Reshaping Convex Polyhedra

Joseph O'Rourke and Costin Vîlcu

July 8, 2021

Contents

Abstract **vii**

Preface **ix**

I Tailoring for Every Body **1**

1 Introduction to Part I **3**
1.1 Alexandrov's Gluing Theorem 4
1.2 Tailoring Examples 5
1.3 Summary of Part-I Results 8

2 Preliminaries **11**
2.1 Cut locus properties 11
2.1.1 Star-Unfolding and Cut Locus 12
2.1.2 Fundamental Triangles 13
2.1.3 Cut Locus Partition 15
2.2 Cauchy's Arm Lemma 15
2.3 A Rigidity Result 18
2.4 Vertex-Merging 20

3 Domes and Pyramids **23**
3.1 Domes 23
3.2 Cube/Tetrahedron Example 24
3.2.1 Slice → G-Domes for Cube/Tetrahedron 25
3.3 Proof: G-dome → Pyramids 25
3.3.1 G-domes → Pyramids for Cube/Tetrahedron 29

4 Tailoring via Sculpting **33**
4.1 Slice → G-domes . . . 33
4.2 Pyramid → Tailoring . . . 35
4.2.1 Notation . . . 35
4.3 Small volume slices . . . 36
4.3.1 Pyramid case . . . 38
4.3.2 General case . . . 43
4.4 Cube/Tetrahedron: Completion . . . 46
4.4.1 Pyramid Removals . . . 46
4.4.2 Pyramid Reductions by Tailoring . . . 46
4.4.3 Seals . . . 49
4.5 Hexagonal Pyramid Example . . . 50
4.6 Tailoring is finer than sculpting . . . 51

5 Pyramid Seal Graph **55**
5.1 Pyramid Digon Removal . . . 55
5.1.1 Notation I . . . 56
5.2 Cone Viewpoint . . . 56
5.2.1 Notation II . . . 56
5.3 Examples . . . 57
5.4 Preliminary Lemmas . . . 60
5.5 Pyramid Seal Graph is a Tree . . . 62
5.5.1 Other Digon Orderings . . . 66

6 Algorithm for Tailoring via Sculpting **71**
6.1 Algorithm 1: slice → g-domes . . . 72
6.1.1 Complexity of sculpting . . . 73
6.2 Algorithm 2: g-dome → pyramids . . . 75
6.3 Algorithm 3: pyramid → digons . . . 76
6.4 Overall Tailoring Algorithm . . . 77

7 Crests **79**
7.1 Examples . . . 79
7.2 Proofs . . . 81
7.3 Algorithm 4: pyramid → crest . . . 89

8 Tailoring via Flattening **91**
8.1 Proofs . . . 92
8.1.1 Digon-tailor $P \to P_{\text{flat}}$ . . . 92
8.1.2 Vertex-merge $Q \to Q_{\text{flat}}$ . . . 92
8.1.3 Scale $Q_{\text{flat}} \to Q^s_{\text{flat}}$ . . . 93
8.1.4 Trim $P_{\text{flat}} \to P^s_{\text{flat}}$ . . . 95
8.1.5 Reverse $P^s_{\text{flat}} \to Q^t$ . . . 95
8.1.6 Theorem: Tailoring via Flattening . . . 96
8.2 Algorithm for Tailoring via Flattening . . . 98

9 Enlarging and P-Unfoldings **101**
9.1 Enlarging and Reshaping . . . 101
9.2 P-unfoldings . . . 102
9.2.1 P-unfoldings and Reshaping . . . 102
9.2.2 P-unfoldings and the WBG Theorem . . . 106

II Vertex Merging and Convexity **109**

10 Introduction to Part II **111**

11 Vertex-Merging Reductions and Slit Graphs **113**
11.1 Slit Graphs for Vertex Mergings . . . 113
11.2 Example: Reductions of Flat Hexagon . . . 116
11.3 Example: Reductions of Cube . . . 118
11.4 Example: Icosahedron . . . 119
11.5 Example: Hexagonal Shape with Cycle . . . 124
11.6 Vertex Merging and Unfoldings . . . 125
11.7 Unfolding Irreducible Surfaces . . . 127
11.7.1 S: Doubly-covered Triangle . . . 128
11.7.2 Net and Overlap . . . 129
11.7.3 S: Isosceles Tetrahedron . . . 130

12 Planar Spiral Slit Tree **131**
12.1 Sequential Spiral Merge . . . 131
12.2 Notation . . . 132
12.3 Algorithm Description . . . 133
12.4 Planar Proof . . . 134

13 Convexity on Convex Polyhedra **139**
13.1 Convex Curves . . . 140
13.2 Notions of Convexity . . . 140
13.3 Ag-convexity . . . 141
13.4 Geodesic Segments and Convex Sets . . . 146
13.5 Relative Convexity . . . 148
13.6 Convex Hull . . . 150
13.7 Relative Convex Hull . . . 154
13.8 Extreme Points . . . 156
13.9 Relative Convex Hull of Vertices . . . 160
13.10Summary of Properties . . . 165

14 Minimal-length Enclosing Polygon **169**
14.1 Properties of the Minimal Enclosing Polygon . . . 169
14.2 Shortening Algorithm . . . 174
14.2.1 Curve-Shortening Flow . . . 175
14.2.2 Algorithm Overview . . . 175
14.2.3 Finding an Enclosing Geodesic Polygon . . . 176
14.2.4 Algorithm for Curve Shortening . . . 179
14.3 AG-convexity and Z . . . 184
14.4 Algorithm for $\mathrm{rconv}(V) = R(W)$ . . . 185

15 Spiral Tree on Polyhedron **187**
15.1 Notation . . . 187
15.2 Icosahedron Example . . . 188
15.3 Spiraling Algorithm for rconv . . . 191
15.4 Proof: Slit Graph is a Tree . . . 195
15.5 Spiraling Algorithm for $Z(V) = \min \ell[V]$ . . . 198

16 Unfoldings via Slit Trees **201**
16.1 Notation . . . 201
16.2 Unfoldings via spiraling algorithms . . . 202
16.2.1 Two Cones . . . 203
16.2.2 Reduction to Cylinder . . . 207
16.2.3 Cube Example . . . 209
16.2.4 Icosahedron Example . . . 211

17 Vertices on Quasigeodesics **215**
17.1 Notation 215
17.2 $\mathcal{Q}_k$ Theorem 216
17.2.1 $|V(Q)| = 1$ 217
17.2.2 $|V(Q)| = 2$ 220
17.2.3 $|V(Q)| = k$, with $3 \leq k \leq n$ 220

18 Conclusions and Open Problems **223**
18.1 Part I 223
18.2 Part II 226

Bibliography **229**

Abstract

Given a convex polyhedral surface P, we define a *tailoring* as excising from P a simple polygonal domain that contains one vertex v, and whose boundary can be sutured closed to a new convex polyhedron via Alexandrov's Gluing Theorem. In particular, a *digon-tailoring* cuts off from P a *digon* containing v, a subset of P bounded by two equal-length geodesic segments that share endpoints, and can then zip closed.

In the first part of this monograph, we primarily study properties of the tailoring operation on convex polyhedra. We show that P can be reshaped to any polyhedral convex surface $Q \subset \text{conv}(P)$ by a sequence of tailorings. This investigation uncovered previously unexplored topics, including a notion of *unfolding of Q onto P*—cutting up Q into pieces pasted non-overlapping onto P.

In the second part of this monograph, we study *vertex-merging* processes on convex polyhedra (each vertex-merge being in a sense the reverse of a digon-tailoring), creating embeddings of P into enlarged surfaces. We aim to produce non-overlapping polyhedral and planar unfoldings, which led us to develop an apparently new theory of convex sets, and of minimal length enclosing polygons, on convex polyhedra.

All our theorem proofs are constructive, implying polynomial-time algorithms.

MSC Classifications

Primary: 52A15, 52B10, 52C45, 53C45, 68U05.

Secondary: 52-02, 52-08, 52A37, 52C99.

Keywords and phrases

convex polyhedron
Alexandrov's Gluing Theorem
polyhedron truncation
cut locus
unfolding polyhedra
net for polyhedron; star-unfolding
vertex-merging
convex sets and convex hulls on convex polyhedra
minimal length enclosing polygon on convex polyhedra

Preface

The research reported in this monograph emerged from exploring a simple question: Given two convex polyhedra P and Q, with Q inside P, can one reshape P to Q by repeatedly "snipping" off vertices? We call this snipping-off operation *tailoring*. A precise definition is deferred to the introductory Chapter 1, but here we contrast it with *vertex truncation*, which slices off a vertex with a plane and replaces it with a new facet lying in that plane. This is, for example, one way to construct the truncated cube: see Fig. 1. Tailoring differs from vertex truncation in two ways: first, it does not slice

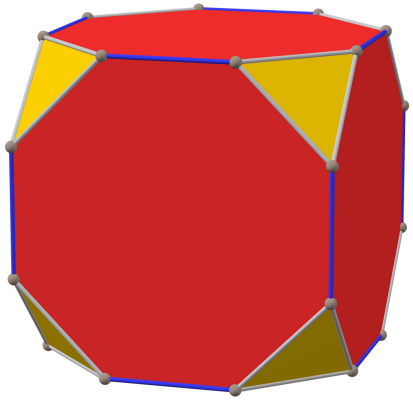

Figure 1: Truncated cube. [Image by Tilman Piesk, Wikipedia].

by a plane but instead uses a digon bound by a pair of equal-length geodesics (again, definition deferred), and second, rather than filling the hole with a new facet, the boundary of the hole is "sutured" closed without the addition of new surface.

Our first experiment started with a paper cube and tailored its 8 vertices, producing the shape shown in Fig. 2. Although not evident from this crude

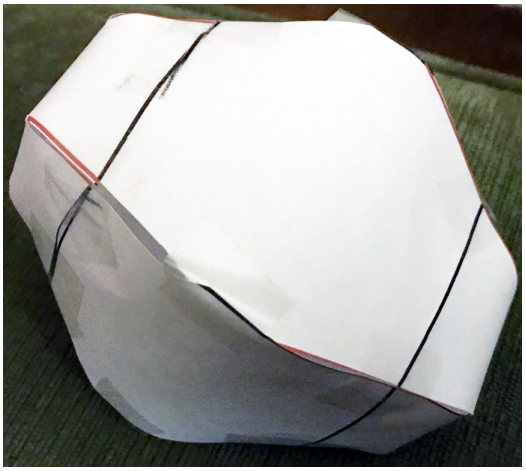

Figure 2: Cube after tailoring 8 vertices.

model, this shape is actually a convex polyhedron of 16 vertices, an implication of Alexandrov's Gluing Theorem. This led us imagine that continuing the process on the 16 vertices might allow reshaping the cube into a roughly spherical polyhedron—"whittling" a cube to a sphere. And indeed, one of our main theorems is that any P can be reshaped to any $Q \subset P$ by a finite sequence of tailorings (Theorem 4.6). This holds whether Q has fewer vertices than P or more vertices than P: see Fig. 3.

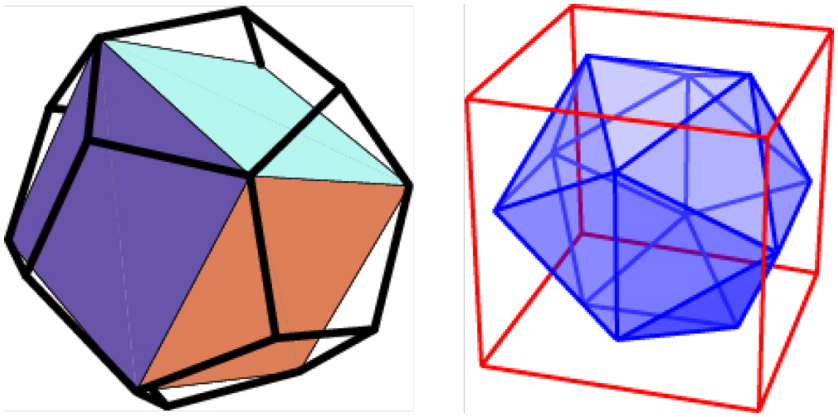

Figure 3: Cube inside dodecahedron. Icosahedron inside cube.

Investigating tailoring in Part I uncovered previously unexplored topics,

including a notion of "unfolding" Q onto P—cutting up Q into pieces pasted non-overlapping onto P. This led us to a systematic study of *vertex-merging*, a technique introduced by Alexandrov that can be viewed as a type of inverse operation to tailoring, a study we carry out in Part II. Here we start with P and gradually enlarge it via vertex-merges until P is embedded onto a nearly flat shape S: a doubly-covered triangle, an isosceles tetrahedron, a pair of cones, or a cylinder. The first two "vertex-merge irreducible shapes" (doubly-covered triangle or isosceles tetrahedron) are the result of vertex-merging in arbitrary order, and can result in disconnecting an n-vertex P into $O(n^2)$ pieces (Corollary 11.2). In contrast we pursue the goal of minimizing the disconnection of P, with the ultimate goal of achieving a planar *net* for P—a nonoverlapping simple polygon.

Toward this end, we partition P into two half-surfaces by a quasigeodesic Q, a simple closed curve on P, convex to both sides.[1] Then we can prove that a particular spiral vertex-merge sequence avoids disconnecting P in each half. (Theorems 15.7 and 15.14.) This approach requires a notion of convexity and convex hull on the surface of a convex polyhedron, apparently novel topics, explored in our longest chapter, Chapter 13.

We can achieve a net for P (Theorem 16.5) assuming the truth of a conjecture concerning quasigeodesics (Open Problem 18.13). This is among the many avenues for future work uncovered by our investigations.

Acknowledgements

We benefitted from discussions with Jin-ichi Itoh and with Anna Lubiw, on some topics of this work.

Costin Vîlcu's research was partially supported by UEFISCDI, project no. PN-III-P4-ID-PCE-2020-0794.

[1] In Part I, Q is the target polyhedron resulting from reshaping P. In Part II, there is no equivalent fixed "target" polyhedron, and we use Q to denote a quasigeodesic.

Part I

Tailoring for Every Body

Chapter 1

Introduction to Part I

Let P and Q be convex polyhedra, each the convex hull of finitely many points in $\mathbb{R}^3$. If $Q \subset P$, it is easy to see that Q can be *sculpted* from P by "slicing P with planes." By this we mean intersecting P with half-spaces each of whose plane boundary contains a face of Q. If $Q \not\subset P$, we can shrink Q until it fits inside. So a *homothet* of any given Q can be sculpted from any given P, where a homothet is a copy possibly scaled, rotated, and translated. Main results of Part I are similar claims but via "tailorings": a homothet of any given Q can be tailored from any given P (Theorems 4.6, 7.5, 8.1).[1]

With some abuse of notation,[2] we will use the same symbol P for a polyhedral hull and its boundary. We define two types of tailoring. A *digon-tailoring* cuts off a single vertex of P along a digon, and then sutures the digon closed. A *digon* is a subset $D(x, y)$ of P bounded by two equal-length geodesic segments that share endpoints x and y. A *geodesic segment* is a shortest geodesic between its endpoints. A *crest-tailoring* cuts off a single vertex of P but via a more complicated shape we call a "crest." Again the hole is sutured closed. We defer discussion of crests to Chapter 7. Meanwhile, we often shorten "digon-tailoring" to simply *tailoring*. Cutting out a digon means excising the portion of the surface between the geodesics, including the vertex they surround.[3] Once removed, the digon hole is closed by naturally identifying the two geodesics along their lengths. This identification is often called "gluing" in the literature, although we also call it "zipping" or

[1] A preliminary version of Part I is [OV20].

[2] Justified by Alexandrov's Gluing Theorem, see ahead.

[3] An informal view (due to Anna Lubiw) is that one could pinch the surface flat in a neighborhood of the vertex, and then snip-off the flattened vertex with scissors.

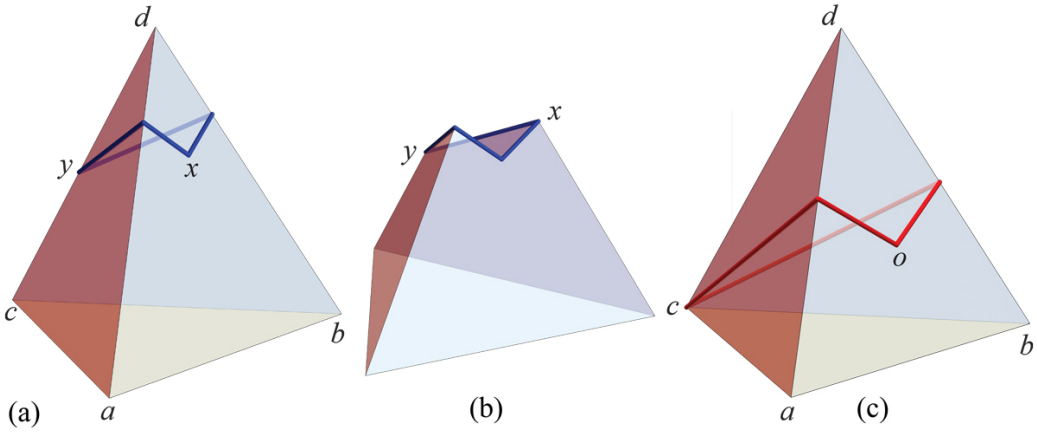

Figure 1.1: (a) A digon $D(x, y)$ on the regular tetrahedron $P = abcd$, surrounding vertex d. (b) Digon excised. (c) A digon $D(o, c)$ with one endpoint at vertex c, the other the centroid o.

"suturing" or "sealing." Fig. 1.1 illustrates digons on a regular tetrahedron P. In (a) of the figure, neither digon endpoint x nor y is a vertex of P; (b) shows P after excision but before sealing. After sealing, both x and y become vertices. (See ahead to Fig. 2.6 for a depiction of the polyhedron that results after sealing.) In (c), one endpoint of the digon coincides with a vertex. Here, after excision and suturing, o becomes a vertex and c remains a vertex (of sharper curvature).

1.1 Alexandrov's Gluing Theorem

Throughout, we make extensive use of Alexandrov's Gluing Theorem [Ale05, p.100], which guarantees that the surface obtained after a tailoring of P corresponds uniquely to a convex polyhedron P'. A precise statement of this theorem, which we will abbreviate to AGT, is as follows.

Theorem AGT. *Let S be a topological sphere obtained by gluing planar polygons (i.e., naturally identifying pairs of sides of the same length) such that at most 2π surface angle is glued at any point. Then S, endowed with the intrinsic metric induced by the distance in $\mathbb{R}^2$, is isometric to a convex polyhedron $P \subset \mathbb{R}^3$, possibly degenerated to a doubly-covered convex polygon. Moreover, P is unique up to rigid motion and reflection in $\mathbb{R}^3$.*

The case of doubly-covered convex polygon is important and we will encountered it often.

Because the sides of the digon are geodesics, gluing them together to seal the hole leaves 2π angle at all but the digon endpoints. The endpoints lose surface angle with the excision, and so have strictly less than 2π angle surrounding them. So AGT applies and yields a new convex polyhedron.

This shows that tailoring is possible and alters the given P to another convex polyhedron. How to "aim" the tailoring toward a given target Q is a long story, told in subsequent sections.

Alexandrov's proof of his celebrated theorem is a difficult existence proof and gives little hint of the structure of the polyhedron guaranteed by the theorem. And as-yet there is no effective procedure to construct the three-dimensional shape of the polyhedron guaranteed by his theorem. It has only been established that there is a theoretical pseudopolynomial-time algorithm [KPD09], achieved via an approximate numerical solution to the differential equation established by Bobenko and Izmestiev [BI08] [O'R07]. But this remains impractical in practice. Only small or highly symmetric examples can be reconstructed, for example [ADO03] for the foldings of a square, and more recently [ALZ20] for polyhedra built from regular pentagons. Figs. 1.3 and 2.6 (ahead) were reconstructed through ad hoc methods.

AGT is a fundamental tool in the geometry of convex surfaces and, at a theoretical level, our work helps to elucidate its implications. While AGT has proved useful in several investigations, our application here to reshaping has, to our knowledge, not been considered before as the central object of study.

1.2 Tailoring Examples

Before discussing background context, we present several examples. Throughout we let xy denote the line segment between points x and y, $x, y \in \mathbb{R}^3$. Also we make extensive use of vertex curvature. The *discrete* (or *singular*) *curvature* $\omega(v)$ at a vertex $v \in P$ is the angle deficit: 2π minus the sum of the face angles incident to v.

Example 1.1. *Let $T = abcd$ be a regular tetrahedron, and let o be the center of the face abd; see Fig. 1.1(c). Cut out the digon on T between c and o "encircling" d, and zip it closed.*

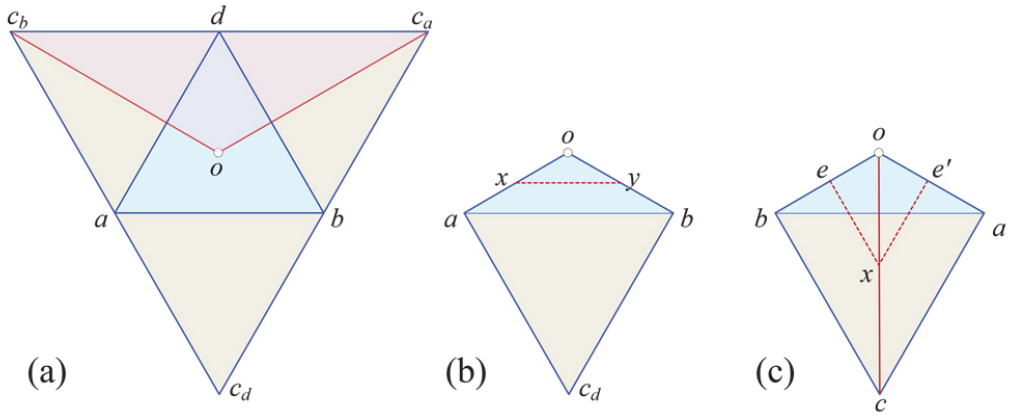

Figure 1.2: Illustrations for Examples 1.1–1.2. In (c), e and e' are joined when the $D(x, y)$ digon is excised.

The unfolding $U(T)$ of T with respect to c is a planar regular triangle $c_a c_b c_d$ with center o. Cutting out that digon from T is equivalent to removing from T the isosceles triangle $oc_a c_b$. See Fig. 1.2(a). We zip it closed by identifying the digon-segments $c_a o$ and $c_b o$, and refolding the remainder of T by re-identifying $c_a b$ and bc_d, and $c_b a$ and ac_d. The result is the doubly-covered kite $K = aobc_d$, shown in Fig. 1.2(b).

Example 1.2. *We further tailor the doubly-covered kite K obtained in Example 1.1. Excising a digon encircling o, between points $x \in oa$ and $y \in ob$, and zipping closed, yields a doubly-covered pentagon, as in Fig. 1.2 (b).*

If instead we excise a digon encircling o, between corresponding points $x, y \in oc$ on different sides of K (see Fig. 1.2 (c)), and seal closed, the result is a non-degenerate pentahedron, illustrated in Fig. 1.3. This is among the rare cases where the polyhedron guaranteed by AGT can be explicitly reconstructed.

Example 1.3. *A digon may well contain several vertices, but for our digon-tailoring we only consider those containing at most one vertex. The limit case of a digon containing no vertex is an edge between two vertices. (This will play a role in Chapter 8.) In this case, gluing back along the cut would produce the original polyhedron, but we can as well seal the cut closed from another starting point. For example, cutting along an edge of an isosceles*

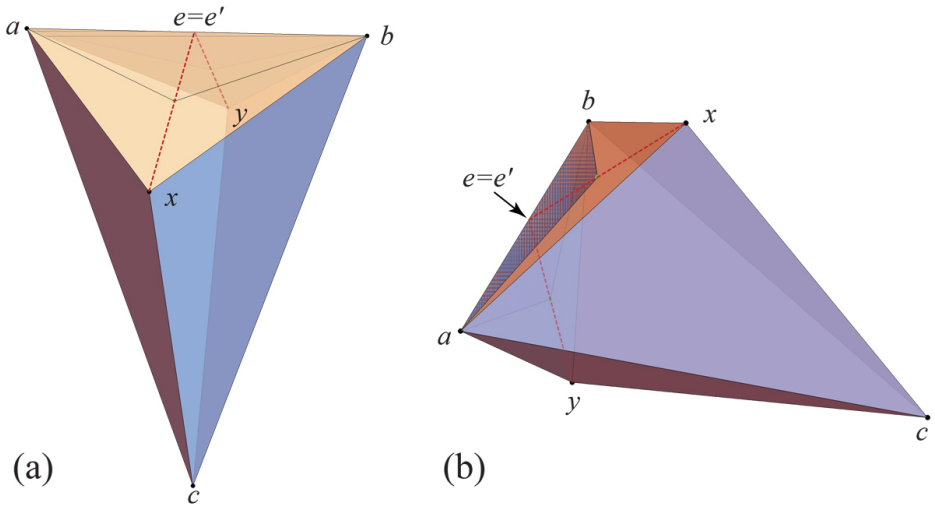

Figure 1.3: The non-degenerate pentahedron obtained in Example 1.2. The point marked $e = e'$ is the join of the points e and e' in Fig. 1.2(c), the midpoint of the edge ab.

tetrahedron (a polyhedron composed of four congruent faces)[4] *and carefully choosing the gluing leads to a doubly-covered rectangle. See Fig. 1.4.*

By AGT, this limit-case tailoring of an empty-interior digon only applies between vertices v_1, v_2 of curvatures $\omega_1, \omega_2 \geq \pi$, because both v_1 and v_2 will be identified with interior points of the edge.[5]

In view of Fig. 1.2(c) and Fig. 1.3, it is clear that, even though tailoring is area decreasing, it is not necessarily volume decreasing.

Let a tailoring step remove vertex v inside digon $D = (x, y)$. In general, neither x nor y is a vertex before tailoring, but they become vertices after removing D, thereby increasing the number of vertices of P by 1. This is illustrated in Fig. 1.1(ab). If one of x or y is a vertex and the other not, then the total number of vertices remains fixed, as in Fig. 1.1(c). And if both x and y are vertices, then the number of vertices of P is decreased by 1. The challenge answered in our work is to direct tailoring to "aim" from one polyhedon P to the target Q, which may have a quite different number of vertices.

[4] Also called a tetramonohedron, or an isotetrahedron.

[5] In [DO07, Sec. 25.3.1] this structure is called a "rolling belt."

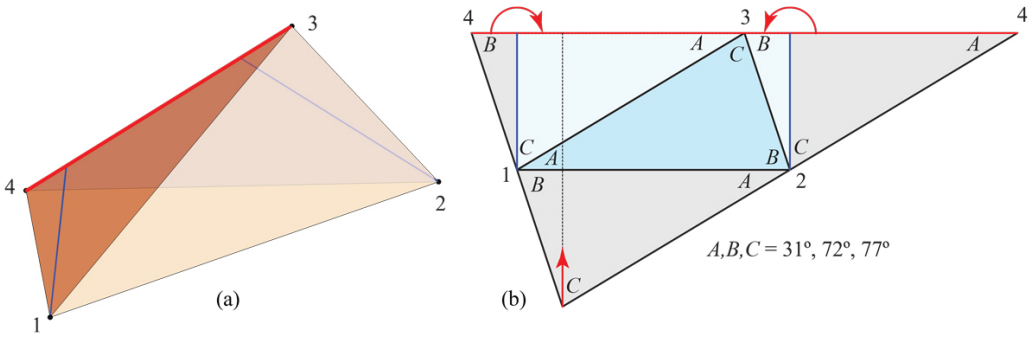

Figure 1.4: (a) An isosceles tetrahedron: All four faces are congruent. Cutting along v_3v_4 and regluing the two halves of that slit differently, creasing at the blue segments, yields a doubly-covered rectangle, as shown in (b), folding as shown.

1.3 Summary of Part-I Results

Here we list our main theorems in Part I, each with a succinct (and at this stage, quite approximate) summary of their claims.

- Theorem 4.6: Q may be digon-tailored from P, tracking a sculpting of P to Q.
- Theorem 7.5: P may be crest-tailored to Q, again tracking a sculpting.
- Theorem 8.1: A different proof of a similar result, that P may be digon-tailored to a homothet of Q, but this time without relying on sculpting.
- Theorems 6.2 and 7.5 and 8.3: Tailoring algorithms have time-complexity $O(n^4)$.
- Theorem 9.1: Reversing tailoring yields procedures for enlarging Q inside P to match P. As a consequence, Q may be cut up and "unfolded" isometrically onto P.

Along the way to our central theorems, we obtain results not directly related to AGT:

- Theorem 2.9: If two convex polyhedra with the same number of vertices match on all but the neighborhoods of one vertex, then they are congruent.

- Theorem 3.2: Every "g-dome" can be partitioned into a finite sequence of pyramids by planes through its base edges.

The above results raise several open problems of various natures, either scattered along the text or presented in the last section of Part I. [Revisit later.]

Finally, we sketch the logic behind the first result listed above, whose statement in Chapter 4 we repeat here:

Theorem 4.6. *Let P be a convex polyhedron, and $Q \subset P$ a convex polyhedron resulting from repeated slicing of P with planes. Then Q can also be obtained from P by tailoring. Consequently, for any given convex polyhedra P and Q, one can tailor P "via sculpting" to obtain any homothetic copy of Q inside P.*

Start with Q inside P, and imagine a sequence of slices by planes that sculpt P to Q. Lemma 4.4 shows how to digon-tailor one such slice, which then establishes the claim that we can tailor P to Q. Theorem 4.6 is achieved by first slicing off shapes we call "g-domes," and then showing in Theorem 3.2 that every g-dome can be reduced to its base by slicing off pyramids, i.e., by vertex truncations. Lemma 4.4 shows that such vertex truncations can be achieved by tailoring. And the proof of Lemma 4.4 relies on the rigidity established by Theorem 2.9. So the path of logic is:

$$\text{plane slice} \;\to\; \text{g-domes} \;\to\; \text{pyramids} \;\to\; \text{digon removals}\,. \qquad \begin{array}{c} \uparrow \\ \text{rigidity} \end{array}$$

Chapter 2

Preliminaries

In this chapter we present basic properties of cut loci on convex polyhedra, the star-unfolding, prove a rigidity result, and describe the technique of vertex-merging. All of these geometric tools will be needed subsequently. The reader might skip this section and return to it as the tools are deployed.

2.1 Cut locus properties

The *cut locus* $C(x)$ *of the point* x on a convex polyhedron P is the closure of the set of points to which there is more than one shortest path from x. This concept goes back to Poincaré [Poi05], and has been studied algorithmically since [SS86] (there, the cut locus is called the "ridge tree"). The cut locus is one of our main tools throughout this work. The next lemma establishes notation and lists several known properties.

Lemma 2.1 (Cut Locus Basics)**.** *The following hold for the cut locus* $C(x)$*:*

(i) $C(x)$ *has the structure of a finite* 1*-dimensional simplicial complex which is a tree. Its leaves (endpoints) are vertices of* P*, and all vertices of* P*, excepting* x *(if it is a vertex) are included in* $C(x)$*. All points interior to* $C(x)$ *of tree-degree* 3 *or more are known as* ramification points *of* $C(x)$*.*[1] *All vertices of* P *interior to* $C(x)$ *are also considered as ramification points.*

[1] In some literature, these points are called "branch points" or "junctions" of $C(x)$.

(ii) Each point y in $C(x)$ is joined to x by as many geodesic segments as the number of connected components of $C(x) \setminus y$. For ramification points in $C(x)$, this is precisely their degree in the tree.

(iii) The edges of $C(x)$ are geodesic segments on P.

(iv) Assume the geodesic segments γ and γ' (possibly $\gamma = \gamma'$) from x to $y \in C(x)$ are bounding a domain D of P, which intersects no other geodesic segment from x to y. Then there is an arc of $C(x)$ at y which intersects D and it bisects the angle of D at y.

Proof. The statements (i)-(ii) and (iv) are well known. The statement (iii) is Lemma 2.4 in [AAOS97]. □

The following is Lemma 4 in [INV12].

Lemma 2.2 (Path Cut Locus)**.** *If $C(x)$ is a path, the polyhedron is a doubly-covered (flat) convex polygon, with x on the rim.*

The following lemma will be invoked in Chapter 5.

Lemma 2.3 (Angle $< \pi$)**.** *Let x be a point on a convex polyhedron P and let y be a ramification point of $C(x)$ of degree $k \geq 2$. Let $e_1, ..., e_k$ be the edges of $C(x)$ incident to y, ordered counterclockwise, and let α_j be the angle of the sector between e_j and e_{j+1} at y (with $k+1 \equiv 1$). Then $\alpha_j < \pi$, for all $j = 1, ..., k$.*

Proof. Let $\gamma_1, ...\gamma_k$ be the geodesic segments from x to y, say with γ_j between e_j and e_{j+1} (again with $k+1 \equiv 1$). See Fig. 2.1. By Lemma 2.1 (iv), e_j is the bisector of $\angle(\gamma_j, \gamma_{j+1})$. Put $2\beta_{j+1} = \angle(\gamma_j, \gamma_{j+1})$. Then the total surface angle θ_y incident to y satisfies $\theta_y = \sum_{j=1}^{k} 2\beta_j \leq 2\pi$ if $k \geq 3$, and $\theta_y < 2\pi$ if $k = 2$. Therefore, when $k \geq 3$, $\sum_{j=1}^{k} \beta_j = \pi$ and so $\alpha_j = \beta_j + \beta_{j+1} < \pi$. And when $k = 2$, $\alpha_1 = \alpha_2 = \beta_1 + \beta_2 < \pi$. □

2.1.1 Star-Unfolding and Cut Locus

Next we introduce a general method for unfolding any convex polyhedron P to a simple (non-overlapping) polygon in the plane. We use this subsequently largely because of its connection to the cut locus.

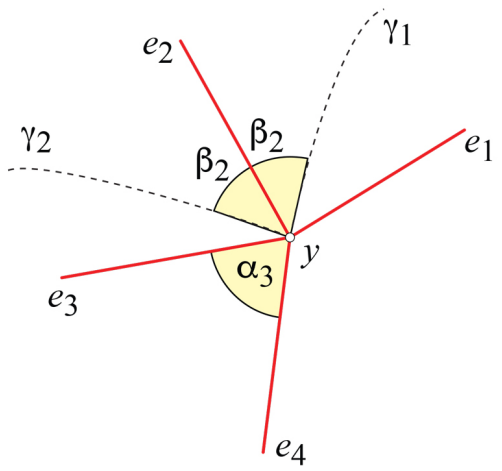

Figure 2.1: Edge e_2 of $C(x)$ bisects $\angle(\gamma_1, \gamma_2) = 2\beta_2$.

To form the *star-unfolding* of a P with respect to x, one cuts P along the geodesic segments (unique if x is "generic") from x to every vertex of P. The idea goes back to Alexandrov [Ale05]; the non-overlapping of the unfolding was established in [AO92], where the next result was also proved. See Fig. 2.2.

Lemma 2.4 (S_P Voronoi Diagram)**.** *Let $S_P = S_P(x)$ denote the star-unfolding of P with respect to $x \in P$. Then the image of $C(x)$ in S_P is the restriction to S_P of the Voronoi diagram of the images of x.*

Notice that several properties of cut loci, presented in the previous section, could easily be derived from Lemma 2.4.

2.1.2 Fundamental Triangles

A geodesic triangle on P (i.e., with geodesic segments as sides) is called *flat* if its curvature vanishes.

Lemma 2.5 (Fundamental Triangles [INV12])**.** *For any point $x \in P$, P can be partitioned into flat triangles whose bases are edges of $C(x)$, and whose lateral edges are geodesic segments from x to the ramification points or leaves of $C(x)$. Moreover, those triangles are isometric to plane triangles, congruent by pairs.*

See Fig. 2.2(b) (and ahead to Fig. 8.1).

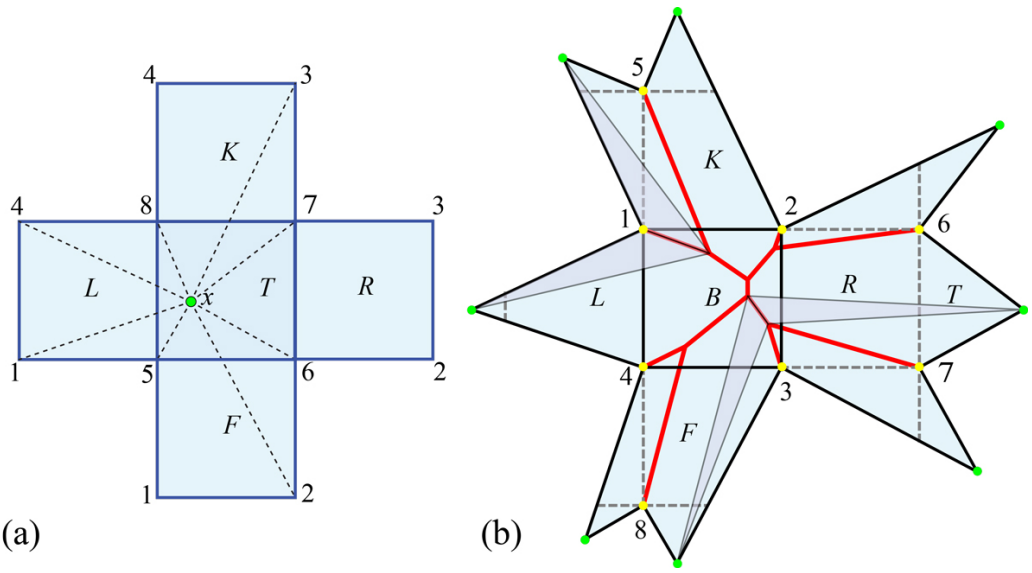

Figure 2.2: (a) Cut segments to the 8 vertices of a cube from a point x on the top face. (No cut is interior to the bottom face.) T, F, R, K, L, B = Top, Front, Right, Back, Left, Bottom. (b) The star-unfolding from x. The cut locus $C(x)$ (red) is the Voronoi diagram of the 8 images of x (green). Two pairs of fundamental triangles are shaded.

2.1.3 Cut Locus Partition

Another tool we need (in Chapter 4) is a cut locus partition lemma, a generalization of lemmas in [INV12]. On a polyhedron P, connect a point x to a point $y \in C(x)$ by two geodesic segments γ, γ'. This partitions P into two "half-surface" digons H_1 and H_2. If we now zip each digon separately closed by joining γ and γ', AGT leads to two convex polyhedra P_1 and P_2. The lemma says that the cut locus on P is the "join" of the cut loci on P_i. See Fig. 2.3.

Lemma 2.6 (Cut Locus Partition)**.** *Under the above circumstances, the cut locus $C(x, P)$ of x on P is the* join *of the cut loci on P_i: $C(x, P) = C(x, P_1) \sqcup_y C(x, P_2)$, where $\sqcup_y$ joins the two cut loci at point y. And starting instead from P_1 and P_2, the natural converse holds as well.*

Proof. Notice first that a straightforward induction and Lemma 2.5 on fundamental triangles shows that the cut locus of x on P_i is indeed the truncation of $C(x, P)$. Therefore, $C(x, P) = C(x, P_1) \sqcup_y C(x, P_2)$.

Assume we start now from P_1 and P_2, having vertices $x_1, y_1 \in P_1$ and $x_2, y_2 \in P_2$ such that

- $\rho_{P_1}(x_1, y_1) = \rho_{P_2}(x_2, y_2)$, where $\rho_{P_i}(\ )$ is the geodesic distance between the indicated points on P_i.
- $\theta_{x_1} + \theta_{x_2} \leq 2\pi$, where θ_x is the total surface angle incident to x, and
- $\theta_{y_1} + \theta_{y_2} \leq 2\pi$.

Then we can cut open P_i along a geodesic segment γ_i from x_i to y_i, $i = 1, 2$, and join the the two halves by AGT, such that x_1, x_2 have a common image x, and y_1, y_2 have a common image y.

Now, all geodesic segments starting at x into H_i remain in H_i, because geodesic segments do not branch. Therefore, H_1 has no influence on $C(x, P_2)$ and H_2 has no influence on $C(x, P_1)$. □

2.2 Cauchy's Arm Lemma

In several proofs, we will need an extension of Cauchy's Arm Lemma, which we now describe.

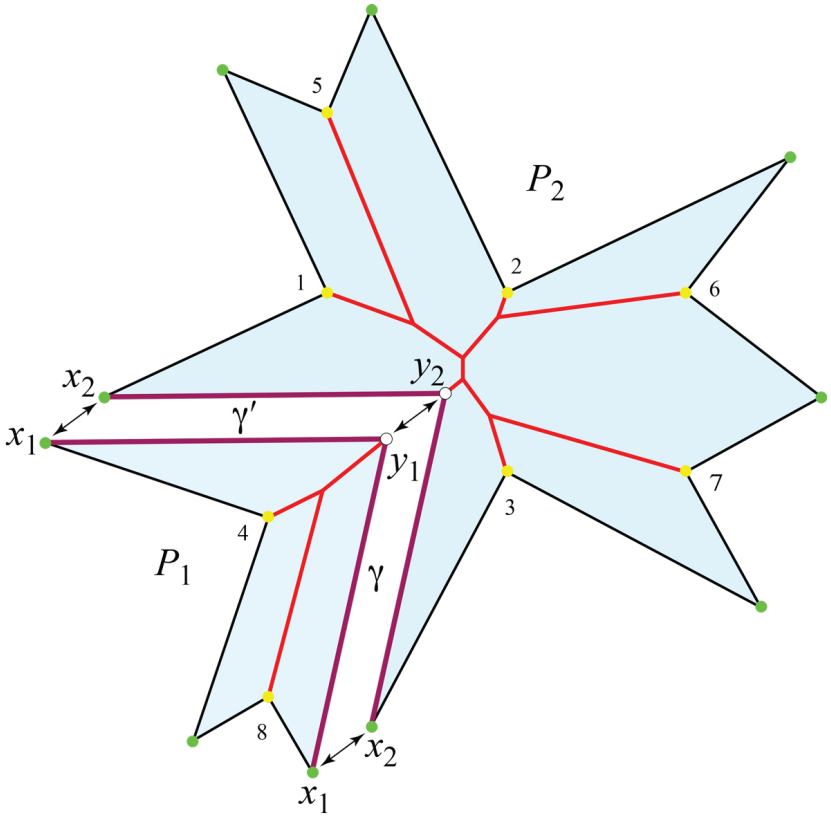

Figure 2.3: Geodesic segments γ and γ' (purple) connect $x{=}x_1{=}x_2$ to $y{=}y_1{=}y_2$. P_1 folds to a tetrahedron, and P_2 to an 8-vertex polyhedron, with x and y vertices in each. P_1 and P_2 are cut open along geodesic segments from x_i to y_i and glued together to form P. Based on the cube unfolding in Fig. 2.2(b).

Let $C = x_1, \ldots, x_n$ be a directed polygonal chain in the plane, with left angles $\theta_i = \angle x_{i-1}x_ix_{i+1}$, where $x_n = x_1$ if C is closed, and θ_1 and θ_n undefined if C is open. Let $d = |x_1x_n|$ be the distance between the endpoints, with $d = 0$ if C is closed. If $\theta_i \leq \pi$ for all i, then C is a *convex chain*. *Cauchy's Arm Lemma* applies to reconfiguration of C while all the edge lengths remain fixed (the edges are *bars*). We use primes to indicate the reconfigured chain. The lemma says that if the θ_i angles are "straightened" while remaining convex, in the sense that $\theta_i \leq \theta_i' \leq \pi$, then the distance $d' = |x_1'x_n'|$ only increases: $d' \geq d$.

For the extension needed later, we reformulate in terms of turn angles $\tau_i = \pi - \theta_i \geq 0$, as described in [O'R01]. Now the straightening condition says that $0 \leq \tau_i' \leq \tau_i$. The extended version of Cauchy's arm lemma says that, as long as, in the reconfiguration, $-\tau_i \leq \tau_i' \leq \tau_i$, then the same conclusion holds: $d' \geq d$. Thus C' might no longer be a convex chain, but its reflexivities are bounded by the original turn angles. One way to interpret $d' \geq d$ is that there is a "forbidden" disk of radius $d = |x_1x_n|$ centered on x_1 that x_n' cannot penetrate. See Fig. 2.4. We summarize in a theorem:

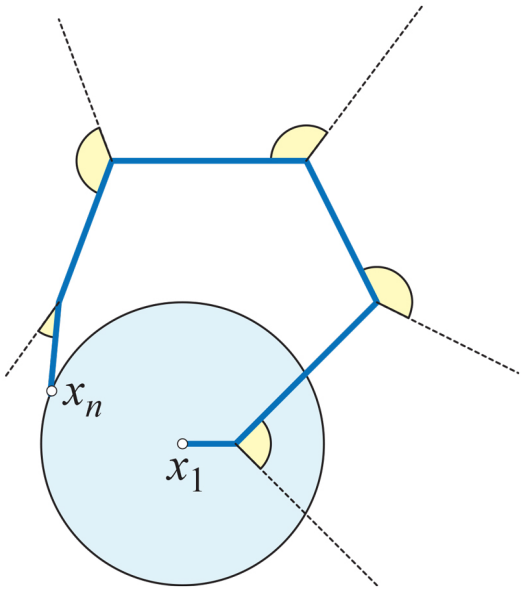

Figure 2.4: Any turns within the indicated $\pm\tau_i$ angle ranges constitute straightening. Then x_n will not penetrate the $|x_1x_n|$ disk.

Theorem 2.7. *(Cauchy's Arm Lemma.) A straightening reconfiguration of*

a planar convex chain C—retaining edge-lengths fixed and confining new turn angles τ_i' to lie in $[-\tau_i, \tau_i]$—only increases the distance between the endpoints: $|x_1'x_n'| \geq |x_1x_n|$.

The next elementary result assures the angle increase, in the frameworks in which we will apply Cauchy's Arm Lemma.

Lemma 2.8. *Consider three rays r_1, r_2, r_3 in $\mathbb{R}^3$, emanating from the point w, and put $\tau_i = \angle(r_i, r_{i+1})$, with $3 + 1 \equiv 1$ mod 3. Then $\theta_1 \leq \theta_2 + \theta_3$.*

Proof. Imagine a unit sphere S centered on w and let $\{s_i\} = r_i \cap S$, and use ρ to indicate spherical distance. Then the claim of the lemma is the triangle inequality for spherical distances: $\rho(s_1, s_2) \leq \rho(s_1, s_3) + \rho(s_2, s_3)$. □

2.3 A Rigidity Result

In this section we present a technical result for later use, which may be of independent interest. The theorem says that two convex polyhedra that are isometric on all but the neighborhoods of one vertex, are in fact congruent. We also show this result cannot be strengthened: two convex polyhedra can differ in the neighborhoods of just two vertices.

Theorem 2.9. *Assume P, Q are convex polyhedra with the same number of vertices, such that there are vertices $p \in P$ and $q \in Q$, and respective small neighborhoods $N_p \subset P$, $N_q \subset Q$ not containing other vertices, and an isometry $\iota : P \setminus N_p \to Q \setminus N_q$. Then P is congruent to Q.*

Proof. The existence of ι on all but neighborhoods of p and q yields, in particular, that the curvatures $\omega_P(p)$ of P at p and $\omega_Q(q)$ of Q at q are equal, to satisfy the curvature sum of 4π (by Gauss-Bonnet).

Take a point $x \in P$ joined to each vertex of P by precisely one geodesic segment, a *generic point* x. Such an x maybe be found in a "ridge-free" region of P [AAOS97]; it is equivalent to the fact that no vertex of P is interior to $C(x)$. Moreover, we may choose x such that $\iota(x)$ has the same property on Q.

Denote by u the ramification point of $C(x)$ neighboring p in $C(x)$, i.e., the ramification point of degree ≥ 3 closest to p. Let v be the similar ramification point of $C(\iota(x))$ neighboring q in $C(\iota(x))$. Since N_p and N_q are small, we

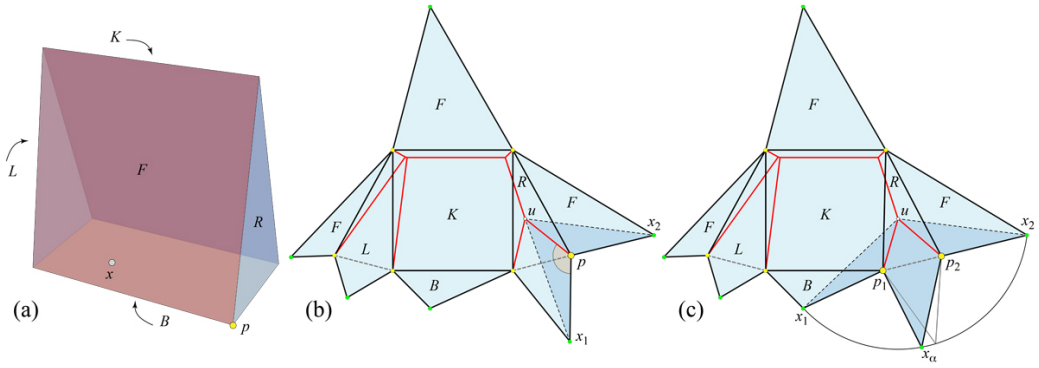

Figure 2.5: (a) A 6-vertex polyhedron P. The F and K faces are unit squares; B is a $1 \times \frac{1}{2}$ rectangle, with $x \in B$. (b) Star-unfolding S_P of P. $\angle x_1 p u$ is marked. (c) Moving x_α on the circle arc moves the bisectors incident to u, and so moves p_1 and p_2. Refolding results in a polyhedron incongruent to (a).

may assume they are disjoint from u and v and all the segments described above.

Star unfold P with respect to x, and Q with respect to $\iota(x)$, and denote by S_P and S_Q the resulting planar polygons. We'll continue to use the symbols p and q, u and v to refer to the corresponding points in S_P and S_Q respectively. Let x_i, $i = 1, 2$ be the images of x surrounding p in S_P, and $\iota(x_i)$ the similar images in S_Q. See Fig. 2.5(a,b).

By hypothesis, we have respective neighborhoods $\bar{N}_p \subset S_P$ and $\bar{N}_q \subset S_Q$ and an isometry $\bar{\iota}$ induced by ι, with $\bar{\iota} : S_P \setminus \bar{N}_p \to S_Q \setminus \bar{N}_q$. Thus in Fig. 2.5(b), all of S_P outside of the wedge (x_1, u, x_2) is identical in S_Q. Therefore the triangles $x_1 u x_2$ and $x_1' v x_2'$ are congruent. Moreover, p lies on the bisector of the angle $\angle x_1 u x_2$, and q lies on the bisector of the angle $\angle x_1' v x_2'$. Since $\angle x_1 p u = \angle x_1' q v = \pi - \frac{1}{2}\omega(p) = \pi - \frac{1}{2}\omega(q)$, p and q are uniquely determined. Consequently, S_P and S_Q coincide, and refolding according to the same gluing identifications leads to congruent P and Q. □

Remark 2.10. *It is perhaps surprising that the above result cannot be extended to claim that isometries excluding neighborhoods of two vertices always imply congruence.*

Proof. If the points $p_1, p_2 \in P$ and $q_1, q_2 \in Q$ do not have a common neighbor in $C(x)$ and $C(\iota(x))$ respectively, the above proof establishes rigidity.

Next we focus on P, and try to find positions for $p_1, p_2 \in P$ determined by the hypotheses. Assume, in the following, that $p_1, p_2 \in P$ have a common degree-3 ramification neighbor u in $C(x)$.

Star unfold P with respect to some $x \in P$, to S_P. The region of S_P exterior to the wedge (x_1, u, x_2) is uniquely determined and identical in S_Q. See Fig. 2.5(c).

Take a point x_α on the circle of center u and radius $|x_1u| = |x_2u|$. We now argue that positions of x_α on this circle allow p_1 and p_2 to vary while maintaining all outside of the (x_1, u, x_2) wedge fixed.

Let $\angle x_\alpha u x_1 = 2\alpha$. On the bisector of that angle incident to u, one can uniquely determine a point p_1 such that $\angle x_1 p_1 u = \pi - \frac{1}{2}\omega(p_1)$. Similarly, one can uniquely determine a point p_2 on the bisector of that angle $\angle x_\alpha u x_2$, such that $\angle x_2 p_2 u = \pi - \frac{1}{2}\omega(p_2)$.

Thus we have identified a continuous 1-parameter family of star-unfoldings, and consequently of convex polyhedra, verifying the hypotheses. □

2.4 Vertex-Merging

Digon-tailoring is, in some sense, the opposite of *vertex-merging*, a technique introduced by A. D. Alexandrov [Ale05, p. 240], and subsequently used later by others, see e.g. [Zal07], [OV14], [O'R20]. We will employ vertex-merging in Chapters 8 and 9, and focus on it in Part II.

Consider two vertices v_1, v_2 of P of curvatures ω_1, ω_2, with $\omega_1 + \omega_2 < 2\pi$, and cut P along a geodesic segment γ joining v_1 to v_2. Construct a planar triangle $T = \bar{v}'\bar{v}_1\bar{v}_2$ of base length $|\bar{v}_1 - \bar{v}_2| = |\gamma|$ and the base angles equal to $\omega_1/2$ and $\omega_2/2$ respectively. Glue two copies of T along the corresponding lateral sides, and further glue the two bases of the copies to the two "banks" of the cut of P along γ. By Alexandrov's Gluing Theorem (AGT), the result is a convex polyhedral surface P'. On P', the points (corresponding to) v_1 and v_2 are no longer vertices because exactly the angle deficit at each has been sutured-in; they have been replaced by a new vertex v' of curvature $\omega' = \omega_1 + \omega_2$. So vertex-merging always reduces the number of vertices of P by one. See Fig. 2.6.

In order to repeat vertex-merging, we need to know when there is a pair of vertices that can be merged. This is answered in the following lemma.

Lemma 2.11. *Every convex polyhedron Q has at least one pair of vertices*

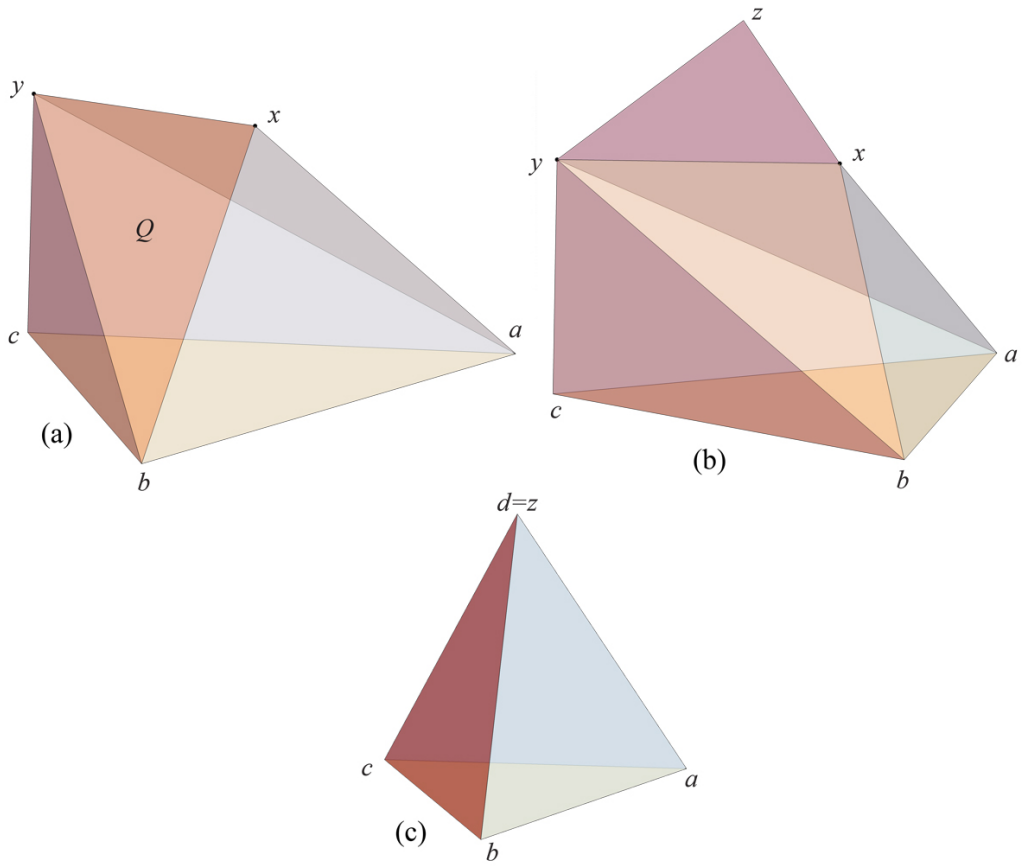

Figure 2.6: (a) Q is a 5-vertex, 6-face polyhedron, symmetric about a vertical plane through xy. Its base abc is an equilateral triangle. (b) Vertex merging x and y by gluing a pair of xyz triangles. (c) The merging reduces Q to a regular tetrahedron (not to same scale). Cf. Fig. 1.1(a,b).

admitting merging, unless it is an isosceles tetrahedron or a doubly-covered triangle.

Recall from Chapter 1 that an *isosceles tetrahedron* T is a tetrahedron with four congruent faces, and each vertex of curvature π. See Fig. 1.4.

Proof. If there is a pair of vertices whose sum of curvatures is strictly less than 2π, then vertex-merging is possible, as just described. So assume that, for any two vertices of Q, their sum of curvatures is at least 2π. In this case, it must be that $n \leq 4$. Indeed, since $\sum_{v \in Q} \omega(v) = 4\pi$ (by the Gauss-Bonnet theorem), if the sum of at least 5 positive numbers is 4π then the smallest two have sum $< 2\pi$.

If $n = 3$, the cut locus of any vertex is a line-segment, by Lemma 2.1(i), so Q is a doubly-covered triangle (see Lemma 2.2).

If $n = 4$ then necessarily all vertex curvatures of Q are π. Indeed, if the sum of 4 positive numbers is 4π then the smallest two have the sum $\leq 2\pi$, with equality if and only if all are π. So Q is an isosceles tetrahedron. $\square$

Chapter 3

Domes and Pyramids

One of our goals in this work, achieved in Theorem 4.6, is to show that if Q can be obtained from P by sculpting, then it can also be obtained from P by tailoring. A key step (Lemma 4.1) repeatedly slices off shapes we call g-domes. Each g-dome slice can itself be achieved by slicing off pyramids, i.e., by suitable vertex truncations. Lemma 4.4 will show that slicing off a pyramid can be achieved by tailoring, and thus leading to Theorem 4.6.

Our main goal in this chapter is to prove that g-domes can be partitioned into "stacked" pyramids, a crucial part of the above process. We illustrate the slice $\rightarrow$ g-domes and g-dome $\rightarrow$ pyramids of the process on a simple example, a tetrahedron Q inside a cube P. Later (in Chapter 4) this example will be completed by tailoring the pyramids.

3.1 Domes

As usual, a *pyramid* P is the convex hull of a convex polygon *base* X, and one vertex v, the *apex* of P, that does not lie in the plane of X. The degree of v is the number of vertices of X.

A *dome* is a convex polyhedron G with a distinguished face X, the *base*, and such that every other face of G shares a (positive-length) edge with X. Domes have been studied primarily for their combinatorial [Epp20], [EL13] or unfolding [DO07] properties. In [Epp20] they are called "treetopes" because removing the base edges from the 1-skeleton leaves a tree, which the author calls the *canopy*.[1] Here we need a slight generalization.

[1] These polyhedra are not named in [EL13].

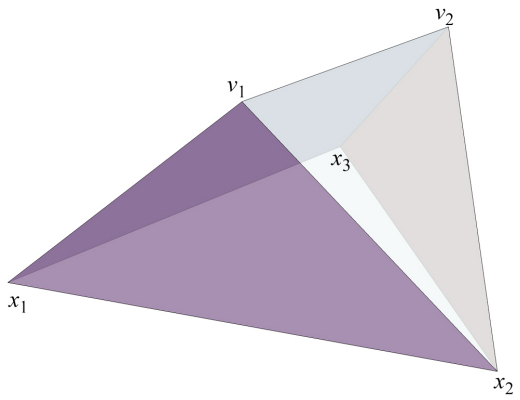

Figure 3.1: A g-dome with base $\triangle x_1x_2x_3$ and top-canopy v_1v_2.

A *generalized-dome*, or *g-dome* G, has a base X, with every other face of G sharing either an edge or a vertex with X. Every dome is a g-dome, and it is easy to obtain every g-dome as the limit of domes. An example is shown in Fig. 3.1, which also shows that removing base edges from the 1-skeleton does not necessarily leave a tree: (v_1, x_2, v_2) forms a cycle. Let us define the *top-canopy* T of a g-dome G as the graph that results by deleting from the 1-skeleton of G all base vertices and their incident edges. In Fig. 3.1 the top-canopy is v_1v_2.

Lemma 3.1. *The top-canopy T of a g-dome G is a tree.*

Proof. If G is a dome, the claim follows, because even including the edges incident to the base X results in a tree, and removing those edges leaves a smaller tree.

If G is not a dome, then slice it with a plane parallel to, and at small distance above, the base. The result is a dome, and we can apply the previous reasoning. □

3.2 Cube/Tetrahedron Example

In this section we start to track an example, tailoring a cube to a tetrahedron. P is a triangulated unit cube. Q is a right tetrahedron in the corner of P; see Fig. 3.2(a). Q can be obtained from P by a single slice by plane $\Pi = p_1p_3p_8$. The algorithm to be described in Chapter 6 will first produce two g-domes G_1

and G_2, and reducing those g-domes will lead to four pyramids P_1, P_2, P_3, P_4. Later (in Chapter 4) we will show how each of those pyramids is reduced by digon-tailoring, finally yielding Q.

3.2.1 Slice → G-Domes for Cube/Tetrahedron

Now we describe at an intuitive level what will be proved in Lemma 4.1: that the slice by Π can be effected by partitioning the sliced portion of P into g-domes.

We start by rotating the plane Π about the edge p_1p_3 until it encounters a face of P; call that plane $\Pi_0 = p_1p_2p_3$. Now continue rotating until we reach Π_1 when (a) The portion of P between Π_0 and Π_1 is a g-dome with base on Π_1, and (b) Any further rotation would render the portion between the planes to a polyhedron no longer a g-dome. $\Pi_1 = p_1p_3p_7p_5$, because any further rotation would isolate the face $p_5p_6p_7$ of P, no longer touching Π_1. Note here the top face of the cube is triangulated with the diagonal p_5p_7. See Fig. 3.2(bc). Continuing to rotate, we reach the original Π, partitioning the g-dome illustrated in Fig. 3.2(d). So G_1 has base $X_1 = p_1p_3p_7p_5$ and top-canopy p_2p_6, and G_2 has base $X_2 = p_1p_2p_3$ and top-canopy p_5p_7. It is clear that removing $G_1 \cup G_2$ would reduce P to Q.

3.3 Proof: G-dome → Pyramids

We now prove the reduction of a g-dome to pyramids, using Fig. 3.3 to illustrate the proof details. Afterward we will return to the cube/tetrahedron example to apply the constructive proof to G_1 and G_2.

Theorem 3.2. *Every g-dome G of base X can be partitioned into a finite sequence of n pyramids P_i with the following properties:*

- *Each P_i has a common edge with X.*
- *Each $G_j = G \setminus \bigcup_{i=1}^{j} P_i$ is a g-dome, for all $j = 1, ..., n$.*
- *The last pyramid P_n in the sequence has the same base X as G.*

The partition into pyramids specified in this theorem has the special properties listed. Without these properties, it would be easier to prove that every g-dome may be partitioned into pyramids. But the properties are needed for the subsequent digon-tailoring described in Chapter 4. In particular, the

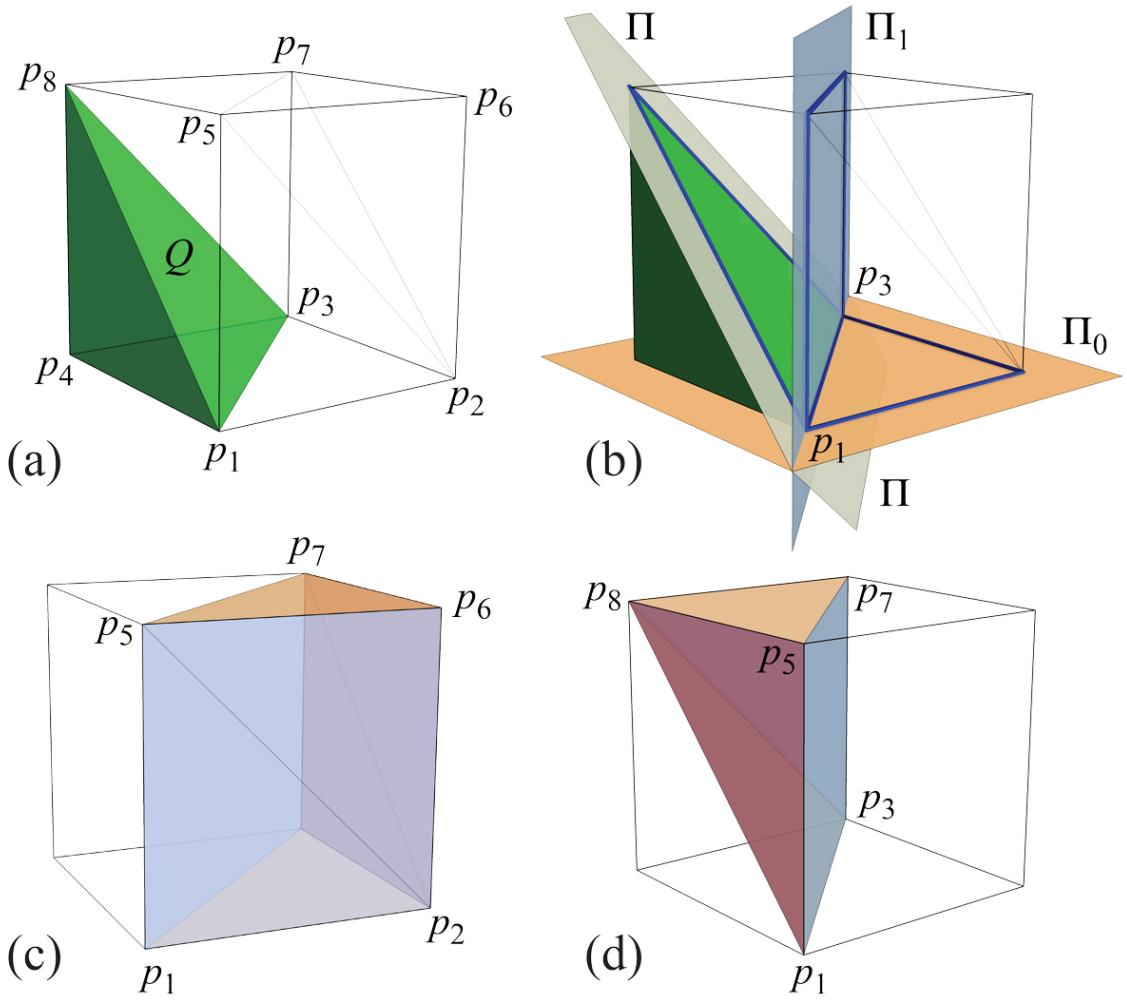

Figure 3.2: (a) Q in corner of cube P. (b) Partitioning sliced portion of P into two g-domes. (c,d) G-domes G_1 and G_2.

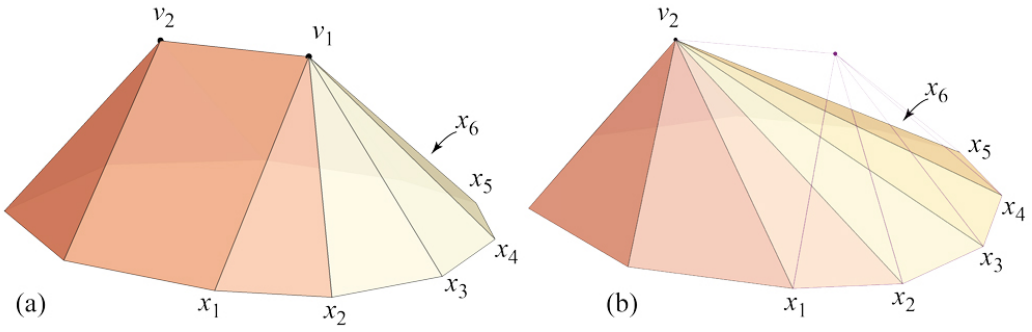

Figure 3.3: (a) A dome G. Vertex v_1 of degree $k_1 = 6$ is adjacent to X, i.e., there are edges v_1x_i. (b) G' after removal of v_1.

pyramids are "stacked" in the sense that each can be sliced-off in sequence, with each slice retaining X unaltered.

The proof is a double induction, and a bit intricate. One induction simply removes one vertex v_1 from the top-canopy. The second induction, inside the first one, reduces the degree of v_1 to achieve removal of v_1, at the cost of increasing the degree of v_2.

Proof. Let m be the number of vertices in the top-canopy T of G. If $m = 1$, G is already a pyramid, and we are finished. So assume G's top-canopy T has at least two vertices. Choose v_1 to be a leaf of T given by Lemma 3.1, and v_2 its unique parent. Let v_1 be adjacent to k vertices of X. If G is a dome, v_1 has degree $k + 1$; if G is a g-dome, then possibly v_1 has degree $k + 2$. Since the later case changes nothing in the proof, we assume for the simplicity of exposition that G is a dome. Our goal is to remove v_1 through a series of pyramid subtractions.

Let the vertices of X adjacent to v_1 be $x_1, x_2, \ldots, x_k$. Let Π_1 be the plane $x_1x_2v_2$. This plane cuts into G under v_1, and intersects the edges v_1x_i, $i \geq 3$, in points a_i. In Fig. 3.4(a), those points are a_3, a_4, a_5, a_6. Remove the pyramid whose apex is v_1 and whose base (in our example) is $x_1, x_2, a_3, a_4, a_5, a_6, v_2$.

We now proceed to reduce the chain of new vertices $a_3, \ldots, a_k$ one-by-one until only a_k remains.

First, with the plane $x_2x_3a_4$, we slice off the tetrahedron whose degree-3 apex is a_3; Fig. 3.4(b). Next, with the plane $x_2x_3a_5$, we slice off the pyramid with apex a_4. Unfortunately, because a_4 has degree-4, this introduces a new vertex a_4'; Fig. 3.4(c). So next we slice with $x_3x_4a_5$ to remove the tetrahedron

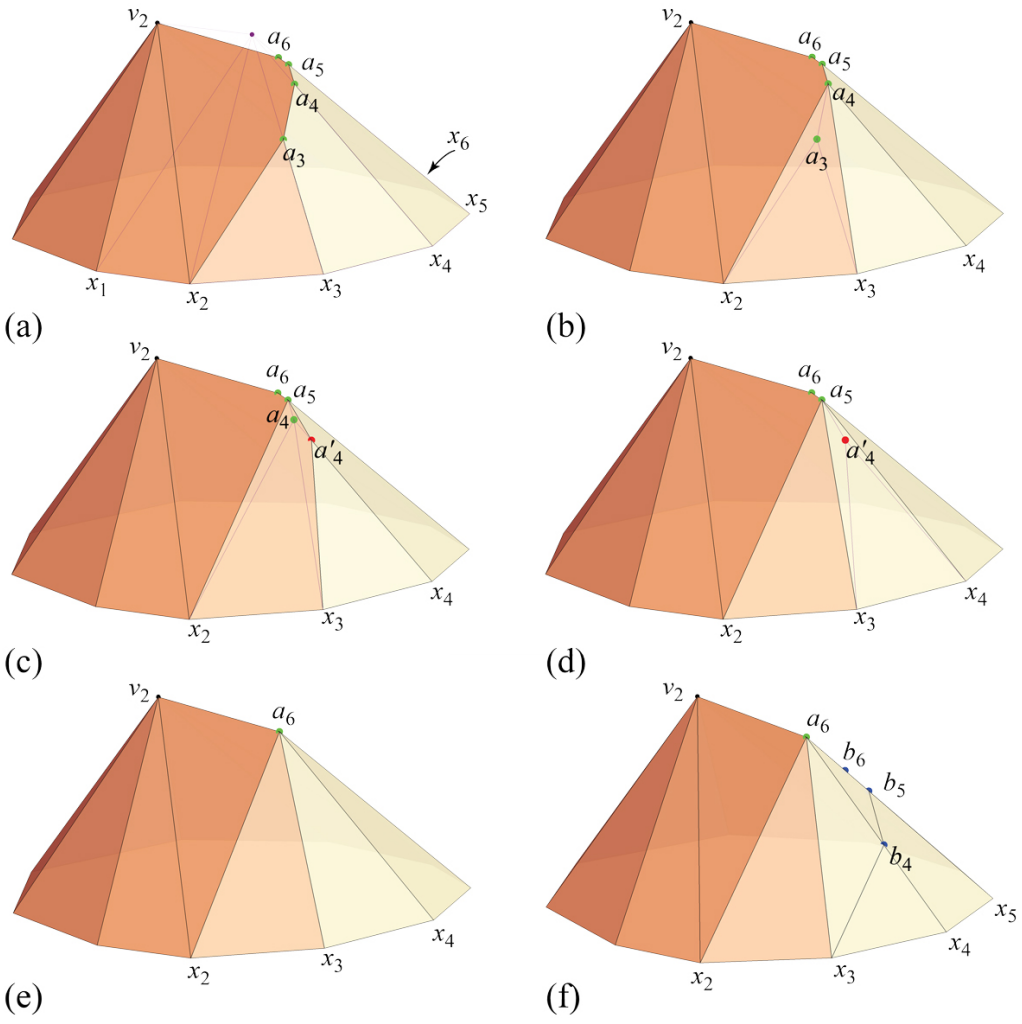

Figure 3.4: (a) After slicing by the plane $\Pi_1 = x_1x_2v_2$. (b) Removing a_3. (c) Removing a_4 creates a'_4. (d) Removing a'_4. (e) a_6 has replaced v_1 with one lower degree. (f) Intersection points with $\Pi_2 = x_2x_3v_2$.

whose degree-3 apex is a_4'; Fig. 3.4(d). Continuing in this manner, alternately removing a tetrahedron followed by a pyramid with a degree-4 apex, we reach Fig. 3.4(e).

Note that v_1 was connected to $k = 6$ vertices of X, but a_6 is only connected to 5: the connection of v_1 to x_1 has in a sense been transferred to v_2. In general, a_k has degree one less than v_1's degree, and the degree of v_2 has increased by 1.

Now we repeat the process, starting by slicing with $\Pi_2 = x_2x_3v_2$, which intersects the a_kx_i edges at $b_4, \ldots, b_k$. We remove the pyramid apexed at a_6 with base (in our example) of $x_2, x_3, b_4, b_5, b_6, v_2$; Fig. 3.4(f). The same methodical technique will remove all but the last new vertex b_k, which replaces a_k but has degree one smaller.

Continuing the process, slicing with $\Pi_i = x_ix_{i+1}v_2$, up to $i = k - 1$, will lead to the complete removal of v_1, as previously illustrated in Fig. 3.3(b), completing the inner induction. Inverse induction on the number of vertices of the g-dome then completes the proof: each step reduces the number of top vertices by 1, and the degree of v_2 increases by $k + 1$, the degree lost at v_1.

With $G_0 = G$, each G_j is the intersection of G_{k-1} with a closed half-space containing a base edge, so it is convex for all $j = 1, ..., n$. Indeed each G_j is a g-dome, because all untouched faces continue to meet X in either an edge or a vertex, and new faces always share an edge with X. □

The partition of a g-dome into pyramids given by Theorem 3.2 is not unique. For our example in Fig. 3.3(a), we finally get the pyramid apexed at v_2 in (b) of the figure, but we could as well have ended with a pyramid apexed at v_1.

We will see in Chapter 4 that one g-dome of $O(n)$ vertices reduces to $O(n^2)$ pyramids of constant size, and $O(n)$ pyramids each of size $O(n)$.

3.3.1 G-domes → Pyramids for Cube/Tetrahedron

We now return to the cube/tetrahedron example, and show that following the proof of Theorem 3.2 partitions the two g-domes G_1 and G_2 into pyramids. We will see that the two g-domes partition into two pyramids each.

The analysis for G_1 is illustrated in Fig. 3.5. Note G_1 has been reoriented so that its base $X_1 = p_1p_3p_7p_5$ is horizontal, and relabeled $x_1x_2x_3x_4$ to match the proof.

We proceed to describe each step, making comparisons to Figs. 3.3 and 3.4.

(a) In (a) of the figure, the base of the g-dome has been oriented horizontally, and the correspondence between the cube labels and the labels used in the proof of Theorem 3.2 are shown. The top-canopy of G_1 is v_1v_2, with v_1 of degree-5 and v_2 of degree-3. Eventually v_1 will be removed and the degree of v_2 increased by 1, as in Fig. 3.3.

(b) The first step is to slice off a pyramid apexed at v_1 by the plane $x_1x_2v_2$. Due to the coplanar triangles, this has the effect of "erasing" the diagonal x_1v_1, as would have occurred if that diagonal had a dihedral angle different from π.

(c) Next, the slice $x_2x_3v_2$ cuts off a pyramid P_1 apexed at v_1, a pyramid we will analyze in Section 4.4.1, leaving the shape in (d).

(d) A new vertex a_3 is created by the slice in (c), with degree-4. This is the "replacement" for v_1 but of smaller degree. This corresponds to a_3 in Fig. 3.4(a), although here there are no further vertices $a_4, \ldots$ along an a-chain.

(e) We again slice with $x_2x_3v_2$, which reduces the degree of a_3 to 3. This a_3 corresponds to b_4 in Fig. 3.4(f).

(f) A final slice by $x_3x_4v_2$ removes a_3, leaving the pyramid P_2 shown. This completes the removal of v_1 in (a), increasing the degree of v_2 from 3 to 4. Because what remains is a pyramid, the processing stops.

So G_1 is partitioned into two pyramids, P_1 and P_2, apexed at p_2 and p_6 respectively.

We perform a similar analysis for the simpler G_2, shown in Fig. 3.6, again reoriented so that its base $p_8p_1p_3$ is horizontal and relabeled $x_1x_2x_3$. Here the top-canopy is v_1v_2, and the first slice $x_1x_2v_2$ reduces G_2 to a pyramid P_2. So G_2 is also partitioned into two pyramids, P_3 and P_4, apexed at p_5 and p_7 respectively. Together G_1 and G_2 are partitioned into four pyramids.

We will complete this example by tailoring the four pyramids in the next chapter.

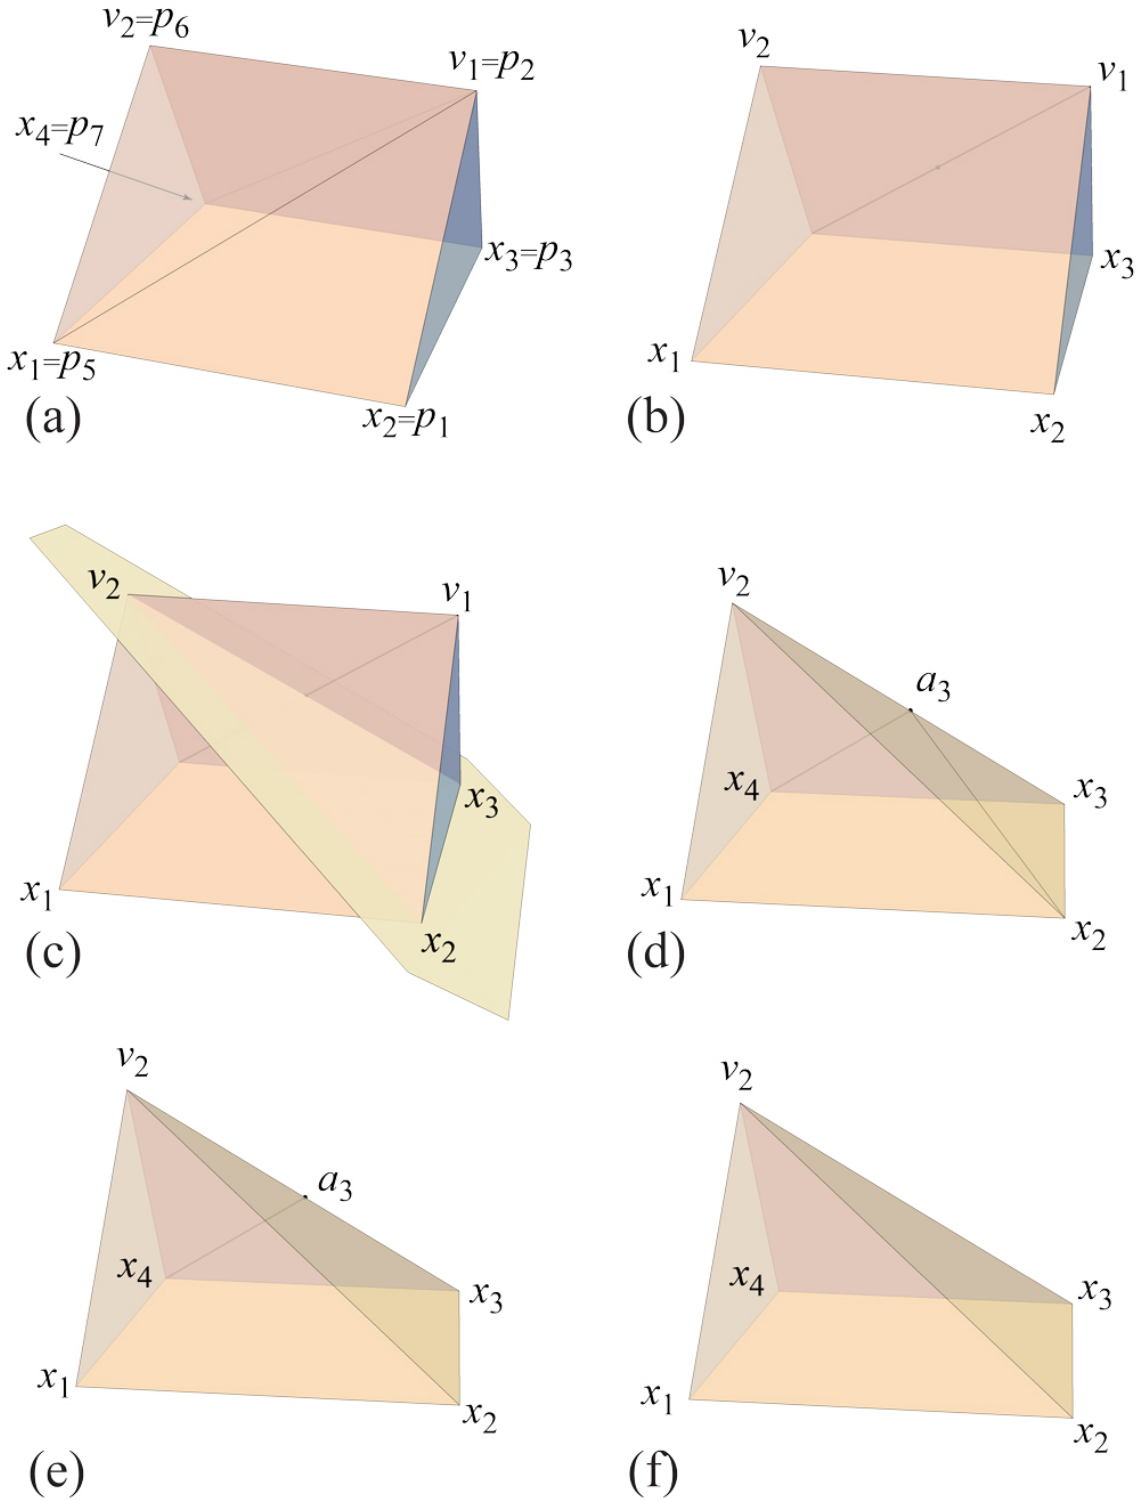

Figure 3.5: (a) g-dome G_1. (b) After slicing by $x_1x_2v_2$. (c,d) Slicing by $x_2x_3v_2$. (e,f) Two further slices remove a_3.

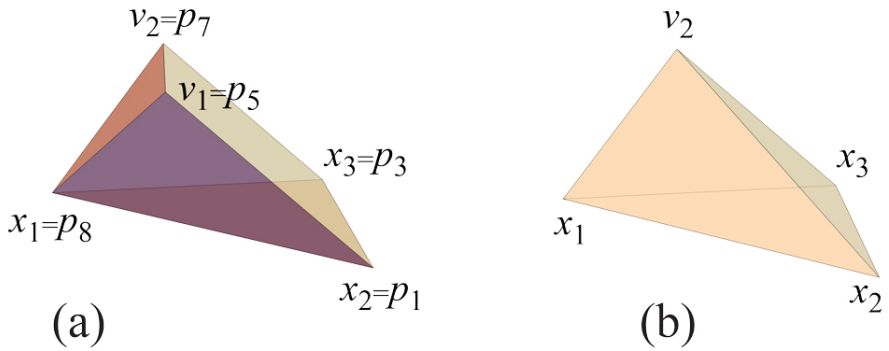

Figure 3.6: (a) G_2 with corresponding cube labels. (b) After slicing with $x_1x_2v_2$.

Chapter 4

Tailoring via Sculpting

In this chapter we complete the proof that one slice of P by plane Π can be tailored to the face of Q lying in Π, following the sequence slice $\rightarrow$ g-domes $\rightarrow$ pyramids $\rightarrow$ tailoring. The previous chapter established the g-domes $\rightarrow$ pyramids link. Here we first prove the relatively straightforward slice $\rightarrow$ g-domes process, and then concentrate on the more complex pyramid $\rightarrow$ tailoring step.

4.1 Slice $\rightarrow$ G-domes

Lemma 4.1. *With $Q \subset P$ and Π a plane slicing P and containing a face F of Q, the sliced-off portion P' can be partitioned into a* fan *of $O(n)$ g-domes $G_1, G_2, \ldots$, a fan in the sense that the bases of the g-domes all share a common edge of F.*

Proof. Assume Π is horizontal, with Q the portion below Π and P' the portion above. Denote by x_i the vertices of Q in Π, $i = 1, \ldots, m$; call the top face of Q with these vertices F. Let e be any edge of F, say $e = x_1x_2$, and let F' be the face of Q sharing e with F. Call the plane lying on F' Π_0.

Now imagine rotating Π_0 about e toward P', noting as it passes through each vertex $v_1, v_2, \ldots$ in that order. For perhaps several consecutive vertices, the portion of P' between the previous and the current plane is a g-dome, but rotating further takes it beyond a g-dome. More precisely, let Π_1 through v_{j_1} be the plane such that, in the sequence

$$v_1, v_2, \ldots, v_{j_1}, v_{j_1+1}, \ldots, v_{j_2}, v_{j_2+1} \cdots ,$$

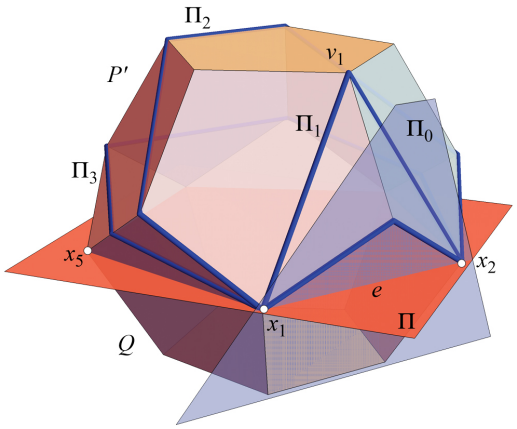

Figure 4.1: Dodecahedron sliced by Π. The polyhedron between each pair Π_i, Π_{i+1} is a g-dome.

the portion G_1 of P' between Π_0 and Π_1, including the vertices $v_1, v_2, \ldots, v_{j_1}$, is a g-dome based on Π_1, but the portion rotating further to include one more vertex, v_{j_1+1}, is not a g-dome. Π_2 through v_{j_2} is defined similarly: the portion G_2 between Π_1 and Π_2 including $v_{j_1+1}, \ldots, v_{j_2}$ is a g-dome based on Π_2, but including v_{j_2+1} it ceases to be a g-dome. Here for each Π_i, Π_{i+1} pair, the base X of the g-dome lies in Π_{i+1}.

Fig. 4.1 illustrates the process. Here Π_1 is through $v_1 = v_{j_1}$, and the last g-dome lies between Π_3 and Π.

So we have now partitioned P' into g-domes $G_1, G_2, \ldots$. □

We will invoke Lemma 4.4 (ahead) to reduce each g-dome to its base by tailoring, in the order $G_1, G_2, \ldots$. This will reduce P' to just the top face F of Q.

Having established the claim of the lemma for one slice, it immediately follows that it holds for an arbitrary number of slices. This was already illustrated in the cube/tetrahedron example in Chapter 3, and we will use it again in Theorem 4.6.

We should note that it is at least conceivable that only a single g-dome is needed above. See Open Problem 18.2.

4.2 Pyramid → Tailoring

The goal of this section is to prove that removal of a pyramid, i.e., a vertex truncation, can be achieved by (digon-)tailoring. We reach this in Lemma 4.4: a degree-k pyramid P can be removed by $k-1$ tailoring steps, each step excising one vertex by removal of and then sealing a digon. The first $k-2$ of these digons each have one endpoint a vertex, and so leave the total number of vertices of P at $k+1$. The $(k-1)$-st digon has both endpoints vertices, and so its removal reduces the number of vertices to the k base vertices.

We start with Lemma 4.2 which claims the result but only under the assumption that the slice plane is close to the removed vertex. Although this lemma is eventually superseded, it establishes the notation and the main idea. Following that, Lemma 4.3 removes the "sufficiently small" assumption of Lemma 4.2, but in the special case of P a pyramid. Finally we reach the main claim in Lemma 4.4, which shows this special case encompasses the general case.

In the following, we use ∂S to indicate the 1-dimensional boundary of a 2-dimensional surface patch S.

4.2.1 Notation

To help keep track of the notation throughout this critical section, we list the main symbols below.

- Initially P and Q are polyhedra with $Q \subset P$, with P sliced by plane Π to truncate vertex v of degree k.
- Later we specialize P to be the pyramid sliced off.
- $X = \Pi \cap P$ is the base of the pyramid, with vertices $\partial X = x_1 x_2 \dots x_k$.
- The lateral faces of P are denoted by L, so $P = X \cup L$.
- $D(x_i, y_i)$ is a digon with endpoints $x_i \in \partial X$ and y_i.
- P_j is the modified pyramid after j digon removals.
- $C(x_i, P_j)$ is the cut locus $C(x_i)$ on P_j.
- The first ramification point of $C(x_i, P_j)$ beyond y_i is a_i.

Notation will be repeated and supplemented within each proof.

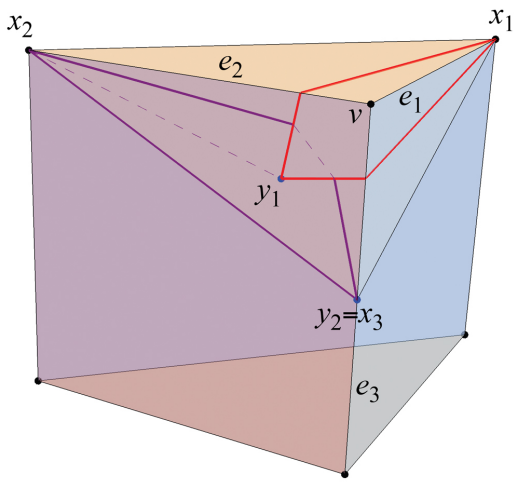

Figure 4.2: P is a prism whose top face is an isosceles right triangle. The truncated vertex v has degree k=3. x_3 is the midpoint of e_3. Digon $D_1 \supset \{v\}$ is shown red; $D_2 \supset \{y_1\}$ is purple. D_1 is sutured closed before D_2 is excised. The last replacement vertex y_{k-1} must be identical to x_k.

4.3 Small volume slices

Lemma 4.2. *Let P be a convex polyhedron, and Q the result obtained by slicing P with a plane Π at sufficiently small distance to a vertex v of P, and removing precisely that vertex. Then Q can be obtained from P by $k-1$ tailoring steps.*

Proof. Let the vertex v to be removed have degree k in the 1-skeleton of P. Let e_i, $i = 1, \ldots, k$, be the edges incident to v, and x_i the intersection of the slicing plane Π with those edges: $\{x_i\} = \Pi \cap e_i$.

We will illustrate the argument with the right triangular prism shown in Fig. 4.2, where $k = 3$ and $\Pi = x_1x_2x_3$. Note that we do not exclude the case when some (or all) of the x_i are vertices of P.

Denote by $\omega_P(x_i)$ $i = 1, \ldots, k$ the curvatures of P at x_i, and by $\omega_Q(x_i)$ the corresponding curvatures of Q. The curvature $\omega_P(v)$ will be distributed to the x_i.

In the figure, $\omega_P(v) = 90°$, the curvatures of the three x_i are $135°, 135°, 0°$ in P, and approximately $156°, 156°, 48°$ in Q. Indeed the increases sum to $90°$: $21° + 21° + 48°$.

The goal now is to excise $k-1$ digons with one end at $x_1, x_2, \ldots, x_{k-1}$, removing precisely the surface angle needed to increase $\omega_P(x_i)$ to $\omega_Q(x_i)$. After digon removals at $x_1, \ldots, x_i$, we call the resulting polyhedron P_i.

Let a digon with endpoints x_i and y_i be denoted $D_i = (x_i, y_i)$. Cut out from P the digon $D_1 = (x_1, y_1)$ containing only the vertex v in its interior, of angle at x_1 equal to $\omega_Q(x_1) - \omega_P(x_1)$. By the assumption that the slice plane Π is sufficiently close to v, the curvature difference is small enough so that D_1 includes only v. Again by the sufficiently-close assumption, we may assume the digon endpoint y_1 lies on the edge of $C(x_1)$ incident to v, prior to the first ramification point of $C(x_1)$. After suturing closed the digon geodesics, y_1 becomes a vertex of curvature $\omega_P(v) - (\omega_Q(x_1) - \omega_P(x_1))$. In the figure, y_1 has curvature $90° - 21° \approx 69°$. In a sense, y_1 "replaces" v.

Next cut out a digon $D_2 = (x_2, y_2)$ containing only the vertex y_1 in its interior, of angle at x_2 equal to $\omega_Q(x_2) - \omega_P(x_2)$. The newly created vertex y_2 "replaces" y_1. Continue cutting out digons $D_i = (x_i, y_i)$ up to $i = k-1$, each D_i surrounding y_{i-1}, and replacing y_{i-1} with y_i.

Because these tailorings have sharpened the curvatures $\omega_P(x_i)$ to match the after-slice curvatures $\omega_Q(x_i)$, it must be that the curvature at the last replacement vertex y_{i-1} is the same as the curvature at x_k: $\omega_P(y_{k-1}) = \omega_Q(x_k)$ (to satisfy Gauss-Bonnet). So now the tailored P_{k-1} matches Q in both the positions of the vertices x_i, $i = 1, \ldots, k-1$, and their curvatures; the only possible difference is the location of y_{i-1} compared to x_k. But the rigidity result, Theorem 2.9, implies that $y_{k-1} = x_k$, and P_{k-1} and Q are now congruent. □

The "sufficiently-small" assumption in the preceding proof allowed us to assume that the digon $D_i = (x_1, y_1)$ endpoint y_1 lay on the segment of $C(x_i)$ incident to v prior to the first ramification point a_1 of $C(x_1)$. Recall that $\omega(x_1) + \omega(y_i) = \omega(v)$, and the further along the segment va_1 that y_1 lies, the larger the digon angle at x_1 is. The procedure would be problematic if the digon angle at x_1 were not large enough even with y_1 at that ramification point a_1. The next lemma removes the sufficiently-small assumption in the special case when P is itself a pyramid, and the vertex truncation reduces P its base, doubly-covered. Following this, we will show that the case when P is a pyramid is the "worst case," and so the general case follows.

4.3.1 Pyramid case

Lemma 4.3. *Let P be a pyramid over base X. Then one can tailor P to reduce it to X doubly-covered, using $k-1$ digon removal steps.*

Proof. We continue to use the notation in the previous lemma, and introduce further notation needed here. Let $L = P \setminus X$ be the lateral sides of the pyramid P; so $P = L \cup X$. After each digon $D_i = (x_i, y_i)$ is removed and sutured closed, the convex polyhedron guaranteed by Alexandrov's Gluing Theorem will be denoted by P_i. We continue to view P_i as $P_i = L_i \cup X_i$, even though already P_1, is in general no longer a pyramid. We will see that all the digon excisions occur on L_i, while X_i remains isometric to the original base X, but no longer (in general) planar.

We will use $C(x_i, P_j)$ to mean the cut locus of x_i on P_j. Regardless of which P_j is under consideration, we will denote by a_i the first ramification point of $C(x_i)$ immediately beyond the vertex y_{i-1} surrounded by the digon $D_i = (x_i, y_i)$.

We need to establish two claims:

Claim (1): The cut locus $C(x_{i+1}, P_i)$ is wholly contained in L_i.

Claim (2): The digon angle α_{i+1} at x_{i+1} to a_{i+1} is large enough to reduce the L-angle at x_{i+1} to its X-angle on the base.

Before addressing the general case of these claims, we illustrate the situation for x_1, referencing Fig. 4.3. The digon $D_1(x_1, y_1)$ surrounding v places y_1 on the va_1 segment of $C(x_1, P) = C(x_1)$. If one imagines y_1 sliding along va_1 from v to a_1, the digon angle at x_1, call it δ_1, increases. To show that y_1 can be placed so that δ_1 is large enough to reduce the angle at x_1 to its angle in X will require a_1 to lie in L (rather than in X).

It turns out that $C(x_1) \subset L$ follows from a lemma in [AAOS97].[1] However, after removing D_1 and invoking Alexandrov's Gluing Theorem, we can no longer apply this lemma. With this background, we now proceed to the general case.

Claim (1): $C(x_{i+1}, P_i) \subset L$. Assume we have removed digons at $x_1, \ldots, x_i$, so that $P_i = X_i \cup L_i$, and L_i contains one vertex y_i, the endpoint of the last

[1] Lem. 3.3: the cut locus is contained in the "kernel" of the star-unfolding which in our case is a subset of L.

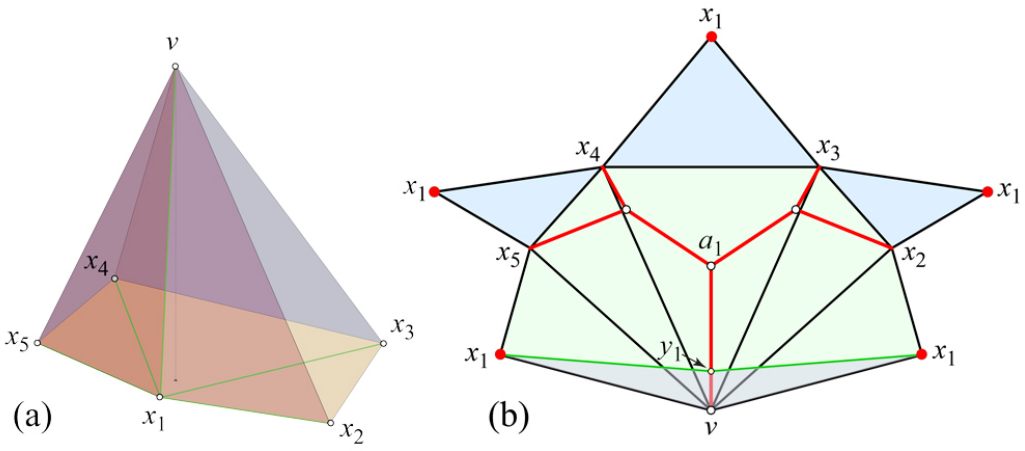

Figure 4.3: (a) A pyramid with pentagonal base X. Shortest paths from x_1 to all vertices are marked green. (b) The star-unfolding with respect to x_1. The triangles from X are blue; those from L are green. $C(x_1) \subset L$ is red. The digon $D_1 = (x_1, y_1)$ is shaded.

digon D_i removed, and X_i contains no vertices. Assume to the contrary of Claim (1) that $C(x_{i+1}, P_i) = C(x_{i+1})$ includes a point z strictly interior to X_i. Because $z \in C(x_{i+1})$, there are two geodesic segments from x_{i+1} to z, call them γ_1^z and γ_2^z. Because X_i contains no vertices, it cannot be that both γ_1^z and γ_2^z are in X_i. Say that γ_1^z crosses L_i. Let $p \in \partial X$ be the first point at which γ_1^z enters X_i, and let $\gamma_1 \subset \gamma_1^z$ be the portion from x_{i+1} to p. See Fig. 4.4.

The geodesic segment γ_1 divides L_i into two parts; let L_i' be the part that does not contain the vertex y_i. Join x_{i+1} to p with a geodesic γ_2 lying in X_i. γ_2 was a shortest path to p in X, but may no longer be shortest in X_i. γ_2 also divides X_i into two parts; let X_i' be the part sharing a portion of ∂X with L_i'.

Now we will argue that $\ell(\gamma_1) \geq \ell(\gamma_2)$, where $\ell(\gamma)$ is the length of γ. This will yield a contradiction, for the following reasons. γ_1 is a shortest geodesic to p, because it extends to γ_1^z, which is a shortest geodesic to z. So γ_2 cannot be strictly shorter than γ_1. Therefore we must have $\ell(\gamma_1) = \ell(\gamma_2)$, which implies that $p \in C(x_{i+1})$. But then γ_1 cannot continue to γ_1^z beyond p, as p is a cut point.

To reach $\ell(\gamma_1) \geq \ell(\gamma_2)$, we will use the extension of Cauchy's Arm Lemma

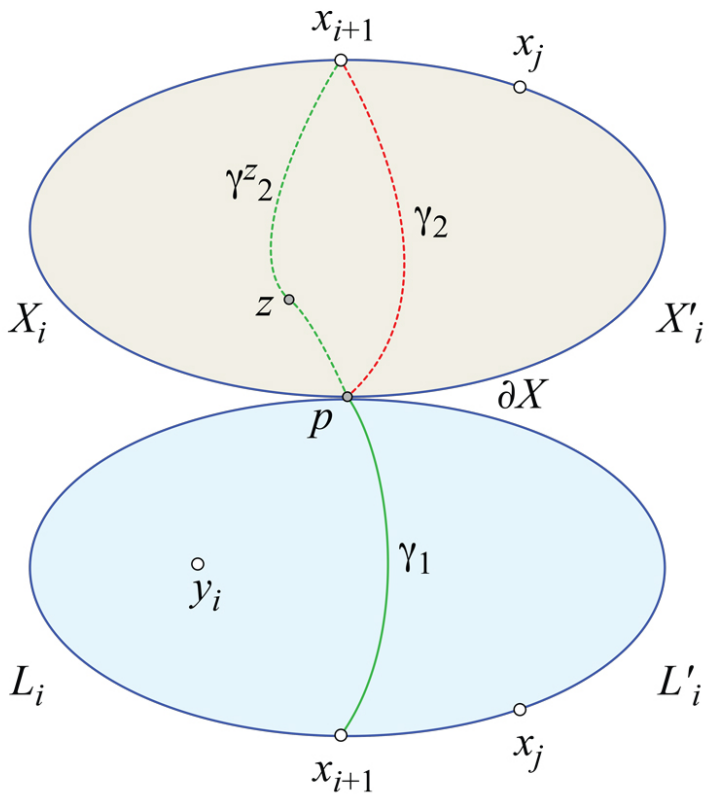

Figure 4.4: X_i is flipped above L_i in this abstract illustration. Cauchy's Arm Lemma ultimately shows that $\ell(\gamma_1) \geq \ell(\gamma_2)$.

described in Chapter 2. Let θ_j^X be the angle at x_j in X_i, and θ_j^L be the angle at x_j in L_i. For $j > i+1$, we know that $\theta_j^L > \theta_j^X$ because P is a pyramid (see Lemma 2.8) and digon removal has not yet reached x_j. We also know that $\theta_j^X \leq \pi$ because X is convex. However, θ_j^L could be nearly as large as $2\pi - \theta_j^X$ if the pyramid P's apex v projects outside the base X.

Let c_X be the planar convex chain in ∂X that corresponds to X_i'; assume c_X is $x_{i+1}, x_{i+2}, \ldots, x_j, \ldots, p$ with angles θ_j^X. (The case where c_X is includes the other part of ∂X, $x_{i+1}, x_i, \ldots, p$ can be treated analoguously.) Then $\ell(\gamma_2)$ is the length of the chord between c_X's endpoints. In order to apply Cauchy's lemma, we rephrase the angles at x_j as turn angles $\tau_j = \pi - \theta_j^X$. The extension of Cauchy's lemma guarantees that, if the chain angles are modified so that the turn angles lie within $[-\tau_j, \tau_j]$, then the endpoints chord length cannot decrease. Roughly, opening (straightening) the angles stretches the chord.

Define c_L as the planar (possibly nonconvex) chain composed of the same vertices x_j that define c_X, but with angles θ_j^L. Because $\theta_j^L \leq 2\pi - \theta_j^X$, the turn angles $\pi - \theta_j^L$ in c_L satisfy

$$\pi - \theta_j^L \geq \pi - (2\pi - \theta_j^X) = -(\pi - \theta_j^X) = -\tau_j \ .$$

Also, because $\theta_j^L > \theta_j^X$,

$$\pi - \theta_j^L \leq \pi - \theta_j^X = \tau_j \ .$$

So the c_L turn angles are in $[-\tau_j, \tau_j]$, and we can conclude from Theorem 2.7 that the c_L endpoints chord length $\ell(\gamma_1)$ is at least $\ell(\gamma_2)$, the c_X endpoints chord length.

We have now reached $\ell(\gamma_1) \geq \ell(\gamma_2)$, whose contradiction described earlier shows that indeed $C(x_{i+1}) \subset L_i$.

Claim (2): $\alpha_{i+1} \geq \theta_{i+1}^L - \theta_{i+1}^X$. Recall that α_{i+1} is the angle at x_{i+1} of the digon from x_{i+1} to a_{i+1}, the first ramification point of $C(x_{i+1})$ beyond the vertex y_i. The claim is that α_{i+1} is large enough to reduce θ_{i+1}^L to θ_{i+1}^X. We establish this by removing a path from $C(x_{i+1})$ and tracking angles, as follows.

From Claim (1), $C(x_{i+1}) \subset L_i$. Let $\rho_{i,i+2}$ be the path in the tree $C(x_{i+1})$ from x_i to x_{i+2}. See Fig. 4.5. Removal of $\rho_{i,i+2}$ from $C(x_{i+1})$ disconnects $C(x_{i+1})$ into the edge $y_i a_{i+1}$, and a series of subtrees T_j. Each T_j shares a point z_j with $C(x_{i+1})$. Let D_j be the digon from x_{i+1} to z_j, and let δ_j be the angle of D_j at x_{i+1}. Finally, let $\Delta_j = \sum_j \delta_j$.

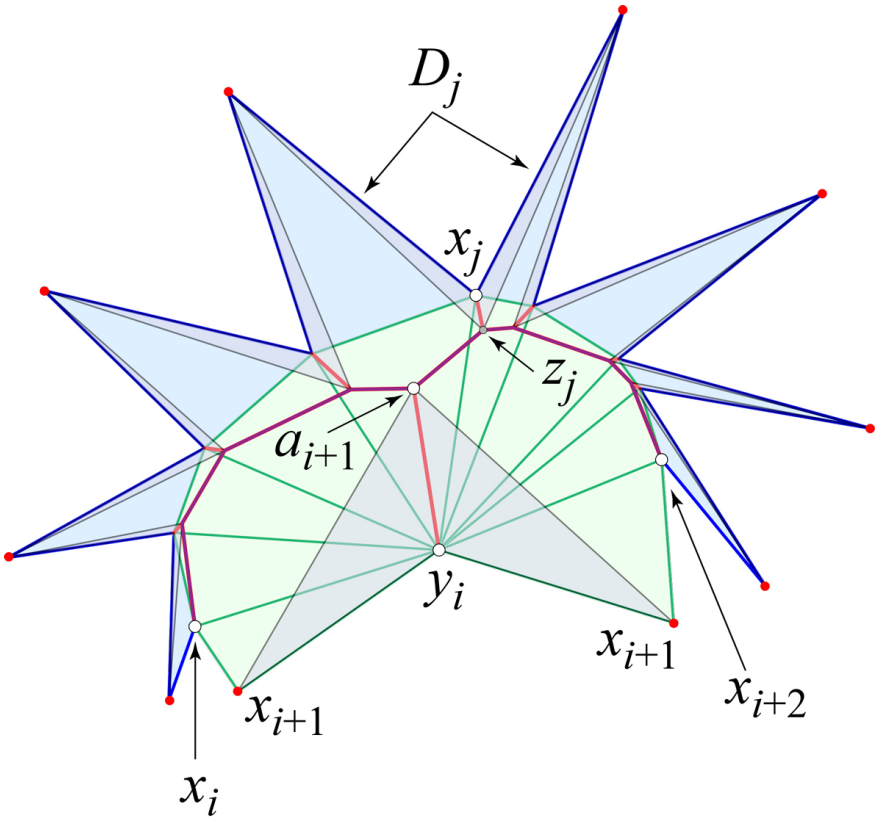

Figure 4.5: Star-unfolding of a pyramid with respect to x_{i+1}. The triangles from X_i are blue; those from L_i are green. Red points are images of x_{i+1}. Digons are shaded. $\rho_{i,i+2}$ is purple; remainder of $C(x_{i+1})$ is red.

Note that all the δ_j angles are in X_i. In contrast, the angle at x_{i+1} in the digon $D_{i+1} = D(x_{i+1}, a_{i+1})$ is in L_i, as illustrated in Fig. 4.5. We defer justifying this claim to later.

Cut off all D_j, and also cut off D_{i+1}. Suture the surface closed; call it $P^* = X^* \cup L^*$. By Lemma 2.6, the cut locus $C(x_{i+1}, P^*)$ is precisely the path $\rho_{i,i+2}$. Therefore, by Lemma 2.2, P^* is a doubly-covered convex polygon, so all angles at x_j are equal above in L^* and below in X^*. In particular, $\theta^{X^*}_{i+1} = \theta^{L^*}_{i+1}$. Now, because Δ angle was removed from θ^X_{i+1}, $\theta^{X^*}_{i+1} = \theta^X_{i+1} - \Delta$. Because α_{i+1} was removed from θ^L_{i+1}, $\theta^{L^*}_{i+1} = \theta^L_{i+1} - \alpha_{i+1}$. Therefore,

$$\theta^L_{i+1} - \alpha_{i+1} = \theta^X_{i+1} - \Delta \leq \theta^X_{i+1} \tag{4.1}$$

$$\alpha_{i+1} \geq \theta^L_{i+1} - \theta^X_{i+1} \tag{4.2}$$

which is Claim (2).

It remains to show that D_{i+1} is in L_i rather than in X_i. Suppose to the contrary that all the angle removal was in X_i. Then $\theta^L_{i+1} = \theta^X_{i+1} - \Delta - \alpha_{i+1}$. So $\theta^L_{i+1} < \theta^X_{i+1}$, which is not possible for P a pyramid. This completes the proof of Claim (2) and the lemma. □

4.3.2 General case

Lemma 4.3 is special in that P sits over a base X. In the general situation, X is the intersection of Q with the truncating slice plane Π, but X is not a face of Q. Rather in general, X is inside Q, the "top" of the portion Q' of Q below Π.

Lemma 4.4. *Let Q be obtained from P by truncating vertex v. Then, if v has degree-k, Q may be obtained from P by $k-1$ tailoring steps, each the excision of a digon surrounding one vertex.*

Proof. Here we argue that the general case is in some sense no different than the special case of P a pyramid just established in Lemma 4.3. In fact, the exact same digon excisions suffice to tailor P to Q.

First we establish additional notation. Let Π be the plane slicing off v above Π, and let $X = \Pi \cap P$. Let the "bottom" part of P be Q', with the final polyhedron $Q = Q' \cup X$. We continue to use L to denote the portion of P above Π, so $P = Q' \cup L$. After removal of digons at $x_1, x_2, \ldots, x_i$, we have $P_i = Q'_i \cup L_i$.

Below it will be important to distinguish between the three-dimensional extrinsic shape of Q'_i and its intrinsic structure determined by the gluings that satisfy Alexandrov's Gluing Theorem. We will use $\bar{Q}'_i$ for the embedding in $\mathbb{R}^3$ and Q'_i for the intrinsic surface, and we will similarly distinguish between $\bar{X}_i$ and X_i. Note that we can no longer assume that $C(x_{i+1}, P_i) \subset L_i$, for the cut locus could extend into Q' (whereas it could not extend into X in Lemma 4.3).

It suffices to show by induction that, on P_i, the following statements hold:
(a) The shortest path γ_{i+1} joining x_{i+1} to y_i is included in L_i.
(b) The ramification point a_{i+1} is still on L_i.
(c) The x_{i+1} angle α_{i+1} of the digon $D_{i+1} = D(x_{i+1}, a_{i+1})$ is larger than or equal to $\omega_Q(x_{i+1}) - \omega_P(x_{i+1})$ (and so sufficient to reduce the curvature to $\omega_Q(x_{i+1})$).

To see (a), assume, on the contrary, that γ_{i+1} intersects Q'_i. Assume, for the simplicity of the exposition, that γ_{i+1} enters Q'_i only once, at x_{i+1}, and exits Q'_i at $p \in \partial X$. Let γ'_{i+1} denote the part of γ_{i+1} between x_{i+1} and p.

We now check (a) for $i = 0$. $Q = Q' \cup X$ and X is planar, hence, because the orthogonal projection of any rectifiable curve onto a plane shortens or leaves its length the same, γ'_1 is longer than or has the same length as its projection γ''_1 onto X. So p is a cut point of x_1 along γ_1, contradicting the extension of γ_1 as a geodesic segment beyond p.

By the induction assumption, all the digon excisions occur on L_i; Q'_i is unchanged. Nevertheless, as part of P_i, neither $\bar{Q}'_i$ nor $\bar{X}_i$ is (in general) congruent to the original $\bar{Q}$ and $\bar{X}$. However, if we consider Q'_i and X_i separate from P_i, we can reshape them so that $\bar{Q}'_i = \bar{Q}$ and $\bar{X}_i = \bar{X}$, precisely because they have not changed. Then $\bar{X}$ is planar and the projection argument used for $i = 0$ works for all i.

Next we check (b) and (c) for $i = 0$. Consider, as in Lemma 4.3, the digon $D_1 = (x_1, y_1)$ with $y_1 \in C(x_1)$, with again a_1 the first ramification point of $C(x_1)$ beyond v. The direction at v of the edge va_1 is only determined by the geodesic segment from x_1 to v, and hence is not influenced at all by Q', because, by $i = 0$ in (a), that segment lies in L.

The ramification point a_1 is joined to x_1 by three geodesic segments, two of them—say γ_1 and γ_2—included in L. The third geodesic γ_3 starts from x_1 towards Q' and finally enters L to connect to a_1. See Fig. 4.6. Because these three geodesics have the same length, the longer γ_3 is, the longer are γ_1 and γ_2, and therefore more distant is a_1 to v. So a_1 is closest to v, and

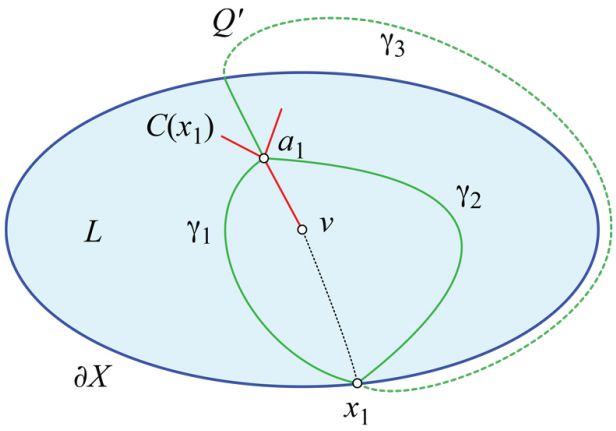

Figure 4.6: Three geodesics to ramification point a_1 on $C(x_1)$. Dashed γ_3 partially in Q'. The same situation holds under the changes: $x_1 \to x_{i+1}$, $a_1 \to a_{i+1}$, $v \to y_i$, $Q' \to Q'_i$, $L \to L_i$.

the segment va_1 shortest, when $Q' = X$ and P is a pyramid. It is when va_1 is shortest that there is the least "room" for y_1 on va_1 to achieve the needed digon angle at x_1, for that angle is largest when y_1 approaches a_1. Therefore, the case when $Q' = X$ and P is a pyramid is the worst case, already settled in Lemma 4.3.

Now we treat the general case for (b) and (c). Again by the induction assumption, all changes to P_i were made on its "upper part" L_i.

Because we ultimately need to reduce L to X, the angle $\omega_Q(x_{i+1}) - \omega_P(x_{i+1})$ necessary to be excised at x_{i+1}, does not depend on Q'_i, only on L_i. Thus the argument used for $i = 0$ carries through. The situation depicted in Fig. 4.6 remains the same, with x_1 replaced by x_{i+1}, v replaced by y_i, and a_1 replaced by a_{i+1}. The ramification point a_{i+1} is closest to y_i, and the segment $y_i a_{i+1}$ shortest, when $Q' = X$ and P is a pyramid. It is when $y_i a_{i+1}$ is shortest that there is the least "room" for y_{i+1} on $y_i a_{i+1}$ to achieve the needed angle excision at x_{i+1}, for that angle is largest when y_{i+1} approaches a_{i+1}. Therefore, the case when $Q' = X$ and P is a pyramid is the worst case, already settled in Lemma 4.3. □

Note that, in the end, the digon removals in Lemma 4.2, and then in Lemma 4.3, also work in the general case, Lemma 4.4.

4.4 Cube/Tetrahedron: Completion

4.4.1 Pyramid Removals

We now return to the cube/tetrahedron example started in Chapter 3. We had reduced the sliced-off portion of the original cube P to four pyramids P_1, P_2, P_3, P_4. Now each of these pyramids needs to be "tailored away" to leave the goal tetrahedron Q. Fig. 4.7 shows that when processed in order, their removal reduces the portion of P "above" Q, leaving the goal tetrahedron Q.

4.4.2 Pyramid Reductions by Tailoring

Now we follow Lemma 4.4 to reduce each of the four pyramids to their bases. We only illustrate this for the first pyramid removed, P_1, with apex p_2 and base $X = p_1p_3p_6$. See Fig. 4.8(a). The base is an equilateral triangle, with edge lengths $\sqrt{2}$, with the apex is connected to the base vertices by unit-length edges. Because the apex p_2 is a cube corner, it has three incident 90° angles, so $\omega(p_2) = 90^\circ$. The three faces incident to p_2 are $45^\circ{-}45^\circ{-}90^\circ$ triangles. The base angles in X are 60°. So each digon must reduce the incident 90° angles by 30° to match 60°. Starting with p_1, the digon geodesics are $\pm 15^\circ$ around the p_1p_2 edge.

The geometry is clearest if we unfold P_1's lateral faces into the plane, as shown Fig. 4.8(b). Excising the first digon and sealing the cut results in a new polyhedron, with the apex p_2 replaced by a new vertex, call it y, of curvature $\omega(y) = 60^\circ$ (because the two digon angles must sum to the 90° curvature at p_2). Unfortunately, we cannot display this new polyhedron because of the difficulty of constructing what AGT guarantees exists.

Lemma 4.4 says that just one more digon needs removal (because p_2 has degree $k = 3$), again $\pm 15^\circ$ this time around the p_3y edge. This removal reduces P_1 to its equilateral triangle base, and, despite the nonconstructive nature of AGT, we know that the full polyhedron is exactly what we illustrated earlier in Fig. 4.7(b).

One further remark on the shape of P_1 after excision of the first digon. If we imagine P_1 standing alone on its base X as illustrated in Fig. 4.8(a), rather than as part of the full cube polyhedron, then it is not difficult to reconstruct the shape of P_1'. It is a flat doubly-covered quadrilateral, with the triangle p_3yp_6 flipped over and joined to the p_3p_6 edge of the equilateral

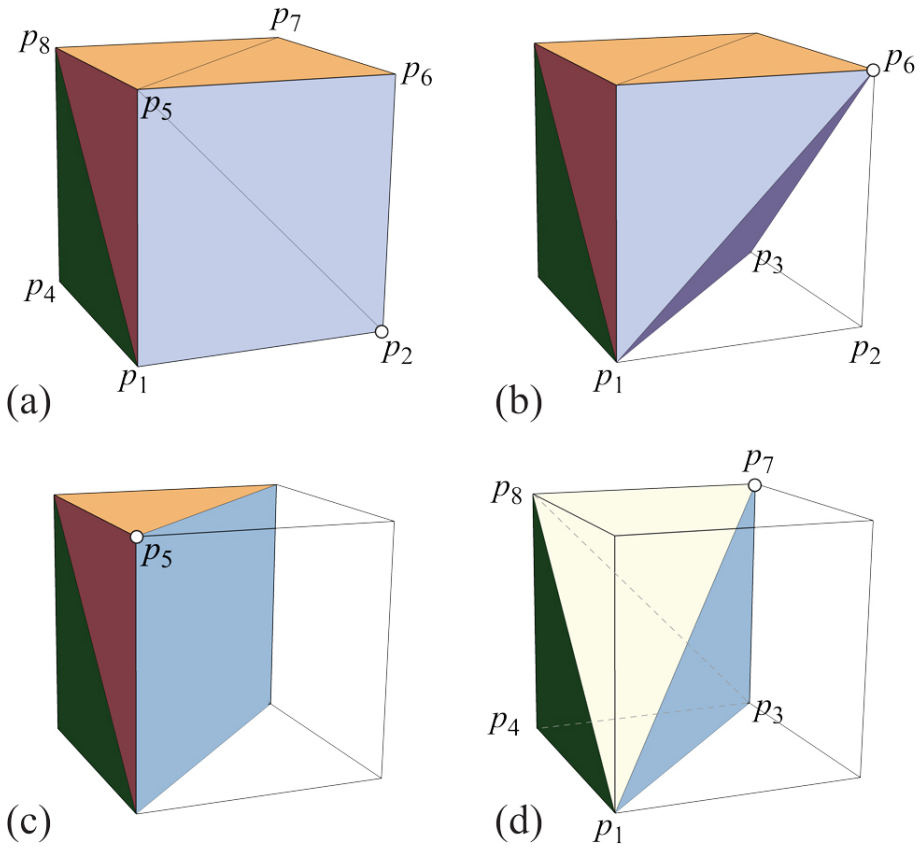

Figure 4.7: Pyramid removals. The apex of the pyramid to be next removed is highlighted. From (a) to (b), the pyramid P_1 with apex p_2 and base $p_1p_3p_6$ is removed. Then P_2 apexed at p_6 is removed, producing (c), then removing P_3 apexed at p_5 leads to (d). The final removal of P_4 apexed at p_7, leaves the tetrahedron $Q = p_1p_3p_4p_8$ previously illustrated in Fig. 3.2(a).

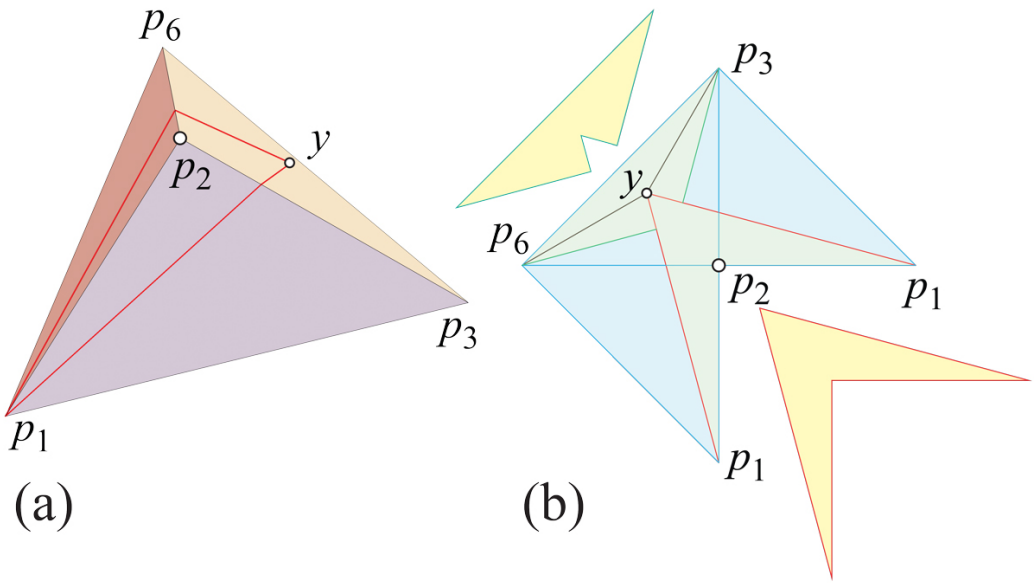

Figure 4.8: (a) Pyramid P_1: apex p_2, base $p_1p_3p_6$. (b) Excision of two digons. Each is bound by geodesics $\pm 15°$ about the edge to p_2. After removing p_2, a new vertex y is created. The removal of y in the second excision flattens the pyramid to its base.

triangle base X. However, with this P_1' piece joined to the full cube, it seems much less straightforward to determine the shape of the full polyhedron.

Each pyramid reduction proceeds in the same manner: $k-1$ digons are excised if the apex has degree-k, and the lateral faces are reduced to the base. So P_2 has apex p_6 and base $p_1p_3p_5p_7$. After removal of three digons, the result is as illustrated in Fig. 4.7(c). P_3's apex p_5 has degree-3, so two digon removals lead to (d). The last pyramid removed, P_4, reduces to the face $p_1p_3p_8$ of Q, completing the tailoring of the cube P to tetrahedron Q.

4.4.3 Seals

After removal of the two digons illustrated in Fig. 4.8(b), P_1 has been reduced to its equilateral triangle base X. Sealing the first digon produces a seal σ_1, which is then clipped to a segment s_1 by the second digon removal, which produces σ_2 along the boundary of X. The seal segments then are as shown in Fig. 4.9(a). Fig. 4.9(b) shows the three seal segments that result by reducing

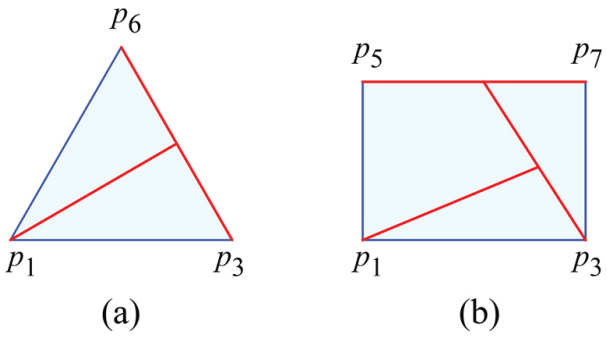

Figure 4.9: (a) Base of P_1. (b) Base of P_2. Seal segments: red.

P_2 to its base. These depictions of the seal graph Σ will be explored in some detail in Chapter 5.

There is an aspect of the seals we are not tracking: As can be seen in Fig. 4.7(b,c), one of the faces of P_2 is the equilateral triangle base of P_1. Since that base is already crossed by seal segment s_1 when P_2 is undergoing digon removal, that segment s_1 will be reflected as a cut in the reduced base of P_2, not depicted in Fig. 4.9(b). We have not attempted to track this complex overlaying of seal cuts in the surface of P. However, we explore seals for one pyramid tailoring in some detail in Chapter 5.

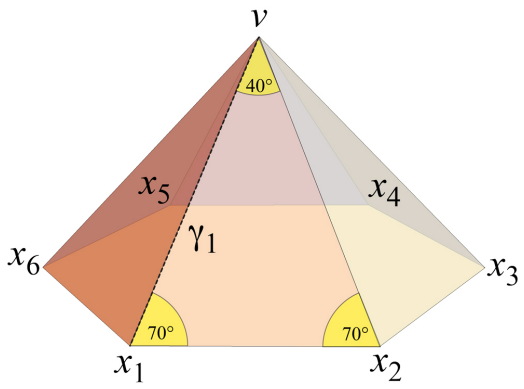

Figure 4.10: Pyramid with regular hexagon base X; all lateral faces congruent triangles.

4.5 Hexagonal Pyramid Example

We now detail a more complex example following Lemma 4.3 to tailor a pyramid P to its base X. We continue to employ the notation used in the lemmas above. The example is shown in Fig. 4.10. X is a regular hexagon, and L consists of $k = 6$ congruent, $70°-70°-40°$ isosceles triangles. The curvature at the apex v is $360° - 6 \cdot 40 = 120°$. The angle at each x_i in X is $120°$ whereas the angle in L is $140°$. So each digon excision must remove $20°$ from x_i. As in the lemmas, we excise the digons in circular order around ∂X.

We display the progress of the excisions on the layout of L in Fig. 4.11(a). Let $v = y_0$ for ease of notation. $D_1 = (x_1, y_1)$ includes the geodesic γ_1 from x_1 to y_0, and locates y_1 on the cut locus segment as described in the lemmas. The digon boundary geodesics each remove $10°$ from the left and right neighborhood of x_1, and meet at y_1 at an angle of $100°$, which is then the curvature at the new vertex: $\omega(y_1) = 100°$. Notice that the digon angles $20° + 100°$ match the curvature $\omega(v) = 120°$ removed, as they must to satisfy Gauss-Bonnet.

One should imagine that D_1 is sutured closed in Fig. 4.11(a), producing L_1, before constructing $D_2(x_2, y_2)$. Let σ_i be the geodesic on L_i that results from sealing D_i closed; σ_i is like a "scar" from the excision. Notice that one of the geodesics bounding D_2 crosses σ_1.

This pattern continues as all $k-1 = 5$ digons are removed, each time replacing vertex y_{i-1} with y_i, flattening the curvature $\omega(y_i) = \omega(y_{i-1}) - 20°$.

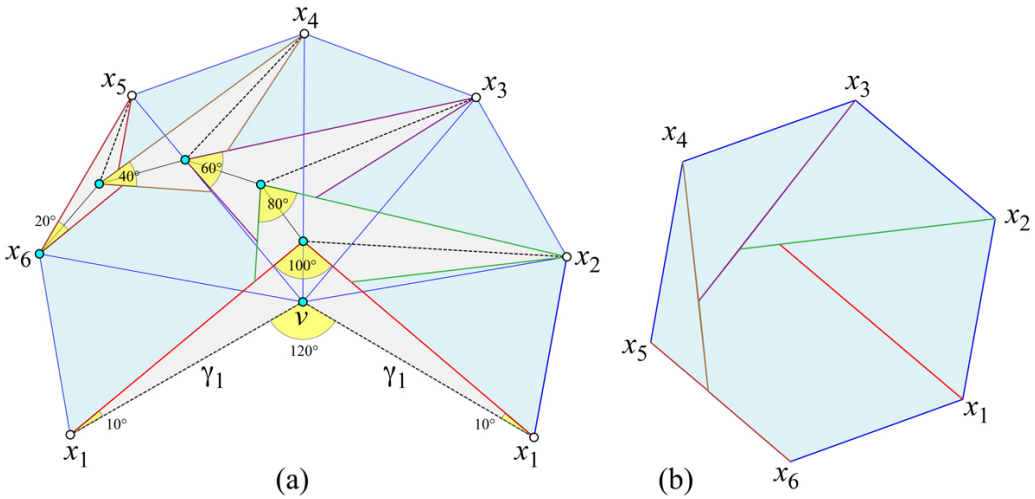

Figure 4.11: (a) Cone L flattened; $\omega(v) = 120°$. Digons $D_i = (x_i, y_{i-1})$ shaded. Dashed lines are geodesics γ_i from x_i to the vertex y_{i-1} (with $y_0 = v$). (b) After excising all digons $D_1, \ldots, D_5$, L_5 is isometric to X. Seals $\tilde{\sigma}_i$ are marked.

Finally, after $D_5(x_5, y_5)$ is removed, y_5 is coincident with x_6. No further digon removal is needed, because D_5 removed 20° from x_6. So now each angle in $L_k = L_5$ at all k=6 vertices is 120°, and L_5 is isometric to a flat regular hexagon, i.e., to X.

This final hexagon is shown in Fig. 4.11(b). The images $\tilde{\sigma}_i$ of the seals are in general clipped versions of σ_i on L_i, clipped by subsequent digon removals. The particular circular order of digon removal followed in this example and the lemmas result in a spiral pattern formed by $\tilde{\sigma}_i$. Other excision orderings, which ultimately would result in the same flat L_{k-1} (effectively proved in Lemmas 4.2–4.4) would create different seal patterns. As mentioned earlier, we study seals in detail in Chapter 5.

4.6 Tailoring is finer than sculpting

In this section we reach one of our main results, Theorem 4.6, which says, roughly, that any polyhedron Q that can be obtained by sculpting P can be obtained by tailoring P. Moreover, Lemma 4.5 shows that polyhedra can be obtained by tailoring that cannot be obtained by sculpting. So, in a sense,

tailoring is finer than sculpting.

Lemma 4.5. *There are shapes P and sequences of tailorings of P that result in polyhedra not achievable by sculpting.*

Proof. We first tailor a regular tetrahedron T as in Example 1.1, resulting in the kite K in Fig. 1.2(b). We now show that K cannot fit inside T, so it couldn't have been sculpted from T. Assume T has edge-length 1. Then its extrinsic diameter is 1 and its intrinsic diameter is $2/\sqrt{3}$ (see, e.g., Theorem 3.1 in [Rou03]). Moreover, the extrinsic diameter of K is precisely the intrinsic diameter of T, and so it cannot fit inside T.

Next we construct a non-degenerate example, a modification of the previous one. Consider a non-degenerate pentahedron F close enough to $K = oac_db$ in Fig. 1.2(b). For example, it could have two vertices close to the vertex a of K. Insert into F the removed digon from T; this is not affected by the new vertex, because it does not interfere with the geodesic segment from c_d to o. We arrive at some surface P close enough to the original tetrahedron T. Therefore, the intrinsic and extrinsic diameters of P and F are close enough to those of K and T, respectively, and the above inequality between the extrinsic diameters of P and F still holds, because of the "close enough" assumption. □

Theorem 4.6. *Let P be a convex polyhedron, and $Q \subset P$ a convex polyhedron resulting from repeated slicing of P with planes. Then Q can also be obtained from P by tailoring. Consequently, for any given convex polyhedra P and Q, one can tailor P "via sculpting" to obtain any homothetic copy of Q inside P.*

Proof. Lemma 4.1 established that one slice leads to domes, Theorem 3.2 showed that each dome leads to pyramids, and Lemma 4.4 showed that each pyramid can be reduced to its base by tailoring. Since this holds for one slice, it immediately follows that it holds for arbitrary slicing.

Concerning the domes → pyramids step, we note that the property that each pyramid P_i has a common edge with X, established in Theorem 3.2, allows reduction of the pyramids in the order that they are obtained in that theorem. After each reduction, the result is still a g-dome, allowing iteration until the original g-dome is reduced to its base.

For the homothet-copy claim of the lemma, shrink Q by a dilation until it fits inside P, and then apply the reductions. □

As we mentioned in the Preface, an informal consequence of this theorem is that P can be "whittled" to e.g., a sphere S:

Corollary 4.7. *For any convex polyhedron P and any convex surface S, one can tailor P to approximate a homothetic copy of S.*

Proof. Bring a homothetic copy of S inside P. Perform a series of slicings of P with planes tangent to S. Any degree of approximation desired can be achieved by increasing the number of plane splicings. Call the result of these slicings Q. Now apply Theorem 4.6. □

Despite this corollary, it does not seem possible to accomplish the reverse: to start with a strictly convex surface and tailor it to a polyhedron. However, one can of course sculpt a surface to a polyhedron.

Chapter 5

Pyramid Seal Graph

5.1 Pyramid Digon Removal

As we have seen in Theorem 4.6, tailoring by tracking sculpting ultimately relies on digon removal reducing pyramids to their bases. We have illustrated such reductions for a few low-degree pyramids in Figs. 4.2, 4.8, and 4.11. In the latter two figures, we displayed the seals $\tilde{\sigma}_i$ on the base X after reduction. It however remains difficult to grasp in detail the digon-removal process for a pyramid P, for at least three reasons:

1. After removing the first digon D_1, P_1 is (in general) no longer a pyramid. The difficulty of computationally realizing the subsequent intermediate shapes P_i, guaranteed by AGT, makes it hard to envision the process.
2. The seals σ_i that result from closing digon D_i cross and clip one another.
3. The process depends on the order in which the digons are removed.

We will continue to circumvent this last difficulty by only studying the natural order of digon removal, anchored at $x_1, x_2, x_3, \ldots$ in counterclockwise order around ∂X. In this section, we introduce a different way to view digon removal that in some sense skirts the first two difficulties.

The process is complex enough to require somewhat extensive notation, which we list in two parts before turning to examples.

5.1.1 Notation I

- P: a pyramid, n vertices around base X.
- Base vertices $x_1, x_2, \ldots, x_n$, in counterclockwise order around ∂X.
- Apex y_0 of degree-n.
- y_i: apex after removing digon D_i.
- D_i: digon from x_i to y_i, surrounding y_{i-1}.
- L_i: (the remaining of the) lateral faces after removing digon D_i. L_0: initial faces before any removals. The apex of L_i is y_i.
- P_i: The polyhedron $L_i \cup X$, guaranteed by AGT. $P = P_0$ is the original, before any digon removal.

We should emphasize that although $P_i = L_i \cup X$, in general X will not be planar in P_i as it is in P_0, and so P_i is not a pyramid, as previously mentioned.

5.2 Cone Viewpoint

Although we do not know the structure of P_i, except at the beginning ($i = 0$) and end ($i = n-1$), when it is P and doubly-covered X respectively, we do know that the lateral faces L_i contain only one vertex, y_i, hence they form a subset of a cone apexed at y_i. Any cone can be cut open along a generator (a ray on the cone from the apex) and laid flat in the plane. Such a layout will have an angle gap of $\omega(y_i)$ at the apex. It is especially useful to cut along $x_i y_{i-1}$ before removing digon D_i. We will provide several examples, after presenting more notation. We emphasize the indices $i-1$, i, and $i+1$ in the following, in an attempt to avoid confusion.

5.2.1 Notation II

- $\bar{L}_{i-1}$: Unfolding of L_{i-1} cut open along $x_i y_{i-1}$. So after removing and sealing digon D_{i-1}, but not yet D_i.

- $\bar{L}_i$: Unfolding of L_i cut open along $x_{i+1}y_i$. So after removing and sealing digon D_i, but not yet D_{i+1}. So $\bar{L}_0$ is L_0 cut open along x_1y_0, and $\bar{L}_1$ is L_1 cut open along x_2y_1.

- $\sigma_i = x_iy_i$ is the i-th seal after suturing closed the digon D_i. We view the seals as directed from x_i to y_i, so that they have distinguished left and right sides. This direction is only used in the proofs; the seals are illustrated as undirected segments in several figures. When the direction plays a role, we use boldface: $\boldsymbol{\sigma}_i$.

- Σ_i: the seal graph after removing digon D_i. $\Sigma_0 = \varnothing$, and $\Sigma_1 = x_1y_1$.

- $s_j \subseteq \sigma_j$, $1 \leq j \leq i$, is the possibly truncated seal segment in Σ_i, on the surface P_j. So, after possibly other truncations, we reach $\tilde{\sigma}_j$, hence the informal inclusion $\tilde{\sigma}_i \subseteq s_j \subseteq \sigma_j$; "informal" because those geodesic segments live in different spaces.

- S_i is the subset of L_i bounded by x_1y_i and x_iy_i, the *sealed region* which we will later prove contains Σ_i.

5.3 Examples

We start with P_1, previously displayed in Fig. 4.8. X is an equilateral triangle, with the apex centered above its centroid. Fig. 5.1 shows the removal of $n-1 = 2$ digons D_1, D_2 that reduce L_0 to the equilateral triangle base X. Images are repeated so that in one row the transition from L_{i-1} to L_i by removal of D_i is evident.

Next is the more complicated P_2 in Fig. 5.2. Here X is a rectangle, and three digons are removed, D_1, D_2, D_3, before reaching X. We should emphasize that any one of these figures could be cut out and closed to a cone. This cone would not be rigid, and its boundary ∂X would not (in general) be planar, as we mentioned earlier.

As a last example, we extract one row illustrating removal of D_5 from a pyramid of degree-12 in Fig. 5.3. We will refer to this figure subsequently.

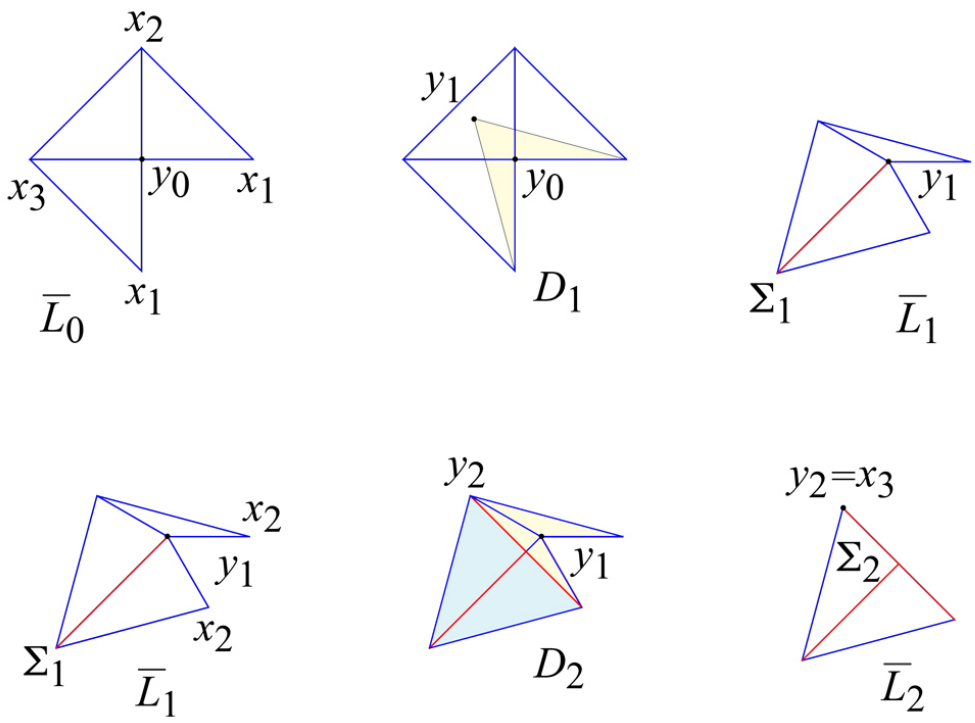

Figure 5.1: P_1. X is an equilateral triangle. Cf. Fig. 4.9(a).

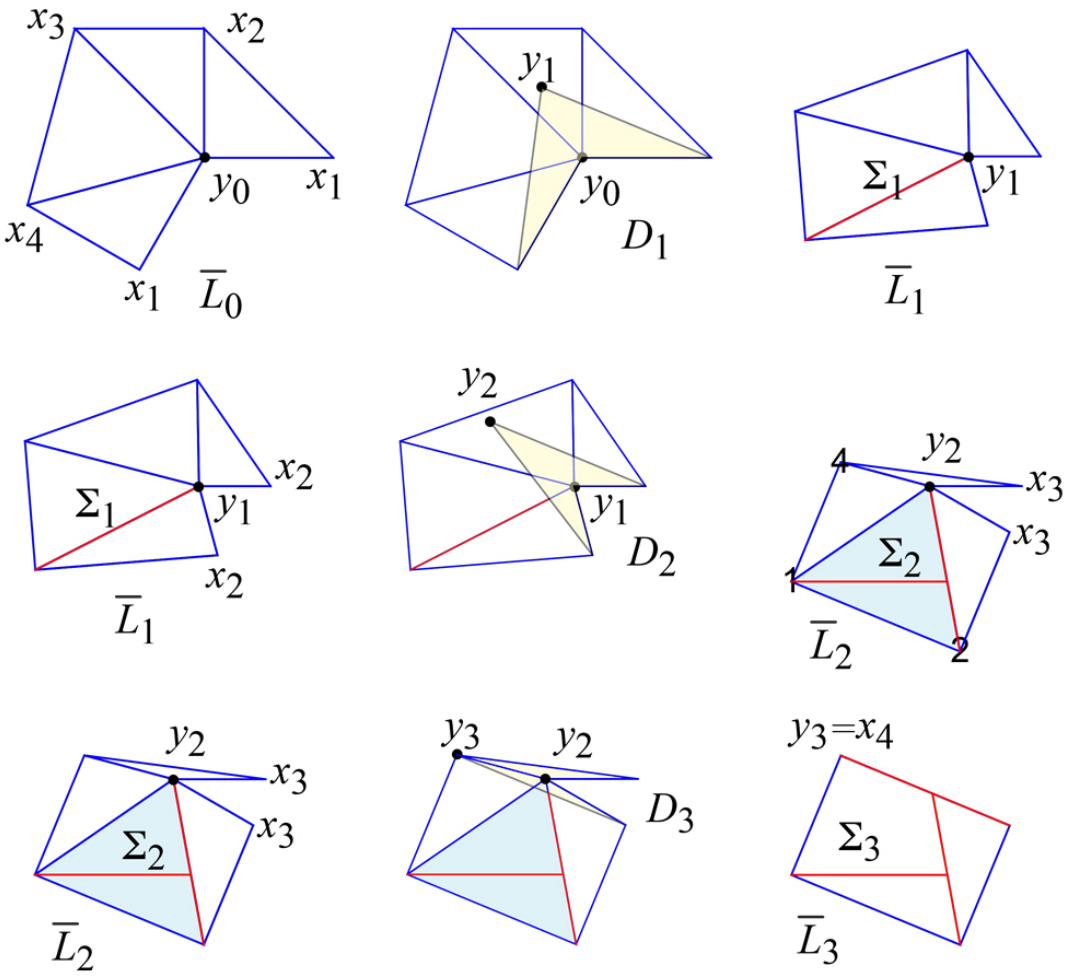

Figure 5.2: P_2. X is a rectangle. Cf. Fig. 4.9(b).

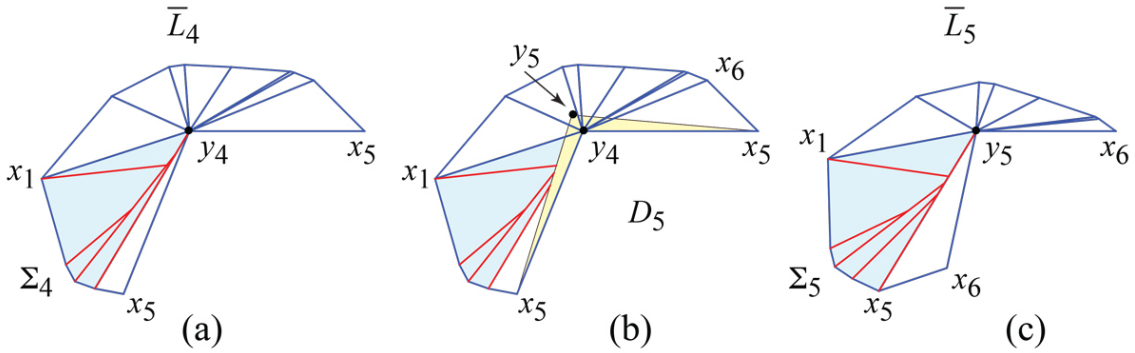

Figure 5.3: Pyramid of n=12 base vertices x_i. (a) $\bar{L}_4$. S_4 (blue). (b) Digon D_5 (yellow). (c) $\bar{L}_5$. Σ_i in red.

5.4 Preliminary Lemmas

The viewpoint just described is simple enough to be implemented, and to allow us to construct the seal graph Σ for any pyramid P. See ahead to Fig. 5.8 for examples. This leads to Theorem 5.3: Σ is a tree. The proof of this claim is somewhat intricate, and presented in Section 5.5. That proof requires two lemmas, both involving the structure of the cut locus, which we present first. The reader might skip these proofs until later.

Lemma 5.1. *After removing digons $D_1, \ldots, D_i$ and closing seals $\sigma_1, \ldots, \sigma_i$, $C(x_{i+1}, P_i)$ includes the path $x_1, \ldots, x_i$, with each node in that path of degree-2.*

For example, in Fig. 5.3, $i = 4$ and $C(x_5, P_4)$ includes x_1, x_2, x_3, x_4.

Proof. We start with the leaf $x_i \in C(x_{i+1})$, and argue that $(x_i, x_{i-1}, \ldots, x_1)$ is a path ρ in $C(x_{i+1})$, i.e., that every point along ρ is of degree ≤ 2. The proof uses techniques detailed in the proof of Lemma 4.3. In particular, Fig. 5.4 below depicts the situation abstractly, similar to Fig. 4.4 in Lemma 4.3.

First, x_i is of degree-1 in $C(x_{i+1})$: The pyramid edge $x_{i+1}x_i$ is the shortest geodesic, unaffected by the digon removals up to D_i. An edge e_i of $C(x_{i+1})$ starts at x_i, and because of the equal angles above on L_i and below on X at x_i, and because e_i is bisecting, initially it starts along the geodesic x_ix_{i-1}. It then either continues to x_{i-1}, or reaches a ramification point.

Suppose the path continues $x_i, x_{i-1}, \ldots, x_j$, but then reaches a ramification point r on $x_j x_{j-1}$. Let ρ_j denote this path up to r. We now analyze this situation and show it is contradictory. Consult Fig. 5.4 throughout.

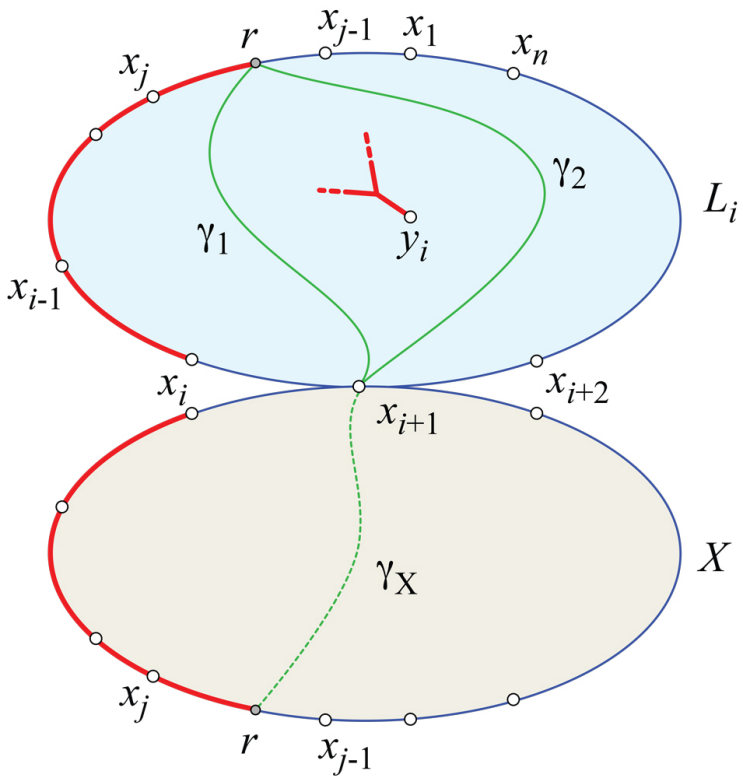

Figure 5.4: Abstract depiction of $x_i, x_{i-1}, \ldots, x_1$ path in $C(x_{i+1})$ (red). X is flipped below; L_i above.

The possible geodesics from x_{i+1} to r are: γ_X on X below, γ_1 and γ_2 on L_i above, and possibly γ_3 both above and below. Note that there can only be the two γ_1 and γ_2 because there is just one vertex y_i on L_i.

First consider γ_1, which, together with γ_X, encloses ρ_j. The planar convex chain along ρ_j, $x_i, x_{i-1}, \ldots, x_j, r$ is congruent above and below, because the angles above and below are equal after digon removals. Thus the chords connecting the endpoints of the chains are equal, and so $|\gamma_1| = |\gamma_X|$.

Next consider γ_2, which, together with γ_X, encloses $x_{i+2}, \ldots, x_n, x_1, \ldots, x_{j-1}$. We again compare the planar convex chain with angles below on X to the chain with angles above on L_i. Because some of the angles above are strictly

larger than their counterparts below, we can apply Cauchy's Arm Lemma just as we did in the proof of Lemma 4.3(Claim (1)) to conclude that $|\gamma_2| > |\gamma_X|$. Therefore γ_2 cannot add to the degree of r in $C(x_{i+1})$.

Finally, a geodesic γ_3 that lies on both L_i and X must have a portion completely above on L_i, to which we may apply the same arm-lemma argument to conclude that $|\gamma_3| > |\gamma_X|$.

Therefore, r is in fact of degree 2, it is not a ramification point, and ρ extends from x_i to x_1 as claimed. □

Lemma 5.2. *Assume the digons $D_1, \ldots, D_i$ have been removed and S_{i-1} is the sealed region containing $\sigma_1, \ldots, \sigma_{i-1}$ (this inclusion will be proven later). Let a_i be the first ramification point of $C(x_i, P_{i-1})$, on the segment $y_{i-1}a_i$. Then $a_i \notin S_{i-1}$.*

First we illustrate the claim of the lemma with the example from Fig. 5.3(b), repeated as Fig. 5.5(a). In order for the lemma to be false, the situation instead must appear as in (b) of the figure, with $a_5 \in S_4$.

Proof. For the purposes of contradiction, consider the situation depicted abstractly in Fig. 5.5(c), with $a_i \in S_{i-1}$. By Lemma 5.1, $x_1, \ldots, x_{i-1}$ is a path of degree-2 nodes in $C(x_i, P_{i-1}) = C(x_i)$. The edge $y_{i-1}a_i$ of $C(x_i)$ bisects the angle formed by the two images of x_i. From a_i, $C(x_i)$ must contain a path ρ that connects to x_i. Because $a_i \in S_{i-1}$, ρ must start with an edge to the right of the line containing $y_{i-1}a_i$.

Note that y_{i-1} is the only vertex on L_{i-1}, so a_i is not a vertex, and therefore has a total surface angle $\theta_{a_i} = 2\pi$. Lemma 2.3 requires a strictly leftward branch at a_i, left of the line containing $y_{i-1}a_i$. Call the path continuing this left branch λ. This path λ is "trapped": It cannot terminate in the interior of S_{i-1} because there are no vertices in that region. It cannot connect to any one of $x_1, \ldots, x_{i-1}$ because those nodes are degree-2. If λ crossed ρ it would form a cycle. Therefore we have reached a contradiction, and so $a_i \notin S_{i-1}$. □

5.5 Pyramid Seal Graph is a Tree

Roughly, we prove in this section that the complete seal graph $\Sigma = \Sigma_n$, and intermediate seal graphs Σ_i, have the structure of a spiraling tree. Spiral trees—slit trees rather than seal trees—will play a significant role in Part II.

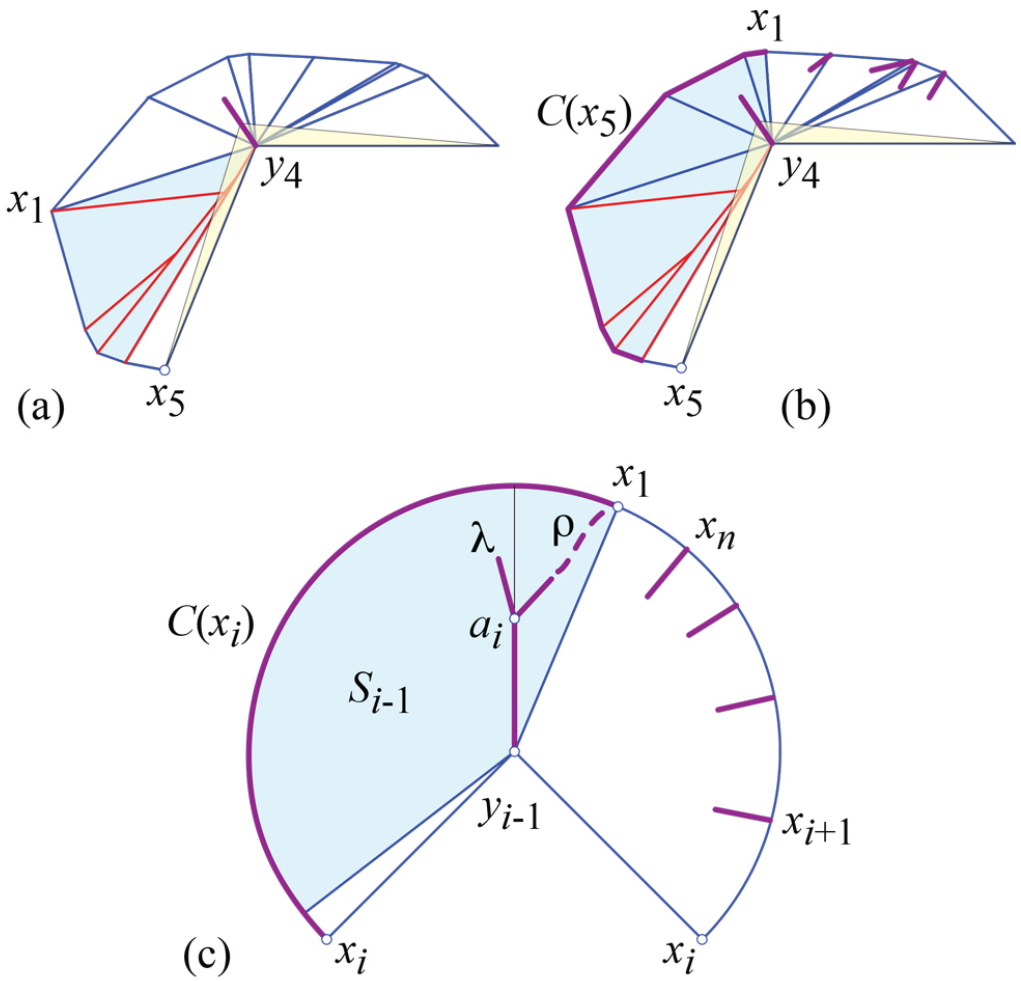

Figure 5.5: (a) Fig. 5.3(b). (b) How a counterexample might appear. (c) Abstract counterexample. Portions of $C(x_i)$ purple.

We repeat some notation here for convenience. The *sealed region* S_i is bounded by the segments x_1y_i, x_iy_i, and the portion of ∂X from x_1 to x_i. Also recall that each seal $\boldsymbol{\sigma}_j = (x_j, y_j)$ is directed from x_j to y_j. The seal graph Σ_i is composed of segments s_j, each a subsegment of σ_j.

Theorem 5.3. *The seal graph Σ_i, after removal of digons $D_1, \ldots, D_i$ in counterclockwise order, has the following properties:*

(1) $\Sigma_i \subset S_i$.

(2) Σ_i *is a directed tree with root* y_i.

(3) Each segment s_j *of* Σ_i *is a (possibly truncated) seal* $\sigma_j = x_jy_j$ *that remains anchored on its* x_j *endpoint; i.e., the truncation is on the* y_i*-end.*

(4) Each leaf x_j *is the start of a directed, convex path* $\boldsymbol{\pi}_j$ *to the root* y_i.

(5) The edges of $\boldsymbol{\pi}_j$ *are portions of seal segments of increasing indices.*

(6) Along $\boldsymbol{\pi}_j$*, lower-indexed edges terminate from the left on higher-indexed edges.*

(7) The last seal, $\boldsymbol{\sigma}_i = x_iy_i$*, the* root segment *of* Σ_i*, has no segments of* Σ_i *incident to its right side.*

The last segment of the complete seal graph $\Sigma = \Sigma_n$ coincides with the edge $x_{n-1}x_n$ of ∂X.

We again refer to Fig. 5.3(ab). When $i = 4$, Σ_4 satisfies the properties, and we seek to re-establish the properties for Σ_5 in (c) of the figure. We enlarge (b) of that figure in Fig. 5.6 to track in this proof. It may also help to consult the complete seal graphs in Fig. 5.8.

Proof. **Induction Basis.** These claims are trivially true for $i = 0$ because $\Sigma_0 = \varnothing$. So assume $i = 1$. L_1 has just had the digon $D_1 = x_1y_1$ excised around y_0. $\bar{L}_1$ is then cut open along x_2y_1. Σ_1 is the single segment $\sigma_1 = x_1y_1$, and all properties are easily verified.

Induction Hypothesis. Assume that all properties hold for Σ_i. The removal of digon $D_{i+1} = x_{i+1}y_{i+1}$ leads to Σ_{i+1}. Now we establish the properties for Σ_{i+1}.

(1) $\Sigma_{i+1} \subset S_{i+1}$. S_{i+1} grows on both sides: by the triangle $x_1y_iy_{i+1}$ (clockwise in Fig. 5.6), and to $x_{i+1}y_{i+1}$ (counterclockwise in the figure). By

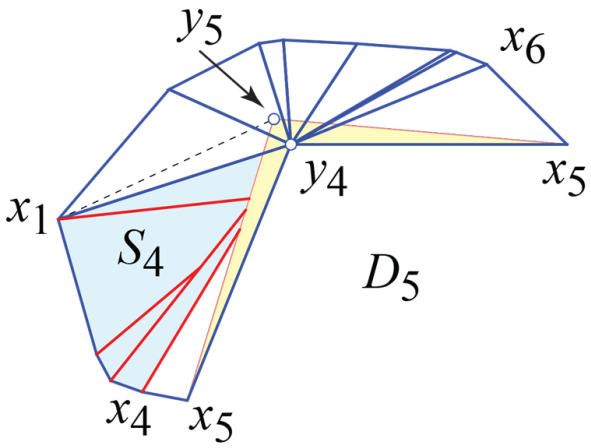

Figure 5.6: $i = 4$, before removal of digon $D_{i+1} = D_5$. Detail from Fig. 5.3(b).

Lemma 5.2, we know that $a_{i+1} \notin S_i$, and because y_{i+1} is on the $y_i a_{i+1}$ segment of $C(x_{i+1})$, we know that $y_{i+1} \notin S_i$. Therefore the triangle $x_1 y_i y_{i+1}$ does in fact grow S_i on the x_1-end. On the x_{i+1}-end, the new seal $\boldsymbol{\sigma}_{i+1} = x_{i+1}y_{i+1}$ is incorporated, and all seal segments right of $\boldsymbol{\sigma}_{i+1}$ are clipped by the removal of digon D_{i+1}. Therefore indeed S_{i+1} expands to include all of Σ_{i+1}.

(2) Σ_{i+1} is a directed tree with root y_{i+1}. We know that Σ_i is a directed tree with root y_i. We first argue that the seal $\boldsymbol{\sigma}_{i+1}$ intersects the segments of Σ_i from right-to-left, re-establishing property (7). Let γ_1 and γ_2 be the left and right geodesics of digon $D_{i+1} = x_{i+1}y_{i+1}$. γ_1 starts within the triangle $\triangle = x_i x_{i+1} y_i$, at an angle $(\alpha_{i+1} + \beta_{i+1})/2$ left of the $x_{i+1}y_i$ edge of $\triangle$. Therefore γ_1 cuts into S_i through the $x_i y_i$ edge of $\triangle$. The segments of Σ_i crossed and clipped by γ_1 are crossed from right-to-left. And since $i + 1$ is the highest indexed seal, the segments of Σ_i that meet $\boldsymbol{\sigma}_{i+1}$ satisfy property (7): lower-indexed segments terminate on the left of $\boldsymbol{\sigma}_{i+1}$.[1] Henceforth, we use σ_{i+1} as determined by γ_1.

By property (8), $\boldsymbol{\sigma}_i = x_i y_i$ is the root segment of Σ_i, which we now

[1]We should mention that this property, that $\boldsymbol{\sigma}_{i+1}$ crosses segments right-to-left, is dependent on the counterclockwise ordering of digon removal. If after removing digon D_i, we next removed some D_j with $j > i+1$, it could be that it is γ_2 rather than γ_1 that clips Σ_i. This would result in seal graphs with a different structure.

know is crossed by $\boldsymbol{\sigma}_{i+1}$ right-to-left, say crossing at point p. Because $y_{i+1} \notin S_i$ by Lemma 5.2, $\boldsymbol{\sigma}_{i+1}$ has the root y_i of Σ_i to its right, and all the leaves x_j of Σ_i to its left.

Now suppose Σ_{i+1} has an undirected cycle Φ; see Fig. 5.7. Because Σ_i has no cycle, any cycle in Σ_{i+1} must have one edge a subsegment of σ_{i+1}. Because all of Σ_i right of $\boldsymbol{\sigma}_{i+1}$ is removed, the cycle Φ must "rest on" the left side of $\boldsymbol{\sigma}_{i+1}$. It is clear Φ cannot rest on the $x_{i+1}p$ portion of $\boldsymbol{\sigma}_{i+1}$. So Φ must rest on the portion of $\boldsymbol{\sigma}_{i+1}$ left of $\boldsymbol{\sigma}_i$, as illustrated. But then, imagining removing $\boldsymbol{\sigma}_{i+1}$, Φ must have formed a cycle $\Phi' \supset \Phi$ in Σ_i, a contradiction. Therefore Σ_{i+1} is indeed a tree.

The remaining properties are now easily established.

(3) Because $\boldsymbol{\sigma}_{i+1}$ clips segments of Σ_i to its right, each segment $s_j \in \Sigma_i$ remains anchored on x_j. And the new segment $\boldsymbol{\sigma}_{i+1}$ is anchored on x_{i+1}.

(4) The directed, convex path $\boldsymbol{\pi}_j$ remains, but now may by shortened where it joins with $\boldsymbol{\sigma}_{i+1}$.

(5) The segments along $\boldsymbol{\pi}_j$ have increasing indices, possibly now including $i+1$.

(6) We earlier established that lower-indexed segments of Σ_{i+1} terminate from the left on higher-indexed segments, possibly now including $\boldsymbol{\sigma}_{i+1}$.

(7) The last and new seal $\boldsymbol{\sigma}_{i+1} = x_{i+1}y_{i+1}$ becomes the root segment incident to the root y_{i+1}, and has no segments incident to its right side.

Finally, it is a consequence of Lemma 4.2 and the rigidity Theorem 2.9 that $y_{n-1} = x_n$ so that the $(n-1)$-st seal coincides with the edge $x_{n-1}x_n$. □

5.5.1 Other Digon Orderings

The proof of Theorem 5.3 depends on removing the digons in the order $x_1, x_2, \ldots, x_{n-1}$ around ∂X. This affects Lemma 5.1's conclusion that $x_1, \ldots, x_i$ is a path in the cut locus, which then affects Lemma 5.2's conclusion that the ramification point a_{i+1} is outside the sealed region S_i. In addition, which side of the removal of digon D_{i+1} clips Σ_i is affected by knowing that x_{i+1} is adjacent to x_i on ∂X. All of these consideration affect the structure of

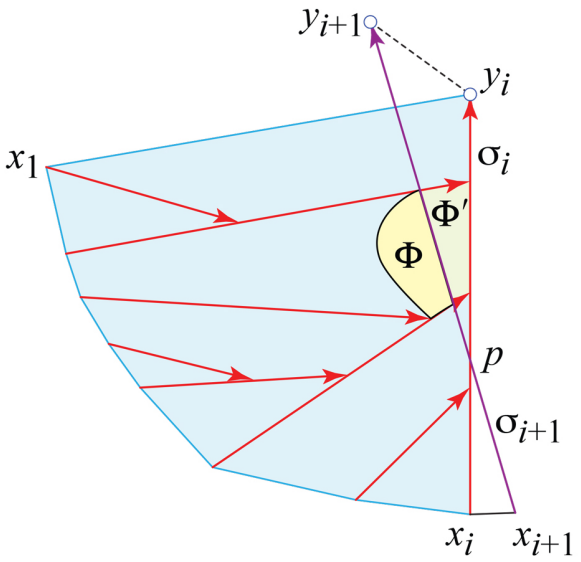

Figure 5.7: Cycle Φ is not possible: removing σ_{i+1} extends Φ to Φ'.

the seal graph. We leave the question of whether the seal graph is a tree for other orderings of digon removals to Open Problem 18.3.

Toward this open problem, we only show here, with the next result, what a degree-4 vertex in Σ must look like.

In the proof below, we simplify the notation of $\tilde{\sigma}_i$ on the base to just σ_i, keeping in mind that σ_i and σ_j formally live in different spaces,.

Lemma 5.4. *If a seal graph Σ for a pyramid L has a degree-4 vertex z, then there exist $i < j < k$ such that σ_i and σ_j end at z, and σ_k passes beyond z. Moreover, the digon excision order is D_k immediately after D_j immediately after D_i.*

Proof. Consider a common point z of σ_i and σ_j, with $i < j$. We may assume $j > i + 1$, since otherwise $\deg z = 3$ in Σ.

Assume first that no other σ_k passes through z. Assume that $\deg z = 4$ in Σ. This implies that the digon D_j crosses σ_i and, since σ_i remains a geodesic after the excision of D_j, σ_i must be orthogonal to both geodesics bounding D_j. Therefore, σ_i creates with those two geodesics two geodesic triangles, both of positive curvature. So each such triangle contains a vertex inside, contradicting that D_j itself contains only one vertex.

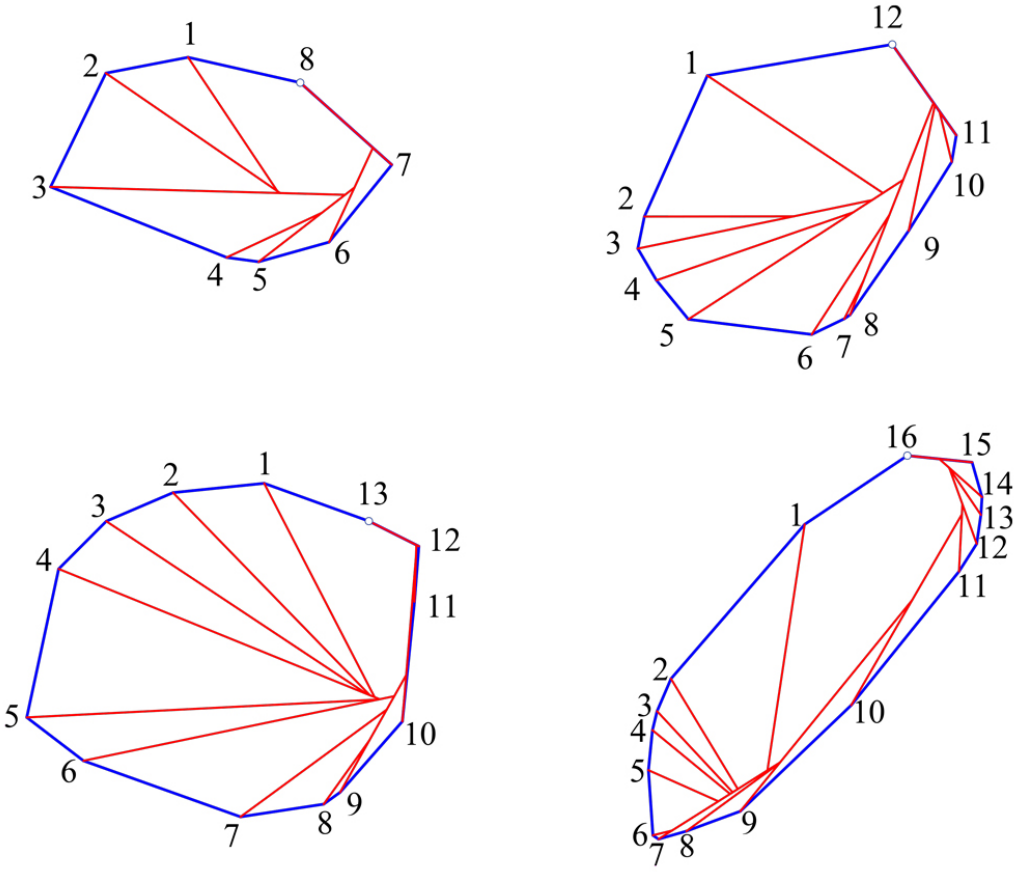

Figure 5.8: Four complete seal graphs.

Assume now that σ_i ends at z, as does σ_k for some $k \notin \{i, j\}$. Notice that we cannot have four σs ending at z, because for the last one arriving—say σ_k—σ_k would create a vertex at z, which will be excised by the digon D_{k+1}, breaking that degree-4 configuration at z.

So we may assume that z belongs only to σ_i, σ_j, and σ_k. We may further assume, without loss of generality, that $i < j < k$.

Notice that both σ_i and σ_j end at z and σ_k does not, because otherwise σ_k would create a vertex on σ_i (or on σ_j) which would be excised by the digon D_{k+1}, breaking that degree-4 configuration at z.

The digon excision order follows: only D_j could surround y_i, so its excision was just after D_i, and similarly for D_j. □

For a particular digon-removal ordering, consider the inverse image of Σ on L, and denote it by χ. χ is a simple geodesic polygon surrounding v. In Fig. 4.11(a), χ is the boundary of the gray region, effectively the union of the digons (but recall that the digons D_i live on different surfaces P_{i-1}; hence "effectively"). Clearly, excising the surface bounded by χ from L all at once achieves the same effect as excising the digons D_i one-by-one.

The region of L bound by a geodesic polygon χ is a particular instances of what we call a *crest*: a subset of L enclosing v whose removal and suitable suturing via AGT will reduce L to X. Note that we allow the boundary of a crest to include portions of ∂L, e.g., χ in Fig. 4.11(a) includes the x_i as well as the edge x_5x_6. In Chapter 7 we will show that it is possible to construct crests directly on L without deriving them from digon removals.

Chapter 6

Algorithm for Tailoring via Sculpting

In this research we have also concentrated on achieving constructive proofs of the theorems, constructive in the sense of leading to finite algorithms. In this chapter we follow Theorem 4.6 to yield an algorithm for achieving the tailoring of P to Q. Throughout we measure computational complexity in terms of n, where $n = \max\{|P|, |Q|\}$ is the number of vertices of the larger of P or Q; so $|P|, |Q| = O(n)$.[1] Our goal for all the algorithms is to achieve polynomial-time complexity, $O(n^k)$, but we have not worked hard to lower k, through, e.g., exploitation of efficient data structures. Instead we are content to leave improvements for future work. We will see that $k = 4$ seems to suffice.

Euler's $V-E+F$ theorem implies that the number of edges and faces of a polyhedron are linearly related to the number of vertices, so all components of P and Q are $O(n)$. Thus for tailoring via slicing, the number of slices of P is $O(n)$: one slice per face of Q. Each slice leads to g-domes, each g-dome to pyramids, and each pyramid is reduced to its base by a series of tailoring steps: digon excisions and suturings.

We analyze the complexity in four parts, the 0-th just the conceptual slicing of P by planes on each face of Q:

- Algorithm 0: Slice P to Q: $O(n)$.

[1] A finer analysis would treat the number of vertices of Q and P independently, say, Q with m vertices. The algorithm steps would be the same, but the complexities would be apportioned differently.

- Algorithm 1: slice → g-domes, following Lemma 4.1.
- Algorithm 2: g-dome → pyramids, following Theorem 3.2.
- Algorithm 3: pyramid → digons, following Lemma 4.4.

Algorithms $0, 1, 2$ operate on the extrinsic 3-dimensional structure of the polyhedra. Algorithm 3 instead processes its calculations on the intrinsic structure of the surface.

6.1 Algorithm 1: slice → g-domes

Algorithm 1: From one slice, $O(n)$ g-domes.

Input : One slice plane Π
Output: $O(n)$ g-domes, a total of at most $O(n)$ vertices.

```
// Consult Fig. 4.1.
```
Let F be the face of Q lying in Π, and e be an edge of F.
Sort vertices angularly about e. // $O(n \log n)$
for $i = 0, 1, 2, \ldots, k$ **do**
  Rotate Π_i about e until the portion swept is not a g-dome.
  `// Following Lemma 4.1.`
  Add to g-domes list.
end
Result: List of $O(n)$ g-domes.

Algorithm 1 follows the proof of Lemma 4.1, which clearly results in at most $O(n)$ g-domes, with the sum of the complexities of the g-domes for any one slice $O(n)$. It can be implemented to run in $O(n \log n)$ time. Therefore, over $O(n)$ slices, each resulting in $O(n)$ g-dome vertices, we have at most $O(n^2 \log n)$ total g-dome complexity: $O(n \log n)$ sorting repeated $O(n)$ times.

However, it could be that a clever choice of ordering of the slices always results in a smaller total complexity, perhaps $O(n \log n)$. Using Ω to indicate a lower bound,[2] the question is: Might each of the $\Omega(n)$ slices from Algorithm 0 lead to g-domes with a total complexity of $\Omega(n^2)$, independent of the order of slicing? Although resolving this question has only a small effect on

[2]For example, a constant fraction of n is $\Omega(n)$.

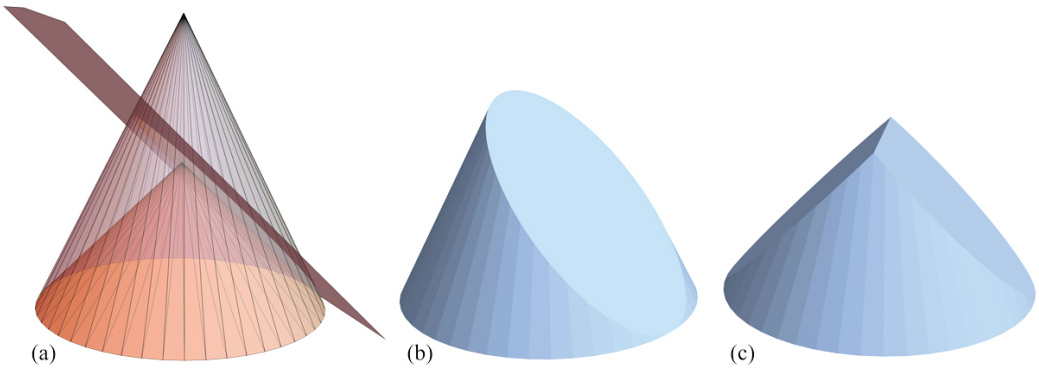

Figure 6.1: $Q \subset P$. The shared base is a regular polygon of n=48 sides. (a) One slice-plane Π_0 lying on a face of Q. (b) After slicing by Π_0. (c) A second, opposite slice $\Pi_{n/2}$.

the overall time-complexity of the tailoring algorithm, we take a detour to explore this issue via an example in the next subsection.

6.1.1 Complexity of sculpting

Example 6.1. *Consider Fig. 6.1, two nested pyramids $Q \subset P$ sharing a common regular polygon base. The combinatorial complexity of slicing P with face planes of Q is order-sensitive: from $\Omega(n^2)$ to $\Omega(n \log n)$.*

Proof. Let the faces of Q ordered in sequence around the base be $F_0, F_1, \ldots, F_{n-1}$, each F_i determining a plane Π_i. Let P_i be the polyhedron after slicing P with planes $\Pi_0, \Pi_1, \ldots, \Pi_i$. Π_0 cuts $n-2$ edges of P, as shown in Fig. 6.1(a,b), and effectively removes half the edges of P from later slices. If P is sliced in the order $i = 1, 2, \ldots, n-2$, following adjacent faces around the base, each plane cuts a diminishing number of the remaining edges. An explicit calculation shows that plane Π_i cuts $\lfloor \frac{n+1-i}{2} \rfloor$ edges of P_{i-1}. And because

$$\sum_{i=1}^{n-2} \left\lfloor \frac{n+1-i}{2} \right\rfloor = \Omega(n^2) \ ,$$

with this plane-slice ordering, $\Omega(n)$ slices each have $\Omega(n)$ vertices, for a total quadratic complexity, $\Omega(n^2)$.

However, if one instead orders the slices in a binary-search pattern, then the total complexity is $\Omega(n \log n)$, as we now show. Let $n = 2^m$ be a power of 2 without loss of generality. The pattern is slicing with planes lying on faces F_i with indices in the order

$$s = \left(0, \frac{n}{2}, \frac{n}{4}, \frac{3n}{4}, \frac{n}{8}, \frac{3n}{8}, \frac{5n}{8}, \frac{7n}{8}, \dots\right) .$$

We partition this sequence s of indices into subsequences, $s = (0, s_1, s_2, \dots, s_m)$, as follows:

$$\begin{aligned} s_1 &= \left(\frac{n}{2}\right) \\ s_2 &= \left(\frac{n}{2^2}, \frac{3n}{2^2}\right) \\ s_3 &= \left(\frac{n}{2^3}, \frac{3n}{2^3}, \frac{5n}{2^3}, \frac{7n}{2^3}\right) \\ &\dots \\ s_k &= \left(\frac{j\,n}{2^k}\right), \text{ where } j = 1, 3, 5, \dots, 2^k - 1 . \end{aligned}$$

Notice that the number of indices in sequence s_k, $|s_k| = 2^{k-1}$. As a check, the total number of indices in s is

$$1 + \sum_{k=1}^{m} |s_k| = 1 + \sum_{k=1}^{m} 2^{k-1} = 2^m = n .$$

One can calculate that the slice at $i = \dfrac{j\,n}{2^k}$ only cuts $\dfrac{n}{2^k}$ edges of P_{i-1}. Only $\dfrac{n}{2^{k-1}}$ edges are "exposed" to Π_i: for example, for $k = 2$ and $i = 1/4$, $n/2$ edges are possibly available for cutting by Π_i, as can be seen in Fig. 6.1(c). However, because of the slant of Π_i, only half of those, $n/4$, are in fact cut by Π_i.

Now we compute the total number of edges sliced by the planes following the sequence s. Because $|s_k| = 2^{k-1}$, and each slice in s_k cuts $\dfrac{n}{2^k}$edges, the total number of cuts over all k is

$$\sum_{k=1}^{m} 2^{k-1} \frac{n}{2^k} = n \sum_{k=1}^{m} \frac{1}{2} = n\,\frac{m}{2} .$$

And since $m = \log n$, the total complexity is $\Omega(n \log n)$, or an average of $\Omega(\log n)$ for each of $\Omega(n)$ slices. □

In the absence of a resolution to this complexity question, we will assume that Algorithm 0 and Algoirthm 1 together result in $O(n)$ g-domes each of $O(n)$ vertices, produced in time $O(n^2)$.

6.2 Algorithm 2: g-dome → pyramids

Algorithm 2: Partition one g-dome to $O(n)$ pyramids.

Input : One g-dome G of $O(n)$ vertices
Output: $O(n)$ pyramids, each of size $O(n)$; and $O(n^2)$ pyramids, size $O(1)$.

```
// Following Theorem 3.2.
```
for *each of* $k = O(n)$ *vertex-degree reductions* **do**
  As in Theorem 3.2, slice with plane: $O(k)$.
  Remove pyramid P of $O(k)$ vertices.
  "Clean-up" by removing k pyramids each of size $O(1)$.
end
Result: List of pyramids.

We follow Theorem 3.2 for partitioning each g-dome G into pyramids. Each vertex v_i of the top-canopy of G is removed, as in Fig. 3.3, until only one remains. Removal of each v_i follows the degree-removal steps illustrated in Fig. 3.4.

Because the sum of the vertex degrees of a g-dome is $2E = O(n)$, the asymptotic complexity of processing a g-dome with many vertices in its top-canopy is no different than it is for just two vertices as in Fig. 3.3. Moreover, we can assume that v_2 has degree-3 while v_1 has degree-k, with $k = O(n)$.

First a plane slice results in a pyramid of k vertices with apex v_1, which is removed (and reduced by Algorithm 3). Next follows a "clean-up" phase that removes $O(k)$ pyramids each of 4 or 5 vertices, so of constant size, $O(1)$.

This is then repeated for the new apex of degree $k-1$: removal of a pyramid of $k-1$ vertices, and cleanup of $O(k-1)$ pyramids of constant size. After iterating through $k, k-1, k-2, \ldots$, the algorithm has sliced off $O(k^2)$

pyramids of constant size, and $O(k)$ pyramids of size $O(k)$. In the worst case $k = O(n)$, for a total complexity of $O(n^2)$.

6.3 Algorithm 3: pyramid $\rightarrow$ digons

Algorithm 3: Tailor one pyramid P to its base X.

Input : A pyramid P of $O(n)$ vertices
Output: $O(n)$ digons whose removal flattens P to X.

```
// Following Lemma 4.4.
// Assume apex degree-k, with k = O(n).
```
for *each of* x_i, $i = 1, 2, \ldots, k$ **do**
  Construct digon $D_i(x_i, y_i)$: Locate y_i.
  Locate y_i by tracing geodesics: $O(k)$.
end
Result: List of $O(n)$ digons

Lastly we concentrate on the cost of removing one pyramid P of $O(n)$ vertices. Following Lemma 4.4, this requires $O(n)$ digon removals. For each digon $D_i(x_i, y_i)$, we need to calculate the location of y_i on $C(x_i)$; then y_i becomes a vertex for the removal of the next digon D_{i+1}. Fortunately, there is no need to compute the cut locus $C(x_i)$.

Let us focus on locating y_i, after the removal of $D_{i-1}(x_{i-1}, y_{i-1})$ the previous iteration. Recall that y_i is the only vertex on L_i, so it is immediate to find the shortest path γ from x_i to y_i. We know the angle $\theta = \omega_Q(x_i) - \omega_P(x_i)$ needed to be removed by D_i, so we know that geodesics γ_1 and γ_2 at angles $\theta/2$ left and right of $\gamma = x_i y_i$ will meet at y_i. Tracing γ_1 and γ_2 over the surface might cross sealed digons $D_1, \ldots, D_{i-1}$, the blueseals σ_i in Fig. 4.11(b). So the cost of computing $\gamma_1 \cap \gamma_2 = y_i$ is $O(n)$.

Thus the complexity of tailoring one pyramid of $O(n)$ vertices is $O(n^2)$.

Note that we do not need the extrinsic 3-dimensional structure of the intermediate polyhedra guaranteed by AGT to perform the calculations, as is evident in the example described in Fig. 4.11.

6.4 Overall Tailoring Algorithm

Putting the three algorithm complexities together, and for succinctness abbreviating $O(n)$ with n and "time-complexity" with "complexity," we have:

(1) n slices $\to$ n g-domes, each of size n: complexity $n^2 \log n$.

(2) 1 g-dome of size n $\to$

  (a) n pyramids, each of size n: complexity n^2.

  (b) $+ n^2$ constant-size pyramids: complexity n^2.

(3) 1 pyramid of size $n \to n^2$ to reduce to base: complexity n^2.

(1) and (2) together have complexity n^3, iterating over each of the n domes. (The $n \log n$ sorting need not be repeated after the g-domes are identified.) (3) repeats n^2 over (potentially) n^2 pyramids: n g-domes resulting in n pyramids each. Thus the total complexity is n^4. We summarize in a theorem:

Theorem 6.2. *Given convex polyhedra P and Q of at most n vertices each, and $Q \subset P$, P can be tailored to Q, tracking a sculpting of Q from P as in Theorem 4.6, in time $O(n^4)$.*

Chapter 7

Crests

In this chapter we revisit the suggestion made at the end of Chapter 5 that the digons to reduce one pyramid to its base could be cut out all at once. As before, let $P = L \cup X$ be a pyramid with base X and lateral sides L. Recall that a *crest* is a subset of L enclosing v whose removal and suitable suturing via AGT will reduce L to X.

We derive a method for identifying a crest that does not rely on digon removals, but rather works directly on a pyramid. This allows us to achieve in Theorem 7.5 reshaping of P to Q by the removal of crests to flatten pyramids. We call this process *crest-tailoring*, in contrast to the *digon-tailoring* explored in Chapter 4. It represents a tradeoff between the simplicity of digons and the number of excisions: for digon-tailoring, $O(n)$ per pyramid, for crests, one per pyramid. We first illustrate the process of identifying a crest on two example pyramids before proving that it always works.

7.1 Examples

Let $P = L \cup X$ with $\partial X = \partial L = (x_1, x_2, \ldots, x_k)$ having vertices x_i. The apex is v, which projects orthogonally to $\bar{v}$. We will describe the procedure for identifying a crest first for $\bar{v} \in X$ and then for $\bar{v} \notin X$. Although the cases initially feel different, the proofs will show that they are nearly the same.

Fig. 7.1 illustrates the case $\bar{v} \in X$. Let $T_i = x_i x_{i+1} v$ and $\bar{T}_i = x_i x_{i+1} \bar{v}$. We proved in Lemma 2.8 that the angle $\bar{\theta}_i$ at x_i in X is strictly smaller than than the angle θ_i on L, the sum of the angles in T_{i-1} and T_i incident to x_i (as long as $|v\bar{v}| > 0$).

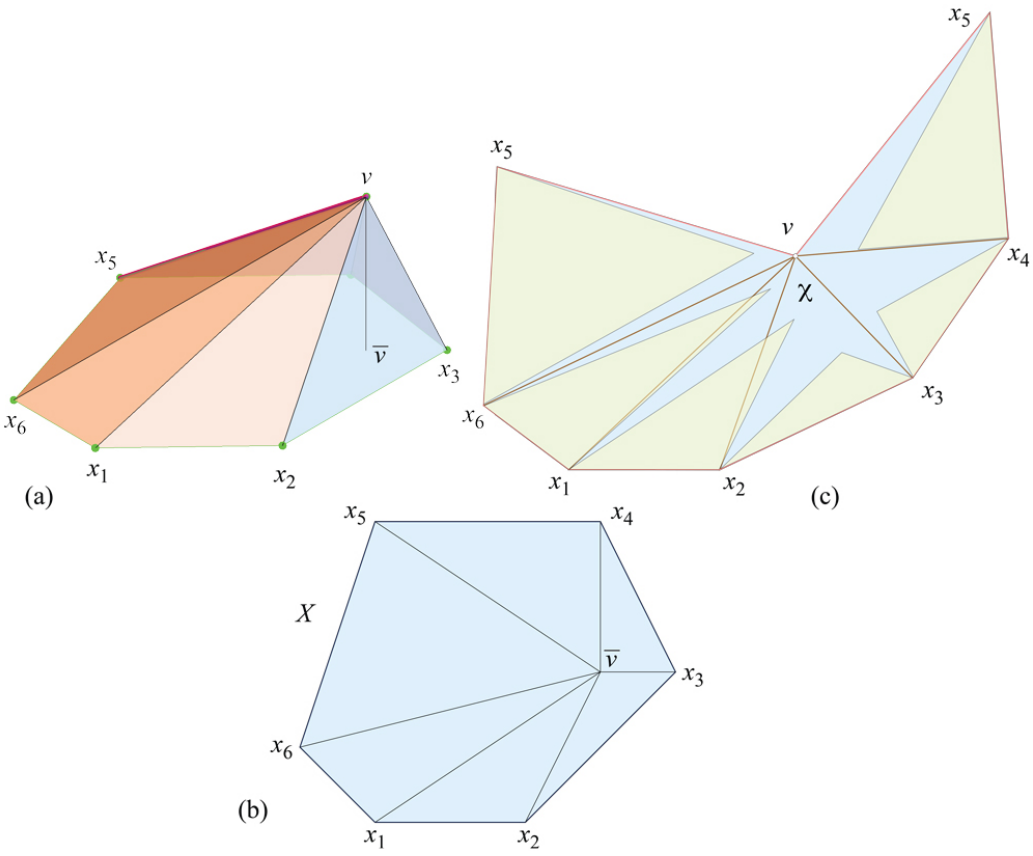

Figure 7.1: (a) Pyramid. (b) X with projected triangles $\bar{T}_i = x_i x_{i+1} \bar{v}$. (c) Flattening of L to $\bar{L}$, in this case by cutting the edge $x_5 v$. The lifted triangles T_i^L are shown yellow. The crest χ is blue.

A key definition is the *lift* of $\bar{T}_i$ onto L. Let $\bar{\alpha}_i$ and $\bar{\beta}_i$ be the base angles of $\bar{T}_i$, at x_i and x_{i+1} respectively. On L, extend geodesic γ_i from x_i at angle $\bar{\alpha}_i$, and extend geodesic γ_{i+1} from x_{i+1} at angle $\bar{\beta}_i$. Let v_i be the point on L at which these geodesics meet; v_i is the image of $\bar{v}$. (We will not establish that indeed these geodesics meet on L until Lemma 7.3.) Then lift$(\bar{T}_i) = T_i^L$ is a geodesic triangle on L isometric to $\bar{T}_i$. Another way to view the lift of $\bar{T}_i$ is to imagine $\bar{T}_i$ rotating about $x_i x_{i+1}$ by the dihedral angle there and pasting it on the inside of L.

Yet another way to view the lift is as follows. L is isometric to a cone and can be flattened by cutting along a generator, i.e., a segment from v to ∂L. Let $\bar{L}$ be a particular flattening, with the cut generator not "near" x_i just for simplicity. Then place a copy of $\bar{T}_i$ on $\bar{L}$ matching $x_i x_{i+1}$. Then refold $\bar{L}$ to L. We will continue to reason with a flattened $\bar{L}$ but remembering that $\bar{L}$ is a representation of L, and so the cut edge is not relevant.

This last layout-viewpoint yields a method to construct a full crest, call it χ. The base X, partitioned into $\bar{T}_i$, can be modified by opening the angle at x_i from $\bar{\theta}_i$ to θ_i. After opening at all x_i, this figure can be superimposed on the flattened $\bar{L}$, matching the boundaries $x_1, \ldots, x_k$. This is illustrated in Fig. 7.1(c). The crest is then the portions of L not covered by the lifted $\bar{T}_i$. It should be clear that cutting out χ and suturing closed the matching edges will reduce L to X, for the $\bar{T}_i$ remaining after removing χ exactly partition X.

We next illustrate the case when $\bar{v} \notin X$; see Fig. 7.2. We perform the exact same process of lifting triangles $\bar{T}_i$ to L, but we clip those triangles to X—i.e., form the polygon $\bar{T}_i \cap X$—as indicated in (c) of the figure. Notice that two triangles, $\bar{T}_3$ and $\bar{T}_4$, are removed by the clipping intersection. Again it should be clear that cutting out χ and suturing closed will reduce L to X.

7.2 Proofs

We will need several geometric properties.

Let $\bar{\theta}_i$ be the angle at x_i in X, and θ_i the angle at x_i on L, the sum of two triangle angles incident to x_i. Then, by Lemma 2.8, $\theta_i > \bar{\theta}_i$.

Lemma 7.1. *Let $T = abc$ be a triangle in $\mathbb{R}^3$, with ab on plane Π and c above that plane. Let $\bar{c}$ be the orthogonal projection of c onto Π, and $\bar{T} = ab\bar{c}$ the projected triangle. Finally, let $T^r = abc^r$ be the triangle T flattened to Π*

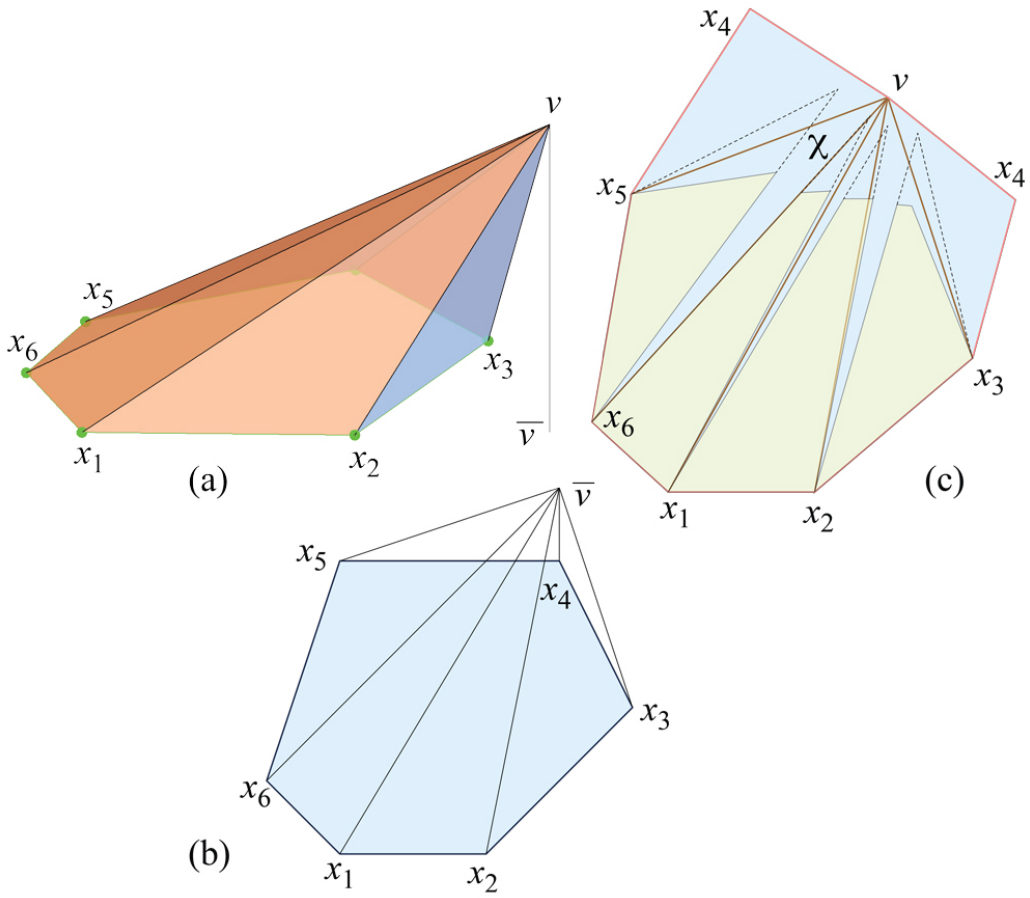

Figure 7.2: Following the same conventions as in Fig. 7.1: (a) Pyramid. (b) X; $\bar{v} \notin X$. (c) $\bar{L}$, yellow clipped lifted triangles T_i^L, and blue crest χ.

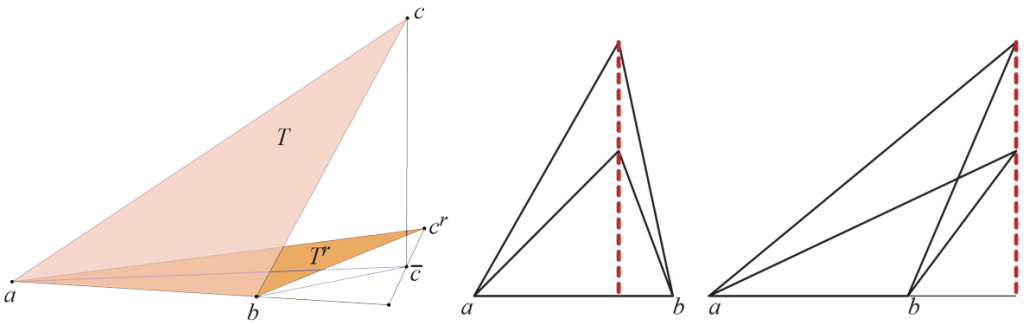

Figure 7.3: Rotated apex c^r and projected apex $\bar{c}$ lie on same altitude.

by rotating about ab. Then c^r and $\bar{c}$ lie on the altitude line perpendicular to the line containing ab.

Proof. See Fig. 7.3. The claim follows from the Theorem of the Three Perpendiculars. Note that T^r is congruent to T. □

The consequence of Lemma 7.1 is that, superimposing a lifted triangle $\bar{T}^L$ on a planar layout $\bar{L}$ of L, the images of v in $\bar{L}$ and $\bar{v}$ of $\bar{T}$, lie along the altitude of T.

The following lemma assumes that $\bar{v} \in X$. The case when $\bar{v} \notin X$ will be treated separately. Let $\bar{L}$ be a planar layout of L, say, cut open at edge x_1v. Let $\tau_i = \pi - (\beta_{i-1} + \alpha_i) = \pi - \theta_i$ be the turn angle at vertex x_i in the layout. Because $\bar{v} \in X$, $\tau_i > 0$, i.e., the planar image of ∂L in $\bar{L}$ is a convex chain, and also convex wrapping around the cut edge x_1v.

Let a_i be the segment altitude of triangle $T_i = x_ivx_{i+1}$ in the layout $\bar{L}$. Let $\nu = 2\pi - \omega(v)$ be the surface angle of P incident to v.

Lemma 7.2. *When $\bar{v} \in X$ and consequentially $\bar{L}$ is a convex chain $x_1, \ldots, x_k$, the following hold (with $k + 1 \equiv 1 \bmod k$):*

(a) The sum of the turn angles, $\sum_i \tau_i = \nu$, the surface angle incident to v.

(b) The angle between a_i and a_{i+1} at v on L is exactly equal to the turn angle τ_i at vertex x_i.

(c) The altitudes occur in order around v, in the sense that a_{i+1} is counterclockwise of a_i around v.

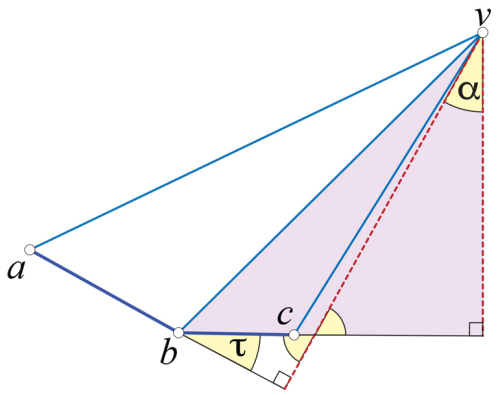

Figure 7.4: Triangles avb and bvc turn τ at b of the convex chain abc. τ is equal to the altitude turn α; angles of similar right triangles marked.

Proof. (a) Viewing the entire layout $\bar{L}$ as a simple polygon, $\sum_i \tau_i = 2\pi$. But we need to distinguish between τ_1 at the edge $x_1 v$ cut to flatten L, and τ_1', the turns at the two images of x_1 in $\bar{L}$:

$$\tau_1' = 2\pi - (\alpha_1 + \beta_k) = \pi + \tau_1 \ .$$

The second anomalous turn in $\bar{L}$ is $\pi - \nu$ at v. So we have

$$\begin{aligned} \tau_1' + \sum_{i=2}^{k} \tau_i + (\pi - \nu) &= 2\pi \ , \\ \sum_{i=1}^{k} \tau_i &= \nu \ . \end{aligned}$$

(b) This can be seen by extending $T_i = x_i v x_{i+1}$ to a right triangle, with right angle at the foot of altitude a_i. See Fig. 7.4.

(c) This follows directly from (b). Note that here we rely on the turns τ_i being positive, i.e., convex. See Fig. 7.5.

□

The consequence of Lemma 7.2 is that the surface angle ν around v is partitioned by the altitudes a_i in order, because $\sum_i \tau_i = \nu$, and the angle between a_i and a_{i+1} is τ_i. Moreover, Lemma 7.1 shows that the apexes of each lifted triangle $\bar{T}_i^L$ lie on those altitudes, at some positive distance from

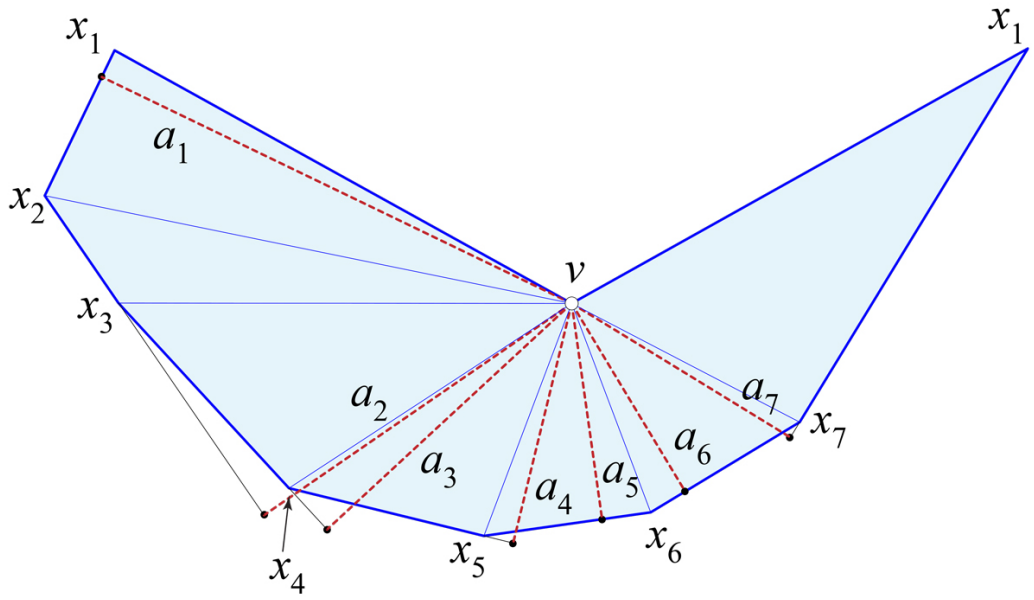

Figure 7.5: A convex chain $x_1, \ldots, x_7$ and the corresponding altitudes $a_1, \ldots, a_7$.

v. Consequently, we can connect those apexes to form a simple geodesic polygon enclosing v on L. Because every turn angle τ_i is strictly less than π, connecting two adjacent apexes along a_i and a_{i+1} will keep v to the same (counterclockwise) side. Call this polygon the *moat* M of P.[1] Fig. 7.6 illustrates the moat for the example in Fig. 7.1(c).

Lemma 7.3. *For the case $\bar{v} \in X$, the lifting of all triangles $\bar{T}_i$ to $\bar{T}_i^L$ onto L has the following properties, (where we shorten "geodesic triangle" to "triangle"):*

(a) Each lifted triangle $\bar{T}_i^L$ fits on L: $\bar{T}_i^L \subset L$.

(b) v does not lie in any triangle $\bar{T}_i^L$.

(c) No lifted triangle self-overlaps, and no pair of triangles overlap.

Proof. (a) Because the apex of the lifted $\bar{T}_i^L$ is on the moat M which surrounds v, $\bar{T}_i^L$ remains on the portion of L outside the moat.

[1] We do not know whether M is always convex, but we only need it to be simple.

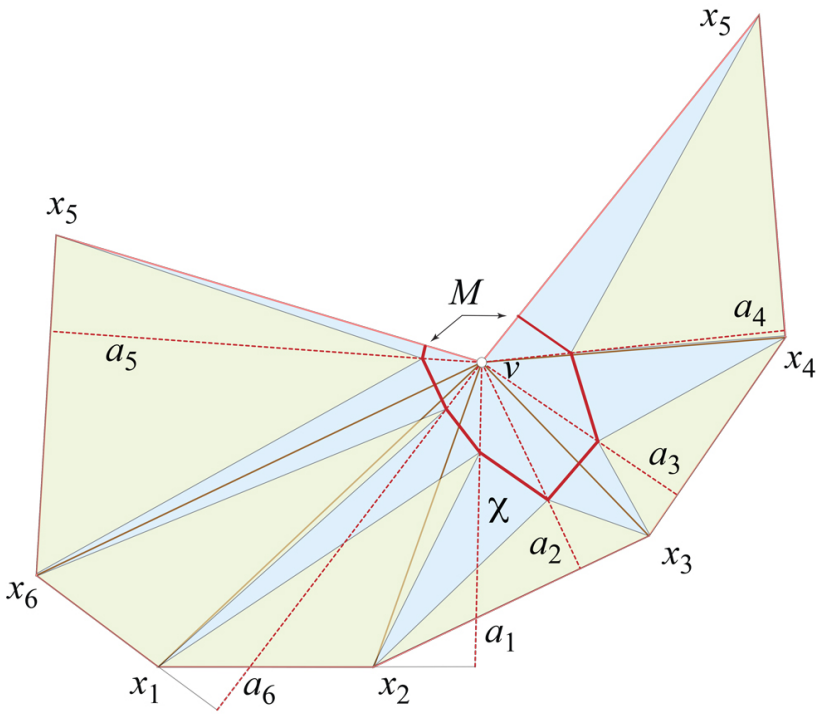

Figure 7.6: The layout from Fig. 7.1(c) shown with moat M and altitudes a_i identified.

(b) Therefore no $\bar{T}_i^L$ includes v.

(c) If we view the overlay of $\bar{L}$ with the opening of ∂X by the angle $\theta_i - \bar{\theta}_i$ at each x_i image, as in Fig. 7.1(c), then Cauchy's Arm Lemma shows that two lifted triangles cannot overlap. Suppose $\bar{T}_i^L$ and $\bar{T}_j^L$ overlap, $i < j$. Then we can identify two points $p_i \in \bar{T}_i^L$ and $p_j \in \bar{T}_y^L$ that coincide in the layout. But in X, p_i and p_j were separated by a positive distance $d = |p_i p_j|$. In X, draw a convex chain from p_i to ∂X, around that boundary, to p_j. The layout opens this chain by the positive angles $\theta_i - \bar{\theta}_i$, and so in the layout, p_i and p_j must be separated further than d, a contradiction.

□

Lemma 7.3 shows that χ, the region of L not covered by the lifted triangles, is indeed a crest.

We now turn to the case $\bar{v} \notin X$. The difficulty here is that ∂L in a layout $\bar{L}$ of L may not be a convex chain, and Lemma 7.2 relies on convexity for the altitudes to connect to v in the same order as the vertices around X. Indeed if v were closer the plane of X in the example in Fig. 7.2(a), then the angle at x_3 would be reflex. In general, a contiguous portion of ∂L could be reflex. Lifting triangles incident to that reflex chain could lead to overlap, violating (c) of Lemma 7.3.

However, as described earlier in Fig. 7.2(c), the crest is formed by clipping the triangles $\bar{T}_i$ to X. Triangles $\bar{T}_3 = x_3 \bar{v} x_4$ and $\bar{T}_4 = x_4 \bar{v} x_5$ in Fig. 7.2(b) fall entirely outside X, and so play no role. The convex portion of ∂L still satisfies Lemma 7.2, so the corresponding altitudes are incident to v in the same order as the vertices along the convex chain. This allows us to define a partial moat M, and then close it off to a simple polygon by a geodesic path surrounding v. This is illustrated in Fig. 7.7.

This renders Lemma 7.3 true for the lifted triangles along the convex chain of ∂L, which are the only ones not clipped entirely away. We summarize in a theorem:

Theorem 7.4. *A crest χ can be constructed as the portion of L not covered by lifted triangles in the case of $\bar{v} \in X$, and clipped lifted triangles in the case $\bar{v} \notin X$, as described above.*

We remark that the same procedure will work for other points $w \in X$ within some neighborhood of $\bar{v} \in X$, resulting in different crests. However, a simple

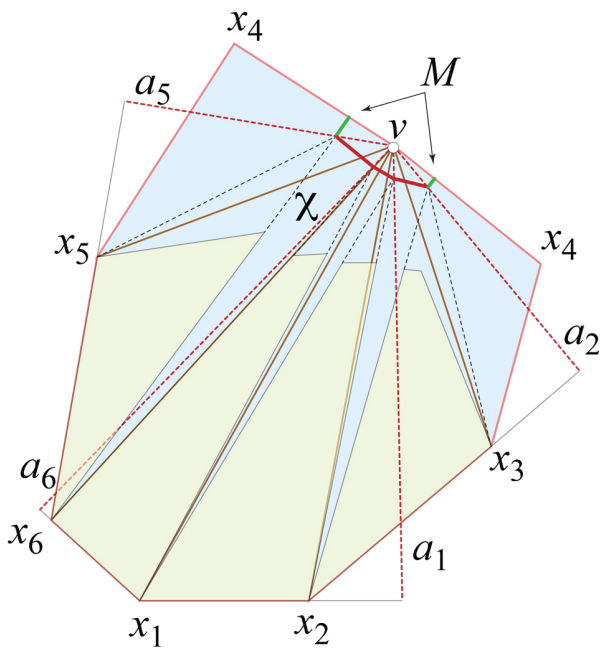

Figure 7.7: The layout from Fig. 7.2(c) shown with moat M and altitudes a_i marked. Note a_3 and a_4 are missing because T_3 and T_4 are clipped as outside X. The green edges mark the closing of the partial moat around v.

example shows that not every point $w \in X$ will produce a crest. Consider $X = x_1x_2x_3$ an equilateral triangle of center o, and v close to o, with vo orthogonal to X. Take $w \in ox_3$ close to x_3. Then the isosceles triangle wx_1x_2 is larger than the isosceles triangle vx_1x_2, so no congruent copy of the former can fit inside L without encompassing v and so self-overlapping.

We have this as a counterpart to Theorem 4.6:

Theorem 7.5. *For any convex polyhedra P and Q, one can crest-tailor P to any homothetic copy of Q inside P, in time $O(n^4)$, where $n = \max\{|P|, |Q|\}$.*

Proof. The lemmas leading to Theorem 4.6 established that ultimately we need to tailor single vertex truncations, i.e., tailor pyramids. So the claim follows from Theorem 7.4. □

7.3 Algorithm 4: pyramid → crest

Algorithm 4: Construct a crest χ that reduces a pyramid P.

Input : A pyramid $P = L \cup X$ of $O(n)$ vertices
Output: Crest χ whose removal flattens P to X.

`// Assume apex` v `degree-`k`, with` $k = O(n)$.
`// See Figs. 7.1 and 7.2.`
for *each of* $\bar{T}_i$, $i = 1, 2, \ldots, k$ **do**
  Compute where $\bar{T}_i$ edges (geodesics) cross x_jv. `//` $O(n)$.
  Clip $\bar{T}_i$ to X. `//` $O(n)$.
end
Result: Crest χ.

The total cost of computing one crest χ is $O(n^2)$, and we believe there are examples with total combinatorial complexity (number of geodesic/edge intersections) of $\Omega(n^2)$. Because this is the same complexity for reducing P to X via digon-tailoring described in Chapter 6, the total time complexity is the same as in Theorem 6.2.

We are assuming that the combinatorial complexity of the crest χ on L determines the time complexity of computing the crest. However, a crest flattened to the plane, not overlaid on $\bar{L}$, has combinatorial complexity $O(n)$—just $k = O(n)$ (possibly clipped) triangles—and can be constructed in $O(n)$

time. Perhaps an implicit representation of shortest paths could suffice for subsequent calculations, as they do in the optimal algorithm for shortest paths on a convex polyhedron [SS08]. Thus it may be that constructing an implicit representation of χ could lead to a lower time-complexity, a question we leave for future work.

Chapter 8

Tailoring via Flattening

In this chapter, we prove a completely different method for tailoring P to Q. The method mixes digon-tailoring steps with vertex-merge steps (Section 2.4). The result is slightly weaker than either tailoring via sculpting (Theorem 4.6) or crest tailoring (Theorem 7.5), weaker in the sense that the homothet of Q obtained could be arbitrarily small. Nevertheless, the proof and algorithm have the advantage of operating entirely intrinsically: the 3D structure of P and Q is never invoked.

The proof depends on the observation that if both P and Q are doubly-covered polygons, it is easy to tailor one to the other: Scale Q to fit in P, and then cut the outline of Q from P. This can be accomplished by a series of digon-tailorings, with each digon bounded by congruent segments on both sides, or by truncations of several vertices at once.

At a high level, Q is reduced to a flat polygon Q_{flat}, P is reduced to a flat polygon P_{flat} and tailored to match Q_{flat}. Finally, the steps used to reduce Q are reversed and applied to the flat remnant of P.

In more detail, the proof (and algorithm that follows the proof) can be summarized in these steps:

(1) Reduce P to P_{flat} by a series of digon tailorings.

(2) Reduce Q to Q_{flat} by a series of vertex-mergings. Q_{flat} is then composed of all of Q's surface, plus surface-inserts from the merges.

(3) Scale Q_{flat} to Q^s_{flat} so that Q^s_{flat} fits inside P_{flat}.

(4) Trim P_{flat} to match Q^s_{flat} via digon-tailoring. Call the result P^s_{flat}.

(5) Reverse the vertex-merging steps that reduced Q to Q_{flat}, but now applied to P^s_{flat}. Each reversal step removes a surface-insert via digon-tailoring.

The end result, call it Q^t, is a polyhedron homothetic to Q, but composed entirely of P-surface.

We now prove these steps, tracking an example that reduces a cube P to the 5-vertex polyhedron Q in Fig. 2.6(a), repeated in Fig. 8.2(a).

8.1 Proofs

We follow the numbered outline above.

8.1.1 Digon-tailor $P \to P_{\text{flat}}$

Recall that Lemma 2.2 showed that, if the cut locus $C(x)$ is a path, then the polyhedron is a doubly-covered convex polygon. We use this lemma to reduce P to P_{flat}.

We will use the cube example from Fig. 2.2 to illustrate the steps. Assume P is non-degenerate, i.e., not flat. Let $x \in P$ be a point joined by unique geodesic segments to all vertices of P, and let ρ be the unique path in $C(x)$ joining a pair of leaves of $C(x)$, i.e., joining the vertices v_i and v_j of P.

Then $C(x) \setminus \rho$ is a finite set of trees T_k. Cut off from P each T_k by excising digons with one endpoint at x, and the other endpoint where T_k joins ρ. In Fig. 8.1(a), ρ connects v_5 and v_7, and separates four trees T_i. After sealing each digon closed, we are left with a polyhedron P_{flat} whose cut locus from the point corresponding to x is precisely ρ (by Lemma 2.2).

If the path ρ in the above proof is chosen to be as long as possible, then P_{flat} has larger surface area than if ρ is short.

8.1.2 Vertex-merge $Q \to Q_{\text{flat}}$

Recall from Section 2.4 that vertex-merging is in a sense the inverse of digon-tailoring. Two vertices are merged along a geodesic γ and additional surface in the form of two congruent triangles is sutured-in along γ. Lemma 2.11 showed that every convex polyhedron Q has at least one pair of vertices that can be merged, unless Q is an isosceles tetrahedron or a doubly-covered

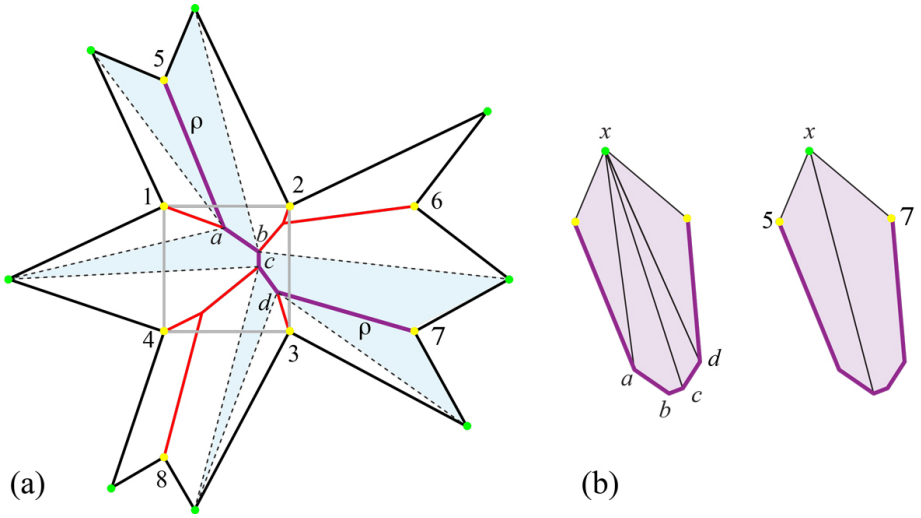

Figure 8.1: Star unfolding of the cube in Fig. 2.2. (a) The path ρ from v_5 to v_7 leaves four trees when removed from $C(x)$. Excising four (white) digons leaves a surface (blue), which when zipped closed folds to a doubly-covered 7-gon, both sides of which are shown in (b).

triangle. So repeatedly vertex-merge Q until it becomes flat, or is reduced to an isosceles tetrahedron.

In our example, Fig. 8.2(a) shows Q, the same polyhedron in Fig. 2.6(a). The first vertex-merge leads to a regular tetrahedron, (c) in the figure. As we discussed in Example 1.3 and illustrated in Fig. 1.4, an isosceles tetrahedron can be reduced to a doubly-covered rectangle by cutting an edge and regluing that edge differently. Here we cut the edge cd of the tetrahedron and reglue it by creasing at the midpoints y_1, y_2 of the cd slit, leading to the rectangle shown in (d).

8.1.3 Scale $Q_{\text{flat}} \to Q^s_{\text{flat}}$

The flattening of P has reduced it in size in the sense that a portion of P's surface area has been excised, while the flattening of Q has augmented it by surface insertions and so has increased Q's surface area. We next select a scale factor $s > 0$ so that Q^s_{flat} can fit inside P_{flat}. Clearly this is always possible. In our example, the scaling might result in Fig. 8.3(a). Note that we should view the top side of Q^s_{flat} on the top side of P_{flat}, and similarly for

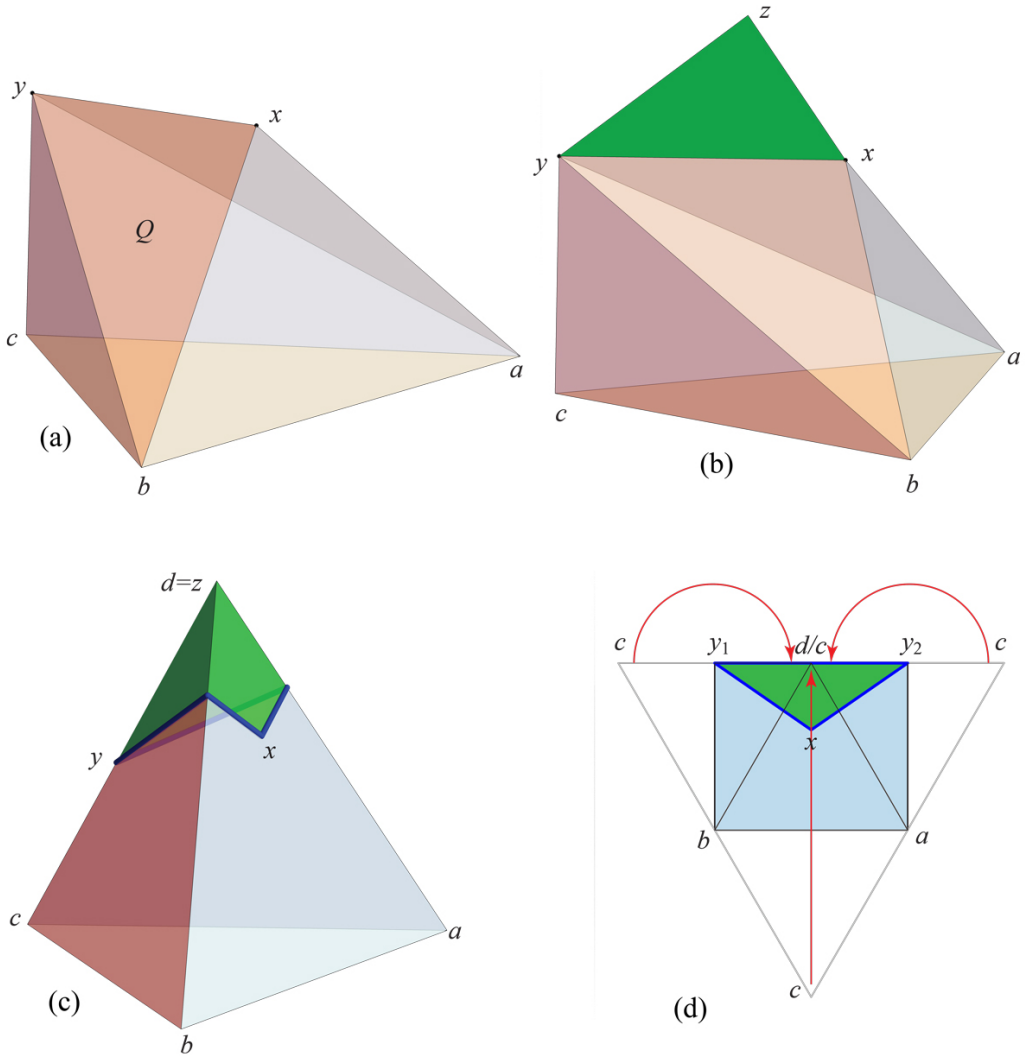

Figure 8.2: (a,b) Vertex-merge of xy. (Repeat of Fig. 2.6.) Green indicates surface-inserts. (c) After first vertex-merge. (d) After second special-case vertex-merge: doubly-covered square. (Images not to same scale.)

the bottom sides.

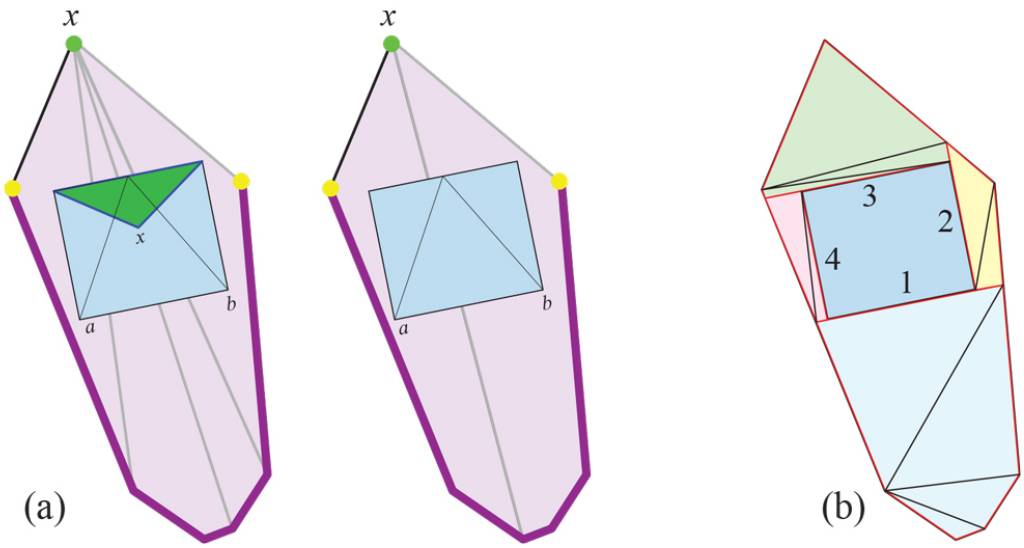

Figure 8.3: (a) Q scaled to Q^s_{flat} to fit in P_{flat}. (b) Trimming digon excisions by extending the edges of P_{flat} in the order indicated.

8.1.4 Trim $P_{\text{flat}} \rightarrow P^s_{\text{flat}}$

As mentioned earlier, it is easy to tailor a doubly-covered convex polygon P_{flat} to match a doubly-covered polygon Q^s_{flat} that fits inside P_{flat}. This is in some sense similar to sculpting $P_{\text{flat}} \rightarrow P^s_{\text{flat}}$. The differences are that (1) here we only work in two dimensions, and (2) this 2D sculpting is an intrinsic operation as well. Fig. 8.3(b) shows one possible sequence of digon tailorings for our example. We extend edge e of P^s_{flat} to a line l_e which intersects P_{flat}, thus separating out a doubly-covered convex polygon, and repeat this for all edges e. Then each of those convex polygons is tailored one vertex at a time, say tracking a triangulation.

8.1.5 Reverse $P^s_{\text{flat}} \rightarrow Q^t$

Finally we reverse the vertex-merge steps that produced Q_{flat}, but applied to P^s_{flat}. In our example, the last vertex-merge was the special-case step that produced the rectangle in Fig. 8.2(d). The reverse step cuts the top edge

y_1y_2 between the two images of y, and re-joins y_1 to y_2, leading to the regular tetrahedron in (c) of the figure. Reversing the first vertex-merge applied to Q excises the inserted surface, green in (c), resulting in Q^t, a polyhedron homothetic to Q (Fig. 8.2(a) and 2.6) but composed entirely of P-surface.

8.1.6 Theorem: Tailoring via Flattening

We have established this theorem:

Theorem 8.1. *For any given convex polyhedra P and Q, one can tailor P "via flattening" so that it becomes homothetic to Q.*

Remark 8.2. *The result Q^t obtained by Theorem 8.1 may be arbitrarily small compared to Q.*

1. *The example of a regular pyramid shows that the area of P_{flat} may be as small as $2/n$ of the original area of P.*

2. *The ratio between the area of Q_{flat} and the area of Q can be arbitrarily large.*

   See Fig. 8.4.

   To see this, consider an isosceles trapezoid Z of base lengths 1 and $1 + 2\varepsilon$, and height h. Its area is $(1 + \varepsilon)h$. Also consider the isosceles triangle T obtained from Z by extending its non-parallel sides until intersecting. An elementary geometry argument provides the height of T, $(1 + 2\varepsilon)h(2\varepsilon)^{-1}$, and the area of T, $(1 + 2\varepsilon)^2h(4\varepsilon)^{-1}$.

   The ratio between the area of the doubles Q_{flat} of T, and Q of Z, is therefore

   $$\frac{(1 + 2\varepsilon)^2}{4\varepsilon(1 + \varepsilon)} = 1 + \frac{1}{4\varepsilon(1 + \varepsilon)}$$

   and can be arbitrarily large for ε arbitrarily small.

The combination of P_{flat} small and Q_{flat} large leads to the arbitrarily-small claim for Q^t with respect to Q.

Algorithm 5: Tailor P to Q via Flattening.

Input : Convex polyhedra P and target Q
Output: A tailored version of P homothetic to Q

// (1) Reduce P to P_{flat}.
Find generic point x. // $O(n^4)$
Compute cut locus $C(x)$. Select path ρ. Digon removals of trees until P_{flat} attained.

// (2) Vertex merge on Q repeatedly.
while $Q \neq T_{isos}$ *and* $|Q| > 3$ **do**
  Identify two vertices v_i and v_j such that $\omega_i + \omega_j < 2\pi$. // $O(n^2)$
  Vertex merge v_i and v_j, reducing Q by one vertex to Q'.
  $Q \leftarrow Q'$.
end

if $Q = T_{isos}$ **then** // Q isosceles tetrahedron
  Special tailor Q to flat rectangle Q_{flat}.
end

// (3) Scale Q_{flat} to Q^s_{flat}.
Find largest inscribed and smallest circumscribed circles. // $O(n)$.
Scale Q_{flat} by radii ratio.

// (4) Trim P_{flat} to P^s_{flat}.
Extend edges of P^s_{flat}, triangulate each cut-off piece. // $O(n^2)$.

// (5) Reverse steps to reduce Q, each applied to P^s_{flat}
foreach *vertex-merging step applied to* Q **do**
  Reverse the step by cutting off the merge vertex.
end

Result: A 3D polyhedron Q^t homothetic to Q.

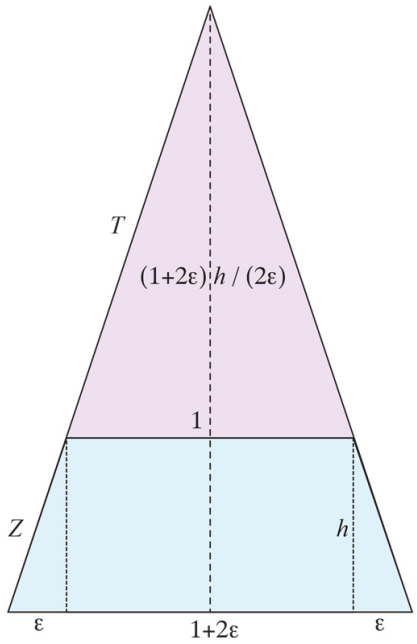

Figure 8.4: The ratio of the areas of T and Z is arbitrarily large.

8.2 Algorithm for Tailoring via Flattening

In this section, we follow the proof of Theorem 8.1 and convert it to a polynomial-time algorithm.

As usual, let $n = \max\{|P|, |Q|\}$ be the combinatorial size of the polyhedra. We now establish an upper bound of $O(n^4)$ on the complexity of implementing the algorithm.

Step (1) is to tailor P to P_{flat} using the cut locus $C(x)$ from a "generic point" x, i.e., one with a unique shortest path to each vertex of P. Although it is possible the need for uniqueness could be avoided, we leave that future work. We know of no way to find a generic x short of computing all the "ridge-free" regions on P, which takes $O(n^4)$ time [AAOS97]. Independent of our work here, it is a interesting question if a generic x can be computed more quickly. We will see this $O(n^4)$ dominates the complexity of the other calculations.

The star-unfolding $S_P(x)$ can be computed in $O(n \log n)$ time using the complex Schreiber-Sharir algorithm [SS08], or in $O(n^2)$ time with the Chen-Han algorithm [CH90, CH96]. $S_P(x)$ only needs to be computed once. With

$S_P(x)$ computed, the cut locus $C(x)$ can be found from the Voronoi diagram of the images of x.

Step (2) is to repeatedly apply vertex-merging to Q until it is reduced to Q_{flat}, when $|Q_{\text{flat}}| \in \{3, 4\}$. Identifying two vertices v_i and v_j such that $\omega_i + \omega_j < 2\pi$ can be achieved in $O(n \log n)$ time just by sorting the curvatures ω_i and selecting the two smallest. From the initial sorting onward, only $O(\log n)$ would be needed to update the list, but we'll see this efficiency is not necessary.

With v_i and v_j selected, the shortest path γ between them needs to be computed. Although there is a complicated optimal $O(n \log n)$ algorithm for computing shortest paths on a convex polyhedron [SS08], that algorithm exploits the three-dimensional structure of the polyhedron, which will not be available to us after the first vertex-merge. As mentioned earlier, there is no effective procedure known to construct the polyhedron guaranteed by AGT. However, we know the intrinsic structure of the polyhedron: its vertices, their curvatures, a triangulation. The algorithm of Chen and Han [CH90, CH96] can compute shortest paths from this intrinsic data in $O(n^2)$ time. Repeating this n times to reach Q_{flat} then can be achieved in $O(n^3)$ time.

The scaling step (3) can be accomplished in linear time, $O(n)$, as follows. The largest circle inscribed in P_{flat} is computed by the linear-time medial axis algorithm [CSW99]; say its radius is r_P. The smallest circle circumscribing Q_{flat} is found in linear-time via Megiddo's algorithm [Meg83]; say its radius is r_Q. Then scale Q_{flat} by $s = r_P/r_Q$.

Trimming P_{flat} to P^s_{flat}, step (4), can be accomplished in many ways. The method we described in Section 8.1.4 can easily be implemented in $O(n^2)$ time by ray-shooting the edge extensions, and then triangulating each convex polygon in linear time. Likely the ray-shooting could be reduced to $O(n \log n)$ time.

Reversing the Q vertex-merging steps, step (5), amounts to digon tailorings cutting of the merged vertices on P. This can easily be accomplished in $O(n \log n)$ time. Keeping track of the considered vertices and employed digons gives in the end a correspondence between Q^t and Q, and thus the 3D structure of Q^t.

So the whole algorithm time-complexity is dominated by the $O(n^4)$ cost of finding a guaranteed generic x.

Because the ridge-free regions are determined by overlaying n cut loci, the regions are delimited by a one-dimensional network of segments. Thus

choosing a random point x on P is generic with probability 1. That still leaves the algorithm requiring $O(n^3)$ time. We believe this time complexity could be improved, perhaps to $O(n^2)$. See Open Problem 18.1.

We repeat Theorem 8.1 with the complexity bound included:

Theorem 8.3. *For any given convex polyhedra P and Q, one can tailor P until it becomes homothetic to Q in time $O(n^4)$, where $n = \max\{|P|, |Q|\}$.*

Finally, we note that our model of reshaping in all cases excises just a single vertex via a tailoring step, or inserts a single vertex via a vertex-merge step. A rather different model, but related to the flattening algorithm above, unfolds P and Q each to nets (non-overlapping planar polygons), and places the unfolded Q inside the unfolded P. For example, Jin-ichi Itoh made the following interesting suggestion:[1] first star-unfold P to S_P and Q to S_Q, shrink S_Q to fit inside S_P, then cut out S_Q from S_P, and refold to obtain a homothet of Q from P.

[1] Personal communication, 2019.

Chapter 9

Enlarging and P-Unfoldings

In this chapter we show how to enlarge a convex polyhedron Q to some $P \supset Q$, via tailoring. Based on this operation, we introduce the notion of a P-unfolding and propose a few methods to accomplish it.

9.1 Enlarging and Reshaping

Previous chapters have established three methods of tailoring P to Q:

- Theorem 4.6: For $Q \subset P$, digon-tailor following a sculpting of P to Q.
- Theorem 7.5: For $Q \subset P$, crest-tailor following a sculpting of P to Q.
- Theorem 8.1: via flattening, digon-tailor P to a (possibly small) homothet of Q.

If we do not have $Q \subset P$, then shrinking Q until it can fit in P leads to a homothet of Q. All three approaches result in the homothet of Q being composed entirely of P-surface. In the tailoring-via-flattening algorithm, Q is enlarged with vertex-merging inserts, but the last steps remove all inserts.

In this short chapter, we explore some results that can be achieved through a mix of vertex-merge enlarging, and tailoring. Part II will explore vertex-merging more thoroughly.

Suppose $Q \subset P$. Then we can enlarge Q to P using any one of the three tailoring algorithms, as follows. We first tailor P to Q, tracking the cuts and digons removed. For crest-tailoring, we cut the boundary of crests. Then, starting with Q, we cut each sealed geodesic or crest-boundary and insert

the earlier removed corresponding digon or crest surface, in reverse order. The result is that Q is enlarged by the surface insertions (of P-surface) in the reversing process until it matches P.

If $Q \not\subset P$, so P and Q are of arbitrary relative sizes, we can reshape P to Q by first enlarging P to some $P' \supset Q$ large enough to enclose Q, and then tailor P' to Q. Here P' is an arbitrary polyhedron midway in the process $P \to P' \to Q$; it only needs to be large enough.

These enlargings can be accomplished within the same $O(n^4)$ time complexity of the tailoring algorithms. And in analogy with Corollary 4.7, enlargings of Q can approximate any target surface S.

We summarize the above discussion in a theorem:

Theorem 9.1. *Let P and Q be convex polyhedra. If $Q \subset P$, Q may be enlarged with surface insertions to P. If $Q \not\subset P$, P can be reshaped to Q with a combination of surface insertions and excisions. Either process can be accomplished in time $O(n^4)$, and can approximate non-polyhedral convex surfaces.*

9.2 P-unfoldings

For $Q \subset P$, we can view the enlarging of Q to P just described as "unfolding" Q onto P. We call this a *P-unfolding of Q*. To our knowledge, this notion has never been considered before. We explore it briefly here, and in more detail in Part II.

9.2.1 P-unfoldings and Reshaping

Unfoldings of convex polyhedra to a plane have been studied extensively; see, e.g., [DO07]. Particular attention has been paid to unfolding to few pieces, to connected unfoldings, and to non-overlapping unfoldings. Often the goal is to achieve a *net*, an unfolding to a single, simply connected, non-overlapping polygon in the plane.

Instead of unfolding a convex polyhedron $Q \subset P$ to a plane, consider the question: Can one cut-up the surface Q so that the pieces may be pasted onto P, non-overlapping, and so form an isometric subset of P? This is a *P-unfolding of Q*, or *an unfolding of Q onto P*. It can be achieved via enlarging: just enlarge Q to P, and then remove all inserted digons or crests. The result is a subset of P isometric to the cut-up Q. Note that all three enlarging

methods retain the entire surface of Q—the enlarging is accomplished by insertions of surface not part of Q.

This viewpoint incidentally yields a different proof of this known result:

Corollary 9.2. *Given convex polyhedra P and $Q \subset P$, the area of Q is smaller than the area of P.*

To justify the use of the term "unfolding," we describe an example of a simply connected P-unfolding embedding that does not achieve a net when unfolded to a plane. Let Q be the classical thin, nearly flat tetrahedron with an overlapping edge-unfolding. See, e.g., [DO07, Fig. 22.8, p. 314]. Take P to be a slightly larger homothet of Q. Then the same edge-cuts that result in overlap in the plane embed Q onto P without overlap. In a sense, it can be easier in some cases to unfold to P than to unfold to a plane.

One focus of interest in the unfolding literature is the "fewest nets" problem [DO07, Prob. 22.1, p. 308]: edge-unfold a polyhedron into the fewest number of nets. Of course one net would be ideal, but it is a long-standing open problem to determine if every convex polyhedron can be edge-unfolded to a net—Dürer's problem [O'R13]. This suggests it could be of interest to minimize the number of disconnected pieces of a P-unfolding of Q, a goal we pursue in Part II. See also Open Problem 18.4.

Example: Hexagon inscribed in a triangle. Let P be a doubly-covered triangle abc, and Q a doubly-covered hexagon inscribed in P, as in Fig. 9.1. Digon-tailoring P to Q is achieved by excising the vertices a, b, c by cutting digons bounded by the hexagon edges e_1, e_2, e_3 respectively, and then sealing the cuts closed. To enlarge Q to P by reversing the process, each of those three edges is slit open and the pair of triangles earlier removed is sutured back along the slit edges. To obtain a P-unfolding, we simply skip inserting the pair of triangles. The result is a "pasting" of the two hexagons inside the two triangular faces of P.

Note that, to reverse a sealed digon, the resulting geodesic γ is slit, but the cut does not delete the endpoints x_1, x_2 of γ, as depicted in Fig. 9.2. In other words, $\gamma \setminus \{x_1, x_2\}$ is doubled to reproduce the digon boundary, but there remains one copy each of the x_1, x_2 endpoints.

Notice that each split e_i is —bluea non-contractible closed curve inside the image of Q, so the embedding of Q onto P in this case is not simply connected: it has three holes.

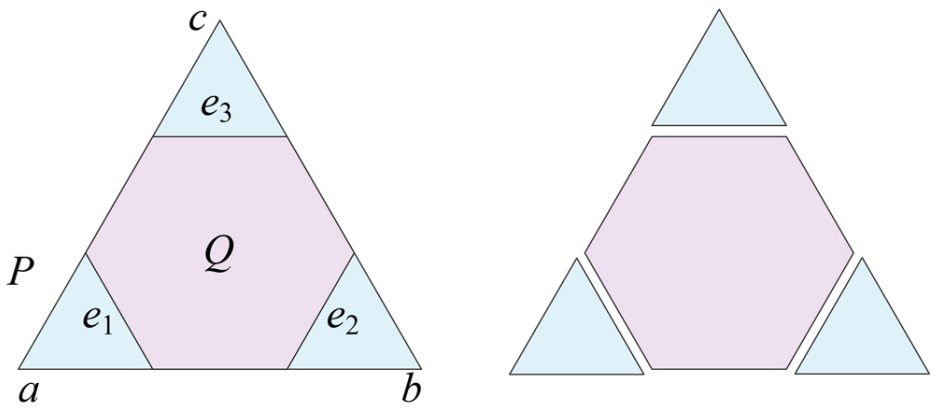

Figure 9.1: Doubly-covered hexagon Q inscribed inside / enlarged to a doubly-covered triangle P.

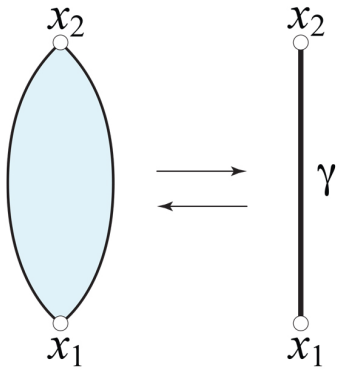

Figure 9.2: Digon x_1x_2 closes to geodesic γ.

The previous example shows that the P-unfolding of Q produced by enlarging is not necessarily simply connected.[1] However, in general, that is indeed the case, as shown by the following result. We next explain the meaning of "in general."

Consider the space $\mathcal{S}$ of all convex surfaces, endowed with the topology induced by the usual Pompeiu-Hausdorff metric. Fix some $P \in \mathcal{S}$. Consider in $\mathcal{S}$ the subset $\mathcal{P} = \mathcal{P}_P^n$ of all polyhedra $Q \subset P$ with precisely n vertices, with the induced topology. Two polyhedra in $\mathcal{P}$ are then close to each other if and only if they have close corresponding vertices. "General" refers to polyhedra Q in an open and dense subset of $\mathcal{P}$.

Theorem 9.3. *For any convex polyhedron P and any $n \in \mathbb{N}$, there exists a subset $\mathcal{Q} = \mathcal{Q}_P^n$ open and dense in $\mathcal{P}$, such that the P-unfolding Q_P of each $Q \in \mathcal{Q}$ is flat (i.e., contains no internal vertices), and is simply connected.*

Proof. Assume we have some convex polyhedron $Q \subset P$ such that Q_P contains an internal vertex v, and so the curvatures of P and Q at v are equal: $\omega_Q(v) = \omega_P(v)$. Slightly alter the position of the vertices of Q, to get $\omega_Q(w) \neq \omega_P(u)$, for any vertices $w \in Q$ and $u \in P$. Of course, this remains valid in a small neighborhood of the new Q.

Assume now we have some convex polyhedron $Q \subset P$ such that Q_P is not simply connected, i.e., Q_P contains a noncontractible curve $\sigma \subset P \cap Q_P$. The Gauss-Bonnet Theorem shows that the total curvature $\Omega_Q(\sigma)$ of Q_P inside σ equals the total curvature $\Omega_P(\sigma)$ of P inside σ: $\Omega_Q(\sigma) = \Omega_P(\sigma)$. We next show that every such Q that violates the theorem can be approximated with polyhedra that do satisfy the theorem.

Slightly alter the position of the vertices of Q, to get a new polyhedron Q' on which the following property (V) is verified. (V): any partial sum of vertex curvatures is different from any partial sum of vertex curvatures in P. This implies that, for any simple closed curve τ on Q', $\Omega_{Q'}(\tau)$ cannot be written as the sum of vertex curvatures of P. Therefore, Q'_P has no curve in common with P, noncontractible in Q'_P. And so Q' does satisfy the theorem.

Since the property (V) is valid on a neighborhood N of Q', it follows that all polyhedra in N do satisfy the theorem. □

[1] Notice that we do not require path-connectivity for the definition of simple connectivity.

9.2.2 P-unfoldings and the WBG Theorem

In this section we show that the P-unfolding question can also be answered by applying the powerful hinged variant of Wallace-Bolyai-Gerwien dissection theorem. However, this will result in a "pseudopolynomial number of pieces" and pseudopolynomial running time [AAC+12].[2]

The Wallace-Bolyai-Gerwien (WBG) Theorem states that any two simple polygons with equal area can be dissected into finitely many simple congruent polygons. It was strengthened in the paper, "Hinged Dissections Exist" [AAC+12], where it is shown that the pieces can be chosen in an arcwise-connected chain, i.e., the dissection can be hinged at vertices along that chain.

Theorem 9.4. *Given convex polyhedra P and Q of at most n vertices each, and $Q \subset P$, a connected P-unfolding of Q can be determined in pseudopolynomial number of pieces and pseudopolynomial running time, following the WBG theorem and the hinge-dissection.*

Proof. We start with a quote from [AAC+12]:

> "One interesting consequence of [our] theorem is that any finite set of polyhedral surfaces of equal surface area have a common hinged dissection: It is known that every polyhedral surface can be triangulated and then vertex-unfolded into a hinged chain of triangles [DEE+03]. Our results (specifically Theorem 6) show how to construct a single hinged chain that can fold into any finite set of such chains, which can then be folded (and glued) into the polyhedral surfaces."

A *vertex-unfolding* of a triangulated polyhedron is a chain of triangle connected at vertices. A vertex-unfolding of a triangulated cube is shown in Fig. 9.3.

Let A_Q and A_P be the areas of Q and P. Triangulate and vertex-unfold both P and Q. The plan is to show that A_Q can "fit" exactly within a sequence of P's triangles, with at most one of P's triangles partitioned into two triangles.

[2] They define pseudopolynomial as follows: "pseudopolynomial means polynomial in the combinatorial complexity (n) and the dimensions of an integer grid on which the input is drawn."

Denote by $T_1, ..., T_m$ the triangles in the chain corresponding to P, and by a_j denote the area of the T_j, $j = 1, ..., m$. So we have $A_Q < A_P = \sum_{j=1}^{m} a_j$. Then there exists $1 \leq h \leq m$ such that $a_1 + ... + a_{h-1} \leq A_Q < a_1 + ... + a_h$. Triangle T_h is then the transitional triangle with respect to area in the sequence of triangles.

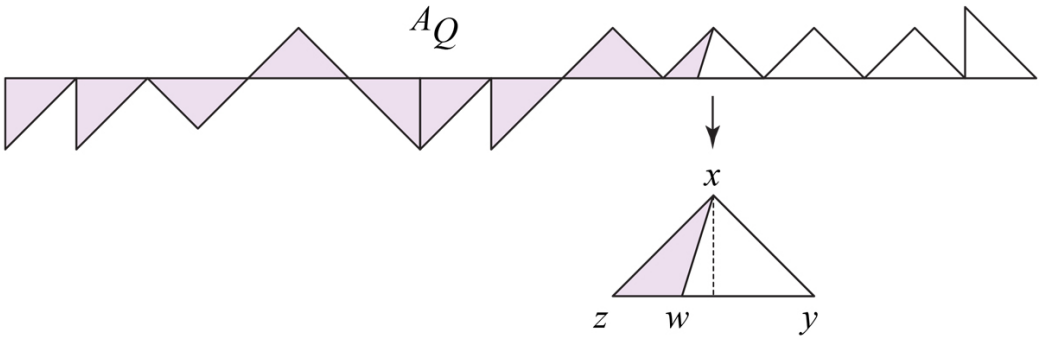

Figure 9.3: Area A_Q (shaded) in vertex-unfolding of a triangulated cube. $T_h = xyz$. Based on Fig.2(a) in [DEE+03].

Now we further dissect the triangle T_h into two parts, one of which has area $c = A_Q - (a_1 + ... + a_{h-1})$. Specifically, let $T_h = xyz$ and choose a point $w \in yz$ such that $|w - y|/|z - y| = c/a_h$. Then the area A' of the triangle xyw is equal to c/a_h of the area of T_h (because both triangles have the same altitude through x).

Therefore, in the chain of triangles obtained from P, we can find an arcwise-connected subchain of total area equal to A_Q, possibly after partitioning one triangle.

Then Theorem 6 in [AAC+12] constructs a single hinged dissection chain of the triangles forming Q into the triangles up to T_h of P. Then the Q-chain triangles can be placed inside the P-chain triangles, achieving a connected P-unfolding of Q.

The "pseudopolynomial number of pieces" and pseudopolynomial running time follows from [AAC+12]. □

Note that although the P-unfolding constructed in this theorem's proof is connected, its interior is not connected.

Part II

Vertex Merging and Convexity

Chapter 10

Introduction to Part II

In the second part of our monograph we focus on the operation of vertex-merging on a convex polyhedron P. Used before as a proof technique by several authors, starting with A.D. Alexandrov, this operation has never before been considered, to our knowledge, as a main object of study. The operation is defined in Chapter 2, Section 2.4, and we employed it in Part I, in Chapters 8 and 9. We explained in Chapter 8 how vertex-merging is in some sense the inverse of digon-tailoring: instead of cutting off a vertex v by excising a digon, vertex-merging slits a geodesic segment connecting two vertices v_1 and v_2, and inserts a pair of triangles that flatten v_1 and v_2 and introduce a new vertex v with curvature $\omega(v) = \omega(v_1) + \omega(v_2)$.

We called the geodesic segment that results from suturing-closed the digon boundary a *seal*. We will call the vertex-merge geodesic segment a *slit*.

A polyhedral surface S that admits no vertex-merging is *irreducible*. By Lemma 2.11, the irreducible polyhedra are isosceles tetrahedra and doubly-covered triangles.

Vertex-merging *reductions*—the processes of repeatedly reducing the number of vertices via merging, from P to an irreducible polyhedron—are studied in Chapter 11 together with their induced *slit graphs*. This can be seen as a theoretical exploration as well as a refinement of a proof idea employed in Chapter 8.

Particularly interesting are the cases when slit graphs are forests of trees, because then the unfolding (i.e., the image) P_S of P onto the irreducible surface S has connected interior. Even in this case, obtaining a net from P_S is not always obvious, as shown in Section 11.7. The overall goal, and the guiding thread, of this Part II, is to identify slit graphs that are forests of

trees.

Our basic algorithmic idea for finding such forests is a certain type of *spiraling* slit tree. We present this first in the planar limit case in Chapter 12. The analysis is based on planar convex hulls of finitely many points, viewed as vertices of zero curvature. The boundary of such a hull is equivalently obtained as the minimal length enclosing polygon of the given set of points.

Next, we aim to extend to convex polyhedra this idea of spiraling slit trees. Toward this goal, we partition the surface into two *half-surfaces*, sharing a simple closed quasigeodesic Q as a common boundary. Every convex polyhedron has at least three such quasigeodesics [Pog49][Pog73].

So we are led first to consider, on convex polyhedra, the notions of convexity and convex hull (in Chapter 13) and then the notion of a minimal length enclosing geodesic polygon (in Chapter 14). Although identical in two dimensions, these notions are not equivalent in our framework. As far as we know, they have not been investigated in detail on convex polyhedra. Given their importance in the plane, we believe they have a certain interest in their own right.

In Chapter 15 we show that the extension to half-surfaces (bounded by Q) of the planar spiraling algorithm works well for both the convex hull and the minimal length enclosing geodesic polygon, of all vertices inside a simple closed quasigeodesic.

The next step, in Chapter 16, is to join the two slit trees, previously obtained for half-surfaces. Thus we obtain unfoldings of the starting polyhedron P onto the union of two cones, or onto to a cylinder. The last case appears if Q contains at most two vertices, when rolling the cylinder on a plane leads to a net of P.

In the penultimate Chapter 17, we prove the existence of a non-empty and open set $\mathcal{Q}_{\leq 2}$ of convex polyhedra R, each polyhedron $R \in \mathcal{Q}_{\leq 2}$ having a simple closed quasigeodesic with at most two vertices. We also conjecture that such a quasigeodesic exists on all P, which would lead to a net of P by rolling.

We end with a list of open problems from both Part I and II in Chapter 18.

Chapter 11

Vertex-Merging Reductions and Slit Graphs

In this chapter we initiate the systematic study of consecutive operations of vertex-merging, already used in Chapter 8. We introduce vertex-merging reductions and their associated slit graphs, and derive their basic properties for later use. Several examples and open problems are discussed. Our main goal, carried out over several chapters and topics, is to relate vertex-mergings to unfoldings and nets. First steps toward this goal are reached in Theorem 11.5 and the discussion that follows in Section 11.6.

11.1 Slit Graphs for Vertex Mergings

Recall from Section 2.4 and the Introduction to Part II that merging of vertices v_1 and v_2 is only possible (by Lemma 2.11) when their curvatures ω_1, ω_2 satisfy $\omega_1 + \omega_2 < 2\pi$. The *vertex-merging irreducible* surfaces (in short, vm-irreducible) are doubly-covered triangles and isosceles tetrahedra.

A *vertex-merging reduction*[1] (in short, a vm-reduction) of a convex polyhedron P is a maximal sequence of consecutive vertex-merging steps starting from P; maximal, in the sense that it reduces P to a vm-irreducible surface S.

Explicitly, during a reduction process we obtain a sequence of convex polyhedral surfaces $P = P_0, P_1, \ldots, P_k = S$, with S vm-irreducible, and

[1]The term "reduction" refers to the number of vertices. The surface area increases at each step.

piecewise isometries $\iota_j : P_j \to P_{j+1}$. Of course, each ι_j is a polyhedral-unfolding of P_j onto P_{j+1}, as is $\iota = \iota_{k-1} \circ \iota_{k-2} \circ \ldots \iota_0 : P \to S$, where $\circ$ is the composition of functions. Put $P_S = \iota(P) \subset S$.

With some abuse, in the following we shall sometimes consider all isometries above as identities, thus identifying P_j and $\iota_j(P_j) \subset P_{j+1}$, for $j = 0, \ldots, k-1$.

The *slit graph for a vertex-merging reduction* $\iota = \iota : P \to S$ is the trace on P of all geodesic segments used during the reduction, called *slits*. Explicitly, at each reduction step we use a geodesic segment γ_j on P_j along which we merge two vertices of P_j. Notice that $\bar{\iota}_j = \iota_{j-1} \circ \ldots \circ \iota_0 : P \to P_j$, hence the inverse $(\bar{\iota}_j)^{-1} : P_j \to P$ maps P_j to P. The slit graph Λ is the union for all $j = 0, \ldots, k-1$ of the slits $\lambda_j = \overline{\iota_j}^{-1}(\bar{\iota}_j(P) \cap \gamma_j)$. Note that only portions of a slit γ_j on P_j might lie on P and so be part of $\Lambda \subset P$.

The importance of studying slit graphs will become apparent later; see Section 11.6 and Chapter 16.

Consider now the vm-reduction $\iota : P \to S$ and its inverse process, from the irreducible surface S to P. Explicitly, at each step of this reverse process, we tailor a digon (the one inserted for vertex-merging). Then the slit graph Λ of P, for the vm-reduction ι, is precisely the seal graph Σ for the corresponding tailoring of S to P, studied in Chapter 5.

The next result shows that the slit graph is composed of as many geodesic segments as steps in the reduction process.

Lemma 11.1. *For any vertex-merging reduction of any convex polyhedron, all slits are non-degenerate geodesic segments.*

Proof. Vertex merging is an intrinsic process. The proof idea below is based on a convenient way of viewing it.

Assume we executed j steps from a reduction process, starting from P and reaching P_j. Consider a geodesic segment γ between vertices $v_1, v_2 \in P_j$ with $\omega_1 + \omega_2 < 2\pi$. Also consider geodesics γ_1, γ_2 starting at v_1 and v_2 such that, together with γ, they bisect the complete angles at v_1 and v_2, respectively.[2] Also consider a neighborhood N of the polygonal path composed by γ_1, γ and γ_2, containing no vertex of P excepting v_1, v_2. See Fig. 11.1(a). Then,

[2]If γ is the unique geodesic segment joining v_1, v_2 then γ_1 contains a geodesic segment included in $C(v_2)$, and similarly for γ_2.

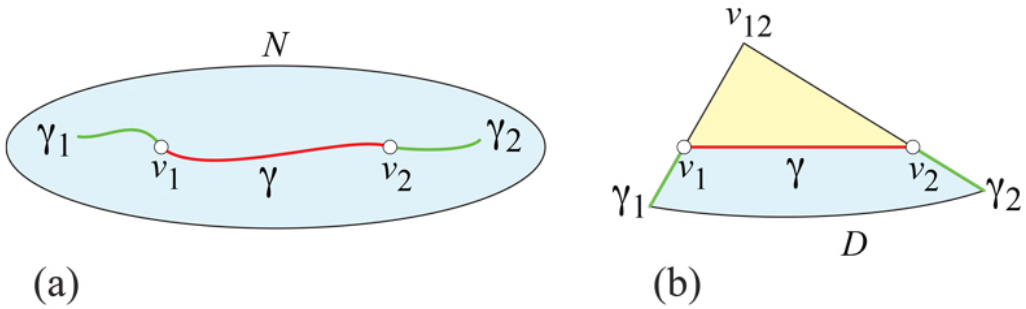

Figure 11.1: (a) Neighborhood N of geodesic γ connecting v_1 to v_2. (b) N folded to a doubly covered domain D.

after merging v_1 and v_2, N is isometric to a doubly covered planar domain D, as in (b) of the figure.

Retain the same notation for D. Merging the vertices v_1 and v_2 can thus be represented in the plane, by the use of D. There, it means extending the edges edges γ_1, γ_2 beyond v_1, v_2, until their intersection at v_{12}. So each newly created edge (say v_1v_{12}) is the extension of an edge of D, which derives from a geodesic on P_j (γ_1 in this case). That is, v_1v_{12} is included in a geodesic which cannot be new from one end to the other.

Applying the previous reasoning to γ, which is already maximal with respect to inclusion, shows that it contains a geodesic subsegment already existing on P_{j-1} and, inductively, we find a (possibly smaller) geodesic subsegment of γ existing on P. □

A nearly immediate consequence of the preceding lemma is this:

Corollary 11.2. *For a polyhedron P of n vertices, the collection of slits (the slit graph) cuts P into at most $O(n^2)$ pieces.*

Proof. Because each slit is a geodesic segment, each pair of slits cross at most once. So the slits are pseudo-segments, subsegments of pseudo-lines. (*Pseudo-lines* are curves each pair of which intersects at most once, where they properly cross.) With $O(n)$ slits, the arrangement of pseudo-segments on the planar surface of P has combinatorial complexity $O(n^2)$, and in particular, has at most $O(n^2)$ cells [AS05]. □

Despite the possibly quadratic complexity of the complement of the slit graph, in the end the slit graph has only 3 or 4 components, as established in the following lemma.

Lemma 11.3. *The slit graph of any vm-reduction process $\iota : P \to S$ has at most $n_S \in \{3, 4\}$ components, where n_S is number of vertices of S.*

Proof. Every slit $\lambda \in \Lambda$ is included in a geodesic γ whose isometric image $\gamma_j \subset P_j$ creates a merge-vertex outside of P. With some abuse, say $\lambda = \gamma_j \cap P$. By Lemma 11.1, λ is a non-degenerate geodesic segment, hence it is included in ∂P_S. But S has at most n_S vertices, and each component of $S \setminus P_S$ has at least one vertex, hence $S \setminus P_S$ has at most n_S components. The boundary of each component is a geodesic polygon Γ (see Figure 11.2(a)), so its inverse image $\iota^{-1}(\Gamma)$ is a connected subgraph of Λ containing λ. Therefore, Λ has at most n_S components. □

11.2 Example: Reductions of Flat Hexagon

Consider a doubly-covered regular hexagon $H = x_1x_2 \dots x_6$. The curvature at each vertex is $4\pi/6 = 2\pi/3$. We describe two different merge sequences.

Assume first that we vertex-merge vertices two-by-two in increasing order of indices, so we flatten the six original vertices and create three new vertices x_{12}, x_{34}, x_{56}, each of curvature $\omega_{12} = \omega_1 + \omega_2 = 4\pi/3$. The resulting surface is a doubly covered equilateral triangle of area 1.5 times larger than the area of H. See Fig. 11.2(a,b). The slit graph in this case is a forest composed of three single-edge trees, every other edge of H. This shows that the slit graph for a vertex-merging reduction is not necessarily connected.

Note that for this example, digon-tailoring the equilateral triangle in (b) to H via sculpting produces the same seals as the digon-tailoring via vertex-merging produces slits.

Notice that we cannot vertex-merge x_{12} with x_j, $3 \le j \le 6$, because the resulting curvature would be exactly 2π, violating the necessary merge condition. Therefore, no matter in which order we vertex-merge the vertices $x_3, \dots, x_6$ in pairs, the resulting slit graph will be a forest of three single-edge trees. Similar reasoning shows that the slit graph for any vertex-merging reduction of H is never connected.

On the other hand, for the example presented earlier in Fig. 2.6, the slit graph is a single-edge tree, hence connected. Moreover, each vm-reduction sequence of a non-isosceles tetrahedron consists of a single vertex-merging operation, yielding a doubly covered triangle. So for these surfaces, the slit graph for a vertex-merging reduction is always connected.

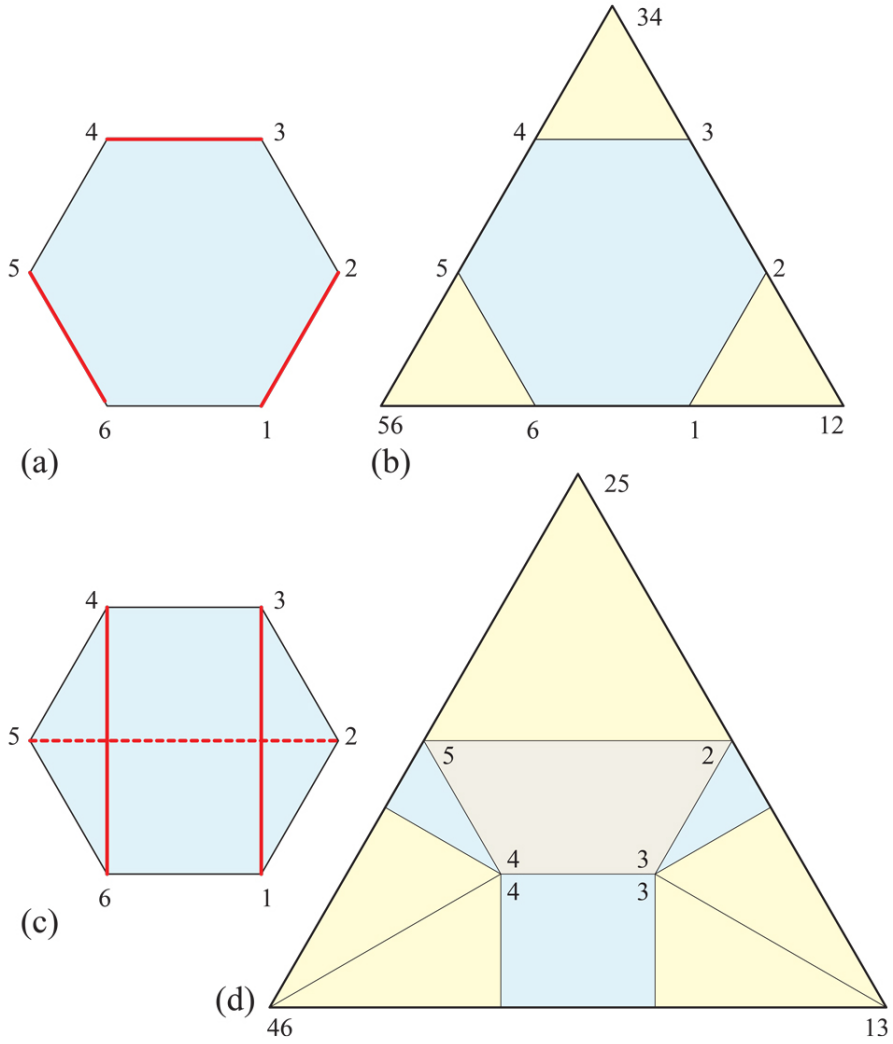

Figure 11.2: (a) Doubly-covered hexagon H with slits marked (red). (b) The result after three merges. (c) Same H but different merge sequence. (d) Final equilateral triangle after vertex-merging (not to same scale as (c)). Blue indicates front surface of the original H, tan back surface, and yellow insertions. The back face is symmetric to the front face. The images of the vertices x_1 and x_6 are on the back face, opposite to v_3 and v_4, respectively.

This leads us to consider the following question, a partial answer of which will be given in the following sections.

For which convex polyhedra does there exist a vm-reduction whose slit graph is connected?

Next we examine a different vertex-merge ordering: merging x_1 and x_3 on the front of H, x_2 and x_5 on the back, and x_4 and x_6 on the front. This produces another doubly covered equilateral triangle, say T,[3] of area 2.5 times larger than the area of H. See Fig. 11.2(c,d).

In particular, the P-unfolding (in the sense of Chapter 9) H_T of H onto T is not simply connected, as for example, the boundary of the v_1v_3 slit is a non-contractible cycle; see Fig. 11.2(d). However, H itself has not been disconnected by the vertex-merging slits.

Clearly for any given P, different vertex-merging reductions of P lead to different surface areas of S. We leave to Open Problem 18.7 exploring bounds on, or how to achieve, the min or max surface area.

11.3 Example: Reductions of Cube

In this section we present two vm-reduction processes for the cube.

The first reduction, illustrated in Fig. 11.3, reduces the cube to a doubly covered square, which is a degenerate isosceles tetrahedron and so vm-irreducible. In any order, slit the top front and back edges, and the bottom front and back edges, and vertex-merge their endpoints. The curvature at each cube corner is $\pi/2$, so the triangle inserts are each isosceles right triangles, a pair of which flatten the vertices to which they are incident. The slit graph is a forest of four single-edge trees.

The second merge example is more intricate; it is illustrated in Fig. 11.4, and will be revisited in Chapter 16. It reduces the cube to a doubly covered triangle, (b) of the figure. We now describe the reduction.

First v_7, v_8 are merged to produce v_{78}, just as in the previous example (although in (a) of the figure the triangle inserts are shown coplanar with the top T of the cube). Next v_{78} is merged with v_5, creating v_{578}. Note that the slit geodesic v_5v_{78} crosses the slit v_7v_8, and so the inserted triangle

[3]The intermediate shapes guaranteed by AGT are, however, not flat as they were in the previous example, but rather have positive volume. Only the final shape is flat.

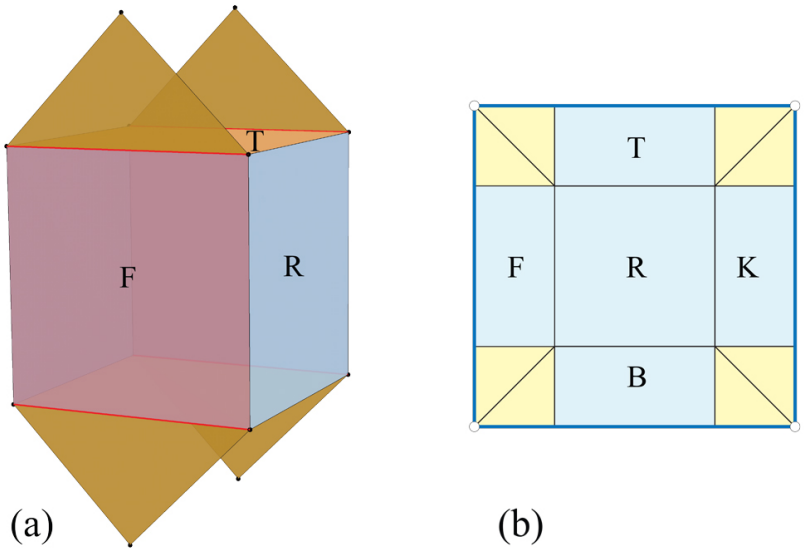

Figure 11.3: (a) A cube and its vm-reduction. (b) The corresponding vm-irreducible surface. The pattern on the back side is similar, except with $R \to L$. Inserted triangles in yellow.

pair separates the top $\triangle v_7 v_8 v_{78}$ (but not its mate underneath). It is not straightforward to track the consequence of each insertion.

The same two merges are mirrored on the cube bottom face B: v_1 and v_2, and v_3 with v_{12}. Finally, v_4 and v_6 are merged. So the five merges have reduced the eight cube vertices to just three: a doubly-covered triangle. It is shown unfolded in Fig. 11.4(b). Note that the three vertices of this triangle, v_{578}, v_{123}, v_{46}, are each surrounded by a polygonal domain of triangle inserts. This follows because all the cube vertices have been flattened.

The slit graph in this example has three components, satisfying Lemma 11.3.

11.4 Example: Icosahedron

We next detail using vertex-merges to reduce an icosahedron to a doubly covered triangle, without disconnecting the icosahedron surface. We will revisit and modify this example in Chapter 16, where further details will be provided.

We label the 12 vertices of P as shown in Fig. 11.5. We will merge in

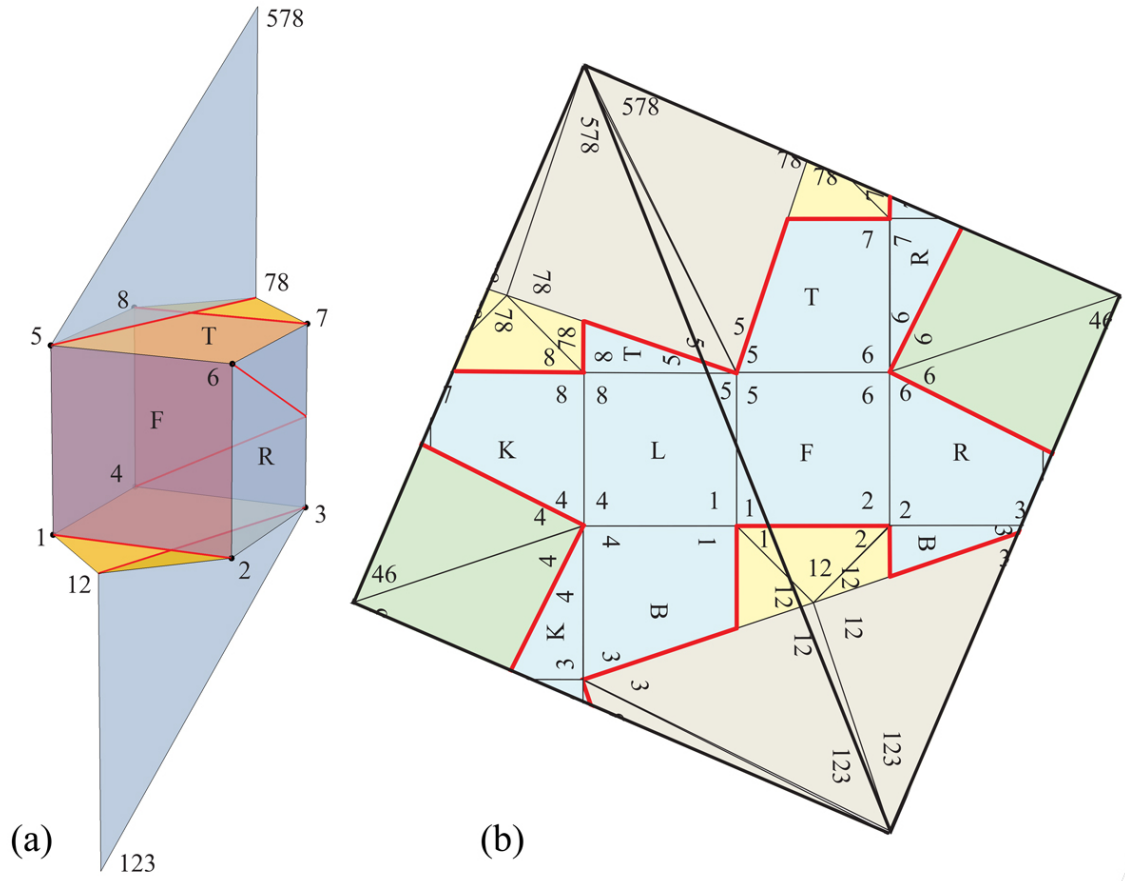

Figure 11.4: (a) Slits in red. Vertex merging: $1+2 \to 12$. $7+8 \to 78$. $5+78 \to 578$. $3+12 \to 123$. $4+6 \to 46$. (b) The doubly covered triangle has vertices at $578, 123, 46$. The triangle is a right isosceles triangle. Cube surface in blue. Image of the slit graph in red.

sequence five vertices, v_1, v_2, v_3, v_4, v_6, a merge affecting three of the triangles incident to the top vertex v_6. Symmetrically, we merge five bottom vertices, $v_{11}, v_{10}, v_9, v_8, v_{12}$. Each of these sequential merges results in a merge vertex of curvature $\frac{5}{3}\pi$, the sum of five $\frac{1}{3}\pi$ curvatures. Finally, we merge v_5 and v_7, cutting across the middle band of triangles, creating a merge vertex of curvature $\frac{2}{3}\pi$. (So the total curvature is $(\frac{5}{3} + \frac{5}{3} + \frac{2}{3})\pi = 4\pi$, satisfying the Gauss-Bonnet theorem.) The resulting doubly covered triangle abc has angles $30°, 30°, 120°$.

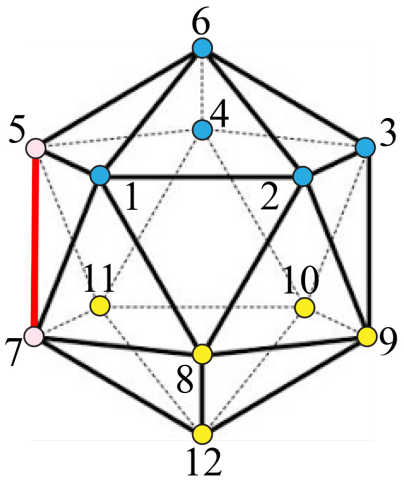

Figure 11.5: Labels i for vertices v_i.

We now detail the top sequential merge of five vertices. Each of the four merges i is accomplished by inserting two copies of a triangle T_i, whose apex is the merge vertex m_i.

$$\begin{aligned} v_1 + v_2 &\to m_1 \ , \ T_1 \\ m_1 + v_3 &\to m_2 \ , \ T_2 \\ m_2 + v_4 &\to m_3 \ , \ T_3 \\ m_3 + v_6 &\to m_4 \ , \ T_4 \ . \end{aligned}$$

(Here we are using '+' and '$\to$' informally to mean 'merge' and 'creating' respectively.) The merge of m_i to the next vertex v_{i+2} slits a geodesic γ_{i+1} from m_i, down one copy of T_i, onto P at the point we call m_i'. This geodesic crosses the previous geodesic cut γ_i from m_{i-1} to v_{i+1} at m_i'. So m_i' is a point

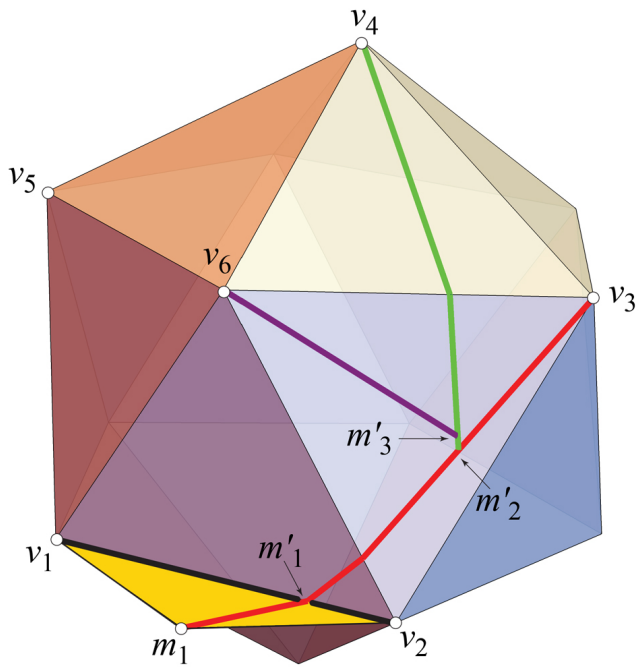

Figure 11.6: Four geodesic slits on P, each entering P at m'_i. v_6 is the top vertex of P.

on P, whereas m_i is not on P, but rather on P_i, the intermediate polyhedron after the i-th merge. In Fig. 11.6 illustrates the four merge cuts on P, with just T_1 shown to illustrate γ_2 crossing from T_1 onto P at m_1'.

Returning to Fig. 11.5, we merge five vertices on the top side of P, and symmetrically five on the bottom side, and two connecting across the "equitorial" band of triangles. These merges reduce the original 12 vertices to 3, so the result is a doubly covered triangle. The slit graph is a forest of three trees.

The resulting doubly covered triangle abc is shown in Fig. 11.7, cut open so that both sides can be seen. It may not be obvious, but the white regions in Fig. 11.7 form a non-overlapping net of the icosahedron, when, for example, the doubly-covered triangle is cut along edges ac and bc as illustrated.

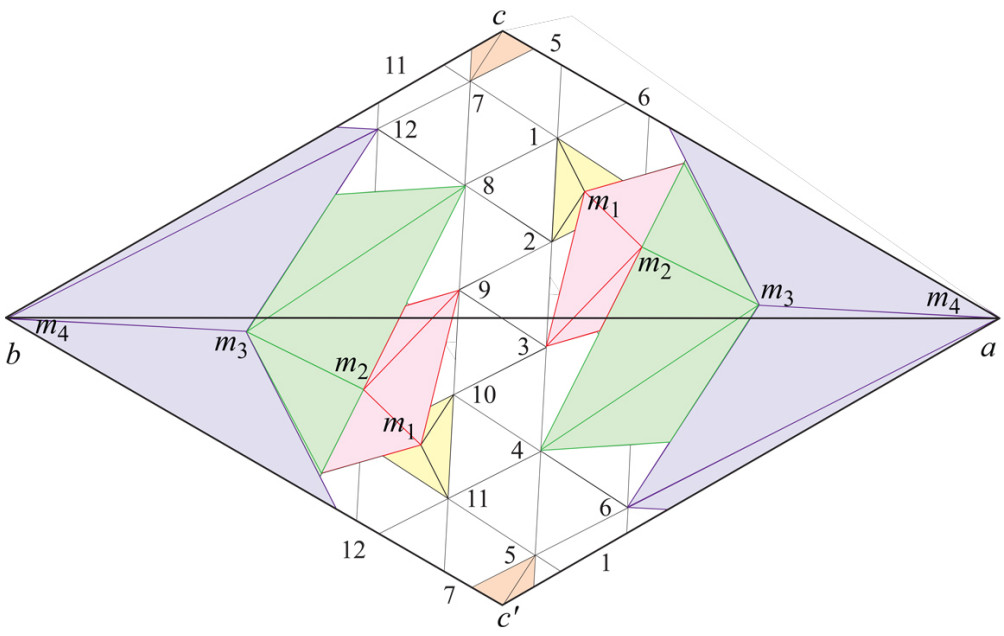

Figure 11.7: Icosahedron surface is white, triangle inserts (merge domains) colored. $\triangle abc'$ is the back side of $\triangle abc$, with $c = c'$ identified.

11.5 Example: Hexagonal Shape with Cycle

The next example shows that the slit graph could have a cycle. P is the convex hull of two similar hexagons in parallel planes, separated by a distance h. Label the top hexagon with vertices $v_1, \ldots, v_6$, and the larger base hexagon $x_1, \ldots, x_6$. See Fig. 11.8. Note that as $h \to 0$, the curvature $\omega(v_i) \to 0$. Let

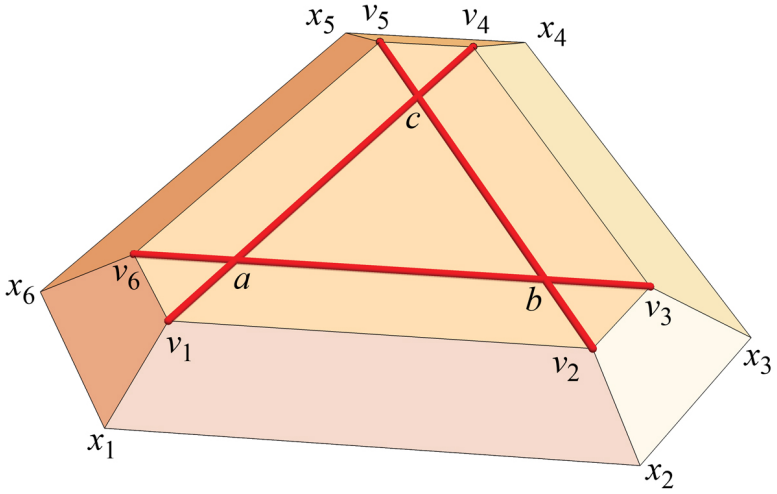

Figure 11.8: Three vertex-merge slits (red) will form a cycle abc.

$h = \varepsilon$ be a small positive height, so $\omega = \omega(v_i)$ is small.

Now we are going to merge $v_1 + v_4$, then $v_2 + v_5$, and finally $v_3 + v_6$, where again $+$ means "merge." The merge of v_1 and v_4 is accomplished by a pair of triangles with angles $\omega/2, \omega/2, \pi - \omega$. With ω small, the inserted digon $D(v_1, v_4)$ is very narrow, akin to a "fat" edge with endpoints v_1 and v_4. Then the geodesic between v_2 and v_5 crosses near c over $D(v_1, v_4)$, but because that is narrow, the geodesic does not look too different from the v_2v_5 edge illustrated in the figure—the geodesic "jags" slightly as it crosses $D(v_1, v_4)$. Again the digon $D(v_2, v_5)$ is narrow, and so does not greatly deviate the geodesic from v_3 to v_4, which now crosses both previously inserted digons, near b and a.

After these three vertex-merges, P_3 is a 9-vertex polyhedron. It should be clear that the roughly triangular region abc in the figure is disconnected by the three slits, which form a cycle in Λ_3. Further slits cannot "repair" the disconnecting cycle, so the final Λ will have at least one cycle.

This leads us to consider the following question: *For which convex polyhedra do there exist vm-reductions whose slit graphs have no cycle?* Although we pursue this question in subsequent chapters, it remains unresolved. See Open Problem 18.8.

11.6 Vertex Merging and Unfoldings

The goal of this section if to relate vm-reductions and their slit graphs to unfoldings.

An unfolding of a convex polyhedron P cuts the surface along a spanning tree of the vertices, producing a polygon U_P when developed in the plane. See, e.g., [O'R13]. When that polygon is *simple*—non-self-intersecting, i.e., non-overlapping—then U_P is called a *net* for P, i.e., an injective embedding of P into the plane.

Consider a reduction process $\iota : P \to S$ of P onto the vm-irreducible surface S. Then $P_S = \iota(P) \subset S$ is an unfolding of P onto S in the sense of Chapter 9. Further unfolding S to the plane (there are in general many ways to do this) provides an unfolding U_P of P in the plane.

We recall here some definitions for subsets related to convex polyhedra.

- A *simple closed curve* is a closed curve without self-intersections (i.e., homeomorphic to a circle).
- A *(geodesic)* polygon is a simple closed curve composed of (geodesic) segments.
- A *domain* is a connected open set.
- A *polygonal domain* is a the closure of a domain whose boundary is a finite union of polygons.
- A *simple polygonal domain* is a polygonal domain with one boundary component.
- By a *simply connected set* we understand a set each path-connected component of which is contractible; so the set itself is not required to be connected.

Lemma 11.4. *Consider a partial vm-reduction $\iota_j : P \to P_j$, for some $1 \leq j \leq k$, and identify, for simplicity, $P = \iota_j(P) \subset P_j$. Every component of $P_j \setminus P$ is a simple polygonal domain containing precisely one vertex. In particular, the conclusion holds for $S \setminus P_S$, where S is the vm-irreducible end result of the reduction process.*

We have seen this lemma verified for $S \setminus P_S$ in the series of examples above: Figs. 11.2(b,d), 11.3(b), 11.4(b), and 11.7.

Proof. The proof is a simple induction on the step m of the vm-reduction sequence.

At step $m = 1$, we merge the first two vertices of P, and $P_1 \setminus P$ consists of precisely the pair of inserted triangles, with one vertex, the shared apex of those triangles, outside P.

Notice that each (partial) vm-reduction ι_j yields a forest F of binary rooted trees, every node of which is a vertex in some P_l, with $1 \leq l \leq j$. (These vertex trees are not to be confused with slit trees.) After a complete reduction, the 3 or 4 vertices of S are the roots of the 3 or 4 vertex-trees that were merged to produce those root vertices. Precisely, the leaves of F are the vertices of P, and two nodes in F have a common parent if and only if they are merged at some step l. If a vertex $v \in P$ has not been merged, it is an isolated node of F. If v has been merged, then it has been flattened: $\omega(v) = 0$.

Assume now that the conclusion holds for the first $m - 1$ steps. At step m we merge, say, the nodes n_1 and n_2 of F, along the geodesic segment γ in P_{m-1}. Note n_i could be a vertex of P, or a vertex in a domain of $P_{m-1} \setminus P$ resulting from a merge. If n_i is a vertex of P, then it is not surrounded by a polygonal domain; for simplicity, we assume in this case that the polygonal domain is $\{n_i\}$ itself. Cutting along γ and inserting the curvature triangles merges the simple polygonal domains around n_1 and n_2 into a larger simple polygonal domain, containing precisely the resulting node n_{12}. □

For example, in Fig. 11.4, when $n_1 = v_5$ and $n_2 = v_{78}$, the domain around v_{78} (yellow in (b) of the figure) is merged with the new curvature triangles (tan) to form one simple polygonal domain containing $n_{12} = v_{578}$. In Fig. 11.7, the domains around triangle vertices a and b are the result of four merges, colored in the figure.

Theorem 11.5. *Consider a reduction process $\iota : P \to P_S \subset S$ of P onto the vm-irreducible surface S, resulting in slit graph Λ.*

- *If Λ is a forest of trees then P_S is a polygonal domain in S.*
- *If Λ is connected then P_S is simply connected.*
- *If Λ is a tree then P_S is a simple polygonal domain in S.*

Proof. Each slit in Λ appears twice in the boundary of P_S, once for each bank. Since P has no boundary, the whole boundary of P_S is produced in this way, hence it is a finite union of segments.

It also follows that each tree component of Λ yields a boundary component of P_S, which is a polygon. Therefore, if Λ is a tree then P_S is a simple polygonal domain in S. And if Λ is a forest of trees then P_S is a polygonal domain.

If a connected component of $P_S \subset S$ is not simply connected then it has several boundary components, impossible if Λ is connected. □

Next we offer a topological viewpoint. View P as a topological sphere with conical-point vertices. A vertex-merge of v_1 and v_2 is a topological-circle hole cut in P, passing through v_1, v_2, which is then filled with a topological disk—the two back-to-back triangles with a new vertex v inside the disk. Two connected vm-slits on P correspond topologically to merging two holes to one hole. A tree of vm-slits is then topologically a single hole, filled with doubly-covered triangles. Thus a forest of trees is topologically a collection of disjoint holes on P. As all holes are bounded by geometric segments, the result is a polygonal domain with polygonal holes.

Assume now that Λ is a tree, so P_S is a simple polygonal domain in S by Theorem 11.5, where S a doubly-covered triangle or an isosceles tetrahedron. With an appropriate unfolding of S, P_S remains connected and thus becomes an unfolding of P in the plane. However, the unfolding may overlap, and so not constitute a net. This is explained in the next section.

11.7 Unfolding Irreducible Surfaces

Assume we have a vm-reduction $\iota : P \to P_S \subset S$ resulting in a tree slit graph Λ so that P_S is a simple polygonal domain of S, by Theorem 11.5. We explore next the question of unfolding S to the plane, preserving connectedness of P_S. We start with S a doubly-covered triangle.

11.7.1 S: Doubly-covered Triangle

Let u_1, u_2, u_3 be the three vertices of S. Each is surrounded by polygonal domains M_1, M_2, M_3 created by repeated triangle inserts. Two examples of these domains are the cube vm-reduction in Fig. 11.4(b) and the icosahedron vm-reduction in Fig. 11.7. Both are displayed with their triangles cut open.

Now we argue that S can be cut open via a spanning tree without disconnecting P. So this provides an unfolding of S that preserves P as a single piece, as in those two figures. Although we can guarantee a single piece, we have not established this piece is a net, avoiding overlap. See Open Problem 18.9.

Topologically, S is a sphere with three conical points / vertices u_i, each surrounded by a *merge-domain* M_i, with the remainder of S the original P: in loose notation, $P = S \setminus (\cup_i M_i)$. Each M_i is topologically a disk, as depicted in Fig. 11.9(a).

There are only two combinatorially distinct spanning slit trees for unfolding S: a path u_1, u_2, u_3 (and its index permutations), or a Y-tree: some point x in P with slits from x to each u_i. (In both Figs. 11.4(b) and 11.7, the slit tree is a path.) For either spanning tree, a slit path ρ must enter M_i to reach u_i, and in the path case, the slit path must also exit M_i.

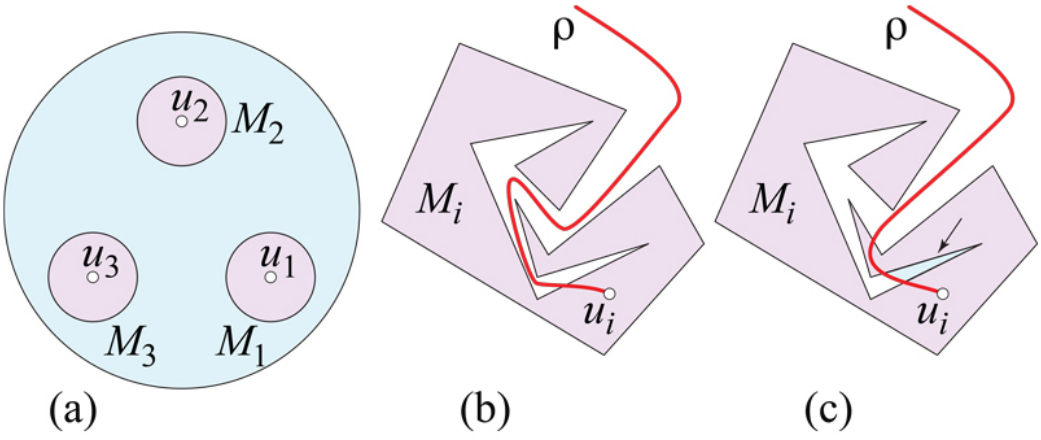

Figure 11.9: (a) Topology of S. (b,c) Slit path ρ to u_i. In (c), the marked region is disconnected from P.

Because each M_i is a disk, there is a ρ that crosses ∂M_i just once on entrance, and in the path case, once again on exit. Fig. 11.9(b) illustrates

such a one-cross ρ. Fig. 11.9(c) shows that, if ρ crossed ∂M_i at more than one point, then a piece of P is disconnected. But because the spanning slit tree can be chosen to avoid the (c) situation, we are guaranteed that we can unfold S to the plane so that $P \subset S$ remains a simply connected polygonal domain in the plane.

11.7.2 Net and Overlap

Let us temporarily ignore the structure of $S = P \cup (\cup_i M_i)$, and just view S as a doubly covered triangle $\triangle$. If, for every spanning slit tree T, the unfolding of $\triangle$ avoids overlap, then we immediately have that the unfolding of P avoids overlap, and so we have found a net for P.

Alas, this is not true for every T, as the example in Fig. 11.10 shows. In (a) the u_2, u_3 slit path spirals around u_3 from the back to the front. The

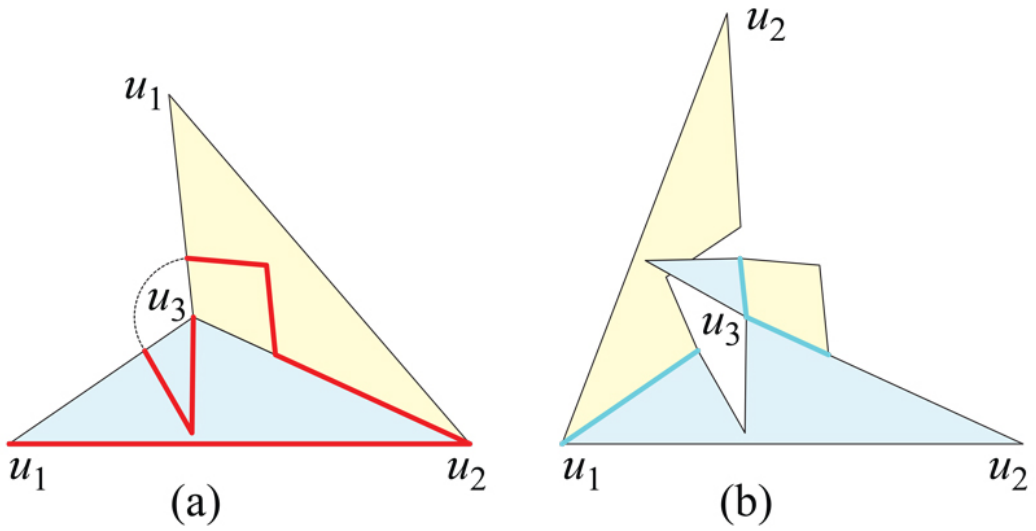

Figure 11.10: (a) Front: blue; back: yellow. Slit path T: (u_1, u_2, u_3). (b) Unfolding after cutting T. Blue segments remain uncut.

unfolding is shown in (b), with the uncut segments highlighted, showing clear overlap.

Although this may seem a contrived example, the merge regions M_i might be created from spiraling slit trees, as we will discuss in the the following chapters, and could be quite complex, as the icosahedron example (Fig. 11.7)

suggests. A natural slit path follows the merge vertices into each M_i region: $\mu = (m_0, m_1, m_2, \ldots, m_i)$, where $m_0 = u_1$. However, we leave it as an open question (Open Problem 18.9) of whether forming a slit tree by following μ for each M_i—or any other slit tree—unfolds P to a net.

11.7.3 S: Isosceles Tetrahedron

Vertex-merging could result in an isosceles tetrahedron S (rather than a doubly covered triangle). This occurred in the first cube reduction, illustrated in Fig. 11.3(b). This does not change the topological picture Fig. 11.9(a) significantly: just a fourth merge-domain disk M_4 is added, and there are two more possible spanning tree structures. Therefore, by similar reasoning, we can ensure a single-piece unfolding of P to the plane. But again, we leave open the question of whether there is always a slit tree that leads to a nonoverlapping single piece, a net for P.

Chapter 12

Planar Spiral Slit Tree

The previous chapter showed that if the slit graph Λ of a vm-reduction is a tree, then we can unfold P to the plane, and possibly to a non-overlapping net. So we have made a goal of finding a vm-reduction ordering that results in Λ a tree. Our overall plan is to partition P into two "halves" via a simple closed quasigeodesic Q, and then to vm-reduce the vertices in each half separately. (As mentioned in Chapter 10, such quasigeodesics always exist.) Let V be the set of vertices inside Q. We will eventually show (in Chapter 15) how to vertex-merge all of V so that the slit graph of that half is a tree. In this chapter we illustrate the main idea by assuming V lies in a plane.

12.1 Sequential Spiral Merge

Let $V = \{v_1, v_2, v_3, \ldots, v_n\}$ be the vertices in the above-half P^+ of polyhedron P. Define a *sequential merge* as a series of vertex-merges where at each step, the previous merge vertex is connected to a vertex of V. Thus the merge region—the union of all the triangle inserts—grows a single connected domain enclosing, in the end, one vertex by Lemma 11.4.

There are many possible sequential merge orders. We choose one that we call a *spiral merge*. See ahead to Fig. 12.5 for why "spiral" is an appropriate term. In the planar situation, we view vertices as having zero curvature, so that P^+ becomes planar. In this context, a slit is a line segment s, wholly in P. Each triangle insert degenerates to just s, and creates a merge vertex somewhere along s. We will show in Theorem 12.1 that a spiral merge

algorithm in two dimensions results in Λ a slit tree.

12.2 Notation

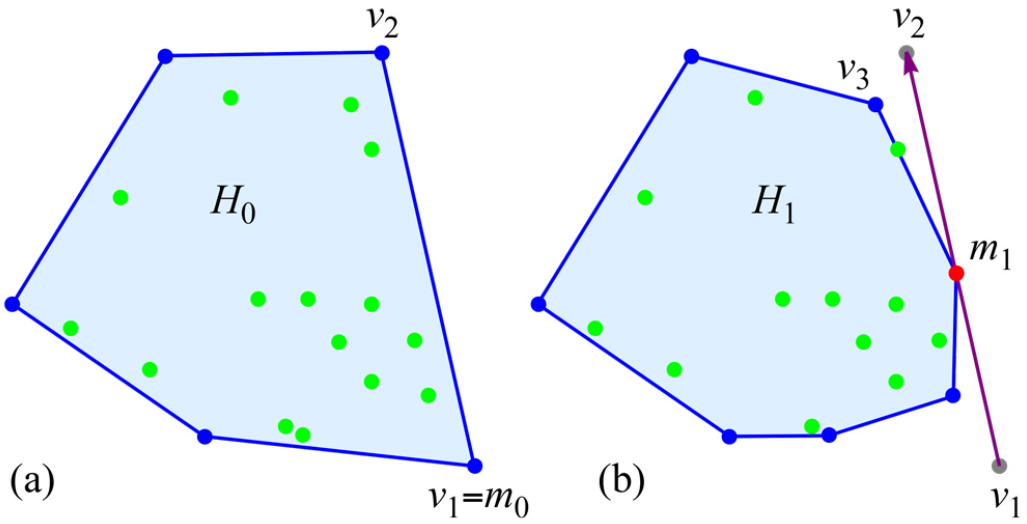

Figure 12.1: v_1 is merged with v_2 to produce m_1. $H_1 = \text{conv}(V_1)$ where $V_1 = V_0 \setminus \{v_1, v_2\} \cup m_1$.

The detailed argument needs considerable notation beyond that introduced earlier, which we gather below for reference.

- $V = V_0 = \{v_1, v_2, v_3, \ldots, v_n\}$: Vertices in the plane. (Later (in Chapter 15) these will be vertices on the surface of the polyhedron P.)
- m_i: merge vertex, the vertex created by merging m_{i-1} with v_{i+1} with $m_0 = v_1$. We sometimes abbreviate this as $m_{i-1} + v_{i+1} \to m_i$.
- $s_i = m_{i-1}v_{i+1}$. The i-th slit/merge segment. The new merge vertex m_i lies on s_i.
- So each slit segment has three labeled points: m_{i-1}, m_i, v_{i+1}.
- v_1 is also given the label m_0, so $s_1 = v_1v_2 = m_0v_2$.

- m_i can lie anywhere along s_i. In Figs. 12.1-12.5, m_i was chosen at a random point on s_i.
- $\Lambda_i = \cup_i s_i$ is the slit graph after the i-th merge. Λ is the full slit graph.
- v_i is called *flattened* if it has already been merged. (In Chapter 15, when v_i is a vertex of positive curvature, the merge will reduce v_i's curvature to zero.)
- V_i is the set of not-flattened vertices remaining after the i-th merge $s_i = m_{i-1}v_{i+1}$. $m_i \in V_i$. $|V_i|$ is the number of vertices in V_i.
- H_i is the convex hull of V_i. We view H_i as a closed region of the plane and ∂H_i its boundary, a convex polygon. H_0 is the convex hull of $V = V_0$.

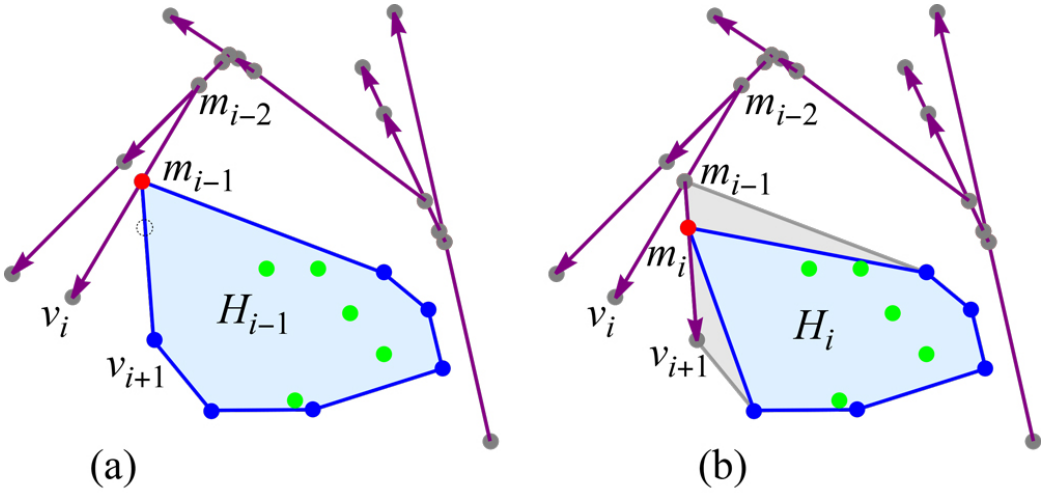

Figure 12.2: General step: i-th merge. (a) $s_{i-1} = m_{i-2}v_i$. $s_{i-1} \cap H_{i-1} = m_{i-1}$. (b) $s_i = m_{i-1}v_{i+1}$. $s_i \cap H_i = m_i$. $s_{i-1} \cap s_i = m_{i-1}$.

12.3 Algorithm Description

First merge. The algorithm starts by selecting any edge of the convex hull H_0 of $V = V_0$, say v_1 and v_2, and merging them via the merge segment

$s_1 = v_1v_2$. In the plane, this amounts to replacing $v_1 = m_0$ and v_2 with with a new vertex $m_0 + v_2 \to m_1$ on the segment s_1. The merge flattens v_1 and v_2, which is why they are removed. $V_1 = V \setminus \{v_1, v_2\} \cup m_1$, and $H_1 = \text{conv}(V_1)$. See Fig. 12.1.

Second merge. Next, m_1 is merged with the first vertex beyond m_1 on H_1. Call this vertex v_3. Note that v_3 might not be the next vertex after v_2 on H_0, because $V_1 \neq V_0$ and so $H_1 \neq H_0$. The merge of m_1 and v_3 introduces a new merge vertex m_2 on $s_2 = m_1v_3$, and both m_1 and v_3 are flattened and removed from V_1, while m_2 is added, to produce V_2.

General step. The i-th merge connects m_{i-1} to v_{i+1}, where v_{i+1} is the next vertex on H_{i-1} beyond m_{i-1}. See Fig. 12.2. Then m_i lies on $s_i = m_{i-1}v_{i+1}$, the vertex set is updated as $V_i = V_{i-1} \setminus \{m_{i-1}, v_{i+1}\} \cup m_i$, and $H_i = \text{conv}(V_i)$. Note that $|V_i| = |V_{i-1}| - 1$, because two vertices are removed (flattened) and one added. But ∂H_i does not bear a similar relation to ∂H_{i-1} because the new hull may wrap around vertices that were strictly interior to H_{i-1} (green in the figures).

Completion. The process continues until all the original v_i vertices are merged, leaving just one vertex $V_{n-1} = \{m_{n-1}\}$. A full trace is illustrated in Fig. 12.3. (The reason there are $n - 1$ merges rather than n is that the first merge flattens two vertices of V, while all subsequent merges flatten two vertices but only one is from V.)

12.4 Planar Proof

We first assume that all points—both vertices in V and merge points m_i—are in *general position* in the sense that no three are collinear. Later we will see that allowing collinearities does not change the main claim that $\Lambda = \cup_i s_i$ is a tree.

The induction hypothesis consists of several claims:

(1) H_i is nested inside H_{i-1}: $H_{i-1} \supset H_i$. See Fig. 12.4.

(2) s_i only intersects H_i in one point, the merge vertex: $s_i \cap H_i = \{m_i\}$. See Fig. 12.2(b).

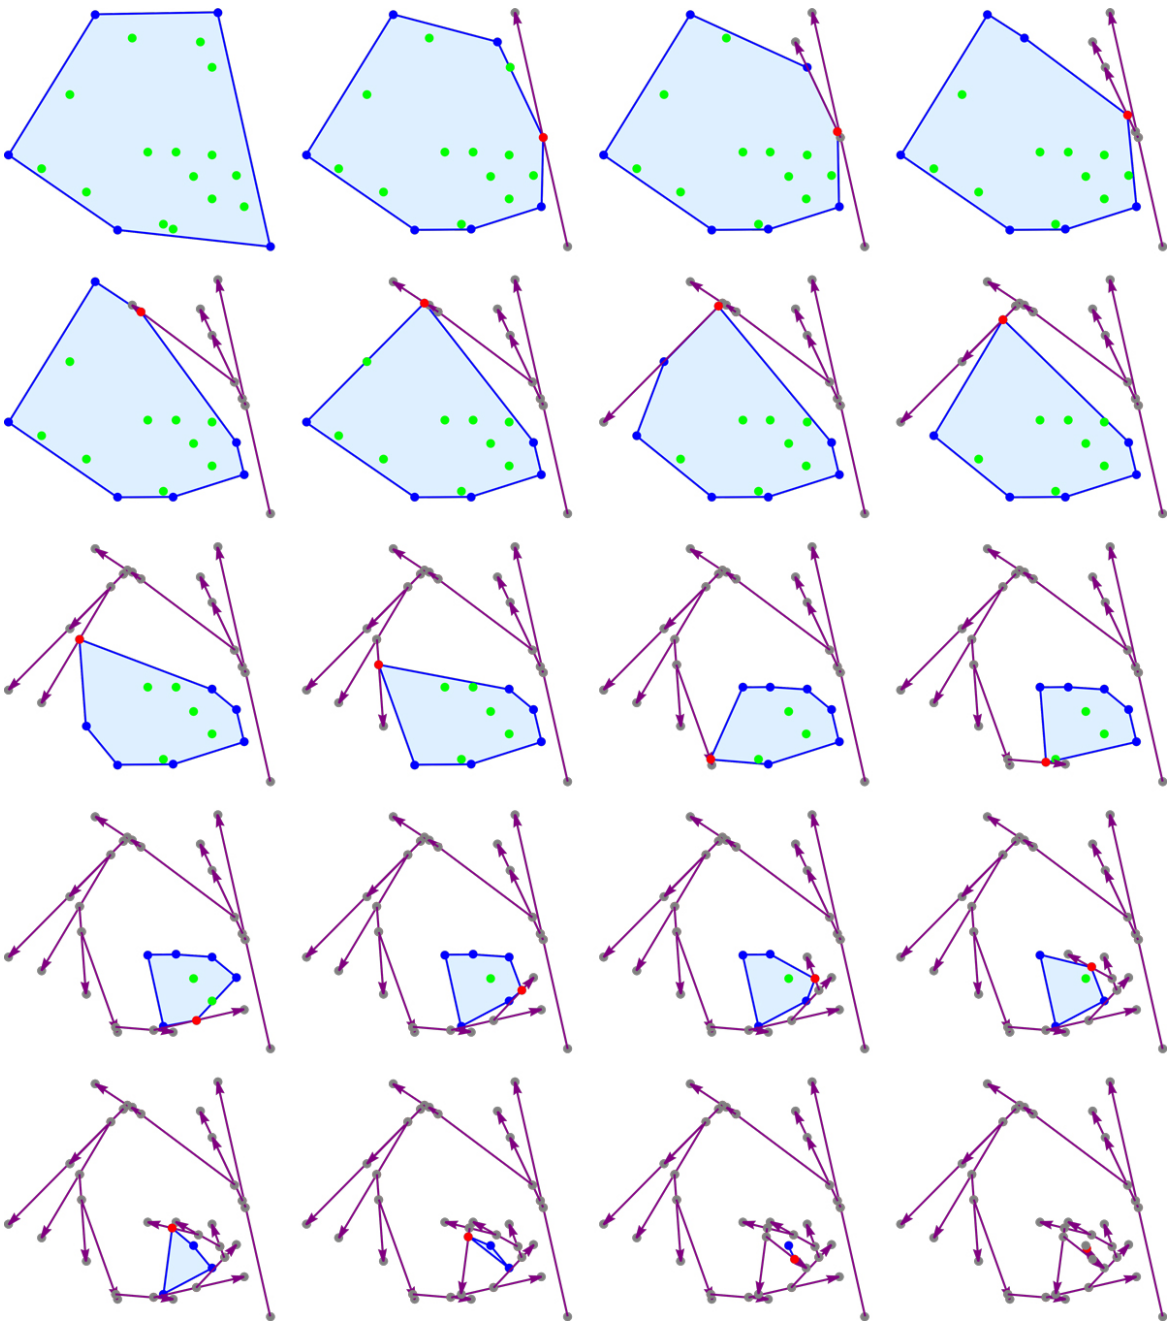

Figure 12.3: Trace of example with $|V| = 20$, from which Figs. 12.1 and 12.2 are details.

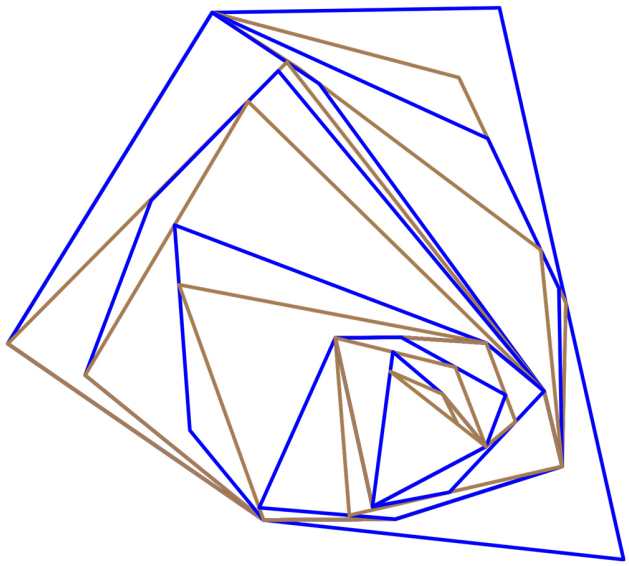

Figure 12.4: Nested H_i for all i, from Fig. 12.3, alternately colored blue and brown.

(3) Consecutive slit segments share just one point, an endpoint of the later segment: $s_{i-1} \cap s_i = \{m_{i-1}\}$. See Fig. 12.2.

(4) The set of vertices reduces by one each iteration: $|V_i| = |V_{i-1}| - 1$, for $i > 1$.

(5) Λ_i is a tree.

Basis. These claims are easy to see for $i = 1, 2$, just by construction according to the algorithm. Neverthless, just for completeness, we run through the basis. Refer to Fig. 12.1.

(1) $H_0 \supset H_1$ because the edge v_1v_2 is replaced by a point m_1 on that edge.

(2) $s_1 \cap H_1 = \{m_1\}$ by construction.

(3) For $i = 1$, there is just one segment s_1. s_2 starts at $m_1 \in s_1$, so indeed $s_1 \cap s_2 = \{m_1\}$.

(4) V_1 has one fewer vertex than does $V = V_0$ because two (v_1 and v_2) are removed and one (m_1) is added.

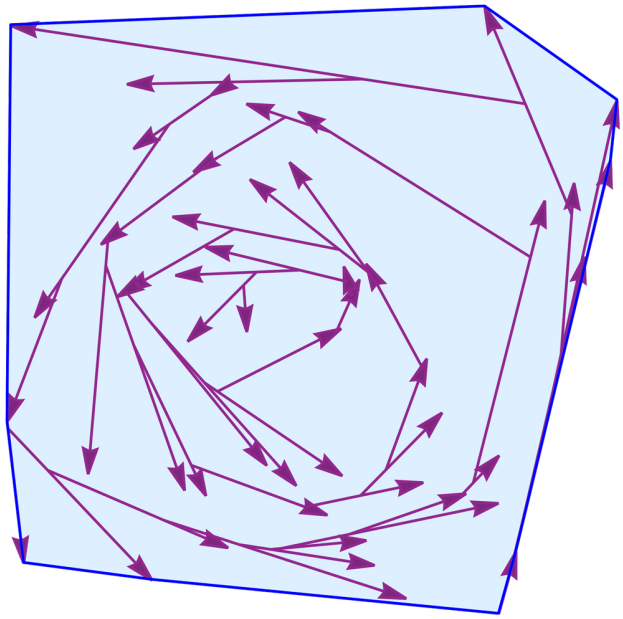

Figure 12.5: Λ for a set of $|V| = 50$ vertices. Only slit segments are shown.

(5) $\Lambda_1 = s_1$ is a single-edge tree.

General Step. Assume the induction hypotheses are satisfied up to $i-1$, and consider the i-th merge. It will help to consult Fig. 12.2, where (a) is established by the induction hypothesis, and (b) reflects the situation to be proved. The algorithm chooses v_{i+1} as the next vertex on H_{i-1}, so $s_i = m_{i-1}v_{i+1}$ is an edge e of ∂H_{i-1}.

(1) $H_{i-1} \supset H_i$ because the edge e is replaced by the merge vertex m_i.

(2) $s_i \cap H_i = \{m_i\}$ by construction.

(3) $s_{i-1} \cap s_i = \{m_{i-1}\}$ by construction.

(4) $|V_i| = |V_{i-1}| - 1$, again because $V_i = V_{i-1} \setminus \{m_{i-1}, v_{i+1}\} \cup \{m_i\}$.

(5) Λ_i is a tree. By (3) above, s_i has an endpoint m_{i-1} on s_{i-1}, and otherwise does not intersect s_{i-1}. It remains to prove that the interaction of s_i with the earlier segments in Λ_i maintains the tree property. In fact, we show that $s_i \cap s_j = \varnothing$ for $j < i-1$.

As we've seen, s_i is an edge e of ∂H_{i-1}, and from (1) we know that $H_j \supset H_{i-1}$ for $j < i-1$. From (2), each s_j intersects H_j in just one point m_j, which is removed from V_j in the next step. So s_j is disjoint from $H_{j+1} = \text{conv}(V_j)$. And so s_j cannot intersect s_i, which is an edge of ∂H_{i-1}.

Therefore, we have shown that s_i just intersects Λ_{i-1} at the one point m_{i-1}, and so maintains the tree structure for Λ_i.

The tree structure is illustrated in Fig. 12.5. It is in a sense a geometric directed binary tree, spiraling from the boundary of H_0 into the center.

Collinearities. If there are collinearities, then several consecutive segments can collinearly overlap. So it could be, for example, that $s_{i-1} = m_{i-2}v_i$ overlaps with $s_i = m_{i-1}v_{i+1}$ in the portion from m_{i-1} to v_i. Thus several slit segments could collinearly overlap with an edge of H_i. If we union each group of collinearly overlapping segments to one segment, then all the properties claimed in the induction hypothesis continue to hold, although different notation would be required to capture the unioning of several collinearly overlapping segments into one.

Theorem 12.1. *For any set V of vertices in the plane, the sequential merging algorithm detailed in Section 12.3 results in a slit graph Λ that is a tree (see Fig. 12.5).*

As the figures suggest, we have implemented this algorithm.

Convexity and, in particular, the convex hull, plays an essential role in this algorithm. Before we can apply the same overall idea to a set V on a three-dimensional P^+ in Chapter 15, we need to explore convexity and convex hulls on P. This is the topic of the next chapter.

Chapter 13

Convexity on Convex Polyhedra

We've set as our goal proving that there is a vm-reduction ordering of the vertices V inside a quasigeodesic Q that results in a slit tree Λ, a goal achieved in Chapter 15. Before we reach that point, we need to develop a clear notion of what constitutes the convex hull of V. This in turn requires a clear notion of convexity on polyhedral convex surfaces, which is our focus in this chapter.

Our investigation into convexity on convex polyhedra is, to our knowledge, the first in this direction, and seems to have a rich potential. We barely touch on some classical convexity results (such as Helly's and Radon's theorems: Examples 13.13 and 13.22), and do not attempt to provide their adaptations in this new theory of convex sets. Instead we focus on developing basic facts that ultimately lead to a characterization of the "relative convex hull" of a set of vertices inside a simple closed quasigeodesic Q (Section 13.10).

We first argue that, for our purposes, the proper notion of a convex set S is one that includes every geodesic segment between points in S, in contrast to, say, including at least one geodesic segment. Defining the convex hull of a set of vertices V as the smallest convex set including V (Section 13.6) leads to some surprising and perhaps undesirable properties. For example, the convex hull of a set inside Q does not always remain inside Q (Example 13.30). These properties lead us to develop a notion of "relative" convexity, and the relative convex hull in Section 13.9. Among our main results is Theorem 13.40, which characterizes the relative convex hull of vertices inside a quasigeodesic Q. This characterization is then employed in Chapter 15.

Because of the many lemmas and details in this long chapter, we offer a concise summary of 38 individual results in Section 13.10.

Throughout, for a set S, we use ∂S for its boundary, and $\mathring{S}$ and $\bar{S}$ for the interior and closure of S respectively.

13.1 Convex Curves

Let C be a simple, closed, curve on the surface of a convex polyhedron P. We will assume C is polygonal, turning at *corners* which in general may or may not be vertices of P (although we'll continue to use labels v_i). View C as directed counterclockwise from above. At each corner v_i of C, let α_i be the surface angle to the left, and β_i the angle to the right. C is called a *convex curve* if $\alpha_i \leq \pi$; all such angles are *convex angles*. Angles β_i strictly greater than π are called *reflex*.

If C also satisfies $\alpha_i \leq \beta_i$, we call it an *$\alpha\beta$-convex curve*. If $\alpha_i < \beta_i$ for all (non-zero curvature) corners, it is a *strictly $\alpha\beta$-convex curve*.

A simple, closed *quasigeodesic* Q is a curve C that is convex to both sides: $\alpha_i \leq \pi$ and $\beta_i \leq \pi$. By a theorem of Pogorelov [Pog49], every convex polyhedron has at least three such quasigeodesics.[1]

In general, quasigeodesics are not $\alpha\beta$-convex to either side. The total curvature of the vertices to either side of Q is $\leq 2\pi$, and only equal to 2π if Q is a geodesic, passing through no vertices, a fact we use later.

A *geodesic polygon* is a polygon whose edges are geodesic arcs. A *geodesic-segment polygon* is a polygon whose edges are geodesic segments, i.e., shortest paths. The distinction between these two types of polygons plays a significant role in this chapter. For brevity, we will frequently abbreviate "geodesic arc" with "geoarc," and "geodesic segment" with "geoseg." Although we are mainly interested in convex sets contained within polygons, a convex curve could include smoothly curved arcs.

13.2 Notions of Convexity

Several notions of convexity have been employed so far on surfaces, see for example [Ale78], [Ban81], [GM01], [Mit16], [Zam91], and the references therein.

[1] We point out here that a different notion of "quasi-geodesic" also exists, see e.g., [BZ21].

We have not made a comprehensive accounting of all the different definitions found in the literature, and instead mention next only three.

A subset S of P is said to be

- *geodesically convex* if, given any two points in S, there is a unique geodesic segment joining them in S; [Udr13].
- *totally convex* if any geodesic which joins two points of S is contained in S; [Ban81].
- *metrically convex* if, given any two points in S, there is at least one geodesic segment joining them in S; [GM01].

We should mention that there is variation in the literature. For example, [Vis18] uses the term "geodesically convex" to mean what we list above as "totally convex."

Geodesic convexity is not suitable for our framework, because an open geodesically convex set could contain no vertex v—two geodesic segments would wrap around v. Since our main concern is with geodesic segments forming a slit graph, total convexity, which focuses on geodesics (not necessarily shortest paths), is also inappropriate.

We will need a notion of a convex hull of a set of points V on P. In analogy with the Euclidean case, we would like to define the convex hull of V to be the intersection of all convex sets containing V. We next indicate why the natural metric convexity, and in particular, its "at least one geodesic segment" criterion, is not the correct version in our context.

Example 13.1. *Let Δ be a doubly-covered triangle $v_1v_2v_3$, with $V = \{v_1, v_2, v_3\}$. The front triangle is a metric-convex set containing V, as is the back triangle. So the metric-convex hull of V consists of the three edges v_1v_2, v_2v_3, v_3v_1. However, one would expect to obtain Δ as the convex hull of its vertices, as we do in Theorem 13.32. Moreover, the situation illustrated by the doubly-covered triangle could be troublesome in some specific situations as well. We address these situations in Section 13.6 below.*

13.3 Ag-convexity

Define S to be *ag-convex*—all-geodesics convex—if every geodesic segment between two points of S is in S. In the following we will focus on ag-convexity, which we henceforth abbreviate as, simply *convexity*.

This same notion of convexity has been used in different contexts, see for example [AKP19] and [LP21], although we have not seen it applied to convex polyhedra. Other types of convexity will be specifically identified to differentiate them from ag-convexity.

Clearly, every convex set is path-connected, but one can say more.

Lemma 13.2 (Classification)**.** *Let S be a closed convex set on P. Then S is either a point, or a geodesic arc, or a simple closed geodesic, or it has interior points.*

Proof. Assume S contains two distinct points, hence it contains a geodesic segment γ. Assume Γ is the maximal (with respect to inclusion) geodesic in S including γ; so it could be an arc, or a simple closed curve. If the image set of Γ is S, we are done. (Here the *image set* of Γ is the set of points on P comprising the curve.)

Assume there exists a point x in S but not on Γ. Join it with an interior point y of Γ, say with the geodesic segment γ_y. Locally around γ_y the set S is flat, so moving y continuously on Γ provides a continuous family of geodesic segments γ_y, all included in S. Therefore S has interior points. □

Example 13.3. *The whole surface P is obviously convex. We also consider the empty set and the single-point sets to be convex.*

Example 13.4. *Notice that the closure of a convex set is not necessarily convex. The interior of a face of a doubly-covered triangle is convex, while the whole face is not, because two points on different edges of the triangle are connected by geodesic segments on both sides of the triangle.*

We'll repeatedly use the following simple fact.

Lemma 13.5 (Local behaviour)**.** *Let $S \subset P$ be a convex set and $x \in S$. Then there exists a small ball B_x around x such that $S \cap B_x$ is isometric either to a planar convex set (if x is not a vertex of P), or to a convex set on a cone of apex x.*

Proof. Just choose B_x around x to include no vertices of P (other than possibly x), and notice that $S \cap B_x$ is still convex on B_x. □

Recall from Section 13.1 that an $\alpha\beta$-convex curve is a *strictly $\alpha\beta$-convex curve* if $\alpha_i < \beta_i$ for all (non-zero curvature) corners. The following characterization of $\alpha\beta$-convexity will be frequently invoked subsequently.

Lemma 13.6 ($\alpha\beta$-convexity)**.**

(i) Let $S \neq P$ be a closed convex subset of P, $S \neq P$, having at least two distinct points. Then ∂S is a convex curve (including the case when S is a geodesic). If, moreover, ∂S is a geodesic polygon, then it is strictly $\alpha\beta$-convex.

(ii) Let S be an open convex subset of P, $S \neq P$, and C a connected component of ∂S. Then C is either a vertex, or a convex curve. If, moreover, C is a geodesic polygon, then it is (not necessarily strictly) $\alpha\beta$-convex.

The conclusion in Lemma 13.6(ii) holds with a similar proof for convex sets which are neither closed nor open, but have interior points. A surface containing such a set is given later in Example 13.31.

Proof. (i) The first claim follows from Lemma 13.5. Indeed, each component of ∂S is a locally-convex curve. If S has interior points then each component of ∂S is a simple closed curve, hence convex.

For the second claim, let v_i be a (non-zero curvature) corner of ∂S at which $\alpha_i \geq \beta_i$. If $\alpha_i = \pi$, then $\beta_i = \pi$, and v_i has zero curvature. So we may assume that $\alpha_i < \pi$.

Let x, y be two points on ∂S close to and on either side of v_i. Then, as shown in Fig. 13.1, the shortest path between x and y inside S is either longer than, or equal in length to, the shortest path outside S. In either case, a geodesic segment between points of S is not inside S, contradicting the fact that S is convex.

(ii) Assume ∂S is not $\alpha\beta$-convex at v, hence $\alpha > \beta$. Then there exist points $x, y \in \partial S$ with v between them, and any geoseg γ joining them is external to S, as in the above argument. There also exist sequences of points $x_n, y_n \in S$ with $x_n \to x$ and $y_n \to y$. Choose geosegs γ_n from x_n to y_n, hence γ_n lies in S. It is well known that the limit of geosegs is a geoseg. Possibly passing to a subsequence, assume $\gamma_n \to \gamma'$, hence γ' is a geoseg joining x to y in $\bar{S}$, which is only possible if $\alpha = \beta$ at v, contradicting our assumption. □

Example 13.7. *To complete Example 13.4, notice that the interior of a face is an open convex set on a doubly covered triangle, but its boundary is not strictly $\alpha\beta$-convex. It remains, nevertheless, $\alpha\beta$-convex.*

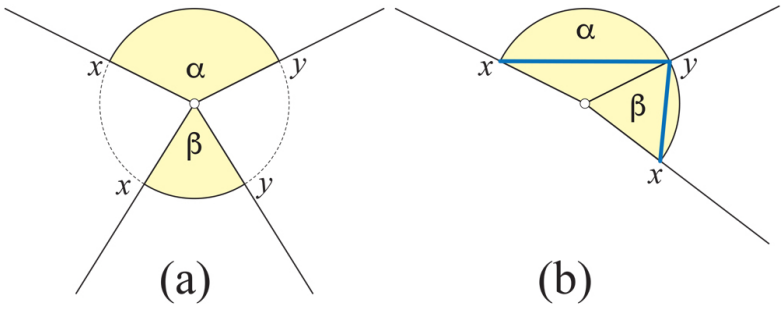

Figure 13.1: (a) $\alpha > \beta$. (b) $|xy|$ is shorter through angle β than through α.

An open convex set may have several vertices as boundary components, as we will see in Example 13.11 below, and at most two convex curves as boundary components, as will be evident in the next result and its proof.

Lemma 13.8 (Simply-connected or cylinder)**.** *Let $S \neq P$ be a closed convex subset of P. Then either S is simply-connected, or it is isometric to a cylinder without lids.*

Recall from Chapter 9 (Section 9.2.1) that in our usage, simply-connectedness does not assume path-connectedness. However, all convex sets are path-connected.

Proof. If S has empty interior, the conclusion follows from Lemma 13.2; so we may assume in the following that S has interior points.

The convexity of S implies that each component of ∂S is a simple closed curve, and since it is locally convex by Lemma 13.6, it is a convex curve.

Let C_1, C_2 be two components of ∂S. There are two possibilities: without loss of generality, C_2 is nested inside C_1, or they are not nested. This latter case violates convexity, because a geodesic segment connecting a point inside C_1 to a point inside C_2 would necessarily include points exterior to both. So let C_2 be nested inside C_1.

The Gauss-Bonnet Theorem implies that the total curvature inside C_1 is $\leq 2\pi$, and similarly for C_2. So the total curvature outside C_1, and also outside C_2, is $\geq 2\pi$. But the exterior of C_2 is strictly inside C_1, and the exterior of C_1 is strictly inside C_2. Therefore, either the region between C_1 and C_2 has zero curvature, and therefore it is isometric to a cylinder without lids, or at least one the above inequalities is strict and we get a contradiction, establishing the conclusion. □

Lemma 13.9 (Vertex point-hole)**.** *Let $S \subset P$ be a convex set and v a vertex in S. Then $S \setminus \{v\}$ is convex.*

Proof. The case $S = \{v\}$ is clear, because we consider the empty set to be a convex set.

Consider distinct points $x, y \in S$, and a geodesic segment γ between them, hence $\gamma \subset S$. If $v \notin \{x, y\}$, $\gamma \subset S \setminus \{v\}$, because no geodesic segment will pass through a vertex.

Assume now that $x = v$. Then each point of $\gamma \setminus \{v\}$ is inside a geodesic subsegment of γ, hence included in $S \setminus \{v\}$. □

Example 13.10. *An example illustrating the cylinder claim of the Lemma 13.8 is provided by P a tall rectangular block, and S the region between two parallel simple closed geodesics close to the middle of the block. Then S is a closed convex set.*

Example 13.11. *Lemma 13.8 does not hold for arbitrary convex sets with interior points. Indeed, the set $P \setminus V$ is convex (see Lemma 13.9), open, but not simply-connected.*

Example 13.12. *More generally, let S be a simply-connected convex set, and v a vertex interior to S. Then $S \setminus \{v\}$ is convex, but not simply-connected. However, the closure of $S \setminus \{v\}$ is convex and simply-connected.*

Example 13.13. *A particular instance of Helly's Theorem states the following: let $S_1, ..., S_n$ be a finite collection of convex subsets of $\mathbb{R}^2$, with $n > 3$. If the intersection of every $h = 3$ such sets is nonempty, then the whole collection has a nonempty intersection.*

The simple example of the four faces of a tetrahedron, each face a closed, convex set, shows that no analogous result holds in our framework.

A more elaborated example places the tetrahedron over a tall right triangular prism. There, the union of the tetrahedron's lateral faces, minus the top vertex, is a convex set. That set, together with the lateral faces, provides another counterexample.

It remains for future work to consider the above Helly-like problem for $h \geq 4$: Open Problem 18.12.

Lemma 13.14 ($S \supseteq Q$)**.** *If the interior of the closed convex set S contains a simple closed quasigeodesic Q, then either ∂S determines with Q a topologically closed cylinder, or $S = P$.*

Proof. Assume the simple closed quasigeodesic Q is included in S.

If Q parallels ∂S (in the sense that they bound a cylinder) then ∂S is itself a simple closed quasigeodesic. Assume this is not the case.

Because Q is interior to S, the sum of curvatures of vertices inside or on Q is at least 2π. If $\partial S \neq \varnothing$ then, since it is convex towards S, the total curvature of the interior of S is $< 2\pi$, by Lemmas 13.8 and 13.6, which violates the Gauss-Bonnet Theorem. □

We don't know if every closed convex set is either included in a half-surface bounded by a simple closed quasigeodesic, or is the whole surface. This is not settled by the previous lemma because there Q is interior to S. This question is Open Problem 18.10. However, the proof of the previous lemma does show that no closed convex set can strictly enclose more than 2π curvature. This picture will be completed by Proposition 13.27 and Theorem 13.32 below.

13.4 Geodesic Segments and Convex Sets

In this section we clarify the behaviour of geosegs between points in, or on the boundary of, or outside of, a convex set. The first part of the next result will be particularly useful later.

Lemma 13.15 (Segment vs. convex set)**.** *Let $S \subset P$ be a convex, $S \neq P$.*

(i) If x, y are interior points of S and γ a geodesic segment joining them, then γ is interior to S (i.e., it does not intersect ∂S).

In particular, adding a boundary point to an open convex set yields a convex set.

(ii) If x is an interior point of S and $y \in \partial S$, then all geodesic segments from x to y are interior to S, possibly excepting at their extremity y.

(iii) If $x, y \in \partial S$ then there is a geodesic segment between them included in the closure $\bar{S} = S \cup \partial S$ of S, possibly excepting its extremities x, y. Other geodesic segments from x to y might not be included in $\bar{S}$.

In particular, adding two boundary points to an open convex set does not necessarily yield a convex set.

(iv) If $x \in \partial S$ is not by itself a connected component of ∂S, then there exists a geodesic segment starting at x and exterior to S.

Proof. (i) Let B_1, B_2 be open intrinsic balls around v_1, v_2 respectively, with $B_1 \cup B_2 \subset S$.

Let $N_\gamma \subset P$ be a tubular neighborhood of the image set of γ. View N_γ as created by translations of γ sufficiently small so that N_γ is included in B_i around v_i, $i = 1, 2$, and so that N_γ is an open set included in S. N_γ is flat everywhere, except at v_1, v_2. Because γ is strictly inside N_γ, γ does not intersect ∂S.

This claim can also be established by assuming that some geodesic segment between interior points touches ∂S, and applying Lemma 13.5 to obtain a contradiction.

(ii) Assume there exists a geodesic segment γ from x to y that goes outside S, so there is a boundary point $z \in \partial S$ on γ between x and y. Assume, moreover, that z is closest to y with these properties.

Notice that such a z exists, i.e., is at positive distance to y. To see this, consider a small ball B_y around y which contains no vertex of P. Then $S \cap B_y$ is a convex set in the plane (if y is not a vertex of P) or on a cone (if y is a vertex), by Lemma 13.5. Denote by γ_z the sub-arc of γ from z to y. The boundary of this small convex set $\partial(S \cap B_y)$ is a convex curve, and the image of γ_z is a segment, with the two intersecting in at most two points (because a segment can intersect a convex curve at most twice). This shows z exists as described.

Denote by A_z the component of ∂S between z and y such that no point of S is inside the lune L bounded by A_z and γ_z. L is then bound by a curve A_z concave towards L and a straight segment γ_z exterior to A_z. The Gauss-Bonnet Theorem implies this is only possible if L contained negative curvature, a contradiction.

(iii) Because $x, y \in \partial S$, there exist a sequence of points $x_n, y_n \in S$ such that $x_n \to x$, $y_n \to y$. Let γ_n be a geodesic segment joining x_n to y_n, $n \in \mathbb{N}$. Then there exists a subsequence of $\{\gamma_n\}_n$, which converges to a geodesic segment γ between x and y. Since all γ_n are in S, their limit is in $\bar{S}$.

(iv) This claim can be obtained from Lemma 13.5. An alternative proof follows.

The space of unit tangent directions T_x at x is a circle, of course of arc-length $\leq 2\pi$.

Assume, in contradiction to the claim that, for any $\nu \in T_x$, the maximal geodesic segment γ^ν starting at x in direction ν intersects $S \cap B_x$, i.e., it has an initial portion of γ^ν is inside S. Choose three directions $\nu_i \in T_x$, $i = 1, 2, 3$, with the angle between adjacent pairs ν_i, ν_j less than π. With

points $x_i \in S$ on γ^{ν_i}, the three flat triangles xx_ix_j are in S, and their union forms an open set around x, contradicting $x \in \partial S$. □

Example 13.4 can be easily elaborated to illustrate all the claims of the above lemma.

Corollary 13.16. *It follows directly from (i) in Lemma 13.15 that the interior of a convex set is still convex.*

Lemma 13.17 (Supporting angle)**.** *Let $S \neq P$ be a closed convex subset of P with interior points, and x a boundary point of S. There exists tangent directions $\mu, \nu \in T_x$ at x of angle $\theta = \theta(\mu, \nu) \leq \pi$ toward S, such that:*

(i) the geodesic segments γ^μ and γ^ν in the directions μ, ν do not intersect the interior of S, locally; and

(ii) for each $\tau \in T_x$ inside θ, the geodesic segment γ^τ in the direction τ does intersect the interior of S, locally. Here, θ is regarded as a subarc of T_x.

Proof. This can established from Lemma 13.5. An alternative way to prove the claims follows.

Let B_x be a small ball around x which contains no vertex of P.

Lemma 13.15 shows that there exists a geodesic segment γ_z connecting x to a point $z \in P \setminus S$. Because $P \setminus S \neq \varnothing$ is an open set, there are other points in the neighborhood of z similarly connected to x by geodesic segments that do not intersect the interior of S. Denote by A_x the maximal open subarc of T_x determined by directions of such geodesic segments γ_z. Let $\mu, \nu \in T_x$ be the extremities of A_x.

The convexity of S implies that for all $\tau \in A_x$, the geodesic segment γ^τ does not intersect $S \cap B_x$. Moreover, $A_x \geq \pi$. This establishes claim (i).

The convexity also implies that $\theta = T_x \setminus A_x \leq \pi$ has the opposite property: for each $\tau \in T_x$ inside θ, the geodesic segment γ^τ intersects the interior of $S \cap B_x$. This establishes Claim (ii). □

13.5 Relative Convexity

For the next result (Lemma 13.18), we need to modify the notion of convexity to *relative convexity*, a variation on a notion of relative convexity we employed in [OV14].

Let $S \subset P$ be a contractible closed convex set with interior points. We proved in the previous section that S is bounded by a closed convex curve $C = \partial S$. We glue to C a tall cylinder L with a base but without a top. This satisfies AGT, because a point $p \in C$ is convex to the S-side and has angle π on the cylinder rim.

Denote by $P^{\#}$ the resulting convex surface. AGT implies that $P^{\#}$ is a polyhedron if and only if ∂S is a geodesic polygon on P. Call S *relatively convex* if its image (also denoted by S) is convex on $P^{\#}$.

Note that C is a quasigeodesic on $P^{\#}$.

There are two considerations that lead us to introduce relative convexity in this section, and the relative convex hull $\mathrm{rconv}(S)$ in Section 13.6. First, the failure of vertex-merging to preserve convexity, Example 13.19 below. Second, the failure of the convex hull of points inside a simple closed quasigeodesic Q to remain inside Q, Example 13.30 below.

In Section 13.9 we shall use a construction similar to the above, but with respect to simple closed quasigeodesics Q instead of convex curves C. The purpose there is to construct the relative convex hull, starting from a non-convex set inside Q. So the term "relative" will refer to either construction, depending on the context, and without confusion.

Lemma 13.18 (Vertex merging)**.** *Let v_1, v_2 be two vertices interior to the convex set $S \subset P$. Merging v_1 and v_2 produces a new polyhedron P' and a set $S' \subset P'$ obtained as the union of S with the two merge triangles T. Then S' is relatively convex on P'.*

Before we prove this lemma, we show that it is false for convexity, without the "relatively" modifier.

Example 13.19 (doubly-covered quadrilateral)**.** *Figure 13.2 shows an example $S \subset P$ where the insertion of merge triangles causes some geodesic segments in the new set S' to leave S', violating convexity. In (a), geodesic segments from any point x on the front face to its image x' on the back face stay within S, whereas after the v_1, v_2 merge (b), there are points x whose shortest path to its back-face image x' crosses ab outside of S'.*

Proof. We know by Lemma 13.15 that the insertion of the triangles T has not affected the boundary of S': $\partial S' = \partial S = C$, where C is a convex curve. In $P'^{\#}$, denote by L the sides of the cylinder glued to C. When L is unfolded flat, it forms a rectangle R.

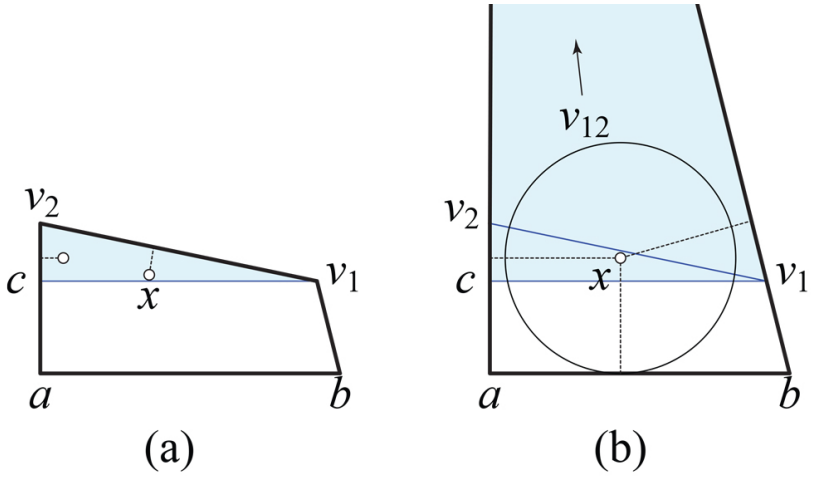

Figure 13.2: (a) P is the doubly-covered quadrilateral abv_1v_2. Convex set $S \subset P$ is the two-sided triangle cv_1v_2, shaded. (b) Merging v_1 and v_2 yields v_{12} far above P. The resulting set S' is cv_1v_{12}.

Let $x, y \in S'$ be two points in S', and γ a geodesic segment between them. We argue that γ cannot cross C. Suppose γ properly enters R at x' and exits at y'. Then the segment $x'y'$ along the top of R is shorter, a contradiction to the assumption that γ is a geodesic segment.

Therefore, S' is convex on $P'^{\#}$, i.e., S' is relatively convex. □

13.6 Convex Hull

We mentioned earlier (Section 13.2) that one reason we are not using metrical convexity is that it leads to an unsatisfactory notion of a convex hull. Here we explore the convex hull, and conclude that we need a variation for our purposes in the next chapter.

In analogy with the Euclidean case, define the *convex hull* $\text{conv}(S)$ of an arbitrary set $S \subseteq P$ as the intersection of all the convex sets that enclose S on P. So, in a familiar sense, it is the smallest convex set with this property.

In the plane, the convex hull of S can equivalently be defined as the set enclosed by a minimal length curve enclosing S. However, as we will show in Section 14.3, the two notions do not coincide in our context. Here we focus on the intersection definition.

Example 13.20. *The convex hull* $\mathrm{conv}(V)$ *of the vertices* V *of a doubly-covered triangle* Δ *is the whole surface* Δ*. Similarly, the convex hull of all the vertices of any convex polyhedron* P *is* P *itself:* $\mathrm{conv}(V) = P$*.*

Example 13.21. *In contrast to the Euclidean situation, for* $S = P \setminus V$ *we have* $\mathrm{conv}(S) = S$*, by Lemma 13.9, because* P *is itself convex, i.e., the vertex holes are not filled-in by the convex hull operation.*

Example 13.22. *A particular instance of Radon's Theorem on (extrinsically) convex sets states that any set of* $r = 4$ *points in* $\mathbb{R}^2$ *can be partitioned into two sets whose convex hulls intersect.*

The simple example of a tetrahedron shows that no analoguos result holds in our framework. Indeed, the convex hull of any two vertices is the corresponding edge, so any $2 : 2$ *partion of the vertices provides disjoint convex hulls. And the convex hull of any three vertices is the corresponding face, so any* $3 : 1$ *partion of the vertices also provides disjoint convex hulls.*

If, as in Example 13.13, we again place the tetrahedron over a tall right triangular prism, then the convex hull of the tetrahedron base vertices is the "roof" minus the fourth vertex.

It remains for future work to study the Radon problem for $r \geq 5$, possibly considering closures of convex hulls:

Also worth studying seems to be the existence of a Carathéodory type theorem in our framework. These questions form Open Problem 18.12.

The following two properties follow immediately from the definition of $\mathrm{conv}(S)$.

Lemma 13.23. *For every* $S \subset P$*,* $\mathrm{conv}(\mathrm{conv}(S)) = \mathrm{conv}(S)$*.*

For every two sets $S, S' \subset P$ *with* $S' \subset S$*,* $\mathrm{conv}(S)' \subseteq \mathrm{conv}(S)$*.*

The next result is particularly useful for computing the convex hull of finite sets.

Lemma 13.24. *For every pair of subsets* S, T *of* P *with* $T \subset S$*, the following holds:*

$$\mathrm{conv}(S) = \mathrm{conv}\left((\mathrm{conv}(S \setminus T)) \cup T\right).$$

In words: If S is partitioned into two parts, T and $S \setminus T$, then the hull of S is the hull of the union of T with the hull of $S \setminus T$.

Proof. The set $\operatorname{conv}((\operatorname{conv}(S \setminus T)) \cup T)$ is convex and includes S, hence it also includes $\operatorname{conv}(S)$.

On the other hand, the convex set $\operatorname{conv}(S)$ includes $\operatorname{conv}(S \setminus T)$, by Lemma 13.23, and also includes T. Hence it includes $\operatorname{conv}((\operatorname{conv}(S \setminus T)) \cup T)$ as well. □

The next simple result follows from Lemma 13.15.

Lemma 13.25. *The convex hull of every open set is open.*

Example 13.26. *Let* $\square$ *be a doubly-covered square with vertices* a, b, c, d*, and let* $S = \{a, b, c\} \subset \square$*. Then* S *is closed and, as we prove below,* $\operatorname{conv}(S) = \square \setminus \{d\}$*, hence* $\operatorname{conv}(S)$ *is not closed. Refer to Fig. 13.3.*

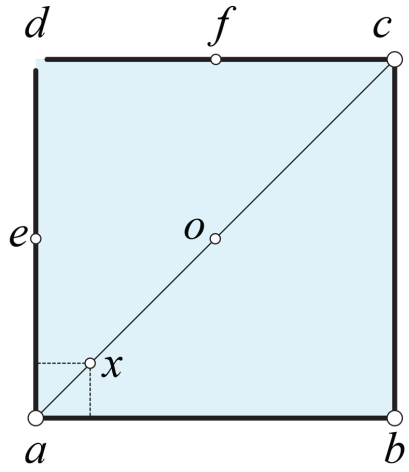

Figure 13.3: $\operatorname{conv}(\{a, b, c\})$ is $\square \setminus \{d\}$.

Clearly, $\operatorname{conv}(S)$ *contains the two triangles abc, hence also the simple closed quasigeodesic* Q *determined by the diagonals ac. Let* $x \in Q$ *be a point close to a on the front side, and* x' *its corresponding point on the back side. Then there are two geodesic segments between* x *and* x'*, one "horizontal" and another one "vertical"—parallel to ab and bc respectively. Therefore, there are points of* $\operatorname{conv}(S)$ *on the edge ad, close to a. Moving* x *continuously towards the center* o *of* $\square$*, shows that the half-edge ae is included in* $\operatorname{conv}(S)$*, where* e *is the midpoint of ad. Similarly,* $\operatorname{conv}(S)$ *contains the half-edge cf, where* f *is the midpoint of cd. Iterating this process establishes that* $\operatorname{conv}(S)$

contains halves of the segments ed and fd, and so on. Therefore $\mathrm{conv}(S)$ *contains the whole edges cd and ad, except missing the corner point d.*

The next result will be useful for proving the last statement in Theorem 13.32.

Proposition 13.27 (Dense hull)**.** *Let* $S \subset P$ *be a closed convex set with interior points, enclosing strictly less than* 2π *curvature. Then the convex hull of its complement* $S' = P \setminus S$ *is dense in* P*, and the convex hull of* ∂S *is dense in* S*.*

Proof. By Lemma 13.25, $\mathrm{conv}(S')$ is open. Then, by Lemma 13.6, the boundary of $\mathrm{conv}(S')$ may consist of vertices and convex curves, convex towards $\mathrm{conv}(S')$. If such a boundary curve existed, it would enclose more than 2π curvature, contradicting the the Gauss-Bonnet Theorem.

For the second claim, notice that $\mathrm{conv}(\partial S) \subset S$. Moreover, a continuity argument proves that $\mathrm{conv}(\partial S)$ has interior points. The remainder is analogous to Lemma 13.6, and to the above argument. □

Example 13.28. *Example 13.26 (Fig. 13.3) can be adapted to illustrate Proposition 13.27. See Fig. 13.4.*

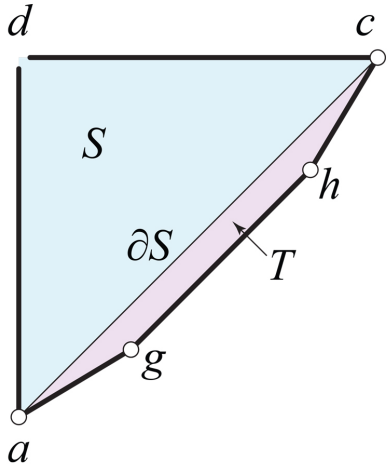

Figure 13.4: $\mathrm{conv}(\partial S)$ is dense in S and $\mathrm{conv}(P \setminus S)$ is dense in $P = S \cup T$.

Let $T = ghca$ *be an isosceles trapezoid with* $gh||ac$*,* $|g - h| < |a - c|$*, and the base angle at* g *almost* π*. Construct, outside* T*, an isosceles right triangle*

$S = acd$, with $|a - d| = |c - d|$. Let P denote the doubly-covered pentagon $T \cup S$.

The (two-sided) subset S of P is convex, and ∂S consists of the two geodesic segments ac. As in the proof of Example 13.26, $\mathrm{conv}(\partial S) = S \setminus \{d\}$ *is dense in S, and therefore* $\mathrm{conv}(P \setminus S) = P \setminus \{d\}$ *is dense in P.*

These examples show that the convex hull has some desirable properties, but, for our purposes, some undesirable properties. The doubly-covered square (Example 13.26) shows that $\mathrm{conv}(V)$ for V a subset of vertices of P could be a set with point holes—so the convex hull of a closed set can be an open set. And in the next section we will show that $\mathrm{conv}(S)$ for a set S is not always the convex hull of the extreme points of S. These considerations lead us to use the notion of relative convexity as first introduced in Section 13.5 above, and the relative convex hull $\mathrm{rconv}(S)$, which will be developed in Section 13.9 below.

Example 13.29. *Let B_h be a rectangular box of height h, and Q a simple closed geodesic on B parallel to its top and bottom faces, at height $h/2$. If h is large enough then Q is itself a convex set, and* $\mathrm{conv}(Q) = Q$*; see Lemma 13.37 and its proof for details. However, for small h, Q is not a convex set, and* $\mathrm{conv}(Q) = B$*. This shows that the convex hull of a set S is sensitive to the surface of P outside of S, and justifies the next section.*

13.7 Relative Convex Hull

The phenomena illustrated in the previous examples lead us to a new notion.

By "relative convex hull" $\mathrm{rconv}(S)$ we mean the intersection of all relative convex sets containing S, with relative convexity as previously defined in Section 13.5. We will further explore this notion in Section 13.9 below.

Example 13.30 (Pyramid)**.** *Let P be a pyramid, the top of a regular octahedron. In Fig. 13.5(a), both C_1 and C_2 are $\alpha\beta$-convex curves. The convex hull of the corners of C_1, is the surface above C_1, including the apex v_5. However, the convex hull of the corners of $C_2 = (v_1, v_2, v_3, v_4)$ is not the surface above C_2, but instead all of P. This is because there are geodesic segments that cross the square base. For example, the shortest path between the midpoints of consecutive edges of C_2 traverse the square base. Fig. 13.5(b) shows the construction of $P^{\#}$ for C_2. Then* $\mathrm{rconv}(C_2)$ *is the surface above C_2, because no shortest paths will enter the cylinder inserted below C_2.*

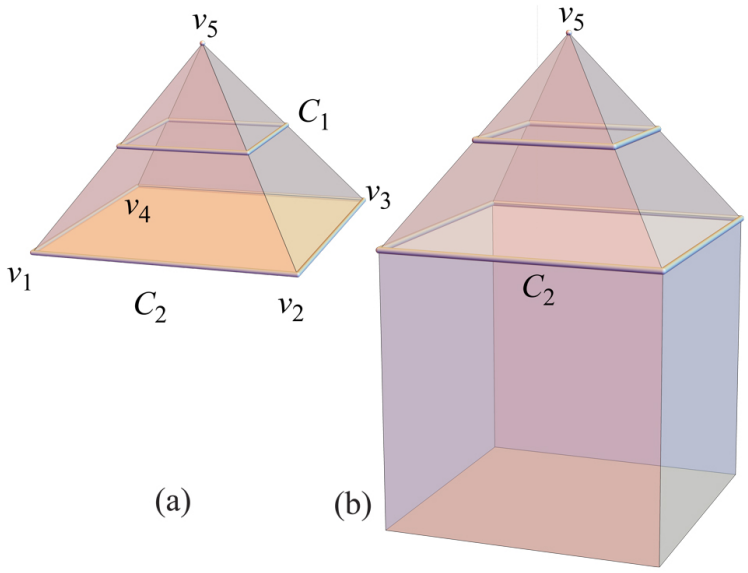

Figure 13.5: (a) $\mathrm{conv}(C_1)$ is the pyramid top, but $\mathrm{conv}(C_2) = P$. (b) $P^{\#}$ for C_2.

Example 13.31. *With the following example, we show that:*

- *neither a digon nor a triangle is necessarily convex;*
- *Neither* $\mathrm{conv}(V)$ *nor* $\mathrm{rconv}(V)$ *is necessarily closed, for* V *the set of all vertices inside a simple closed quasigeodesic;*
- *the exceptional case in Theorem 13.32 below may well appear for the closure of* $\mathrm{rconv}(V)$*, which in our case is convex but not the convex hull of its extreme points.*

Start with the double of the quadrilateral $D = abcd$*, where* $Q = aba$ *is a simple closed quasigeodesic. See Figure 13.6. We could insert rectangles below* Q*, so that everything happens above* Q*.*

The construction of D *follows: start with an arbitrary line-segment* ab*, and denote by* m *the mid-point of* ab*. Consider an isosceles triangle* $d'ab$*, and a point* d *slightly to the left of* d'*. Take* c *on* $d'b$*.*

Consider, moreover, ae *perpendicular to* $d'b$*, with* e *on* $d'b$*.*

Put $V = \{a, c, d\}$*. Then* $\mathrm{rconv}(V)$ *is the half-surface bounded by the two segments* ae*, without those segments (but including* a*). To see this, let* x_1

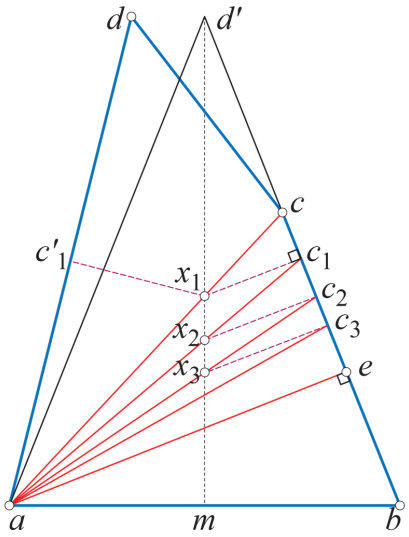

Figure 13.6: $V = \{a, c, d\}$, $\partial\,\mathrm{rconv}(V) = aea$, but $ae \setminus \{a\} \notin \mathrm{rconv}(V)$.

be the intersection point between ac and $d'm$, and c_1 the foot of x_1 onto $d'b$. Then the geodesic segment from x_1 to its "opposite" point goes through c_1, by construction. So c_1 belongs to $\mathrm{rconv}(V)$*, and therefore ac_1 is included in* $\mathrm{rconv}(V)$*. (Notice that $|x_1 - c_1| < |x_1 - c_1'|$.)*

Now let x_2 be the intersection point between ac_1 and $d'e$, and c_2 the foot of x_2 onto $d'b$. Then the geodesic segment from x_2 to its "opposite" point goes through c_2, by construction. So c_2 belongs to $\mathrm{rconv}(V)$*, and therefore ac_2 is included in* $\mathrm{rconv}(V)$*.*

Iterate, and pass to the limit.

It follows (as one can prove by running the above procedure backward) that all points above ae belong to $\mathrm{rconv}(V)$*, but not ae (see Lemma 13.36).*

Clearly, the digon aca is not convex. Also, consider that if we "split" the vertex a into vertices a', a'', then neither is the triangle $a'a''c$ convex.

13.8 Extreme Points

An *extreme point* of a convex set $S \subseteq P$ is a point in S that is not interior to any geodesic segment joining two points of S. For example, every vertex is an extreme point for every convex set containing it.

Roughly, the Krein–Milman Theorem states that every compact convex set is the extrinsic convex hull of its extreme points. The closest analogy for ag-convexity we could prove is the following theorem, which will be invoked later.

Theorem 13.32. *Let S be a closed convex subset of P with ∂S a closed curve, enclosing strictly less than 2π total curvature. Then either S is the relative convex hull of its extreme points, or ∂S contains a geodesic arc which is not a geodesic segment. In the former case, the interior extreme points of S are all vertices.*

We do not know if the above theorem statement is still true without the "relative" modifier; Open Problem 18.11.

Proof. **Case 1.** If $S \subset P$ has empty interior then the conclusion follows from Lemma 13.2, which shows that S is either a point, a geodesic arc, or a simple closed geodesic.

Case 2. Assume $S \subset P$ has non-empty interior. The proof proceeds by induction over the number of vertices of P interior to S.

For the base case, assume S has exactly one interior vertex v. By Lemma 13.6, ∂S is locally isometric to a planar convex curve. If ∂S contains a geodesic arc which is not a geodesic segment, we have established the lemma claim.

So assume now that ∂S contains a geodesic segment, maximal with respect to inclusion, between its corners v_i, v_{i+1}. Then we get a geodesic-segment triangle vv_iv_{i+1}, which is the convex hull of its extreme points. This is valid for all such geodesic-segment arcs in ∂S.

If ∂S contains an arc A locally isometric to a planar strictly convex curve, then each point of A is an extreme point, and the conclusion follows.

For the general case, again if ∂S contains a geodesic arc that is not a geodesic segment, we are finished. So now we prove that S is the relative convex hull of its extreme points.

Assume that S has at least two interior vertices, say v_1 and v_2. Because they are inside Q, $\omega_1 + \omega_2 < 2\pi$. Let γ be a geodesic segment joining them in S.

Merge v_1 and v_2 to v_{12} along γ to obtain a new surface P', hence where P' is P cut open along γ and the union T of two twin triangles is inserted. The set S with T inserted yields S', a relatively convex set of P' (Lemma 13.18).

Because $\partial S' = \partial S$ by Lemma 13.15, S' is closed, because S is closed.

By the induction hypotheses, S' is the relative convex hull of its extreme points E' on $P'^{\#}$. Clearly, the only extreme point of S' in T is v_{12}. Put $E'' = E' \setminus \{v_{12}\}$ and, on P, $E = E'' \cup \{v_1, v_2\}$. It suffices to show that (i) E is the set of extreme points of S, and (ii) $S = \mathrm{conv}(E)$.

(i) Since E'' is included in the part of S' isometric to S, each point in E'' in also extreme for S, as are the vertices v_1, v_2.

(ii) One can easily see that each convex set on P' containing E' corresponds, via digon-tailoring T, to a convex set on P containing E, and conversely. So the intersection of all such sets on P' corresponds to the intersection of their correspondents on P, and conversely. Consequently, $S = \mathrm{rconv}(E)$.

The last claim of the theorem—that the interior extreme points of S are all vertices—follows from Proposition 13.27 and its proof. □

Note that throughout the induction, the boundary of all the sets remains fixed at ∂S.

The next example illustrates the exceptional situation of Theorem 13.32.

Example 13.33 (Extreme Points)**.** *Consider a "tall" triangular prism with top abc and base $a'b'c'$. Now add a point v on the line containing bb', above b, and let P be the boundary of the extrinsic convex hull of $\{a', b', c', a, c, v\}$. See Fig. 13.7. Then $Q = abc$ is a simple closed quasigeodesic and the "roof" above Q—all the surface of P above Q—is our convex set. (Here we need the prism to be sufficiently tall, but for clarity the figure is more squat.) Take points $x, y \in P$ with x on ab close to a and y on bc close to c. Then the arc of Q from x to y through b is not a geodesic segment.*

Note that b is not an extreme point of S, because b is interior to geodesic segments connecting points on ab and bc that are close to b. The only extreme points of S are a, c and v, and so S is not the convex hull of its extreme points. This is the exception in Theorem 13.32.

Example 13.34 (Simple closed geodesic)**.** *For the simple closed geodesic Q on B_h, the rectangular box of height h considered in Example 13.29,* $\mathrm{conv}(Q) = Q$ *holds, and there are no extreme points on Q.*

Denote by $\mathrm{ext}(S)$ the extreme points of the convex set S.

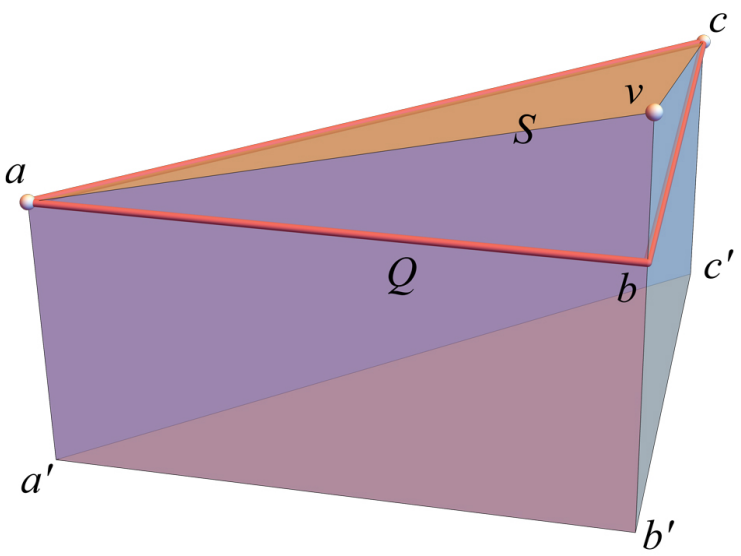

Figure 13.7: The region above the simple closed quasigeodesic Q is convex, but is not the convex hull of its extreme points a, c, v.

Lemma 13.35. *For any set $S \subset P$,* $\text{ext}(\text{conv}(S)) \subset S$.

Notice that not all extreme points of the closure of $\text{conv}(S)$ are in S, as Example 13.26 shows.

Proof. Assume there exists a point x in $\text{ext}(\text{conv}(S)) \setminus S$. Then, since x is not interior to any geodesic segment joining two points of $\text{conv}(S)$ (because it is extreme), the set $\text{conv}(S) \setminus \{x\}$ is still convex (see the proof of Lemma 13.9). So we have $S \subset \text{conv}(S) \setminus \{x\} \subset \text{conv}(S)$ and $\text{conv}(S) \setminus \{x\} \neq \text{conv}(S)$, contradicting the minimality of $\text{conv}(S)$. □

Lemma 13.36. *For any $S \subset P$, if* $\text{conv}(S)$ *contains a boundary geoarc which is not a geoseg, then S also contains that arc.*

Proof. Assume the contrary be true, and let $A \subset \text{conv}(S)$ be a maximal (with respect to inclusion) such geoarc not in S. So $A \subset \partial(\text{conv}(S)) \setminus S$. Remove A from $\text{conv}(S)$, but leave its extremities, and denote by K the resulting set. It follows, just as in the proof of Lemma 13.9, that K is a convex set including S and strictly included in $\text{conv}(S)$, a contradiction to the inclusion-minimality of $\text{conv}(S)$. □

Example 13.31 also illustrates this result. Indeed, in Fig. 13.6, the arc aea is a boundary geoarc of $\text{rconv}(\{a, c, d\})$, but not a geoseg. And that arc is not included in $\{a, c, d\}$, hence neither is it included in $\text{rconv}(\{a, c, d\})$.

13.9 Relative Convex Hull of Vertices

We return to the notion of *relative convexity* and the *relative convex hull*, used in the proof of Theorem 13.32. For our purposes we only consider the relative convex hull of vertices V that fall to one side of (i.e., *inside*, or above) a simple closed quasigeodesic Q.

Notation.

(1) Q quasigeodesic. Orient Q counterclockwise (ccw), and let P^+ be the half-surface to the left of (above) Q. $P^\#$ is P^+ union the unbounded cylinder to the right of (below) Q.

(2) V: Set of (positive curvature) vertices inside or on Q.

(3) For C any simple closed curve in P^+, let $R(C)$ be the region of P^+ to the left of C. So $\partial R(C) = C$. This is well-defined because C is a simple curve. And because it is oriented ccw, $R(C) \subset P^+$.

Given Q oriented so that V is to Q's left, define a polyhedron $P^\#(Q) = P^\#$ as in Section 13.5: Cut P along Q and insert below Q a sufficiently tall cylinder; for example, a cylinder of height equal to $\text{diam}(P)$. Because Q is convex to both sides, AGT implies that $P^\#$ is a convex polyhedron.

Note that Q is strictly $\alpha\beta$-convex on $P^\#$, because $\beta_i = \pi$ at all corners[2] of Q, and corners have positive curvature.

Now we define the *relative convex hull* $\text{rconv}(V) = H$ of V to be the intersection of all the convex sets that enclose V on $P^\#$. This is the same notion earlier explored in Section 13.7, but here we focus on vertices inside Q.

Lemma 13.37. *Let V be a set of points on P, inside a simple closed quasigeodesic Q. The relative convex hull $\text{rconv}(V) = H$ of V can be obtained from only employing geodesic segments, and consequently convex sets, inside Q (instead of constructing $P^\#$).*

[2]Of course, the vertices of Q with $\alpha = \pi$ are flattened when passing to $P^\#$.

Proof. Let $x, y \in H$ be two points in H, and γ a geodesic segment between them. We argue that γ cannot cross Q.

When the cylinder attached below $P^{\#}$ is unfolded flat, it forms a rectangle R. Suppose γ properly enters R at x' and exits at y'. Then the segment $x'y'$ along the top of R is shorter, a contradiction to our assumption. (Note $x'y'$ is a portion of Q, and it is possible that H shares that portion of Q.) □

A *node* n of a geodesic polygon N is a point of N interior to no geodesic subarc of N. Call it a *g-node* if at least one side of N incident to n is not a geoseg (instead a geoarc), and a *gs-node* otherwise.

Recall that Lemma 13.6 showed that, if S is closed and convex, then ∂S is strictly $\alpha\beta$-convex; and if S has interior points then ∂S is $\alpha\beta$-convex but not necessarily strictly. We now explore to what extent there is a converse to this lemma, a result we need to compute $\mathrm{rconv}(S)$.

Lemma 13.38 ($\alpha\beta$ converse)**.** *Let S be a set satisfying these conditions: (1) $S \subset R(Q)$ contains V, (2) ∂S is a geodesic polygon, and (3) any geoarcs of ∂S which are not geosegs, are not included in S.*

Then S is relatively convex if and only if ∂S is strictly $\alpha\beta$-convex at each gs-node, and $\alpha\beta$-convex at each g-node.

Proof. If S is relatively convex, the statement is covered by the $\alpha\beta$-convexity Lemma 13.6. The remainder of the proof establishes the converse: That if ∂S is $\alpha\beta$-convex at the gs- and g-nodes as above, then S is relatively convex.

Consider first two boundary points $x, y \in \partial S$ (the roles x and y will play, will be decided later) and let γ' be a geoseg joining them, with γ' not included in $\bar{S} = S \cup \partial S$. We now show that there is a subarc γ of γ' (possibly $\gamma = \gamma'$) that lies completely external to the interior $\mathring{S}$ of S.

Notice that γ' intersects ∂S finitely many times. Otherwise, since ∂S is formed by geoarcs, hence locally by geosegs, we could find arcs A of ∂S arbitrarily small, having common extremities with subarcs of γ'. But those arcs would be ramifying geosegs, impossible.

Therefore, possibly replacing γ' with a subarc γ, we may assume that γ lies completely external to the interior $\mathring{S}$ of S.

Let F be the region of P bounded by γ and the arc $\partial^* S$ of ∂S from x to y, so that F doesn't include S. Then $\mathring{F}$ contains no vertex, by hypothesis.

Now we are going to identify a geoarc g from x to y inside $\bar{S}$. The boundary of S is, in particular, a convex curve, hence we may construct with

AGT the double $S^{\#}$ of $\bar{S}$. Let g be a geoseg on $S^{\#}$ from x to y. It is included in a half-surface of $S^{\#}$, because geosegs do not branch. So we may assume that g is in $\bar{S}$.

We prove next that g is shorter than or equal to γ: $\ell(g) \leq \ell(\gamma)$. Notice that the cut locus $C(x)$ of x on $S^{\#}$ doesn't intersect g, which is a geoseg on $S^{\#}$. On $S^{\#}$, put $\{c_1, \ldots, c_k\} = C(x) \cap \partial\bar{S}$. Each point c_j is the extremity of two digons of $S^{\#}$ at x (one on each copy of $\bar{S}$). Remove all those digons from $\bar{S} \subset S^{\#}$ and zip close the result by AGT. We get a doubly covered flat surface, by Lemma 2.2. Denote by S^* the part of it bounded by $\partial^* S$ and g which doesn't include ∂S.

We compare now F and S^*, which are both flat (hence planar) polygons. Clearly, the angles α^* of $\partial^* S$ towards S^* are at most π, and they are smaller or equal to those towards F, denoted by β^*. (For convenience, we suppress the indices of α^* and β^*.) By the extension of Cauchy's Arm Lemma (Theorem 2.7). we obtain that g is shorter than or equal to γ.

Next we will consider separately the cases of x or y (or both) belongs to S. That case-distinction and the non-strict inequality will imply convexity, via strictness or non-strictness of $\alpha\beta$-convexity.

Case 1. x and y in S. The equality case appears if $\alpha^* = \beta^*$ everywhere, meaning that $\partial^* S$ is outside S. So if $x, y \in \partial S \cap S$ the inequality $\ell(g) < \ell(\gamma)$ is strict because of the strictness of $\alpha\beta$-convexity, hence all geosegs between them are in S.

Case 2. Consider now two points $x', y' \in \mathring{S}$. From these two points we will obtain x and y. Assume γ' exits S. Because ∂S is a convex curve towards S, γ' must intersect it at least two times. Choosing two such points x, y, the above argument implies the ramification of γ' assuming $|g| = |\gamma|$, hence not strict $\alpha\beta$-convexity, impossible. So γ' lies inside $\mathring{S}$.

Case 3. The case $x \in \partial S$ and $y' \in \mathring{S}$ can be treated analogously, completing the proof. □

Example 13.39. *The statement of Lemma 13.38 is not necessarily true if S does not include all of V: $V \setminus S \neq \varnothing$.*

(a) To see this, take the quadrilateral $H = abcd$ formed by three equilateral triangles obc, ocd, oda, and let P' be the double of H. Let P be obtained from P' by extending it below by long rectangles. See Figure 13.8(a). $Q = aobo'a$ is a simple closed quasigeodesic on P (with o' "opposite" to

o). Let $S = R(oco'do)$ be the closed set constituted by $\triangle ocd$ and $\triangle o'cd$. Observe that the boundary of S is not strictly $\alpha\beta$-convex: both the angles α and β at c are $120°$ (and similarly at d). Then by Lemma 13.6(i), which established that closed S are strictly $\alpha\beta$-convex, S is not convex. Indeed, there is a geoseg outside S from x to x' on the back, for x on oc (and similarly for points on od). And for the same reason, there are geosegs from o to o' outside S, namely oeo' and ofo'.

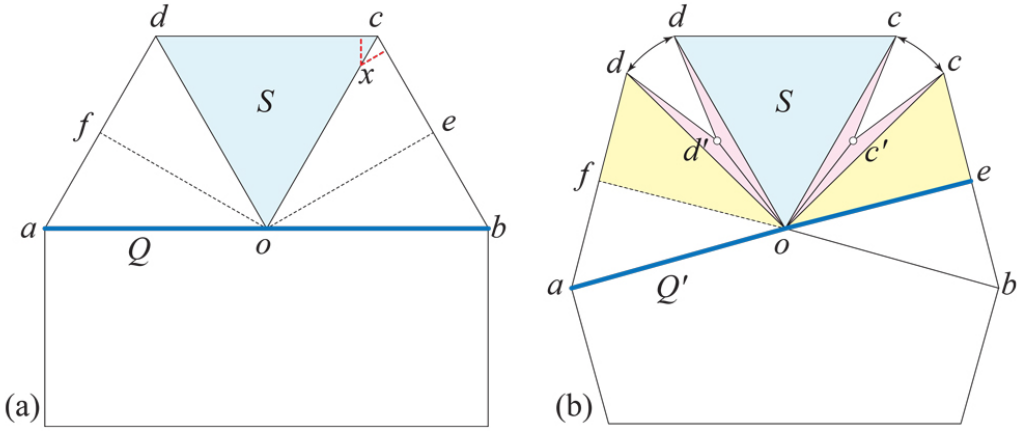

Figure 13.8: S is not convex. (a) Lemma 13.38 doesn't apply because ∂S is not strictly $\alpha\beta$-convex at c and d. (b) After insertion of digons, ∂S is strictly $\alpha\beta$-convex, but Lemma 13.38 does not hold in this case because S doesn't include c' and d'.

(b) Now we show that even altering the example so that every gs-node on ∂S is strictly $\alpha\beta$-convex, does not always permit the converse conclusion that S is convex. Modify P by inserting two digon doubled triangles along oc and od; see Figure 13.8(b). Each triangle has base angles $7.5°$, and so adds $15°$ to c, and d, and adds a total of $30°$ of angle at o. Cut out $30°$ at o below o, so that o is again flat. Now the Q used above is no longer a quasigeodesic, but $Q' = aoeo'a$ is a simple closed quasigeodesic. The effect of the digons is two-fold. First, c and d now are both strictly $\alpha\beta$-convex, and ∂S remains strictly $\alpha\beta$-convex at o and at o'. So S is closed, ∂S is a geoseg polygon, strictly $\alpha\beta$-convex at every gs-node (and it has no g-nodes). So S satisfies all but one condition of Lemma 13.38, which would allow concluding that S is convex. The

missing condition is that $V \setminus S \neq \varnothing$: S does not include the two digon vertices—call them c' and d'—which are now part of V. And again the geosegs oeo' and ofo' lie outside of S, showing that indeed S is not convex.

(c) *Next, alter S to include the digon vertices c' and d': $S' = R(oc'co'dd'o)$. So now $V \setminus S = \varnothing$. But ∂S is not $\alpha\beta$-convex at c' and d': they both have $\alpha = \beta = 165°$. So again the preconditions for Lemma 13.38 are not satisfied, and S' is not convex.*

(d) *If $S = R(oeo'fo)$, then o is a g-node of ∂S, which is strictly $\alpha\beta$-convex at o: $\alpha = 150°, \beta = 210°$. However, the geoarcs oeo' and $o'fo$ must not be included in S for Lemma 13.38 to apply. However, S is nevertheless relatively convex.*

(e) *If $S = Q = aoeo'a$ in Figure 13.8(b), then ∂S is not strictly $\alpha\beta$-convex at o ($\alpha = \beta = 180°$), but again the lemma does not apply because ∂S contains the geoarc oeo'. Again, S is relatively convex.*

The next result gathers the proven facts about $\mathrm{rconv}(V)$.

Theorem 13.40 ($\mathrm{rconv}(V)$)**.** *Let V be the set of all (at least three) vertices of P, inside a simple closed quasigeodesic Q. V might also include some, but not all, vertices on Q, such that $\sum_{v\in V} \omega_v < 2\pi$.*

Then $H = \mathrm{rconv}(V)$ is simply-connected and its boundary ∂H is a geodesic polygon, strictly $\alpha\beta$-convex at gs-nodes and $\alpha\beta$-convex at g-nodes, on $P^{\#}$. Moreover, $\mathrm{ext}(H) = V$.

It will be established in Theorem 14.12 that $H = \mathrm{rconv}(V)$ can be constructed in time $O(n^5 \log n)$ where $n = |V|$.

Proof. By Lemma 13.35, $\mathrm{ext}(\mathrm{rconv}V) \subset V$, and clearly $V \subset \mathrm{rconv}V$, hence $\mathrm{ext}(H) = V$. Therefore, each point in $\partial\,\mathrm{rconv}(V) \setminus V$ is interior to a geoarc, so $\partial rconv(V)$ is a geodesic polygon. The $\alpha\beta$-convexity follows now from Lemma 13.38.

The simply-connectedness follows from the algorithmic construction in Section 14.4, justified by Lemma 13.38. □

Example 13.41. *Notice that the claim that $\mathrm{ext}(H) = V$ may not hold when V does not include all vertices of P inside Q. This can be seen by adapting Example 13.26, Fig. 13.3, as follows.*

Imagine there is another vertex a' close to a, above a, and almost flat. (This vertex a' is added so that $|V| \geq 3$.) Then, in the upper-left half-surface S bounded by the simple closed quasigeodesic $Q = aca$, take $V = \{a, a', c\}$ (i.e., skip one vertex d inside Q). Then $H = \text{conv}(V) = S \setminus \{d\}$, by Proposition 13.27. Therefore, H is not simply-connected and ∂H is not a geodesic-segment polygon.

13.10 Summary of Properties

We attempt to succinctly summarize below the facts about convexity established in this chapter, either in lemmas and theorems, or via examples. The goal is to help the reader interested in applications of this theory. Please note that, unless otherwise specified, "convexity" means "ag-convexity."

Ag-Convexity

(1) S closed convex $\implies$ Either point, arc, geodesic, or has interior points. Lemma 13.2.

(2) Closure $\bar{S}$ of convex $S \not\implies \bar{S}$ convex. Example 13.4.

(3) S closed convex $\implies \partial S$ a convex curve, and if a geodesic polygon, then $\alpha\beta$-convex. Lemma 13.6

(4) S closed convex $\implies S$ encloses $\leq 2\pi$ curvature.

(5) S closed convex $\implies$ simply-connected, or cylinder. Lemma 13.8.

(6) S convex and $v \in S$ a vertex $\implies S \setminus \{v\}$ convex. Lemma 13.9.

(7) S convex and not closed $\not\implies$ simply-connected. Example 13.12.

(8) No Helly type theorem for $h = 3$. Example 13.13.

(9) S closed convex and includes $Q \implies$ cylinder, or $S = P$. Lemma 13.14.

Geodesic Segments and Convex Sets

(10) S convex $\implies$ γ connecting interior points x_1, x_2 does not meet ∂S. Lemma 13.15.

(11) S convex, x interior to S and $y \in \partial S$ $\implies$ all geodesic segments from x to y are interior to S, possibly excepting their extremity y. Lemma 13.15.

(12) S convex, $x, y \in \partial S$ then there is a geodesic segment between them included in the closure $\bar{S}$ of S. Lemma 13.15.

(13) S convex; $x \in \partial S$ is not itself a component of ∂S then there exists a geodesic segment starting at x and exterior to S. Lemma 13.15.

(14) S convex $\implies$ the interior of S is convex. Corollary 13.16.

(15) $S \neq P$ convex with interior points, $x \in \partial S$ $\implies$ supporting angle at x. Lemma 13.17.

Relative Convexity

(16) S convex and v_1, v_2 merge $\implies$ S' relatively convex. Lemma 13.18.

(17) S convex and v_1, v_2 merge $\not\Longrightarrow$ S' convex. Example 13.19.

Convex Hull Properties: $\mathrm{conv}(S)$

(18) V vertices of P $\implies$ $\mathrm{conv}(V) = P$. Example 13.20.

(19) $\mathrm{conv}(P \setminus V) = P \setminus V$. Example 13.21.

(20) No Radon type theorem for $r = 4$. Example 13.22.

(21) For $S' \subset S$, $\mathrm{conv}(S') \subseteq \mathrm{conv}(S)$. Lemma 13.23.

(22) For $T \subset S$, $\mathrm{conv}(S) = \mathrm{conv}\left((\mathrm{conv}(S \setminus T)) \cup T\right)$. Lemma 13.24.

(23) S open $\implies$ $\mathrm{conv}(S)$ open. Lemma 13.25.

(24) S closed $\not\Longrightarrow$ $\mathrm{conv}(S)$ closed. Example 13.26.

(25) S closed, convex, with interior points, enclosing $< 2\pi$ curvature $\implies$ $\mathrm{conv}(P \setminus S)$ dense in P, $\mathrm{conv}(\partial S)$ dense in S. Proposition 13.27.

(26) For box B_h and Q: $\mathrm{conv}(Q) = Q$ or $\mathrm{conv}(Q) = B$, depending on h. Example 13.29.

Relative Convex Hull: $\mathrm{rconv}(S)$

(27) Convexity and relative convexity can differ. Example 13.30.

(28) Neither a digon nor a triangle is necessarily convex; $\mathrm{conv}(V)$ is not necessarily closed; the closure of $\mathrm{rconv}(V)$ is not the convex hull of its extreme points. Example 13.31.

Extreme Points: $\mathrm{ext}(S)$

(29) S closed convex with ∂S enclosing $< 2\pi$ curvature. Either $S = \mathrm{rconv}(\mathrm{ext}(S))$, or S contains a non-segment geoarc. Theorem 13.32. Exception: Examples 13.31, 13.33.

(30) S closed, convex and $\mathrm{conv}(S) \neq \mathrm{conv}(\mathrm{ext}(S))$. Example 13.33.

(31) S closed convex with $\mathrm{ext}(S) = \varnothing$. Example 13.34.

(32) $\mathrm{ext}(\mathrm{conv}(S)) \subset S$. Lemma 13.35.

Not all extreme points of the closure of $\mathrm{conv}(S)$ are in S. Example 13.26.

(33) $\partial\mathrm{conv}(S)$ contains a non-segment geoarc only if S does so. Lemma 13.36.

Relative Convex Hull of Vertices

(34) $S \subset R(Q)$, points V in S: $\mathrm{rconv}(S)$ can be constructed without $P^{\#}$. Lemma 13.37.

(35) $S \subset R(Q)$, $V \subset S$, ∂S $\alpha\beta$-convex $\implies$ S relatively convex. Lemma 13.38.

(36) Lemma 13.38 is not true if $V \setminus S \neq \varnothing$. Example 13.39.

(37) V all vertices inside Q, + some on Q $\implies$ $H = \mathrm{rconv}(V)$ simply-connected, ∂H a geodesic polygon, $\alpha\beta$-convex on $P^{\#}$. $\mathrm{ext}(H) = V$. Theorem 13.40.

(38) $\text{ext}(H) \neq V$ if V does not include all vertices inside Q. Example 13.41.

Chapter 14

Minimal-length Enclosing Polygon

In this chapter we explore the minimial-perimeter polygon Z enclosing a set of vertices V in or on a quasigeodesic Q. We derive its key properties, and provide a polynomial-time algorithm for constructing it (Section 14.2). We then show that this polygon is not always the same as $\partial\,\mathrm{rconv}(V)$ (Example 14.10). The algorithm for Z also works to construct $\partial\,\mathrm{rconv}(V)$ with minor modifications (Section 14.4).

In the next chapter we will show that either of these notions can support an algorithm that ensures that sequential vertex-merging cuts do not disconnect P^+.

14.1 Properties of the Minimal Enclosing Polygon

Notation. We repeat some previous notation and introduce new notation to be used subsequently. We again use the abbreviation *geoseg* to mean "geodesic segment" and *geoarc* to mean a simple geodesic that may or may not be a geoseg.

(1) Q quasigeodesic. Orient Q counterclockwise (ccw), and let P^+ be the closed half-surface to the left of (above) Q. $P^\#$ is P^+ union the unbounded cylinder to the right of (below) Q.

(2) V: Set of (positive curvature) vertices inside or on Q.

(3) For C any simple closed curve in P^+, let $R(C)$ be the closed region of P^+ to the left of C ("R" for region). So $\partial R(C) = C$. This is well-defined because C is a simple curve. And because it is oriented ccw, $R(C) \subset P^+$.

(4) For C any curve on P, $\ell(P)$ denotes its length.

Definition of $\min \ell$**.** For a finite set V of points in the plane, $\text{conv}(V)$ is precisely the set bounded by the minimal length curve enclosing V. We adapt this property for vertices $V \subset P^+$. Toward that end, we define $Z = Z(V) = \min \ell[V]$ to be the minimal-length geodesic polygon enclosing V. Here *enclosing* is to the left when Z is oriented ccw. We must also stipulate that "enclosing" means that the interior of the region between Z and Q is empty of vertices. Otherwise a clockwise traversal of the boundary of any triangle $\triangle \subset P^+$ would enclose V to its left, and in fact would enclose all of $P \setminus \triangle$. Also, for a simple closed geodesic enclosing V, its position is not fixed by the above requirements, in the sense that it could be "slid" parallel to itself while remaining a geodesic of the same length. In such a case, we define Z to be "lifted" to a position that touches a vertex.

Each edge of the geodesic polygon Z is a geoarc but not necessarily a geoseg. See, e.g., Fig. 14.2(b). We view polygons, and so Z, as the boundary, not the region enclosed, which is denoted by $R(Z)$.

Here we concentrate on properties of Z and algorithms to construct Z. We stress that all the following considerations take place on $P^\#$ and not on P. For simplicity, we will omit in the rest of this section the modifier "relative," which would refer precisely to $P^\#$.

We first establish three basic properties of $Z = Z(V) = \min \ell[V]$.

Lemma 14.1 ($\alpha\beta$-convex)**.** *Z is strictly $\alpha\beta$-convex.*

Proof. At vertices $v \in Z \cap Q$, $\beta = \pi$ on $P^\#$. On the other hand, at vertices $v \in Z \setminus Q$ with $\beta < \pi$ one could further shorten Z, a contradiction to minimality.

Since $\beta \geq \pi$ and $\alpha + \beta < 2\pi$ at vertices, $\alpha < \pi \leq \beta$. □

Lemma 14.2 (Z simple)**.** *Z is simple, i.e., non-self-crossing.*

Proof. Suppose Z is not simple. At each self-crossing, clip off the portions of the geodesics toward the inside. The resulting outermost geodesic polygon still encloses all of V, and is shorter. See Fig. 14.1. □

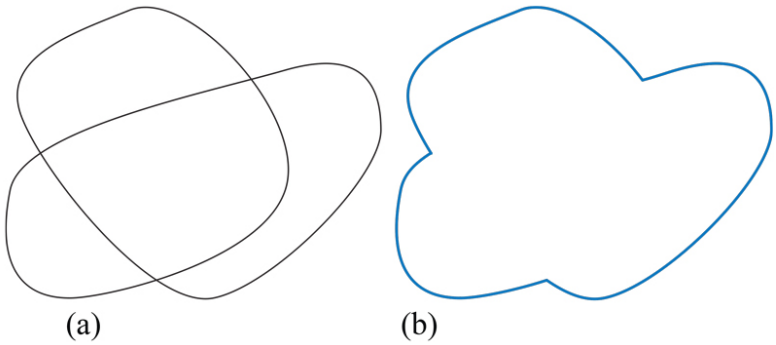

Figure 14.1: (a) Non-simple polygon. (b) Outermost arcs.

Lemma 14.3 (Z Unique)**.** *Let Z be a shortest geodesic polygon enclosing V. Z is unique, up to isometry.*

Proof. Assume there is another geodesic polygon Z' also enclosing all of V and of the same length: $\ell(Z) = \ell(Z')$. Then either one is enclosed in the other, or the two polygons cross at two or more points.

Case 1. Assume Z' is nested inside Z.

Case 1a: First assume Z and Z' are disjoint. Because all vertices are on or interior to Z', Z must be a simple closed geodesic. Let A be the annulus between Z and Z'; A contains no vertices. Intrinsically, A is a subset of a cylinder with lower circular rim Z, and upper boundary Z'. Thus we can project any point $p \in Z'$ orthogonally onto Z. See Fig. 14.2(a). Thus $\ell(Z) \leq \ell(Z')$, with equality when Z and Z' are parallel. In either case, we can "lift" Z parallel to itself until it touches a vertex in V. As mentioned earlier, we choose this isometric version of Z as $\min \ell[V]$.

Case 1b: Assume Z and Z' share some portion of their boundaries. Again because all vertices are on or inside Z', this means that there must be a geodesic arc γ of Z, outside of Z', between two shared vertices v_1 and v_2. See Fig. 14.2(b). Let B be the region of P bounded by γ and the edges E of Z' between v_1 and v_2. Then (a direct induction shows that) B is isometric to a plane polygon, with one straight edge γ, and E a polygonal chain. Therefore E is strictly longer than γ, and so $\ell(Z) < \ell(Z')$.

Case 2. Neither of Z or Z' is inside the other. So they must cross at least twice, with always the outer arc a geodesic γ, and the inner sequence of edges

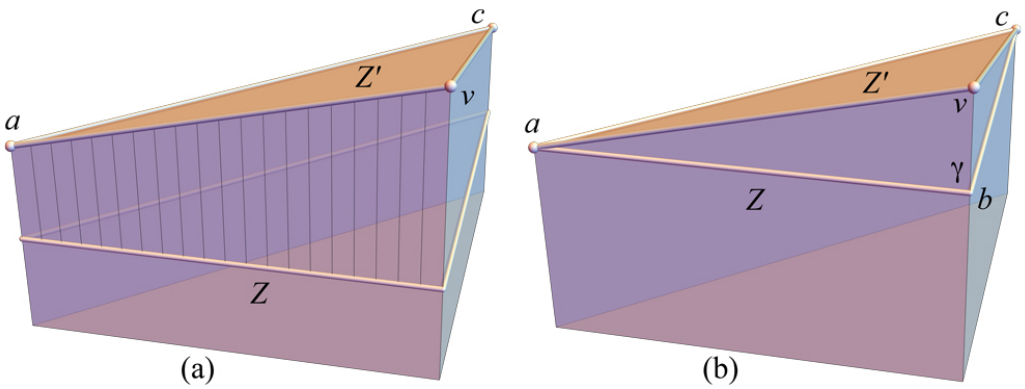

Figure 14.2: $Z' = avc$. (a) Edge av projected onto Z. (b) γ with ab, bc forms a planar polygon.

possibly containing vertices. Call such a region a lune L. Now the argument in Case 1b applies: L is isometric to a planar polygon with straight edge γ, and so the outer arc of L is strictly shorter than the inner arc. Now form Z'' by following all outer arcs, and otherwise the common portions of the two polygons. Z'' still encloses all of V, but it is strictly shorter than Z and Z'. This contradiction to the assumption that both Z and Z' are shortest establishes the lemma. □

This next lemma is a counterpart to the $\alpha\beta$-converse lemma, Lemma 13.38 in Chapter 13. In the following, a *pinned node* v of a geoarc polygon G is a flat point which behaves like a vertex, in the sense that G is not allowed to be shorten at v towards its interior.

Lemma 14.4. *Assume that all nodes of a geodesic polygon $G \subset R(Q)$ are either vertices or pinned, and that $V = V(P^+) \subset R(G)$. Then $G = Z(V)$ if and only if $\beta \geq \pi$ at each node of G.*

Proof. That $\beta \geq \pi$ for $Z(V)$ was established in Lemma 14.1.

In the following we show that $G = Z(V)$, using reasoning similar to that used to prove uniqueness in Lemma 14.3.

In Case 1 of that lemma's proof, one of G or Z is nested inside the other. Case 1a: Assume the two polygons are disjoint and Z is inside G. Then the lemma argument leads to $\ell(G) \leq \ell(Z)$, but because Z is shortest, we must have $\ell(G) = \ell(Z)$; and so $G = Z$. If instead G is inside Z, then recall that Z must be a simple closed quasigeodesic. If Z and G are not identical, then

the arc of Z between touching points of G is a short-cut of G, contradicting the assumption that G has no further short-cuts.

In Case 1b, G and Z share some portion of their boundaries. If Z is inside G, the lemma argument shows that $\ell(G) < \ell(Z)$, contradicting the minimality of Z. If G is inside Z, then consider the region B in the proof of Lemma 14.3, bound by γ and the edges E of G between v_1 and v_2. B is isometric to a plane polygon, and so some vertex along E must be convex. We have thus identified a vertex on G with a right angle $\beta < \pi$. So G can be short-cut there, contracting the assumption that no more G-short-cuts were possible.

In Case 2, neither is nested in the other, so they must cross. This leads to a polygon shorter than either G or Z, a contradiction to the minimality of Z. □

Example 14.5. *The assumption of pinned nodes in Lemma 14.4 is clearly necessary. Here we show that the second assumption, that $V \subset R(G)$, is also necessary.*

Consider an equilateral triangle abd' inscribed in the rectangle $abce'$, and choose $d \in ce'$ between d' and e', and $e \in ae'$ such that $ed \perp bd$. See Figure 14.3. Extend the sides $e'a$ and cb sufficiently beyond a and b respectively, and let L be the double of the extended $abcde$. Let $V = \{b, c, d\}$. Note that $e \notin V$.

On L, $Q = aba$ is a simple closed quasigeodesic, and the geodesic polygon $G = bdb$ has each of its angles β towards Q at least π. However, $e \notin R(G)$ and $Z(V) = Q$, because $|a - b| = |b - d'| < |b - d|$.

We prove here the following result, which will be employed in the next chapter.

Lemma 14.6. *For any point $p \notin V$, $R(Z(V)) \subseteq R(Z(V \cup \{p\}))$.*[1]

Proof. The argument is similar to the one proving the uniqueness of Z.

Set $V' = V \cup \{p\}$. Clearly, $R(Z(V)) = R(Z(V'))$ if $p \in R(Z(V))$. So assume that p is strictly exterior to $R(Z(V))$.

Put $I = R(Z(V)) \cap R(Z(V'))$ and assume there exists $q \in R(Z(V)) \setminus R(Z(V'))$. If the lemma were to hold, then $I = R(Z(V))$ and $R(Z(V)) \setminus$

[1] We believe that also $\ell(Z(V)) \leq \ell(Z(V \cup \{p\}))$, but this inequality is not important in the following.

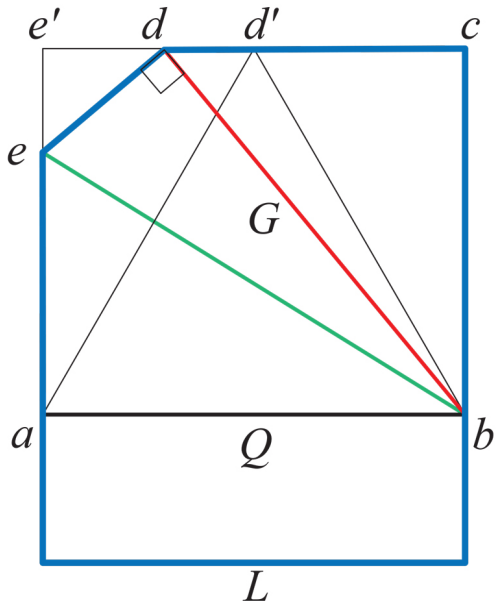

Figure 14.3: Lemma 14.4 is optimal. $V = \{b, c, d\}$, $G = bdb \neq Z = Q$.

$R(Z(V')) = \varnothing$, so there could be no such q. We now prove that the existence of q leads to a contradiction.

Since $V \subset I$ and q is outside I, q lies on an geoarc $A \subset Z(V)$ with extremities on $Z(V')$. Those extremities determine an arc $A' \subset Z(V')$ such that $A \cup A'$ bounds a flat polygonal domain L.

Because of the length minimality of both $Z(V')$ and $Z(V)$, $\ell(A') = \ell(A)$.

On the other hand, if A' is not a geoarc, the flat polygon bounded by $A \cup A'$ has the side A necessarily shorter than A', a contradiction. And if A' is a geoarc, L is a flat digon, contradicting the Gauss-Bonnet Theorem. □

14.2 Shortening Algorithm

Our goal in this section is to provide a polynomial-time algorithm to construct $\min \ell[V]$.

14.2.1 Curve-Shortening Flow

Here we digress to highlight the intuition behind curve-shortening. The famous Gage-Hamilton-Grayson (GHG) theorem says that,

> If a smooth simple closed curve undergoes the curve-shortening flow, it remains smoothly embedded without self-intersections. It will eventually become convex, and once it does so it will remain convex. ... The curve will then converge ... to a "round point," [a vanishingly small circle]. [Wik21]

First assume that all of V is strictly inside Q: Q and V are disjoint. At every vertex u of Q strictly convex to its left, replace u with a small arc of a circle, smoothly joined to the incident edges of Q. This is now a smooth simple closed convex curve C, a smooth approximation of Q, to which the GHG theorem applies. Shortening C will either result in a simple closed geodesic, which can no longer be shortened, or it will hit one or more vertices in V.

Consider these vertices *pinned*, partitioning C into sections between pinned vertices. A recent result [ALT12] established that "open curves with fixed endpoints evolving ... do not develop singularities, and evolve to geodesics." Thus the arcs of C between pinned vertices shorten to geodesics. The result is our shortest enclosing geodesic polygon Z.

Although there is recent work on discrete curve-shortening, e.g., [AN19] and [EHPN20], it seems it could be difficult to turn this smooth curve-shortening viewpoint into a formal proof or a finite algorithm. So we take a different approach.

14.2.2 Algorithm Overview

We now turn to a discrete, finite algorithm, which we will see is in some sense the reverse of the curve-shortening flow. Imagining Q as roughly equitorial, the flow as described above shortens Q "upward" to Z, whereas the algorithm to be described starts with G "above" Z and shortens it "downward" to $Z = \min \ell[V]$.

The algorithm consists of two steps. The first step constructs a geodesic polygon G enclosing V. The second step shortens G to Z.

14.2.3 Finding an Enclosing Geodesic Polygon

Our goal here is to construct a geodesic polygon G that encloses all vertices in V.

It is tempting to build G from the subgraph of the 1-skeleton of P induced by the vertices in V. However, this subgraph is not always connected, as illustrated in Fig. 14.4. As it seems difficult to base an algorithm on this subgraph, we pursue a different approach.

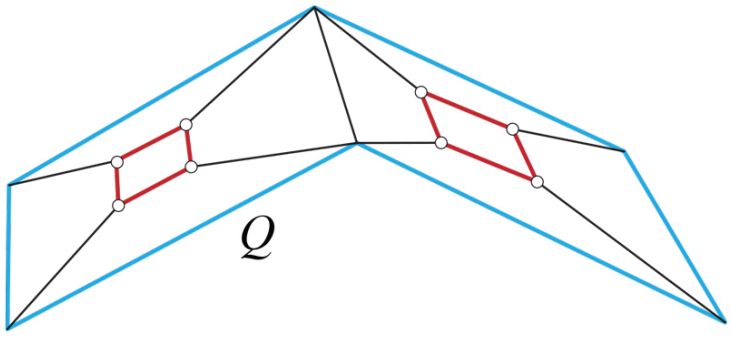

Figure 14.4: Q (blue) encloses V, the 8 marked vertices. The subgraph V induces is the two disconnected red quadrilaterals.

We start by computing all geosegs between pairs of vertices in V, and then identify those geoseg subsets ("edgelets") on the outer boundary of the resulting planar subdivision.

A *planar subdivision* is a data structure $\mathcal{R}$ that maintains a subset of the plane partitioned into regions. Often called a Planar Straight-Line Graph (PSLG) when its edges are straight-line segments, $\mathcal{R}$ represents an embedding of a planar graph, and maintains the incidence relations between its regions and edges. Such data structures support updates (for us, the addition of new geosegs) in $O(\log n)$ time per intersection.

The algorithm described below has time complexity at worst $O(n^5 \log n)$, although a careful analysis might lead to $O(n^4 \log n)$. Our focus is more to show it is polynomial-time, as opposed to optimizing efficiency.

We first describe the three stages of the algorithm at a high-level.

(1) Calculate all geosegs between pairs of vertices in V. Let Γ be the collection of these geosegs.

(2) Intersect all geosegs in Γ, incrementally building a planar subdivision data structure $\mathcal{R}$ on $P^+ = R(Q)$. As each $g \in \Gamma$ crosses the others already processed, the surface is partitioned into convex regions, call them *g-faces* bounded by *edgelets*. Each g-face is flat and isometric to a planar convex polygon; an edgelet is a (generally short) subsegment of a geoseg g, between two points where other geosegs cross g. G-faces do not necessarily lie in one face of P, but they are always empty of strictly interior vertices, or internal edgelets. Fig. 14.5 illustrates a simple such a partition; Fig. 14.7 below is a bit more complicated. As the intersections are calculated, the data structure $\mathcal{R}$ is updated to record ccw orientation of each newly created g-face. The process parallels the incremental algorithm for constructing an arrangement of lines; see Fig. 14.6.

(3) Each interior edgelet is shared by exactly two g-faces, recorded in $\mathcal{R}$. A edgelet e shared with just one g-face must be on the boundary of the outer face of the planar subdivision. Connecting all the outer-boundary edgelets yields a geodesic polygon G that encloses all vertices in V.

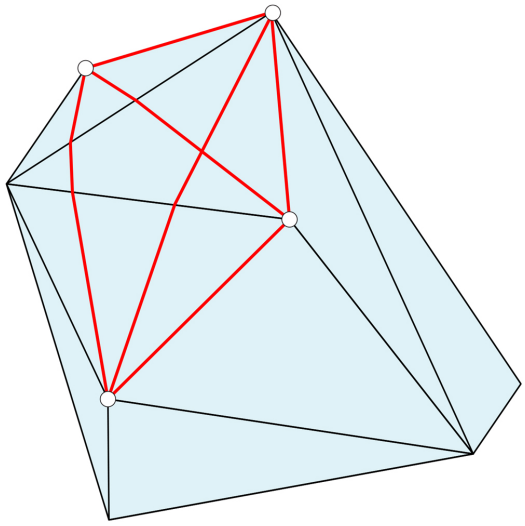

Figure 14.5: V is the four marked vertices. Geosegs (red) form four g-faces.

Next we discuss some computational details for each of the three stages.

(1) The total number of geosegs between vertices of P is $O(n^2)$, each of which can be found in $O(n \log n)$ by the algorithm in [SS08]. Although

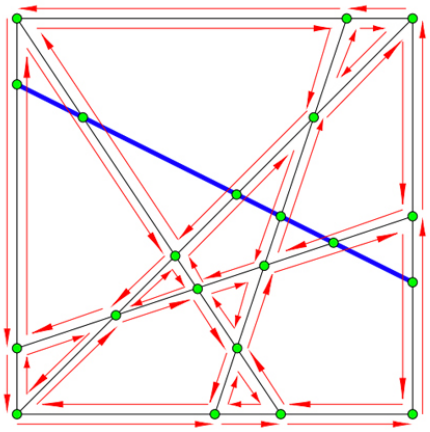

Figure 14.6: Inserting a line into a line-arrangement data structure. (Image from Jan Verschelde Lecture Notes: `http://homepages.math.uic.edu/~jan/mcs481/arrangements.pdf`)

two vertices may have more than one geoseg connecting them, the number of geosegs from one vertex v and the $n-1$ other vertices is still $O(n)$. This claim can be proved by induction using the cut locus partition (Lemma 2.6 in Chapter 2. For another proof idea, notice that each extra pair of geosegs between two vertices must have a vertex separating them. And there are only n vertices.

(2) The total combinatorial size of the data structure $\mathcal{R}$ is $O(n^4)$, the same size as an arrangement of $\binom{n}{2}$ lines. That each of the $O(n^4)$ edgelets may cross several edges of P is accounted for in the implicit data structure of [SS08].

(3) We note that the sides of G are not necessarily full geosegs between vertices. G in the example shown in Fig. 14.7 includes two edgelets v_1x and v_3x whose containing geosegs are not wholly on the outer boundary. This example also shows that G is not necessarily convex: it is reflex at x.
The complexity of any face, including the outer face, of an arrangement of segments is just slightly more than linear in the number of segments. Therefore, since G is built from $O(n^2)$ geosegs, the total

number of edgelets on G is $O(n^2\alpha(n))$, where α is the inverse Ackermann function [CEG+93].

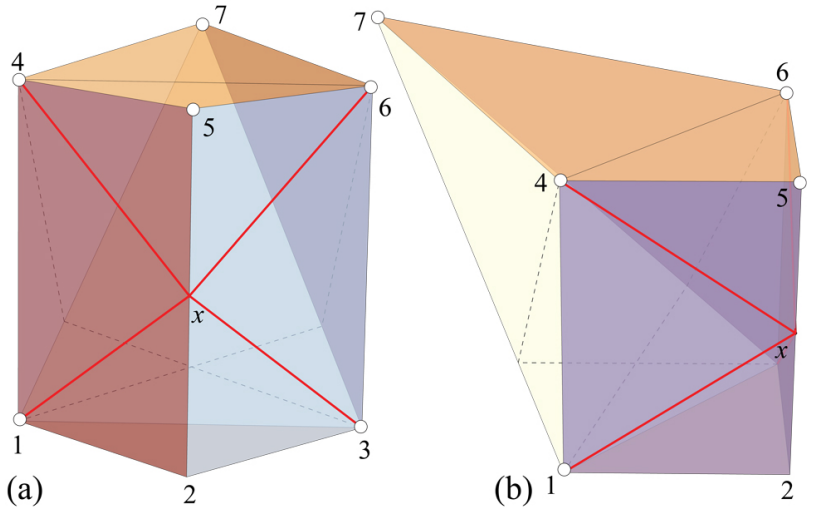

Figure 14.7: V is all the vertices except excluding v_2. $Q = v_1v_2v_3$. $G = v_1xv_3v_1$. g_{16} and g_{35} (red) on the front are geosegs. The dashed paths around the back are longer and so not geosegs.

A more formal presentation is displayed below as Algorithm 6.

There are two special cases. When V is a single vertex v, we allow $G = \{v\}$. And when $|V| = 2$, say, $V = \{a, b\}$, G is the ccw loop aba.

We have established this lemma:

Lemma 14.7. *Let P have n vertices, and V a set of vertices in or on quasi-geodesic Q (possibly including all vertices of Q). Then a geodesic polygon G enclosing all of V can be constructed in $O(n^5 \log n)$ time.*

14.2.4 Algorithm for Curve Shortening

We now describe an algorithm for shortening $G = G_0$ to $Z = \min \ell[V]$. We will repeatedly shorten G_0 to $G_1, G_2, \ldots, G_k = G$, until $G_k = Z$ for some $k \leq n = |V|$. We next describe one shortening step.

Case 1. $|V \cap G_j| \geq 3$. Let v_{i-1}, v_i, v_{i+1} be three consecutive nodes of G_j (for some j), and let α_i and β_i be the angles at v_i left and right of G_j. If

Algorithm 6: Algorithm to find an enclosing geodesic polygon

Input : n vertices V inside or on quasigeodesic Q, on $P^{\#}$.
Output: Geodesic polygon G enclosing V.

`// Calculate all geosegs between pairs of vertices:` $O(n^3 \log n)$.
foreach $\binom{n}{2}$ *pairs of vertices in* V **do**
  Compute geosegs v_i to v_j, each in $O(n \log n)$ time [SS08].
end
Let Γ be the set of these $O(n^2)$ geosegs.

`// Compute g-faces and edgelets:` $O(n^5 \log n)$.
foreach *geoseg* $g \in \Gamma$ **do**
  Intersect g with previously built data structure $\mathcal{R}$.
  Create g-face regions, orient each ccw, bounded by edgelets.
  $O(n \log n)$ per geoseg insertion.
end

`// Identify outer boundary:` $O(n^4 \log n)$.
foreach *edgelet* e **do**
  Check in data structure $\mathcal{R}$ if e is shared between two g-faces.
  If not, e is on the boundary of the outer face of the subdivision:
  Then add e to G.
end
return G.

$v_i \in Q$, or if $\beta_i \geq \pi$, we take no action—G_j should not or can not be shortened at v_i. If however $v_i \notin Q$ and $\beta_i < \pi$, then we can locally shorten G_j in a neighborhood of v_i. But rather than locally shorten, we shorten by replacing the sequence (v_{i-1}, v_i, v_{i+1}) by (v_{i-1}, v_{i+1}). This is a shortening because the triangle inequality holds for geodesic triangles on a convex polyhedron. Let $\gamma = v_{i-1}v_{i+1}$.

Before proceeding with the description, we re-examine the example in Fig. 14.8. In the example previously shown in Fig. 14.2, we have $G_0 =$

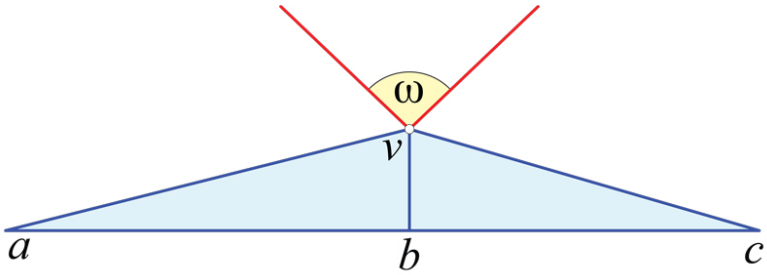

Figure 14.8: Flattening the region between G and Q in Fig. 14.2(b). $\alpha \approx 94°$, $\beta \approx 150°$, $\omega \approx 116°$.

(a, v, c), and $Q = Z$ as illustrated. Here $\beta < \pi$, and flattening the region below v leads to the planar triangle, short-cut by geoarc $\gamma = abc$.

Returning to the shortening step, the region of P^+ between G_j and Q is empty of vertices. So the path γ does not cross any portion of G_j, nor can it cross Q, because the $P^\#$ construction would require γ to cross below and return above Q, violating the fact that geodesics do not branch. So the triangle $v_{i-1}v_iv_{i+1}$ is between G_j and Q, and empty of vertices. By the triangle inequality, γ is shorter than $|v_{i-1} - v_i| + |v_i - v_{i+1}|$.

Another way to view this construction is to imagine a point p initially at v_i, sliding down the edge $v_{i-1}v_i$ until it reaches v_{i-1}. Then the line segment pv_{i+1} never encounters a vertex during this motion, and leads to $pv_{i+1} = v_{i-1}v_{i+1} = \gamma$.

Applying one of these shortening steps moves a node that was originally on G—either a vertex or an intersection of edgelets—to the interior of G_j. Because the combinatorial size of G is bounded by $O(n^2\alpha(n))$, there can be no more than that nearly quadratic number of steps.

When no further shortening steps are possible, G_k is a geodesic polygon with $\beta_i \geq \pi$ at every vertex v_i. By Lemma 14.4 then $G_k = Z = \min \ell[V]$.

Before passing to the other cases, we notice that the total curvature ω of $V \subset R(Q)$ is at most 2π.

Case 2. $|V \cap G_j| = 2$ and G_j is a digon with endpoints v_1 and v_2. Let β_i be the exterior angles at v_i, and let $\tau_i = \beta_i - \pi$ be the turn angles at v_i. Gauss-Bonnet requires that $\omega + \tau = 2\pi$. If both $\beta_i < \pi$, then $\tau_1 + \tau_2 = \tau < 0$, and $\omega + \tau = 2\pi$ cannot be satisfied. If both $\beta_i \geq \pi$, the algorithm halts. So let $\beta_2 < \pi$.

Because Q is a convex curve, we can merge inside $R(Q)$ all vertices it encloses to obtain one apex a of a cone Υ. Clearly, the vertex-free region between Q and G_j is isometric to a subset of Υ. Unfold Υ by cutting it open along the cone generator av_1. Then there exists a geodesic loop at v_1 enclosing $R(G_j)$ and strictly shorter than G_j.

A special case is when $|V| = |V \cap G_j| = 2$ and $\alpha_i = 0$. For example, in Fig. 14.9, with $V = \{c, d\}$, $G = cdc$, and short-cutting at d leads to $Z = cfc$.

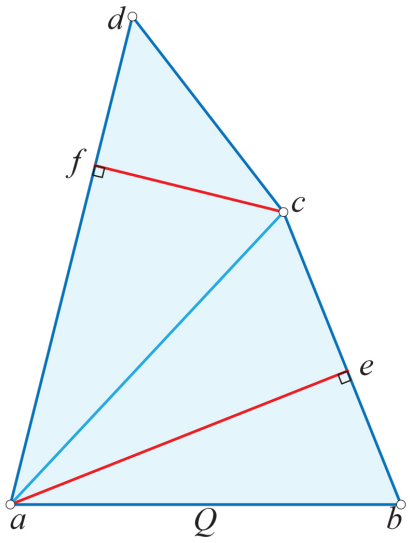

Figure 14.9: Doubly-covered quadrilateral. If $V = \{a, c, d\}$, $G = aca$, $Z = aea$. If $V = \{c, d\}$, $G = cdc$, $Z = cfc$. Vertex c is nonconvex in the doubled $\triangle acd$.

Case 3. $|V \cap G_j| = 1$. Then G_j is a geodesic loop at v, where Gauss-Bonnet implies that necessarily $\beta \geq \pi$. Therefore the algorithm halts.

This completes the description of the shortening algorithm.

A few remarks on the shortening algorithm:

- The short-cut edge $\gamma = v_{i-1}v_{i+1}$ is a geoarc, not a geoseg, because if γ were a geoseg, it would have been incorporated into the construction of $G = G_0$. So every shortening step inserts a geoarc. Later (in Chapter 15), algorithms will vertex-merge along such a geoarc (beyond a geoseg).

- The algorithm relies on the property that the region of P^+ between G and Q is empty of vertices. All the shortenings insert arcs in this vertex-free region. For that reason, V must include all vertices inside Q, i.e., the algorithm may not work if V is a proper subset of the vertices inside Q. (However, the construction of G still works, but the shortening may not.) So in this sense, it is not a general "convex hull" algorithm, but rather a construction tailored to our needs.

- The algorithm also works for all vertices inside a convex curve C, because doubling $R(C)$ via AGT transforms C into a quasigeodesic.

- Even when V includes all vertices on Q, it may be that $G \neq Q$, if Q contains geoarcs. But then the shortening algorithm will result in $Z = Q$.

- Viewing P^+ as the upper half-surface, one can view Q as the rim of the base of P^+, Z a closed convex curve "above"/inside Q, and G a closed curve "above"/inside Z. The shortening algorithm works by, in some sense "area growing" G down to Z, even though each growing step is in fact length-shortening. In contrast, the curve-shortening flow described in Section 14.2.1 "raises" Q until it matches Z.

We have now established this theorem:

Theorem 14.8 (Z-Algorithm)**.** *The minimum-length geodesic polygon enclosing V, $Z = Z(V) = \min \ell[V]$, can be constructed in polynomial-time, specifically in time $O(n^5 \log n)$ where $n = |V|$.*

14.3 AG-convexity and Z

To parallel the notation $Z = Z(V) = \min \ell[V]$, define $W = W(V) = \partial\,\mathrm{rconv}(V)$.

Lemma 14.9. *The set $R(Z)$ enclosed by $Z = Z(V)$ is a relatively ag-convex set containing V, but it is not necessarily equal to* $\mathrm{rconv}(V)$.

Proof. The relative convexity of $R(Z)$ follows directly from the strict $\alpha\beta$-convexity of Z and Lemma 13.38. The second claim is established by the next example. □

Example 14.10. *It could be that $Z(V) \neq W(V)$.*

Let P be the octahedron illustrated in Fig. 14.10. The shape is symmetric front-to-back and top-to-bottom. Triangles acd and $ac'd'$ are congruent $45° - 45° - 90°$ triangles, and triangles acc' and add' are congruent equilateral triangles. Q (red) is $aba'ea$, and $V = \{a, c, d, a'\}$. Let $G = acda'$ as illustrated. Then $\alpha_i < \beta_i$ at each of the four vertices / gs-nodes. By Lemma 13.38, $\mathrm{rconv}(V)$ *is the union of the two triangles acd and $a'cd$. However, because β at both c and d is strictly less than π, the algorithm would short-cut there, leading to $Z = aba'ea = Q$. So $Z(V) \neq W(V)$.*

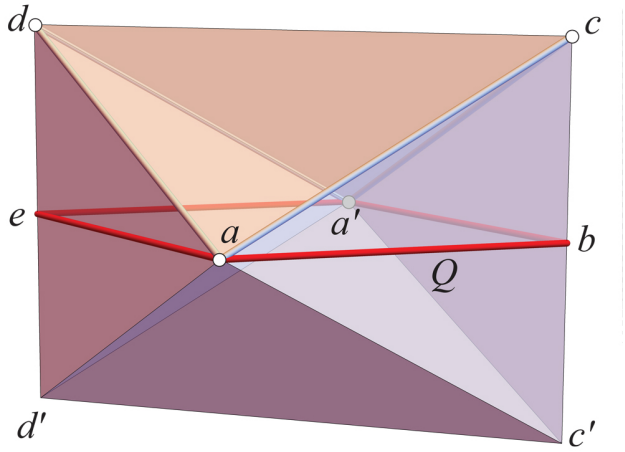

Figure 14.10: $Q = aba'ea$, $V = \{a, c, d, a'\}$. $Z = aba'ea = Q$. $W(V) = aca'da$. $Z(V) \neq W(V)$.

It follows, directly from Lemma 14.9 and the definition of rconv(V), that $R(W) \subseteq R(Z)$, with $Z = Z(V)$ and $W = W(Z)$. So of the two generalizations of the planar convex hull, rconv(V) is "tighter" than $\min \ell[V]$.

14.4 Algorithm for rconv$(V) = R(W)$

An algorithm to construct $W = W(V)$ can be obtained by minor modification of the algorithm for $Z = Z(V)$. The first part, calculation of a geodesic polygon G enclosing V, is identical. The second part, short-cutting G, can be followed with just different criteria of when to short-cut. We will continue to call it a "short-cut" even though it is no longer aimed at length-shortening.

Just as the construction of Z relies on Lemma 14.4, the construction of W relies on the $\alpha\beta$-converse Lemma 13.38. Assuming the preconditions of these lemmas are satisfied, then the short-cutting decisions of the two algorithm are as follows.

Z: If G_j is such that, for every v on G_j, $\beta \geq \pi$, then $G_j = Z$. So we short-cut whenever $\beta < \pi$.

W: If G_j is such that, for every v on G_j,

(a) if v is a gs-node, then v is strictly $\alpha\beta$-convex. So we short-cut whenever $\alpha \geq \beta$. For positive curvature v, this implies $\beta < \pi$.

(b) if v is a g-node, then v is $\alpha\beta$-convex. So we short-cut whenever $\alpha > \beta$. And this implies $\beta < \pi$.

then $G_j = W$.

We next explain in some detail how the short-cutting conditions for W imply $\beta < \pi$, as claimed in (a) and (b) above. Because $\alpha + \beta = 2\pi$, we have: $\alpha \geq \beta$ implies that $\beta \leq \pi$, and $\alpha > \beta$ implies that $\beta < \pi$. Below we abbreviate "the Z-algorithm" and "the W-algorithm" with just Z and W, for readability.

Positive Curvature v. If v has positive curvature, then $\alpha \geq \beta$ implies that $\beta < \pi$ (because if $\beta = \pi$, then $\alpha = \beta = \pi$ and v is flat). So, for positive curvature v, whenever W short-cuts, $\beta < \pi$ and also Z short-cuts. And for a gs-node, if $\alpha < \beta$, then W does not short-cut independent of β, but Z will short-cut if $\beta < \pi$. And for a g-node, if $\alpha \leq \beta$, then W does not short-cut independent of β, but Z will short-cut if $\beta < \pi$.

Flat v. For flat v, if $\alpha = \beta = \pi$, then short-cutting effectively leaves a geodesic through v, and so neither algorithm takes action. For flat v and $\alpha < \beta$, neither algorithm short-cuts: Z because $\beta > \pi$; W because strictly $\alpha\beta$-convex. For flat v and $\alpha > \beta$, both algorithms short-cut: Z because $\beta < \pi$; W because not $\alpha\beta$-convex. We have now established this result:

Lemma 14.11. *The W-algorithm's short-cuts are a subset (or equal) to the Z-algorithm's short-cuts.*

This accords with our conclusion from Lemma 14.9 that $R(W) \subseteq R(Z)$.

The arguments that the W-algorithm's short-cuts are possible, and that the algorithm halts after at most n steps, are identical to those for the Z-algorithm, because every W-algorithm short-cut is a Z-algorithm short-cut, by Lemma 14.11. So we achieve a result parallel to Theorem 14.8:

Theorem 14.12 (W-Algorithm)**.** *The boundary $W(V)$ of* $\mathrm{rconv}(V)$ *can be constructed in polynomial-time, specifically in time $O(n^5 \log n)$ where $n = |V|$.*

Chapter 15

Spiral Tree on Polyhedron

In the previous chapters we have established the planar model for our spiral tree (based on convex hulls), and the extensions to convex polyhedra of the planar notions of convex hull/minimal length enclosing polygon. In this chapter we continue our program and prove that the spiraling idea works as well for vertex-merging in polyhedral half-surfaces bounded by simple closed quasigeodesics, with respect to either convex hull, or enclosing geodesic polygon.

Although in some sense the 3D algorithm follows the 2D algorithm in Chapter 12 closely, there are several significant differences. One is that, in 2D, no triangles are inserted along each slit, whereas in 3D these insertions change P to a different (larger) polyhedron P_i at each step. Another difference is that the 2D slit segments were all geosegs, whereas in 3D some of the slits are geoarcs. These and other differences make the proofs in this chapter somewhat intricate.

15.1 Notation

The notation is a bit complex, so we list the central symbols here for later reference. It may help to refer ahead to Fig. 15.2 to illustrate the definitions below.

- $V = V_0 = \{v_1, v_2, v_3, \ldots, v_n\}$: Vertices on the surface P.
- P_i: Polyhedron after the i-th vertex-merge.

- m_i: merge vertex, the vertex created by merging m_{i-1} with v_{i+1}, with $m_0 = v_1$.
- $g_i = m_{i-1}v_{i+1}$. The i-th slit/merge geoarc.
- v_1 is also given the label m_0, so $g_1 = v_1v_2 = m_0v_2$.
- v_i is called *flattened* if it has already been merged; the merge will reduce v_i's curvature to zero.
- V_i is the set of not-flattened vertices on P_i remaining after the i-th merge along $g_i = m_{i-1}v_{i+1}$. $m_i \in V_i$. $|V_i|$ is the number of vertices in V_i.
- T_i^2 represents the pair of triangles inserted along g_i in a vertex-merge, and T_i refers to the one of the pair crossed by g_{i+1}.
- $H_i = \text{rconv}(V_i)$ is the relative ag-convex hull of V_i, defined in Section 13.7. We view H_i as a region of P_i and let ∂H_i denote its boundary. H_0 is the convex hull of $V = V_0$.
- With some abuse of notation, we identify objects on P_i with their image on P and vice-versa.

  For an object X_i, we put $\tilde{X}_i \subset P$ for the image on P of $X_i \subset P_i$. So, $\tilde{H}_i = H_i \cap P$, and $\tilde{g}_i = g_i \cap P$, for all i.
- Occasionally, we may also use notation from Chapter 11. For example, as in that chapter, we also identify objects on P_{i-1} with their image on P_i and vice-versa.
- $\tilde{\Lambda}_i = \cup_i \tilde{g}_i$ is the slit graph after the i-th merge. $\tilde{\Lambda}$ is the full slit graph.

15.2 Icosahedron Example

Before we describe and prove the spiral algorithm in Section 15.3, we illustrate its application to an icosahedron P. This repeats our discussion in Section 11.4 but following a spiral sequence of vm-reductions. In Chapter 11 the endpoint of the reductions was the doubly-covered triangle shown in Fig. 11.7. Here the endpoint of the reductions will be a half-cylinder, described in the next chapter (Fig. 16.9).

We label the 12 vertices of P as shown in Fig. 15.1. Q is a simple closed geodesic around the "equator" of the icosahedron, and V the six vertices above Q, $v_1, v_2, v_3, v_4, v_5, v_6$. The vertices are merged in that order, as indicated in Fig. 15.1. (This differs from the order in Section 11.4, which merged v_1, v_2, v_3, v_4, v_6, leaving v_5 unmerged.)

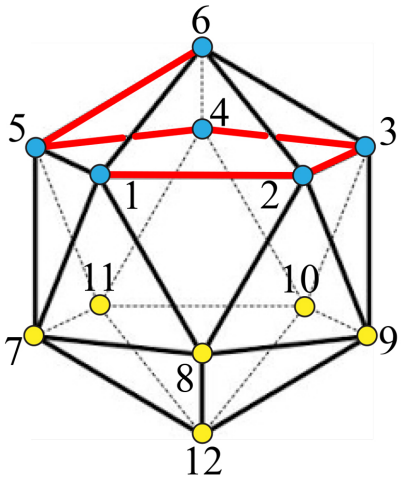

Figure 15.1: Labels i for vertices v_i. Sequential merge path in red. Cf. Fig. 11.5.

Each of the five merges i is accomplished by inserting two copies of a triangle T_i, whose apex is the merge vertex m_i. By convention $m_0 = v_1$. For $i = 1, \ldots, 5$,

$$m_{i-1} \ + \ v_{i+1} \ \to \ m_i \ , \ T_i \text{ along } g_i \ .$$

Here m_5 is special for the icosahedron in that T_5 is actually an infinite parallelogram sending m_5 off to $z = +\infty$, because the sum of the curvatures of $v_1, \ldots, v_6$ is exactly 2π.

Now we describe the steps of the algorithm, referring to Fig. 15.2 throughout.

With $V = V_0$ on $P = P_0$, the initial convex hull containing V_0 is $\text{rconv}(V_0) = H_0 = \tilde{H}_0$, with $\partial H_0 = v_1v_2v_3v_4v_5$. See Fig. 15.2(b).

The first step, i=1, merges two consecutive vertices of ∂H_0, say v_1 and v_2 (as illustrated). So the slit segment is $g_1 = v_1v_2 = m_0v_2$, where $v_1 = m_0$. Doubled triangles T_1^2 are inserted along g_1, which flatten v_1 and v_2, and introduces a new vertex m_1 at the apex of the triangles. The result is a new

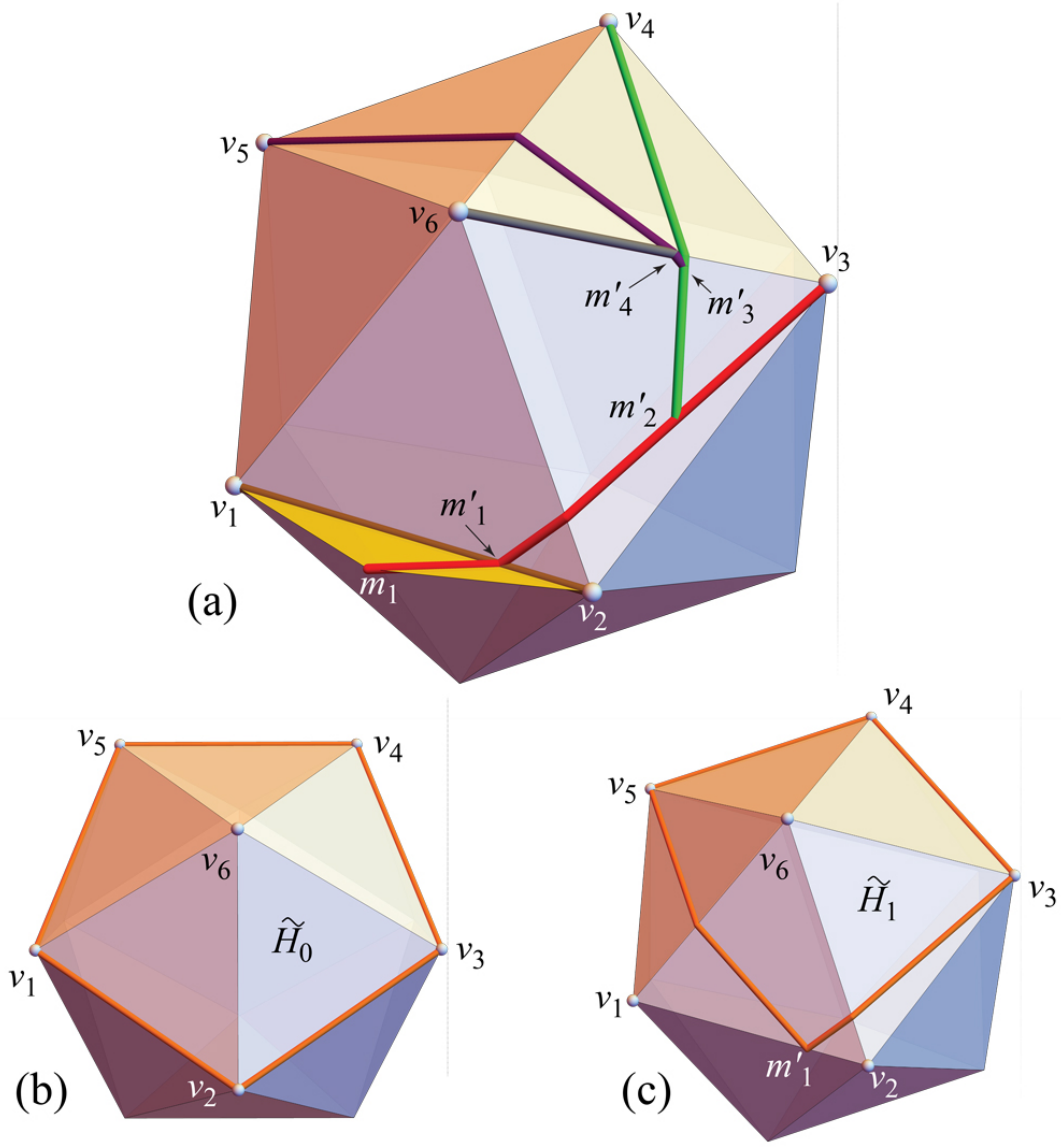

Figure 15.2: (a) Five geodesic slits γ_i on P, each entering P from T_i at m'_i. (b,c) $\tilde{H}_0$ and $\tilde{H}_1$. Cf. Fig. 11.6.

polyhedron P_1. We cannot easily illustrate P_1 (because of the nonconstructive nature of AGT), but it can be imagined from the (yellow) triangle T_1 shown attached to v_1v_2 in Fig. 15.2(a).

On P_1, $V_1 = \{m_1, v_3, v_4, v_5, v_6\}$. $H_1 = \text{rconv}(V_1)$, with $\partial H_1 = m_1v_3v_4v_5$, with m_1v_3 a geodesic from the apex m_1 of T_1, to m_1' where it enters P, and then on P to v_3. $\tilde{H}_1$ has boundary $\partial\tilde{H}_1 = m_1'v_3v_4v_6$, as shown in Fig. 15.2(c). Notice that $\tilde{H}_1 \subset \tilde{H}_0$ on P, but $H_1 \not\subset H_0$ on P_1: $m_1 \notin H_0$.

P_2 is obtained by inserting T_2^2 triangles along the slit g_2 connecting m_1 to v_3, red in Fig. 15.2(a). These triangles flatten m_1 and v_3, leaving $V_2 = \{m_2, v_4, v_5, v_6\}$, and $H_2 = \text{rconv}(V_2)$, with $\partial H_2 = m_2v_4v_5$ and $\partial\tilde{H}_2 = m_2'v_4v_5$.

The process continues, with the final set of slits as depicted in Fig. 15.2(a). Note that, on P, each slit $\tilde{g}_i$ starts from a point m_{i-1}' on $\tilde{g}_{i-1}$ and ends at v_{i+1}.

15.3 Spiraling Algorithm for rconv

First Step. As described above in the icosahedron example, $H_0 = \text{rconv}(V)$ is computed via the algorithm described in Section 14.4. Two vertex endpoints of an edge of ∂H_0, $v_1 = m_0$ and v_2, are merged along $g_1 = m_0v_2$. Triangles T_1^2 are inserted along g_1, flattening v_1 and v_2 and introducing a merge vertex m_1. Now $V_1 = (V \setminus \{v_1, v_2\}) \cup \{m_1\}$ and $H_1 = \text{rconv}(V_1)$.

General Step. We now describe the general step, referring to Fig. 15.3. Suppose step $i-1$ has been completed. $H_{i-1} = \text{rconv}(V_{i-1})$ is on P_{i-1}, with $V_{i-1} = \{m_{i-1}, v_{i+1}, \ldots\}$. The merge vertex m_{i-1} is a vertex of ∂H_{i-1}, and v_{i+1} is the next vertex of that hull boundary, counterclockwise.

For the i-th step, we consider two cases:

Case 1: $m_{i-1} \neq v_{i+1}$. Then $g_i = m_{i-1}v_{i+1}$ is an edge of ∂H_{i-1}. Triangles T_i^2 are inserted along g_i, flattening m_{i-1} and v_{i+1}, and forming P_i. V_i loses these two vertices and gains the new merge vertex m_i. $H_i = \text{rconv}(V_i)$.

Case 2: $m_{i-1} = v_{i+1}$, hence ∂H_{i-1} is a geodesic loop. We have seen that this case can occur. For example, in Fig. 15.4, (a, e, a) is a geodesic loop. (See Example 13.31 for the proof.) In such a case, ∂H_{i-1} consists only of the vertex m_{i-1}, and a geodesic loop at m_{i-1}. Clearly we cannot execute vertex-merging along such a loop, so we need another strategy for this exceptional case.

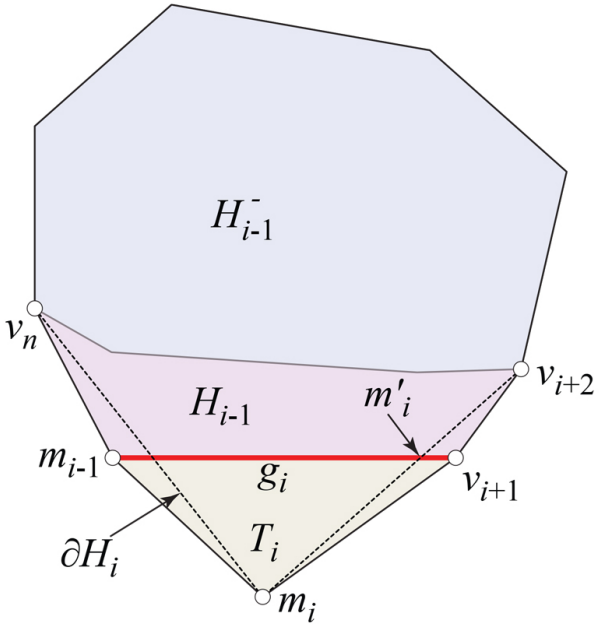

Figure 15.3: H_{i-1} and H_i on P_i.

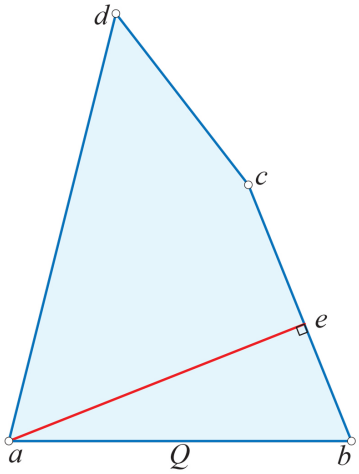

Figure 15.4: Double-sided quadrilateral, $Q = (a, b, a)$. The path (a, e, a) is a geodesic loop. For $V = \{a, c, d\}$, $\partial \operatorname{rconv}(V) = (a, e, a)$.

Let $V'_{i-1} = V_{i-1} \setminus \{m_{i-1}\}$. Construct $H'_{i-1} = \text{rconv}V'_{i-1}$, and let p be the point on $\partial H'_{i-1}$ closest to m_{i-1}, say $p \in \gamma$, where γ is a geoarc of $\partial H'_{i-1}$. Notice that either p is a vertex of V'_{i-1}, or there is a unique geoseg from m_{i-1} to p, which in particular is orthogonal to γ at p.

If p is a vertex of V'_{i-1}, take $g_i = m_{i-1}p$ and notice that g_i intersects neither ∂H_{i-1} nor $\partial H'_{i-1}$, other than at m_{i-1} and at p, respectively, by construction.

Assume now that p is a flat point of H'_{i-1}, hence interior to γ. Then take v_{i+1} to be either endpoint-vertex of γ (or the only vertex on γ if γ is a geodesic loop). Next we show there exists a geoarc $g_i = m_{i-1}v_{i+1}$ intersecting neither ∂H_{i-1}, nor $\partial H'_{i-1}$, other than at m_{i-1} and at v_{i+1}, respectively.

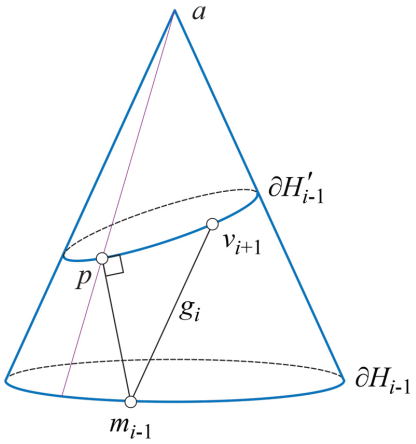

Figure 15.5: On the cone Υ, $g_i = m_{i-1}v_{i+1}$ crosses neither ∂H_{i-1} nor $\partial H'_{i-1}$.

Notice first that V'_{i-1} is strictly interior to H_{i-1} (because ∂H_{i-1} only contains the one vertex m_{i-1}), hence so is $H'_{i-1} = \text{rconv}(V'_{i-1})$ (by Lemma 13.15, the interior of a convex set is convex). Therefore, $\partial H_{i-1} \cap \partial H'_{i-1} = \varnothing$.

Because both ∂H_{i-1} and $\partial H'_{i-1}$ are convex curves, we can merge all vertices of V'_{i-1} to obtain one apex a of a cone Υ. It follows that both ∂H_{i-1} and $\partial H'_{i-1}$ "live on Υ," in the sense that the closed region A of P_{i-i} they bound is isometric to a subset of Υ containing a. (For details see [OV14].) Refer to Fig. 15.5.

Unfold A by cutting it open along the cone generator through ap. Also denote by A the result. Then (the images of) ∂H_{i-1} and $\partial H'_{i-1}$ are convex planar curves in A, and $m_{i-1}p \perp \gamma \subset \partial H'_{i-1}$. Then we see that $g_i = m_{i-1}v_{i+1}$ again intersects neither ∂H_{i-1} nor $\partial H'_{i-1}$ except at its endpoints.

Lemma 15.1. *The new merge vertex m_i created at step i is on ∂H_i (as opposed to strictly interior to ∂H_i).*

Proof. The algorithm constructs H_i by removing the two flattened vertices m_{i-1} and v_{i+1} and adding the new merge vertex m_i. Formally,

$$H_i = \text{rconv}\,(\,(V_{i-1} \setminus \{m_{i-1}, v_{i+1}\}) \cup \{m_i\})\ .$$

Let $H_{i-1}^- = \text{rconv}(V_{i-1} \setminus \{m_{i-1}, v_{i+1}\})$, i.e., the relative convex hull of V_{i-1} with the flattened vertices removed, but the new merge vertex m_i not yet added. See Fig. 15.3. The geoarc $g_i = m_{i-1}v_{i+1}$ is strictly exterior to H_{i-1}^-, because it is included in ∂H_{i-1}. Now attach the doubled T_i^2 triangle to g_i, with triangle apex m_i. It is clear that m_i is also strictly exterior to H_{i-1}^-: some points on a geoseg from m_i to a point in H_{i-1}^- must be exterior to H_{i-1}^-. So $m_i \notin H_{i-1}^-$. Therefore, with $H_i = \text{rconv}(H_{i-1}^- \cup m_i)$, it must be that m_i is a vertex of ∂H_i.

For the exceptional geodesic-loop case of the algorithm, the conclusion follows by construction. □

Notice that H_{i-1} is not necessarily closed, which in some circumstances leads to ∂H_{i-1} being not strictly $\alpha\beta$-convex, but it remains nevertheless $\alpha\beta$-convex. See Lemma 13.6 and Example 13.7 concerning $\alpha\beta$-convexity.

Lemma 15.2 (Visibility)**.** *Let γ be the geoarc connecting the new merge vertex m_i created at step i to the next vertex v_{i+2} to be merged. Then γ crosses the geoarc g_i at a point $\{m_i'\} = g_i \cap \gamma$. In a sense, v_{i+2} is "visible" to m_i through g_i.*

Proof. We use the $\alpha\beta$-convexity property of ∂H_{i-1}: the angles $\alpha_{i+1}, \beta_{i+1}$ at the v_{i+1} endpoint of the geoarc $g_i = m_{i-1}v_{i+1}$ satisfy $\alpha_{i+1} \le \beta_{i+1}$.

We aim to show that, after insertion of the double triangle along g_i, the new angle at v_{i+1}, $\angle m_i v_{i+1} v_{i+2}$, is convex, which proves the claim of the lemma.

So we aim to prove that $\alpha_{i+1} + \frac{1}{2}\omega_{i+1} \le \pi$; the factor $\frac{1}{2}$ appears because that is the angle of T_i's corner at v_{i+1}. To simplify the derivation, we suppress the index $i+1$; so the goal is $\alpha + \frac{1}{2}\omega \le \pi$. Because the triangle insertion flattens v_{i+1}, we know that $\omega = 2\pi - (\alpha+\beta)$. So we get equivalent inequalities:

$$\begin{aligned} \alpha + \tfrac{1}{2}\omega &\le \pi \\ \alpha + \pi - \tfrac{1}{2}\alpha - \tfrac{1}{2}\beta &\le \pi \\ \tfrac{1}{2}\alpha &\le \tfrac{1}{2}\beta \end{aligned}$$

which holds at v_{i+1} by $\alpha\beta$-convexity.

The same argument applies to the other endpoint m_{i-1} of g_i. □

This lemma verifies that the drawing in Fig. 15.3 is a correct depiction. If this lemma did not hold, then $\gamma = m_i v_{i+2}$ would not necessarily cross g_i.

15.4 Proof: Slit Graph is a Tree

The next lemma is the counterpart to the nesting property of the 2D spiral algorithm, previously illustrated in Fig. 12.4.

Lemma 15.3 (Nesting)**.** $\tilde{H}_i \subset \tilde{H}_{i-1}$.

It may be useful to keep in mind the situation on P_i in Fig. 15.3, and the $\tilde{H}_i$ examples on P in Fig. 15.2(bc), which shows $\tilde{H}_1 \subset \tilde{H}_0$. With some abuse of notation in the proof, we will identify objects on P_{i-1} with their image on P_i and vice-versa.

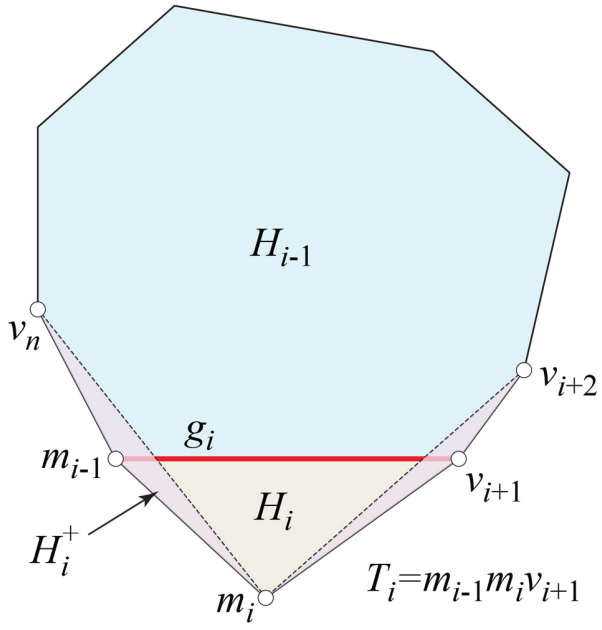

Figure 15.6: $H_i \subset H_i^+$.

Proof. The proof is in two parts. We show first that $\tilde{H}_i \subset \tilde{H}_{i-1}$. Set $V_{i-1}^- = V_{i-1} \setminus \{m_{i-1}, v_{i+1}\}$, i.e., V_{i-1} without the two vertices that will form the next vertex-merge insertion. To keep track of the notation, we refer to Fig. 15.6 (a variation of Fig. 15.3) and this display, where we use "$-$" to indicate "missing" vertices among the three m_{i-1}, m_i, v_{i+1}:

$$\begin{aligned} & V_{i-1}^- = \{v_{i+2}, \ldots, v_n\} \\ H_{i-1} = \mathrm{rconv}(V_{i-1}) \ , \ & V_{i-1} = \{m_{i-1}, -, v_{i+1}\} \cup V_{i-1}^- \\ H_i = \mathrm{rconv}(V_i) \ , \ & V_i = \{-, m_i, -, \} \cup V_{i-1}^- \\ H_i^+ = \mathrm{rconv}(V_i^+) \ , \ & V_i = \{m_{i-1}, m_i, v_{i+1}\} \cup V_{i-1}^- \end{aligned}$$

Set H_i^+ as above. Then as Fig. 15.6 shows, $H_i \subset H_i^+$ and thus $\tilde{H}_i \subset \tilde{H}_i^+$.

Denote by $T_i = \triangle m_{i-1} m_i v_{i+1}$ one of the two inserted triangles along g_i, the one sharing g_i in common with H_{i-1}.

For the second part of the proof, we next show that, on P_i, $H_i^+ = H_{i-1} \cup T_i$. First, notice that $H_{i-1} \cup T_i \subset H_i^+$ (by Lemma 13.24 on the hull of a set partition), so $\mathrm{rconv}(H_{i-1} \cup T_i) \subset H_i^+$. Second, the $\alpha\beta$-convexity of ∂H_{i-1} at m_{i-1} and v_{i+1} on P_{i-1} shows that, on P_i, the two sides of T_i incident to m_i either extend two boundary arcs of H_{i-1} (if the $\alpha\beta$-convexity is not strict), or they make with the respective boundary arcs angles $< \pi$ towards $H_{i-1} \cup T_i$ (if the $\alpha\beta$-convexity is strict); see Lemma 15.2 (Visiblity). Therefore, $H_{i-1} \cup T_i$ is a convex set (by Lemma 13.38) containing $V_{i-1}^- \cup \{m_{i-1}, m_i, v_{i+1}\}$, hence $H_{i-1} \cup T_i \supset \mathrm{rconv}(V_{i-1}^- \cup \{m_{i-1}, m_i, v_{i+1}\}) = H_i^+$.

Finally, since $H_i^+ = H_{i-1} \cup T_i$ and the image on P of T_i is a subarc of $\tilde{g}_i$, we have $\tilde{H}_i^+ = \tilde{H}_{i-1}$.

In conclusion, $\tilde{H}_i \subset \tilde{H}_i^+ = \tilde{H}_{i-1}$. □

Notice that, unlike the planar case where g_i only intersects H_i in one point, $g_i \cap H_i = \{m_i\}$ (Fig. 12.2(b)), on polyhedra g_i intersects H_i on a sub-segment (Fig. 15.3).

Recall that, in our notation, g_{i-1} is a curve (a geoarc) on P_{i-2} and g_i is a curve on P_{i-1}.

Lemma 15.4 ($g_i \cap g_{i-1}$)**.** *On P_{i-1}, g_i intersects the two images ("banks") of g_{i-1} at precisely one point. Consequently, $\tilde{g}_i$ has an endpoint on $\tilde{g}_{i-1}$ and no other common point.*

Proof. On P_{i-1}, g_i joins the point m_{i-1} to a vertex v_{i+1}. But m_{i-1} is inside the inserted region (a pair of triangles) T_{i-1}, while v_{i+1} is outside T_{i-1}. Remember that the boundary of T_{i-1} consists of two copies of g_{i-1}, say g'_{i-1} and g''_{i-1}. So g_i intersects those two copies, $g'_{i-1} \cup g''_{i-1}$, at least once.[1]

Assume, for the sake of contradiction, that there are two distinct points

$$x', x'' \in g_i \cap (g'_{i-1} \cup g''_{i-1}) \ .$$

Assume first that $\{x'\} = g_i \cap g'_{i-1}$ and $\{x''\} = g_i \cap g''_{i-1}$. It follows that g_i exits T_{i-1} and enters it again, and so to reach v_{i+1} it has to exit T_{i-1} again, crossing $g'_{i-1} \cup g''_{i-1}$ once more. Therefore, we may assume next that both x' and x'' arise from g_i intersecting one g_{i-1} image, say g'_{i-1}. Hence g_i and g'_{i-1} determine a geodesic digon, which necessarily contains a vertex $v \in V$ (by the Gauss-Bonnet Theorem). Since there is no vertex both outside H_{i-1} and inside H_{i-2}, and there is no vertex both outside H_{i-2} and inside H_{i-1} other than m_{i-1}, a contradiction is obtained.

Therefore, g_i intersects the two images of g_{i-1} at precisely one point, and thus we get that $|\tilde{g}_{i-1} \cap \tilde{g}_i| = 1$. □

Lemma 15.5. *The set of vertices reduces by one each iteration:* $|V_i| = |V_{i-1}| - 1$*, for* $i > 1$*.*

Proof. Clearly $|V_i| = |V_{i-1}| - 1$, because $V_i = V_{i-1} \setminus \{m_{i-1}, v_{i+1}\} \cup \{m_i\}$. □

Lemma 15.6. $\tilde{\Lambda}_i$ *is a tree.*

Proof. By Lemma 15.4 above, $\tilde{g}_i$ has an endpoint on $\tilde{g}_{i-1}$, and otherwise does not intersect $\tilde{g}_{i-1}$. It remains to prove that the interaction of $\tilde{g}_i$ with the earlier segments in $\tilde{\Lambda}_i$ maintains the tree property. In fact, we show that $\tilde{g}_i \cap \tilde{g}_j = \varnothing$ for $j < i - 1$.

By our sequential-merge choice, g_i is an edge $e = m_{i-1}v_{i+1}$ of ∂H_{i-1}, and so $\tilde{g}_i$ is an edge $\tilde{e}$ of $\partial \tilde{H}_{i-1}$. For example, in Fig. 15.2(bc), $\tilde{g}_1 = v_1v_2$ is an edge of $\tilde{H}_0$, and $\tilde{g}_2 = m'_1v_3$ is an edge of $\tilde{H}_1$.

From Lemma 15.3 (Nesting) we know that $\tilde{H}_j \supset \tilde{H}_{i-1}$ for $j < i - 1$.

Therefore, $\tilde{g}_j \cap \tilde{H}_{i-1}$ is either the empty set, or an edge of $\tilde{H}_{i-1}$, so it cannot intersect other edges of $\tilde{H}_{i-1}$ excepting those adjacent to it (by Theorem 13.40, which established the simplicity of rconv boundaries).

[1] With yet another abuse of notation, we use the same symbol g for an arc $g : I \to P$ (with $I \subset \mathbb{R}$ an interval) and its geometrical image on P.

Since the edge $\tilde{g}_{i-1} \cap \tilde{H}_{i-1}$ is between $\tilde{g}_j \cap \tilde{H}_{i-1}$ (not empty by Lemma 11.1) and $\tilde{g}_i$ ("between" with respect to the circular order of sides in $\partial \tilde{H}_{i-1}$), $\tilde{g}_j$ cannot intersect $\tilde{g}_i$ for $j < i-1$.

For example, returning to Fig. 15.2, when $i = 3$ and $j = 1$, we have $\tilde{H}_{i-1} = \tilde{H}_2$, $\tilde{g}_i = \tilde{g}_3 = m_2' v_4$, $\tilde{g}_j = \tilde{g}_1 = v_1 v_2$, and indeed $\tilde{g}_1$ does not intersect $\tilde{g}_3$.

Therefore, we have shown that $\tilde{g}_i$ just intersects $\tilde{\Lambda}_{i-1}$ at the one point, and so maintains the tree structure for $\tilde{\Lambda}_i$. □

Putting together the above lemmas yields the following theorem, one of our primary goals.

Theorem 15.7. *Let Q be a simple closed quasigeodesic on the convex polyhedron P, and V the set of vertices of P enclosed by Q, to either side of Q. Then the sequential vertex-merging algorithm detailed in Section 12.3 results in a slit graph $\tilde{\Lambda}$ that is a tree.*

The consequence of this theorem is that the slits do not disconnect a half -surface of P, i.e., it remains simply connected. This will be expanded upon in the following chapter.

15.5 Spiraling Algorithm for $Z(V) = \min \ell[V]$

The spiraling approach, developed in Section 15.3 into Algorithm 1 based on the sets $H_i = \text{rconv}(V_i)$, can be also based on the sets $M_i = R(Z(V_i))$. And, because it can be that $W(V) \neq Z(V)$ (Example 14.10), in these cases we get a different but analogous Algorithm 2.

The general description of Algorithm 2 is precisely the same as for Algorithm 1, as it only uses the geodesic polygon ∂H_i which can be replaced by Z. So we will retain the same notation, but always remember to replace H_i by M_i.

For the icosahedron example in Fig. 15.2, Section 11.4, $\tilde{H}_i = \tilde{M}_i$ for all i, and the resulting slit tree is the same.

Recall that Fig. 14.10 is an example where $Z \neq W$. Using the Z-algorithm, the first merge would be along the aba' geoarc, whereas the W-algorithm would start with, say, the ac geoseg. So the two vm-reductions would not be identical, and would result in different slit trees.

Lemma 15.8. *The new merge vertex m_i created at step i is on Z_i (as opposed to strictly interior to Z_i).*

Proof. Similar to that for Lemma 15.1. The algorithm constructs M_i by removing the two flattened vertices m_{i-1} and v_{i+1} and adding the new merge vertex m_i. Formally,

$$M_i = R\left(Z\left(\left(V_{i-1} \setminus \{m_{i-1}, v_{i+1}\}\right) \cup \{m_i\}\right)\right) \ .$$

Let $M_{i-1}^- = R(Z(V_{i-1} \setminus \{m_{i-1}, v_{i+1}\}))$; see Fig. 15.3 with $N \to M$. The geoarc $g_i = m_{i-1}v_{i+1}$ is strictly exterior to M_{i-1}^-, because it is included in Z_{i-1}. Now attach the doubled T_i^2 triangle to g_i, with triangle apex m_i. It is clear that m_i is also strictly exterior to M_{i-1}^-. Therefore, with $M_i = R(Z(M_{i-1}^- \cup m_i))$, it must be that m_i is a vertex of Z_i.

For the exceptional geodesic-loop case of the algorithm, the conclusion follows by construction. □

Lemma 15.9 (Visibility)**.** *Let γ be the geoarc connecting the new merge vertex m_i created at step i to the next vertex v_{i+2} to be merged. Then γ crosses the geoarc g_i at a point $\{m_i'\} = g_i \cap \gamma$. In a sense, v_{i+2} is "visible" to m_i through g_i.*

Proof. As in the proof of Lemma 15.2, we only make use of the $\alpha\beta$-convexity property of Z_{i-1}, which in this case is strict. □

Lemma 15.10 (Nesting)**.** $\tilde{M}_i \subset \tilde{M}_{i-1}$.

With some abuse of notation in the proof, we will identify objects on P_{i-1} with their image on P_i and vice-versa.

Proof. Similar to that for Lemma 15.3.

The proof is in two parts. We show first that $\tilde{M}_i \subset \tilde{M}_{i-1}$. Set $V_{i-1}^- = V_{i-1} \setminus \{m_{i-1}, v_{i+1}\}$, i.e., V_{i-1} without the two vertices that will form the next vertex-merge insertion. To keep track of the notation, we refer to Fig. 15.6 (a variation of Fig. 15.3) with $N \to M$, and this display, where again we use "$-$" to indicate "missing" vertices among the three m_{i-1}, m_i, v_{i+1}:

$$\begin{aligned} V_{i-1}^- &= \{v_{i+2}, \ldots, v_n\} \\ M_{i-1} = R(Z_{i-1}) \ , \ V_{i-1} &= \{m_{i-1}, -, v_{i+1}\} \cup V_{i-1}^- \\ M_i = R(Z_i) \ , \ V_i &= \{-, m_i, -, \} \cup V_{i-1}^- \\ M_i^+ = R(Z(V_i^+)) \ , \ V_i^+ &= \{m_{i-1}, m_i, v_{i+1}\} \cup V_{i-1}^- \end{aligned}$$

Set M_i^+ as above. Then as Fig. 15.6 shows, $M_i \subset M_i^+$ (by Lemma 14.6 applied twice) and thus $\tilde{M}_i \subset \tilde{M}_i^+$.

Denote by $T_i = \triangle m_{i-1} m_i v_{i+1}$ one of the two inserted triangles along g_i, the one sharing g_i in common with M_{i-1}.

For the second part of the proof, we next show that, on P_i, $M_i^+ = M_{i-1} \cup T_i$.

First, notice that $M_{i-1} \cup T_i \subset M_i^+$. Second, the strict $\alpha\beta$-convexity of $Z(V_{i-1})$ at m_{i-1} and v_{i+1} on P_{i-1} shows that, on P_i, the two sides of T_i incident to m_i make with the respective boundary arcs angles $< \pi$ towards $M_{i-1} \cup T_i$, hence $> \pi$ on the other side; see Lemma 15.9 (Visiblity).

Therefore, the boundary of $M_{i-1} \cup T_i$ is precisely the minimal length enclosing polygon of V_i^+ (by Lemma 14.4).

Finally, since $M_i^+ = M_{i-1} \cup T_i$ and the image on P of T_i is a subarc of $\tilde{g}_i$, we have $\tilde{M}_i^+ = \tilde{M}_{i-1}$. In conclusion, $\tilde{M}_i \subset \tilde{M}_i^+ = \tilde{M}_{i-1}$. □

The proofs for the next four results are identical to the proofs or their counterparts in Section 15.4 (Lemma 15.4 and the following results).

Lemma 15.11 ($g_i \cap g_{i-1}$)**.** *On P_{i-1}, g_i intersects the two images ("banks") of g_{i-1} at precisely one point. Consequently, $\tilde{g}_i$ has an endpoint on $\tilde{g}_{i-1}$ and no other common point.*

Lemma 15.12. *The set of vertices reduces by one each iteration: $|V_i| = |V_{i-1}| - 1$, for $i > 1$.*

Lemma 15.13. *$\tilde{\Lambda}_i$ is a tree.*

Putting together the above lemmas, we obtain a counterpart to Theorem 15.7.

Theorem 15.14. *Let Q be a simple closed quasigeodesic on the convex polyhedron P, and V the set of vertices of P enclosed by Q, to either side of Q. Then the sequential vertex-merging algorithm, based on $Z = \min \ell[V]$ and detailed in Section 12.3, results in a slit graph $\tilde{\Lambda}$ that is a tree.*

Application of these slit trees to unfoldings of P is discussed in the next chapter.

Chapter 16

Unfoldings via Slit Trees

In this chapter we gather together results from Chapter 11 concerning the connectivity structure of slit graphs, and the spiral slit trees just obtained in Chapter 15, to draw conclusions about unfolding convex polyhedra P via vertex-merging.

Depending on circumstances we will detail, vertex-mergings result in P embedded in a doubly-covered triangle (or isosceles tetrahedron), in a pair of joined cones, or in a cylinder. The slits may leave P in one piece, or cut into several pieces. In some situations, we can develop P to the plane as a net.

16.1 Notation

We first recall previous notation to be used in this chapter, and introduce some new notation.

- Q is a simple closed quasigeodesic on P.
- P^+ and P^- are the closed half-surfaces bounded by Q. Because closed, both of these half-surfaces include Q.
- $V(Q) = \{q_1, \ldots, q_k\}$ is set of vertices on Q.
- ε is either of $+, -$.
- We refer to the Z-algorithm (minimum enclosing polygon) and the W-algorithm (∂ rconv), summarized in Theorems 14.8 and 14.12 respectively.

- $\tilde{\Lambda}^\varepsilon$ is the full slit tree on P^ε, obtained from a set V of vertices containing all vertices interior to P^ε, and possibly some vertices in $V(Q)$, via either the Z- or W-spiral algorithm.
- α_i and β_i are the angles incident to q_i in P^+ and in P^- respectively, i.e., above and below Q.
- The curvature of P at q_i is $\omega_i = 2\pi - (\alpha_i + \beta_i)$.
- The *partial curvatures* at q_i toward P^+ and P^- are $\omega_i^+ = \pi - \alpha_i$ and $\omega_i^- = \pi - \beta_i$, respectively. Hence $\omega_i = \omega_i^+ + \omega_i^- < 2\pi$.
- Ω^ε is the total curvature of all interior vertices of P^ε.

The Gauss-Bonnet Theorem gives, for $\varepsilon \in \{+, -\}$,

$$\Omega^\varepsilon + \sum_{i=1}^{k} \omega_i^\varepsilon = 2\pi \ , \tag{16.1}$$

$$\Omega^+ + \Omega^- + \sum_{i=1}^{k} \omega_i = 4\pi \ . \tag{16.2}$$

16.2 Unfoldings via spiraling algorithms

Assume in the following that V contains *all* vertices of P enclosed by Q, to either side, including those on Q: $V \supset V(Q)$.

Next we see that the spiraling algorithm works fine in this case, too.

The total curvature of V may be $> 2\pi$ on P, but on $P^\#$ it is precisely 2π, because all angles β_i are π on $P^\#$.

Choose two consecutive vertices on Q, q_1, q_2. If there are no such vertices, then Q is either a closed geodesic or a geodesic loop, handled later.

Merge q_1 and q_2 along the geoarc g_{12} of Q joining them (either arc, if there are two), to produce $m_1 \in V_1 \subset P_1$. This merging inserts a double triangle of base g_{12} and base angles $\omega_i^+/2$ at v_i. Call this a *partial merging*.

By Lemmas 15.1 and 15.8, m_1 is interior to P_1. Further merging m_1 to another vertex $q_3 \in Q$ would force $\tilde{\Lambda}^+$ to have a leaf on Q. Therefore, if $k = |V(Q)| > 2$ then $\tilde{\Lambda}^+$ intersects Q at g_{12} and at each q_j with $\alpha_j < \pi$, for $j = 3, \ldots, k$.

For example, return to Fig. 15.2(a), repeated below as Fig. 16.6. Suppose that $Q = v_1v_2v_3v_4v_5$ is the pentagon of edges surrounding v_6. Then $\tilde{\Lambda}^+$ has leaves at v_3, v_4, v_5.

Similar considerations lead to the the conclusion that $\tilde{\Lambda}^-$ also intersects Q at g_{12} (if it starts by merging q_1 and q_2) and at each q_j with $\beta_j < \pi$, for $j = 3, \ldots, k$.

Continuing with the example in Fig. 16.6, if $\tilde{\Lambda}^-$ also had leaves at, say, v_3 and v_4, then a piece of P, bounded on P^+ by $v_3m_2'v_4$, and similarly bounded on P^-, would be disconnected from the remainder of P.

The above discussion of the way the trees $\tilde{\Lambda}^\varepsilon$ connect to one another along Q leads to several main results.

We see two main options: merge the vertices strictly inside P^ε, or, in addition, include all vertices on Q via partial merging.

16.2.1 Two Cones

We start with merging all vertices inside Q.

Lemma 16.1. *Let P be a convex polyhedron and Q a simple closed quasi-geodesic on P. Merging all vertices strictly inside Q to either side, along the slit trees $\tilde{\Lambda}^\varepsilon$ ($\varepsilon = +, -$), unfolds P onto the union $\mathcal{U}$ of two cones, each of base Q, glued together along Q. The unfolding of P onto $\mathcal{U}$ can be decomposed into two simple geodesic polygons, one to each side of Q, sharing Q.*

Proof. Merging all vertices inside P^ε reduces all those vertices to one merge vertex, the apex of a cone. Because Q is unaffected by the merges, it remains shared by the cones. Theorem 11.5 (that $\tilde{\Lambda}^\varepsilon$ a tree implies that P^ε is a simple polygonal domain) guarantees the P unfolding is a simple geodesic polygon to each side of Q. □

Notice that the image of Q on $\mathcal{U}$ above is not necessarily planar.

A special case of this lemma is that, if there is just one vertex q_1 on Q, so Q is a geodesic loop, then $\mathcal{U}$ has three vertices—two cone apexes and the unmerged $q_1 \in Q$—and so is a doubly-covered triangle.

Another special case of this lemma is if Q is a simple closed geodesic; i.e., there is no vertex on Q. We defer this case to Theorem 16.6 below.

Example: Spiral Q on Box. We illustrate Lemma 16.1 with an example based on Fig. 2 in [DHK21], reproduced in Fig. 16.1. Its interesting feature is that the illustrated simple closed quasigeodesic spirals around the box, arbitrarily many times as L grows large. For simplicity, in our version,

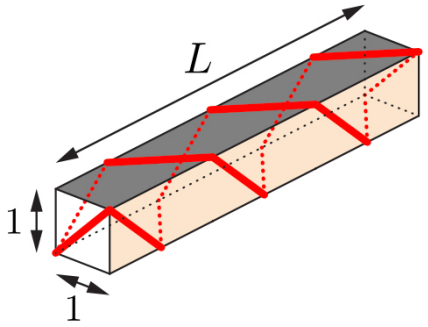

Figure 16.1: Detail from Fig. 2 in [DHK21].

Fig. 16.2(a), the quasigeodesic Q spirals just one turn, but the example could be extended to many turns.

Our quasigeodesic $Q = v_1v_3v_5v_7$ passes through four vertices, and encloses two vertices, v_4 and v_6, strictly to its left. Fig. 16.2(b) shows an unfolding of P^+, the half-surface left of Q. P^- (not shown) is symmetrical to the right of Q.

Merging the two vertices v_4 and v_6 produces a cone apexed at the merge vertex m_{46}. For display purposes, it is flattened to the plane as a doubly-covered triangle in (c) of the figure. One can see that the portion of P to the left of Q is indeed a simple geodesic polygon, as claimed by Lemma 16.1.

Merging v_2 and v_8 on P^- produces a second cone apexed at m_{28}. The two cones are glued together along Q, which still contains four vertices. So $\mathcal{U}$ is a polyhedron of 6 vertices, an octahedron.

It is tempting to develop each cone to the plane to achieve nets of P^ε. This works in the example just presented, but not always, as we now detail.

Let $\mathcal{U}^+$ be the upper cone containing P^+. P^+ is topologically an annulus, with $\partial P^+ = Q \cup Q'$. Parametrize points $p(t)$ on the upper boundary Q', $t \in [0, 1]$, with $p(0) = p(1)$. Define $\phi(t)$ to be the angular turn of the segment from $ap(0)$ to $ap(t)$ about the apex a on the surface of $\mathcal{U}^+$. Then with $\phi(0) = 0$, we have that $\phi(1) = \alpha$, for the segment makes one full turn

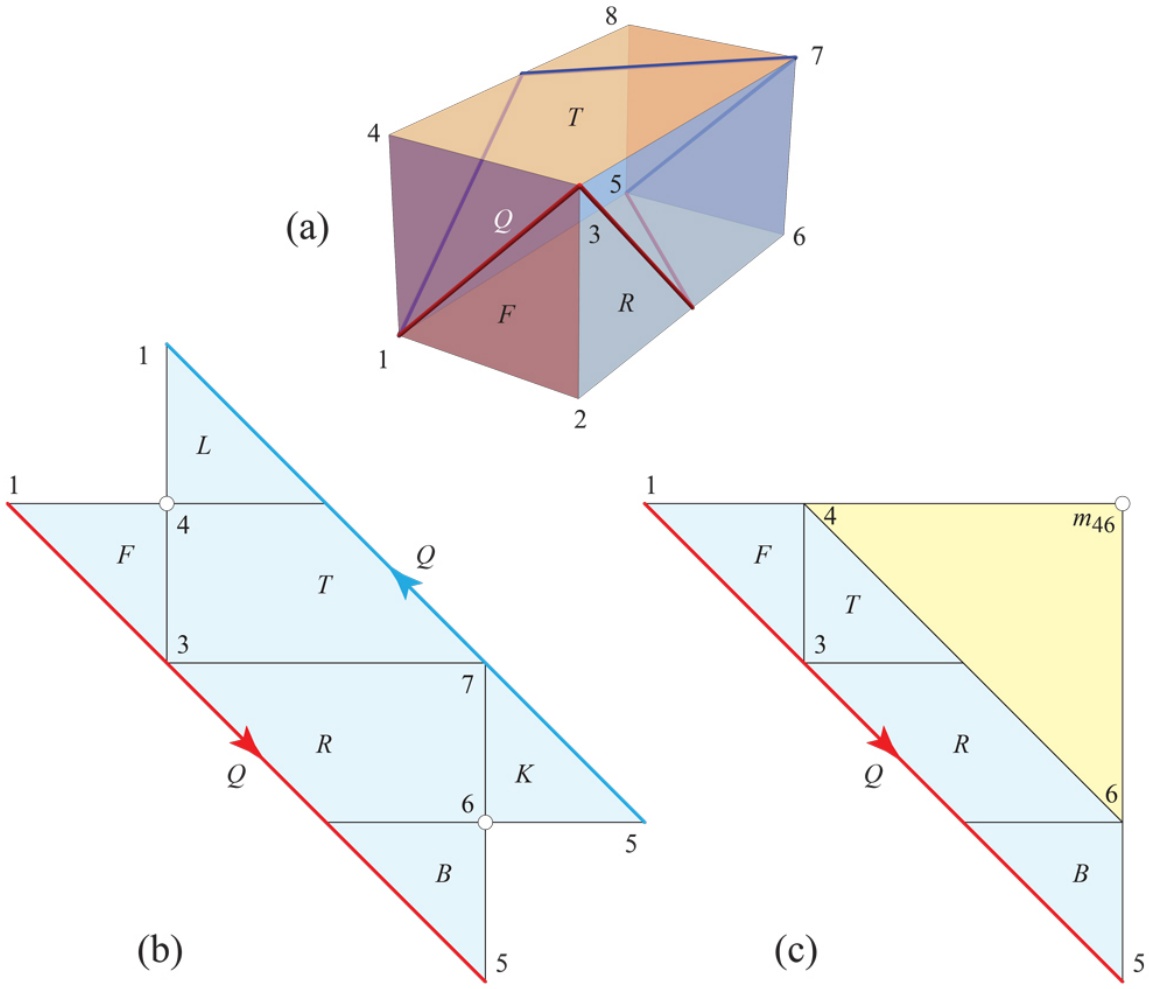

Figure 16.2: (a) A $1 \times 1 \times 2$ box. (b) The surface P^+ to the left of Q. Note the two images of v_1 are identified, as are the images of v_5. $\omega(v_4) = \omega(v_6) = \pi/2$. (c) After merging v_4 and v_6, a cone, here flattened to a doubly-covered triangle.

to return to $p(0)$. However, it could be that if Q' spirals around the cone, intermediate positions represent turns greater than α. See Fig. 16.3.

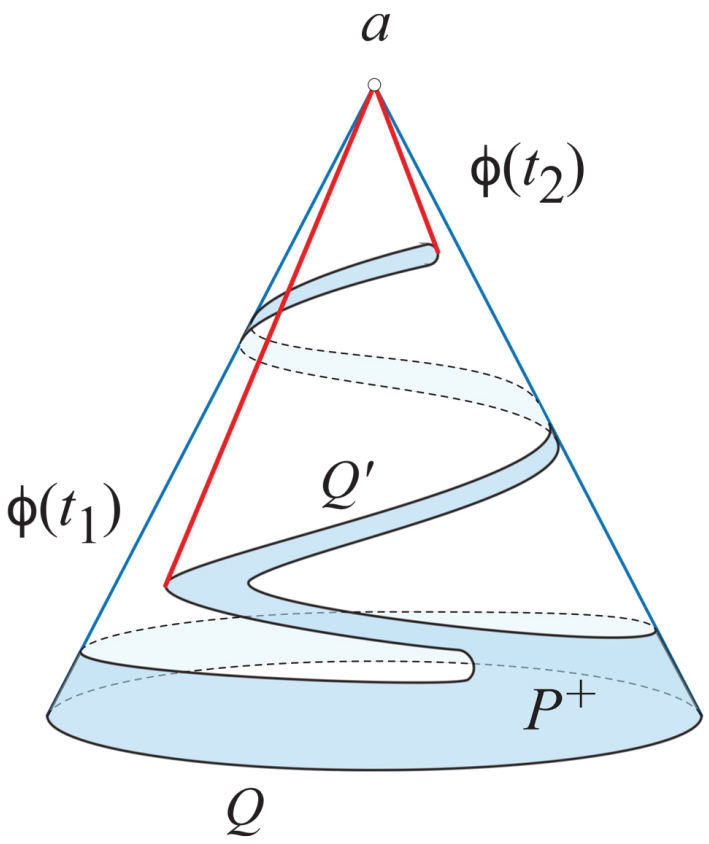

Figure 16.3: $\phi(t_2) - \phi(t_1) > \alpha$.

Let $\phi_{\max} = \max_{t_1<t_2} |\phi(t_2) - \phi(t_1)|$. Then $k = \lceil \phi_{\max}/\alpha \rceil$ represents the number of full turns around the cone that might be needed to fully develop P^+ to the plane. No overlap in the development can occur unless $k\alpha > 2\pi$, so that the imprint of the cut-open cone cycles around a more than once. Therefore we have this result.

Proposition 16.2. *If $k\alpha \le 2\pi$, where $k = \lceil \phi_{\max}/\alpha \rceil$, then P^+ develops to the plane without overlap.*

In the absence of the limit detailed in this proposition, it is indeed possible for the development of P^+ to overlap. This was established with a nontrivial

example in [OV14, Fig. 14] by a curve that spirals around a four times. Building this curve into P^+ would establish an overlapping development of P^+.

Another very special case occurs when α evenly divides 2π. For then, even if the development spirals around a more than 2π, the imprint of the cut-open cone lays exactly on top of its earlier-rolled image. So any overlap in the development would have been mirrored by overlap on the cone. But P^+ does not overlap on $\mathcal{U}^+$.

Proposition 16.3. *If α evenly divides 2π, then P^+ develops without overlap.*

16.2.2 Reduction to Cylinder

We return now to the second option: including all vertices on Q via partial merging.

Lemma 16.4. *With P and Q as before, let there be $k = |V(Q)|$ vertices on Q. Merging all vertices of P, along the trees $\tilde{\Lambda}^\varepsilon$, including the vertices V via partial merges as described above, unfolds P onto a cylinder $\mathcal{C}$. The unfolding of P onto $\mathcal{C}$ is the union of at most $k-2$ simple geodesic polygons, joined circularly at $k-2$ vertices of Q.*

Proof. The reason the final surface is a cylinder is that (a) all vertices are merged; none remain, and (b) the curvature to each side is exactly 2π: Eq. 16.1. Although Theorem 11.5 again guarantees the P unfolding to each side of Q is a simple geodesic polygon, as we've seen, the coinciding leaves of $\tilde{\Lambda}^\varepsilon$ partitions P into at most $k-2$ pieces. □

In contrast to the two-cones case, rolling the cylinder $\mathcal{C}$ on the plane cannot cause overlap.

We've already mentioned that $k = |V(Q)| = 1$ is a special case of Lemma 16.1, and $k \le 2$ plays a special role in Lemma 16.4. We highlight the latter in the following theorem.

Theorem 16.5. *Let P be a convex polyhedron and Q be a simple closed quasigeodesic on P containing $k = |V(Q)|$ vertices. If either*

(1) $k \le 2$, or

(2) if the total number of non-π angles α_i and β_i (i.e., $< \pi$ angle to both sides) is at most two,

then the unfolding of P onto $\mathcal{C}$ (Lemma 16.4) is a simple geodesic polygon. In this case, rolling $\mathcal{C}$ onto a plane develops P to a net.

Proof. Lemma 16.4 established the specialness of $k \leq 2$: then the unfolding of P onto $\mathcal{C}$ is a single, simple geodesic polygon.

When either of the angles α_i or β_i incident to q_i is equal to π, there is no need to partially merge to q_i from one side or the other. The two slit trees can have a common leaf only at a vertex q_i having to both sides angles strictly less than π. This reduces the number of pieces from $k - 2$. □

Returning to Fig. 16.6 with $Q = v_1v_2v_3v_4v_5$, $\alpha_i = \frac{2}{3}\pi$ and $\beta_i = \pi$, so there are no angles $< \pi$ to both sides, and Theorem 16.5 applies. We will indeed see below (Fig. 16.9) this leads to a net for the icosahedron.

We believe that every convex poyhedron has a a simple closed quasigeodesic containing at most one vertex, see Open Problem 18.13. If this would be true, the above result would provide vertex-merging nets for all convex polyhedra. We discuss this question of the number of vertices on a simple closed quasigeodesic in Chapter 17.

The case $k = 0$ (i.e., Q is a simple closed geodesic) is considered next. In this case, P contains a region R isometric to a cylinder; assume R is maximal with respect to inclusion. Then the two boundary components of R are simple closed quasigeodesics, denoted by Q^ε. We then denote by P^ε the caps of P bounded by Q^ε outside R, and apply the previous considerations. Precisely, we unfold both P^ε onto the same cylinder, which in this case contains R, and further roll the cylinder to obtain a net of P.

The case $k = 0$ is also special in that we can obtain an unfolding of P onto an isosceles tetrahedron (as opposed to the doubly covered triangles obtained in Lemma 16.1 for $k = 1$). In this case, Q is a simple closed geodesic. Then the merging processes end with two vertices of total curvature 2π, to each side of Q. If both of those vertices (to the same side of Q) have curvature π, we have reached half of an isosceles tetrahedron. Otherwise, suitably choosing a partial merge between the two vertices results in one of them having curvature π, and a new merge vertex also of curvature π. Globally, we obtain $\mathcal{U}$ an isosceles tetrahedron, the special case of vm-irreducible surfaces. This discussion is consistent with the above argument: once we have an

isosceles tetrahedron, we can obtain the cylinder by cutting open a pair of opposite edges.

Theorem 16.6. *If P contains a simple closed geodesic then, unfolding it as described above provides a net.*

We next provide two examples illustrating this theorem.

16.2.3 Cube Example

To illustrate Theorem 16.6, we revisit the cube example in Section 11.3 (Fig. 11.4). There we sequentially-merged three vertices on the top face, v_7, v_8 and then v_5 with v_{78}, and symmetrically three on the bottom face.

Instead now we merge all four vertices on the top face, and then on the bottom face. So, for a quasigeodesic Q around the middle of the cube, we are merging all vertices inside P^+, and all vertices inside P^-, leading to a cylinder.

We perform the same two initial merges on the top vertices, resulting again in v_{578}. Let $a = m'_1$ and $b = m'_2$, the two points on the top face where the geodesic segments from the merge vertices $m_1 = v_{12}$ and $m_2 = v_{578}$ enter the top face of the cube.

Note that the angle incident to the merge vertex v_{578} is 90°, and the angle incident to the as-yet unmerged top vertex v_6 is 270°. So a merge of v_{578} with v_6 would not produce a triangle pair, because the sum of their curvatures is 2π. However, we can imagine a merge resulting in a pair of unbounded parallelograms. If we cut the surface along the geodesic segment v_6v_{578} (of length $2\sqrt{2}$ for a unit cube) and insert twin parallelograms, the result is a cylinder unbounded above. See Fig. 16.4. Note the two 45° angles inserted at v_6 flattens that vertex, and the insertion of the two 135° angles flattens v_{578}. Symmetric merges on the bottom-face vertices leads to an unbounded cylinder in both directions.

Unrolling the cylinder unfolds the cube to a non-overlapping net: Fig. 16.5.

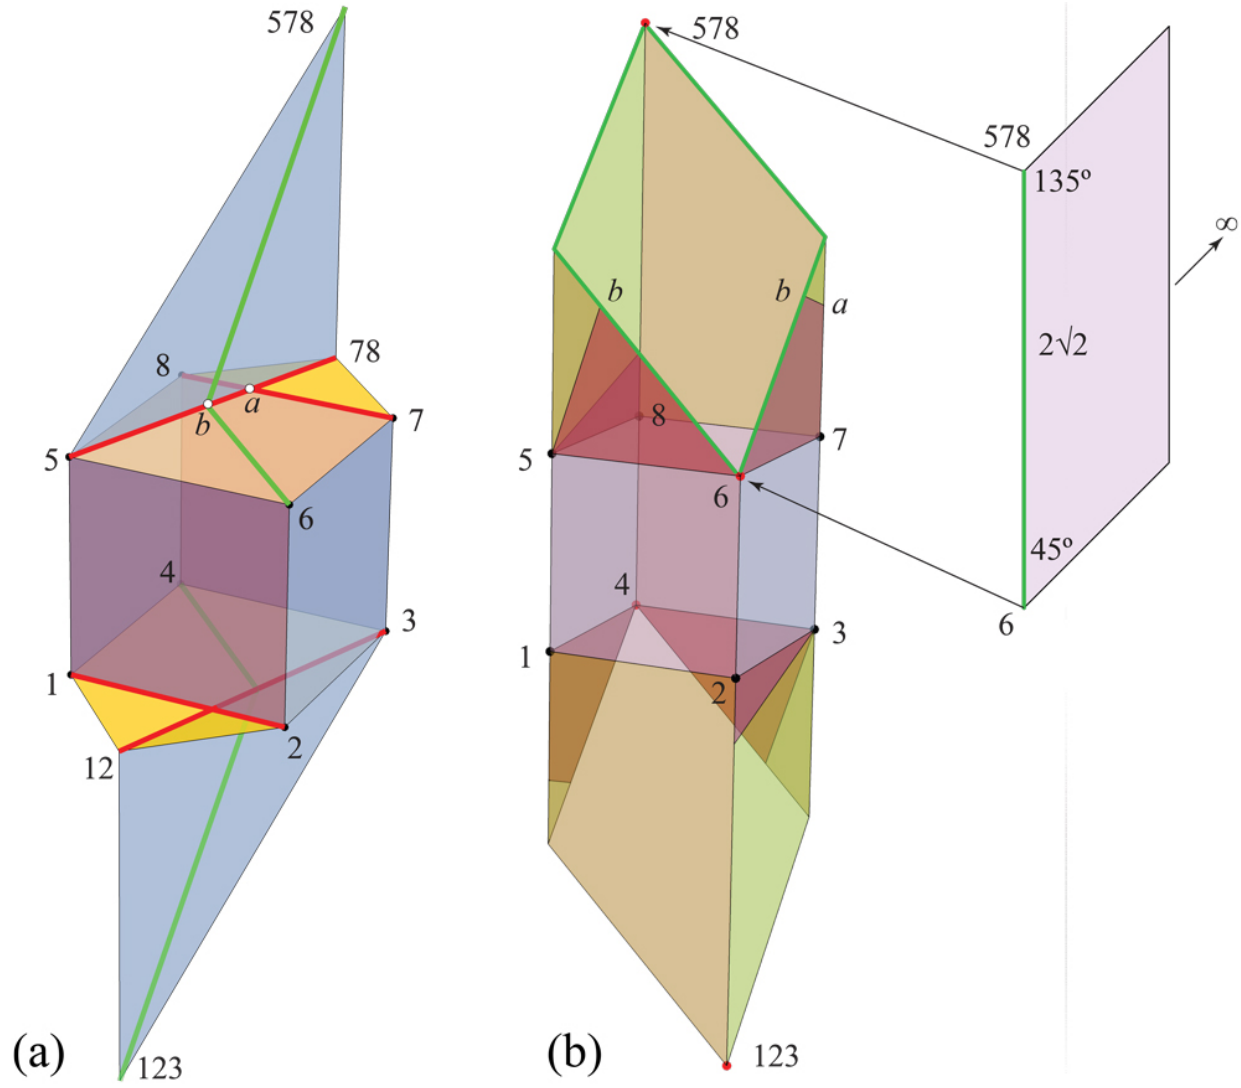

Figure 16.4: (a) Cutting along the v_6v_{578} geodesic segment (green), and inserting double parallelograms (b) leads to an unbounded cylinder above, and similarly below. In (b) the yellow regions are inserted merge triangles; pink regions pieces of the top and bottom cube faces.

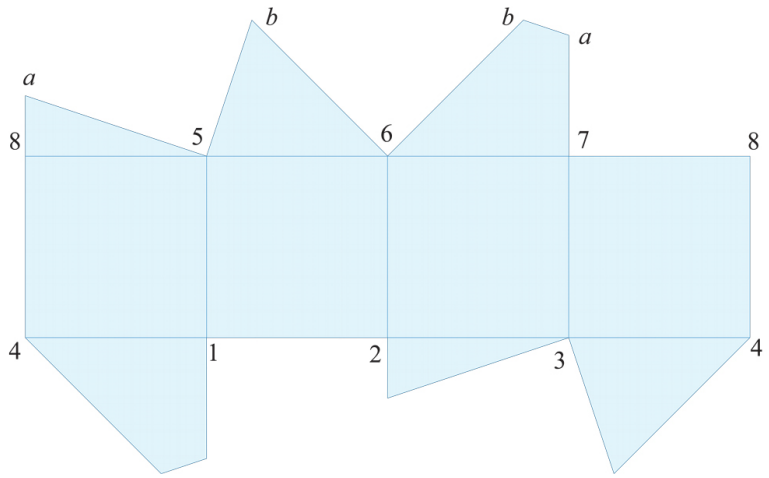

Figure 16.5: Unfolding of cube to a net by rolling the cylinder on the plane.

16.2.4 Icosahedron Example

For a second example of Theorem 16.6, we revisit and continue the example of the icosahedron from Sections 11.4 and 15.2. Fig. 16.6 below repeats for easy reference Fig. 15.2(a) showing the cuts $\tilde{g}_i$.

The triangles inserts T_1, T_2, T_3, T_4 are shown in Fig. 16.7. Note the half angles at the merge vertex m_i apexes of the triangles are $120°, 90°, 60°, 30°$ respectively. The fifth merge vertex m_5 would have zero angle, and represents the infinite parallelogram that sends m_5 to infinity.

Inserting the doubled triangles T_i^2 along the cuts g_i produces the layout shown in Fig. 16.8. Here we only show the top half of the icosahedron.

The assembly closes to a half-cylinder shown in Fig. 16.9. The infinite parallelogram T_5 with angles 60° and 120° (not shown) would attach at v_6 and m_4 respectively. Joining the two symmetric half-cylinders together and rolling on the plane produces a net for the icosahedron, the white faces in Fig. 16.8.

Earlier we discussed the quasigeodesic $v_1v_2v_3v_4v_5$, call it Q'. Note that Q' can be viewed as the Q in Fig. 16.8 slid upward parallel to itself until it touches those vertices. So both lead to the same embedding on the cylinder $\mathcal{C}$.

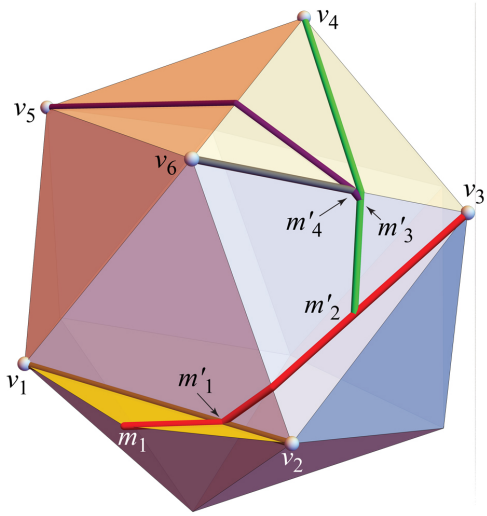

Figure 16.6: The vertex-merge cuts on P. Repeat of Fig. 15.2(a).

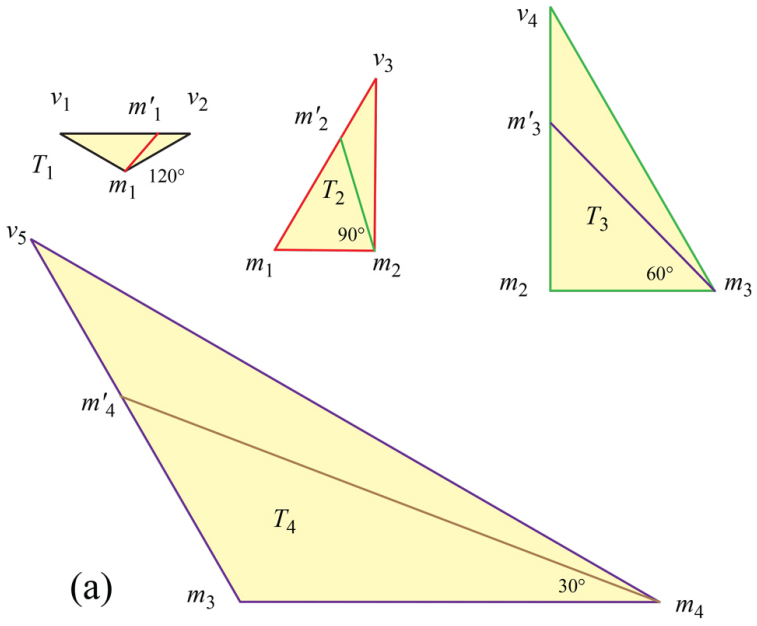

Figure 16.7: The four triangles T_i. Each is doubled to T_i^2 and inserted along g_i.

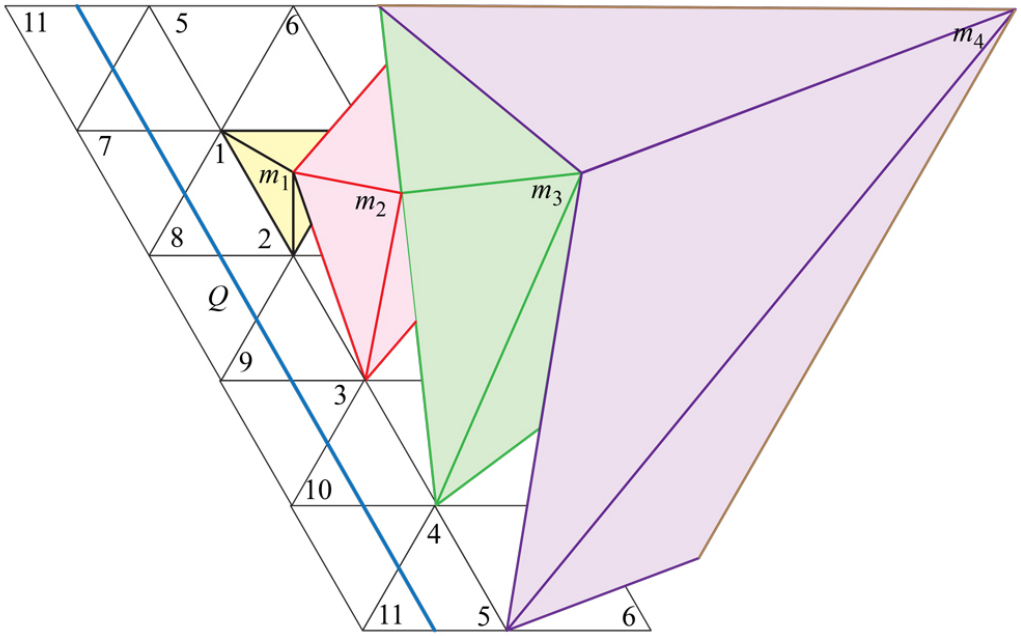

Figure 16.8: Unfolding of the top half of the icoshedron with T_i^2 inserted (shaded polygonal domains), $i = 1, 2, 3, 4$. Icosahedron faces are white; Q is blue.

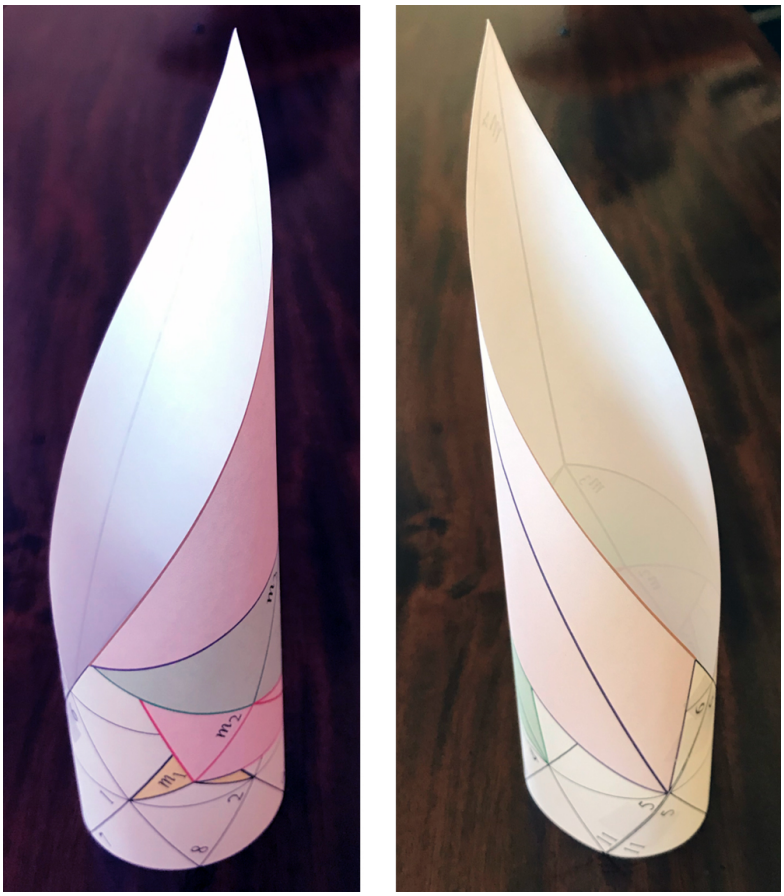

Figure 16.9: The half-cylinder obtained from Fig. 16.8 by joining the two images of $v_{11}v_5v_6$.

Chapter 17

Vertices on Quasigeodesics

Theorem 16.5 demonstrated the importance in our context of the number of vertices on a quasigeodesic. If, as we conjecture in Open Problem 18.13, every convex polyhedron P has a quasigeodesic Q containing at most one vertex, then the vertex-merging described in that theorem leads to an unfolding of P to a cylinder $\mathcal{C}$ and then to a net for P. In this chapter, we prove that the space $\mathcal{P}_n$ of convex polyhedra of n vertices contains infinitely many polyhedra with quasigeodesics through one or two vertices, or in fact through any number k of vertices, $1 \leq k \leq n$. "Infinitely many" is expressed as a set in $\mathcal{P}_n$ with non-empty interior. Precise definitions are given below.

17.1 Notation

We continue to use P for a convex polyhedron and Q for a quasigeodesic. As in the previous chapter, $V(Q) = \{q_1, \ldots, q_k\}$ is the set of vertices on Q; hence $k = |V(Q)|$. The spaces of polyhedra are defined as follows:

- For any $n \in \mathbb{N}$, let $\mathcal{P}_n$ be the space of all convex polyhedra in $\mathbb{R}^3$ with n vertices, with the topology induced by the usual Pompeiu-Hausdorff metric. Two polyhedra in $\mathcal{P}$ are then close to each other if and only if they have close respective vertices. (See Chapter 9.)

- For all $0 \leq k \leq n$, $\mathcal{Q}_k = \mathcal{Q}_k(n)$ denotes the subset of $\mathcal{P}_n$ such that every $P \in \mathcal{Q}_k$ has a simple closed quasigeodesic Q through exactly k vertices.

17.2 $\mathcal{Q}_k$ Theorem

The goal of this chapter is to prove the following result.

Theorem 17.1.

(1) $\mathcal{Q}_0$ has empty interior in $\mathcal{P}_n$.

(2) $\mathcal{Q}_k$ has non-empty interior in $\mathcal{P}_n$, for all $1 \leq k \leq n$.

If our conjecture in Open Problem 18.13 is true, then $\mathcal{Q}_0 \cup \mathcal{Q}_1 = \mathcal{P}_n$.

It is not surprising that $\mathcal{Q}_0$ is rare, as it requires a partition of vertex curvatures into two halves of exactly 2π each. Indeed, (1) is a known result [Gru91], [Gal03]. As far as we are aware, (2) is new.

The proof of the second part of Theorem 17.1 will follow directly from several lemmas. Our approach is to identify a polyhedron P in $\mathcal{Q}_k$ for each k, and then "fatten" P via the next lemma.

Lemma 17.2. *Assume the convex polyhedron P has a simple closed quasigeodesic Q with $|V(Q)| = k \geq 1$ and*

$$\alpha_i < \pi, \ \ \beta_i < \pi, \ \ \forall\, 1 \leq i \leq k. \qquad (*)$$

Then, all polyhedra P' sufficiently close to P in $\mathcal{P}_n$ have such a quasigeodesic.

Proof. The complete angles at the vertices of P depend continuously on the vertex positions in $\mathbb{R}^3$. Also, the geoarcs between two vertices remain separated from other vertices by positive distances, so they also depend continuously on the vertex positions, for small perturbations of the vertices of P. Therefore, the strict inequalities $\alpha_i < \pi$, $\beta_i < \pi$ everywhere also remain strict for small perturbations of the vertices of P.

In other words, there exists a neighborhood of P, such that each polyhedron P' in this neighborhood still has a quasigeodesic Q' with the same number of vertices as Q: $|V(Q)| = |V(Q')|$. And this is true for all Q on P. □

In view of Lemma 17.2, in order to prove the second part of Theorem 17.1, it suffices to find examples of polyhedra of n vertices with Q satisfying $|V(Q)| = k$ and Eq. $(*)$, for each $1 \leq k \leq n$.

17.2.1 $|V(Q)| = 1$

Following the plan just articulated, we will identify doubly-covered convex polygons in $\mathcal{Q}_1$, and then apply Lemma 17.2.

Recall that the *width* of a convex polygon P is the shortest distance between parallel supporting lines. The characterization (a) in the lemma below has long been known. We also need (b), which we could not find in the literature.

Lemma 17.3. *The width of P is always achieved by (a) one of the supporting lines flush with an edge e of P, and (b) the other line including a vertex that projects onto e.*

Proof. Claim (a) is Theorem 2.1 in [HT88], and known earlier. Their proof uses a lemma (their Lemma 2.1) that says the following. Let u and v be two points in the plane, and parallel lines L_u and L_v through each respectively. Then there exists a "preferred direction of rotation" of the lines, maintaining parallelism and maintaining contact with u and v, that diminishes the separation between the lines. They call this the PDR lemma. In the special case when the lines are orthogonal to vu, then both directions of rotation reduce separation.

Claim (b) relies on this PDR lemma. But as they skip proving their lemma, we include a proof here. Let $ab = e$ be the edge of P with supporting line L_{ab} including e. Let the parallel line be L_c, including only vertex c, with the separation between the lines width w. Assume, in contradiction to claim (b), that c does not project onto ab. The situation is then as depicted in Fig. 17.1, with ac playing the role of uv in the PDR lemma.

Rotate L_{ab} and L_c about a, clockwise in the figure—this is the preferred direction. Then L_c pivots around point p, and is no longer supporting P. Now move it (down in the figure) until it again contacts c. Call the new lines L'_{ab} (red) and L'_c (green), and let w' be their separation. Note that $h^2 = w^2 + x^2$, and $h^2 = (w')^2 + (x')^2$. The rotation ensures that $x' > x$, and therefore $w' < w$, contradicting the assumption that the width is w.

If c instead does project into ab, as per claim (b), say to point $q \in ab$, then the line L_{ab} is blocked from the preferred direction of rotation because it would penetrate P at q.

If c is the endpoint of an edge cd parallel to ab, then at least one of the four vertices $\{a, b, c, d\}$ must project onto the opposite edge, so claim (b) holds. □

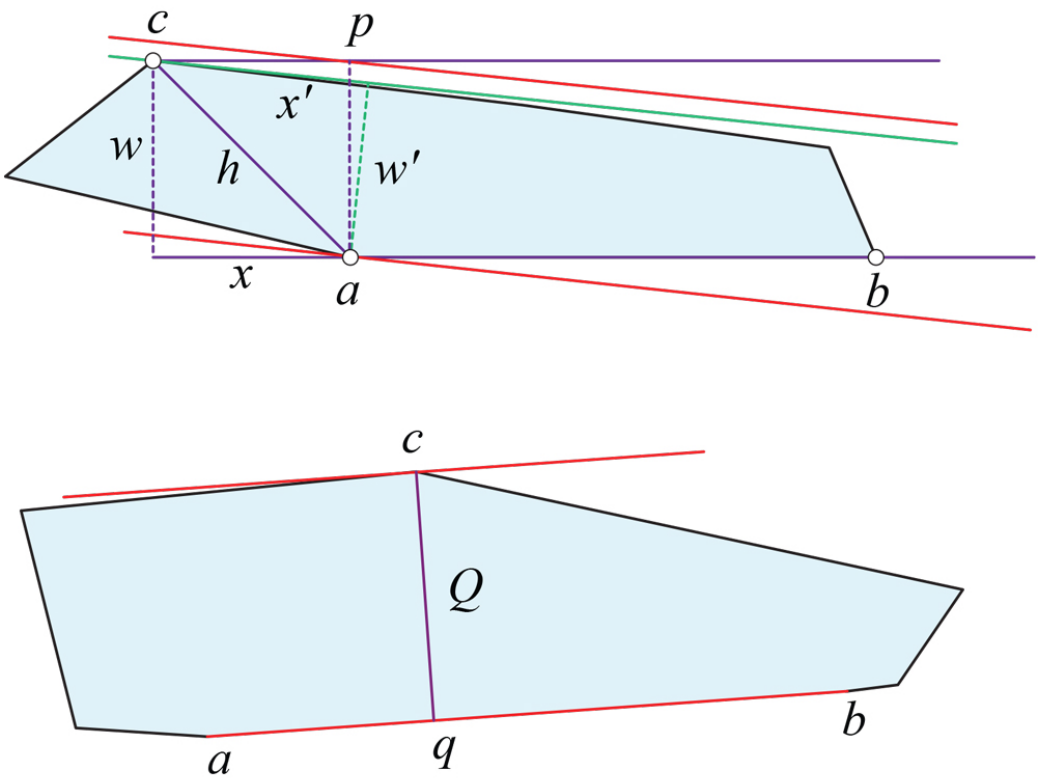

Figure 17.1: (a) Vertex c projects outside ab. (b) cq realizes the width.

Lemma 17.4. *Every doubly-covered convex polygon P with no parallel edges has a quasigeodesic through one vertex.*

Proof. Let Q be the segment from a vertex c orthogonal to an edge ab at a point q, as guaranteed by Lemma 17.3 and illustrated in Fig. 17.1(b). The angles on the doubled polygon at q are both π, and the left and right angles at c are both $< \pi$ (strictly less than because of the no-parallel-edges assumption). Therefore $Q = cq$ is a quasigeodesic.

If $q = a$, i.e., if c projects to an endpoint of ab, then the PDR lemma shows that c and ab could not have realized the width; see Fig. 17.2(a). Both directions of rotation diminish the separation, with one direction blocked by the flush edge. Therefore, q must project into the interior of ab, and Q includes just the one vertex c. □

Although not needed for Theorem 17.1, for completeness we also handle parallel edges.

Lemma 17.5. *If the width of a convex polygon P is achieved on parallel edges, then the doubly covered P has a simple closed geodesic.*

Proof. Let the width of P be realized by edges ab and cd. If the projection of cd onto ab is a positive-width interval, then a closed geodesic can be achieved by a pair of points, one on each edge, strictly interior to this interval.

If instead the projection interval is a point, then that point corresponds to a vertex-to-vertex projection, say c to b in Fig. 17.2(b). But in this circumstance, again the PDR lemma shows that $|cb|$ cannot have been the width, as rotation of the supporting lines toward the $< \pi/2$ side of b and c (counterclockwise in the figure) reduces the separation between the lines. □

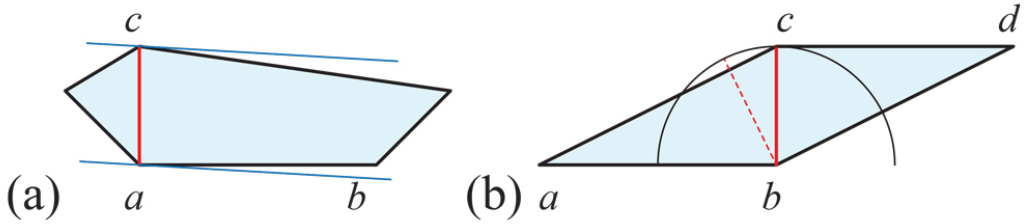

Figure 17.2: Neither ca (a) nor bc (b) realizes the width.

Note that the previous two lemmas together ensure that every doubly-covered convex polygon has a quasigeodesic through at most one vertex.

The following lemma is not needed for our main goal, but might possibly help in resolving Open Problem 18.13.

Lemma 17.6. *The width-quasigeodesic identified in Lemmas 17.4–17.5 is the shortest quasigeodesic for P.*

Proof. Call the width-quasigeodesic Q_w, and the width w. So $\ell(Q_w) = 2w$. First we show there is no shorter 1-vertex quasigeodesic. Suppose Q were such a quasigeodesic through v. It must be orthogonal to an edge e of P, and then the line parallel to e through v is a supporting line for P. Therefore, if Q were shorter than Q_w, the width of P would not be w.

Suppose now there is a shorter diagonal-quasigeodesic, connecting vertices a and b. Then lines orthogonal to ab through a and b must be supporting lines, to satisfy the angle constraints needed for the quasigeodesic to be convex to both sides at a and b. This again would contradict Q_w realizing the width.

A 0-vertex quasigeodesic Q must cross parallel edges of P. It could be slid left or right until it touches a vertex, in which case it is now a 1-vertex quasigeodesic, and the argument above applies. □

17.2.2 $|V(Q)| = 2$

Lemma 17.7. *Every doubly covered polygon has a simple closed quasigeodesic Q with $|V(Q)| = 2$*

Proof. Each extrinsic diameter of a convex polyhedron P,

$$\operatorname{diam}(P) := \max_{x,y \in P} |x - y|,$$

is realized between two vertices of P, see for example Proposition 1 in [IRV20]. Because of its length maximality, such a diameter provides, for doubly covered polygons, the desired quasigeodesics through 2 vertices. □

17.2.3 $|V(Q)| = k$, with $3 \leq k \leq n$

For these cases, we construct particular examples of polyhedra P_n of n vertices each of which has at least one quasigeodesic through k vertices.

Lemma 17.8. *For every $3 \leq k \leq n$ there exist doubly covered n-gons having simple closed quasigeodesics Q_k with $|V(Q_k)| = k$ and satisfying Eq. $(*)$.*

Proof. Consider a regular k-gon $R_k = v_1 \dots v_k$, where $3 \leq k \leq n$. Figure 17.3 illustrates the case $k = 6$.

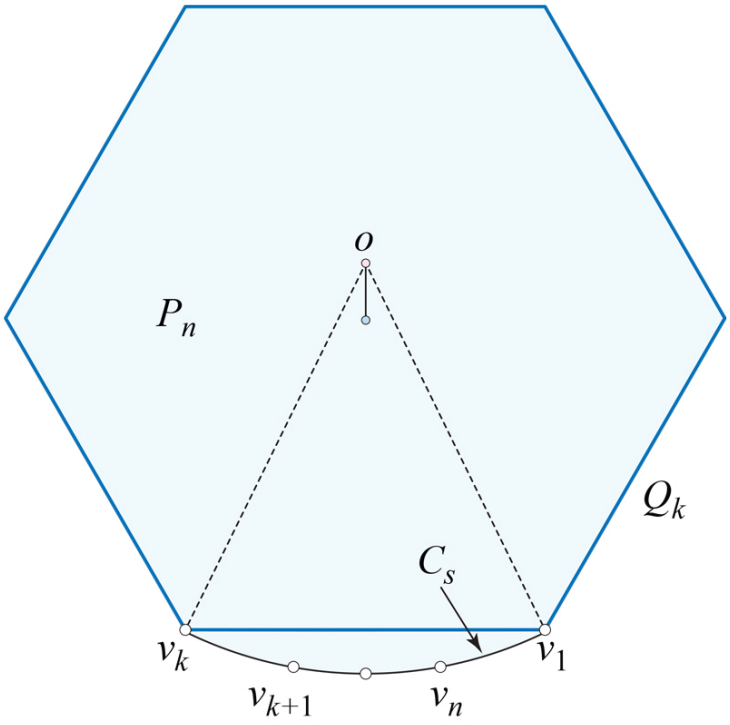

Figure 17.3: Q_6 is the blue hexagon. C_s is centered at o above the hexagon center.

Also consider the circle C through v_1 and v_k, of centre o on the same side of v_1v_k as R_k. On the short arc C_s of C determined by v_1 and v_k, choose points $v_{k+1}, \dots, v_n$, to increase the polygon to n vertices, and denote by P_n the doubly covered polygon $v_1 \dots v_n, v_1$.

When the center o is beyond the center of R_k as illustrated, the measure (length) of $m(C_s)$ satisfies $m(C_s) < 2\pi/k$. It follows that P_n has two simple closed quasigeodesics Q_k with k vertices, both corresponding to the polygon $v_1 \dots v_k, v_1$, one including edge v_1v_k on the front, and one on the back.

Indeed, for Q_k

- the angles α_i are each the vertex angles of R_k, equal to $\frac{(k-2)}{k}\pi < \pi$, for $1 \leq i \leq k$.
- the angles $\beta_i = \alpha_i$ for $1 < i < k$.
- For $i = 1$, β_1 is α_1 plus twice $\angle v_k v_1 v_n$, and $\angle v_k v_1 v_n < m(C_s)/2$. Similarly for $i = k$. Therefore, for $i \in \{1, k\}$,

$$\beta_i < \alpha_i + m(C_s) < \frac{k-2}{k}\pi + \frac{2}{k}\pi < \pi \ .$$

Thus Q_k is a quasigeodesic satisfying Eq. $(*)$ of Lemma 17.2. □

Returning to Theorem 17.1, we have identified polygons with $|V(Q)| = k$, for $1 \leq k \leq n$, and applying Lemma 17.2 creates $\mathcal{Q}_k$ with nonempty interiors. This establishes Theorem 17.1.

Chapter 18

Conclusions and Open Problems

Our work leaves open several questions of various natures, most of which have been mentioned in the text in some form. Here we list open problems and indicate possible directions for future research. The exposition roughly follows the order we treated the subjects, with the comments on both parts joined in this final chapter.

18.1 Part I

In the first part of this work, we mainly studied properties of the tailoring operation on convex polyhedra.

- We have presented three methods for tailoring, of very different flavors. On one hand, the methods of tailoring with sculpting given in Chapters 4 and 7—digon-tailoring or crest-tailoring—seem appropriate for local tailoring, and can produce any Q inside P.

  On the other hand, the method of tailoring via flattening presented in Chapter 8 is purely intrinsic, in that it doesn't need the spatial structure of P and Q to work. The surfaces can be given, in this case, as a collection of polygons glued together as in AGT. But it has the disadvantage of being "non-economical," in the sense that it discards a lot of P's surface area.

  Even with "surface-removal optimal" tailoring, we could be forced to lose

almost all the surface area of P to reshape to Q, for example if approximating by tailoring a sphere inscribed in a very long convex surface.

All three methods can be reversed to enlarge surfaces, where the results are the same, but not-requiring the spatial structure might be a clear advantage.

- We made little attempt to optimize our algorithm complexities, resting content with polynomial-time upper bounds. Likely several algorithms could be improved, or lower bounds established. Especially notable is this problem, which dominates the complexity of the tailoring-via-flattening algorithm, Theorem 8.3.

  Open Problem 18.1. *Given a polyhedron P of n vertices, find a cut-locus generic point x in less than $O(n^4)$ time.*

  Such a generic point has a unique geodesic segment (shortest path) to each vertex of P.

- Recall that Lemma 4.1 established that, with a slice along the plane of a face X of $Q \subset P$, the portion of P sliced-off can be partitioned into a fan of g-domes. Then Theorem 3.2 showed how to reduce a g-dome to its base by removing pyramids.

  Rather than repeating this for each face X of Q, it is conceivable that a single base X and a single g-dome suffice.

  Open Problem 18.2. *Let X be a "base" face of P, and D a fixed g-dome over X, interior to P. Is it possible to partition P into pyramids and D, with planes through the edges of X? After each sectioning we remove the sliced pyramid.*

  If g-dome D is replaced with an arbitrary interior polyhedron, then the answer is easily seen to be NO. Without the restriction to planes through base edges, the answer is YES as shown by Lemma 4.1.

- As discussed in Section 5.5.1, the proof of Theorem 5.3—that the seal graph Σ for a pyramid is a tree—depends on the ordering of digon removal.

  Open Problem 18.3. *Is the seal graph for a pyramid a tree for other orderings of digon removal?*

- We have shown in Chapter 9 that different unfoldings of $Q \subset P$ onto P exist. Clearly, the P-unfolding obtained via tailoring depends on the order of tailoring operations. The P-unfolding of Q does not necessarily have connected interior, which suggests this problem:

  Open Problem 18.4. *Is there some method and/or orderings of tailoring operations that would render interior-connected a P-unfolding of arbitrary $Q \subset P$? A less ambitious goal would be to minimize the number of interior-connected pieces.*

  Notice that Theorem 9.3 established simply connectedness in general. But in our usage, simply connected does not imply connected, e.g., the union of several disjoint disks is a disconnected set but simply connected. In Part II we pursue this question for particular P, for example, a doubly-covered triangle.

- It seems that at least a part of the present work could apply to 1-*polyhedra*. These are polyhedra whose faces are (congruent to) geodesic polygons on the unit sphere. They can approximate convex surfaces with curvature bounded below by 1 (in the sense of A. D. Alexandrov), just as convex polyhedra can approximate ordinary convex surfaces [AZ67], [IRV15].

  Open Problem 18.5. *How much of the presented results can be carried over and applied to* 1*-polyhedra?*

- One could also define tailoring for general (i.e., not necessarily convex) polyhedra, of arbitrary topology. Of course, the methods we developed here apply locally to convex caps. Globally, a necessary condition for Q to be tailored from a homothetic copy of P is to have the same topology as P. Our Theorem 4.6 might suggest this is also sufficient, but that is not true. By Alexandrov's Gluing Theorem (AGT), tailoring a convex polyhedron always produces a convex polyhedron, never a nonconvex polyhedron homeomorphic to the sphere but having negative curvature at some vertex. There is as yet no counterpart to AGT for nonconvex polyhedra. Therefore, in the general framework, tailoring could be a much subtler topic.

  Open Problem 18.6. *Is there a type of tailoring operation that permits global reshaping of non-convex polyhedra?*

18.2 Part II

In the second part of this work, we mainly studied the vertex-merging (as opposed to tailoring) processes on convex polyhedra, with the aim of producing "nice" polyhedral and planar unfoldings. Toward that aim, we needed to develop a theory of convex sets on convex polyhedra.

- For any convex polyhedron P, there clearly exist vertex-merging reductions of P that increase the surface area by the least, or the most amount.

  Open Problem 18.7. *Find upper and lower bounds on the added surface area for vertex-merging reductions of P. Find the reduction yielding the minimal, and respectively maximal area.*

- General relationships between the properties of the slit graph and those of the corresponding unfolding of P are presented in Theorem 11.5. Finding particular vertex-mergings with "good" behaviour seems to be a difficult task, only partly fulfilled in this study.

  Open Problem 18.8. *Does there exist for every convex polyhedron P a vertex-merging reduction onto a vm-irreducible surface whose slit graph has no cycle? If not, for which P are there such vm-reductions?*

- A positive answer to the above question raises another one, related to our discussion in Section 11.7.

  Open Problem 18.9. *For any simple unfolding P_S of P onto a vm-irreducible surface S, does there exists an unfolding of S in the plane which results in a net for P?*

- Our attempt to answer Open Problem 18.8 is based on a spiral-merging idea. In order to formalize our spiraling algorithms, we needed to develop a theory of convex sets on convex polyhedra. We only proved some basic results in this theory, and much remains for future study. In particular, the next question is related to Lemma 13.14 for the case when $S \supseteq Q$.[1]

  Open Problem 18.10. *Is it the case that every closed convex subset S of a convex polyhedron is either included in a half-surface bounded by a simple closed quasigeodesic Q, or is the whole surface?*

[1] Recall that in Part II, we use Q to denote a simple closed quasigeodesic, in contrast to Part I's use of Q to represent the target polyhedron.

- Of the classical results in the combinatorial theory of convex sets bearing a name, we only adapted to our framework the Krein–Milman Theorem (that a convex set is the convex hull of its extreme points); see Theorem 13.32. Perhaps the statement of that result, even though suitable to our purpose, could be improved.

  Open Problem 18.11. *Is Theorem 13.32 concerning extreme points and the relative convex hull, still true without the "relative" modifier?*

- Examples 13.13 and 13.22 show that the precise planar versions of Helly's and Radon's theorems are false on the surface of a convex polyhedron. It seems worth studying if there are versions of these theorems, and of Carathéodory's theorem, in our context.

  Open Problem 18.12. *Are there versions of Helly's, Radon's, and Carathéodory's theorems that hold on the surface of a convex polyhedron, using ag-convexity?*

  The Radon point of any four points in the plane is their geometric median, the point that minimizes the sum of distances to the other points [CEM+96]. Therefore, it could also be of some interest to study such points in our framework. Note, for example, on a doubly covered square there are two such points, the centers of the two faces.

- Theorem 16.5 established that if a polyhedron P has a quasigeodesic including at most two vertices, vm-reduction can lead to a net for P. One might view the reduction process as moving P closer and closer to planarity, as P_i has fewer and fewer vertices. Finally P is on a cylinder which rolls to a net.

  Little seems known about the structure of quasigeodesics. Aside from Pogorelov's original nonconstructive proof [Pog49], [Pog73], only a pseuopolynomial-time algorithm is available which finds a possibly self-crossing quasigeodesic [KPD09]. It is not difficult to find a plane Π through one vertex whose slice curve $C = \Pi \cap P$ halves the curvature as needed. But C is not necessarily convex [O'R03], and so in general is not a quasigeodesic.

  Open Problem 18.13. *We conjecture that every convex polyhedron either has a simple closed geodesic, or a simple closed quasigeodesic through just one vertex.*

Note that this conjecture is verified for doubly-covered convex polygons, by Lemmas 17.4 and 17.5.

Bibliography

[AAC+12] Timothy G. Abbott, Zachary Abel, David Charlton, Erik D. Demaine, Martin L. Demaine, and Scott Duke Kominers. Hinged dissections exist. *Discrete Comput. Geom.*, 47(1):150–186, 2012.

[AAOS97] Pankaj K. Agarwal, Boris Aronov, Joseph O'Rourke, and Catherine A. Schevon. Star unfolding of a polytope with applications. *SIAM J. Comput.*, 26:1689–1713, 1997.

[ADO03] Rebecca Alexander, Heather Dyson, and Joseph O'Rourke. The convex polyhedra foldable from a square. In *Proc. 2002 Japan Conf. Discrete Comput. Geom.*, volume 2866 of *Lecture Notes in Comput. Sci.*, pages 38–50. Springer-Verlag, 2003.

[AKP19] Stephanie Alexander, Vitali Kapovitch, and Anton Petrunin. Alexandrov geometry: Preliminary version no. 1. arXiv:1903.08539, 2019.

[Ale78] Stephanie Alexander. Local and global convexity in complete Riemannian manifolds. *Pacific J. Math.*, 76(2):283–289, 1978.

[Ale05] Aleksandr D. Alexandrov. *Convex Polyhedra*. Springer-Verlag, Berlin, 2005. Monographs in Mathematics. Translation of the 1950 Russian edition by N. S. Dairbekov, S. S. Kutateladze, and A. B. Sossinsky.

[ALT12] Paul T. Allen, Adam Layne, and Katharine Tsukahara. The dirichlet problem for curve shortening flow. *arXiv preprint arXiv:1208.3510*, 2012.

[ALZ20] Elena Arseneva, Stefan Langerman, and Boris Zolotov. A complete list of all convex polyhedra made by gluing regular pentagons. *J. Info. Processing*, 28:791–799, 2020.

[AN19] Sergey Avvakumov and Gabriel Nivasch. Homotopic curve shortening and the affine curve-shortening flow. *arXiv preprint arXiv:1909.00263*, 2019.

[AO92] Boris Aronov and Joseph O'Rourke. Nonoverlap of the star unfolding. *Discrete Comput. Geom.*, 8:219–250, 1992.

[AS05] Pankaj K. Agarwal and Micha Sharir. Pseudo-line arrangements: Duality, algorithms, and applications. *SIAM J. Comput.*, 34(3):526–552, 2005.

[AZ67] Aleksandr D. Alexandrov and Victor A. Zalgaller. *Intrinsic Geometry of Surfaces*. American Mathematical Society, Providence, RI, 1967.

[Ban81] Victor Bangert. Totally convex sets in complete Riemannian manifolds. *J. Differential Geom.*, 16(2):333–345, 1981.

[BI08] Alexander I. Bobenko and Ivan Izmestiev. Alexandrov's theorem, weighted Delaunay triangulations, and mixed volumes. *Ann. Inst. Fourier (Grenoble)*, 58(2):447–505, 2008.

[BZ21] Thomas Bendokat and Ralf Zimmermann. Efficient quasi-geodesics on the Stiefel manifold. arXiv:2105.07017, 2021.

[CEG+93] Bernard Chazelle, Herbert Edelsbrunner, Leonidas Guibas, Micha Sharir, and Jack Snoeyink. Computing a face in an arrangement of line segments and related problems. *SIAM J. Comput.*, 22(6):1286–1302, 1993.

[CEM+96] Kenneth L. Clarkson, David Eppstein, Gary L. Miller, Carl Sturtivant, and Shang-Hua Teng. Approximating center points with iterative Radon points. *Internat. J. Comput. Geom. Appl.*, 6(03):357–377, 1996.

[CH90] Jindong Chen and Yijie Han. Shortest paths on a polyhedron. In *Proc. 6th Ann. Symp. Comput. Geom.*, pages 360–369, 1990.

[CH96] Jindong Chen and Yijie Han. Shortest paths on a polyhedron, Part I: Computing shortest paths. *Internat. J. Comput. Geom. Appl.*, 6(02):127–144, 1996.

[CSW99] Francis Chin, Jack Snoeyink, and Cao An Wang. Finding the medial axis of a simple polygon in linear time. *Discrete Comput. Geom.*, 21(3):405–420, 1999.

[DEE+03] Erik D. Demaine, David Eppstein, Jeff Erickson, George W. Hart, and Joseph O'Rourke. Vertex-unfoldings of simplicial manifolds. In A. Bezdek, editor, *Discrete Geometry*, pages 215–228. Marcel Dekker, 2003.

[DHK21] Erik D. Demaine, Adam C. Hesterberg, and Jason S. Ku. Finding closed quasigeodesics on convex polyhedra. *Discrete Comput. Geom.*, 2021. arXiv:2008.00589. SoCG2020. To appear *Discrete Comput. Geom.*

[DO07] Erik D. Demaine and Joseph O'Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, 2007. `http://www.gfalop.org`.

[EHPN20] David Eppstein, Sariel Har-Peled, and Gabriel Nivasch. Grid peeling and the affine curve-shortening flow. *Exp. Math.*, 29(3):306–316, 2020.

[EL13] David Eppstein and Maarten Löffler. Bounds on the complexity of halfspace intersections when the bounded faces have small dimension. *Discrete Comput. Geom.*, 50(1):1–21, 2013.

[Epp20] David Eppstein. Treetopes and their graphs. *Discrete Comput. Geom.*, pages 1–31, 2020.

[Gal03] Gregorii Aleksandrovich Gal'perin. Convex polyhedra without simple closed geodesics. *Regul. Chaotic Dyn.*, 8(1):45–58, 2003.

[GM01] Clara I. Grima and Alberto Márquez. *Computational Geometry on Surfaces*. Kluwer Academic, Dordrecht, 2001.

[Gru91] Peter Gruber. A typical convex surface contains no closed geodesic. *J. Reine Angew. Math.*, 416:195–205, 1991.

[HT88] Michael E. Houle and Godfried T. Toussaint. Computing the width of a set. *IEEE Trans. Pattern Anal. Mach. Intell.*, 10(5):761–765, 1988.

[INV12] Jin-ichi Itoh, Chie Nara, and Costin Vîlcu. Continuous flattening of convex polyhedra. In *Computational Geometry*, volume 7579, pages 85–97. Lecture Notes in Comput. Sci., 2012.

[IRV15] Jin-ichi Itoh, Joël Rouyer, and Costin Vîlcu. Moderate smoothness of most alexandrov surfaces. *Internat. J. Math.*, 26(04):1540004, 2015.

[IRV20] Jin-ichi Itoh, Joël Rouyer, and Costin Vîlcu. Some inequalities for tetrahedra. *Beitr. Algebra Geom.*, pages 1–11, 2020. `https://doi.org/10.1007/s13366-020-00549-w`.

[KPD09] Daniel Kane, Gregory N. Price, and Erik D. Demaine. A pseudopolynomial algorithm for Alexandrov's theorem. In *Workshop Algorithms Data Struct.*, pages 435–446. Springer, 2009.

[LP21] Alexander Lytchak and Anton Petrunin. About every convex set in any generic Riemannian manifold. arXiv:2103.15189, 2021.

[Meg83] Nimrod Megiddo. Linear-time algorithms for linear programming in R^3 and related problems. *SIAM J. Comput.*, 12(4):759–776, 1983.

[Mit16] Octavian Mitrea. Geodesic convexity types in Riemannian manifolds. arXiv:1611.08643, 2016.

[O'R01] Joseph O'Rourke. An extension of Cauchy's arm lemma with application to curve development. In *Proc. 2000 Japan Conf. Discrete Comput. Geom.*, volume 2098 of *Lecture Notes in Comput. Sci.*, pages 280–291. Springer-Verlag, 2001.

[O'R03] Joseph O'Rourke. On the development of the intersection of a plane with a polytope. *Comput. Geom.*, 24(1):3–10, 2003.

[O'R07] Joseph O'Rourke. Computational geometry column 49. *Internat. J. Comput. Geom. Appl.*, 38(2):51–55, 2007. Also in *SIGACT News*, **38**(2): 51–55(2007), Issue 143.

[O'R13] Joseph O'Rourke. Dürer's problem. In Marjorie Senechal, editor, *Shaping Space: Exploring Polyhedra in Nature, Art, and the Geometrical Imagination*, pages 77–86. Springer, 2013.

[O'R20] Joseph O'Rourke. Vertex-transplants on a convex polyhedron. In *Proc. 32st Canad. Conf. Comput. Geom.*, 2020. To appear.

[OV14] Joseph O'Rourke and Costin Vîlcu. Conical existence of closed curves on convex polyhedra. *Comput. Geom.*, 47:149–163, 2014.

[OV20] Joseph O'Rourke and Costin Vîlcu. Tailoring for every body: Reshaping convex polyhedra. `https://arxiv.org/abs/2008.01759`, August 2020.

[Pog49] Aleksei V. Pogorelov. Quasi-geodesic lines on a convex surface. *Mat. Sb.*, 25(62):275–306, 1949. English transl., *Amer. Math. Soc. Transl.* 74, 1952.

[Pog73] Aleksei V. Pogorelov. *Extrinsic Geometry of Convex Surfaces*, volume 35 of *Translations of Mathematical Monographs*. Amer. Math. Soc., Providence, RI, 1973.

[Poi05] Henri Poincaré. Sur les lignes géodésiques des surfaces convexes. *Trans. Amer. Math. Soc.*, 6:237–274, 1905.

[Rou03] Joël Rouyer. Antipodes sur un tétraedre régulier. *J. Geom.*, 77(1-2):152–170, 2003.

[SS86] Micha Sharir and Amir Schorr. On shortest paths in polyhedral spaces. *SIAM J. Comput.*, 15:193–215, 1986.

[SS08] Yevgeny Schreiber and Micha Sharir. An optimal-time algorithm for shortest paths on a convex polytope in three dimensions. *Discrete Comput. Geom.*, pages 500–579, 2008. Twentieth Anniversary Volume.

[Udr13] Constantin Udriste. *Convex functions and optimization methods on Riemannian manifolds*, volume 297. Springer Science & Business Media, 2013.

[Vis18] Nisheeth K. Vishnoi. Geodesic convex optimization: Differentiation on manifolds, geodesics, and convexity. arXiv:1806.06373, 2018.

[Wik21] Wikipedia. Curve-shortening flow. `https://en.wikipedia.org/wiki/Curve-shortening_flow`, March 2021.

[Zal07] V. A. Zalgaller. An isoperimetric problem for tetrahedra. *J. Math. Sci. (N.Y.)*, 140(4):511–527, 2007.

[Zam91] Tudor Zamfirescu. Baire categories in convexity. *Atti Sem. Mat. Fis. Univ. Modena*, 39:139–164, 1991.