Steinitz's Theorem

Will Dowling   Mario Tomba   Dylan Fridman

Abstract

Steinitz's theorem is a beautiful result that completely characterizes the combinatorial structure of convex 3-polytopes, commonly known as convex polyhedra, in terms of their corresponding graphs. In this expository paper, we present a relatively self-contained proof using graph-theoretic tools.

1 Introduction

Steinitz's theorem is a beautiful result that completely characterizes the combinatorial structure of convex 3-polytopes, commonly known as convex polyhedra. From now on, we will mean *convex* 3-polytope any time we say polytope. The statement of the theorem is as follows:

Theorem 1.1 (Steinitz's Theorem [6]). *The function mapping a* 3*-polytope to its corresponding graph is a bijection onto the set of* 3*-connected, simple, planar graphs.*

Even though the statement itself is remarkably simple, it is certainly nontrivial. Some of the reasons that make the statement even more interesting are the fact that no similar statement has been found in higher dimensions, and that no simple proof of the theorem is known. In our paper, we present a graph-theoretical proof of Steinitz's theorem that mostly follows Ziegler's exposition in [10], which establishes a bijection between the set of convex 3-polytopes and the set of simple, planar, 3-connected graphs. However, before diving into the proof, we introduce the concepts from graph theory that are required to understand both the statement and the proof.

Definition 1.1. *Let G be a graph. G is* ***simple*** *if it contains no loops and no multiple edges between two vertices. G is* ***planar*** *if it can be drawn on the plane in such a way that no two edges cross each other. G is* **d−*connected*** *if the graph obtained by removing any $d-1$ vertices of G is connected.*

Throughout our paper, we will be interested in 2-connected and 3-connected graphs. The picture below shows two graphs that are both simple and planar, one of them being 2-connected and the other being 3-connected.

Figure 1. (a) A 2-connected graph that is not 3-connected. (b) A 3-connected graph.

With these few notions from graph theory, we can start the proof.

2 The graph of a convex 3-polytope is simple, planar, and 3-connected

We begin the proof by creating a map that produces the graph of any given 3-convex polytope. We denote this map by Π and we define it as follows: given any convex 3-polytope P, draw a sphere that contains P and, from a point p in the interior of P, (1) use radial projection from p with the vertices of P onto the sphere and (2) move the projected points on the sphere so that they are all contained in the upper hemisphere. As an example, consider Π applied to the tetrahedron:

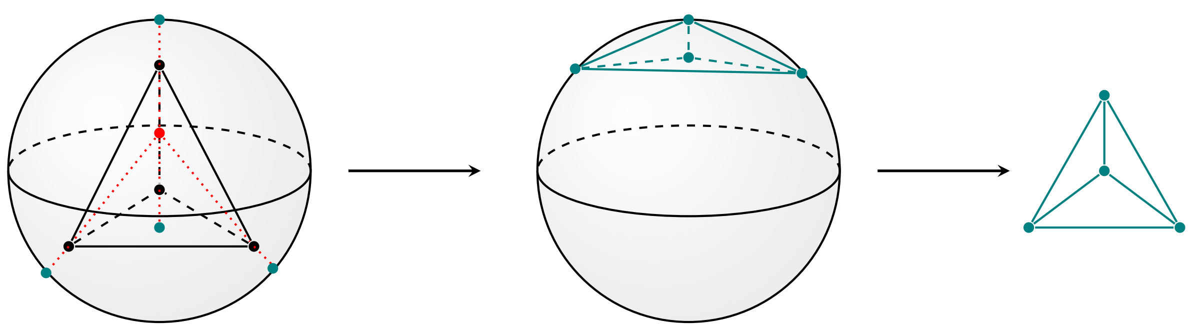

Figure 2. An interior point and the radial projective lines in red. The projected points in green.

As we can see, Π applied to the tetrahedron yields the complete graph K_4. This is easier to see when we project the graph embedded on the upper hemisphere vertically onto the sphere's equatorial plane (which we may identify with $\mathbb{R}^2$). The identification of Π applied to the tetrahedron with K_4 will be important when we show that Π is surjective.

Now note that from the definition of Π, it is clear that the graph of any 3-polytope is simple, since a 3-polytope does not have any loops or multiple edges, and Π is simply projecting. Similarly, Π yields a planar graph, since the projection onto the sphere from an interior point guarantees that no edges cross. Therefore, we only need to carefully establish the 3-connectedness of the graph of a 3-polytope. This statement is known as Balinski's theorem, and it states that the graph of a convex $d-$polytope is $d-$connected [1]. We will only prove the 3-dimensional case of Balinski's theorem, but the proof for the $d-$dimensional case follows from the same argument using induction. We introduce a few more definitions before giving the proof.

Definition 2.1. *A set of vertices X in a graph G* ***separates*** *$x, y \in V(G)$ if every path from x to y contains an element of X. We then say that X is an* **x, y—*separating set*.**

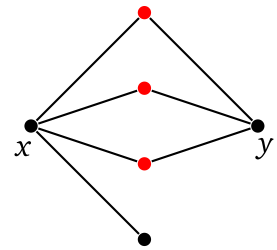

Figure 3. The set of vertices in red is an x, y—separating set

Definition 2.2. *Let P be a polytope and let x be a vertex of P. The* ***link*** *of x, denoted by G_x, is the graph induced by the set of edges of P that are not incident to x but that lie on a face that contains x.*

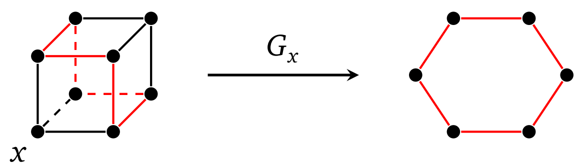

Figure 4. The link of a vertex in the cube

We now have defined all the concepts necessary to prove Balinski's theorem. The following lemma will be fundamental in the proof of Balinski's theorem.

Lemma 2.1. *If P is a 3-polytope and x is any vertex in P, then G_x is isomorphic to the graph of a 2-polytope.*

Proof: Consider the edges $e_1, \ldots, e_k$ incident to x labeled in a way such that e_i and e_{i+1} belong to the same face in P. Let $u_1, \ldots, u_k$ be the other endpoints of $e_1, \ldots, e_k$, respectively. Since the faces of a 3-polytope are polygons, their graphs are cycles. Thus, for each $1 \le i \le k-1$, there is a path from u_i to u_{i+1} that does not contain x but goes through every vertex distinct to x in the face for each i. The same applies to u_k and u_1. Denote these paths from u_i to u_{i+1} by $L_{i,i+1}$. Now note that G_x is the subgraph induced by the set of edges $E(L_{1,2}) \cup E(L_{2,3}) \cup \ldots \cup E(L_{k-1,k}) \cup E(L_{k,1})$, where $E(L)$ is the set of edges in the graph L. Thus, G_x is a cycle, so it is isomorphic to the graph of a polygon. □

To better understand the proof of the lemma above, one may refer to the following picture:

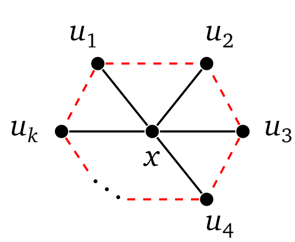

Figure 5. The dashed red lines represent the paths between the vertices adjacent to x. These induce G_x

We need one more result, which is a well-known theorem from graph theory.

Theorem 2.2 (Menger, 1927). *Let G be a graph and let x, y be two non-adjacent vertices in G. Then, the minimum size of a $x, y-$separating set is equal to the maximum number of vertex-disjoin paths from x to y.*

We do not give the proof of Menger's theorem since it is a standard result in graph theory and it would deviate from our topic. However, see [8] for a reference. We are now prepared to state Balinski's theorem for dimension $d = 3$. We follow a proof by Pineda [5].

Theorem 2.3 (Balinski, 1961). *The graph of a 3-polytope is 3-connected.*

Proof: Let P be a 3-polytope and let $G = G(P)$ be its graph. If G is a complete graph, then we are done since G has at least four vertices. Suppose then that G is not the complete graph, and let $y, z \in V(G)$. If y and z are adjacent, then the only way to disconnect them is to remove one (or both) of them. Now note that since P is a 3-polytope, then y and z have degree at least three, and each of the faces containing y and z are polygons, so their graphs are cycles. Therefore, deleting y or z, or both, does not disconnect G. Then consider the case where y and z are not adjacent vertices. Let A be a minimum $y, z-$separating set. By Menger's theorem, there are $|A|$ vertex-disjoint paths from y to z, and each of these paths must contain exactly one vertex from A. This follows because A is a separating set, so every path from y to z contains a vertex from A, so if there are $|A|$ vertex-disjoint paths, each must contain only one vertex from A. Let $L = (u_1, \ldots, u_k)$ be one such path, where $u_1 = y$ and $u_k = z$, and let $u_j \in V(L) \cap A$ be the unique vertex in L contained in A for some $1 < j < k$. Write $u_j = x$. By the lemma above, G_x is isomorphic to the graph of a polygon, and, in particular, G_x is 2-connected. Now note that since every vertex in G has degree at least three, the neighbors of x in G must be vertices in G_x. Therefore, $u_{j-1}, u_j \in V(G_x)$. Now note that $A \backslash \{x\}$ must separate u_{j-1} and u_j. If it did not, then there would be a path $(u_{j-1}, v_1, \ldots, v_n, u_j)$ where $v_1, \ldots, v_n \notin A$, but then $(u_1, \ldots, u_{j-1}, v_1, \ldots, v_n, u_j, \ldots, u_k)$ would be a path from y to z not containing elements from A, which contradicts the assumption that A is a $y, z-$separating set. Hence, since G_x is 2-connected and $A \backslash \{x\}$ separates two vertices in G_x, we have $|A \backslash \{x\}| \geq 2$, which implies $|A| \geq 3$. This establishes the 3-connectedness of G. $\square$

3 Π is injective

In this section, we will show that the projection map Π is injective. To do so, we first need to introduce the notion of graph embedding. Informally, a graph embedding on a given surface is a way to "draw" that graph on that surface. We can formalize this notion as follows:

Definition 3.1. *We say (G, φ, Σ) is an **embedding** of the graph $G = (V, E)$ on the surface Σ if $\varphi : G \to \Sigma$ such that*

- *For every edge $e \in E$, we have that $\varphi(e)$ is an arc on Σ with endpoints $\varphi(v)$ and $\varphi(w)$ where $v, w \in V$ are the vertices incident to e.*
- *The interior of $\varphi(e)$ is disjoint from $\varphi(e')$ and $\varphi(v)$ for any $e' \in E \setminus \{e\}$ and $v \in V$.*

For our purposes, it suffices to consider embeddings on the plane $\mathbb{R}^2$ and the sphere S^2. Note that we can now define what it means for a graph to be planar more rigorously: G is planar if and only if it can be embedded into $\mathbb{R}^2$. Also, note that any embedding into a 2-manifold will partition the manifold into topological disks, which we call **faces**.

Remark 3.1. *G is planar if and only if it can be embedded into S^2. It is clear that if we can embed G on $\mathbb{R}^2$, then we can embed it into S^2 since $S^2 \setminus \{p_0\} \simeq \mathbb{R}^2$. Also, if we can embed G on S^2, we can imagine moving our embedded graph so that it lies entirely on one of the hemispheres of S^2. Then, projecting onto the equatorial plane gives an embedding into $\mathbb{R}^2$.*

Example 1. *The following embeddings are equivalent on the sphere but not on the plane:*

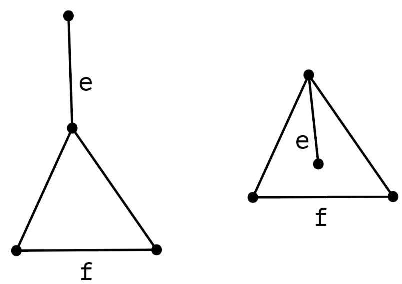

Figure 6. Two embeddings that are equivalent on the sphere but not on the plane.

In the plane, it is clear that the embeddings are distinct since the vertex of degree 1 lies inside a triangle on only one of the embeddings. In the sphere, we could imagine taking the left embedding, fixing e while dragging f through the other hemisphere until it lies "above" e.

Note that if G is the embedded graph resulting from projecting a polytope P into S^2 as described in the definition of Π, then G contains all of the combinatorial structure of P: we know that P, as a combinatorial object, is completely determined by its faces, edges, and vertices with their respective containments. Clearly, this is all information we can obtain from the graph embedding.

Hence, we conclude that Π is actually injective into the collection of 3-connected, planar, simple graph *embeddings*. It turns out that 3-connected, planar, simple graphs have a unique embedding on S^2 up to homeomorphism, a well-known result by Whitney (see [9]). Clearly, this tells us that Π is indeed injective. Thus, the remainder of this section is devoted to establishing Whitney's Theorem.

Theorem 3.2 (Whitney's Theorem). *Any 3-connected, simple, planar graph has a unique embedding on S^2 up to homeomorphism.*

Example 2. *Consider the following distinct embeddings of the same graph on S^2:*

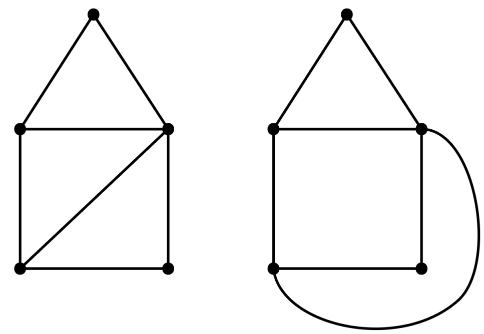

Figure 7. Two distinct embeddings of a simple, planar graph that fails to be 3-connected.

We know that these are distinct embeddings on S^2 because the embedding to the right has a face of degree 4 while the one to the left does not.

Definition 3.2. *Let φ be an embedding of G into S^2 ($\mathbb{R}^2$). Then, we say that an embedding ψ of G into S^2 ($\mathbb{R}^2$) is the **mirror** of φ if, for any vertex $v \in V$, the order of the edges incident to v according to ψ is the same as φ but reversed.*

Intuitively, two graph embeddings are mirrors of each other if you can obtain one of them by flipping the other one like a pancake.

Remark 3.3. *If φ and ψ are homeomorphic embeddings into S^2, then either they are equivalent or they are mirrors of each other.*

Definition 3.3. *A connected subgraph $C \subset G$ is a **cycle** iff every vertex of C has degree 2 in C.*

Lemma 3.4. *Let G be a planar graph and let $C \subset G$ be a cycle. If $G \setminus C$ is connected, then $\varphi(C)$ is the boundary of a face of G under any embedding φ of G into S^2.*

Proof: Suppose that $G \setminus C$ is connected and let φ be an embedding of G into S^2. By the Jordan Curve Theorem, we know that $S^2/\varphi(C)$ consists of two connected components A and B. Since $G \setminus C$ is connected, $\varphi(G \setminus C)$ must be contained entirely in one of the components of $S^2/\varphi(C)$ which we can assume WLOG to be A. Then, $\varphi(C) = \partial B$ and B must be a face of the embedding of G by φ. □

Proof (Whitney's Theorem): Let G be a 3-connected planar graph and suppose by contradiction that we have two distinct embeddings φ and ψ of G onto S^2 that are not mirrors of each other. Then, there must exist a cycle $C \subset G$ such that $\varphi(C)$ is the boundary of a face but $\psi(C)$ is not. By Lemma 3.4, this tells us that $G \setminus C$ is disconnected. Hence, we can write $G = A \cup B \cup C$ where A, B, and C partition the vertices of G and there are no edges between A and B.

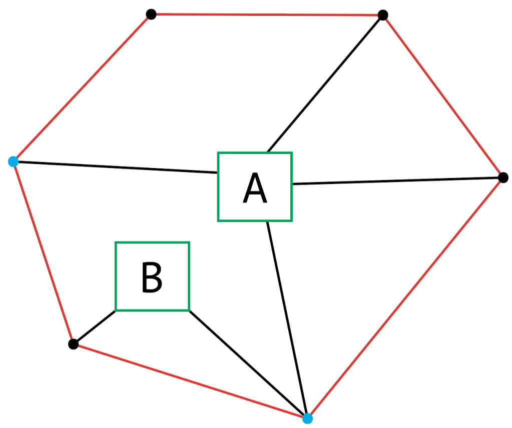

Figure 8. Removing the blue vertices disconnects the graph, contradicting its 3-connectivity.

Since φ embeds G into S^2 and has $\varphi(C)$ as the boundary of a face, we can assume WLOG that $\varphi(C)$ is the boundary of the outer face if we consider φ as an embedding into $\mathbb{R}^2$. Then, there must exist edges e_1 and e_2 adjacent to both A and C such that B is contained in the interior of $\varphi(e_1) \cup \varphi(e_2) \cup \varphi(T) \cup \varphi(S)$ for some $T \subset C$ and some $S \subset \partial A$. Thus, removing the vertex incident to e_1 in C and the vertex incident to e_2 in C disconnects A and B, contradicting the 3-connectivity of G. We conclude that there exists only one embedding of G up to homeomorphism. $\square$

This visual proof of Whitney's Theorem is due to Marc Culler [2].

4 Π is surjective

It remains to show that the map Π is surjective onto the set of simple, 3-connected planar graphs. In particular, given such a graph G, we will show how to construct a polytope P such that $\Pi(P) = G$.

The basic strategy is to perform a series of n reductions on G to yield K_4, which (as mentioned above) corresponds to the tetrahedron under Π. First, we need to define a reduction on graphs that preserves 3-connectedness since K_4 is the smallest 3-connected graph. Then we will see: *(a)* if G' obtained from G by such a reduction is the image of a 3-polytope P' under Π, then G is the image of a 3-polytope P; and *(b)* every 3-connected planar graph is reducible to K_4 in this way, allowing us to inductively construct a 3—polytope P with graph G. The complexity of this proof structure (for Steinitz himself proved the result in this way, though now two other general proof methods are known) is one of the most striking characteristics of Steinitz's theorem, and the structure is no less beautiful if the details are at times tedious (for which reason some such details, mainly Proposition 4.4, will be assumed rather than proven).

Suppose we have an arbitrary 3-connected planar graph G.

Definition 4.1. *Let Δ be K_3, and let Y $K_{1,3}$, the 3-star. Then a ΔY* ***reduction*** *consists of the following two steps:*

(1) Choose a subgraph of G equal either to Δ with no internal vertices (so $G \setminus \Delta$ is connected) or to Y with center of degree exactly 3 (so it is not contained in a larger star). Then replace Δ by Y or Y by Δ on the same vertices. (To be more precise, we will call the former a Δ-to-Y operation, which requires the addition of a new vertex, and the latter a Y-to-Δ operation, which removes a vertex.)

(2) Delete any multiple edges and contract any series edges created by (1).

Step (2) is where the size of G is reduced. A Δ-to-Y reduction and a Y-to-Δ reduction can each reduce the size of G by 0, 1, 2, or 3, so that there are eight classes into which these reductions fall. See Figure 9 for two examples.

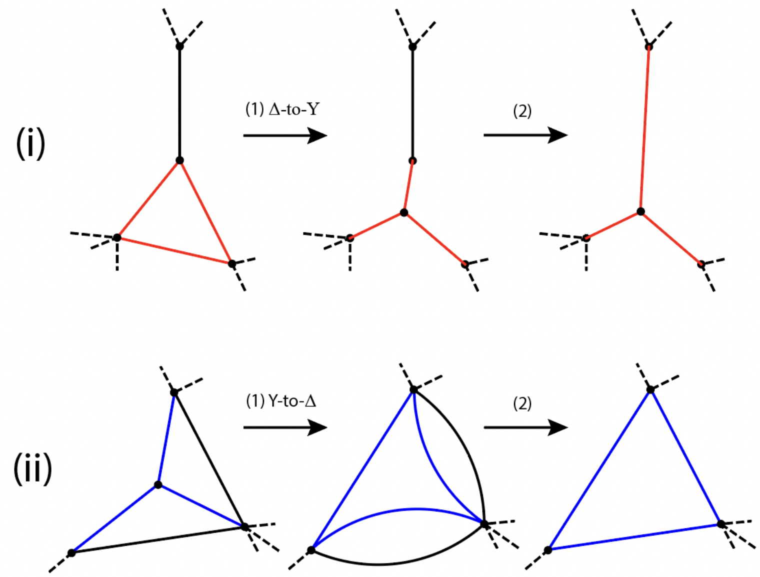

Figure 9. (i) shows a Δ-to-Y reduction that reduces the size of G by 1; (ii) shows a Y-to-Δ reduction that reduces the size of G by 2. Dotted edges are unaffected by the reductions.

From these it is clear that the magnitude of reduction by a Y-to-Δ reduction is equal to the number of edges (3 are possible) that already exist in G between the leaves of Y–each of these results in a multiple edge–and the magnitude of reduction by a Δ-to-Y reduction is equal to the number of vertices of Δ with degree exactly three–each of these results in a series edge. We get the following lemma easily:

Lemma 4.1. *ΔY reductions preserve 3-connectedness.*

Proof: An equivalent definition of 3-connectedness to the one given in Definition 1.1 is that, for any two vertices $u, v \in G$, there are at least 3 disjoint uv-paths. Only one of the 3 (or more) disjoint uv-paths in G can use an edge in the subgraph replaced by a ΔY reduction (or they would meet at a vertex), and it is clear that the reductions replace any subpath contained in Δ or Y by a suitable new one. $\square$

Now we know by what means we can reduce G to K_4. How can we then 'reverse' this reduction to build up P from the tetrahedron such that $\Pi(P) = G$? The following lemma tells us:

Lemma 4.2. *Let G' be obtained from G by a ΔY reduction. Then if G' is the graph of a 3-polytope, so is G.*

Before we begin the proof, recall that, by Lemma 3.4, each of the regions in the graph G' (when drawn with no crossing edges, which is always possible for a planar graph) represents a face of the 3-polytope P' which has G' as its graph. Thus, by a slight abuse of language, we can refer unambiguously to these regions in G' as faces. In Figure 11, for instance, we will call the triangular region of the right-hand graph the 'Δ face.' Further, we will say a point a is 'below' a face f of P' if a and $P' \setminus f$ are in the same half-space defined by the plane on which f lies, and we will say b is 'above' f if it is in the other half-space.

Proof: Suppose P' is the 3-polytope whose graph is G'. We will prove the lemma by describing the construction on P' that gives a new 3-polytope P with graph G. Each of the eight classes of reduction requires a slightly different construction.

Suppose G', the graph of 3-polytope P', is obtained from G by a Δ-to-Y reduction which reduces the size of G by 1. (This is the first example in Figure 9.) Take the plane defined by two of the three leaves of the Y in G' produced by the reduction (points a and b) and a point on the third edge (whose leaf we do not use) that is not one of the vertices (point c). Cut off the vertex represented by the center of the Y by this plane; the convex hull P of the result is clearly a 3-polytope with graph G. This reversal is shown in Figure 10.

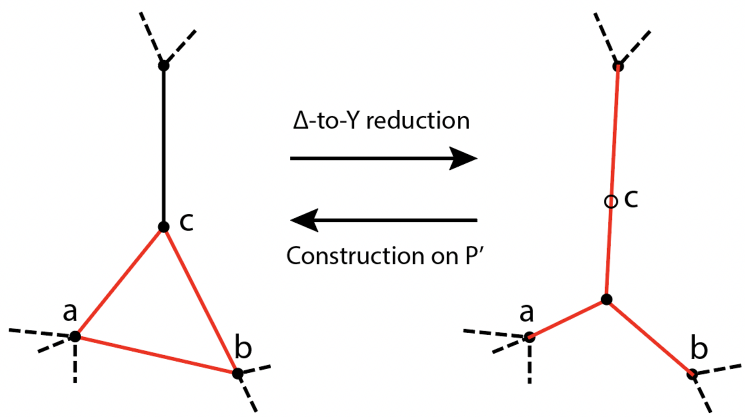

Figure 10. Reversing a Δ-to-Y reduction that reduces the size of G by 1, by cutting a vertex off of P' by the plane through a, b, c.

All Δ-to-Y reductions can be reversed by cutting off the vertex representing the Y they produce by a suitable plane.

Now suppose that G', the graph of 3-polytope P', is obtained from G by a Y-to-Δ reduction which reduces the size of G by 2. (This is the second example in Figure 9.) Take the convex hull of P' with a point p that is above the Δ face, lies on the plane on which face A is situated, and below faces B and C. The result is a 3-polytope P, and it has graph G. This reversal is shown in Figure 11.

All Y-to-Δ reductions can be reversed by taking the convex hull of P' with a suitable point in this way.

There is a problem, however, in reversing a Y-to-Δ reduction which preserves the size of G. For we want, similarly to the example described above, to take the convex hull of P' with the

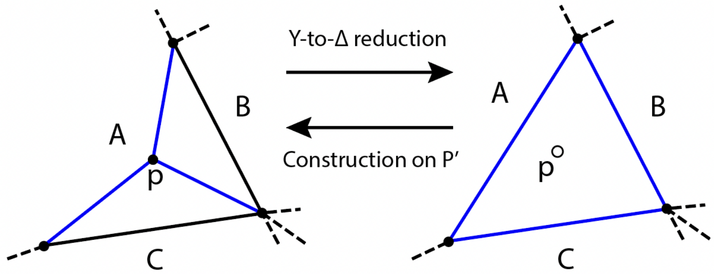

Figure 11. Reversing a Y-to-Δ reduction that reduces the size of G by 2, by taking the convex hull of P' with p on the same plane as A, above the Δ face, and below B and C.

intersection of the three planes given by faces A, B, C. If this convex hull is to be a 3-polytope, the planes must intersect at a point above the Δ face, but this is not necessarily true: they could be parallel or intersect at a point below the Δ face. (We do not have this problem for any of the other classes of Y-to-Δ reductions because the strictest requirement for the selected point is that it lies on a line intersecting the Δ face, so we can always choose a point that will yield a 3-polytope.) This problem is resolved as follows: fix the Δ face in place, and select a point on its interior. Extend a projective ray from this point to all vertices in P' except for those on the Δ face; push each vertex outward along this ray by the same scalar multiple of the ray's length. (Notice that this preserves the combinatorial properties of P', if not the geometric ones. We can imagine this process as blowing the 3-polytope up like a balloon.) Some suitably large scalar shifts the intersection point onto (or generates one on) the correct side of P', and we can proceed as above.

(Note that K_4 is never obtained directly from this problematic reduction. Clearly, the described solution would not work for any selection of three faces of the tetrahedron. More generally, we can say that there are no shared edges among three faces with edges on the Δ face obtained by this particular reduction. The existence of an edge shared by any two of these faces would mean that the graph on which the reduction was performed had a vertex of degree 2, which is impossible in a 3-connected graph.) □

Corollary 4.3. *Any graph that is reducible to K_4 by a series of ΔY reductions is the graph of a 3-polytope.*

Notice how much more involved it is to find the polytope of a given graph than to find the graph of a given polytope. Lemma 4.2 is especially valuable because it not only tells us that the construction of a polytope with ΔY-reducible graph G is theoretically possible but also shows us how to perform it in actuality. See Figure 12 for an example of the whole process–graph reduction and 3-polytope construction–beginning with a ΔY reducible graph G whose polytope is by no means intuitively clear.

Armed with Corollary 4.3, the final piece of the puzzle is the following proposition:

Proposition 4.4. *Every 3-connected planar graph is reducible to K_4 by a series of ΔY reductions.*

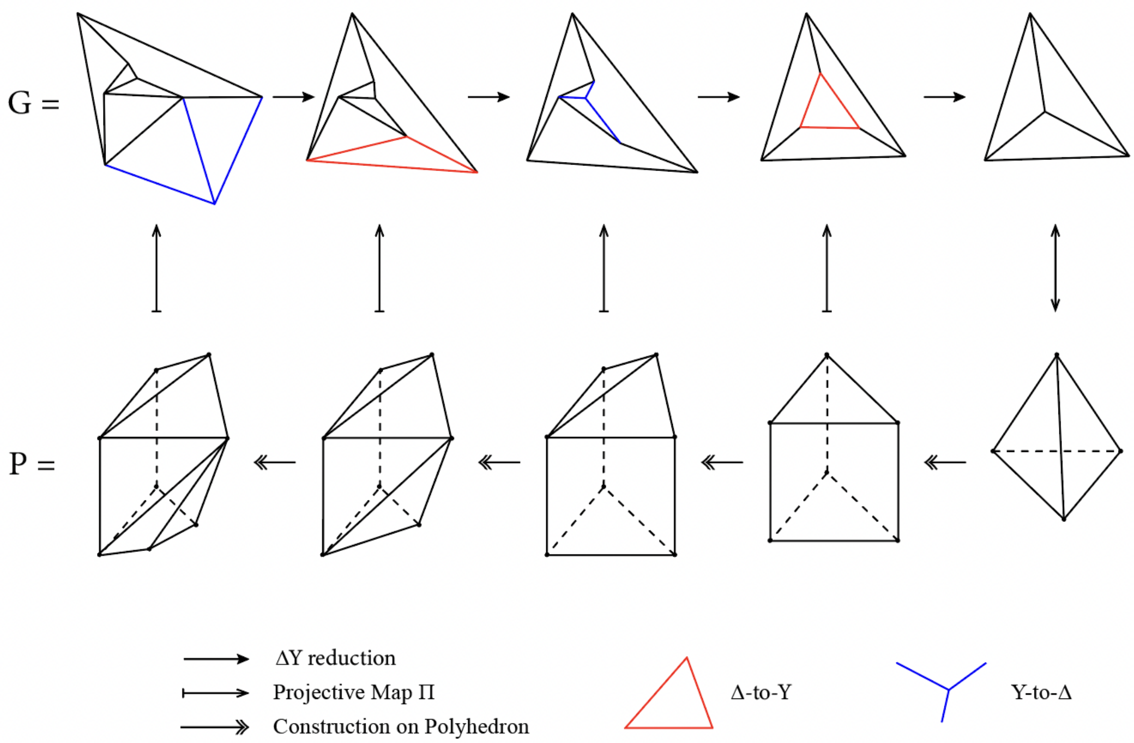

Figure 12. Constructing a 3-polytope P with given 3-connected planar graph G.

We could prove this proposition in a number of ways. For instance, since there are 6 classes of ΔY reduction that actually do reduce the size of G and two that do not, it suffices by induction to show that for any 3-connected graph on more than four vertices either one of the size-reducing classes is applicable, or the two size-preserving ones can be applied so as to make this true. It is tedious to show this and requires some very complex graph-theoretical concepts. See [3, Chapter 13], or Steinitz's own proof in [7] for details. We could also show: *(1)* any 3-connected minor of a ΔY-reducible graph is ΔY-reducible as well; *(2)* every 3-connected planar graph is the minor of a grid graph; and *(3)* every grid graph is ΔY-reducible to K_4. In combination, these prove the proposition. None of these lemmas has a very interesting proof; see [10, Chapter 4] if interested. A proof of the proposition will not be shown here.

This completes our proof of Steinitz's Theorem: we have shown that the projective map $\Pi : \{\text{convex 3-polytopes}\} \to \{\text{simple, 3-connected planar graphs}\}$ is both injective and surjective, and so bijective.

References

[1] M. L. Balinski. On the graph structure of convex polyhedra in n-space. *Pacific J. Math.*, 11:431–434, 1961.

[2] Marc Culler. Planar graphs. http://homepages.math.uic.edu/~culler/talks/planargraphs.pdf. Accessed: 2023-03-05.

[3] Branko Grünbaum. *Convex Polytopes*, volume 221 of *Graduate Texts in Mathematics, 2nd edition*. Springer-Verlag, New York, 2003.

[4] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*, volume 10. JHU press, 2001.

[5] Guillermo Pineda-Villavicencio. A new proof of balinski's theorem on the connectivity of polytopes. *Discrete Mathematics*, 344(7):112408, 2021.

[6] Ernst Steinitz. *Polyeder und Raumeinteilungen*. Encyclopädie der mathematischen Wissenschaften. 1922.

[7] Ernst Steinitz and Hans Rademacher. *Vorlesungen Über die Theorie der Polyeder*. Julius Springer Verlag, Berlin, 1934.

[8] Douglas B. West. *Introduction to graph theory*. Pearson, 2018.

[9] Hassler Whitney. Congruent graphs and the connectivity of graphs. *Hassler Whitney Collected Papers*, pages 61–79, 1992.

[10] Günter M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

[11] Günter M Ziegler. Convex polytopes: extremal constructions and f-vector shapes. *arXiv preprint math/0411400*, 2004.