Turning Orthogonally Convex Polyhedra into Orthoballs

Fidel Barrera-Cruz*, Therese Biedl†, Martin Derka†, Stephen Kiazyk†, Anna Lubiw†, Hamideh Vosoughpour†

Abstract

As a step towards characterizing the graphs of orthogonally convex polyhedra, we show that for any *simple* orthogonally convex polyhedron there is an *orthoball* that is equivalent in the sense that it has the same graph and the same face normals. An *orthoball* is a simple orthogonally convex polyhedron with a point inside that sees the whole interior (informally, it is "round"). The consequence for reconstructing polyhedra from graphs is that if we start from a 3-regular planar graph labelled with face normals, and wish to find a corresponding orthogonally convex polyhedron, then we can restrict our search to orthoballs.

1 Introduction

A famous result of Steinitz [6] characterizes the graphs of 3D convex polyhedra as precisely the 3-connected planar graphs (see e.g. [8]). This gives a polynomial time algorithm to recognize such graphs. There are also polynomial time algorithms to reconstruct a polyhedron from a given graph [5]. Recently, Eppstein and Mumford [3] characterized the graphs of some classes of 3D orthogonal polyhedra. They restrict attention, as we will, to *simple* orthogonal polyhedra: polyhedra with the topology of a sphere, with simply-connected faces, and with exactly three mutually-perpendicular axis-parallel edges meeting at every vertex. They characterized the graphs of simple orthogonal polyhedra as well as two subclasses, and they gave polynomial time recognition and reconstruction algorithms. See below for further details.

Eppstein and Mumford leave open the question of characterizing and efficiently recognizing the graphs of 3D [simple] orthogonally convex polyhedra. This seems difficult, but might be easier with extra information added to the graph. While the face orientations (i.e., $\{X, Y, Z\}$) are unique up to rotation, one could add the desired face normals from the set $\{X^+, X^-, Y^+, Y^-, Z^+, Z^-\}$. We call this the *face-labelled graph.* Even with such extra information the questions of characterizing or recognizing the graphs of 3D [simple] orthogonally convex polyhedra remain open.

Our main result is that, even with face labels, the graphs of simple orthogonally convex polyhedra are all graphs of "ball-shaped" orthogonal polyhedra. More precisely, an *orthoball* is a simple orthogonally convex polyhedron with a point inside that sees the whole interior and we prove:

Theorem 1 *Let $\mathcal{P}$ be a simple orthogonally convex polyhedron. Then there exists an orthoball $\mathcal{B}$ with the same face-labelled graph as $\mathcal{P}$.*

The 2D analogue of this result is quite obvious, see Figure 1. Our proof of Theorem 1 involves modifying a simple orthogonally convex polyhedron until it is an orthoball. The modification does not preserve edge directions (see also Figure 5b and c).

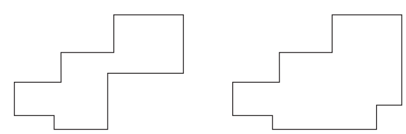

Figure 1: A 2D orthogonal polygon (left) and an equivalent 2D orthoball (right).

Background. Eppstein and Mumford [3] characterized the graphs of simple orthogonal polyhedra and of two subclasses: *corner polyhedra*, which have only one X^-, one Y^-, and one Z^- face; and *xyz polyhedra*, which have at most two vertices in any axis-parallel line. The graphs of the latter class are precisely the cubic bipartite 3-connected planar graphs.

Biedl and Genç [1, 2] have studied the problem of reconstructing an orthogonal polyhedron from information that includes not only the graph but also edge lengths, facial angles, or dihedral angles.

2 Definitions

An *α-line* (for $\alpha \in \{X, Y, Z\}$) denotes a line parallel to the α-axis. An *α-plane* (in 3D) is a plane that is perpendicular to the α-axis. For any geometric object $\mathcal{O}$ where the α-coordinate is the same throughout, we write $\alpha(\mathcal{O})$ to denote that fixed coordinate.

All polygons in this paper are required to have a nonempty interior that is simply connected. An *orthogonal polygon* has axis-aligned edges. An orthogonal 2D

---

*Combinatorics & Optimization.

†David R. Cheriton School of Computer Science.

For all authors: University of Waterloo, Waterloo, ON N2L 3G1, Canada, {fbarrera, biedl, mderka, skiazyk, alubiw, hvosough}@uwaterloo.ca. The third author was supported by NSERC Vanier CGS.

polygon is *orthogonally convex* if its intersection with any axis-aligned line is either empty or a line segment. An *orthogonal polyhedron* $\mathcal{P}$ has only axis-aligned face-normals. An orthogonal polyhedron $\mathcal{P}$ is *orthogonally convex* if the intersection (or "cross-section") of $\mathcal{P}$ with any axis-aligned plane is either empty or an orthogonally convex 2D polygon. (Recall that the latter are required to be connected.) An *orthoball* $\mathcal{B}$ is an orthogonally convex polyhedron that is *star-shaped*, i.e., it contains a point o such that all points in $\mathcal{B}$ are visible from o (where two points are visible if the straight line between them remains in the polyhedron).

An orthogonal polyhedron $\mathcal{P}$ is *simple* if it has the topology of a sphere, all faces are simply-connected, and any vertex is incident to exactly three pairwise perpendicular edges. The underlying graph of a simple orthogonal polyhedron is planar, 2-connected, 3-regular, and bipartite. By Heawood's theorem (see [7]), such an embedded graph has a unique 3-coloring of the faces, and hence a unique labelling (up to rotation) of the faces as X-faces, Y-faces or Z-faces. However, the graph does not determine the face-normals, i.e., whether the outward normal of (say) an X-face is in the positive or negative direction along the X-axis. The *face-labelled graph* of an orthogonal polyhedron $\mathcal{P}$ is the underlying graph of $\mathcal{P}$ with each face labelled by its outward face normal from the set $\{X^+, X^-, Y^+, Y^-, Z^+, Z^-\}$.

A *realization* of a face-labelled graph is a 3D orthogonal polyhedron with that face-labelled graph. Note that when all vertices in the graph have degree 3 we can always perturb edge lengths in any given realization to ensure no two faces are co-planar (this fact will simplify later constructions). This condition does not hold in the presence of vertices of degree 4 or higher.

3 Equivalent Properties

There exist many equivalent characterizations of orthogonally convex polygons. We list here a few that we need later (see also [4] for many related results.) See the appendix for a full proof.

Theorem 2 *The following are equivalent for an orthogonal (2D) polygon P:*

1. *P is orthogonally convex.*
2. *For any two points p, q in P, there exists a path from p to q inside P that is XY-monotone [4].*
3. *For any axis-aligned rectangle R, the intersection $R \cap P$ is empty or connected.*
4. *For any axis-aligned rectangle R, the intersection $R \cap P$ is empty or an orthogonally convex polygon.*

One can easily show similar characterizations for orthogonally convex polyhedra in 3D. Note that we do not require them to be simple. See the appendix for a full proof.

Theorem 3 *The following are equivalent for an orthogonal polyhedron $\mathcal{P}$:*

1. *$\mathcal{P}$ is orthogonally convex.*
2. *For any two points p, q in $\mathcal{P}$, there exists a path from p to q inside $\mathcal{P}$ that is XYZ-monotone.*
3. *For axis-aligned box $\mathcal{B}$, the intersection $\mathcal{B} \cap \mathcal{P}$ has at most one connected component.*
4. *For axis-aligned box $\mathcal{B}$, the intersection $\mathcal{B} \cap \mathcal{P}$ is either empty or an orthogonally convex polyhedron.*

For our later transformation from orthogonally convex polyhedra to orthoballs, it will be convenient to know equivalent characterizations of orthoballs as well. Again, we do not require simplicity for this theorem.

Theorem 4 *The following are equivalent for an orthogonal polyhedron $\mathcal{P}$:*

1. *$\mathcal{P}$ is an orthoball.*
2. *The* link-distance *of $\mathcal{P}$ is at most 2, i.e., for any two points $p, q \in \mathcal{P}$ there exists a point $o \in \mathcal{P}$ such that the line segments $\overline{po}$ and $\overline{oq}$ are within $\mathcal{P}$.*
3. *For any $\alpha \in \{X, Y, Z\}$, no α^--face has a greater α-coordinate than any α^+-face.*

Proof. (1) $\Rightarrow$ (2) holds because we can use as point o a point that can see all of the star-shaped polyhedron.

To show (2) $\Rightarrow$ (3), let f^+ be (say) an X^+-face and f^- be an X^--face f^-. Let p^+ and p^- be points in the relative interior of f^+ and f^- respectively. Let o be a point such that the line segments $\overline{p^+o}$ and $\overline{p^-o}$ are inside $\mathcal{P}$. Since p^+ is in the interior of f^+, this is possible only if $X(o) \leq X(p^+) = X(f^+)$ and similarly $X(o) \geq X(f^-)$. Hence $X(f^-) \leq X(f^+)$.

To see (3) $\Rightarrow$ (1), let X_o be an X-coordinate between the maximum X-coordinate of an X^--face and the minimum X-coordinate of an X^+-face. Similarly define Y_o and Z_o. We claim that point $o := (X_o, Y_o, Z_o)$ can see all of $\mathcal{P}$. To prove this, let p be any point in $\mathcal{P}$, and assume wlog that it is in the $(+,+,+)$-quadrant of o. If $\overline{op}$ is not in $\mathcal{P}$, then it must intersect an α^--face in its interior, for some $\alpha \in \{X, Y, Z\}$. But by choice of o no such face can exist in the $(+,+,+)$-quadrant of o, and so $\overline{op}$ belongs to $\mathcal{P}$. □

Finally, we will need the fact that projections preserve orthogonal convexity. The proof is in the appendix.

Lemma 5 *If $\mathcal{P}$ is an orthogonally convex polyhedron and $\mathcal{Q}$ is its projection to a coordinate plane then $\mathcal{Q}$ is an orthogonally convex polygon.*

4 Main Result

In this section we prove our main result, Theorem 1. The idea is to modify any simple orthogonally convex polyhedron $\mathcal{P}$ to be an orthoball. We use the fact that $\mathcal{P}$ is an orthoball if and only if no α^--face appears before an α^+-face in α-coordinate ordering (Property 3 of Theorem 4). If this property fails, there must be two consecutive out-of-order faces. We show how to exchange the relative order of these two faces, while maintaining orthogonal convexity, the same face-labelled graph, and the same planes of all other faces. By repeatedly applying this construction over all axes, we can "bubblesort" the faces so that $\mathcal{P}$ becomes an orthoball.

Let us formalize the idea. Assume $\mathcal{P}$ is a simple orthogonally convex polyhedron. Define the *α-face sequence* $f_1, f_2, \ldots, f_k$ of polyhedron $\mathcal{P}$ on dimension $\alpha \in \{X, Y, Z\}$ to be its α-faces listed in order of strictly increasing α-coordinates. That is, for every f_i, f_{i+1}, $\alpha(f_i) < \alpha(f_{i+1})$. (Recall that in Section 2, we established that every pair of α-faces can be assumed to be not co-planar.) Two consecutive α-faces are called a *bad pair* if they face opposite directions (i.e., one + and one − face), and the α-coordinate of the − face is greater than the α-coordinate of the + face. See Figure 2a for an example of a bad pair.

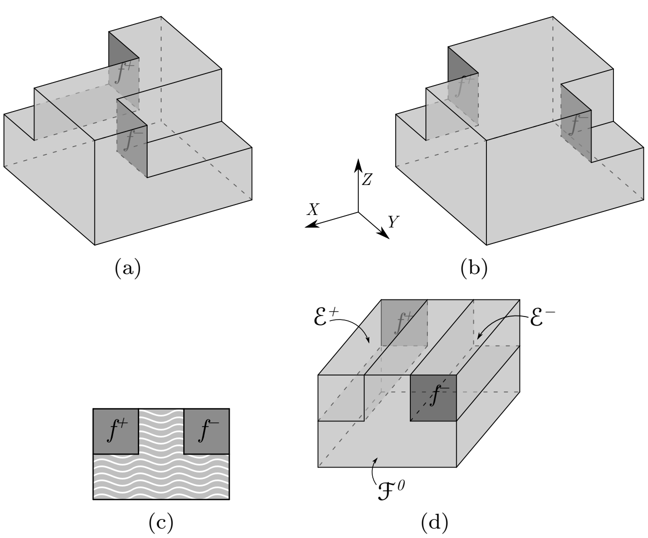

Figure 2: A bad pair and the result of swapping them: (a) the bad pair f^+, f^- in $\mathcal{P}$; (b) the result of the swap; (c) the polygon $f^0 \cup f^+ \cup f^-$ where f^0 is shown with the wavy pattern; (d) the polyhedron $\mathcal{F}^0$ and the extrusions $\mathcal{E}^+$ and $\mathcal{E}^-$.

A bad pair can be "fixed" by a *swap* operation that exchanges the α-coordinates of the vertices of f_i with the α-coordinates of the vertices of f_{i+1} (see Figure 2b). Lemma 7 formalizes this swap operation and shows that it does the right thing, i.e., produces a valid simple orthogonally convex polyhedron with the same face-labelled graph and the same face sequences except for the exchange of f_i and f_{i+1}. We begin with some notation to be used throughout this section.

Notation. Let f^+ and f^- be a bad pair of α-faces in $\mathcal{P}$; for ease of description let us assume that $\alpha = X$. We can express $\mathcal{P}$ as a union of three orthogonal polyhedra: $\mathcal{F}^+$, the points with X-coordinate greater than or equal to $X(f^-)$; $\mathcal{F}^-$, the points with X-coordinate less than or equal to $X(f^+)$; and $\mathcal{F}^0$, the closure of the set of points of $\mathcal{P}$ with X-coordinate strictly between $X(f^-)$ and $X(f^+)$. All three polyhedra are orthogonally convex by Property 4 of Theorem 3. Observe that $\mathcal{F}^0$ has the same cross-section in any X-plane—call this polygon f^0, and note that it is orthogonally convex. See Figure 2c. We will use the notation f^+ and f^- not only for the bad faces, but also for the polygons resulting from projecting them onto an X-plane—the distinction will be clear from the context. Observe that the cross-section of $\mathcal{P}$ at $X(f^-)$ consists of $f^0 \cup f^-$ and that these polygons are therefore internally disjoint. Similarly, the cross-section of $\mathcal{P}$ at $X(f^+)$ consists of $f^0 \cup f^+$ and these polygons are internally disjoint.

Lemma 6 *Polygons f^0, f^+, and f^- have the following properties:*

1. *$f^- \cup f^0 \cup f^+$ is orthogonally convex.*
2. *f^+ and f^- are internally disjoint.*
3. *f^+ and f^- do not share a line segment.*
4. *f^+ and f^- intersect in at most two vertices.*

Proof. Assume that $\alpha = X$. (1) Observe that $f^- \cup f^0 \cup f^+$ is the X-projection of the orthogonally convex polyhedron formed by slicing $\mathcal{P}$ by two X-planes just before f^+ and just after f^- and taking the middle piece. Therefore, by Lemma 5, it is orthogonally convex.

(2) An internal point of polygons f^+ and f^- corresponds to an axis-parallel line ℓ that exits $\mathcal{P}$ on face f^+ and enters on face f^-, contradicting the assumption that $\mathcal{P}$ is orthogonally convex (see Figure 3).

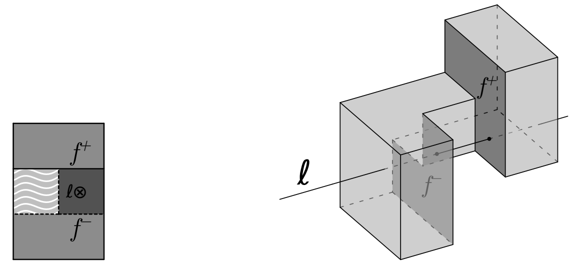

Figure 3: f^+ and f^- cannot have an interior point in common. The left side shows $f^- \cup f^0 \cup f^+$ where f^0 is shown with the wavy pattern.

(3) If polygons f^+ and f^- share a line segment then the polygon f^0 includes the shared line segment but not the points in small neighbourhoods to either side (see Figure 4), contradicting the fact that f^0 is orthogonally convex.

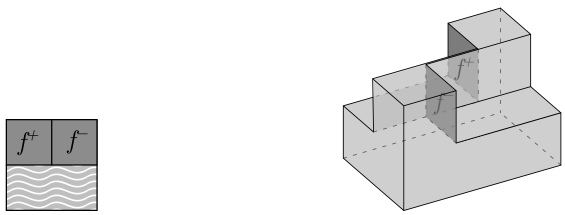

Figure 4: f^+ and f^- cannot share a line segment. The left side shows $f^- \cup f^0 \cup f^+$ where f^0 is shown with the wavy pattern.

(4) Suppose polygons f^+ and f^- have 3 shared vertices corresponding to edges, say, e_1, e_2, e_3, between faces f^+ and f^-. Consider the polygon f^0. Edges e_1, e_2, e_3 project to vertices of f^0. The polygons f^+ and f^- also have the projections of e_1, e_2, e_3 as vertices. But the three polygons f^0, f^+, and f^- are internally disjoint. We claim that this violates planarity. Select three points F, A and B in the interior of f^0, f^+ and f^- respectively, and connect them to all of e_1, e_2, e_3 by disjoint paths interior to the respective polygons. This creates a planar $K_{3,3}$ with the two partitions formed by e_1, e_2, e_3 and F, A, B. Contradiction. □

Lemma 7 *Let $\mathcal{P}$ be a simple orthogonally convex polyhedron, and let f^+ and f^- be a bad pair of α-faces of $\mathcal{P}$ for some $\alpha \in \{X, Y, Z\}$. Then there exists a simple orthogonally convex polyhedron $\mathcal{P}'$ with the same face-labelled graph as $\mathcal{P}$ and the same face-sequences except that f^+ and f^- have been exchanged.*

Proof. Assume w.l.o.g. that $\alpha = X$. We defined a swap to exchange the α-coordinates of the vertices of f^+ with the α-coordinates of the vertices of f^-. We first express this in terms of the change to the solid body $\mathcal{P}$. Face f^- moves from $X(f^-)$ to $X(f^+)$ sweeping out a volume $\mathcal{E}^-$, the extrusion of f^-. Similarly, face f^+ moves from $X(f^+)$ to $X(f^-)$ sweeping out a volume $\mathcal{E}^+$, the extrusion of f^+. Observe that $\mathcal{E}^+$ and $\mathcal{E}^-$ are internally disjoint because f^+ and f^- are, by Lemma 6. We will define $\mathcal{P}'$ to be $\mathcal{P} \cup \mathcal{E}^+ \cup \mathcal{E}^-$. Our proof will justify that this does in fact correspond to a swap. Recall that $\mathcal{F}^-$, $\mathcal{F}^0$, and $\mathcal{F}^+$ were defined at the beginning of the section.

We will prove three things: (1) For every face (except f^+, f^-) in $\mathcal{P}$ there exists a face in $\mathcal{P}'$ with the same supporting plane and face normal; (2) Face-edge incidences have not changed, which implies that $\mathcal{P}'$ has the same face-labelled graph as $\mathcal{P}$; (3) $\mathcal{P}'$ is orthogonally convex. Since $\mathcal{P}$ is simple, (2) also implies that $\mathcal{P}'$ is simple because simplicity of a polyhedron depends only on its face lattice (the vertices, edges, faces and their incidences). Also, (1) implies that face-sequences remain the same except for the swap of the two consecutive faces f^+ and f^-, since faces retain their supporting planes and hence their order.

For step (1), we claim that any face (except f^+, f^-) contains a vertex v that is neither on f^+ nor f^-. Assume for contradiction that for some face $f \neq f^+, f^-$ all vertices belong to f^+ and f^-. Then the X-projection of f is inside both polygons f^+ and f^-, and so f^+ and f^- share at least an edge, contradicting Lemma 6. So f contains a vertex v not in f^+ or f^-. The X-coordinate of v is outside the interval $[X(f^+), X(f^-)]$ (since f^+ and f^- are consecutive in the X-face sequence) and hence adding $\mathcal{E}^+$ and $\mathcal{E}^-$ does not change the polyhedron near v. Therefore, in particular, all incident faces of v (which include f) are unchanged near v and retain their supporting planes and face normals.

For step (2), we consider different kinds of edges. First, let e be an edge where neither endpoint belongs to f^+ or f^-. Then neither endpoint changes coordinates, and so e is unchanged and adjacent to the same faces as before. Next, let e be an edge where exactly one endpoint belongs to f^+ or f^-. The other endpoint of e does not change coordinates, and so the two incident faces at e (which both exist at this endpoint) remain the same. Next, let e be an edge where both ends are on, say, f^+ (the case of f^- is similar). Then e is incident to f^+, and the extrusion $\mathcal{E}^+$ of f^+ contains two edges corresponding to e. When adding $\mathcal{E}^+$ to $\mathcal{P}$ hence e disappears, but is replaced by its copy, which is incident to the new location of f^+ as well as the (extended or retracted) other face that was incident to e.

The final (and most complicated) kind of edge to consider for (2) is an edge $e = (v^+, v^-)$ where endpoint v^+ belongs to f^+ and endpoint v^- belongs to f^-. This implies that e is an X-edge and its X-projection belongs to both polygons f^+ and f^-; by Lemma 6 there are at most two such edges.

Case 1: There exists exactly one edge e with ends incident to f^+ and f^-. Consider $f^- \cup f^0 \cup f^+$; see also Figure 5a. Since all three polygons meet the projection of e, but f^+ and f^- share no edge, the facial angle of f^+ at v^+ is 90°, as is the facial angle of f^- at v^-, and the dihedral angle at e. Hence up to symmetry, the situation at e looks (locally, i.e., in a small neighbourhood around e) exactly as in Figure 5b.

Consider the four extruded quadrants at e. Exactly one is occupied by $\mathcal{F}^0$. Two of them belong to $\mathcal{E}^+$ and $\mathcal{E}^-$ and become filled in $\mathcal{P}'$. The final quadrant remains empty in $\mathcal{P}'$, resulting in a new edge e' in $\mathcal{P}'$. While the facial angles at the ends of e' are different from those at e, an inspection of the resulting structure (see Figure 5c)

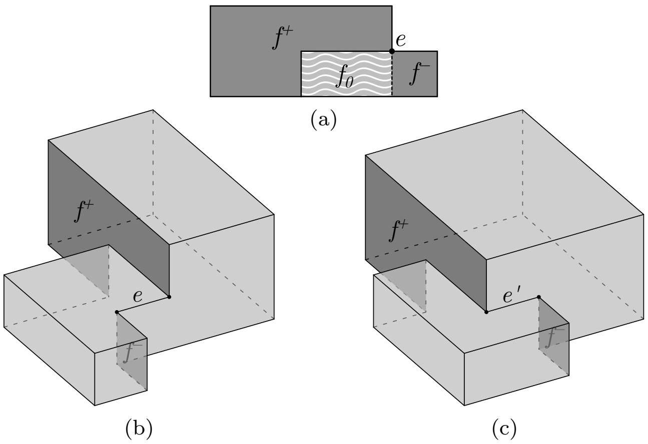

Figure 5: Case 1 of Lemma 7: one edge e has ends in f^+ and f^-.

shows that the ends of e' have the same adjacent faces, in the same clockwise order, as the ends of e.

Case 2: There exist two edges e_1, e_2 with ends incident to f^+ and f^-. Consider again $f^- \cup f^0 \cup f^-$; see Figure 6a, 6b. The two points that e_1 and e_2 project to can be separated by a (horizontal or vertical) line. Suppose there exists a Y-coordinate Y_c such that $Y(e_2) < Y_c < Y(e_1)$; we may pick Y_c such that no vertex of $\mathcal{P}$ has this Y-coordinate.

Now cut $\mathcal{P}$ with the $(Y = Y_c)$-plane into two polyhedra $\mathcal{P}_1$ and $\mathcal{P}_2$. By construction $\mathcal{P}_i$ contains e_i (for $i = 1, 2$), and hence contains parts of f^- and f^+ as a bad pair. So we can apply the swap of this bad pair to each $\mathcal{P}_i$ separately. Since each of them has only one edge which begins and ends at the faces of the bad pair, Case 1 applies and the graph of $\mathcal{P}_i$ remains unchanged. It only remains to argue that the two resulting polyhedra $\mathcal{P}'_1$ and $\mathcal{P}'_2$ can be re-combined. By construction the $(Y = Y_C)$-plane intersects no vertices, so it only intersects Y-edges of $\mathcal{P}$. For any Y-edge not in f^+ or f^-, the coordinates of its endpoints remain unchanged during the swaps in $\mathcal{P}_1$ and $\mathcal{P}_2$, so the pieces can simply be recombined. For any Y-edge that belongs to f^+, both endpoints increase X-coordinates by $X(f^-) - X(f^+)$, regardless of which one of $\mathcal{P}_1, \mathcal{P}_2$ they belong to. Hence again the pieces can be re-combined. Similarly one argues for a Y-edge in f^-. So in summary, for each of the two parts of $\mathcal{P}$ the transformation maintains the graph, and the two graphs can be re-combined and hence give the same graph as in $\mathcal{P}$. This finishes the proof of (2).

It remains to show that $\mathcal{P}'$ is orthogonally convex. We will use Property 2 of Theorem 3 and show that any two points $p, q \in \mathcal{P}'$ are joined by an XYZ-monotone path in $\mathcal{P}'$. If both points are in $\mathcal{P}$, then we already know the property holds. The case where both points are in the middle portion $\mathcal{F}^0 \cup \mathcal{E}^+ \cup \mathcal{E}^-$ is also fine because this polyhedron is the extrusion of $f^- \cup f^0 \cup f^+$, which

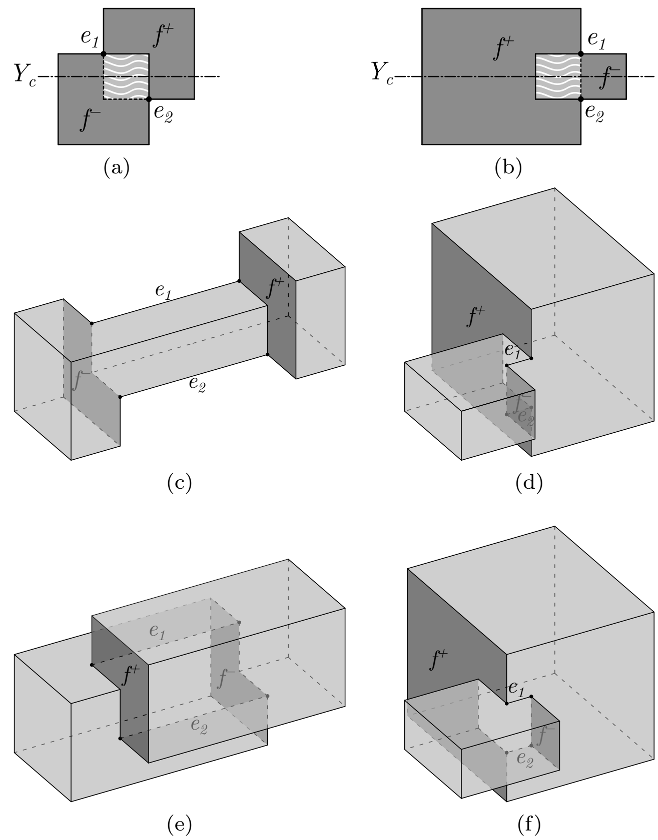

Figure 6: Case 2 of Lemma 7: two edges have ends in f^+ and f^-. These two edges may or may not belong to a common Y-plane or Z-plane.

is orthogonally convex by Lemma 5. So suppose that $p \in \mathcal{E}^+$ (the case where $p \in \mathcal{E}^-$ is analogous.) Let p' be the projection of p in the X-direction to f^+. We will separate into two cases depending whether q is in $\mathcal{F}^-$ or $\mathcal{F}^+$.

Case a: $q \in \mathcal{F}^-$. We get an XYZ-monotone path in $\mathcal{P}'$ from p to q by taking segment pp' followed by an XYZ-monotone path in $\mathcal{P}$ from p' to q.

Case b: $q \in \mathcal{F}^+$. Take an XYZ-monotone path from q to p' in $\mathcal{P}$. Cut it where it crosses the X-plane of p into subpaths σ_1 and σ_2. Let σ'_2 be the projection of σ_2 to the X-plane of p. Observe that $\sigma_1\sigma'_2$ is an XYZ-monotone path from q to p as it is formed by removing X-segments of $\sigma_1\sigma_2$. It remains inside $\mathcal{P}'$ because the YZ-projection between p and p' is uniform. □

With Lemma 7 in hand, proving Theorem 1 is easy, by induction on the number of face-pairs that consist of an α^+-face preceeding an α^--face in some α-face sequence. If there are no such pairs, then we have an orthoball by Property 3 of Theorem 4. If there is such an "out-of-order" pair, then there is a consecutive out-of-order pair, i.e., a bad pair, and we can apply Lemma 7 to exchange this bad pair without affecting any other pairs.

5 Restrictions and Possible Extensions

In this section we remark on some of the restrictions we placed on our polyhedra and graphs.

Degree 3 Vertices. Our construction requires that all vertices in the face-labelled graph have degree 3. Without this requirement, there exist face-labelled graphs that can be realized as orthogonally convex polyhedra, but not as orthoballs. See the example in Figure 7, and observe that the degree 4 vertex v prevents us from performing a swap to create an orthoball. But degree 4 (or higher) vertices do not always prevent realizations of face-labelled graphs as orthoballs. Exact characterization of such graphs would be interesting.

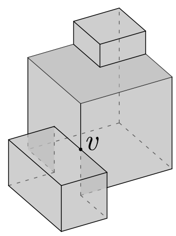

Figure 7: Orthogonally convex polyhedron, but not realizable as an orthoball due to degree-4 vertices.

Simply Connected Faces. We suspect our construction still works when faces need not be simply connected. In the presence of the degree-3 constraint, the only additional possibility is faces with “holes”, i.e., a disconnected graph. We conjecture that any such orthogonally convex polyhedron can be realized as an orthoball as follows. First, no face can be incident to three connected components of the face-labelled graph. Hence the connected components form a tree-structure (with an arc if and only if two connected components share a face.) Say an α^c-face (for some $\alpha \in \{X, Y, Z\}$ and $c \in \{+, -\}$) is incident to two components. Then by orthogonal convexity one of the two subgraphs separated by this face cannot contain an $\alpha^{\overline{c}}$-face, where $\overline{c}$ is the opposite of c. It follows that at most one connected component can have all six face-types, and vice versa, one such component must exist; call this the *root-component* and root the tree at it.

Our idea is to study each rooted sub-tree separately (after replacing the α^c-face that separates it from its parent in the tree by an $\alpha^{\overline{c}}$-face) and to build from it an *ortho-pyramid*, i.e., an orthoball whose α-projection is exactly that $\alpha^{\overline{c}}$-face. The face between the subtree and its parent must be an extreme face in the realization of the parent, so we can then paste each ortho-pyramid onto the appropriate extreme face of the parent component to obtain an orthoball even for disconnected graphs.

Constraints on Edge Directions or Facial Angles. One cannot constrain edge directions or facial angles to remain unchanged in the underlying graph while converting an orthogonally convex polyhedron into an orthoball. Refer to Figure 5. If an edge has one endpoint on each of a pair of bad faces f^+, f^-, then maintaining the edge direction would force the order of f^+, f^- and hence preclude conversion into an orthoball. Furthermore, since facial angles are determined by edge directions, they also cannot be preserved during such a transformation.

6 Conclusions and Future Work

We have shown that the [face-labelled] graphs of simple orthogonally convex polyhedra are the same as the [face-labelled] graphs of simple orthoballs. Can this be used to characterize these graphs?

Acknowledgements. The question of reconstructing an orthogonal polyhedron from a graph with specified face normals was posed by Alla Sheffer.

References

[1] T. C. Biedl and B. Genç. Cauchy's theorem for orthogonal polyhedra of genus 0. In *European Symposium on Algorithms (ESA)*, volume 5757 of *LNCS*, pages 71–82. Springer, 2009.

[2] T. C. Biedl and B. Genç. Stoker's theorem for orthogonal polyhedra. *Int. J. Comput. Geometry Appl.*, 21(4):383–391, 2011.

[3] D. Eppstein and E. Mumford. Steinitz theorems for orthogonal polyhedra. In *Symposium on Computational Geometry*, pages 429–438, 2010.

[4] G. J. E. Rawlins and D. Wood. Ortho-convexity and its generalizations. In G. T. Toussaint, editor, *Computational morphology: a computational geometric approach to the analysis of form*, Machine intelligence and pattern recognition, pages 137–152. North-Holland, 1988.

[5] G. Rote. Realizing planar graphs as convex polytopes. In M. Kreveld and B. Speckmann, editors, *Graph Drawing*, volume 7034 of *LNCS*, pages 238–241. Springer Berlin Heidelberg, 2012.

[6] E. Steinitz. Polyeder und Raumeinteilungen. *Encyclopädie der mathematischen Wissenschaften, Band 3 (Geometries)*, pages 1–139, 1922.

[7] M.-T. Tsai and D. B. West. A new proof of 3-colorability of eulerian triangulations. *Ars Mathematica Contemporanea*, 4(1), 2011.

[8] G. M. Ziegler. *Lectures on Polytopes*. Springer, 1995.

Appendix

Theorem 2 *The following are equivalent for a connected orthogonal (2D) polygon P:*

1. *P is orthogonally convex.*
2. *For any two points p, q in P, there exists a path from p to q inside P that is XY-monotone [4].*
3. *For any axis-aligned rectangle R, the intersection $R \cap P$ is empty or connected.*
4. *For any axis-aligned rectangle R, the intersection $R \cap P$ is empty or an orthogonally convex polygon.*

Proof. To prove (1) $\Rightarrow$ (2), recall that P is connected and hence any two points p, q in P can be connected by a curve inside P. Let C be the shortest such curve. We claim that C is XY-monotone. Assume for contradiction that C is not X-monotone. Then some X-coordinate X_0 is visited twice by C, i.e., there are two points p_1, p_2 with X-coordinate X_0 that belong to C, but the sub-curve $C[p_1, p_2]$ between p_1 and p_2 is not the line segment between p_1 and p_2. By (1) the $(X = X_0)$-line intersects P in an interval, so since $p_1, p_2 \in C \subset P$ the line segment $\overline{p_1 p_2}$ belongs to P. If we now replace $C[p_1 p_2]$ by this line segment, we again get a curve from p to q inside P, and it is shorter than C. Contradiction.

To see (2) $\Rightarrow$ (3), let R be any axis-aligned rectangle and let p and q be any points in $R \cap P$. Then the XY-monotone path from p to q inside P necessarily stays inside R by monotonicity, so p and q are connected by a path in $R \cap P$.

Now we show that (3) $\Rightarrow$ (4). Observe first that (3) trivially implies (1) since an α-line is a degenerate axis-aligned rectangle. Let R be any axis-aligned rectangle. There is nothing to show if $R \cap P$ is empty. If it is not, then by (3) $R \cap P$ is connected. Let ℓ be any α-line, for $\alpha \in \{X, Y\}$. Then $\ell \cap P$ is an interval by (1) and $\ell \cap R$ in an interval since R is a rectangle. Since the intersection of intervals is an interval, therefore $\ell(R \cap P)$ is an interval. So $R \cap P$ is orthogonally convex.

Finally (4) $\Rightarrow$ (1) is trivial since we can take R to be a rectangle that encloses all of P. $\square$

Theorem 3 *The following are equivalent for an orthogonal polyhedron $\mathcal{P}$:*

1. *$\mathcal{P}$ is orthogonally convex.*
2. *For any two points p, q in $\mathcal{P}$, there exists a path from p to q inside $\mathcal{P}$ that is XYZ-monotone.*
3. *For axis-aligned box $\mathcal{B}$, the intersection $\mathcal{B} \cap \mathcal{P}$ has at most one connected component.*
4. *For axis-aligned box $\mathcal{B}$, the intersection $\mathcal{B} \cap \mathcal{P}$ is either empty or an orthogonally convex polyhedron.*

Proof. To prove (1) $\Rightarrow$ (2), recall that $\mathcal{P}$ is connected and hence any two points p, q in $\mathcal{P}$ can be connected by a curve inside $\mathcal{P}$. We may assume that the curve is polygonal, i.e., consists of a finite number of line segments. For any polygonal curve C, define $||C||_1$ to be the sum of the L_1-distances of its segments, i.e., if C is defined by points $q_1, q_2, \ldots, q_n$, then set $||C||_1 = \sum_{i=1}^{n-1} |X(q_{i+1}) - X(q_i)| + |Y(q_{i+1}) - Y(q_i)| + |Z(q_{i+1}) - Z(q_i)|$. Now let C be a curve from p to q within $\mathcal{P}$ that minimizes $||C||_1$. We claim that C is XYZ-monotone. Assume for contradiction that C is not X-monotone. Then some X-coordinate X_0 is visited twice by C, i.e., there are two points p_1, p_2 with X-coordinate X_0 that belong to C, but the sub-curve $C[p_1, p_2]$ between p_1 and p_2 is not entirely within the $(X = X_0)$-plane π. By (1) the intersection $\pi \cap \mathcal{P}$ is orthogonally convex, and so there exists a YZ-monotone curve C' from p_1 to p_2 within $\pi \cap \mathcal{P}$. Observe that C' is XYZ-monotone and has the same endpoints $C[p_1, p_2]$, which implies $||C'||_1 \leq ||C[p_1, p_2]||_1$. In fact, the inequality is strict since $C[p_1, p_2]$ leaves plane π while C' does not. Replacing $C[p_1, p_2]$ by C' hence gives a shorter curve, a contradiction.

The proofs of (2) $\Rightarrow$ (3), (3) $\Rightarrow$ (1), and (4) $\Rightarrow$ (1) are the same as in the previous proof (except one dimension higher) and are left to the reader. Finally we must show (1)&(3) $\Rightarrow$ (4). Let $\mathcal{B}$ be any axis-aligned box. There is nothing to show if $\mathcal{B} \cap \mathcal{P}$ is empty. If it is not, then by (3) $\mathcal{B} \cap \mathcal{P}$ is connected. For $\alpha \in \{X, Y, Z\}$ and any α-plane π, set $\pi \cap \mathcal{P}$ is orthogonally convex by (1) and $\pi \cap \mathcal{B}$ in an axis-aligned rectangle. By the previous theorem the intersection of an axis-aligned rectangle with an orthogonally convex polygon is empty or orthogonally convex. So $\pi(\mathcal{B} \cap \mathcal{P}) = (\pi \cap \mathcal{B}) \cap (\pi \cap \mathcal{P})$ is empty or orthogonally convex as desired. $\square$

Lemma 5 *If $\mathcal{P}$ is an orthogonally convex polyhedron and $\mathcal{Q}$ is its projection to a coordinate plane then $\mathcal{Q}$ is an orthogonally convex polygon.*

Proof. We use the result that orthogonal convexity—in 2D or 3D—is equivalent to having orthogonal monotone paths between all pairs of vertices (Theorem 2, Property 2 for 2D, and Theorem 3, Property 2 for 3D). Suppose $\mathcal{Q}$ is the X-projection of $\mathcal{P}$. Any points p, q in $\mathcal{Q}$ are projections of points p', q' in $\mathcal{P}$. There is an XYZ-monotone path between p' and q' in $\mathcal{P}$, and this projects to a YZ-monotone path between p and q in $\mathcal{Q}$. $\square$