

OPTIMAL TRIANGULATION OF POLYGONS

CHRISTOPHER J. BISHOP

ABSTRACT. How do we cut a polygon into triangles that are all as “round” as possible, e.g., minimizing the maximum angle used? In this paper, we compute the optimal upper and lower angle bounds for triangulating an N -gon P with Steiner points, sharpening the 1960 theorem of Burago and Zalgaller that every polygon has an acute triangulation. For any polygon, we show both the upper and lower bounds can be computed in linear time from the list of interior angles of the polygon. We also show that both types of optimal bound are usually attained by some finite triangulation of the polygon (but sometimes they cannot both be attained by a single triangulation). We do not address the interesting problem of finding efficient triangulations that attain the optimal angle bounds; even in some simple cases, our construction gives many more triangles than are actually needed. The exceptional polygons where the optimal bounds can only be approximated, but not attained, are easily described: if and only if every interior angle is an integer multiple of 60° , and some pair of sides has irrational length ratio. We also show that the optimal angle bounds for polygonal triangulations are the same as for triangular dissections. This implies, in a stronger form, a 1984 conjecture of Gerver. Although the statements of our results involve only Euclidean geometry, the proofs depend on conformal and quasiconformal techniques.

Date: first draft February 21, 2021; revised Sept. 21, 2025.

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

Key words and phrases. acute triangulation, Steiner points, dissections, Schwarz-Christoffel formula, quasiconformal mappings, Delaunay triangulation.

The author is partially supported by NSF Grant DMS 2303987.

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 [46], [47], and algorithms for minimizing the maximum angle (again without Steiner points) are given in [5] and [30]. Here we solve the analogous problems when Steiner points are permitted.

In this case, Burago and Zalgaller [19] proved in 1960 that every planar polygon P has an acute triangulation (all angles $< 90^\circ$). This is the best possible if we want a uniform 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? How do the angle bounds for triangulations differ from those for dissections? (A dissection of a polygonal region P is a decomposition of P into finitely many triangles with disjoint interiors, but a triangulation also requires the simplex condition: adjacent triangles intersect in either a vertex of both triangles or a full edge of both triangles. See Figure 4.)

Using ideas involving conformal and quasiconformal mappings, we shall answer each of these questions. Theorem 1.1 below will characterize when a polygon P has a ϕ -triangulation (i.e., a triangulation with all angles $\leq \phi$), in terms of ϕ -admissible labelings of the vertices of P with positive integers, and the curvatures of such labelings. Before stating the theorem we define these two terms.

We start with admissibility. Let V_P be the set of vertices of P , and let $|V_P|$ denote the size of this set. In a ϕ -triangulation, all the angles will lie in the interval $I(\phi) = [180 - 2\phi, \phi]$. (This holds since the angles in a triangle sum to 180° .) A positive integer L is called admissible for a vertex $v \in V_P$ if L values can be chosen in $I(\phi)$, so that they sum to the interior angle θ_v of P at v . In other words, L is admissible for v if the interior angle at v can be divided into L sub-angles that all lie in the required interval. A labeling $L : V_P \rightarrow \mathbb{N}$ is admissible if each $L(v)$ is an admissible integer for v .

Given a triangulation, we get a labeling $L : V_P \rightarrow \mathbb{N}$ by setting $L(v)$ to be the number of triangles containing the vertex v (we call this the “face-degree” of v). Since each angle of a ϕ -triangulation is in $I(\phi)$, we must have $\theta_v \in L(v) \cdot I(\phi)$, so this labeling is clearly ϕ -admissible. Thus having a ϕ -admissible labeling is a necessary condition for having a ϕ -triangulation, but the converse is not true without some extra conditions, discussed below.

Next, we define the curvature of a triangulation \mathcal{T} of P . Let $V_{\mathcal{T}}$ denote the vertices of \mathcal{T} , and note that $V_P \subset V_{\mathcal{T}}$. We define $\partial\mathcal{T} = V_{\mathcal{T}} \cap P$ to be the boundary vertices of \mathcal{T} , and let $\text{int}(\mathcal{T}) = V_{\mathcal{T}} \setminus \partial\mathcal{T}$ denote the interior vertices. Note that $\partial\mathcal{T} \setminus V_P$ is just the set of Steiner points on the boundary. For $v \in \partial\mathcal{T}$, we define its discrete curvature as $\kappa(v) = 3 - L(v)$, and for an interior vertex we set $\kappa(v) = 6 - L(v)$. See Figure 1. Using these definitions, Euler’s formula applied to a triangulation can be rewritten to look like the Gauss-Bonnet formula (for details see Section 14):

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

This value is called the curvature of the triangulation, and is denoted $\kappa(\mathcal{T})$.

The term “curvature” for this quantity serves to strengthen the comparison between Euler’s formula and the Gauss-Bonnet theorem. It also arises from the observation that if the triangles were all equilateral¹, then an interior vertex v having degree six corresponds to an angle sum of 360° around v , as happens in the Euclidean plane (a zero curvature surface). A degree less than six gives an angle sum $< 360^\circ$, as happens in spherical geometry (positive curvature) and degrees larger than six give angle sums $> 360^\circ$, as in hyperbolic geometry (negative curvature).

The curvature of a labeling L of V_P is defined as

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

Note that by (1.1), we have

$$(1.2) \quad \kappa(L) = \kappa(\mathcal{T}) + \sum_{v \in \partial\mathcal{T} \setminus V_P} \kappa(v).$$

As noted earlier, the second sum is over boundary Steiner points.

¹Riemann surfaces constructed by glueing together Euclidean equilateral triangles are an intensely studied topic with important connections to complex analysis, probability, algebraic number theory, Galois theory and Grothendieck’s theory of *dessins d’enfants*. See [17], [65], and their references.

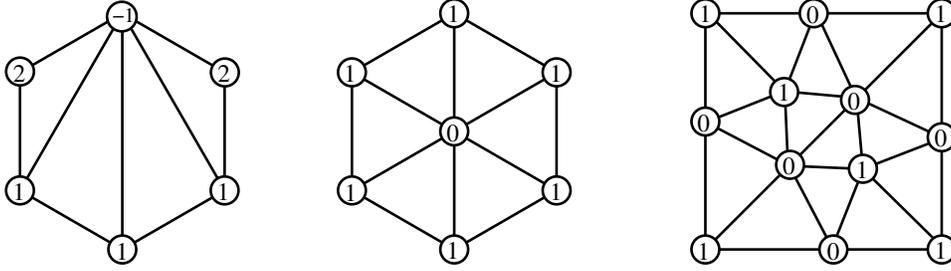


FIGURE 1. Vertices are labeled with their discrete curvatures; the triangulations have curvatures 0, 0, 2 respectively. The left and center show that allowing Steiner points can lower the maximum angle (here, from 120° to 60°). The center and right triangulations have the optimal upper angle bounds for a hexagon (60°) and a square (72°).

Given a triangulation \mathcal{T} of a polygon P , we noted above that we can label each of its vertices by the number of triangles that contain that vertex. We say this is the labeling associated to the triangulation. For an interior vertex v , this number is the same as the degree of v , i.e., the number of edges with endpoint v . However, for a boundary vertex v , the number of faces touching v is one less than the number of edges. I am not aware of a standard term for the number of bounded faces touching a vertex in a planar graph, so we will call this the “face-degree” of the vertex, although sometimes we will simply use the term “degree” for brevity, especially when there is little chance of confusion (e.g., for interior vertices). Notice that an interior vertex (necessarily Steiner) must have degree at least three, and that a Steiner vertex on the boundary must have face-degree at least two.

We would like to know which labelings of V_P are the restrictions to V_P of the labeling of a ϕ -triangulation of P . We noted above that 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. Moreover, it must satisfy the following constraints, depending on ϕ (these will be explained more carefully in Section 14).

- (1) For any $\phi < 90^\circ$, we have $\kappa(L) \leq \kappa(\mathcal{T})$. This follows from (1.2) since vertices $v \in \partial\mathcal{T} \setminus V_P$ (boundary Steiner points), must have face-degree ≥ 3 , and hence $\kappa(v) \leq 0$.
- (2) If $\phi < 72^\circ$ then $\text{int}(\mathcal{T})$ has no interior vertices of degree ≤ 5 , so $\kappa(\mathcal{T}) \leq 0$. By the previous case, we also have $\kappa(L) \leq \kappa(\mathcal{T}) \leq 0$.

- (3) Finally, if $\phi < \frac{5}{7} \cdot 90^\circ \approx 64.2857^\circ$ then every vertex in $\text{int}(\mathcal{T})$ has degree exactly 6 and every vertex in $\partial\mathcal{T} \setminus V_P$ has face-degree 3, so $\kappa(L) = \kappa(\mathcal{T}) = 0$.

Remarkably, these elementary necessary conditions are also sufficient.

Theorem 1.1. *For $60^\circ < \phi \leq 90^\circ$, a polygon P has a ϕ -triangulation if and only if*

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

It is sometimes convenient to restate the theorem using slightly different notation. Fix a polygon P and 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 ϕ -admissible labeling. Note that $\mathcal{K}(\phi)$ is either $\{\infty\}$ or a non-empty interval of integers. Hence $\mathcal{K}(\phi)$ has a unique closest element to 0, denoted $\kappa(\phi)$ (this either an integer, possibly 0, or ∞). The three conditions in Theorem 1.1 can be restated as $\kappa(\phi) < \infty$ (finite curvature), $\kappa(\phi) \leq 0$ (non-positive curvature) and $\kappa(\phi) = 0$ (zero curvature) respectively. We should write $\mathcal{K}(\phi, P)$ and $\kappa(\phi, P)$, since these quantities also depend on P , but in this paper P is usually fixed and clear from context.

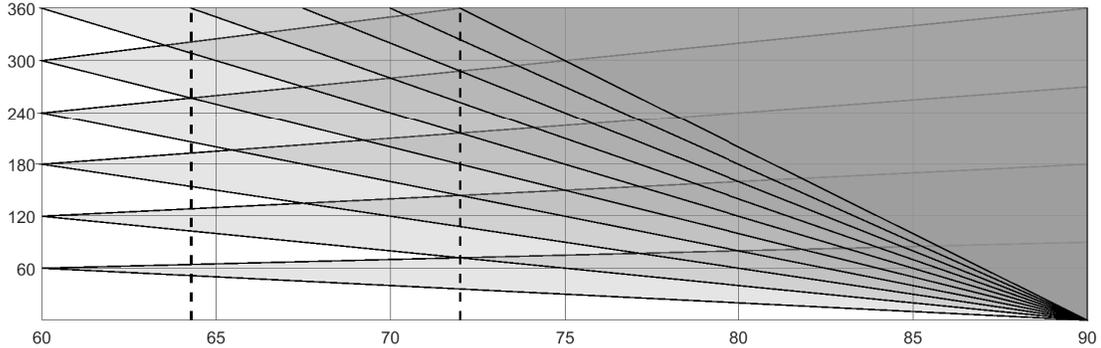


FIGURE 2. The horizontal axis is ϕ . Each shaded triangle is a plot of $k \cdot I(\phi)$ for $k = 1, 2, \dots$. The dashed vertical lines indicate where the transitions occur in Theorem 1.1, i.e., $\phi = \frac{5}{7} \cdot 90^\circ$, and $\phi = 72^\circ$.

Figure 2 explains the idea of an admissible labeling graphically. It plots $\cup_k k \cdot I(\phi)$ vertically above each value of ϕ , and the result is a union of shaded triangles (one for each k). A polygon 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 $72^\circ \leq \phi \leq 90^\circ$, $\cup_k k \cdot I(\phi) = [180^\circ - 2\phi, 360^\circ]$ so having an admissible labeling only depends on the size of θ_{\min} , and is equivalent to having a ϕ -triangulation. For $\phi < 72^\circ$, having a ϕ -triangulation requires other conditions as well, involving which shaded triangles the angles of P lie in.

It turns out that applying Theorem 1.1 is often fairly simple. In the remainder of the introduction we will discuss a number of corollaries and examples that follow from it. The first consequence is that the optimal angle bound is usually attained. For a simple polygon P , define

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

Corollary 1.2. *If $\Phi(P) > 60^\circ$, then $\Phi(P)$ is attained by some finite triangulation.*

Proof. Suppose P has ϕ_n -admissible labelings L_n for a sequence $\phi_n \searrow \Phi(P) < 90^\circ$. By removing finitely many elements, we may assume the whole sequence lies in the same case of Theorem 1.1 that $\Phi(P)$ does, and also that $\phi_0 < 90^\circ$. Since $\phi_n \leq \phi_0$, L_n is also ϕ_0 -admissible, and hence $L_n(v)(180^\circ - 2\phi_0) \leq 360^\circ$. Therefore $L_n(v)$ is uniformly bounded, depending only on ϕ_0 . There are only finitely many labelings satisfying such a bound, so some labeling L occurs along an infinite subsequence $\{n_k\}$. If $\theta_v \in L(v) \cdot I(\phi_{n_k})$ for every k , then $\theta_v \in L(v) \cdot I(\phi)$, since $I(\phi) = \cap_k I(\phi_{n_k})$. Hence L is $\Phi(P)$ -admissible, and hence a $\Phi(P)$ -triangulation exists by Theorem 1.1. It is a finite triangulation since all triangulations are finite by definition. \square

Note that even though the number of triangles in the ϕ_n -triangulations in the above proof might tend to infinity, the fact that L is a $\Phi(P)$ -admissible labeling of the vertices V_P implies that a $\Phi(P)$ -triangulation exists. In other words, the triangulation we get from Theorem 1.1 need not be the limit of the ϕ_n -triangulations associated to L_n .

When $\Phi(P) = 60^\circ$, the preceding argument fails because of the strict inequality in Case 3 of Theorem 1.1. It is easy to see that $\Phi(P) = 60^\circ$ implies P must be a 60° -degree polygon, i.e., every interior angle is an integer multiple of 60° . See Figure 3. Attaining $\Phi(P) = 60^\circ$ means that P has a triangulation by equilateral triangles, and all these triangles must have the same side length ℓ . Hence any two sides of P

have lengths that are both integer multiples of ℓ , and thus they have rational length ratio. The converse is also easy to prove, so we have the following result.

Corollary 1.3. *If $\Phi(P) = 60^\circ$, then P is a 60° -polygon and the bound is attained if and only if any two side lengths of P have rational ratio.*

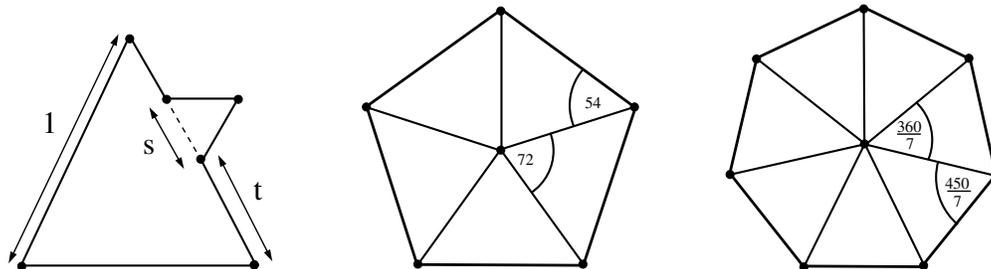


FIGURE 3. The 60° -polygon on the left satisfies $\Phi(P) = 60^\circ$, but this bound is not attained unless both s and t are rational. The other pictures show where the special angles in Theorem 1.1 come from: these angles are forced by interior vertices of degree five or seven.

As noted earlier, Burago and Zalgaller [19] proved that $\Phi(P) < 90^\circ$ for any polygon P . Theorem 1.1 implies a sharp, explicit improvement of this.

Corollary 1.4. *If the minimal interior angle of a polygon P satisfies $\theta_{\min} \leq 36^\circ$, then $\Phi(P) = 90^\circ - \frac{1}{2}\theta_{\min}$. If $\theta_{\min} > 36^\circ$ then $\Phi(P) \leq 72^\circ$.*

This is proved in Section 9. The first estimate is best possible, since if a polygon P has an interior angle θ at vertex v , then any triangulation of P contains a triangle T with angle $\leq \theta$ at v . Since the angles of T sum to 180° , T must also contain an angle $\geq 90^\circ - \theta/2$. The fact that $\Phi(P) \leq 72^\circ$ if $\theta_{\min} \geq 36^\circ$, strengthens a result of Gerver for dissections (Theorem 2, [37]). A triangular dissection of P is a finite collection of closed triangles that exactly covers P and its interior, and so that the triangles have pairwise disjoint interiors. The edges of adjacent triangles in a dissection need not match up exactly; if they do, then we have a triangulation. See Figure 4. A ϕ -dissection is a triangular dissection with maximum angle $\leq \phi$. In 1984, Gerver [37] 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 are sufficient for a ϕ -dissection to exist if $\phi > 60^\circ$. Combining Gerver's result and Corollary 1.2 gives the following result.

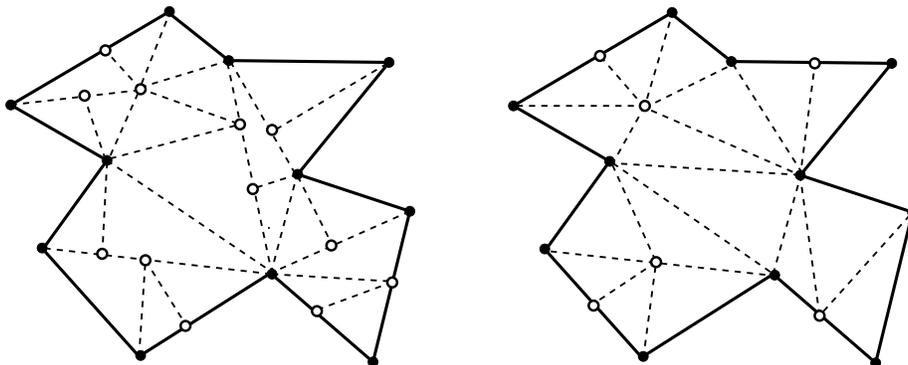


FIGURE 4. On the left is a dissection of a polygon and on the right a triangulation. The white dots are the Steiner points. Despite triangulations being more restrictive than dissections, the optimal upper angle bounds are the same for both types of decomposition.

Corollary 1.5. *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.*

Thus $\Phi(P) = \mathcal{M}(P) := \inf\{\phi : P \text{ has a } \phi\text{-dissection}\}$ for every polygon P .

The final conclusion is surprising (at least to the author). Since dissections satisfy less stringent conditions than triangulations do, one might expect $\Phi(P) > \mathcal{M}(P)$ for some polygons P , but no such gaps occur. Figure 3 shows (for irrational s) a 60° -polygon that has no equilateral triangulation, but that does have an equilateral dissection. However, Tutte [63] proved that a convex 60° -polygon has no equilateral dissection unless all the length ratios are rational. See also Theorem 4 of [43].

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 [25].

Corollary 1.6. *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.*

Finding triangulations with good angle bounds has a long history and many applications, e.g, see [12] or [67] for lists of algorithms, such as the finite element method, that work better with well formed meshes. Burago and Zalgaller's theorem from [19]

was an element of their polyhedral version of the Nash embedding theorem, but it long remained unknown in the western computational geometry literature. The first reference to it that I am aware of is [40] in 2004. In 1988, Baker, Grosse and Rafferty [4] 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., [6], [7], [8], [12], [29], [35], [45], [49], [51], [56]. For a recent survey, see Chapter 29 of [38]. In 2002 Maehara [50] 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 [66]. A simpler approach was given by Saraf in [57].

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. Triangulating the square with optimal angles is discussed by Gerver [37] and Eppstein [31], but I do not know of other cases in the literature. One possible reason may be the close connection to conformal mappings described below; it is hard to see how our proof of Theorem 1.1 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, e.g., 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 large angle $\geq 90^\circ - \theta/2$. Thus the number of elements of an angle-optimal triangulation of an N -gon is not bounded by any fixed function of N , independent of the geometry. This paper does not address the question of finding efficient triangulations attaining the bound $\Phi(P)$, but we will show computing $\Phi(P)$ itself is “easy”.

Corollary 1.7. *$\Phi(P)$ can be computed in time $O(|V_P|)$.*

In 1992 Edelsbrunner, Tan and Waupotitsch [30] gave a $O(N^2 \log N)$ algorithm for minimizing the maximum angle of a polygonal triangulation without using Steiner points, and their result has not yet been improved, so far as I (or Edelsbrunner, personal communication) know. Thus finding the optimal angle bound over the infinite dimensional space of Steiner triangulations is faster than than the best known

algorithms for finding the optimal bound over the finitely many triangulations that don't use Steiner points.

A similar situation regarding Steiner points occurs when decomposing a polygon into convex pieces: Chazelle and Dobkin [21] obtain a faster solution using Steiner points than is obtained by Keil and Snoeyink [41] without Steiner points. Their result, and the one given in this paper, runs counter to the reasonable expectation that an optimization problem becomes harder by introducing Steiner points. However, one of the referees pointed out that in some situations, allowing Steiner points removes difficult constraints, and thus can make a problem easier, citing the fact that integer programming problems are sometimes solved by first considering a less constrained linear program.

To illustrate the relative ease of computing $\Phi(P)$, we will calculate it for the following examples in Section 11.

Corollary 1.8. *For the regular N -gon, the sharp upper bound is $\Phi_N = 72^\circ$ except when $N = 3, 6, 7, 8, 9$; then $\Phi_N = 60^\circ, 60^\circ, \frac{5}{7} \cdot 90^\circ, 67.5^\circ, \text{ and } 70^\circ$ respectively.*

Corollary 1.9. *Any axis-parallel polygon P has $\Phi(P) = 72^\circ$. There is an optimal triangulation with all interior vertices of degree six, except for two of degree five.*

Given P , let $\theta_{\min}, \theta_{\max}$ be the minimum and maximum interior angles of P . If a sequence of polygons is inscribed in a fixed smooth curve γ , then their angles tend uniformly to 180° as the maximum edge length tends to zero. Thus the following result applies whenever P is inscribed in a smooth curve and it has sufficiently short edges (depending on the smooth curve).

Corollary 1.10. *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 have degree six).*

Note that 72° is frequently the upper bound. This is partially explained by the fact that the set N -gons with this optimal bound contains an open set in the space \mathfrak{P}_N of all simple N -gons (we put a topology on \mathfrak{P}_N by thinking of it as a subset of $\mathbb{C}^N = \mathbb{R}^{2N}$). For example, Corollary 1.10 shows the set of polygons with $\theta_{\min} > 144^\circ$ is such an open set. The following result is proven in Section 13.

Corollary 1.11. *The map $P \rightarrow \Phi(P)$ is continuous. Thus $\{P \in \mathfrak{P}_N : \Phi(P) \leq \phi\}$ is closed in \mathfrak{P}_N , as is $\{P \in \mathfrak{P}_N : \Phi(P) = \phi\}$. For N sufficiently large, the latter set has non-empty interior iff $\phi = \frac{5}{7} \cdot 90^\circ$ or $\phi = 72^\circ$; otherwise it has co-dimension ≥ 1 .*

Suppose we wish to maximize the minimal angle of a triangulation, instead of minimizing the maximum angle? As noted earlier, the Delaunay triangulation of a point set maximizes the minimal angle without Steiner points; with Steiner points, angles arbitrarily close to 60° can be achieved for point sets. The constrained Delaunay triangulation does the same for triangulating polygons without Steiner points (see [46], [47]), and an algorithm using only interior Steiner points is presented in [52]. The methods of this paper can be used to maximize the minimum angle when triangulating a polygon P with arbitrary Steiner points. To state the result we introduce some more notation, analogous to that used in Theorem 1.1.

For $0 < \phi < 60^\circ$ we define $\tilde{I}(\phi) = [\phi, 180^\circ - 2\phi]$; any triangle having smallest angle ϕ must have all its angles inside $\tilde{I}(\phi)$. Define a labeling L to be ϕ -lower-admissible if $\theta_v \in L(v) \cdot \tilde{I}(\phi)$ where θ_v is the interior angle of P at $v \in V_P$. The curvature $\kappa(L)$ is defined just as before, and $\tilde{\mathcal{K}}(\phi)$ is the set of curvatures of ϕ -lower-admissible labelings. Also as before, $\tilde{\kappa}(\phi)$ 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 with all angles $\geq \phi$ and we define $\tilde{\Phi}(P)$ to be the supremum of ϕ so that P has a ϕ -lower-triangulation. The following result and its corollaries are proven in Section 15 (only minor changes to the proof of Theorem 1.1 are needed).

Theorem 1.12. *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 $\tilde{\kappa}(\phi) < \infty$,
- (2) $\frac{1}{7} \cdot 360^\circ < \phi \leq 54^\circ$, and $\tilde{\kappa}(\phi) \geq 0$,
- (3) $54^\circ < \phi < 60^\circ$, and $\tilde{\kappa}(\phi) = 0$.

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

Corollary 1.13. *$\tilde{\Phi}(P)$ can be computed in time $O(|V_P|)$.*

Corollary 1.14. *If P has a ϕ -lower-triangulation, then it also has an acute ϕ -lower-triangulation.*

Corollary 1.15. *If $\theta_{\min} \leq 45^\circ$, then $\tilde{\Phi}(P) = \theta_{\min}$.*

Comparing Corollaries 1.4 and 1.15, we see that if P is a polygon with $\theta_{\min} = 45^\circ$, then $\Phi(P) \leq 72^\circ$ and $\tilde{\Phi}(P) = 45^\circ$. 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 $< 45^\circ = \tilde{\Phi}(\phi)$. Thus we have the following result.

Corollary 1.16. *There exist polygons so that any triangulation attaining the optimal upper angle bound $\Phi(P)$, does not achieve the optimal lower angle bound $\tilde{\Phi}(P)$.*

Gerwer [37] used conformal maps to transfer dissections from the unit disk to a polygon. In this paper we will associate to each polygon P a 60° -polygon P' , together with a conformal map f between their interiors that defines a bijection between their vertices (but in some cases, P will be replaced by a subdomain obtained by cutting a few “slits” in P). We then transfer a nearly equilateral triangulation from P' to P using f . More precisely, we map the triangulation vertices from P' to P , and connect the images 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 5. The triangulation labeling shown in Figure 5 is 72° -admissible and has curvature 0, so $\kappa(72^\circ) = 0$. Moreover, one can check that $\kappa(\phi) > 0$ for $\phi < 72^\circ$, hence $\Phi(P) = 72^\circ$ by Theorem 1.1. As the mesh in Figure 5 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).

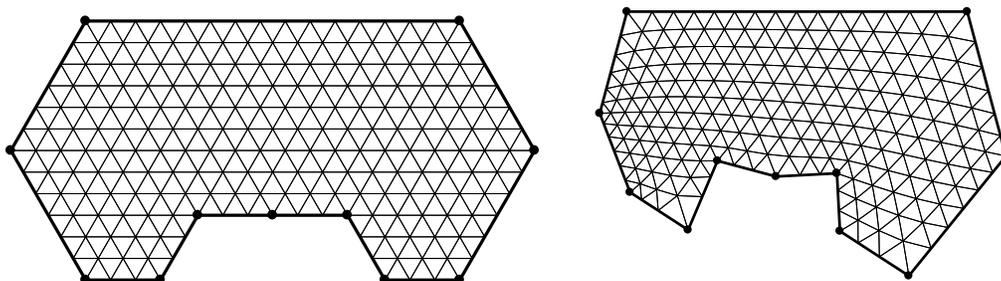


FIGURE 5. An equilateral triangulation of P' (left) and its conformal image P (angles of P : $270^\circ, 126^\circ, 96^\circ, 126^\circ, 105^\circ, 105^\circ, 144^\circ, 144^\circ, 80^\circ, 262^\circ, 162^\circ$). The maximum angle here is $\approx 73.5205^\circ$ and approaches 72° as the mesh gets finer; 72° can be attained by modifying a sufficiently fine triangulation near some of the vertices.

However, this scheme does not always work without some modifications. An equilateral triangulation of P' must have curvature zero, and gives a labeling of P' that has curvature zero. Hence the corresponding labeling of P should have curvature zero, but this is not always possible. For example, suppose P is a pentagon with five equal angles of 108° . Corollary 1.8 implies that $\Phi(P) = 72^\circ$. However, the only 72° -admissible labels for a 108° -vertex are $\{2, 3\}$, and this implies $\mathcal{K}(\phi) = \{1, \dots, 6\}$. and hence that $\kappa(\phi) = 1 > 0$. This holds even if we add extra 180° -vertices to P (their possible labels are $\{3, 4, 5\}$). Thus, any 72° -triangulation of P has positive curvature and so has some interior vertex of degree 5 (degree ≤ 4 would imply an angle $\geq 90^\circ$).

We get around this topological difficulty by cutting slits in P and allowing two adjacent boundary edges of P' to map onto either side of a slit. The vertex of P' common to these two sides of P' becomes an interior vertex of P (the tip of the slit), and this vertex need not have face-degree six. See Figure 6. For a polygon P , up to $|\kappa(\Phi(P))|$ slits are needed, introducing interior vertices of degree five ($\kappa > 0$) or seven ($\kappa < 0$). In Section 10 we will prove the following sharp bounds for the number of such “exceptional” vertices.

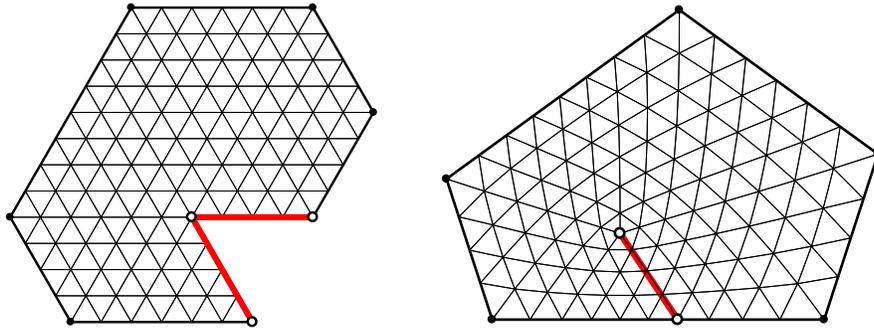


FIGURE 6. The edges adjacent to the 300° -vertex of P' are mapped to the two sides of the slit inside P . An interior vertex v of degree 5 is created. The slit is slightly curved to make the triangles on either side match up, although this is not easily visible (the actual slit lies slightly above the chord between its endpoints; total bending is about 3°).

Corollary 1.17. *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 $72^\circ \leq \phi < 90^\circ$, then $\max(0, \kappa(\phi))$ vertices have degree five,*

- (2) if $\frac{5}{7} \cdot 90^\circ \leq \phi < 67.5^\circ$, then $-\kappa(\phi)$ interior vertices have degree seven,
- (3) For $67.5^\circ \leq \phi \leq 72^\circ$ there are $-\kappa(\phi)$ vertices of curvature -1 , and these may each be chosen to be either in the interior (degree seven) or on the boundary (face-degree four).

The scheme outlined above encounters a number of difficulties in choosing the correct P' , estimating angles of conformal push forwards of triangles, and merging different triangulations. We overcome these problems using ideas from complex analysis. We list a few of these ideas here, giving details later.

- **Conformal welding:** When we map boundary edges of P' to a slit in P , the images of certain triangles in P' must match up across the slit in P , so that the image will be a triangulation. This requires that arclength measure on each of the two boundary segment maps to identical measures 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., [9], [39], [54], and in our case it can be solved explicitly. The slit in P is generally not a line segment, but is slightly curved.

- **Riemann surfaces:** In cases where we introduce an interior vertex of degree seven, P' will need to have a boundary vertex of face-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 7.

- **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.

- **Quasiconformal mappings:** Transferring a nearly equilateral triangulation from P' to P will allow us to approximate the optimal bounds, but to actually attain them, we 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. We merge the triangulations by using a partition of unity to interpolate from one

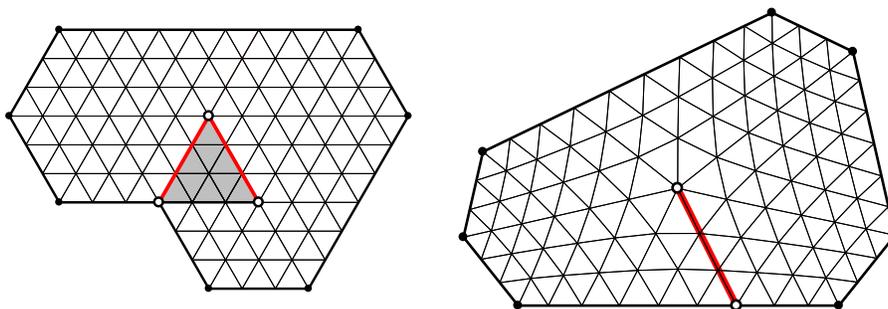


FIGURE 7. We triangulate an equal-angle heptagon using a Riemann surface with a 420° -vertex to insert a degree 7 interior vertex into the triangulation. The degree seven vertex is the image of the “top” corner of the shaded triangle. The surface on the left “wraps around” this point, and the self-overlapping part of the surface is the shaded triangle. The “left” and “right” sides of this shaded triangle in P' map to the two sides of the slit in P . See Figure 21 for an alternate rendering of shaded triangle.

conformal map to the other. The result is a quasiconformal map, and we shall use standard properties of such maps to control the interpolated triangulations.

This paper arose from a related question asked by Florestan Brunck: given a planar triangulation with lower angle bound θ , does it have an acute refinement with upper angle bound strictly less than 90° and only depending on θ ? In [15] I proved the answer is yes, by showing that every planar straight line graph (PSLG) Γ has a uniformly acute triangulation \mathcal{T} . This means that there is a universal $\theta_0 > 0$, independent of Γ , so that Γ has a conforming triangulation with all angles inside the fixed interval $[\theta_0, 90^\circ - \theta_0/2]$, except for triangles that contain vertices v of Γ that have an interior angle $\theta_v < \theta_0$ (these are unavoidable). The exceptional triangles are isosceles with angles in $[\theta_v, 90^\circ - \theta_v/2]$. The proof in [15] uses a compactness argument, and does not give an explicit value for θ_0 , but [14] shows we can take $\theta_0 = 30^\circ$ in the special case of polygons, i.e., every polygon has a triangulation with all angles in $[30^\circ, 75^\circ]$, except for one isosceles triangle at each vertex of P that has angle $< 30^\circ$. The current paper reduces the interval to $[36^\circ, 72^\circ]$, by allowing a bounded number of exceptional triangles in a neighborhood of each vertex of P that has angle $\leq 36^\circ$.

Section 2 gives an overview of the proof Theorem 1.1, and later sections provide the details. We end the paper with some open questions and remarks. I thank Joe

Mitchell and Herbert Edelsbrunner for their helpful comments on an earlier version of this paper, and for their encouragement. Several figures are drawn using Toby Driscoll's SC-Toolbox package for MATLAB [27], an updated version of code by Nick Trefethen [61] for computing Schwarz-Christoffel maps. I thank Toby for his assistance with the toolbox. The results in this paper were announced at SODA 2022, and the referees judging submissions provided many helpful suggestions. I am also very thankful to several other anonymous referee who provided numerous suggestions to improve the mathematics and exposition. Their very thorough and thoughtful reports were extremely helpful, and this paper greatly benefited from their heroic efforts.

2. OVERVIEW OF THE PROOF OF THEOREM 1.1

As noted in the introduction, given an N -gon P we want to define an associated 60° -polygon P' , and to transfer a (nearly) equilateral triangulation from P' to P via a conformal map. The simplest case is when this map induces a bijection between the vertices of P' and P , although later we will map P' to a polygon P with some slits removed, and allow some vertices of P' to map to interior points of P . In either case, if the triangles in P' are small enough, the conformal map will not distort their angles very much except near the vertices of P' , where the map may have singularities. To minimize the angle distortion near these vertices, we would like each interior angle of P' to be close to the corresponding angle of P . However, taking the closest multiple of 60° will not always work, as this may violate the requirement that the interior angles of P' must sum to $180^\circ(N - 2)$. In order to motivate the technicalities in the next few sections, in this section we briefly discuss how we build P' , although the actual construction is delayed until Sections 9 and 10.

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 the left side of Figure 8.

It will be convenient to also consider the larger 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 $[60^\circ - \epsilon, 60^\circ + \epsilon]$

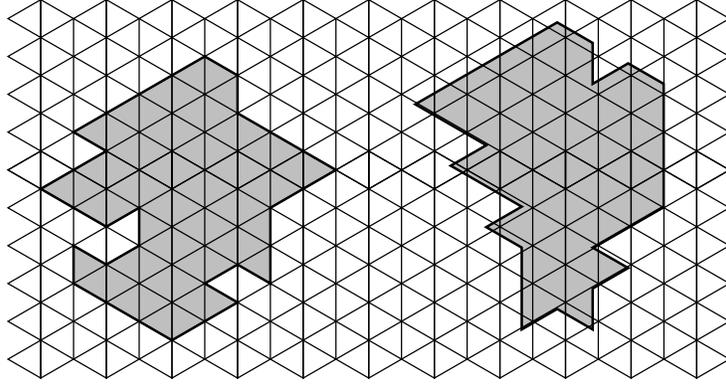


FIGURE 8. On the left is a equilateral-grid polygon, and on the right is a 60° -polygon.

and that each vertex of P has a neighborhood in which the triangulation elements are actually equilateral (this extra condition 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 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 will arise as Schwarz-Christoffel images of the disk in Section 9 when all the angle parameters are multiples of 60° , but the conformal map is not globally 1-to-1.

Suppose that $f : \Omega' \rightarrow \Omega$ is a conformal mapping between the interiors of two polygons P' and P that induces a bijection between vertices of P' and vertices of P . (Below, we will sometimes use P to refer to both the boundary curve and the interior domain, instead of using Ω for the latter; the intended meaning should always be clear from context.) In Corollary 3.3 we will prove that f only slightly perturbs the angles of sufficiently small triangles in Ω' , unless they are near vertices of P' . If v'_k is a vertex of P' with angle ψ_k , so that v'_k maps to a vertex v_k of P with angle θ_k , then Lemma 6.1 will imply that any small triangle close enough to v'_k will have its interior angles distorted by at most a factor of $\max_k(\theta_k/\psi_k, \psi_k/\theta_k)$. To minimize overall angle distortion we want to choose P' so that this factor is as close as possible to 1.

The triangulations of P we construct will have most of their angles between 36° and 72° ; the exceptions will only occur near vertices of P that have angle less than

36° . Angles of P larger than 72° will be subdivided by the triangulation to give smaller angles in the interval $[36^\circ, 72^\circ]$, and we want these sub-angles to each map to 60° under the conformal map between P and P' .

In our construction, small angles in P (say $\leq 72^\circ$) will correspond to angles of size 60° in P' . We will want each larger angle θ in P to correspond to an angle $\psi = L(v) \cdot 60^\circ$ in P' that satisfies

$$(2.1) \quad \frac{3}{5}\psi \leq \theta \leq \frac{6}{5}\psi.$$

This allows each angle $\theta \in (0, 360^\circ]$ to be associated to at least one multiple of 60° (and often to several). The restrictions imposed by (2.1) are summarized by Table 1. For example, if P has a vertex v with interior angle $\theta = 135^\circ$ the corresponding vertex v' in the 60° -polygon P' can have angle either 120° or 180° .

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 $\Phi(P) \leq 72^\circ$.

Unfortunately, there are polygons P whose angles $\{\theta_k\}$ cannot be mapped to the angles $\{\psi_k\}$ of any 60° -polygon P' , in a way that satisfies both the constraints of Table 1 and the requirement (valid for all N -gons): $\sum_k \psi_k = (N - 2)180^\circ = \sum_k \theta_k$. 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 $480^\circ > 360^\circ$. However, we can “fix” the angle discrepancy $\sum \psi_k \neq \sum \theta_k$ by adding extra vertices to the edges of P . Note that $\sum \psi_k$ is always a multiple of 60° (since each term is) and $\sum \theta_k = (N - 2)180^\circ$ is clearly a multiple of 180° . Thus the difference $\sum \theta_k - \sum \psi_k$ is a multiple of 60° .

First suppose $\sum_k \psi_k < \sum_k \theta_k$. We add a new vertex v of angle 180° in the interior of an edge of P , and we assign the corresponding vertex v' in P' the angle 240° . See Figure 9. Doing this increases the angle sum $\sum \theta_k$ by 180° but increases the

angle sum $\sum \psi_k$ by 240° , thus 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 under the conformal map). Thus, at best, this method gives ϕ -triangulations with $\phi \geq 67.5^\circ$. This suffices for Case 1 of Theorem 1.1 but a more elaborate construction is needed for Cases 2 and 3.

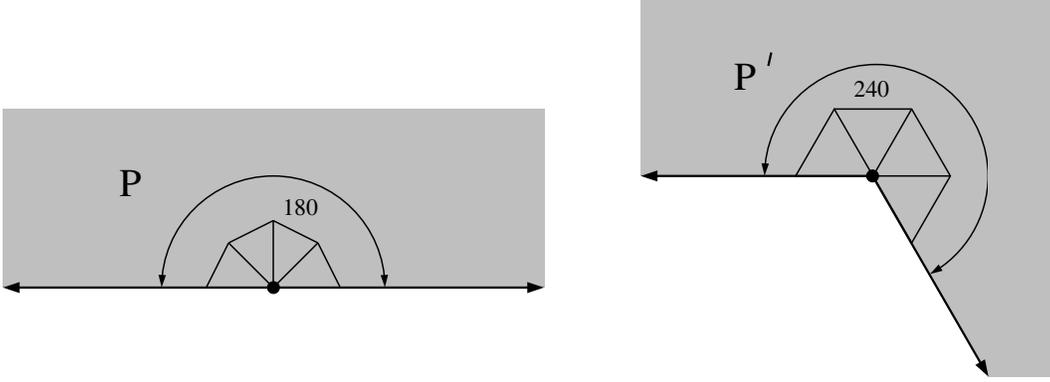


FIGURE 9. Our first trick for increasing the ψ -sum relative to the θ -sum is to pair a 180° -vertex in P with a 240° -vertex in P' . The conformal map locally looks like $z^{3/4}$, and maps a 60° sub-angle to 45° . A triangle containing this angle also contains an angle ≥ 67.5 .

In order to prove Case 2 of Theorem 1.1 we will need to lower the discrepancy between $\sum \psi_k$ and $\sum \theta_k$, while only introducing angles smaller than $\frac{5}{7} \cdot 90^\circ \approx 64.2857$ (instead of 67.5° as in the previous paragraph). We will use a triangulation of a slit half-plane based on transferring an equilateral triangulation from a polygonal Riemann surface that has a 420° interior angle. The idea is shown in Figure 10; 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 21. R has a 1-1 projection onto a sector of angle 240° , except for the darker triangle where the projection is 2-1. Traversing the boundary, we encounter interior 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.

The two segments adjacent to the 420° -vertex are mapped to a (curved) slit in P where the interior angles are 60° , 360° and 120° . The worst distortion comes from

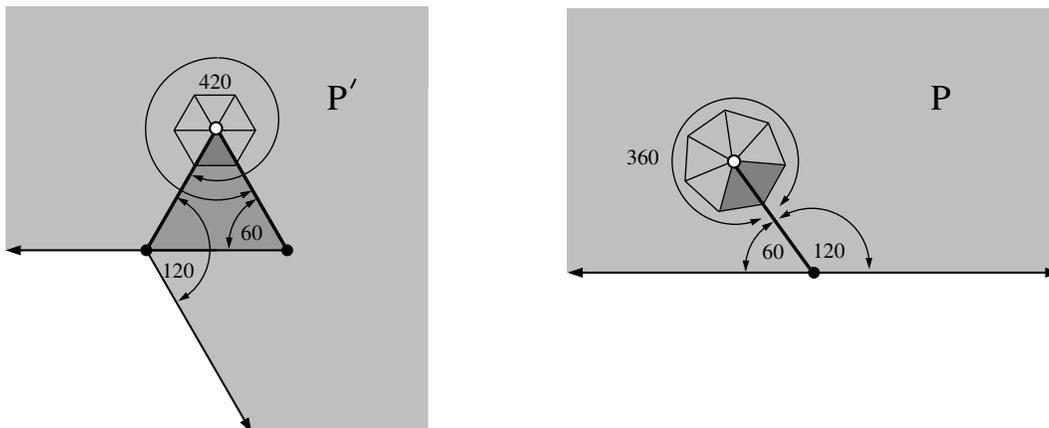


FIGURE 10. We cut a slit in the upper half-plane at angle 60° (right hand picture). The angles we observe tracing the outline of the slit in P are: 60° , 360° and 120° . The triangulation near this slit will correspond to an equilateral triangulation of a Riemann surface R with angles 60° , 420° and 120° , pictured at left. Near infinity the map $P' \rightarrow P$ looks like $z^{3/4}$.

mapping the 420° -vertex v' in P' to the 360° -vertex $v \in P$ (the tip of a slit). Locally the map looks like $z^{6/7}$ near v' , and it maps each 60° sub-angle to $\frac{6}{7} \cdot 60^\circ = \frac{4}{7} \cdot 90^\circ \approx 51.4286$. Since the angles of a triangle sum to 180° , a triangle in P containing v must also contain an angle $\geq \frac{5}{7} \cdot 90^\circ \approx 64.2857$. This 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. See Figure 22. This requires that the conformal map sends the length measures on the two segments to identical measures on the slit (but not necessarily length measure). We will explain how to choose the slit in Section 8.

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 11. Traversing the boundary of the slit half-plane we encounter interior 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. Exactly one degree five interior vertex is created in the triangulation, located at the tip of the slit. See also Figure 19.

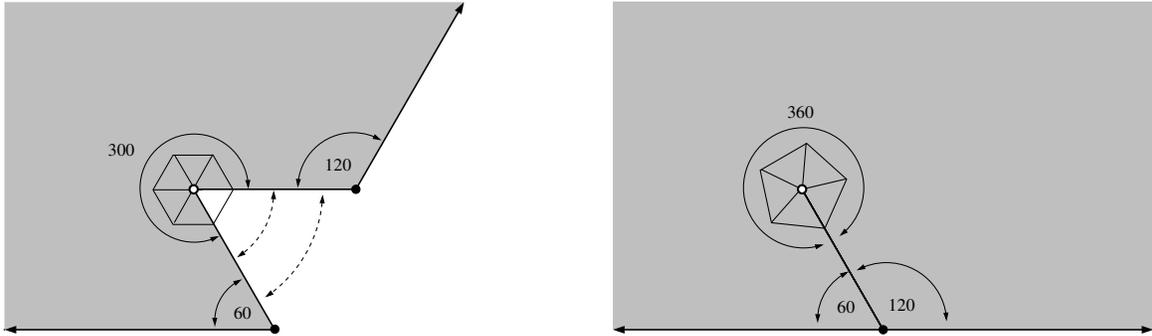


FIGURE 11. Take a 120° -sector and cut out an equilateral triangle (left). Conformally map this region to a slit half-plane (right). The two edges adjacent to the 300° -vertex are mapped to opposite sides of the (slightly) curved slit, and points on these edges equidistant from this vertex are identified. This implies that an equilateral triangulation on the left pushes forward to a triangulation on the right.

Sections 3-8 will provide more details of the various constructions sketched above. The proof of Theorem 1.1 and its 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 it 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 that is not homeomorphic to the 2-sphere or conformally equivalent to the plane. This is true for all the Riemann surfaces 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 an 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*

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

Proof. By replacing g by a scalar multiple of itself, we may assume that $\operatorname{Re}(g) < 1$ 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 $\tau^{-1}(D(0, \frac{1}{2})) \subset D(0, 2)$. \square

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 sense. 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 12. 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.

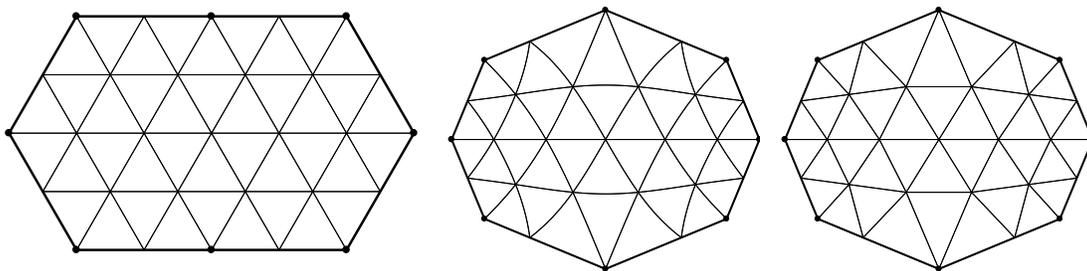


FIGURE 12. An equilateral triangulation (left), the actual conformal images of the triangles (center), and the pushed forward triangles (right) where vertex images are connected by segments.

Lemma 3.2. *If f is conformal on the disk $D(z, r)$, if $\delta > 0$ is sufficiently small, and if $T = \Delta ABC \subset D(z, \delta r)$ is a triangle, then $f^*T = \Delta f(A)f(B)f(C)$ 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 [36]), says that

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

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

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

if $|z| \leq 1/4$. Therefore, inside $D(0, \delta)$ any segment S is mapped to a smooth curve all of whose tangents are within 8δ of 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 . \square

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 $\text{diam}(T) < \delta$.*

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 enough. \square

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 $\mu = h_{\bar{z}}/h_z$

satisfies $\|\mu\|_\infty < 1$ (recall $h_z = \frac{1}{2}(\frac{\partial h}{\partial x} - i\frac{\partial h}{\partial y})$ and $h_{\bar{z}} = \frac{1}{2}(\frac{\partial h}{\partial x} + i\frac{\partial h}{\partial y})$). See [2] or [48] for the basic properties of such maps. (Although general quasiconformal maps need not be differentiable everywhere that they are defined, in our particular applications they will be.) 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) sends circles to ellipses of eccentricity at most K (the eccentricity of an ellipse is the length ratio of the major to minor axis).

One important 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 μ may change, but its absolute value does not. Another useful fact about K -quasiconformal maps $\mathbb{C} \rightarrow \mathbb{C}$ 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 $\{z : |z| < 1/\epsilon\}$ into \mathbb{C} that fixes 0 and 1, then $|f(z) - z| \leq \delta|f(0) - f(1)|$ for all $z \in \mathbb{D}$.*

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 $\{z : |z| < n\}$ into \mathbb{C} with $K_n \rightarrow 1$, 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, since it fixes 0 and 1. This contradicts the assumption that every map moves some point in \mathbb{D} by at least δ . \square

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 and a point $z \in \Omega$ the Riemann mapping theorem states that there is a 1-to-1, onto holomorphic map f from the disk $\mathbb{D} = \{|z| < 1\}$ to Ω so that $f(0) = z$ and $f'(0) > 0$. If the boundary $\partial\Omega$ is locally connected, then Caratheodory's theorem says that f extends continuously to $\mathbb{T} =$

$\partial\mathbb{D}$, and if $\partial\Omega$ is a Jordan curve (e.g., a simple polygon) then the extension is a homeomorphism. In the latter case, given a Borel set $E \subset \partial\Omega$, we define the harmonic measure of E as $\omega(z, E, \Omega) = |f^{-1}(E)|/2\pi$, where $|f^{-1}(E)|$ denotes the Lebesgue length of $f^{-1}(E) \subset \mathbb{T}$. (For general simply connected domains, the conformal map f need not have a continuous boundary extension, but it always has radial boundary limits almost everywhere, and these can be used to define harmonic measure for any simply connected domain. However, we will not need this more general approach here.)

More intuitively, the harmonic measure of a boundary set E is the probability that a Brownian motion started at z first hits $\partial\Omega$ in the set E . Harmonic measure is also the value at z of the harmonic function u on Ω that has boundary value 1 on E and zero elsewhere (appropriately defined). The text [36] of Garnett and Marshall is a comprehensive treatment of conformal maps, their boundary values, harmonic measure, and it includes all the results we shall use below.

An infinite sector of angle θ is a region congruent to $S(\theta) = \{re^{i\phi} : r > 0, -\theta/2 < \phi < \theta/2\}$. 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 $\theta \in (0, 360^\circ]$. For $r > 1$, let $A(r) = \{z : 1/r < |z| < r\}$. An annular sector is a region congruent to $S(\theta, r) = S(\theta) \cap A(r)$. Given a sector $S(\theta)$, 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 and Ω_2 are simply connected domains so that for $k = 1, 2$, the regions $\Omega_k \cap A(r)$ both have connected components equal to $S(\theta, r)$. Suppose $F : \Omega_1 \rightarrow \Omega_2$ is the conformal map that fixes z_0 and z_1 (defined above). For any $\epsilon > 0$ there is an r_0 so that if $r > r_0$, then $\sup_{S(\theta, 2)} |F(z) - z| < \epsilon$.*

In other words, if we think of Ω_1 and Ω_2 as a family of domains depending on the parameter r , then the corresponding map $F : \Omega_1 \rightarrow \Omega_2$ (which also depends on r) tends uniformly to the identity on $S(\theta, 2)$ as $r \nearrow \infty$.

Proof of Lemma 4.2. It is enough to take $\Omega_2 = S(\theta, r)$; in general, the conformal map $\Omega_1 \rightarrow \Omega_2$ can be written as the composition of two conformal maps $\Omega_1 \rightarrow S(\theta, r) \rightarrow \Omega_2$, so it is enough to show each of these is close to the identity.

Let $f : \Omega_1 \rightarrow \mathbb{D}$ be the conformal map sending z_0 to 0 and z_1 to i . See Figure 13. By standard estimates, e.g., Theorem IV.6.2 of [36], the harmonic measure of the circular arcs of $S(\theta) \cap \partial A(r)$ 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. The images of these curves are crosscuts of \mathbb{D} (i.e., arcs in the open disk \mathbb{D} with both endpoints on the boundary $\partial\mathbb{D} = \mathbb{T}$). By the Beurling projection theorem (e.g., Theorem II.9.2 of [36]) these crosscuts both 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 $F = f^{-1} \circ g \circ f : \Omega_1 \rightarrow S(\theta, r)$ is conformal, fixes z_0 and z_1 , and tends to the identity uniformly on $S(\theta, 2)$ as $r \nearrow \infty$. \square

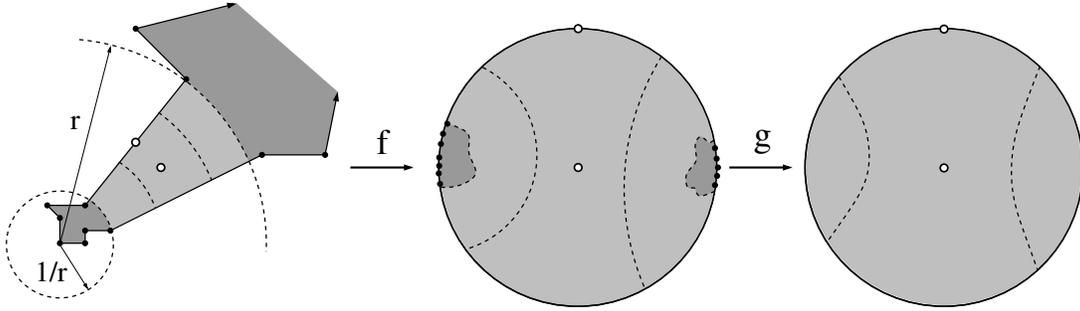


FIGURE 13. The conformal map from Ω to $S(\theta, r)$ can be written as a composition $F = f^{-1} \circ g \circ f$ where $f : \Omega \rightarrow \mathbb{D}$ and g is a map of $\mathbb{D} \setminus f(\Omega \setminus S(\theta))$ to \mathbb{D} . The components of $f(\Omega \setminus S(\theta))$ have small diameter as $r \nearrow \infty$, so both g and $F = f^{-1} \circ g \circ f$ are close to identity maps.

Lemma 4.3. *Suppose Ω_1 and 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 $\sup_{S(\theta, 3)} |f_1 - f_2| < \epsilon$. If I is either one of the two radial segments of $\partial S(\theta, 3)$, suppose that f_1, f_2 both map I into the same line. Let $\eta : [0, \infty) \rightarrow [0, 1]$ be increasing and 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 and f_2 can be extended by reflection across the radial boundary segments of $\partial S(\theta, 3)$ 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 $\mu = \partial_{\bar{z}}g/\partial_zg$. First,

$$\begin{aligned} \partial_zg(z) &= f'_1(z)(1 - \eta(|z|)) - f_1(z)\partial_z\eta(|z|) + f'_2(z)\eta(|z|) + f_2(z)\partial_z\eta(|z|) \\ &= f'_1(z) + (f'_2(z) - f'_1(z))\eta(|z|) + (f_2(z) - f_1(z))\partial_z\eta(|z|) \\ &= f'_1(z) + O(\epsilon) \simeq 1, \end{aligned}$$

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 $|\mu_g| = |\partial_{\bar{z}}g/\partial_zg| = O(\epsilon)$ as desired. \square

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 $f : \Omega_1 \rightarrow \Omega_2$ is conformal and $f(0) = 0$, $f'(0) = 1$. Let $\eta : [0, \infty) \rightarrow [0, 1]$ be smooth with $\eta(r) = 0$ if $r \leq 1/2$ and $\eta(r) = 1$ if $r \geq 2$. 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 $|z| \leq 2$. However, this is the Distortion Theorem for conformal maps, e.g., Theorem I.4.5 and Equation (I.4.17) of [36]. \square

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 (independent of ϵ). Also recall that we need this result for both planar polygonal regions and Riemann surfaces with polygonal boundaries.

Lemma 5.1. *Every 60° -surface R has nearly equilateral triangulations. In particular, $\Phi(R) = 60^\circ$.*

Proof. We will 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 sides of the usual equilateral grid in \mathbb{C} (with one side of each triangle parallel to the vertical axis). Rescale the grid to have 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 14. 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.

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 15.

Let $F : P' \rightarrow P$ 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

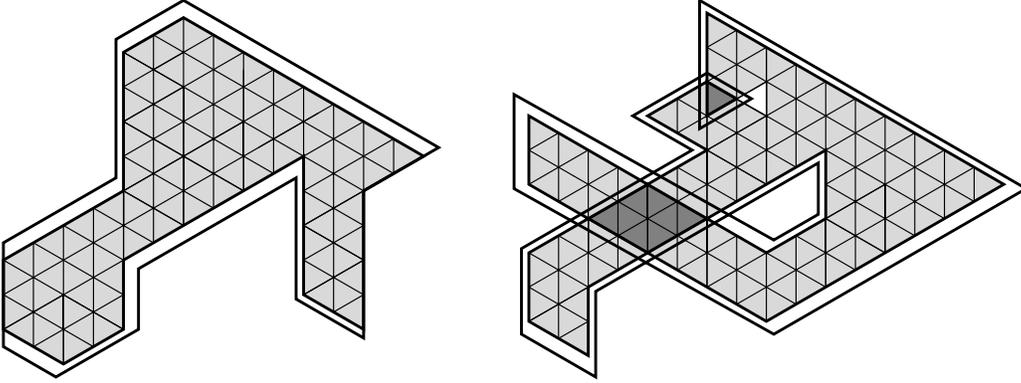


FIGURE 14. Two examples of approximating a 60° surface (white) by a union of equilateral triangles (shaded). On the left is simple polygon and on the right a Riemann surface. Overlaps of the surface with itself have darker shading.

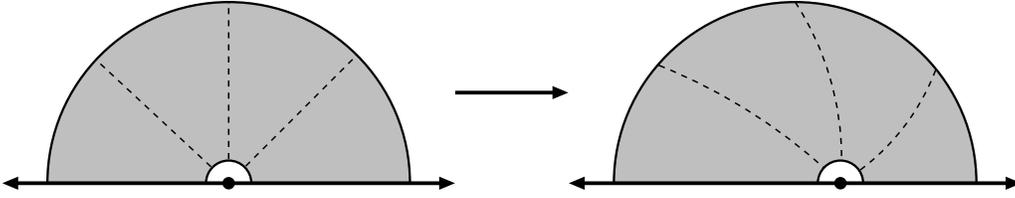


FIGURE 15. Pre-vertices of P and P' are paired and can be identified by a quasiconformal map Φ with small dilatation supported inside an annular region surrounding each pair.

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 $\theta = k \cdot 60^\circ$. Set $\alpha = 3/k$. Then $z \rightarrow z^\alpha$ 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 origin. By the Schwarz reflection principle, 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. \square

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 $\Omega_1 \rightarrow \Omega_2$ 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 SECTORS

As before, $S(\theta)$ denotes the infinite sector of angle θ with the 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 16. This triangulation extends to a triangulation of the right half-plane, and in this section we record a computation that gives angle bounds for images of this triangulation under conformal maps of the form $z \rightarrow z^\alpha$. See Figure 17 for two examples.

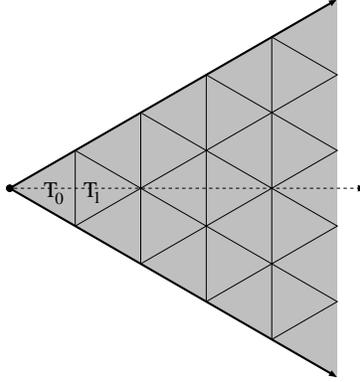


FIGURE 16. The sector $S(60^\circ)$ and its equilateral triangulation \mathcal{G} .

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 . If τ is the interior angle of some angle in T , and τ' is the corresponding angle of the image triangle f^*T , then $|\tau - \tau'| \leq |\alpha - 1| \cdot \theta$, where θ is the angle subtended by T from the origin.*

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 . \square

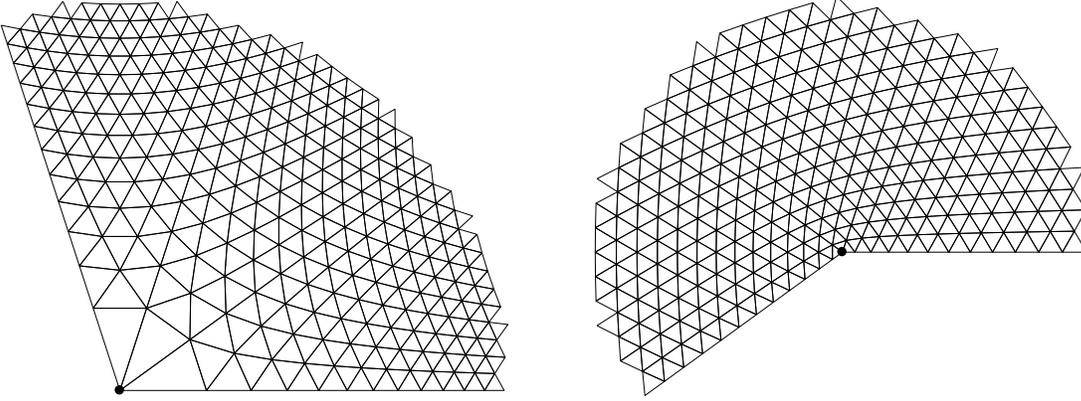


FIGURE 17. Image of the equilateral grid in the upper halfplane under the maps $z^{3/5}$ and $z^{6/5}$. These are the extreme values that keep images of 60° angles between 36° and 72° .

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 $(\alpha - 1)30^\circ$. Since $\phi \geq 60^\circ$, these angles are bounded by

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

The pushed forward triangle f^*T_0 is isosceles with vertex angle $\phi = \alpha \cdot 60^\circ$, 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 (recall $30^\circ = \frac{1}{2}60^\circ \leq \phi/2$),

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

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

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

as desired. The triangle T_0 maps to an isosceles triangle with angle ϕ at the origin, and two base angles equal to $90^\circ - \phi/2$; these are the endpoints of the desired interval. 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 $< 30^\circ$. Therefore 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

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

For brevity, we will refer to the “image of an equilateral triangulation of a θ -sector under a power map z^α ” more simply as a “sector triangulation” or “sector mesh”. When we need to be more precise, we refer to this as an (θ, α) -sector mesh.

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.

Lemma 6.3. *Suppose $f : P' \rightarrow P$ is a conformal map between polygons that maps vertices to vertices. Suppose $f(V') = v$ where w 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*

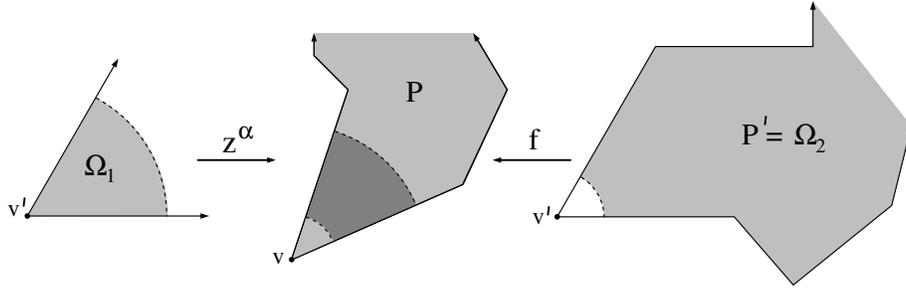


FIGURE 18. Proof of Lemma 6.3. Near each vertex v of P the conformal push forward of the nearly equilateral mesh on P' gives the correct angle bounds in the limit as the triangulation gets finer. To get the correct angle bounds with a finite triangulation, we merge $f^*\mathcal{T}$ with a sector mesh near each vertex.

is a neighborhood U of v and a triangulation \mathcal{S} of P that equals $f^*\mathcal{T}$ outside U and so that every triangle of \mathcal{S} touching U has all angles $\leq \max(\theta/k, 90^\circ - \theta/2k)$.

Proof. Let Ω_1 be the infinite sector with vertex v' (the same as the chosen vertex of P') and angle θ , so that Ω_1 agrees with P' in neighborhood of v' . We apply Lemma 4.3 with this Ω_1 and taking $\Omega_2 = P'$. Choose an annular region $\lambda A(r) + v'$ around v' , by first choosing r large and then λ small, as shown in Figure 18. By 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

Next we discuss the “120°-trick”² that decreases $\sum \psi_k$ (the sum of the interior angles of P') by 60° relative to $\sum \theta_k$ (the angle sum for P). This involves adding a vertex v to an edge of P (hence this vertex has interior angle 180°), so that v is paired with a vertex v' of angle 120° in P' . However, near v the triangulation of P will not be an image of a triangulation of P' near v' . Instead, we cut a slit in P with v as one endpoint, and near this slit we use a triangulation that is a conformal image of a “notched 120°-sector”, described below. The basic idea was introduced in Sections 1 and 2, but here we describe the construction more precisely.

²A referee suggested that “300°-trick” would be more natural, but I have stuck with my original name since the method involves a 120° sector.

Consider the region Ω shown on the left in Figure 11. This is a 120° -sector with an equilateral triangle at the origin removed. Translate the region so the 300° -vertex is at the origin. We apply a branch of $z^{6/5}$, and the 300° angle at the origin becomes 360° , and the two finite segments in $\partial\Omega$ adjacent to the origin become identified with a radial slit in the image. The two infinite rays in $\partial\Omega$ map to the boundary curve of a simply connected region Ω' . See lower left in Figure 19. 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 $\approx 2.75^\circ$). 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 possibly near either end of the slit.

Actually, near v there is very little distortion. Near this point, P minus the slit looks like the union of a small 60° -sector and a 120° -sector; the shared edge is slightly curved, but the angles at v are exactly 60° and 120° , since these are conformal images of actual truncated sectors with these same angles. Therefore the triangulation near v is close to equilateral (if the triangulation is fine enough), and hence the angles are all less than 72° , as desired.

In a neighborhood of the interior tip of the slit, angles are bounded above by $72^\circ + \epsilon$ if the mesh is fine enough, but might exceed 72° due to distortion of the conformal map. However, we can replace the mesh near the tip by the standard $(300^\circ, 6/5)$ -sector triangulation using Lemma 4.4. This gives the 72° bound for the whole triangulation of the upper half-plane.

We use this triangulation of the half-plane only in a small neighborhood of v . In an annulus around v (and much larger than the slit we have constructed) we merge this mesh with the $(120^\circ, 3/2)$ -sector triangulation, and on a larger annulus this sector mesh can be merged with the conformal image of a nearly equilateral mesh of P' moved to P by a conformal map between these domains. This is made precise in the following lemma.

Lemma 7.1. *Suppose $f : P' \rightarrow P$ 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*

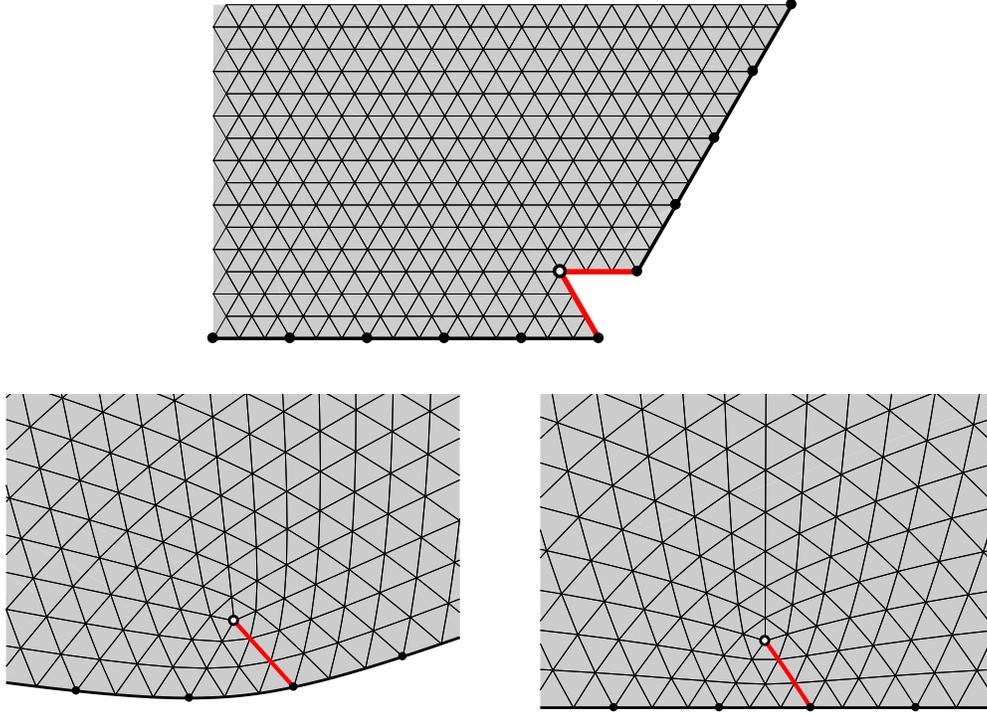


FIGURE 19. The cut 120-sector (top) is mapped to a simply connected region Ω' (lower left) by a branch of $z^{6/5}$. This is then mapped to the upper half-plane by a conformal map. Because the power map identifies points according to distance to zero (the white dot), the push forward of the triangulation is still a triangulation.

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 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 $\leq 72^\circ$.

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 20. \square

8. THE 420° -TRICK

There are (at least) two methods we can use to increase the ψ -sum for P' by 60° with respect to 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

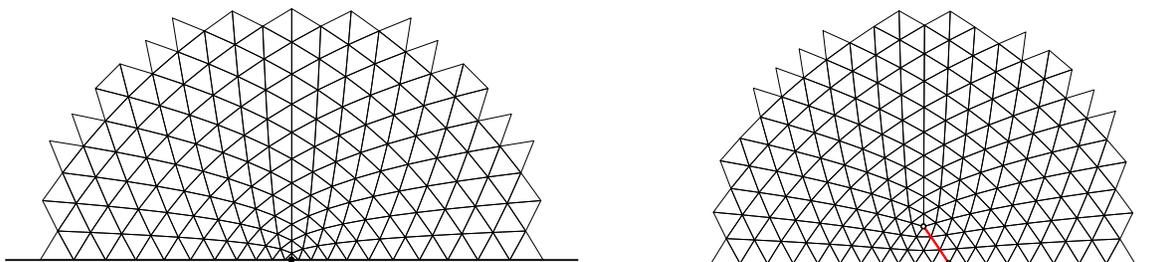


FIGURE 20. The left shows the $(120^\circ, 3/2)$ -sector mesh of the upper half-plane. The right shows the triangulation coming from the “120°-trick”. These two meshes can be merged using quasiconformal interpolation as described in the text.

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 is mapped to a triangle with an angle of size 45° at v . The opposite angles in the image triangles are $67.5^\circ < 72^\circ$, 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 that introduces triangulation angles no larger than $\frac{5}{7} \cdot 90^\circ \approx 64.2857^\circ$. This is precisely the angle bound we get if a 420° -vertex $v' \in P'$ is mapped to a 360° -vertex $v \in P$. We accomplish this by considering a non-planar Riemann surface.

The idea is illustrated in Figure 21. 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 with interior angles 60° , 420° and 120° .

We can conformally map R to a slit upper half-plane in two steps as illustrated in Figure 22 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

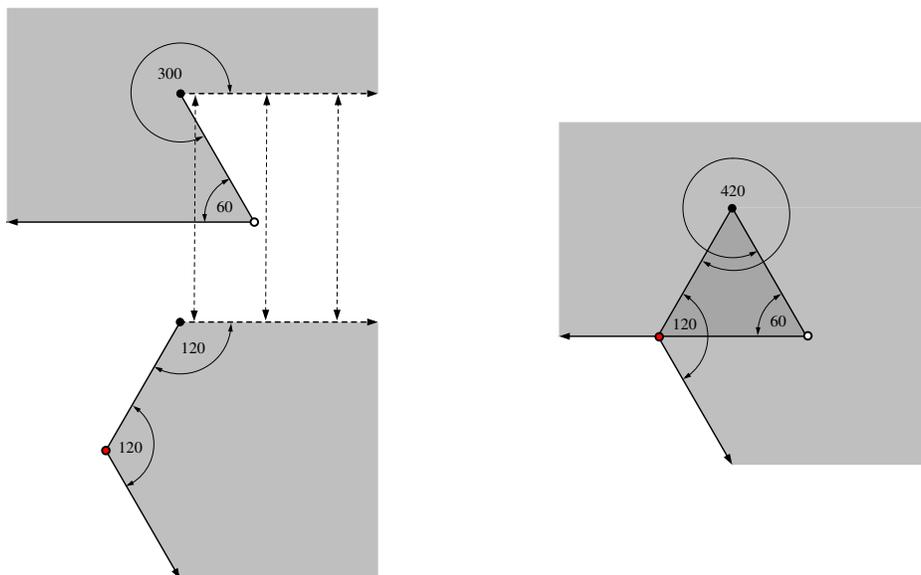


FIGURE 21. Building a 60° -surface with an 420° interior angle. The surface can be constructed from two planar domains glued along the dashed edges of each as illustrated on the left. The darker triangle indicates where the surface has two sheets over the plane.

of R will be mapped to triangulations of the upper half-plane. Near the tip of the slit, we replace the triangulation constructed above with sector triangulation, just as in the previous section. Near v , the triangulation constructed above has all angles close to 60° (if the triangulation is fine enough), again for the same reasons as in Section 7.

9. SUFFICIENCY IN THEOREM 1.1: CASE 1

In this section, we finally build the model 60° -polygon P' that we have discussed throughout the paper. In order to simplify the discussion, we use a rather artificial P' that gives the correct angle bounds, but that does not “look like” P and that often creates a triangulation with many more triangles than necessary. In Section 16 we will discuss why this happens, and how more efficient model polygons might be constructed.

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 triangulation in a neighborhood

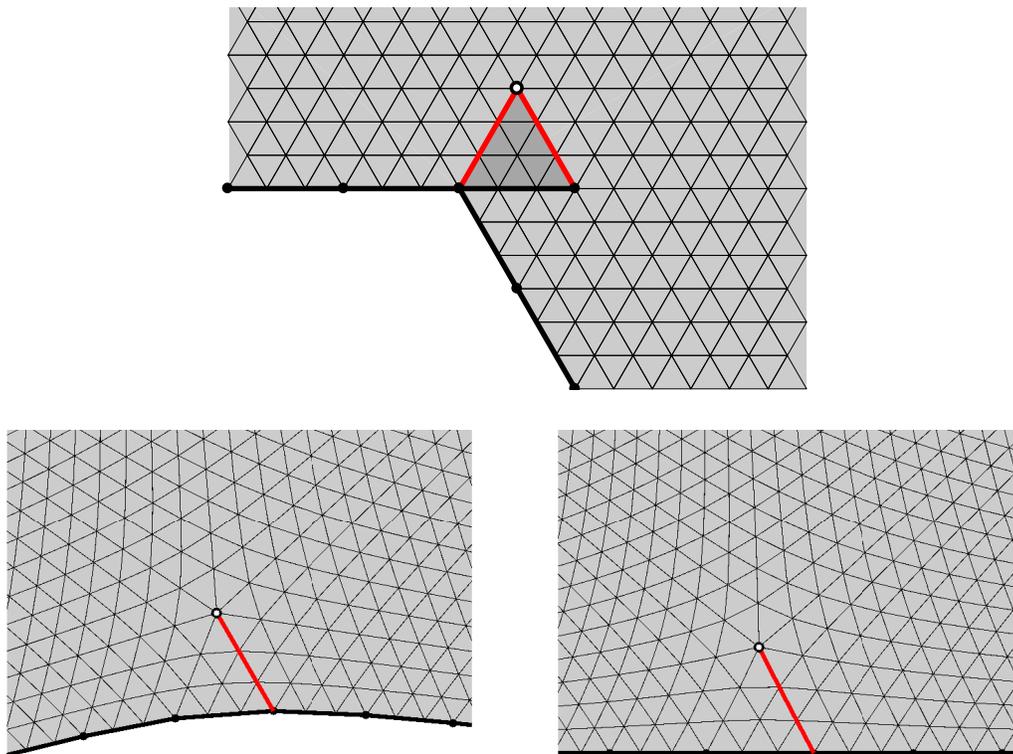


FIGURE 22. The Riemann surface R can be conformally mapped to a slit domain Ω by a branch of $re^{i\theta} \rightarrow re^{i6\theta/7}$. This map is followed by a conformal map to the upper half-plane that bends the the straight slit to an analytic arc. (But it looks quite straight; the tangent directions only change by about 1° along the arc.)

of a few boundary points, but the 420° -trick is not needed until we give the proof for 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 $\alpha_j \cdot 180^\circ$. 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

$$(9.1) \quad f(z) = A + C \int_0^z \prod_{j=1}^N \left(1 - \frac{w}{z_j}\right)^{\alpha_j-1} dw,$$

for some appropriate choice of constants A, C . (The integral is a line integral on the segment from 0 to z , but it would give the same value for any rectifiable curve in \mathbb{D}

connecting 0 to z , by Cauchy's theorem.) See e.g., [28], [53], [62]. The formula was discovered independently by Christoffel in 1867 [22] and Schwarz in 1869 [59], [58]. For other references and a brief history, see Section 1.2 of [28]. Given a polygon P , the angles are known, but the prevertices must be solved for. See [10] for a provably convergent algorithm for finding the Schwarz-Christoffel parameters, and for citations to other methods used in practice.

Given N distinct points $\mathbf{z} = \{z_1, \dots, z_N\}$ on the unit circle and N real values $\{\alpha_1, \dots, \alpha_N\}$ summing to $N - 2$, Formula (9.1) defines a holomorphic function on the disk that maps each component of $\mathbb{T} \setminus \mathbf{z}$ to a line segment, with the segments meeting at $f(z_k)$ making interior angle $\alpha_k \cdot 180^\circ$. The map given by (9.1) is always locally 1-1 on \mathbb{D} (since f' never vanishes), but it 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 , consider a conformal map f of its interior to the unit disk, \mathbb{D} . The N vertices of P map to N distinct points $\mathbf{z} = \{z_1, \dots, z_N\}$ on the unit circle \mathbb{T} . We then want to choose N real values $\psi_k \in Z = \{60^\circ, 120^\circ, 180^\circ, 240^\circ, 300^\circ\}$ so that $\sum_k \psi_k = 180(N - 2)$. If such a choice is possible, we then set $\alpha_k = \psi_k/180$ and apply the Schwarz-Christoffel formula to get a map $g : \mathbb{D} \rightarrow P'$. Then $g \circ f : P \rightarrow P'$ 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 .

For the general case, 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 each is assigned to have angle $\psi_v = 120^\circ$ in P' . Assign angle 180° to all other vertices of P' , so the ψ -angle sum is $6 \cdot 120^\circ + 180^\circ N = (M - 2)180^\circ$. Applying Schwarz-Christoffel gives a 60° -hexagon, as in Figure 23.

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

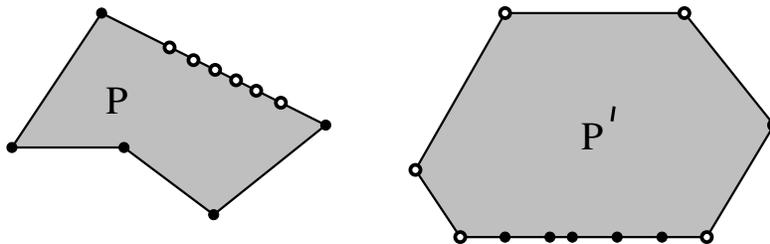


FIGURE 23. In the proof of Case 1 of Theorem 1.1 we can assume P' is planar. The first step is to choose six “artificial” vertices on one edge of P (left) and make these correspond to six 120° -vertices in P' (right).

The new vertices associated to v will all 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 will 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 listed below and illustrated in Figure 24. Suppose v is an original vertex of P with interior angle θ_v :

- (i) if $0 < \theta_v \leq 72^\circ$, set $\psi_v = 60^\circ$ and add two vertices each with angle 240° ,
- (ii) if $72^\circ < \theta_v \leq 144^\circ$, set $\psi_v = 120^\circ$ and add one vertex with angle 240° ,
- (iii) if $144^\circ < \theta_v \leq 216^\circ$, set $\psi_v = 180^\circ$ and add no associated vertices,
- (iv) if $216^\circ < \theta_v \leq 288^\circ$, set $\psi_v = 240^\circ$ and add one vertex with angle 120° ,
- (v) if $288^\circ < \theta_v \leq 360^\circ$, 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 25. In particular, P' is not self-intersecting and so it 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).

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. In these neighborhoods, the angles are all between $36^\circ - \epsilon$ and $72^\circ + \epsilon$, where ϵ can be made as small as we wish by taking

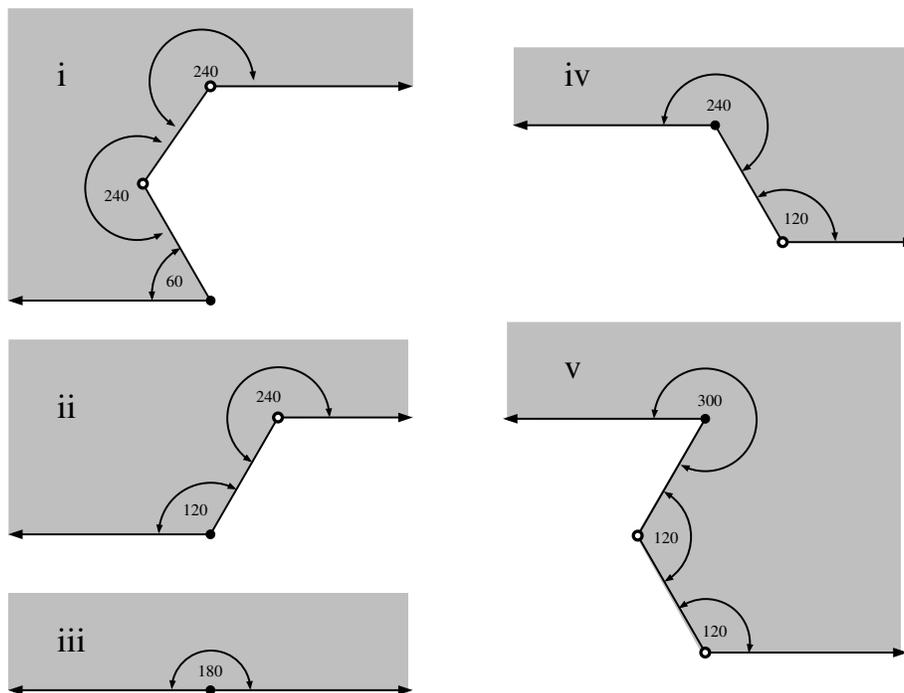


FIGURE 24. The five cases for assigning image angles and associated vertices. In each case the black indicates the image of the original vertex and the white dots the new associated vertices. These arcs represent small subarcs of the 60° -polygon P' . The associated vertices all map to 180° vertices that we add to the edges of P .

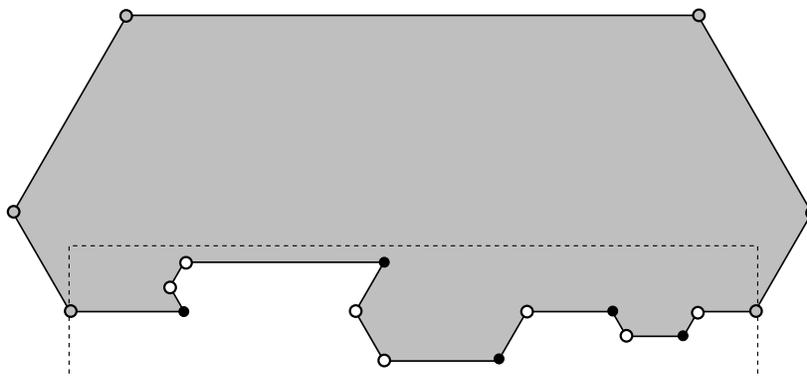


FIGURE 25. The 60° -polygon P' . The six gray points were added at the beginning, the black points are the images of the original vertices, and the white points are the associated vertices.

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 sharper bounds $\geq 36^\circ$ and $\leq 72^\circ$ 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. \square

Note that we have also proven Corollary 1.4. The upper bound in Corollary 1.4 implies every angle in the triangulation is at least $\min(\theta_{\min}, 36^\circ)$. However, the proof of Theorem 1.1 actually proves the following, slightly stronger, lower bound:

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

This formula is somewhat clearer when graphed, as in Figure 26. The first term follows directly from Corollary 1.4. 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 $144^\circ \leq \theta \leq 162^\circ$ can be divided into three angles of size $\theta/3$ and larger angles can be divided into three or more angles $\geq 54^\circ$. 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., $\theta_{\min} = 80^\circ$ implies 40° is attained by any 72° -triangulation of P . If $\theta_{\min} \geq 162^\circ$, then both 72° and 54° must be attained in an optimal triangulation.

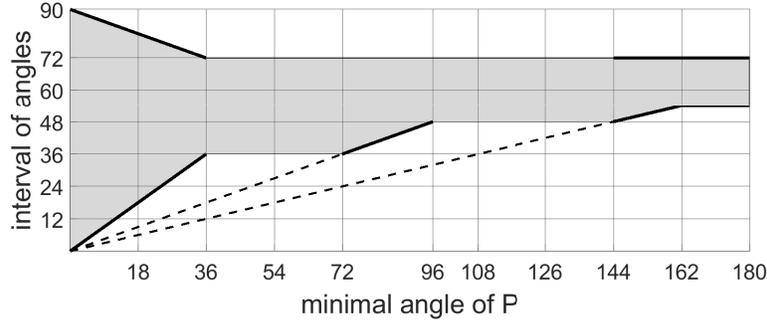


FIGURE 26. If P has minimal angle $0 < \theta_{\min} < 180^\circ$ (plotted on the horizontal axis), then P has a triangulation with all angles inside the shaded interval above θ_{\min} .

10. 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 given in the previous section, so that it avoids using the 120° -trick in Case 2 (this would force an angle $\geq 72^\circ$), and to avoid both the 120° -trick and the 420° -trick in Case 3 (the latter forces angles $\geq \frac{5}{7} \cdot 90^\circ$).

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,

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

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

Proof of Corollary 1.17 (number of slits). 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 $< 67.5^\circ$, since using a 240° sector mesh will work for larger angles. The number of times we need to apply the 420°-trick is $(\sum_P \theta_v - \sum_{P'} \psi_v)/60^\circ = -\kappa(\phi)$, so this is the number of degree 7 interior vertices.

One degree 5 interior 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. \square

11. SOME EXAMPLES

Proof of Corollary 1.8 (regular polygons). 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 $\Phi_3 = \Phi_6 = 60^\circ$. Otherwise, there are N interior angles of size $\theta_N = 180^\circ - 360^\circ/N \geq 90^\circ$. Since these are all $> 36^\circ$, Corollary 1.4 says $\Phi_N \leq 72^\circ$, so we only have to check whether an even smaller bound is possible.

For $N \geq 10$, the interior angles are $\geq 144^\circ$, and hence the ϕ -admissible labels for $\phi < 72^\circ$ are at least 3. Then any ϕ -admissible label satisfies $\kappa(L) \geq 6 > 0$, so we must be in Case 1, so $\Phi_N = 72^\circ$. (This case also follows from Corollary 1.10).

For $N = 9$, the interior angle is $\theta_9 = 140^\circ$. If $\phi < 70^\circ$ 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 $\phi = 70^\circ$ 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 $\Phi_9 = 70^\circ$ and this must be the sharp upper bound.

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

For $N = 7$, we have $\theta_7 = \frac{5}{7} \cdot 180^\circ \approx 128.5714$, and apply the argument to $\Phi_7 = \frac{5}{7} \cdot 90^\circ \approx 64.2857$. 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 $\phi \leq 72^\circ$ 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., $\Phi_4 \geq 72^\circ$. By our remark at the beginning of the proof, we must have equality. Exactly the same argument works for $N = 5$, but using $\kappa(L) \geq 6 - 5 \cdot (3 - 2) = 1 > 0$. \square

Proof of Corollary 1.9 (axis-parallel). By Corollary 1.4, $\phi(P) \leq 72^\circ$. 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 $\phi < 90^\circ$ 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 $\phi < 90^\circ$ 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 interior vertices of degree 5 (the rest are degree 6). \square

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 1.11.

Proof of Corollary 1.10 (large angles). If all angles of P are $\geq 144^\circ$, then for any $\phi < 72^\circ$, any ϕ -admissible labeling is at least 3 at every vertex, and hence has curvature ≥ 6 . By Corollary 1.4, 72° is therefore the sharp angle bound. If the minimum angle of P is $\geq 162^\circ$, 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 $[54^\circ, 72^\circ]$. If we also assume that all angles are $\leq 216^\circ$, then taking all labels equal to 3 is 72° -admissible, and has curvature 6, so we need only six applications of the 120° -trick. \square

12. COMPUTING $\Phi(P)$ IN LINEAR TIME

Before starting the proof that the optimal upper angle bound for an n -gon can be computed in time $O(n)$, I should mention that one of the referees of this paper objected to the length of proof given below, stating “In theory and practice, sorting takes linear time. . . . Only if you impose an absurd restriction like demanding an algorithm that fits into the decision tree model of computing do you need $O(n \log n)$ time to sort n integers. . . . Invoking the impractical median-of-medians algorithm (rather than quickselect³) adds another layer of absurdity, because now you’re restricting yourself to the deterministic decision tree model of computing. It’s like a theoretical circus trick where you prove you can accomplish the stunt without a net (or a real computer). A cute trick, but nobody should implement it.”

The author agrees no one should to implement it. Indeed, in Section 16, I report computing the distribution of angle bounds for a billion random choices of angles; this code used sorting, as it is easier to write and faster to run than the method in the proof. However, it seems worthwhile to the author to introduce a “cute trick” in order to remove a theoretical factor of $\log n$, much in the spirit of Chazelle’s very intricate proof in [20] that triangulating a polygon takes linear time, despite there being quite simple $O(n \log n)$ algorithms for the same task. In this paper, we do not address the question of finding efficient triangulations attaining the optimal bounds, so the most likely practical impact of our results is to provide a fast test that can be applied to existing triangulation methods to measure how close they are to optimal. Anyone employing such a test should attend to the referee’s remarks above.

Proof of Corollary 1.7. First check whether every angle of P is a multiple of 60° . If so, then $\Phi(P) = 60^\circ$. If not, find the smallest angle θ_{\min} of P . If $\theta_{\min} \leq 36^\circ$, then $\Phi(P) = 90 - \theta_{\min}/2$ by Corollary 1.4. If $\theta_{\min} \geq 144^\circ$ then $\Phi(P) = 72^\circ$ by Corollary 1.10. All this can be done in time $O(|V_P|)$.

Otherwise, we may assume $36^\circ < \theta_{\min} < 144^\circ$. We claim that computing $\Phi(P)$ reduces to finding (recall that $\kappa(\phi)$ was defined following Theorem 1.1):

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

³the quickselect algorithm chooses the k th smallest element of a list, reducing expected time from $O(n \log n)$ to $O(n)$, but is $O(n^2)$ in the (very rare) worst cases, see e.g., [1], [33]

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

Since $I(90^\circ) = [0^\circ, 90^\circ]$, 4 is a 90° -admissible label for any vertex of any polygon, so we always have $\kappa(90^\circ) < \infty$. 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. See Figure 2. If P has interior angle θ at v , We simply compute the smallest ϕ so that the point (ϕ, θ) is in one of the finitely many shaded triangles shown in that figure. This, in turn, consists of checking where the horizontal line at height θ hits the boundary of each triangle and taking the smallest such value.

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

We may therefore assume that $\kappa(90^\circ) = 0$. Thus ϕ_0 is well defined and $\phi_0 \leq 90^\circ$.

First, if $\kappa(\phi_\infty) = 0$, then $\phi_\infty = \phi_0$ and $\Phi(P) = \phi_\infty$. So assume $\kappa(\phi_\infty) \neq 0$; thus $\phi_\infty < \phi_0$.

Second, if $\kappa(\phi_\infty) > 0$, then $\kappa(\phi)$ is a decreasing, non-negative function of ϕ and

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

Finally, if neither of the previous cases hold, then $\kappa(\phi_\infty) < 0$, and this implies $\kappa(\phi)$ is non-positive and increasing, and we have:

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

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

We have now reduced finding $\Phi(P)$ to computing ϕ_0 , assuming that $\kappa(72^\circ) = 0$, and hence that $\phi_0 \leq 72^\circ$. 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 $\phi \in (60^\circ, 72^\circ]$ where θ_v lies on the boundary of one of the triangles in Figure 2. There are 10 possible triangles, so at most 10 possible ϕ 's for each θ_v . In fact, Figure 2 shows at most 5 triangles can be hit by a single horizontal line of height $\theta \in [0, 360^\circ]$ over the interval $[60^\circ, 72^\circ]$. This gives a set $X_1 \subset (60^\circ, 72^\circ]$ of size $\leq O(|V_P|)$ that contains all possible ϕ values where the integer valued step functions 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 we could find ϕ_0 by searching for the smallest ϕ value so that $\kappa(\phi) = 0$. This method takes time $O(|V_P| \log |V_P|)$ due to the sorting. (However, as noted before the proof, if you believe sorting takes linear time, you may stop reading the proof at this point.)

Recall that we are now assuming $\kappa(\phi_\infty) \neq 0$ and $\kappa(72^\circ) = 0$. 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 [18]. (See also [3] 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 $\phi_0 \in X_1 \cap [\phi_1, 72^\circ]$. 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

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

13. CONTINUITY OF Φ

Proof of Corollary 1.11. 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 , so that the vertices of P are sent to the vertices of P_n . This creates a triangulation of P_n whose angles are within $O(\delta)$ of the corresponding angles of \mathcal{T} . Thus $\limsup_n \Phi(P_n) \leq \Phi(P) + O(\delta)$ for any $\delta > 0$. Hence $\limsup_n \Phi(P_n) \leq \Phi(P)$.

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 .

Recall from the remark preceding Corollary 1.11 in Section 1 that $\mathfrak{P}_N \subset \mathbb{R}^{2N}$ denotes the set of simple N -gons. If $\phi > 60$, then $\mathcal{E}(N, \phi) = \{P \in \mathfrak{P}_N : \Phi(P) \leq \phi\}$ is the same as the set of N -gons that have ϕ -triangulations (for $\phi = 60^\circ$ the latter is a subset of the former), so it is a closed set. Similarly for $\mathcal{F}(N, \phi) = \{P \in \mathfrak{P}_N : \Phi(P) = \phi\}$. Corollary 1.10 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 1.9 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 27. Note P has eight 60° -vertices, one 300° vertex and a large number of $180^\circ + \epsilon$ vertices, where ϵ is 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. Thus in this range $\kappa(\phi) = 6 - 8(3 - 1) - 1(3 - 5) = 6 - 16 + 2 = -8 < 0$. Therefore, we are in

Case 2 of Theorem 1.1 and $\frac{5}{7} \cdot 90^\circ$ is the sharp bound. A small perturbation of P has the same labels, so P is in the interior of $\mathcal{F}(N, 72^\circ)$.

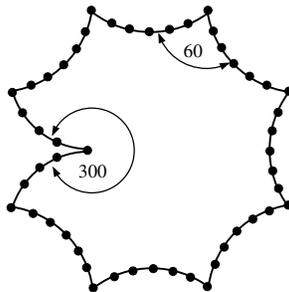


FIGURE 27. This polygon has optimal angle bound $\phi = \frac{5}{7} \cdot 90^\circ$, and so does every small perturbation, showing the set of such polygons has non-empty interior.

Finally, we want to check that these are the only angles that can give an open set. For any angle $\phi \in (60^\circ, 90^\circ]$ with $\phi \neq \frac{5}{7} \cdot 90^\circ, 72^\circ$, having $\Phi(P) = \phi$ means that one or more of the angles of P lies on the boundary of a shaded triangle in Figure 2 (otherwise we could decrease ϕ slightly and still meet the criteria of Theorem 1.1). Thus 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

14. NECESSITY IN THEOREM 1.1

For the sake of completeness, we give the details of the proof of necessity that was sketched just before Theorem 1.1. This direction is due to Gerber [37] and we will also restate our curvature conditions using his notation from [37].

Proof. Fix $60^\circ < \phi < 90^\circ$ 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 in \mathcal{T} containing v . Let F (for faces) be the number of triangles in \mathcal{T} , and let E be 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 $\phi \geq 72^\circ$, then (see Figure 2)

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

so having a ϕ -admissible labeling is equivalent to $\theta_{\min} \geq 180^\circ - 2\phi$. Thus Corollary 1.4 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).

This is the notation used by Gerver in [37]. 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 a labeling $L : V_P \rightarrow \mathbb{N}$ as above, we have $\sum_k k r_k = \sum_{v \in V_P} L(v)$ 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 these two inequalities are either both equalities or are both strict simultaneously.

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),
- $2E = 3F - \sum_k r_k - \sum_k s_k$ (triangle sides counted in two ways),

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

$$(14.1) \quad \sum_k (6 - k)q_k + \sum_k (3 - k)r_k + \sum_k (3 - k)s_k = 6,$$

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

$$(14.2) \quad \sum_{v \in \text{int}(\mathcal{T})} \kappa(v) + \sum_{v \in V_P} \kappa(v) + \sum_{v \in \partial\mathcal{T} \setminus V_P} \kappa(v) = 6.$$

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 (14.1) and (14.2) is non-positive, i.e., $\sum_{v \in \text{int}(\mathcal{T})} \kappa(v) + \sum_{v \in V_P} \kappa(v) \geq 6$.

In Case 2 of Theorem 1.1, $\phi < 72^\circ$, so every interior vertex has degree ≥ 6 . This implies $q_1 = \dots = q_5 = 0$ and thus the first term in (14.2) is also non-positive, i.e., $\kappa(L) = 6 - \sum_{v \in V_P} \kappa(v) \leq 0$, as desired.

In Case 3, $60^\circ < \phi < \frac{5}{7}90^\circ$, so every interior vertex of \mathcal{T} has degree six and every vertex in $\partial\mathcal{T} \setminus V_P$ has face-degree three. Thus $\sum_{v \in \text{int}(\mathcal{T})} \kappa(v) = \sum_{\partial\mathcal{T} \setminus V_P} \kappa(v) = 0$ and so (14.2) implies $\kappa(L) = 0$, as desired. \square

15. MAXIMIZING THE MINIMUM ANGLE

We recall some notation from Section 1. For $0 < \phi < 60^\circ$ we define $\tilde{I}(\phi) = [\phi, 180^\circ - 2\phi]$; any triangle having smallest angle ϕ has all its angles inside $\tilde{I}(\phi)$. Define a labeling L to be ϕ -lower-admissible if $\theta_v \in L(v) \cdot \tilde{I}(\phi)$ where θ_v is the angle of P at $v \in V_P$. See Figure 28. The curvature $\kappa(L)$ is the same as before, and $\tilde{\mathcal{K}}(\phi)$ is the set of curvatures of ϕ -lower-admissible labelings. We set $\tilde{\kappa}(\phi)$ 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 $\geq \phi$.

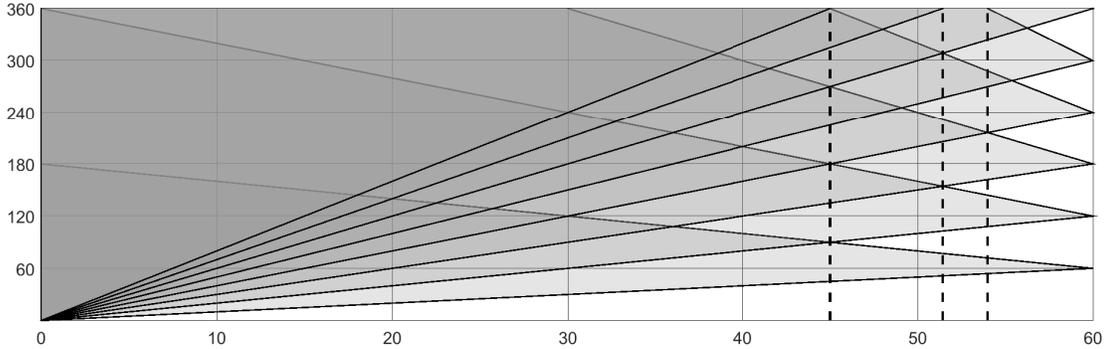


FIGURE 28. P has an ϕ -lower-admissible labeling iff all its angles lie in the intersection of the union of shaded triangles with the vertical line above ϕ . For $\phi \leq 45^\circ$, this only depends on the size of the smallest angle. The dashed lines indicate $\phi = 45^\circ$ and the two transition angles in Theorem 1.12.

Proof of Theorem 1.12. First consider necessity of the stated conditions. By definition, having a ϕ -lower-triangulation means that a lower-admissible labeling exists; 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 face-degree 3, so (1.1) becomes

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

If we also have $\phi > \frac{1}{7} \cdot 360^\circ$, 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 $\phi > 54^\circ$, 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 $\phi = \frac{1}{7} \cdot 360^\circ \approx 51.4286$, and the 120° -trick is eliminated at $\phi = 54^\circ$. \square

Proof of Corollary 1.13. The linear time calculation of $\tilde{\Phi}(P)$ is very similar to the calculation of $\Phi(P)$ described earlier, so we only note a few changes to the proof of Corollary 1.7. We define

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

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

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

- if $\tilde{\phi}_\infty \leq 54^\circ$, then $\tilde{\Phi}(P) = \tilde{\phi}_\infty$;
- if $\tilde{\phi}_\infty > 54^\circ$, then $\tilde{\Phi}(P) = 54^\circ$.

Otherwise, $\tilde{\kappa}(\tilde{\phi}_\infty) < 0$, and $\tilde{\kappa}(\phi)$ is decreasing and non-positive and

- if $\tilde{\phi}_\infty \leq \frac{4}{7} \cdot 90^\circ$, then $\tilde{\Phi}(P) = \tilde{\phi}_\infty$;
- if $\tilde{\phi}_\infty > \frac{4}{7} \cdot 90^\circ$, then $\tilde{\Phi}(P) = \tilde{\phi}_0$.

Thus we are reduced to calculating $\tilde{\phi}_0$ and $\tilde{\phi}_\infty$. Each of these can be done in linear time, just as in the proof Corollary 1.7 (the logic is the same, although the formulas are slightly different since $I(\phi)$ is different from $\tilde{I}(\phi)$). \square

Proof of Corollary 1.14. If $\phi > 45^\circ$, then every triangulation angle is $\leq 180^\circ - 2\phi < 90^\circ$, as desired. For $\phi \leq 45^\circ$, the triangles are always acute, except possibly near a vertex v of P . At each vertex $v \in P$, we choose the largest possible ϕ -lower-admissible label $L(v)$ for v . Then $L(v) \cdot \phi \leq \theta_v < (1 + L(v)) \cdot \phi$, so

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

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

Proof of Corollary 1.15. See Figure 28. Note that if $\phi \leq 45^\circ$, then $\bigcup_k k \cdot \tilde{I}(\phi) = [\phi, \infty)$, so P has a ϕ -lower-admissible labeling iff $\theta_{\min} \geq \phi$. \square

16. QUESTIONS AND REMARKS

We observed in the introduction that the number of elements of an angle-optimal triangulation of an N -gon will not satisfy a polynomial bound in N in general.

Question 16.1. *Can we find a ϕ -triangulation with the minimal number of elements (or even within a bounded multiplicative factor of minimal), in time comparable to the number of elements output?*

Question 16.2. *Can we estimate, to within a bounded factor, the smallest number of triangles needed for an angle optimal triangulation in terms of the geometry of P ?*

Question 16.3. *Is computing the minimal number of triangles required by an angle-optimal triangulation NP hard?*

Our proof of Theorem 1.1 often 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 $O(\inf_z \inf_e \omega(z, e)^{-2})$ 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 29. Here we have chosen a P' that mimics the overall shape of the rectangle, and we obtain $O(r)$ triangles, at the cost of introducing many more degree five interior 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 $\Phi(P) = 72^\circ$? What if we try to also optimize the number of singular points (interior vertices with degree $\neq 6$)?

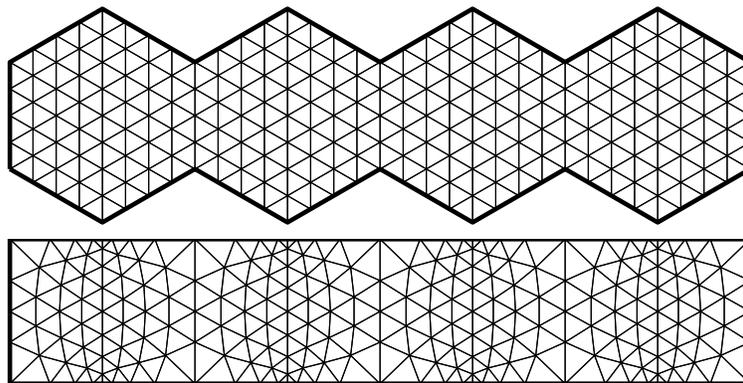


FIGURE 29. A different choice of P' leads to a number of triangles that is within a bounded factor of optimal. In this figure we have to make the grid finer and apply the 120° -trick to get rid of certain boundary vertices with angle 90° .

One of the referees of this paper noted that “The computer graphics and finite element mesh generation literature on frame fields contains some results on the necessity of ”singular” vertices with degrees other than the ideal (four for quadrilateralizations), and bounds on the number of them. They also discuss annihilating pairs of vertices, one with too-small degree with one with too-large degree. This may be relevant to mention.” Frame fields are smooth vector fields on the domain that give a local orientation for mesh elements (usually quadrilateral or hexahedral) to conform to, but singular vertices arise for topological reasons, just as discussed in Corollary 1.17. “Annihilating singularities” refers to replacing a portion of the mesh containing equal amounts of negative and positive curvature (in our notation) with a submesh with curvature zero (no singular vertices). Finding a precise connection between our bounds and results in the frame field literature seems like an extremely interesting problem, but one we leave for future investigations. The interested reader might start by looking at the papers [23], [24], [26] and their references. Alternate results concerning the number of necessary singular vertices for quadrilateral meshes are given in [13] and [34].

At several points in the proof we need to take the equilateral triangulation “fine enough”. This makes the triangulation elements small everywhere, not just at the location being discussed. Thus the construction is not local: requirements at any

point can effect the triangulation everywhere. Is there a more localized version of the construction? Can we build an optimal-angle triangulation that respects local feature size? Figure 30 shows a way of conformally manipulating a triangulation based on a square grid to change the size of the triangles, while still keeping some control of the angles. Can such an idea be used to create triangulations with both good angle bounds and triangle size adapted to local features?

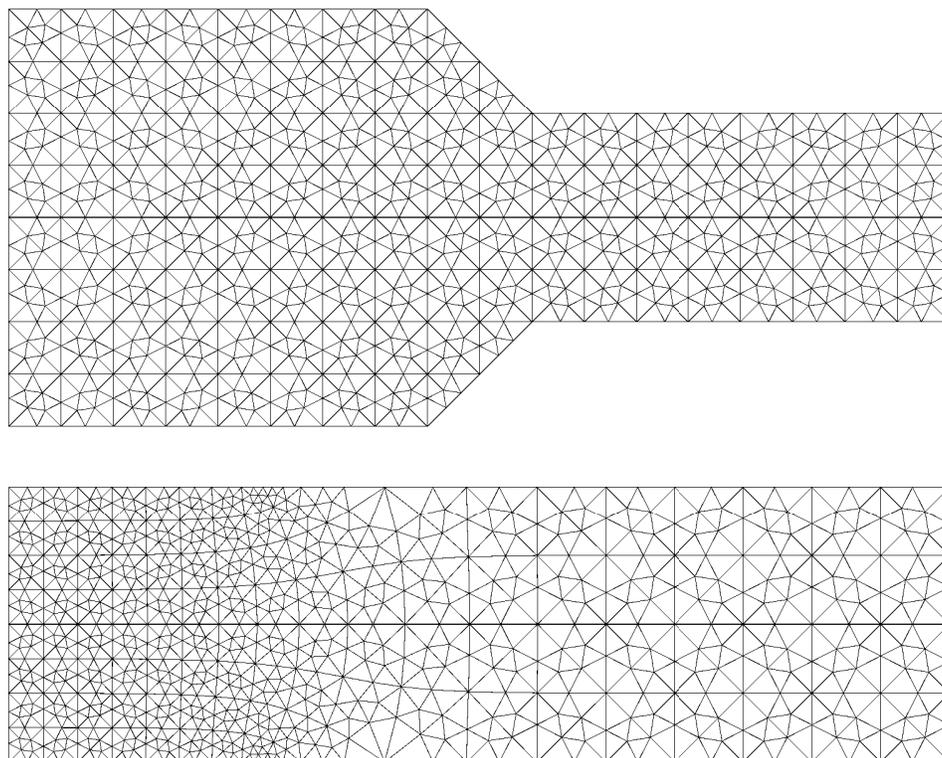


FIGURE 30. The mesh of the funnel region is transferred to the rectangle by the conformal map. This allows for a change of scale in the mesh elements. Can this idea be adapted to the construction in this paper to give more efficient triangulations where element sizes adapt smoothly to local features?

Can the angle-optimal triangulation results for polygons given in this paper be extended to planar straight line graphs (PSLGs)?

Problem 16.3. *Compute the optimal upper angle bound, $\Phi(\Gamma)$, for triangulating a given PSLG Γ .*

Conjecture 16.3. *If a PSLG Γ has minimal angle θ , then $\Phi(\Gamma) \leq 90^\circ - \min(36^\circ, \theta)$.*

In [15] I prove this, but with 36° replaced by some $\theta_0 > 0$ given by a compactness argument. Elementary examples show that 36° is best possible (e.g., any triangulation of a square has an angle $\geq 72^\circ$). Conjecturing $\theta_0 = 36^\circ$ for PSLGs may be too optimistic, but any explicit bound would be very significant.

Does Corollary 1.5 have analogs for other types of meshes? For example, does a quadrilateral dissection with certain angle bounds imply the existence of a quadrilateral mesh with the same bounds? Weaker but controlled bounds?

We have shown in this paper that triangulations minimizing the maximum angle exist (except in a few, well defined cases). However, it is unknown if a polygon with vertices in general position must have a minimal weight Steiner triangulation, i.e., a triangulation that minimizes the total edge length. Such a triangulation need not exist if we allow three co-linear vertices, [16]. An algorithm for approximating the infimum to within a constant factor is given in David Eppstein's paper [32]. Are there polygons where approximating the minimal weight triangulation requires the number of triangles to tend to ∞ ? If not, prove a bound on the number of triangles needed.

A referee of an earlier draft of this paper stated "The quasiconformal deformations used in this paper are also philosophically related to Thurston's complex hyperbolic moduli space of triangulations. I would have liked to see the connections between these two papers elucidated." William Thurston's paper [60] deals with characterizing combinatorial triangulations of the 2-sphere that have non-negative curvature (i.e., every vertex has degree at most six), by interpreting them as lattice points in a higher dimensional complex hyperbolic space defined using cone metrics on the sphere. The current paper constructs specific planar embeddings of triangulations whose curvature properties are forced by the desired geometry, e.g., Corollary 1.17 says precisely how many vertices of degree other than six are needed in an angle optimal triangulation. If we glue two copies of the polygon along the boundaries, we get a triangulation of the sphere with specified curvatures. When these are all positive, can the results in [60] help construct an optimal triangulation faster? Can they help find an optimal triangulation with fewest elements (or nearly fewest)? Do Thurston's ideas extend to cover the case when negative curvature vertices are needed?

Many problems similar to those studied in this paper remain wide open in higher dimensions.

Question 16.4. *Do polyhedra in \mathbb{R}^3 have non-obtuse tetrahedralizations of polynomial size (all dihedral angles $\leq 90^\circ$)?*

Question 16.5. *Given a polyhedra in \mathbb{R}^3 , compute the minimal possible maximal dihedral angle of a finite tetrahedralization. Similarly for maximizing the minimum dihedral angle.*

These are both probably very difficult. Even finding any acute tetrahedralization (all dihedral angles $< 90^\circ$) of a cube in \mathbb{R}^3 was open until about fifteen years ago (the smallest known example uses 1,370 pieces, [64]) and there is no acute decomposition for the cube in \mathbb{R}^4 , [44]. My own papers related to 2-dimensional meshing (e.g., [10], [11], [12]) use a thick/thin decomposition of a polygon that is analogous to the thick/thin decomposition of a hyperbolic manifold, in order to partition the polygon into two types of regions where Euclidean and hyperbolic geometry respectively are used to create the mesh. Can we use analogous ideas in \mathbb{R}^3 ? Can one create a 3-manifold out of a polyhedron, run a Ricci flow on it (as in Perelman’s proof of Thurston’s geometrization conjecture) to decompose it into pieces with geometric structure and then utilize the “natural” geometries on the different pieces to define meshes? An intermediate problem between dimensions 2 and 3 is to consider angle optimal refinements of triangulated surfaces.

Corollary 1.11 showed that the the set of N -gons with $\Phi(P) = 72^\circ$ contains an open set, if N is sufficiently large. Thus this set should have positive measure with respect to any measure on the space of N -gons that is mutually 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? The theory of Schramm-Loewner Evolutions (or SLE for brevity; see [42] and [55] among many other references) provides a satisfactory method of constructing random Jordan curves as solutions of a certain PDE using Brownian motion as input data, but I am not aware of how to efficiently compute uniform samples of random n -gons. Figure 31 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; perhaps $N = 10$ is too small for this open set to be non-empty. What is the smallest N for which it is non-empty?

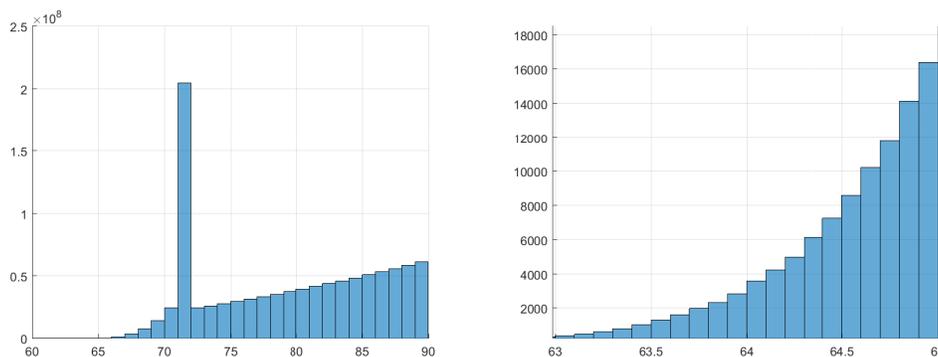


FIGURE 31. The distribution of optimal upper bounds over 10^9 random samples as described in the text. On the left is a histogram based on 1° bins. The spike at 72° is evident. On the right is an enlargement near 64° using $.1^\circ$ bins. No spike at $\frac{5}{7} \cdot 90^\circ \approx 64.26^\circ$ is visible.

Data Availability Statement: there is no data associated to this paper.

REFERENCES

- [1] Quickselect. Webpage <https://en.wikipedia.org/wiki/Quickselect>, Accessed: September 20, 2025.
- [2] 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.
- [3] 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.
- [4] B.S. Baker, E. Grosse, and C.S. Rafferty. Nonobtuse triangulation of polygons. *Discrete Comput. Geom.*, 3(2):147–168, 1988.
- [5] 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.
- [6] 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).

- [7] 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).
- [8] M. Bern and P. Plassmann. Mesh generation. In *Handbook of computational geometry*, pages 291–332. North-Holland, Amsterdam, 2000.
- [9] C.J. Bishop. Conformal welding and Koebe’s theorem. *Ann. of Math. (2)*, 166(3):613–656, 2007.
- [10] C.J. Bishop. Conformal mapping in linear time. *Discrete Comput. Geom.*, 44(2):330–428, 2010.
- [11] C.J. Bishop. Optimal angle bounds for quadrilateral meshes. *Discrete Comput. Geom.*, 44(2):308–329, 2010.
- [12] C.J. Bishop. Nonobtuse triangulations of PSLGs. *Discrete Comput. Geom.*, 56(1):43–92, 2016.
- [13] C.J. Bishop. Quadrilateral meshes for PSLGs. *Discrete Comput. Geom.*, 56(1):1–42, 2016.
- [14] C.J. Bishop. Uniformly acute triangulations of polygons. *Discrete Comput. Geom.*, 70(4):1571–1592, 2023.
- [15] C.J. Bishop. Uniformly acute triangulations of PSLGs. *Discrete Comput. Geom.*, 70(3):1090–1120, 2023.
- [16] C.J. Bishop. Minimal weight Steiner triangulation need not exist. 2024. preprint.
- [17] C.J Bishop and L. Rempe. Non-compact Riemann surfaces are equilaterally triangulable. to appear in *Inventiones Mat.*
- [18] 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.
- [19] Yu.D. Burago and V.A. Zalgaller. Polyhedral embedding of a net. *Vestnik Leningrad. Univ.*, 15(7):66–80, 1960.
- [20] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6(5):485–524, 1991.
- [21] B. Chazelle and D.P. Dobkin. Optimal convex decompositions. In *Computational geometry*, volume 2 of *Mach. Intelligence Pattern Recogn.*, pages 63–133. North-Holland, Amsterdam, 1985.
- [22] E.B. Christoffel. Sul problema della temperature stazonaire e la rappresetazione di una data superficie. *Ann. Mat. Pura Appl. Serie II*, pages 89–103, 1867.
- [23] G. Coiffier and E. Corman. The method of moving frames for surface global parametrization. *ACM Trans. Graph.*, 42(5), September 2023.
- [24] E. Corman and K. Crane. Symmetric moving frames. *ACM Trans. Graph.*, 38(4), 2019.
- [25] H.T. Croft, K.J. Falconer, and R.K. Guy. *Unsolved problems in geometry*. Problem Books in Mathematics. Springer-Verlag, New York, 1991.
- [26] D. Desobry, F. Protais, .N Ray, E. Corman, and D. Sokolov. Frame fields for cad models. In *Lecture Notes in Computer Science*, volume 13018, pages 421–434. 2022.
- [27] T.A Driscoll. Algorithm 843: Improvements to the Schwarz-Christoffel Toolbox for MATLAB. *ACM Transactions on Mathematical Software (TOMS)*, 31(2):239–251, 2005.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] D. Eppstein. Acute square triangulation. Webpage <https://www.ics.uci.edu/~eppstein/junkyard/acute-square/>, Accessed: January 2, 2021.

- [32] 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.
- [33] D. Eppstein. Blum-style analysis of quickselect. Accessed: September 20, 2025. <https://en.wikipedia.org/wiki/Quickselect>.
- [34] D. Eppstein, M.T. Goodrich, E. Kim, and R. Tamstorf. Motorcycle graphs: canonical quad mesh partitioning. In *Proceedings of the Symposium on Geometry Processing, SGP '08*, page 1477–1486, Goslar, DEU, 2008. Eurographics Association.
- [35] H. Erten and A. Üngör. Computing acute and non-obtuse triangulations. In *CCCG 2007, Ottawa, Canada*. 2007.
- [36] J.B. Garnett and D.E. Marshall. *Harmonic measure*, volume 2 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2005.
- [37] J.L. Gerver. The dissection of a polygon into nearly equilateral triangles. *Geom. Dedicata*, 16(1):93–106, 1984.
- [38] 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].
- [39] D.H. Hamilton. Conformal welding. In *Handbook of complex analysis: geometric function theory, Vol. 1*, pages 137–146. North-Holland, Amsterdam, 2002.
- [40] J. Itoh and T. Zamfirescu. Acute triangulations of the regular icosahedral surface. *Discrete Comput. Geom.*, 31(2):197–206, 2004.
- [41] 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.
- [42] Antti Kemppainen. *Schramm-Loewner evolution*, volume 24 of *SpringerBriefs in Mathematical Physics*. Springer, Cham, 2017.
- [43] R. Kenyon. Tilings of convex polygons. *Ann. Inst. Fourier (Grenoble)*, 47(3):929–944, 1997.
- [44] E. Kopczyński, I. Pak, and P. Przytycki. Acute triangulations of polyhedra and \mathbb{R}^N . *Combinatorica*, 32(1):85–110, 2012.
- [45] S. Korotov and J. Staňdo. Nonstandard nonobtuse refinements of planar triangulations. In *Conjugate gradient algorithms and finite element methods*, Sci. Comput., pages 149–160. Springer, Berlin, 2004.
- [46] 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.
- [47] D.T. Lee and A.K. Lin. Generalized Delaunay triangulation for planar graphs. *Discrete Comput. Geom.*, 1(3):201–217, 1986.
- [48] 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.
- [49] 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.
- [50] H. Maehara. Acute triangulations of polygons. *European J. Combin.*, 23(1):45–55, 2002.
- [51] 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.

- [52] S.A. Mitchell. Approximating the maxmin-angle covering triangulation. volume 7, pages 93–111. 1997. Fifth Canadian Conference on Computational Geometry (Waterloo, ON, 1993).
- [53] Z. Nehari. *Conformal mapping*. Dover Publications Inc., New York, 1975. Reprinting of the 1952 edition.
- [54] S. Rohde. On conformal welding and quasicircles. *Michigan Math. J.*, 38:111–116, 1991.
- [55] Steffen Rohde and Oded Schramm. Basic properties of SLE. *Ann. of Math. (2)*, 161(2):883–924, 2005.
- [56] 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.
- [57] S. Saraf. Acute and nonobtuse triangulations of polyhedral surfaces. *European J. Combin.*, 30(4):833–840, 2009.
- [58] 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, [59], pp. 84-101.
- [59] H.A. Schwarz. *Gesammelte Mathematische Abhandlungen*. Springer, Berlin, 1890.
- [60] William P. Thurston. Shapes of polyhedra and triangulations of the sphere. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 511–549. Geom. Topol. Publ., Coventry, 1998.
- [61] L.N. Trefethen. Numerical computation of the Schwarz-Christoffel transformation. *SIAM J. Sci. Statist. Comput.*, 1(1):82–102, 1980.
- [62] 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.
- [63] W.T. Tutte. The dissection of equilateral triangles into equilateral triangles. *Proc. Cambridge Philos. Soc.*, 44:463–482, 1948.
- [64] E. VanderZee, A.N. Hirani, V. Zharnitsky, and D. Guoy. A dihedral acute triangulation of the cube. *Comput. Geom.*, 43(5):445–452, 2010.
- [65] V. A. Voevodskii and G. B. Shabat. Equilateral triangulations of Riemann surfaces, and curves over algebraic number fields. *Dokl. Akad. Nauk SSSR*, 304(2):265–268, 1989.
- [66] L. Yuan. Acute triangulations of polygons. *Discrete Comput. Geom.*, 34(4):697–706, 2005.
- [67] 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