OPTIMAL ANGLE BOUNDS FOR STEINER TRIANGULATIONS OF POLYGONS

CHRISTOPHER J. BISHOP

Abstract. For any simple polygon P we compute the optimal upper and lower angle bounds for triangulating P with Steiner points, and show that these bounds are attained (except in one special case). The bounds for an N-gon are computable in time $O(N)$; this adds to a short list of optimization problems that are faster to solve with Steiner points than without them. In general, the sharp upper and lower bounds cannot both be attained by a single triangulation, although this does happen in some cases. For example, we show that any polygon with minimal interior angle θ has a triangulation with all angles in the interval $I = [\theta, 90^\circ - \min(36^\circ, \theta)/2]$, and for $\theta \leq 36^\circ$ both bounds are best possible. I compute the bounds for various examples, e.g., all regular polygons, all axis-parallel polygons and all sufficiently close approximations to a smooth curve. Surprisingly, we prove the optimal angle bounds for polygonal triangulations are the same as for triangular dissections. The proof of this verifies, in a stronger form, a 1984 conjecture of Gerver.

Date: February 21, 2021; revised Oct 11, 2021.

1991 *Mathematics Subject Classification.* Primary: 68U05, Secondary: 30C35, 52C20, 30C62 .

Key words and phrases. acute triangulation, Steiner points, dissections, Schwarz-Christoffel formula, conformal welding, quasiconformal mappings, discrete curvature, Delaunay triangulation, Gerver's conjecture.

The author is partially supported by NSF Grant DMS 1906259.

1. Statement of results

It is a problem of long-standing theoretical and practical interest to triangulate a polygon with the best possible bounds on the angles used. For example, the constrained Delaunay triangulation famously maximizes the minimal angle if no additional vertices (called Steiner points) are allowed [30], [31], and algorithms for minimizing the maximum angle (again without Steiner points) are given in [4] and [19]. Here we solve the analogous problems when Steiner points are permitted.

In this case, Burago and Zalgaller [12] proved in 1960 that every planar polygon P has an acute triangulation (all angles $< 90°$). This is best possible if we want an angle bound independent of P, and there is now a large collection of theorems, heuristics and applications involving acute triangulations, but several fundamental questions have remained open. What are the optimal upper and lower angle bounds for triangulating a given polygon P with Steiner points? Are these bounds attained or can they only be approximated? Can both bounds be attained simultaneously? How regular are the corresponding triangulations? How do the angle bounds for triangulations differ from those for dissections? Using ideas involving conformal and quasiconformal mappings, we shall answer each of these questions. However, the emphasis of this paper is on computing the optimal angle bounds; finding efficient triangulations that attain these bounds remains an interesting open question. See Section 16.

We start with some notation. Suppose $\mathcal{T}$ is a triangulation of P. Let V_P be the vertex set of P, $V_{\mathcal{T}}$ the vertex set of $\mathcal{T}$, $\partial\mathcal{T} = V_{\mathcal{T}} \cap P$ the boundary vertices of $\mathcal{T}$, and let $\text{int}(\mathcal{T}) = V_{\mathcal{T}} \setminus \partial\mathcal{T}$ denote the interior vertices. Label each $v \in V_{\mathcal{T}}$ with the number, $L(v)$, of triangles in $\mathcal{T}$ that have v as a vertex. For $v \in \partial\mathcal{T}$, we define its **discrete curvature** as $\kappa(v) = 3 - L(v)$, and for an interior vertex we set $\kappa(v) = 6 - L(v)$. Using these definitions, Euler's formula applied to a triangulation can be rewritten to look like the Gauss-Bonnet formula:

$$\sum_{v \in \text{int}(\mathcal{T})} \kappa(v) = 6 - \sum_{v \in \partial\mathcal{T}} \kappa(v). \tag{1.1}$$

We define this common value to be $\kappa(\mathcal{T})$, the curvature of the triangulation.

For $\phi > 0$, a ϕ-triangulation of P is one with all angles at most ϕ. For $\phi \in [60°, 90°]$ define the interval $I(\phi) = [180 - 2\phi, \phi]$; any ϕ-triangulation must have all of its angles

in $I(\phi)$. Let $|V_P|$ be the number of vertices in P, and for $v \in V_P$, let θ_v denote the interior angle of P at v. A labeling $L : V_P \to \mathbb{N} = \{1, 2, \dots\}$ is called ϕ-admissible if $\theta_v \in L(v) \cdot I(\phi)$ for every $v \in V_P$. The curvature of a labeling L is defined as

$$\kappa(L) = \sum_{v \in V_P} L(v) - (3|V_P| - 6) = 6 - \sum_{v \in V_P} (3 - L(v)).$$

If a labeling L of V_P comes from a ϕ-triangulation $\mathcal{T}$ of P with $\phi < 90^\circ$, then it is automatically ϕ-admissible and satisfies $\kappa(L) \le \kappa(\mathcal{T})$, since vertices of $\partial\mathcal{T} \setminus V_P$ must have degree ≥ 3. If $\phi < 72^\circ$ then $\text{int}(\mathcal{T})$ has no vertices of degree ≤ 5, so (1.1) implies $\kappa(L) \le \kappa(\mathcal{T}) \le 0$. See Figure 1. Similarly, if $\phi < \frac{5}{7} \cdot 90^\circ \approx 64.2857^\circ$ then every vertex in $\text{int}(\mathcal{T})$ has degree 6 and every vertex in $\partial\mathcal{T} \setminus V_P$ has degree 3, so $\kappa(L) = \kappa(\mathcal{T}) = 0$. Remarkably, these elementary necessary conditions are also sufficient.

Theorem 1.1. *For $60^\circ < \phi \le 90^\circ$, a polygon P has a ϕ-triangulation iff*

1. *$72^\circ \le \phi < 90^\circ$ and there is some ϕ-admissible labeling L of V_P,*
2. *$\frac{5}{7} \cdot 90^\circ \le \phi < 72^\circ$, and there is a ϕ-admissible labeling with $\kappa(L) \le 0$,*
3. *$60^\circ < \phi < \frac{5}{7} \cdot 90^\circ$, and there is a ϕ-admissible labeling with $\kappa(L) = 0$.*

For $60^\circ < \phi < 90^\circ$ define $\mathcal{K}(\phi)$ be the set of possible values of $\kappa(L)$ over all ϕ-admissible labelings of V_P; we set $\mathcal{K}(\phi) = \infty$ if there is no admissible labeling. Note that $\mathcal{K}(\phi)$ is either ∞ or a non-empty interval of integers. Hence $\mathcal{K}(\phi)$ has a unique closest element to 0, denoted $\kappa(\phi)$ (possibly 0 or ∞). The three conditions in Theorem 1.1 can be restated as $\kappa(\phi) < \infty$, $\kappa(\phi) \le 0$ and $\kappa(\phi) = 0$ respectively. (We should write $\mathcal{K}(\phi, P), \kappa(\phi, P)$, since these quantities also depend on P, but in this paper P is usually fixed and clear from context.)

Theorem 1.1 fails for $\phi = 60^\circ$: Figure 1 shows that a polygon can satisfy $\kappa(60^\circ) = 0$, but have no equilateral triangulation. However, if $\kappa(60^\circ) = 0$, then P can be triangulated using only angles $\le 60^\circ + \epsilon$ for any $\epsilon > 0$; see Lemma 5.1. This is the only case where the optimal bound need not be attained by any triangulation. Note that $\kappa(60^\circ) = 0$ if and only if P is a 60°-polygon, i.e., all angles are multiples of 60°. Such a polygon has an equilateral triangulation iff all of its edge lengths are integer multiples of a single value (the length of the triangulation edges).

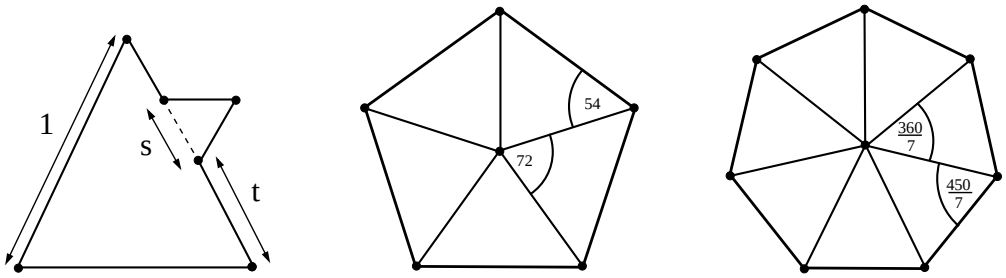

FIGURE 1. The left polygon satisfies $\kappa(60^\circ) = 0$, but has no equilateral triangulation unless both s and t are rational. The other pictures show where the special angles in Theorem 1.1 come from: these angles are forced by interior vertices of degree five or seven.

For a simple polygon P, and let $\theta_{\min}$ and $\theta_{\max}$ denote the minimum and maximum interior angles of P and define

$$\Phi(P) = \inf\{\phi : P \text{ has a } \phi\text{-triangulation }\}.$$

Theorem 1.1 makes it easy to compute $\Phi(P)$. To illustrate this, we list a few corollaries of the result and its proof. (Here and below we also given the Section number where the proof of each result can be found.)

Corollary 1.2. *[§13] $\Phi(P)$ can be computed in time $O(|V_P|)$.*

Corollary 1.3. *(§9) For any polygon, $\Phi(P) \le 90^\circ - \min(\theta_{\min}, 36^\circ)/2$. If $\theta_{\min} \le 36^\circ$, then $\Phi(P) = 90^\circ - \frac{1}{2}\theta_{\min}$ and P has a triangulation $\mathcal{T}$ with all angles in $[\theta_{\min}, 90^\circ - \theta_{\min}/2]$.*

Corollary 1.4. *[§11] In Theorem 1.1, if a ϕ-triangulation $\mathcal{T}$ exists, then it can be constructed so that all interior vertices have degree six except that*

(1) *if $72^\circ \le \phi < 90^\circ$, then $\max(0, \kappa(\phi))$ vertices have degree five,*
(2) *if $\frac{5}{7} \cdot 90^\circ \le \phi < 67.5^\circ$, then $-\kappa(\phi)$ vertices of have degree seven,*

For $67.5^\circ \le \phi \le 72^\circ$ there are $-\kappa(\phi)$ vertices of curvature -1, but these vertices may each be chosen either in the interior (degree seven) or on the boundary (degree four).

Corollary 1.5. *[§12] For the regular N-gon P_N, $\Phi(P_N) = 72^\circ$ except when $N = 3, 6, 7, 8, 9$; then $\Phi(P_N) = 60^\circ$, 60°, $\frac{5}{7} \cdot 90^\circ$, 67.5°, and 70° respectively.*

Corollary 1.6. *[§12] $\Phi(P) = 72^\circ$ for any axis-parallel polygon P, and is attained by a triangulation with two interior vertices of degree five (the rest are degree six).*

Corollary 1.7. *[§12] If $\theta_{\min} \geq 144^\circ$, then $\Phi(P) = 72^\circ$. If $\theta_{\min} \geq 162^\circ$, then every triangulation angle may be taken in $[54^\circ, 72^\circ]$. If $144^\circ \leq \theta_{\min} \leq \theta_{\max} \leq 216^\circ$, then the triangulation may be chosen with six interior vertices of degree five (the rest are degree six).*

Corollary 1.7 applies whenever P is inscribed in a smooth curve γ, and has sufficiently short edges (depending on γ). Note that 72° is frequently the sharp bound. This is not a coincidence; the set of N-gons with $\Phi(P) = 72^\circ$ contains an open set in the space $\mathcal{P}_N$ of all simple N-gons (we put a topology on $\mathcal{P}_N$ by thinking of it as a subset of $\mathbb{R}^{2N}$), and 72° is one of only two values with this property.

Corollary 1.8. *[§14] The map $\Phi : \mathcal{P}_N \to \mathbb{R}$ is continuous. Thus $\{P \in \mathcal{P}_N : \Phi(P) \leq \phi\}$ is closed in $\mathcal{P}_N$ for any $\phi \geq 60^\circ$, as is $\{P \in \mathcal{P}_N : \Phi(P) = \phi\}$. For large enough N, the latter set has non-empty interior if $\phi = \frac{5}{7} \cdot 90^\circ$ or $\phi = 72^\circ$; for other values of ϕ it always has co-dimension ≥ 1.*

Next we discuss some consequences that relate triangulations of a polygon to dissections. A triangular dissection covers P and its interior by finitely many closed triangles with disjoint interiors. The edges of adjacent triangles need not match up exactly; if they do, then we have a triangulation. See Figure 2. A ϕ-dissection is a triangular dissection with maximum angle $\leq \phi$. In 1984 Gerver [24] showed that the conditions in Theorem 1.1 are necessary if P has a $(\phi + \epsilon)$-dissection for every $\epsilon > 0$, and he conjectured that they were sufficient for a ϕ-dissection to exist if $\phi > 60^\circ$. Theorem 1.1 strengthens this by proving that a ϕ-triangulation exists.

Corollary 1.9. *For a polygon P and $\phi \in (60^\circ, 90^\circ]$, the following are equivalent:*

1. *For every $\epsilon > 0$, P has a $(\phi + \epsilon)$-dissection,*
2. *P has a ϕ-dissection,*
3. *P has a ϕ-triangulation.*

This is surprising (at least to the author). Since dissections satisfy much less stringent conditions than triangulations do, one might expect a gap between the best possible angle bounds in these two cases, but such a gap does not occur. Given a ϕ-dissection, can one obtain a ϕ-triangulation from it, e.g., as a refinement?

The conditions in Theorem 1.1 only depend on the set of angles of P, not on their ordering around P, nor on the side lengths of P. This solves Problem C7 of [15].

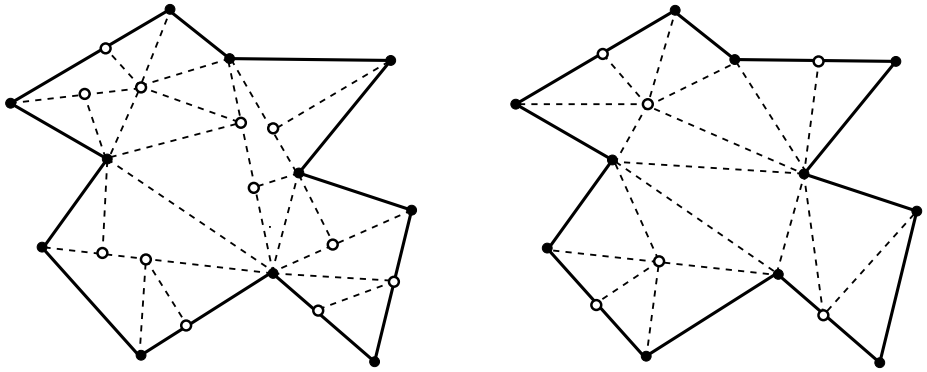

FIGURE 2. On the left is a dissection of a polygon and on the right a triangulation. The white dots the Steiner points. Despite triangulations being more restrictive than dissections, the optimal upper angle bounds are the same for both types of decomposition.

Corollary 1.10. *If $\phi > 60^\circ$ and P, P' are N-gons with the same set of angles (possibly in different orders around the boundary) then P has a ϕ-triangulation (or a ϕ-dissection) if and only if P' does.*

As noted earlier, the Delaunay triangulation of a point set maximizes the minimal angle needed to triangulate the point set without Steiner points. The constrained Delaunay triangulation does the same for polygonal triangulations without Steiner points (see [30], [31]), and an algorithm using only interior Steiner points is presented in [36]. The methods of this paper can be used to maximize the minimum angle for triangulating a polygon P with arbitrary Steiner points.

Next we consider optimal lower angle bounds for a triangulation. For $0 < \phi < 60^\circ$ we define $\widetilde{I}(\phi) = [\phi, 180^\circ - 2\phi]$; any triangle having smallest angle ϕ must have all its angles inside $\widetilde{I}(\phi)$. Define a labeling L to be ϕ-lower-admissible if $\theta_v \in L(v) \cdot \widetilde{I}(\phi)$ where θ_v is the angle of P at $v \in V_P$. The curvature $\kappa(L)$ is defined just as before, and $\widetilde{\mathcal{K}}(\phi)$ is the set of curvatures of ϕ-lower-admissible labelings. Also as before, $\widetilde{\kappa}(\phi)$ is the element of this set closest to 0 (equal to ∞ if no ϕ-lower-admissible labeling exists). A ϕ-lower-triangulation means a triangulation will all angles $\geq \phi$ and we define $\widetilde{\Phi}(P)$ to be the supremum of ϕ so that P has a ϕ-lower-triangulation.

Theorem 1.11. *[§15] For $0 < \phi < 60^\circ$, a polygon P has a ϕ-lower-triangulation iff*

(1) $0 < \phi \leq \frac{1}{7} \cdot 360^\circ \approx 51.4286^\circ$ *and* $\widetilde{\kappa}(\phi) < \infty$,
(2) $\frac{1}{7} \cdot 360^\circ < \phi \leq 54^\circ$, *and* $\widetilde{\kappa}(\phi) \geq 0$,

(3) $54^\circ < \phi < 60^\circ$, *and* $\widetilde{\kappa}(\phi) = 0$.

As before, there are a variety of consequences that follow; we list a few here.

Corollary 1.12. *[§15]* $\widetilde{\Phi}(P)$ *can be computed in time* $O(|V_P|)$.

Corollary 1.13. *[§15] If* P *has a* ϕ*-lower-triangulation then it also has an acute* ϕ*-lower-triangulation.*

Corollary 1.14. *[§15] If* $\theta_{\min} \le 45^\circ$, *then* $\widetilde{\Phi}(P) = \theta_{\min}$.

Comparing Corollaries 1.3 and 1.14, we see that if P is a polygon with $\theta_{\min} = 45^\circ$, then $\Phi(P) \le 72^\circ$ and $\widetilde{\Phi}(P) \ge 45^\circ$. If P has at least one angle $\theta \in (72^\circ, 90^\circ)$, then any triangulation of P that attains the optimal upper bound $\Phi(P)$ must subdivide θ and hence has an angle strictly less than $45^\circ \le \widetilde{\Phi}(\phi)$.

Corollary 1.15. *There exist polygons so that any triangulation attaining the optimal upper angle bound* $\Phi(P)$*, does not achieve the optimal lower angle bound* $\widetilde{\Phi}(P)$.

In this paper we will give detailed proofs of Theorem 1.1 and its corollaries. The proofs of Theorem 1.11 and its consequences are very similar and usually we only sketch the necessary changes. The idea behind both results is to associate to each polygon P a model polygon P', and then to transfer a nearly equilateral triangulation from P' to P using a conformal map. We map the triangulation vertices from P' to P and connect them by segments in P; we call these the "pushed forward" triangles (conformal images of the triangles themselves would have curved sides). A simple example is shown in Figure 3. The labeling shown in Figure 3 is 72°-admissible and has curvature 0, so $\kappa(72^\circ) = 0$. Moreover, the reader can check that $\kappa(\phi) > 0$ for $\phi < 72^\circ$, hence $\Phi(P) = 72^\circ$. As the mesh in Figure 3 gets finer, the largest angle tends to 72°, and we will show later that this limiting bound can be attained by modifying a sufficiently fine triangulation near the vertices (see Lemma 6.3).

We want P' to have a (nearly) equilateral triangulation, so we will assume it is a 60°-polygon. Note that such a triangulation of P' has all interior vertices of degree six. The angle of P' at the vertex corresponding to $v \in V_P$ is given by $60^\circ \cdot L(v)$, where L is a labeling of V_P. Since the angles of P' sum to $(|V_P| - 2) \cdot 180^\circ$, a short calculation shows that we must have $\kappa(L) = 0$. Given such a labeling of V_P, we will

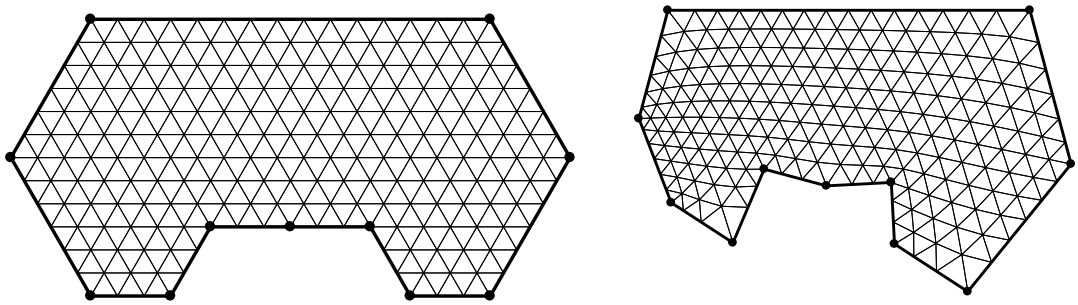

FIGURE 3. An equilateral triangulation of P' (left) and its conformal image P (angles: 270°, 126°, 96°, 126°, 105°, 105°, 144°, 144°, 80°, 262°, 162°). The maximum angle used here is $\approx 73.5205°$ and approaches 72° as the mesh gets finer; 72° is sharp by Theorem 1.1 and can be attained by modifying a sufficiently fine triangulation near the vertices.

use the Schwarz-Christoffel formula to define P' and f. If we transfer a sufficiently fine and nearly equilateral triangulation from P' to P using f, the image will be close to a ϕ-triangulation of P if the labeling L is ϕ-admissible. Thus to start the construction it seems that we need a ϕ-admissible labeling of P with zero curvature.

However, such a labeling need not exist. For example, suppose P is a pentagon with five equal angles of 108°. Theorem 1.1 (in particular, Corollaries 1.5 and 1.10) implies that $\Phi(P) = 72°$. However, the only 72°-admissible labels for a 108°-vertex are $\{2, 3\}$, so $\mathcal{K}(\phi) = \{1, \ldots, 6\}$ and so $\kappa(\phi) = 1 > 0$. This holds even if we add extra 180°-vertices to P (their possible labels are $\{3, 4, 5\}$). Therefore, any 72°-triangulation of P has positive curvature and thus at least one interior vertex of degree five (degree ≤ 4 implies an angle $\geq 90°$).

How can such a triangulation be a conformal image of a triangulation of P' that only has interior vertices of degree six? The answer is that we can choose f to conformally map the interior of P' into a subdomain of P obtained by cutting a slit in P. Two adjacent edges of P' are mapped to the two sides of the slit, and some boundary vertices of P' become interior vertices of P, thus the topology of the triangulation changes. See Figure 4. In general, up to $|\kappa(\Phi(P))|$ slits are used, introducing vertices of either degree five ($\kappa > 0$) or degree seven ($\kappa < 0$).

This scheme encounters a number of difficulties, that we will overcome using ideas from complex function theory. We list a few here, giving details later.

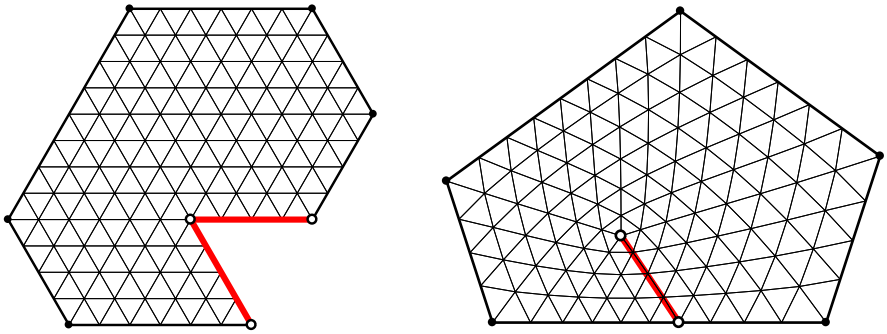

FIGURE 4. The edges adjacent to the 300°-vertex of P' are mapped to the two sides of the slit inside P. An interior vertex v of degree 5 is created. The slit is slightly curved to make the triangles on either side match up, although this is not easily visible (the actual slit lies slightly above the chord between its endpoints).

• **Conformal welding:** When we map boundary edges of P' to a slit in P, the images of certain boundary triangles in P' must match up across the slit in P, so that the image is a triangulation and not a dissection. This is only possible if the shape of the slit is carefully chosen so that arclength on each boundary segment maps to the same measure on the slit, i.e., we need $f(z) = f(w) \Rightarrow |f'(z)| = |f'(w)|$. This is a special case of a conformal welding problem, e.g., [8], [26], [38], and in our case it can be solved explicitly. The slit is generally not a segment, but may be close to one, as in Figure 4. If the slit were straight, then the Schwarz reflection principle would imply it lies on a line of symmetry for the triangles, which is not true here.

• **Riemann surfaces:** In cases where we introduce an interior vertex of degree seven, P' will need to have a boundary vertex of degree seven, i.e., P' has interior angle 420° at some vertex. Thus we necessarily consider "polygons" P' that are actually Riemann surfaces and not planar regions. See Figure 5.

• **Distortion estimates:** A conformal map f preserves interior angles infinitesimally, but to control our triangulation angles we shall need angle distortion estimates at positive scales, with bounds depending on the size of the triangle, its distance to the nearest vertex, and the ratio of the corresponding angles in P and P' at that vertex. Here we make use of the classical distortion theorems for conformal maps.

• **Harmonic measure:** Transferring a nearly equilateral triangulation from P' to P will allow us to approximate the optimal bounds, but to actually attain them, we

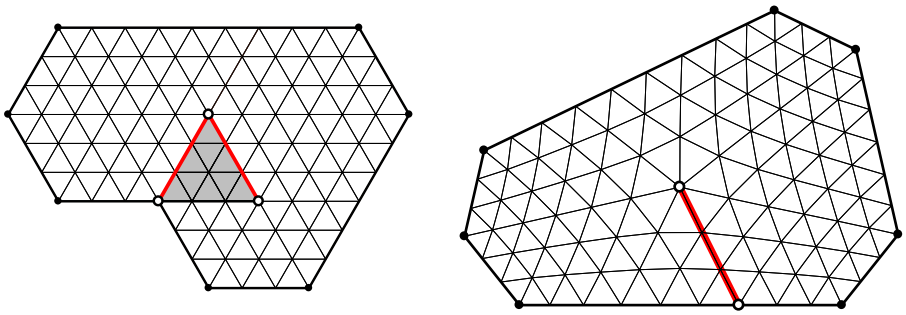

FIGURE 5. We triangulate an equal-angle heptagon using a Riemann surface with a 420°-vertex to insert a degree 7 vertex into the triangulation. The self-overlapping part of the surface is shaded.

need to also use triangulations of infinite sectors that arise as images of an equilateral triangulation of a 60°-sector under a power map z^α. Thus in some regions of P we utilize triangulations arising from two different conformal maps, and we use harmonic measure estimates to bound the difference between these conformal maps.

• **Quasiconformal mappings:** Given two triangulations of a region arising from different, but close, conformal maps, we merge the triangulations by using a partition of unity to interpolate from one conformal map to the other. The result is a quasiconformal map, and we shall use standard estimates on the angle distortion of such maps to control the angles of the interpolated triangulations. We also use such estimates to prove that any 60°-polygon has a nearly equilateral triangulation.

Once all the technical problems have been solved, we obtain a ϕ-triangulation of P. The "bad" triangles (e.g., those that are more than 1° from being equilateral) form disjoint islands that are separated by a sea of nearly equilateral triangles. Computing the optimal angle bound for P reduces to computing the worst bound near each vertex of P. This is why the optimal angle bounds can be computed in linear time: conformal mapping reduces the problem for an N-gon to $O(N)$ independent calculations each taking time $O(1)$. In 1992 Edelsbrunner, Tan and Waupotitsch [19] gave a $O(N^2 \log N)$ algorithm for minimizing the maximum angle of a polygonal triangulation without using Steiner points (their result has not been improved, so far as I know). Thus Corollary 1.2 says that it is actually faster to compute the optimal angle bound when allowing Steiner points. However a triangulation achieving the optimal bound might not be computable in polynomial; see the remarks below. This

paper does not address the question of finding efficient triangulations attaining the optimal bounds. See Section 16 for some ideas on how this might be done.

A similar situation regarding Steiner points occurs when decomposing a polygon into convex pieces: Chazelle and Dobkin [13] obtain a faster solution using Steiner points than is obtained by Keil and Snoeyink [28] without Steiner points. This result, and the one given in this paper, run counter to the general expectation that an optimization problem becomes harder by introducing Steiner points. For example, it is unknown whether a minimum-weight Steiner triangulation (smallest total edge length) even exists for every polygon with vertices in general position (it is known that it need not exist if we allow some co-linear vertices, [9]), and it is even challenging to simply approximate the optimal length to within a constant factor, e.g., [21].

Finding triangulations with good angle bounds has a long history and many applications, e.g, see [10] or [46] for lists of algorithms, such as the finite element method, that work better with well formed meshes. Corollary 1.3 gives an explicit bound for the acute triangulation theorem of Burago and Zalgaller [12] mentioned earlier. Their result was an element of their polyhedral version of the Nash embedding theorem, but it long remained unknown in the computational geometry literature. The first reference to it that we are aware of is [27] in 2004. In 1988 Baker, Grosse and Rafferty [3] independently proved that every polygon has a non-obtuse triangulation (all angles $\leq 90^\circ$). This led to a large literature on algorithms for finding triangulations in various settings with guaranteed angle bounds, e.g., [5], [6], [7], [10], [18], [22], [29], [33], [35], [39]. For a recent survey, see Chapter 29 of [25]. In 2002 Maehara [34], showed every non-obtuse triangulation can be converted to an acute one (with a comparable number of triangles), giving an alternate proof of the Burago-Zalgaller result (see also Yuan's paper [45]). A simpler approach was given by Saraf in [40].

Despite much effort devoted to finding triangulations with good geometry and optimal complexity, finding triangulations with optimal geometry has attracted less attention, at least when Steiner points are allowed. One case that has been considered is triangulating the square with optimal angles, a problem discussed by Gerver in [24] and by Eppstein in [20]. One possible reason that Theorem 1.1 has been overlooked is the close connection to conformal mappings; at least it is difficult to see how our proof could have been discovered using purely discrete geometric ideas.

Another reason may be the traditional focus on complexity. If the size of the triangulation is bounded by a function of $N = |V_P|$, independent of the geometry, then 90° is the best possible upper angle bound. For example, if a $1 \times R$ rectangle with $R \gg 1$ is triangulated by $O(1)$ triangles, then there must be a small angle $\theta = O(1/R)$, and hence some angle $\geq 90^\circ - \theta/2$. Thus the complexity of an angle-optimal triangulation of an N-gon is not polynomial in N. Even so, the sharp angle bounds proven here are fast to compute and provide a benchmark against which other triangulation methods can be compared.

Section 2 gives an overview of the proof Theorem 1.1, and later sections provide the details. We end with some calculations, questions and remarks. I thank Joe Mitchell and Herbert Edelsbrunner for their encouraging and helpful comments on an earlier version of this paper. I also thank the anonymous referees who read the paper for SODA 2022 for numerous suggestions and corrections. Several figures are drawn using Toby Driscoll's SC-Toolbox package for MATLAB [16], an improved version of an earlier algorithm of Nick Trefethen [44] for computing Schwarz-Christoffel maps. I thank Toby for his assistance with the toolbox.

2. Overview of the proof

The basic idea is quite simple: we introduce a class of polygons that have "nearly equilateral" triangulations (all angles close to 60°) and use conformal maps to transfer these triangulations to general polygons.

We will say that a simple polygon is an equilateral grid-polygon if its edges are contained in a grid of the plane consisting of congruent equilateral triangles, and its vertices are vertices of the grid. These are exactly the simple polygons that have a triangulation by equilateral triangles. See Figure 6.

It will be convenient to enlarge this class to the class of 60°-polygons, whose interior angles are all multiples of 60°. We will say that a polygon P has nearly equilateral triangulations if for any $\epsilon > 0$ it has a triangulation with all angles in $[60^\circ - \epsilon, 60^\circ + \epsilon]$ and that each vertex of P has a neighborhood in which the triangulation elements are actually equilateral (this is needed to attain the desired angle bounds, instead of just approximating them). We will prove that every 60°-polygon has nearly equilateral triangulations; see Lemma 5.1. This lemma includes 60°-surfaces, i.e., simply

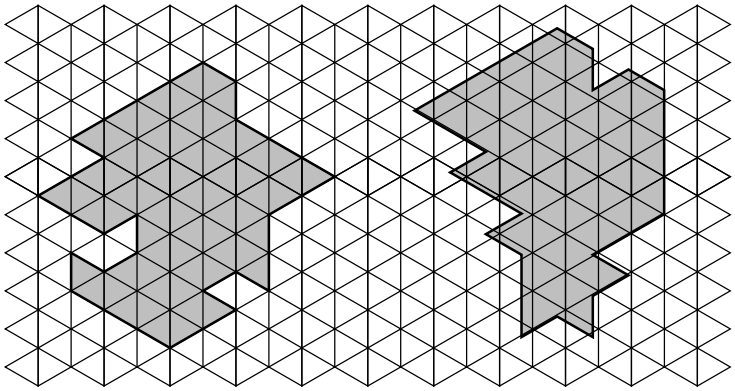

FIGURE 6. On the left is a equilateral-grid polygon, and on the right is a 60°-polygon.

connected Riemann surfaces R obtained by identifying 60°-polygons along matching edges. The boundary of R projects into the plane, possibly with self-intersections. Such surfaces can arise as Schwarz-Christoffel images of the disk (see Section 10) when all the angles are multiples of 60°, but the map is not globally 1-to-1.

Suppose $f : \Omega' \to \Omega$ is a conformal mapping between the interiors of two polygons P' and P, and that f induces a bijection between vertices of P' and vertices of P. (Below, we will often use P to refer to both the boundary curve and the interior domain, instead of using Ω for the latter; the meaning should always be clear from context.) Then f will only slightly perturb the angles of sufficiently small triangles in Ω', unless they are near vertices of P' (see Corollary 3.3). If v' is a vertex of P' with angle ψ that maps to a vertex v of P with angle θ, then any small triangle close enough to v' will have its interior angles distorted by at most θ/ψ (see Lemma 6.1).

The triangulations we construct will have all their angles between 36° and 72°, except for some triangles near vertices of P that have angle less than 36°. Larger angles of P will be subdivided by the triangulation to give new angles that are all in the interval $[36°, 72°]$, and these sub-angles should each map to 60° under the conformal map from P to P'. In order to have this work out correctly, we need an angle θ in P to correspond to an angle $\psi = L(v) \cdot 60°$ in P' that satisfies

$$\frac{3}{5}\psi \le \theta \le \frac{6}{5}\psi. \tag{2.1}$$

The restrictions imposed by (2.1) are summarized by Table 1 and Figure 7. For example, if P has a vertex v with interior angle $\theta = 135°$ the corresponding vertex v'

in the 60°-polygon P' must have angle either 120° or 180°. Any other choice means that the triangles containing v' in the nearly equilateral triangulation of P' map to triangles with angles either less than 36° or larger than 72°.

θ range	allowable ψ
0–72	60
72–108	120
108–144	120, 180
144–180	180, 240
180–216	180,240, 300
216–288	240, 300, 360
288–360	300, 360

TABLE 1. Given an angle θ of P, this table gives the possible corresponding ψ's in P' needed to attain $\Phi(P) \leq 72°$. Note that angles $\leq 36°$ in P will always give angles $\geq 72°$ in the triangulation.

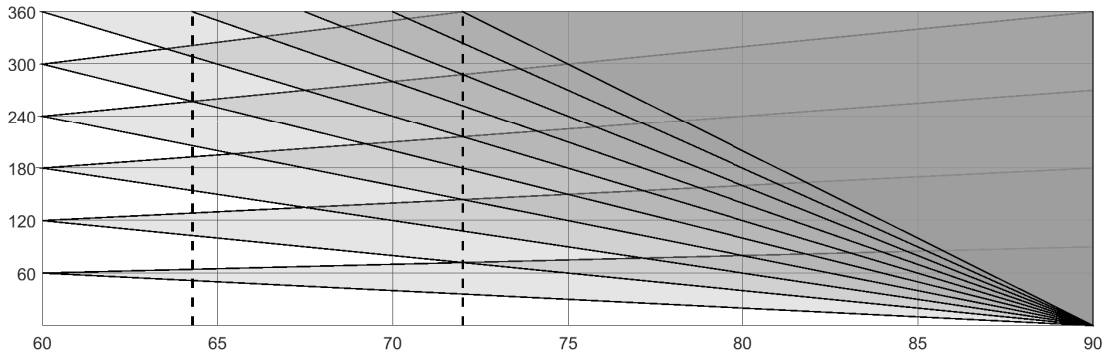

FIGURE 7. P has a ϕ-admissible labeling iff all its angles lie in the union of shaded triangles, on the vertical line through ϕ. The dashed vertical lines indicate where the transitions occur in Theorem 1.1, i.e., $\phi = \frac{5}{7} \cdot 90°$, and $\phi = 72°$.

Figure 7 plots $\cup_k k \cdot I(\phi)$ vertically above each value of ϕ. The result is a union of shaded triangles. P has a ϕ-admissible labeling if and only if all its angles lie in the intersection of the shaded region and the vertical line through ϕ. For $72° \leq \phi \leq 90°$, $\cup_k k \cdot I(\phi) = [180° - 2\phi, \phi]$ so this condition only depends on the size of $\theta_{\min}$, and is equivalent to having a ϕ-triangulation. For $\phi < 72°$, having a ϕ-triangulation requires other conditions as well, involving which triangles the angles of P lie in.

Unfortunately, there are some polygons P whose vertices cannot be put into 1-1 correspondence with the vertices of a 60°-polygon so that they satisfy the restrictions in Table 1 and Figure 7. In general, the interior angles $\{\theta_1, \ldots, \theta_N\}$ of an N-gon must satisfy $\sum_k \theta_k = (N-2)180°$. When we assign image angle values $\{\psi_1, \ldots, \psi_N\}$ using (2.1) or Table 1, we need to have $\sum_k \psi_k = (N-2)180°$, but this is sometimes impossible. For example, if P is a square, then each of its four 90° angles would have to be assigned angle 120° in P', giving an angle sum $480° > 360°$. We can "fix" the angle discrepancy by adding extra vertices to the edges of P.

First suppose $\sum_k \psi_k < \sum_k \theta_k$. We add a new vertex v of angle 180° in an edge of P, and assign the corresponding vertex v' in P' the angle $240° \le \frac{6}{5} \cdot 180°$. See Figure 8. Doing this increases the angle sum $\sum \theta_k$ by 180° but increases the angle sum $\sum \psi_k$ by 240°, decreasing the gap between them by 60°. Doing this several times we can clearly make the two sums match, as desired. Four equilateral triangles in P' touch v', and they are mapped to four triangles in P touching v. they have angle 45° at v and the opposite angles are approximately 67.5° (some distortion may occur). Hence using this method can only give ϕ-triangulations with $\phi \ge 67.5°$. This is adequate to prove Case 1 of Theorem 1.1 but a more elaborate construction is needed to prove Cases 2 and 3.

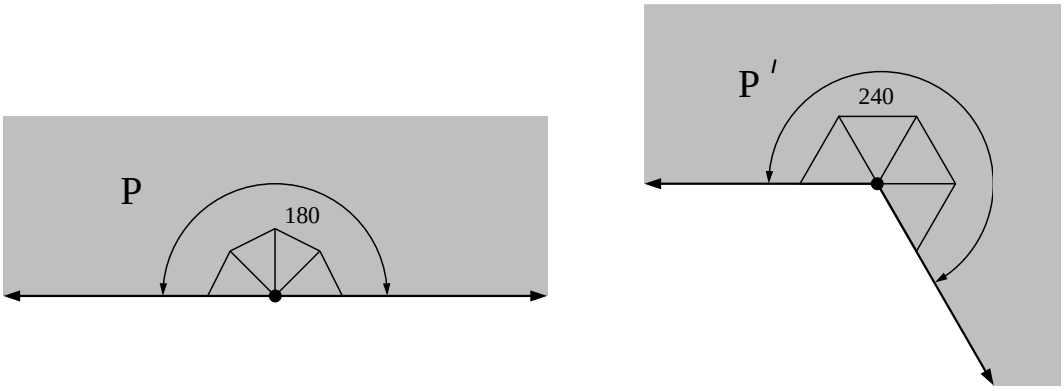

FIGURE 8. Our first trick for increasing the ψ-sum relative to the θ-sum is to pair a 180°-vertex in P with a 240°-vertex in P'. The conformal map locally looks like $z^{3/4}$, and maps a 60° sub-angle to 45°. A triangle containing this angle also contains an angle ≥ 67.5.

In Case 2 of Theorem 1.1 we want to get the angle 67.5° down to $\frac{5}{7} \cdot 90° \approx 64.2857$. We will do this by using a triangulation of a slit half-plane based on transferring an

equilateral triangulation from a polygonal Riemann surface that has a 420° angle in its boundary. The idea is shown in Figure 9; the details will be given in Section 8. The Riemann surface R is built by attaching two planar domains as shown on the left side of Figure 20. R has a 1-1 projection onto a sector of angle 240°, except for the darker triangle where it is 2-1. Traversing the boundary, we encounter angles 60°, 420° and 120°. The 420°-vertex belongs to seven triangles in R and will map to a degree seven interior vertex in the final triangulation.

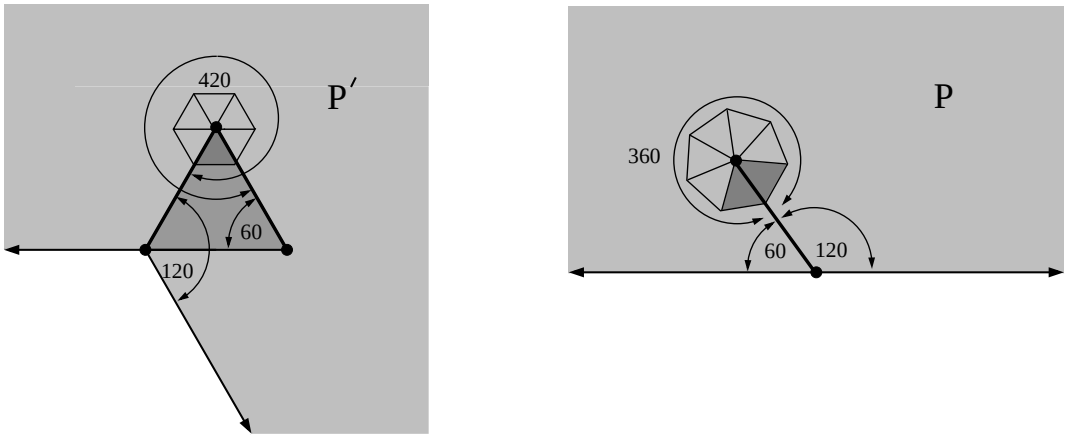

FIGURE 9. We cut a slit in the upper half-plane at angle 60°. This models a neighborhood of a 180°-vertex on the boundary of P. The angles we observe tracing the outline of the slit are: 60°, 360° and 120°. The triangulation near this slit will correspond to an equilateral triangulation of a Riemann surface R with angles 60°, 420° and 120°, pictured at right. R is a 1-1 cover of a 240°-sector, except for the darker triangle where it is 2-1.

The two segments adjacent to the 420°-vertex are mapped to a slit in P where the angles are 60°, 180° and 120°. The worst distortion comes from mapping the 420°-vertex in P' to the 360°-vertex in P (the tip of a slit). Locally the map looks like $z^{6/7}$, that maps each 60° sub-angle to $\frac{6}{7} \cdot 60° = \frac{4}{7} \cdot 90° \approx 51.4286$. A triangle with this angle must also contain an angle $\geq \frac{5}{7} \cdot 90° \approx 64.2857$, which is where this angle in Theorem 1.1 comes from. Note that the two finite boundary segments of R are both mapped to the slit in P, so triangulation edges along these two sides of R must map to matching edges along the slit. This requires that the conformal map sends

the length measures on the two segments to the same measure on the slit (but not necessarily length measure). See Figure 21.

If $\sum_k \psi_k > \sum_k \theta_k$ we use a slightly easier variant of the 420°-trick that we call the "120°-trick". This involves mapping a slit half-plane to a 120°-sector with a triangle removed, as shown in Figure 10. Traversing the boundary of the slit half-plane we encounter angles 60°, 360°, and 120°, but traversing the boundary of the modified 120°-sector we encounter 60°, 300° and 120°, so the ψ-sum decreases by 60° relative to the θ-sum. As in the 420°-trick, the shape of the slit can be chosen so that points on the two identified segments are paired according to their distance from the 300°-vertex. Then an equilateral triangulation of the modified sector maps to a triangulation of the half-plane. Note also that exactly one degree five vertex is created in the triangulation, located at the tip of the slit. See also Figure 18.

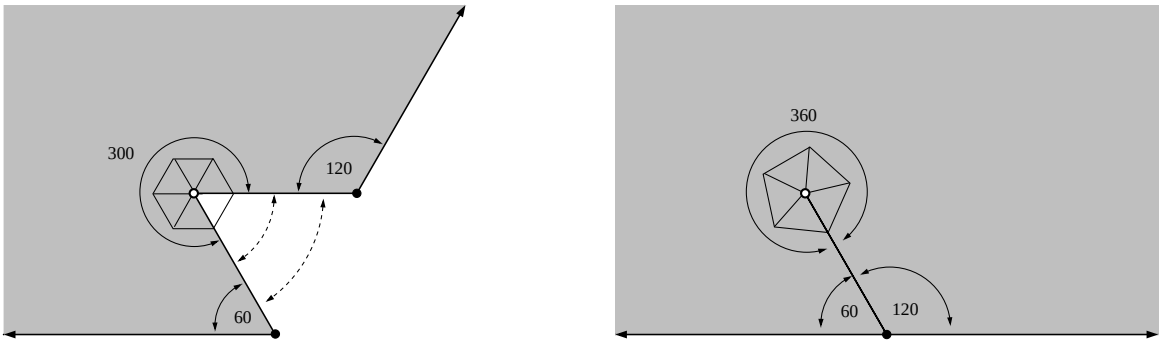

FIGURE 10. Take a 120°-sector and cut out an equilateral triangle (left). Conformally map this region to a slit half-plane (right). The two edges adjacent to the 300°-vertex are mapped to opposite sides of the (slightly) curved slit, and points on these edges equidistant from this vertex are identified. This implies that an equilateral triangulation on the left pushes forward to a triangulation on the right.

Sections 3-8 will provide the details of the various constructions sketched above. The proof of Theorem 1.1 and is consequences starts in Section 9.

3. The distortion theorem

We let $D(z,r) = \{w : |z-w| < r\}$, $\mathbb{D} = D(0,1)$ and $\mathbb{T} = \partial\mathbb{D} = \{z : |z| = 1\}$. In this paper, a conformal map always refers to a 1-to-1 holomorphic mapping. Thus is is angle and orientation preserving. By the Riemann mapping theorem, there is a

conformal map from $\mathbb{D}$ onto any proper, simply connected subdomain of the plane, in particular, the interior of any bounded polygon. The same holds for any simply connected Riemann surface this is not homeomorphic to the 2-sphere or conformally equivalent to the plane; this holds for all the Riemann surface we shall consider in this paper. The conformal map is unique if we specify the image of 0 and of one boundary point. For a conformal map onto the region Ω bounded by a N-gon P, the preimages of the vertices are called the prevertices.

We start by recalling a fact from complex analysis that we will use below (this is a special case of the Borel-Carathéodory theorem).

Lemma 3.1. *If g is holomorphic on $\mathbb{D}$ and $g(0) = 0$ then*

$$\max_{|z|<1/2} |g(z)| \leq 2 \max_{|z|<1} |\mathrm{Re}(g(z))|.$$

Proof. Without loss of generality we may assume that $\mathrm{Re}(g) < 1$ on $\mathbb{D}$, so that g maps the disk into the half-plane $H = \{x + iy : x < 1\}$. If $\tau(z) = z/(2 - z)$ is a Möbius transformation that takes H to $\mathbb{D}$ and fixes 0, then by the Schwarz lemma, $|\tau(g(z))| \leq |z|$. Therefore g maps $\mathbb{D}$ into $\tau^{-1}(D(0, \frac{1}{2})) \subset D(0, 2)$. □

By definition, a conformal map f on a domain Ω preserves angles infinitesimally. We will need to show they almost preserve angles in triangulations in the following way. When we map a triangulation by a conformal map, we don't take the image $f(T)$ of each triangle T; this would have curved sides. If $T = \Delta ABC \subset \Omega$ has vertices A, B, C, we define the pushed forward triangle $f^*(T) = \Delta f(A)f(B)f(C)$, i.e., the triangle with vertices $f(A), f(B), f(C)$. See Figure 11. We wish to prove that pushing forward small triangles in Ω by conformal maps alters the angles by as little as we wish, except near the vertices.

Lemma 3.2. *If f is a conformal map on a disk $D(z, r)$, $\delta > 0$ is sufficiently small, and $T = \Delta ABC$ is triangle inside $D(z, \delta r)$, then the triangle $f^*T = \Delta f(A)f(B)f(C)$ has angles that are within $O(\delta)$ of the corresponding angles of T.*

Proof. Since pre or post-composing f by a similarity does not change the angles of f^*T, we may assume $z = 0$, $r = 1$, $f(0) = 0$ and $f'(0) = 1$. The distortion theorem for conformal maps (e.g., Theorem I.4.5 of [23]), says that

$$\frac{1 + |z|}{1 + |z|^3} \leq |f'(z)| \leq \frac{1 + |z|}{1 - |z|^3}.$$

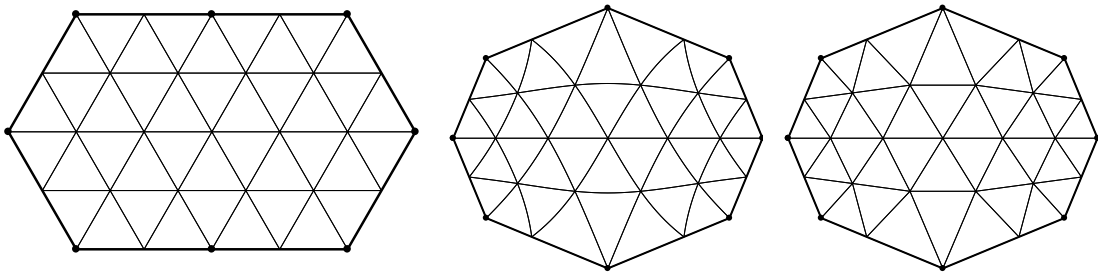

FIGURE 11. An equilateral triangulation (left), the actual conformal images of the triangles (center), and the pushed forward triangles (right) where vertex images are connected by segments.

Thus $\log|f'(z)| \le \log(1+|z|) - \log(1-|z|^3)$ and Lemma 3.1 implies

$$|\arg(f'(z))| \le 2\log(1+2|z|) + 2\log\frac{1}{1-8|z|^3} \le 8|z|,$$

if $|z| \le 1/4$. Therefore, inside $D(0,\delta)$ any segment S is mapped to a smooth curve all of whose tangents are within 8δ being parallel of S. Hence the chord connecting the endpoints of $f(S)$ is within 8δ of being parallel to S, and so the angles of f^*T are within 16δ of the corresponding angles of T. □

Next we quantify the fact that, except near the corners, a conformal map between polygons alters triangle angles very little.

Corollary 3.3. *Suppose f is conformal map between the interiors of two polygons P and P' that maps vertices to vertices. Suppose V_P is the vertex set of P, Ω is the interior of P, and that $\{D_v\}_{v\in V_P}$ are disjoint disks around each vertex v. Define $\Omega_0 = \Omega \setminus \cup_{v\in V_P} D_v$ and suppose $T \subset \Omega_0$ is a triangle. Then for every $\epsilon > 0$ there is a $\delta > 0$ so that f changes the angles of T by less than ϵ if* $\operatorname{diam}(T) < \delta$.

Proof. Let $r > 0$ be the distance from Ω_0 to V_P and let $s > 0$ the minimal distance between any two connected components of $P \setminus \cup_{v\in V_P} D_v$. If $z \in \Omega_0$ and $t = \frac{1}{2}\min(r,s)$, then we claim f extends to be conformal on the disk $D_1 = D(z,t)$. If $D_1 \subset \Omega_0$ this is obvious. Otherwise, D_1 does not hit any vertex of P and hits at most one edge e of P. Suppose $w \in D_1 \cap e$. Then $D_1 \subset D_2 = D(w,2t)$. Since the edge f of P maps to an edge of P', the Schwarz reflection principle says f can be extended to be conformal on all of D_2, hence on D_1, proving the claim.

Therefore, if $T \subset \Omega_0$ has diameter δ, then it is contained in a disk of radius t where f is conformal and hence angles of T are distorted by at most $O(\delta/t)$. Since t is fixed, this is as small as we wish, if δ is small. □

4. Approximation in finite sectors

A quasiconformal map h of the plane is a homeomorphism of the plane to itself that is absolutely continuous on almost all lines and whose complex dilatation $\mu = h_{\overline{z}}/h_z$ satisfies $\|\mu\|_\infty < 1$ (recall $h_z = \frac{1}{2}(\frac{\partial h}{\partial x} - i\frac{\partial h}{\partial y}$ and $h_{\overline{z}} = \frac{1}{2}(\frac{\partial h}{\partial x} + i\frac{\partial h}{\partial y})$. See [1] or [32] for the basic properties of such maps. One important such property is that pre or post composing a K-quasiconformal map by a conformal map gives another K-quasiconformal map: a simple calculation with the chain rule shows the phase of the dilatation may change, but its absolute value does not. A quasiconformal map h is called K-quasiconformal if $\|\mu\|_\infty \leq k = (K-1)/(K+1)$. More geometrically, at almost every point h is differentiable and its derivative (which is a real linear map) send circles to ellipses of eccentricity at most K (the eccentricity of an ellipse is the ratio of the major to minor axis). An important property of K-quasiconformal maps $\mathbb{C} \to \mathbb{C}$ is that they form a compact family when normalized to fix two points, usually taken to be 0 and 1. Also, a 1-quasiconformal map is conformal, so a 1-quasiconformal map of the plane to itself must be linear. These facts imply the following.

Lemma 4.1. *For any $\delta > 0$, there is a $\epsilon > 0$ so that if f is a $(1+\epsilon)$-quasiconformal map from $\{|z| < 1/\epsilon\}$ into $\mathbb{C}$, then $|f(z) - z| \leq \delta |f(0) - f(1)|$ on $\mathbb{D}$.*

Proof. Suppose not. Since pre and post-composing by similarities (these are conformal) does not change the quasiconformal bound, we may assume there is a sequence of K_n-quasiconformal maps on $\{|z| < n\}$ into $\mathbb{C}$ with $K_n \to 1$, that fix both 0 and 1, but so that each map moves some point of $\mathbb{D}$ by more than δ. By compactness of normalized K-quasiconformal maps, any subsequence has a uniformly convergent (on $\mathbb{D}$) subsequence to a conformal linear map, that must be the identity. This contradicts the assumption that every map moves some point by at least δ. □

Using the Law of Sines, we can deduce that under the same hypotheses the push forward of any triangle $T \subset \mathbb{D}$ has its angles changed by at most a factor of $O(\delta)$.

Given a simply connected domain Ω, a point $z \in \Omega$ and a Borel set $E \subset \partial\Omega$, the harmonic measure, $\omega(z, E, \Omega)$ is the Lebesgue length of $f(E) \subset \mathbb{T} = \partial\mathbb{D}$, where f is

a conformal map $f : \Omega \to \mathbb{D}$ with $f(z) = 0$. More intuitively, it is the probability that a Brownian motion started at z first hits $\partial\Omega$ in the set E, and is value at z of the harmonic function u on Ω that has boundary value 1 on E and zero elsewhere (appropriately defined). The text [23] of Garnett and Marshall is a comprehensive treatment of this topic, and includes all the results we shall use below.

An infinite sector is a region congruent to $S(\theta) = \{re^{\phi} : r > 0, -\theta/2 < \phi < \theta/2\}$. The boundary consists of two infinite rays meeting at its vertex and the angle of the sector is the interior angle θ made by these rays. A finite sector is a region congruent to $S(\theta) \cap D(0,t)$ for some $t > 0$ and $\theta \in (0, 360°]$. For $r > 1$, let $A(r) = \{z : 1/r < |z| < r\}$. An annular sector is a region congruent to $S(\theta, r) = S(\theta) \cap A(r)$. Let $z_0 = 1$ be the center of the circular arc $\gamma = S(\theta) \cap \mathbb{T}$ and let $z_1 = e^{i\theta/2}$ be one endpoint of γ.

Lemma 4.2. *With notation as above, suppose Ω_1, Ω_2 are simply connected domains so that for $k = 1, 2$, $\Omega_k \cap A(r)$ both have connected components equal to $S(\theta, r)$. Suppose $F : \Omega_1 \to \Omega_2$ is the conformal map that fixes z_0 and z_1. As $r \to \infty$, F converges uniformly to the identity on $S(\theta, 2)$.*

Proof. It is enough to take $\Omega_2 = S(\theta, r)$; in general, the conformal map $\Omega_1 \to \Omega_2$ can be written as the composition of conformal maps $\Omega_1 \to S(\theta, r) \to \Omega_2$, so it is enough to show each of these is close to the identity.

Let $f : \Omega_1 \to \mathbb{D}$ be the conformal map sending z_0 to 0 and z_1 to i. See Figure 12. By standard estimates, e.g., Theorem IV.6.2 of [23], the harmonic measure of the circular arcs of $S(\theta, r) \cap \partial A(r)$ with respect to z_0 is $O(r^{-\pi/\theta})$. This tends to zero as $r \nearrow \infty$. Since harmonic measure is a conformal invariant, the images of these two curves under f have the same harmonic measures with respect to 0. By the Beurling projection theorem (e.g., Theorem II.9.2 of [23]) these image curves are crosscuts of $\mathbb{D}$ that have Euclidean diameter $O(r^{-\pi/\theta})$. Therefore the conformal map g from $\mathbb{D} \setminus f(\Omega \setminus S(\theta, r))$ to $\mathbb{D}$ that fixes 0 and i converges to the identity on $f(S(\theta, 2))$ as $r \nearrow \infty$. Thus $F = f^{-1} \circ g \circ f : \Omega_1 \to S(\theta, r)$ is conformal, fixes z_0 and z_1, and tends to the identity uniformly on $S(\theta, 2)$ as $r \nearrow \infty$. □

Lemma 4.3. *Suppose Ω_1, Ω_2 are simply connected domains and that both $\Omega_1 \cap A(3)$ and $\Omega_2 \cap A(3)$ have connected components equal to $S(\theta, 3)$. Suppose f_k is conformal on Ω_k, $k = 1, 2$ and $\sup_{S(\theta,3)} |f_1 - f_2| < \epsilon$. Suppose that f_1, f_2 both map each radial*

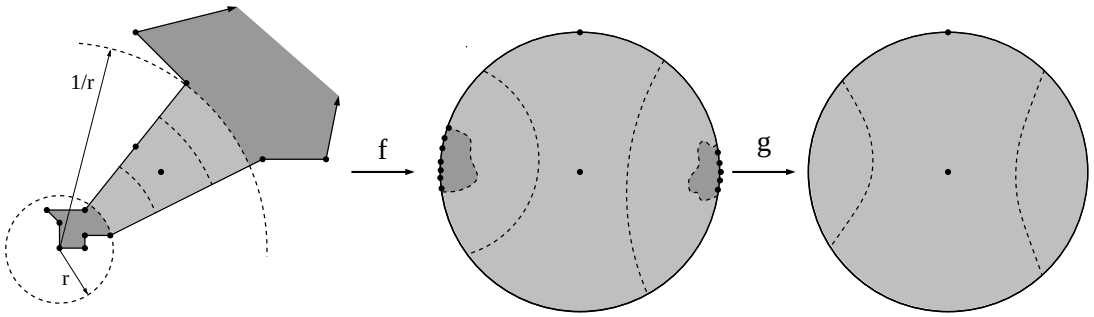

FIGURE 12. The conformal map from Ω to $S(\theta, r)$ can be written as a composition $F = f^{-1} \circ g \circ f$ where $f : \Omega \to \mathbb{D}$ and g is a map of $\mathbb{D} \setminus f(\Omega \setminus S)$ to $\mathbb{D}$. The components of $f(\Omega \setminus S(\theta))$ have small diameter as $r \nearrow \infty$, so both g and $F = f^{-1} \circ g \circ f$ are close to identity maps.

segment of $\partial S(\theta, 3)$ *into the same line. Let* $\eta : [0, \infty) \to [0, 1]$ *be smooth with* $\eta(r) = 0$ *if* $r < 1/2$ *and* $\eta(r) = 1$ *if* $r > 2$. *Then*

$$g(z) = f_1(z)(1 - \eta(|z|)) + f_2(z)\eta(|z|)$$

is quasiconformal on $\Omega_1 \cup \Omega_2$ *with complex dilatation bounded by* $O(\epsilon)$.

Proof. By the assumption on the boundaries, both f_1, f_2 can be extended by reflection across the radial boundary segments $\partial S(\theta, 3)$ and hence the Cauchy estimates are valid with a uniform radius at every point of $\partial S(\theta, 2)$. Therefore $|f_k'(z)| \simeq 1$ on $S(\theta, 2)$ for $k = 1, 2$. It also follows that

$$|f_2'(z) - f_1'(z)| = O(|f_2(z) - f_1(z)|) = O(\epsilon).$$

Next we estimate the dilatation $\mu = \partial_{\bar{z}} g / \partial_z g$. First,

$$\begin{aligned} \partial_z g(z) &= f_1'(z)(1 - \eta(|z|)) - f_1(z)\partial_z \eta(|z|) + f_2'(z)\eta(|z|) + f_2(z)\partial_z \eta(|z|) \\ &= f_1'(z) + (f_2'(z) - f_1'(z))\eta(|z|) + (f_2(z) - f_1(z))\partial_z \eta(|z|) \\ &= f_1'(z) + O(\epsilon) \simeq 1, \end{aligned}$$

because $|\eta|, |\nabla \eta| = O(1)$ (recall η is a fixed smooth function). Since $\partial_{\bar{z}} f = 0$ for holomorphic functions f,

$$|\partial_{\bar{z}} g(z)| = |(f_2(z) - f_1(z))\partial_{\bar{z}} \eta(|z|)| = O(\epsilon).$$

Thus $|\mu_g| = |\partial_{\bar{z}} g / \partial_z g| = O(\epsilon)$ as desired. □

The same argument proves the following, slightly simpler, version where we assume the maps are defined on a common disk, instead of a sector. This version will be used to merge triangulations near tips of slits in the 120°-trick (see Section 7) and the 420°-trick (see Section 8) .

Lemma 4.4. *Suppose Ω_1, Ω_2 are simply connected domains and $D(0,r) \subset \Omega_1 \cap \Omega_2$ for some $r > 4$. Suppose $f : \Omega_1 \to \Omega_2$ is conformal and $f(0) = 0$, $f'(0) = 1$. Let $\eta : [0,\infty) \to [0,1]$ be smooth with $\eta(r) = 0$ if $r \le 1/2$ and $= 1$ if $r \ge 2$. Then*

$$g(z) = z(1 - \eta(|z|)) + f(z)\eta(|z|)$$

is quasiconformal on $\Omega_1 \cup \Omega_2$ with complex dilatation bounded by $O(1/r)$.

Proof. The proof is almost identical to the proof of Lemma 4.3, except that we have to verify that $|f(z) - z| = O(1/r)$ if $|z| \le 2$. However, this is the Distortion Theorem for conformal maps, e.g., Theorem I.4.5, Equation (I.4.17) of [23]. □

5. 60°-surfaces have nearly equilateral triangulations

As noted in the introduction, not every 60°-polygon P has an equilateral triangulation, but in this section we will prove that they all have nearly equilateral triangulations. Recall that this means that for any $\epsilon > 0$, there is a triangulation of P with all angles within ϵ of 60° and that all angles are equal to 60° in some neighborhood of each vertex (the neighborhood may depend on the triangulation). Also recall that we need this result for both planar polygonal regions and Riemann surfaces with polygonal boundaries.

Lemma 5.1. *Every 60°-surface P has nearly equilateral triangulations.*

Proof. We refer to the boundary curve of a 60° Riemann surface R as P, although it need not be a simple planar polygon. However, it does project to a planar N-gon, possibly with self-intersections. By rotating, we may assume the sides of this curve are parallel to the usual equilateral grid in $\mathbb{C}$ (with one side of each triangle parallel to the vertical axis). Rescale the grid to have a small side length δ, and consider the union of triangles compactly contained in the surface R that project onto elements of the δ-grid. When δ is sufficiently small, this union forms a simply connected sub-surface R' with boundary curve P' that is also a 60°-surface and also

has N vertices. Each edge of P' corresponds to a parallel edge of P. See Figure 13. If we fix conformal map of $\mathbb{D}$ to R, then by taking δ small enough, we can choose a conformal map onto R' so that the corresponding conformal prevertices on $\mathbb{T}$ (recall these the preimages of a conformal map from the unit disk to the surface P') can be chosen to approximate the prevertices of P as closely as we wish.

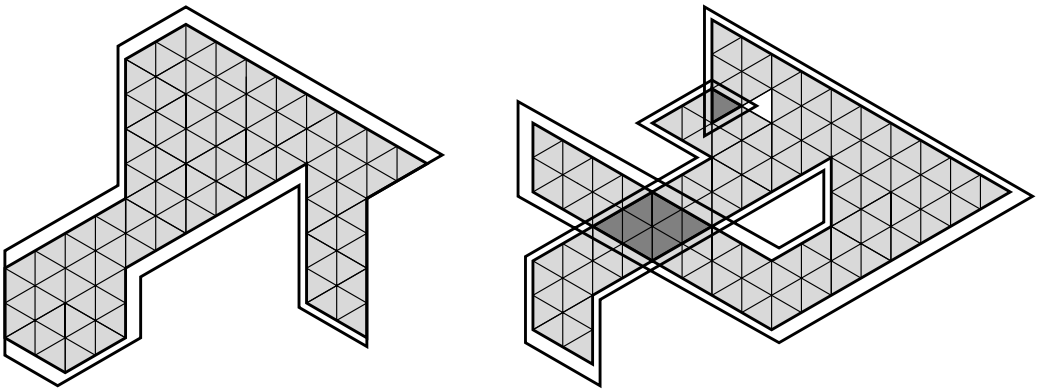

FIGURE 13. Two examples of approximating a 60° surface (white) by a union of equilateral triangles (shaded). On the left is simple polygon and on the right a Riemann surface. Overlaps of the surface with itself have darker shading.

In particular, if the minimal distance between prevertices of P on $\mathbb{T}$ is $d > 0$, then for any $\epsilon > 0$ we may assume P' has exactly one prevertex within $\epsilon \cdot d/10$ of each prevertex of P. This implies we can define a $1 + O(\epsilon)$-quasiconformal map Φ of the unit disk to itself that sends the prevertices of P' to the prevertices of P. Moreover, we may choose ϕ to be the identity near 0 and conformal except in the regions

$$A_x = \{z \in \mathbb{D} : \epsilon \cdot d/4 \leq |z - x| \leq d/4\}$$

around each prevertex x of P'. See Figure 14.

Let $F : P' \to P$ be the composition of the conformal map from P' to $\mathbb{D}$, followed by Φ, followed by the conformal map of $\mathbb{D}$ to P. This is a quasiconformal map from P' to P with dilatation bounded by $O(\epsilon)$ and mapping each vertex P' to a vertex of P with equal angle. Suppose $v = v' = 0$ are corresponding vertices in P and P', both with angle $\theta = k \cdot 60°$. Set $\alpha = 3/k$. Then $z \to z^\alpha$ maps the angle θ to 180°, so $G(z) = (F(z^\alpha))^{1/\alpha}$ maps a half-disk centered at the origin conformally to a region also bounded partly by a segment through the orgin. By the Schwarz reflection principle,

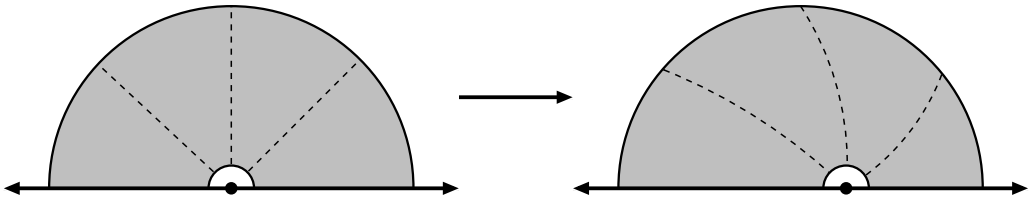

FIGURE 14. Pre-vertices of P and P' are paired and can be identified by a quasiconformal map Φ with small dilatation supported inside an annular region surrounding each pair.

G extends conformally to a disk around the origin, and hence setting $w = z^\alpha$, we have $F(w) = G(w^{1/\alpha})^\alpha$. This implies $F(w) = cw + O(w^2)$ near the origin, for some $c \neq 0$. Hence, after a linear rescaling, we can use Lemma 4.3 to replace the push forward triangulation under F by an equilateral triangulation in a neighborhood of the vertex v. Doing this for every vertex of P' completes the proof Lemma 5.1. □

We will need the estimate from the final part of the argument again later, so we record it as a lemma.

Lemma 5.2. *Suppose $0 < \theta \leq 360^\circ$ and that Ω_1, Ω_2 are simply connected domains such that $\Omega_1 \cap D(0,1) = \Omega_2 \cap D(0,1) = S(\theta) \cap D(0,1)$ and f is a conformal map $\Omega_1 \to \Omega_2$ so that $f(0) = 0$. Then $f(z) = cz + O(z^2)$ on $D(0, 1/2)$ for some $c \neq 0$.*

6. Triangulations of infinite sectors

As before, $S(\theta)$ denotes the infinite sector of angle θ with positive real half-line as its axis of symmetry. Note that $S(60^\circ)$ comes with a natural equilateral triangulation $\mathcal{G}$ as shown in Figure 15. This triangulation can obviously be extended to a triangulation of the right half-plane, and in this section, we record a computation that gives angle bounds images of this triangulation under conformal maps of the form $z \to z^\alpha$. See Figure 16 for two examples.

Lemma 6.1. *Consider the grid $\mathcal{G}$ of unit equilateral triangles in $S(60^\circ)$. Let $0 < \alpha \leq 2$. Suppose $T = \Delta ABC \in \mathcal{G}$ and $f^*T = \Delta f(A)f(B)f(C)$, where f is a branch of z^α defined on T. Then the interior angles of f^*T differ from the corresponding angles of T by at most $|\alpha - 1| \cdot \theta$ where θ is the angle subtended by T from the origin.*

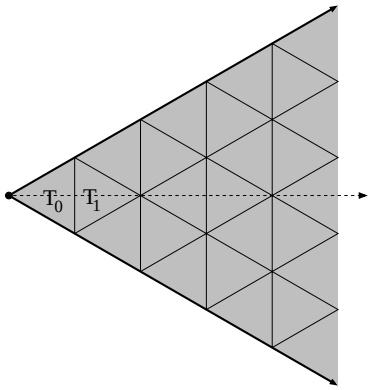

FIGURE 15. The sector $S(60°)$ and its equilateral triangulation $\mathcal{G}$.

Proof. Consider the angle $\angle f(A)f(B)f(C)$, i.e., the angle between the segments $\overline{f(A)f(B)}$ and $\overline{f(C)f(B)}$. By Rolle's theorem

$$\arg[f(B) - f(A)] = \arg[B - A] \cdot \arg f'(z),$$

for some $z \in \overline{f(A)f(B)}$. Thus the change in angle $|\angle ABC - \angle f(A)f(B)f(C)|$ is at most $|\arg(f'(z)) - \arg(f'(w))|$ where z, w are on $\overline{f(A)f(B)}$ and $\overline{f(C)f(B)}$ respectively. Since $\arg f'(z) = (\alpha - 1)\arg(z)$, the difference $|\arg(f'(z)) - \arg(f'(w))|$ for $z, w \in T$ is no bigger than $|\alpha - 1|$ times the angle subtended by T. □

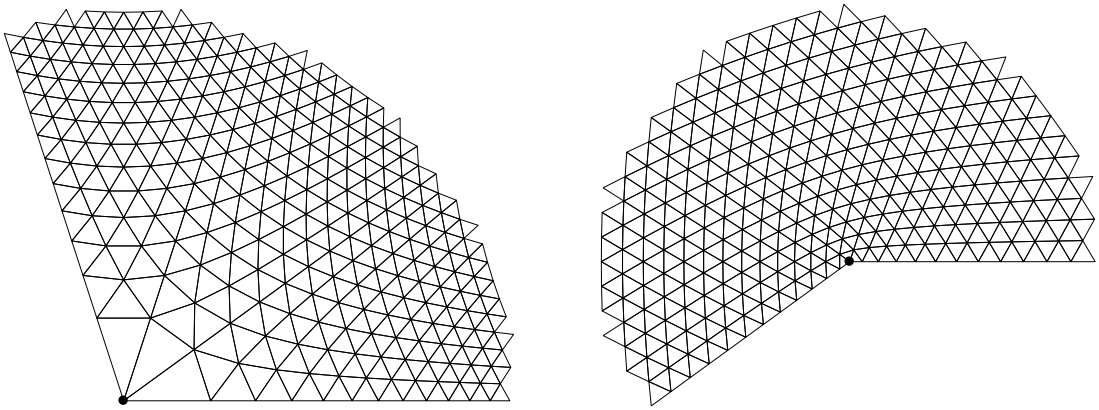

FIGURE 16. Image of the equilateral grid in the upper halfplane under the maps $z^{3/5}$ and $z^{6/5}$. These are the extreme values that keep images of 60° angles between 36° and 72°.

Corollary 6.2. *Suppose $0 < \phi < 90^\circ$. The sector $S(\phi)$ has a triangulation with all angles in $[180^\circ - 2\phi, \phi]$ if $\phi \geq 60^\circ$, and in $[\phi, 90^\circ - \phi/2]$ if $\phi < 60^\circ$.*

Proof. First suppose $\phi \geq 60^\circ$ and set $\alpha = \phi/60^\circ$. Consider the image of the triangular grid $\mathcal{G}$, defined above, under the power map $f(z) = z^\alpha$. Note that $\alpha > 1$ and hence $|\alpha - 1| = \alpha - 1$. Let $T_0 \in \mathcal{G}$ be the triangle containing the origin, and let T_1 be the unique triangle in $\mathcal{G}$ that shares an edge with T_0. All other triangles in $\mathcal{G}$ subtend angle $\leq 30^\circ$, so their angles can increase by at most $(\alpha - 1)30^\circ$, and so these angles are bounded by

$$60^\circ + (\alpha - 1)30^\circ = 30^\circ + \alpha \cdot 30^\circ = 30^\circ + \frac{\phi}{60^\circ} \cdot 30^\circ = 30^\circ + \phi/2 \leq \phi,$$

since $\phi \geq 60^\circ$. The pushed forward triangle f^*T_0 is isosceles with vertex angle $\phi = \alpha \cdot \theta$, so its angles are in $[90^\circ - \phi/2, \phi] \subset [180^\circ - 2\phi, \phi]$. T_1 is divided into two right triangles by the real axis (each with angles 30° and 60°) and each sub-triangle subtends angle $\leq 30^\circ$. Apply the previous lemma to these sub-triangles. The 30° angle has image angle ψ bounded above and below by

$$60^\circ - \phi/2 = 30^\circ - (\alpha - 1)30^\circ \leq \psi \leq 30^\circ + (\alpha - 1)30^\circ = \phi/2.$$

Doubling the upper bound gives $2\psi \leq \phi$ as the bound for the angle of f^*T_1 opposite the common edge with f^*T_0, as desired. The other two angles are equal by symmetry and bounded above by

$$\frac{1}{2}(180^\circ - 2\psi) = 90^\circ - \psi \leq 30^\circ + \phi/2 \leq \phi,$$

since $30^\circ = \frac{1}{2}60^\circ \leq \phi/2$.

For $\phi < 60^\circ$ we have $|\alpha - 1| = 1 - \alpha$, so the push forward f^*T of a general triangle $T \in \mathcal{G} \setminus \{T_0, T_1\}$ has angles bounded by

$$60^\circ + (1 - \alpha)30^\circ = 90^\circ - \phi/2 \leq 180^\circ - 2\phi,$$

as desired. Triangle T_0 works out just as before. We split T_1 in half as above. Since $\alpha < 1$ the argument of $f'(z) = \alpha z^{\alpha - 1}$ becomes more negative as we move away from the real axis along the hypotenuse; thus the hypotenuse maps to a concave down curve γ. By conformality, γ meets the real axis at 30°, and thus the segment connecting its endpoints meets the axis at angle $< 30^\circ$. Thus the corresponding angle ψ of f^*T_1 is $< 60^\circ$. By symmetry the angles at the other two vertices of f^*T_1 are equal and

bounded above by

$$\frac{1}{2}(180^\circ - 2\psi) \leq 90^\circ - (30^\circ - (1-\alpha)30^\circ) = 90^\circ - \phi/2. \quad \square$$

Note that if $T \in \mathcal{G}$ is at distance d from the origin, then it subtends angle $O(1/d)$ so f^*T has angles bounded by $60^\circ + O(1/d)$. Thus the pushed forward triangles converge to equilateral as we move away from the vertex of the sector. The following lemma allows us to modify a triangulation pushed forward by a conformal map between polygons to agree with a sector triangulation, as described above, near each vertex. This will be used to show that the bounds in Theorem 1.1 can be attained.

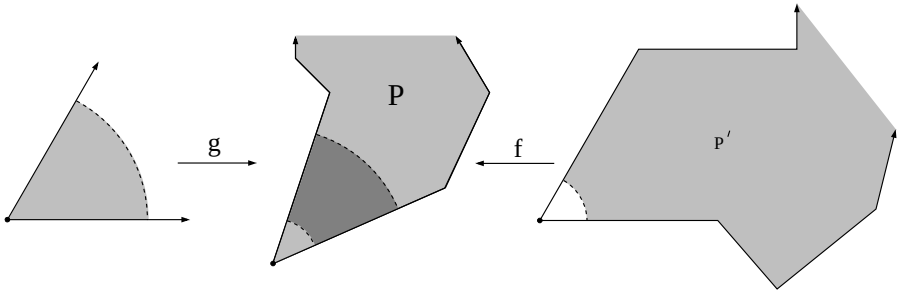

FIGURE 17. Proof of Lemma 6.3. Near each vertex v of P the conformal push forward of the nearly equilateral mesh on P' gives the correct angle bounds in the limit as the triangulation gets finer. To get the correct angle bounds at a positive scale we replace $f^*\mathcal{T}$ with a sector mesh, that we know has the correct angle bounds at all scales.

Lemma 6.3. *Suppose $f : P' \to P$ is a conformal map between polygons that maps vertices to vertices. Suppose $f(v') = v$ where v' is a vertex of P' and v is a vertex of P, with angles $\psi = k \cdot 60^\circ$ and θ respectively. Suppose $\mathcal{T}$ is a nearly equilateral triangulation of P' and $f^*\mathcal{T}$ the image triangulation. If $\mathcal{T}$ is fine enough, then there is a neighborhood U of v and a triangulation $\mathcal{S}$ of P that equals $f^*\mathcal{T}$ outside U and every triangle of $\mathcal{S}$ touching U has all angles bounded by* $\max(\theta/k, 90^\circ - \theta/2k)$.

Proof. We apply Lemma 4.3 where Ω_1 is the infinite sector with vertex v and angle θ that matches P in neighborhood of v, and $\Omega_2 = P$. By choosing an annular region $\lambda A(r) + v$ around v by first choosing r large and then λ small, as shown in Figure 17, and normalizing the sector map correctly, we can assume the two maps are close on $\lambda S(\theta, 3)$. The merged mesh agrees with the sector mesh close to v and with the image mesh $f^*\mathcal{T}$ away from v, as desired. $\square$

7. The 120°-trick

In this section we provide the details of the "120°-trick" for triangulating the upper half-plane in a way that uses maximum angle 72°, and near infinity looks like the push forward under $z^{3/2}$ of the standard equilateral mesh of a 120°-sector. This involves cutting a slit in P, as discussed in Sections 1 and 2.

Consider the region Ω shown on the left in Figure 10. This is a 120°-sector with an equilateral triangle at the origin removed. We translate the picture so the 300°-vertex is at the orgin. If we then apply a branch of $z^{6/5}$, the 300° angle becomes 360°, and the two finite segments in $\partial\Omega$ adjacent to it become identified with a radial slit in the image. The two rays in $\partial\Omega$ map to the boundary curve of a simply connected region Ω'. See lower left in Figure 18. By the Riemann mapping theorem, Ω' can be mapped to the upper half-plane, and the slit maps to a curved arc, meeting the real line at angle 60° (the slit looks quite straight since the tangents at the two endpoints differ by only $\approx 2.75°$).

Since the power map identifies points on the two segments adjacent to w that are equidistant from w, any equilateral triangulation of Ω will push forward to triangulation of the upper half-plane. If the triangulation is fine enough, then all the pushed forward triangles will be nearly equilateral, except near the corners and tip of the slit. However, near the point v where the slit joins the real line, P looks like the union of a small 60°-sector and a 120°-sector and the map to P' sends each of these to subdomains of P' that contain and are contained in small sectors of the same angles. Thus by Lemma 5.2, the conformal maps restricted to each finite sector is approximately linear and the image of the equilateral triangulation of P' is close to equilateral in P near v. This leaves only the tip of slit. In a small neighborhood of the tip, angles are bounded above by $72° + \epsilon$ if the mesh is fine enough, but might exceed 72°. However, we can replace the mesh in a neighborhood of the tip by the standard mesh of a 360°-sector using Lemma 4.4. This gives the 72° bound.

Lemma 7.1. *Suppose $f : P' \to P$ is a conformal map between polygons that maps vertices to vertices. Suppose $f(v') = v$ where v' is a vertex of P' and v is a vertex of P, with angles* 120° *and* 180° *respectively. Suppose $\mathcal{T}$ is a nearly equilateral triangulation of P' and $f^*\mathcal{T}$ the image triangulation. If $\mathcal{T}$ is fine enough, then there is a*

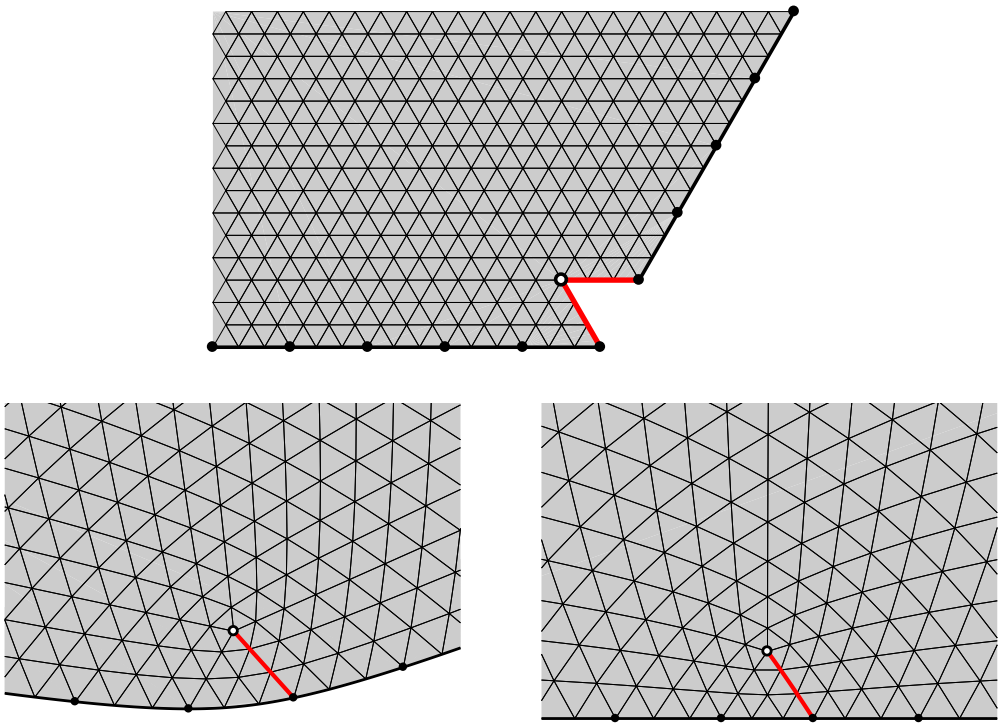

FIGURE 18. The cut 120-sector (top) is mapped to a simply connected region (lower left) by a branch of $z^{6/5}$. This is then mapped to the upper half-plane by a conformal map. Because the power map identifies points according to distance to zero (the white dot), the push forward of the triangulation is still a triangulation.

neighborhood U of v and a triangulation $\mathcal{S}$ of P that equals $f^*\mathcal{T}$ outside U and every triangle of $\mathcal{S}$ touching U has all angles $\leq 72°$.

Proof. The proof is the same as for Lemma 6.3, except that now we use the mesh coming from the 120°-trick. See Figure 19. □

8. The 420°-trick

There are two things we can do to increase the ψ-sum for P' by 60° with respect the θ-sum for P. The first is to introduce a 180°-vertex v in an edge of P and add a corresponding 240°-vertex v' to P'. This clearly increases the ψ-sum by an extra 60° relative to the θ-sum. The angle at v' is subdivided into four equilateral triangles by the nearly equilateral triangulation, and each of these are mapped to four angles

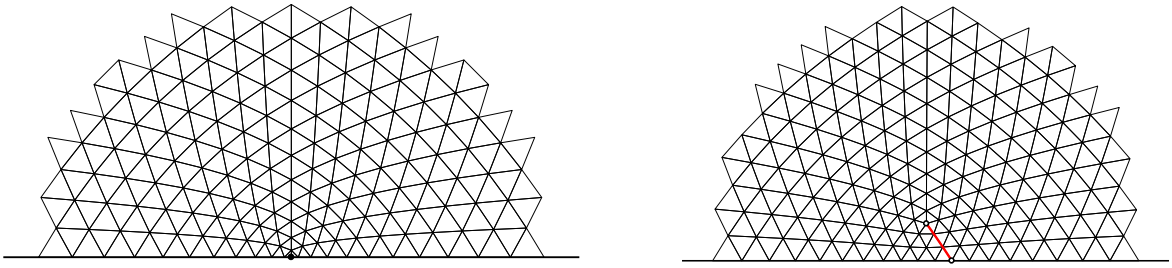

FIGURE 19. The left shows the equilateral triangulation of a 120°-sector pushed forward to the half-plane by $z^{3/2}$. The right shows the triangulation coming from the "120°-trick". These two meshes can be merged using quasiconformal interpolation as described in the text.

of size 45° at v. The opposite angles in the image triangles are $67.5° < 72°$, so this construction will be enough for proving Case 1 of Theorem 1.1.

However, in order to handle Case 2 of Theorem 1.1 we need another "trick" that can add 60° to the ψ-sum relative to the θ-sum, but introduces triangulation angles no larger than $\frac{5}{7} \cdot 90° \approx 64.2857°$. This is precisely the angle bound we get if a 420°-vertex $v' \in P'$ is mapped to a 360°-vertex $v \in P$. The 360° vertex v can occur as the end vertex of a slit in P, but how do we get a 420°-vertex in P'? We do this by considering a non-planar Riemann surface.

The idea is illustrated in Figure 20. Consider the two planar regions shown on the left side of the figure and define a Riemann surface by identifying them along the dashed ray. This creates a simply connected Riemann surface R with single boundary curve that is the union of two infinite rays, two finite segments and has three corners of 60°, 420° and 120°.

We can conformally map R to a slit upper half-plane in two steps as illustrated in Figure 21 so that the two segments of ∂R that are adjacent to the 420° angle are identified with the slit, and length measure on these segments is pushed forward to the same measure on the slit. Translate the 420°-vertex to the origin and apply a branch of $z^{6/7}$ defined on R. This maps R to a simply connected planar domain Ω with a straight slit; the two segments of ∂R are identified with this slit in the correct way, and the two infinite rays are mapped to disjoint, unbounded arcs on $\partial\Omega$. The domain Ω can then be conformally mapped to a half-plane. Thus equilateral triangulations of R will be mapped to triangulations of the upper half-plane.

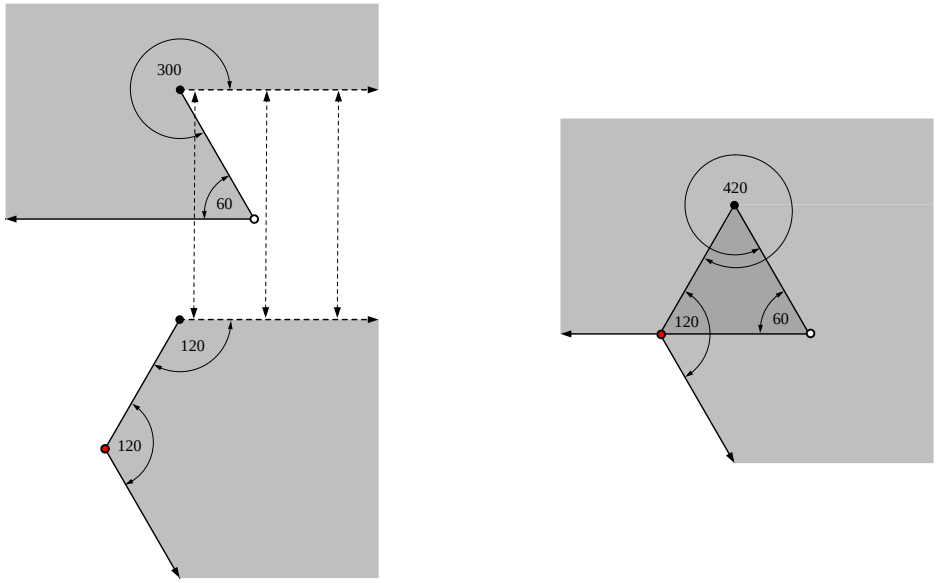

FIGURE 20. Here a 180° vertex in P corresponds to a 420° in P'. This is obtained by making P' a Riemann surface instead of a planar domain. The surface can be constructed from two planar domains glued along the dashed edges of each as illustrated on the left. The darker triangle indicates where the surface has two sheets over the plane.

9. Necessity in Theorem 1.1

Proof. For the sake of completeness, we include the details of the proof of necessity that was sketched preceding the statement of Theorem 1.1. We also restate our conditions in the form conjectured by Gerver in [24].

Fix $60° < \phi < 90°$ and suppose $\mathcal{T}$ is a ϕ-triangulation of P. Let V_P be the vertex set of P and $|V_P|$ the number of points in V_P (in general we let $|X|$ denote the number of elements in a set X). As before, for $v \in V_P$, let $L(v)$ denote the number of triangles containing v. Let F (for faces) be the number of triangles in $\mathcal{T}$ and E the number of edges. Let $V_{\mathcal{T}}$ be the vertices of the triangulation $\mathcal{T}$. If an interior angle θ of P is subdivided into k sub-angles that all lie in $I(\phi)$, then $\theta \in k \cdot I(\phi)$. Thus an ϕ-triangulation of P gives an admissible labeling of V_P. This proves necessity in Case 1. Moreover, if $\phi \geq 72°$, then (see Figure 7)

$$\bigcup_{k \geq 1} k \cdot I(\phi) = [180° - 2\phi, \infty),$$

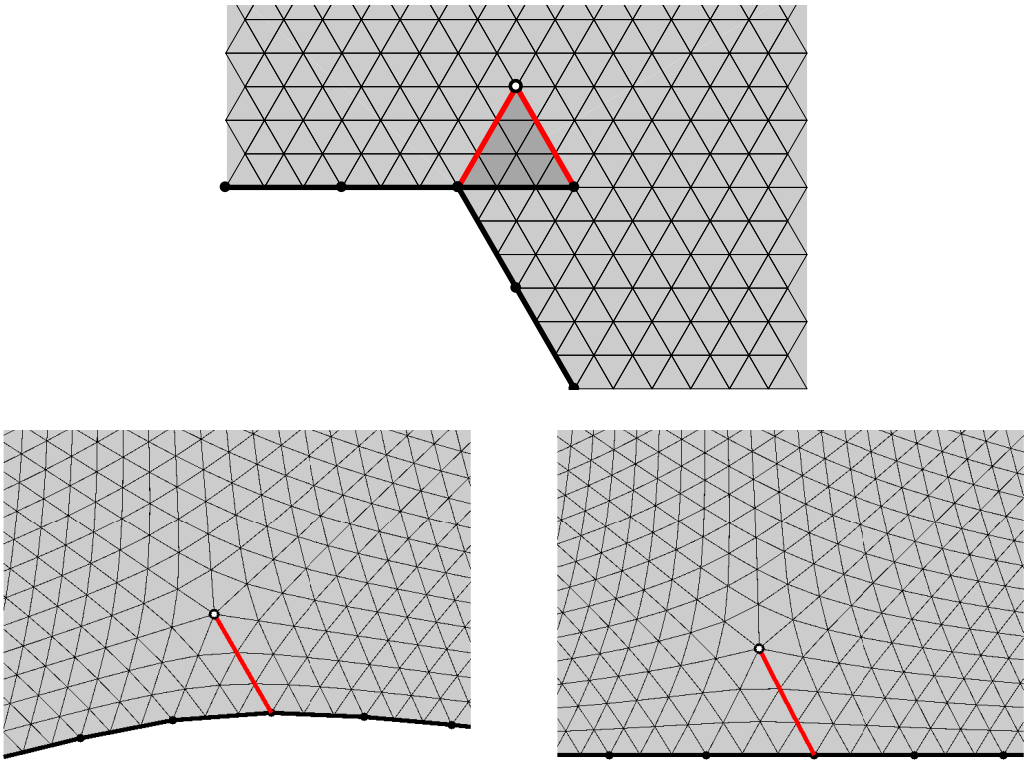

FIGURE 21. The Riemann surface R can be conformally mapped to a slit domain Ω by a branch of $re^{i\theta} \to re^{i6\theta/7}$. This map is followed by a conformal map to the upper half-plane that bends the the straight slit to an analytic arc. (but it looks quite straight; the tangent directions only change by about 1° along the arc).

so having a ϕ-admissible labeling is equivalent to $\theta_{\min} \geq 180^\circ - 2\phi$. Thus Corollary 1.3 and Case 1 in Theorem 1.1 are equivalent.

Let $V_{\mathcal{T}}(k) = \{v \in V_{\mathcal{T}} : L(v) = k\}$, i.e., the vertices that are in k triangles. Set

- $r_k = |V_{\mathcal{T}}(k) \cap V_P|$ (vertices of P in k triangles),
- $q_k = |V_{\mathcal{T}}(k) \setminus P|$ (interior vertices of $\mathcal{T}$ in k triangles),
- $s_k = |V_{\mathcal{T}}(k) \cap (P \setminus V_P)|$ (boundary vertices of $\mathcal{T}$ that are not vertices of P).

With this notation, $|V_P| = \sum_k r_k$, $|\text{int}(\mathcal{T})| = \sum_k q_k$, and $|\partial\mathcal{T}| = \sum_k s_k + \sum_k r_k$. Given L as above, we have $\sum_k kr_k = \sum_{v \in V_P} L(v)$ and hence

$$\sum_k (3-k) r_k = 3|V_P| - \sum_v L(v) = 6 - \kappa(L).$$

Therefore $\kappa(L) \leq 0$ if and only if $\sum_k (3-k) r_k \geq 6$, and equality holds simultaneously. The latter is the notation used by Gerver in [24].

We have the relations

- $F - E + |V_{\mathcal{T}}| = 1$ (Euler's formula),
- $|V_{\mathcal{T}}| = \sum_k q_k + \sum_k r_k + \sum_k s_k$ (every vertex is in V_P, $P \setminus V_P$ or $V_{\mathcal{T}} \setminus P$),
- $3F = \sum_k k q_k + \sum_k k r_k + \sum_k k s_k$ (triangle corners counted in two ways),
- $2E = 3F - \sum_k r_k - \sum_k s_k$ (triangles sides counted in two ways),

Combining these four equations and eliminating F, $|V_P|$, E and $|V_{\mathcal{T}}|$ we get

$$\sum_k (6-k) q_k + \sum_k (3-k) r_k + \sum_k (3-k) s_k = 6, \tag{9.1}$$

in Gerver's notation, or equivalently using discrete curvatures:

$$\sum_{v \in \mathrm{int}(\mathcal{T})} \kappa(v) + \sum_{v \in V_P} \kappa(v) + \sum_{v \in \partial\mathcal{T} \setminus V_P} \kappa(v) = 6. \tag{9.2}$$

This is the discrete Gauss-Bonnet formula, (1.1). For acute triangulations we have $s_1 = s_2 = 0$, and hence the third term in (9.1) and (9.2) is non-positive, i.e., $\sum_{v \in \mathrm{int}(\mathcal{T})} \kappa(v) + \sum_{v \in V_P} \kappa(v) \geq 6$.

In Case 2 of Theorem 1.1, $\phi < 72°$, so every interior vertex has degree ≥ 6. This implies $q_1 = \cdots = q_5 = 0$ and thus the first term in (9.2) is also non-positive, i.e., $\kappa(L) = 6 - \sum_{v \in V_P} \kappa(v) \leq 0$, as desired.

In Case 3, $60° < \phi < \frac{5}{7}90°$, so every interior vertex of $\mathcal{T}$ has degree six and every vertex in $\partial\mathcal{T} \setminus V_P$ has degree three. Thus $\sum_{v \in \mathrm{int}(\mathcal{T})} \kappa(v) = \sum_{\partial\mathcal{T} \setminus V_P} \kappa(v) = 0$ and so (9.2) implies $\kappa(L) = 0$, as desired. □

10. Sufficiency in Theorem 1.1: Case 1

Proof. We want to show that any polygon P (possibly after adding extra vertices) can be conformally mapped to a 60°-polygon P', with the restrictions on the angles given by Table 1. We will use the 120°-trick to "fix" the pushed forward triangulation in a neighborhood of a few boundary points, but the 420°-trick is not needed until the proof of Cases 2 and 3 in the next section.

The Schwarz-Christoffel formula gives a conformal map f of the disk onto a polygonal region P in terms of two types of data. First are the angles of P: suppose

$V_P = \{v_j\}_1^N$ are the vertices of P and the interior angle at v_j is $\alpha_j \cdot 180°$. Second, suppose f maps $z_j \in \mathbb{T}$ to $v_j \in P$; these points are called the prevertices or Schwarz-Christoffel parameters of f. Then the conformal map f is given by

$$f(z) = A + C \int^z \prod_{j=1}^{N} (1 - \frac{w}{z_j})^{\alpha_j - 1} dw, \tag{10.1}$$

for some appropriate choice of constants A, C. See e.g., [17], [37], [43]. The formula was discovered independently by Christoffel in 1867 [14] and Schwarz in 1869 [42], [41]. For other references and a brief history, see Section 1.2 of [17]. Given a polygon P, the angles are known, but the prevertices must be solved for.

Given N distinct points $\mathbf{z} = \{z_1, \ldots, z_N\}$ on the unit circle and N real values $\{\alpha_1, \ldots, \alpha_N\}$ summing to $N - 2$, Formula (10.1) defines a locally 1-1 holomorphic function on the disk that maps each component of $\mathbb{T} \setminus \mathbf{z}$ to line segment, with the segments meeting at $f(z_k)$ making interior angle $\alpha_k \cdot 180°$. The map given by (10.1) is always locally 1-1 on $\mathbb{D}$, but need not be globally 1-1 in general. In this case, the image is a Riemann surface with an obvious projection onto the plane. For the proof of Case 1 of Theorem 1.1 we can arrange for the image to be a planar 60°-polygon, although in the proof of Cases 2 and 3, given in the next section, the image is allowed to be a non-planar 60°-surface (this occurs when we apply the 420° trick).

Given an N-gon P, we take some conformal map f of its interior to the unit disk, $\mathbb{D}$. The N vertices of P map to N distinct points $\mathbf{z} = \{z_1, \ldots, z_N\}$ on the unit circle $\mathbb{T}$. We then want to choose N real values $\psi_k \in Z = \{60°, 120°, 180°, 240°, 300°\}$ so that $\sum_k \psi_k = 180(N-2)$. If this is possible, we then set $\alpha_k = \psi_k/180$ and apply the Schwarz-Christoffel formula to get a map $g : \mathbb{D} \to P'$. Then $g \circ f : P \to P'$ is the desired map. However, as noted in Sections 1 and 2, such a choice of angles ψ_k may not be possible without adding extra vertices to P.

First choose six interior points of some edge of P. This creates an M-gon with $M = N + 6$. These are 180°-vertices in P and are assigned to have angle $\psi_v = 120°$ in P'. Assign angle 180° to every other vertex of P', so the ψ-angle sum is $6 \cdot 120° + 180°N = (M-2)180°$. Applying Schwarz-Christoffel gives a a 60°-hexagon, as in Figure 22.

Next we modify the angle assignments to get a P' that approximates this hexagon. Let $L : V_P \to \mathbb{N}$ be a ϕ-admissible labeling of the vertices of P. For $v \in V_P$, assign

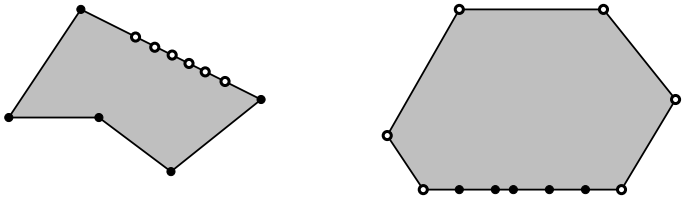

FIGURE 22. In the proof of Case 1 of Theorem 1.1 we can assume P' is planar. The first step is to choose six "artificial" vertices on one edge of P and make these correspond to six 120°-vertices in P'.

angle $60^\circ \cdot L(v)$ to the corresponding vertex v' of P'. In order to adjust the angle sums, for each vertex v of P we define either 0, 1 or 2 "associated vertices". The new vertices will be in the edge of P that begins with v (P has the counterclockwise orientation with the domain interior on the left) and may be taken as close to v as we wish. The vertices associated to v on P have angle 180° and the corresponding vertices associated to v' in P' have angle either 120° or 240°. These angles are assigned so that P' leaves the last associated vertex in the same direction as it entered v'. This implies that the part of P' near v' approximates a straight line. The rules for making the assignments are simple and illustrated in Figure 23. Suppose v is an original vertex of P with interior angle θ_v:

(i) if $0 < \theta \leq 72^\circ$, set $\psi_v = 60^\circ$ and add two vertices each with angle 240°,
(ii) if $72^\circ < \theta \leq 144^\circ$, set $\psi_v = 120^\circ$ and add one vertex with angle 240°,
(iii) if $144^\circ < \theta \leq 216^\circ$, set $\psi_v = 180^\circ$ and add no associated vertices,
(iv) if $216^\circ < \theta \leq 288^\circ$, set $\psi_v = 240^\circ$ and add one vertex with angle 120°,
(v) if $288^\circ < \theta \leq 360^\circ$, set $\psi_v = 300^\circ$ and add two vertices of angle 120°.

If the vertices associated to v are close enough to v, then the image arc is close to a line. See Figure 24. In particular, P' is not self-intersecting and so is a 60°-polygon. By Lemma 5.1, P' has nearly equilateral triangulations. In the remainder of the proof we will take this triangulation as fine as is needed (but only finitely many conditions are involved, so we finish with a positive grid size).

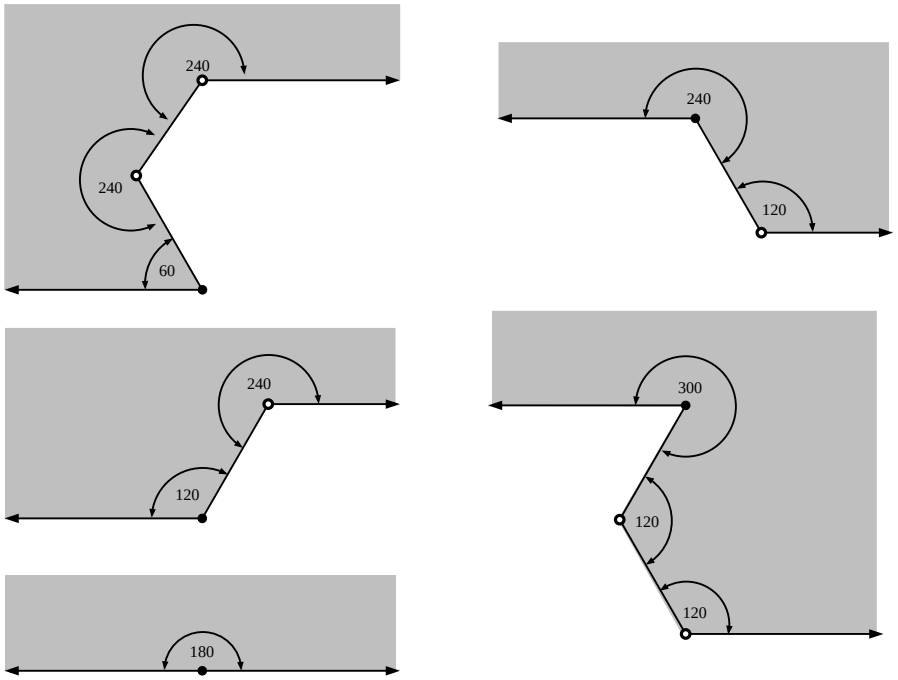

FIGURE 23. The five cases for assigning image angles and associated vertices. In each case the black indicates the image of the original vertex and the white dots the new associated vertices. These arcs illustrate small subarcs of the 60°-polygon P'. The associated vertices map to 180° vertices that we add to the edges of P.

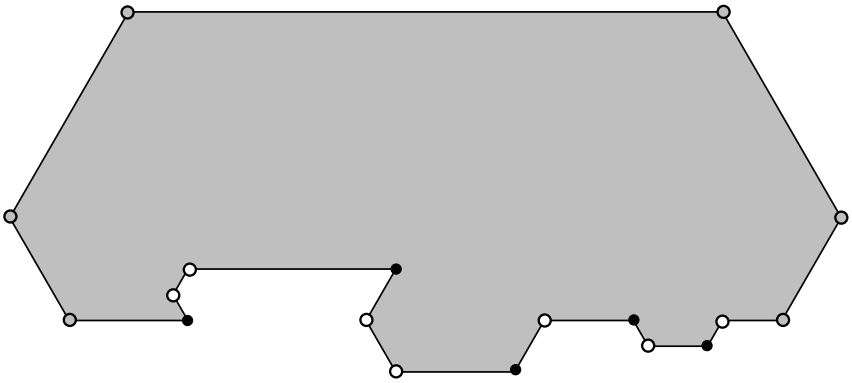

FIGURE 24. The 60°-polygon P'. The six gray points were added at the beginning, the black points are the images of the original vertices and the white points are the associated vertices.

Each original vertex with angle $\theta_v \geq 36^\circ$ was assigned an image angle ψ_v in the allowable range from Table 1. Thus transferring a nearly equilateral triangulation of P' gives a triangulation of P with all angles between 36° and 72° except possibly in small neighborhoods of these vertices, where the angle bounds are with all angles between $36^\circ - \epsilon$ and $72^\circ + \epsilon$, and ϵ can be made as small as we wish be taking the triangulation fine enough. In a neighborhood of each such vertex, we may use Lemma 6.3 to replace the pushed forward triangulation with a sector triangulation for which the bounds $36^\circ, 72^\circ$ hold.

For each original vertex with interior angle $\theta_v < 36^\circ$ the same argument applies, except that now we get the bounds in the interval $I(\theta_v) = [\theta_v, 90^\circ - \theta_v/2]$. Again, we may use interpolation to locally replace the pulled back triangulation (which might only approach the desired bounds as the triangulation gets finer), with a sector triangulation satisfying the desired bounds.

Next consider the associated vertices with angles $> 120^\circ$; Cases (i) and (ii) above. In Case (i) each 240°-vertex is hit by 4 equilateral triangles and so each 60° sub-angle is mapped to an angle of size $180^\circ/4 = 45^\circ$. In this case, the interpolation with a sector triangulation isn't needed; with small enough distortion, the angles are already inside $[36^\circ, 72^\circ]$. Case (ii) is the same, except there is only one associated vertex.

Finally we consider Cases (iv) and (v). Here we only use image angles of size 120°. Such an angle is divided into two 60° angles that are mapped to 90° by the conformal map: too large. We use Lemma 7.1 to interpolate between the conformal image triangulation and the triangulation of the half-plane coming from the 120°-trick. This gives a triangulation of with angles in $I(36^\circ) = [36^\circ, 72^\circ]$ in a neighborhood of each associated vertex. This completes the proof of Case 1 of Theorem 1.1. □

Note that we have also proven Corollary 1.3. The upper bound in Corollary 1.3 implies every angle in the triangulation is at least $\min(\theta_{\min}, 36^\circ)$. However, the proof of Theorem 1.3 actually proves the following, slightly stronger, lower bound:

$$\max\left(\min(\theta_{\min}, 36^\circ), \min(\theta_{\min}/2, 48^\circ), \min(\theta_{\min}/3, 54^\circ)\right).$$

This formula is somewhat clearer when graphed, as in Figure 25. The first term follows directly from Corollary 1.3. The second term holds because any angle $72^\circ \leq \theta \leq 96^\circ$ can be divided into two angles of size $\theta/2$ and larger angles can be divided into two or more angles $\geq 48^\circ$. Similarly, the third term arises because any angle

$144^\circ \le \theta \le 162^\circ$ can be divided into three angles of size $\theta/3$ and larger angles can be divided into three or more angles $\ge 54^\circ$.

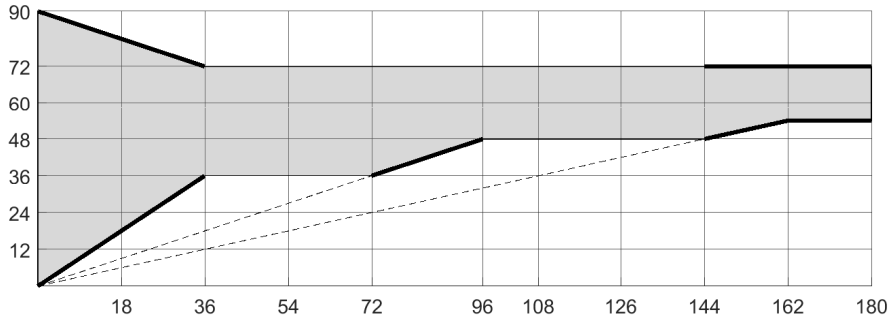

FIGURE 25. If P has minimal angle $0 < \theta_{\min} < 180^\circ$ (plotted on the horizontal axis), then P has a triangulation with all angles inside the interval above $\theta_{\min}$. For example, if $\theta_{\min} = 108^\circ$ the interval is $[48^\circ, 72^\circ]$. Bold lines indicate values that must be attained by any upper-optimal triangulation, e.g., if $\theta_{\min} = 80^\circ$ implies 40° is attained by any 72°-triangulation of P.

11. Sufficiency in Theorem 1.1: Cases 2 and 3

Proof. We have already proven sufficiency in Case 1. To prove it in the other two cases, we just have to modify the construction in the proof of Case 1 to avoid using the 120°-trick in Case 2 (this forces an angle $\ge 72^\circ$), and to avoid both the 120°-trick and the 420°-trick in Case 3 (the latter forces angles $\ge \frac{5}{7} \cdot 90^\circ$).

Suppose L is a ϕ-admissible labeling of V_P so that $\kappa(\phi) = \kappa(L)$ (i.e., choose L to minimize $|\kappa(L)|$ among admissible labelings). As before, let θ_v denote the angle of P at vertex v, and for each vertex v in P, suppose $\psi_v = L(v) \cdot 60^\circ$ is the tentative corresponding angle of v' in P'. As we have noted before,

$$\begin{aligned} \sum_v \theta_v &= (|V_P| - 2)180^\circ = 60^\circ(3|V_P| - 6) \\ &= 60^\circ \left(3|V_P| - 6 - \sum_v L(v) + \sum_v L(v) \right) = 60^\circ \cdot \kappa(L) + \sum_v \psi_v. \end{aligned}$$

Thus in Case 2 ($\kappa(L) \le 0$) we only need to introduce 180°-vertices on P that correspond to 240°-vertices in P'. If $\phi \ge 67.5^\circ$, we can do this by replacing the pushed forward triangulation from P' by the 240°-sector triangulation. If $\frac{5}{7} \cdot 90^\circ \le \phi < 67.5^\circ$

then we replace it with the triangulation of the half-plane obtained by the 420°-trick. This proves sufficiency in Case 2.

Finally, if $\kappa(L) = 0$, then no extra vertices or "tricks" are needed. We simply use a fine enough triangulation pushed forward by the conformal map from P' to P and replace it in a neighborhood of each vertex by the appropriate sector mesh. □

Proof of Corollary 1.4. By construction, all the interior vertices of our triangulation have degree 6, except for the single degree 7 vertex used in each application of the "420°-trick" and the vertex of degree 5 used in the "120°-trick". The first is only used in Case 2 of Theorem 1.1, and, indeed, is only needed for angles $< 67.5°$, since using a 240° sector mesh will work for larger angles. The number of times we need to apply the 420°-trick is $(\sum_P \theta_v - \sum_{P'} \psi_v)/60° = -\kappa(\phi)$, so this is the number of degree 7 vertices.

One degree 5 vertex is created in each application of the 120°-trick. This occurs only in Case 1 and only if $\sum_P \theta_v < \sum_{P'} \psi_v$; otherwise we can use 240°-vertices in P' to make up the gap. Thus $\max(0, \kappa(\phi))$ degree 5 vertices are created. □

Proof of Corollary 1.9. As noted earlier, Gerver proved the conditions in Theorem 1.1 are implied by the weaker condition that for every $\epsilon > 0$, the polygon P has a $(\phi+\epsilon)$-dissection. He defines q_k to be the number of interior vertices of the triangulation that are shared with k triangles and are not interior to any edges of the triangulation. Also s_k is the number of vertices of the triangulation that are shared by k triangles and that lie on the interior of an edge of P or of an edge of some triangle. The proof then proceeds as before. See [24] for details. □

12. Some examples

Proof of Corollary 1.5. We compute the optimum upper bound Φ_N for triangulating the regular N-gon. The cases $N = 3, 6$ are trivial; these have equilateral triangulations so $\Phi_3 = \Phi_6 = 60°$. Otherwise, there are N interior angles of size $\theta_N = 180° - 360°/N \geq 90°$. Since these are all $> 36°$, Corollary 1.3 says $\Phi_N \leq 72°$, so we only have to check whether an even smaller bound is possible.

For $N \geq 10$, the interior angles are $\geq 144°$, and hence the ϕ-admissible labels for $\phi < 72°$ are at least 3. Then any ϕ-admissible label satisfies $\kappa(L) \geq 6 > 0$, so we must be in Case 1, so $\Phi_N = 72°$. (This case also follows from Corollary 1.7).

For $N = 9$, the interior angle is $\theta_9 = 140°$. If $\phi < 70°$ then $\theta_9 \notin 2 \cdot I(\phi)$, so all the ϕ-admissible labels are ≥ 3, which implies $\kappa(L) \geq 6$; thus such ϕ's don't work. However, for $\phi = 70°$ we can take six labels equal to 2 and three equal to 3; this labeling satisfies $\kappa(L) = 6 - 6(3-2) - 3(3-3) = 0$. Thus Case 2 holds for $\Phi_9 = 70°$ and this must be the sharp upper bound.

For $N = 8$, the same argument holds by considering $\theta_8 = 135°$ and $\phi = 67.5°$. Below this value, admissible labels have positive curvature, but for this value we obtain a zero curvature label by taking six labels 2 and two of 3. Thus $\Phi_8 = 67.5°$.

For $N = 7$, we have $\theta_7 = \frac{5}{7} \cdot 180° \approx 128.5714$, and apply the argument to $\Phi_7 = \frac{5}{7} \cdot 90° \approx 64.2857$. The only Φ_7-admissible label for θ_7 is 2, and the curvature of this labeling is $6 - 7(3-2) = -1$. Therefore there is a Φ_7-triangulation, and this is sharp since there are no zero curvature labelings for smaller ϕ's.

For $N = 4$ the internal angles are 90° so the only ϕ-admissible labels for $\phi \leq 72°$ are ≥ 2, and hence any such ϕ-admissible label L satisfies $\kappa(L) \geq 6 - 4 \cdot (3-2) = 2 > 0$. Thus Cases 2 and 3 of Theorem 1.1 can't hold and we must be in Case 1, i.e., $\Phi_4 \geq 72°$. By our remark at the beginning of the proof, we must have equality. Exactly the same argument works for $N = 5$, but using $\kappa(L) \geq 6 - 5 \cdot (3-2) = 1 > 0$. □

Proof of Corollary 1.6. By Corollary 1.3, $\phi(P) \leq 72°$. An axis-parallel polygon only has angles of 90°, 180° or 270°, so suppose P has n_{90}, n_{180} and n_{270} angles of each size. Then $|V_P| = n_{90} + n_{180} + n_{270}$ and the angle sum formula for polygons implies $n_{270} = n_{90} - 4$. Thus $2n_{270} = |V_P| - 4 - n_{180}$. Any ϕ-admissible labeling for $\phi < 90°$ gives the 90°-vertices label 2, the 180°-vertices labels ≥ 3 and the 270°-vertices labels ≥ 4. Thus the smallest admissible label sum is

$$2n_{90} + 3n_{180} + 4n_{270} = 2n_{90} + 3n_{180} + |V_p| - 4 - n_{180} + 2n_{270} = 3|V_P| - 4,$$

so $\kappa(L) \geq 2$ for any $\phi < 90°$ and any ϕ-admissible labeling L. Thus we must be in Case 1 and 72° is the sharp angle bound. Moreover, there is a 72°-admissible labeling (the minimal one) with $\kappa(L) = 2$, so we can apply the 120°-trick twice and construct a triangulation with exactly two vertices of degree 5 (the rest are degree 6). □

The proof above still holds if the angles are perturbed slightly (the bounds on the labels don't change), so the axis-parallel N-gons are in the interior of the set of all N-gons with sharp bound 72°, e.g., Corollary 1.8.

Proof of Corollary 1.7. If all angles of P are $\geq 144°$, then for any $\phi < 72°$, any ϕ-admissible labeling is at least 3 at every vertex, and hence has curvature ≥ 6. By Corollary 1.3, $72°$ is therefore the sharp angle bound. If the minimum angle of P is $\geq 162°$, then every angle of P can be split into three or more angles between $54°$ and $72°$ and the triangulation only uses interior vertices of degree 5 and 6. Thus all angles in the triangulation can be taken in $[54°, 72°]$. If we also assume that all angles are $\leq 216°$, then taking all labels equal to 3 is $72°$-admissible, and has curvature 6, so we need only six applications of the $120°$-trick. □

13. Computing $\Phi(P)$ in linear time

Proof of Corollary 1.2. First check whether every angle of P is a multiple of $60°$. If so, then $\Phi(P) = 60°$. If not, find the smallest angle $\theta_{\min}$ of P. If $\theta_{\min} \leq 36°$, then $\Phi(P) = 90 - \theta_{\min}/2$ by Corollary 1.3. If $\theta_{\min} \geq 144°$ then $\Phi(P) = 72°$ by Corollary 1.7. All this can be done in time $O(|V_P|)$.

Otherwise we may assume $36° < \theta_{\min} < 144°$, and we claim that computing $\Phi(P)$ reduces to finding

$$\phi_\infty = \inf\{\phi \in [60°, 90°] : \kappa(\phi) < \infty\},$$

$$\phi_0 = \inf\{\phi \in [60°, 90°] : \kappa(\phi) = 0\}.$$

Since $I(90°) = [0°, 90°]$, 4 is a $90°$-admissible label for any vertex of any polygon, so we always have $\kappa(90°) < \infty$. This implies ϕ_∞ is well defined. Computing ϕ_∞ in time $O(|V_P|)$ is easy: for each $v \in V_P$, we find the smallest ϕ so that v has a ϕ-admissible label (time $O(1)$ per angle) and then take the maximum of these $|V_P|$ values.

Increasing an $90°$-admissible label gives another $90°$-admissible label, so $\mathcal{K}(90°)$ is a half-infinite interval of the form $\{k \geq k_0\}$. Thus $\kappa(90°) \geq 0$ for any polygon. If $\kappa(90°) > 0$, then $\kappa(\phi) > 0$ for all $\phi < 90°$ and hence $\Phi(P) = \max(72°, \phi_\infty)$.

We may therefore assume that $\kappa(90°) = 0$. Thus ϕ_0 is well defined and $\phi_0 \leq 90°$. If $\kappa(\phi_\infty) = 0$, then $\phi_\infty = \phi_0$ and $\Phi(P) = \phi_\infty$. So assume $\kappa(\phi_\infty) \neq 0$; thus $\phi_\infty < \phi_0$. If $\kappa(\phi_\infty) > 0$, then $\kappa(\phi)$ is a decreasing, non-negative function of ϕ and

- if $\phi_\infty \geq 72°$, then $\Phi(P) = \phi_\infty$;
- if $\phi_\infty < 72°$ and $\phi_0 \geq 72°$, then $\Phi(P) = 72°$;
- if $\phi_0 < 72°$, then $\Phi(P) = \phi_0$.

Otherwise, $\kappa(\phi_\infty) < 0$. Then $\kappa(\phi)$ is non-positive and increasing, and we have:

• if $\phi_\infty \geq \frac{5}{7} \cdot 90°$, then $\Phi(P) = \phi_\infty$;
• if $\phi_\infty < \frac{5}{7} \cdot 90°$ and $\phi_0 \geq \frac{5}{7} \cdot 90°$, then $\Phi(P) = \frac{5}{7} \cdot 90°$;
• if $\phi_0 < \frac{5}{7} \cdot 90°$, then $\Phi(P) = \phi_0$.

Compute $\kappa(72°)$. If this is non-zero, then $\phi_0 > 72°$. In every case where this holds (the first two bullets in each triple), $\Phi(P)$ is determined from ϕ_∞ alone.

We have now reduced finding $\Phi(P)$ to computing ϕ_0, assuming that $\kappa(72°) = 0$, and hence that $\phi_0 \leq 72°$. Note that $\kappa(\phi) = \min(\kappa^+(\phi), 0) + \max(\kappa^-(\phi), 0)$, where

$$\kappa^-(\phi) = -3|V_P| + 6 + \sum_v \inf\{k \in \mathbb{N} : \theta_v \in k \cdot I(\phi)\},$$

$$\kappa^+(\phi) = -3|V_P| + 6 + \sum_v \sup\{k \in \mathbb{N} : \theta_v \in k \cdot I(\phi)\},$$

are the smallest and largest elements of $\mathcal{K}(\phi)$. For a single value of ϕ, each of these can be computed in time $O(|V_P|)$, hence so can $\kappa(\phi)$. Moreover, given each angle θ_v of P, we can compute the values of $\phi \in (60°, 72°]$ where θ_v lies on the boundary of one of the triangles in Figure 7. There are 10 possible triangles, so at most 10 possible ϕ's for each θ_v. In fact, Figure 7 shows at most 5 triangles can be hit by a single horizontal line of height $\theta \in [0, 360°]$ over the interval $[60°, 72°]$. This gives a set $X_1 \subset (60°, 72°]$ of size $\leq 5|V_P|$ that contains all possible ϕ values where κ^- and κ^+ can have jumps. If we were to sort X_1, then we could compute these functions at every jump point by summing jumps from left to right, and then find ϕ_0 by searching for the smallest ϕ value so that $\kappa(\phi) = 0$. This takes time $O(|V_P| \log |V_P|)$ due to the sorting, but ϕ_0 can be computed even faster.

Recall that we are now assuming $\kappa(\phi_\infty) \neq 0$ and $\kappa(72°) = 0$. Find the median value ϕ_1 of X_1; this can be done in time $O(|X_1|)$ by the median-of-medians algorithm in [11]. (See also [2] for some history and updates of this algorithm.) Compute $\kappa(\phi_1)$ in time $O(|X_1|)$. If $\kappa(\phi_1) = 0$, then $\phi_0 \in X_1 \cap [\phi_\infty, \phi_1]$ and otherwise $\phi_0 \in X_1 \cap [\phi_1, 72°]$. In either case, we now know that ϕ_0 is contained in a subinterval X_2 of X_1 with at most $\frac{1}{2}|X_1| + 1$ elements. Thus $|X_2| \leq \frac{3}{4}|X_1|$ if $|X_1| \geq 4$. We can construct X_2 in time $O(|X_1|)$ by comparing each element of X_1 to ϕ_1.

We inductively create intervals $X_1 \supset X_2 \supset X_3 \supset \dots$ containing ϕ_0 (by induction, κ will be non-zero at the leftmost point of X_n and zero at the rightmost point). Given X_n, we find its median ϕ_n in time $O(|X_n|)$ by the median-of-medians algorithm. We can compute $\kappa(\phi_n)$ in time $O(|X_n|)$ because we already know κ at the leftmost

endpoint of X_n and we can sum the jumps between this point and ϕ_n by examining each element of X_n. If $\kappa(\phi_n) = 0$ we let X_{n+1} be the elements of X_n that are $\leq \phi_n$, and otherwise the elements of X_n that are $\geq \phi_n$. Clearly X_{n+1} has at most $\frac{1}{2}|X_n|+1 \leq \frac{3}{4}|X_n|$ elements and takes $O(|X_n|)$ work to compute. Stop when $|X_n| \leq 4$, and compute κ at all these values to find ϕ_0. The total time taken is

$$O(|X_1| + |X_2| + \dots) = O(|X_1|)(1 + 3/4 + (3/4)^2 + \dots) = O(|V_P|). \quad \square$$

14. Continuity of Φ

Proof of Corollary 1.8. Suppose P is an N-gon and $\{P_n\}$ is a sequence of N-gons converging to P. This means the ordered list of vertices converges in $\mathbb{R}^{2N}$.

First suppose P has a ϕ-triangulation $\mathcal{T}$. Choose ϵ so small that any two distinct vertices of T are at least distance ϵ apart. For any $\delta > 0$ choose n so large that all the interior vertices of $\mathcal{T}$ are contained inside P_n, and each boundary vertex of $\mathcal{T}$ is within $\delta \cdot \epsilon$ of a point on P_n. Move the boundary points of $\mathcal{T}$ to these points on P_n. This creates a triangulation of P_n whose angles are within $O(\delta)$ of the corresponding angles of $\mathcal{T}$. Thus $\limsup_n \Phi(P_n) \leq \Phi(P)$.

Conversely, suppose $\phi = \liminf_n \Phi(P_n)$. Fix $\epsilon > 0$ and suppose $\mathcal{T}_n$ is a $(\phi + \epsilon)$-triangulation of P_n. Passing to a subsequence, if necessary, we assume $\phi = \lim_n \Phi(P_n)$ and $\Phi(P_n) < \phi + \epsilon$ for all n. As in the proof of Lemma 5.1, we may construct a quasiconformal map f_n from the interior of P_n to the interior of P with dilatation tending uniformly to zero as $n \nearrow \infty$. By Lemma 4.1, the push-forward of $\mathcal{T}_n$ under f_n is a triangulation with maximum angle $\leq \phi + 2\epsilon$ if n is large enough (it may also be necessary to choose $\mathcal{T}_n$ sufficiently fine). Thus P has $(\phi + 2\epsilon)$-triangulations for every $\epsilon > 0$, so $\Phi(P) \leq \liminf_n \Phi(P_n)$. Thus Φ is continuous at P.

If $\phi > 60$, then $\mathcal{E}(N, \phi) = \{P \in \mathcal{P}_N : \Phi(P) \leq \phi\}$ is the same as the set of N-gons that have a ϕ-triangulations (for $\phi = 60^\circ$ the latter is a subset of the former), so it is a closed set. Similarly for $\mathcal{F}(N, \phi) = \{P \in \mathcal{P}_N : \Phi(P) = \phi\}$. Corollary 1.7 shows that the $\mathcal{F}(N, 72^\circ)$ contains every polygon with all angles $\geq 144^\circ$; this is an open set for $N \geq 10$. Our remark following the proof of Corollary 1.6 shows that every axis-parallel polygon is also in the interior of $\mathcal{F}(N, 72^\circ)$ for $N \geq 4$.

To see that $\mathcal{F}(N, \frac{5}{7}90^\circ)$ has interior for N large enough, consider the polygon P in Figure 26. Note P has eight 60°-vertices, one 300° vertex and a large number of

$180^\circ + \epsilon$ vertices, where ϵ may be as small as we wish. For $\phi \leq \frac{5}{7} \cdot 90^\circ$, the only ϕ-admissible labels for the 60°, 180° and 300° vertices are 1, 3 and 5 respectively (see Figure 7). Thus in this range $\kappa(\phi) = 6 - 8(3-1) - 1(3-5) = 6 - 16 + 2 = -8 < 0$. Therefore, we must be in Case 2 of Theorem 1.1 and $\phi = \frac{5}{7} \cdot 90^\circ$ is the sharp bound. Any small perturbation of P has the same vertex labels, so P is in the interior of $\mathcal{F}(N, 72^\circ)$.

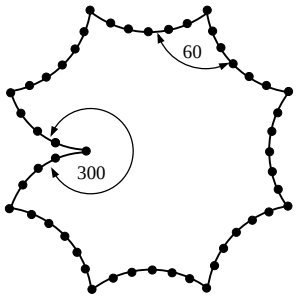

FIGURE 26. This polygon has optimal angle bound $\phi = \frac{5}{7} \cdot 90^\circ$, and so does every small perturbation, showing the set of such polygons has non-empty interior.

Finally, for any angle $\phi \in (60^\circ, 90^\circ] \setminus \{\frac{5}{7} \cdot 90^\circ, 72^\circ\}$, having sharp bound ϕ means that one or more of the angles of ϕ lies on the boundary of a shaded triangle in Figure 7 (otherwise we could decrease ϕ slightly and still meet the criteria of Theorem 1.1). This means three adjacent points on P form an angle from a certain finite set of possibilities depending only on ϕ. Therefore $\mathcal{F}(N, \phi)$ has co-dimension at least 1. $\square$

15. Minmax versus Maxmin

We recall some notation from Section 1. For $0 < \phi < 60^\circ$ we define $\widetilde{I}(\phi) = [\phi, 180^\circ - 2\phi]$; any triangle having smallest angle ϕ has all its angles inside $\widetilde{I}(\phi)$. Define a labeling L to be ϕ-lower-admissible if $\theta_v \in L(v) \cdot \widetilde{I}(\phi)$ where θ_v is the angle of P at $v \in V_P$. See Figure 27. The curvature $\kappa(L)$ the same as before, and $\widetilde{\mathcal{K}}(\phi)$ is the set of curvatures of ϕ-lower-admissible labelings. We set $\widetilde{\kappa}(\phi)$ to be the element of this set closest to 0 (equal to ∞ if no lower-admissible labeling exists). A ϕ-lower-triangulation means a triangulation will all angles $\geq \phi$.

Proof of Theorem 1.11. First consider necessity of the stated conditions. By definition, having a ϕ-lower-triangulation means that a lower-admissible labeling exists;

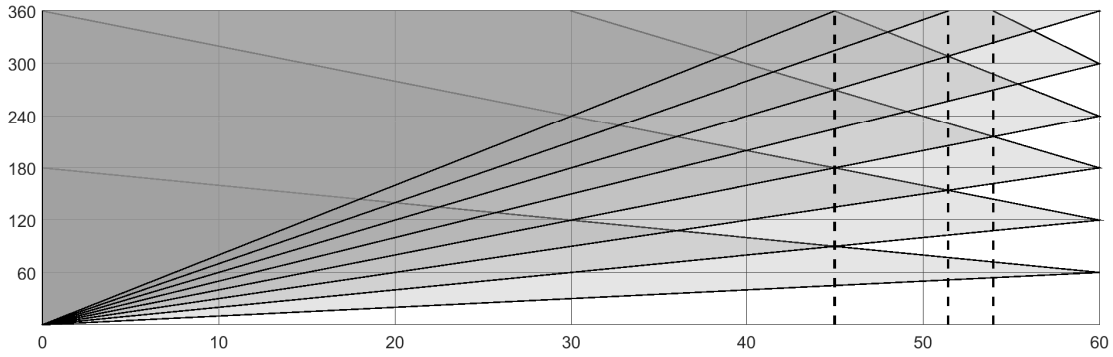

FIGURE 27. P has an ϕ-lower-admissible labeling iff all its angles lie in the intersection of the union of shaded triangles with the vertical line above ϕ. For $\phi \leq 45°$, this only depends on the size of the smallest angle. The dashed lines indicate $\phi = 45°$ and the two transition angles in Theorem 1.11.

this is Case 1. If every angle of the triangulation is greater than 45°, then every angle is also strictly less than 90°. This implies every vertex in $\partial\mathcal{T} \setminus V_P$ has degree 3, so (1.1) becomes

$$\sum_{v \in \text{int}(\mathcal{T})} \kappa(v) + \sum_{v \in V_P} \kappa(v) = 6. \tag{15.1}$$

If we also have $\phi > \frac{1}{7} \cdot 360°$, then there are no interior vertices of degree seven, so the first term in (15.1) is non-negative. Thus $\kappa(L) = 6 - \sum_{v \in V_P} \kappa(v) \geq 0$, as desired. If $\phi > 54°$, then every $v \in \text{int}(\mathcal{T})$ has degree six, so $\kappa(L) = 0$ as desired.

To prove sufficiency, we simply follow the proof of Theorem 1.1, except that in this case the 420°-trick is the first one eliminated at $\phi = \frac{1}{7} \cdot 360° \approx 51.4286$, and the 120°-trick is eliminated at $\phi = 54°$. □

Proof of Corollary 1.12 . The linear time calculation of $\widetilde{\Phi}(P)$ is very similar to the calculation of $\Phi(P)$ described earlier, so we only note a few changes to the proof of Corollary 1.2. We define

$$\widetilde{\phi}_\infty = \sup\{\phi \in [60°, 90°] : \widetilde{\kappa}(\phi) < \infty\}$$

$$\widetilde{\phi}_0 = \sup\{\phi \in [60°, 90°] : \widetilde{\kappa}(\phi) = 0\}.$$

Since $\widetilde{I}(0) = [0°, 180°]$ it is easy to check that $\widetilde{\kappa}(0°) = 0$, so $0 < \widetilde{\phi}_0 \leq \widetilde{\phi}_\infty$. If $\widetilde{\kappa}(\widetilde{\phi}_\infty) = 0$, then $\widetilde{\phi}_0 = \widetilde{\phi}_\infty$, and this common value equals $\widetilde{\Phi}(P)$. If $\widetilde{\kappa}(\widetilde{\phi}_\infty) > 0$, Then

$\widetilde{\kappa}(\phi)$ is increasing and non-negative and
• if $\widetilde{\phi}_\infty \le 54^\circ$, then $\widetilde{\Phi}(P) = \widetilde{\phi}_\infty$;
• if $\widetilde{\phi}_\infty > 54^\circ$, then $\widetilde{\Phi}(P) = 54^\circ$.
Otherwise, $\widetilde{\kappa}(\widetilde{\phi}_\infty) < 0$, and $\widetilde{\kappa}(\phi)$ is decreasing and non-positive and
• if $\widetilde{\phi}_\infty \le \frac{4}{7} \cdot 90^\circ$, then $\widetilde{\Phi}(P) = \widetilde{\phi}_\infty$;
• if $\widetilde{\phi}_\infty > \frac{4}{7} \cdot 90^\circ$, then $\widetilde{\Phi}(P) = \widetilde{\phi}_0$.
Thus we are reduced to calculating $\widetilde{\phi}_0$ and $\widetilde{\phi}_\infty$. Each of these can be done in linear time, just as in the proof Corollary 1.2 (the logic is the same, although the formulas are slightly different since $I(\phi)$ is different from $\widetilde{I}(\phi)$). □

Proof of Corollary 1.13. If $\phi > 45^\circ$, then every triangulation angle is $\le 180^\circ - 2\phi < 90^\circ$, so we are done.. For $\phi \le 45^\circ$, the angles created in the proof is always acute, except possible near a vertex v of P. If v has angle θ_v we choose the largest possible ϕ-lower-admissible label $L(v)$ for v. This means $L(v) \cdot \phi \le \theta_v < (1 + L(v)) \cdot \phi$, so

$$\phi \le \theta_v / L(v) < (1 + 1/L(v)) \cdot \phi \le 2 \cdot \phi \le 90^\circ.$$

Thus angles in a neighborhood of each vertex are in the correct range, and interior vertices of degree five or seven only introduce angles $\ge \frac{4}{7} \cdot 90^\circ > \phi$ and $\le 72^\circ$. □

Proof of Corollary 1.14 . See Figure 27. Note that if $\phi \le 45^\circ$, then $\bigcup_k k \cdot \widetilde{I}(\phi) = [\phi, \infty)$, so P has a ϕ-lower-admissible labeling iff $\theta_{\min} \ge \phi$. □

16. Questions and remarks

As noted at the beginning, this paper is concerned with computing the optimal angle bounds for triangulating any N-gon, not with finding efficient triangulations that attain these bounds. We have observed that the number of elements of such triangulations will not satisfy a polynomial bound in N in general, but can we find a ϕ-triangulation with the minimal number of elements, say in time comparable to the number of elements output (recall this need not be polynomial in the number of vertices)? Can we even estimate the smallest number of triangles needed in terms of the geometry of P? What if the vertices of P have integer coordinates in $[0, N]^2$?

Our proof of Theorem 1.1 does not give anything near the optimal number of triangles. Each edge e of the polygon P has a harmonic measure $\omega(z, e)$ that depends on a choice of base point z in the interior of P. This is the point that is mapped to

the origin by our conformal map to the disk, and $\omega(z, e)$ is the length of the image of e on the unit circle (usually normalized so the circle has length 1). Our construction will generally then use $O(\inf_z \inf_e \omega(e)^{-2})$ triangles. In the case of a $1 \times r$ rectangle, $\omega(z, e) \simeq \exp(-\pi r/2)$ for at least one of the short ends e, no matter how we choose the base point z, so our proof gives an exponential number of triangles as a function of r. However, it is easy to see by a direct construction that only $O(r)$ triangles are needed to achieve the optimal angle bound 72°. See Figure 28. Here we have chosen a P' that mimics the overall shape of the rectangle, and obtain $O(r)$ triangles, at the cost of introducing many more degree five vertices into the triangulation when we apply the 120°-trick. Does choosing a 60°-polygon P' to mimic P always give a nearly optimal number of triangles, at least if the optimal angle bound ϕ is $\geq 72°$?

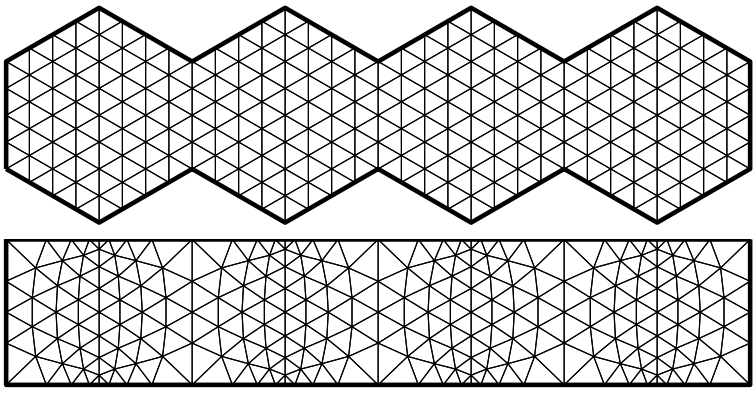

FIGURE 28. A different choice of P' leads to a number of triangles that is within a bounded factor of optimal. In this figure we have to make the grid finer and apply the 120°-trick to get rid of certain boundary vertices with angle 90°.

The Wallace-Bolyai-Gerwien theorem says two polygons P_1, P_2 of the same area are scissors-congruent, i.e., they each have triangular dissections into isometric triangles. If two polygons have the same set of angles, do they have a common dissection attaining $\Phi(P_1) = \Phi(P_2)$? In general, can we take the largest angle in a mutual dissection to be $\max(\Phi(P_1), \Phi(P_2))$? If not, how can we compute the optimal upper angle bound for consistent dissections of two polygons?

Does every planar straight line graph with minimal interior angle 36° have a 72° conforming triangulation? What about a polygon with holes? A triangulated surface?

If a PSLG has minimum angle θ, does it have conforming triangulation with all angles in the interval $J(\theta) = [\theta, 90^\circ - \min(\theta, 36^\circ)/2]$? It is true that there is a $\theta_0 > 0$ so that every PSLG has a conforming triangulation will all angles in $J(\theta_0)$ except for triangles that contain a vertex v of the PSLG where there is an angle $\theta < \theta_0$; such triangles are isosceles and have their angles in $J(\theta)$. We call this a uniformly acute triangulation. In particular, every triangulation with minimal angle θ has an acute refinement with maximum angle $\leq 90^\circ - \min(\theta, \theta_0)/2$, answering a questions of Florestan Brunck. When the PSLG is a simple polygon P, there is an uniformly acute triangulation with $\theta_0 = 30^\circ$ (this gives a weaker angle bound than the current paper, but it is stronger in the sense that only triangles touching a vertex of P have angles outside the interval $J(30^\circ)$). The optimal value of θ_0 remains unknown for both polygons and PSLGs.

Corollary 1.8 showed the that set of N-gons with $\Phi(P) = 72^\circ$ contains an open set, and thus it should have positive measure with respect to any measure on the space of N-gons that is absolutely continuous with respect to volume measure on $\mathbb{R}^{2N}$. What are natural examples of such measures, i.e., what is a random polygon? Figure 29 show the result of computing the optimal upper bound for a billion random lists of numbers summing to $1440 = 8 \cdot 180$. This is meant to simulate a random 10-gon, but the lists were generated by choosing ten random numbers in $[0, 1]$ and renormalizing to get the correct sum. This will sometimes generate angles $> 360^\circ$ and since side lengths are not accounted for, the lists don't correctly represent the angles of simple polygons. However, the predicted mass at 72° is clearly visible (about 20% of the total mass), but the predicted peak at $\frac{5}{7} \cdot 90^\circ$ is not; presumably its probability is too small to stand out here.

References

[1] L. V. Ahlfors. *Lectures on quasiconformal mappings*. The Wadsworth & Brooks/Cole Mathematics Series. Wadsworth & Brooks/Cole Advanced Books & Software, Monterey, CA, 1987. With the assistance of Clifford J. Earle, Jr., Reprint of the 1966 original.

[2] A. Alexandrescu. Fast deterministic selection. In *16th Symposium on Experimental Algorithms*, volume 75 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. 24, 19. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2017.

[3] B.S. Baker, E. Grosse, and C.S. Rafferty. Nonobtuse triangulation of polygons. *Discrete Comput. Geom.*, 3(2):147–168, 1988.

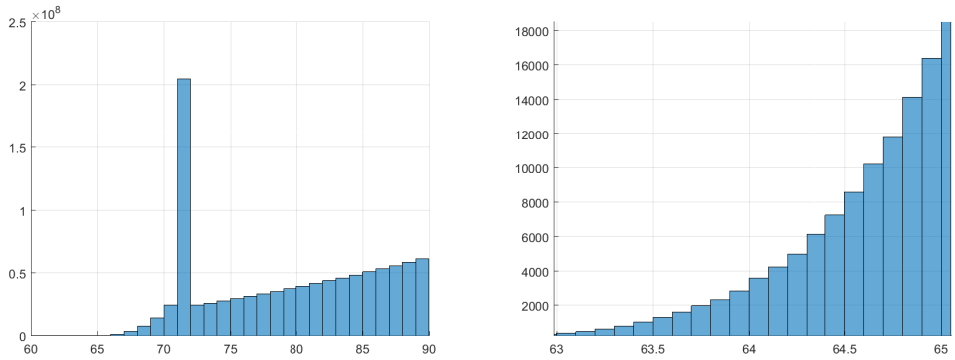

FIGURE 29. The distribution of optimal upper bounds over 10^9 random samples as described in the text. On the left is a histogram based on 1° bins. The spike a 72° is evident. On the right is an enlargement near 64° using .1° bins. No spike at $\frac{5}{7} \cdot 90° \approx 64.26°$ is visible.

[4] M. Bern, H. Edelsbrunner, D. Eppstein, S. Mitchell, and T. S. Tan. Edge insertion for optimal triangulations. In *LATIN '92 (São Paulo, 1992)*, volume 583 of *Lecture Notes in Comput. Sci.*, pages 46–60. Springer, Berlin, 1992.

[5] M. Bern, D. Eppstein, and J. Gilbert. Provably good mesh generation. *J. Comput. System Sci.*, 48(3):384–409, 1994. 31st Annual Symposium on Foundations of Computer Science (FOCS) (St. Louis, MO, 1990).

[6] M. Bern, S. Mitchell, and J. Ruppert. Linear-size nonobtuse triangulation of polygons. *Discrete Comput. Geom.*, 14(4):411–428, 1995. ACM Symposium on Computational Geometry (Stony Brook, NY, 1994).

[7] M. Bern and P. Plassmann. Mesh generation. In *Handbook of computational geometry*, pages 291–332. North-Holland, Amsterdam, 2000.

[8] C. J. Bishop. Conformal welding and Koebe's theorem. *Ann. of Math. (2)*, 166(3):613–656, 2007.

[9] C.J. Bishop. Minimal weight Steiner triangulation need not exist. 2007. Preprint.

[10] C.J. Bishop. Nonobtuse triangulations of PSLGs. *Discrete Comput. Geom.*, 56(1):43–92, 2016.

[11] M. Blum, V. Pratt, R.E. Tarjan, R.W. Floyd, and R.L. Rivest. Time bounds for selection. *J. Comput. System Sci.*, 7:448–461, 1973.

[12] Yu. D. Burago and V. A. Zalgaller. Polyhedral embedding of a net. *Vestnik Leningrad. Univ.*, 15(7):66–80, 1960.

[13] Bernard Chazelle and David P. Dobkin. Optimal convex decompositions. In *Computational geometry*, volume 2 of *Mach. Intelligence Pattern Recogn.*, pages 63–133. North-Holland, Amsterdam, 1985.

[14] E.B. Christoffel. Sul problema della tempurature stazonaire e la rappresetazione di una data superficie. *Ann. Mat. Pura Appl. Serie II*, pages 89–103, 1867.

[15] H.T. Croft, K.J. Falconer, and R.K. Guy. *Unsolved problems in geometry*. Problem Books in Mathematics. Springer-Verlag, New York, 1991.

[16] T.A Driscoll. Algorithm 843: Improvements to the Schwarz-Christoffel Toolbox for MATLAB. *ACM Transactions on Mathematical Software (TOMS)*, 31(2):239–251, 2005.

[17] T.A. Driscoll and L.N. Trefethen. *Schwarz-Christoffel mapping*, volume 8 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2002.
[18] H. Edelsbrunner. Triangulations and meshes in computational geometry. In *Acta Numerica, 2000*, volume 9 of *Acta Numer.*, pages 133–213. Cambridge Univ. Press, Cambridge, 2000.
[19] H. Edelsbrunner, T.S. Tan, and R. Waupotitsch. An $O(n^2 \log n)$ time algorithm for the minmax angle triangulation. *SIAM J. Sci. Statist. Comput.*, 13(4):994–1008, 1992.
[20] D. Eppstein. Acute square triangulation. Webpage https://www.ics.uci.edu/ eppstein/junkyard/acute-square/, Accessed: January 2, 2021.
[21] D. Eppstein. Approximating the minimum weight triangulation. In *Proceedings of the Third Annual ACM-SIAM Symposium on Discrete Algorithms (Orlando, FL, 1992)*, pages 48–57. ACM, New York, 1992.
[22] H. Erten and A. Üngör. Computing acute and non-obtuse triangulations. In *CCCG 2007, Ottawa, Cananda*. 2007.
[23] J.B. Garnett and D.E. Marshall. *Harmonic measure*, volume 2 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2005.
[24] J. L. Gerver. The dissection of a polygon into nearly equilateral triangles. *Geom. Dedicata*, 16(1):93–106, 1984.
[25] J.E. Goodman, J. O'Rourke, and C.D. Tóth, editors. *Handbook of discrete and computational geometry*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton, FL, 2018. Third edition of [MR1730156].
[26] D. H. Hamilton. Conformal welding. In *Handbook of complex analysis: geometric function theory, Vol. 1*, pages 137–146. North-Holland, Amsterdam, 2002.
[27] J. Itoh and T. Zamfirescu. Acute triangulations of the regular icosahedral surface. *Discrete Comput. Geom.*, 31(2):197–206, 2004.
[28] M. Keil and J. Snoeyink. On the time bound for convex decomposition of simple polygons. *Internat. J. Comput. Geom. Appl.*, 12(3):181–192, 2002.
[29] S. Korotov and J. Stańdo. Nonstandard nonobtuse refinements of planar triangulations. In *Conjugate gradient algorithms and finite element methods*, Sci. Comput., pages 149–160. Springer, Berlin, 2004.
[30] C.L. Lawson. *Software for C^1 surface interpolation*, pages ix+388. Academic Press [Harcourt Brace Jovanovich Publishers], New York, 1977. Publication of the Mathematics Research Center, No. 39.
[31] D. T. Lee and A. K. Lin. Generalized Delaunay triangulation for planar graphs. *Discrete Comput. Geom.*, 1(3):201–217, 1986.
[32] O. Lehto and K. I. Virtanen. *Quasiconformal mappings in the plane*. Springer-Verlag, New York-Heidelberg, second edition, 1973. Translated from the German by K. W. Lucas, Die Grundlehren der mathematischen Wissenschaften, Band 126.
[33] J. Y. S. Li and H. Zhang. Nonobtuse remeshing and mesh decimation. In *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing*, pages 235–238, Aire-la-Ville, Switzerland, Switzerland, 2006. Eurographics Association.
[34] H. Maehara. Acute triangulations of polygons. *European J. Combin.*, 23(1):45–55, 2002.
[35] E.A. Melissaratos and D.L. Souvaine. Coping with inconsistencies: a new approach to produce quality triangulations of polygonal domains with holes. In *SCG '92: Proceedings of the eighth annual symposium on computational geometry*, pages 202–211, New York, NY, USA, 1992. ACM.
[36] S.A. Mitchell. Approximating the maxmin-angle covering triangulation. volume 7, pages 93–111. 1997. Fifth Canadian Conference on Computational Geometry (Waterloo, ON, 1993).

[37] Z. Nehari. *Conformal mapping*. Dover Publications Inc., New York, 1975. Reprinting of the 1952 edition.
[38] S. Rohde. On conformal welding and quasicircles. *Michigan Math. J.*, 38:111–116, 1991.
[39] J. Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (Austin, TX, 1993)*, pages 83–92, New York, 1993. ACM.
[40] S. Saraf. Acute and nonobtuse triangulations of polyhedral surfaces. *European J. Combin.*, 30(4):833–840, 2009.
[41] H.A. Schwarz. Confome abbildung der oberfläche eines tetraeders auf die oberfläche einer kugel. *J. Reine Ange. Math.*, pages 121–136, 1869. Also in collected works, [42], pp. 84-101.
[42] H.A. Schwarz. *Gesammelte Mathematische Abhandlungen*. Springer, Berlin, 1890.
[43] L. N. Trefethen and T.A. Driscoll. Schwarz-Christoffel mapping in the computer era. In *Proceedings of the International Congress of Mathematicians, Vol. III (Berlin, 1998)*, number Extra Vol. III, pages 533–542 (electronic), 1998.
[44] L.N. Trefethen. Numerical computation of the Schwarz-Christoffel transformation. *SIAM J. Sci. Statist. Comput.*, 1(1):82–102, 1980.
[45] L. Yuan. Acute triangulations of polygons. *Discrete Comput. Geom.*, 34(4):697–706, 2005.
[46] C.T. Zamfirescu. Survey of two-dimensional acute triangulations. *Discrete Math.*, 313(1):35–49, 2013.

C.J. Bishop, Mathematics Department, Stony Brook University, Stony Brook, NY 11794-3651

Email address: bishop@math.stonybrook.edu