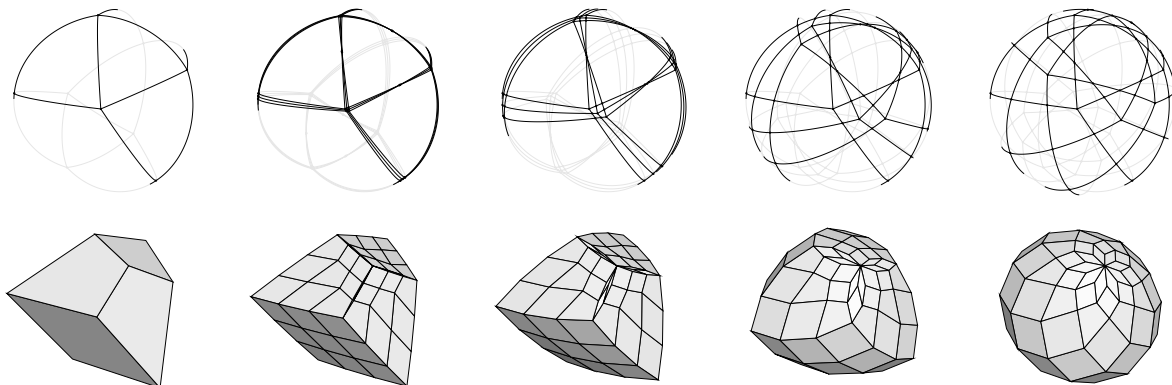


**Figure 6** Three identical enneahedra, rotating at different speeds. In the first frame, when all the copies are aligned, the blend is simply an enneahedron three times as large.

$\alpha = 0$   $\alpha = 0.25$   $\alpha = 0.5$   $\alpha = 0.75$   $\alpha = 1$
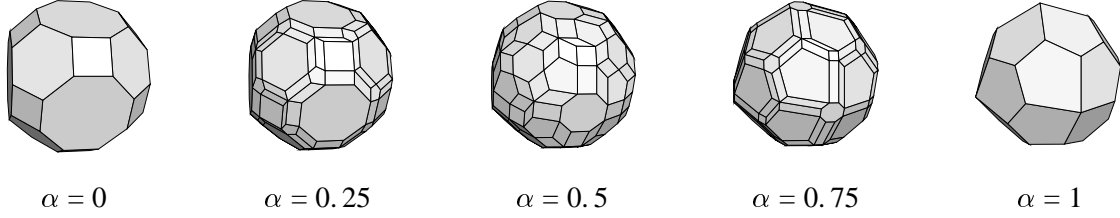
**Figure 7** A linear combination $(1 - \alpha)P_1 + \alpha P_2$, where $P_1$ is a truncated cuboctahedron and $P_2$ is a pentagonal icositetrahedron. The dual networks of each combination are all identical, except for the weights on each arc.

## A    Appendix: Dualization gives a polyhedron

Mathematicians will recognize our "blend" as a 3D Minkowski sum. Our method is apparently a new technique for compting it, but one which, so far, we have only been able to apply to a restricted class of polyhedra.

Not all networks successfully dualize to form polyhedra. When given a non-dualizable network, the dualization procedure described in Section 2 will fail. At some point, the procedure will attempt to gather together all the dual edges incident on some face $F$, but find that two or more of them have already been attached to the fixed complex, in such a way that their free endpoints do not already meet. Since these edges can't be moved, there is no way to complete the gather for that face.

We now sketch a proof that when the network in question is constructed by overlaying $n$ networks obtained from acceptable polyhedra, the overlay network dual is in fact a simple convex polyhedron. As indicated in Section 3, blendable polyhedra have all edges tangent to a unit sphere. By the Koebe-Andreev-Thurston Circle Packing Theorem, any such polyhedron has a dual polyhedron with edges tangent to the sphere at the same places, but orthogonal to them. [4, p. 117] Because of this property, if the dual is projected onto the sphere to give a network of arcs, the normals to the planes of the arcs give the directions of the original polyhedron. These normals are used below to construct the blend.

We imagine constructing the dual in two steps: first we construct the dual's *topology* in the form of a graph (actually, a digraph) representing the vertices and edges. We then show that, for an overlay of dualizable networks, a 3D position can be assigned to each vertex in a way that makes the difference along an edge equal to the normal vector described in the dualizing procedure above.

Figure 8 illustrates how edges of the polyhedron $P$ are constructed. For each face $F$ of the original network, the dual will contain a vertex $\tilde{F}$. Where an edge $e$ joins two faces $F_1$ and $F_2$, the dual will contain a directed edge $\tilde{e}$ joining $\tilde{F}_1$ and $\tilde{F}_2$. Each dual edge $\tilde{e}$ will be labeled with a 3D vector — a vector normal to the plane, $e$, containing the arc endpoints and the origin. The orientation of $\tilde{e}$ will be used to encode which of the two possible unit normals was used. The directed edge should point toward the vertex corresponding to the face in the direction of the normal. It would be equally valid to choose the other normal to the plane — this would result only in flipping the orientation of the edge and negating its label.

Next we'll assign a position to each vertex, so that the label on each edge is equal to the difference between the positions of the two endpoints. Begin by fixing any single vertex by giving it an arbitrary $(x, y, z)$ position. Then, for each edge $\tilde{e}$ joining a fixed vertex $\tilde{F}_i$ with an unfixed vertex $\tilde{F}_j$, fix $\tilde{F}_j$ at the position of $\tilde{F}_i$ plus the label of $\tilde{e}$ if the edge is from $\tilde{F}_i$ to $\tilde{F}_j$, or the position of $\tilde{F}_i$ minus the label if the edge is from $\tilde{F}_j$ to $\tilde{F}_i$. Repeat this procedure until there are no more unfixed vertices.

This procedure starts at the first fixed vertex $\tilde{f}_0$ and propagates a set of positions out along chains of edges.
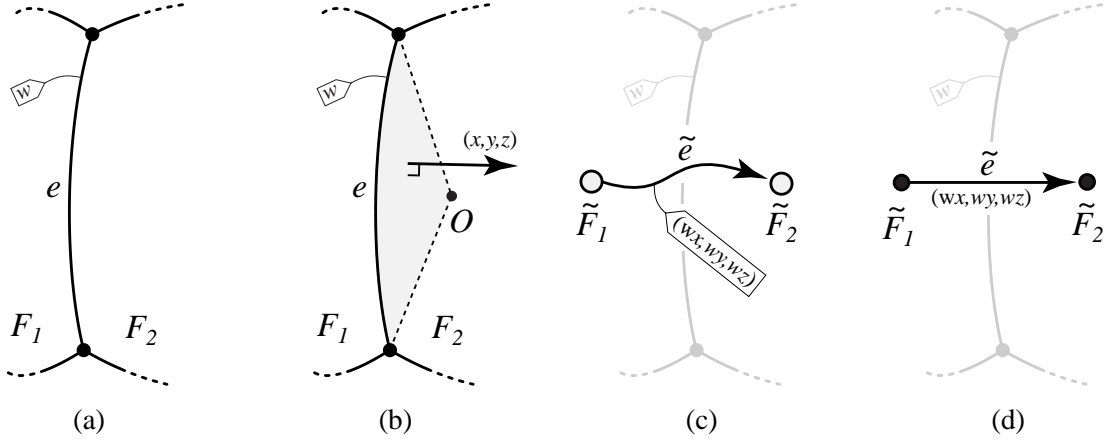
**Figure 8** Constructing one edge of the dual for given spherical network. Each arc $e$ separating faces $F_1$ and $F_2$ is labeled with a real number $w$ (part (a)). We consider the normal vector $(x, y, z)$ to the plane containing $e$ and the sphere's center $O$ (part (b)). The dual is first constructed as a digraph (part (c)), with each edge $\tilde{e}$ labeled with the normal of the corresponding plane times the original arc's weight. The orientation of $\tilde{e}$ records which of the two possible normal vectors was chosen. Finally, we assign a position to each vertex $\tilde{F}_i$ of the dual polyhedron, so that the vector between neighboring vertices is equal to its label (part (d)).

We will argue that if the network has been constructed as an overlay of other dualizable networks, then the set of positions assigned will be consistent, that is, there is no vertex which is assigned to two different positions by following two different chains of edges. This is equivalent to saying that the net sum of the labels around any loop of edges is zero. (If vertex $\tilde{F}_j$ is assigned to one position by following path $P$ from vertex $\tilde{F}_i$, and to another position via path $P'$ from $\tilde{F}_i$, then path $P$ plus the reverse of path $P'$ will be a loop from $\tilde{F}_i$ to $\tilde{F}_j$ and back with a nonzero sum.) The direction of edges is not considered when forming loops — a loop can cross an edge "against the arrow". However, when crossing an edge this way, the label should be subtracted from the running sum, rather than added to it. (Remember that reversing the direction of an edge is equivalent to negating its label.)

Suppose that network $D$ has been constructed by overlaying networks $D_1, D_2, \ldots, D_n$, and that $P$ is the dual of $D$. Choose any loop in $P$, starting at a vertex $\tilde{F}_i$ and following any chain of edges that eventually returns to $\tilde{F}_i$. From the construction of $P$, this loop corresponds to a closed loop of neighboring faces in $D$. Every hop from face to face in $D$ corresponds to a hop in one or more of the $D_i$ source networks. Since the chain returns to the starting face of $D$, it must correspond to a closed chain of faces in each of the $D_i$s. Each $D_i$ contributes its own loop of segments which sums to zero, so the total loop must sum to zero.

Now we know that all the edges of $P$ can be joined up in a consistent way, we will argue that $P$ is a polyhedron. Around each vertex $V$ of $D$ is a cycle of faces which contributes a cycle of edges in $P$. These edges form a planar face in $P$, because by construction each of the differences is orthogonal to the radius $OV$ in $D$, so the path in $P$ can not leave the plane orthogonal to $OV$ that it starts in. Furthermore, $P$ is a simple planar graph because it is topologically dual to $D$, which is a planar graph because it lies on a sphere. $P$ is convex because an outward normal for each face can be constructed in the direction of the radius $OV$ as described above. These determine convex dihedral angles at each edge of $P$ because they correspond to edges of $D$ which are less than semicircles. ∎