Discrete & Computational Geometry (2023) 70:1090–1120
https://doi.org/10.1007/s00454-023-00524-x

Uniformly Acute Triangulations of PSLGs

Christopher J. Bishop[1]

Received: 18 August 2021 / Revised: 16 June 2022 / Accepted: 2 July 2022 /
Published online: 8 July 2023

Abstract
We show that any PSLG Γ has an acute conforming triangulation $\mathcal{T}$ with an upper angle bound that is strictly less than 90° and that depends only on the minimal angle occurring in Γ. In fact, all angles are inside $[\theta_0, 90^\circ - \theta_0/2]$ for some fixed $\theta_0 > 0$ independent of Γ, except for triangles T containing a vertex v of Γ where Γ has an interior angle $\theta_v < \theta_0$; then T is an isosceles triangle with angles in the sharpest possible interval $[\theta_v, 90^\circ - \theta_v/2]$.

Keywords Acute triangles · Nonobtuse triangles · Gabriel points · Conforming triangulations · Delaunay triangulations

Mathematics Subject Classification 68U05 · 52B55 · 68Q25

1 Introduction

Every planar straight line graph (PSLG) has an acute conforming triangulation, e.g., [2, 6, 11, 15, 17]. However, no upper angle bound strictly less than 90° holds for all PSLGs. If a PSLG Γ has an interior angle θ at a vertex v, then any conforming triangulation $\mathcal{T}$ of Γ has a triangle T containing v with angle $\le \theta$. Since the angles of T sum to 180°, T also has an angle $\ge 90^\circ - \theta/2$. Therefore vertices of Γ with small angles force triangles in $\mathcal{T}$ with large angles. Can we obtain a uniform bound strictly less than 90° except for triangles containing such vertices? Is there an acute

Editor in Charge: János Pach

Data sharing not applicable to this article as no datasets were generated or analyzed during the current study. The author is partially supported by NSF Grant DMS 1906259.

Christopher J. Bishop
bishop@math.stonybrook.edu

[1] Mathematics Department, Stony Brook University, Stony Brook, NY 11794-3651, USA

angle bound for the whole triangulation that only depends on the minimal angle of Γ?

Theorem 1.1 *There is a $\theta_0 > 0$ so that any planar straight line graph Γ has a conforming triangulation with every angle in $[\theta_0, 90^\circ - \theta_0/2]$ except for triangles containing a vertex v of Γ where Γ has an angle $< \theta_0$. If the minimal interior angle of Γ at v is $\theta_v < \theta_0$, then every triangle containing v is isosceles with angles in $[\theta_v, 90^\circ - \theta_v/2]$. In particular, if Γ has minimal interior angle θ, then it has a triangulation with all angles $\leq 90^\circ - \min(\theta, \theta_0)/2$.*

The last claim quantifies the 1960 result of Burago and Zalgaller [11] that any polyhedral surface has an acute triangulation. The precise definition of a PSLG, of an interior angle, and of a conforming triangulation will be given in Sect. 2.

Theorem 1.1 was originally motivated by the special case of triangulations. Recall that a refinement of a triangulation $\mathcal{T}_1$ is another triangulation $\mathcal{T}_2$ of the same region so that each element of $\mathcal{T}_1$ is a union of elements of $\mathcal{T}_2$. A ϕ-triangulation is one in which every angle of every triangle is at most ϕ. Similarly for a ϕ-refinement.

Theorem 1.2 *There is a $\theta_0 > 0$ so that any planar triangulation $\mathcal{T}_1$ with lower angle bound $\theta > 0$ has a ϕ-refinement $\mathcal{T}_2$ with $\phi = 90^\circ - \min(\theta, \theta_0)/2$. We can choose $\mathcal{T}_2$ so that each triangle in $\mathcal{T}_1$ contains at most $O(1)$ triangles of $\mathcal{T}_2$, where the bound depends only on θ.*

The first statement is immediate from Theorem 1.1 and the second follows from its proof. This result answers a question of Florestan Brunck: in [9], he uses Lemma 13.1 of this paper (a consequence of Theorem 1.2), together with a result from [10], to transfer triangulations like the ones given in this paper from the Euclidean plane to a general Riemannian triangle complex.

The value θ_0 in Theorem 1.1 is given by a compactness argument and is not explicit, but in the special case of simple closed polygons, a more concrete construction using conformal mappings shows that Theorem 1.1 holds with $\theta_0 = 30^\circ$. See [8]. This implies that every polygon P with minimal angle θ has a ϕ-triangulation with $\phi = 90^\circ - \min(\theta, 30^\circ)/2$. A sharper version is given in [7]: P has a ϕ-triangulation with $\phi = 90^\circ - \min(\theta, 36^\circ)/2$. Moreover, [7] shows how to compute the optimal angle bound for any particular polygon in linear time. The analogous problem for PSLGs is open and probably much harder. Indeed, finding the sharp (or even an explicit) value for θ_0 in Theorem 1.2 for PSLGs already seems like a challenging problem, given the difficulties encountered in this paper merely to get below 90°.

Remark. Acute triangulations come with complexity bounds, but uniformly acute triangulations do not. For example, every n-gon P has an acute triangulation where the number of triangles is $O(n)$ [13, 16], and every PSLG has an acute conforming triangulation with $O(n^{5/2})$ elements [6]. It is known that $O(n^2)$ triangles are sometimes required for PSLGs, and this lower bound is conjectured to be sharp. To see that there are no such bounds for uniformly acute triangulations, consider a $1 \times r$ rectangle R where r is large. If R is triangulated with a bounded number M of triangles, one of them has diameter $\geq r/M$ but is contained in a strip of width one, so it has an angle bounded by $O(M/r)$. If M is bounded and r tends to infinity, this means this triangle

has an angle close to zero, and hence it has another angle $\geq 90^\circ - O(M/r)$. Thus uniformly acute triangulations cannot satisfy a complexity bound independent of the geometry.

A more detailed outline of the proof of Theorem 1.1 is given in Sect. 3, but the main idea is as follows. Every PSLG has an acute triangulation, and acute triangulations remain acute under small perturbations. Thus the space of PSLGs is covered by open neighborhoods, on each of which the angle bound required for triangulation is strictly less than 90° (depending on the neighborhood). If the space of PSLGs were compact, we would be done: a finite number of these neighborhoods would cover it, and we would simply take the worst of the finitely many bounds. However, PSLGs do not form a compact space, and we need to seek compactness elsewhere. Starting with a general PSLG, we first reduce to the case when each face is bounded by a simple polygon. We then take a set of disjoint disks centered at each vertex of the PSLG (called "protecting disks") and inscribe a polygon in their boundaries. The "unprotected" faces outside these polygons now have interior angles bounded uniformly below. By adding extra vertices to the edges of unprotected faces we can make these faces "thick" in the sense of [4]. Thick n-gons form a compact set; this is the polygonal version of Mumford compactness for Riemann surfaces [14]. By results in [3], the unprotected faces can be meshed by quadrilaterals with bounded eccentricity and with angles bounded strictly above 0° and below 180°, i.e., the quad-mesh elements are drawn from a compact family. However, the quad-meshes need not be consistent between different faces. By methods from [1, 6] the quadrilaterals are each acutely triangulated in a way that is consistent with the triangulations of adjacent quadrilaterals, and compactness of the family of quadrilaterals implies we get a uniform angle bound strictly less than 90°. The final step is to extend the triangulations inside the protecting disks, keeping a uniform angle bound except for triangles that touch the vertex v at the center of the disk. This extension is done by an explicit, but somewhat tedious, construction, and it completes the proof of Theorem 1.1.

As usual, if A and B are quantities that both depend on some parameter, we write $A = O(B)$ if there is a constant $M < \infty$ so that $A \leq MB$ for all parameter values. If $A \leq MB$ and $B \leq MA$ for all parameters, we say A and B are comparable with constant M. If this holds for some $M < \infty$ we write $A \simeq B$.

I thank Joe Mitchell and Florestan Brunck for several helpful conversations about the results in this paper. I also thank the two anonymous referees, whose numerous thoughtful comments greatly improved the exposition.

2 Planar Straight Line Graphs

In this section we recall some definitions and notation from [5]. A *planar straight line graph* Γ (or *PSLG* from now on) is a compact subset of the plane $\mathbb{R}^2$, together with a finite set $V = V(\Gamma) \subset \Gamma$ (called the vertices of Γ) such that $E = \Gamma \setminus V$ is a finite union of disjoint, bounded, open line segments (called the edges of Γ). The vertex set V includes both endpoints of every edge, and may include other points as well (i.e., isolated points of the PSLG). Although a PSLG really consists of the pair (Γ, V), we

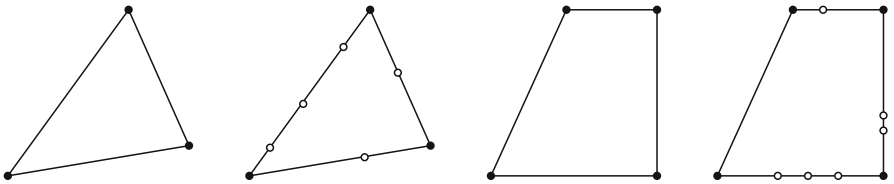

Fig. 1 A triangle, a triangular shaped octagon, a trapezoid and a trapezoid-shaped decagon. The black dots are the corners and the white dots are the interior edge vertices

will invariably refer to it as just Γ, with V implied by context. In several parts of the construction we will obtain a new PSLG by adding vertices to edges of a previous one, e.g., $\Gamma = [-1, 1]$ with one edge and $\{-1, 1\}$ as vertices, is made into a new PSLG Γ' with two edges and three vertices by adding $\{0\}$ to the vertex set.

A *polygon* or *polygonal curve* is a PSLG consisting of a sequence of vertices $z_1, \ldots, z_n$ and open edges $(z_1, z_2), \ldots, (z_n, z_1)$. A *polygonal path* or *arc* is a similar list of vertices, but with edges $(z_1, z_2), \ldots, (z_{n-1}, z_n)$; the last is not connected back to the first. A polygon is *simple* if the vertices are all distinct and the edges are pairwise disjoint.

A *triangle* is a simple polygon with three vertices (hence three edges). We say a simple polygon P has a *triangular shape* if there is a triangle T so that P is obtained by adding vertices to the edges of T. See Fig. 1. Similarly, a *quadrilateral* Q is a simple polygon with four vertices. We say a simple polygon P has a *quadrilateral shape* (or is *quad-shaped*), if P is obtained by extra adding vertices to the edges of a quadrilateral Q. The four vertices of Q will be called the *corners* of P. These are the only vertices of P where the interior angle is not $180°$, unless Q happens to be triangular shaped. The other vertices of P will be called *interior edge vertices*.

A quadrilateral Q is called θ-*nice* if the interior angles are all in $[90° - \theta, 90° + \theta]$, e.g., a $0°$-nice quadrilateral is a rectangle. When $\theta = 30°$ we simply call Q "nice". Given $0 \le \theta < 90°$, a θ-*trapezoid* is a θ-nice quadrilateral Q with two parallel sides (called the base sides). A θ-*isosceles trapezoid* Q is a θ-trapezoid where the two base sides have a common perpendicular bisecting line, called the axis of the trapezoid, and Q is symmetric with respect to the axis. An (L, H, θ)-*isosceles trapezoid* has shorter base length L, height H (the distance between base sides) and angles exactly $90° \pm \theta$. Note that in this case, the two non-base sides have length $H \sec \theta$.

A *face* of Γ is any of the bounded, open connected components of $\mathbb{R}^2 \setminus \Gamma$. Every PSLG has a unique unbounded complementary component that we sometimes call the *unbounded face*, but a PSLG may or may not have faces. We say that a bounded, connected open set Ω is a *polygonal domain* if it is the face of some PSLG (informally, $\partial\Omega$ is a finite union of points and line segments). It is a *simple polygonal domain* if its boundary is a simple polygon. For a simple polygon, it is clear what adjacent edges means. For the boundary of a general polygonal domain Ω, we say that two edges are *adjacent* if they share an endpoint and if for any $\epsilon > 0$ the interiors of the edges can be joined by a crosscut of Ω of length $< \epsilon$. A *crosscut* of a domain Ω is a Jordan arc whose interior is contained in Ω, and whose endpoints are on the boundary of Ω.

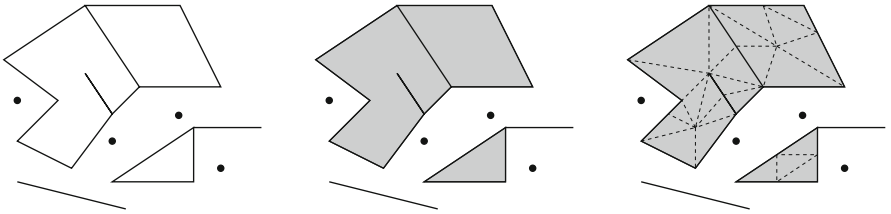

Fig. 2 A PSLG, its polynomial hull, and a conforming triangulation

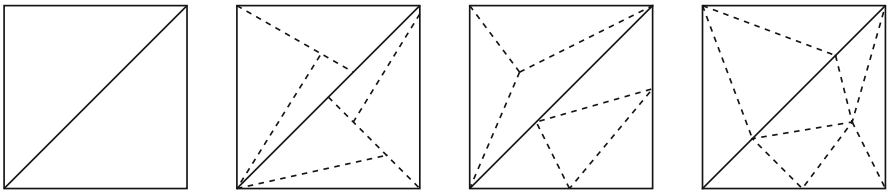

Fig. 3 A PSLG Γ with two faces, a triangular dissection of Γ, a triangulation of each face, and a conforming triangulation of Γ. As we move left to right, the type of subdivision becomes more restrictive, and so the optimal upper angle bound potentially becomes larger. Does it actually increase?

The union of faces is called the *interior* of the PSLG. The *polynomial hull* of a PSLG Γ is the union of Γ and its interior, and is denoted $\mathrm{PH}(\Gamma)$. The name comes from complex analysis, where the polynomial hull of a compact set K is defined as

$$\mathrm{PH}(K) = \left\{ z \in \mathbb{C} : |p(z)| \leq \sup_{w \in K} |p(w)| \text{ for all polynomials } p \right\}.$$

This agrees with our definition in the case K is a PSLG. See Fig. 2.

A *refinement* (also called a *sub-division*) of a PSLG Γ is a PSLG Γ' such that $V(\Gamma) \subset V(\Gamma')$ and $\Gamma \subset \Gamma'$. Informally, Γ' is obtained from Γ by adding new vertices and edges and by subdividing existing edges. A *mesh* of Γ is a refinement Γ' of Γ such that $\Gamma' \subset \mathrm{PH}(\Gamma)$ and every face of Γ' is a simple polygonal domain. Note that we allow the addition of new vertices (called *Steiner points*) when we mesh a PSLG. A mesh Γ' is called a *triangulation* if every face of Γ' is a triangle and is called a *quadrilateral mesh* or *quad-mesh* if every face is a quadrilateral.

A mesh in which each face is triangular shaped (but not necessarily a triangle) is a *triangular dissection*. Thus a triangulation of each face of a PSLG automatically gives a triangular dissection of the PSLG. See Fig. 3.

A triangulation of faces is a triangulation of the PSLG if it is consistent between faces, i.e., any two intersecting triangles, even from different faces, either intersect at a single point that is a vertex of both triangles or intersect in a segment that is a full edge of both triangles. When this holds for triangles from a sub-collection $\mathcal{F}$ of faces, we say the dissection is $\mathcal{F}$-*consistent*. We will consider several different types of decompositions of a PSLG into triangles:

- triangular dissections (no consistency required),
- triangulations of faces (consistency within faces, but not across edges of the PSLG),

- $\mathcal{F}$-consistent triangulations (consistency within all faces and between members of $\mathcal{F}$),
- triangulations (consistency within and between all faces).

To emphasize the difference between a dissection of a PSLG, a triangulation of its faces, and a triangulation of the PSLG, we will sometimes call the latter a "conforming triangulation" or "consistent between all faces". Similar definitions and terminology holds for quadrilateral dissections, quad-meshes of faces, and quad-meshes of the whole PSLG. Every triangular or quadrilateral dissection of a PSLG can be refined to a conforming triangulation or quad-mesh respectively, with polynomial growth in complexity. See [5] and [6].

Theorem 1.1 of [3] says that a simple polygon with all interior angles $\geq 60°$ has a mesh by nice quadrilaterals. If the polygon has n sides, then this mesh can be taken with $O(n)$ elements, but there is no uniform bound on their *eccentricity*, i.e., the ratio of an element's longest edge length to its shortest edge length. We say a quad-mesh is a (θ, E)-*quad-mesh* if every element is a θ-nice quadrilateral, with eccentricity $\leq E$. We say that Γ' is a (θ, E, δ)-*quad-dissection* if each quad-shaped face is θ-nice, has eccentricity $\leq E$ and each edge e of a quad-shaped face Q has length $\ell(e) \geq \delta \cdot \operatorname{diam}(Q)$. Similarly, a (θ, E, δ)-*marked quadrilateral* is a θ-nice quadrilateral with eccentricity E, together with a finite collection of points on its edges so that the resulting segments all have length $\geq \delta \cdot \operatorname{diam}(E)$. Note this is a special case of a quad-shaped polygon, defined earlier. Later we will use the fact that for fixed values of (θ, E, δ), such marked quadrilaterals form a compact set (at least when normalized to have diameter 1 and contain the origin).

If v is a vertex of Γ, then an angle of Γ at v is an angle between two edges of Γ that have a common endpoint at v. The angle is an *interior angle* if the two edges bound the same face of Γ and are adjacent with respect to that face. Note that a PSLG may have more than one interior angle at v, but a simple polygon has exactly one at each vertex. In the special case that the vertex v is the endpoint of a single edge of Γ, and v is in the interior of $\operatorname{PH}(\Gamma)$, we set the interior angle at v to be $360°$.

Since it is enough to prove Theorem 1.1 for each connected component of $\operatorname{PH}(\Gamma)$, we may assume from here on that $\operatorname{PH}(\Gamma)$ is connected (just take the sub-PSLG consisting of all edges and vertices in this component). By the following result we may also reduce to the case that each face is bounded by a simple polygon.

Lemma 2.1 [5, Lem. 3.1] *If Γ is a PSLG with n vertices such that $\operatorname{PH}(\Gamma)$ is connected, then by adding at most $O(n)$ new edges and vertices, we can find a connected refinement Γ_1 of Γ so that every face of Γ_1 is a simple polygon and any interior angle of a face that is less than $60°$ was already an interior angle of a face of Γ (i.e., such angles are not subdivided).*

3 Outline of the Proof

The proof of Theorem 1.1 will be given in several steps, each described in its own section. Here we summarize the steps to give an overview of the proof.

In what follows, "disk" will always mean a circular disk $D(x, r) = \{y : |y-x| < r\}$. The boundary of such a disk is the circle $\partial D = \{y : |y - x| = r\}$. Given a disk $D = D(x, r)$ and $\lambda > 0$ we let $\lambda D = D(x, \lambda r)$ and $\lambda \partial D = \partial(\lambda D)$. For example, $2D$ is the disk concentric with D and of twice the radius. We call this the double of D.

Assume we start with a PSLG Γ_1 with n vertices, and that each face of Γ_1 is bounded by a simple closed polygon. The nth step of the construction takes a PSLG Γ_n and produces a new PSLG Γ_{n+1} (usually a refinement of Γ_n and always a refinement of Γ_1 and Γ_2).

- **Step 1:** Suppose $\theta_0 > 0$. A value of θ_0 will be fixed at the end of Sect. 12. For each vertex v of Γ_1, let θ_v be the minimum of θ_0 and the minimum interior angle of Γ_1 at v. For each vertex v of Γ_1 we take a small enough disk D_v centered at v so that the doubles of all the disks are pairwise disjoint. We inscribe polygons in ∂D_v and $(1 - s_v)\,\partial D_v$ (for some $0 < s_v < 1/2$ depending on θ_v), and connect v to the vertices on ∂D_v by radial segments. This gives a new PSLG that we call Γ_2. The faces of Γ_2 inside D_v are either isosceles triangles with vertex angles at v between θ_v and $2\theta_v$, or they are θ-isosceles trapezoids with the same bounds on θ. These are called the *protected faces* of Γ_2. The remaining faces are all bounded simple polygons with all interior angles $> 90°$ and are called the *unprotected faces*. These collections of faces will be denoted $\mathcal{P}$ and $\mathcal{U}$ respectively.
- **Step 2:** The unprotected faces of Γ_2 consist of a collection of simple polygons for which all the interior angles are $\geq 90°$. We add extra vertices to the edges of Γ_2 to make a PSLG Γ_3 whose faces are all simple polygons and $(1/2)$-thick. The latter condition means that if e, f are two edges on the face of the PSLG then either e and f share a vertex or $\operatorname{dist}(e, f) \geq \min(\operatorname{diam}(e), \operatorname{diam}(f))/2$, where $\operatorname{dist}(e, f) = \inf\{|x - y| : x \in e,\ y \in f\}$. In other words, the distance between non-adjacent edges is at least half the length of the shorter edge.
- **Step 3:** We show that each thick polygonal face of Γ_3 can be meshed using nice quadrilaterals that have uniformly bounded eccentricity. This gives Γ_4.
- **Step 4:** By sliding vertices of Γ_4 that lie on (interiors of) edges of Γ_3, we obtain a PSLG Γ_5 that is a (θ, E, δ)-quadrilateral mesh of each of the unprotected faces of Γ_3. Here we can take any $\theta > 30°$, say $\theta = 40°$; $E < \infty$ and $\delta > 0$ will be fixed independent of the PSLG.
- **Step 5:** We will define a Gabriel cover of Γ_5 by disks and add the centers and tangency points of these disks to Γ_5 to get a PSLG Γ_5'. Using a method of Bern et al. [1], we show that each quadrilateral face Q of Γ_5 can be non-obtusely triangulated and the boundary vertices of the triangulation are exactly the vertices of Γ_5' that lie on ∂Q. This gives a PSLG Γ_6 that is an nonobtuse triangular dissection of Γ_2 and is $\mathcal{U}$-consistent (consistent between the unprotected faces).
- **Step 6:** We show that the nonobtuse triangulation of each quadrilateral Q in Step 5 can be refined to an acute triangulation by twice bisecting each boundary segment of the nonobtuse triangulation. This gives a PSLG Γ_7 that is an acute triangular dissection of Γ_2 and is also $\mathcal{U}$-consistent.
- **Step 7:** Since the set of quadrilaterals and boundary vertices is compact, we will show the acute triangulations in Step 6 can be chosen with uniform upper angle bound $90° - \eta$, if $\eta > 0$ is small enough. This may alter some of the triangulations in

Step 6, so we obtain a PSLG Γ_8 that is a triangular dissection of Γ_2, is $\mathcal{U}$-consistent, and has an upper angle bound strictly less than 90° inside the unprotected faces.

- **Step 8:** We extend the mesh into the protecting disks in two steps. First we extend the Gabriel cover in Step 5 to cover the boundaries of the protected isosceles trapezoid faces of Γ_2. This also extends the acute triangulations into these trapezoids. The shorter base side of each trapezoid is divided into four segments with length ratios bounded by 2. This gives a new triangulation Γ_9 that is now consistent for all faces except for the protected triangles touching vertices of Γ_1.
- **Step 9:** The final step is to explicitly construct acute triangulations of the isosceles triangles touching vertices of Γ_1. This will reduce to finding acute triangulations of isosceles trapezoids with angles $0 \leq \theta < \theta_1$, for some positive θ_1. These triangulations must have exactly one boundary edge on the shorter base side e of the trapezoid, exactly four specified edges on the other base, and the vertices on the other two sides (radial with respect to v) must have distances from e that do not depend on θ. The last condition implies the triangulations for neighboring trapezoids agree along common edges, and so we get a triangulation of all of $\mathrm{PH}(\Gamma)$, finishing the proof. We will show that this can be done if $\theta_1 > 0$ is small enough, by giving an explicit construction for rectangles and then using the fact that acute triangulations are stable under small perturbations. This gives a PSLG, Γ_{10}, that satisfies Theorem 1.1 with $\theta_0 = \min(\eta, \theta_1)$.

The remaining sections add details to the discussion above, and verify the claims we have made.

4 Step 1: Protecting the Vertices

Let V be the set of vertices of the PSLG Γ_1. Without loss of generality we may assume the polynomial hull is connected, every face of Γ_1 is bounded by a simple polygon, and each vertex of Γ_1 is in the polynomial hull. Let $\{D_v\} = \{D(v, r_v)\}$ be a collection of disjoint disks centered at the points $v \in V$ so the doubles are pairwise disjoint (as before, "double" means a concentric disk with twice the radius). Then for each such disk $\mathrm{PH}(\Gamma_1) \cap D_v$ is a finite union of closed sectors of the form $\{z : \psi_1 \leq \arg(z - v) \leq \psi_2,\ |z - v| \leq r_v\}$. Consider one such sector.

This sector is divided into finitely many sub-sectors by the edges of Γ_1. Recall that $\theta_0 > 0$ has been fixed. Eventually θ_0 will be chosen quite small, but for now we just assume that $\theta_0 \leq 5°$. Let ψ_v denote the minimal angle at v formed by these sub-sectors. Let $\theta_v = \min(\theta_0, \psi_v)$. Leave each sub-sector with angle $\theta < 2\theta_v$ alone, but subdivide larger angles into further sub-sectors with angles satisfying $\theta_v \leq \theta < 2\theta_v \leq 10°$ by adding edges from v to ∂D_v. This can be done either by recursively bisecting each sector until the angles satisfy the desired estimate, or by dividing a sector of angle ψ into $\lfloor \psi/\theta_v \rfloor$ equal sub-sectors. Finally, we inscribe one set of polygons where these radial segments hit $\{|z - v| = r_v\} = \partial D_v$, and another where it hits $\{|z - v| = (1 - s_v) r_v\}$, where $s_v = (5/4) \sin \theta_v$. Add the radial segments and add the inscribed polygonal arcs to Γ_1 to get a new PSLG Γ_2. See Fig. 4.

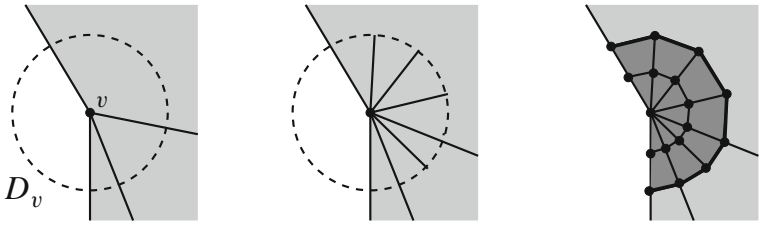

Fig. 4 The solid lines in the left picture are a part of Γ_1 near v and the shaded region is part of $\mathrm{PH}(\Gamma_1)$. At each vertex v of Γ_1 we add edges from v to ∂D_v that cut the sectors of $\mathrm{PH}(\Gamma_1)$ into sectors with angles between θ_v and $2\theta_v$. The rightmost picture shows Γ_2. The thick polygon arc in the right picture is the inscribed arc between the protected faces (dark gray) and the unprotected faces (lighter gray)

The definition of s_v ensures that the trapezoids formed between the two polygonal arcs have uniformly bounded eccentricity. To see this, rescale so $r_v = 1$; then the radial sides of the trapezoid each have length $s_v = (5/4)\sin\theta_v \leq (5/4)\sin 5^\circ < 1/5$. The distance d between the radial sides of such a trapezoid is equal to the length of the shorter base side of the trapezoid. Using the fact that $\sin(x/2) > (\sin x)/2$ for $0 \leq x \leq 90^\circ$ we can estimate:

$$d = 2(1 - s_v)\sin\frac{\theta}{2} \geq 2(1 - s_v)\sin\frac{\theta_v}{2} > \left(1 - \frac{1}{5}\right)\sin\theta_v \geq \frac{16}{25}s_v.$$

The distance between the two base sides of the trapezoid is at least

$$s_v\cos\frac{\theta}{2} \geq s_v\cos\theta_v \geq s_v\cos 5^\circ > \frac{4}{5}s_v \geq \frac{d}{2}.$$

Thus the distance between either pair of opposite sides of the trapezoids we create is always at least half the length of the shorter side of the pair (this means that the trapezoids are $(1/2)$-thick, a concept to be defined in the next section).

The faces of Γ_2 are either isosceles triangles with vertex angles between θ_v and $2\theta_v$, isosceles trapezoids adjacent to these triangles, or are bounded by simple polygons with all interior angles larger than 90° (removing the triangular faces containing the vertices of Γ_1 and the adjacent trapezoids makes all the angles in the remaining faces $> 90^\circ$). We call the isosceles triangles containing v and the adjacent trapezoids the protected faces of Γ_2 and the other faces the unprotected faces, and denote these two collections by $\mathcal{P}$ and $\mathcal{U}$ respectively. The two types of faces are separated by polygonal arcs (or possibly a closed polygonal curve for vertices v in the interior of $\mathrm{PH}(\Gamma_1)$) inscribed on ∂D_v. We will call these the protecting arcs of Γ_2. Our construction will produce a series of refinements $\{\Gamma_n\}$ of Γ_2. Faces of Γ_n that are subsets of unprotected faces of Γ_2 will also be called unprotected. The unprotected faces will all be triangulated with a uniform angle bound strictly less than 90°. Only the protected faces will eventually contain triangles touching the vertices of Γ and with angles that depend on the corresponding θ_v.

5 Step 2: Making the Faces Thick

The paper [4] introduced the idea of decomposing the interior of a polygonal region into thick and thin parts and showed this could be accomplished in linear time (linear in the number of vertices of the polygon). The thin pieces come in two types: parabolic and hyperbolic (the names come from a precise analogy with the types of thin parts of a Riemann surface or a hyperbolic manifold). For a simple polygon P, the parabolic thin parts correspond to certain neighborhoods of each vertex, so a simple n-gon has exactly n parabolic thin parts. Each hyperbolic thin part corresponds to a pair of non-adjacent edges $e,\ f$ of P that are very close to each other compared to their diameters. The polygon P is called "thick" if no hyperbolic thin parts occur.

The precise definition of a hyperbolic thin part in terms of extremal length is given in [4], but for our purposes here, it suffices to consider a simpler sufficient condition. A simple polygon P will be ϵ-thick, if given any two non-adjacent edges $e,\ f$ of P, we have

$$\operatorname{dist}(e,\ f) \geq \epsilon \cdot \min(\operatorname{diam}(e), \operatorname{diam}(f)). \tag{5.1}$$

As before, $\operatorname{dist}(e,\ f) = \inf\{|x - y| : x \in e,\ y \in f\}$. An edge e is called ϵ-thick if

$$\operatorname{dist}(e,\ f) \geq \epsilon \cdot \operatorname{diam}(e) \tag{5.2}$$

for every edge f not adjacent to e. Clearly if each edge of a PSLG is ϵ-thick, then every boundary edge of every face of the PSLG is ϵ-thick.

Lemma 5.1 *We can add vertices to the edges of* Γ_2 *to create a PSLG* Γ_3 *so that every face is a* $(1/2)$*-thick polygon.*

Proof All the vertices of Γ_2 are either vertices of Γ_1 or vertices of the protecting arcs. By construction, all the edges of the protecting faces are $(1/2)$-thick: triangles are trivially $(1/2)$-thick (since there are no non-adjacent edges), and we proved the trapezoids are $(1/2)$-thick in the previous section.

Each remaining edge e of Γ_2 connects vertices of inscribed arcs around different vertices of Γ_1. Cut this edge into segments with length between L and $L/2$, where L is the minimal length of a protecting edge adjacent to e. See Fig. 5. Then each of the new edges is $(1/2)$-thick, and (5.1) is satisfied by all pairs of edges. □

6 Step 3: Nice Quad-Meshes of Thick Polygons

As noted earlier, [5, Thm. 1.1] says that any polygon with all interior angles $\geq 60°$ has a nice quadrilateral mesh (i.e., all angles in $[60°, 120°]$). That result gives a linear bound on the number of mesh elements, but this bound forces some quadrilaterals to have large eccentricity. However, all the large eccentricity elements occur inside the hyperbolic thin parts of P. If we assume there are no such thin parts, there will be no high eccentricity quadrilaterals. We make this precise as follows. (Recall that

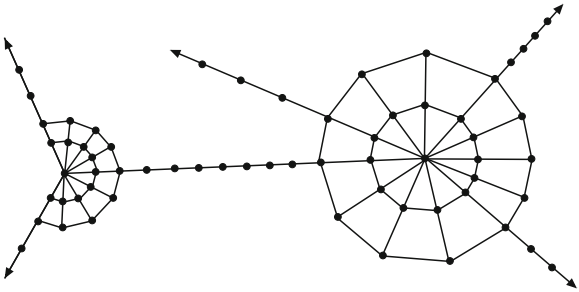

Fig. 5 We add vertices to the edges outside the protected faces to make all the faces (1/2)-thick

parameterized quantities A, B are comparable with constant M if $1/M \leq B/A \leq M$ for all parameters values.)

Lemma 6.1 *Suppose $\epsilon > 0$. Every ϵ-thick polygon with all angles larger than 60° has a nice quad-mesh whose elements have uniformly bounded eccentricity. Moreover, if a quadrilateral Q in the mesh has a side s lying on an edge e of P, then s and e have comparable lengths (with a constant that may depend on ϵ, but not on s, e, Q, or P).*

Proof The existence of a quad-mesh with the given angle bounds is the statement of [5, Thm. 14.1]. The uniform bound on eccentricity follows from the proof of that theorem. The proof divides the interior of P into three types of regions: thick parts, parabolic thin parts and hyperbolic thin parts. By assumption, no hyperbolic thin parts occur in our case. All the quadrilaterals used in the thick parts have uniformly bounded eccentricity. The only quadrilaterals used to mesh the parabolic thin parts that do not have uniformly bounded eccentricity are the kite shaped elements that meet a vertex v of the polygon. If the interior angle of P is small, this kite has large eccentricity, but we have assumed that all interior angles of P are $\geq 60^\circ$, so this does not occur. Thus for thick polygons with angles $\geq 60^\circ$, all elements of the corresponding nice quad-mesh have bounded eccentricity.

The proof of [5, Thm. 14.1] also shows that each boundary segment of P is covered by at most $N = O(1)$ mesh elements (N may depend on ϵ). Since each element has uniformly bounded eccentricity $\leq M$, side lengths of adjacent quadrilaterals differ by at most M^2 and, by induction, two quadrilaterals that meet the same edge of P have diameters differing by at most M^N. Since each edge e of P is covered by a bounded number of quadrilateral sides that all have comparable length, each of these sides must have length comparable to the length of e. This proves the lemma. □

The proof in [5] shows that the number of quadrilaterals used in Lemma 6.1 is $O(n)$, where n is the number of vertices of P and the multiplicative constant depends only on ϵ. This will be used in the proof of Theorem 1.2.

Let Γ_4 be the refinement of Γ_3 obtained by nicely quad-meshing each of the unprotected faces of Γ_3. Since Γ_3 and Γ_2 have the same faces, this also gives a nice quad-mesh of the unprotected faces of Γ_2.

7 Step 4: Perturbing the Quadrilaterals

We now have a PSLG Γ_4 that gives a 30°-nice, bounded eccentricity quad-mesh of each face of Γ_3, but the mesh need not be consistent between different faces. Refining a quad-mesh of faces to be consistent between faces is a difficult problem; this is dealt with in [5]. However, here we only need a much weaker form of consistency: we want any edge e of our PSLG to have length that is comparable to the diameter of any quad-shaped face Q that contains e in its boundary. In other words, we want to perturb Γ_4 (if necessary) to be a (θ, E, δ)-quad-dissection of the faces of Γ_2.

Lemma 7.1 *For each $\theta > 30°$ there is an $\delta > 0$ so that the following holds. We can slide the vertices of Γ_4 that lie in the interior of edges of Γ_3 along those edges to get a new PSLG Γ_5 that is combinatorially the same as Γ_4 (i.e., the underlying abstract graphs are the same) but is a $(\theta, 2E, \delta)$-quad-dissection.*

Proof We know that Γ_4 is a refinement of Γ_3, that every edge e of Γ_4 lies on some edge f of Γ_3, and that e and f have comparable lengths, say within a factor $M < \infty$. Let $\ell(e)$ denote the Euclidean length of an edge e. Choose $\delta \ll 1/M$ and divide f into segments of length between $\delta\ell(f)$ and $2\delta\ell(f)$. Call the resulting set of endpoints V'. Take each boundary vertex of Γ_4 and move it to the closest element of V'. Note that vertices of Γ_3 are not moved since these points are already in V', and that any other vertex of a mesh element Q is moved by less than $O(\delta \cdot \text{diam}(Q))$. Hence the interior angles of Q are changed by less than θ if δ is small enough, depending only on θ. Similarly, the eccentricity of any quadrilateral changes by a small amount if δ is small enough. □

Fix some value of $\theta > 30°$, say $\theta = 10°$, apply the previous lemma to our PSLG Γ_4 and let Γ_5 be the new PSLG. Note that Γ_5 is a refinement of Γ_3 (but not necessarily of Γ_4; we have may have replaced edges of Γ_4 that are incident on interior points of edges of Γ_3).

8 Step 5: Gabriel Coverings and Nonobtuse Triangulation

Although our quad-mesh Γ_5 of the faces of Γ_2 is not consistent between different faces, we will show that we can triangulate each quadrilateral so that the triangulations are consistent within and between all the unprotected faces of Γ_2, i.e., we get a $\mathcal{U}$-consistent triangulation of Γ_2 (recall that $\mathcal{U}$ was our collection of unprotected faces, and $\mathcal{U}$-consistent means that intersecting triangles in or between such faces meet in a vertex or full edge of each). In this section, we use an idea of Bern, Mitchell, and Ruppert to create consistent nonobtuse triangulations (all angles $\leq 90°$), and in the next section we use an idea of Yuan to make the triangulation acute.

Given a point set V and two points $v, w \in V$, the segment vw is called a Gabriel edge if it is the diameter of an open disk containing no points of V (see [12]). This is special case of a Delaunay edge, that needs only be the chord of a disk missing V. A Gabriel cover of a PSLG Γ is a finite collection $\mathcal{D}$ of disks centered at points of Γ so that

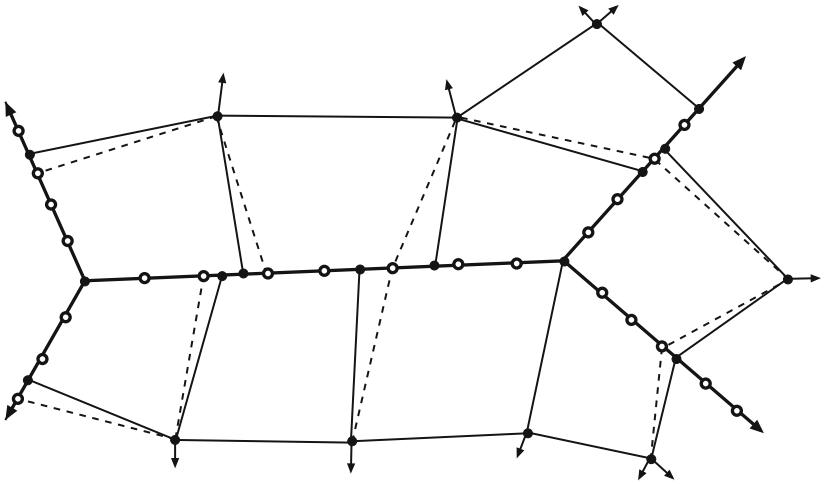

Fig. 6 Here Γ_3 is shown in thick solid lines, Γ_4 with the thinner solid lines, and Γ_5 using the dashed lines. By sliding vertices of Γ_4 along edges of Γ_3 we get the new PSLG Γ_5 whose vertices on Γ_3 are well separated, i.e., two distinct vertices on the same edge f of Γ_3 are at least $\delta\ell(f)$ apart. Γ_5 is combinatorially the same as Γ_4, and corresponding angles change by only a small amount. The white dots represent the set V' in the proof of Lemma 7.1

- the interiors of the disks are disjoint,
- the closures cover Γ,
- each vertex v of Γ is the center of some disk,
- if $D \in \mathcal{D}$ is not centered at a vertex v of Γ then $D \cap \Gamma$ is a single segment.

In particular, $D \cap \Gamma$ is always either a union of segments, all of the same length, meeting a vertex of Γ, or it is a diameter of D contained in a single edge of Γ. See Fig. 7. We say a Gabriel cover is M-Gabriel if for any disk D in the cover that hits an edge e of Γ, we have $\operatorname{diam}(D) \leq \operatorname{diam}(e) \leq M \cdot \operatorname{diam}(D)$. In other words, the cover divides edges into subsegments of comparable lengths. Because faces of Γ_5 are uniformly θ-nice, have bounded eccentricity, and adjacent edges have comparable diameters and edge lengths, it is easy to check that Γ_5 has an M-Gabriel cover for some $M < \infty$, and that only a uniformly bounded number of points need to be added to each edge.

Given a Gabriel cover of Γ_5, add the centers of all disks and the tangent points between all pairs of adjacent disks to Γ_5 to obtain a refinement Γ_5'. In [1] Bern et al. proved that there is a triangulation of each quadrilateral face Q of Γ_5 whose boundary vertices are exactly the vertices of Γ_5' that lie on ∂Q. Taking these triangulations of each face gives a PSLG Γ_6 that is a $\mathcal{U}$-consistent non-obtuse triangulation of the faces of Γ_2, i.e., it is consistent between all the unprotected faces. We need a slight strengthening of their result, so we will describe their construction in detail, first recalling a few definitions from [1].

Suppose we start with a Gabriel cover of a quadrilateral, e.g., the light gray disks in Fig. 8. For each pair for adjacent disks in the cover, we add a small disk tangent to both of them; see the white gray disks in Fig. 8. These disks are taken to be disjoint. We

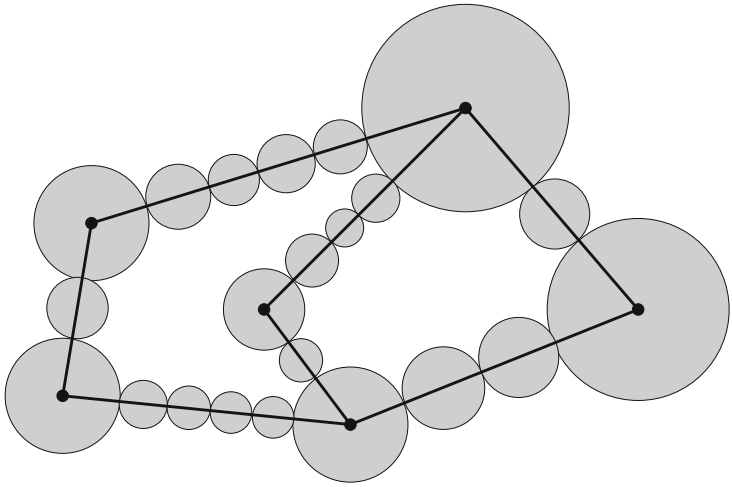

Fig. 7 A Gabriel cover of a PSLG

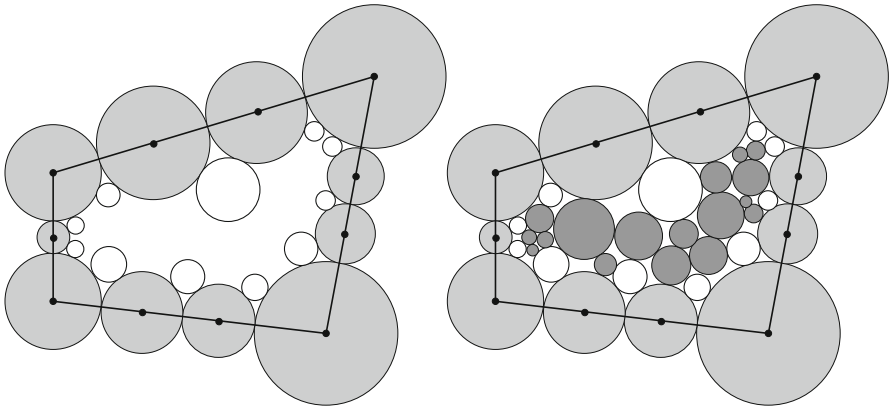

Fig. 8 Given a Gabriel cover of a quadrilateral (light gray), we add disjoint disks that are tangent to adjacent pairs of the cover (white disks). These guarantee that the boundary of the quadrilateral is covered by the augmented regions of 3-gaps. This will later ensure that the boundary of the quadrilateral is covered by hypotenuses of right triangles in our triangulation. We then add more tangent disks (dark gray) as in [1] until only 3-gaps and 4-gaps remain

then add more disks (dark gray in Fig. 8) until the remaining regions are all bounded by either three or four circular arcs. These are called 3-gaps and 4-gaps respectively. That this can be done is proven in [1]: given a region bounded by $k > 4$ arcs, they show we can remove a disk so the remaining regions are all bounded by at most $k - 1$ arcs. They prove [1, Lem. 1] that the total number of disks used is linear in the number of Gabriel disks covering the boundary of the quadrilateral. This observation will be used in the proof of Theorem 1.2.

Given a 3-gap or 4-gap region R we define the associated augmented region R^+ by joining the centers of the tangent disks defining the gap. See Fig. 9. These form

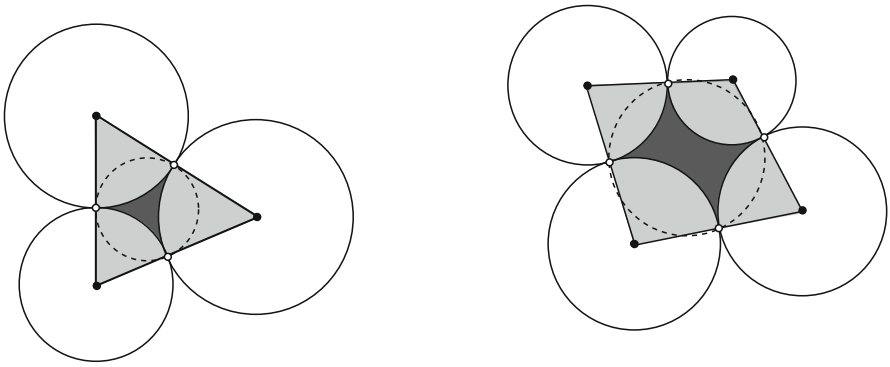

Fig. 9 Augmented regions for a 3-gap and 4-gap are shown in gray (the gap region is dark gray). A lemma in [1] says the four tangent points of a 4-gap all lie on a circle (dashed in the picture)

triangles and quadrilaterals that mesh the original quadrilateral. Using a case-by-case analysis Bern, Mitchell, and Ruppert show that each augmented region can be nonobtusely triangulated so that the only boundary vertices are at corners of the augmented region or at the tangent points between disks. Thus the triangulation of different augmented regions are consistent and give a non-obtuse triangulation of the original quadrilateral Q, and the boundary vertices of this triangulation on ∂Q are exactly the set $V \cap \partial Q$. Hence the triangulations of adjacent quadrilaterals are consistent and we get a nonobtuse triangulation of the unprotected faces of Γ_2.

In order to convert this nonobtuse triangulation into an acute triangulation, we are going to need a little more information about the triangulations of the augmented regions than is given in [1]. The original version of their result, and the modifications we need are as follows.

Theorem 8.1 (Bern–Mitchell–Ruppert [1]) *Suppose R is a* 3*-gap or* 4*-gap and let R^+ be the corresponding augmented region. Then R^+ can be triangulated by at most* 28 *right triangles so that no new vertices are added to the boundary of R^+. Alternatively, we can use nonobtuse triangles that satisfy the following conditions:*

(i) *Any two right triangles that share an edge, must share an edge of the same type* (*either leg-leg or hypotenuse-hypotenuse*). *When two triangles share a leg, the same vertex of that segment is the* 90° *angle of both triangles.*
(ii) *Any right triangle that has a side on ∂R^+ has only one side on ∂R^+ and that side is its hypotenuse.*
(iii) *The boundary of an augmented* 3*-gap is covered by hypotenuses of right triangles in the triangulation.*

Proof The first part is proven in [1]. The alternative conclusions (i)–(iii) will be needed in Sect. 9 when we want to refine our nonobtuse triangulation to an acute triangulation, and they require only a few minor changes (we will add a few extra right triangles and convert a few right triangles to acute triangles). We sketch these alterations below.

The first difference is in triangulating 3-gaps. In [1], the center of the inscribed circle of R^+ is connected to the centers of the circles and to the points of tangencies

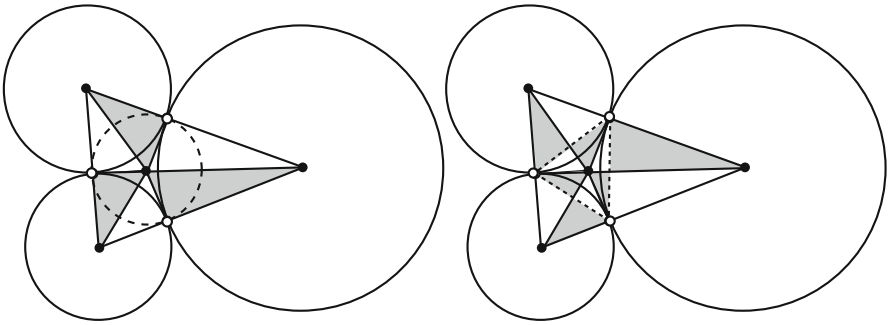

Fig. 10 The triangulation of 3-gaps in [1] uses six right triangles, but has legs on ∂R^+. This is shown on the left. We modify this by adding the dashed triangle (as shown on the right) to form twelve right triangles with only hypotenuses on ∂R^+

between the circles. This gives six right triangles whose legs lie on the boundary of R^+. However, (ii) and (iii) do not hold. To fix this, we add the three chords connecting the points of tangency, and get twelve right triangles so that only hypotenuses lie on ∂R^+. Now conditions (i)–(iii) of the lemma hold. See Fig. 10.

Before triangulating the 4-gaps, recall that the four vertices all lie on a single circle, denoted C_*. This is [1, Lem. 3]. See Fig. 9. That lemma also states that the angle measure of the four boundary arcs sums to 2π (the angles correspond to the interior angles of the augmented quadrilateral). Thus at most one of the arcs can be reflex, i.e., have angle measure $> \pi$.

In [1] the 4-gaps are split into three cases. The first case occurs when:

- R is centered, that is, the center of the circle C_* passing through the four points of tangency is inside the convex hull of these points, and
- none of the arcs in ∂R is reflex.

If both these conditions hold, then the triangulation given in [1] has the desired properties: R^+ is divided into kites by connecting the center of C_* to the four tangent points. A kite is a strictly convex quadrilateral that is symmetric with respect to at least one of its diagonals. Any kite is divided into four right triangles by adding its diagonals and all four hypotenuses are on the boundary of the kite. See Fig. 11.

The second case occurs when:

- R is not centered, that is, the center of the circle C_* passing through the four points of tangency is outside the convex hull of these points, and
- none of the arcs in ∂R is reflex.

The authors of [1] show that a fifth disk can be added, tangent to two opposite circles, creating two new centered 4-gaps (one possibly self-intersecting) and such that the union W of the two augmented regions can be written as a union of seven kites, and each is triangulated by its diagonals. This causes the boundary of W to contain only hypotenuses and for all adjacencies to be of matching type. Thus no changes to the argument in [1] are needed in this case. See Fig. 12.

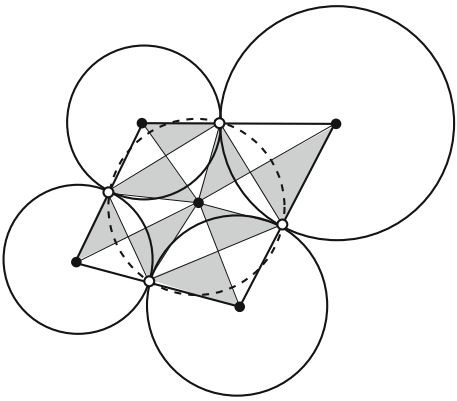

Fig. 11 If R is centered and every arc is non-reflex, then R^+ is a union of sixteen right triangles with the desired properties

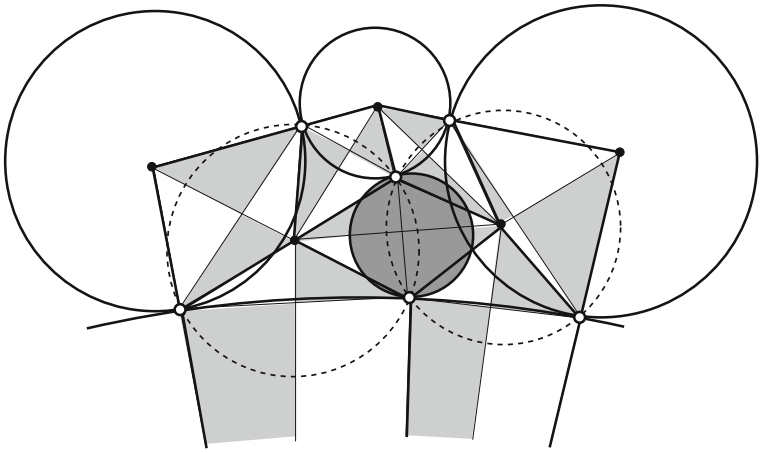

Fig. 12 A non-centered 4-gap is split into two centered 4-gaps by the shaded disk. It is proven in [1] that such a disk exists and we use the same triangulation as given there

The third case is when

- one of the four arcs of the 4-gap is reflex.

Suppose C_2 is the circle with the reflex arc and it is opposite C_4. The authors of [1] insert a new disk C_5 centered on the segment connecting the centers of C_2 and C_4 and tangent to both. This creates two new 4-gaps, neither with a reflex arc, since they both contain an arc of measure π. However, the new disk may intersect one of the other two. See Fig. 13. Note that C_5 is smaller than either C_1 or C_3. The fact that the new disk is on the segment connecting the centers of C_2 and C_4 means that if it intersects, say, C_3, then the common chord of C_5 and C_3 separates the centers of C_5 and C_3. The proof of this is left to the reader. Similarly if it intersects C_1. This chord-separated property implies the associated augmented region can be triangulated by sixteen right triangles as in Fig. 14.

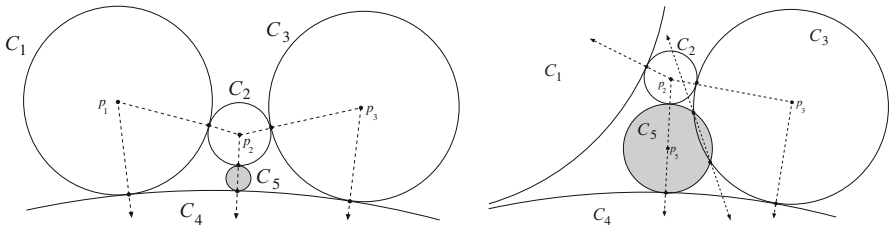

Fig. 13 Case 3. If R reflex, then add an extra disk (shaded) centered on the segment connecting the non-reflex circle to its opposite. If this disk intersects one of the two other disks bounding the 4-gap (as shown on the right), the line containing the common chord must separate the centers of the overlapping circles

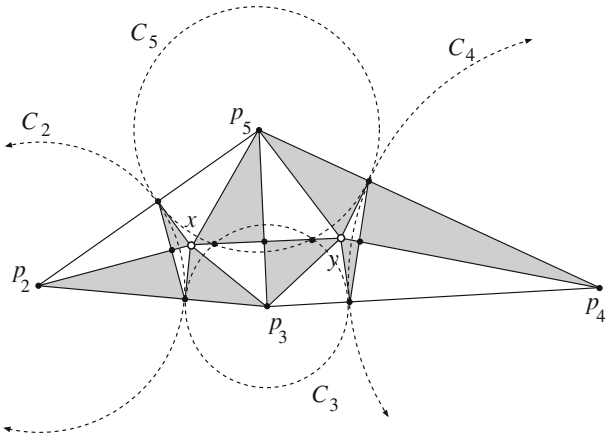

Fig. 14 Case 3 continued. A non-reflex but overlapping 4-gap. This triangulation with 16 right triangles is given in [1]. However, four triangles have legs on the boundary of the augmented region. We fix this by moving some interior vertices, as in Fig. 15

Some of the right triangles in Fig. 14 have legs on ∂R^+. We will fix this by making these triangles acute. Suppose the circles are C_2, C_3, C_4, C_5 with C_3, C_5 intersecting, and let p_2, p_3, p_4, p_5 denote their centers. The tangent line between C_2 and C_3, the tangent line between C_2 and C_5 and the line through the two intersection points of C_3 and C_5 all intersect in one point x, one of the white dots of degree six in Fig. 14. Let y be the corresponding point for circles C_3, C_4, C_5.

Both x, y lie on the same side as p_3 of the line through p_2, p_4. Slide x by a small amount $\epsilon > 0$ on the ray from p_2 through x away from p_2. This gives a new point x'. See Fig. 15. Similarly, slide y along the ray from c_4 through y, keeping the line through x', y' parallel to the line through x, y. For small $\epsilon > 0$ this makes the four right triangles with legs on ∂R^+ acute and leaves all the remaining triangles right. Thus (i) and (ii) of the lemma are now satisfied. □

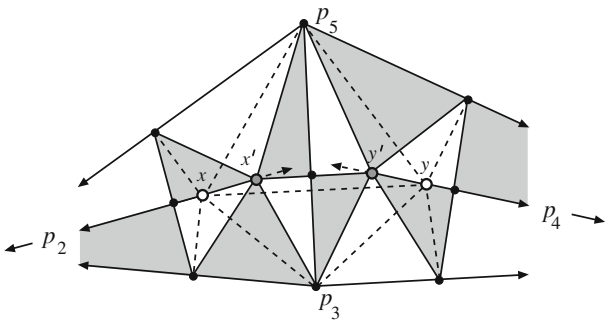

Fig. 15 Case 3 completed. This shows an enlargement and slight alteration of Fig. 14. The dashed lines indicate the previous triangulation and the solid lines (and shaded triangles) show the new one. We move the white points x, y to the gray points x', y': x' is on the line through p_2 and x, y' is on the line through p_4 and y, and the line $x'y'$ is parallel to xy. Any small enough motion makes four of the right triangles acute and the remaining twelve right triangles satisfy conditions (i) and (ii) of Theorem 8.1

9 Step 6: From Nonobtuse to Acute Triangulation

The argument so far produces a nonobtuse triangulation of each quadrilateral. We next convert it to an acute triangulation using an idea of Yuan from [16]. The quarter points of a segment refer to the midpoint of the segment and the midpoints of each resulting half of the segment. If T is a right triangle, the two vertices with angles $< 90°$ will be called the acute vertices of T.

Theorem 9.1 *Let T be a right triangle. Form a 12-gon P by adding the quarter points of each edge. This polygon has a triangulation by 24 triangles using exactly these boundary vertices. There are 22 acute and two right triangles. The two right triangles T_1, T_2 contain the two acute angles of T and their hypotenuses lie on the hypotenuse of T. For any small enough $\epsilon > 0$, the entire triangulation can be made acute by sliding the vertices of T_1, T_2 that lie on the legs of T by ϵ away from the acute vertices of T. In particular, the boundary vertices of this acute triangulation lying on the hypotenuse of T are exactly the quarter points of the hypotenuse.*

Proof The proof is basically a series of pictures; see Fig. 16. Divide T into a rectangle and two right triangles by connecting the midpoint of the hypotenuse to the midpoints of the legs. Then repeat this in the two triangles. Acutely triangulate the large rectangle as shown. Then move the marked interior vertices as shown in Fig. 16. This makes all the triangles acute except for the two containing the acute angles of T. These are T_1, T_2 and they can be made acute by sliding the two vertices as described in the theorem (see Fig. 16). □

In the final step of the proof we could also slide the two vertices of T_1, T_2 that lie on the hypotenuse of T by ϵ towards the acute vertices of T. However, to maintain consistency of our mesh between quadrilaterals, we want to leave vertices on the hypotenuse of T fixed. This is why we have stated Yuan's theorem as we have.

Corollary 9.2 *Suppose we have a nonobtuse triangulation of a quadrilateral Q such that*

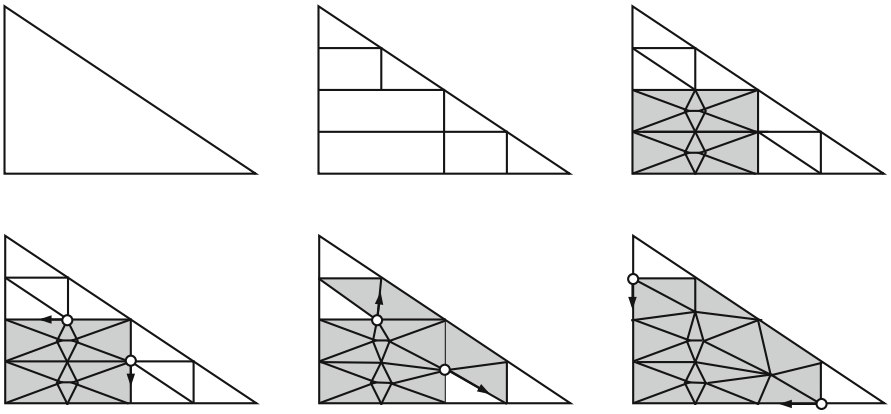

Fig. 16 The proof of Yuan's theorem. The shaded triangles in each picture are acute. Moving the white dots in the indicated directions makes some of the incident right triangles acute; a small enough move keeps all acute triangles acute. We only move two points on the boundary (last step). Both are on legs of the original right triangle T and both move towards the right angle of T

- *adjacent right triangles only share edges of the same type (leg or hypotenuse),*
- *when adjacent right triangles share a leg, the same endpoint is the right corner of each triangle,*
- ∂Q *is covered by hypotenuses of right triangles of the triangulation.*

Then the triangulation has an acute refinement with 24 *times as many triangles. The segments of the original triangulation on* ∂Q *are twice bisected by the segments of the new triangulation.*

Proof First subdivide every triangle by connecting the midpoints of each side. Sub-triangles of right triangles are right and sub-triangles of acute triangles are acute. Then repeat the midpoint subdivision on each resulting triangle. This gives a nonobtuse refinement of the original nonobtuse triangulation. To make it an acute triangulation we "fix" all the right triangles by sliding vertices as described in Theorem 9.1. If the edge where we slide is shared with an acute triangle, choose ϵ so small that it remains acute. If the edge is shared with another right triangle, the shared edge must be of the same form (leg or hypotenuse), so sliding in one direction fixes both triangles at once. No points on ∂Q are moved since only hypotenuses occur on ∂Q. □

A similar statement can be made for general simple polygons, not just quadrilaterals. Combining this with Lemma 8.1 gives the following.

Lemma 9.3 *Suppose* Q *is a quadrilateral with a Gabriel cover of its boundary. Then there is an acute triangulation of* Q *whose boundary segments are obtained by twice bisecting the boundary segments of the nonobtuse triangulation given in Lemma 8.1.*

We have now proven that there is an acute refinement Γ_7 of Γ_6 that is a triangular dissection of Γ_2 and is consistent for all the unprotected faces of Γ_2.

10 Step 7: Compactness Gives Uniformly Acute Triangulations

By the construction so far, each unprotected face of Γ_5 is quadrilateral shaped and is acutely triangulated to give Γ_7, and the triangulations are consistent between different faces. In this section, we want to observe that the angles used in these triangulations can be taken to be bounded strictly below 90°, independent of Γ. This is essentially because the quadrilateral faces of Γ_5 and the triangulation vertices on boundaries of these quadrilaterals come from a compact family of possibilities. We explain this more carefully below.

First recall that each quadrilateral shaped face of Γ_5 is θ-nice for some fixed θ, say $\theta = 40°$ (i.e., the angles at the corners are bounded between $50° = 90° - \theta$ and $130° = 90° + \theta$). Each such quadrilateral also has bounded eccentricity, say less than $E < \infty$. For the rest of this section, we assume all the quadrilaterals we refer to have these angle and eccentricity bounds. If we also normalize such quadrilaterals to have diameter 1 and to contain the origin, then clearly they form compact family, i.e., the limit of any sequence of such normalized quadrilaterals is another such quadrilateral. To define the topology on quadrilaterals more precisely, we can think of a quadrilateral as defined by the location of its four corners, and hence as a point of $\mathbb{R}^8$; thus the set of normalized quadrilaterals described is a closed, bounded set in $\mathbb{R}^8$.

The boundaries of the quadrilateral shaped faces of Γ_5 may contain other vertices of Γ_5 beside their corners, but these vertices are δ-separated, i.e., there is a uniform $\delta > 0$ so that any two vertices of the boundary of a face Q are at least $\delta \cdot \operatorname{diam}(Q)$ apart. Thus each face forms a quadrilateral with a bounded number of δ-separated marked points on its boundary. Since the number of marked points on each quadrilateral is uniformly bounded by $N = \lfloor 4 \cdot \operatorname{diam}(Q)/\delta \rfloor$, we can think of the collection of marked quadrilaterals with $4 \le m \le N$ vertices as a subset of $\mathbb{R}^{2m}$. Again, it is easy to see that, normalized as above, these marked quadrilaterals form a compact family with respect to the topology on $\mathbb{R}^{2m}$, since a limit of δ-separated points will also be δ-separated. Since there are only a finite number of possible values of m, the collections of all marked quadrilaterals we are considering is also compact.

The next step of the construction added more points to the edges of Γ_5 to obtain Γ_5', and we triangulated the quad-shaped faces using these boundary points to get Γ_6. These new points were the centers and tangencies of a M-Gabriel cover by disks. Because M was bounded in terms of θ, E, and δ, the new points were again well separated, possibly with δ replaced by smaller constant δ'. These points formed the boundary vertices of a non-obtuse triangulation of each face. The final step was to divide each of the edges of Γ_6 into four equal subsegments, and we proved that each quadrilateral face could be acutely triangulated using exactly these vertices on its boundary. The equal subdivision means the quadrilateral faces are still marked with $(\delta'/4)$-separated points. Just as above, these marked quadrilaterals, normalized to have diameter 1 and contain the origin, form a compact family.

Given an acute triangulation of a marked quadrilateral, the same abstract graph still gives an acute triangulation if we move the triangulation boundary vertices by a small enough distance. Thus if a particular marked quadrilateral Q can be triangulated with maximum angle $90° - \eta$, there is a neighborhood of Q in the collection of marked quadrilaterals that can be triangulated with all angles $< 90° - \eta/2$. Since

we have shown that every marked quadrilateral can be acutely triangulated, such open neighborhoods cover the whole set of normalized marked quadrilaterals. Since the space of marked triangles is compact, it is covered by a finite number of such neighborhoods. Taking the maximum of a finite set of numbers each strictly less than 90° gives a number strictly less than 90°, so we have proved the following.

Lemma 10.1 *Suppose* $0 \le \theta < 90°$, $E < \infty$, *and* $\delta > 0$ *are fixed. Suppose* Q *is a* θ*-nice quadrilateral and with eccentricity* $\le E$ *and the edges of* Q *are divided into edges that each have length at least* $\delta \cdot \operatorname{diam}(Q)$. *Also assume that these points are the quarter points of a sub-division of* ∂Q *into segments corresponding to a Gabriel cover of* ∂Q. *Then* Q *has an acute triangulation with exactly the marked points as the boundary vertices of the triangulation, and so that every angle is* $\le 90° - \eta$ *where* $\eta > 0$ *depends only on* θ, E, *and* δ.

Let Γ_8 be the resulting PSLG (the compactness argument above might replace the triangulation of some faces of Γ_7 by perturbed versions of approximating faces). This completes the proof that the unprotected faces of Γ_2 can all be triangulated with a maximum angle that is strictly below 90° and is independent of the PSLG. In the final two steps, we triangulate the protected faces, where we know that some angles arbitrarily close to 90° may be needed (depending on the angles of the original PSLG at the protected vertices).

11 Step 8: Meshing the Protected Trapezoids

We wish to extend the triangulation of the unprotected faces of Γ_1 into the protected faces. Recall from Step 1 that the region between ∂D_v and $(1 - s_v)\partial D_v$ is meshed by isosceles trapezoids Q that each have angle θ_Q satisfying $\theta_v \le \theta_Q \le 2\theta_v$. The shorter edge of such a trapezoid Q has length $h_Q = 2(1 - s_v)r_v \sin(\theta_Q/2)$. We place a disk of radius $h_v = (2/3)(1-s_v)r_v \sin(\theta_v/2)$ centered at each endpoint of the shorter base edge of Q. These disks are dark gray in Fig. 17. Since the radius depends on v but not on Q, choices for adjacent trapezoids agree. Around the midpoint of the shorter base of Q we put a disk of radius $(h_Q - 2h_v)/2$, that covers the remaining segment. These disks are light gray in Fig. 17.

Note that since $\theta_v \le \theta_Q \le 2\theta_v$ and using the fact that $\sin(x/2) \ge (\sin x)/2$ for small enough x (recall $\sin x$ is concave down on $[0, 90°]$), we have

$$h_v = \frac{2}{3} s_v r_v \sin\frac{\theta_v}{2} \le \frac{2}{3} s_v r_v \sin\frac{\theta_Q}{2} \le \frac{h_Q}{3}$$

and

$$h_v = \frac{2}{3} s_v r_v \sin\frac{\theta_v}{2} \ge \frac{2}{3} s_v r_v \sin\frac{\theta_Q}{4} \ge \frac{2}{6} s_v r_v \sin\frac{\theta_Q}{2} = \frac{h_Q}{6}.$$

We then add disjoint disks to the non-base sides of each Q to form an M-separated Gabriel cover of the whole quadrilateral mesh; $M < \infty$ can be taken independent of Q if Q satisfies the conditions above. These are also light gray in Fig. 17. Since Q

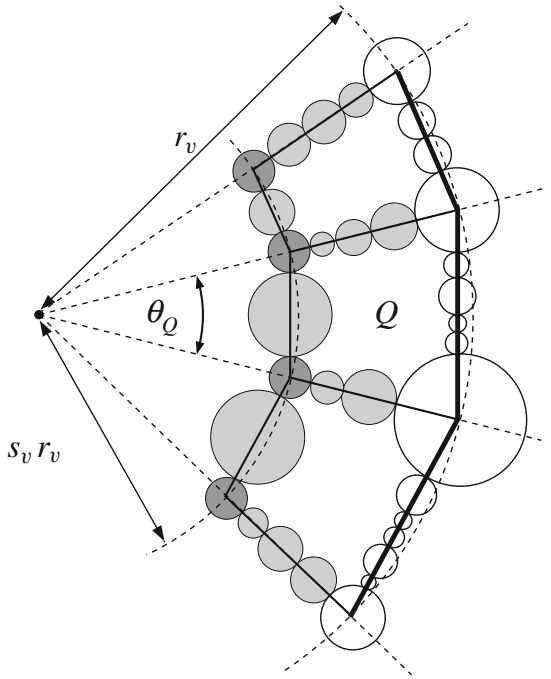

Fig. 17 The boundary arc between the protected and unprotected faces is the union of the longer bases of the protected isosceles trapezoids (thick line). The Gabriel cover of Γ_5 previously constructed restricts to a Gabriel cover of these arcs (white disks). We cover the two opposite corners by disks of same size (dark gray disks), and place one disk between each adjacent pair of these corner disks and complete the Gabriel cover of the radial edges using $O(1)$ disks (light gray)

has bounded eccentricity, the triangulation of Q can be done with a uniform bound strictly below 90°.

We now have the desired acute triangulation Γ_9 that is a refinement of Γ_8 and is a uniformly acute triangulation of all the unprotected faces of Γ_2 and the protected trapezoidal regions of Γ_2. The final step will be to extend it into the protected triangular faces of Γ_2. This reduces to triangulating an isosceles triangle with vertex angle between θ_v and $2\theta_v$ so that there are three mesh vertices on the base side that divide that side into four subsegments each with length between $1/6$ and $1/3$ of the total edge length. Vertices may be added to the radial sides adjacent to v, but the distances of these points to v must be independent of θ, so that triangulations of different faces will "mesh".

12 Step 9: Meshing the Protected Triangles

To extend the triangulation into the protecting triangles containing v, it suffices to prove the following result. Recall from Sect. 2 that a (L, H, θ)-isosceles trapezoid has shorter base length L, height H (the distance between base sides), and angles

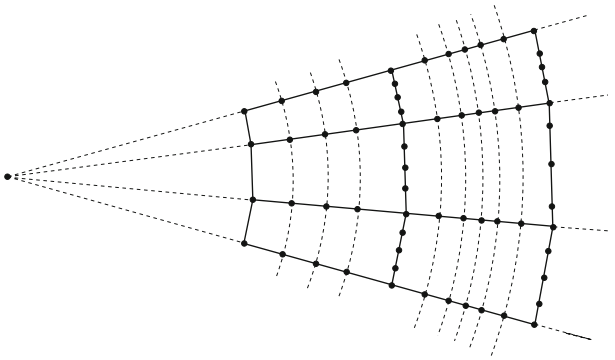

Fig. 18 The proof of Lemma 12.1 comes in two steps: acutely triangulating a trapezoid to interpolate between an undivided base and a base divided into four equal pieces and then interpolating between this and an edge symmetrically, but unevenly divided

equal to $90° \pm \theta$. The radial sides have length $H \sec\theta$, so a $(L, H\cos\theta, \theta)$-isosceles trapezoid has radial sides of fixed length H as θ varies.

Lemma 12.1 *There is an $H_1 > 0$ and $\theta_1 > 0$ so that the following holds. If $1/6 \le a \le 1/3$, $1 \le L \le 2$, and $0 \le \theta \le \theta_1$ then any $(L, H_1\cos\theta, \theta)$-isosceles trapezoid has an acute triangulation so that*

(i) *there are no boundary vertices on the shorter base side e,*
(ii) *boundary vertices on the radial sides are symmetrically placed and their distances from e are independent of L and θ,*
(iii) *the mesh divides the other base side of length $L' = s + 2H_1\sin(\theta/2)$ into four subsegments of lengths (top to bottom) aL', $(1/2 - a)L'$, $(1/2 - a)L'$, aL'.*

The upper angle bound is strictly less than $90°$ and is independent of a, L, and θ within the given ranges.

Figure 18 illustrates the basic idea: we break the proof into two steps. Each isosceles triangle with vertex v will be divided into a smaller isosceles triangle and two isosceles trapezoids. On the first trapezoid (the one adjacent to the triangle), we interpolate between a single segment on the shorter base edge and four equally sized segments on the other base edge. On the second trapezoid, the triangulation interpolates between the three evenly spaced points on the shorter base and the three points given in Lemma 12.1 on the other base. The idea is to do this explicitly for rectangles (the case $\theta = 0$) and then use the stability of acute triangulations to deduce the case when $\theta > 0$ is sufficiently small. In order to join the triangulations of the trapezoids to be a triangulation of their union, we need them to have the same boundary points on the radial sides; hence condition (ii) in the lemma.

We will deal with each sub-problem separately. We first show we can solve it for rectangles ($\theta = 0$) and then use the stability of acute triangulations under small perturbations to solve it for sufficiently small θ.

Lemma 12.2 *There is a $0 < H_2 < \infty$ so that for every $s \in [1, 4]$, the rectangle $R_2(s) = [0, H_2] \times [-s, s]$ can be acutely triangulated so that the only boundary*

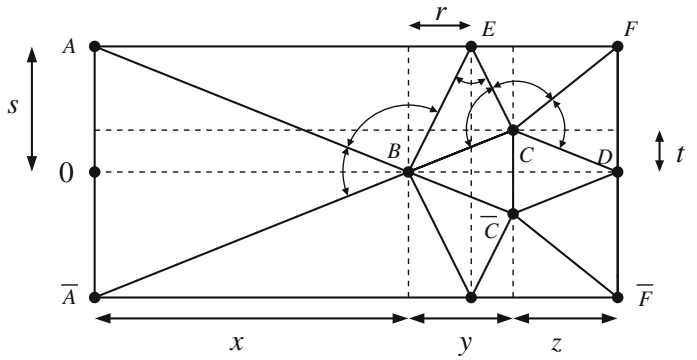

Fig. 19 The abstract graph of our acute triangulation in Lemma 12.2. This figure does not have accurate angles or side lengths, but is drawn to show the labels more clearly; the actual mesh is drawn in Fig. 20. The six marked angles might become obtuse for some parameter values, so need to be checked by computation

vertices of the mesh are the four corners of R_2, the midpoint of the right side and symmetrically placed points on the top and bottom sides; the first coordinates of these points do not depend on s. Every mesh has the same abstract underlying graph and the vertices move continuously with s, so by compactness, the maximum angle used is strictly below $90°$ *with a bound independent of* s.

Proof The proof is simply a picture (Fig. 19) and a calculation: we draw the triangulation and compute all the angles. The figure is symmetric with respect to the real axis and the vertices in the upper half-plane are given by (using complex notation)

$$A = is, \quad B = x, \quad C = x + y + it, \quad D = x + y + z,$$
$$E = x + r + is, \quad F = D + is.$$

There are 11 triangles in the mesh, hence 33 interior angles to check. By symmetry, fifteen of these in the lower half-plane are equal to angles in the upper half-plane. Thus there are really only eighteen distinct angle values. Eight of the remaining ones, like $\angle BAE$, $\angle BA\overline{A}$, $\angle AEB$, are clearly proper sub-angles of $90°$ angles and hence they each have an upper bound that is strictly below $90°$, if the parameters vary over a compact set. In particular, if we fix values of r, t, x, y, z and restrict s to the interval $[1, 4]$, we obtain such bounds.

The angles $\angle CB\overline{C}$ and $\angle CD\overline{C}$ are fixed independently of s and are acute if $t < \min(y, z)$, which will be the case for our choices below.

This leaves only six "ambiguous" angles that might possibly be $\geq 90°$ for some parameter values. Each of these is easily computed in terms of the given parameters r, s, t, x, y, z as follows:

$$\angle AB\overline{A} = 2 \arctan \frac{s}{x},$$
$$\angle ECB = 180° - \arctan \frac{y}{t} - \arctan \frac{y - r}{s - t},$$

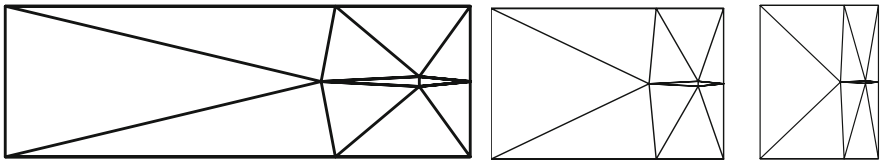

Fig. 20 The abstract triangulation from Fig. 19 but drawn with the parameter values r, t, x, y, z given in the text and $s = 1, 2, 4$

$$\begin{aligned}
\angle FCD &= 180 - \arctan\frac{z}{t} - \arctan\frac{z}{s-t},\\
\angle ABE &= 180^\circ - \arctan\frac{s}{x} - \arctan\frac{s}{r},\\
\angle BEC &= 180^\circ - \arctan\frac{s-t}{y-r} - \arctan\frac{s}{r},\\
\angle ECF &= 180 - \arctan\frac{s-t}{z} - \arctan\frac{s-t}{y-r}.
\end{aligned}$$

We are interested in finding fixed values of r, t, x, y, z so that all these angles are $< 90^\circ$ for all values of $1 \le s \le 4$. After testing various values of r, t, x, y, z, I settled on using

$$r = .18906, \quad t = .065824, \quad x = 4.1945309, \quad y = 1.3035010, \quad z = .675759.$$

I do not claim that these values minimize the maximum angle used, only that all thirty-three of the corresponding angles are strictly less than 90°, for r, t, x, y, z fixed as above and for all $s \in [1, 4]$. See Fig. 20 for the shape of the mesh for three values of s. Once the values of r, t, x, y, z are fixed, it is easy to check that the first three angles listed above are increasing as functions of s, and the last three are decreasing. Thus we only need to check angles when $s \in \{1, 4\}$. Calculating the angles for $s = 1$ shows the largest is $\angle ABE \approx 87.2933^\circ$, and for $s = 4$ the largest is $\angle AB\overline{A} \approx 87.2952^\circ$. Indeed, checking all 33 angles in the configuration gives the same maximum. Thus the lemma holds with $H_2 = x + y + z \approx 6.1738$. □

For the reader's convenience, plots of the six ambiguous angles for $1 \le s \le 4$ are shown in Fig. 21, as are plots of all the angles in the triangulation.

Corollary 12.3 *If $H_3 = 3H_2/2$, then for every $s \in [1, 4]$, the rectangle $R_3(s) = [0, H_3] \times [-s, s]$ can be acutely triangulated so that the only boundary vertices of the mesh are the four corners of R_3, three evenly spaced on the right side, and three symmetrically placed points on the each of top and bottom sides (the first coordinates of these points are independent of s).*

Proof Shrink the rectangle R_1 (and the triangulation from Lemma 12.2) by a factor of 2 and add two copies of this to the right side of R_1, as illustrated in Fig. 22. □

Corollary 12.4 *If $H_4 = 2H_2$, then for every $s \in [1, 4]$, the rectangle $R_4(s) = [0, H_4] \times [-s, s]$ can be acutely triangulated so that the only boundary vertices of the mesh are*

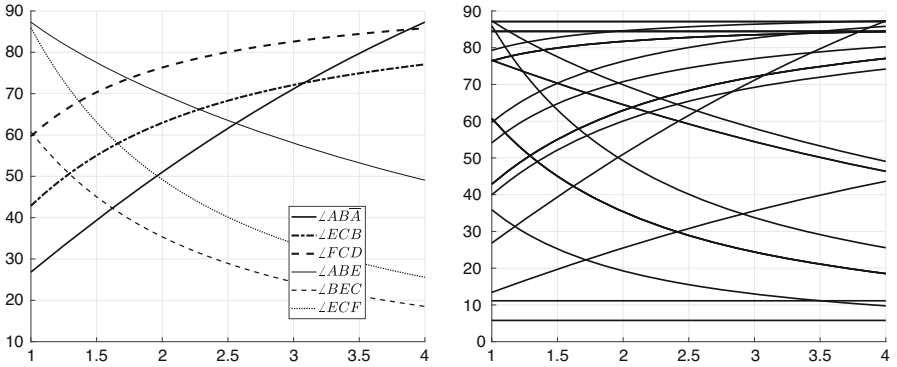

Fig. 21 On the left are the six possibly-not-acute angles in Fig. 19 plotted for $1 \le s \le 4$ and with the fixed values of r, t, x, y, z given in the text. On the right, all 18 distinct angles are plotted. The maximum in both plots is ≈ 87.2952 taken by $\angle AB\bar{A}$ at $s = 4$

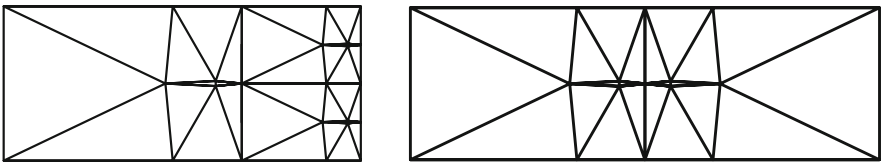

Fig. 22 Proof of Corollaries 12.3 and 12.4. R_2 consists of three copies of R_1 (left) and R_3 is two copies of R_1 (right). Drawn with $s = 2$

the four corners of R and three symmetrically placed points on each of the top and bottom sides, whose first coordinates do not depend on s. The left and right sides of R_4 are not subdivided by the triangulation.

Proof Just reflect the rectangle and triangulation in Lemma 12.2 over the right-hand edge of the rectangle. See Fig. 22. □

Corollary 12.5 *If $s \in [1, 2]$ and $\theta > 0$ is small enough, then every $(2s, H_4, \theta)$-trapezoid (not necessarily isosceles) has an acute triangulation with ten boundary vertices: the four corners of the trapezoid and three points on each radial side. The latter points are symmetric with respect to the axis and their first coordinates are independent of s and θ.*

Proof We just proved this for rectangles (the case $\theta = 0$), in Corollary 12.4. Since acute triangulations are stable under small enough perturbations it also holds for trapezoids with angles close enough to 90°. More precisely, If we normalize so one base edge of the trapezoid Q is the vertical segment $[-s, s]$ on the y-axis, and the other base edge lies on the vertical line $x = H_4$, then Q is the perturbation of the rectangle $R_4 = [0, H_4] \times [-s, s]$ by the map

$$(x, y) \mapsto (x, y + xy \tan \psi_1), \qquad y > 0,$$

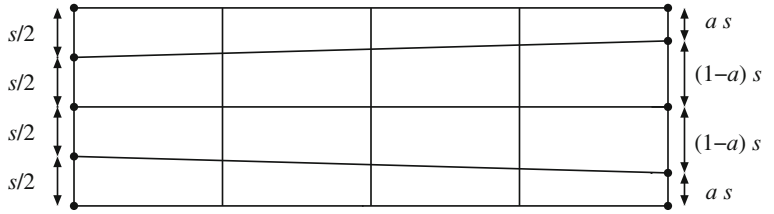

Fig. 23 This is a quadrilateral mesh of the rectangle $R_5(s)$. The left side is subdivided into four equal length segments. The right side is subdivided into segments of length $a, 1-a, 1-a, a$ with $1/3 \le a \le 2/3$

$$(x, y) \mapsto (x, y - xy \tan \psi_2), \qquad y < 0,$$

if the top and bottom edges of Q make angles $\psi_1, \psi_2 \le \theta$ with the horizontal, respectively. If θ is small enough (depending on the upper angle bound in Lemma 12.2), then the perturbed triangulation vertices will form an acute triangulation of Q, with an angle bound that tends to the angle bound in Lemma 12.2 as θ tends to zero. □

The fact that the first coordinates of the mesh points on the top and bottom sides of the trapezoid are independent of the parameters s and θ means that these trapezoids can be "stacked" to form a consistent triangulation. This will be used in the next, and final, step of the proof. Consider Fig. 23. We will assume $1/3 \le a \le 2/3$. If the number of columns in this quadrilateral mesh of the rectangle is large enough, then the quadrilateral mesh elements are themselves close to rectangles. So if we apply Corollary 12.5, we can place an acute triangulation inside each quadrilateral and obtain a triangulation of the whole rectangle. The following makes this idea precise.

Corollary 12.6 *There is an $H_5 > 0$ so that for every $s \in [1, 2]$, and every $1/3 \le a \le 2/3$, the rectangle $R_5(s) = [0, H_5] \times [-s, s]$ can be acutely triangulated so that the only boundary vertices of the mesh are:*

- *the four corners of R_5,*
- *a finite number of symmetrically placed points on the top and bottom sides of R_5 that do not depend on s or a, and*
- *three points on the right side of R_5 that include the midpoint of that side and points that cut each of remaining two segments of length s into subsegments of length as and $(1-a)s$, as shown in Fig. 23.*

Proof We assume $H_5 = m H_4/3$ for some positive integer m. Draw vertical segments in R_5 at the points cutting $[0, H_5]$ into m equal intervals of length $H_4/3$, with endpoints $x_k = k H_5/(3m)$ for $k = 0, \ldots, m$. Connect the three points on the left side of R_5 by line segments to the corresponding three points on the right side of R_5. If m is large enough, then these lines are as close to horizontal as we wish. Thus R_5 is decomposed into $4m$ quadrilaterals that are each θ-trapezoids, for θ as small as we wish.

Suppose Q is one of these quadrilaterals, with left vertical side length t. Then Q us a perturbation of a $H_4/3 \times t$ rectangle, which is clearly similar (by a dilation) to a $H_4 \times 3t$ rectangle. Since $1/3 \le a \le 2/3$, t is bounded between $s/3$ and $2s/3$. Using

this, and the fact that $1 \le s \le 2$, we deduce that

$$1 \le 3 \cdot \frac{s}{3} \le 3t \le 3 \cdot \frac{2s}{3} \le 4.$$

Therefore, by Corollary 12.5, Q has an acute triangulation with no subdivision of either the left or right sides. Moreover, the mesh vertices on the top and bottom edges have been chosen independent of t and the angles of the quadrilaterals. This implies the triangulations of adjacent quadrilaterals match up. Therefore $R_5(s)$ has an acute triangulation for each $1 \le s \le 2$. □

Proof of Lemma 12.1 We have already done this for rectangles, e.g., $\theta = 0$ with $H_1 = H_2 + H_5$. By the stability of acute triangulations it also holds for small enough perturbations. More precisely, if we perturb the rectangle $R_1(s) = [0, H_1] \times [-s, s]$

$$(x, y) \mapsto \left(x \cos\theta, y + x \frac{y}{s} \sin\theta \right),$$

then $R_1(s)$ is mapped to a θ-isosceles trapezoid and the map is isometric on the top and bottom edges, so that the triangulation vertices on these edges preserve their distance from the left side. This proves the lemma for θ_1 small enough. □

Thus if $\theta_0 \le \min(\eta, \theta_1)$, where η is from Lemma 10.1 and θ_1 is form Lemma 12.1, then Theorem 1.1 holds.

13 Proof of Theorem 1.2

In the introduction we mentioned that [9] makes use of a consequence of Theorem 1.2. The precise statement is the following.

Lemma 13.1 *In Theorem 1.2, $\mathcal{T}_2$ can be chosen to satisfy the following additional condition. For each triangle $T = \Delta ABC$ in $\mathcal{T}_1$ and each vertex D of $\mathcal{T}_2$ in the interior of T the angle $\angle DAB$ is greater than $\theta_1 > 0$, where θ_1 depends only on θ.*

Proof Suppose T is an element of $\mathcal{T}_1$ and $T' \subset T$ is an element of $\mathcal{T}_2$. By Theorem 1.2 every triangle in $\mathcal{T}_2$ has all angles bounded strictly below 90° (depending on θ) and hence every angle of T' is also bounded uniformly away from zero (also depending only on θ). The Law of Sines implies that all three edges of T' have comparable lengths, and these lengths are each comparable to the diameter of T'. Similarly, the distance of each vertex of T' to the opposite side of T' is comparable to the diameter of T'. Thus it suffices to show $\operatorname{diam}(T') \simeq \operatorname{diam}(T)$.

Our remarks above imply that triangles in $\mathcal{T}_2$ that share an edge must have comparable diameters. Since only a bounded number (at most $2\pi/\theta$) of the triangles can hit the same vertex, this implies that all triangles of $\mathcal{T}_2$ that share a common vertex also have comparable diameters. Hence triangles of $\mathcal{T}_2$ that are connected by a chain of k triangles in $\mathcal{T}_2$ also have comparable diameters (with constant depending on θ and k). Thus the condition in the lemma holds iff the number of elements of $\mathcal{T}_2$ inside T is

uniformly bounded, depending only on θ. This is the second part of Theorem 1.2 (to be proven next). □

Finally, we come to the proof of Theorem 1.2.

Proof of Theorem 1.2 The first claim follows immediately from Theorem 1.1. The second claim comes from examining the proof of Theorem 1.1: we will show that if we start with a triangulation, then each face is replaced by a bounded number of faces (depending on the minimum angle θ) in each of the nine steps of the construction. We will briefly review each of the steps. Unless otherwise stated, "comparability" below will mean with a constant depending only on θ.

Since we start with a triangulation Γ_1, every face is already a simple polygon, so Lemma 2.1 need not be applied. In Step 1, we are assuming Γ_1 is a triangulation with minimal angle $\theta > 0$, so the argument in the proof of Lemma 13.1 shows that adjacent edges of Γ_1 have comparable lengths. Thus any vertex v has comparable distance to each adjacent vertex. This means that we can take the disk D_v in Step 1 to have diameter that is comparable to the distance from v to any of the adjacent vertices. This choice of D_v is the only change we make to the previous proof. The number of protecting triangles and trapezoids at each vertex is bounded by $O(1/\theta)$ since each subtends an angle $\geq \theta_v \geq \theta$ from v. (Alternatively, in Sect. 4 one could have chosen to subdivide the angle ψ into $\lfloor \psi/\theta_v \rfloor$ equal sub-angles).

In Step 2, the number of vertices we need to add to each unprotected face is bounded in terms of θ. This holds because after removing the disks D_v the distance between the remaining segments on the boundary of a triangle is comparable to the diameter of the triangle, so only a bounded number of points need to be added to make the face $(1/4)$-thick. In Step 3, we quad-mesh each unprotected face by quadrilaterals, and each face uses at most $O(1)$ quadrilaterals, where the constant is linear in the number of boundary segments of the face, and we just saw that this is bounded in terms of θ. Step 4 only perturbs the quadrilaterals and does not increase the number inside any original unprotected face. Step 5 adds only $O(1)$ Gabriel points to each edge of Γ_5 (this constant is independent of θ). The construction of Bern et al. produces a number of triangles that is comparable to the number of Gabriel points on the boundary of each quadrilateral, and Yuan's construction in Step 6 only increases this by a bounded factor. The compactness argument in Step 7 does not increase the number of components; at worst it replaces the number of triangles in one quadrilateral by the number found in a different one, and this has been bounded by the argument so far.

The final two Steps, 8 and 9, extend the triangulation into the protected triangles and trapezoids by an explicit construction, and it follows directly from the construction that the number of sub-triangles created inside D_v is bounded in terms of a lower bound for θ_v, i.e., it is bounded by a constant that depends only on θ. This completes the proof.

References

1. Bern, M., Mitchell, S., Ruppert, J.: Linear-size nonobtuse triangulation of polygons. Discrete Comput. Geom. **14**(4), 411–428 (1995)

2. Bern, M., Shewchuk, J.R., Amenta N.: Triangulations and mesh generation. In: Handbook of Discrete and Computational Geometry, 3rd ed., pp. 763–785. CRC Press, Boca Raton (2017)
3. Bishop, C.J.: Optimal angle bounds for quadrilateral meshes. Discrete Comput. Geom. **44**(2), 308–329 (2010)
4. Bishop, C.J.: Conformal mapping in linear time. Discrete Comput. Geom. **44**(2), 330–428 (2010)
5. Bishop, C.J.: Quadrilateral meshes for PSLGs. Discrete Comput. Geom. **56**(1), 1–42 (2016)
6. Bishop, C.J.: Nonobtuse triangulations of PSLGs. Discrete Comput. Geom. **56**(1), 43–92 (2016)
7. Bishop, C.J.: Optimal triangulation of polygons (2021). https://www.math.stonybrook.edu/~bishop/papers/opttri.pdf
8. Bishop, C.J.: Uniformly acute triangulations of polygons. Discrete Comput. Geom. (2023). https://doi.org/10.1007/s00454-023-00525-w
9. Brunck, F.: Acute triangulations of spherical and hyperbolic triangle complexes. Preprint (2022)
10. Brunck, F.: Iterated medial triangle subdivision in surfaces of constant curvature. Discrete Comput. Geom. (2023). https://doi.org/10.1007/s00454-023-00500-5
11. Burago, Yu.D., Zalgaller, V.A.: Polyhedral embedding of a net. Vestnik Leningrad. Univ. **15**(7), 66–80 (1960). (in Russian)
12. Gabriel, K.R., Sokal, R.R.: A new statistical approach to geographic variation analysis. Syst. Zool. **18**(3), 259–278 (1969)
13. Maehara, H.: Acute triangulations of polygons. Eur. J. Combin. **23**(1), 45–55 (2002)
14. Mumford, D.: A remark on Mahler's compactness theorem. Proc. Am. Math. Soc. **28**, 289–294 (1971)
15. Saraf, Sh.: Acute and nonobtuse triangulations of polyhedral surfaces. Eur. J. Combin. **30**(4), 833–840 (2009)
16. Yuan, L.: Acute triangulations of polygons. Discrete Comput. Geom. **34**(4), 697–706 (2005)
17. Zamfirescu, C.T.: Survey of two-dimensional acute triangulations. Discrete Math. **313**(1), 35–49 (2013)