Continuously Flattening Polyhedra Using Straight Skeletons

Zachary Abel
MIT Department of Mathematics
77 Massachusetts Ave.
Cambridge, MA 02139, USA
zabel@mit.edu

Erik D. Demaine
Martin L. Demaine
MIT CSAIL
32 Vassar St., Cambridge
MA 02139, USA
{edemaine,mdemaine}@mit.edu

Jin-ichi Itoh
Faculty of Education
Kumamoto University
Kumamoto, 860-8555, Japan
j-itoh@kumamoto-u.ac.jp

Anna Lubiw
Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, N2L 3G1, Canada
alubiw@uwaterloo.ca

Chie Nara
Liberal Arts Education Center
Aso Campus, Tokai University
Aso, Kumamoto, 869-1404, Japan
cnara@ktmail.tokai-u.jp

Joseph O'Rourke
Department of Computer Science
Smith College
Northampton, MA 01063, USA
orourke@cs.smith.edu

ABSTRACT

We prove that a surprisingly simple algorithm folds the surface of every convex polyhedron, in any dimension, into a flat folding by a continuous motion, while preserving intrinsic distances and avoiding crossings. The flattening respects the straight-skeleton gluing, meaning that points of the polyhedron touched by a common ball inside the polyhedron come into contact in the flat folding, which answers an open question in the book *Geometric Folding Algorithms*. The primary creases in our folding process can be found in quadratic time, though necessarily, creases must roll continuously, and we show that the full crease pattern can be exponential in size. We show that our method solves the fold-and-cut problem for convex polyhedra in any dimension. As an additional application, we show how a limiting form of our algorithm gives a general design technique for flat origami tessellations, for any spiderweb (planar graph with all-positive equilibrium stress).

Categories and Subject Descriptors: F.2.2 [Nonnumerical Algorithms and Problems]: Geometrical problems and computations

General Terms: Theory, Algorithms

Keywords: flattening, folding, straight skeleton, medial axis, fold-and-cut, tessellations, origami

1. Introduction

When you step on a cardboard cereal box to flatten it, some natural folds appear that are like the folds of a paper bag—each narrow side of the box has a fold down the middle joining a triangle at the bottom. See Figure 1. To capture these natural folds for any 3D polyhedron, Demaine et al. [15] examined the mapping (or "gluing") of the surface induced by the flattening. Imagine piercing the flat folding with a perpendicular line that does not go through any creases. The line pierces successive layers of the flat folding, alternately entering and exiting the polyhedron, and each entering point is glued to the next exiting point. The natural flattening of the cardboard box generalizes to the *straight-skeleton gluing* which, for a convex polyhedron, glues points together precisely when there is a ball inside the polyhedron touching those points. Demaine et al. [15] conjectured that any convex 3D polyhedron has a flattening that respects the straight-skeleton gluing. This is Open Problem 18.3 in Demaine and O'Rourke's book [16]. Our main result is that every convex polyhedron can be flattened according to the straight-skeleton gluing. Furthermore, the flattening process is continuous.

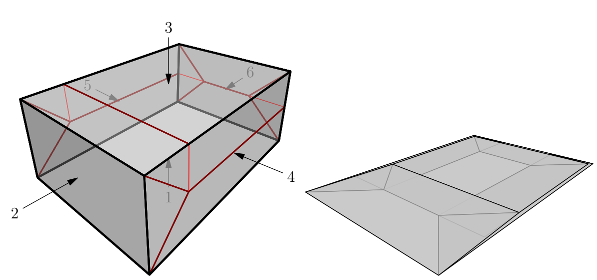

Figure 1. Flattening a box.

What does it mean to flatten a polyhedron? Cauchy's rigidity theorem for convex polyhedra and the Bellows theorem for nonconvex polyhedra preclude any flattening process that keeps faces rigid, even if a finite set of additional creases are added; see [16]. Thus, to flatten a polyhedron, we must allow the surface of the polyhedron to move more flexibly. One approach is to ignore the folding process and just prove the existence of a final flat folded state. Bern and Hayes [9] took this approach and used the disk-packing method to prove that every polyhedral surface homeomorphic to a sphere or disk has a flat folded state. The folds are not particularly natural, and it is unknown whether the flat folded state can be reached by a continuous folding process.

Recently, Itoh, Nara, and Vilcu [22] proved that every convex polyhedron can be flattened via a continuous flattening process that repeatedly "pinches off" a vertex of the polyhedron to form a doubly covered triangle. Their method does not yield a straight-skeleton gluing. In particular, the vertex they pinch off acquires two incident folds, whereas the straight-skeleton gluing always produces at least as many folds as incident edges at every vertex. Also, it is not easy to compute the folds with their method. After pinching off a vertex, they appeal to Alexandrov's theorem [5] to show that the result is again a convex polyhedron Q. Their continuous folding process uses Q, and the pinched-off doubly covered triangle is folded around Q. However, computing Q is difficult, and can only be done approximately [10, 23].

Our Results. We show that any convex polyhedron is flattened according to the straight-skeleton gluing by the following simple process. Pick any total order on the planes that contain the faces of the polyhedron. Move each plane parallel to itself at constant speed towards the interior of the polyhedron. When a point of the surface becomes incident to a lower-numbered plane, the point "sticks" to that plane and moves with it, until the point becomes incident to an even lower-numbered plane and so on. We call this process *orderly squashing*. See Figures 2 and 3.

Several properties are immediately clear: orderly squashing is a continuous process, and the surface does not penetrate itself. It is also easy to show that orderly squashing flattens the polyhedron—in particular, the polyhedron is flat when all points are stuck to the first plane. What we must prove is that orderly squashing does not stretch or compress the surface.

The flexing of the surface during orderly squashing is quite limited—just that a straight crease/fold moves (or "rolls") across the surface. To visualize this, imagine applying wallpaper to a wall. The line of contact between the wallpaper and the wall rolls along the wallpaper.

Orderly squashing, as described above, does not keep the surface of the convex polyhedron convex. In order to apply induction, we must generalize to *positive hyperplane arrangements*. We also generalize to any dimension.

Let $H_1, H_2, \ldots, H_n$ be a sequence of hyperplanes in d-dimensional Euclidean space. Let H_i^+ be the halfspace to one side of hyperplane H_i. Let u_i be the unit normal vector of hyperplane H_i, directed towards H_i^+. Let *positive cell* P_i be the intersection of the positive halfspaces $H_1^+, H_2^+, \ldots, H_i^+$ and let P_0 be the whole space. Observe that $P_i \supseteq P_{i+1}$.

Let *face* F_i be the intersection of hyperplane H_i with the intersection of the previous halfspaces P_{i-1}. Then the union $F = F_1 \cup F_2 \cup \ldots \cup F_n$ is a *positive hyperplane arrangement*. See Figures 4 and 5 for examples in two and three dimensions. We will assume throughout that $u_1 \neq u_2$, i.e., H_1 and H_2 are not parallel and directed the same way. We allow $u_1 = -u_2$.

Observe that any convex polyhedron can be extended to a positive hyperplane arrangement—take the hyperplanes containing the faces in any order and choose the positive side of each hyperplane towards the inside of the polyhedron.

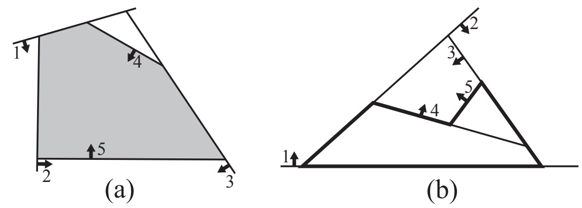

Figure 4. Examples of 2D positive hyperplane arrangements: (a) P_5 is shaded; (b) a nonconvex polygonal subset is shown in bold.

Our main result is that any bounded subset of a positive hyperplane arrangement can be flattened, and furthermore, can be flattened continuously. The flattening motion is the obvious extension of that described above: every hyperplane H_i moves in its positive normal direction at constant speed, and when a point of face F_i becomes incident to a lower numbered hyperplane $H_j, j < i$, the point "sticks" to that hyperplane and moves with it until the point becomes incident to an even lower numbered hyperplane and so on. When points of F_i join a lower numbered hyperplane H_j, we say that [part of] F_i is *folded onto* H_j. We call the whole process *orderly squashing* of the positive hyperplane arrangement. Note that if H_1 and H_2 are parallel and directed the same way (i.e. $u_1 = u_2$) then the arrangement will never fold onto H_1; this is the reason for our assumption that $u_1 \neq u_2$. An example is shown in Figure 5.

Orderly squashing is a continuous transformation. It is easy to show that any bounded subset of the arrangement will eventually fold entirely onto the first hyperplane. The content of our main result is to show that orderly squashing is an isometry.

THEOREM 1. *Let S be a bounded subset of a d-dimensional positive hyperplane arrangement F in which $u_1 \neq u_2$. Then orderly squashing folds S onto H_1 in finite time, behaves isometrically on all of F, and respects the straight-skeleton gluing in positive cell P_n.*

We also study the creases produced by orderly squashing. The *primary creases* on F_i are the boundaries that separate regions of F_i according to the first lower-ordered hyperplane they glue to. The other ("secondary") creases are the ones that form on F_i when a region of F_i has folded onto face $F_j, j < i$ and then F_j acquires a primary crease. We show that the size of the primary crease structure is $O(n^{\lfloor (d+2)/2 \rfloor})$ and that this bound is tight in the worst case. In particular, in 3D the number of primary creases is $O(n^2)$. We give an algorithm to find the primary creases in time $O(n^{\lfloor (d+2)/2 \rfloor})$. For a convex polyhedron we have the freedom to choose any ordering of its hyperplanes. We describe an ordering that reduces the size of the primary crease structure to $O(n^{\lfloor (d+1)/2 \rfloor})$. The total number of creases is exponential in the worst case even for a convex polygon in 2D.

Orderly squashing is not only conceptually simple but also practical to run on real 3D examples. We implemented the general algorithm in Python, which allowed us to automatically generate the images in this paper. The main algorithmic task, namely computing the primary crease pattern

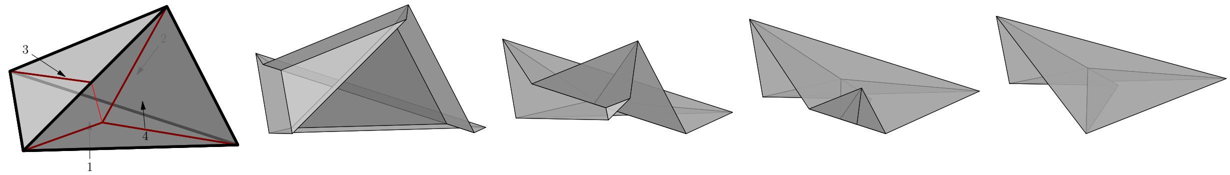

Figure 2. Orderly squashing of a tetrahedron. The left pane shows the ordering of faces and the crease pattern, with primary creases in thick dark red and other creases in thin light red. The other panes show an animation of the flattening. Face 2, originally at the back, ends up on top.

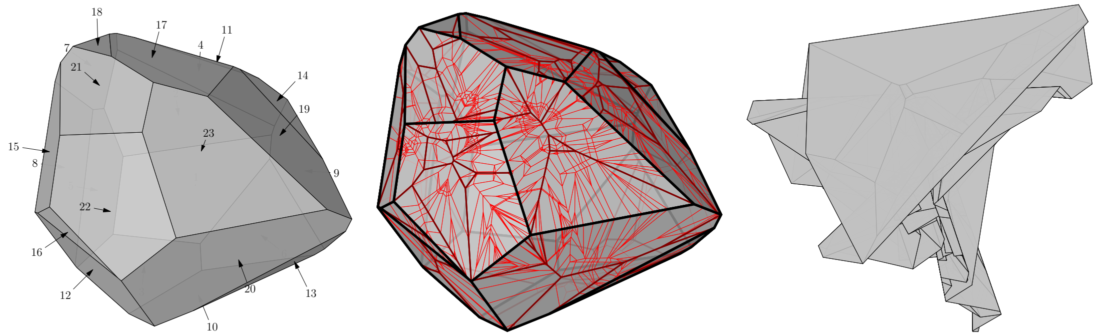

Figure 3. Ordered flattening of a polyhedron: (left) the ordering of faces; (centre) the crease pattern, with primary creases in thick dark red, others in thin light red; (right) the flattened state. Face 2, originally at the back, ends up on top. More details of this example are shown in Figure 11.

of a positive hyperplane arrangement, reduces to computing the straight skeleton of convex polyhedra (specifically the *positive cells* of the arrangement). This in turn reduces to computing the intersection of 4-dimensional half-spaces as explained in Section 2, which is dual to computing convex hulls. Our orderly squash implementation calls Qhull [7] as a subroutine for this step. While the code makes no attempt at optimization, it is still practical on decently-sized examples: on a laptop featuring a 2.90GHz Intel Core i7-3520M processor and 4GB RAM, each image in this paper was computed in no more than 15 seconds, and we computed the 642,295 faces in the overall crease pattern of a randomly generated 3D hyperplane arrangement of 64 planes in just over 70 minutes. More benchmarking information is presented in a forthcoming long version of this paper.

The straight skeleton is a natural approach to polyhedron flattening because of a connection to the *fold-and-cut* problem [16]. Given a d-dimensional polyhedron, the fold-and-cut problem asks for a flat folding of d-space through $(d+1)$-space and back into d-space that maps the surface of the polyhedron to a $(d-1)$-dimensional hyperplane, and maps the inside and outside of the polyhedron to opposite sides of the hyperplane. In particular, restricting the folding to its action on the surface yields a flattening of the polyhedron. The 2-dimensional fold-and-cut problem has been solved using the straight skeleton as a crease pattern [16], which leads to a polygon flattening that respects the straight-skeleton gluing, even for nonconvex polygons. We show how to use orderly squashing in $(d+1)$-dimensions to solve the fold-and-cut problem for any convex d-dimensional polyhedron.

As a further application, we use orderly squashing to provide a simple algorithm for folding planar origami tessellations. Through a subtle limiting process of applying orderly squashing to successively shallower scalings of a polyhedral surface, we show that the result is a well-defined, flat origami fold. This technique applies to an arbitrary *spiderweb* (that is, a planar graph having an positive equilibrium stresses):

THEOREM 2. *Let $G = (V, E)$ be a spiderweb with a positive equilibrium stress $w : E \to \mathbb{R}_{>0}$. Infinitesimal orderly squashing results in a flat-foldable crease pattern containing G augmented with a pleat along each edge e with width proportional to $w(e) \cdot |e|$ such that all remaining creases lie within a small neighborhood of the vertices.*

Lang and Bateman [24] recently provided a different method to fold spiderwebs into origami tessellations using twist folds, but our method is distinguished by not modifying the original pattern G before determining the crease pattern; G itself is among the set of creases. The spiderweb constraint arises naturally in both techniques.

Our paper is organized as follows. Section 2 contains background on the straight skeleton and medial axis. In Section 3 we prove the main theorem stated above. Our results on computing and counting the creases are in Section 4. Section 5 contains a polynomial-time algorithm to test if a bounded subset of a hyperplane arrangement is a subset of a positive hyperplane arrangement. The results on fold-and-cut are in Section 6 and the results on tessellations are in Section 7. Due to space limitations, some proofs are deferred to the final long version of this paper.

2. The Straight Skeleton and the Medial Axis

The *medial axis* of a polyhedron P is the set of points inside P that have at least two closest points on the boundary of P [6]. For a convex polyhedron, the medial axis and the straight skeleton are the same. For bounds on the size and the complexity of computing the medial axis, see [20, 1, 13, 30]. For definitions and results on straight skeletons, see [4, 3, 19, 21, 8].

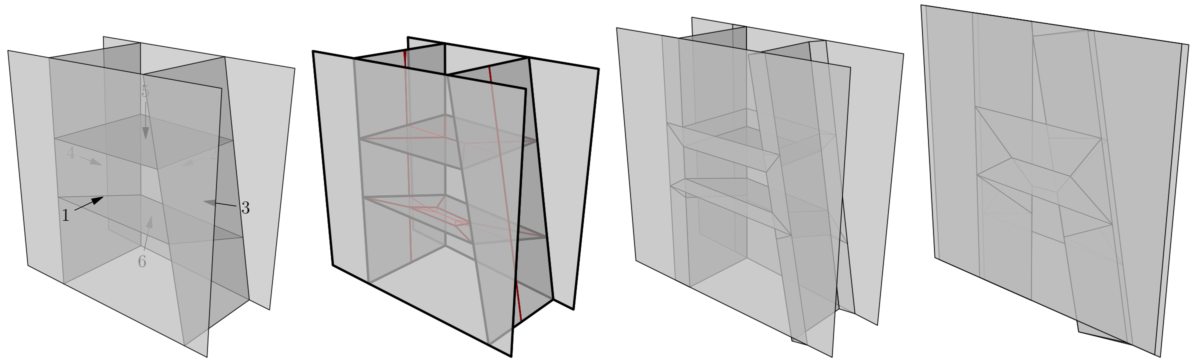

Figure 5. Orderly squashing of a 3D positive hyperplane arrangement: (from left to right) the ordering of hyperplanes; the crease pattern, with primary creases in thick dark red and others in thin light red; two frames of the flattening.

We need one crucial result, which is that the medial axis of a convex polyhedron with n facets in d dimensions has worst-case size $\Theta(n^{\lceil d/2 \rceil})$ and can be computed in the same time bound. Although we have not found an explicit statement of this in the literature (except in one technical report [26]), it is well-known to experts, and follows from general arguments of Edelsbrunner and Seidel [18], which we now summarize. The bounds are proved using a transformation from the medial axis of a convex polyhedron of n facets in d dimensions to the convex hull of n points in $d+1$ dimensions. We first describe the 2D case for ease of visualization. Given a convex n-gon P in the xy plane, construct planes in 3D through the edges of P, where each plane has slope 1 in the z direction and has the interior of P below the plane. The intersection of the lower half-spaces of the n planes is a convex polyhedron whose edges and vertices, when projected to the xy plane, form the medial axis of P. (This is known as the "roof" structure for the straight skeleton [3].) Finding the intersection of n halfspaces in 3D is the dual problem to finding the convex hull of n points in 3D.

In d-dimensional space each of the n facets of the input polyhedron is lifted to a hyperplane in $d+1$ dimensions of slope 1 in the new dimension. The medial axis is the projection of the polyhedron formed by intersecting the halfspaces below the n hyperplanes. In dual space, we want the convex hull of n points in $d+1$ dimensions (see [17, 28]). Such a convex hull has size $O(n^{\lfloor (d+1)/2 \rfloor}) = O(n^{\lceil d/2 \rceil})$ [17, 28] and can be computed in time $O(n^{\lceil d/2 \rceil} + n \log n)$, as shown by Chazelle [11], or by Seidel's earlier randomized algorithm [27]. We note that the $O(n \log n)$ term dominates only for $d = 2$, but in that case, the cyclic order of the input polygon saves the work of sorting—the medial axis of a 2D polygon can be found in linear time [2, 12].

3. Orderly Squashing

In this section we will prove Theorem 1. When only two hyperplanes are involved, the result is straightforward. In general, most of the interactions during orderly squashing involve just two hyperplanes; it is only at a measure-0 subset of F where three or more hyperplanes interact, so by continuity, the collapse should be isometric everywhere. Our proof is a formalization of this idea.

We use the notation from Section 1. For each point $v \in F$ and for each $t \geq 0$, define $h_t(v)$ as the position to which v moves during the orderly squash after each hyperplane has moved by a distance t along its positive normal. Imagine the hyperplanes move at a rate of 1 unit of distance per unit of time, so t also represents the duration of the squash so far. The resulting function $(t, v) \mapsto h_t(v)$ is continuous. It is also easily seen to be piecewise linear.

Next, we define a measure of when a point on F_n will hit another part of F during orderly squashing. For a point $v \in F_n$ and a hyperplane H_i with $u_i \neq u_n$, let $c_i(v)$ be the time at which v would hit hyperplane H_i if no other hyperplanes were present. That is, $c_i(v)$ is the radius of the sphere in $H_n^+ \cap H_i^+$ that is tangent to H_n at v and is tangent to H_i. Note that if v' is the point of tangency of this sphere with H_i, i.e., the reflection of v through this bisecting hyperplane, then v' is precisely the point on H_i that v will make contact with (by symmetry). The function $c_i(v)$ is a linear function on F_n. If $u_i = u_n$ then the hyperplanes H_i and H_n will never meet, so define $c_i(v) = \infty$.

We are now ready to prove Theorem 1:

Proof Proof of Theorem 1. Suppose S is a bounded subset of positive hyperplane arrangement F. We will prove by induction on n that orderly squashing folds S onto H_1 in finite time, behaves isometrically on all of F, and respects the straight-skeleton gluing in every positive cell P_i. Let S_i be the part of S on face F_i.

Note that the behaviors of $S_1 \cup \cdots \cup S_{n-1}$ and $F_1 \cup \cdots \cup F_{n-1}$ are the same whether H_n is present or not. Therefore we can assume by induction that orderly squashing is isometric on $F_1 \cup \cdots \cup F_{n-1}$ and folds any bounded subset of $F_1 \cup \cdots \cup F_{n-1}$ onto H_1 in finite time.

First we show that S_n will eventually fold onto H_1. During orderly squashing, each point $v \in S_n$ will first stick to some other hyperplane after $\min_{1 \leq i \leq n-1} c_i(v)$ seconds. We claim that this minimum is finite, for otherwise $u_n = u_i$ for all i, in which case $u_2 = u_1$ contrary to our assumption. The function $\min_{1 \leq i \leq n-1} c_i(v)$ is a continuous, piecewise-linear function of v and is therefore bounded above on the bounded region S_n. All of S_n will eventually stick to (a bounded subset of) $F_1 \cup \cdots \cup F_{n-1}$. This bounded subset eventually folds onto H_1 by induction.

Next, we show that orderly squashing respects the straight-skeleton gluing of P_n. Suppose the sphere with center $a \in \mathbb{R}^k$ and radius r lies inside P_n and is tangent to it at points v, w (it may have other points of tangency as well). We must show that orderly squashing pairs v and w. Because each hyperplane H_i does not pierce the sphere, the perpendicular distance d_i from a to H_i is at least r. After t seconds for $0 \le t \le r$, the distance from a to the new location of H_i is $d_i - t \ge r - t$, so the sphere with radius $r - t$ around a is not pierced by a hyperplane at time t. So points v and w will stick to each other at time r at the position a, as required.

Finally, we show that orderly squashing behaves isometrically on F_n. Fix a time $T \ge 0$. We must show that, for any rectifiable path $\gamma : [0,1] \to F$ on F, the path $h_T \circ \gamma$ obtained after T seconds of orderly squashing has the same length as γ.

We first decompose F_n based on which hyperplanes are first hit during orderly squashing during time interval $[0,T]$—we call this decomposition the *primary crease pattern* of F_n. Specifically, define $F_n[0]$ as those points $v \in F_n$ that do not stick to another hyperplane during time interval $[0,T)$ (but may precisely at time T), i.e., such that $c_i(v) \ge T$ for each $1 \le i \le n-1$. Similarly, for each $1 \le i \le n-1$, let $F_n[i]$ be the set of points v that stick to hyperplane H_i first during this time interval, i.e., $c_i(v) \le c_j(v)$ for all $1 \le j \le n-1$, and $c_i(v) \le T$. These regions $F_n[0], F_n[1], \ldots, F_n[n-1]$ exhaust F_n, and as they are defined by closed and linear constraints, they are closed, convex, polygonal regions (and possibly unbounded, degenerate, or even empty).

Fix a rectifiable path γ on F. Divide γ into subpaths, each lying on one F_i. Because h_t is continuous it suffices to consider each subpath. By assumption, h_T behaves isometrically on $F_1 \cup \cdots \cup F_{n-1}$ so we may assume that γ lies on F_n. In fact we may further subdivide the path and assume that γ lies on a single region $F_n[i]$ for some i. If $i = 0$, then $h_T \circ \gamma$ is simply a translation of γ and therefore has the same length.

Now suppose γ is on $F_n[i]$ for some $i < n$. If we reflect γ across the H_n-H_i bisecting hyperplane to a new path γ' on H_i, then each point of γ sticks to its corresponding point on γ' in no more than T seconds (by definition of $F_n[i]$). It follows that $h_T \circ \gamma = h_T \circ \gamma'$. By our inductive hypothesis, h_T is an isometry on F_i, so $h_T \circ \gamma'$ has the same length as γ'. But γ' and γ have the same length by symmetry, which proves the claim. □

4. Creases, Events, and Algorithms

We use the notation from Section 1. The *crease structure* is the subdivision of each F_i according the complete sequence of hyperplanes that a point sticks to. The *primary crease structure* is the subdivision of each F_i according to the first hyperplane that a point sticks to.

Before turning to our analysis of creases, we mention that it is possible to take a more "time-centered" view of orderly squashing, similar in flavour to the definition of the straight-skeleton. During orderly squashing the configuration changes combinatorially at a discrete set of events. This can be captured by *shrinking the positive hyperplane arrangement*: We ignore what happens to a point p of F_i after it sticks to a lower numbered hyperplane. Each hyperplane moves at constant speed in its positive normal direction, following the rule that face F_i is always the intersection of hyperplane H_i with the intersection of the previous halfspaces, P_{i-1}. For hyperplanes in general position, events correspond to vertices of the crease structure, although in degenerate situations events may correspond to higher-dimensional faces of the crease structure.

Our analysis of creases will not use the notion of shrinking. Rather, we use two main ingredients: the relationship between orderly squashing and medial axes of the positive cells of the arrangement (from the previous section), and the relationship between medial axis and convex hull (described in Section 2). Hyperplanes $H_j, j > i$ are irrelevant to the crease structure on F_i. By Theorem 1, orderly squashing on $H_1, \ldots, H_i$ respects the straight-skeleton gluing in P_i. This means that the primary crease structure on F_i is the projection onto F_i of F_i's cell of the medial axis of P_i. Note that when we are dealing with a subset S of F, it may happen that a primary crease falls on $F_i - S_i$.

THEOREM 3. *For orderly squashing of a positive hyperplane arrangement of n hyperplanes in $\mathbb{R}^d$, the primary crease structure has size $O(n^{\lfloor (d+2)/2 \rfloor})$ and can be computed in that time bound.*

PROOF. As discussed above, the primary crease structure on F_i is the projection on F_i of F_i's cell in the medial axis of P_i. As discussed in Section 2, the medial axis of P_i is the projection of the intersection of n halfspaces in $(d+1)$-dimensions which has size $O(n^{\lfloor (d+1)/2 \rfloor})$. To find one cell of the medial axis, we restrict the intersection to one halfplane, i.e., we go down one dimension, and the size is $O(n^{\lfloor d/2 \rfloor})$. For the total size of the primary crease structure we sum over $i = 1, \ldots, n$ which gives $O(n^{\lfloor (d+2)/2 \rfloor})$.

We use Chazelle's convex-hull algorithm [11] to compute the primary crease structure on F_i in time $O(n^{\lfloor d/2 \rfloor} + n \log n))$. The $n \log n$ term is only relevant for $d = 2, 3$ and it is easy to see that we can preprocess the input to avoid repeating that term for each i which allows us to compute the primary crease structure for $d = 2, 3$ in time $O(n^2)$. Thus in general the primary crease structure can be computed in time $O(n^{\lfloor (d+2)/2 \rfloor})$. □

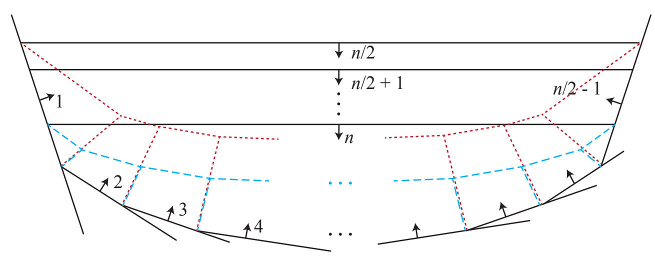

Figure 6. A 2D positive hyperplane arrangement with primary crease structure of size $\Theta(n^2)$. Each of the horizontal lines $n/2, \ldots, n$ interacts with each of the lines $1, \ldots, n/2 - 1$. Dashed blue lines show the medial axis of P_n; dotted red lines show the medial axis of $P_{n/2}$.

The bound in Theorem 3 is tight for $d = 2$ as shown by the example in Figure 6 that has a crease structure of size $\Theta(n^2)$. For general dimension d, we can construct a similar tight example by placing hyperplanes $1, 2, \ldots, \frac{n}{2}$ so that the last hyperplane's cell in the medial axis has size $\Theta(n^{\lfloor d/2 \rfloor})$ (which is possible since the convex-hull bounds are tight) and then placing hyperplanes $\frac{n}{2} + 1, \ldots, n$ parallel to hyperplane $\frac{n}{2}$.

For convex polyhedra, we have the freedom to choose the hyperplane ordering for orderly squashing. In this case the bound given above can be reduced, at least for even n:

THEOREM 4. *For any convex polyhedron in $\mathbb{R}^d$ there is an ordering of its n facets so that orderly squashing of the resulting positive hyperplane arrangement gives a primary crease structure of size $O(n^{\lfloor (d+1)/2 \rfloor})$.*

PROOF. To define the ordering, take the medial axis of the convex polyhedron. Define H_n to be the facet whose cell in the medial axis is the smallest. Remove H_n and repeat.

The size of the primary crease structure is the sum over $i = n, \ldots, 1$ of the size of the cell of F_i in the medial axis of P_i. The medial axis of P_i has size at most $kn^{\lfloor (d+1)/2 \rfloor}$ for some constant k (by bounds in Section 2). Thus, its smallest cell has size at most $kn^{\lfloor (d-1)/2 \rfloor}$. Summing over i, we get a bound of $kn^{\lfloor (d+1)/2 \rfloor}$. □

The example in Figure 7 shows that the bound of Theorem 4 is tight for $d = 3$. The primary crease structure has size $\Theta(n^2)$ regardless of the ordering of faces, because every front face interacts with every back face. The example is based on Held's example of a 3D polyhedron with a medial axis of size $\Theta(n^2)$ [20].

Computing the ordering promised in Theorem 4 requires n medial axis computations which takes time $O(n^{\lfloor (d+3)/2 \rfloor})$. However, as in Seidel's randomized convex-hull algorithm [27], a random ordering of the hyperplanes will give a primary crease structure of expected size $O(n^{\lfloor (d+1)/2 \rfloor})$ with an expected time bound to match.

Although the above theorems give nice bounds on the size of the primary crease structure, the size of the total crease structure can grow exponentially. Figure 8 shows a 2D positive hyperplane arrangement with $2n+1$ lines whose orderly squashing has $\Theta(2^n)$ creases. On the left is the basic 'Π' building block, with the property that line 3 folds double on line 1 during orderly squashing. The plan is to repeat this so that part of line $2n+1$ folds double on line $2n-1$, and then that part of line $2n-1$ folds double on line $2n-3$ and so on. The i-th Π consists of lines $2i-1$ (the "roof") and $2i, 2i+1$ (the "legs"). The Π's alternate between vertical and horizontal. The configuration for $n = 4$ is shown in the right hand side of the figure. The length of the doubled portion of the i-th Π is δ_i (shown in red in the figure), with $\delta_n = 1$ and $\delta_i = \frac{\delta_{i+1}}{2} = 1/2^{(n-i)},\ i = n-1, \ldots, 1$. We define d_i to be half the distance between the legs of the i-th Π, and we set $d_n = 1$ and $d_i = n - i + 1 + \delta_i,\ i = 1, \ldots, n-1$. In the i-th Π there are two time-points of interest: time t_i^0 when the doubled portion from the $(i+1)$-st Π begins to fold onto the roof; and time t_i when doubling completes. Observe that $t_i^0 = d_i - \delta_i$, and $t_i = d_i + \delta_i$. We must verify that $t_i \le t_{i-1}^0$, i.e. that doubling completes in the i-th Π before the doubled portion starts folding onto the roof in the $(i-1)$-st Π. For $i = n$ we have $t_n = 2$ and $t_{n-1}^0 = d_{n-1} - \delta_{n-1} = 2\frac{1}{2} - \frac{1}{2}$. For $i = 2, \ldots, n-1$, we have

$$t_{i-1}^0 - t_i = (n-i+2+\delta_{i-1}-\delta_{i-1}) - (n-i+1+2\delta_i) = 1 - 2\delta_i \ge 0.$$

In the full version of this paper we give an example of a convex polygon in 2D where the number of creases is $\Theta(2^{n/2})$. Unlike the example above, the polygon example requires n-bit numbers for the lines. We conjecture that bounding the bit complexity of the input reduces the number of creases. In particular, we do not believe that an exponential blow-up in the number of creases will easily occur in natural examples.

For general dimensions, we prove the following upper bound:

THEOREM 5. *For orderly squashing of a positive hyperplane arrangement of n hyperplanes in $\mathbb{R}^d$, the number of $(d-1)$-dimensional faces of the crease structure is $O(2^n)$.*

PROOF. We classify points on F by the full sequence of hyperplanes they stick to during the collapse. Each such sequence is a subsequence of $H_n, \ldots, H_1$, so there are 2^n such subsequences, and thus at most 2^n subregions of F. □

For $d = 2, 3$ this implies that the total size of the crease structure is $O(2^n)$. Computing the total crease structure is a straight-forward enhancement of the method described in Theorem 3: Compute the total crease structure $\mathcal{C}_{n-1}$ for $H_1, \ldots, H_{n-1}$. Compute the cell C of F_n in the medial axis of P_n. Consider the face B_i of C that bisects H_n and H_i. Project the relevant portion of $\mathcal{C}_{n-1}$ from H_i to B_i to F_n.

5. Characterizing Subsets of Positive Hyperplane Arrangements

Our theorem says that any bounded subset of a positive hyperplane arrangement can be flattened via orderly squashing. Our main aim was to flatten convex polyhedra, but the result is more general—for example see the nonconvex polygon in Figure 4. In this section we characterize bounded sets that are subsets of positive hyperplane arrangements.

Let S be a bounded set given to us as a subset of a hyperplane arrangement. Let S_i be the part of S on H_i. We seek an ordering of the hyperplanes $H_1, H_2, \ldots, H_n$ and a choice of halfspace H_k^+ for each H_k so that $S_i \cup \cdots \cup S_n$ is contained in $P_i = \cap_{j=1}^i H_j^+$ for all $1 \le i \le n$. Call this a *positive halfspace ordering*. The first halfspace H_1^+ must satisfy $S \subseteq H_1^+$. Among the given halfplanes, let T^+ be any halfspace such that $S \subseteq T^+$. We claim (in the lemma below) that if there is a positive halfspace ordering then there is a positive halfspace ordering beginning with T^+. Therefore, to find a positive halfspace ordering we simply find a first halfspace T^+ that contains S, if one exists, and then recurse on $S - T$. This is a polynomial-time algorithm.

Note that our main theorem required the first two hyperplanes to have different unit normals, but this condition is moot in the current situation because we can simply add a new hyperplane H_0 that is not parallel to any other hyperplane and such that S lies to one side H_0^+ of the hyperplane.

LEMMA 1. *Suppose S is a subset of a positive hyperplane arrangement $H_1^+, H_2^+, \ldots, H_n^+$. Suppose that $T = H_i$ for some i, and that S is contained in one of T's halfspaces, T^+. Then there is a positive halfspace ordering for S that starts with T^+.*

PROOF. Removing T from the sequence $H_1^+, \ldots, H_n^+$ yields a positive halfspace ordering of $S - T$. □

6. *d*-Dimensional Fold and Cut

In this section we show that orderly squashing can be used to solve the fold-and-cut problem for convex d-dimensional polyhedra. Furthermore, our method provides an explicit folding process, which had only been known previously for the case of convex polygons in the plane [14].

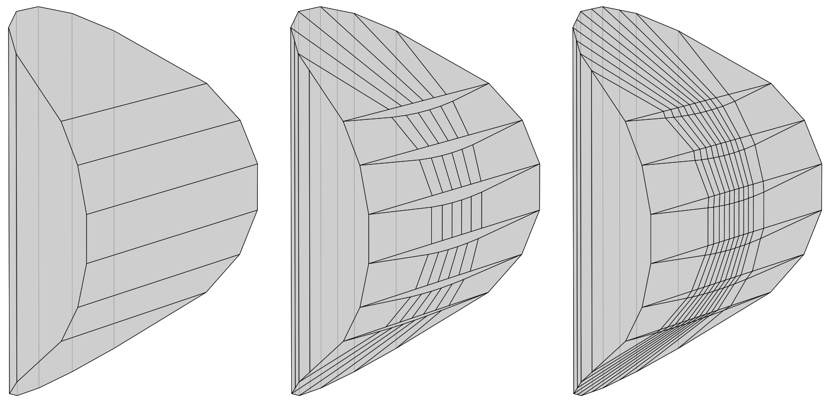

Figure 7. A convex polyhedron where any face ordering produces $\Theta(n^2)$ primary creases: (left) the polyhedron; (middle) the primary creases for one face ordering; (right) all creases for that face ordering.

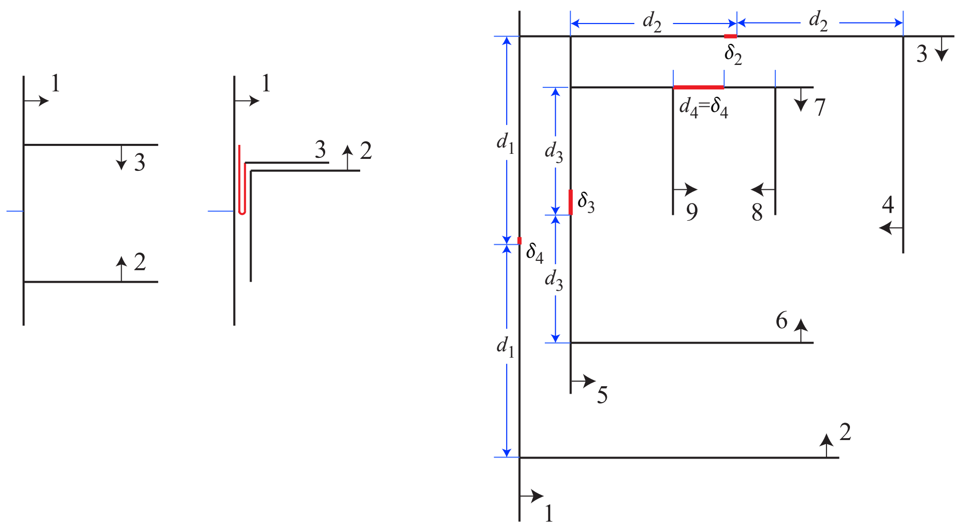

Figure 8. Exponentially many creases for orderly squashing of a 2D positive hyperplane arrangement.

Theorem 6. *Given a convex polyhedron P contained in a bounded subspace U of d-space, there is a continuous folding process that takes U through $(d{+}1)$-space and back to d-space such that the surface of P maps to a $(d-1)$-dimensional hyperplane, and the inside and outside of P map to opposite sides of the hyperplane.*

Proof. We first describe the case where $U = P$, i.e., we are not concerned with the outside of P. Suppose P has k faces $f_1, \ldots, f_k$. We construct a positive hyperplane arrangement of $k{+}1$ hyperplanes in $(d{+}1)$-dimensions. Hyperplane H_{k+1} is the hyperplane containing P, with arbitrary positive side. Hyperplane H_i, $1 \le i \le k$ contains face f_i and is perpendicular to H_{k+1}, and has P in its positive half-space. P is a subset of this positive hyperplane arrangement. Apply orderly squashing. See Figure 9.

P will end up flattened onto H_1, which is a d-dimensional hyperplane. Let B, the *base hyperplane* be the initial H_{k+1}. We will regard the base as fixed. Observe that once a point of P sticks to one of the H_i's, $1 \le i \le k$, its distance from the base remains constant for the remainder of the folding process, because all the H_i's are perpendicular to B. Also observe that the time until a point of P sticks to one of the H_i's is equal to its distance to the boundary of P. Consequently, for any $d \ge 0$, all the points inside P at distance d from the surface of P end up on hyperplane H_1 at distance d from the base. In particular, the surface of P maps to the $(d-1)$-subspace where H_1 intersects B, and the inside of P maps to one side of this subspace.

To handle the general case where $U - P$ is nonempty, simply offset each plane H_i (for $1 \le i \le k$) along its negative normal by a fixed amount s to trace a larger polyhedron P'; choose s large enough so that P' encloses U. Because the surface of P consists exactly of the points at distance s from the boundary of P', thus the result follows from the above argument. □

7. Folding Origami Tessellations

By applying orderly squashing to a shallow, almost planar polyhedron, the result is an almost-planar folding resembling a planar tesselation with crimps along the edges; the intermediate stages of Figure 11 provide a good example. Shallower original polyhedra result in folds that are closer to planar, and in this section we show that an appropriate

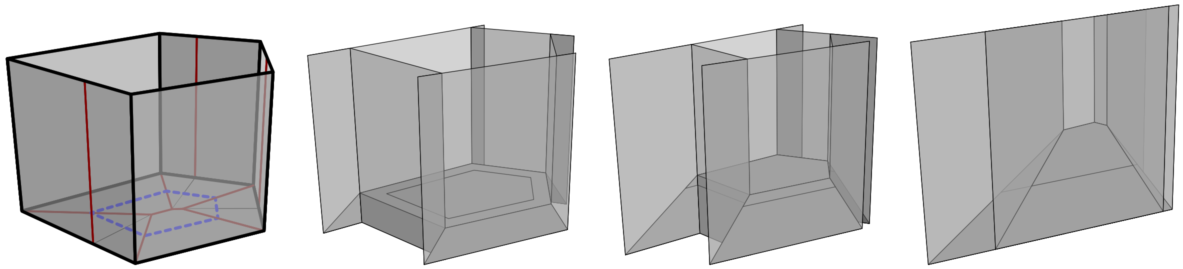

Figure 9. Solving fold and cut in 2D via orderly squashing in 3D: (from left to right) the polygon we want to fold and cut (dashed blue) and the extra planes with the crease pattern; three frames of the flattening.

limiting process, obtained by applying orderly squashing to polyhedra that limit to flat, indeed results in planar tessellation folds, as in Figure 10. This may be accomplished for any planar *spider-web* graph, i.e., any graph that has a convex polygonal lifting.

Our *Spiderweb Folding Algorithm* method proceeds as follows: begin with a plane spiderweb G with a strictly convex polyhedral lifting L, i.e., each vertex $v_i = (x_i, y_i)$ of G is assigned a z-coordinate z_i such that each face of G lifts to a planar polygon in 3D and such that all of G lifts to a convex surface. If we instead use the lifting L^ϵ that uses z-coordinates $\epsilon \cdot z_i$ for some small $\epsilon > 0$, the lifting is still convex but very shallow, and orderly squashing on L^ϵ resembles a folding of the original plane graph G. In the limit as ϵ tends toward 0, we recover the desired folding of G. This limit, however, is subtle: on a shallower lift (ϵ approaching 0), orderly squashing must proceed for a longer time (distance approaching infinity) because the dihedral angle between neighboring planes is very near 180°. We show that orderly squashing L^ϵ for a distance of $1/\epsilon$ properly balances these two effects and results in a well-defined limit as ϵ approaches 0.

There is another, more informal way to interpret this tessellation-folding process. Again begin with a plane graph $G = (V, E)$, and imagine continuously translating each face f_i along velocity vector w_i *in the same plane*, inserting additional crimps and creases near G's edges as necessary to maintain the paper's integrity. Before even considering how these additional creases are chosen, what constraints can be place on the velocities w_i? If faces f_i and f_j share an edge e, then to prevent the paper from shearing or ripping along an edge e, it must be true that faces f_i and f_j move directly toward each other (relatively), that is, the relative velocity $w_j - w_i$ is orthogonal to e in the direction pointing from f_j to f_i. (In this case, their shared crimp will have width $|w_i - w_j|/2$ after one second of motion.) This orthogonality condition means precisely that the vectors w_i define an orthogonal embedding of the dual graph of G, and in particular that G is a spiderweb [29]. So the spiderweb constraint does indeed arise naturally. (Note that the same spiderweb condition shows up in Lang and Bateman's related origami tessellation technique using twist-folds [24].)

By the Maxwell-Cremona correspondence [31], velocities w_i as above have a corresponding convex polygonal lifting described as follows: if $w_i = (a_i, b_i)$, then the lifting of face f_i has normal $(-a_i, -b_i, 1)$. We show that by applying the Spiderweb Folding Algorithm to this lifting, the resulting tessellation indeed has a crimp of width $|w_i - w_j|/2$ along the edge joining faces f_i and f_j, as long as these widths are small enough to avoid more global interaction. For larger widths, the Algorithm still provides a valid flat-folding, but the "fat" crimps now interact in more complicated ways, as shown in Figure 10, right.

8. Conclusion

Orderly squashing provides a surprisingly simple way to continuously flatten any convex polyhedron, or more generally any positive hyperplane arrangement, while respecting the straight-skeleton gluing. While flattening 3D polyhedra is our primary concern, the algorithm works just as well in any dimension. It would be interesting to generalize our computer implementation from 3D to 4D, and watch (in projection) the continuous flattening of, say, a tesseract.[1]

A major open problem is to generalize our result to flatten (even noncontinuously) any *nonconvex* polyhedron while respecting the straight-skeleton gluing. For nonconvex 3D polyhedra, the *straight-skeleton gluing* is defined as follows [15, 16]. From a point p on the surface, shoot a ray into the interior perpendicular to the face. When the ray hits a plane of the straight skeleton, it reflects through the plane. When the ray exits the polyhedron, we pair the exit point with p.

References

[1] P. K. Agarwal, B. Aronov, and M. Sharir. Computing envelopes in four dimensions with applications. *SIAM Journal on Computing*, 26(6):1714–1732, 1997.

[2] A. Aggarwal, L. Guibas, J. Saxe, and P. Shor. A linear-time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete & Computational Geometry*, 4(1):591–604, 1989.

[3] O. Aichholzer and F. Aurenhammer. Straight skeletons for general polygonal figures in the plane. In *Proceedings of Computing and Combinatorics, Second Annual International Conference, (COCOON '96)*, volume 1090 of *Lecture Notes in Computer Science*, pages 117–126. Springer, 1996.

[4] O. Aichholzer, F. Aurenhammer, D. Alberts, and B. Gärtner. A novel type of skeleton for polygons. *Journal of Universal Computer Science*, 1(12):752–761, 1995.

[5] A. D. Alexandrov. *Convex Polyhedra*. Springer, 2005. First published in Russian, 1950.

[1] According to Mrs. Whatsit [25], "there is such a thing as a tesseract."

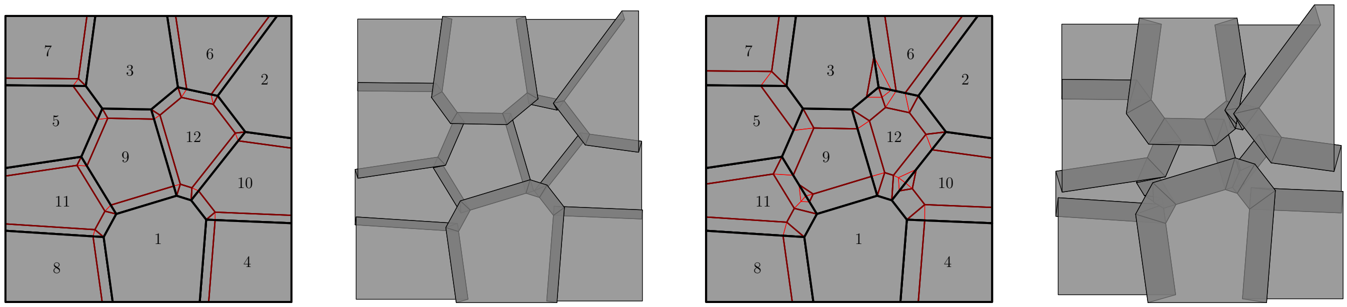

Figure 10. Origami tessellations made from a spiderweb graph through a limiting orderly squash. Small offsets (left) result in elegant tessellations, while large offsets (right) have more interaction.

[6] D. Attali, J.-D. Boissonnat, and H. Edelsbrunner. Stability and computation of medial axes—a state-of-the-art report. In T. Möller, B. Hamann, and R. Russell, editors, *Mathematical foundations of scientific visualization, computer graphics, and massive data exploration*, Mathematics and Visualization, pages 109–125. Springer, 2009.

[7] C. B. Barber, D. P. Dobkin, and H. Huhdanpaa. The Quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software*, 22(4):469–483, 1996.

[8] G. Barequet, D. Eppstein, M. T. Goodrich, and A. Vaxman. Straight skeletons of three-dimensional polyhedra. In *European Symposium on Algorithms (ESA 2008)*, volume 5193 of *Lecture Notes in Computer Science*, pages 148–160. Springer, 2008.

[9] M. Bern and B. Hayes. Origami embedding of piecewise-linear two-manifolds. *Algorithmica*, 59(1):3–15, 2011.

[10] A. I. Bobenko and I. Izmestiev. Alexandrov's theorem, weighted Delaunay triangulations, and mixed volumes. *Annales de l'institut Fourier*, 58(2):447–506, 2008.

[11] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10(1):377–409, 1993.

[12] F. Y. L. Chin, J. Snoeyink, and C. A. Wang. Finding the medial axis of a simple polygon in linear time. *Discrete & Computational Geometry*, 21(3):405–420, 1999.

[13] T. Culver, J. Keyser, and D. Manocha. Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design*, 21(1):65–98, 2004.

[14] E. D. Demaine and M. L. Demaine. Computing extreme origami bases. Technical Report CS-97-22, Department of Computer Science, University of Waterloo, May 1997.

[15] E. D. Demaine, M. L. Demaine, and A. Lubiw. Flattening polyhedra. Unpublished manuscript, 2001.

[16] E. D. Demaine and J. O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, 2007.

[17] H. Edelsbrunner. *Algorithms in Computational Geometry*. Springer New York, 1987.

[18] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1(1):25–44, 1986.

[19] D. Eppstein and J. Erickson. Raising roofs, crashing cycles, and playing pool: Applications of a data structure for finding pairwise interactions. *Discrete & Computational Geometry*, 22(4):569–592, 1999.

[20] M. Held. On computing Voronoi diagrams of convex polyhedra by means of wavefront propagation. In *Sixth Canadian Conference on Computational Geometry*, pages 128–133, 1994.

[21] S. Huber and M. Held. A fast straight-skeleton algorithm based on generalized motorcycle graphs. *International Journal of Computational Geometry & Applications*, 22(05):471–498, 2012.

[22] J.-i. Itoh, C. Nara, and C. Vîlcu. Continuous flattening of convex polyhedra. In *Computational Geometry*, volume 7579 of *Lecture Notes in Computer Science*, pages 85–97. Springer, 2012.

[23] D. Kane, G. N. Price, and E. D. Demaine. A pseudopolynomial algorithm for Alexandrov's Theorem. In *Proceedings of the 11th Algorithms and Data Structures Symposium (WADS 2009)*, volume 5664 of *Lecture Notes in Computer Science*, pages 435–446, 2009.

[24] R. J. Lang and A. Bateman. Every spider web has a simple flat twist tessellation. In *Origami^5: Proceedings of the 5th International Conference on Origami in Science, Mathematics and Education*, pages 449–454. A K Peters, Singapore, July 2010.

[25] M. L'Engle. *A Wrinkle in Time*. Farrar, Straus, and Giroux, 1962. Also a TV movie.

[26] P. Mitra and L. Hafer. Efficient computation of the medial axis. Technical Report CMPT TR 98-07, Simon Fraser University, 1998.

[27] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete & Computational Geometry*, 6(1):423–434, 1991.

[28] R. Seidel. Convex hull computations. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry, 2nd edition*, pages 495–512. CRC Press, Inc., 2004.

[29] B. Servatius and W. Whiteley. Constraining plane configurations in computer-aided design: Combinatorics of directions and lengths. *SIAM Journal on Discrete Mathematics*, 12(1):136–153, 1999.

[30] M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete & Computational Geometry*, 12(1):327–345, 1994.

[31] W. Whiteley. Motions and stresses of projected polyhedra. *Structural Topology*, 7:13–38, 1982.

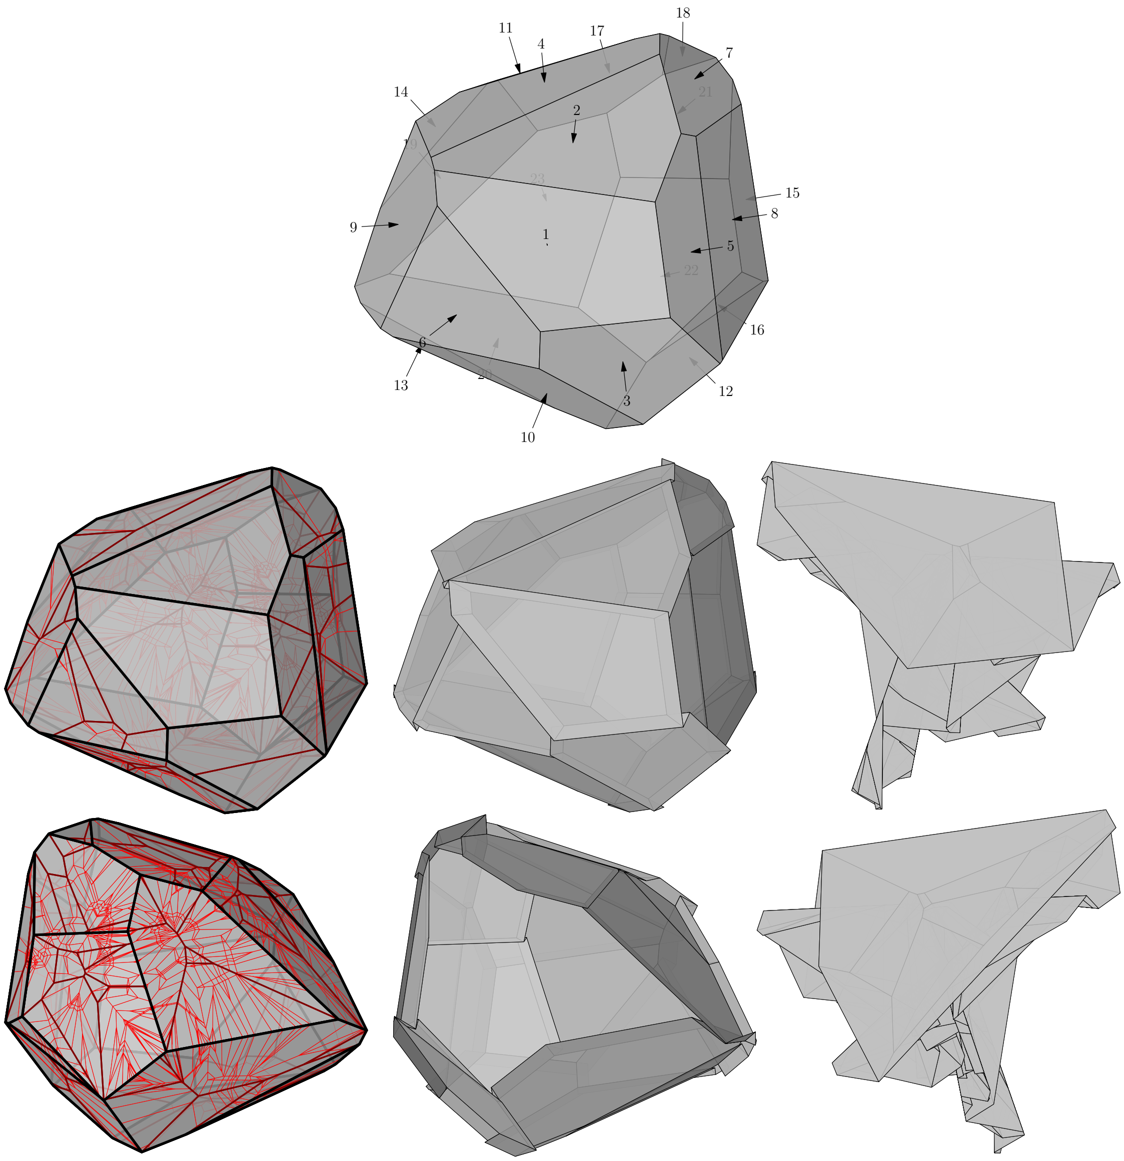

Figure 11. More details of the orderly squashing of the polyhedron from Figure 3: (top) the ordering of front faces; (middle row) front view of the creases and two stages of the flattening; (bottom row) back view of the creases and two stages of the flattening. Primary creases are shown in thick dark red and other creases in thin light red.