CONFORMAL MAPS AND OPTIMAL MESHES

Christopher J. Bishop
Stony Brook University

Duke University, Feb. 24, 2014

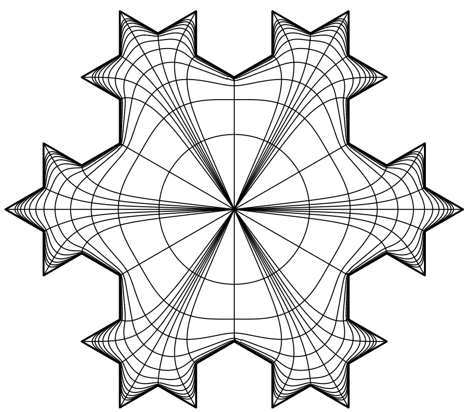 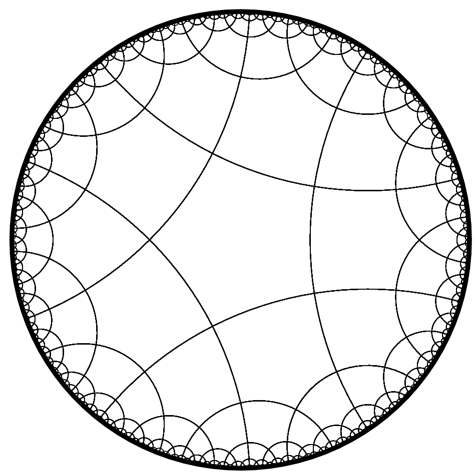 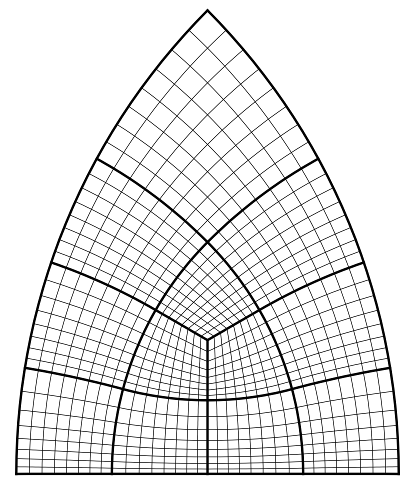

`www.math.sunysb.edu/~bishop/lectures`

Riemann Mapping Theorem: If $\Omega \subsetneq \mathbb{R}^2$ is simply connected, then there is a conformal map $f : \mathbb{D} \to \Omega$.

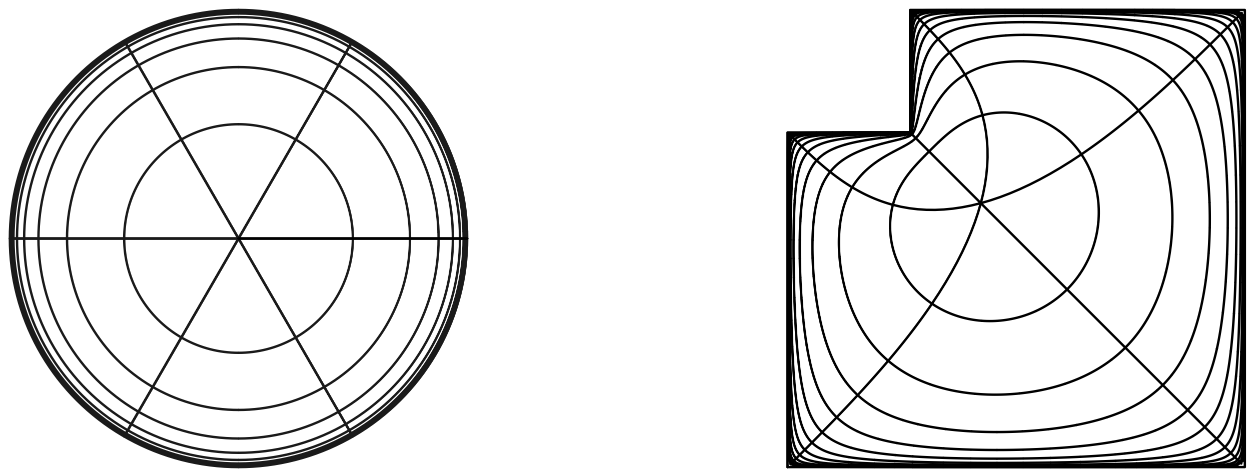

Conformal = angle preserving

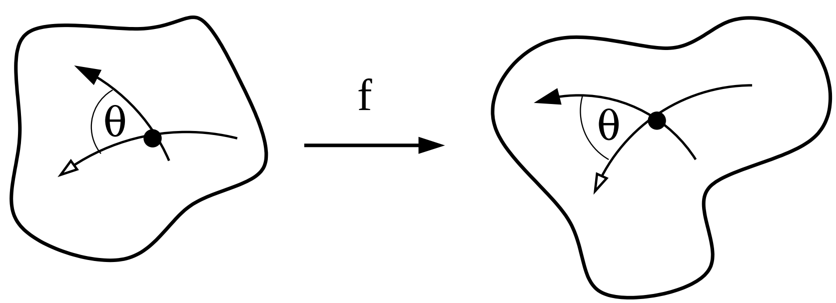

Fact: Conformal map to a polygon is determined by the points on circle that map to the vertices.

Problem: given n-gon, compute these n points.

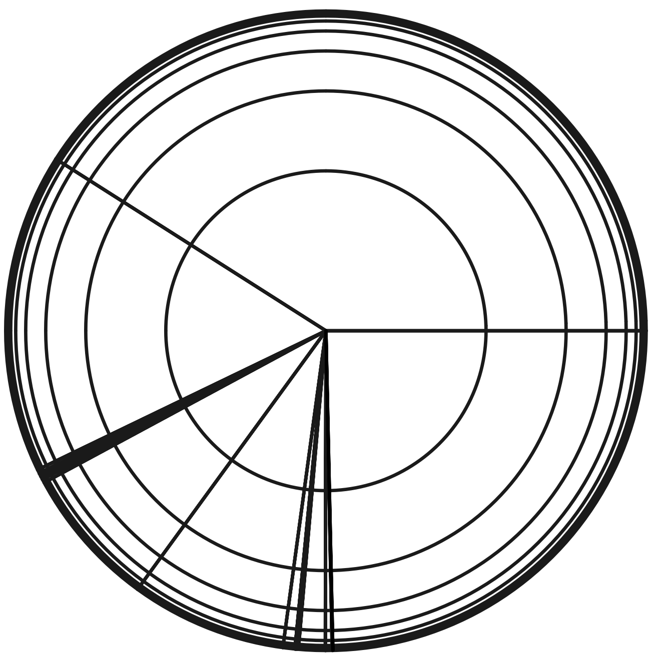

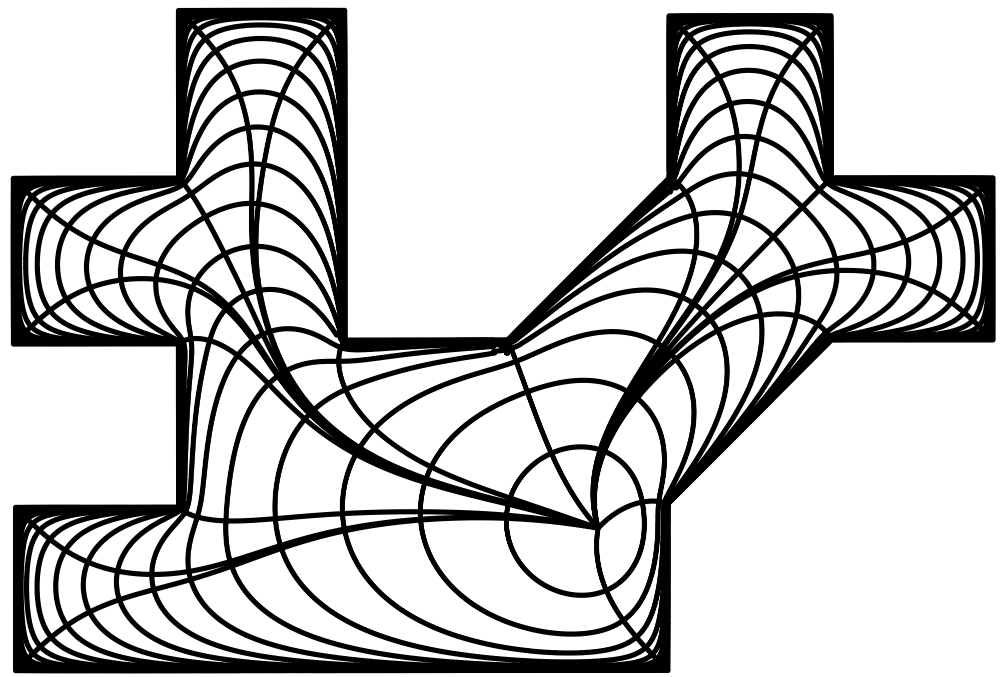

Schwarz-Christoffel formula (1867):

$$f(z) = A + C \int^{z} \prod_{k=1}^{n} (1 - \frac{w}{z_k})^{\alpha_k - 1} dw,$$

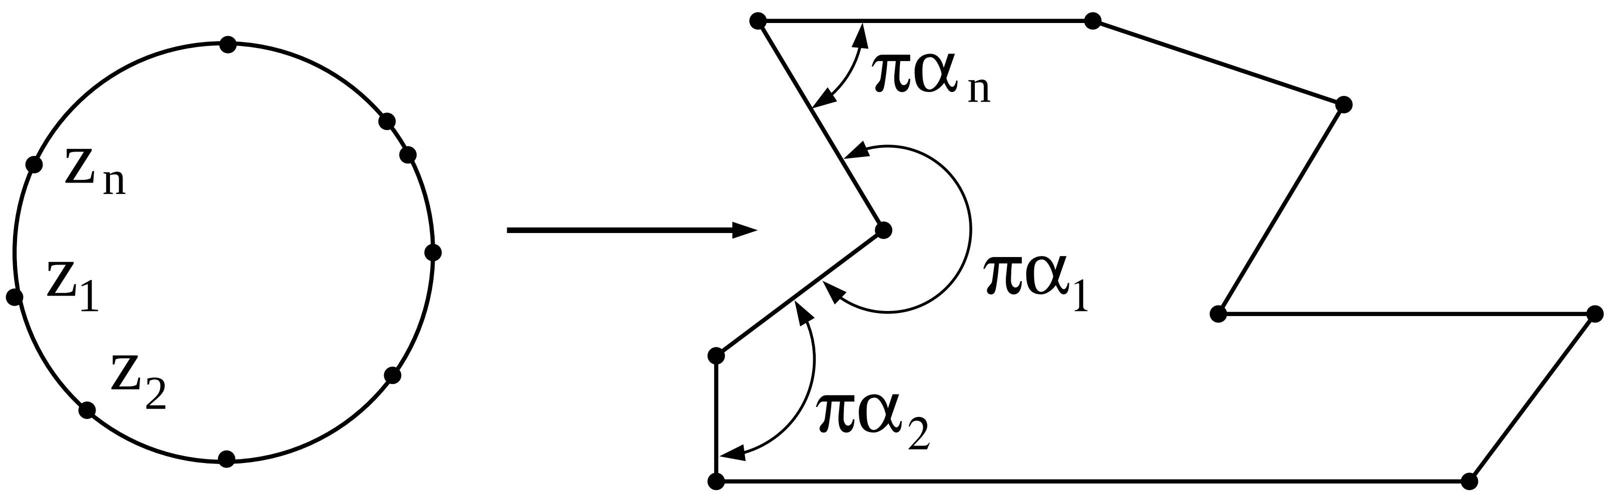

α's known

z's unknown (= **SC-parameters** = **pre-vertices**)

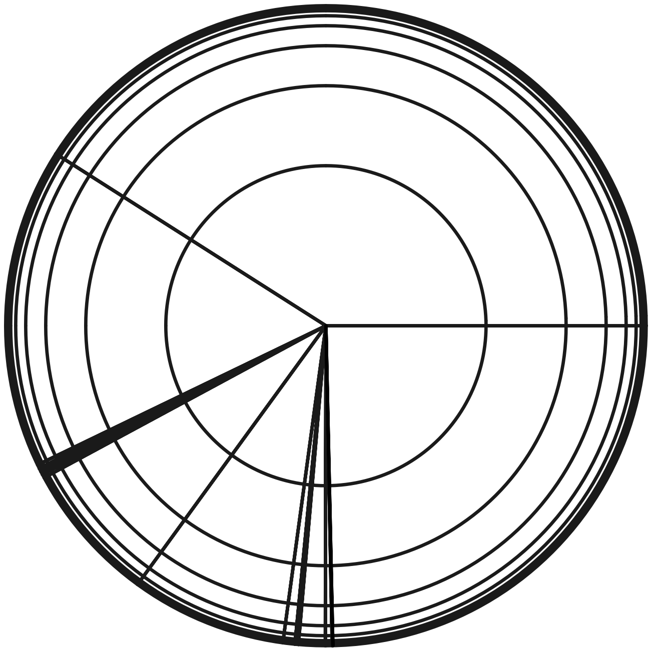

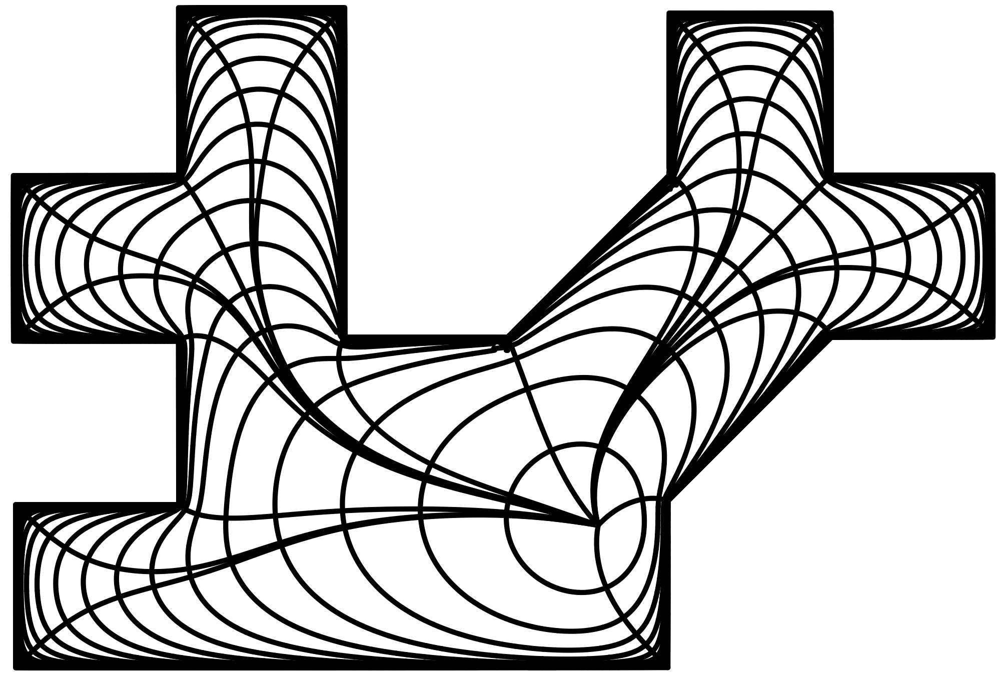

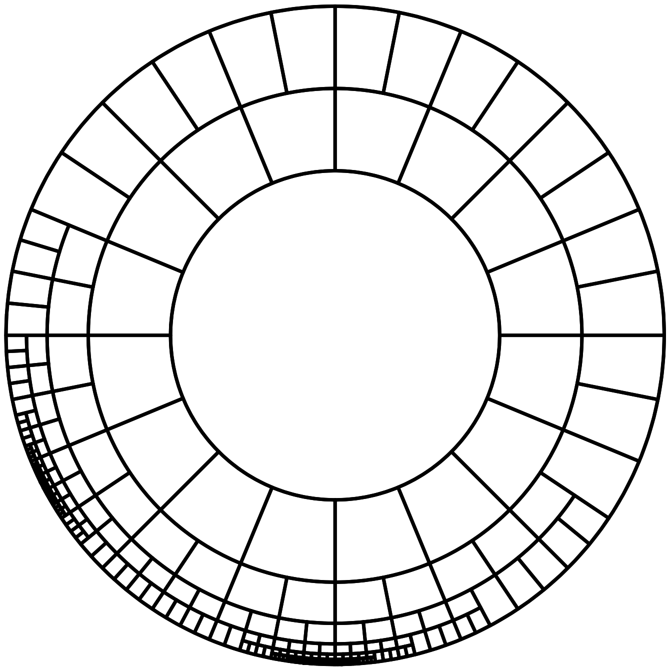

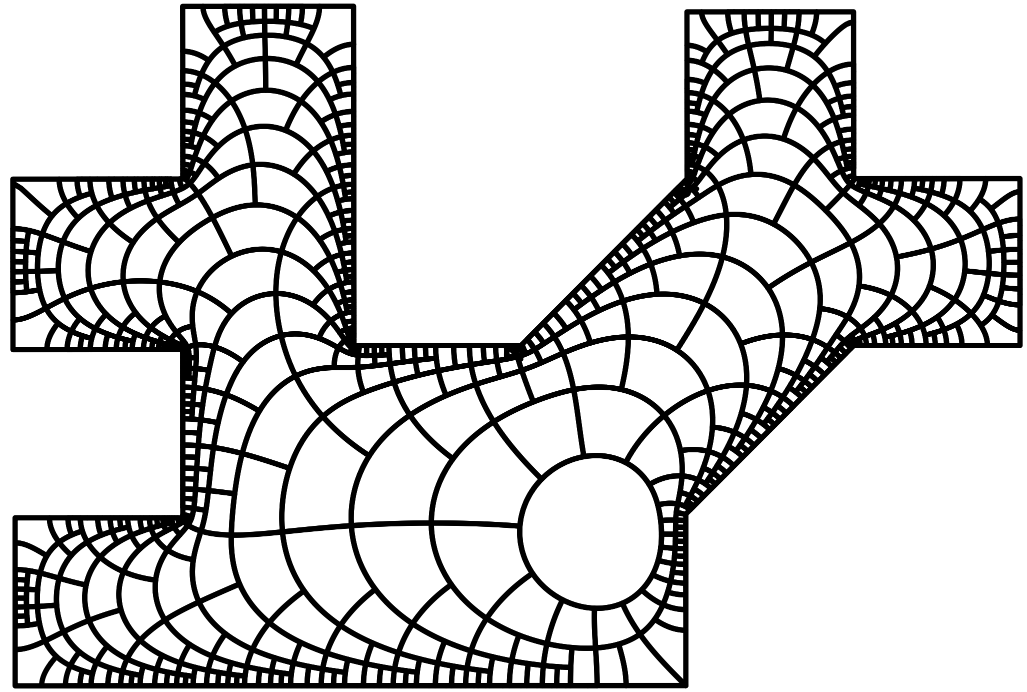

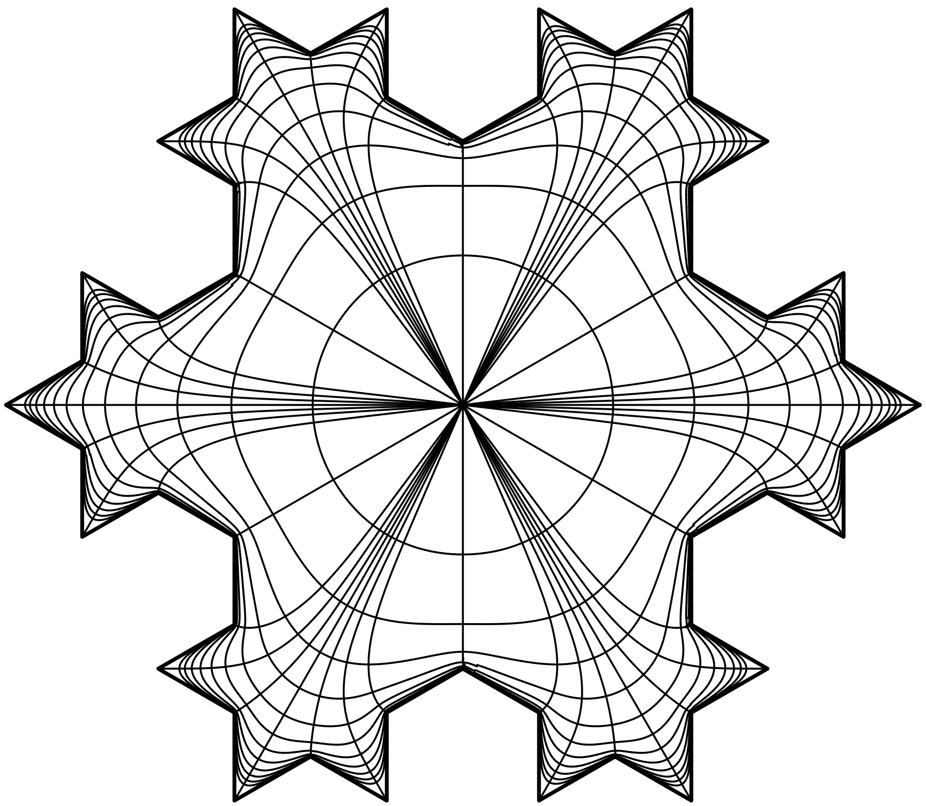
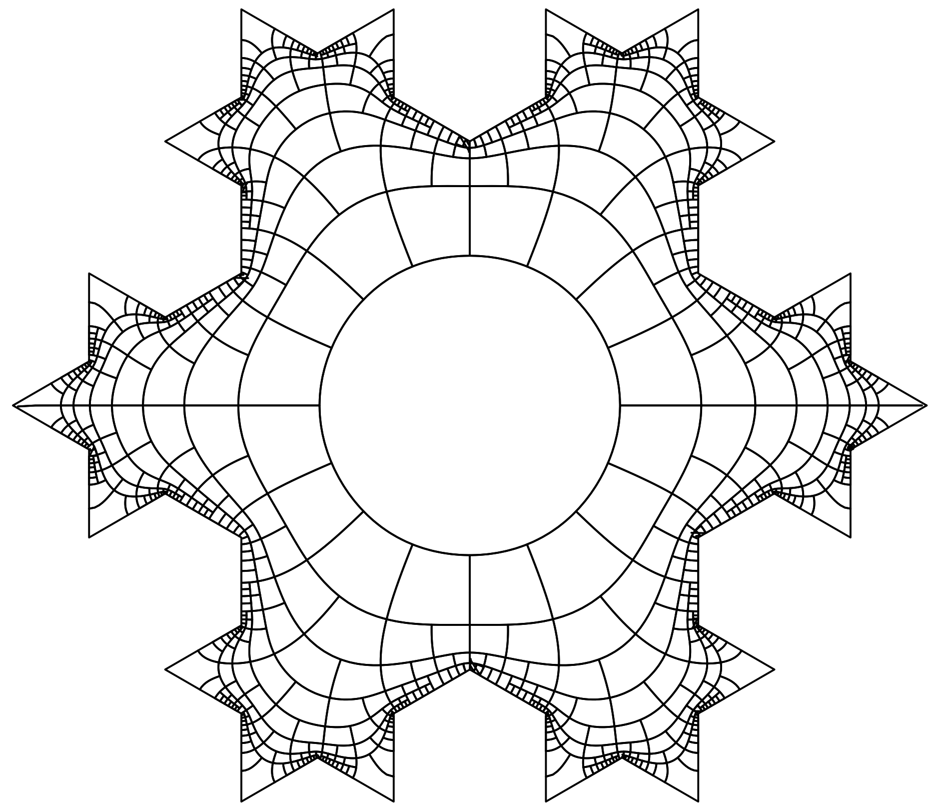

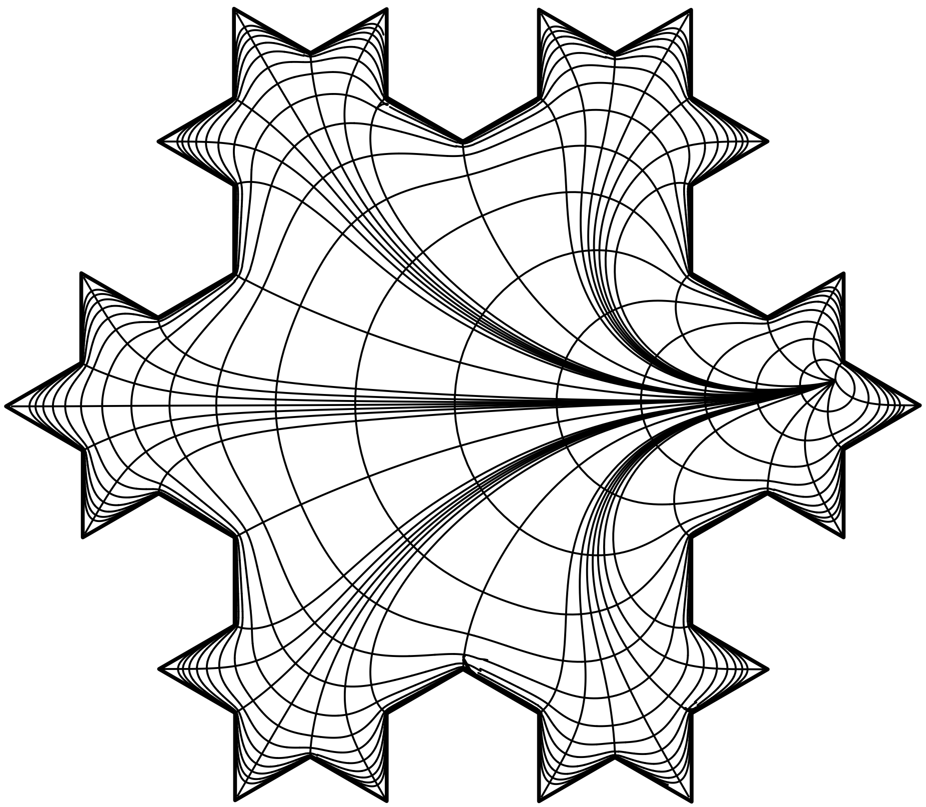

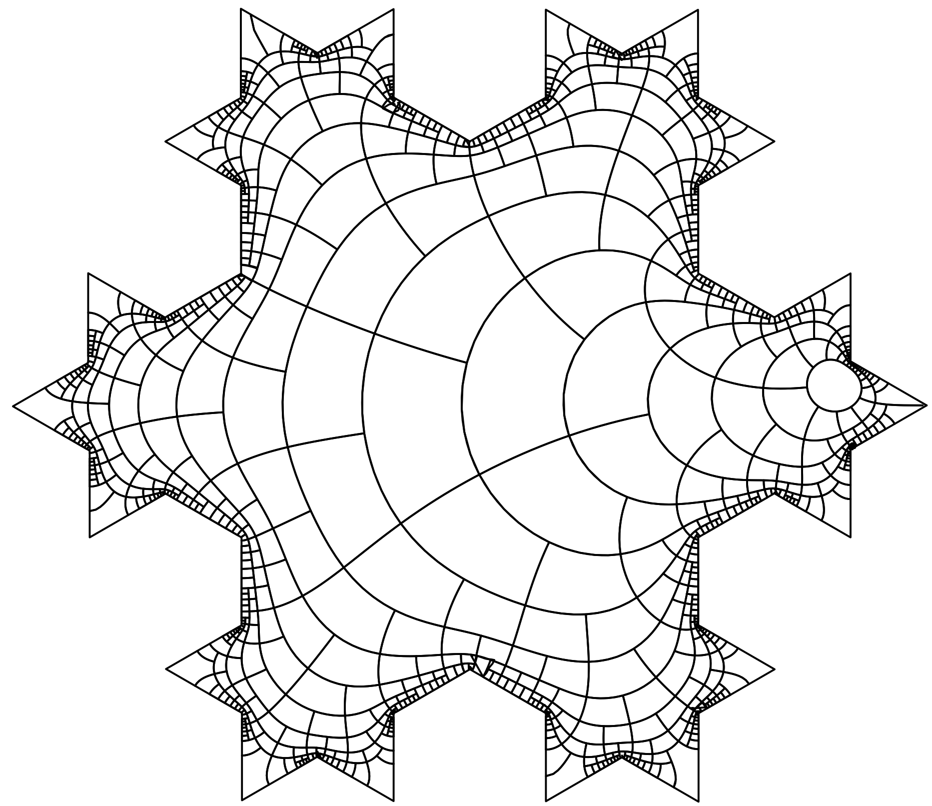

How fast can we compute the pre-vertices?

How fast can we compute the pre-vertices?

Theorem: Can compute SC-parameters in time $C_\epsilon \cdot n$.

How fast can we compute the pre-vertices?

Theorem: Can compute SC-parameters in time $C_\epsilon \cdot n$.

ϵ = error in quasiconformal sense.

Approximate pre-vertices can be mapped to true pre-vertices by $(1+\epsilon)$-QC map of disk.

Implies uniform approximation.

How fast can we compute the pre-vertices?

Theorem: Can compute SC-parameters in time $C_\epsilon \cdot n$.

ϵ = error in quasiconformal sense.

Approximate pre-vertices can be mapped to true pre-vertices by $(1+\epsilon)$-QC map of disk.

Implies uniform approximation.

$C_\epsilon = O(\log \frac{1}{\epsilon} \log \log \frac{1}{\epsilon})$.

Quasiconformal maps: bounded angle distortion

Roughly, a K-QC map multiplies angles by $\leq K$.

More precisely: tangent map sends circles to ellipses a.e.

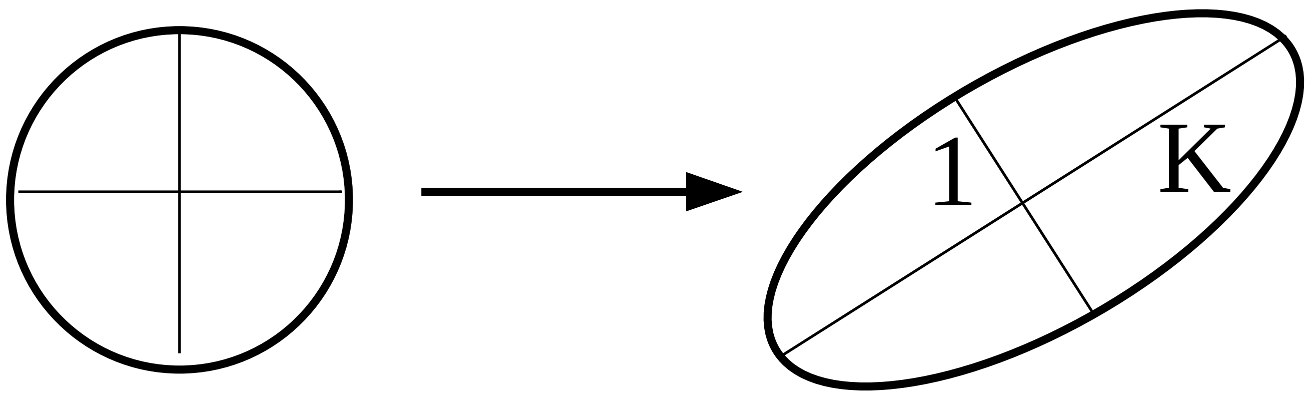

1-QC = conformal

complex dilatation $\mu = \mu_f = f_{\overline{z}}/f_z$.

f is K-QC iff $|\mu| \leq k = \frac{K-1}{K+1}$.

Measurable Riemann Mapping Theorem:
Given QC map $g : \mathbb{D} \to \Omega$, there is there is a QC $f : \mathbb{D} \to \mathbb{D}$ so that $g \circ f$ is conformal.

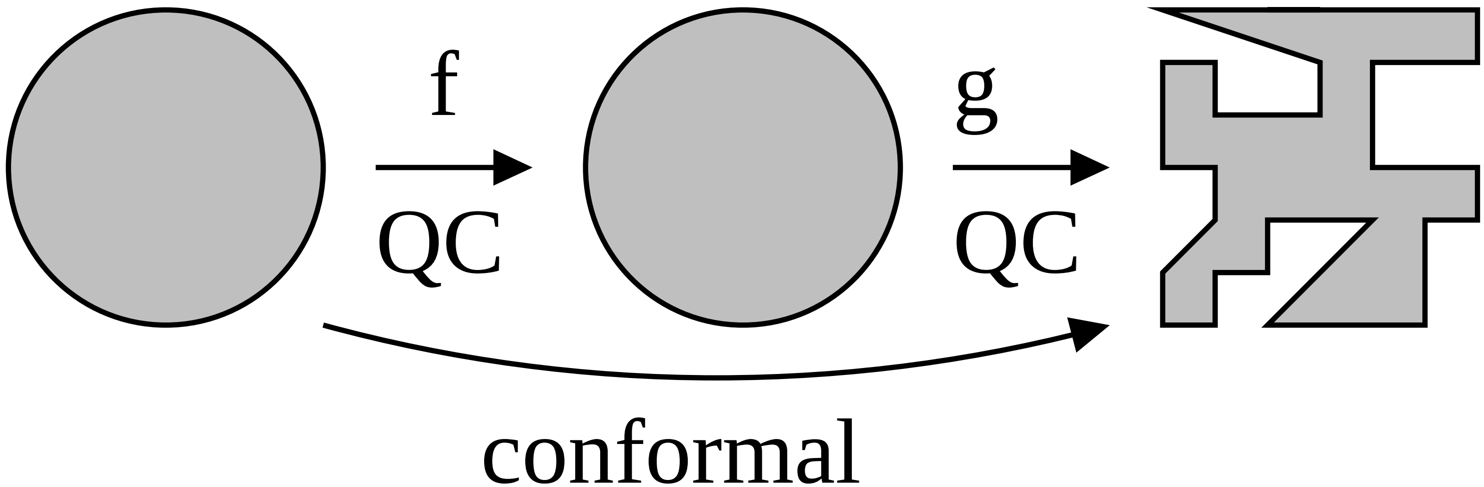

Beltrami equation $f_{\overline{z}} = \mu f_z$ has solution.

Several proofs. One gives f as power series of singular integral operators applied to μ_g.

Method:

- Find initial QC map to polygon.
- Approximate Beltrami solution to lower QC constant
- Iterate

RIETMANN RIETMANN

RIETMANN RIETMANN

RIETMANN RIETMANN

RIETMANN RIETMANN

RIETMANN RIETMANN

Need:

Fast way to approximate Beltrami solution (multipole)

Estimate of radius of convergence (proof of MRMT)

Initial guess inside convergence region (continuation)

Initial guess bounded distance from answer

Rest of talk focuses only on last point.

Need initial guess that is **fast** to compute and guaranteed **close** to correct answer.

Need initial guess that is **fast** to compute and guaranteed **close** to correct answer.

- **fast** comes from computational geometry.

- **close** comes from hyperbolic geometry.

Medial axis:

centers of disks that hit boundary in at least two points.

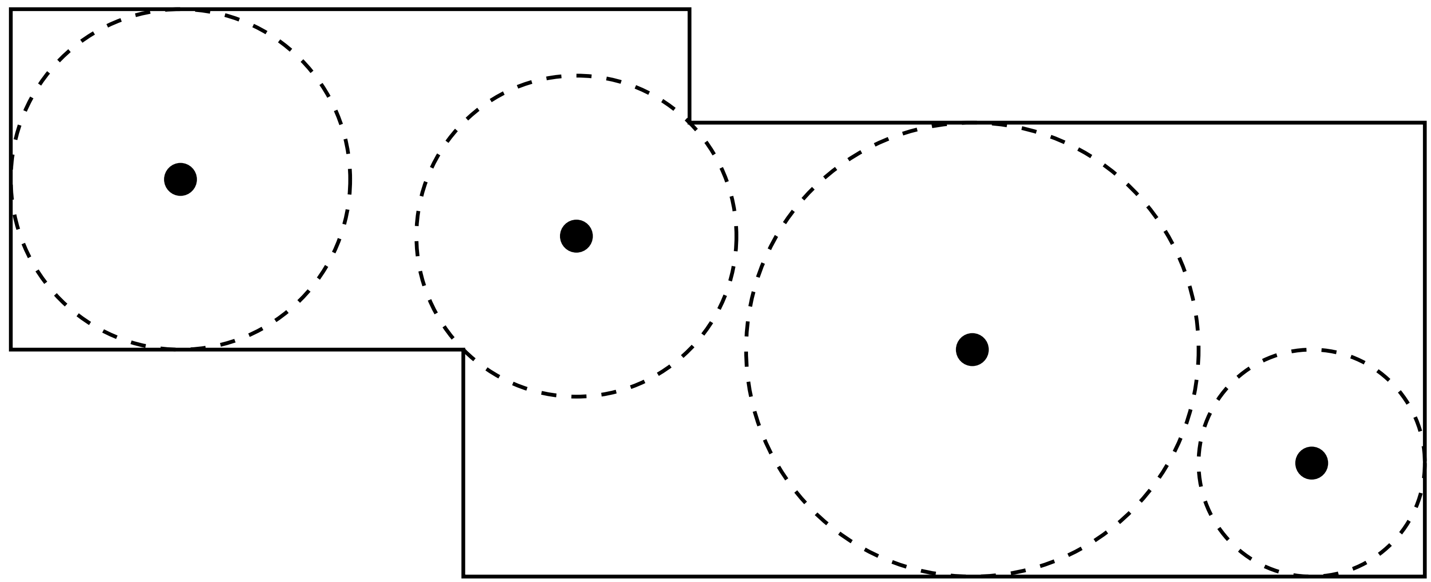

Medial axis:
centers of disks that hit boundary in at least two points.

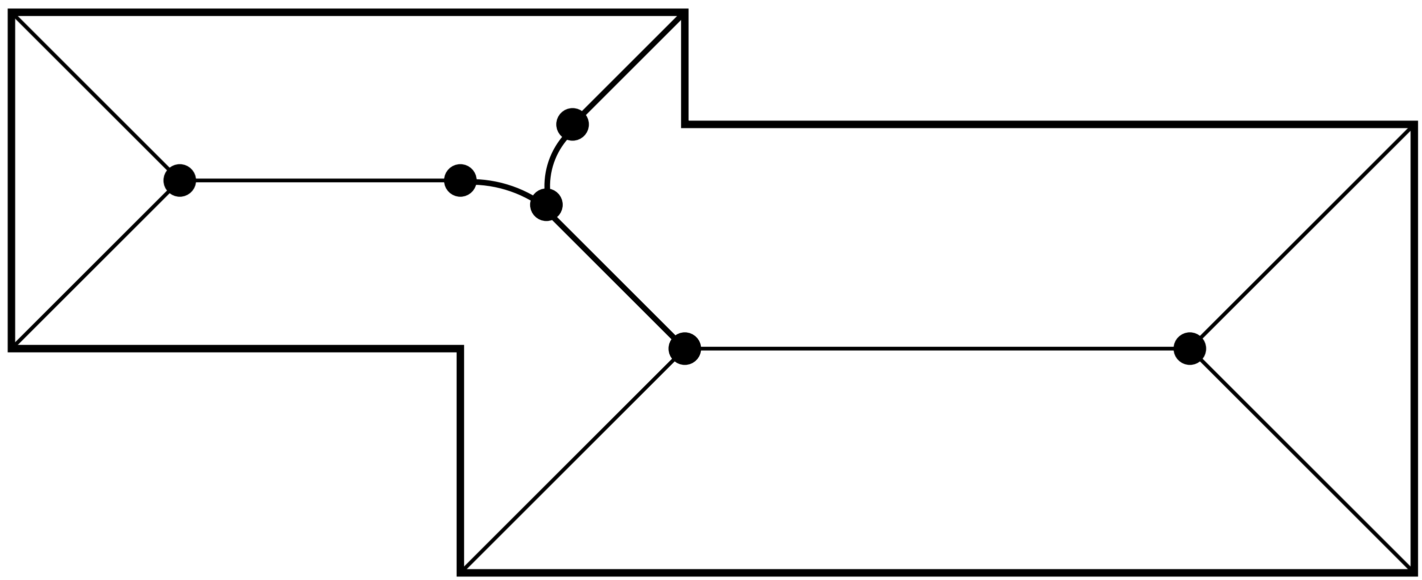

Medial axis of a polygon is a finite tree.

Computable in O(n), Chin-Snoeyink-Wang.

Medial axis:
centers of disks that hit boundary in at least two points.

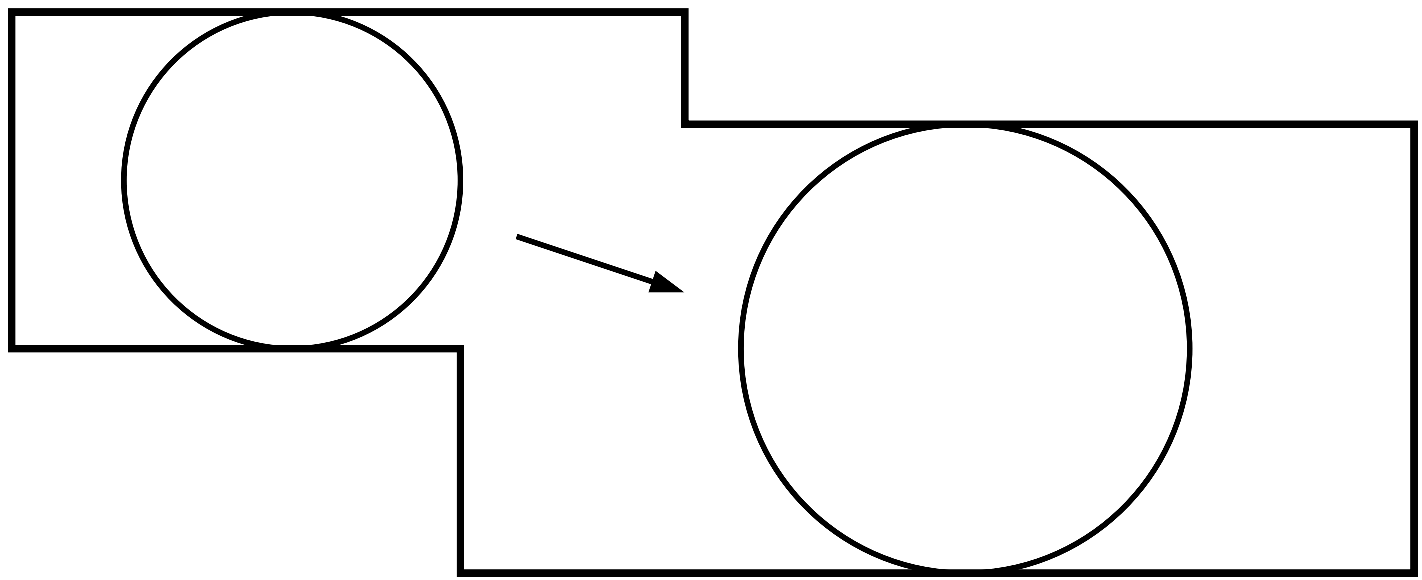

There is a “natural” choice of conformal map between any two medial axis disks.

A **Möbius transformation** is a map of the form

$$z \to \frac{az+b}{cz+d}.$$

A **Möbius transformation** is a map of the form

$$z \to \frac{az+b}{cz+d}.$$

Exactly the 1-1, conformal self-maps of $S^2 = \mathbb{R}^2 \cup \{\infty\}$.

A **Möbius transformation** is a map of the form

$$z \to \frac{az+b}{cz+d}.$$

All conformal maps between disks have this form.

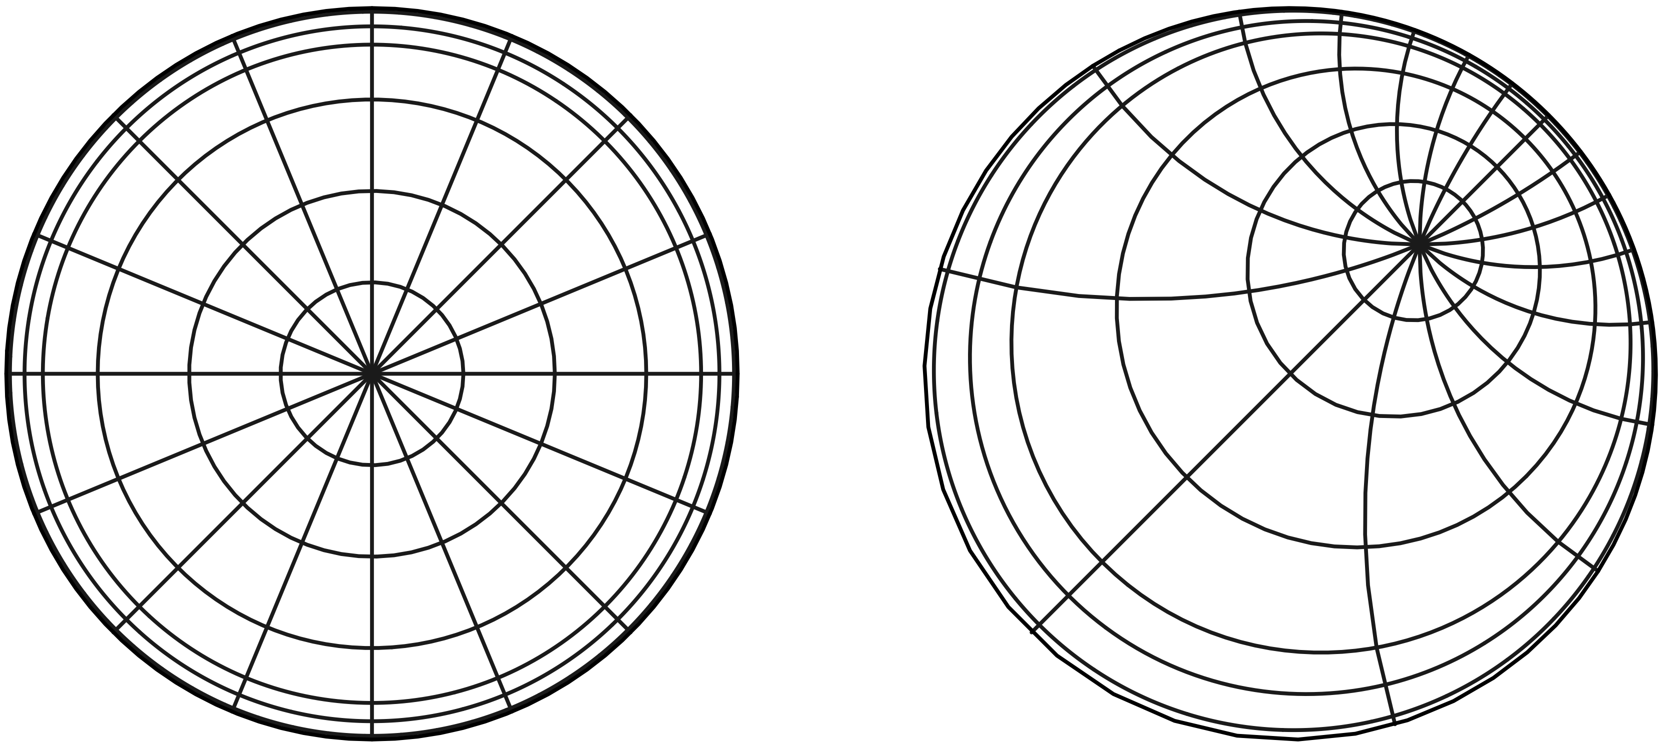

A **Möbius transformation** is a map of the form

$$z \to \frac{az+b}{cz+d}.$$

Form a group under composition.

Uniquely determined by images of 3 distinct points.

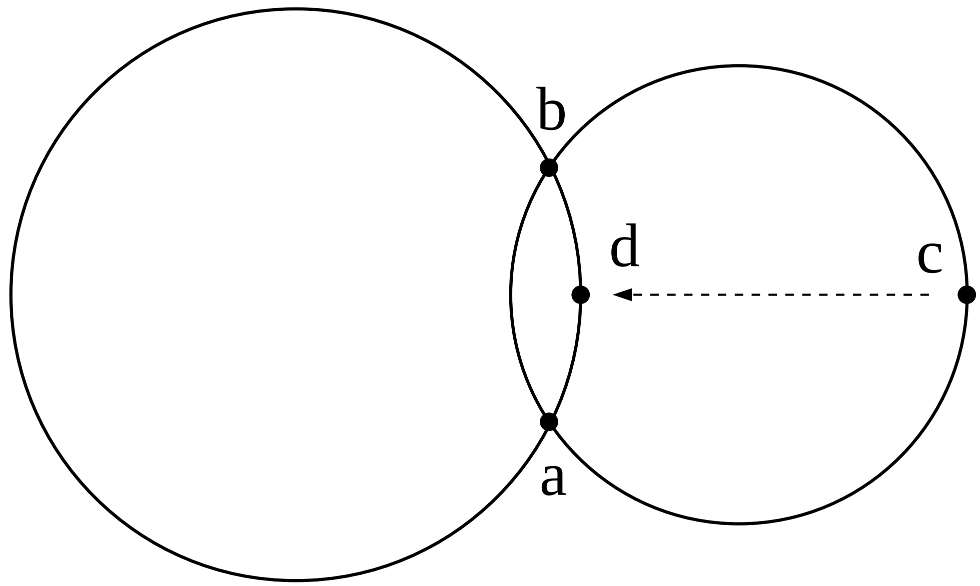

Fix intersection points a, b and map $c \to d$ as shown.

Determines unique Möbius map between disks.

Part of 1-parameter symmetric family fixing a, b.

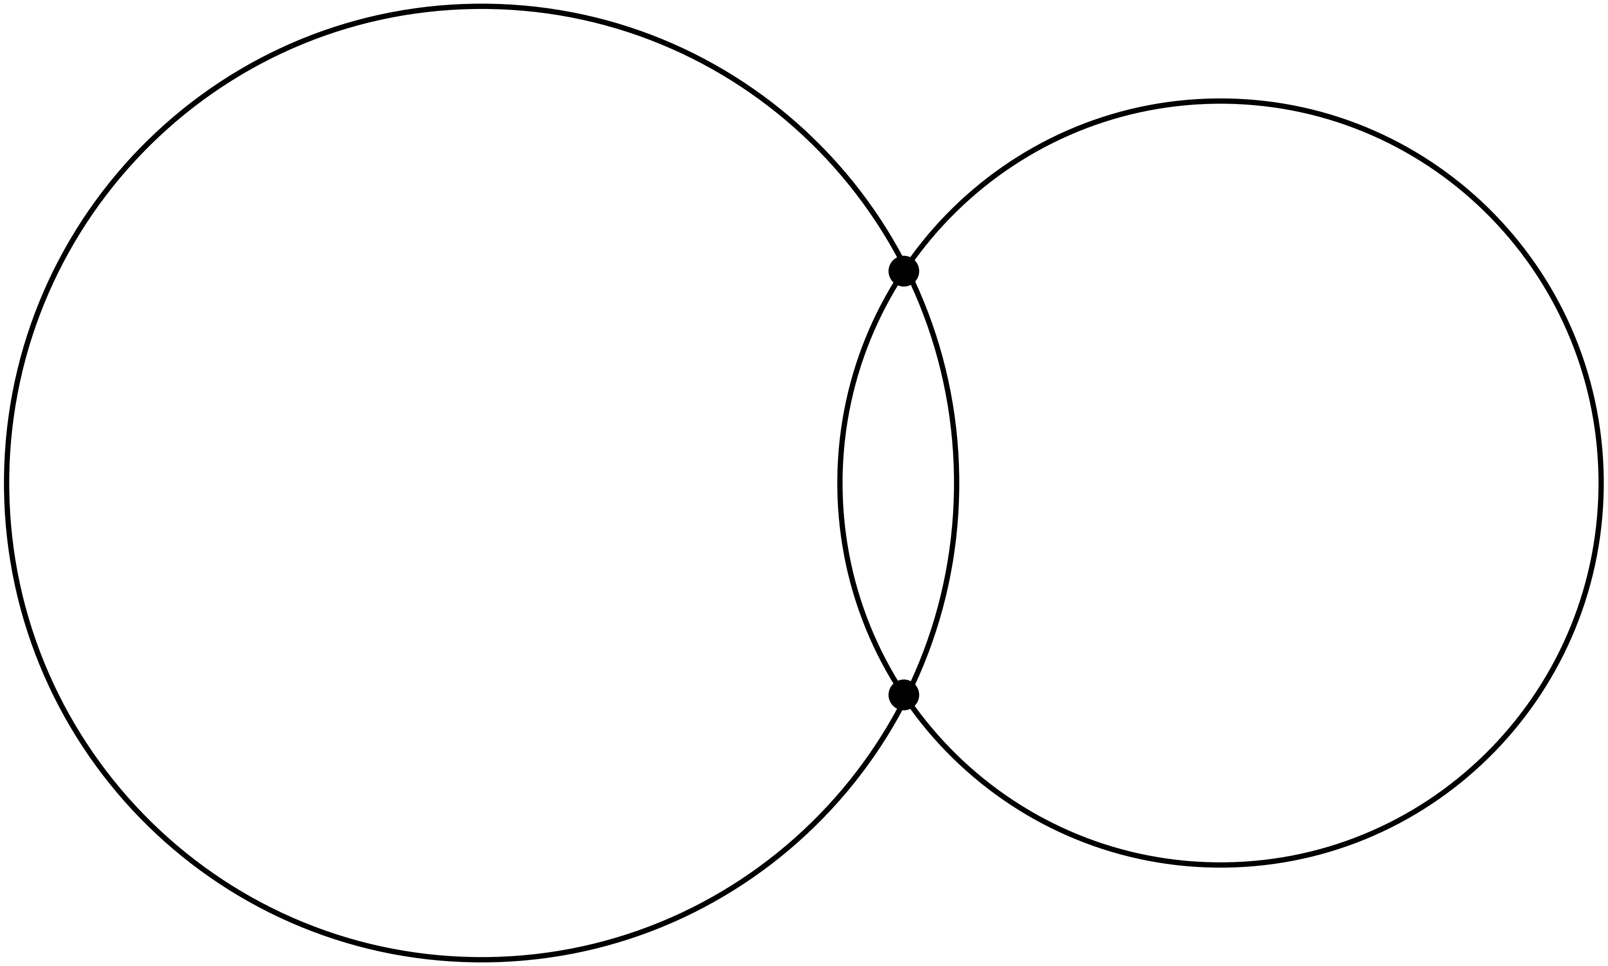

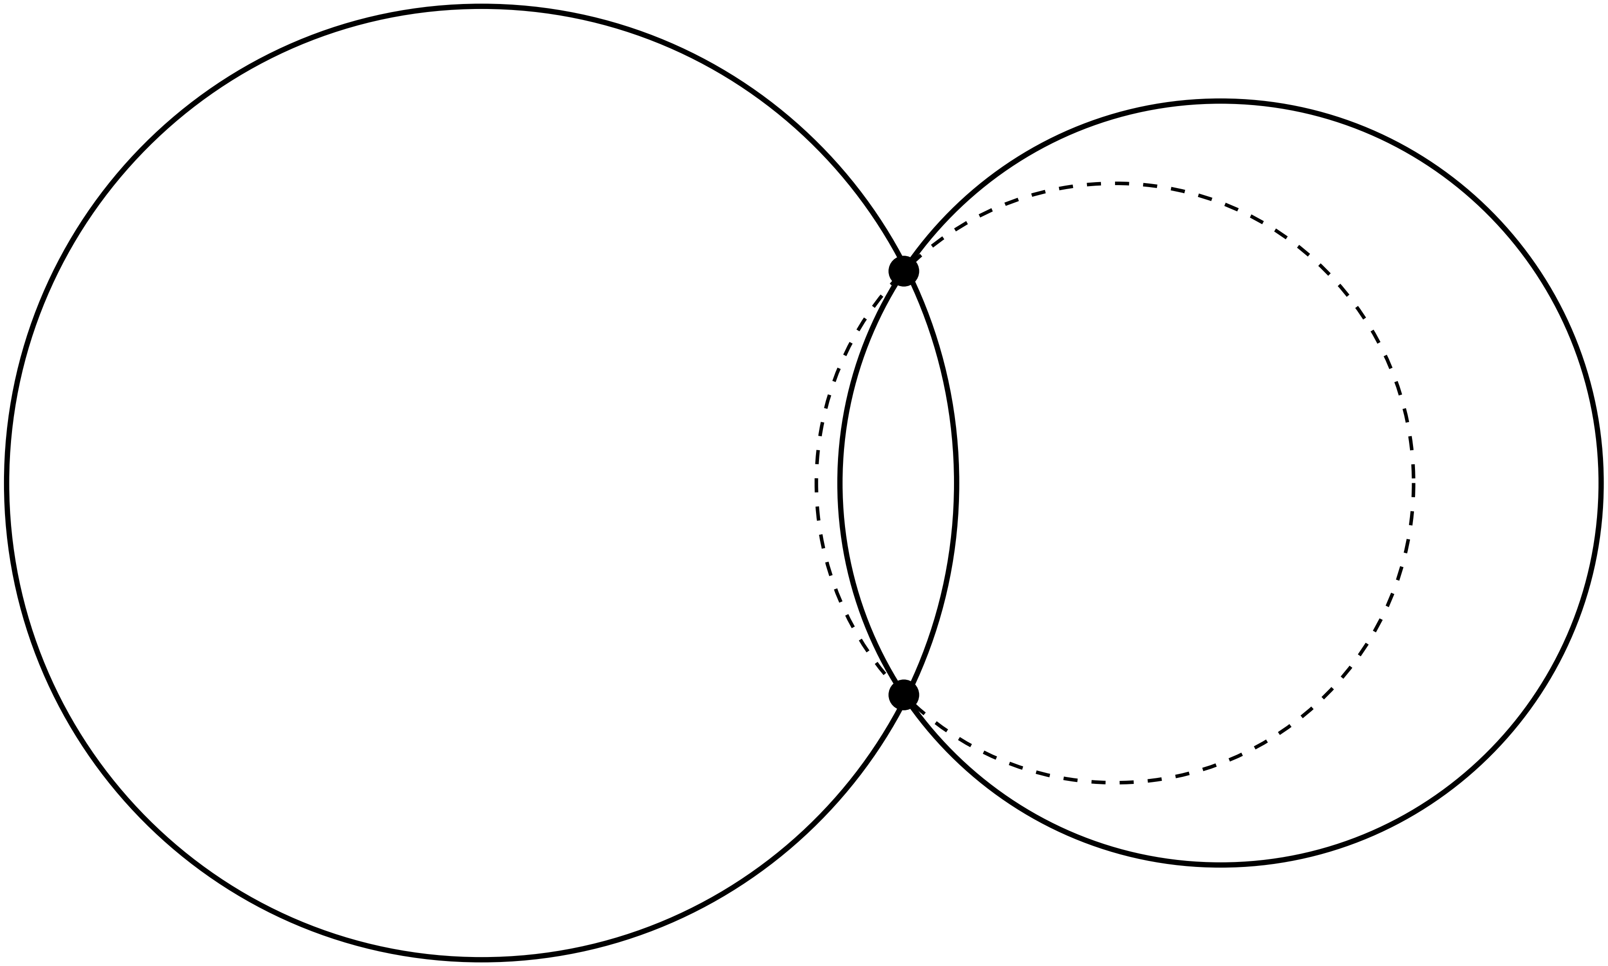

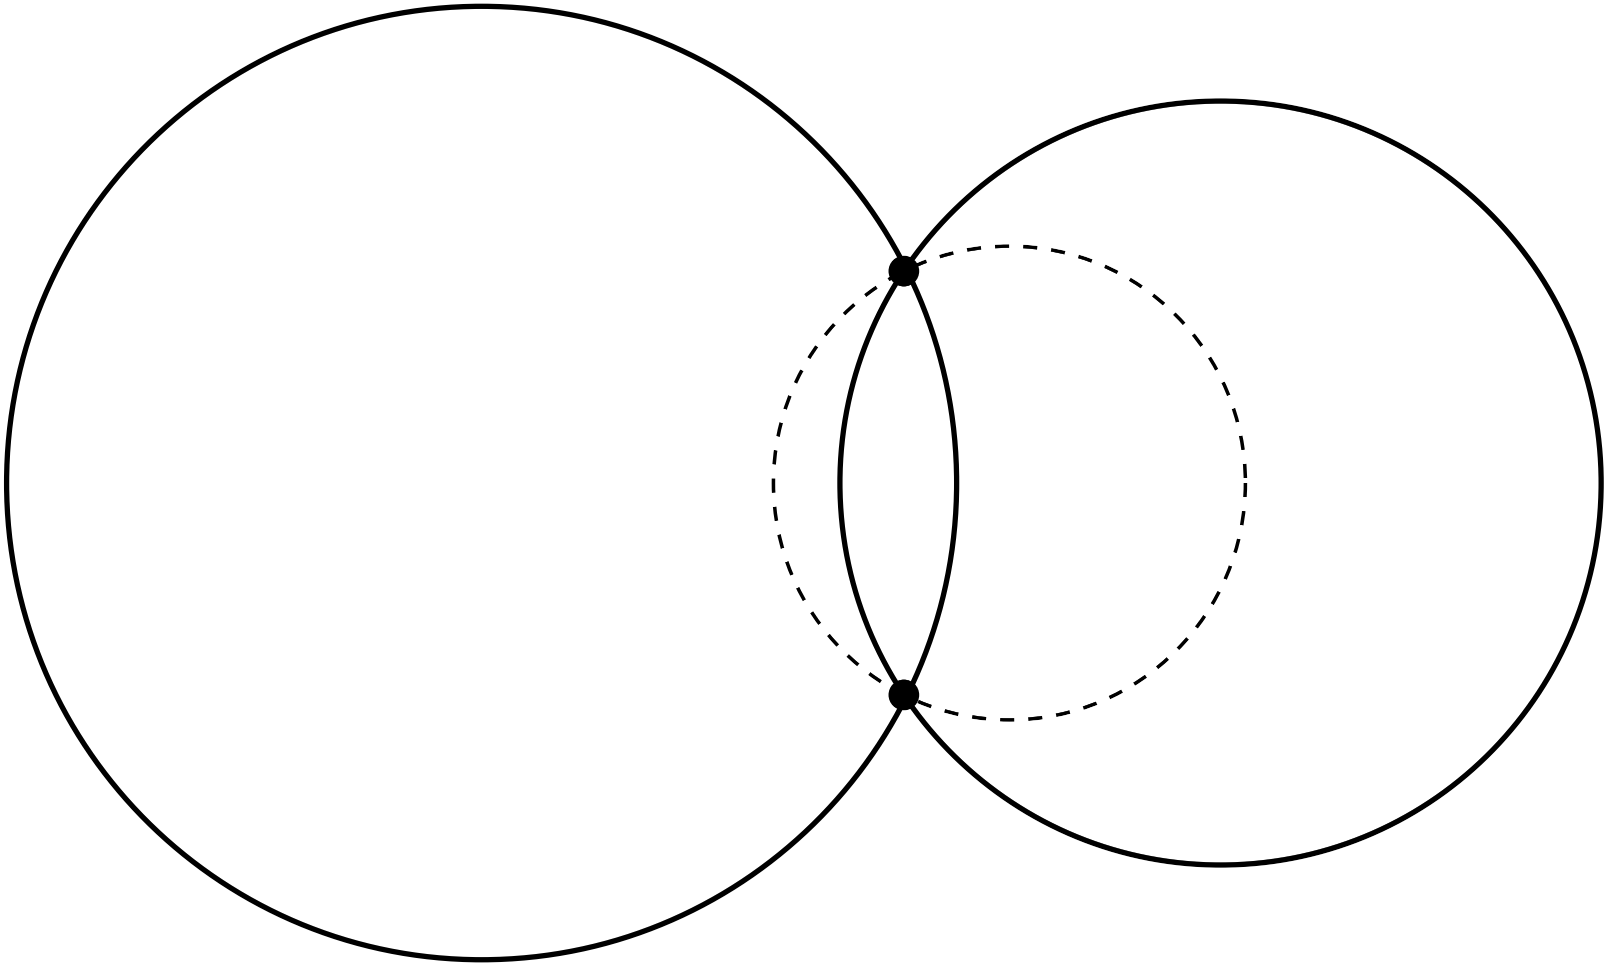

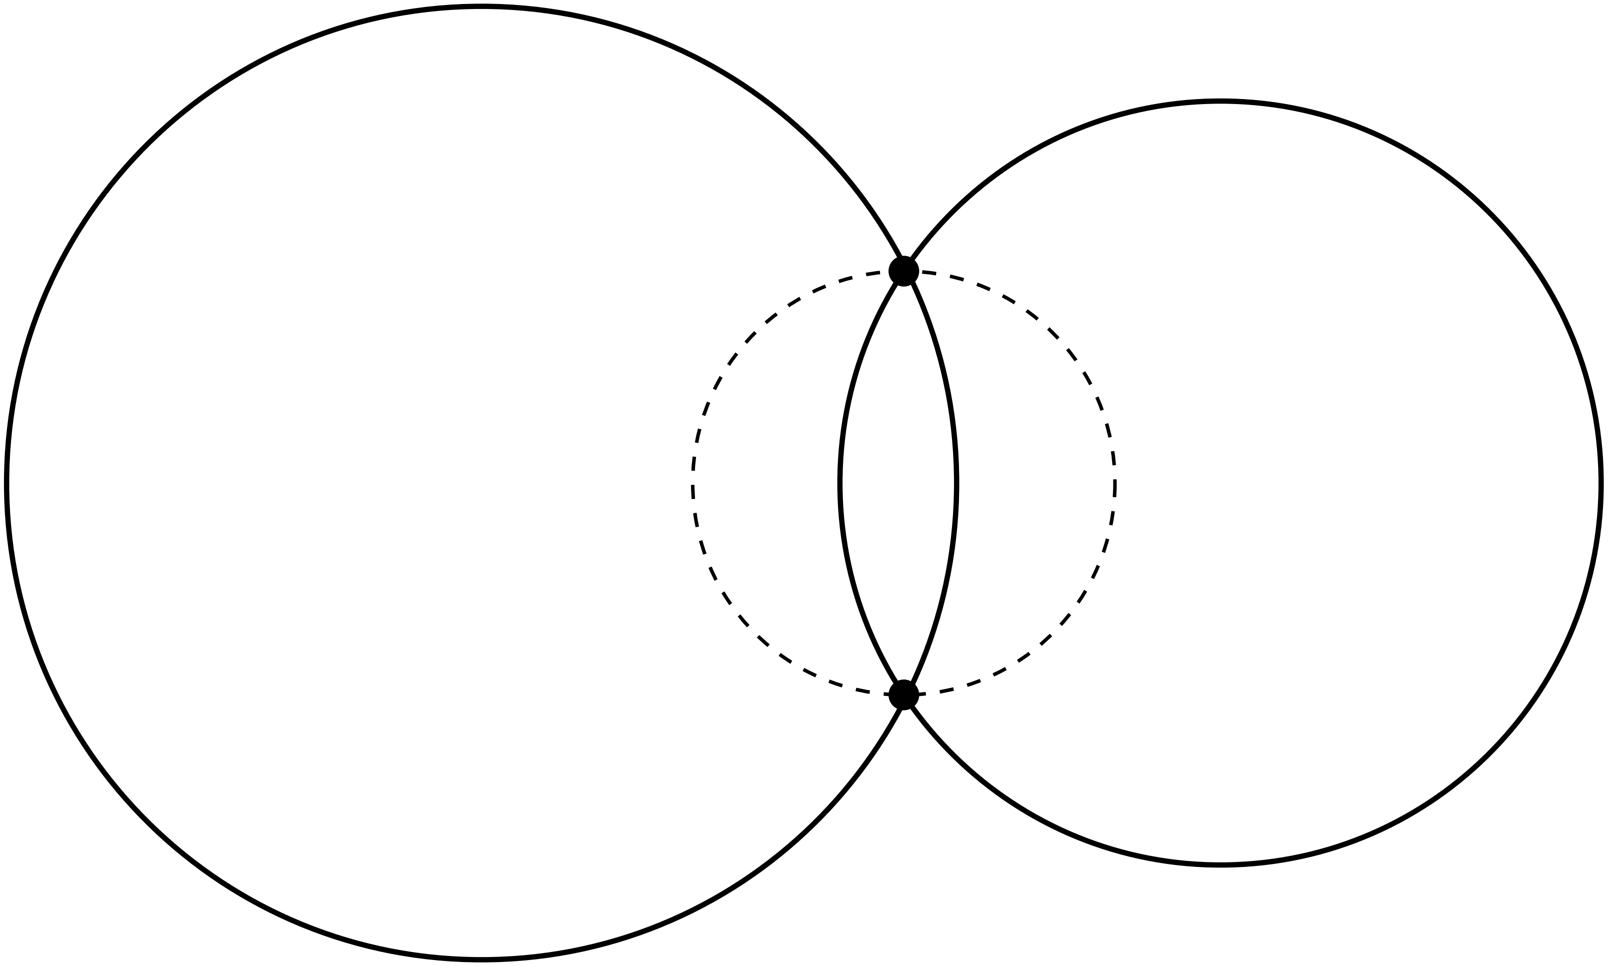

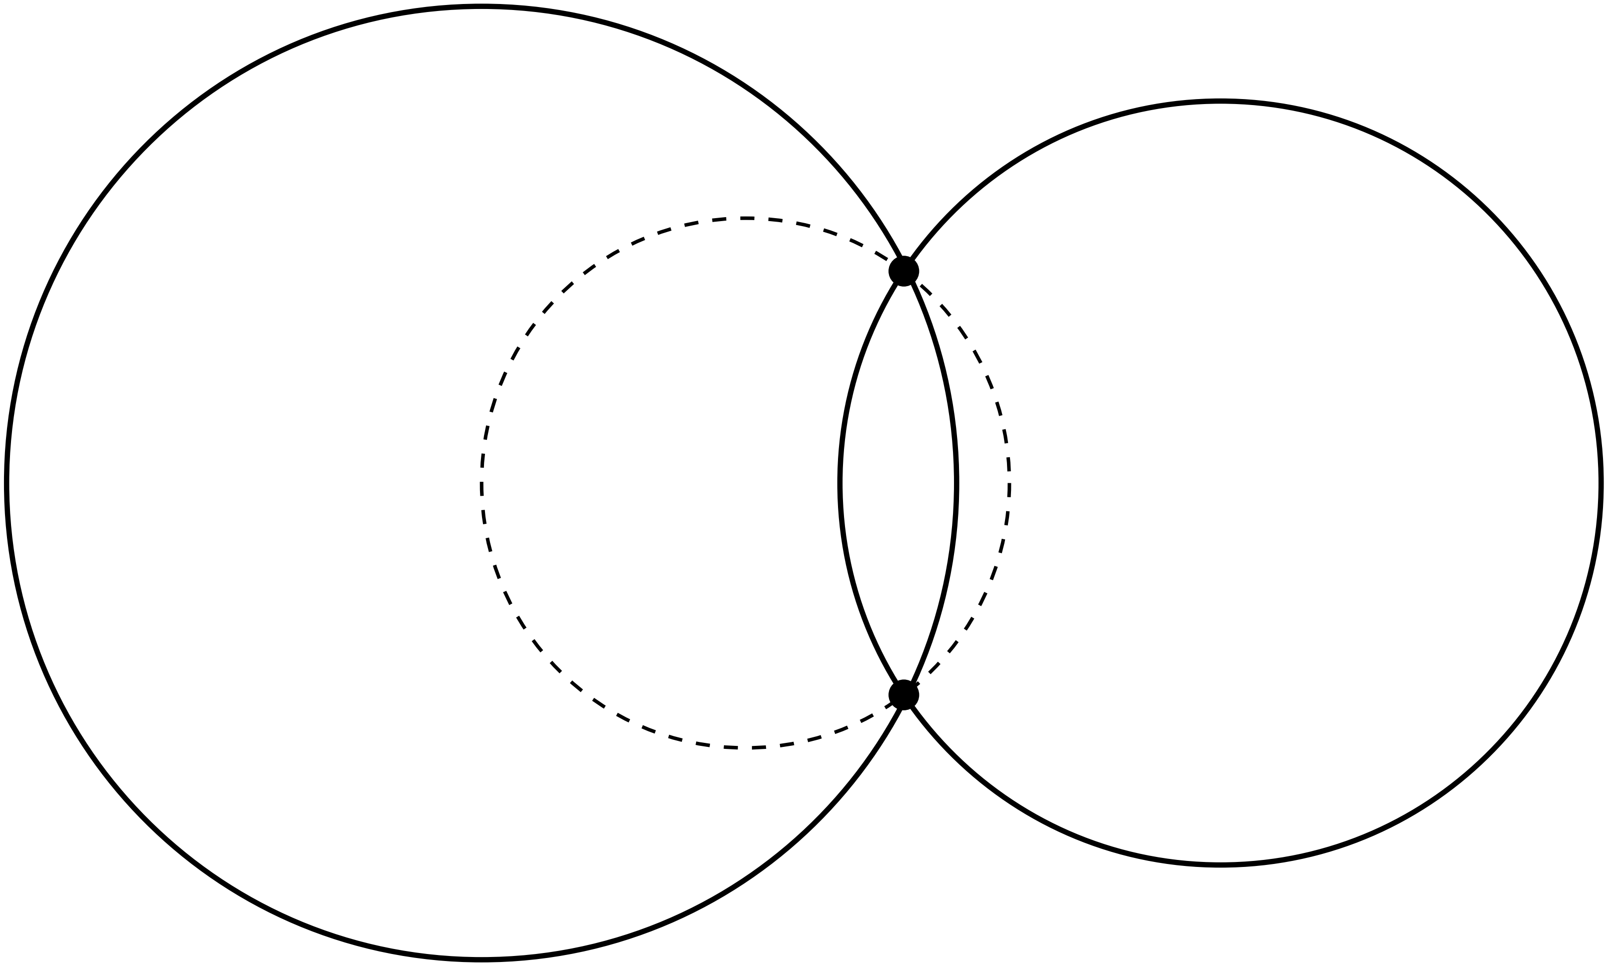

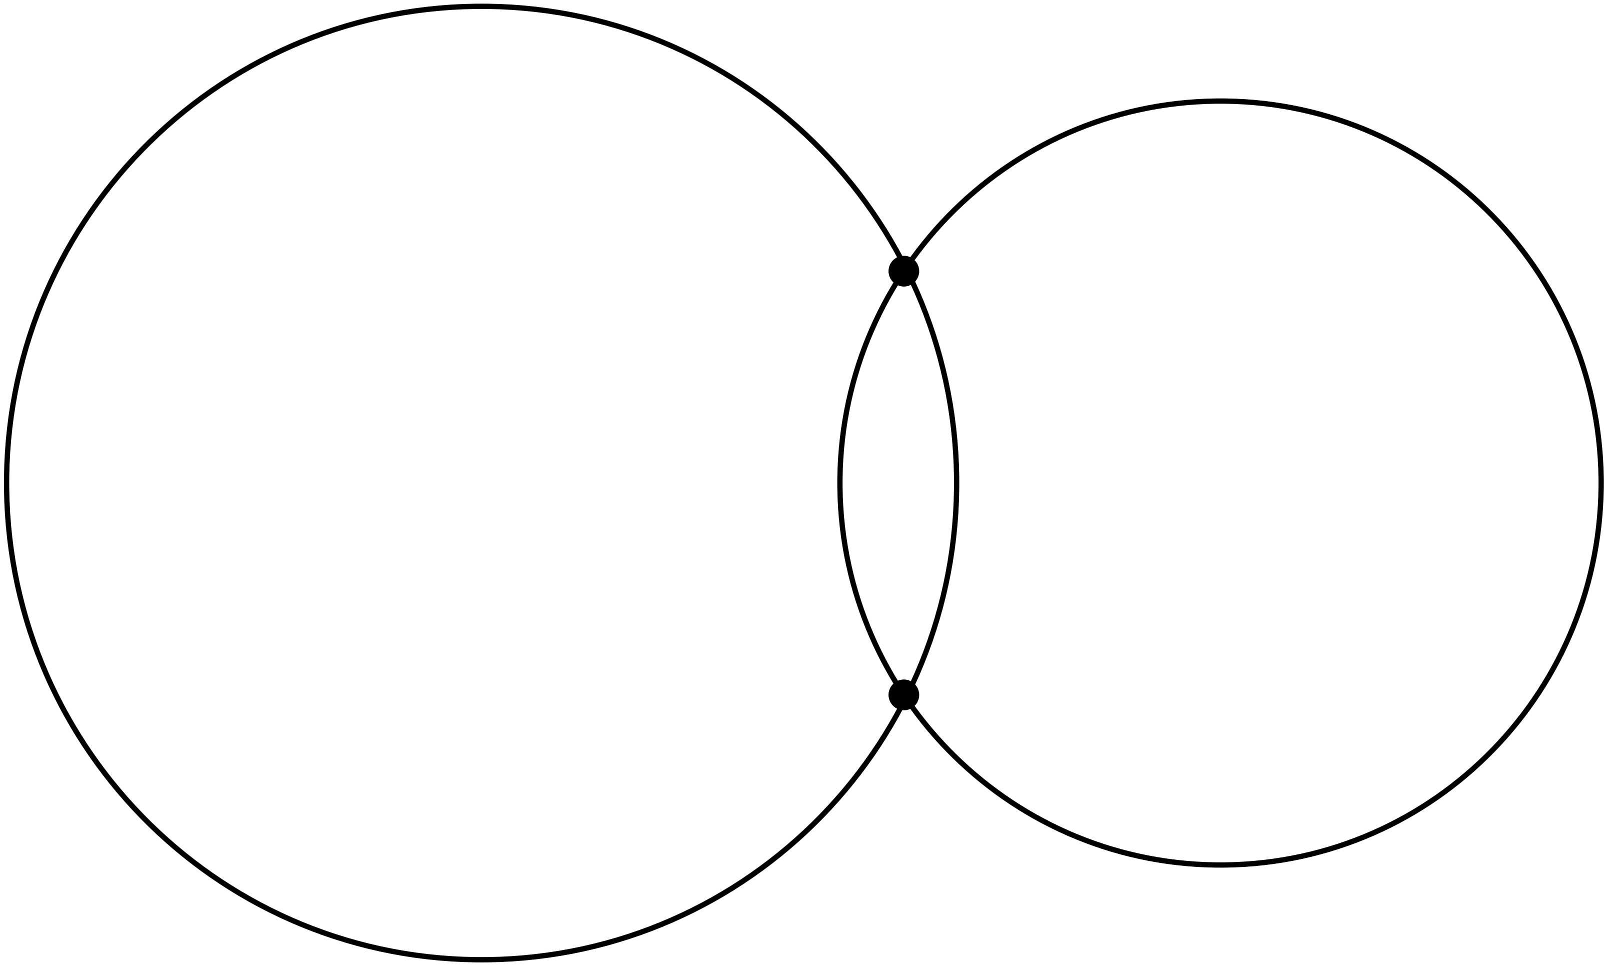

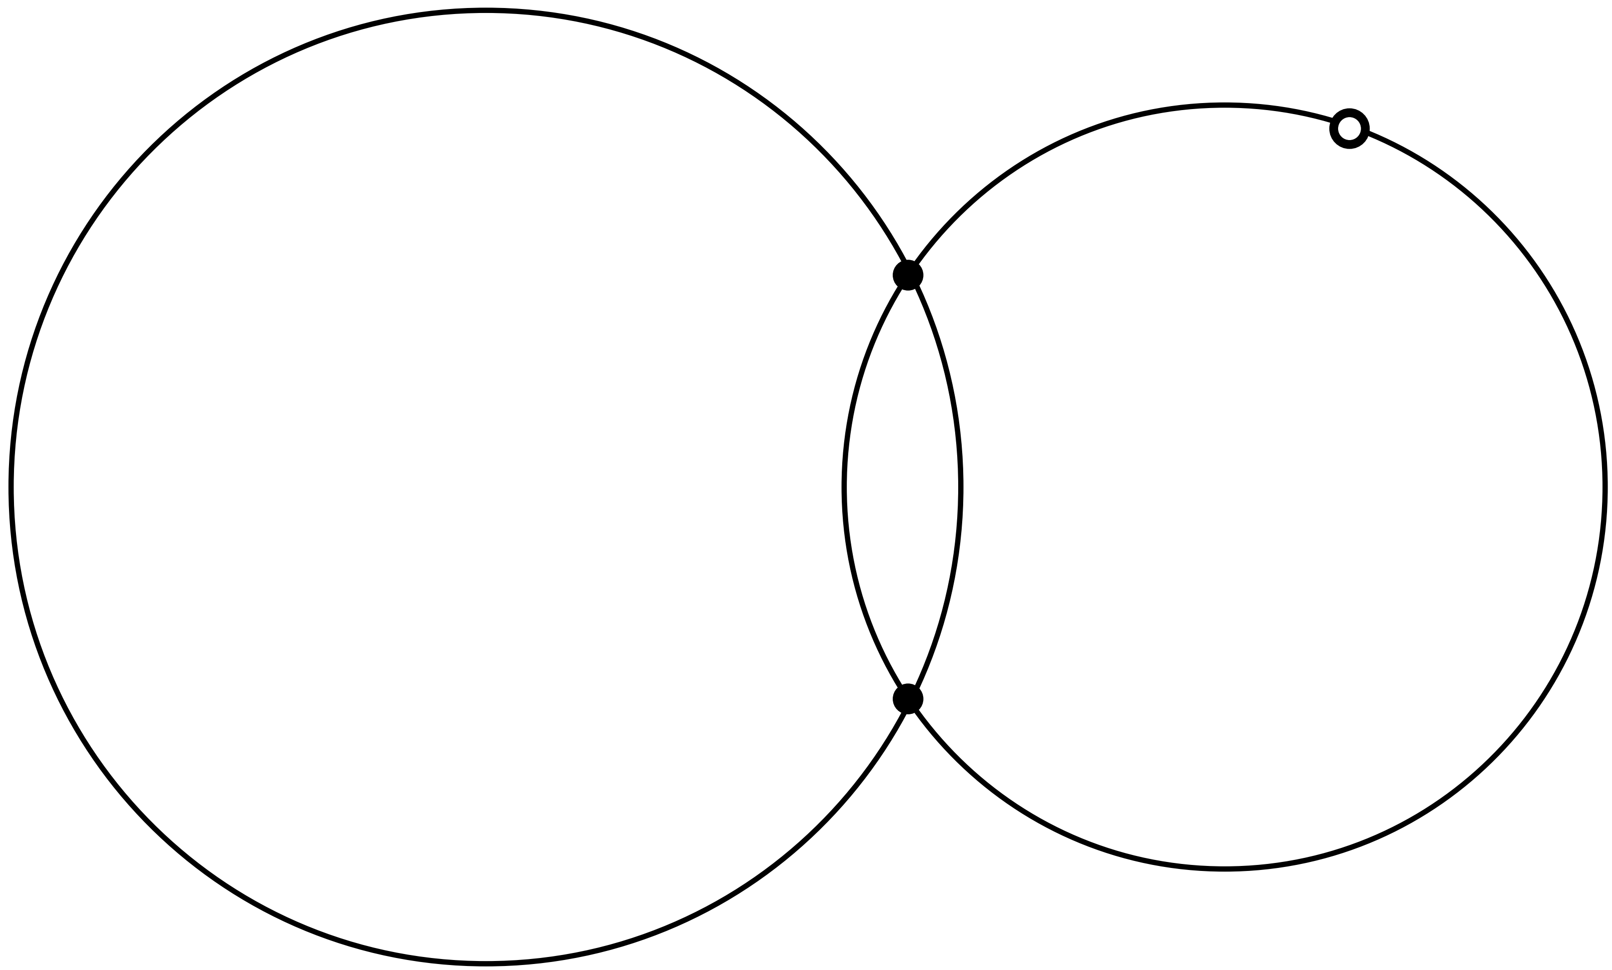

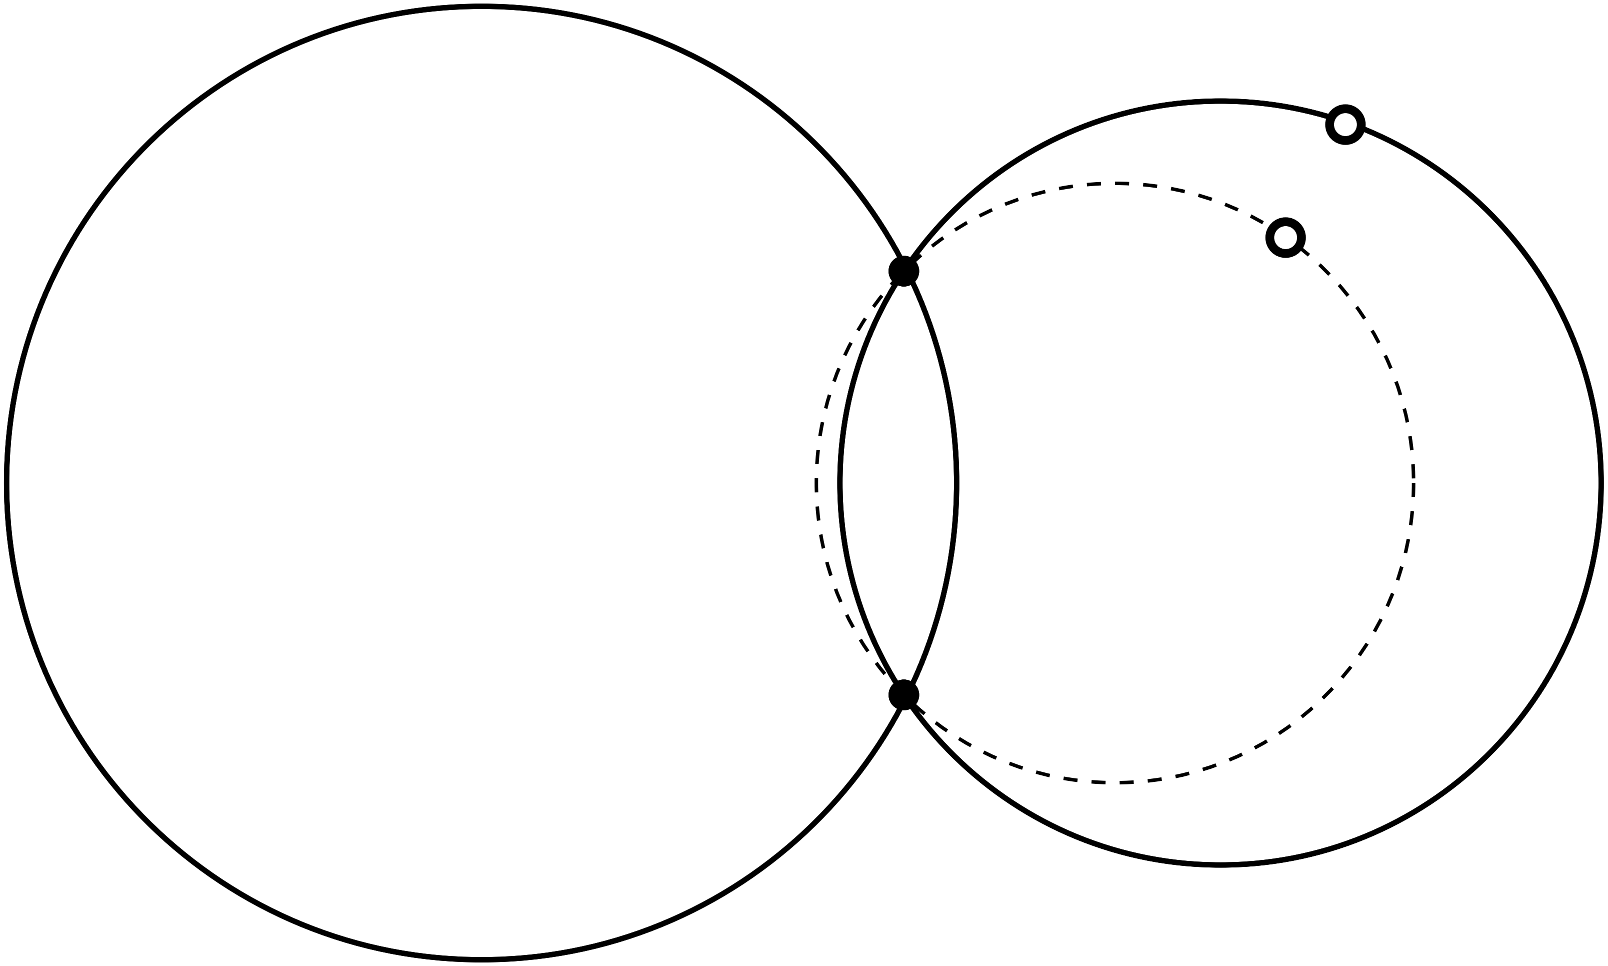

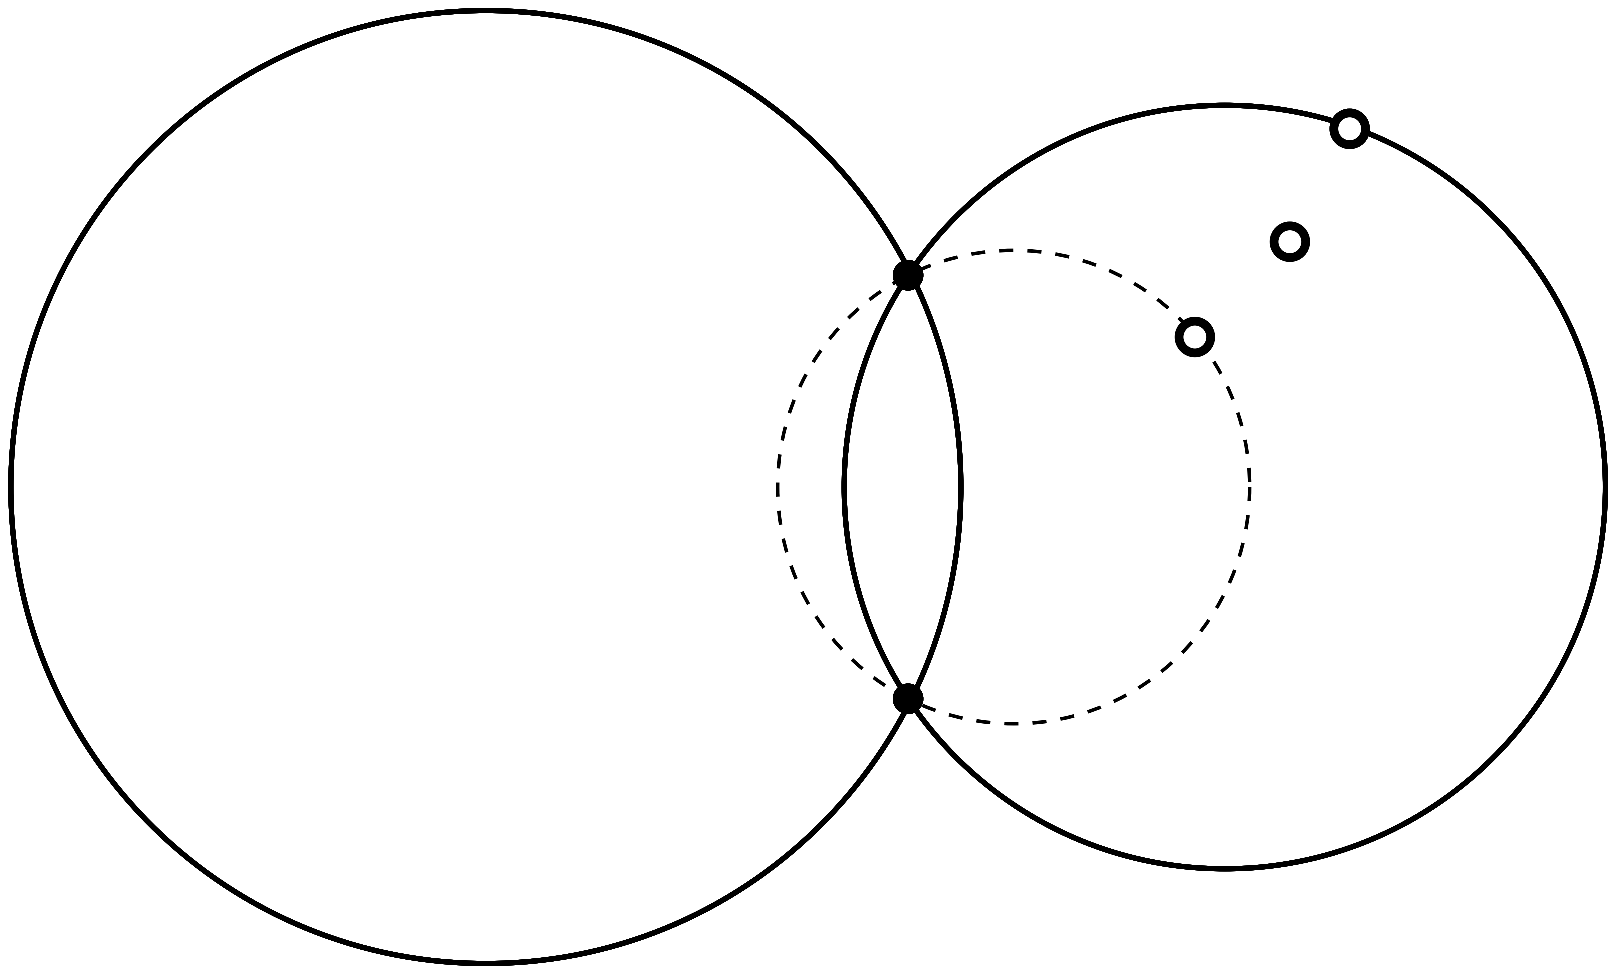

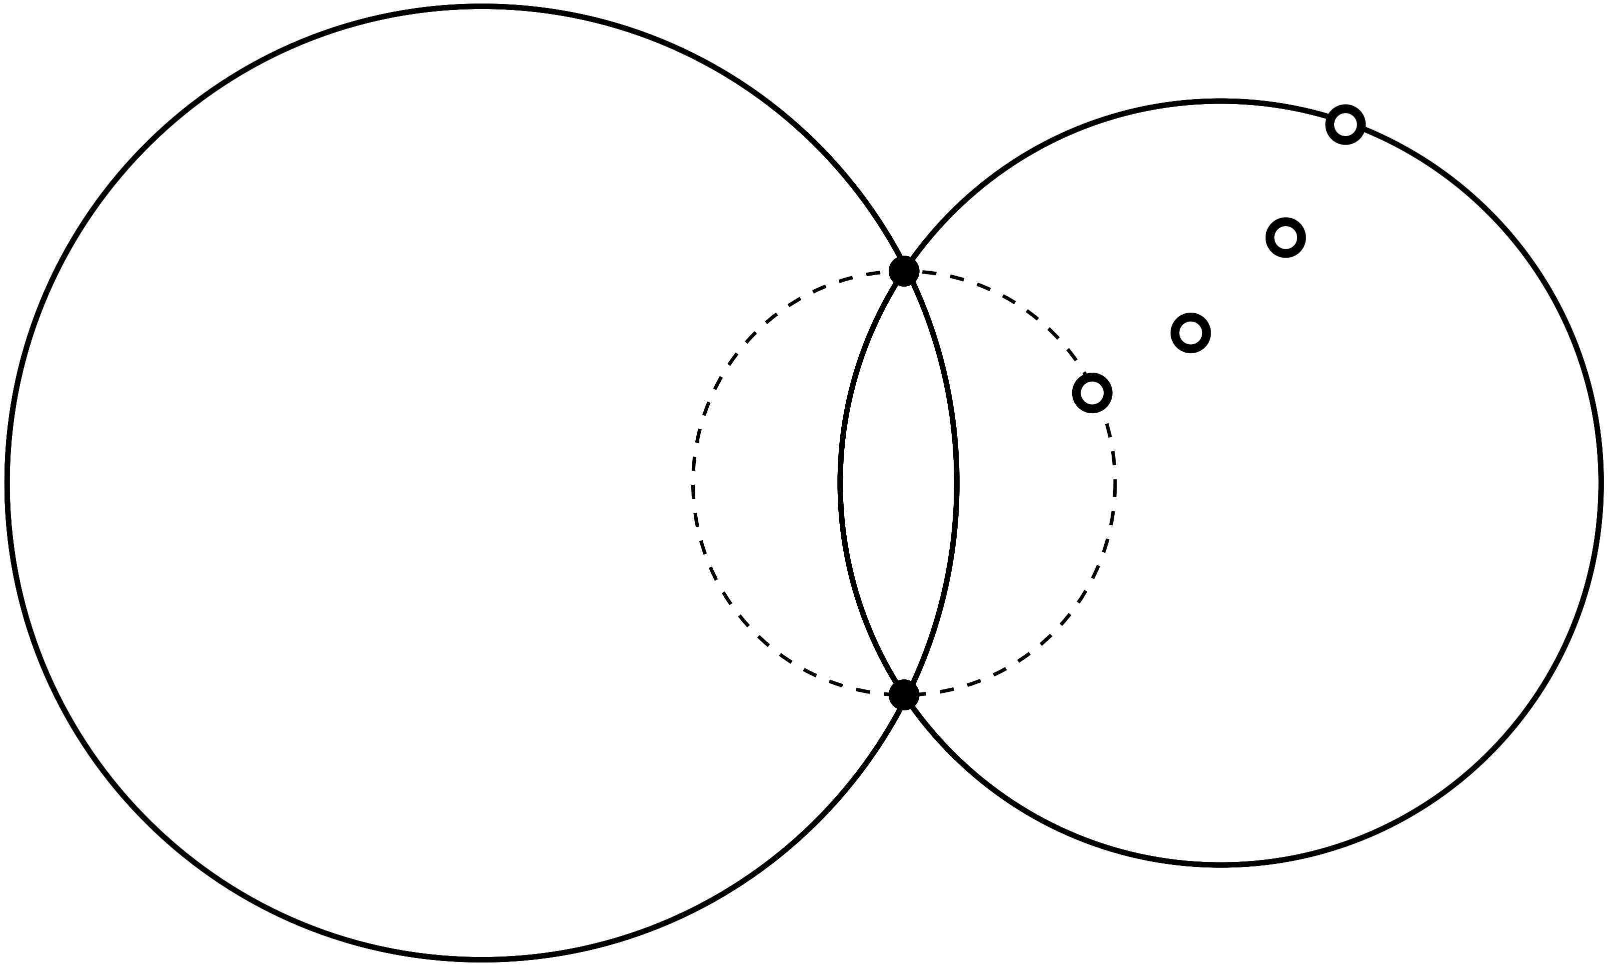

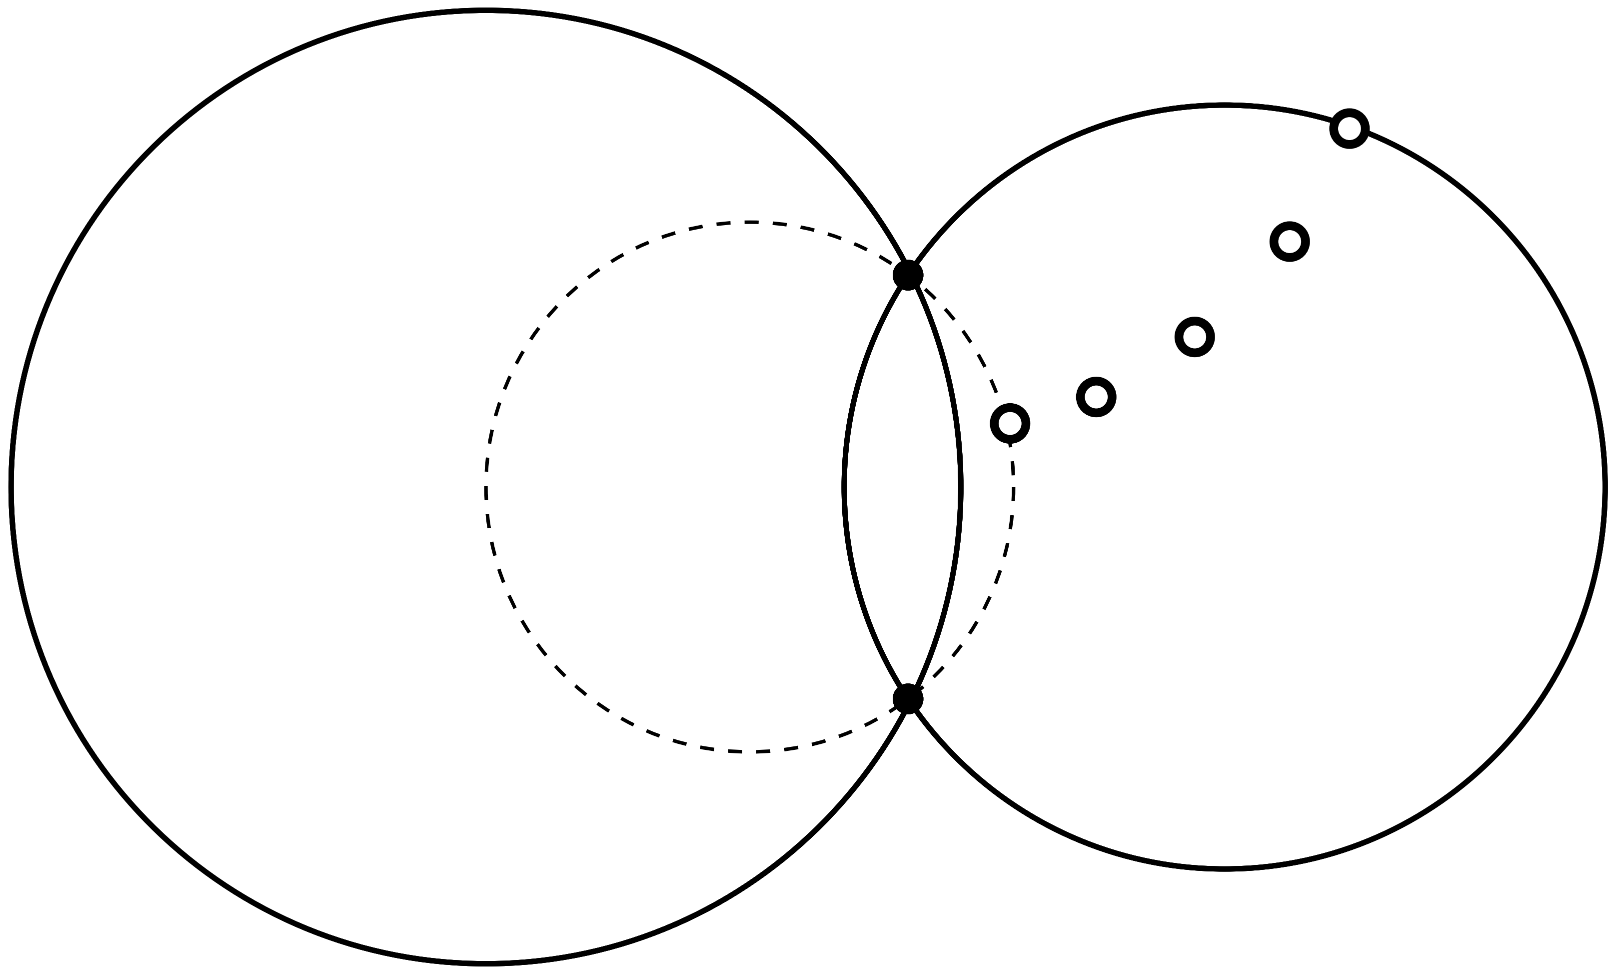

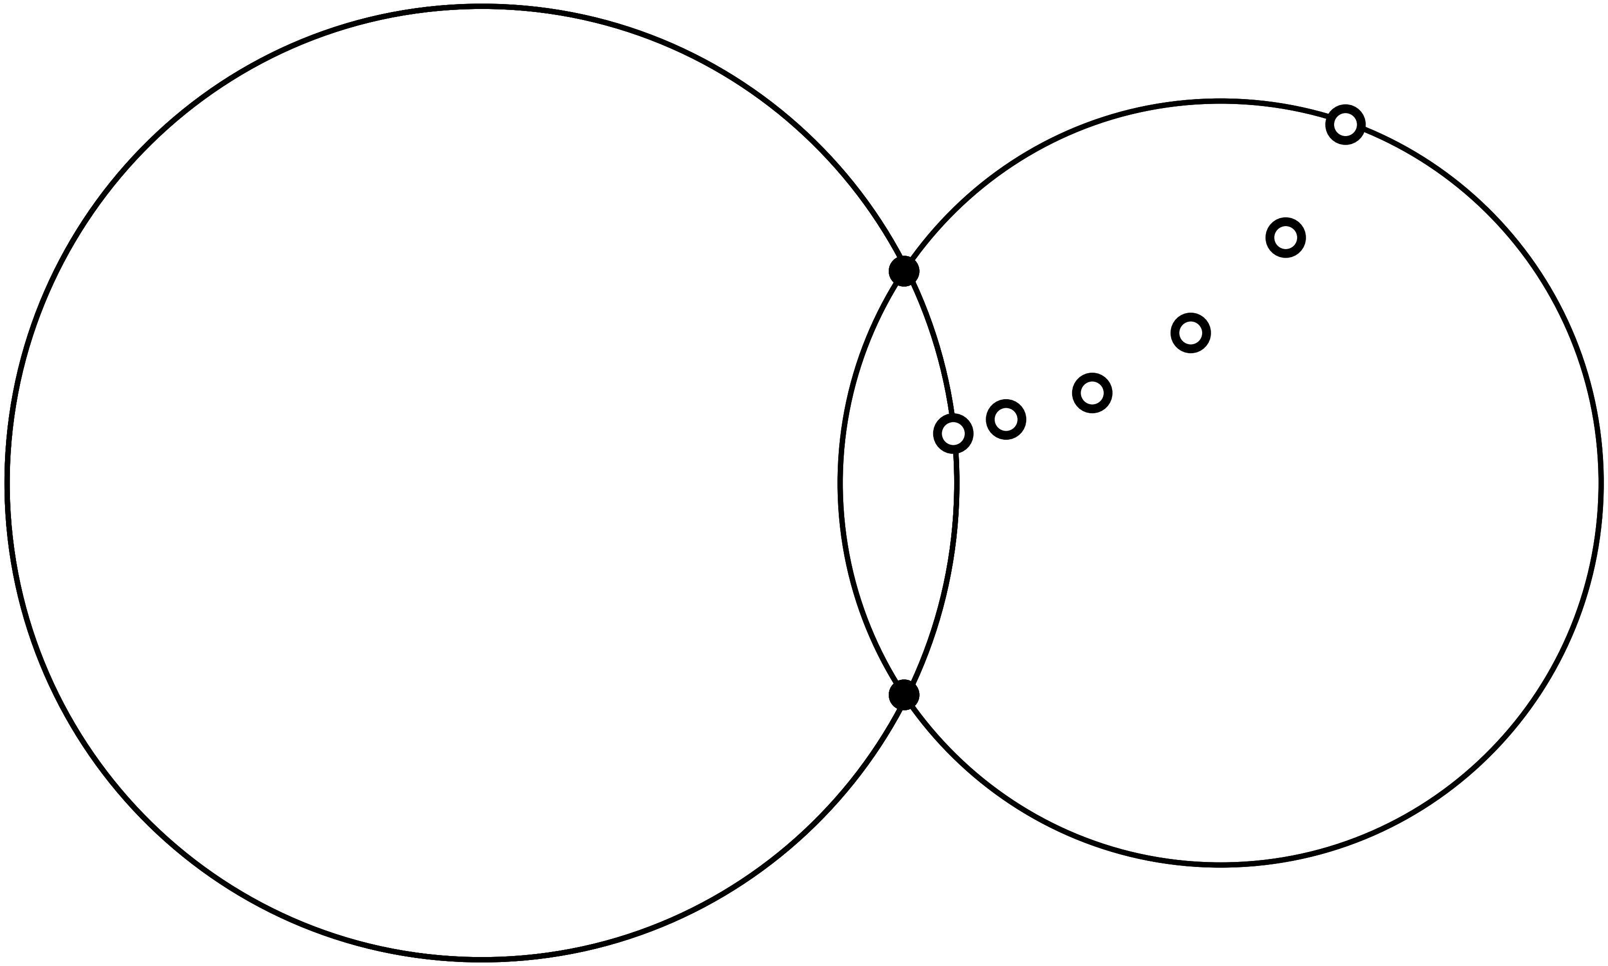

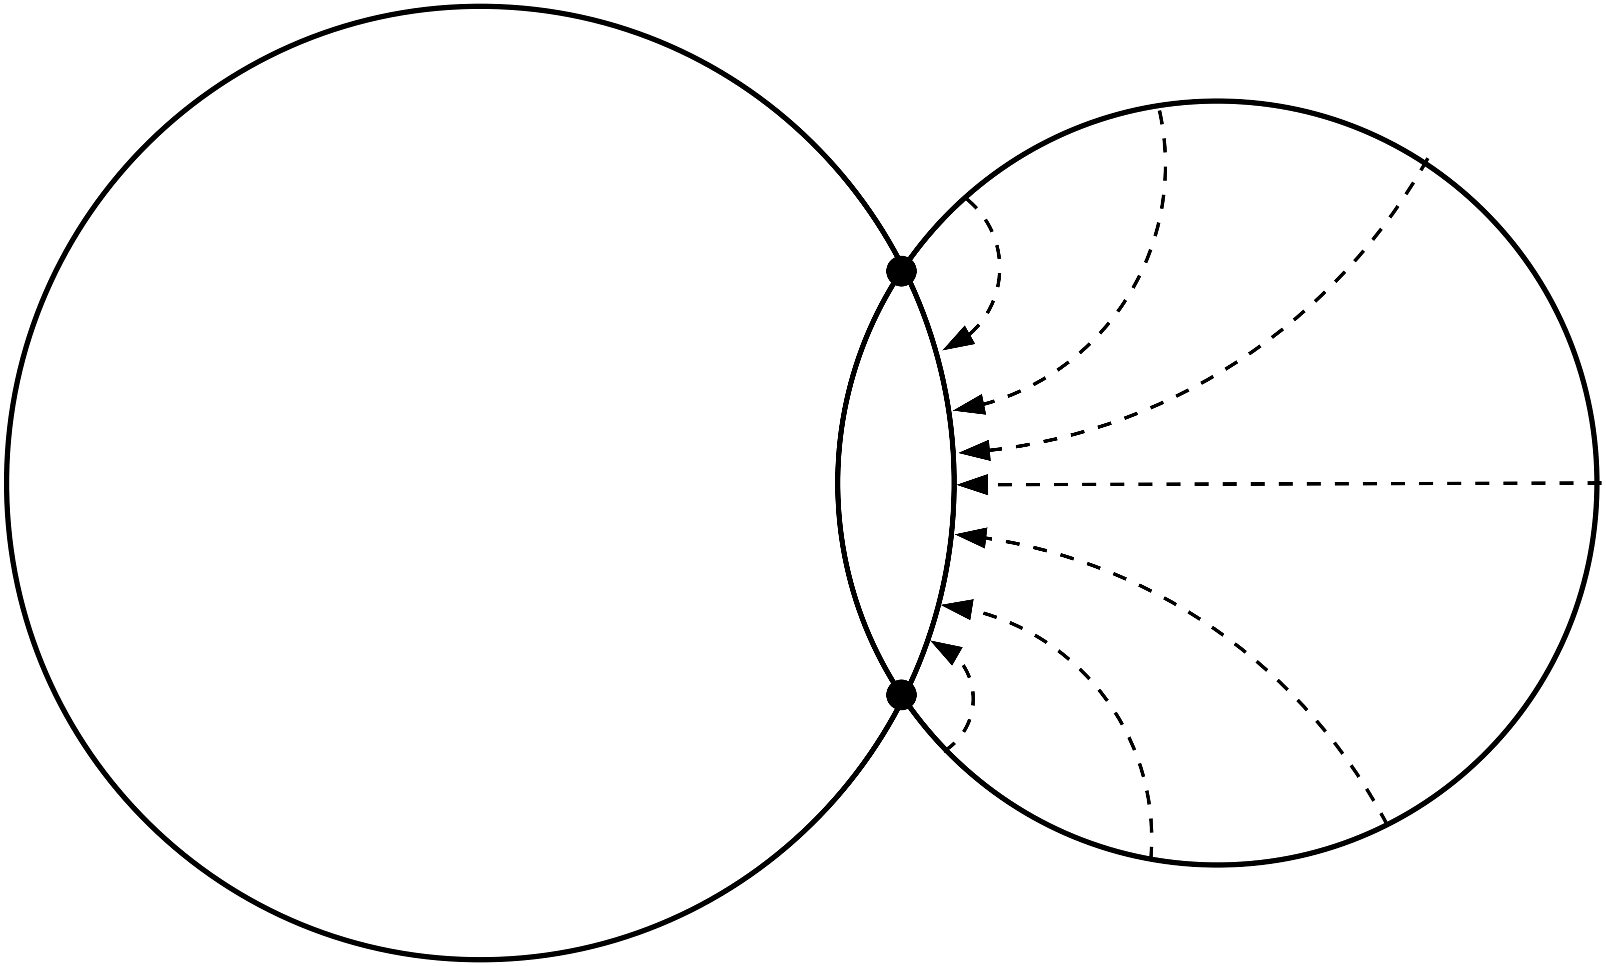

Points follow circular paths, perpendicular to boundary.

Compose along chain to map first disk to last.

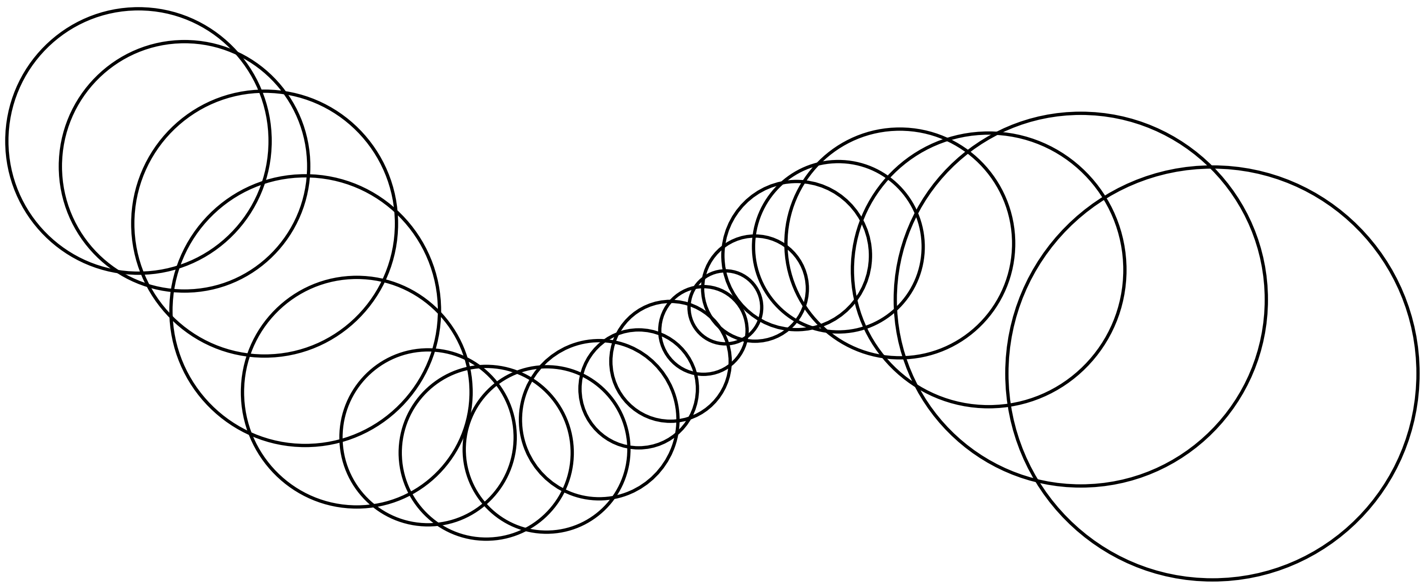

Composition is Möbius (maps form group).

Compose along chain to map first disk to last.

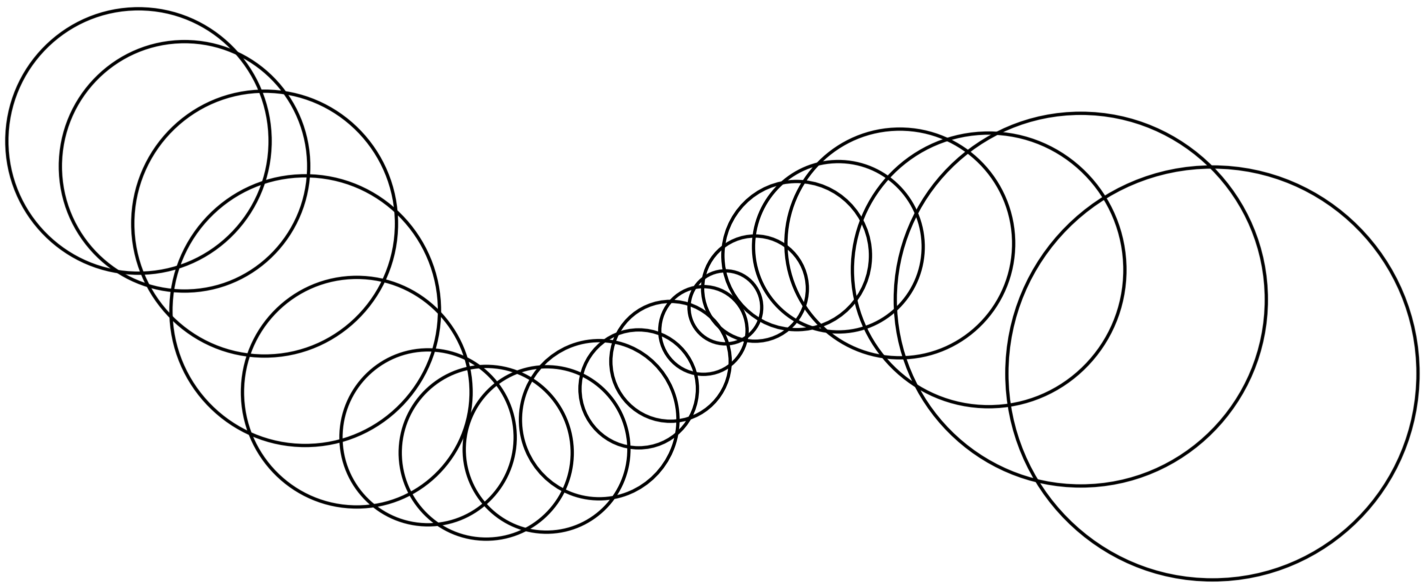

Composition is Möbius (maps form group).

Take limit to associate Möbius map to path of circles.

How does this give a map from polygon P to a circle C?

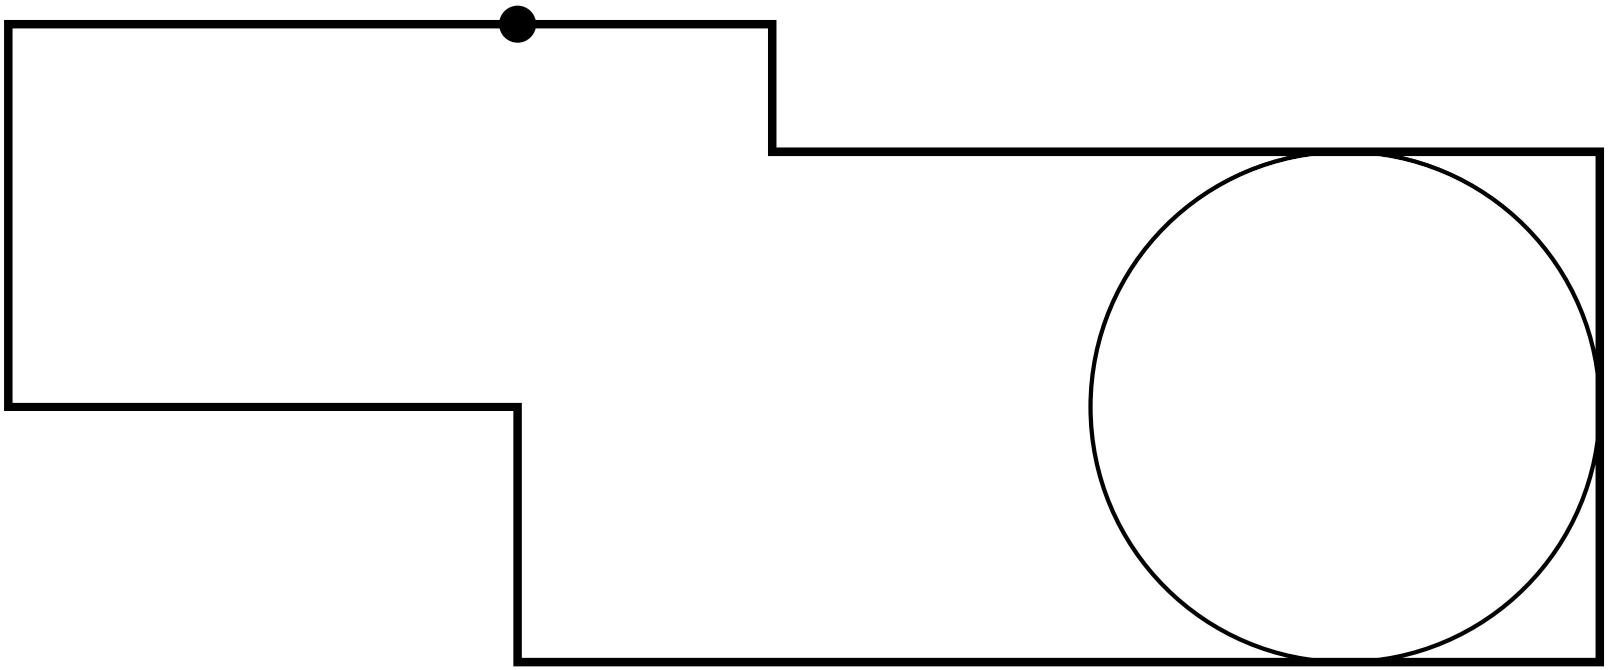

- Fix a “root” MA disk D. C is boundary of D.

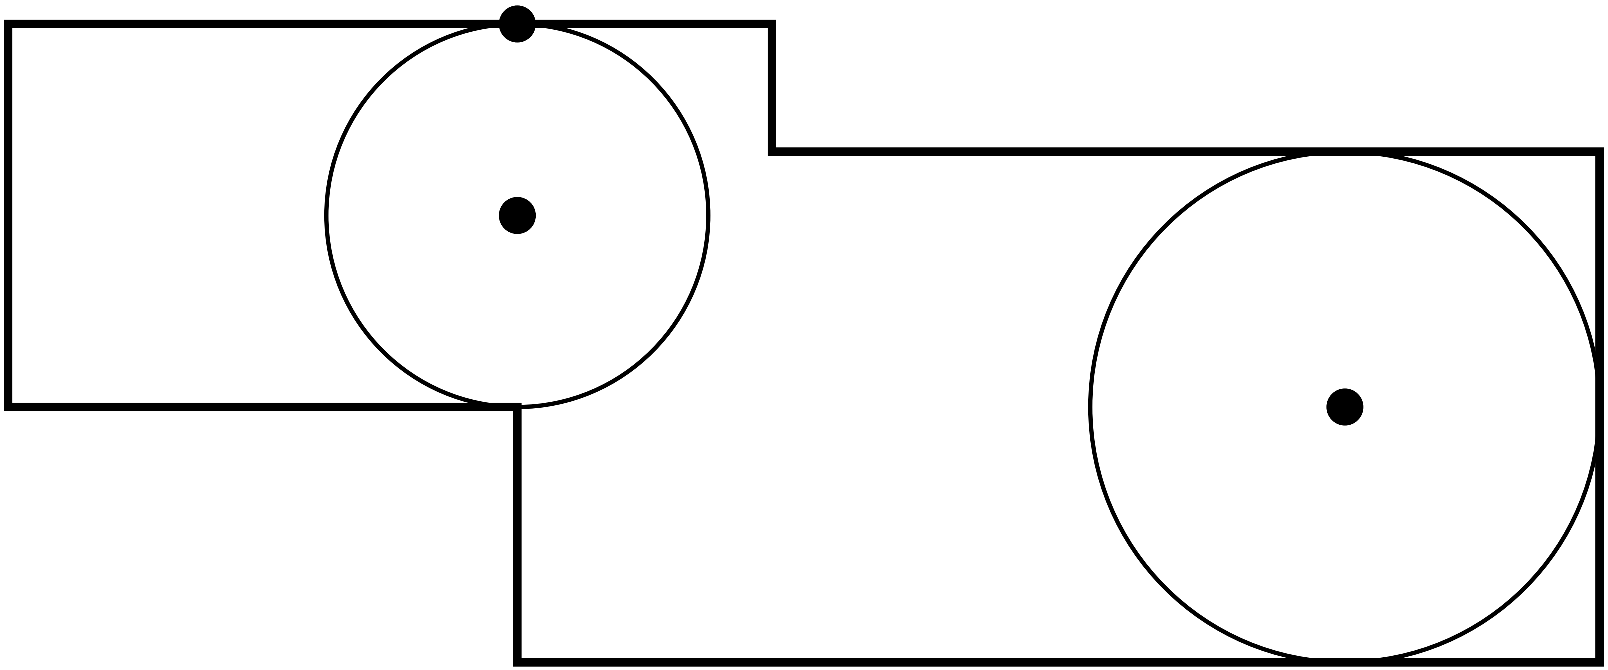

- For any $z \in P$, take MA disk D_z touching z.

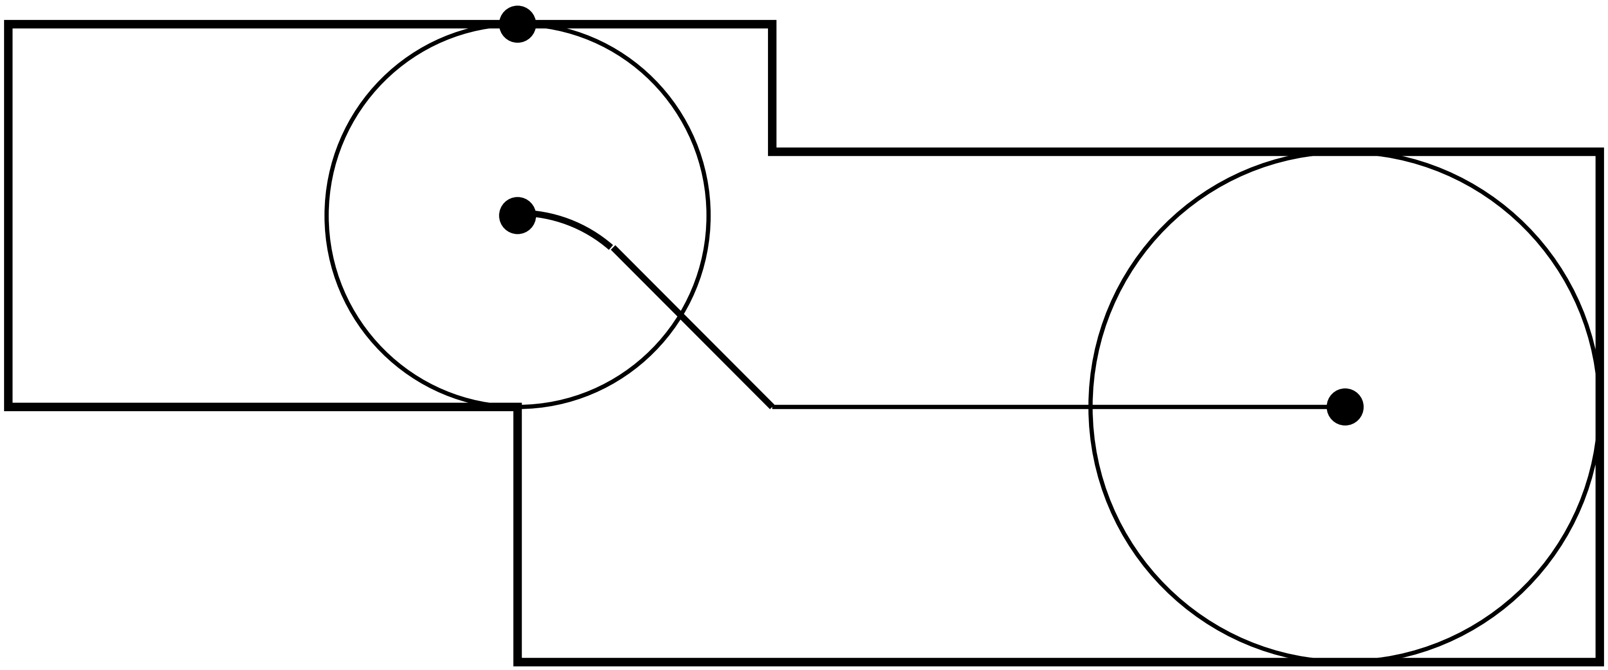

- Connect D_z to D on MA. Path map sends z into C.

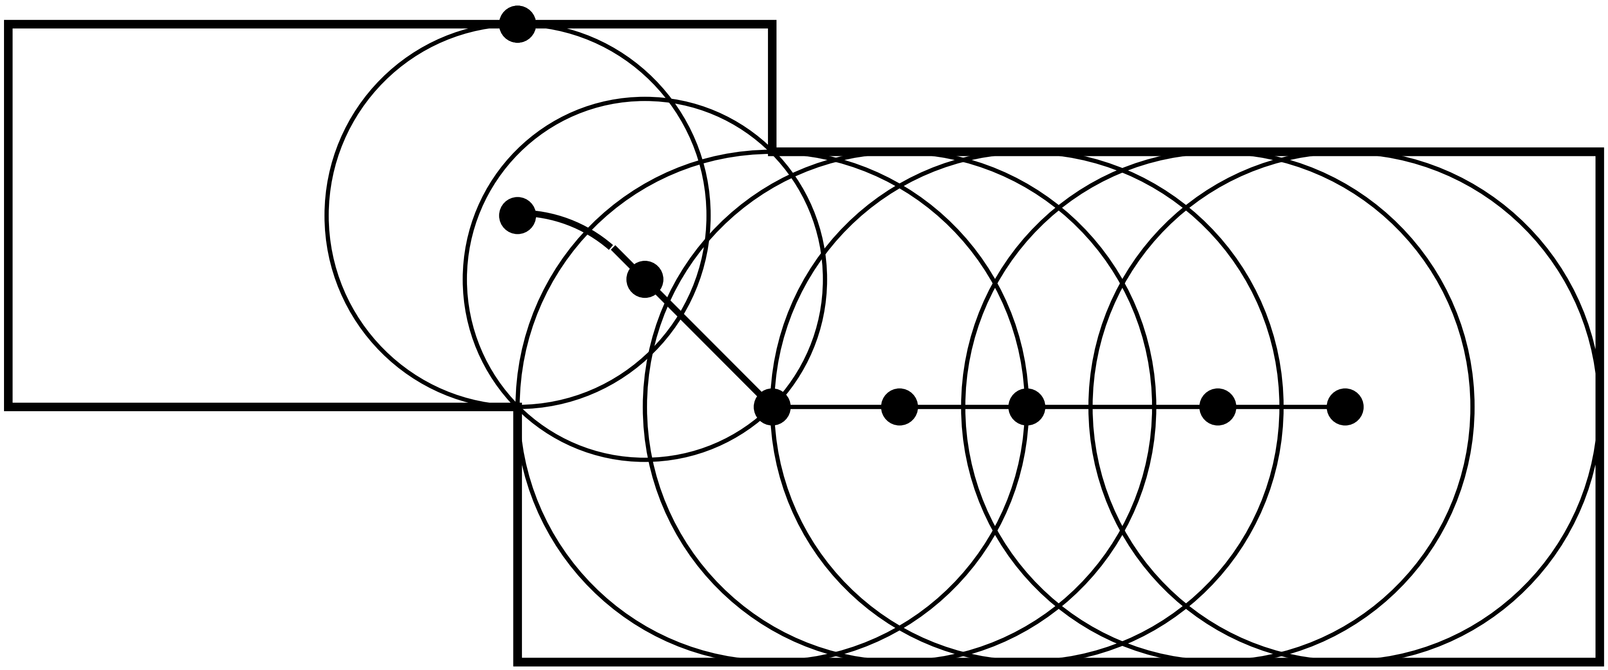

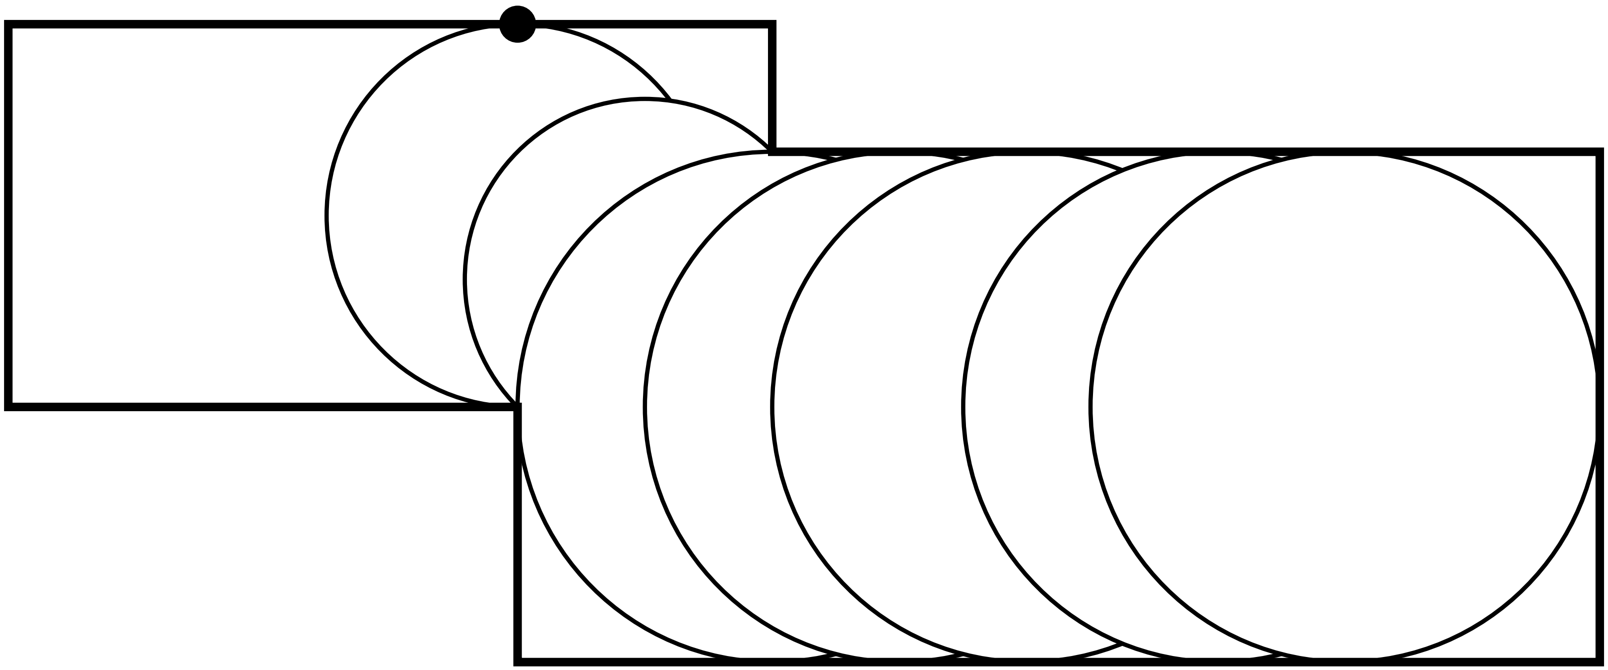

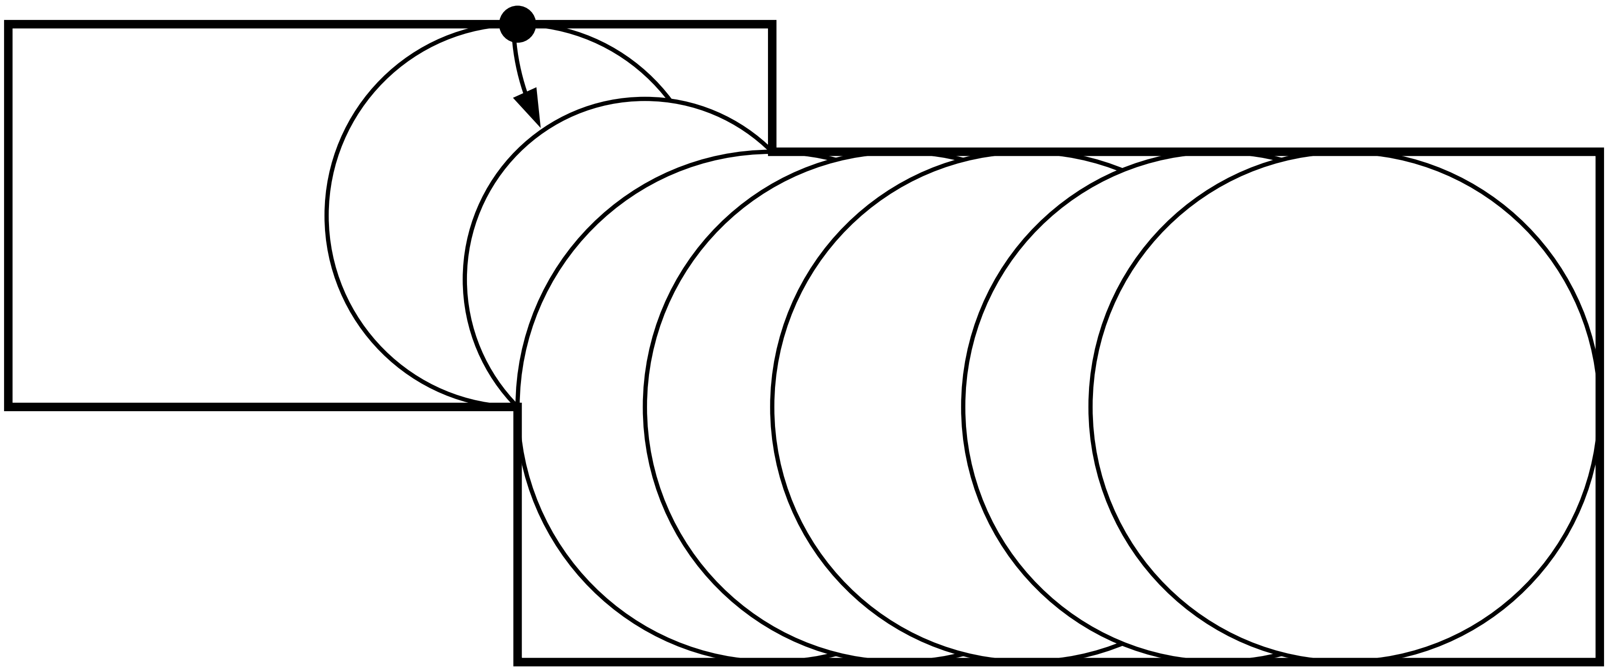

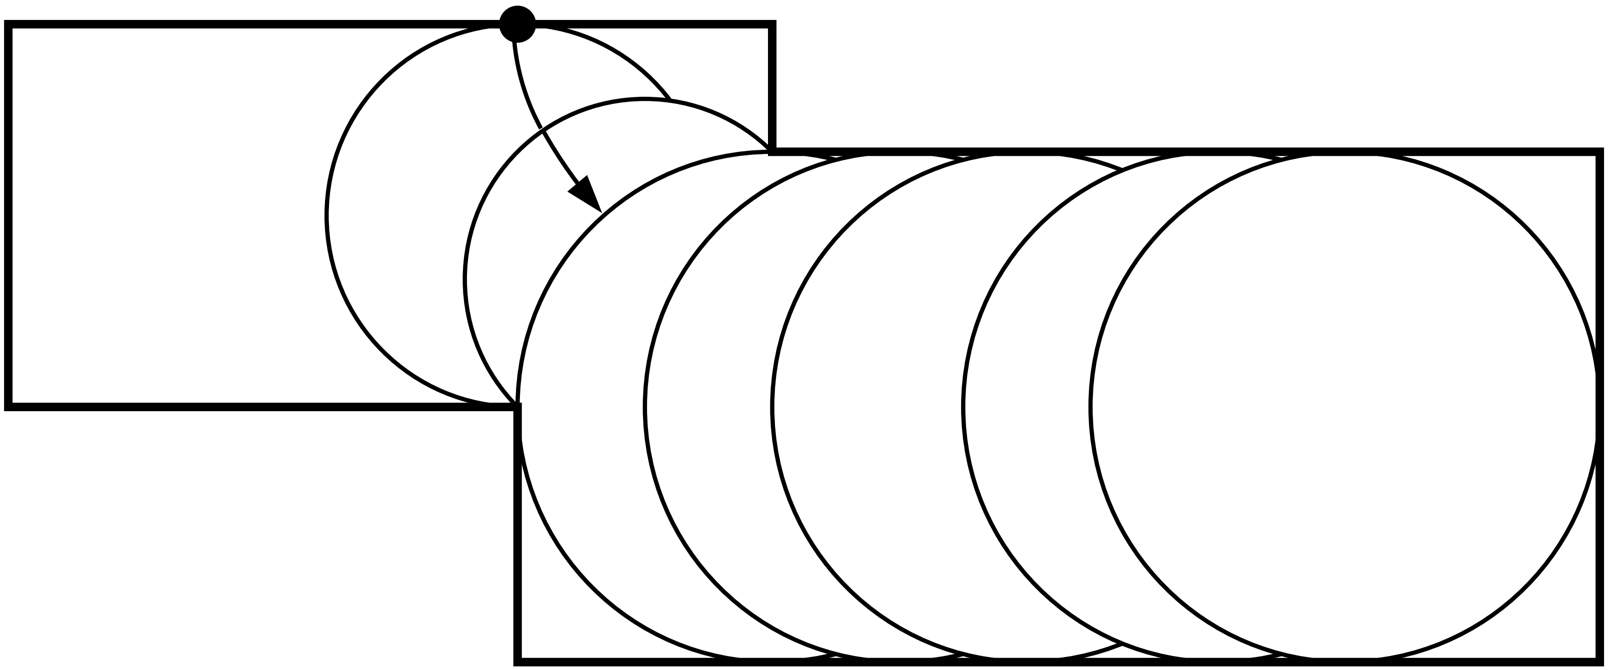

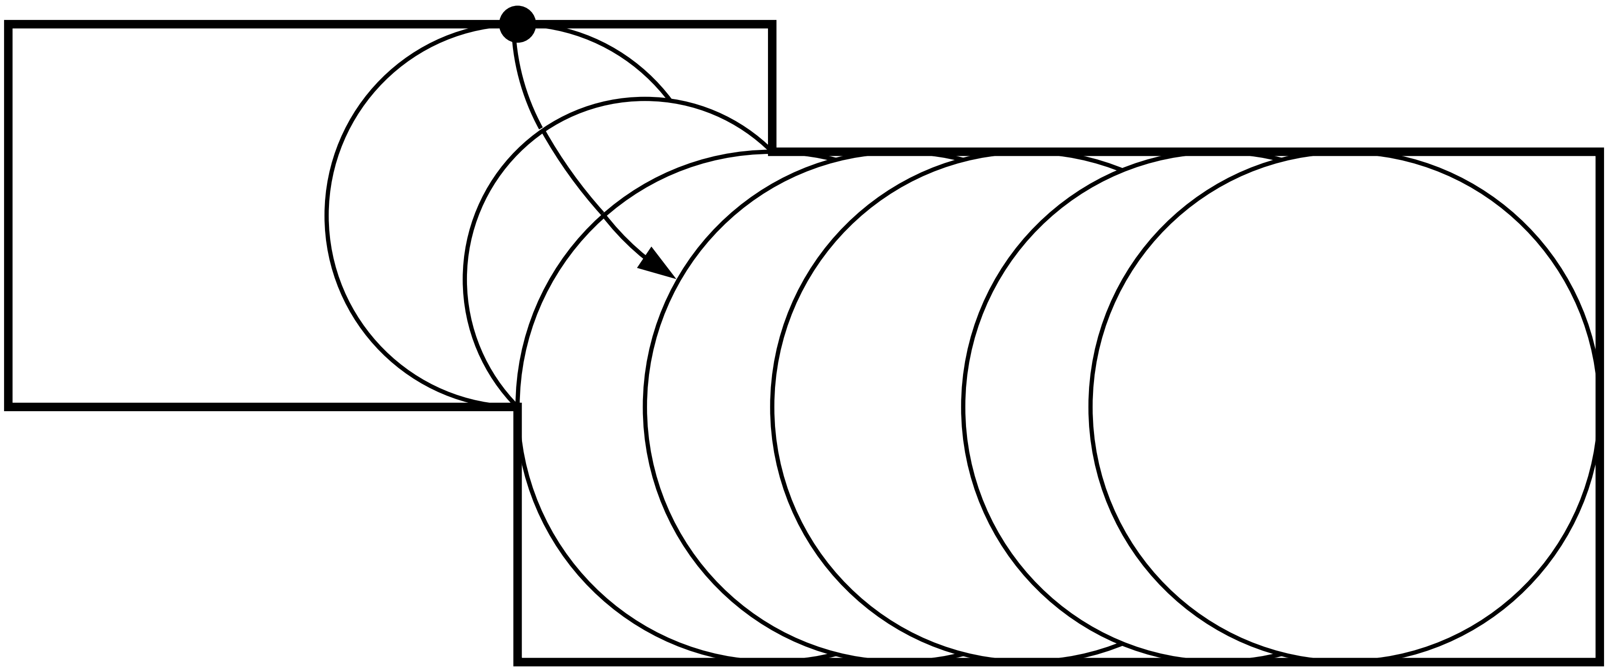

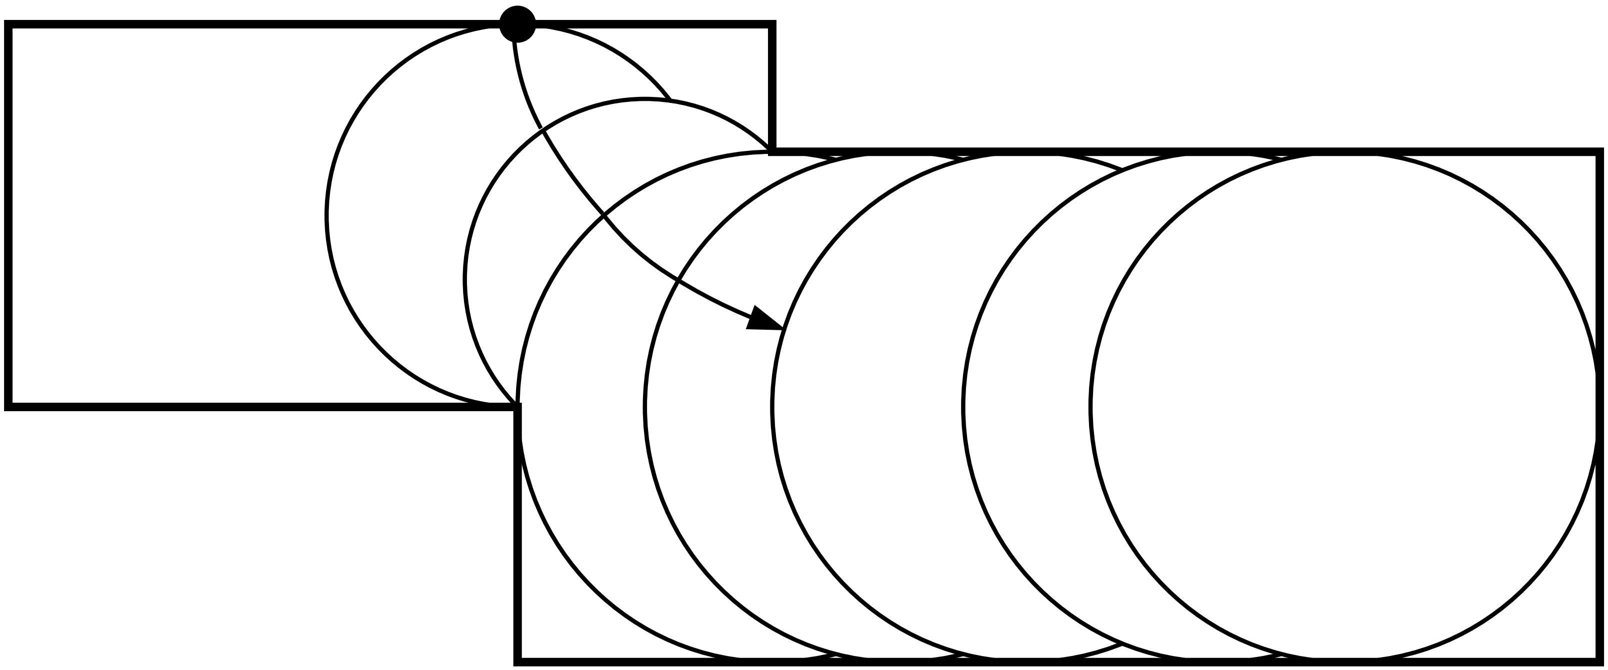

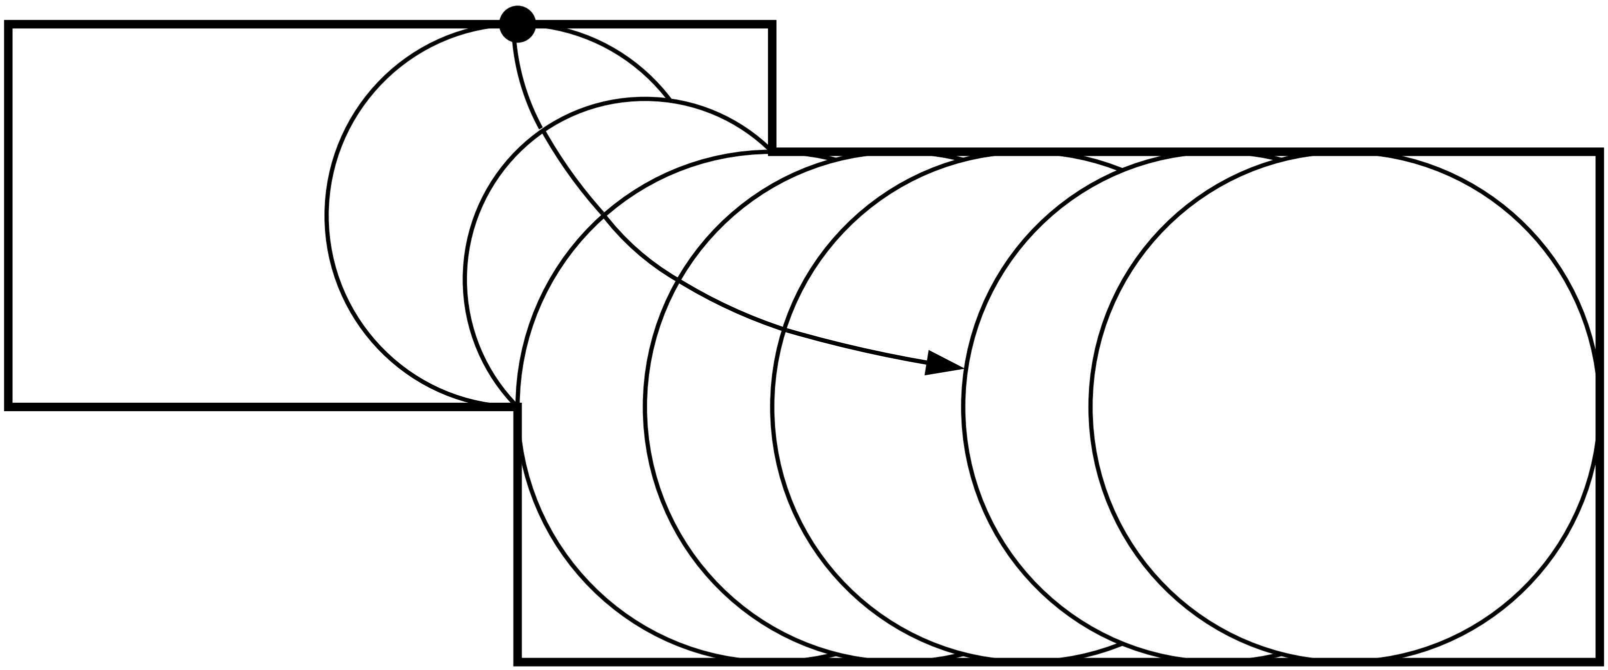

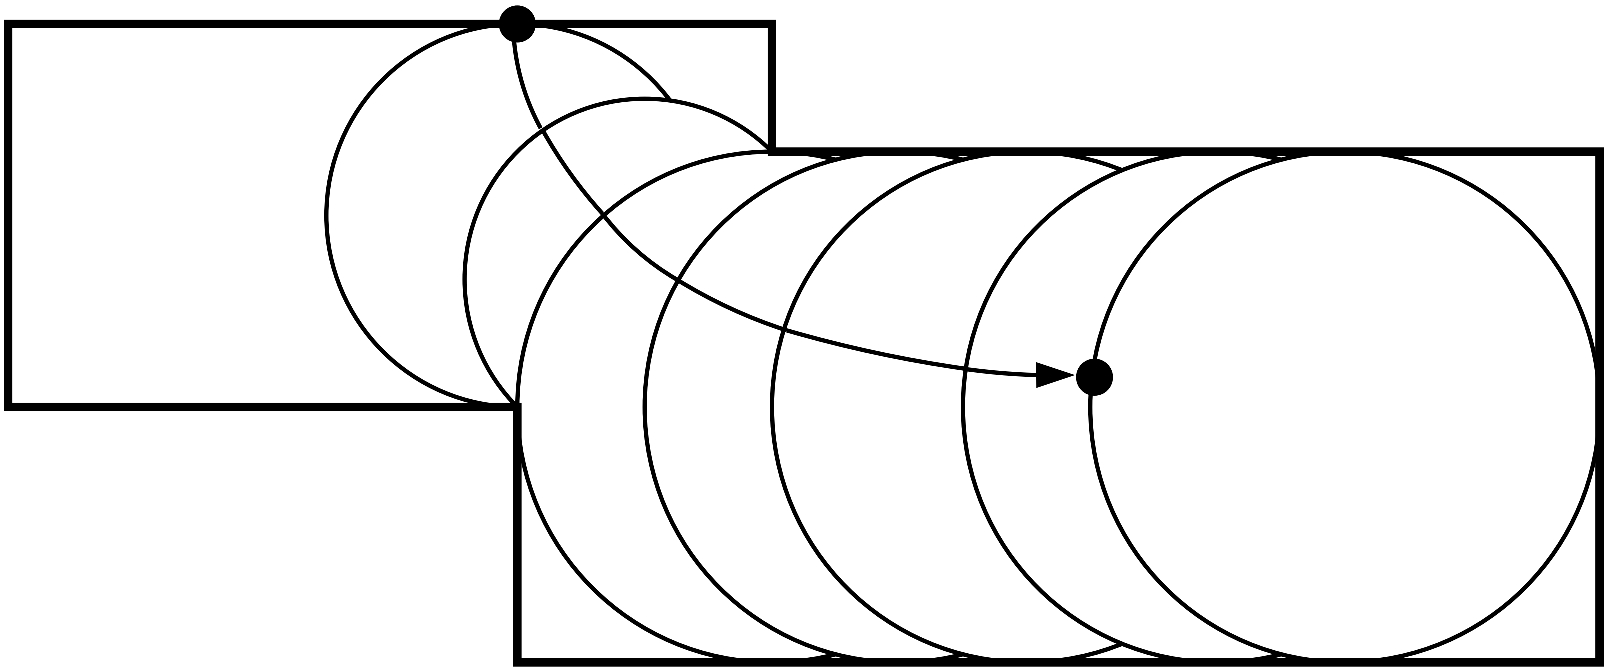

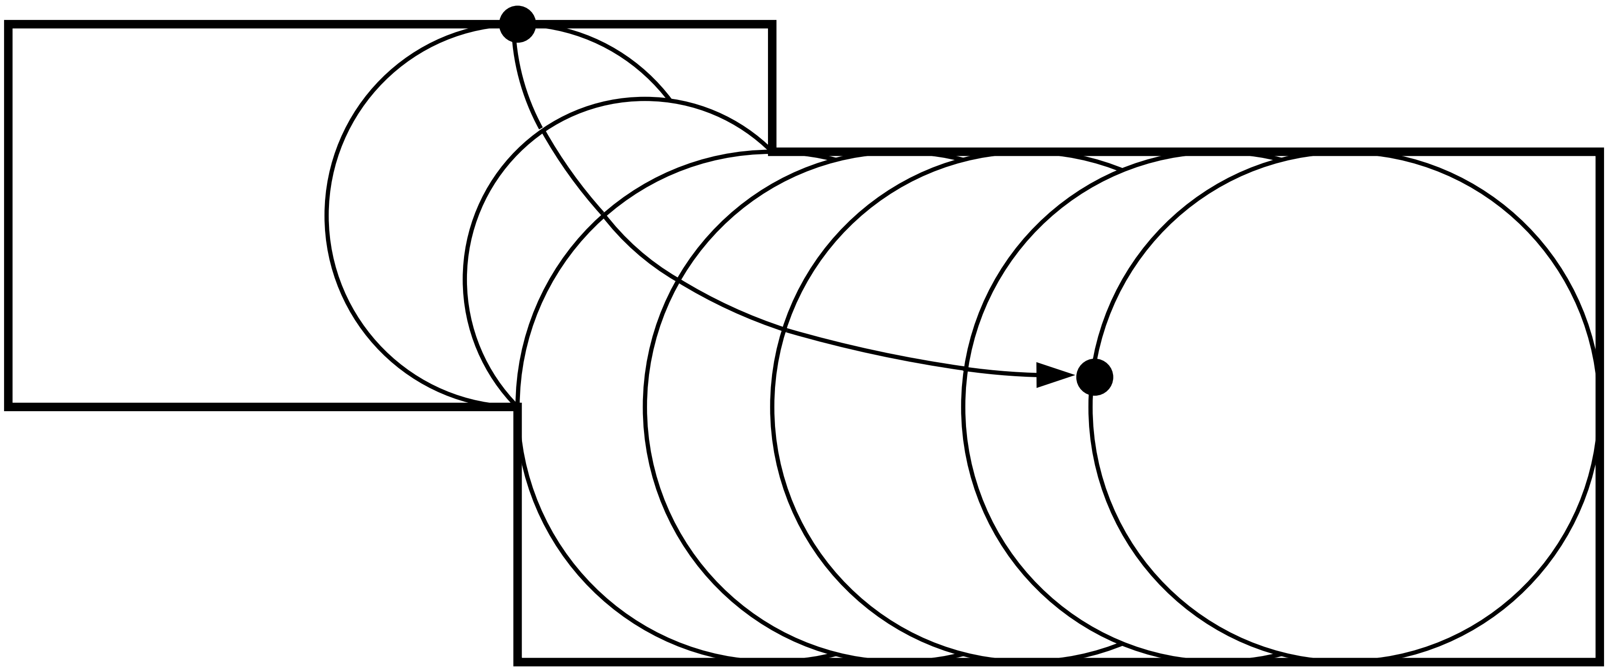

We discretize only to draw picture.

Limiting map has **formula** in terms of medial axis.

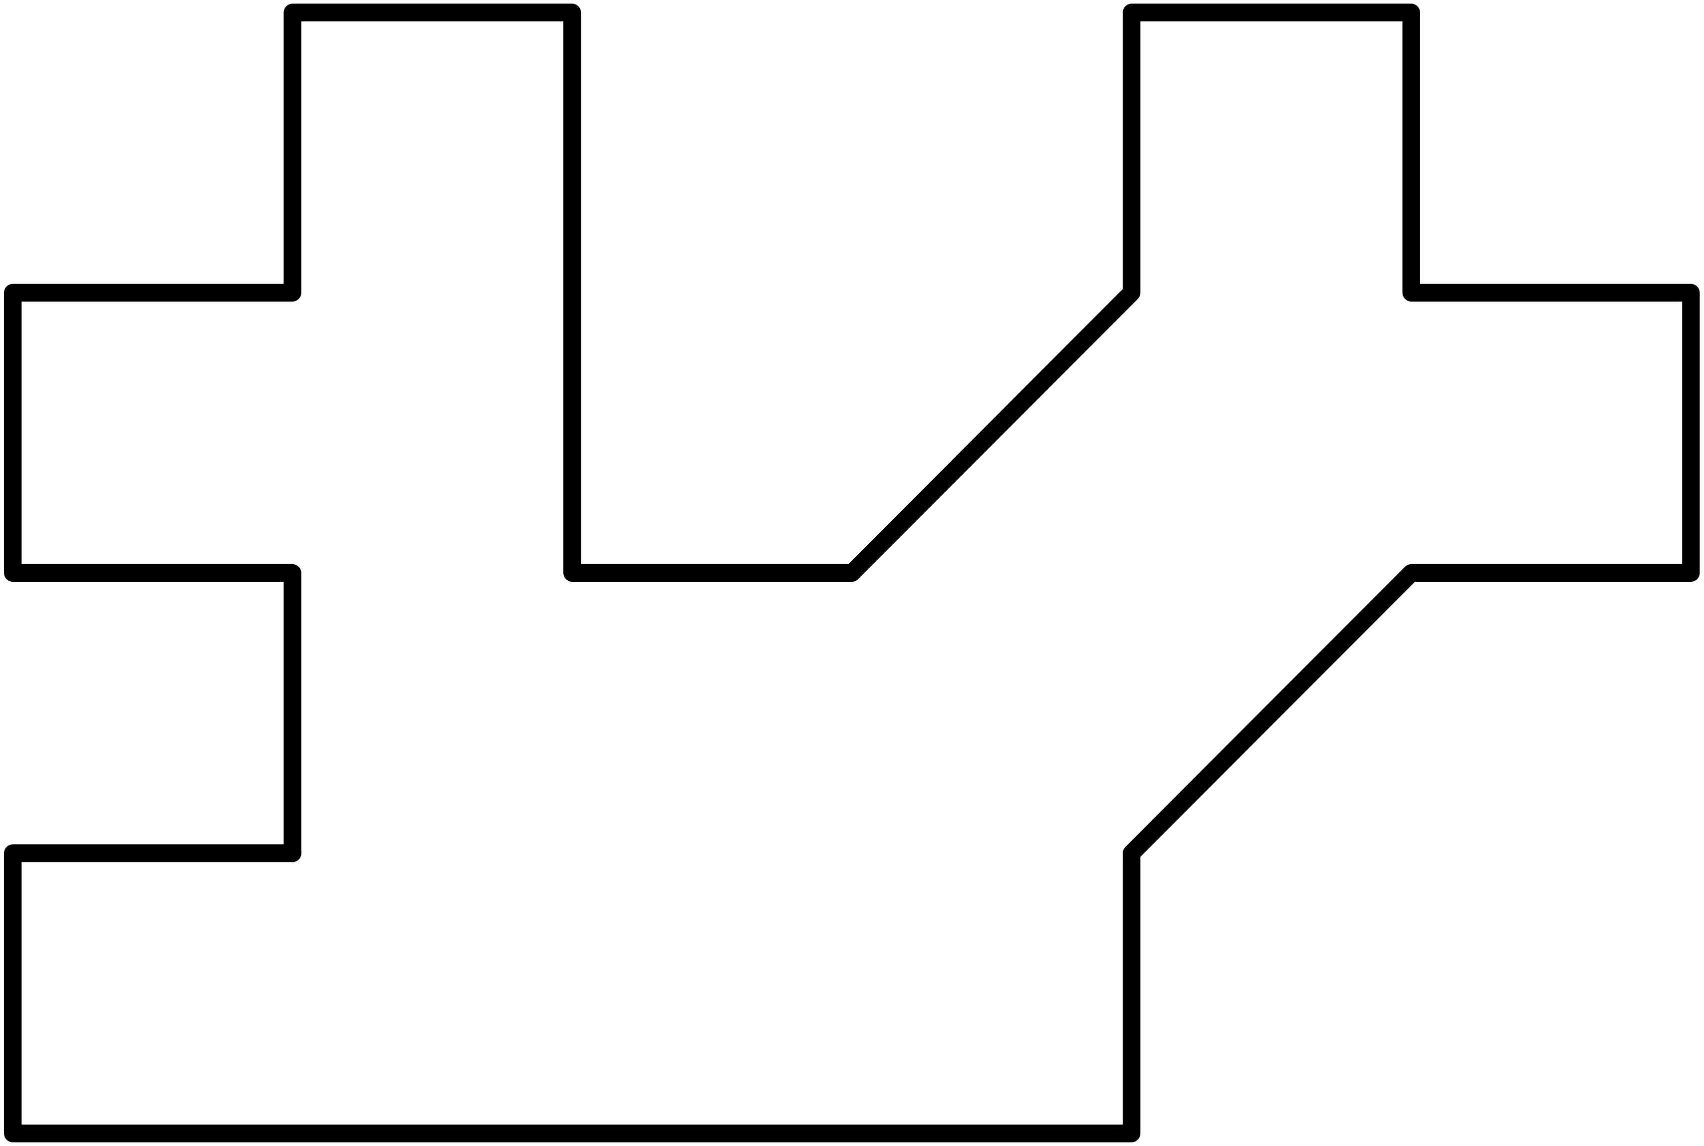

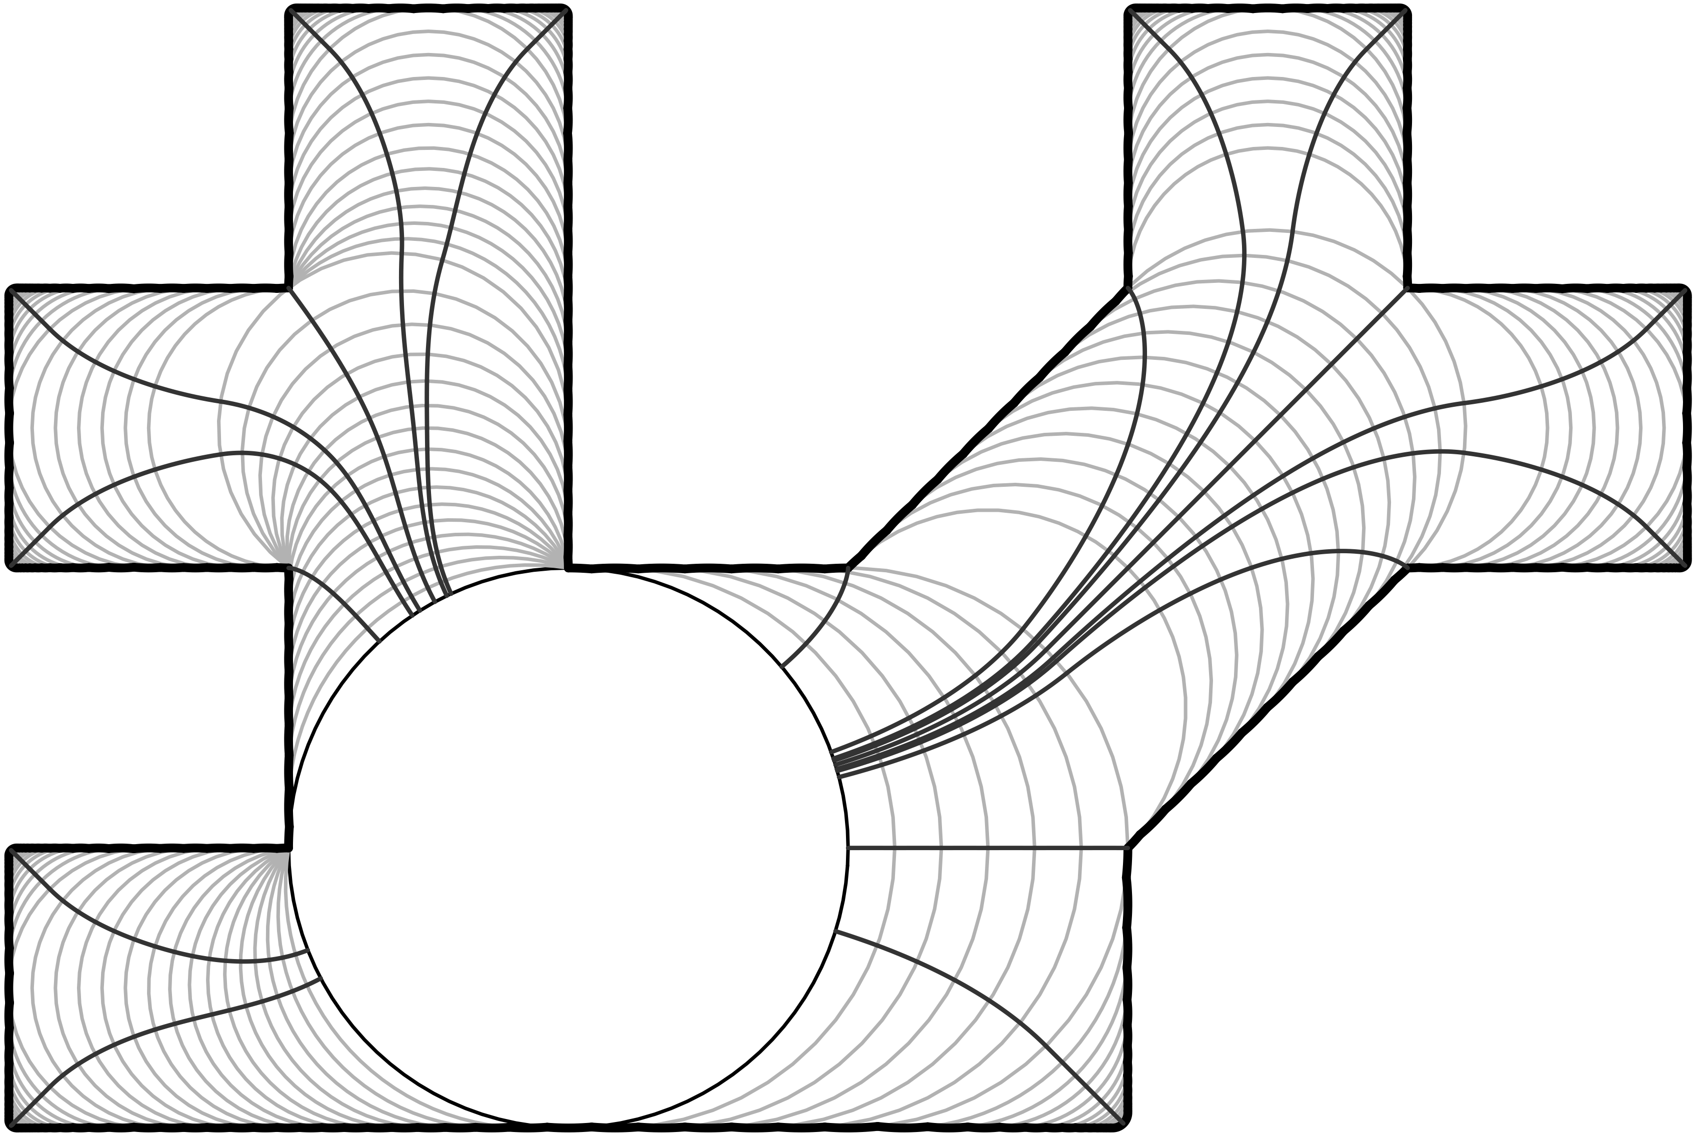

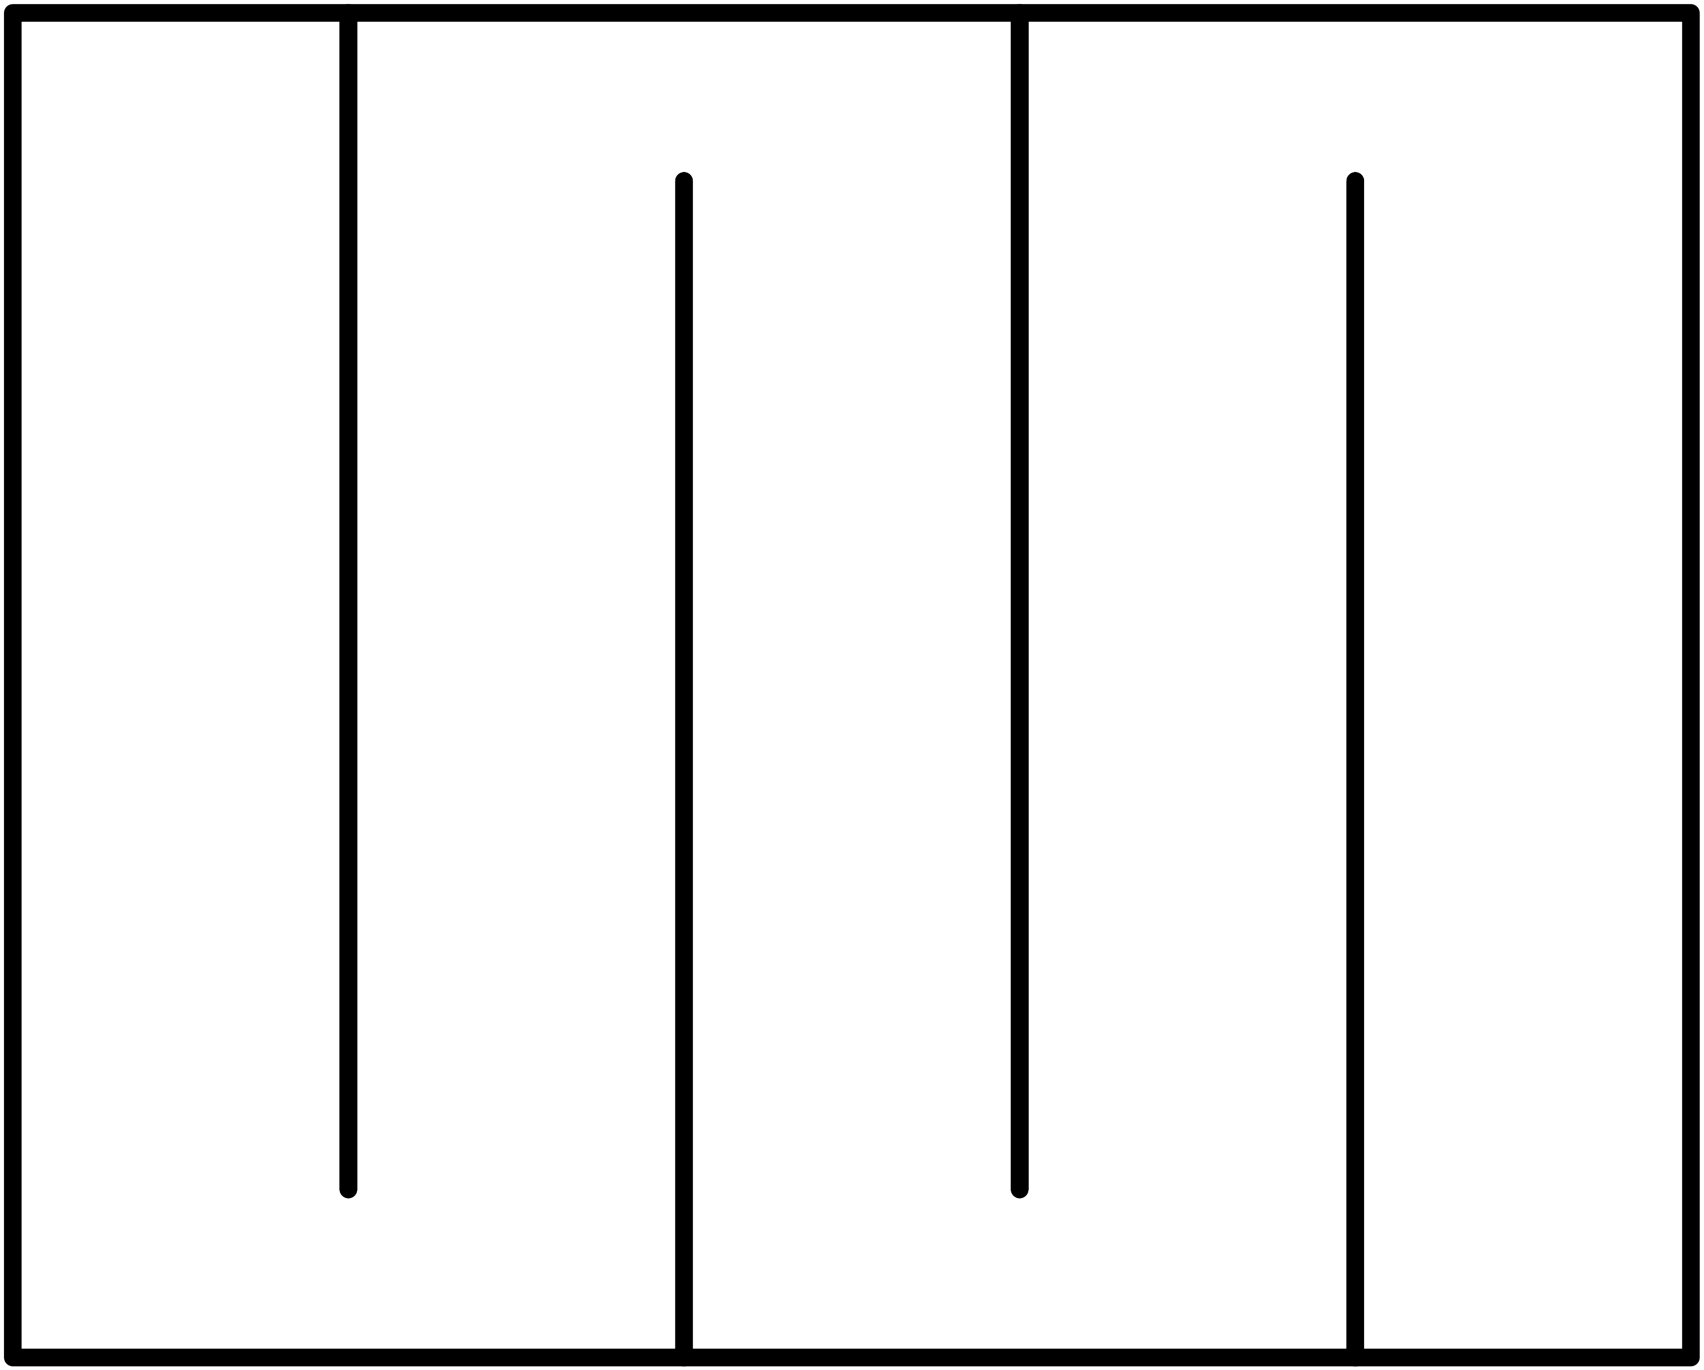

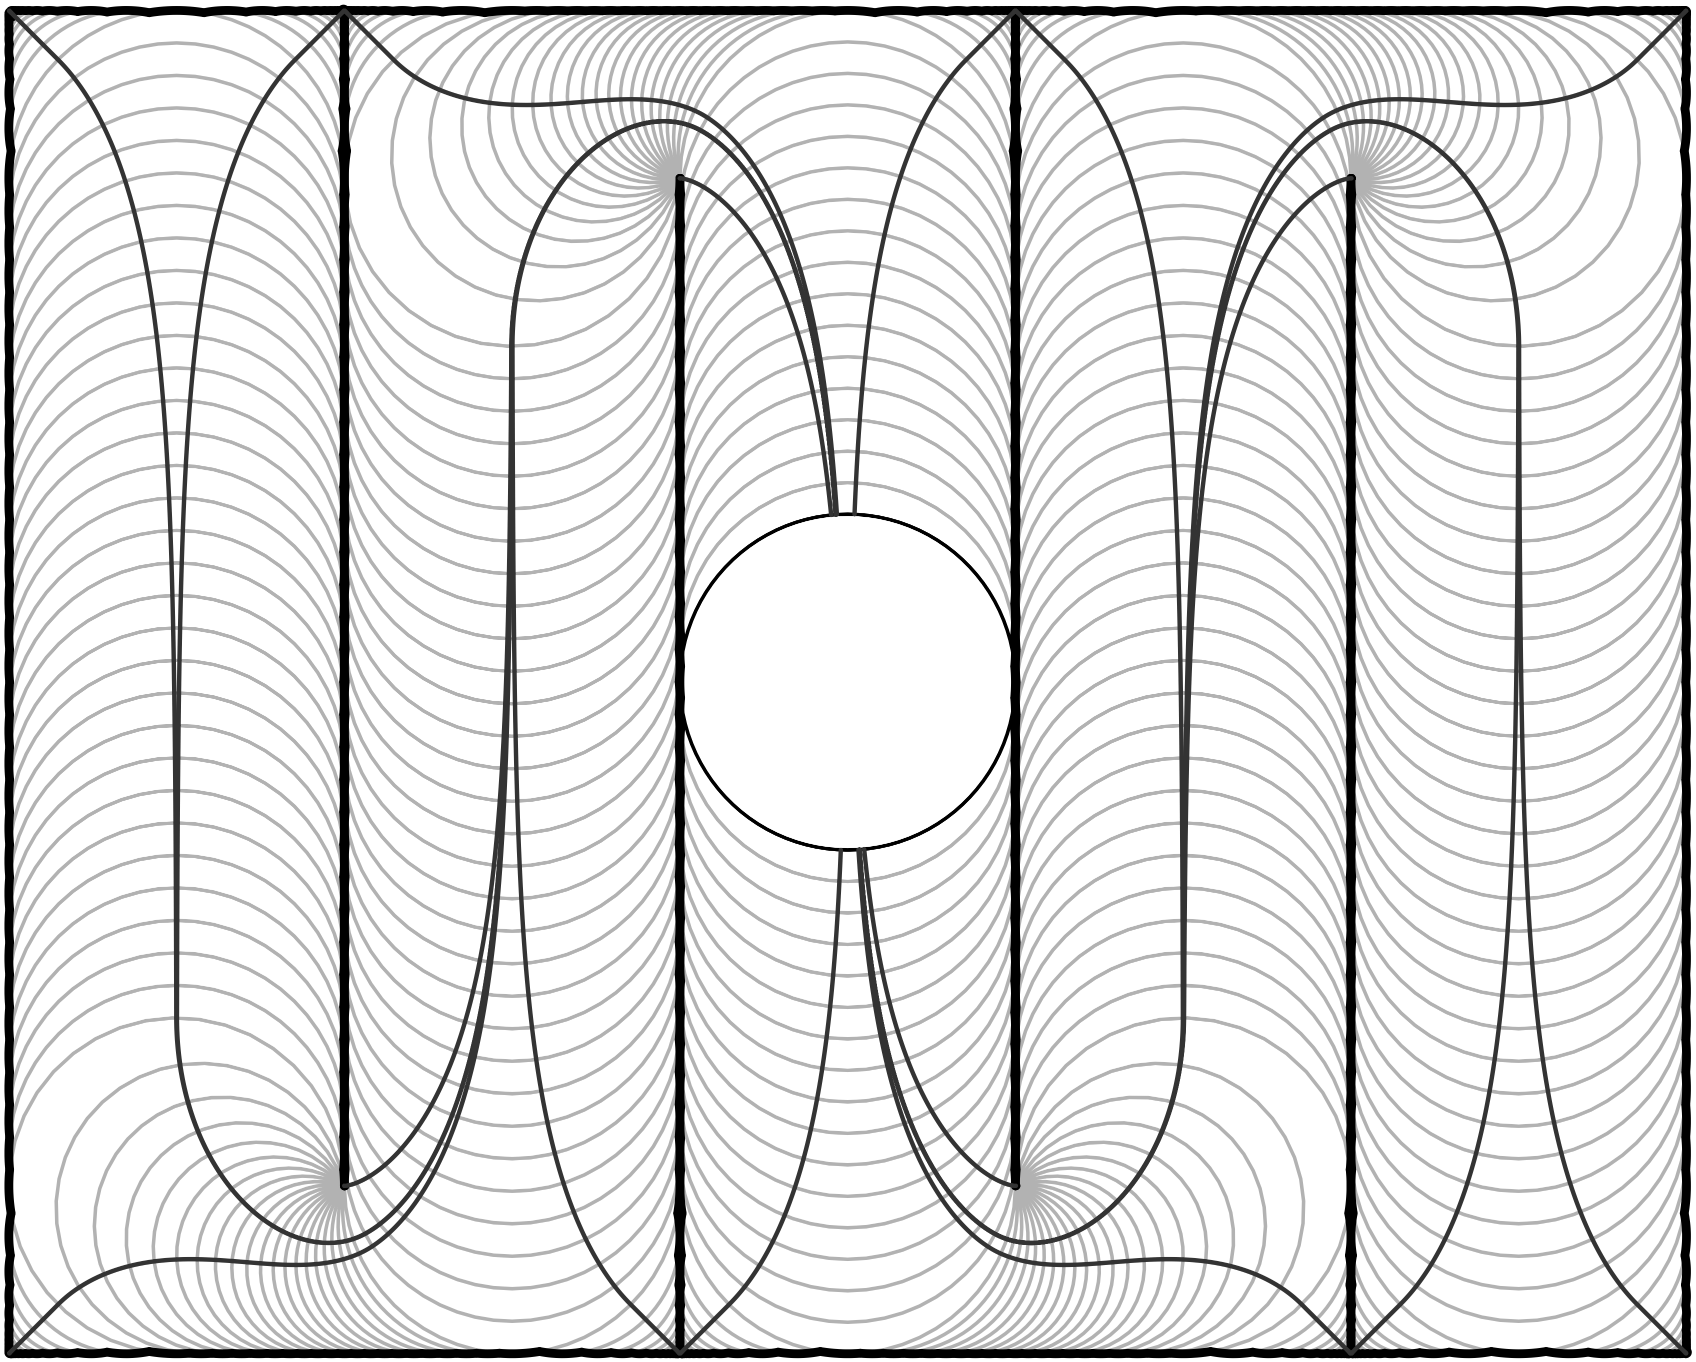

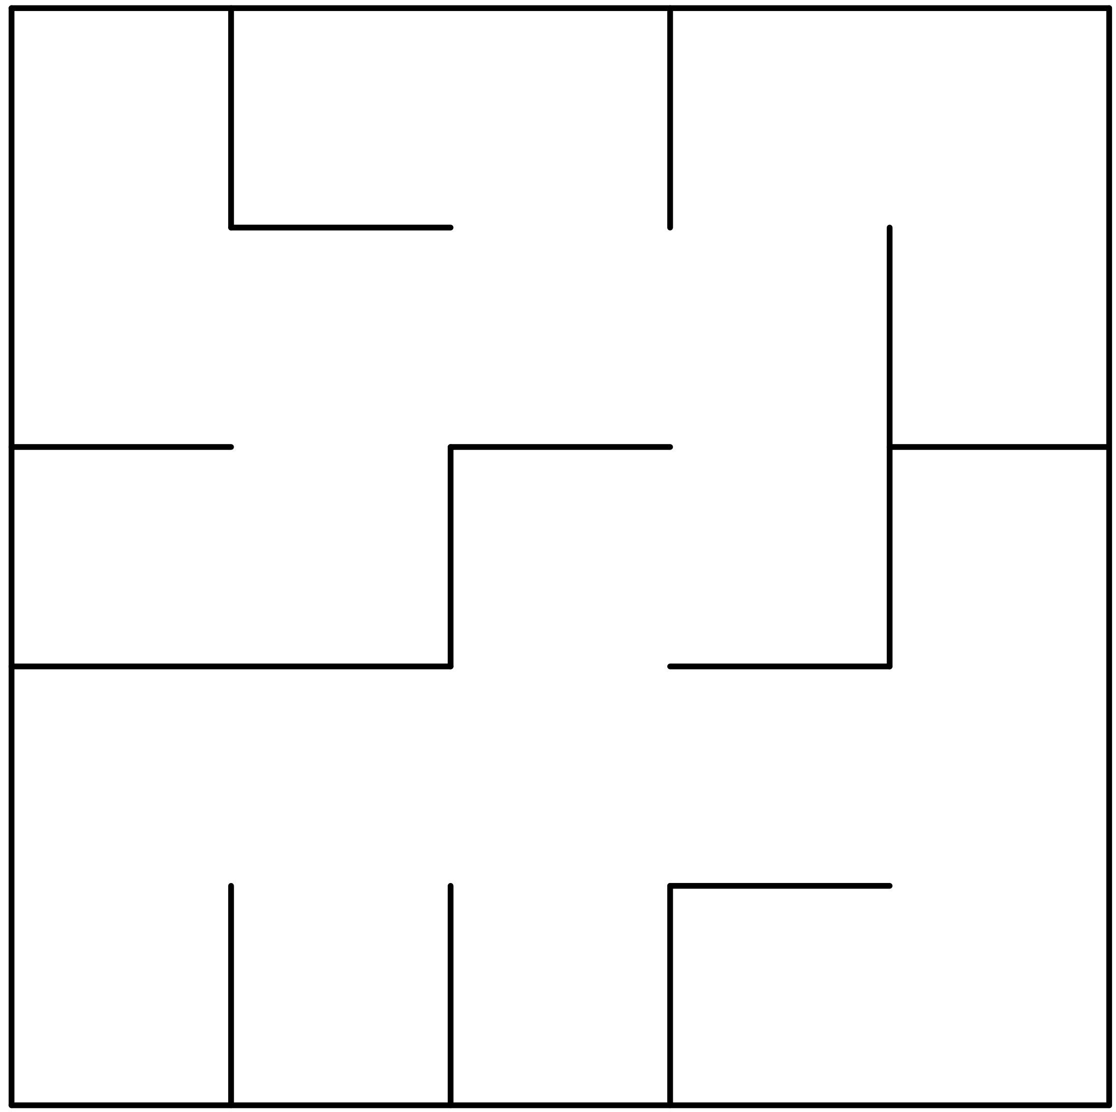

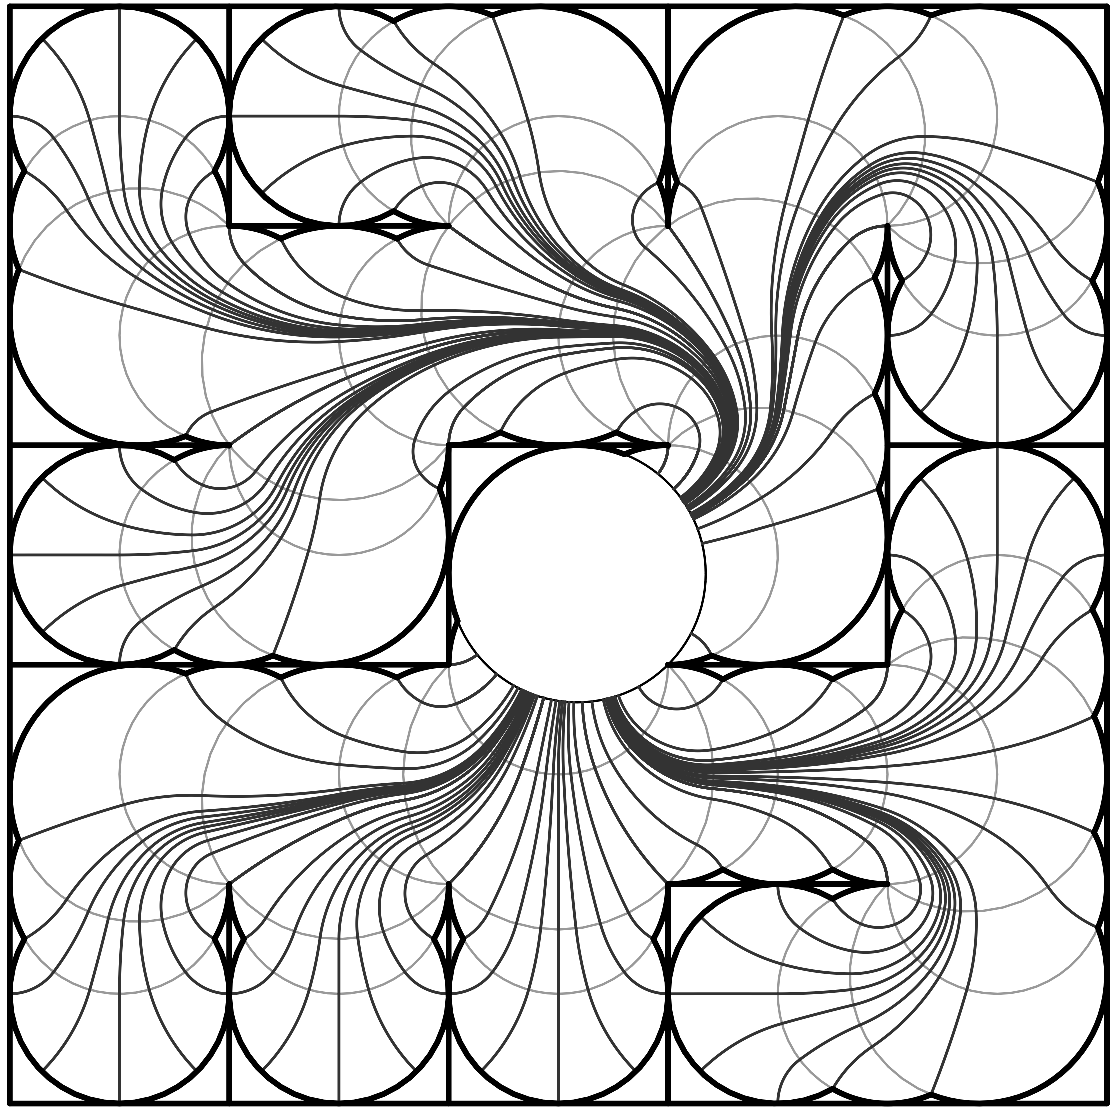

Theorem: Mapping all n vertices takes $O(n)$ time.

Uses linear time computation of MA (Chin-Snoeyink-Wang) and book-keeping with cross ratios.

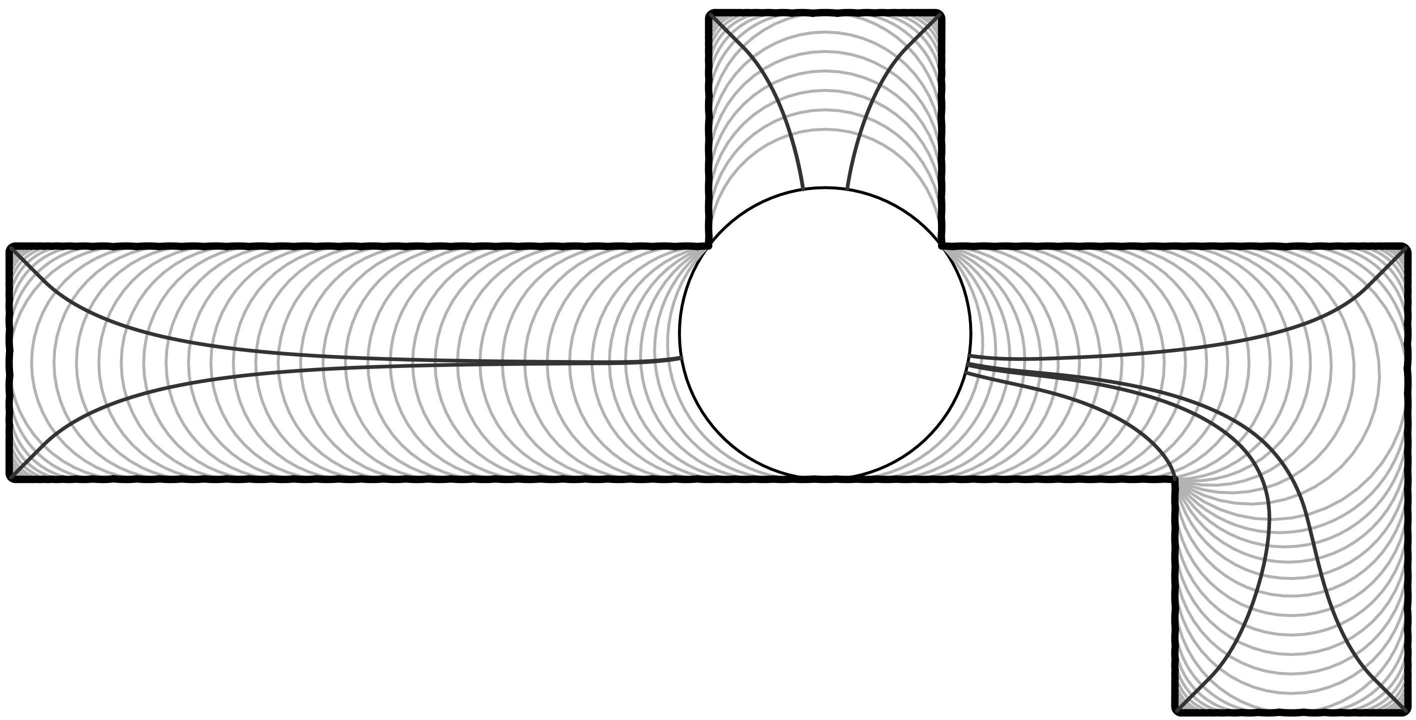

How close is medial axis map to conformal map?

How close is medial axis map to conformal map?

Use “MA-parameters” in Schwarz-Christoffel formula.

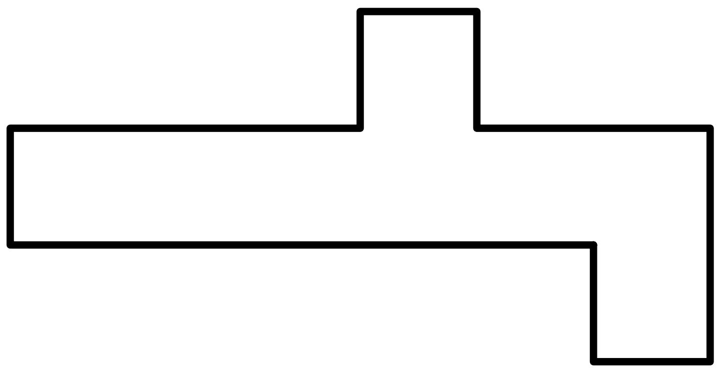

Target Polygon

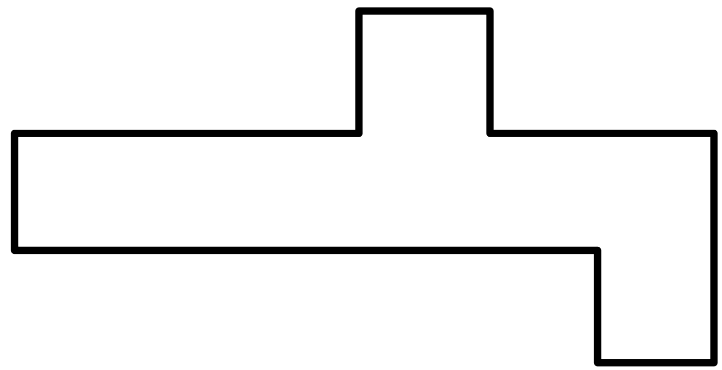

MA Parameters

How close is medial axis map to conformal map?

Use “MA-parameters” in Schwarz-Christoffel formula.

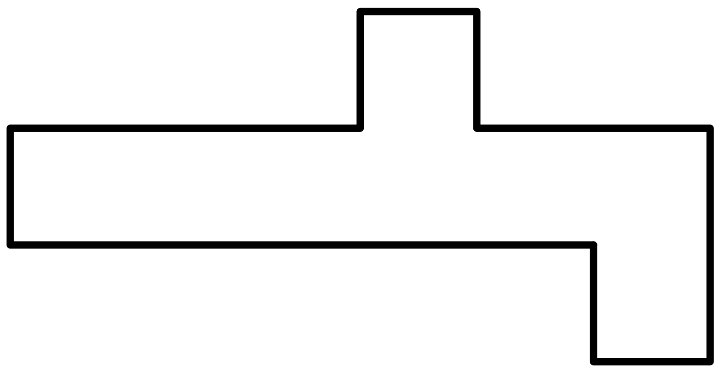

Target Polygon

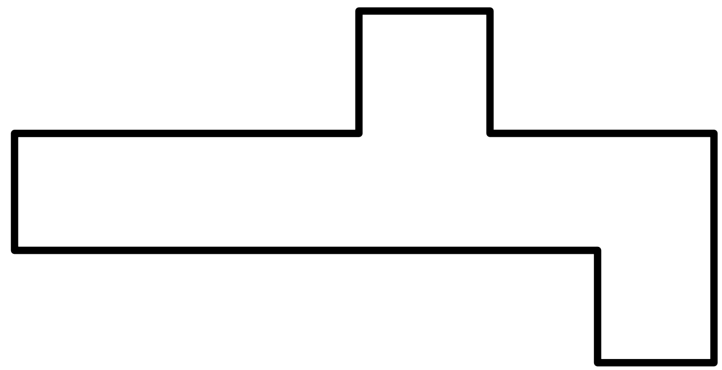

MA Parameters

Looks pretty close. How do we measure error?

How close is medial axis map to conformal map?

Use "MA-parameters" in Schwarz-Christoffel formula.

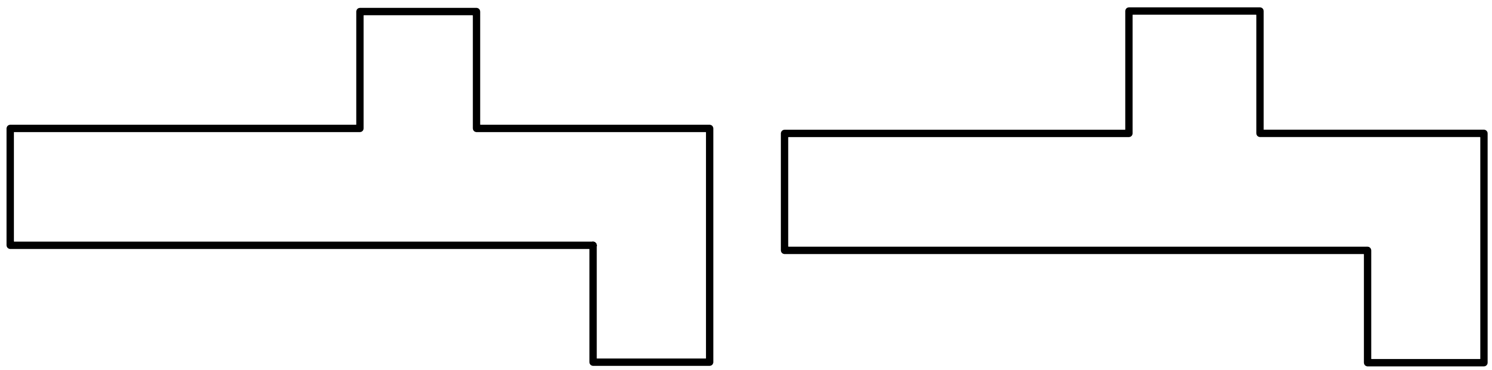

Target Polygon MA Parameters

Looks pretty close. How do we measure error?

By minimal K-QC map mapping vertices to vertices.

Any example gives upper bound, e.g., piecewise linear.

Upper bound for QC-distance:

- Find compatible triangulation
- Compute affine map between triangles
- Take largest K.

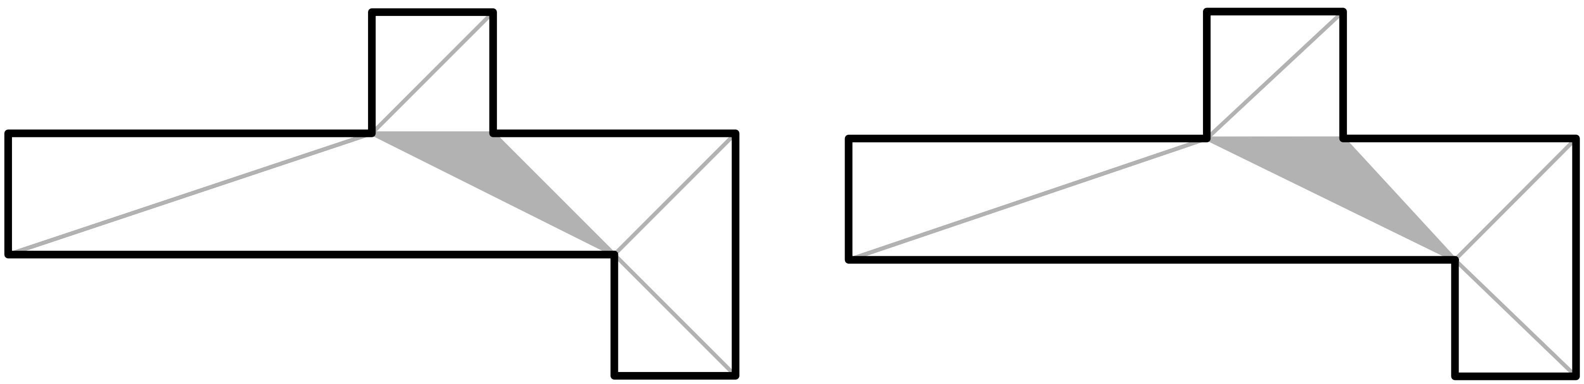

The most distorted triangle is shaded. Here $K = 1.24$.

We can bound QC-distance of our approximation to the correct answer without knowing the correct answer.

Chris Green 2012 Stony Brook Thesis.

- Computes polygon from parameter guess
- Computes PL map to target
- Transfers piecewise constant dilatation to disk
- Approximately solves Beltrami equation
- Updates parameters using resulting map $\mathbb{D} \to \mathbb{D}$
- Repeat

Proves quadratic convergence to correct parameters if intial guess is close enough that guessed polygon and target polygon have compatible triangulations.

Theorem: Medial axis map is always K-QC, $K < 8$.

Theorem: Medial axis map is always K-QC, $K < 8$.

MA-map extends to QC map of polygon interior to disk.

Why is this theorem true?

Theorem: Medial axis map is always K-QC, $K < 8$.

MA-map extends to QC map of polygon interior to disk.

Why is this theorem true?

Short answer: hyperbolic geometry in 3-space

Theorem: Medial axis map is always K-QC, $K < 8$.

MA-map extends to QC map of polygon interior to disk.

Why is this theorem true?

Longer answer: because MA-map has conformal extension to different region with same boundary.

Other region is surface in 3-space.
The surface can be mapped conformally to disk.
Medial axis map = boundary values of this map.

The **dome** of a domain is upper envelope of all hemispheres where base disk is a medial axis disk of Ω.

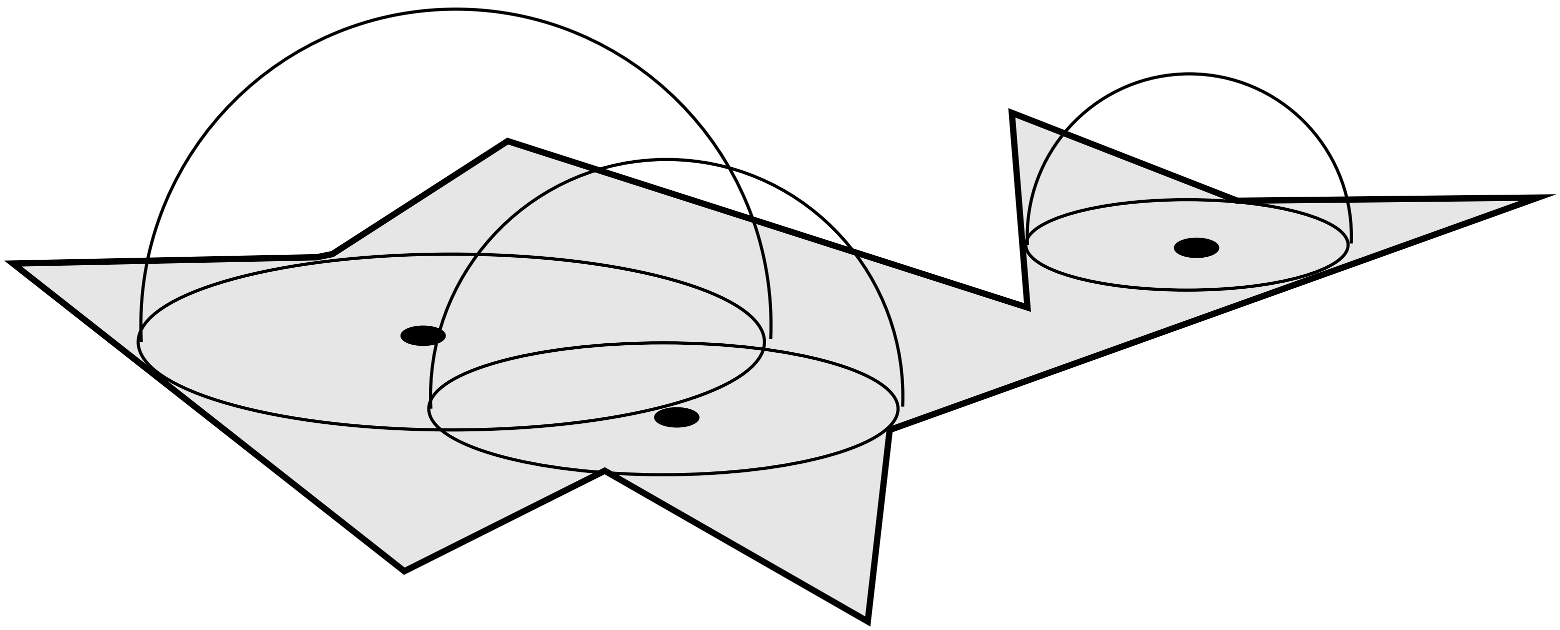

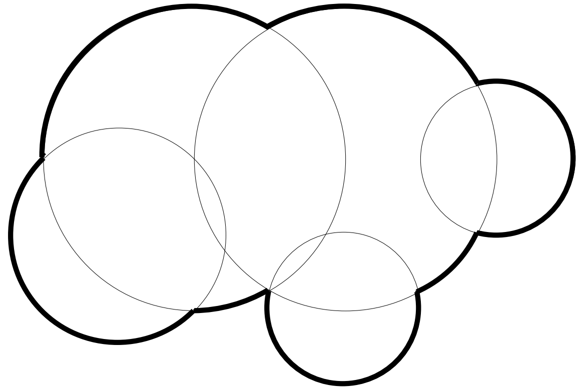

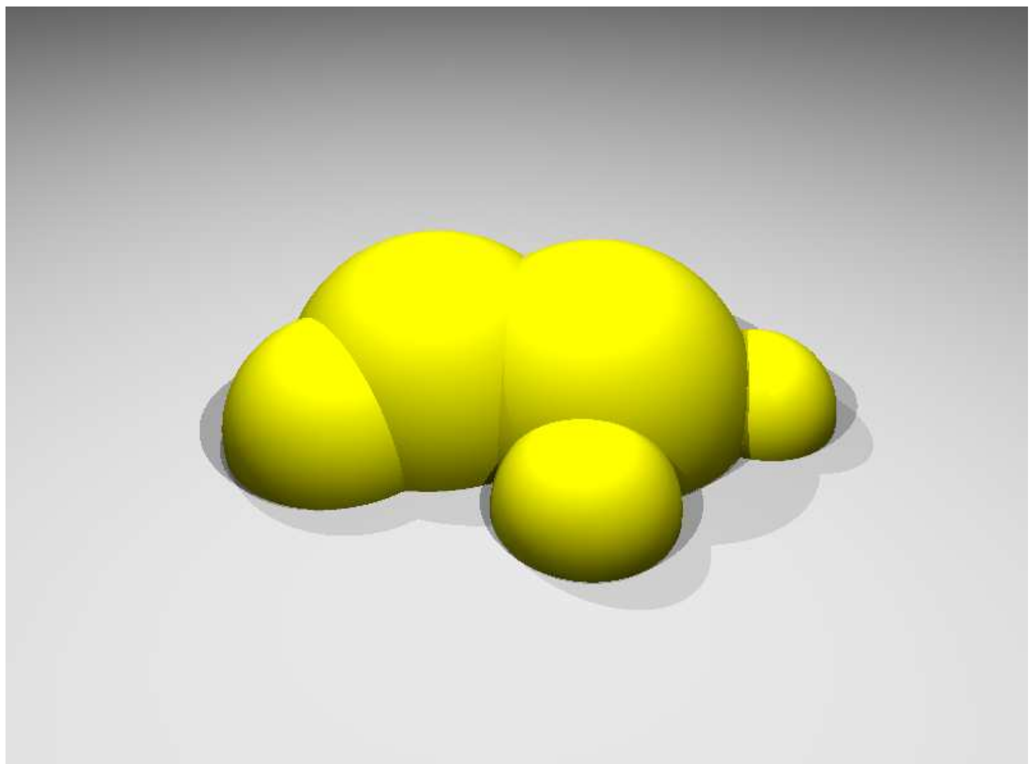

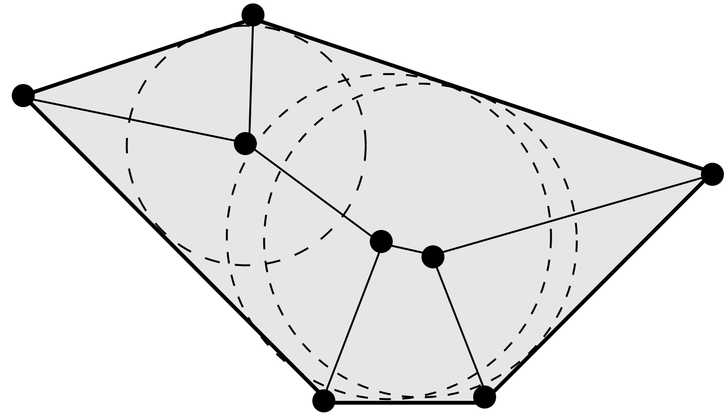

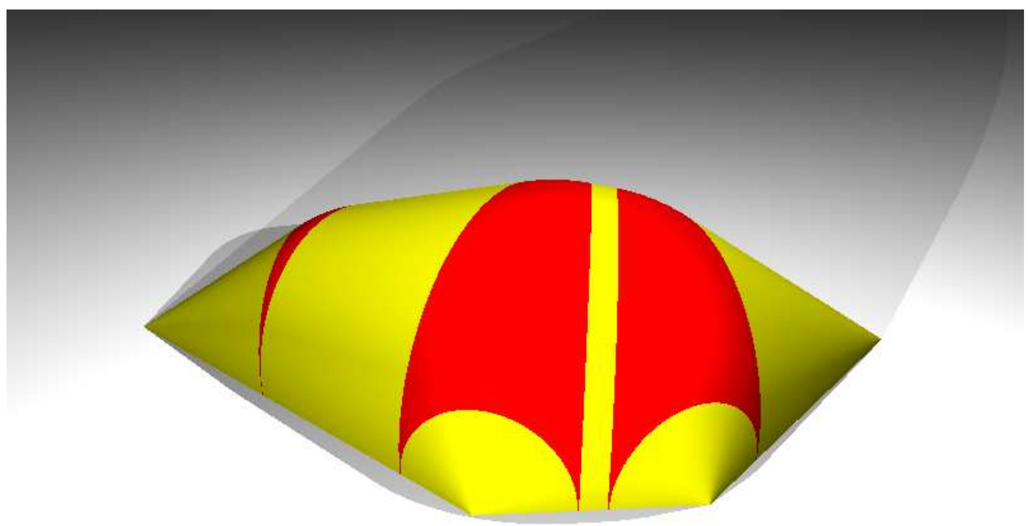

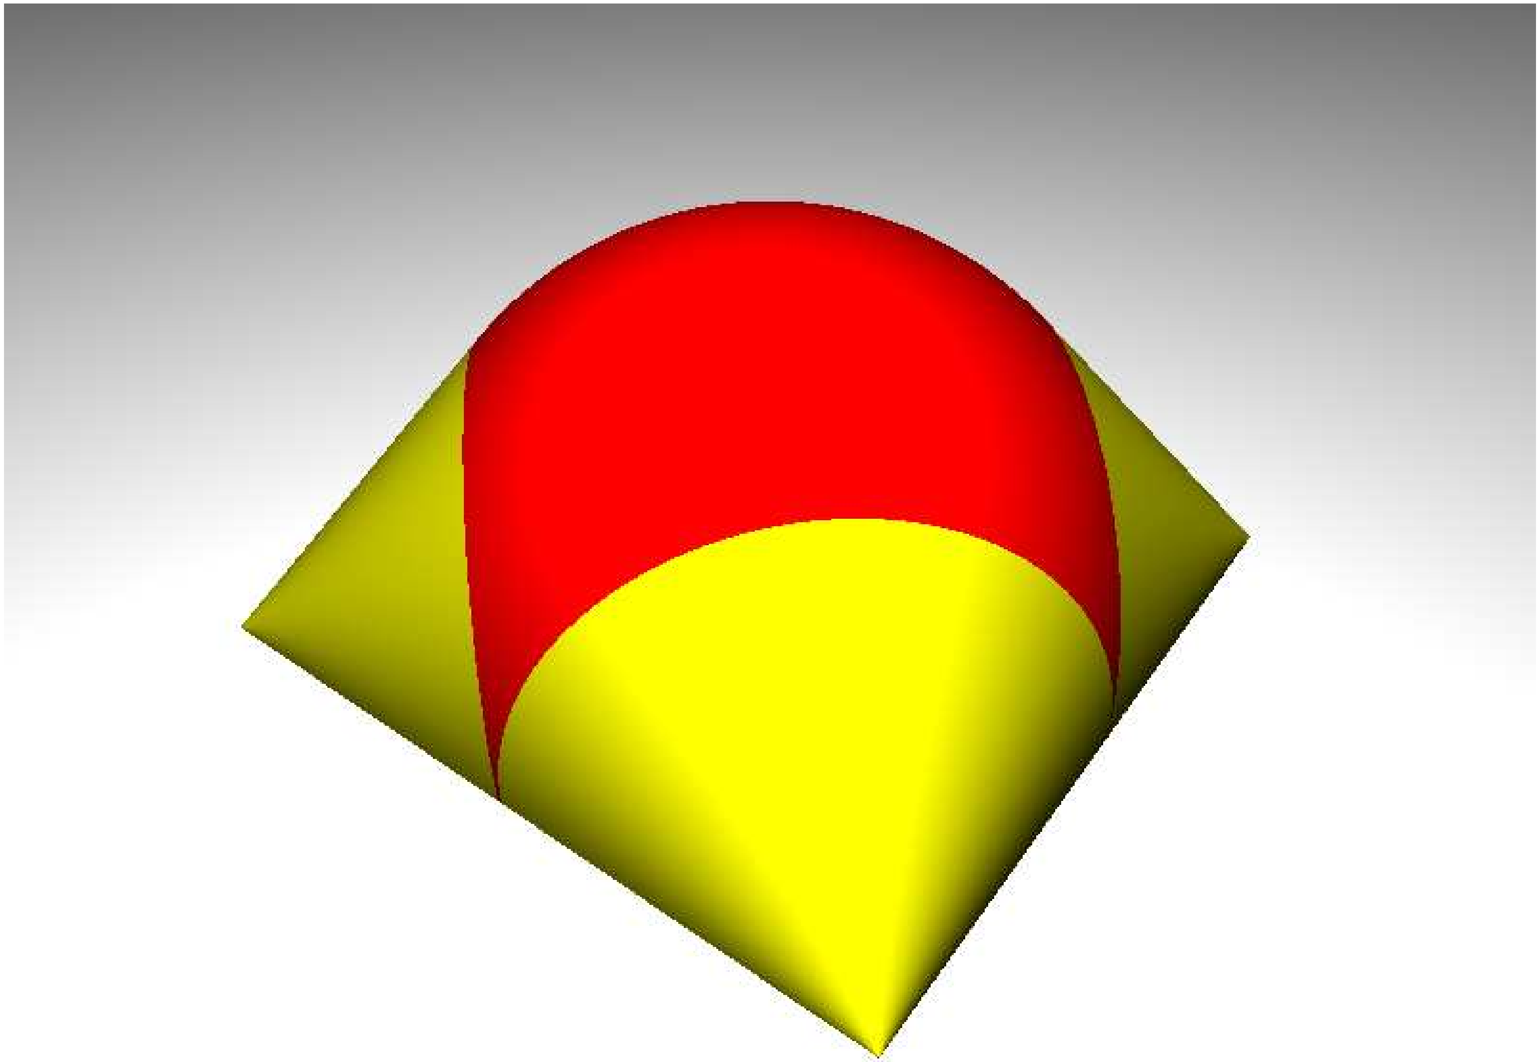

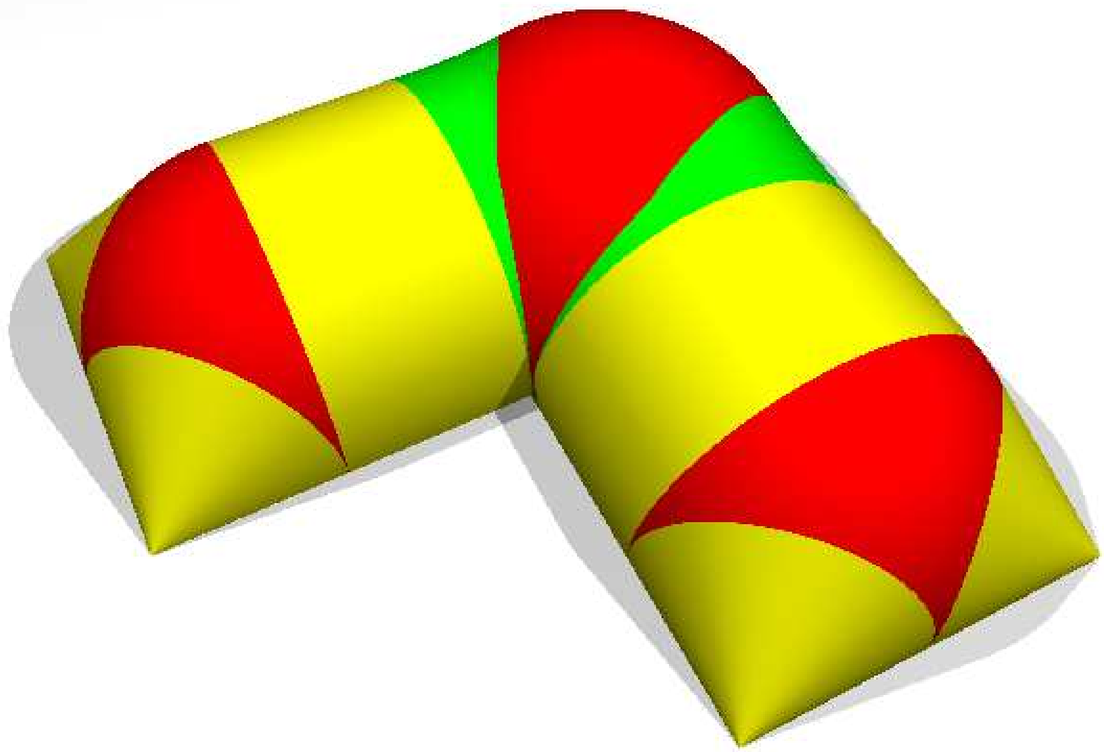

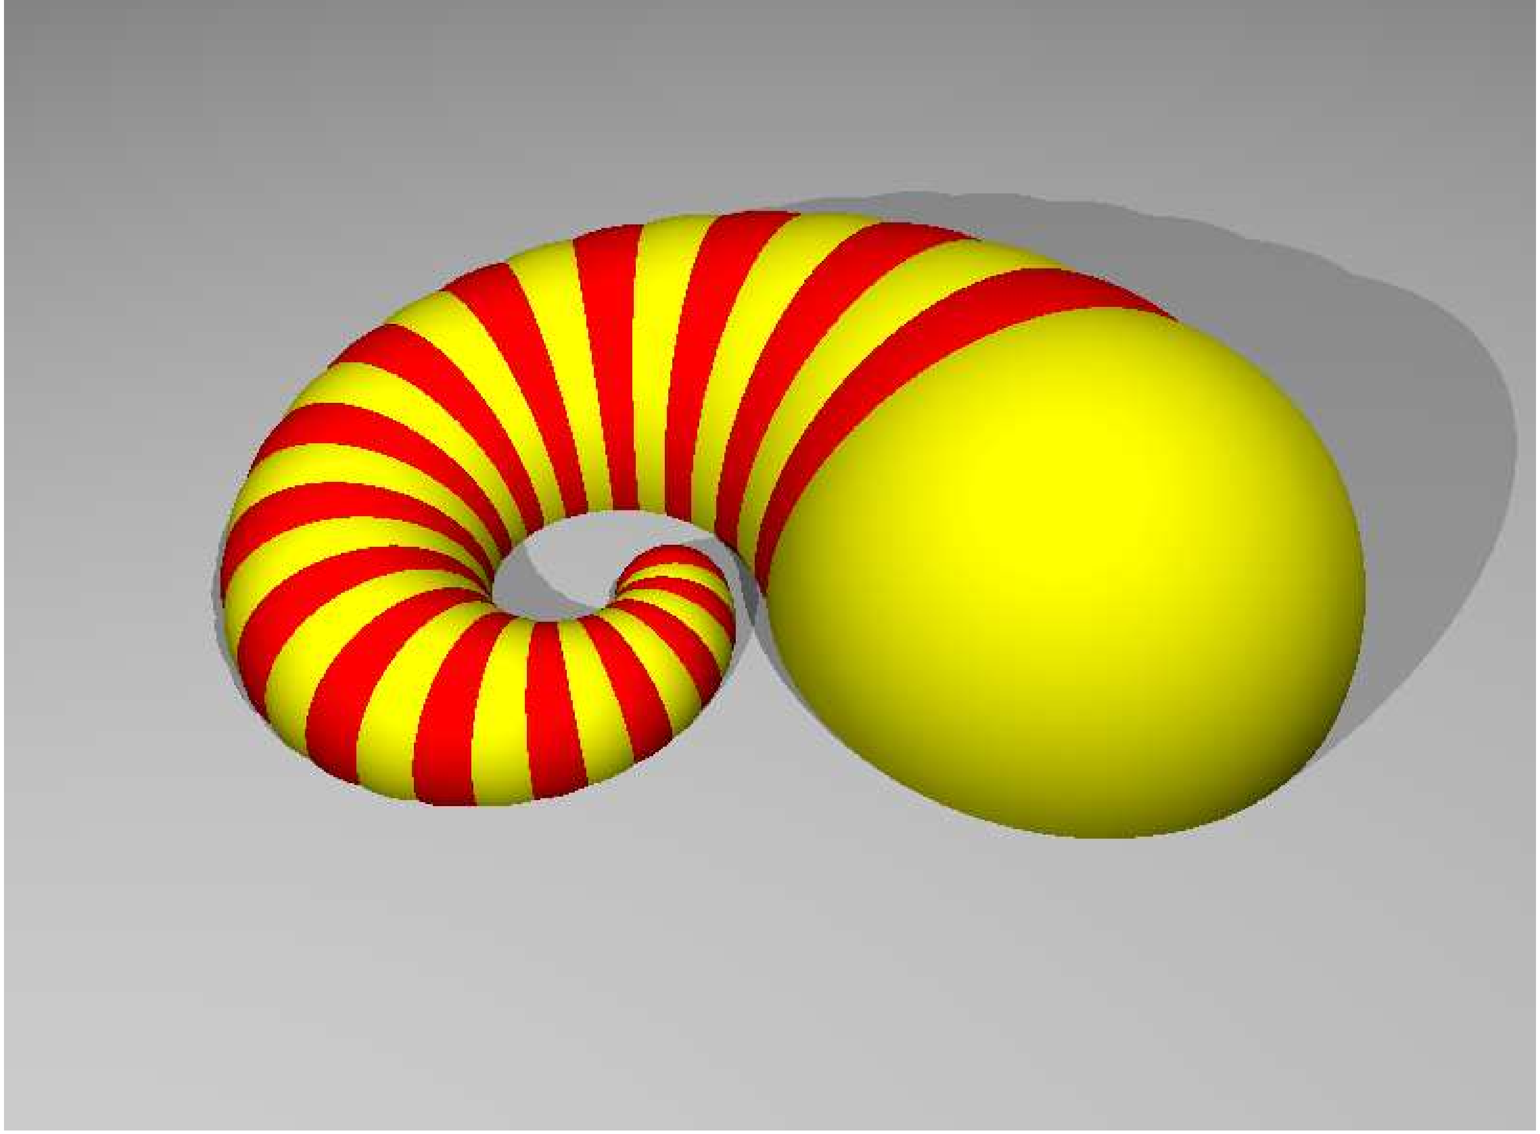

Every dome has conformal map to disk by “flattening”.

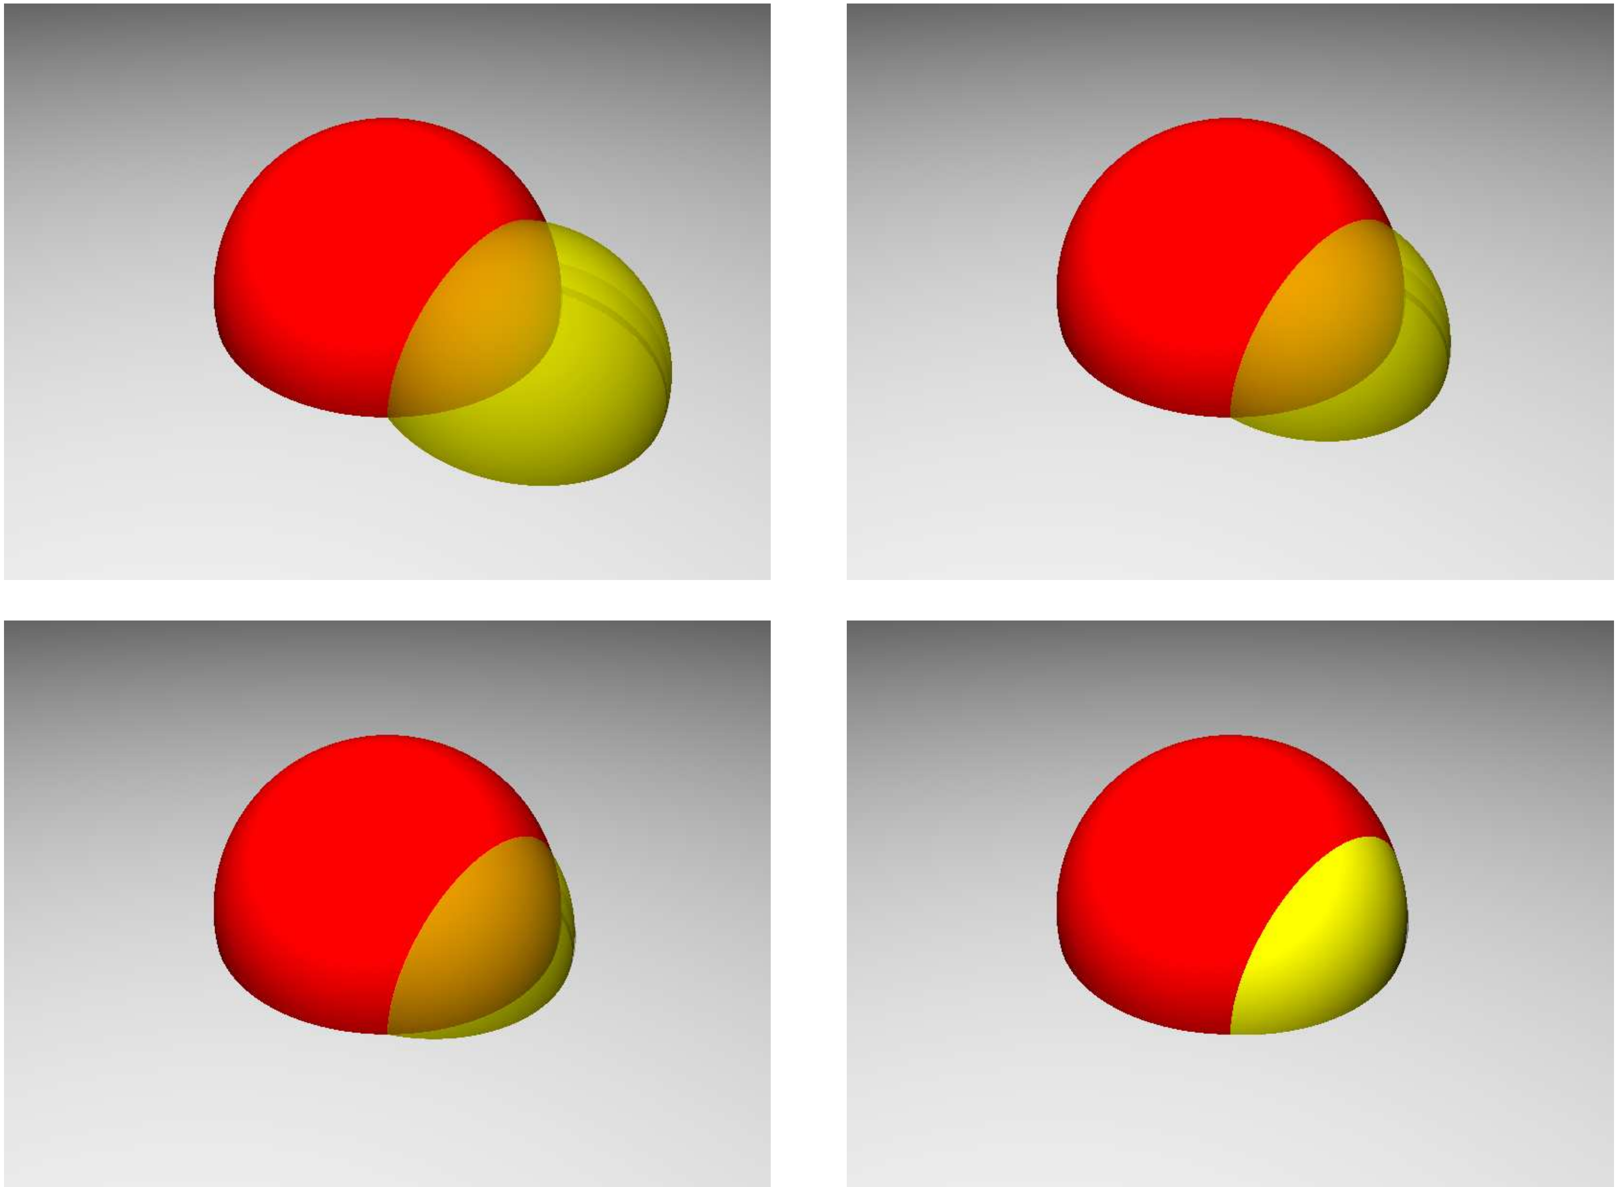

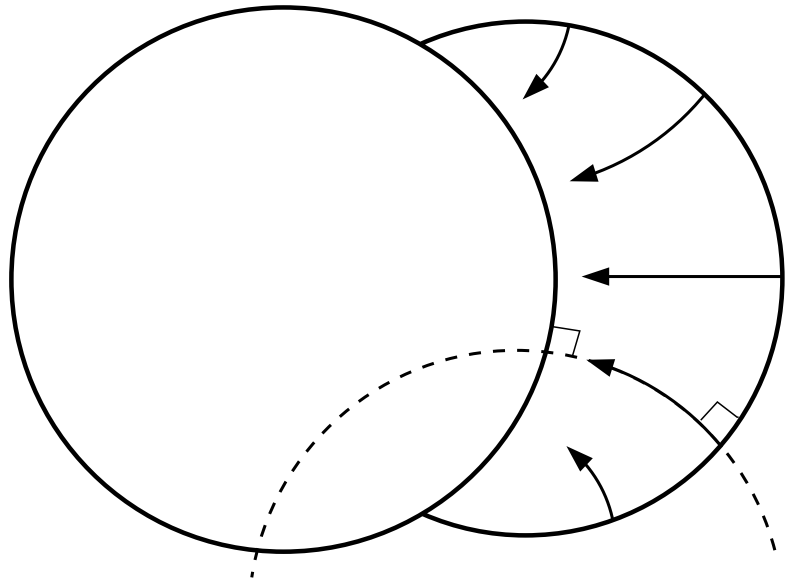

Medial axis map = boundary of flattening map (iota)

= boundary of conformal map of dome to hemisphere

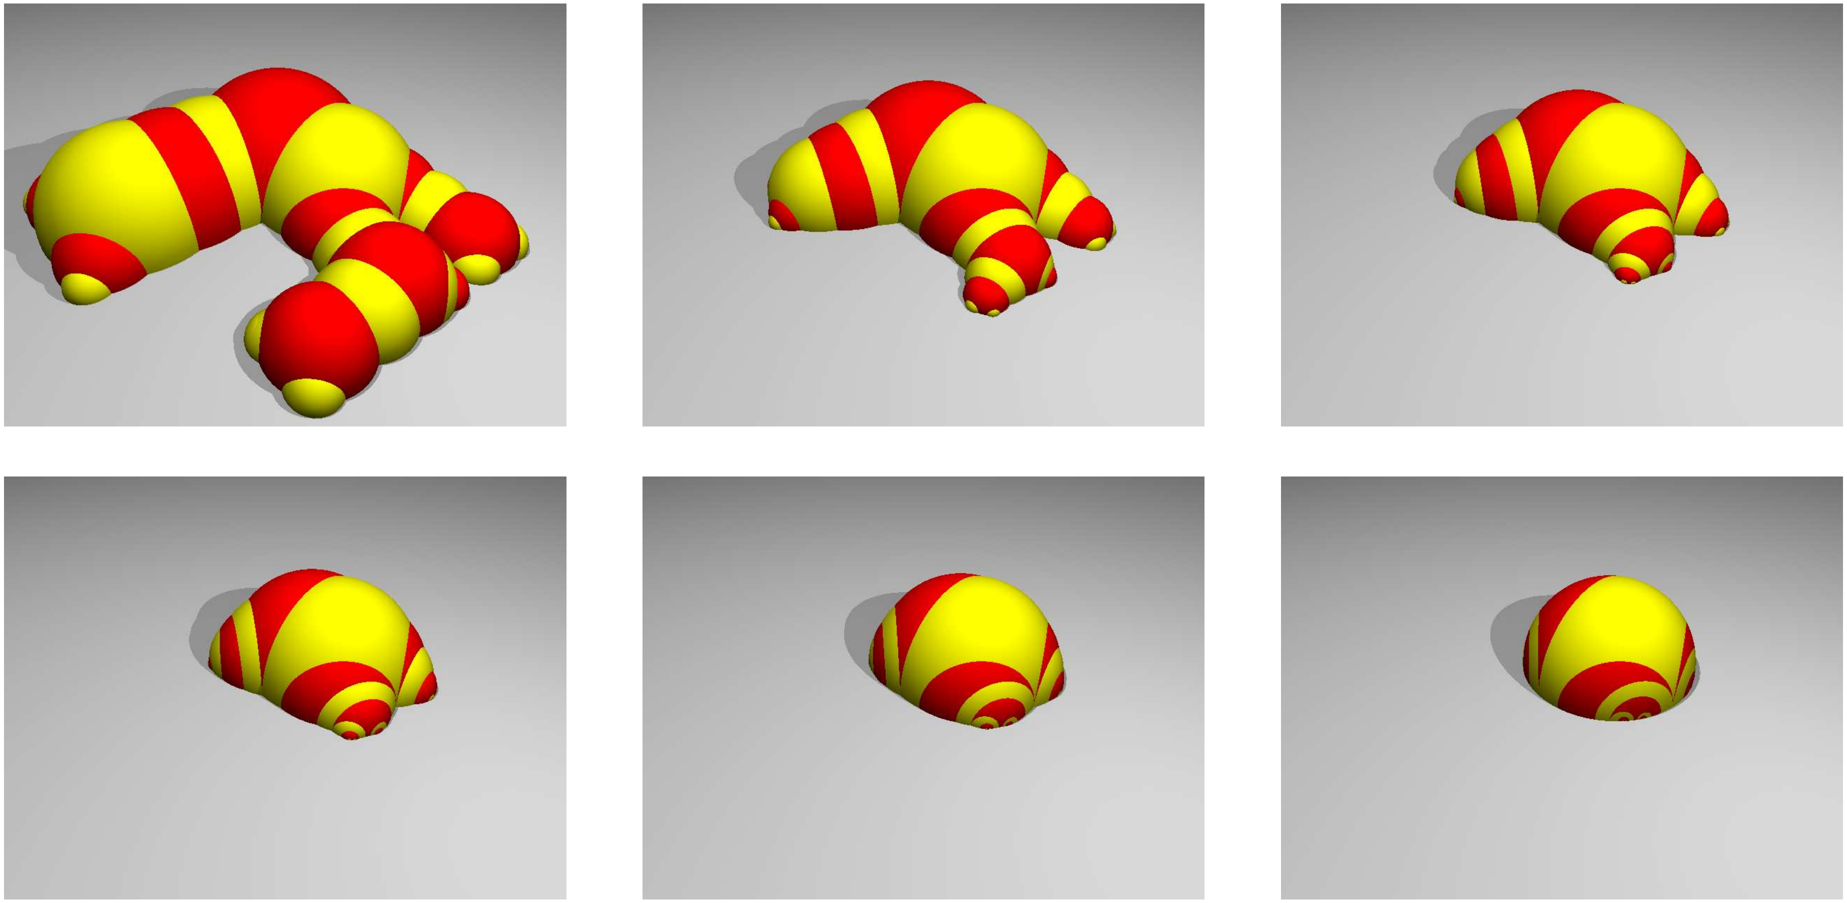

Map dome to hemisphere by makings sides flush.

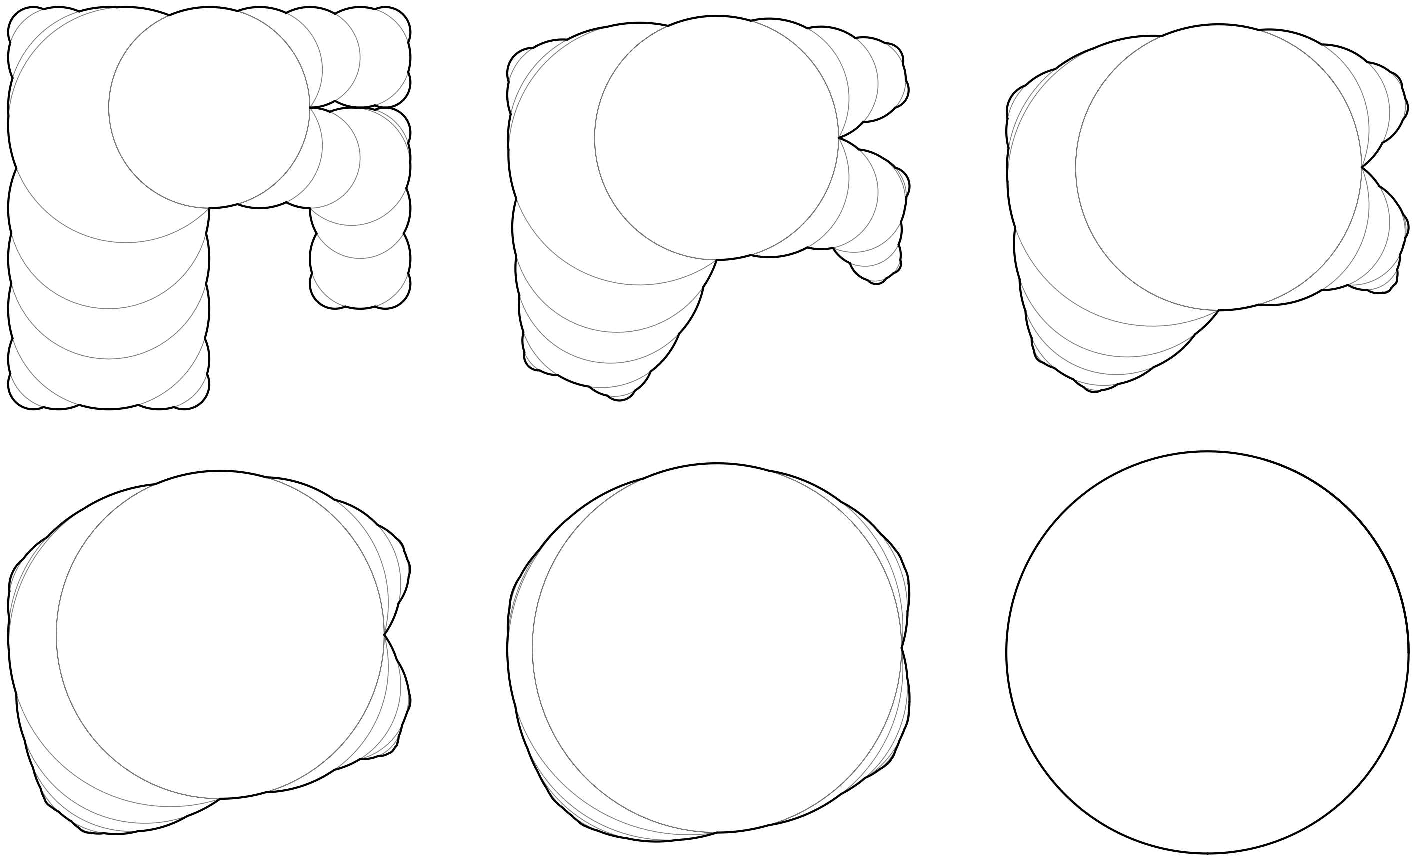

Angle scaling family: flatten bends simultaneously

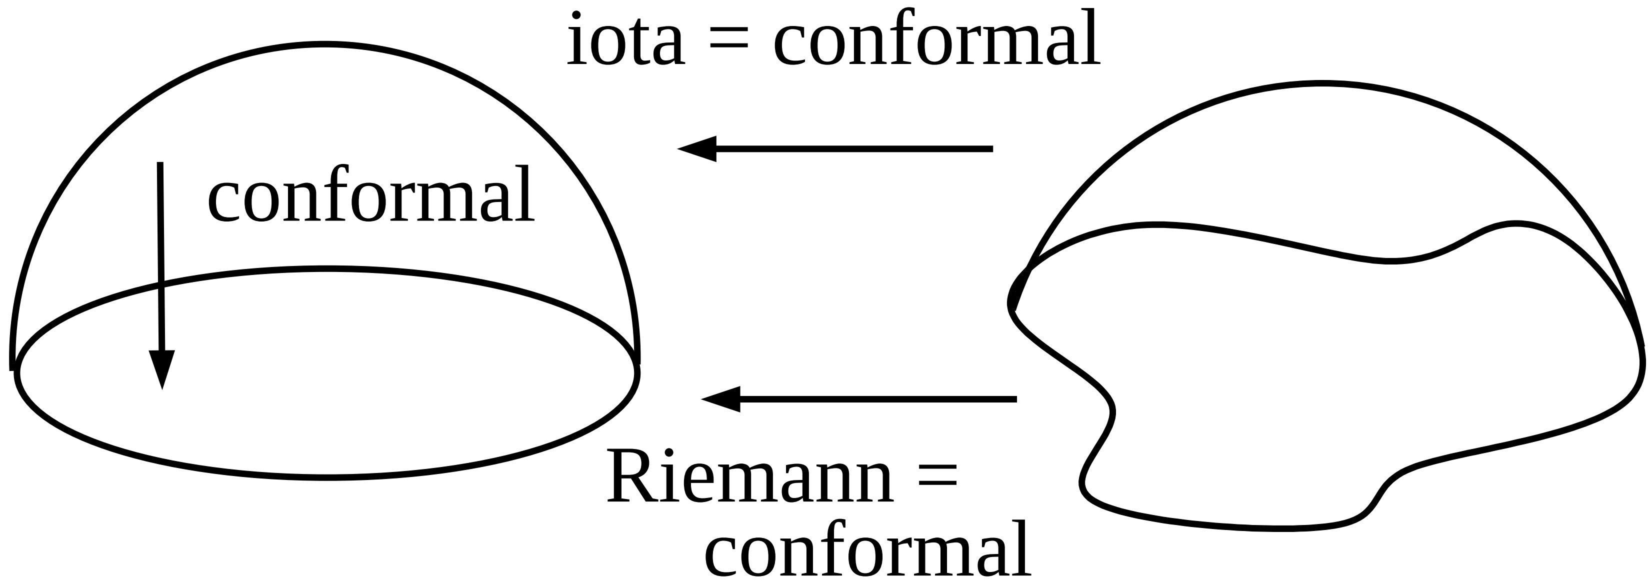

Iota = conformal from dome to disk.

Medial axis flow = boundary values of iota

Riemann map = conformal from base to disk

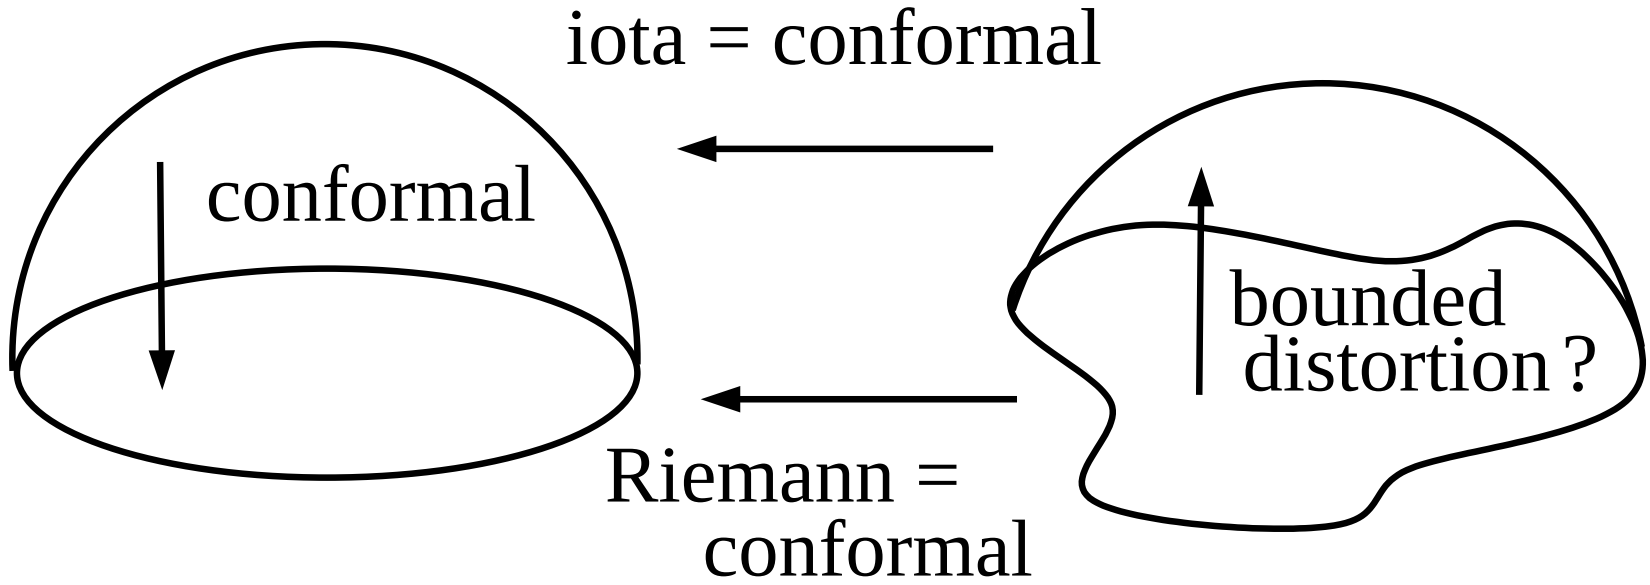

MA-map from polygon to circle has QC extension
iff
there is QC map from base to dome fixing boundary.

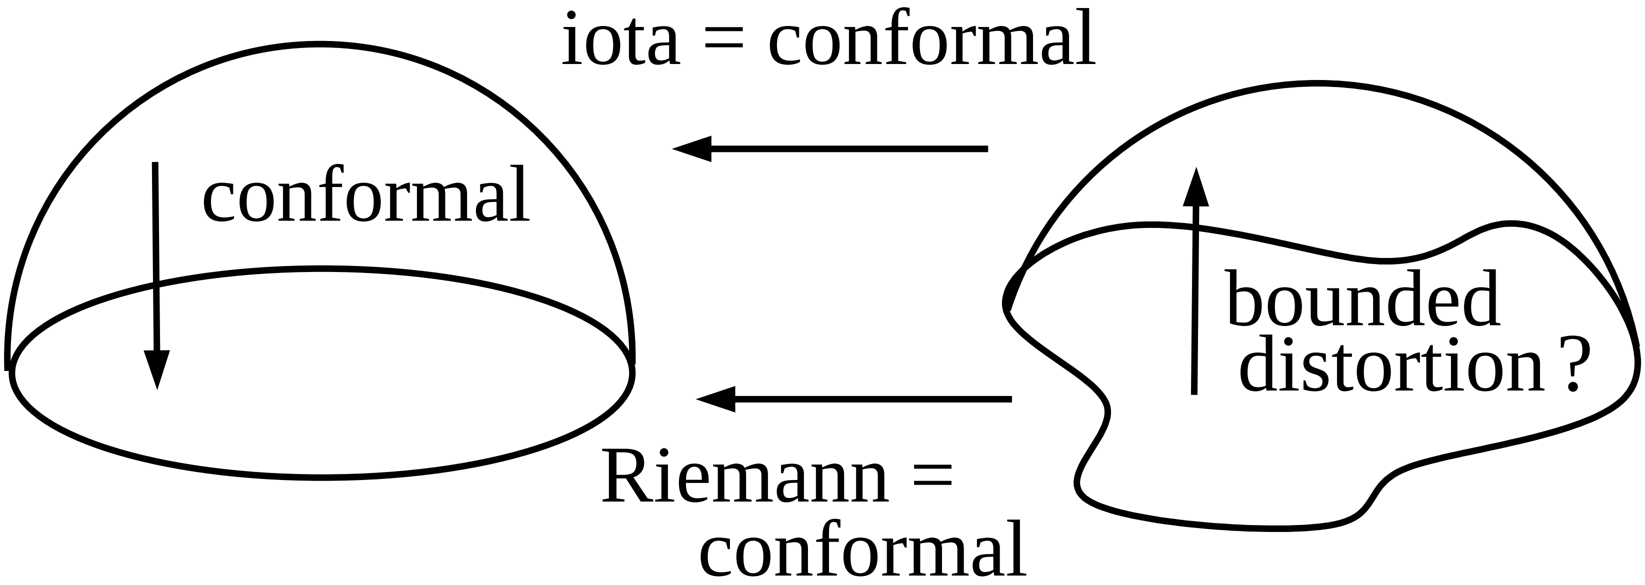

MA-map from polygon to circle has QC extension
iff
there is QC map from base to dome fixing boundary.

2nd claim is Dennis Sullivan's **convex hull theorem**.

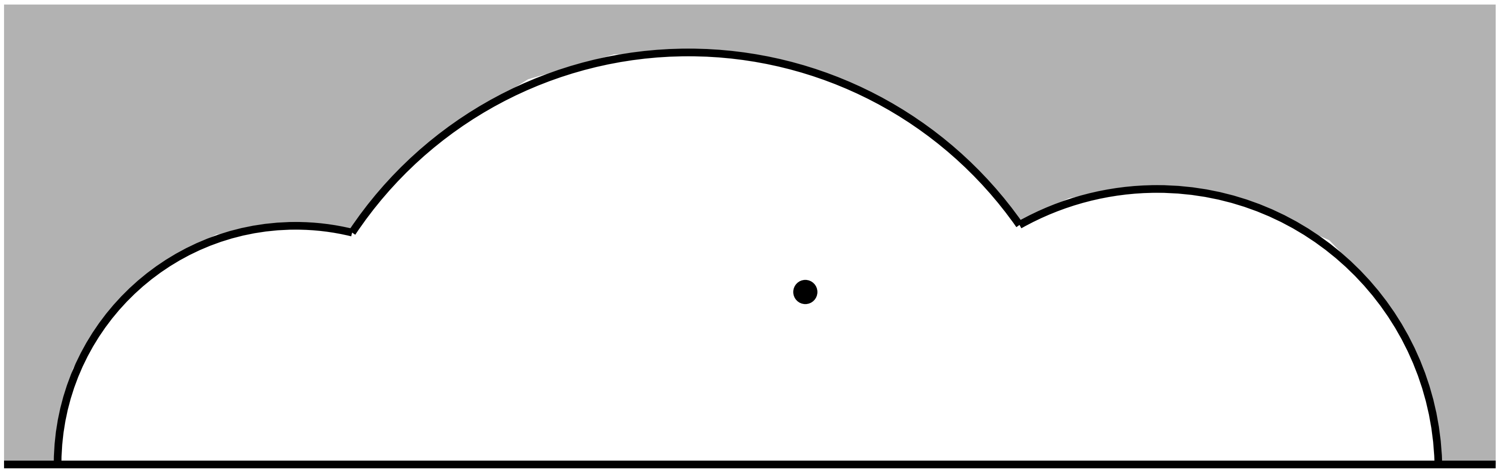

Region below dome is union of hemispheres

Hemispheres = hyperbolic half-spaces.

Region above dome is hyperbolically convex.

Consider nearest point retraction onto this convex set.

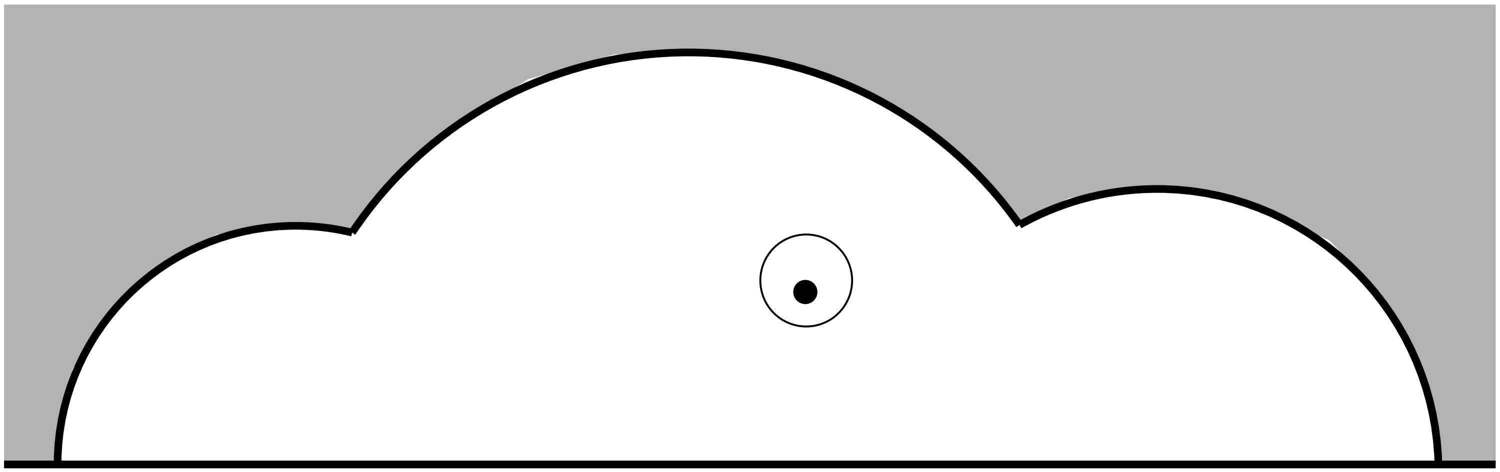

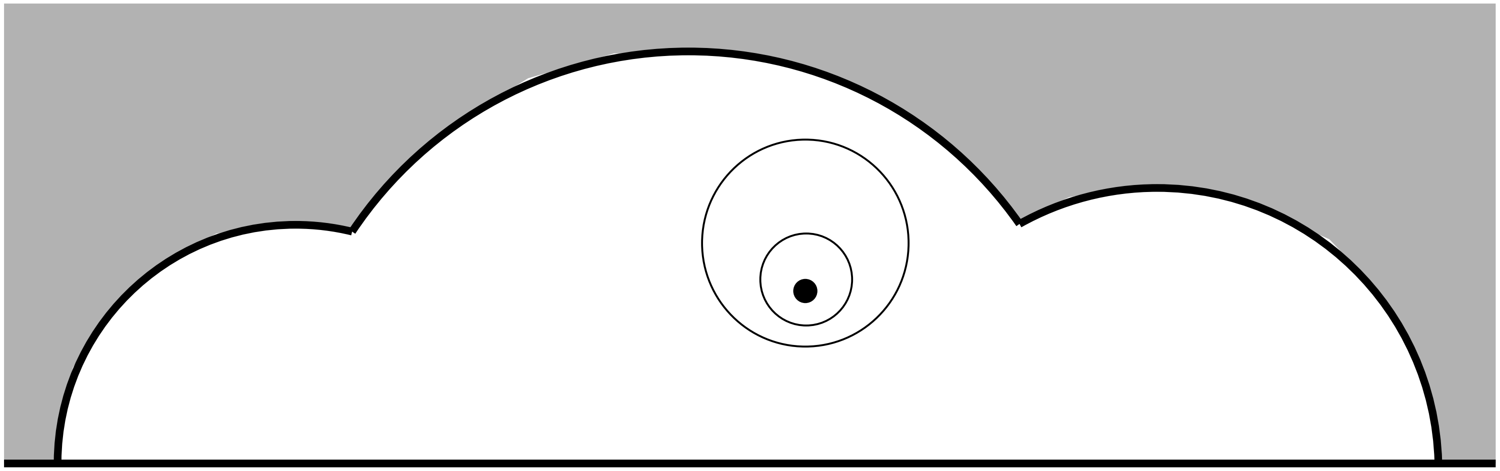

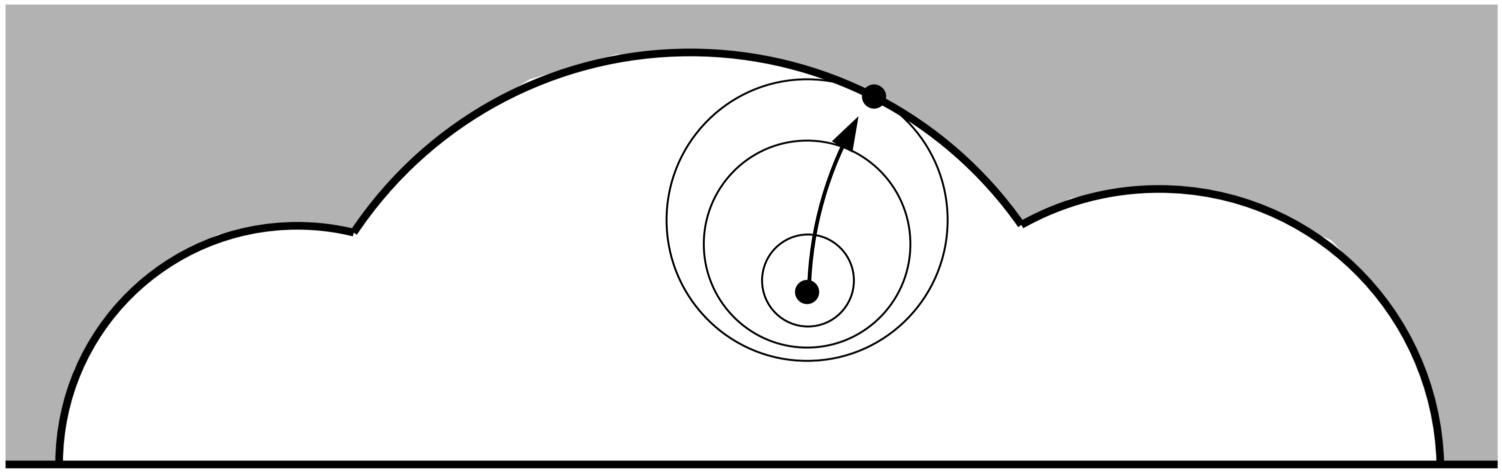

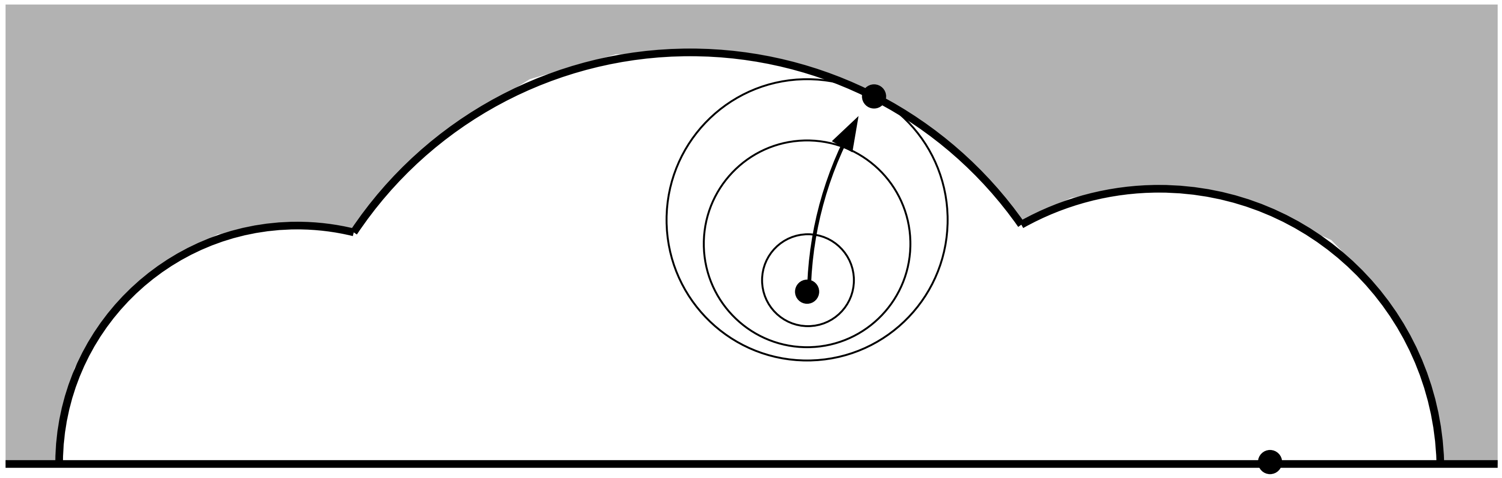

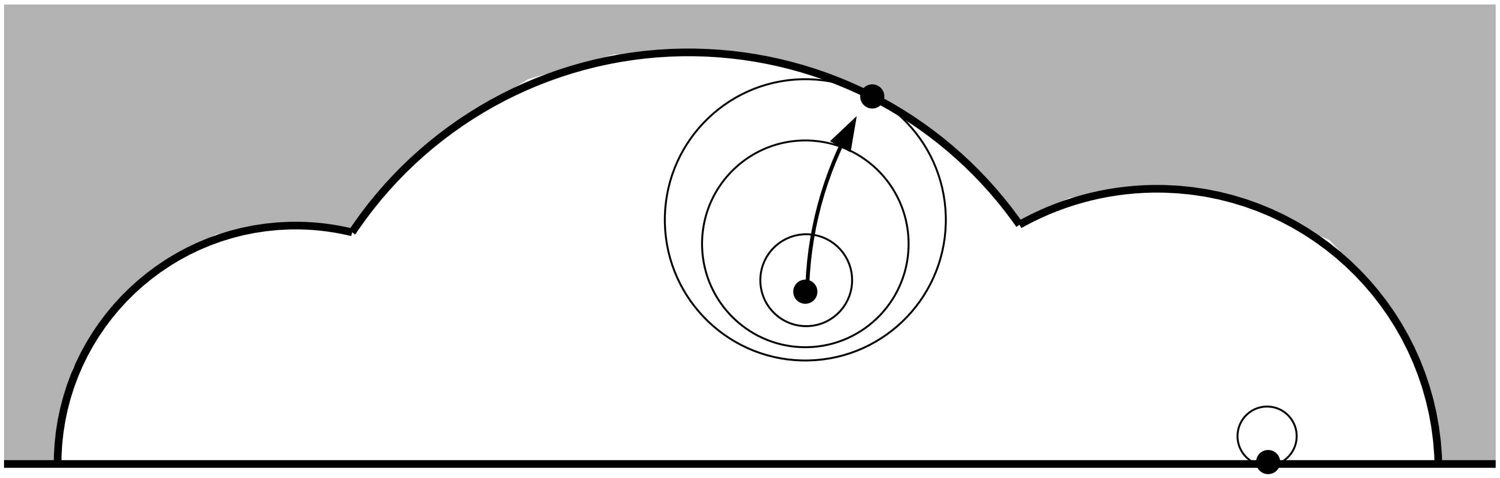

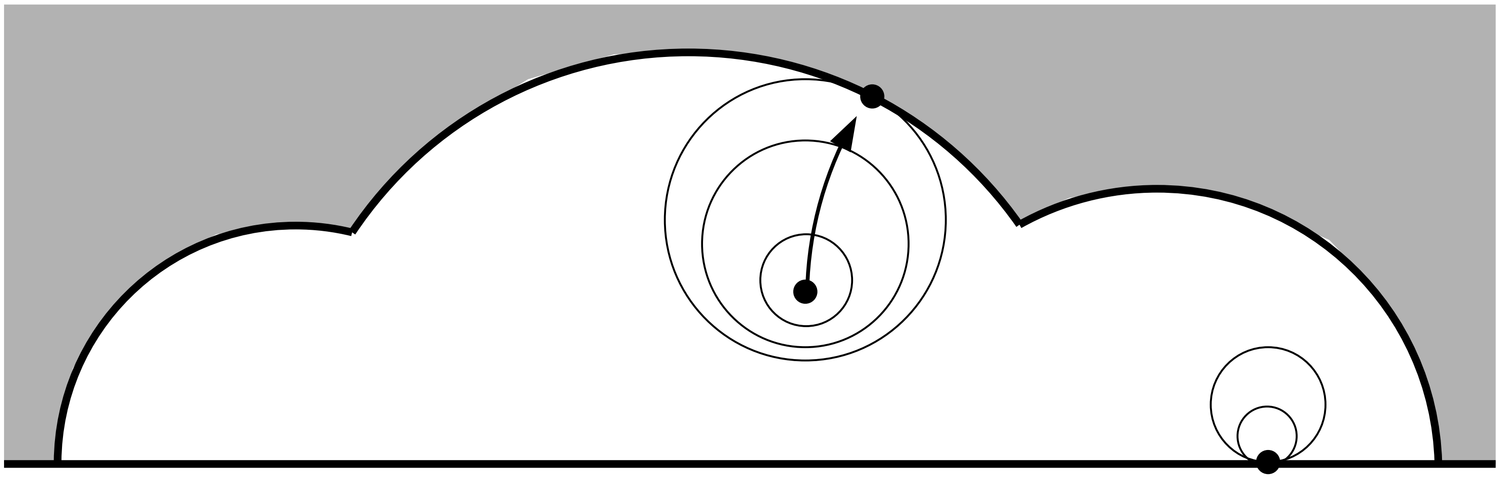

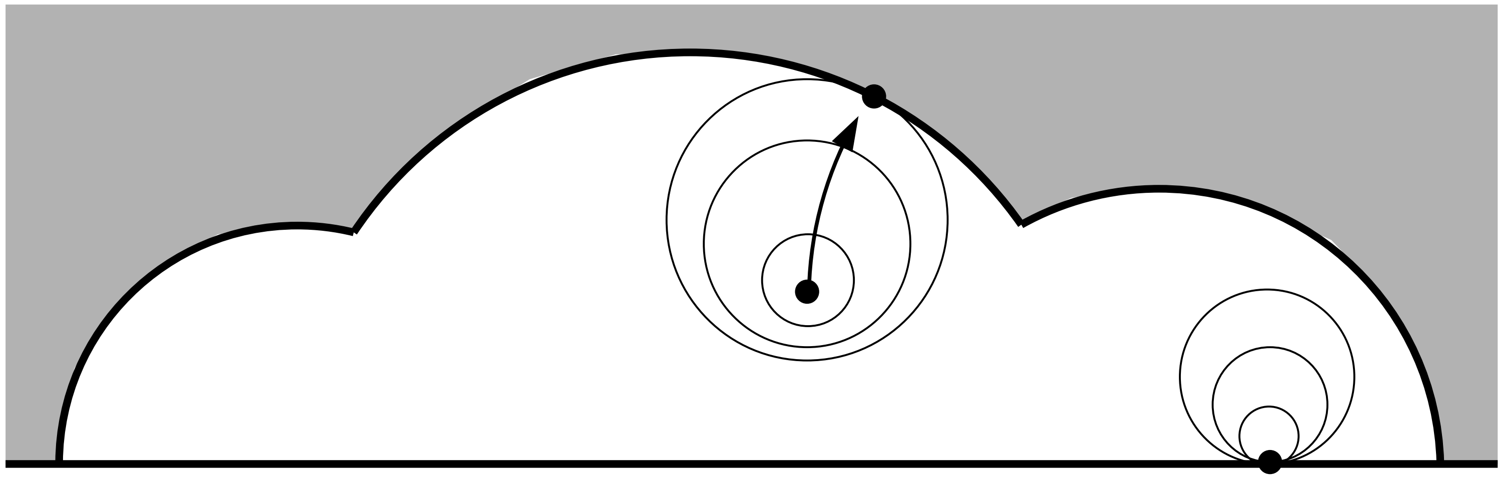

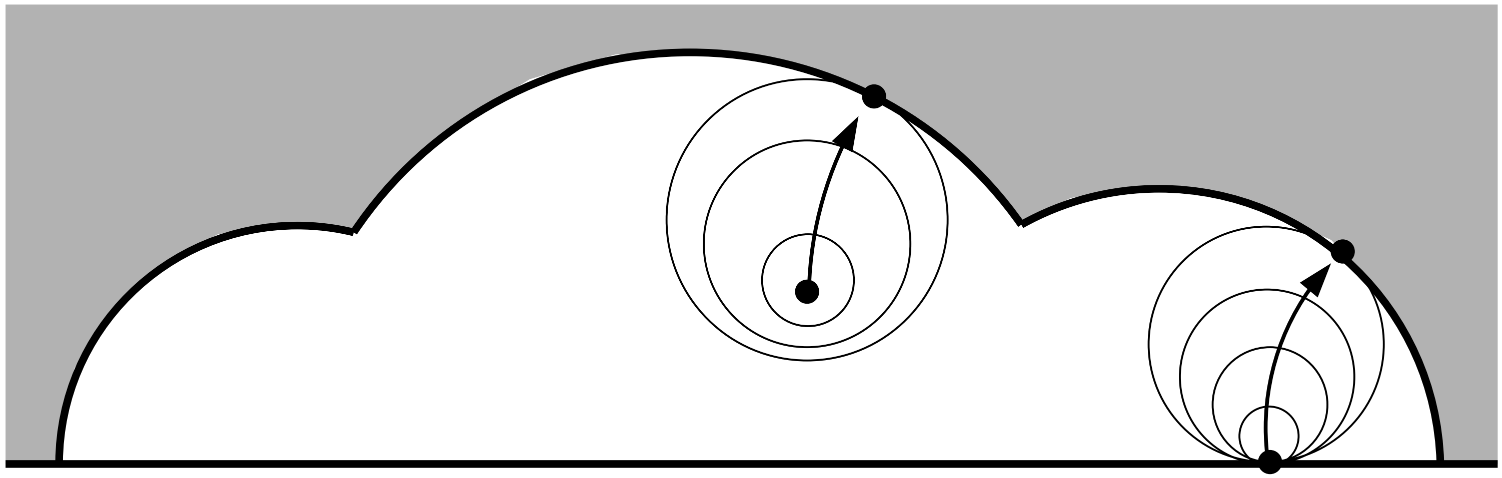

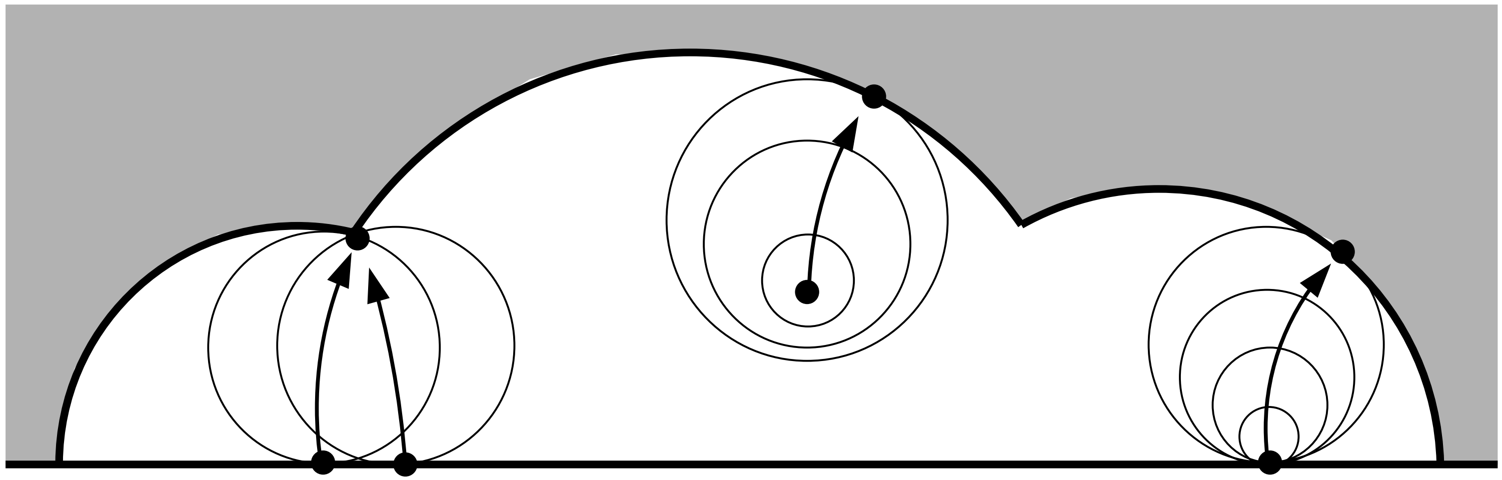

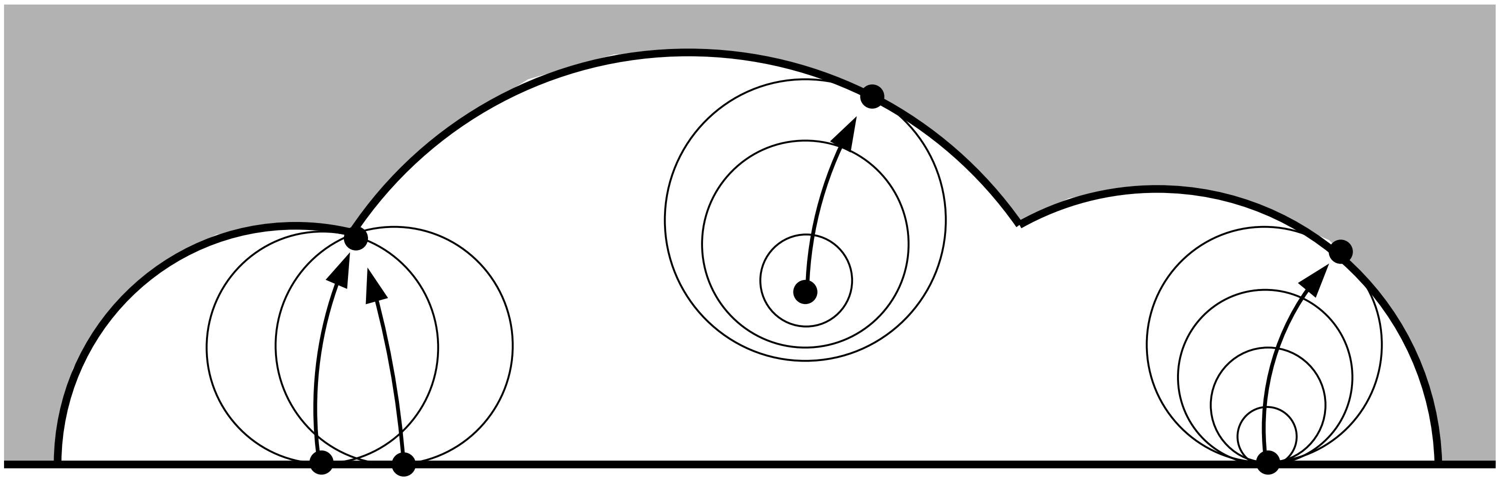

Nearest point retraction R is a **quasi-isometry**

$$\frac{1}{A} \leq \frac{\rho(R(x), R(y))}{\rho(x, y)} \leq A, \quad \text{if } \rho(x, y) \geq B.$$

i.e., R is bi-Lipschitz on large scales.

Metrics are hyperbolic metrics on Ω and S.

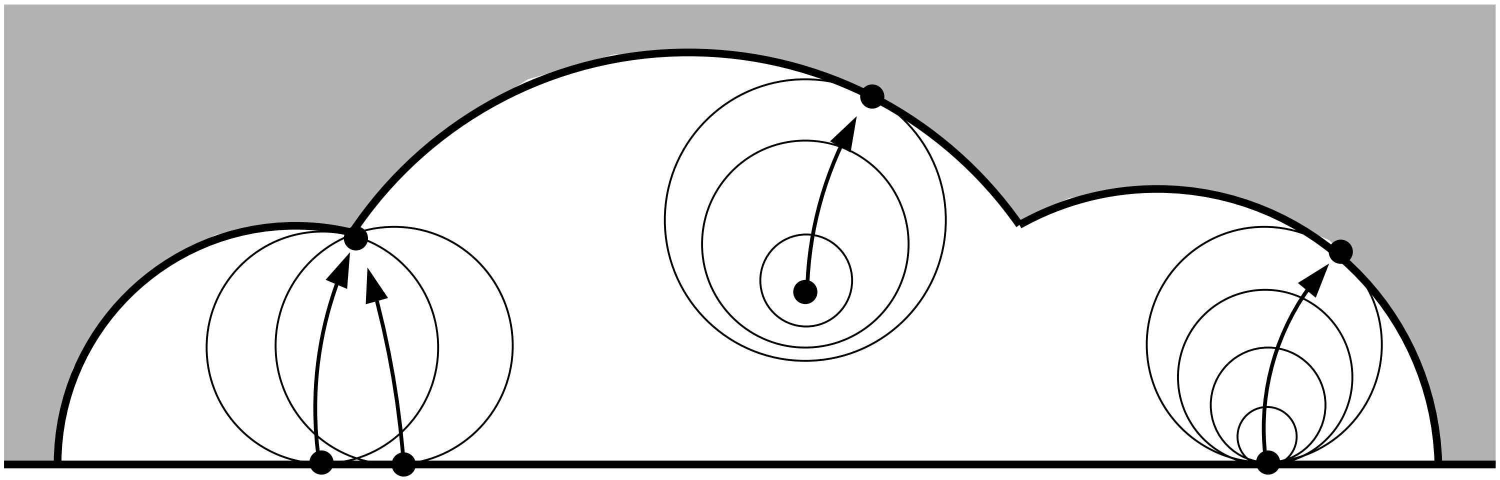

“Smoothing” gives QC map that fixes boundary points.

Hence medial axis map approximates Riemann map.

Dennis Sullivan, David Epstein and Al Marden, C.B.

The factorization theorem: Riemann map can be factored as $f = h \circ g$ where

- $g : \Omega \to \mathbb{D}$ is Lipschitz in Euclidean path metrics,
- $h : \mathbb{D} \to \mathbb{D}$ is biLipschitz in hyperbolic metric

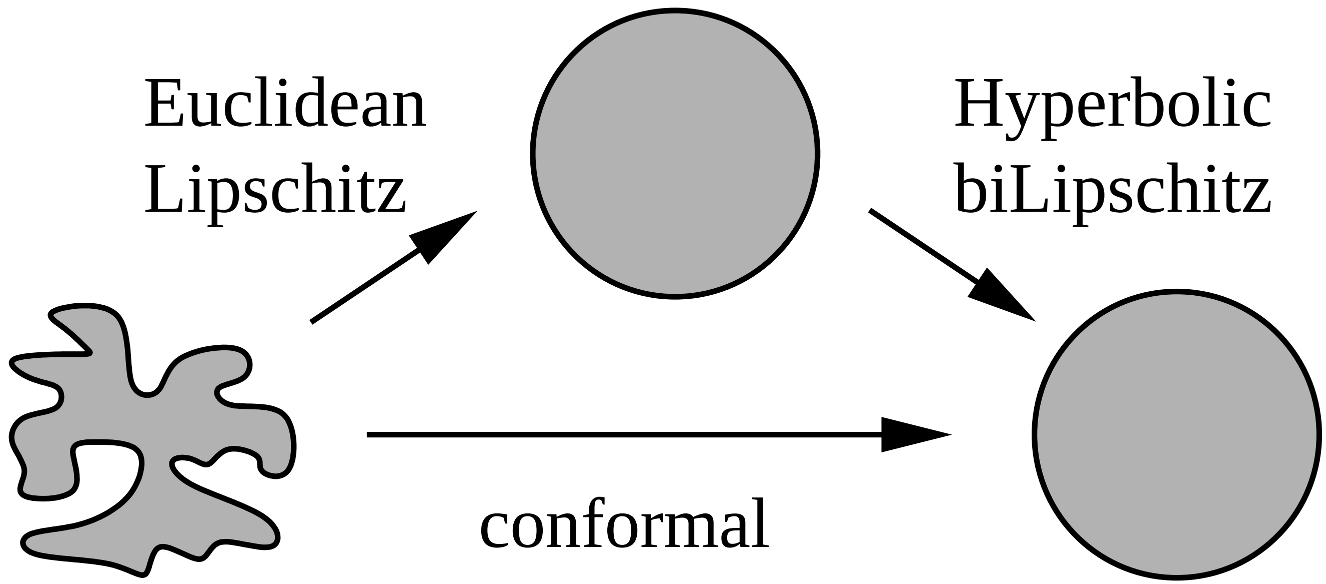

Cor: Any s.c. domain can be mapped 1-1, onto a disk by a contraction for the internal path metric.

Application of conformal mapping to CG: Quadrilateral meshes

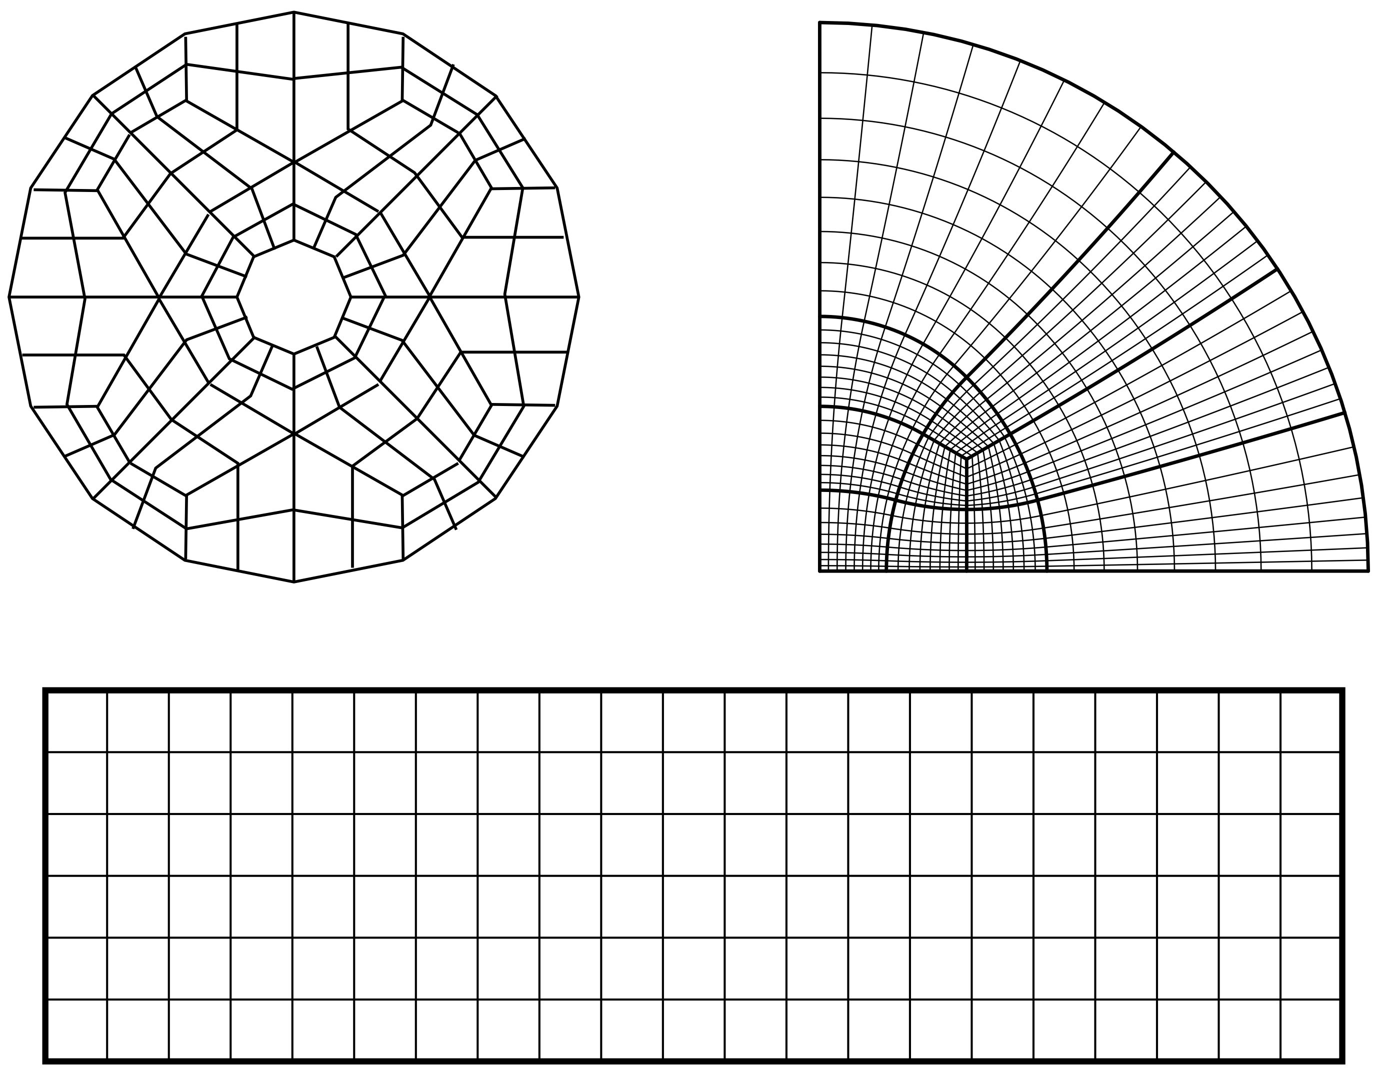

- n-gons have $O(n)$ quad mesh with angles $\leq 120^\circ$.

Bern and Eppstein, 2000. $O(n \log n)$ work.

- Regular hexagon shows 120° is sharp.

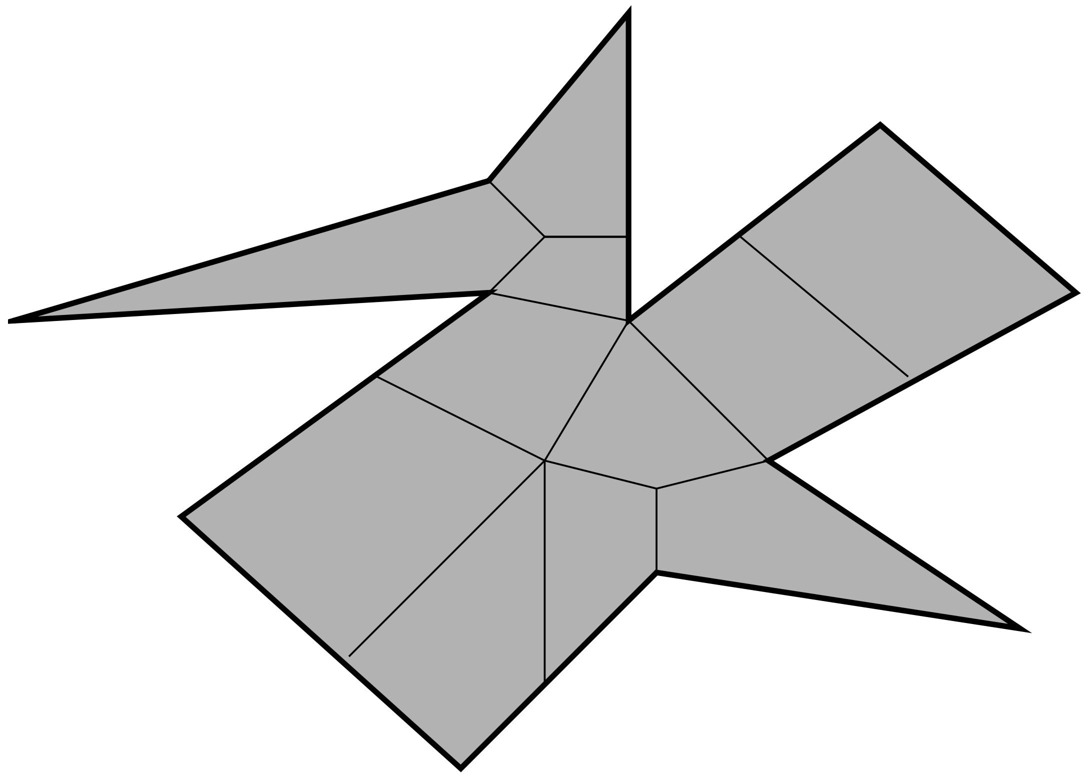

- n-gons have $O(n)$ quad mesh with angles $\leq 120^\circ$.

Bern and Eppstein, 2000. $O(n \log n)$ work.

- Regular hexagon shows 120° is sharp.

Can we bound angles from below?

- n-gons have $O(n)$ quad mesh with angles $\leq 120°$.

Bern and Eppstein, 2000. $O(n \log n)$ work.

- Regular hexagon shows $120°$ is sharp.

Can we bound angles from below?

Theorem: Every n-gon has $O(n)$ quad mesh with all angles $\leq 120°$ and new angles $\geq 60°$. $O(n)$ work.

Original angles $< 60°$ remain unchanged. $60°$ is sharp.

Idea of proof: transfer mesh from disk

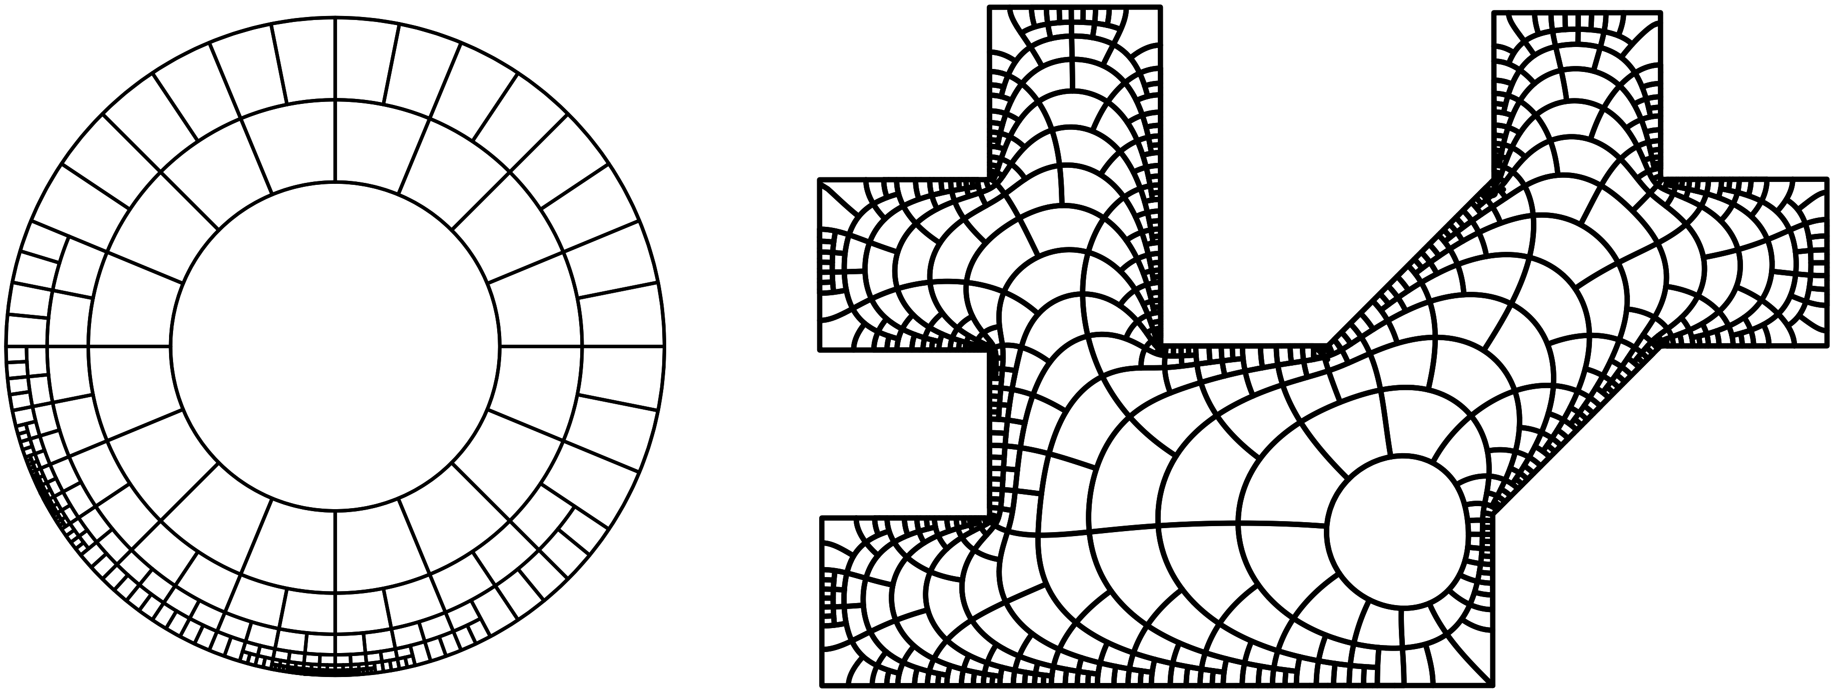

Idea of proof: transfer mesh from disk

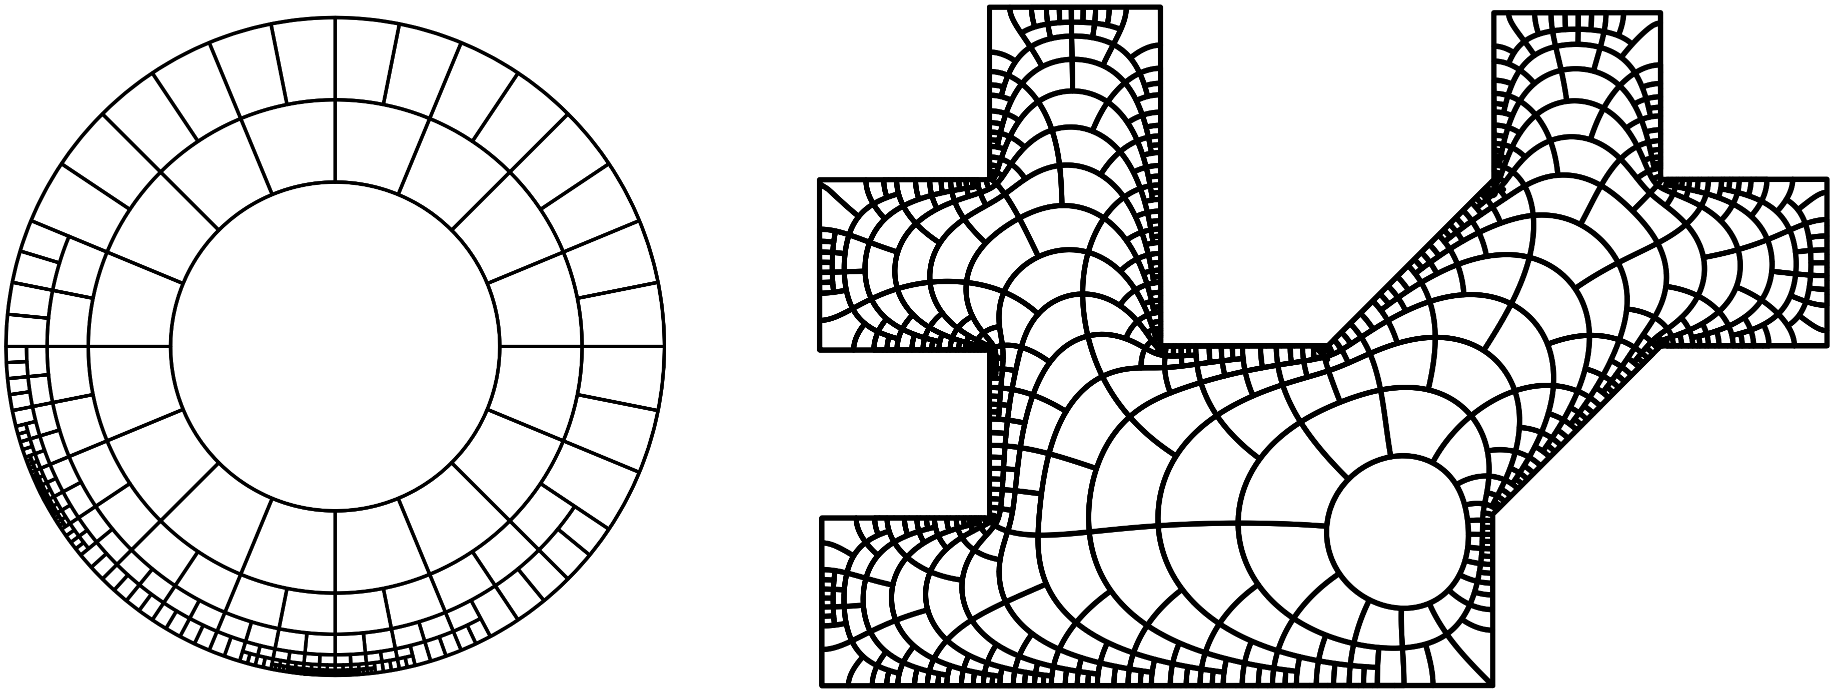

But number of boxes depends on geometry (not just n).

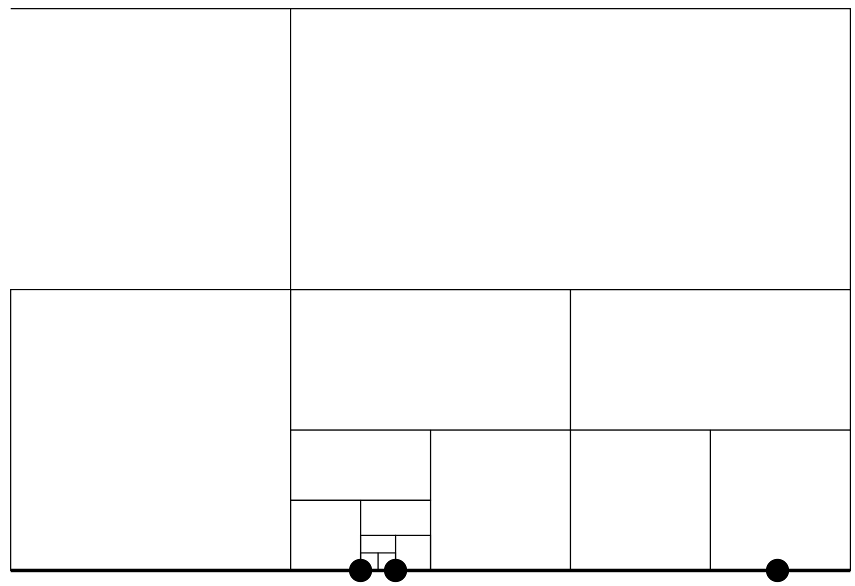

Subdivide until boxes separate points.

Gives $\gg n$ boxes if pre-vertices cluster tightly.

A **cluster** is a set $S \subset V$ so that

$$\operatorname{diam}(S) \ll \operatorname{dist}(S, V \setminus S).$$

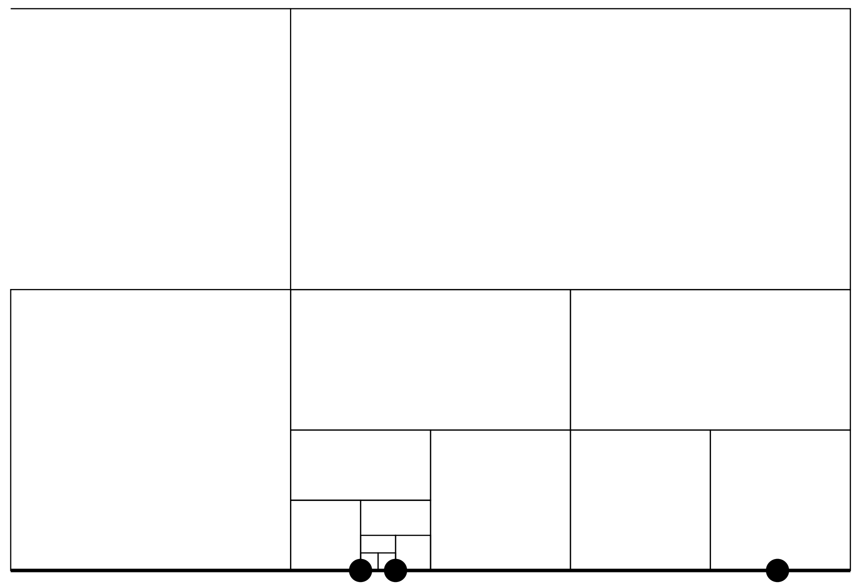

Subdivide until boxes separate points.

Gives $\gg n$ boxes if pre-vertices cluster tightly.

Fact: can find all clusters in $O(n)$ using medial axis.

“Cover” each cluster by a half-annulus.

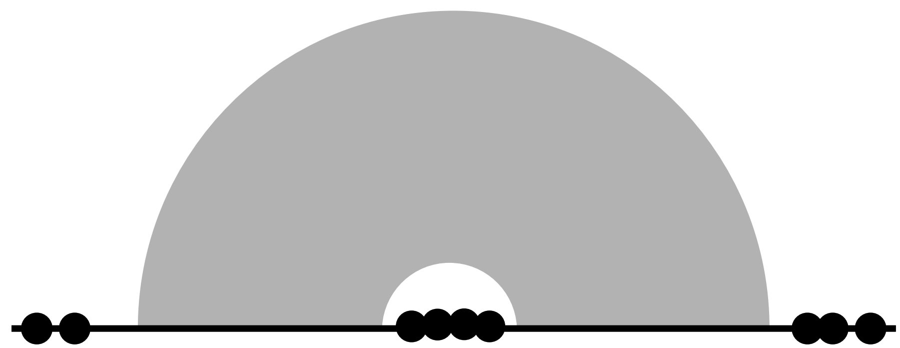

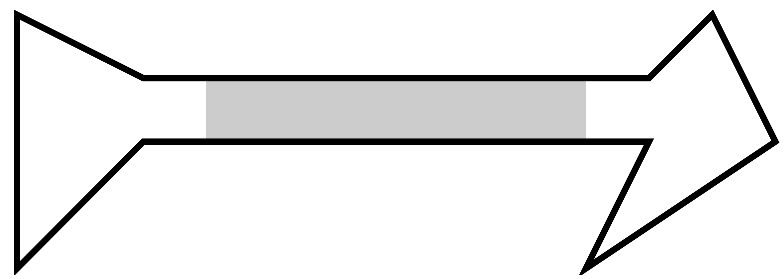

Maps to a **hyperbolic thin piece** dividing polygon.

Surface **thin part** is union of short non-trivial loops.

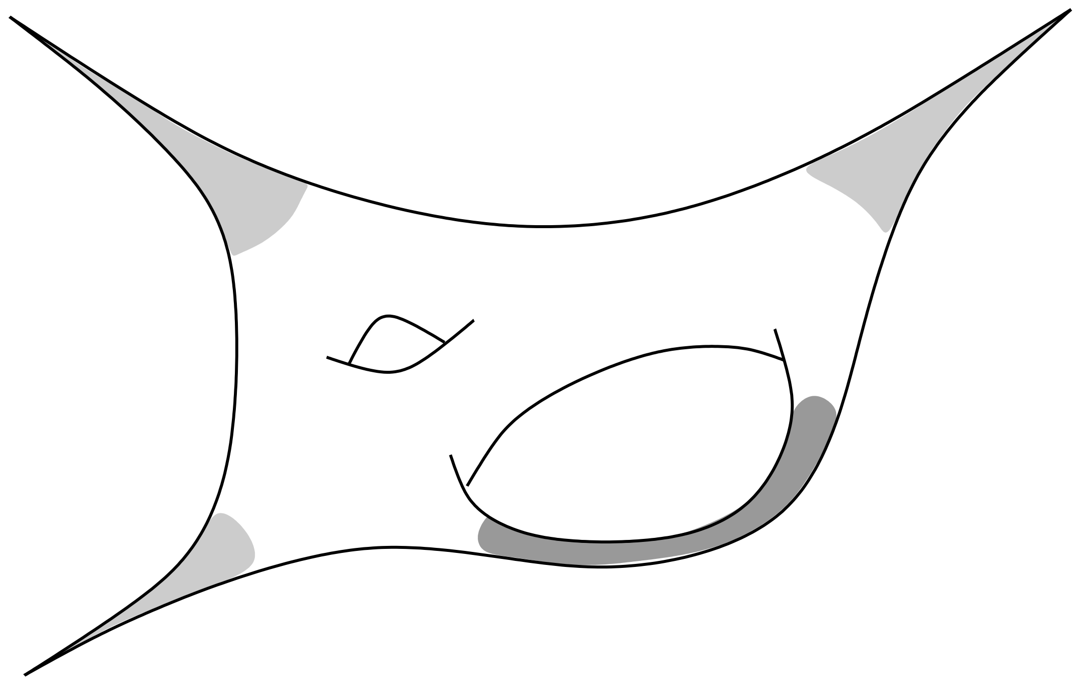

parabolic = puncture,   hyperbolic = handle

Thick and Thin parts of a polygon

Thin part is union of short curves between edges.

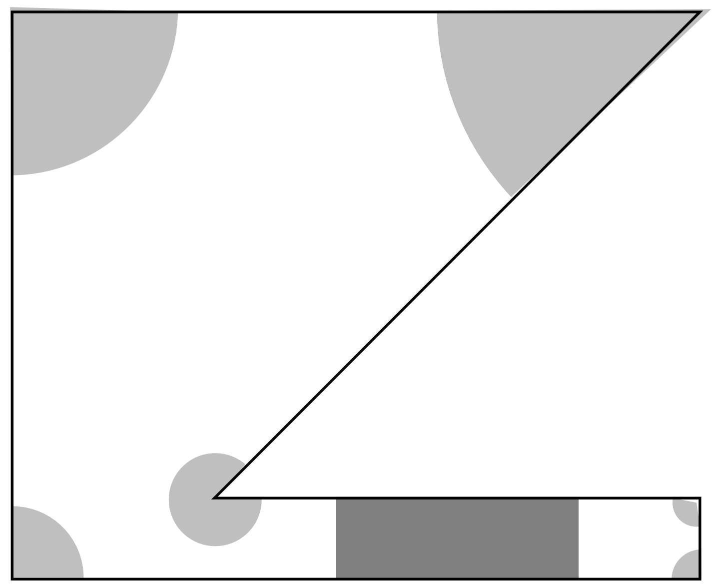

Parabolic = adjacent,   Hyperbolic = non-adjacent

Rough idea: sides I, J so $\mathrm{dist}(I, J) \ll \min(|I|, |J|)$.

Thick parts = remaining components (white)

More examples of hyperbolic thin parts.

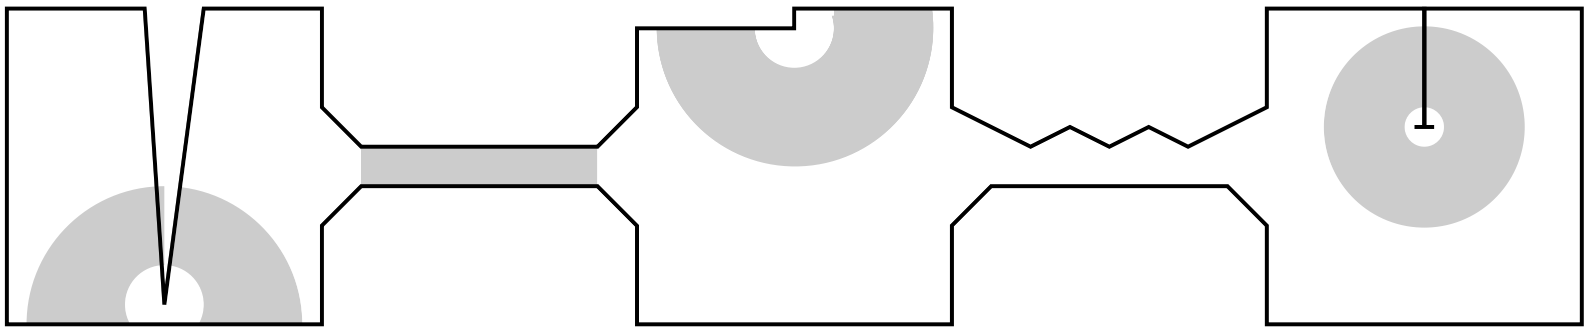

Inside thick regions (white) conformal pre-vertices are well separated on circle (no clusters).

Implies good estimates for conformal map.

Idea for quad mesh theorem:

• Decompose polygon into $O(n)$ thick and thin parts.

• Mesh thin parts "by hand"; explicit construction with $O(1)$ pieces each. Some pieces may be long and narrow.

• Transfer mesh on disk to thick parts by conformal map. $O(n)$ pieces. Every quad is "roundish".

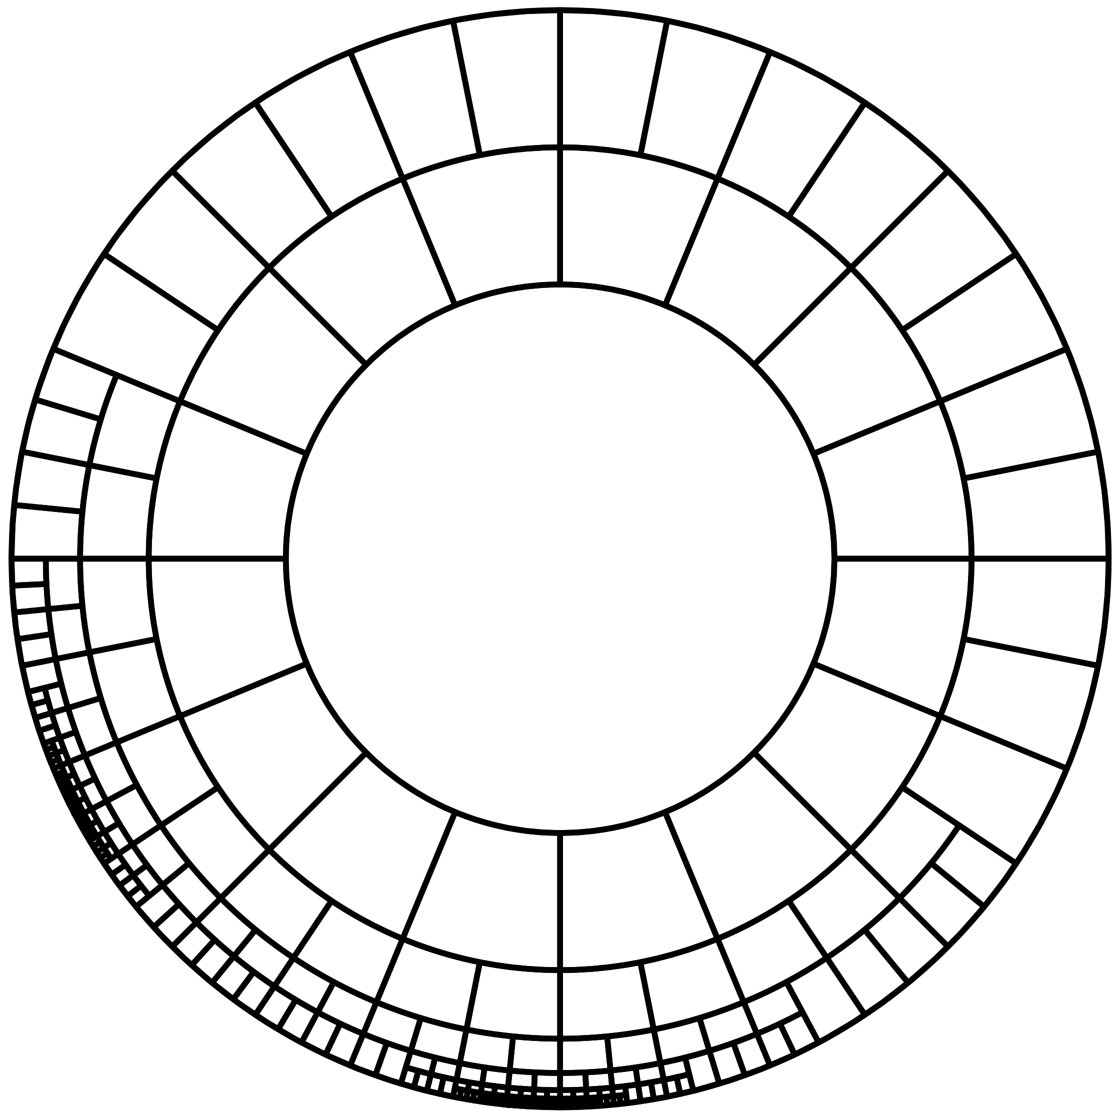

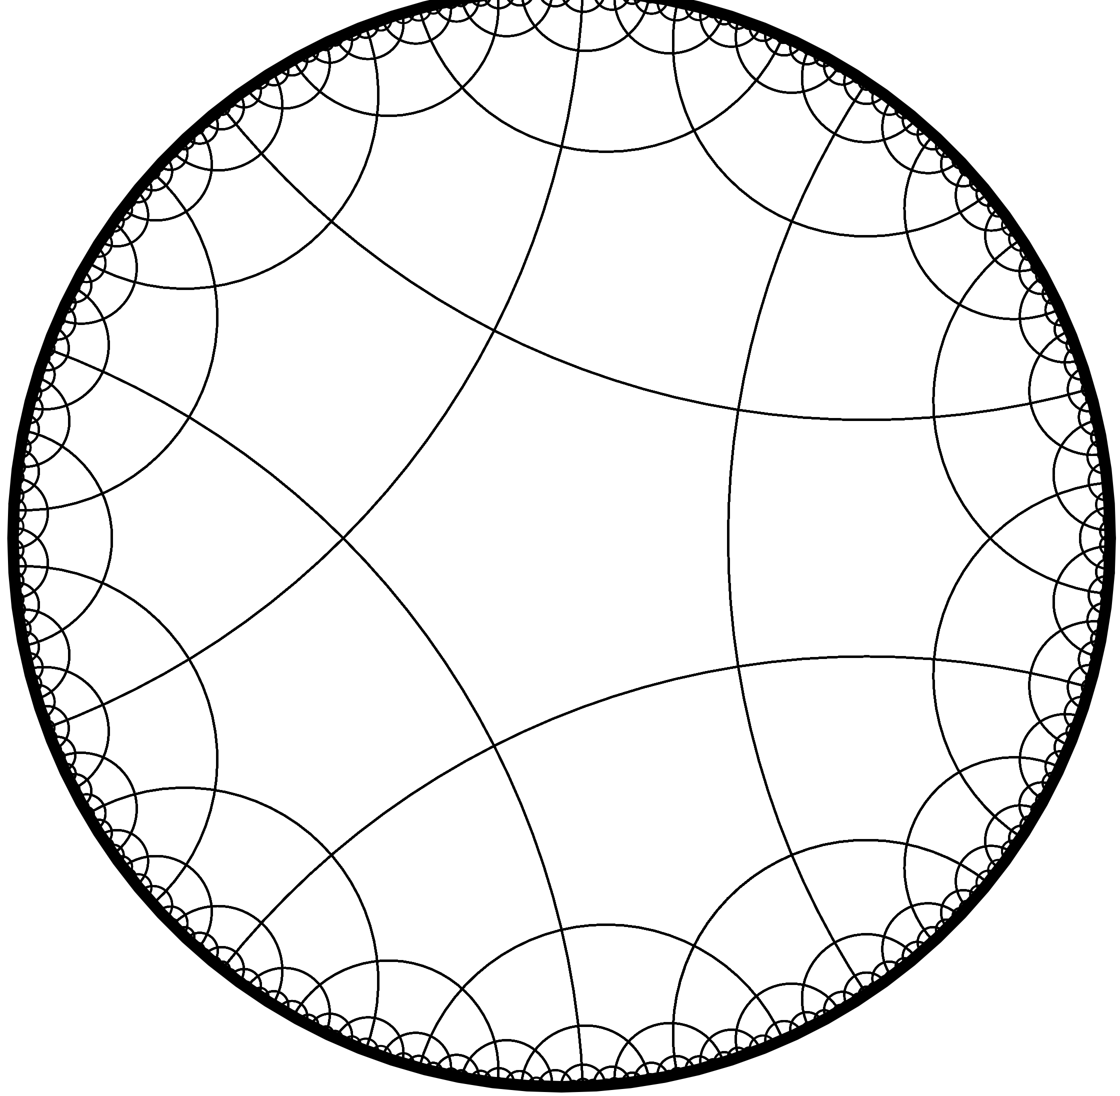

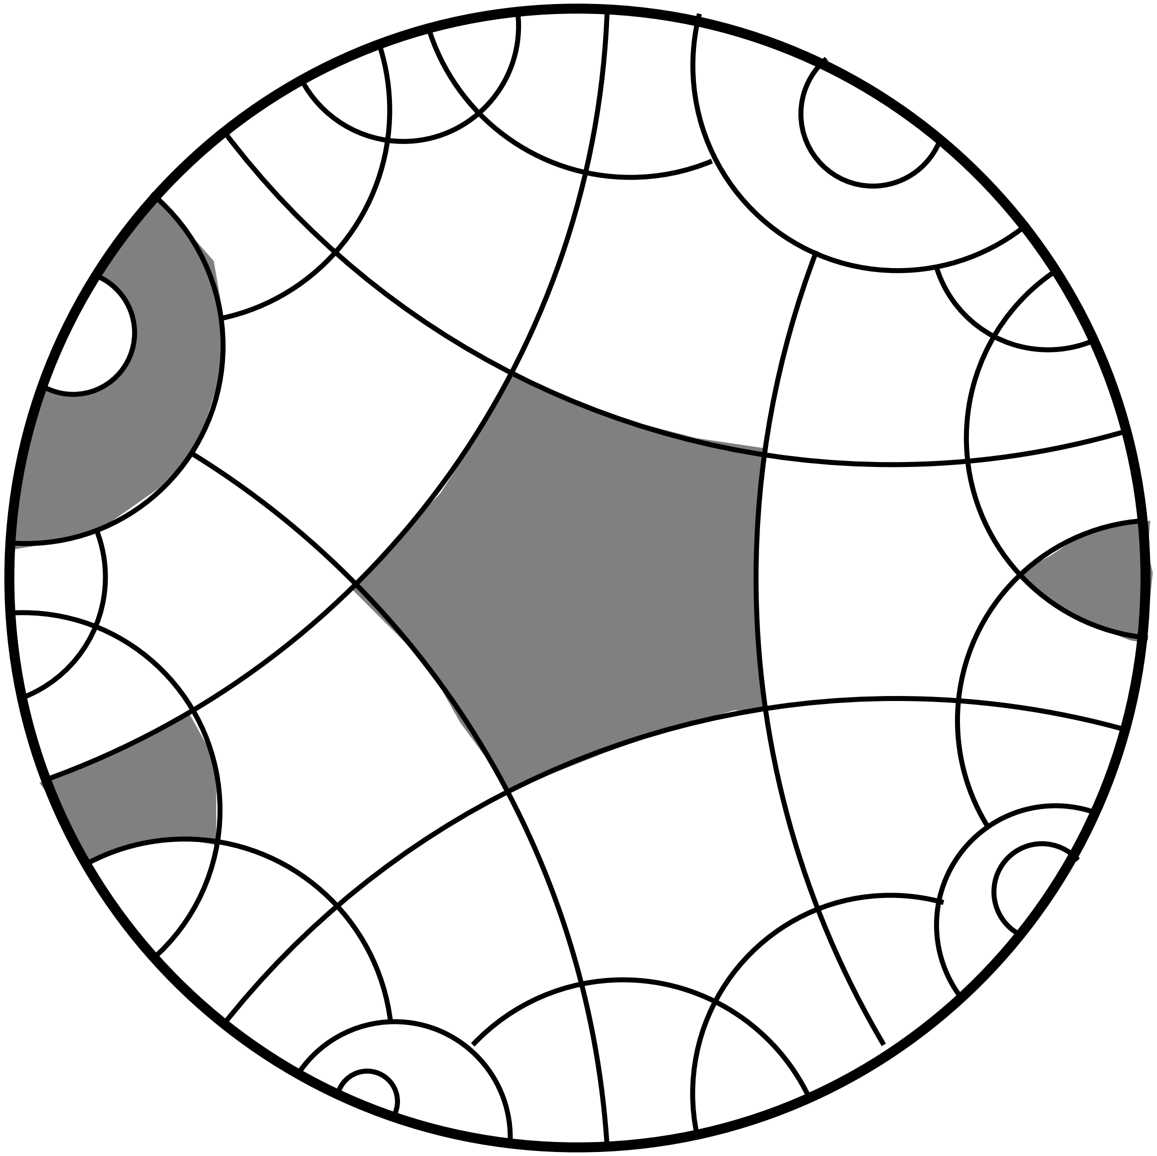

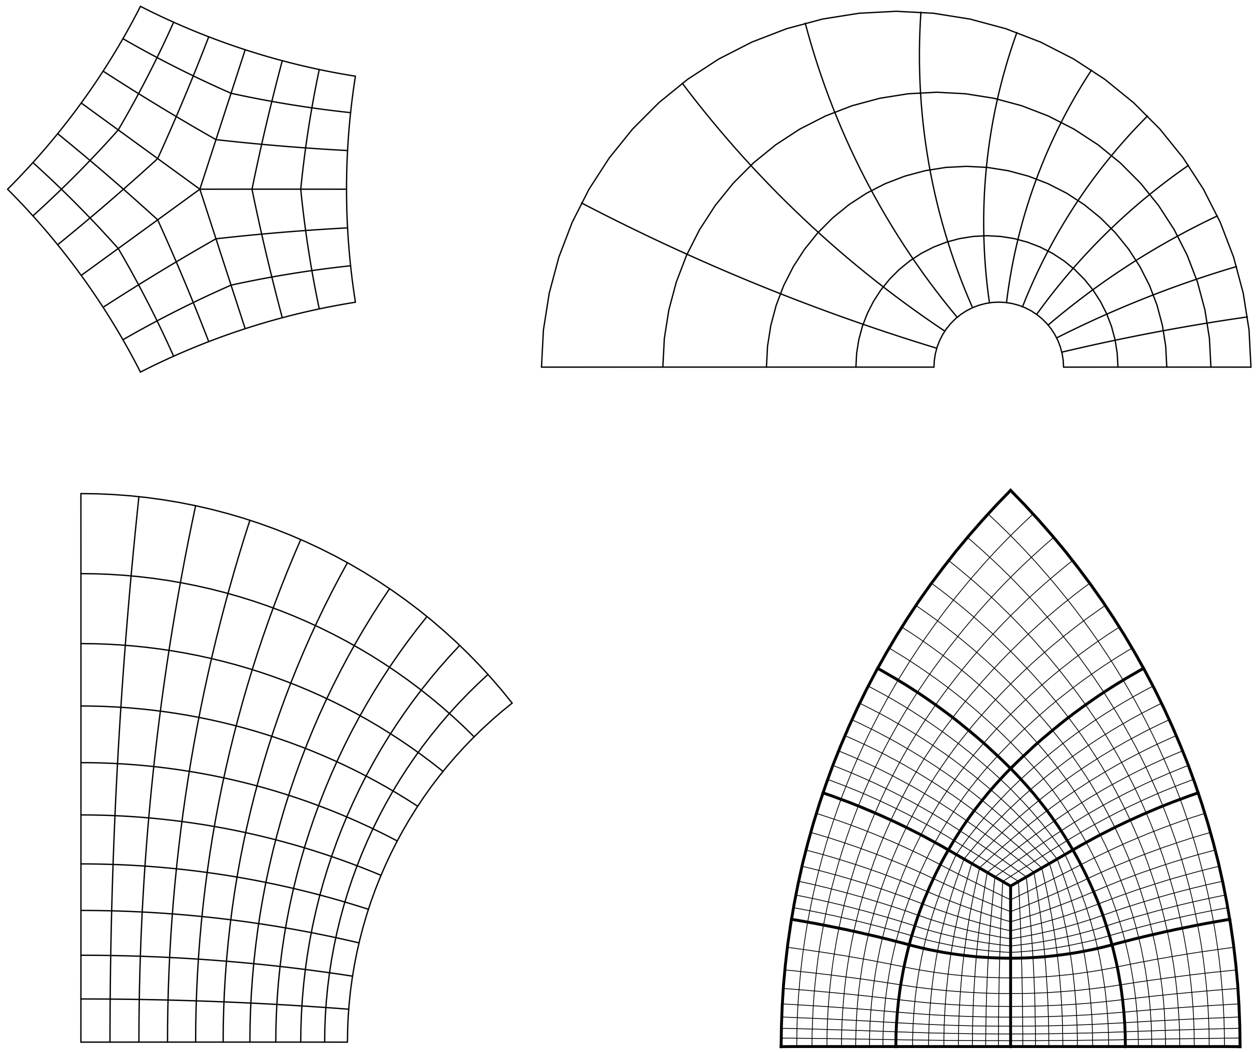

Meshes designed to match along common edges.

A **Planar Straight Line Graph** (PSLG) is a finite point set plus a set of disjoint edges between them.

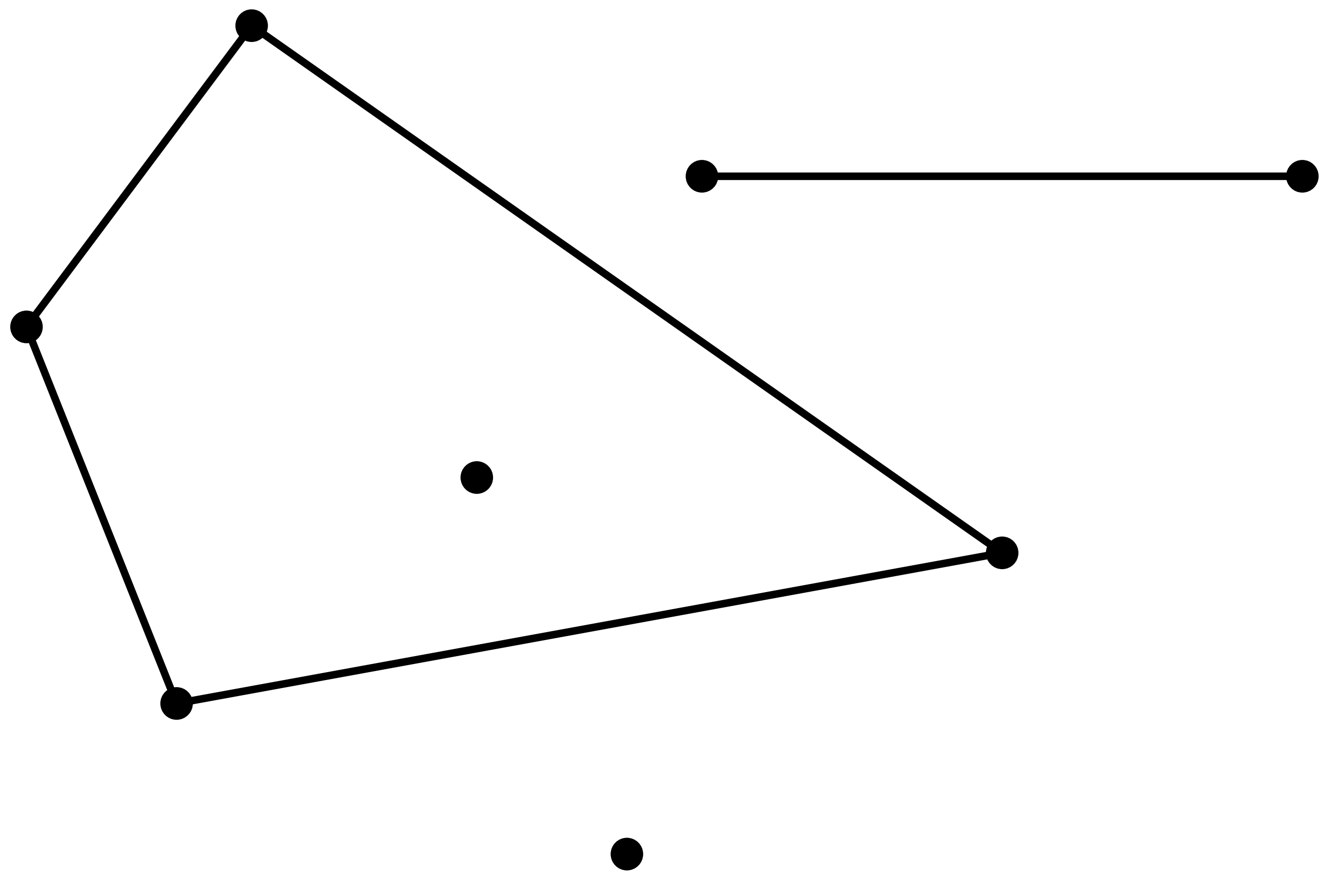

A **Planar Straight Line Graph** (PSLG) is a finite point set plus a set of disjoint edges between them.

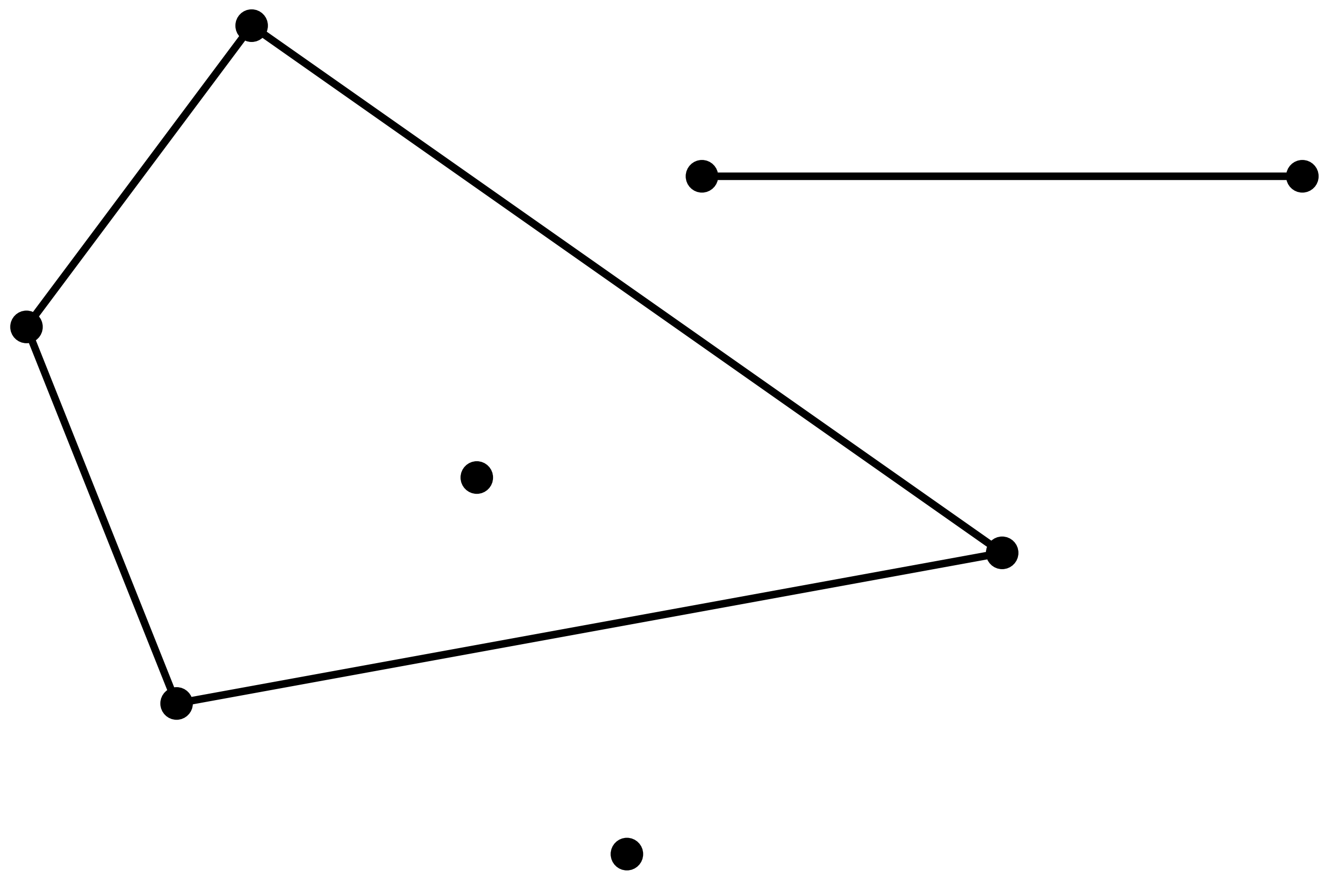

Size = number of vertices = n.

A **Planar Straight Line Graph** (PSLG) is a finite point set plus a set of disjoint edges between them.

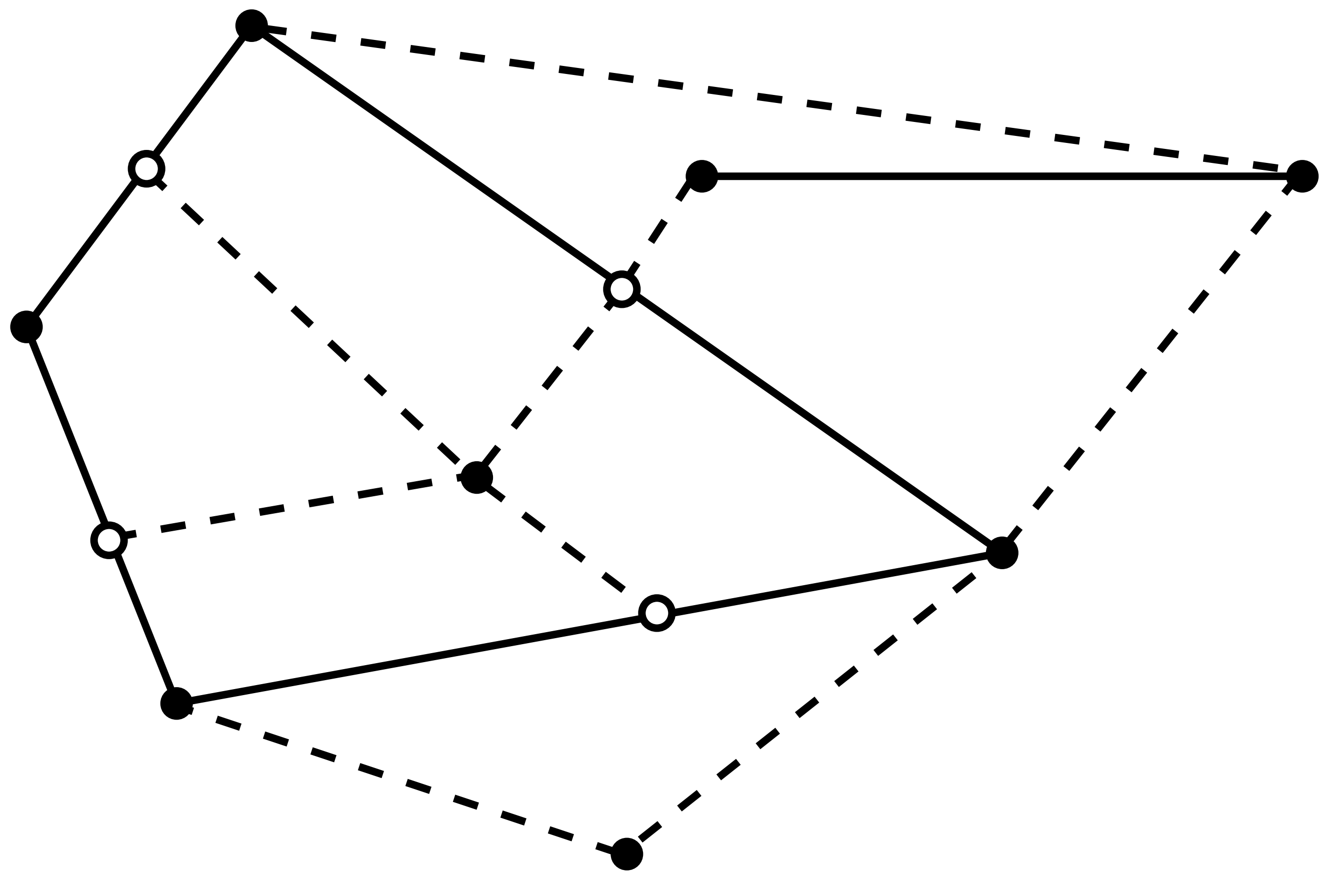

Mesh convex hull conforming to PSLG.

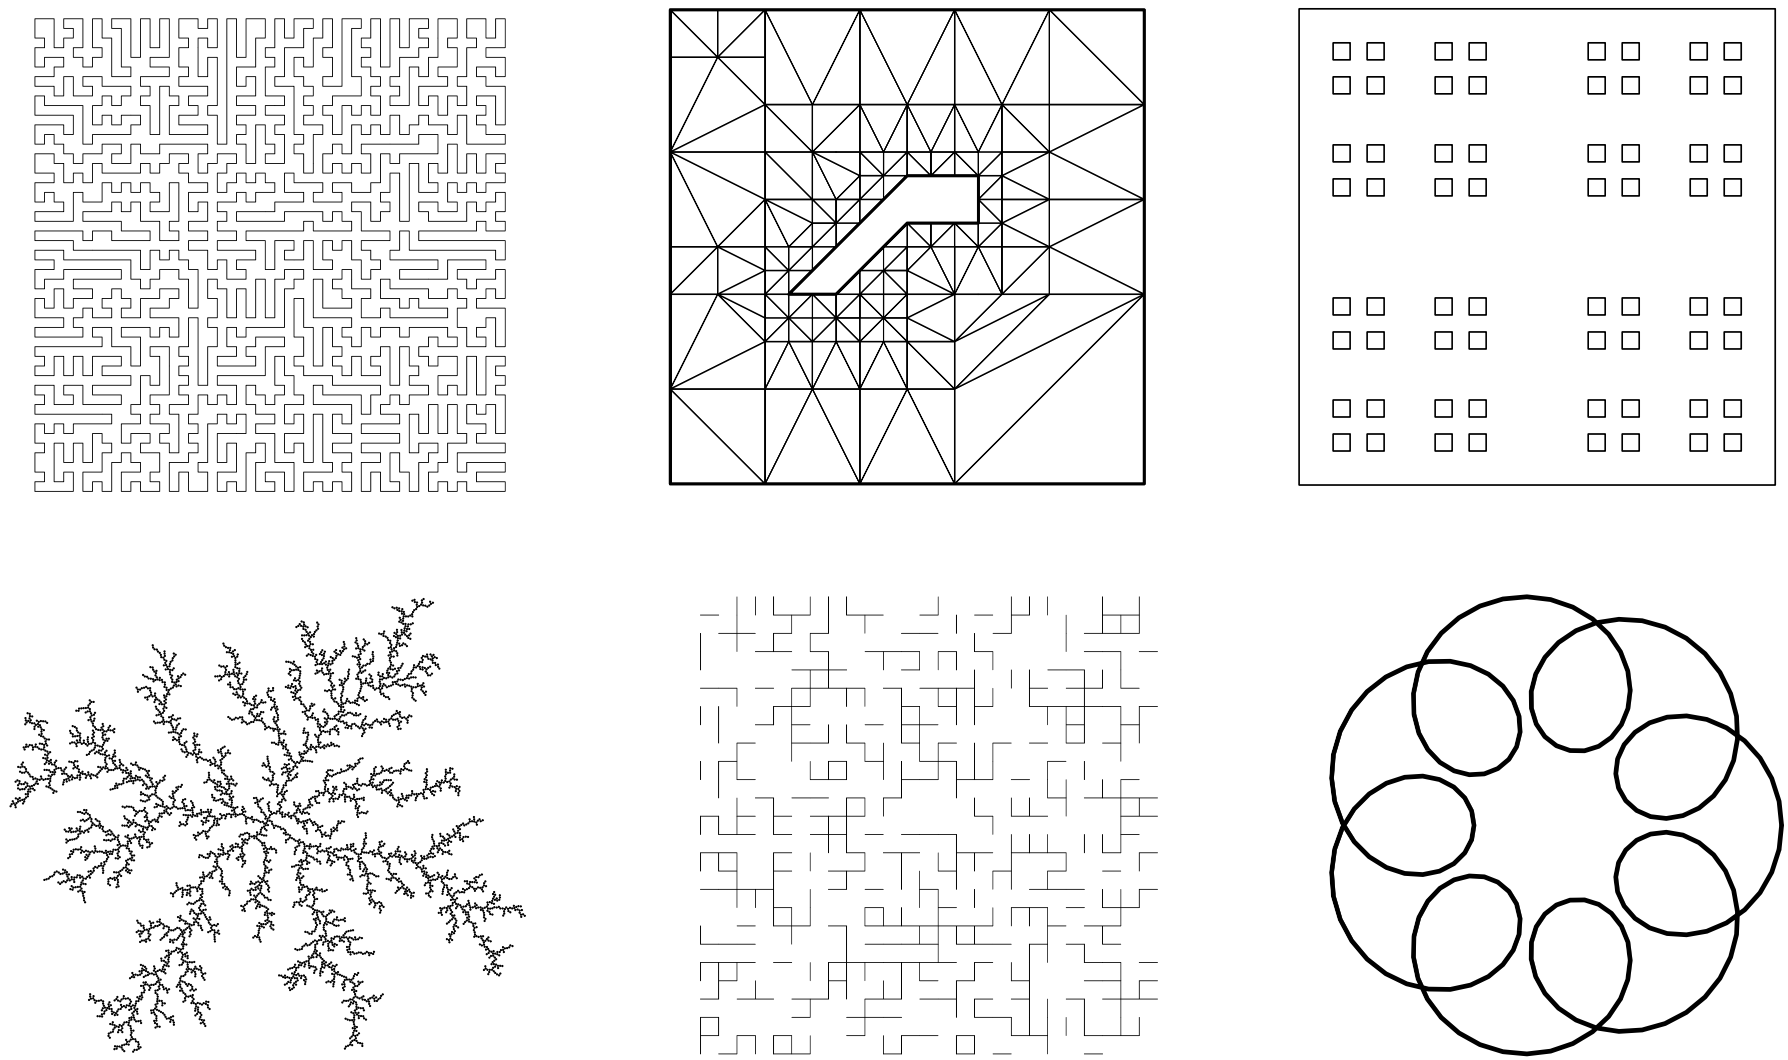

More PSLGs

Theorem: Every PSLG of size n has $O(n^2)$ quad mesh with all angles $\leq 120°$ and all new angles $\geq 60°$.

Theorem: Every PSLG of size n has $O(n^2)$ quad mesh with all angles $\leq 120^\circ$ and all new angles $\geq 60^\circ$.

Angles and complexity sharp.

All but $O(n/\epsilon)$ vertices have angles in $[90^\circ - \epsilon, 90^\circ + \epsilon]$.

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180°$ sometimes requires n^2 elements.

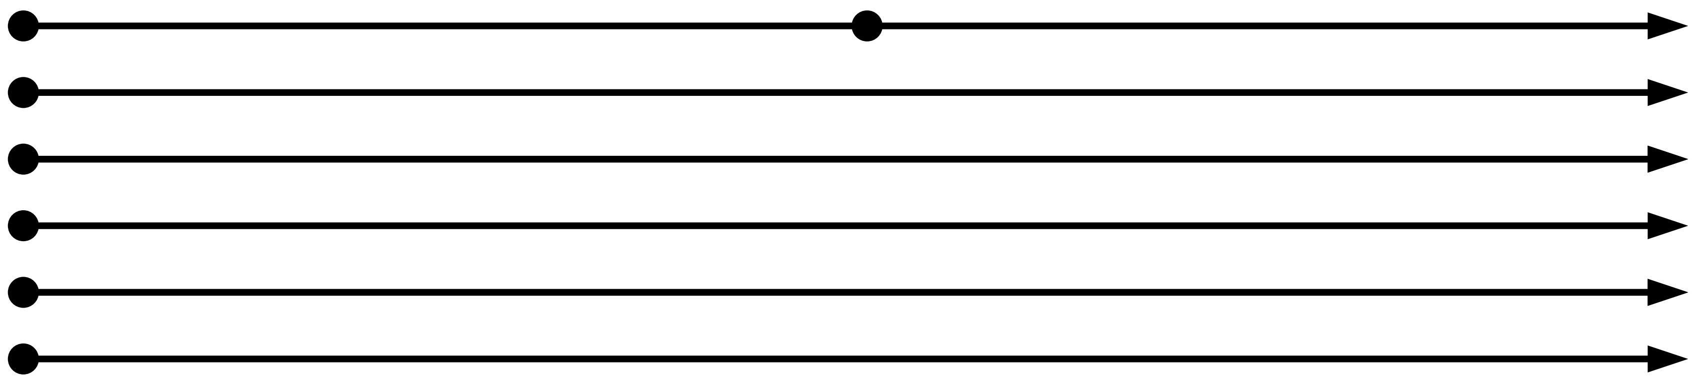

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

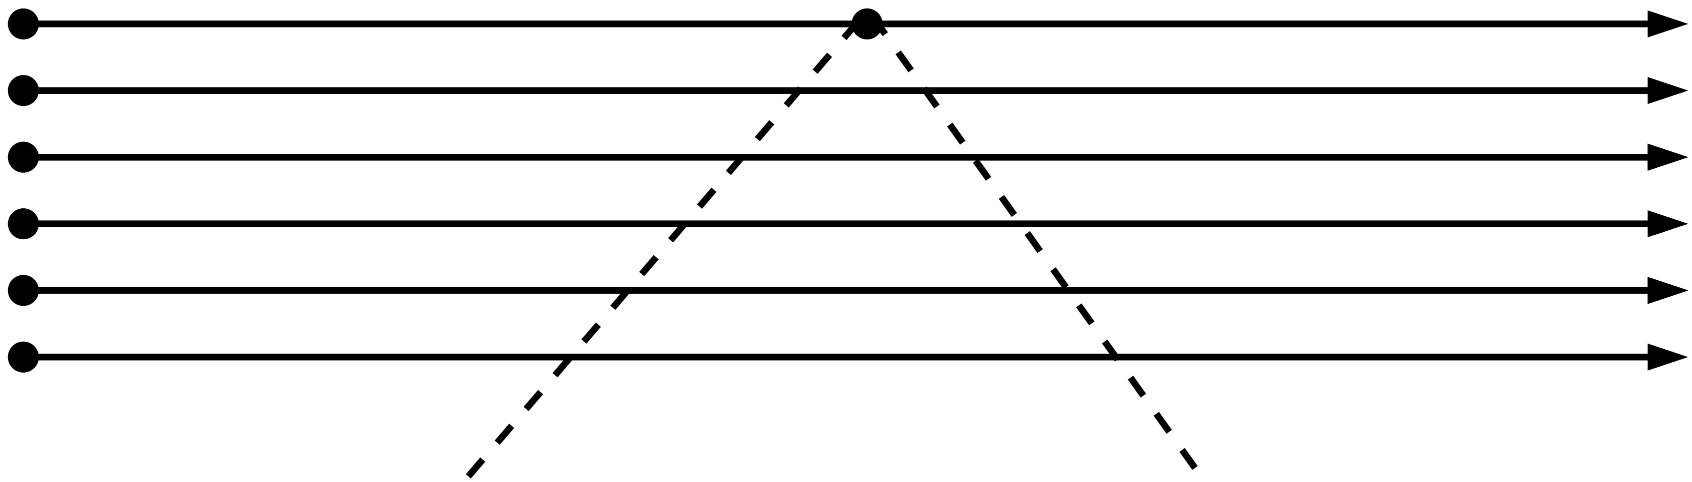

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

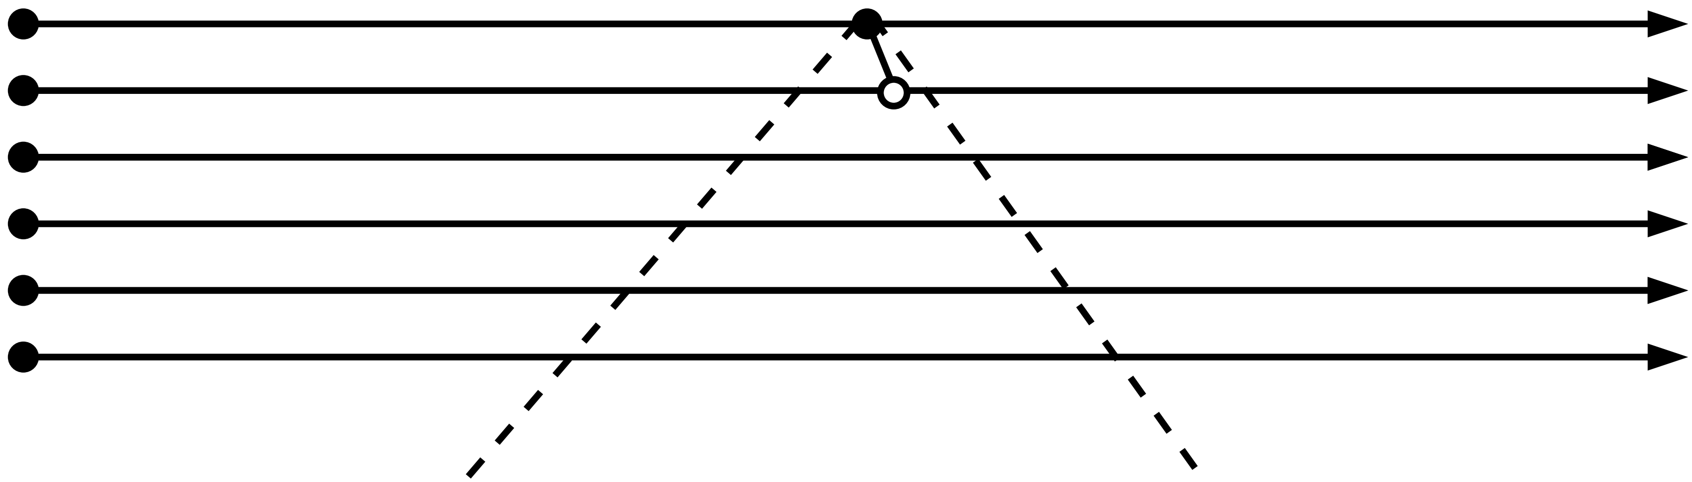

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

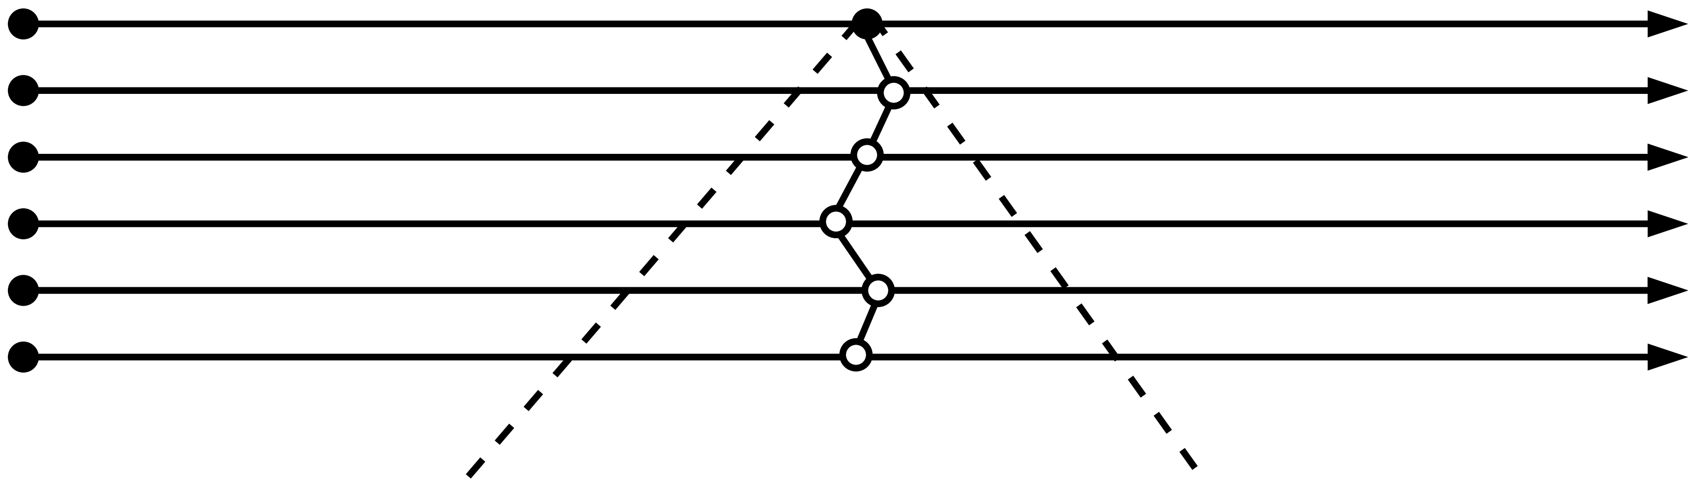

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180°$ sometimes requires n^2 elements.

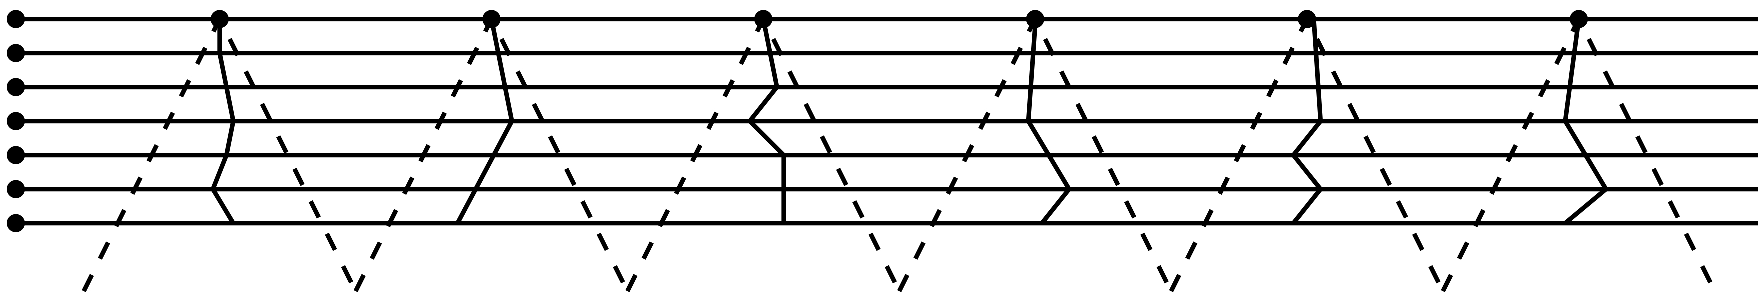

Convert quadrilaterals to triangles by adding diagonals.

Convert quadrilaterals to triangles by adding diagonals.

Corollary: Every PSLG has a $O(n^2)$ triangulation with maximum angle $\leq 120°$.

Convert quadrilaterals to triangles by adding diagonals.

Corollary: Every PSLG has a $O(n^2)$ triangulation with maximum angle $\leq 120°$.

Previous: S. Mitchell 1993 (157.5°), Tan 1996 (132°).

Can we improve the 120° upper bound?

Can we get a positive lower bound on angles?

Convert quadrilaterals to triangles by adding diagonals.

Corollary: Every PSLG has a $O(n^2)$ triangulation with maximum angle $\leq 120°$.

Previous: S. Mitchell 1993 (157.5°), Tan 1996 (132°).

Can we improve the 120° upper bound? **Yes**

Can we get a positive lower bound on angles? **No**

Convert quadrilaterals to triangles by adding diagonals.

Corollary: Every PSLG has a $O(n^2)$ triangulation with maximum angle $\leq 120°$.

Previous: S. Mitchell 1993 ($157.5°$), Tan 1996 ($132°$).

Can we improve the $120°$ upper bound? **Yes**

Can we get a positive lower bound on angles? **No**

(At least not with complexity depending only on n.)

No lower angle bound. For $1 \times R$ rectangle

$$\text{number of triangles} \gtrsim R \times (\text{smallest angle})$$

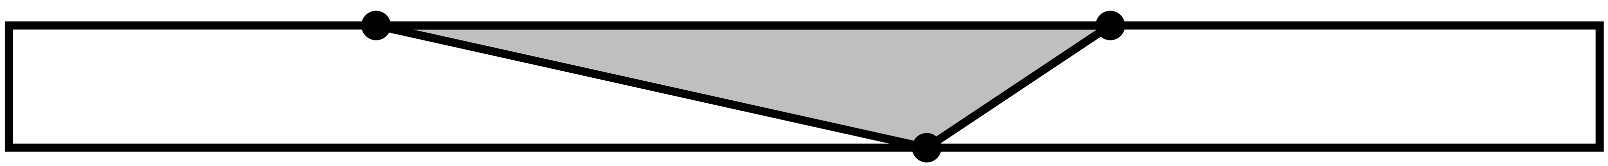

So, bounded complexity $\Rightarrow$ no lower angle bound.

Upper bound $< 90^\circ$ implies lower bound > 0:

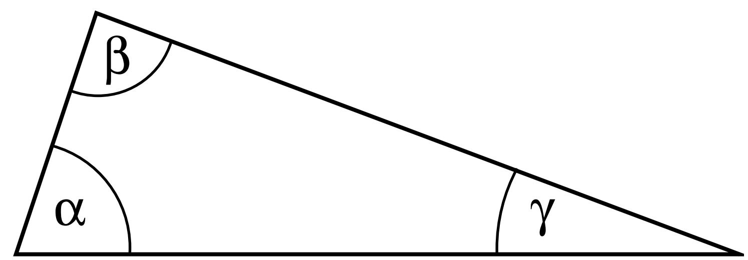

$$\alpha, \beta < (90^\circ - \epsilon) \Rightarrow \gamma = 180^\circ - \alpha - \beta \geq 2\epsilon.$$

Upper bound $< 90^\circ$ implies lower bound > 0:

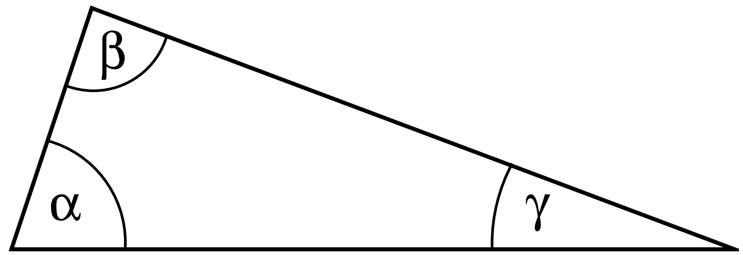

$$\alpha, \beta < (90^\circ - \epsilon) \Rightarrow \gamma = 180^\circ - \alpha - \beta \geq 2\epsilon.$$

So 90° is best uniform upper bound we can hope for.

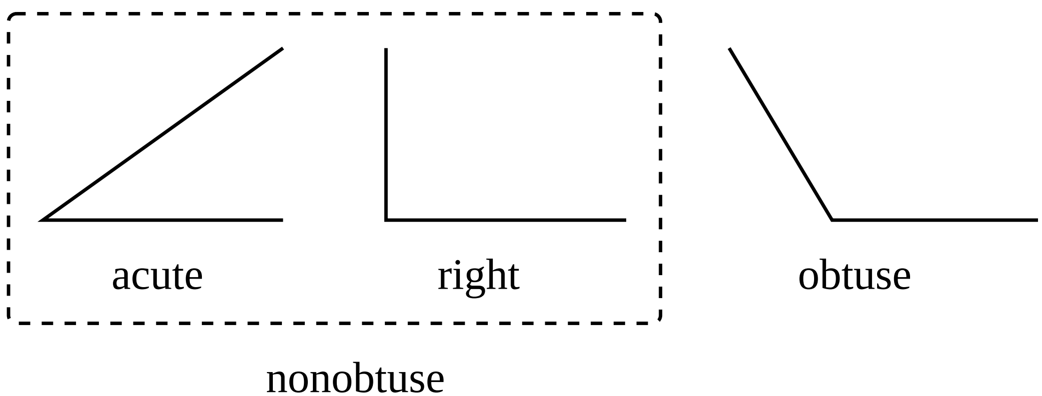

Is NOT="non-obtuse triangulation" possible?

Upper bound $< 90^\circ$ implies lower bound > 0:

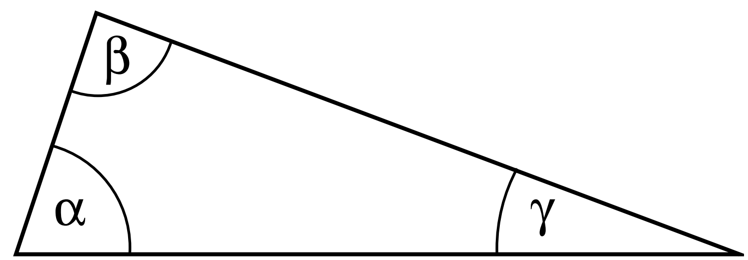

$$\alpha, \beta < (90^\circ - \epsilon) \Rightarrow \gamma = 180^\circ - \alpha - \beta \geq 2\epsilon.$$

So 90° is best uniform upper bound we can hope for.

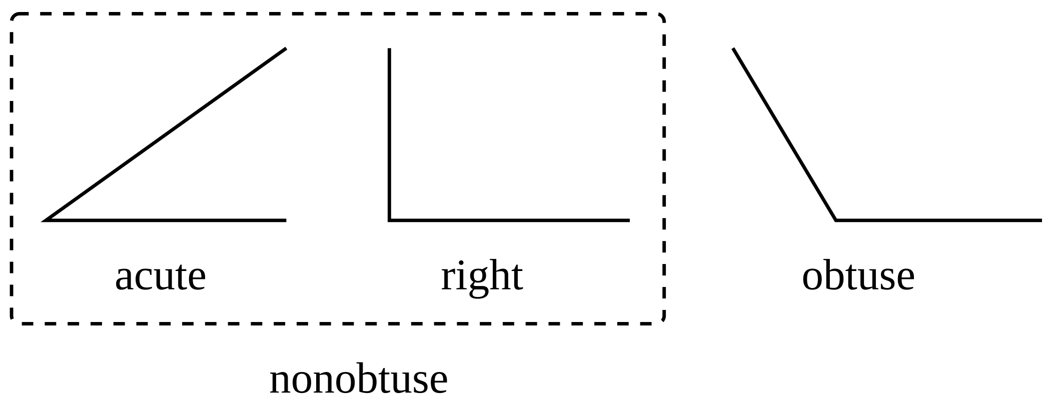

Is NOT="non-obtuse triangulation" possible? **Yes**

Brief history of NOTs:

- Always possible: Burago, Zalgaller 1960.
- Rediscovered: Baker, Grosse, Rafferty, 1988.
- $O(n)$ for points sets: Bern, Eppstein, Gilbert 1990
- $O(n^2)$ for polygons: Bern, Eppstein 1991
- $O(n)$ for polygons: Bern, Mitchell, Ruppert 1994

Numerous applications, heuristics.

Applications of nonobtuse triangulations:

- Discrete maximum principle (Ciarlet, Raviart, 1973)
- Convergence of finite element methods (Vavasis, 1996)
- Fast marching method (Sethian, 1999)
- Meshing space-time (Ungör, Sheffer, 2002)
- Machine learning (Salzberg, et. al., 1995)

Special case: triangulate of both sides of Γ (e.g., interface between two different materials).

Obvious strategy: mesh one side then other. But must re-mesh to include new vertices, ...

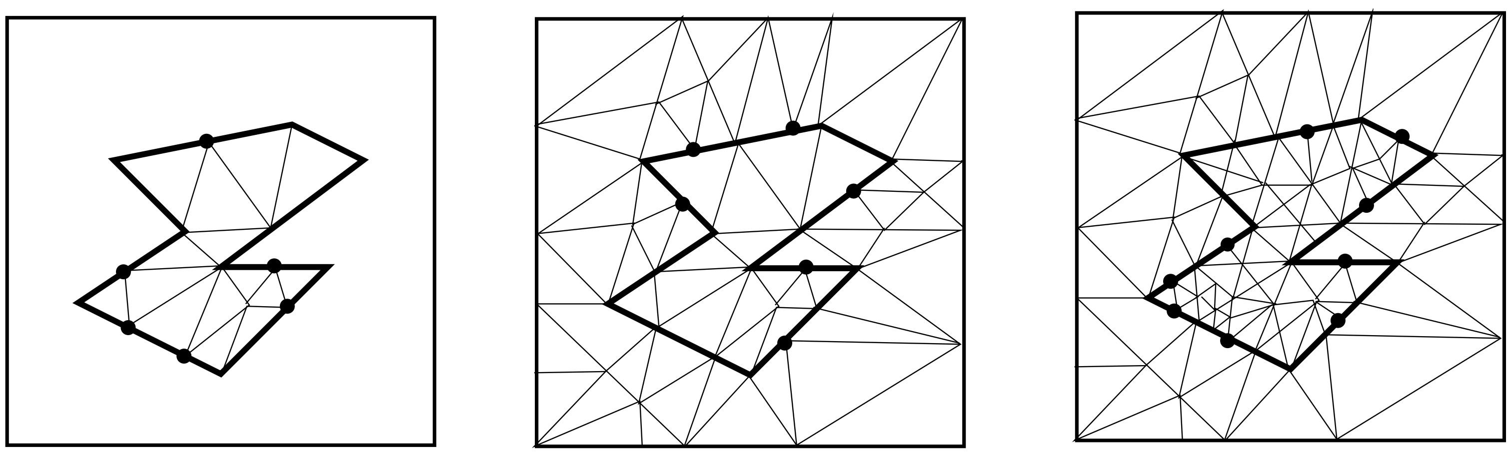

Does a polynomial number of points suffice?

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

$O(n^2)$ is best known lower bound, so a gap remains.

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

$O(n^2)$ is best known lower bound, so a gap remains.

Equivalent: Every planar triangulation has $O(n^{2.5})$ non-obtuse refinement.

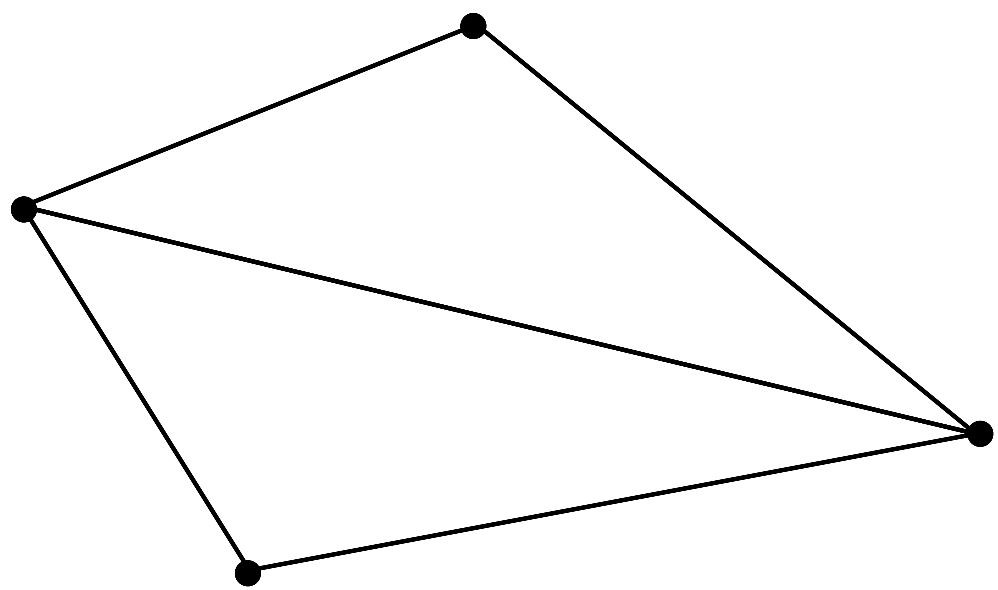

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

$O(n^2)$ is best known lower bound, so a gap remains.

Equivalent: Every planar triangulation has $O(n^{2.5})$ non-obtuse refinement.

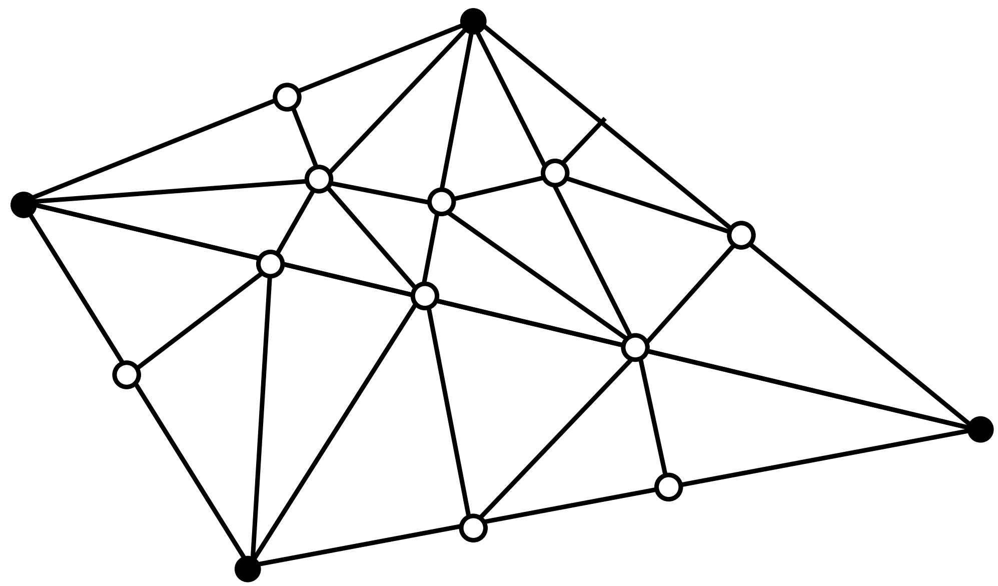

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

$O(n^2)$ is best known lower bound, so a gap remains.

Equivalent: Every planar triangulation has $O(n^{2.5})$ non-obtuse refinement.

I will sketch proof of an even more special case:

Theorem: Every triangulation of a simple n-gon has a nonobtuse refinement with $O(n^2)$ elements.

This improves $O(n^4)$ result of Bern and Eppstein (1992).

The segment $[v, w]$ is a **Gabriel** edge if it is the diameter of a disk containing no other points of V.

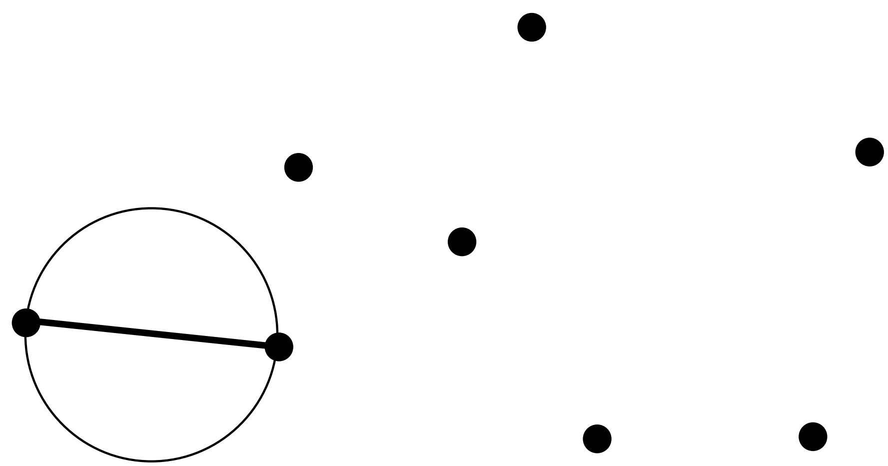

Gabriel edge.

The segment $[v, w]$ is a **Gabriel** edge if it is the diameter of a disk containing no other points of V.

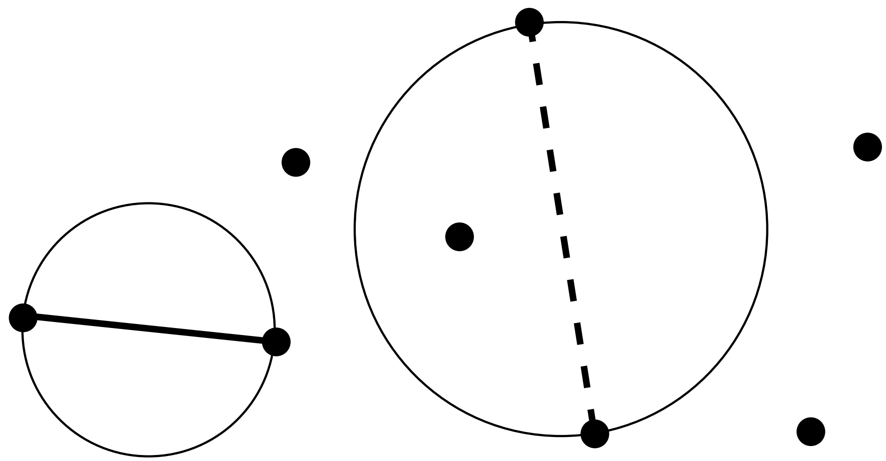

Not a Gabriel edge.

The segment $[v, w]$ is a **Gabriel** edge if it is the diameter of a disk containing no other points of V.

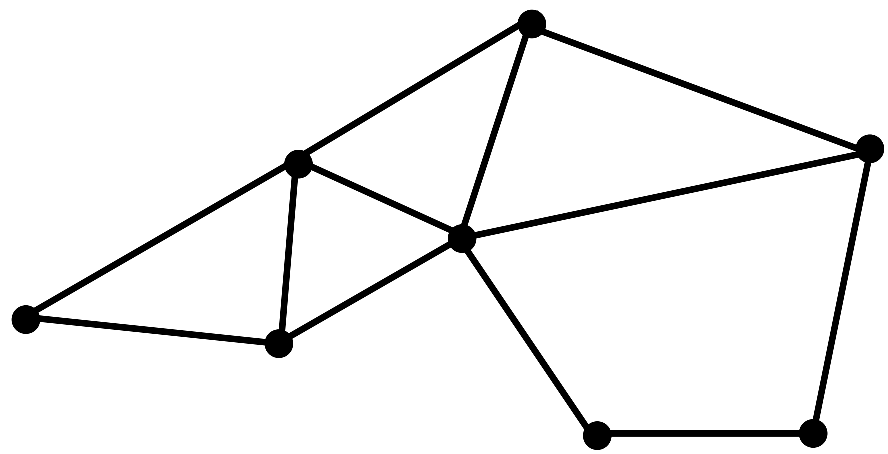

Gabriel graph contains the minimal spanning tree.

Gabriel and Sokol, *A new statistical approach to geographic variation analysis*, Systematic Zoology, 1969.

Gabriel edges $\Rightarrow$ non-obtuse triangulation

Bern-Mitchell-Ruppert (1994)

BMR Lemma: Add k vertices to sides of triangle (at least one per side) so all edges become Gabriel, then add all midpoints. Resulting polygon has a $O(k)$ NOT, with no additional vertices on boundary.

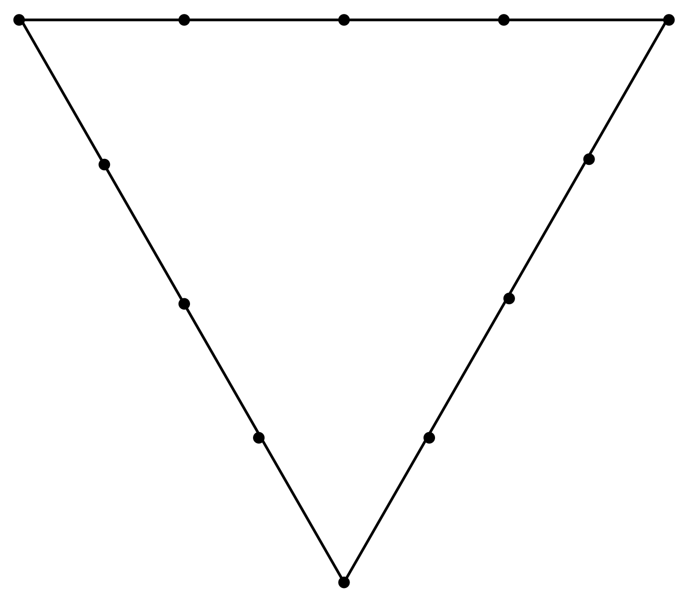

Gabriel edges $\Rightarrow$ non-obtuse triangulation

Bern-Mitchell-Ruppert (1994)

BMR Lemma: Add k vertices to sides of triangle (at least one per side) so all edges become Gabriel, then add all midpoints. Resulting polygon has a $O(k)$ NOT, with no additional vertices on boundary.

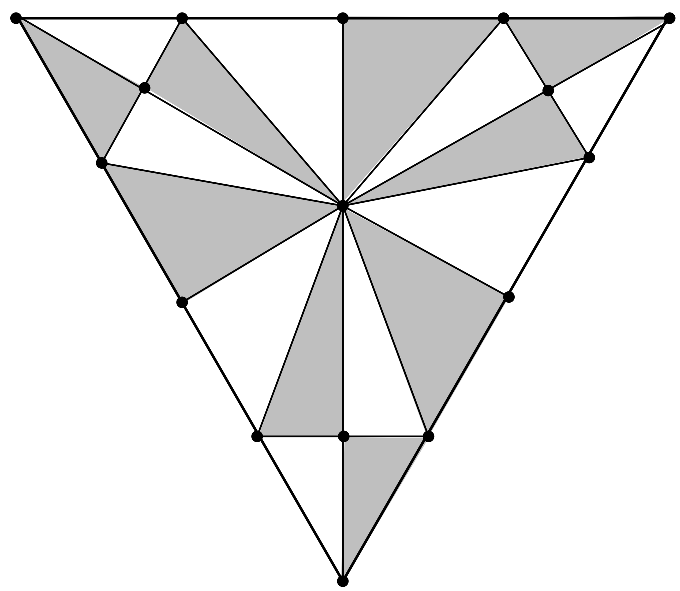

Gabriel edges $\Rightarrow$ non-obtuse triangulation

Bern-Mitchell-Ruppert (1994)

BMR Lemma: Add k vertices to sides of triangle (at least one per side) so all edges become Gabriel, then add all midpoints. Resulting polygon has a $O(k)$ NOT, with no additional vertices on boundary.

Building a NOT for a PSLG:

- Replace PSLG by triangulation of itself.
- Add vertices to make all edges Gabriel.
- Apply BMR lemma.

Theorem: Given a triangulation, we can add $O(n^{2.5})$ points to the edges so that every edge is Gabriel.

Construct Gabriel points:

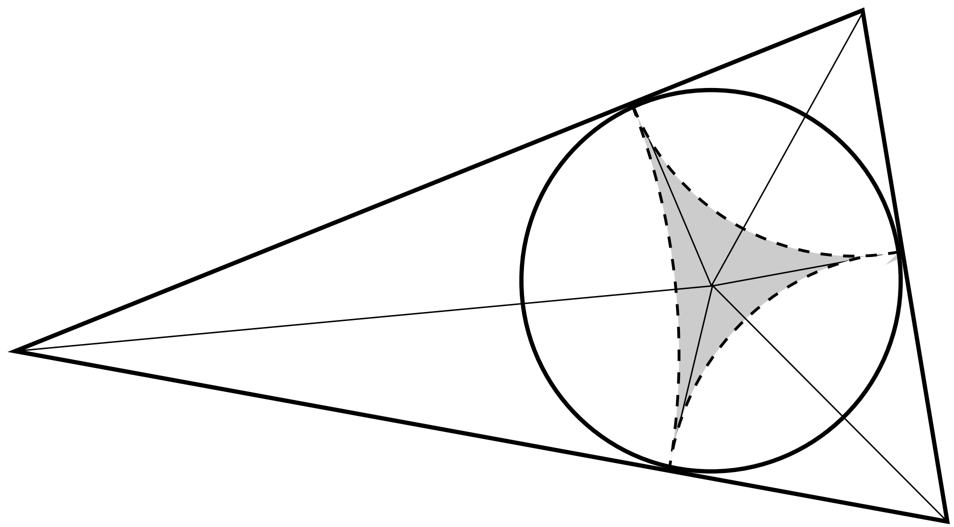

Break every triangle into thick and thin parts.

Thin parts = corners, Thick part = central region

Construct Gabriel points:

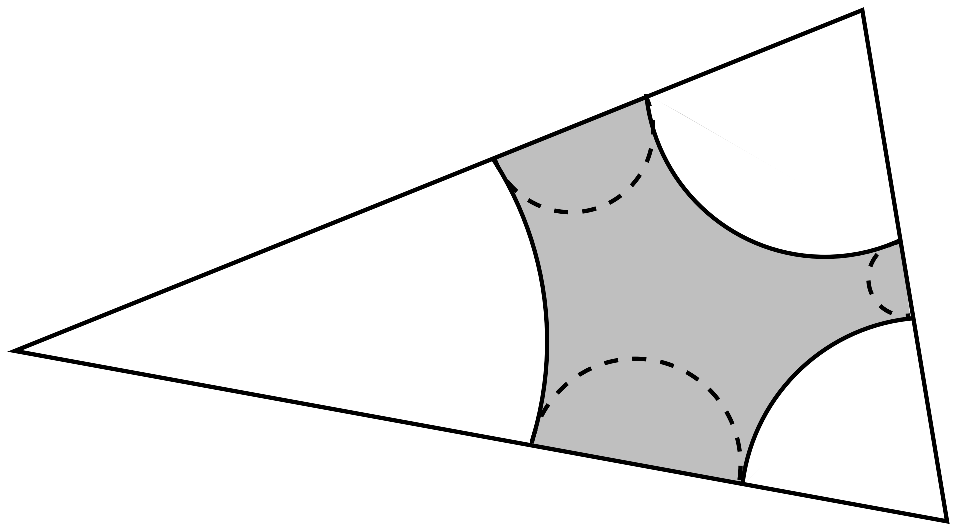

Advantageous to increase thick part.

Thick sides are base of half-disk inside triangle.

Construct Gabriel points:

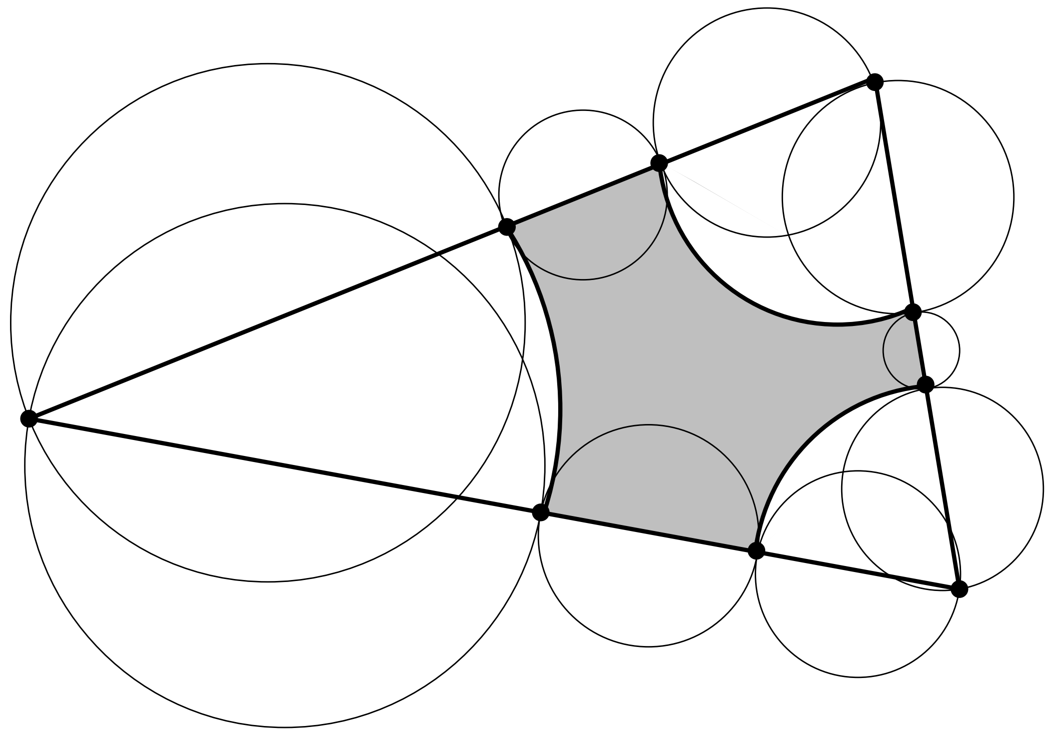

Then vertices of thick part give Gabriel edges.

Construct Gabriel points:

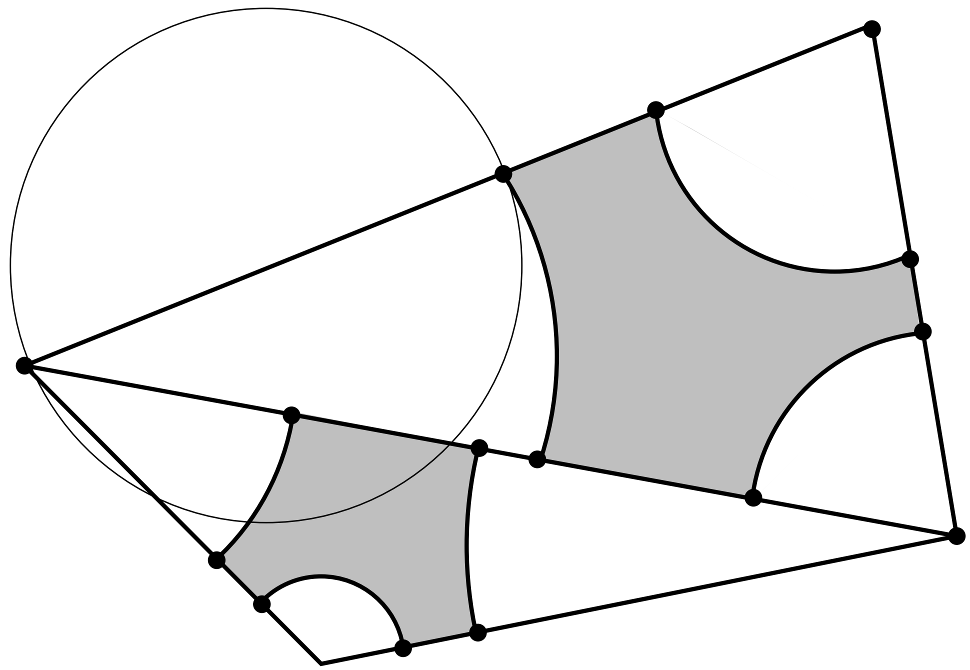

But, adjacent triangle can make Gabriel condition fail.

Construct Gabriel points:

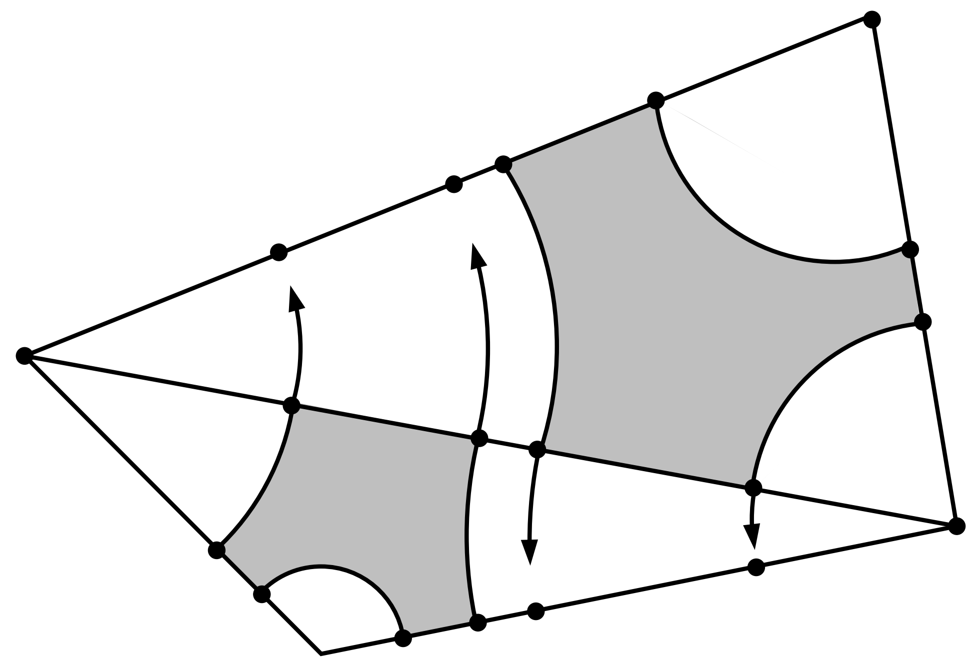

Idea: "Push" vertices across the thin parts.

Construct Gabriel points:

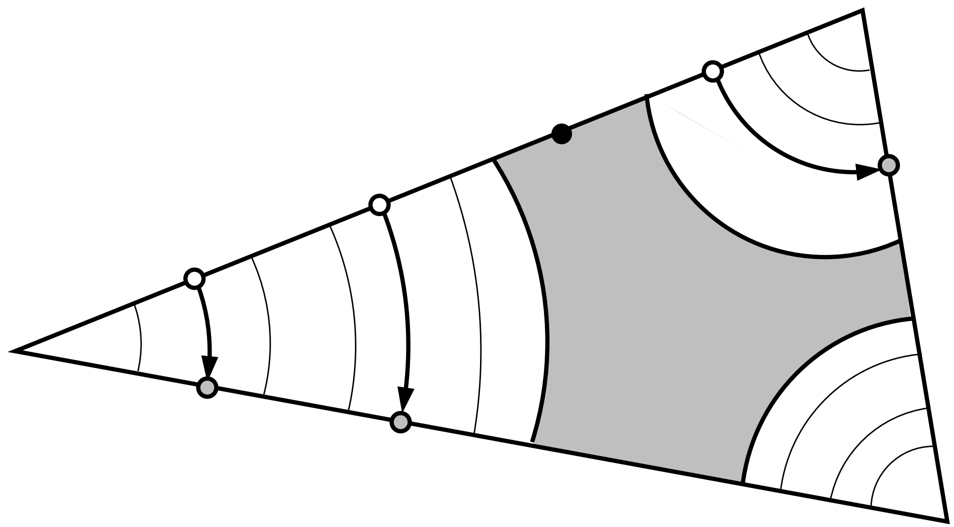

Thin parts foliated by circles centered at vertices.

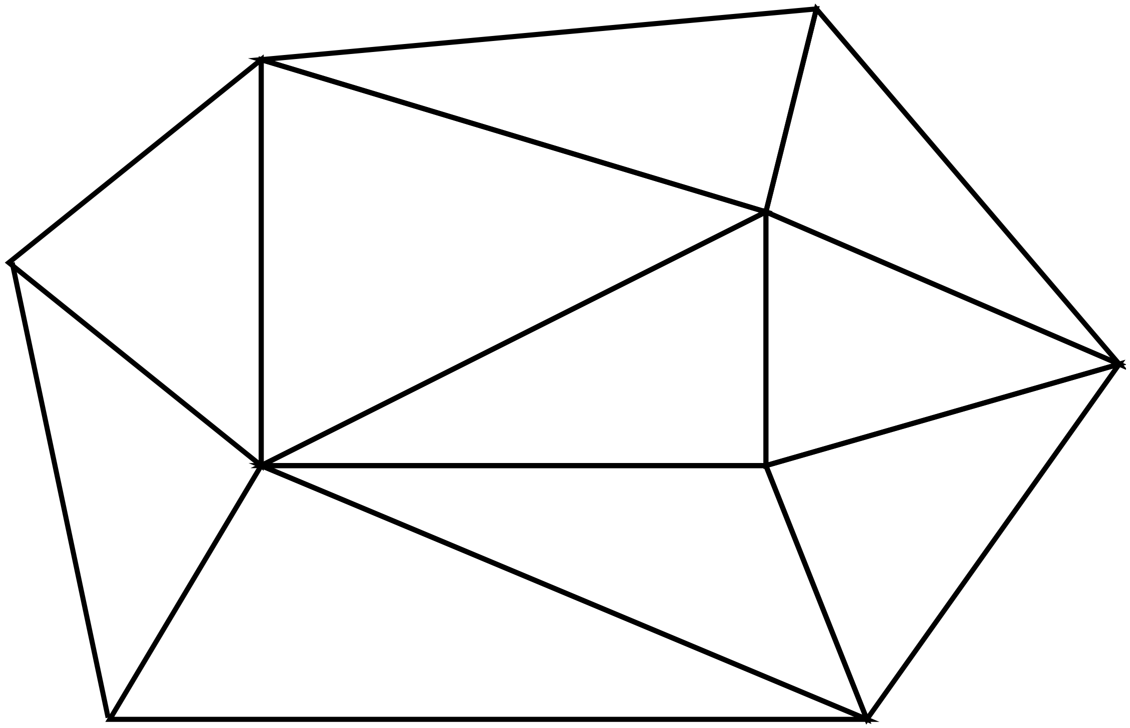

- Start with any triangulation.

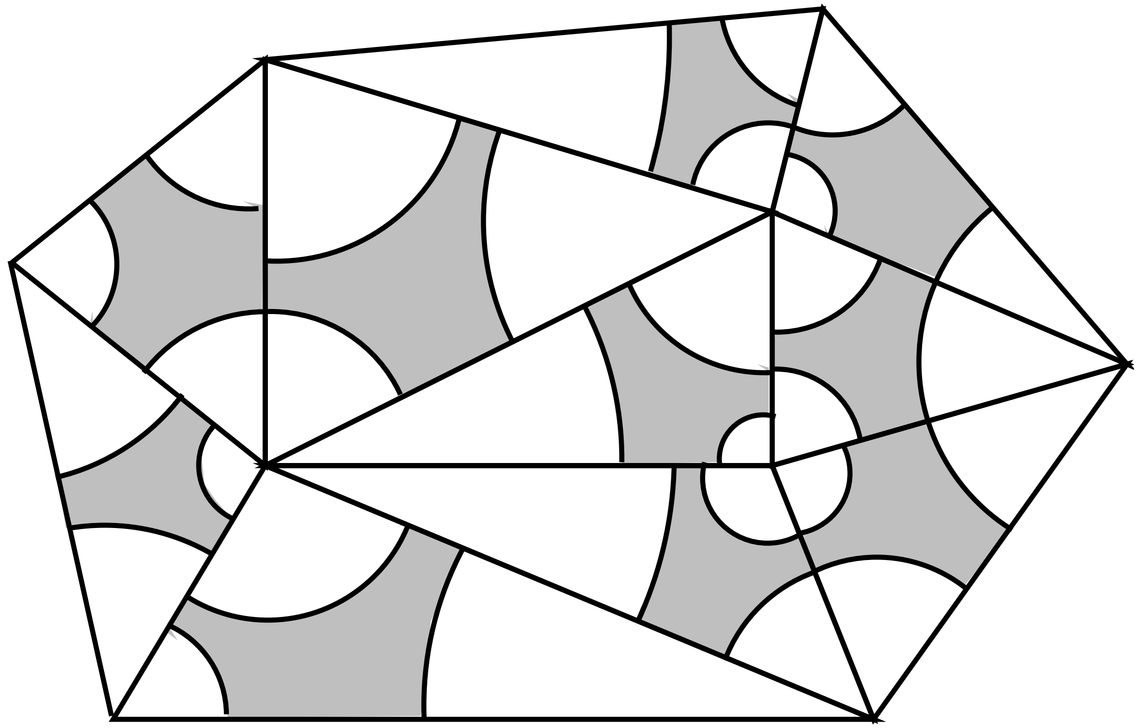

- Start with any triangulation.
- Make thick/thin parts.

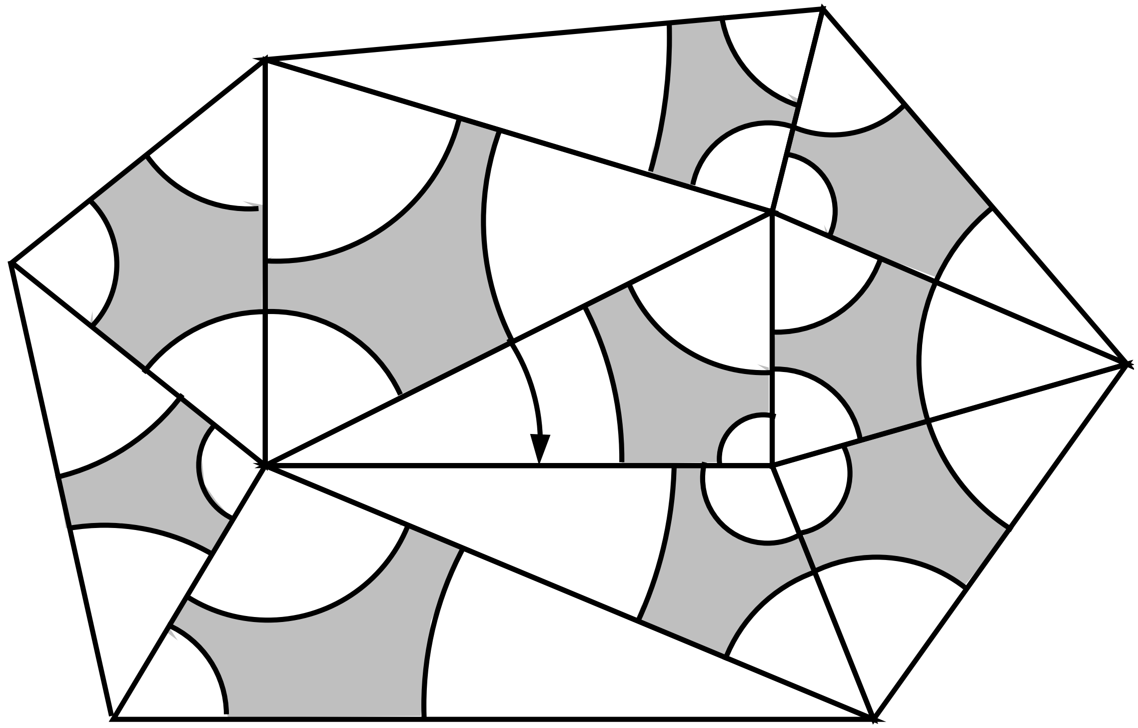

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

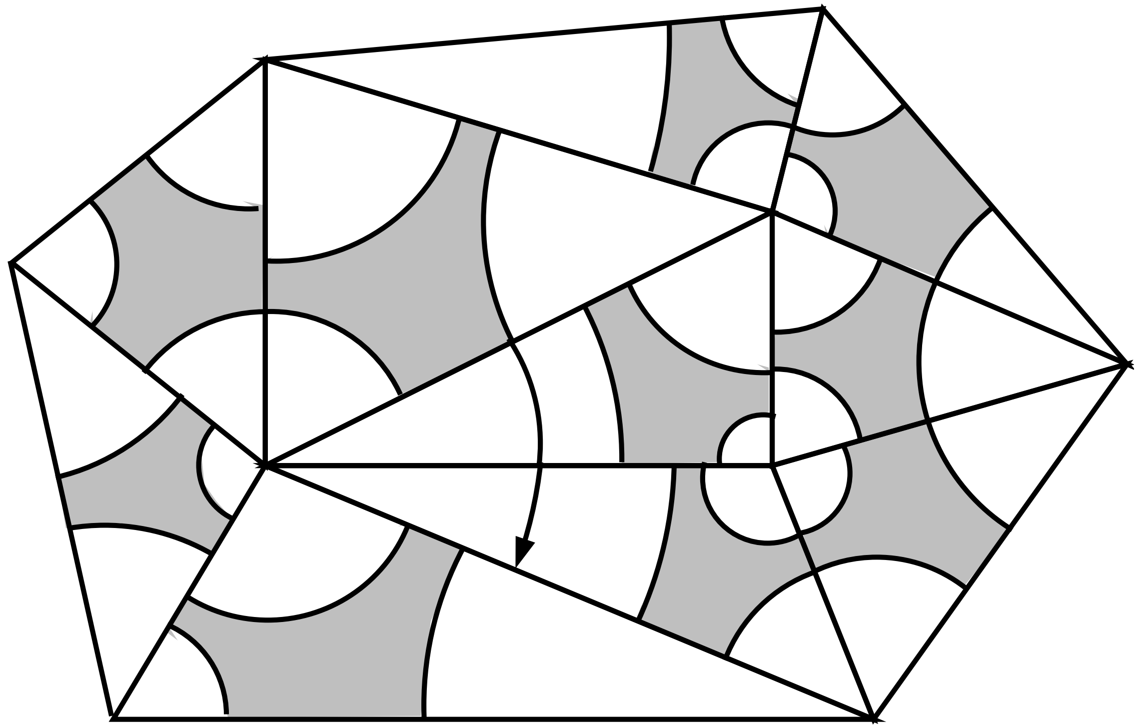

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

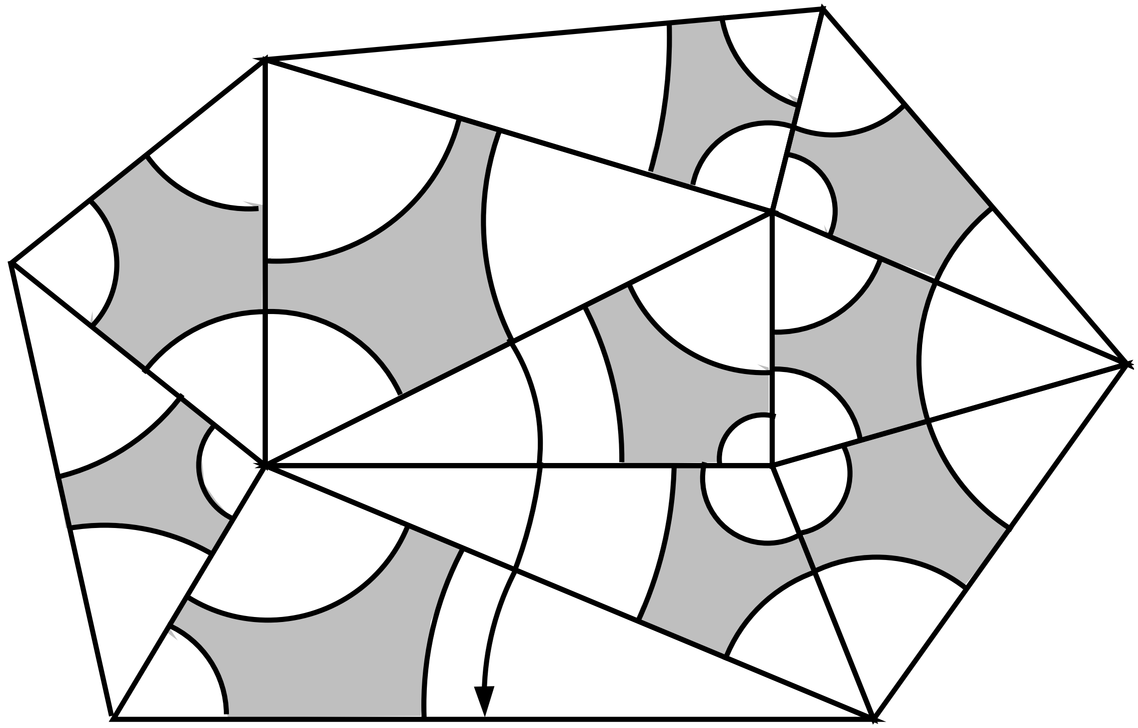

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

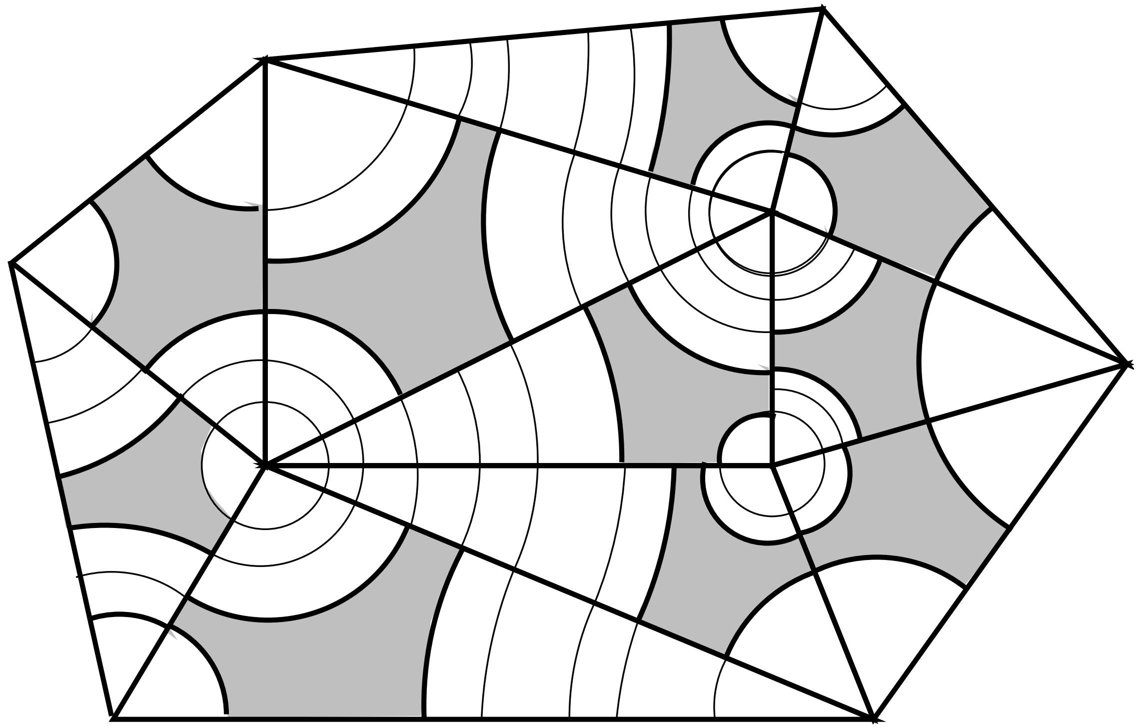

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

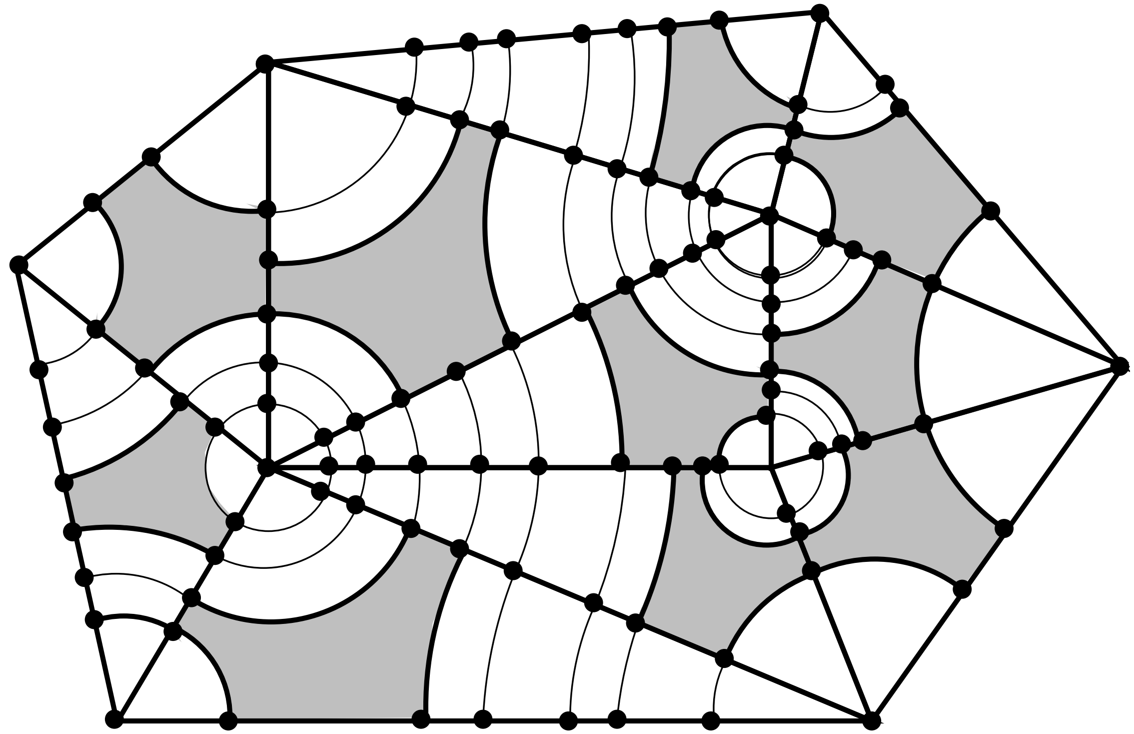

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.
- Intersections satisfy Gabriel condition. Why?

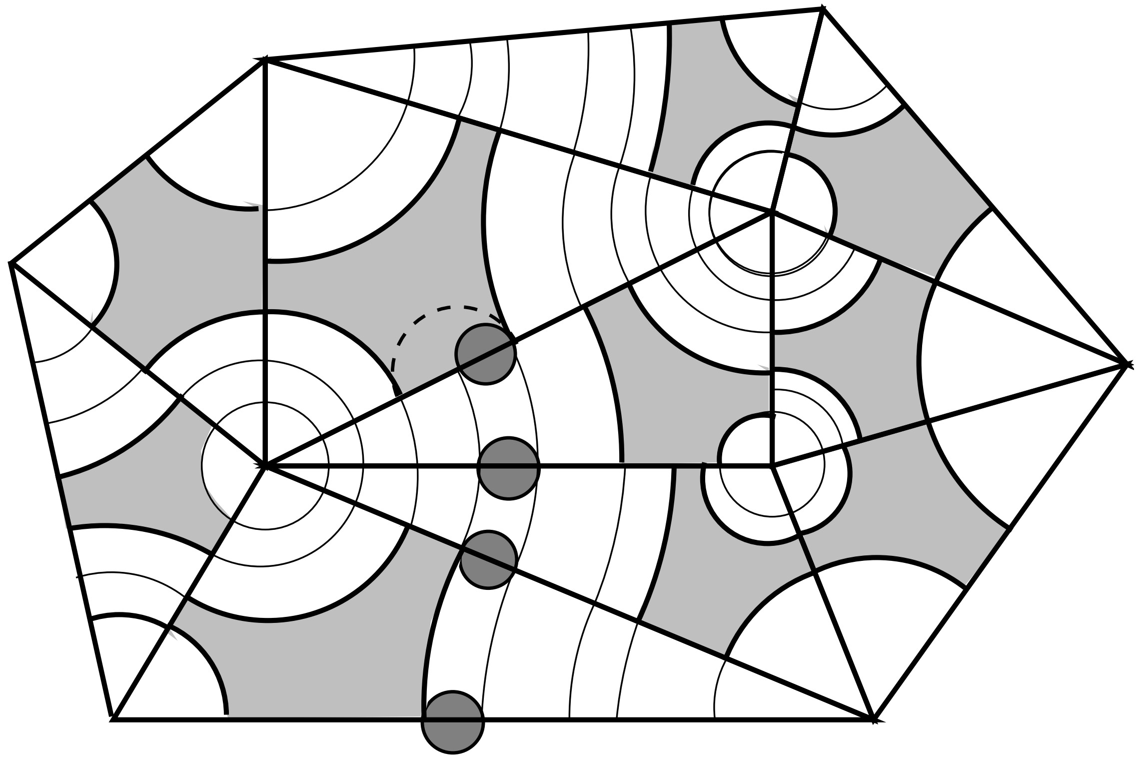

Tube is “swept out” by fixed diameter disk.

Disk lies inside tube or thick part or outside convex hull.

In triangulation of a n-gon, adjacent triangles form tree.

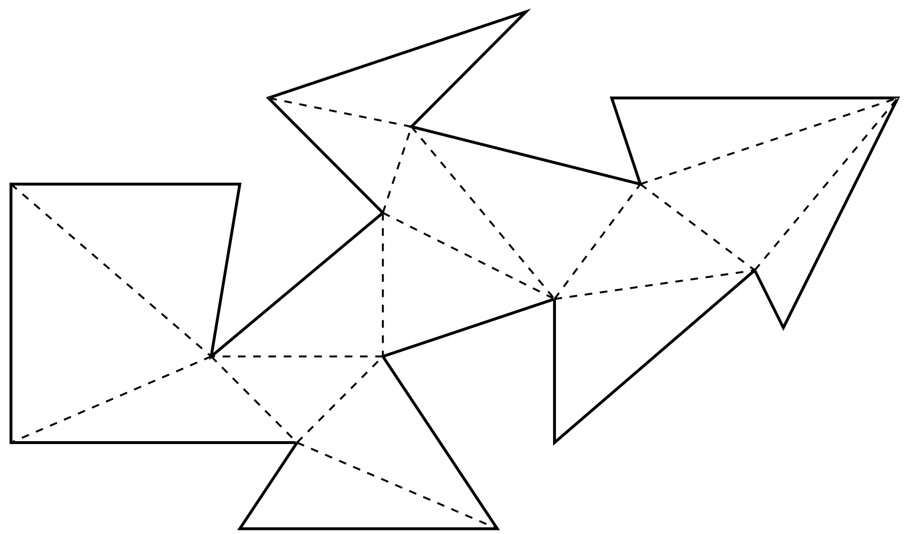

In triangulation of a n-gon, adjacent triangles form tree.

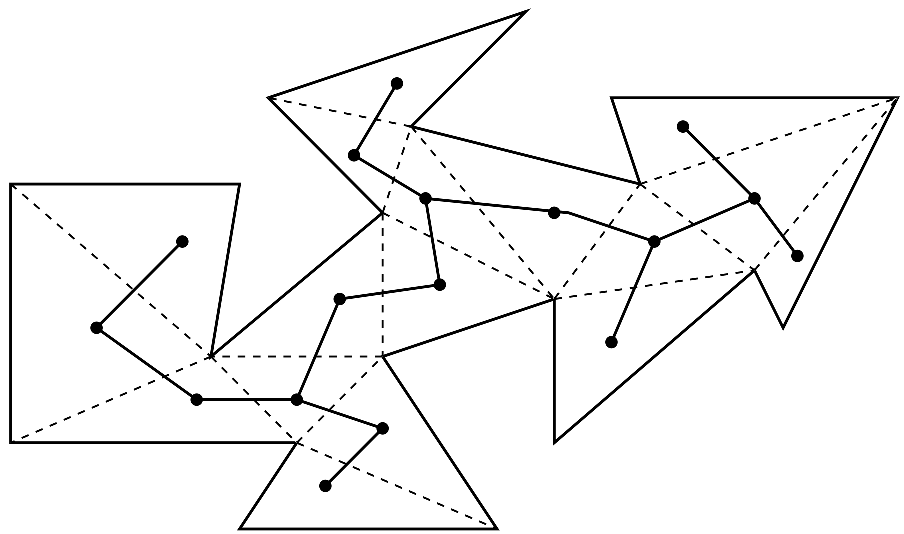

In triangulation of a n-gon, adjacent triangles form tree.

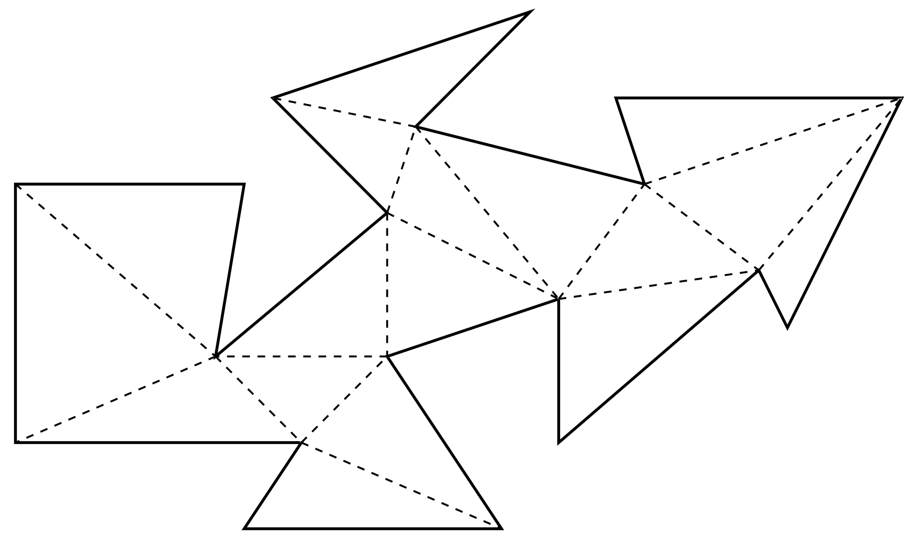

Hence paths never revisit a triangle.
$6n$ starting points, so $O(n^2)$ points are created.

General case requires intricate construction to terminate “sprialing paths” that return often to same triangle.

Large polynomial bound is easy, but 2.5 is harder.

Cor: Any PSLG has a conforming Delaunay triangulation of size $O(n^{2.5})$.

Edelsbrunner, Tan (1993) proved this with $O(n^3)$.

Questions:

- Implementable?
- Average versus worst case bounds?
- Improve 2.5 to 2?
- 3-D meshes? The eightfold way? Ricci flow?
- Other applications for thick/thin pieces?
- Applications of Mumford-Bers compactness?
- Best K for medial axis map? $(2.1 < k < 7.82)$
- Can we do better than medial axis map?