HARMONIC MEASURE, CONFORMAL MAPS AND OPTIMAL MESHES

Christopher J. Bishop
Stony Brook University

Simons Collaboration on Algorithms & Geometry
February 28, 2020

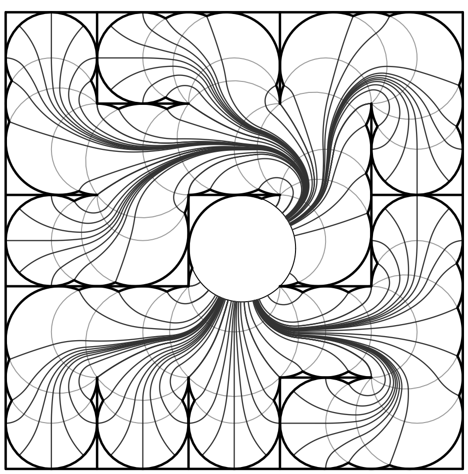
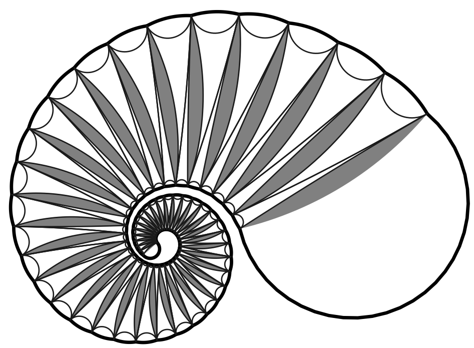
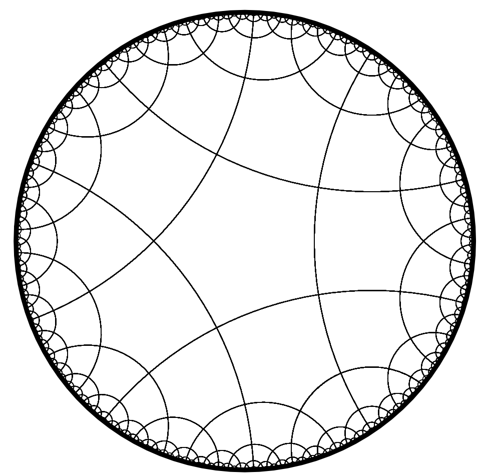
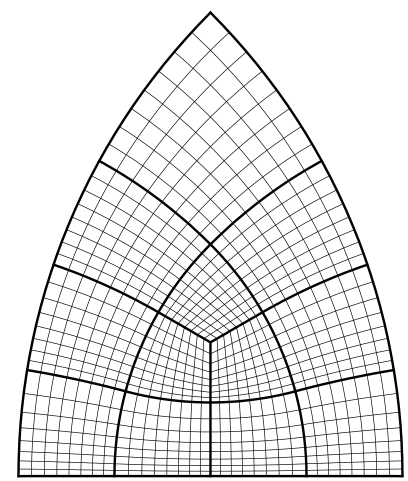

www.math.sunysb.edu/~bishop/lectures

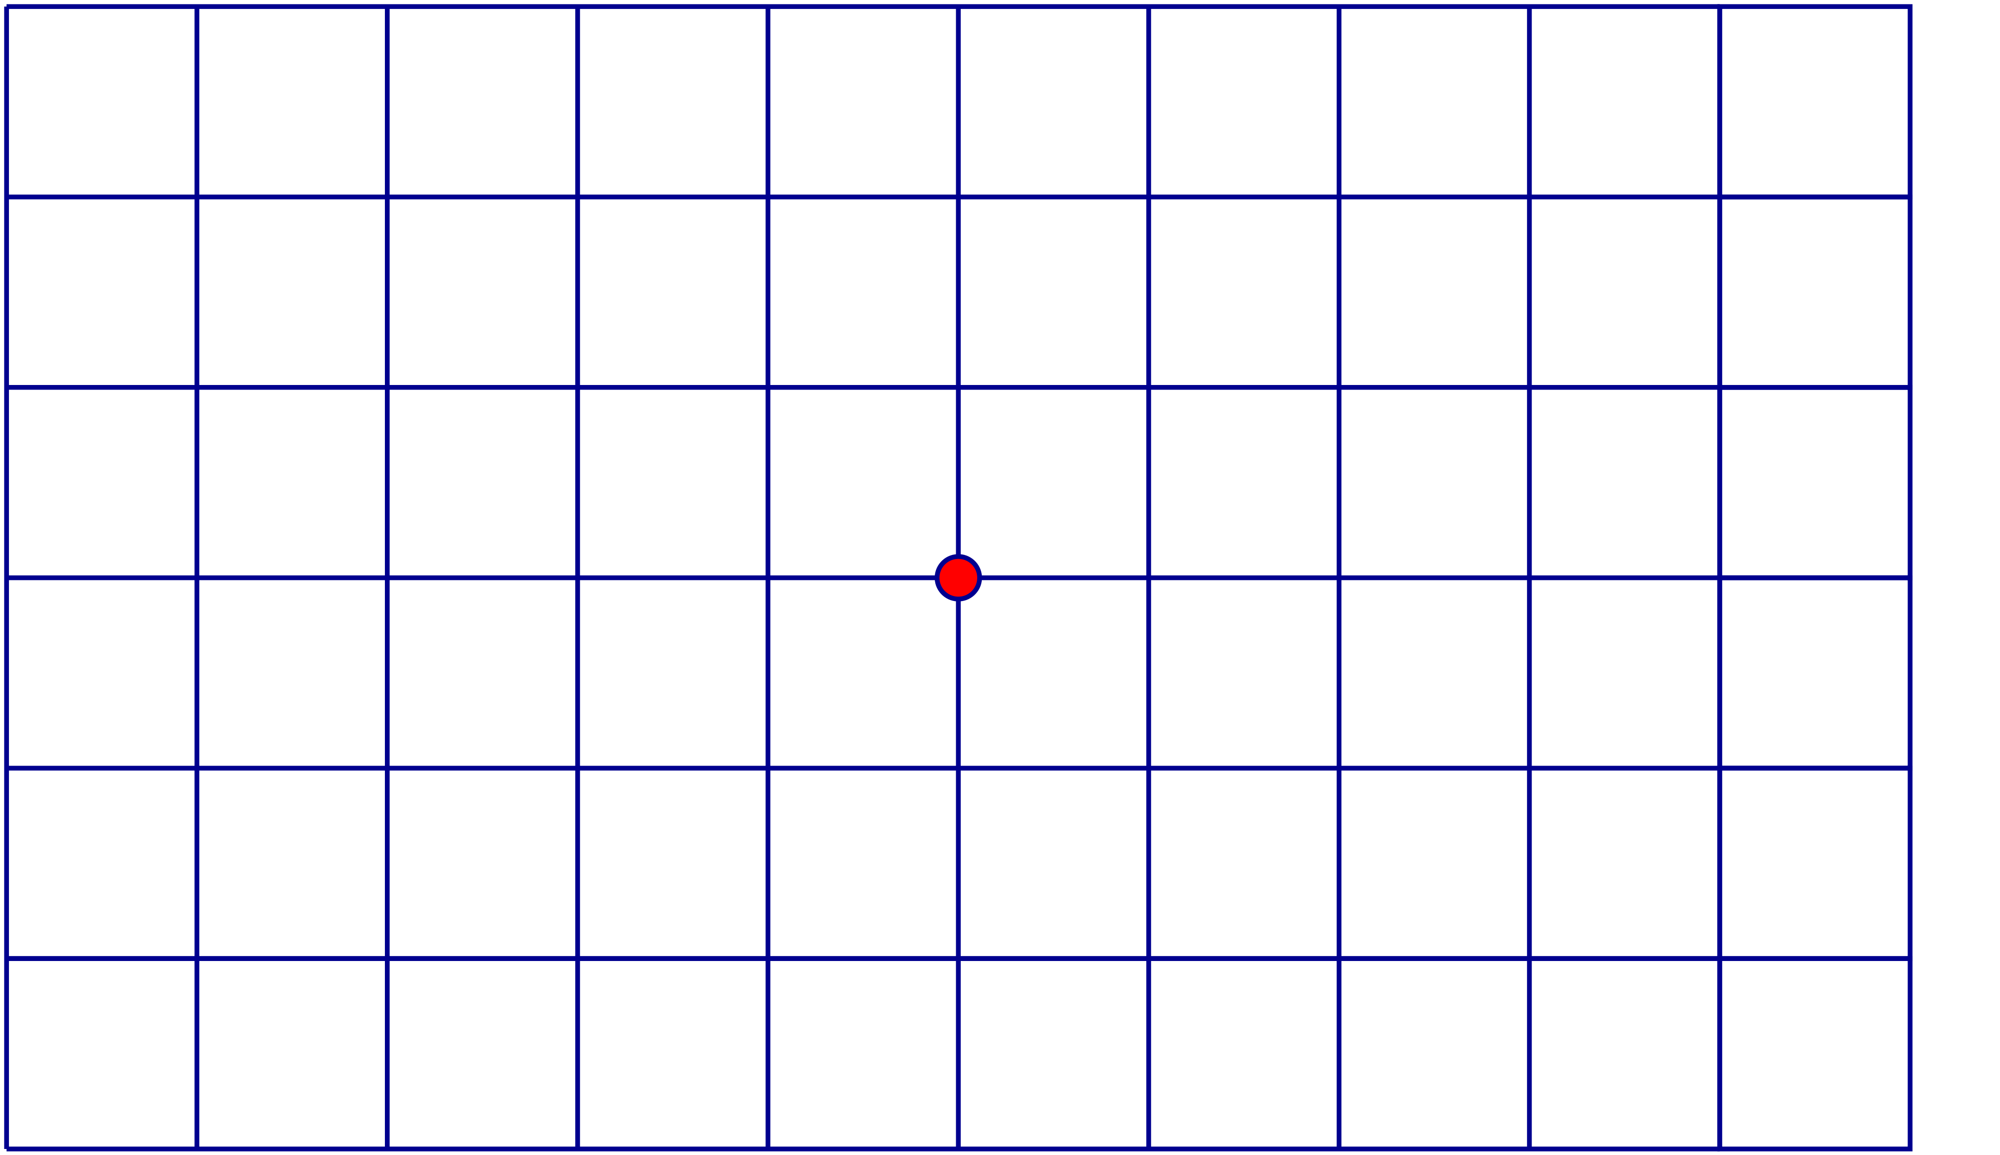

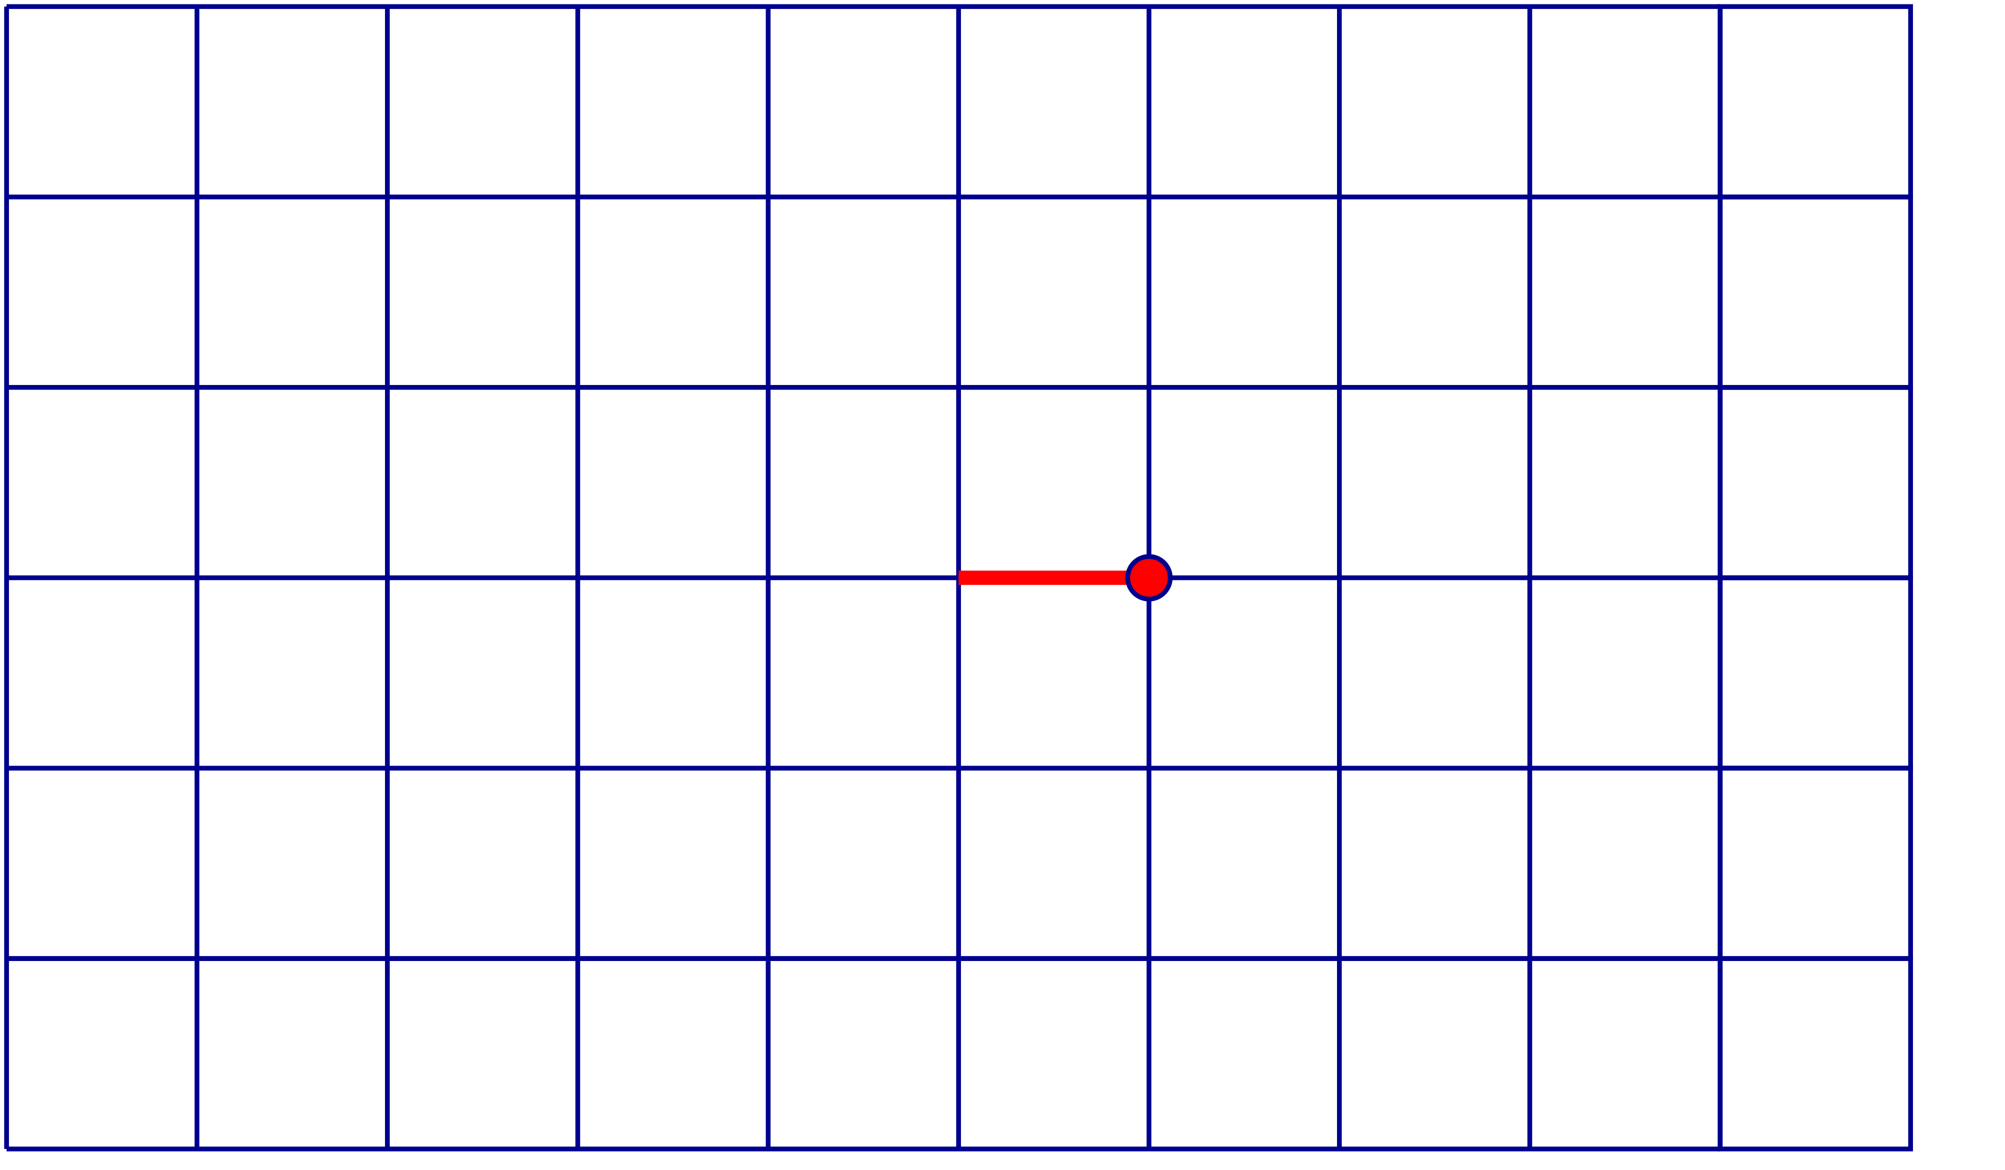

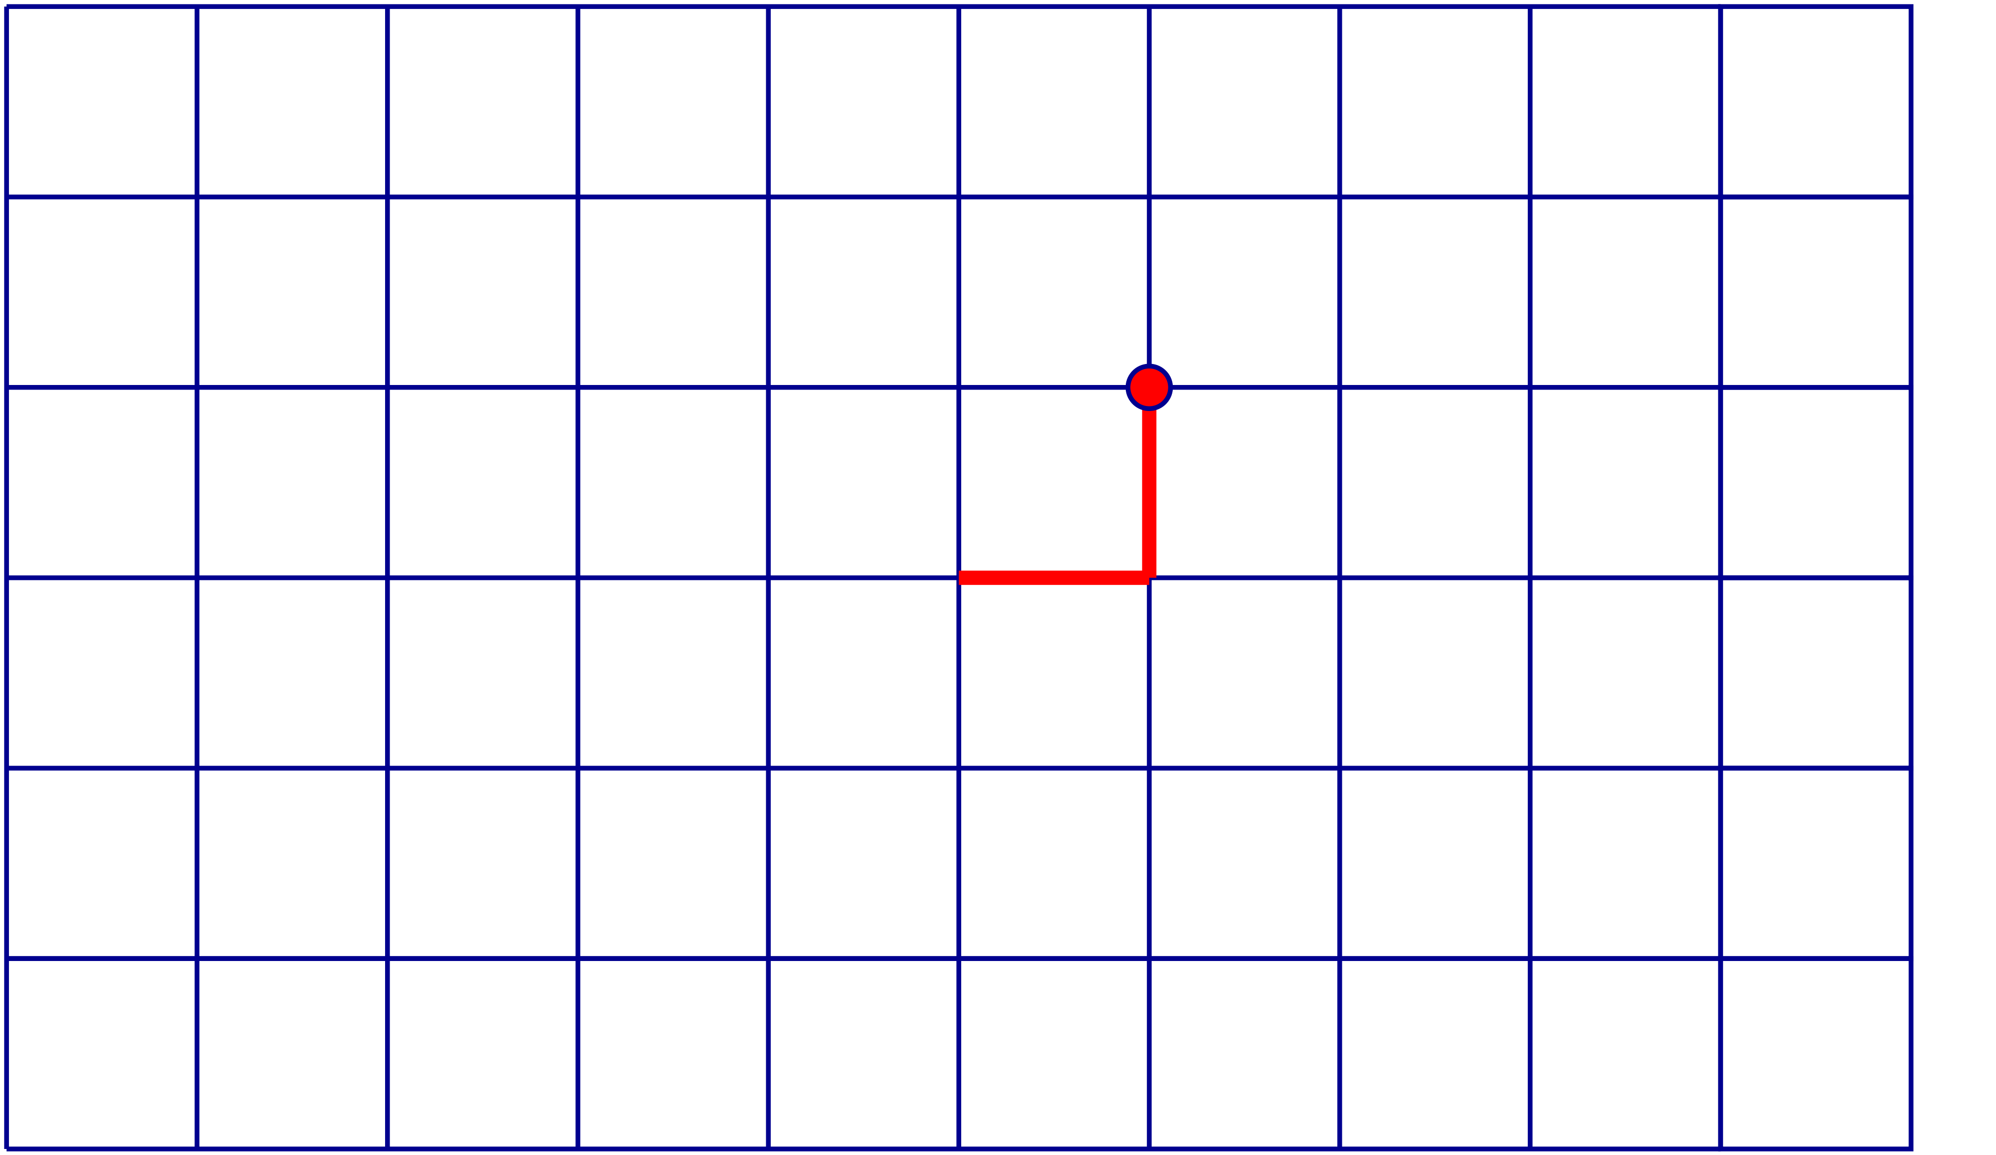

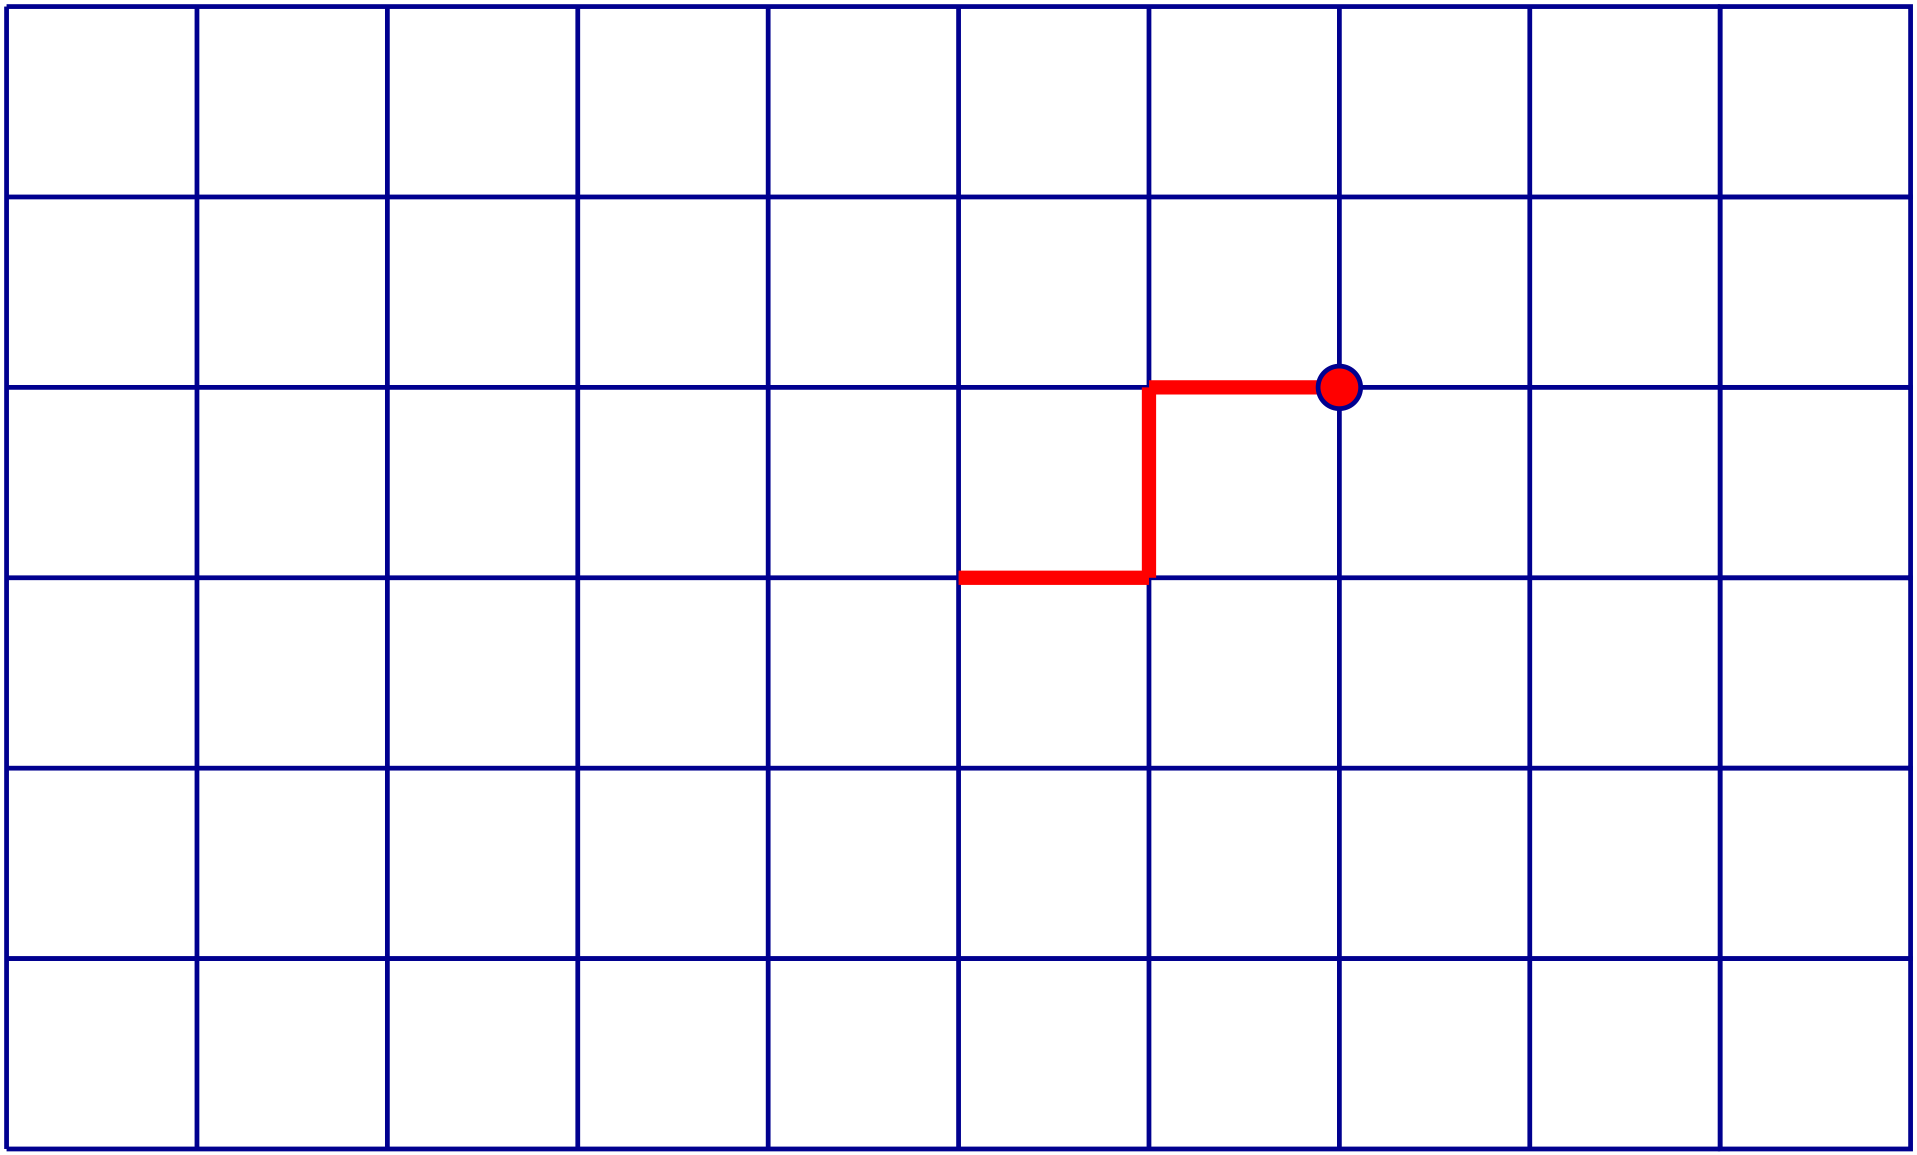

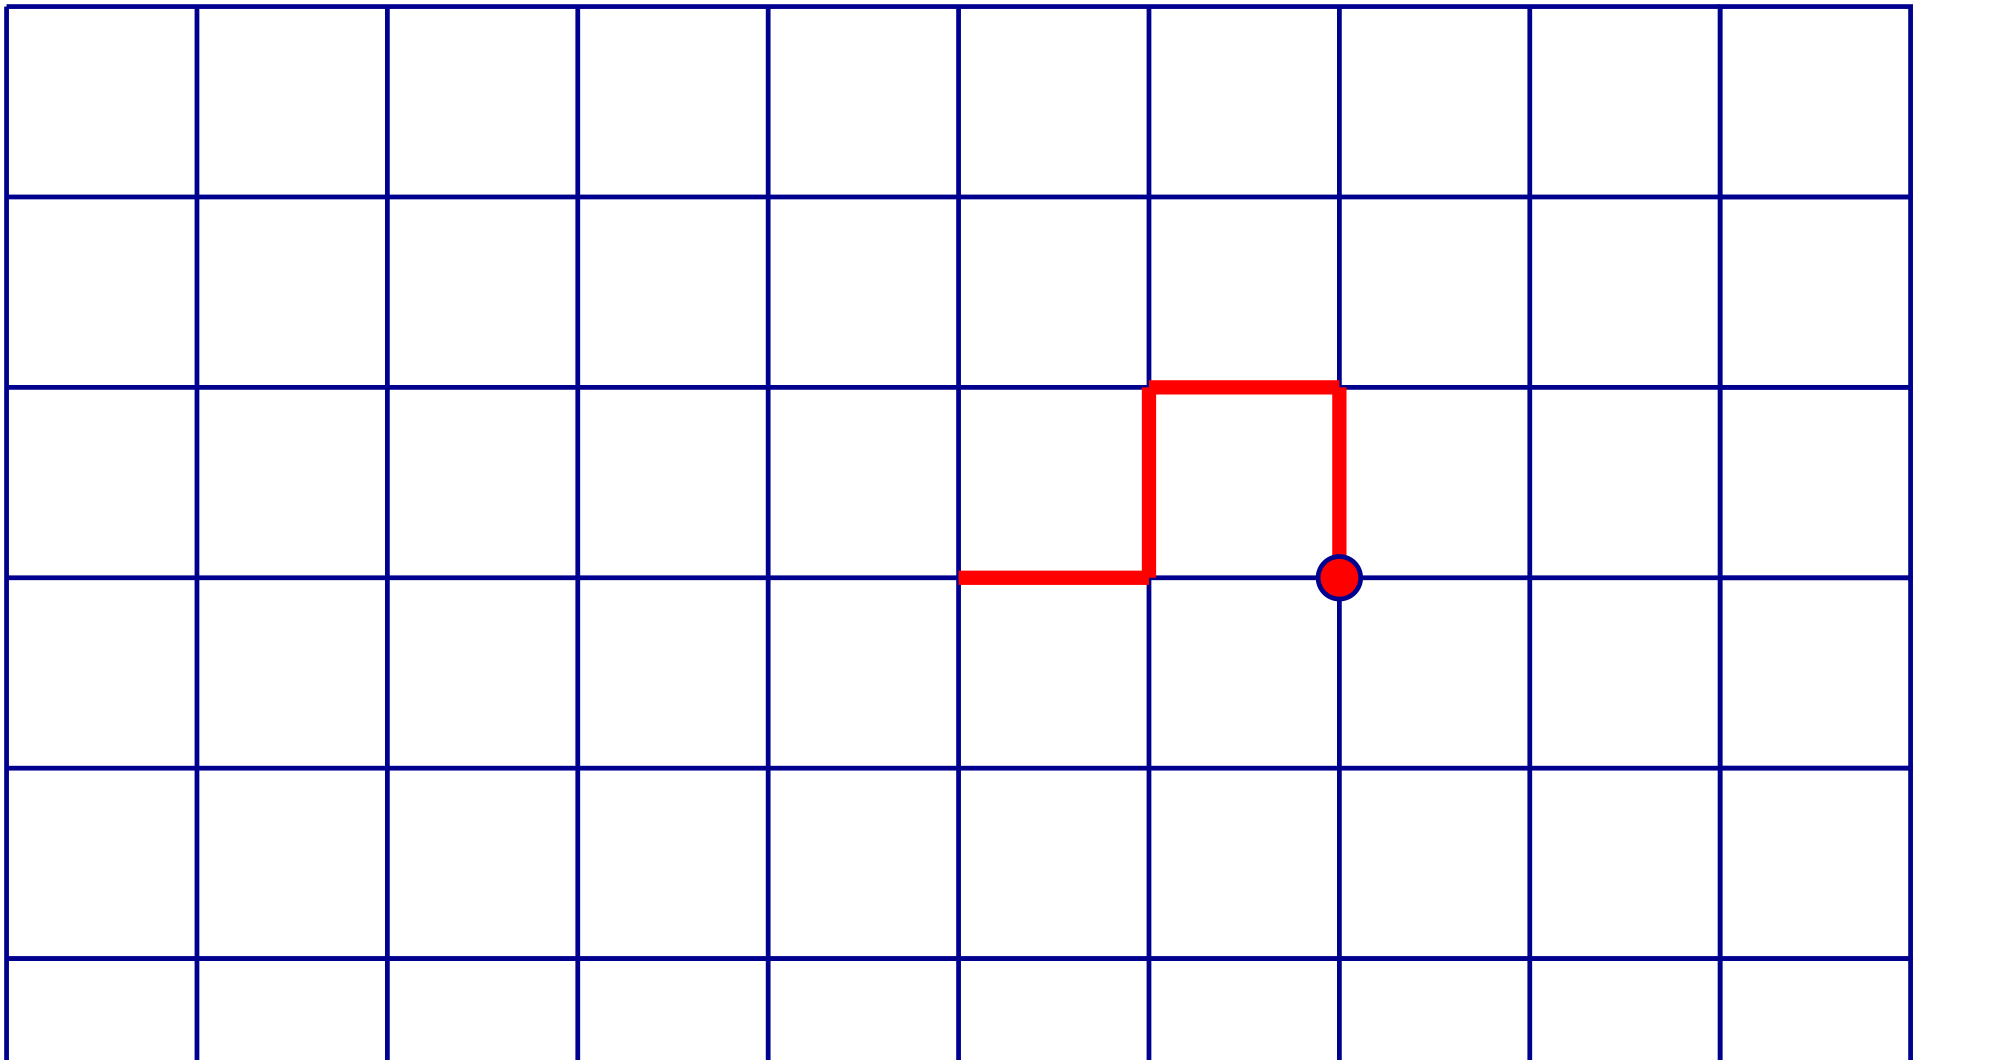

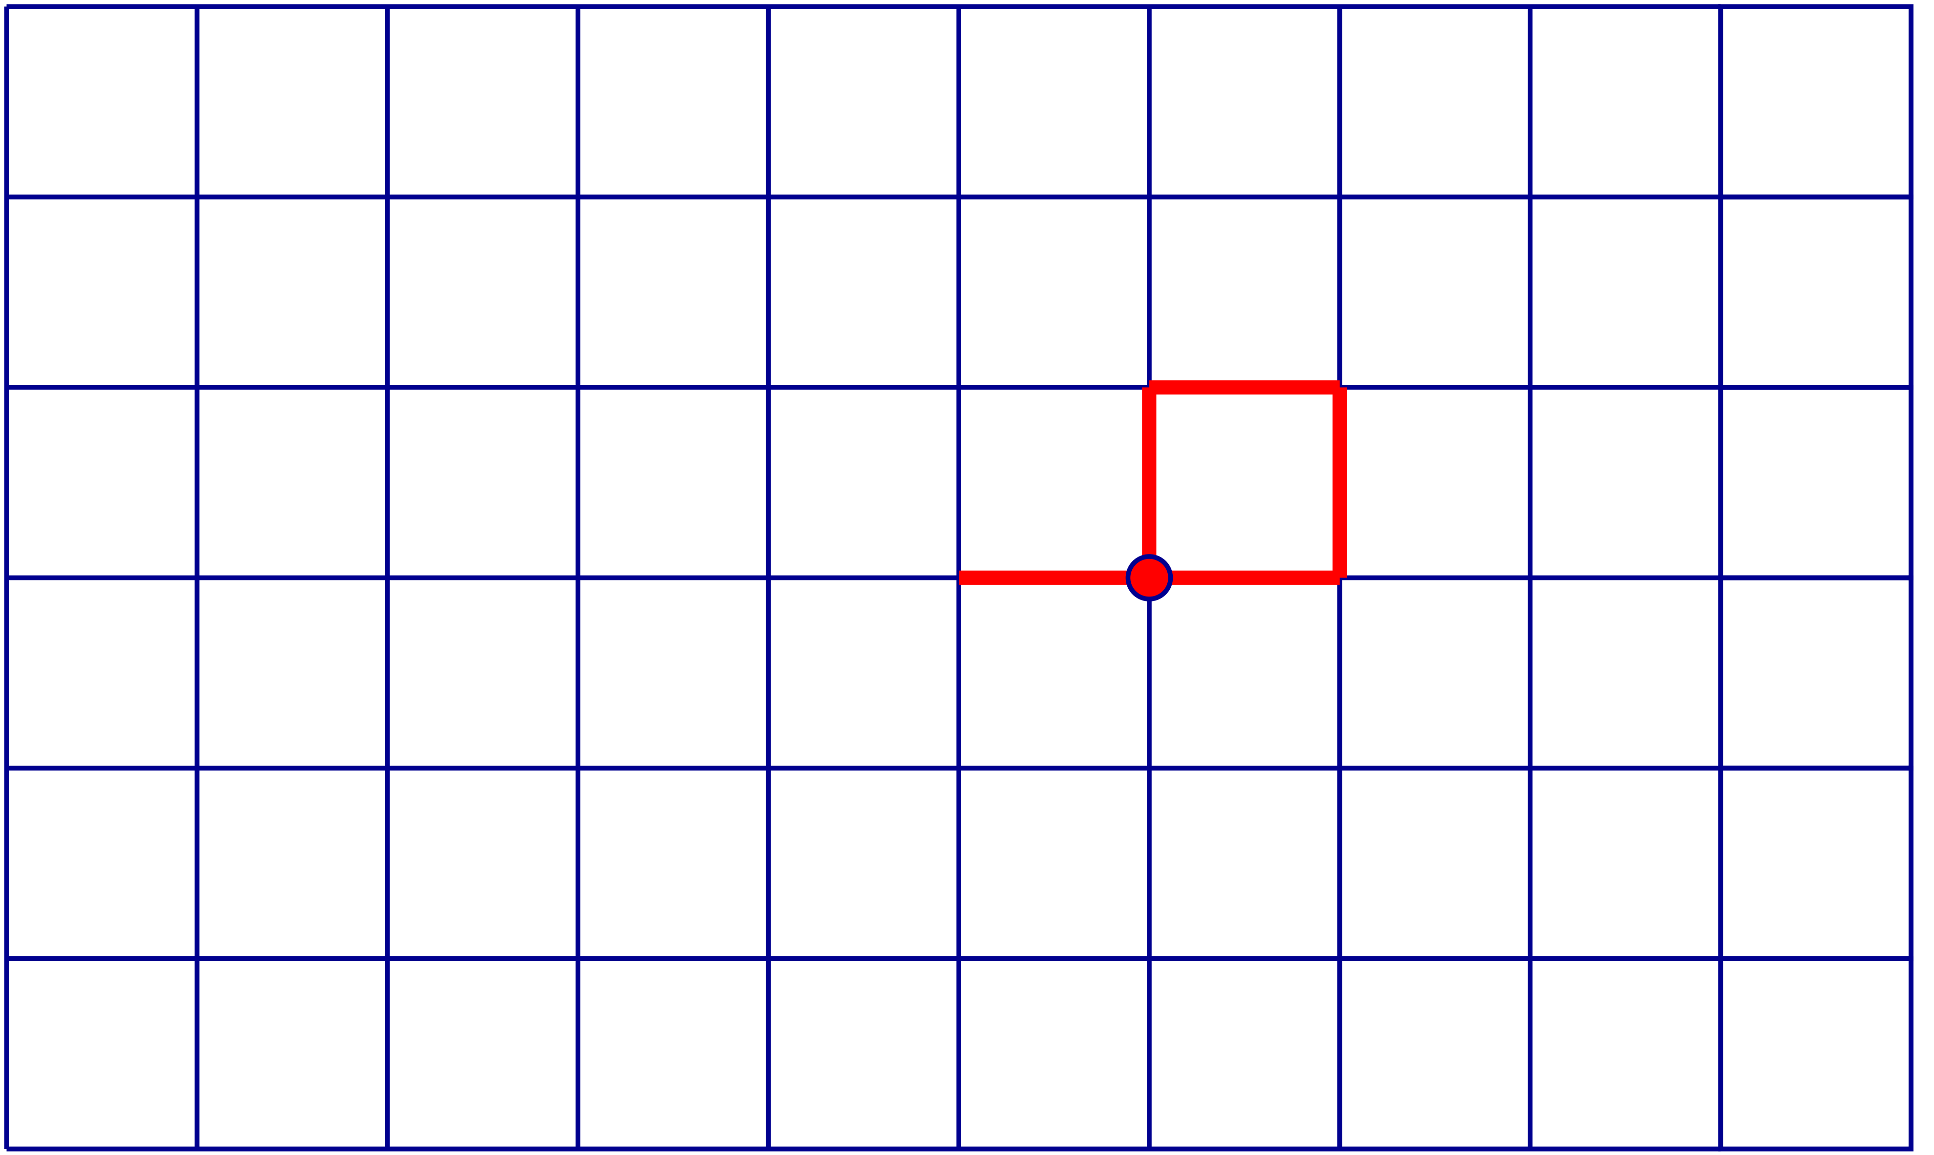

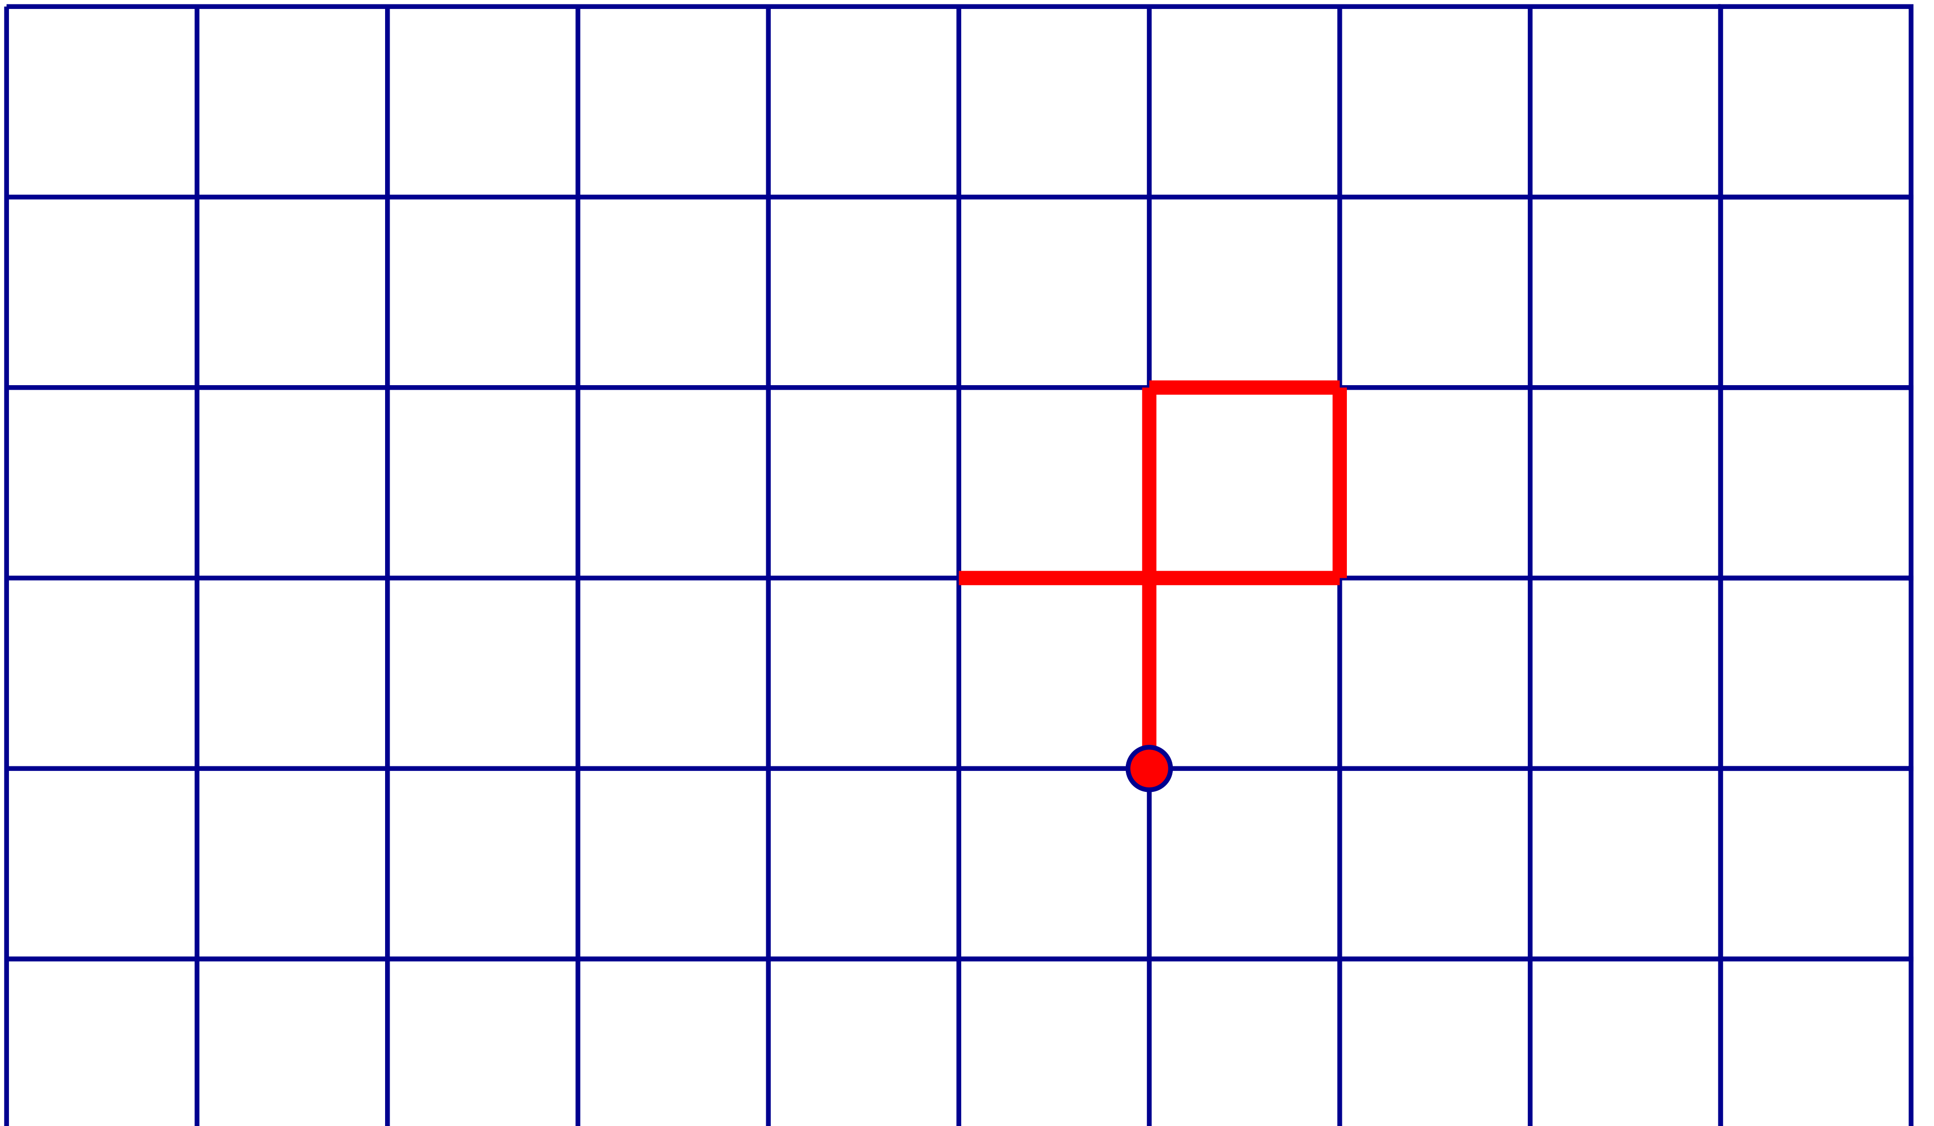

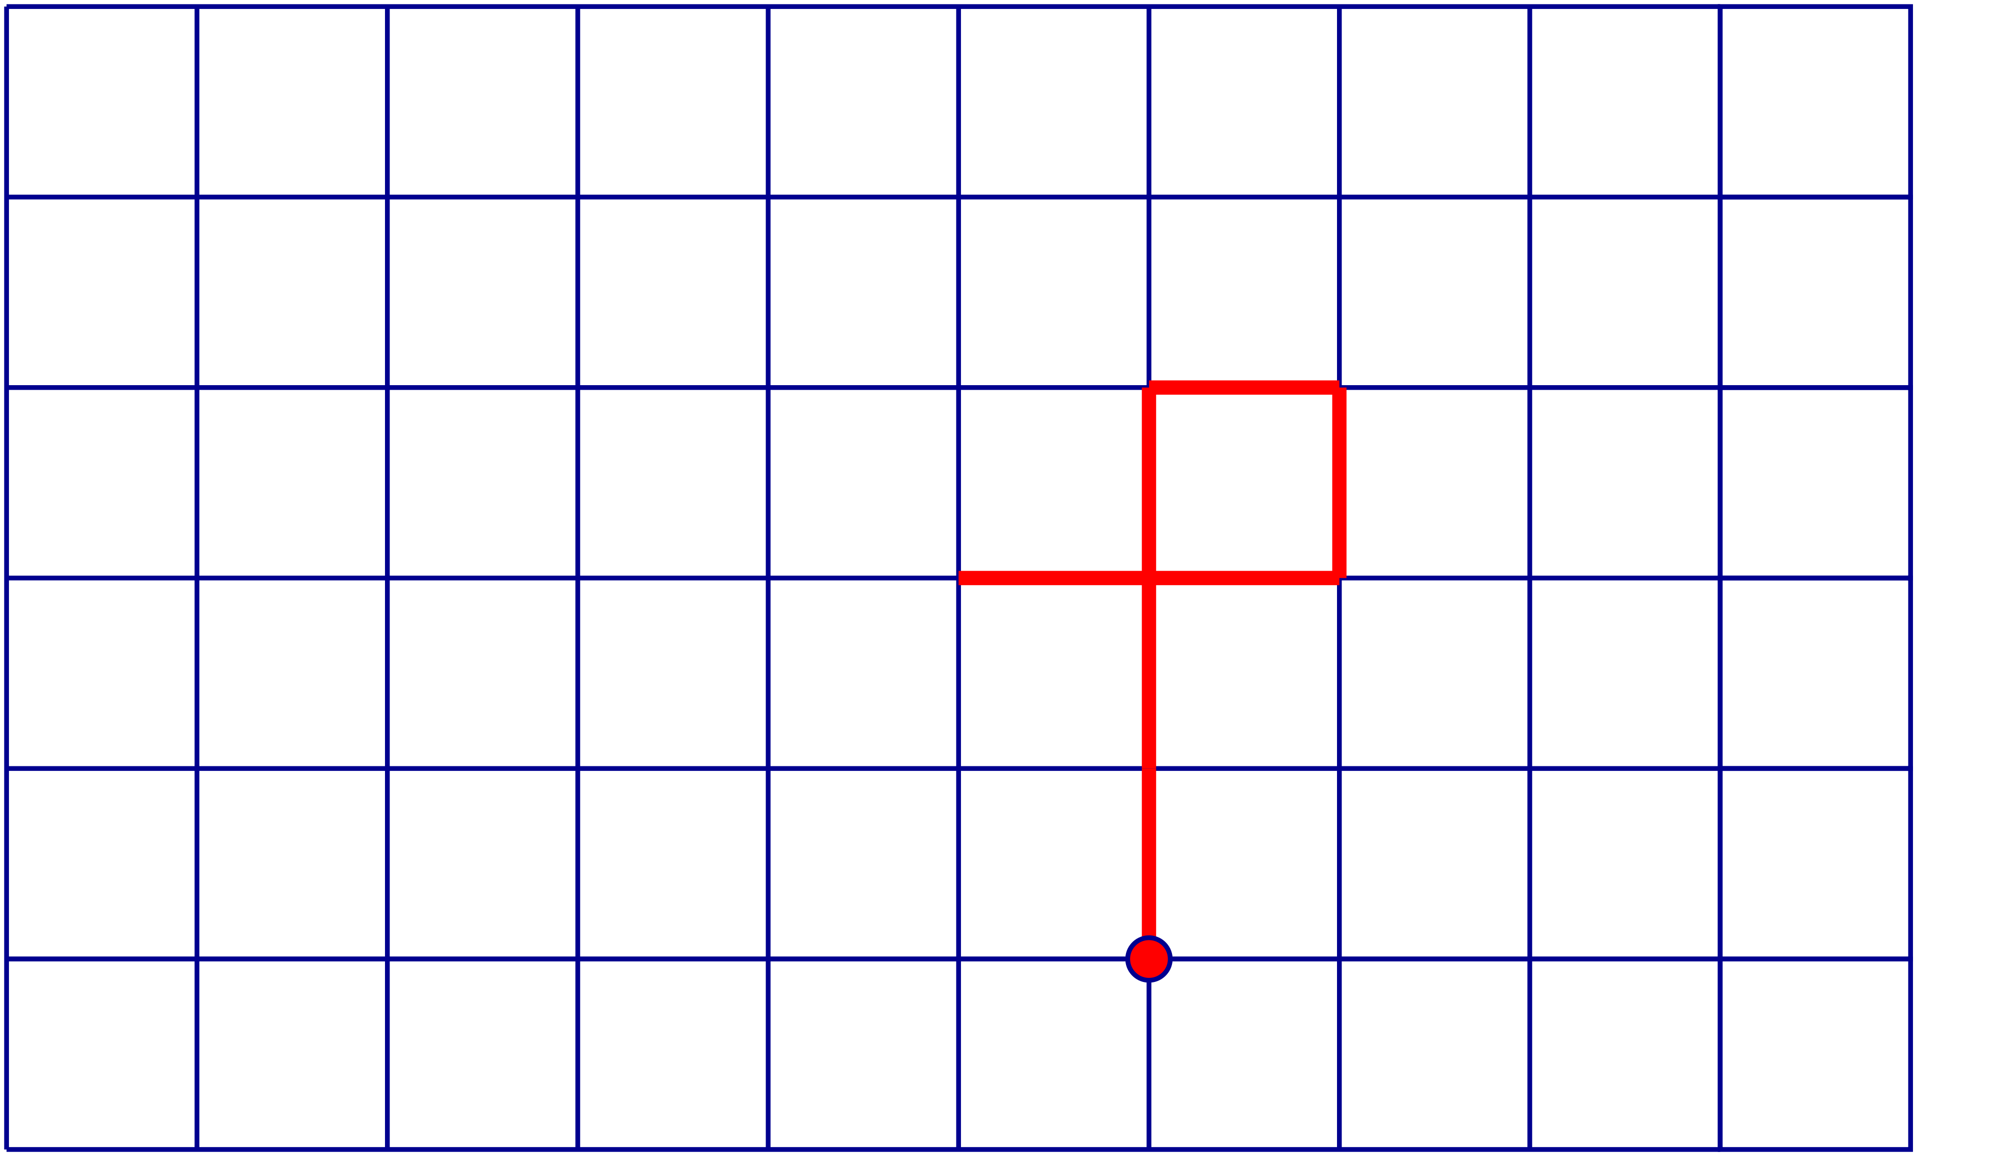

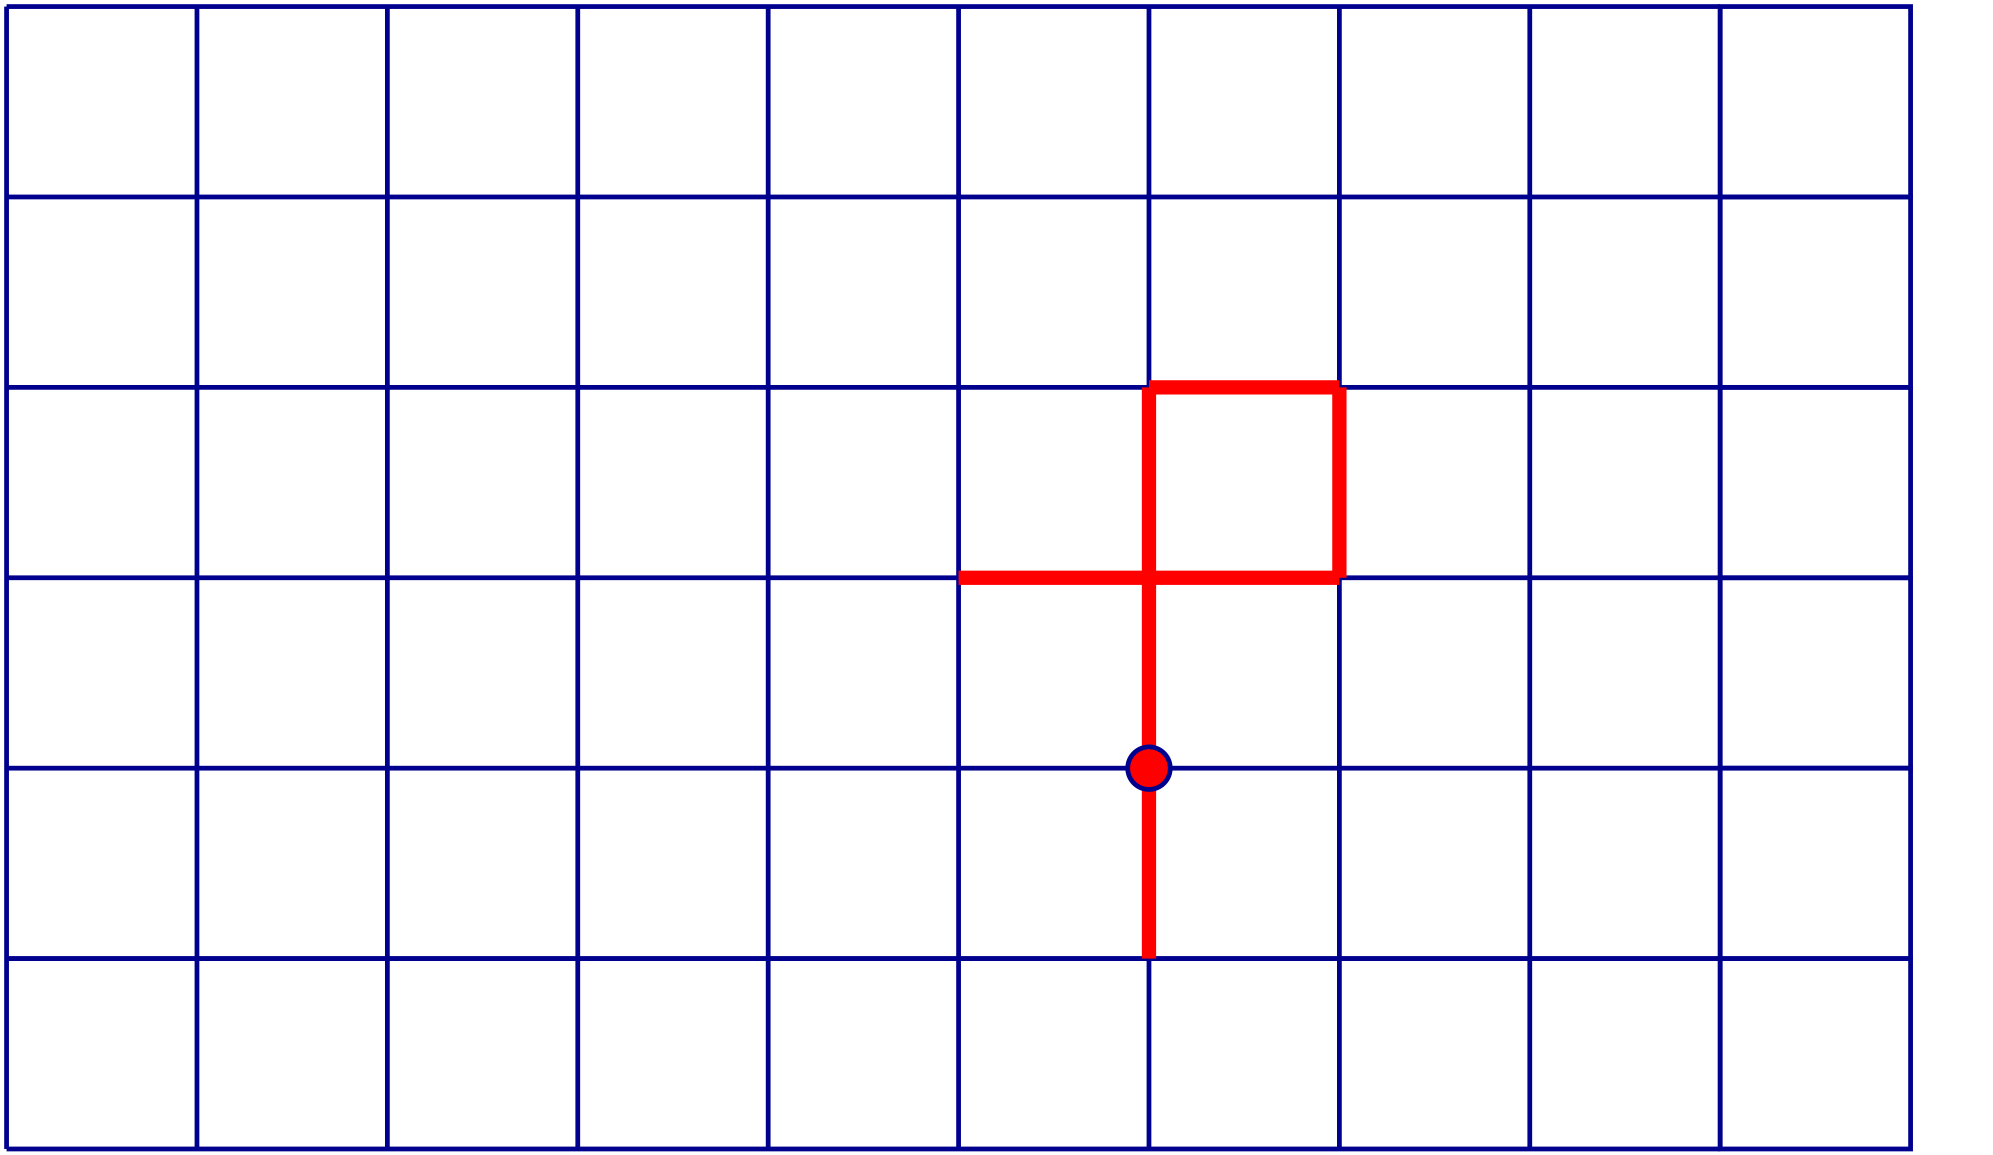

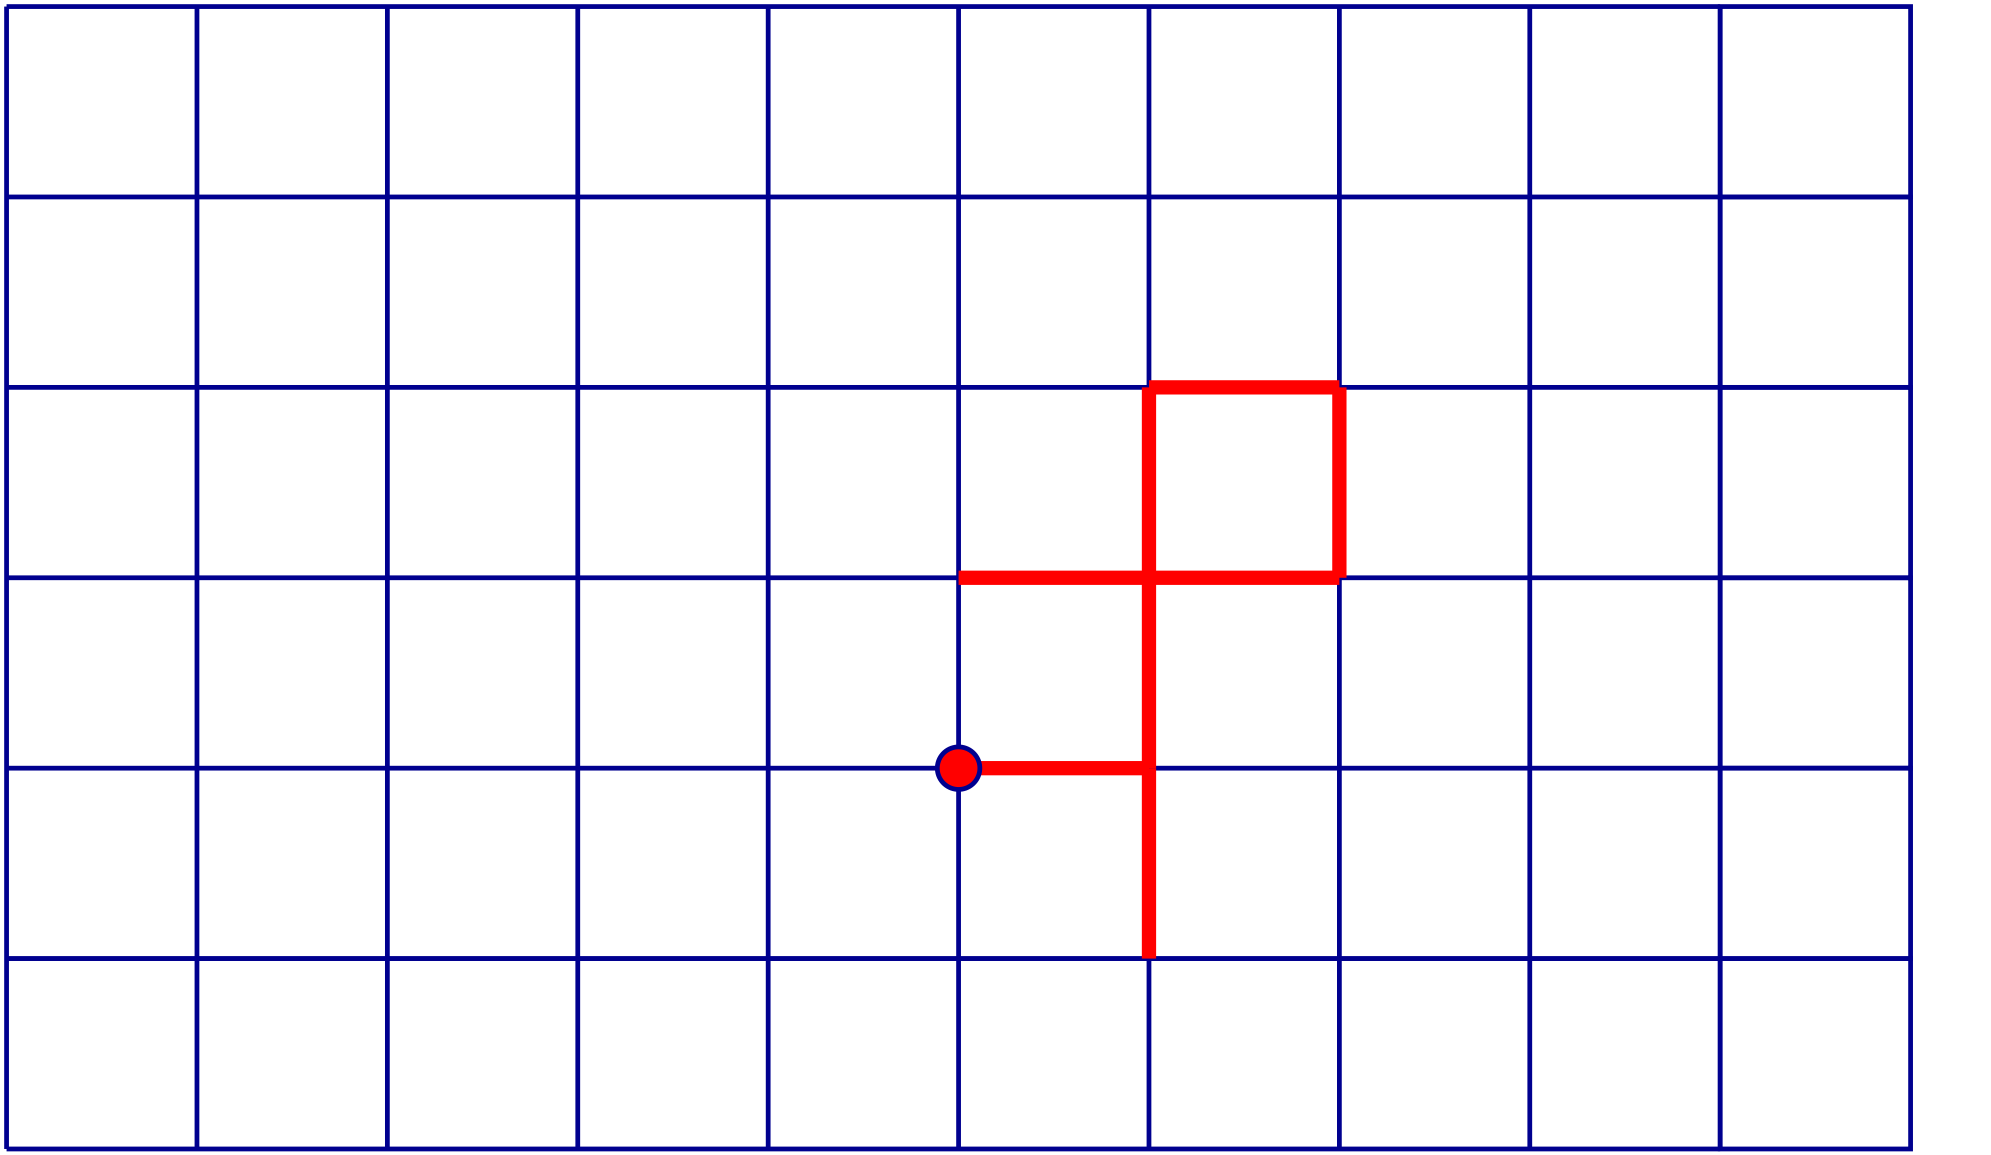

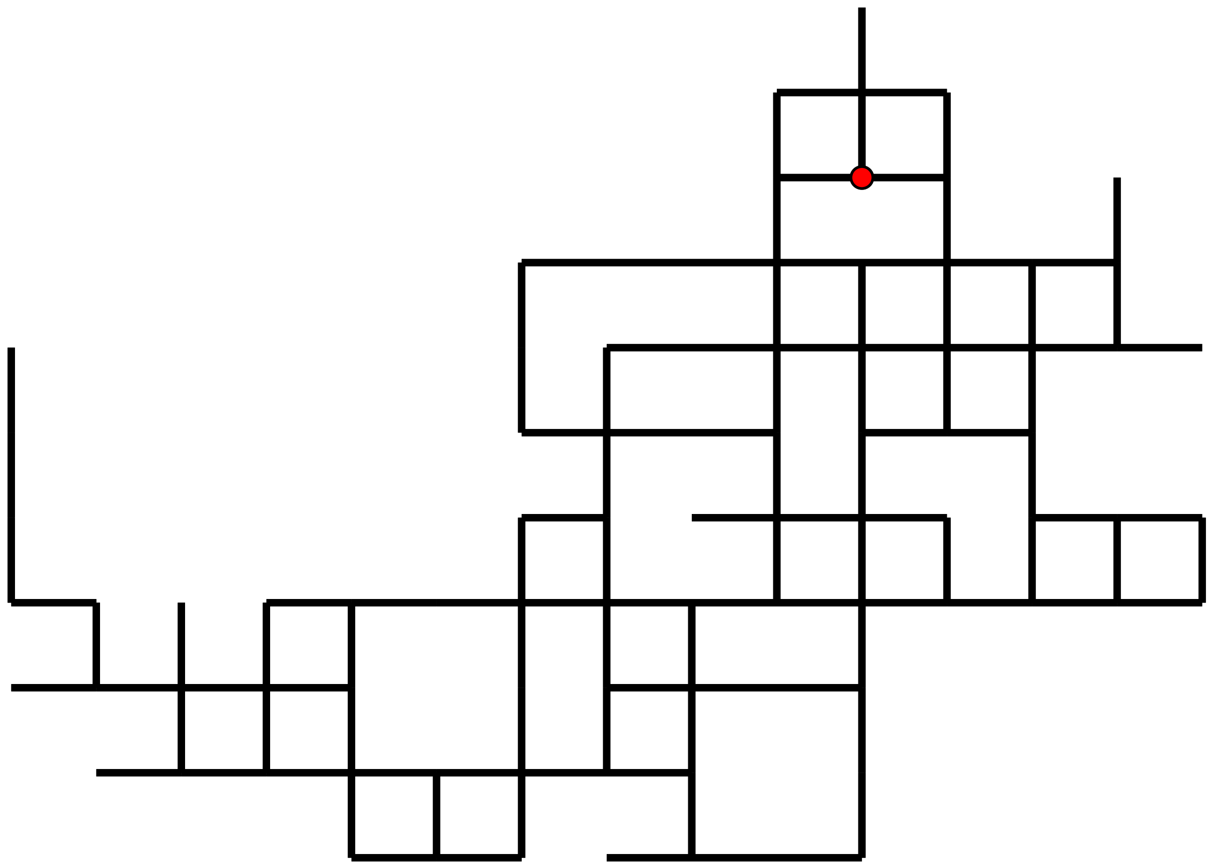

200 step random walk.

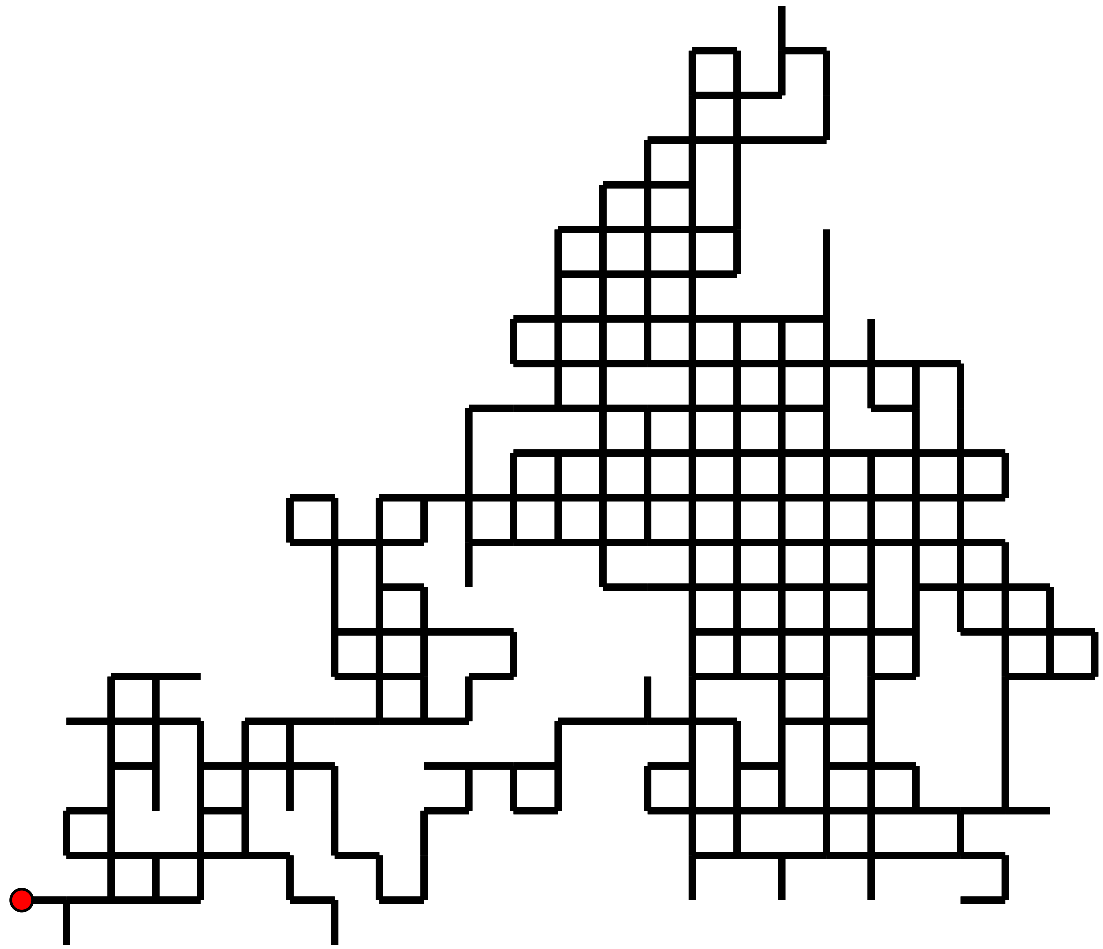

1000 step random walk.

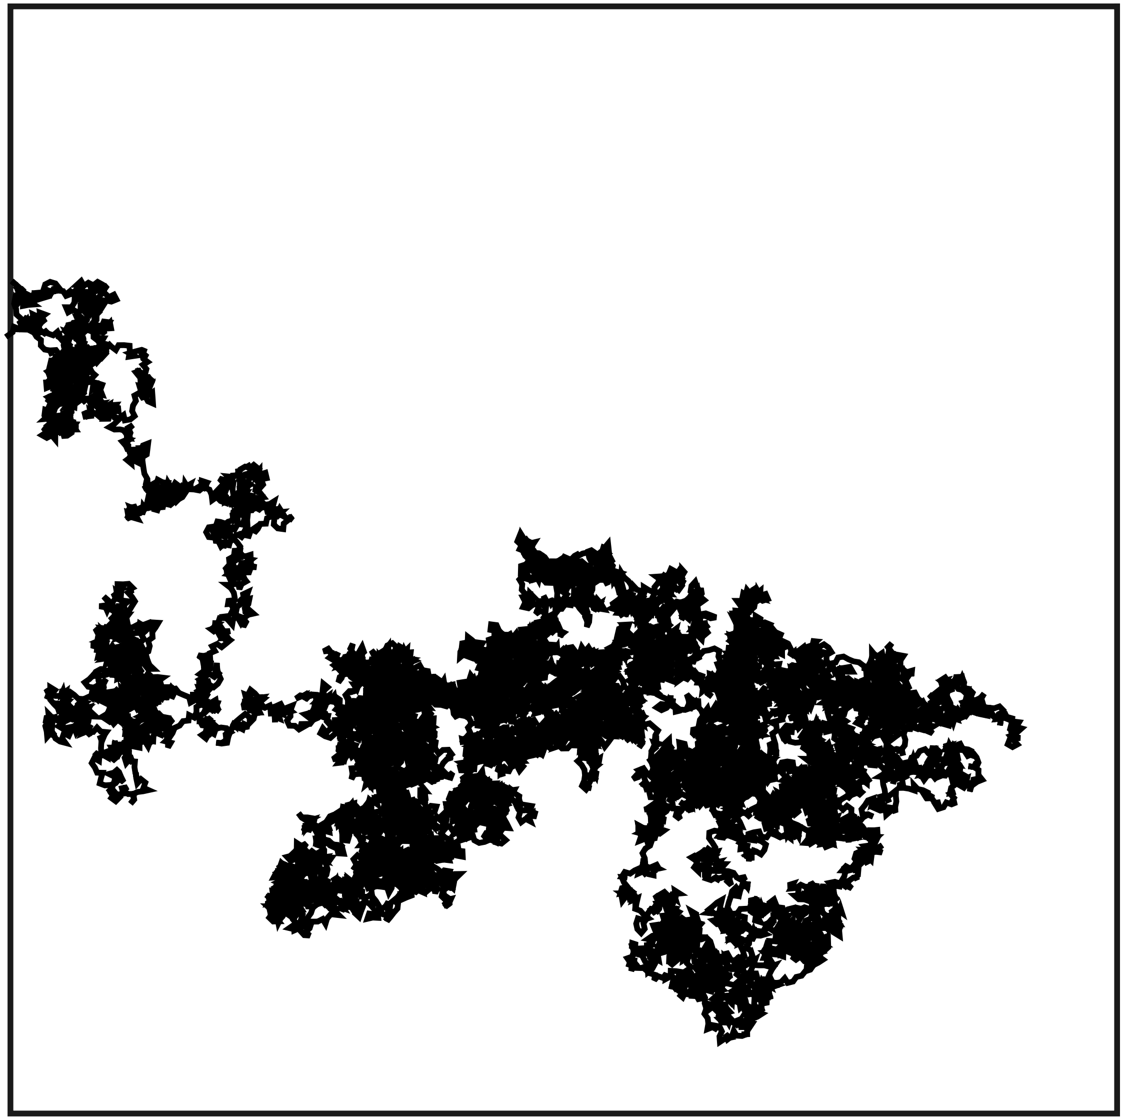

Get Brownian motion in limit.

Harmonic measure = hitting distribution of Brownian motion

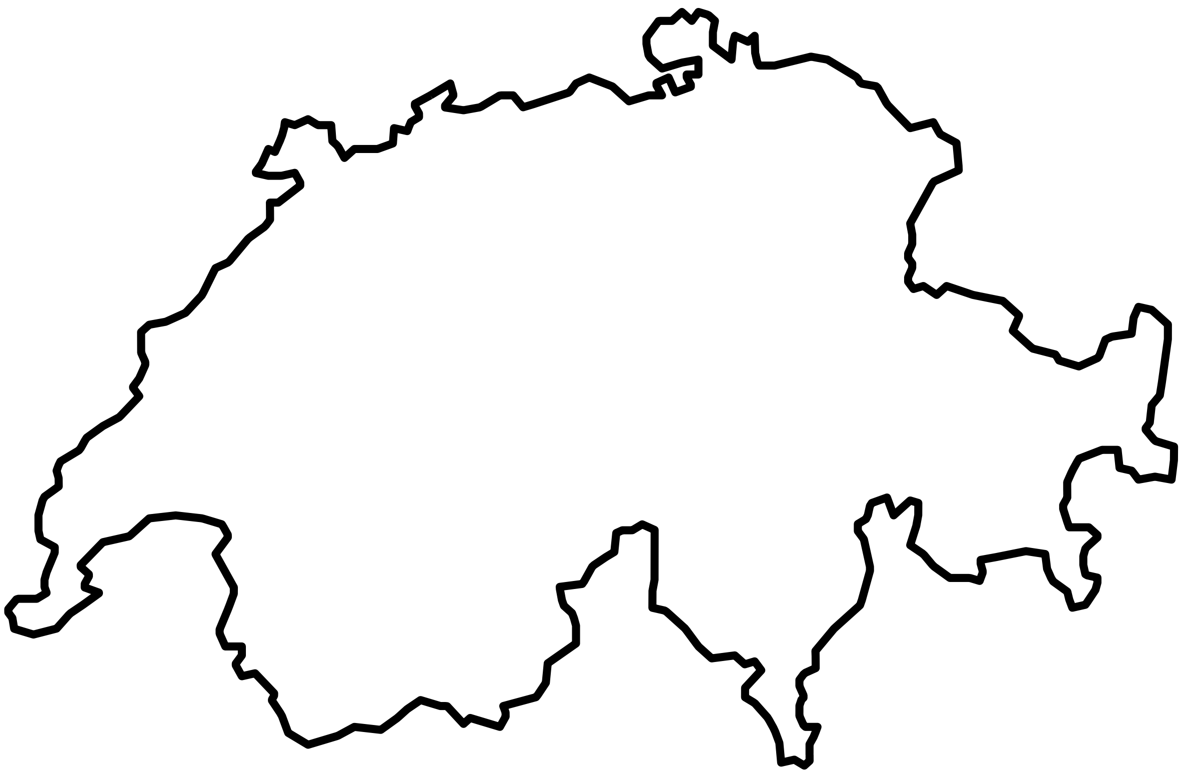

Suppose Ω is a planar Jordan domain.

Harmonic measure = hitting distribution of Brownian motion

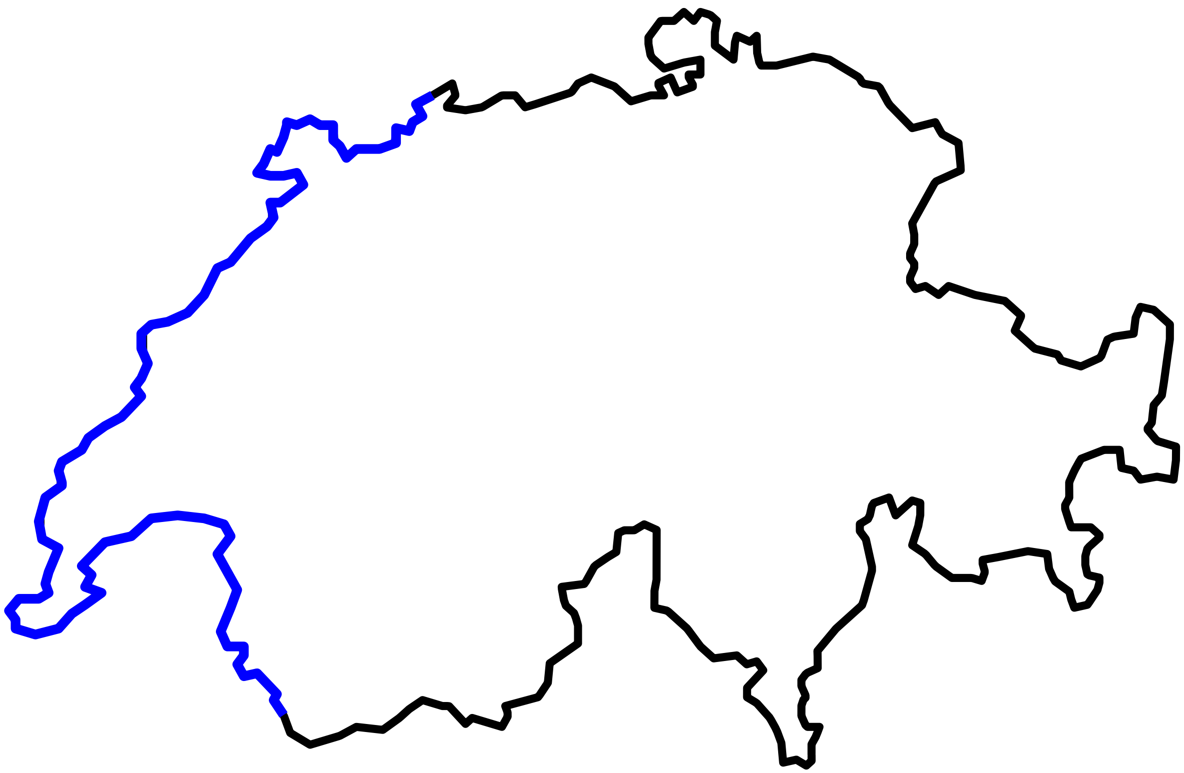

Let E be a subset of the boundary, $\partial\Omega$.

Harmonic measure = hitting distribution of Brownian motion

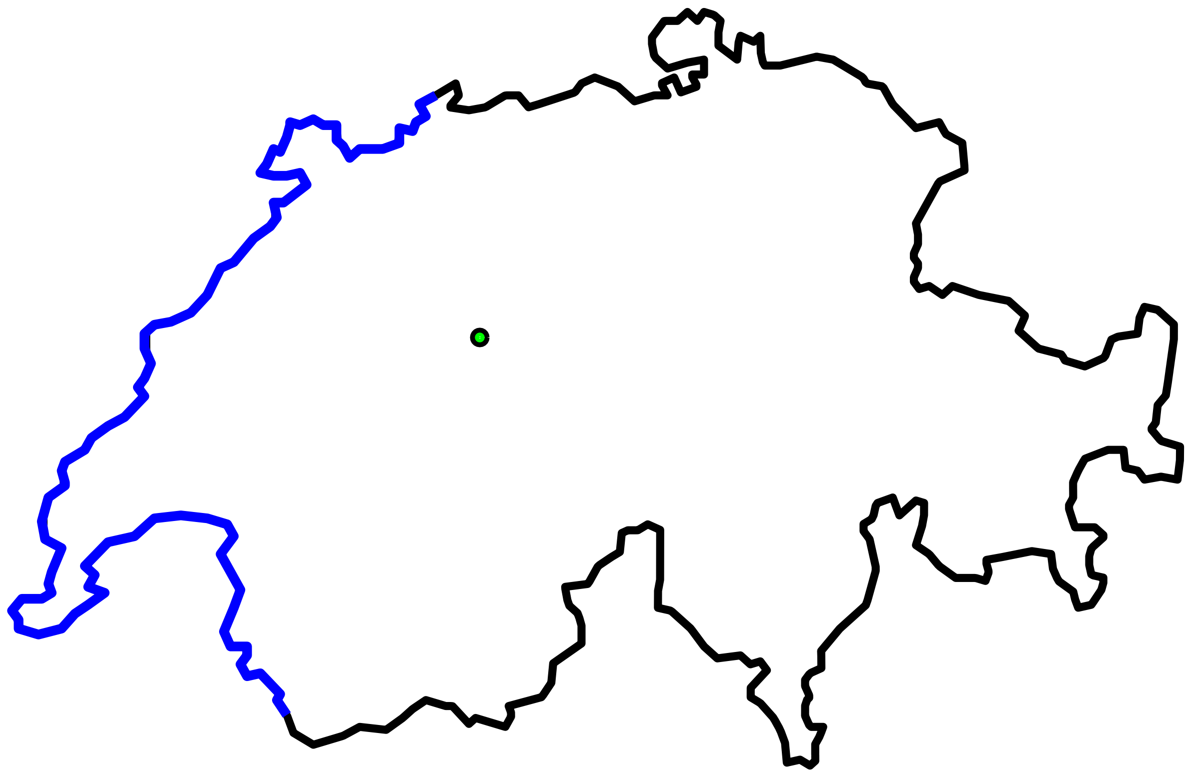

Choose an interior point $z \in \Omega$.

Harmonic measure = hitting distribution of Brownian motion

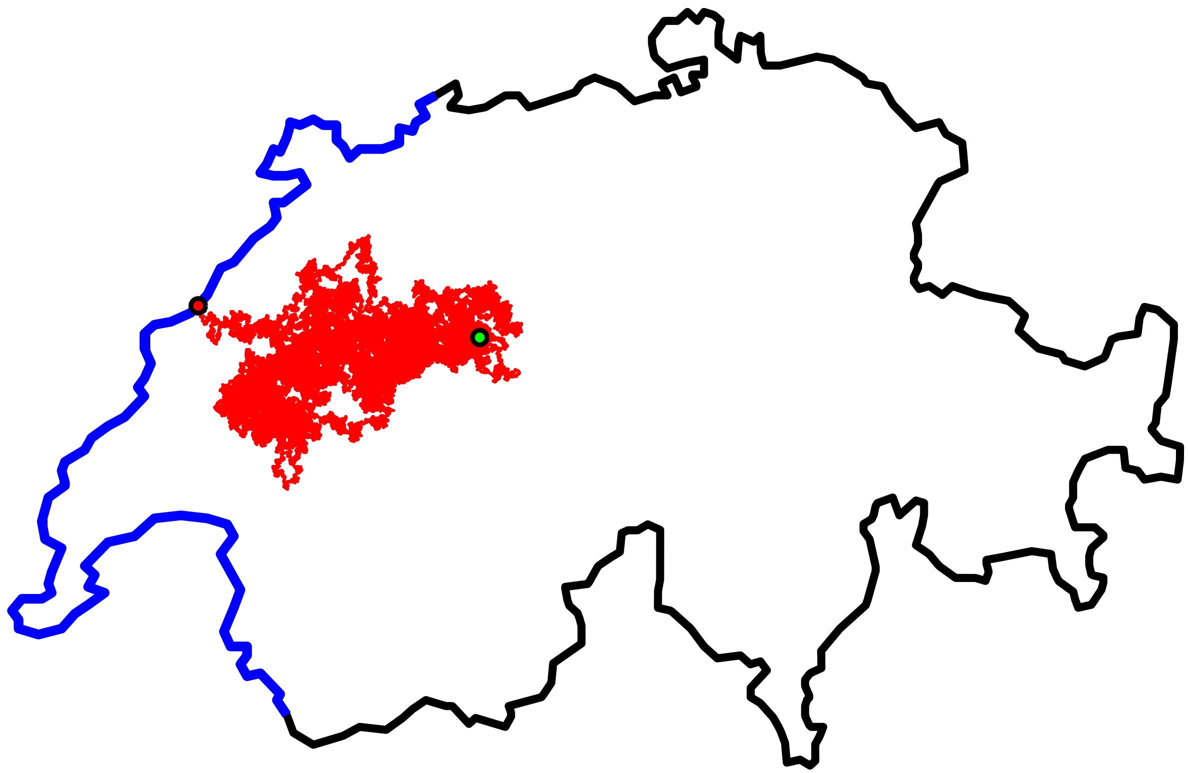

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

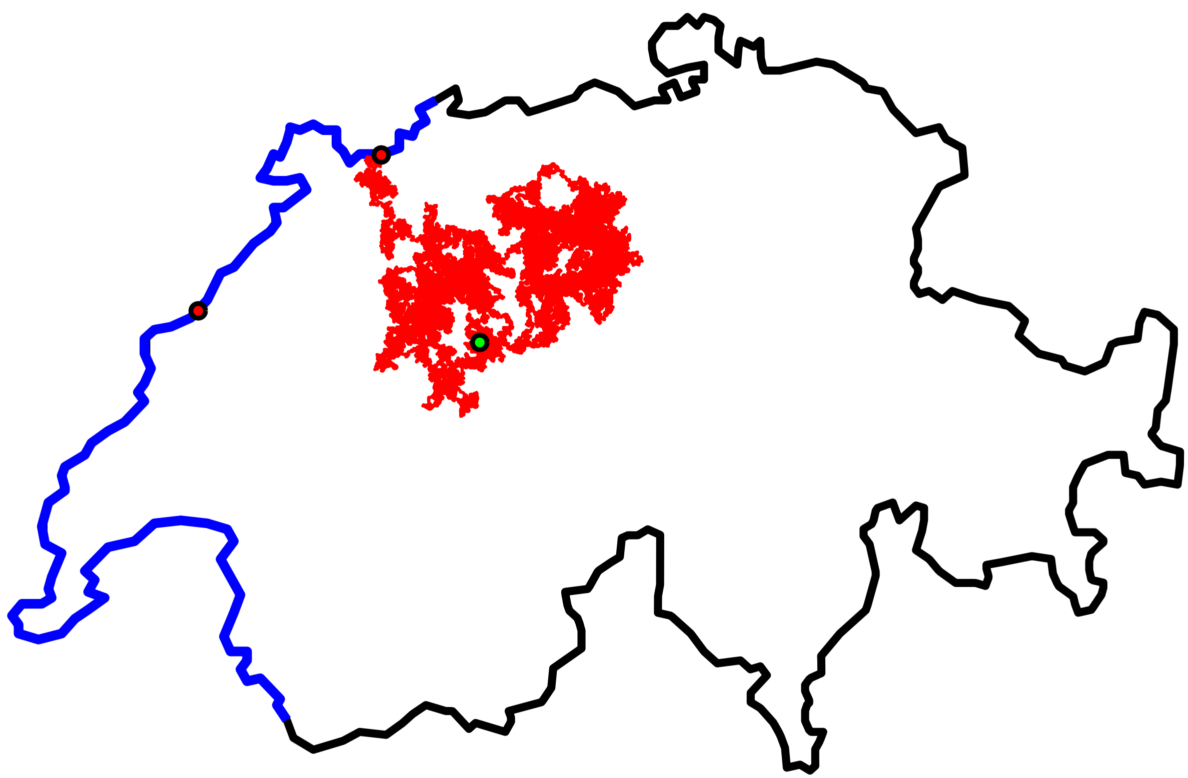

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

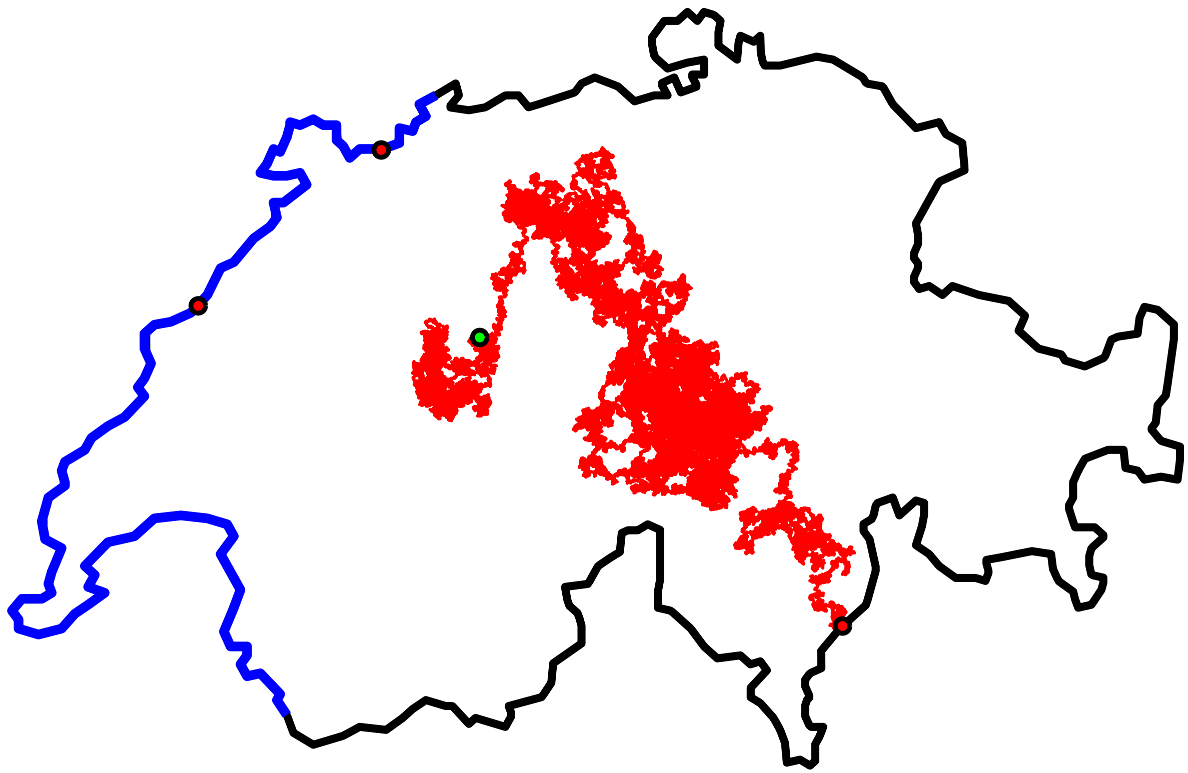

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

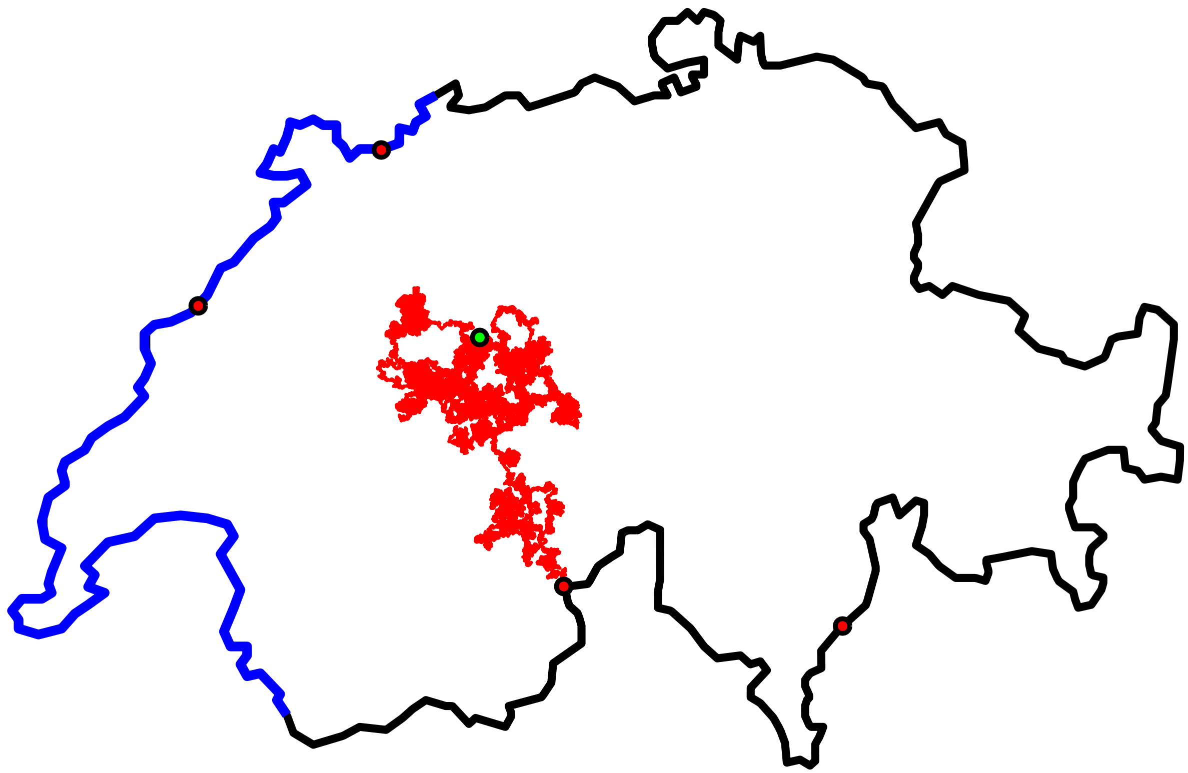

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

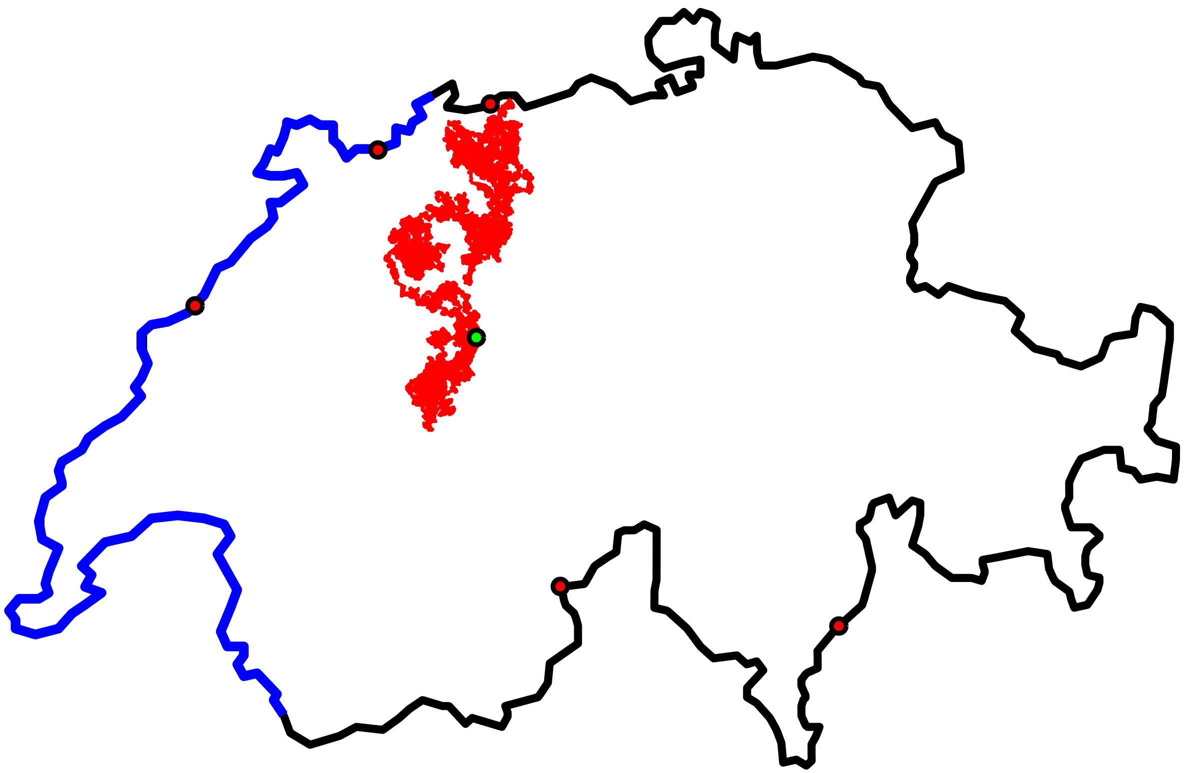

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

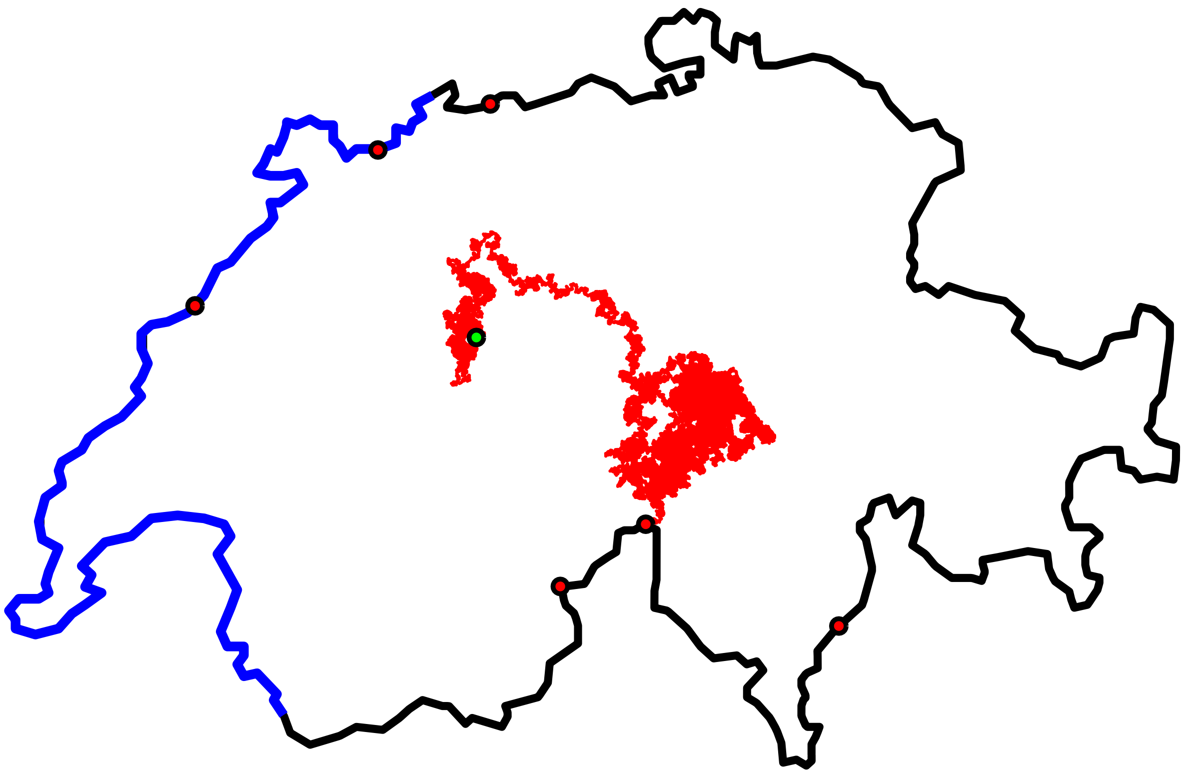

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

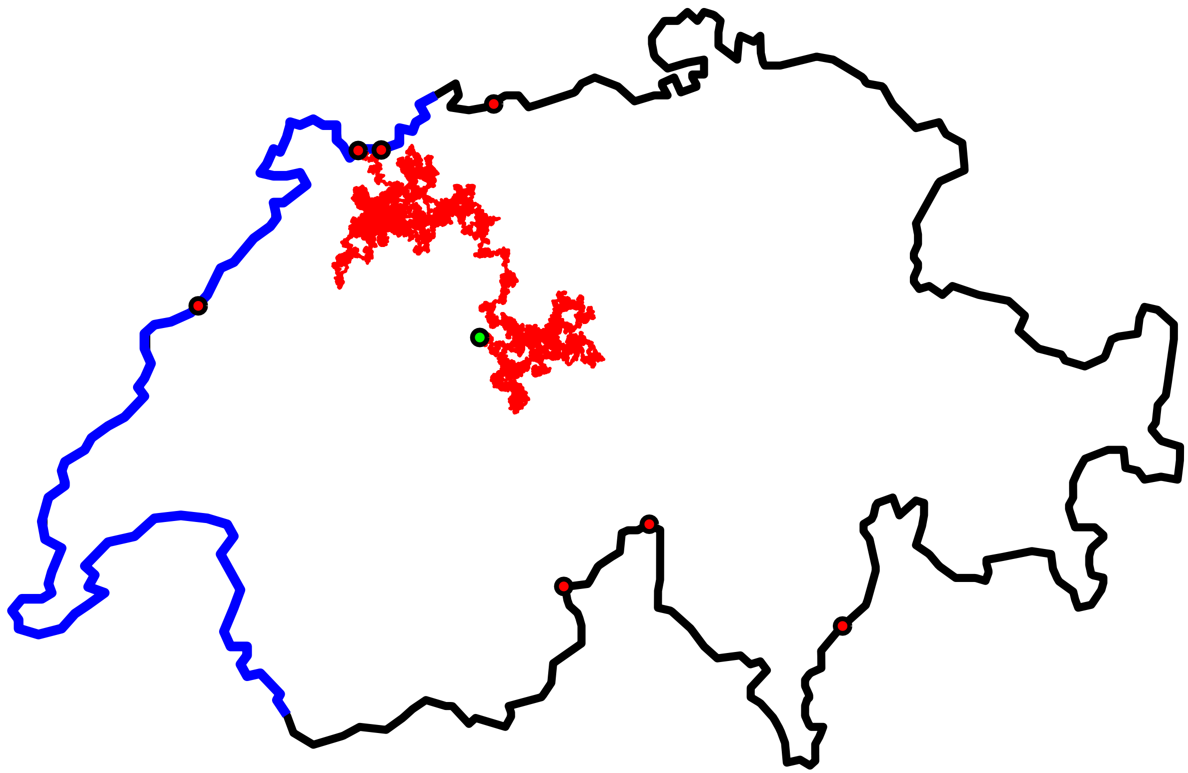

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

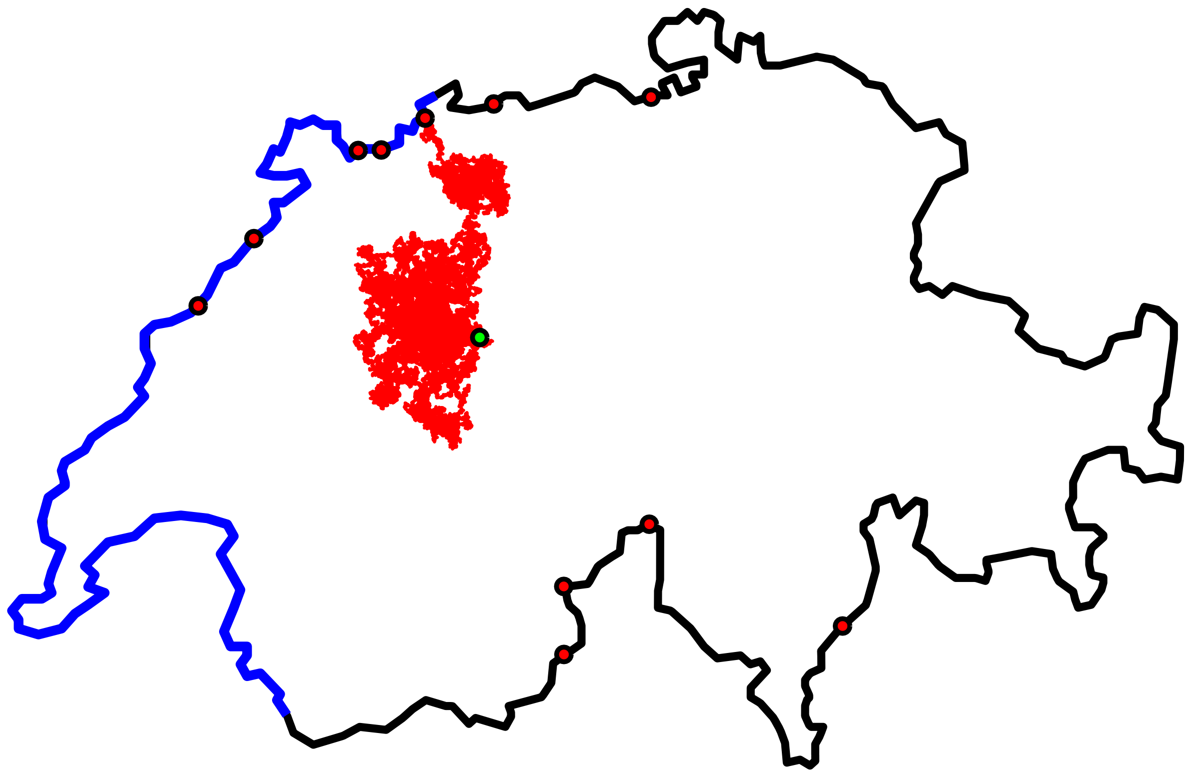

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

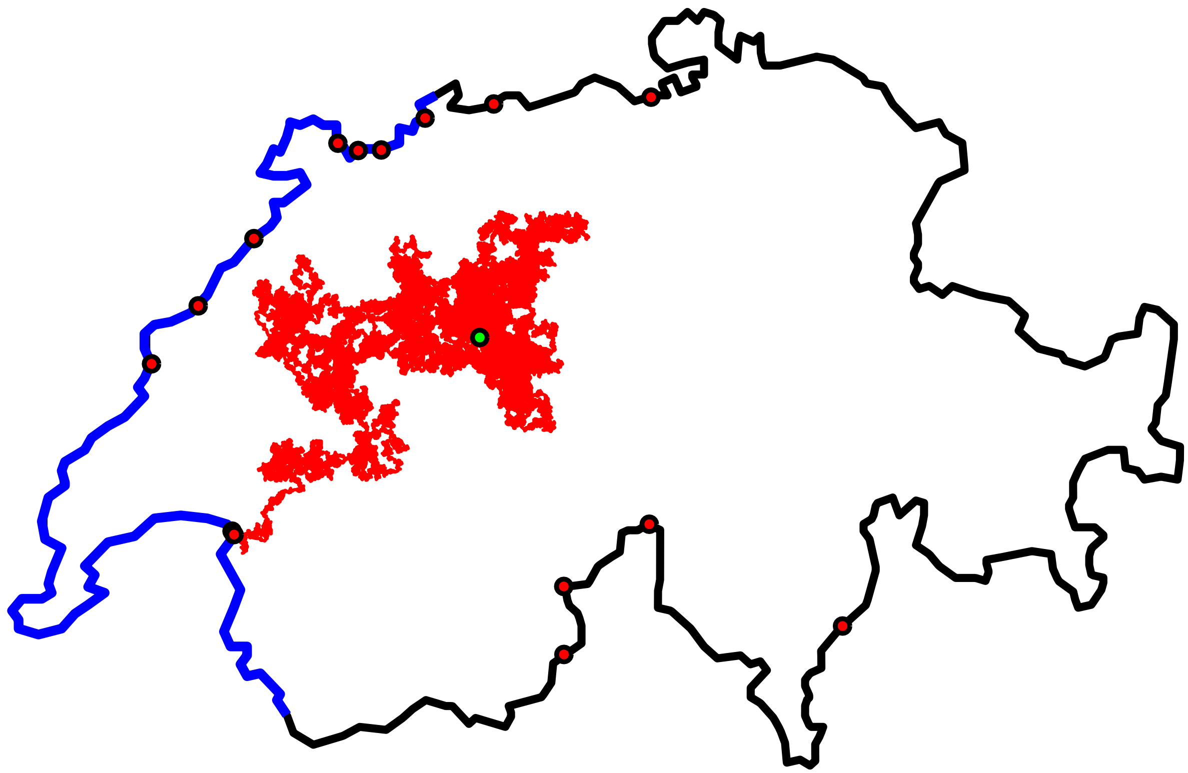

$\omega(z, E, \Omega)$ = probability a particle started at z first hits $\partial\Omega$ in E.

Harmonic measure = hitting distribution of Brownian motion

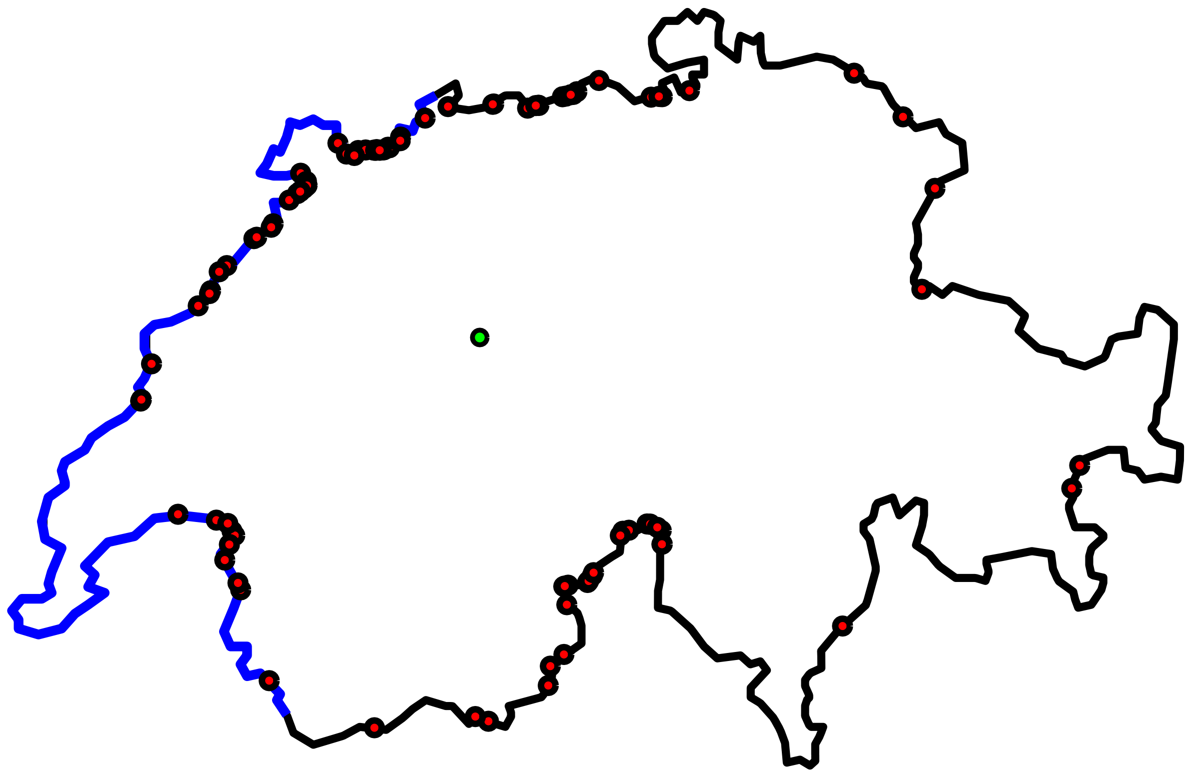

$\omega(z, E, \Omega) \approx 35/100$. What if we move starting point z?

Harmonic measure = hitting distribution of Brownian motion

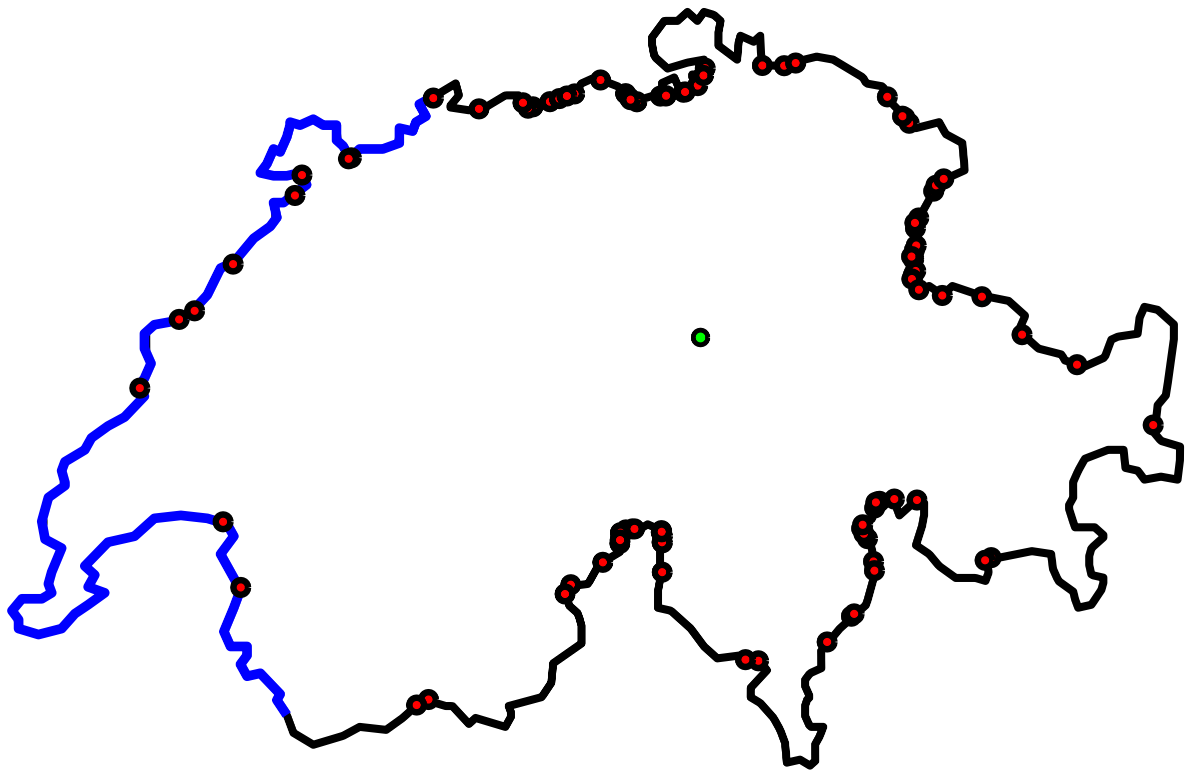

Different z gives $\omega(z, E, \Omega) \approx 10/100$.
Slow to compute each path ($\simeq \epsilon^{-2}$ steps on ϵ-grid). Can we go faster?

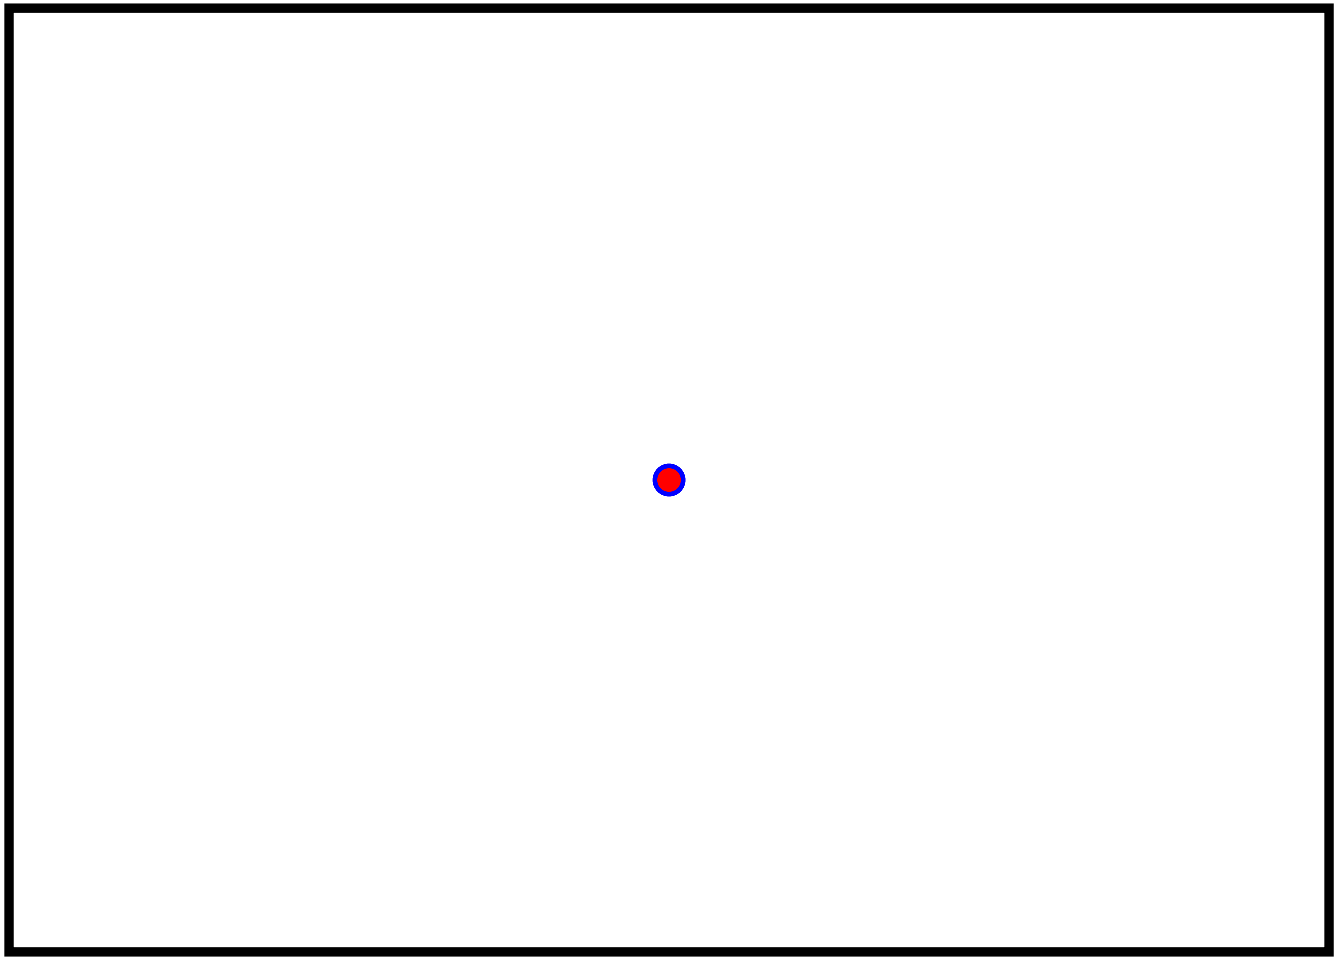

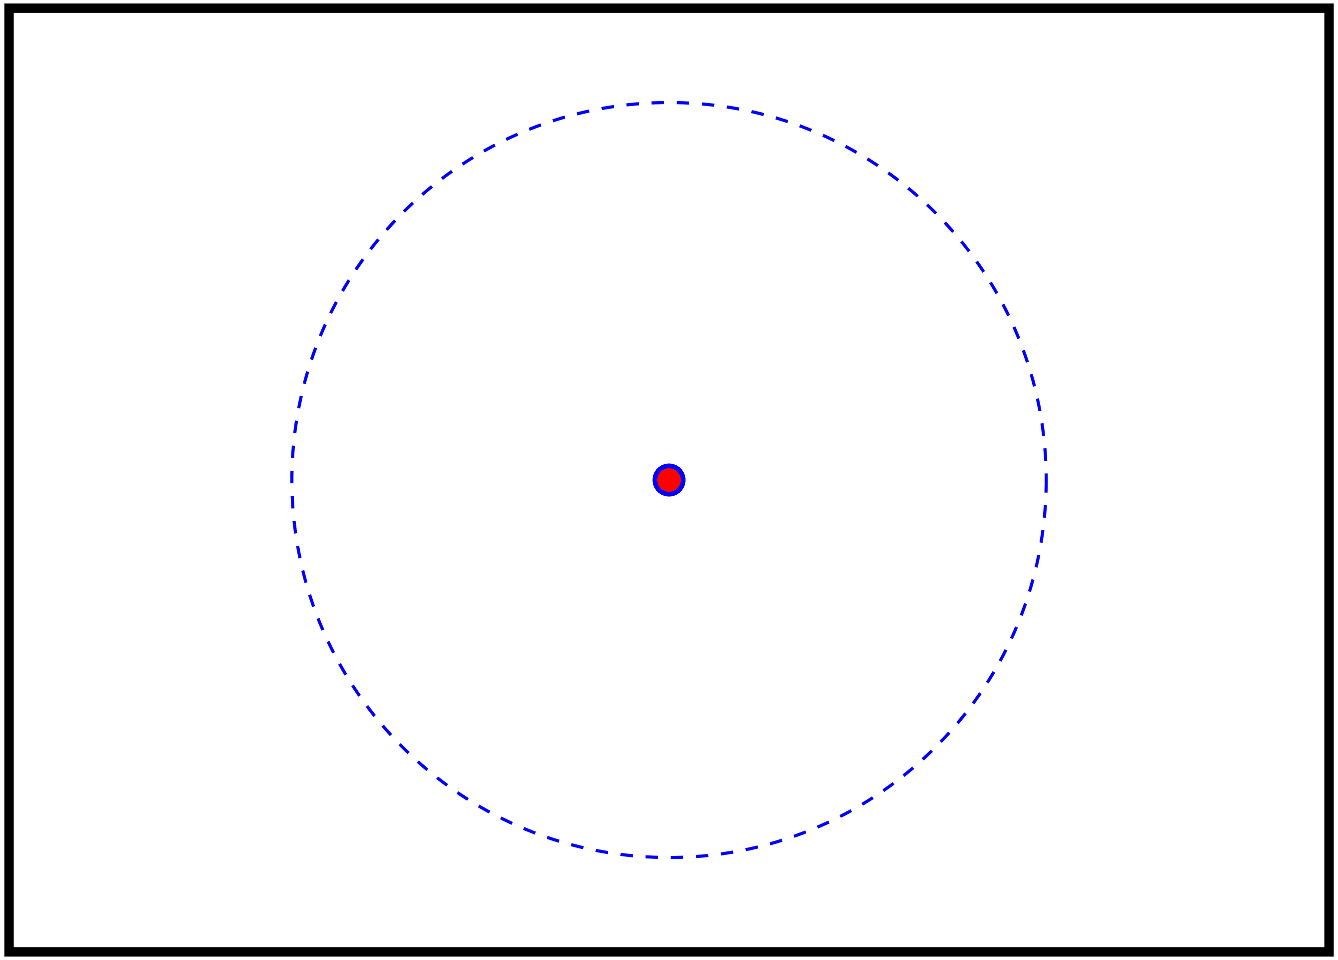

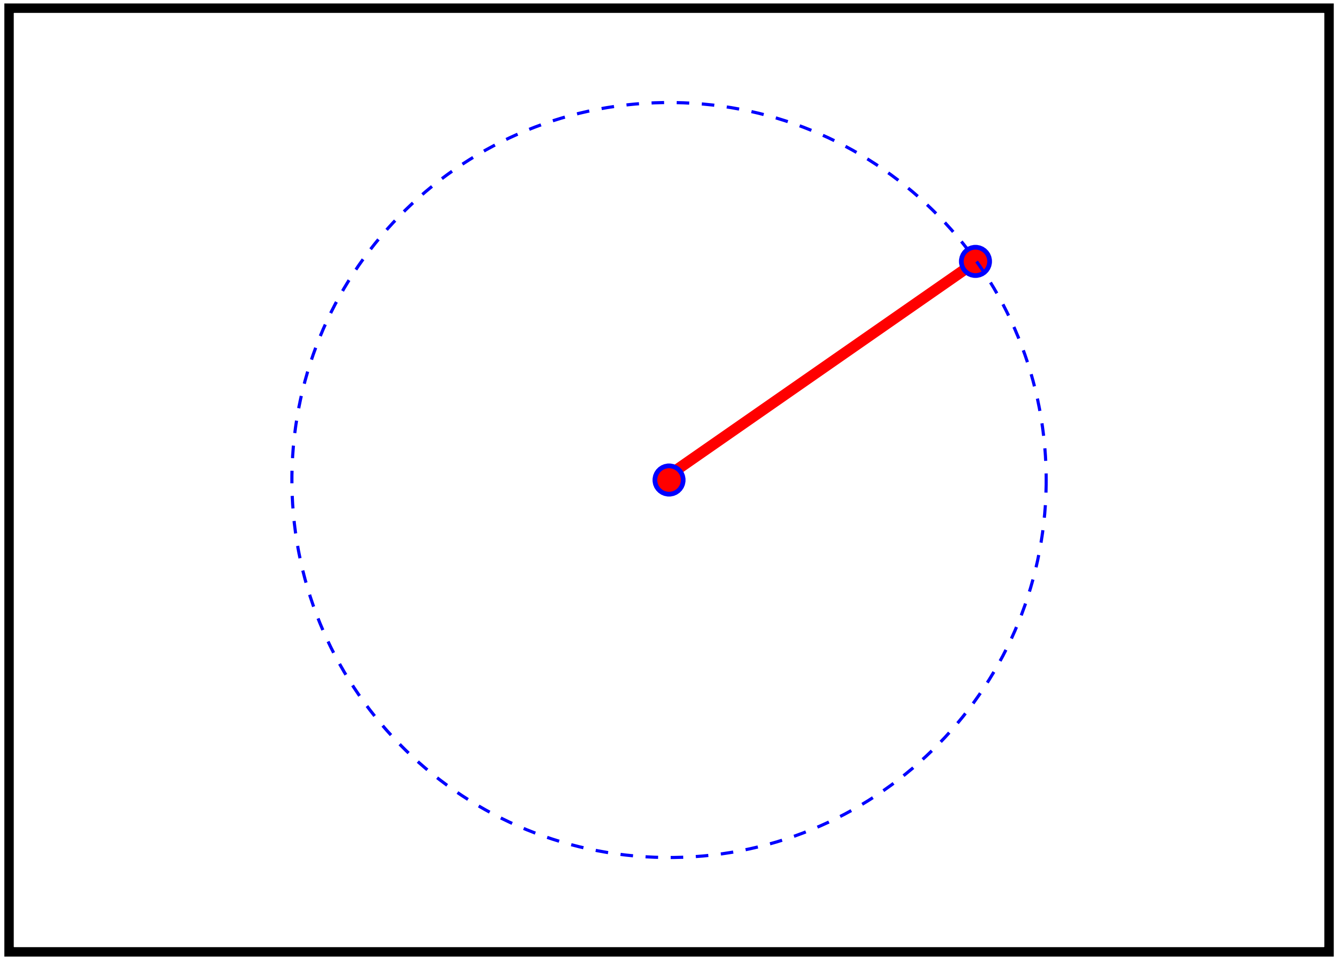

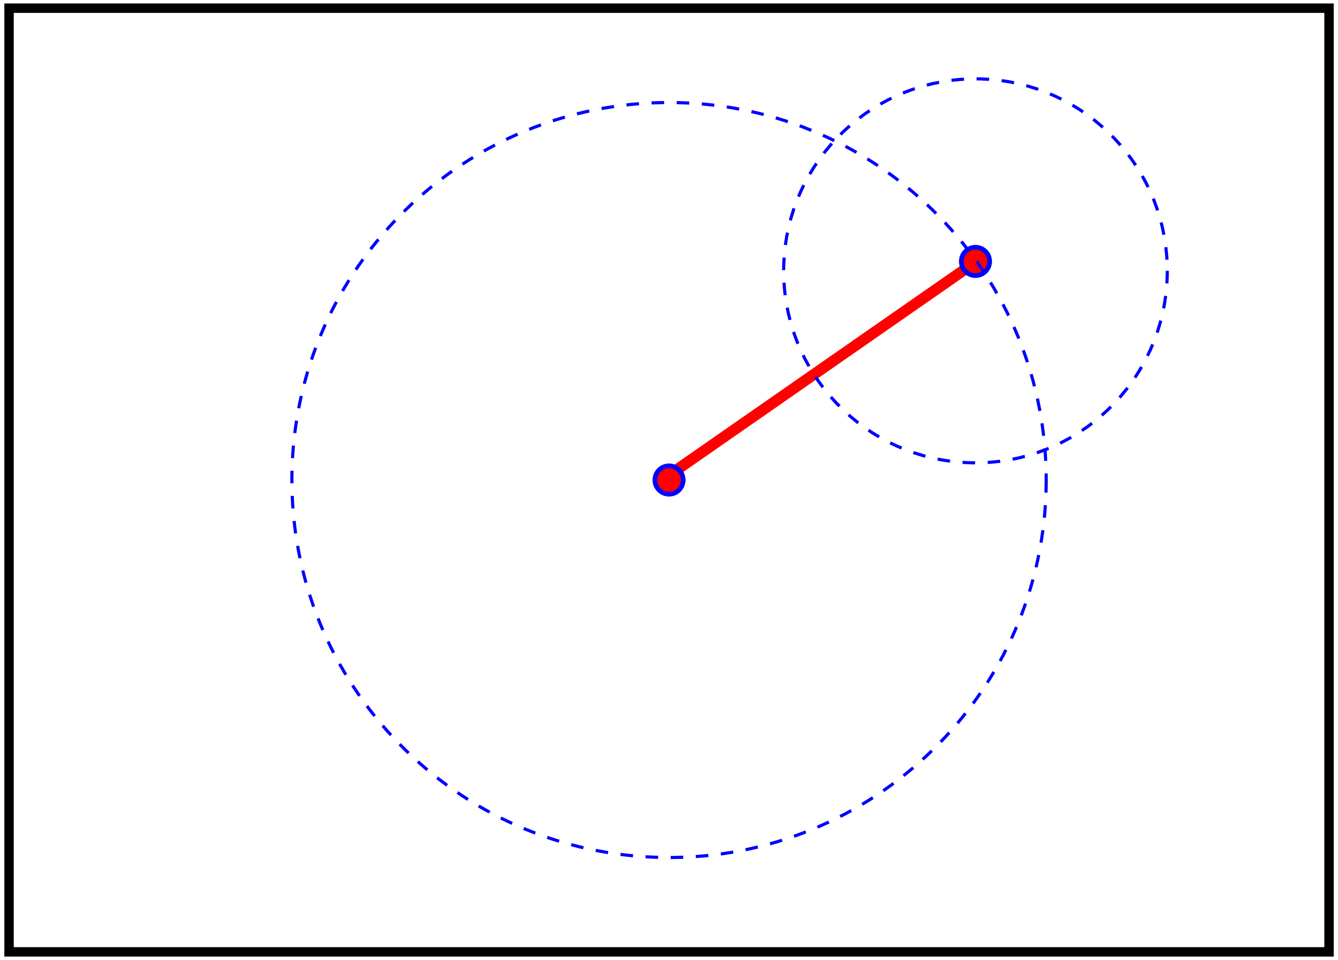

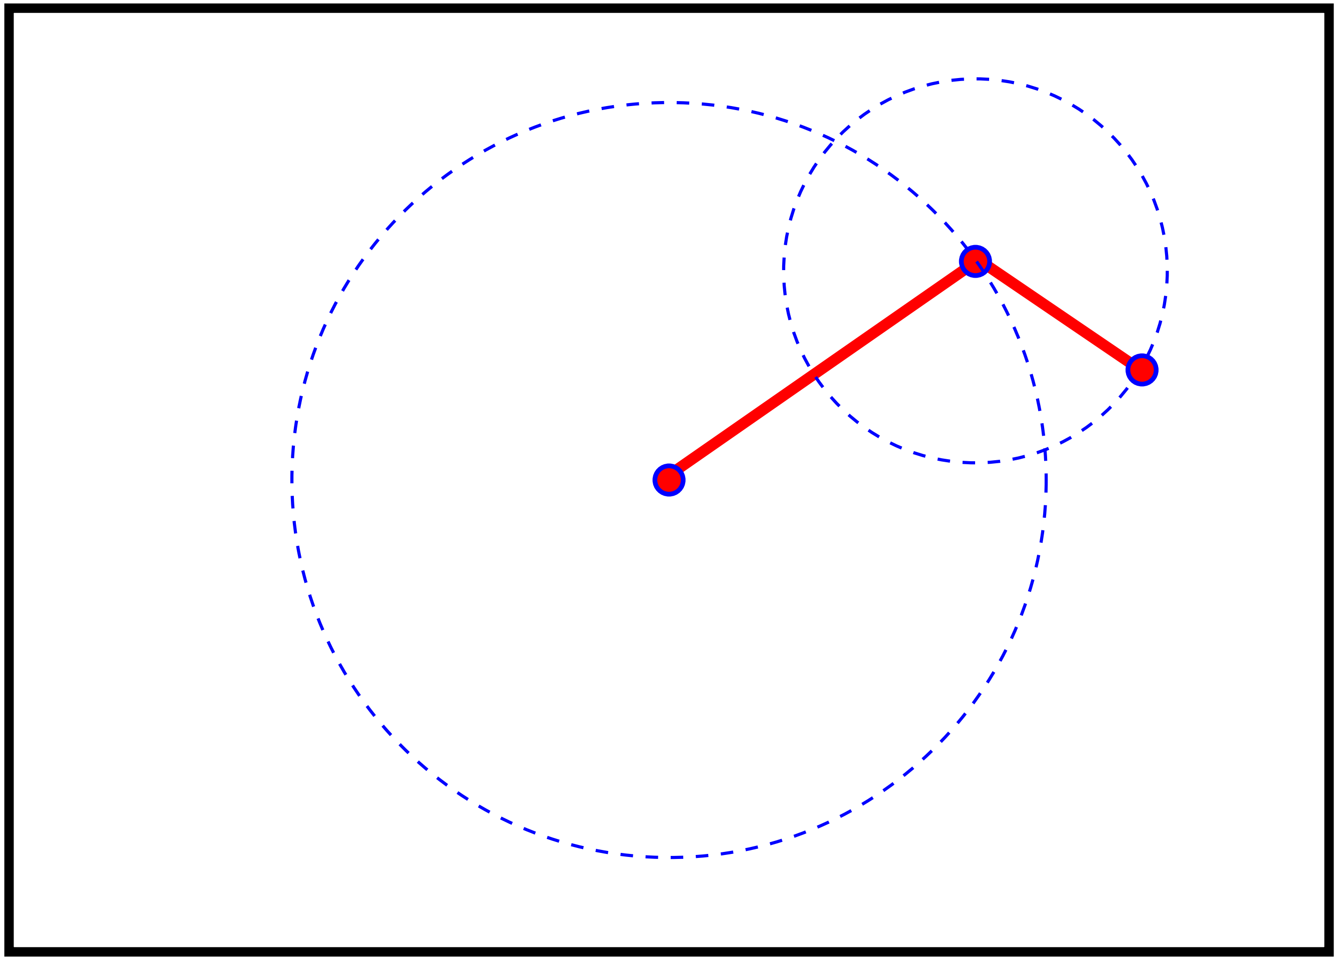

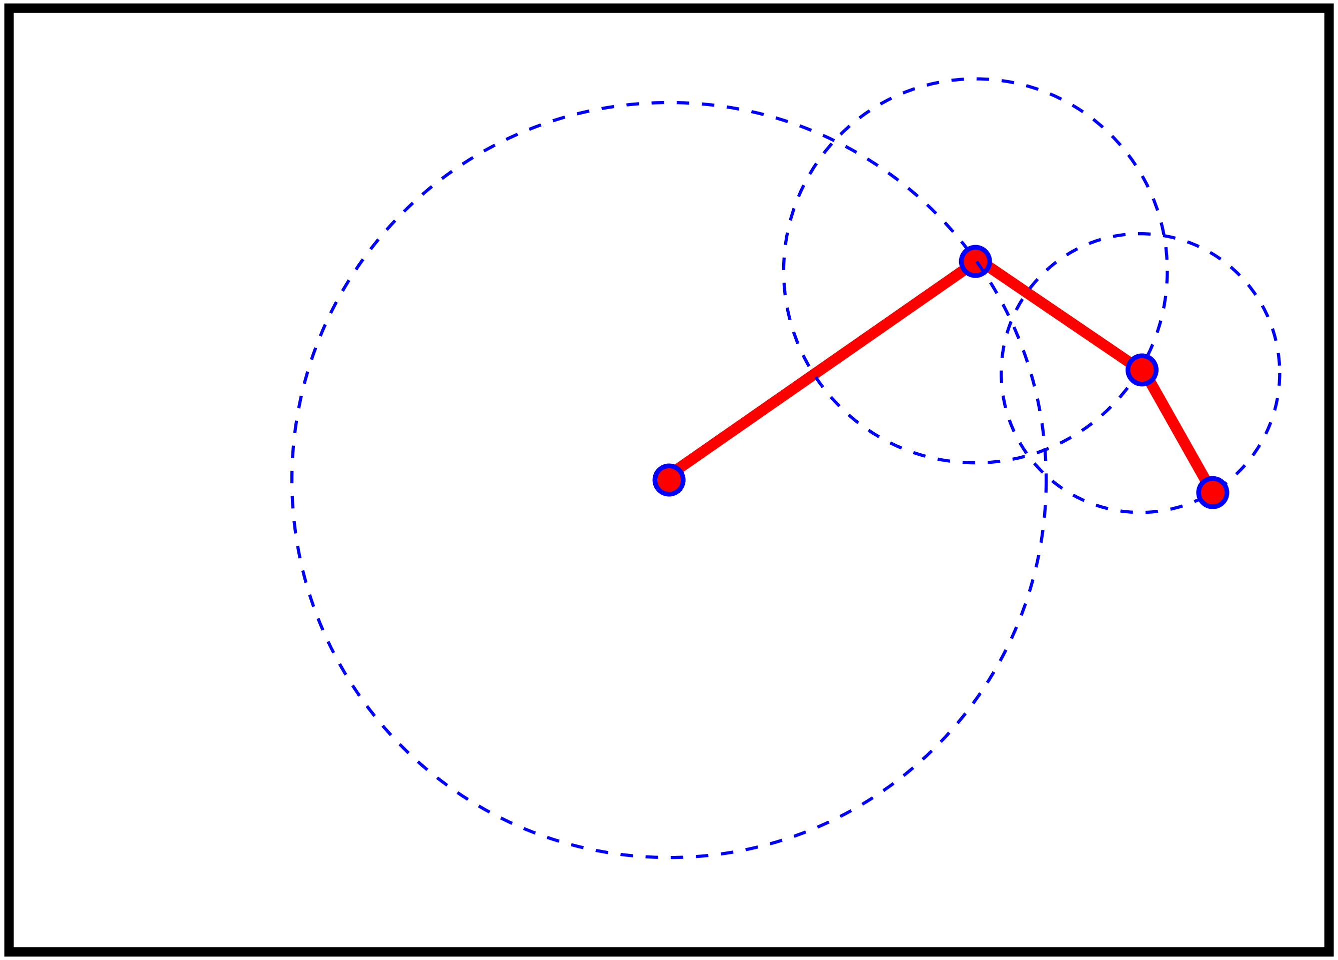

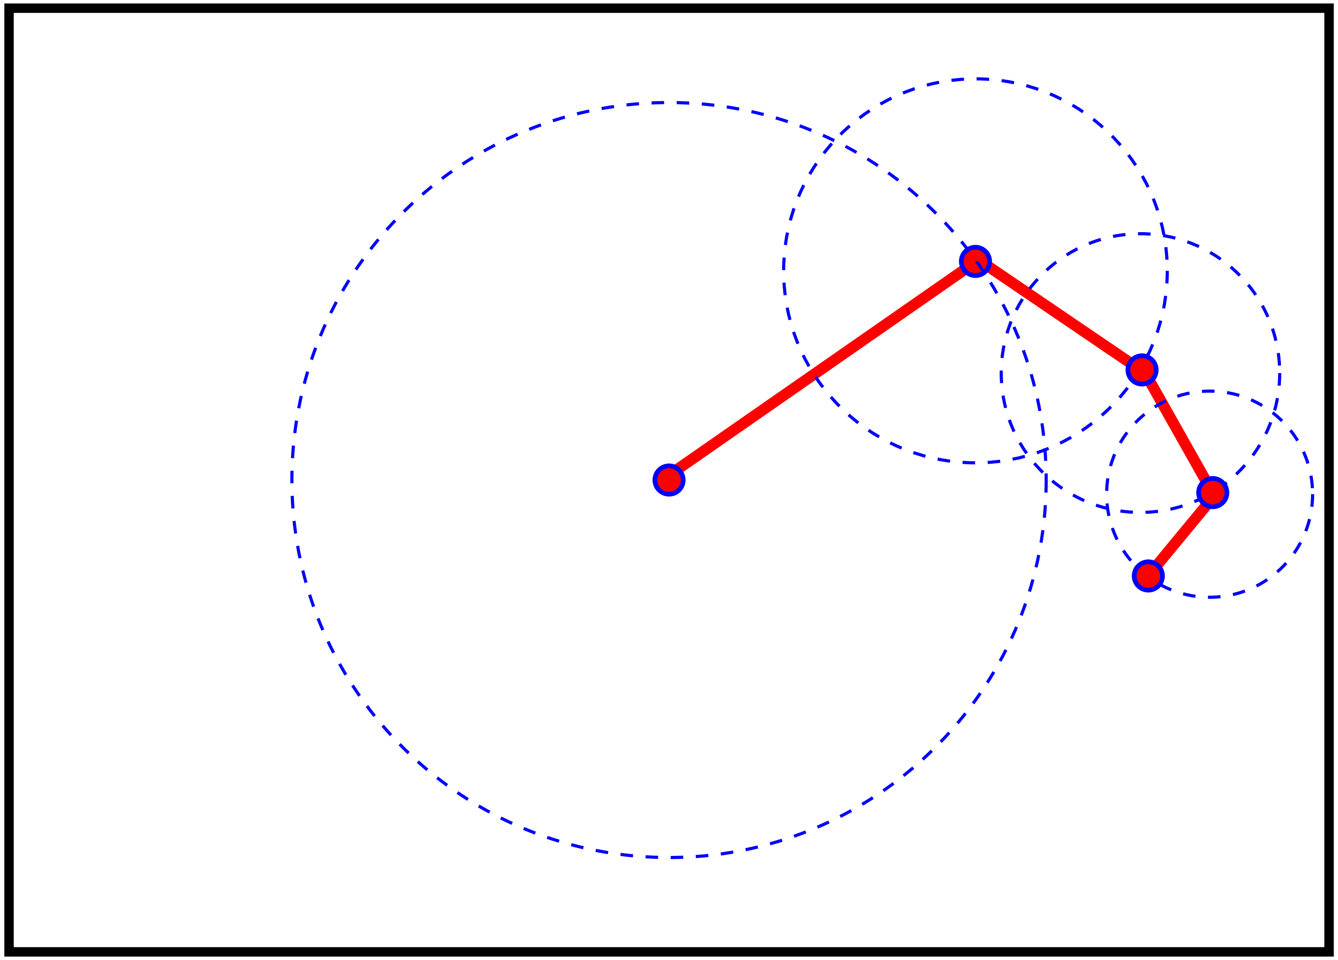

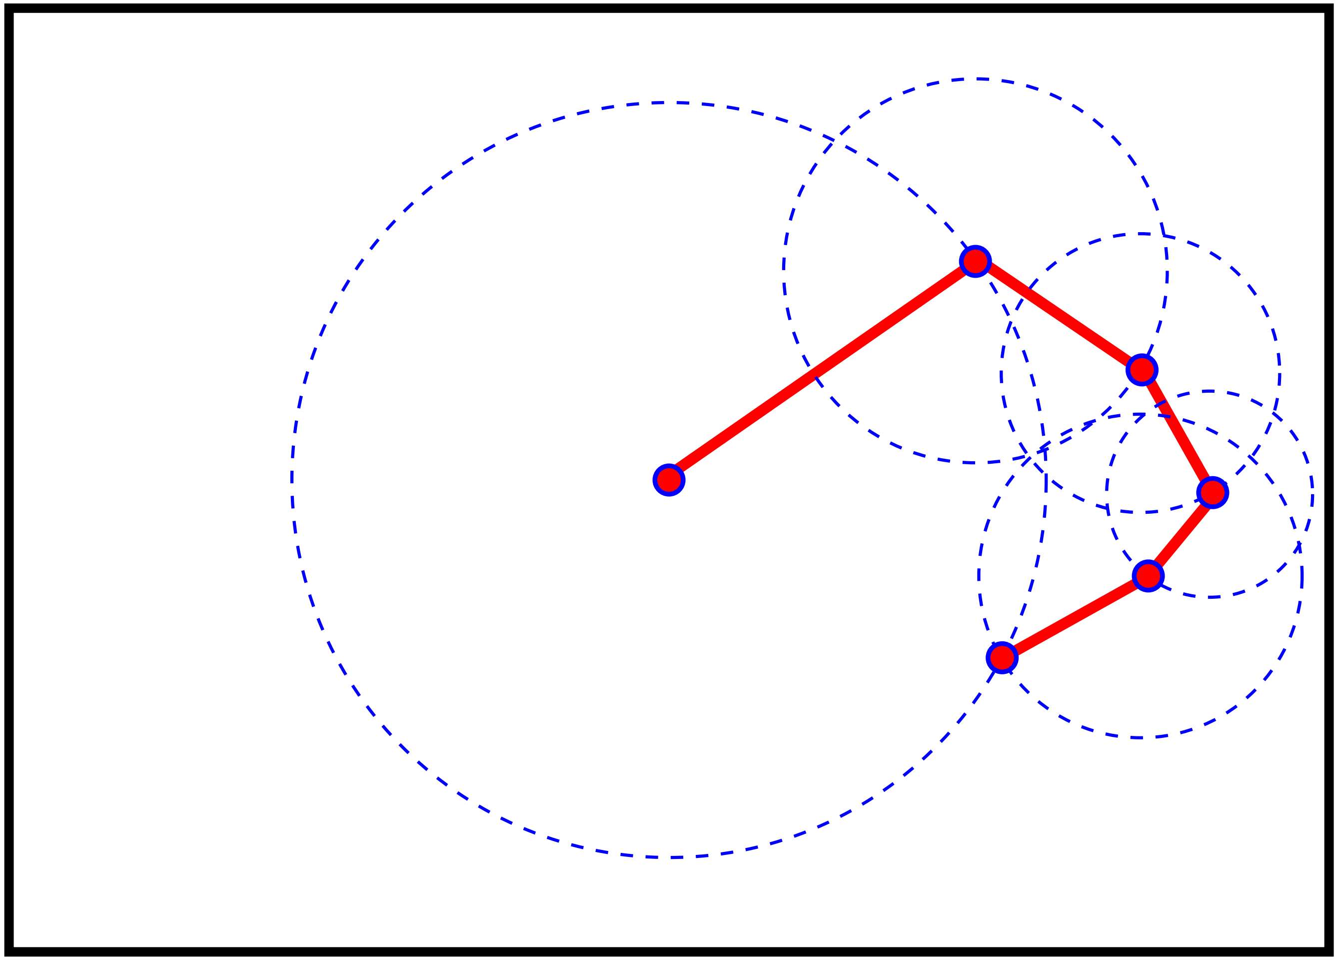

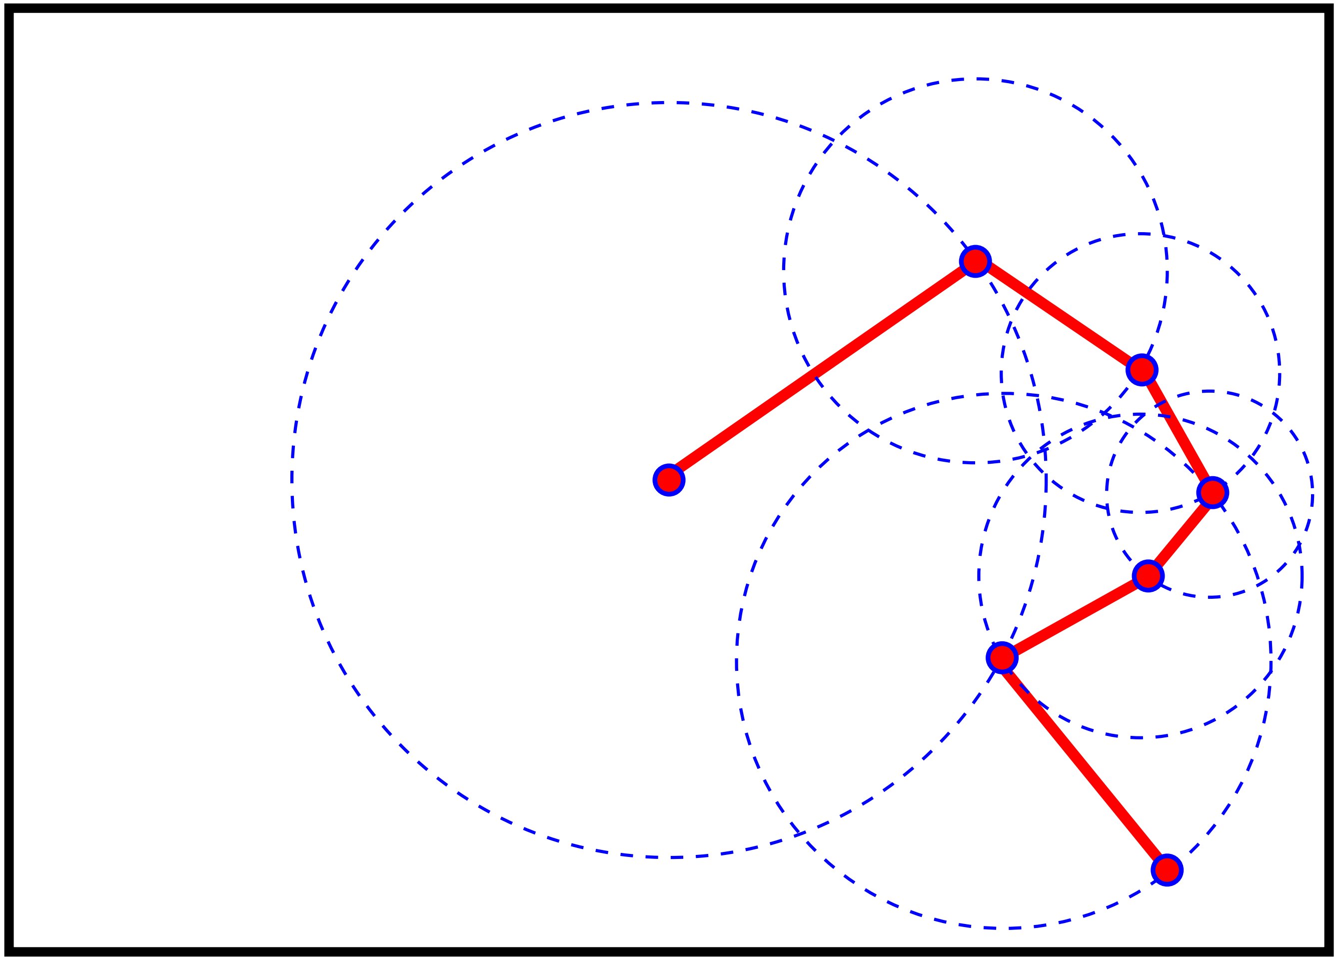

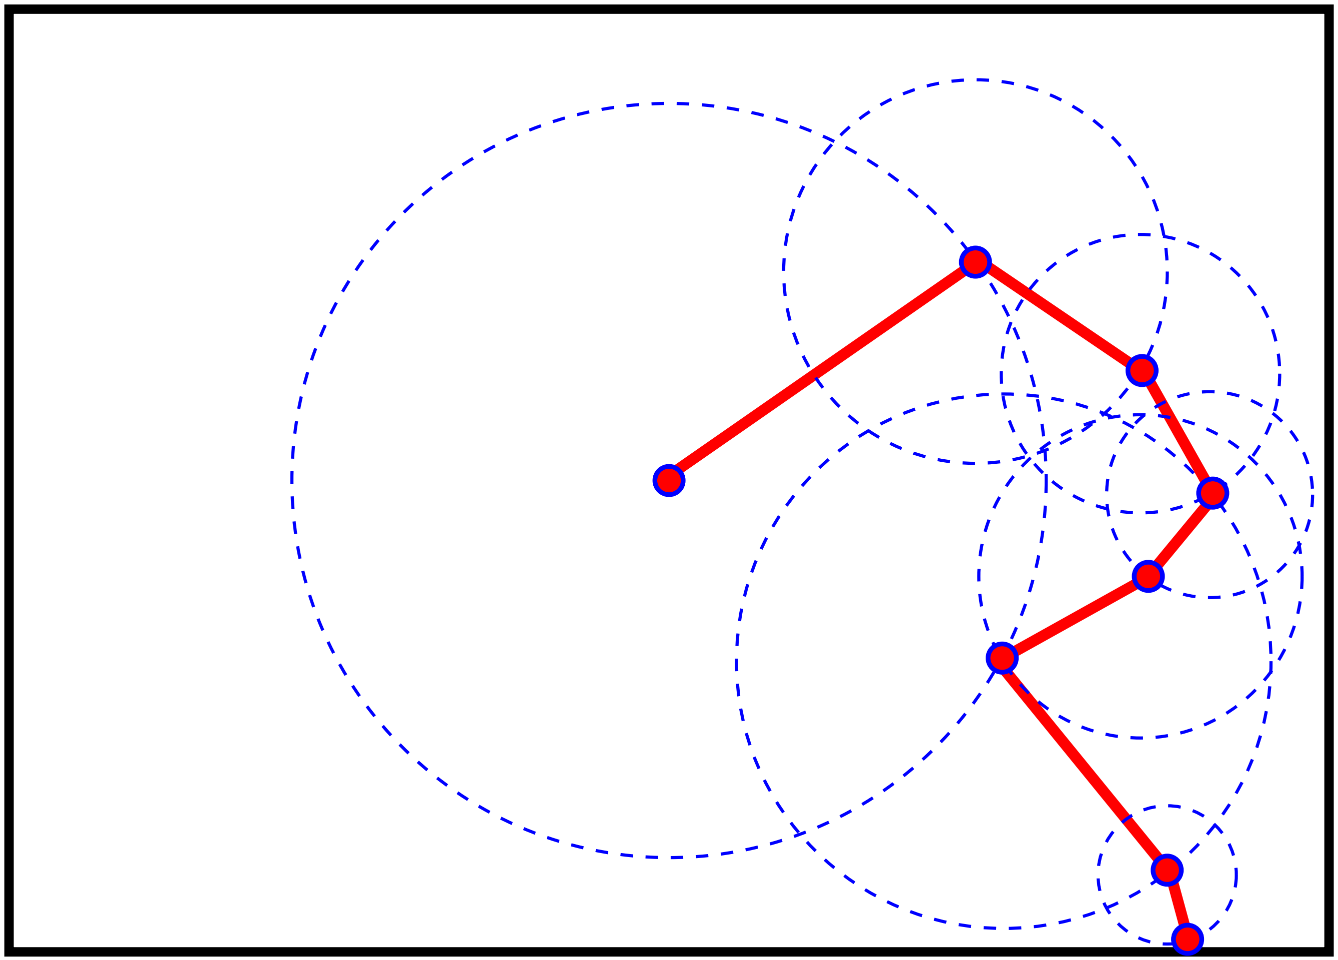

This walk will get within ϵ of boundary in about $O(\log 1/\epsilon)$ steps.
Takes about ϵ^{-2} steps on an ϵ-grid.

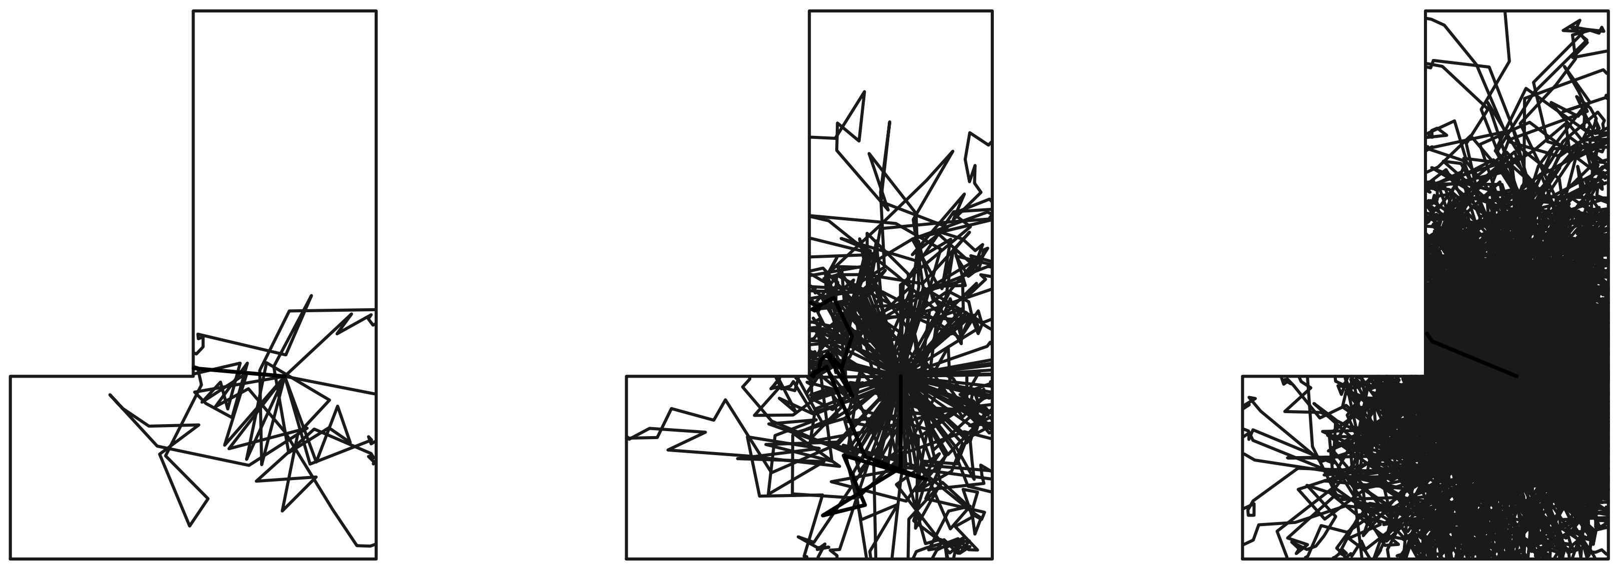

10, 100 and 1000 steps of walk inside a polygon. This illustrates the difficulty Brownian motion has in penetrating narrow corridors.

This is called the crowding phenomena (name to be explained later).

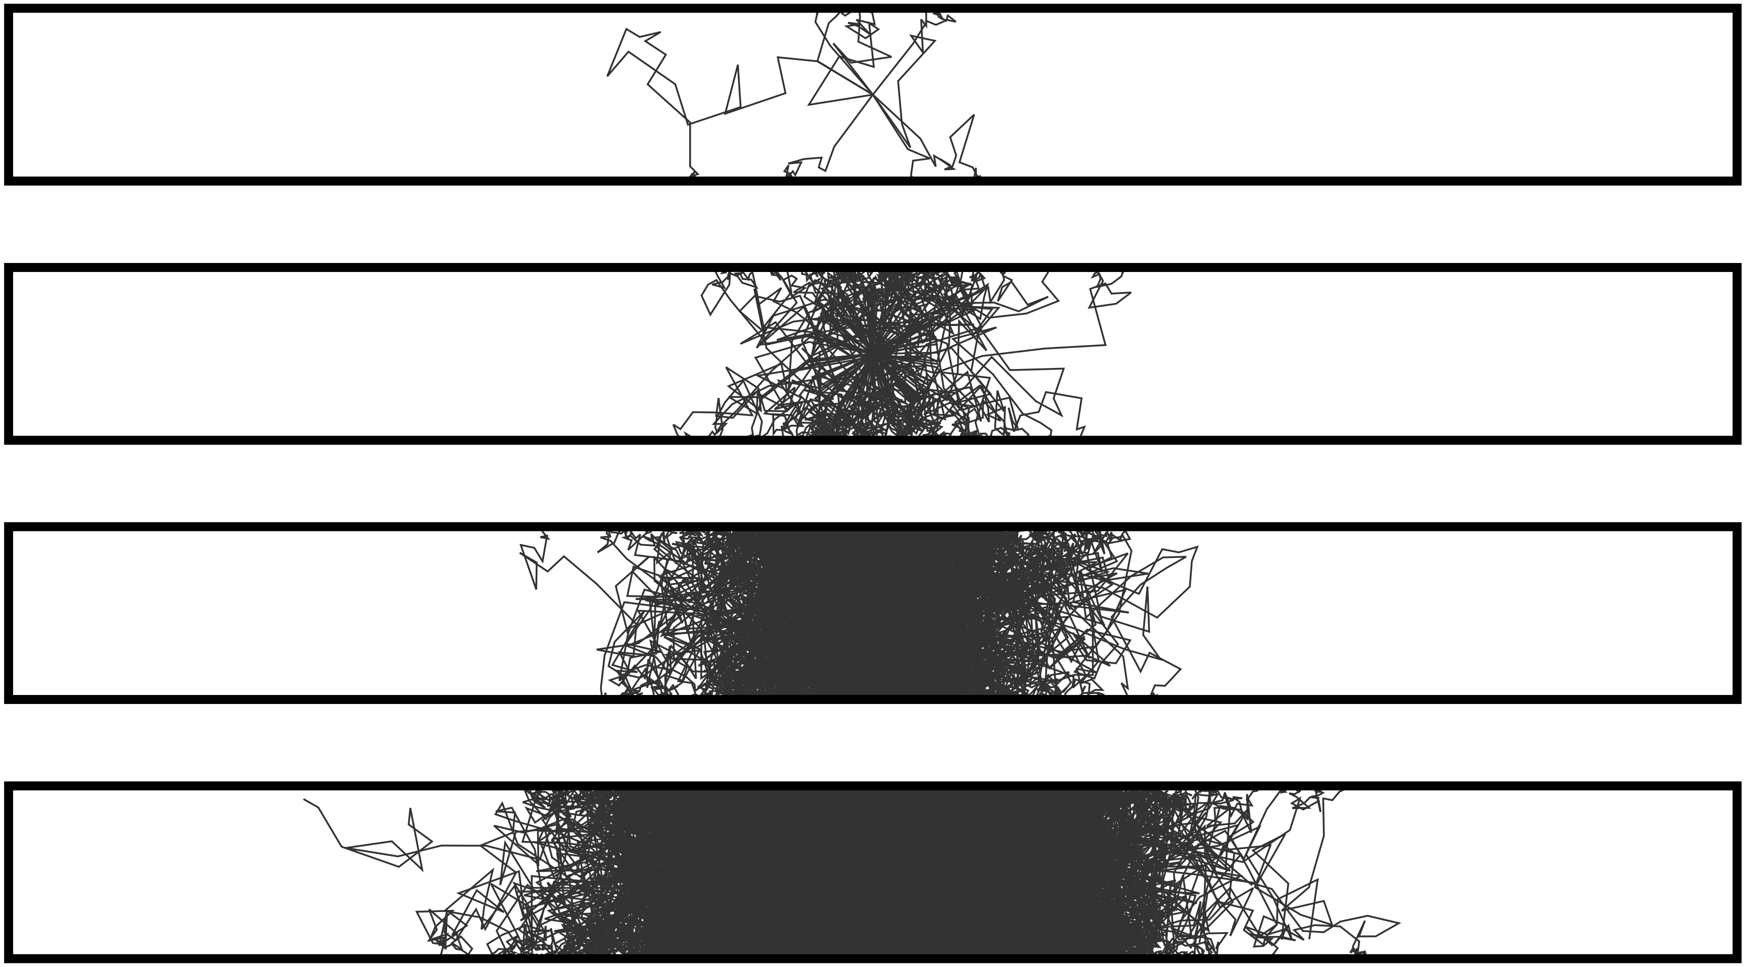

10, 100, 1000 and 10000 walks inside a 1×10 polygon.

Illustrates exponential difficulty of penetrating corridors.

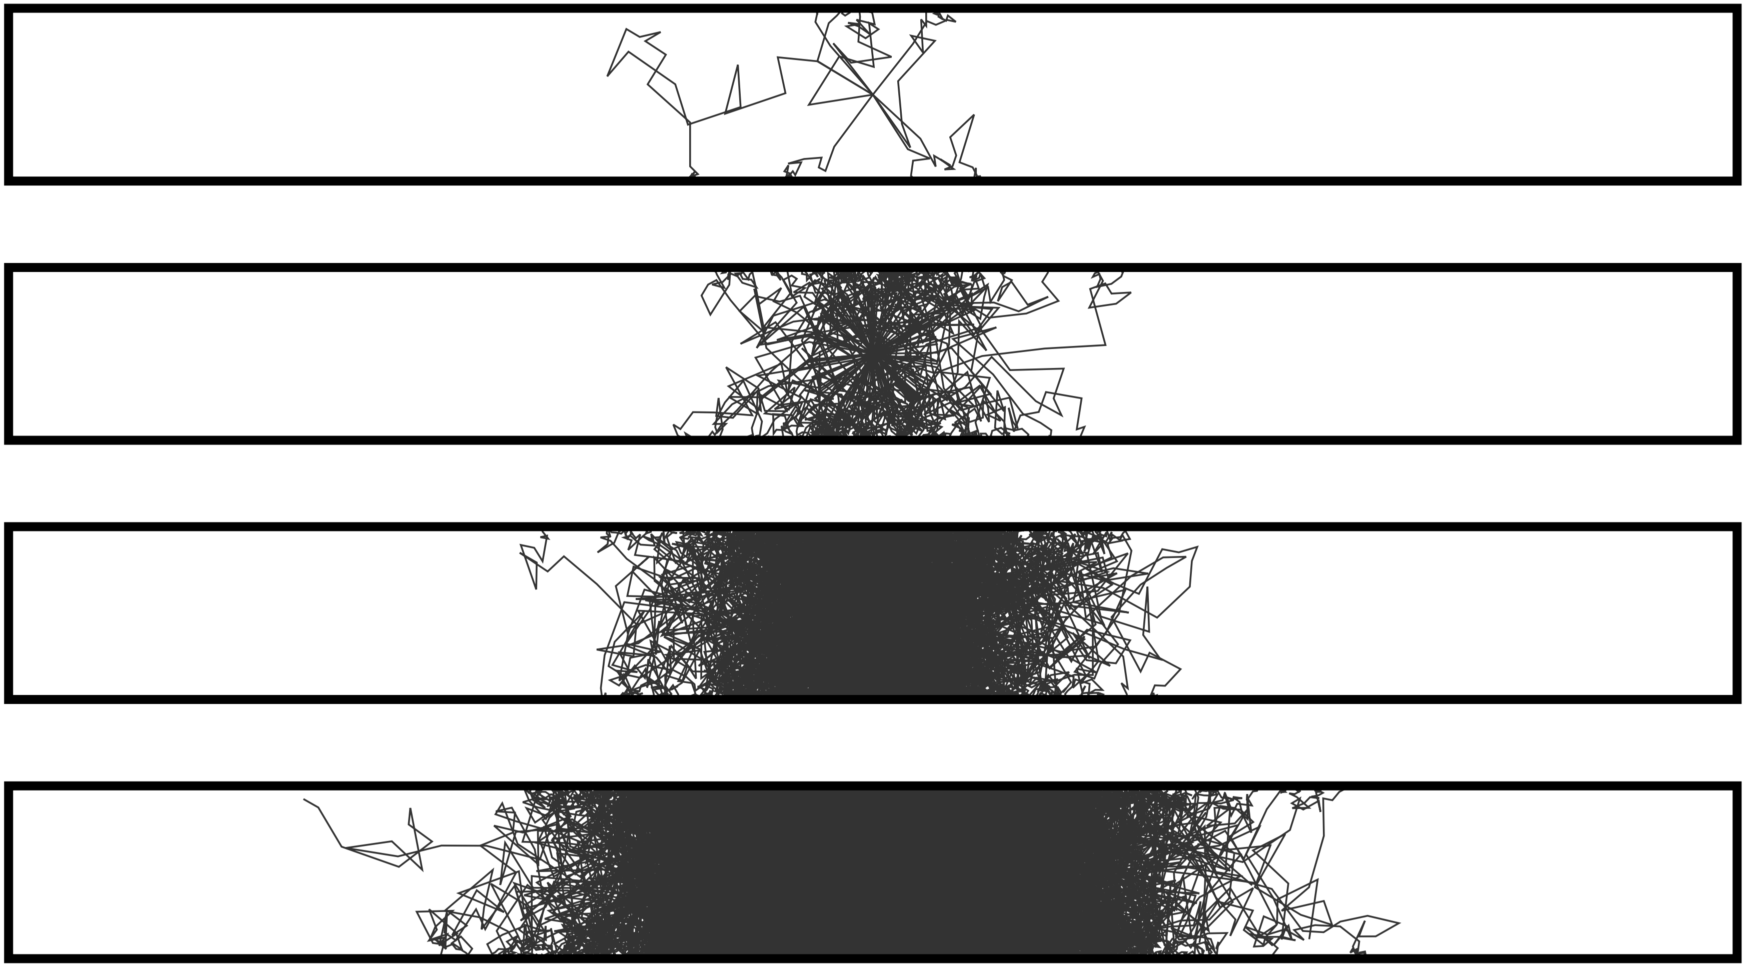

Harmonic measure of short sides of $1 \times R$ rectangle is $\approx \exp(-\pi R/2)$.

Consider harmonic measure as a function of starting point z.

Path started at z is equally likely to hit anywhere on circle around z.

Then restart until it hits domain boundary.

Harmonic measure at z is average of harmonic measures on small circle.

Harmonic measure satisfies the mean value property.

$\Rightarrow$ Harmonic measure is a harmonic function of z.

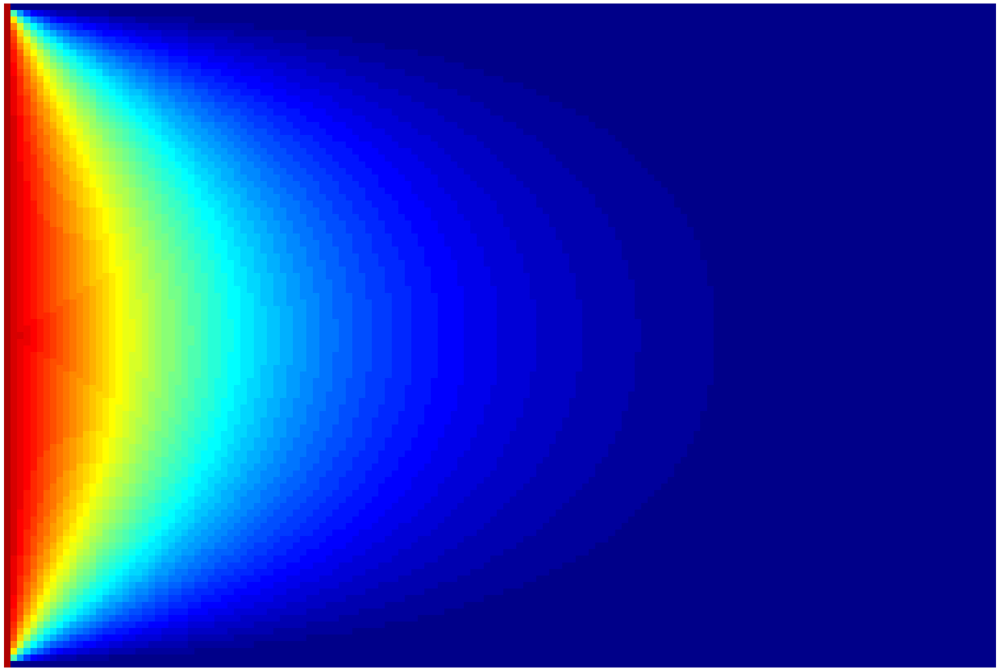

Harmonic measure of $E \subset \partial\Omega$ is the harmonic function of base point with boundary values 1 on E and 0 elsewhere.

Harmonic functions are conformally invariant.

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

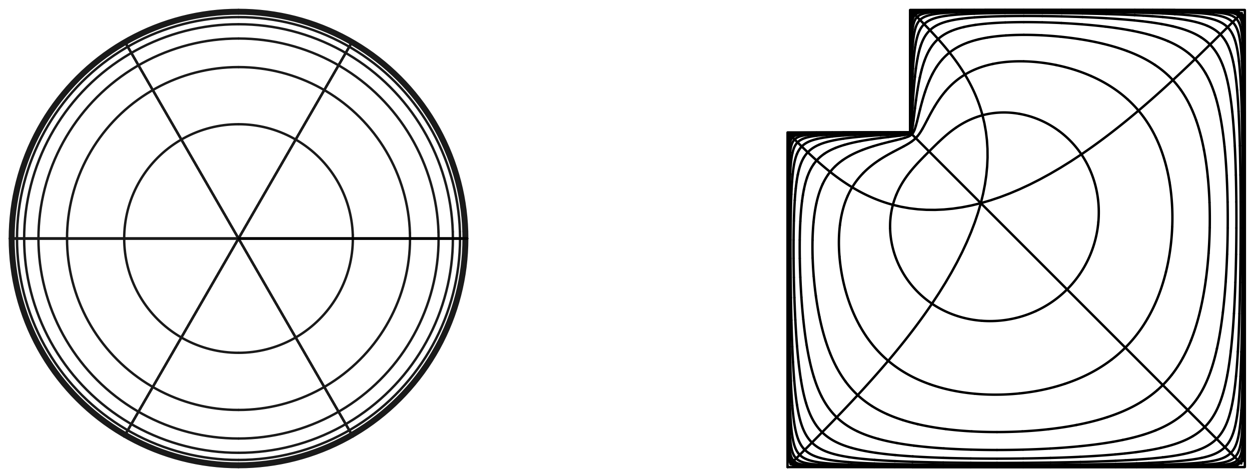

Conformal = angle preserving

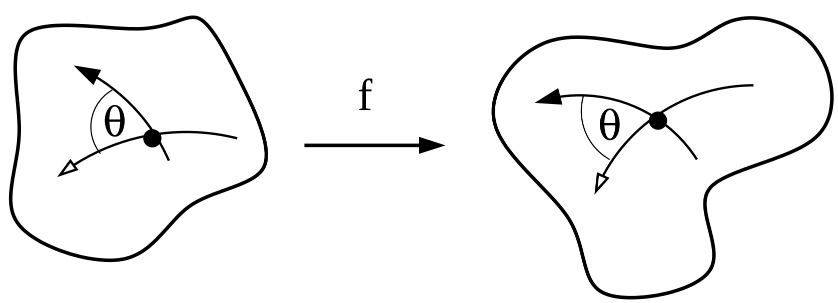

Georg Friedrich Bernhard Riemann
Stated RMT in 1851

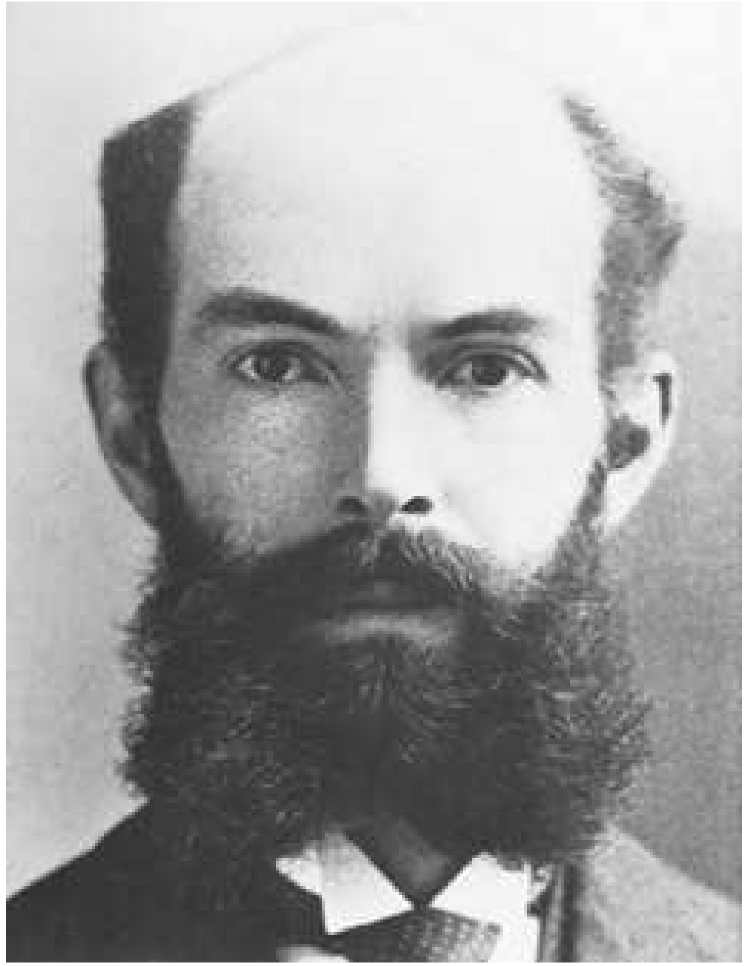

William Fogg Osgood
First proof of RMT
Transactions of AMS, vol. 1, 1900

The proof of Osgood represented, in my opinion, the "coming of age" of mathematics in America. Until then, . . . the mathematical productivity in this country in quality lagged behind that of Europe, and no American before 1900 had reached the heights that Osgood then reached.

J.L. Walsh, "History of the Riemann mapping theorem", Amer. Math. Monthly, 1973.

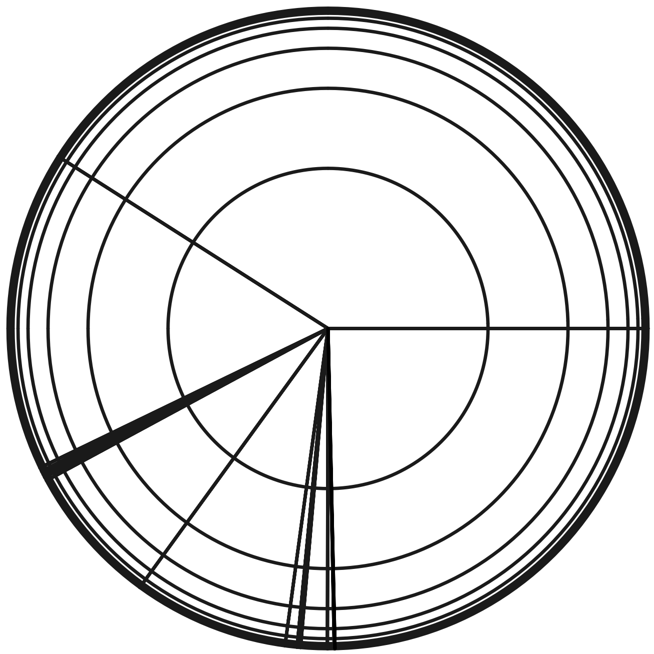

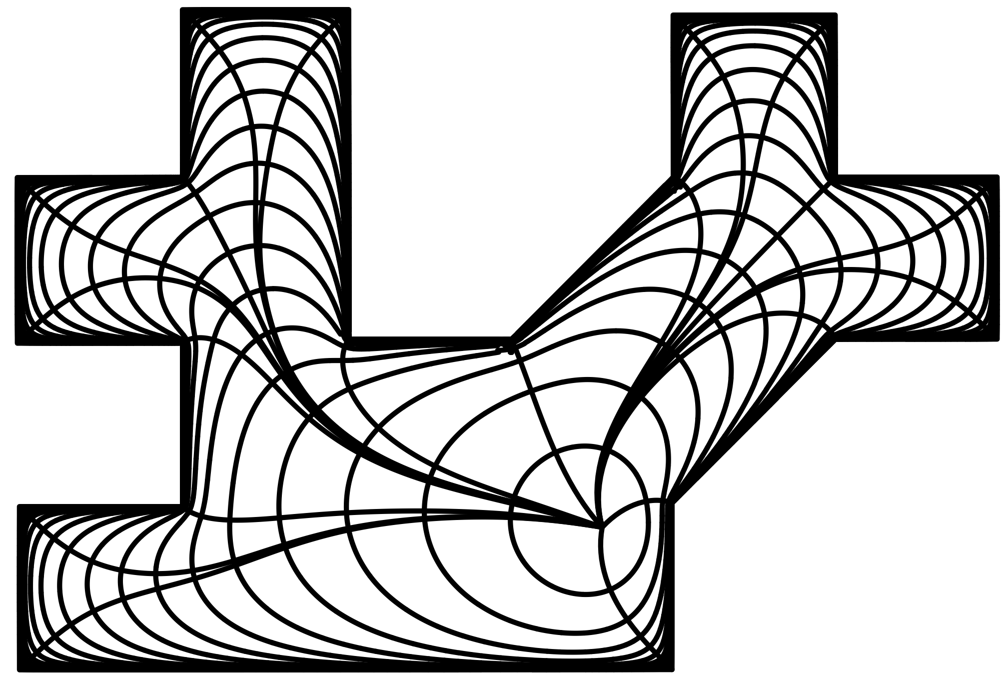

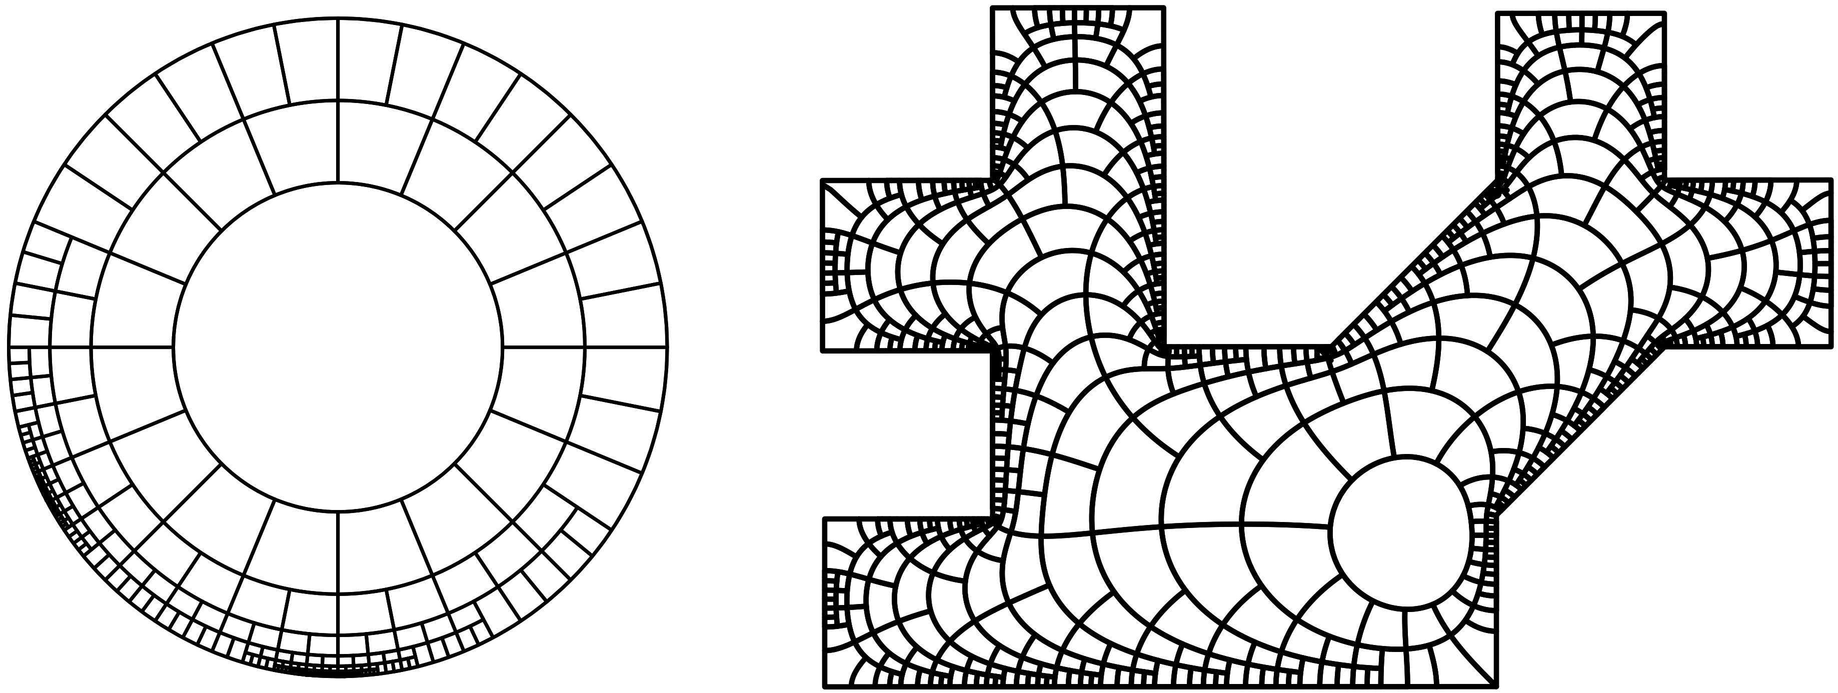

Numerical conformal mapping:

- Koebe
- Theodorsen
- Fornberg
- Wegman
- Gaier
- Symm
- Kerzman-Stein
- Integral equations via fast multipole, Rokhlin
- Circle packing, Sullivan, Rodin, Stephenson
- `CRDT`, Driscoll and Vavasis
- `SCToolbox`, Trefethen, Driscoll
- `ZIPPER`, Marshall

Standard textbook proof of RMT by normal families gives an algorithm:

- Choose similarity map of Ω into D.
- Choose image boundary point closest to origin.
- Move this point to origin by Möbius transformation σ.
- Take square root (well defined on simply connected region).
- Move $\sqrt{\sigma(0)}$ back to 0 by Möbius transformation.
- Repeat.

$d_n = \mathrm{dist}(\partial\Omega_n, 0)$

Lemma: If $d_0 \geq 1/2$ then $d_n = 1 - O(1/n)$.

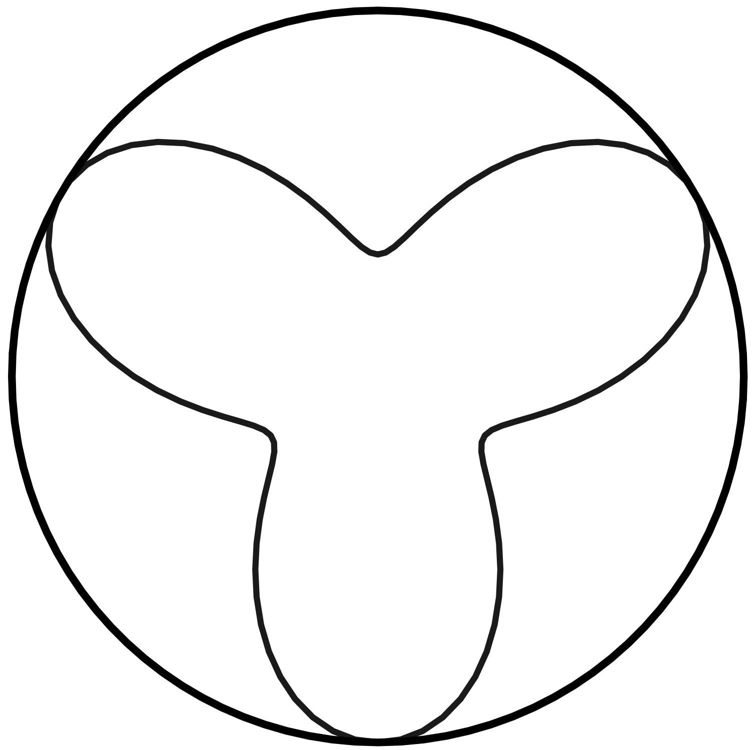

0th iteration

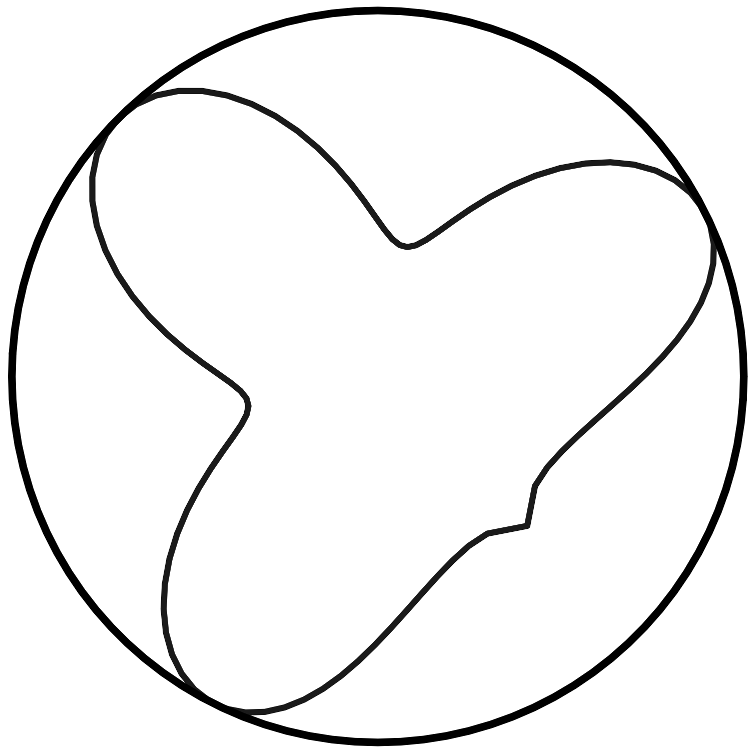

1th iteration

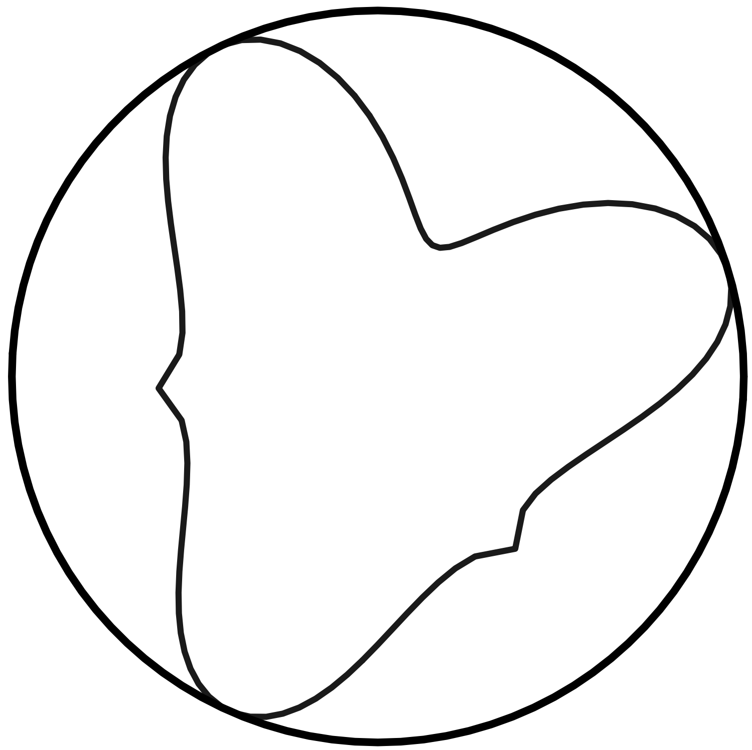

2nd iteration

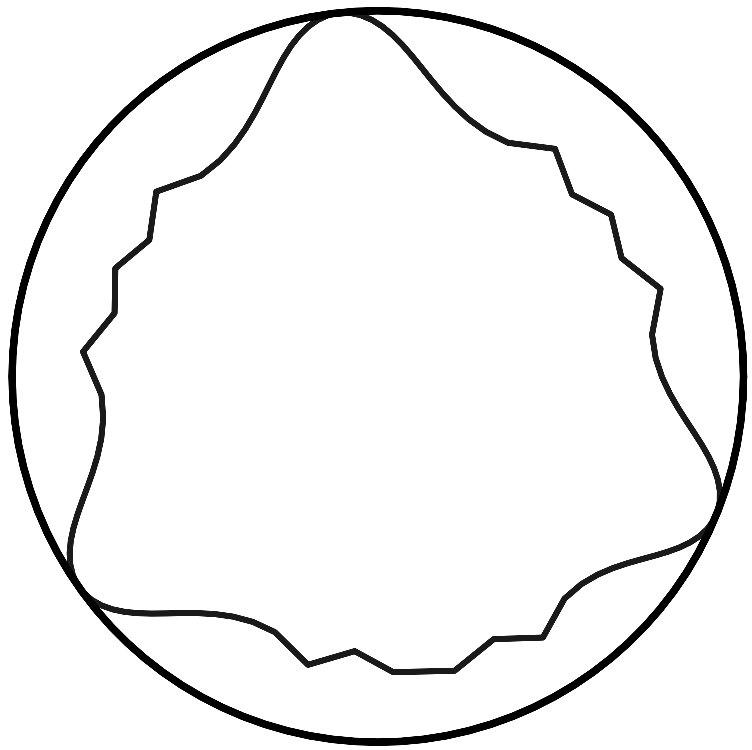

10th iteration

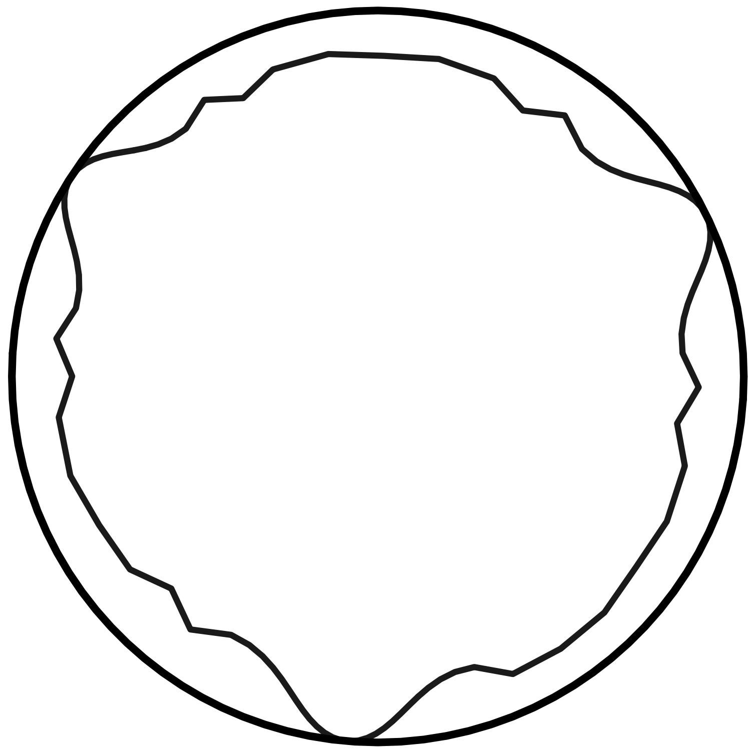

20th iteration

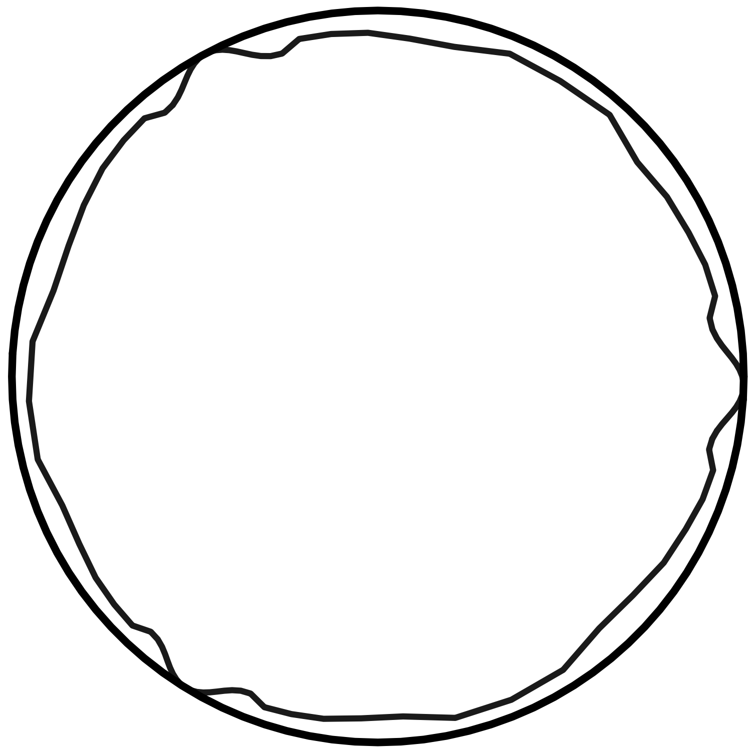

50th iteration

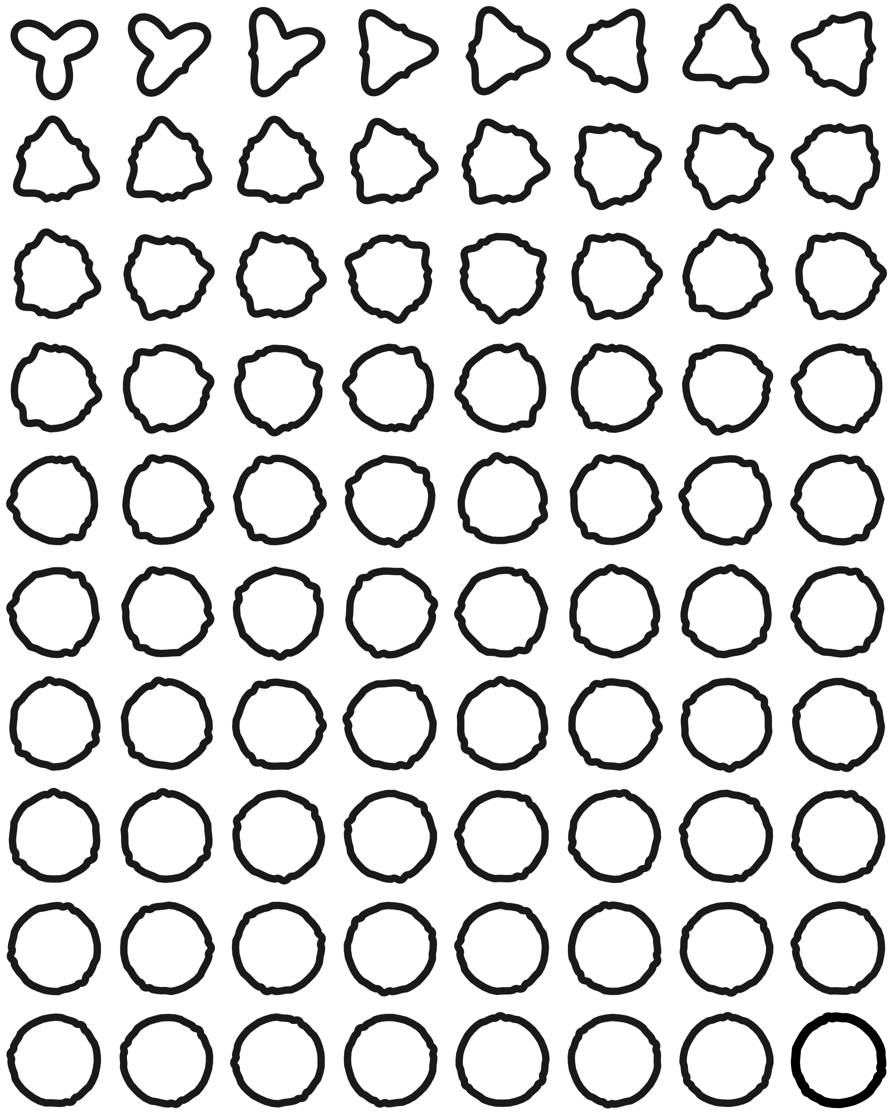

First 80 iterations. Converges, but is slow.

We want to measure computational complexity of conformal mapping.

Want finite input and finite output.

We will consider conformal maps between polygons and unit disk.

Input = vertices = n-tuple of complex numbers (plus base point)

Output = harmonic measures of edges = n-tuple of reals summing to 1
= conformal prevertices = n-tuple on circle

Schwarz-Christoffel formula (1867):

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

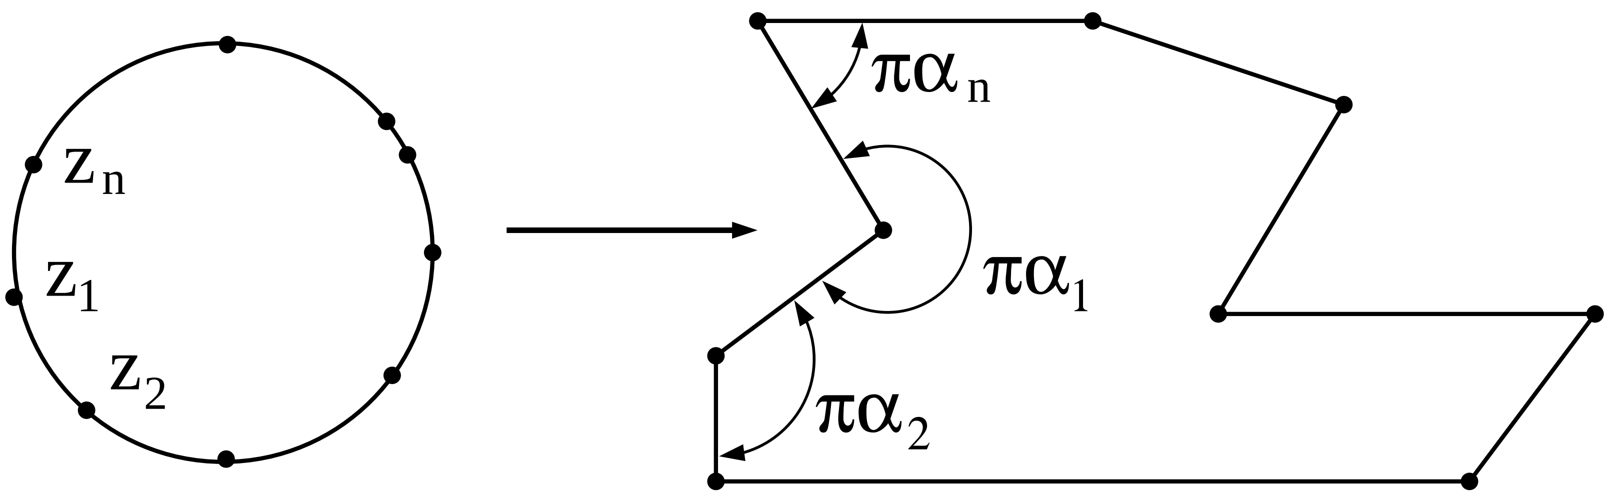

α's known

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

Schwarz-Christoffel formula (1867):

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

Christoffel

Schwarz

Problem: given n-gon, compute the SC-parameters.

How fast can we do this?

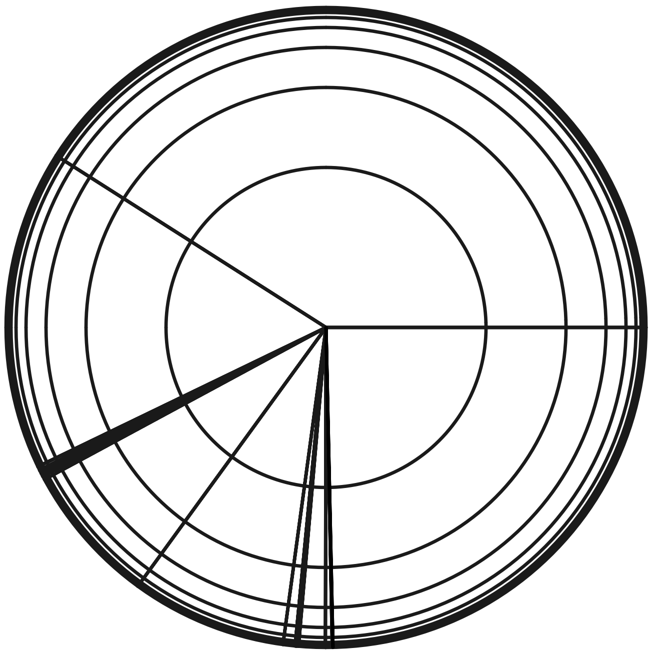

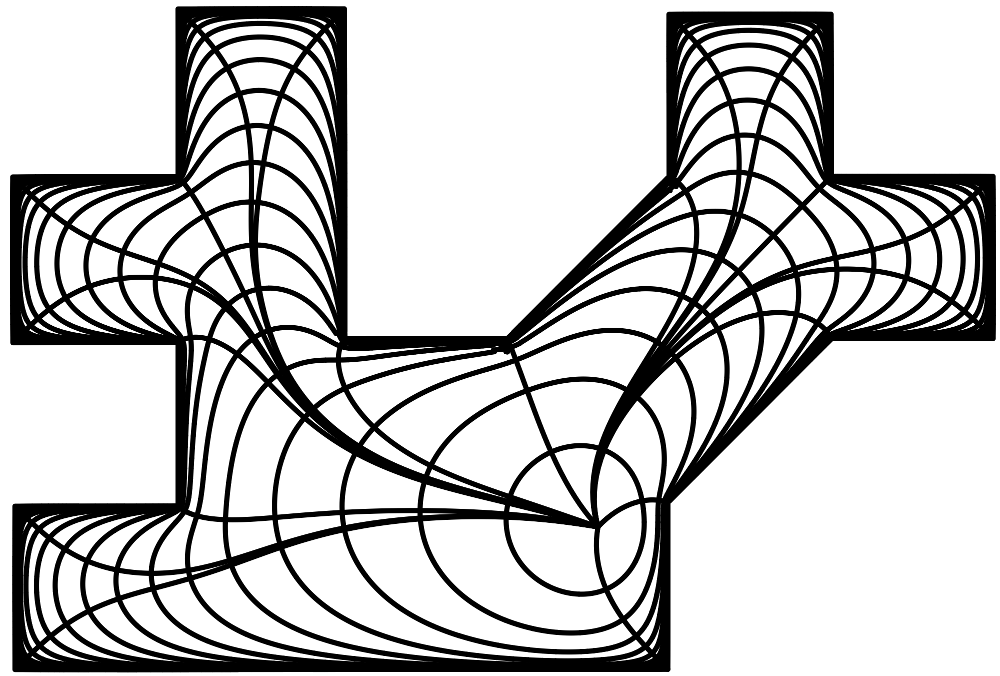

Davis' method:

- Guess initial parameter values (all equal?).
- Plug guesses in SC-formula, get a polygon (correct angles, wrong lengths).
- If an edge is too long, shorten corresponding parameter arc.
- If too long, lengthen the gap.

$$\text{new gap} = \frac{\text{old gap} \times (\text{target side}/\text{guessed side})}{\text{sum of new gaps}}$$

Does Davis' method always converge?

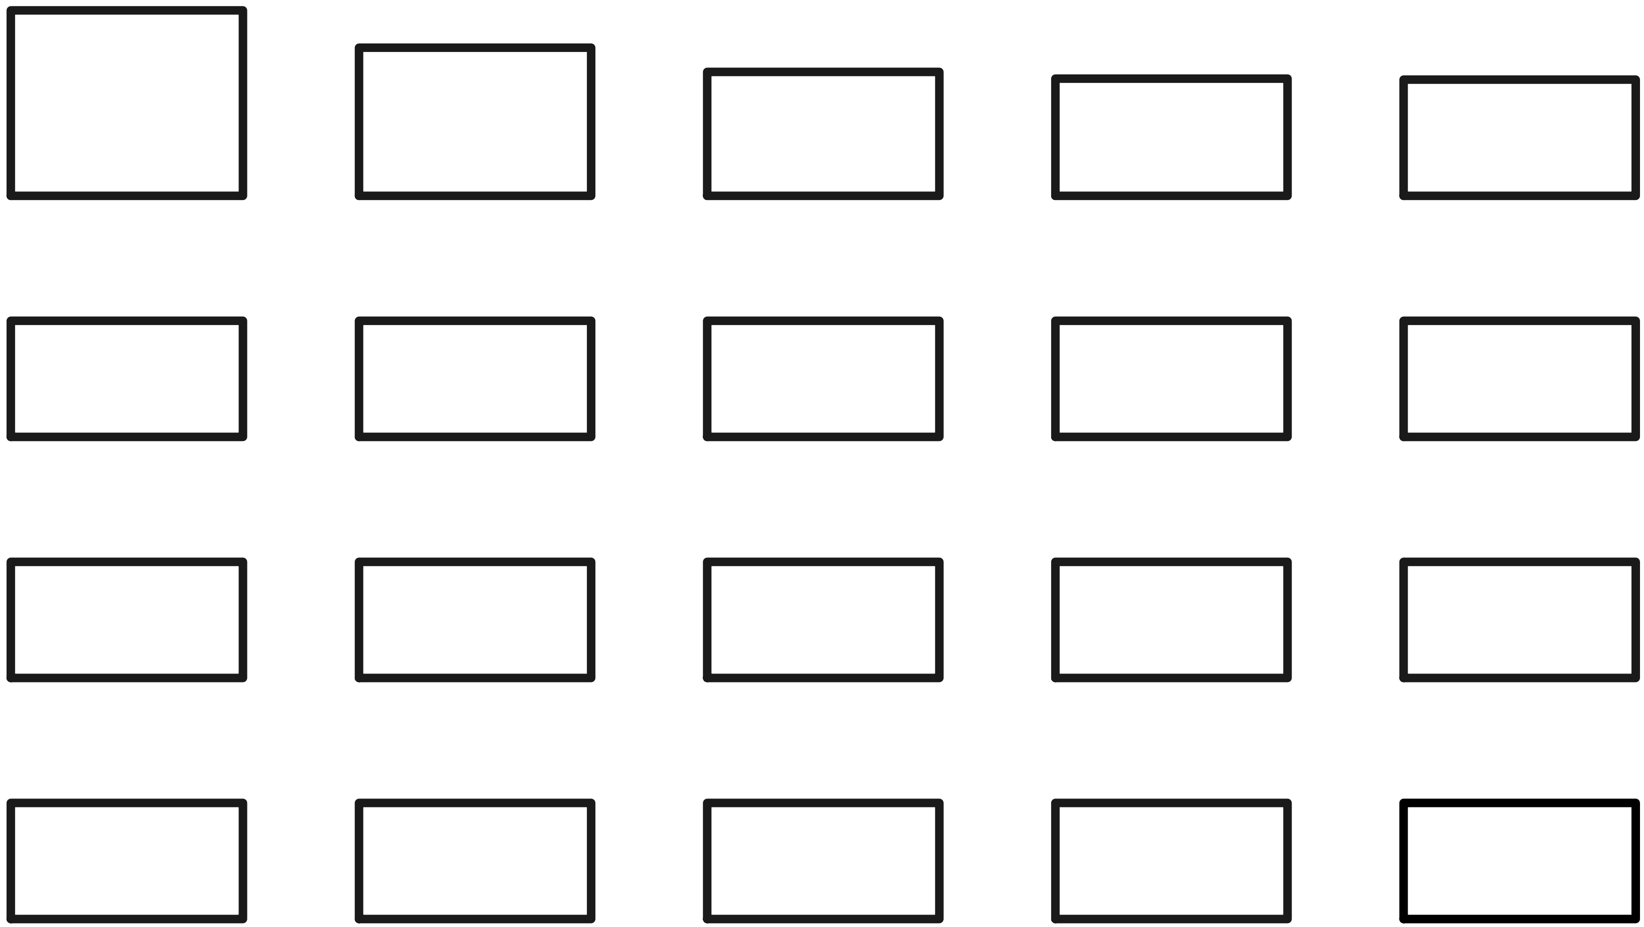

20 iterations of Davis' method - QC error.

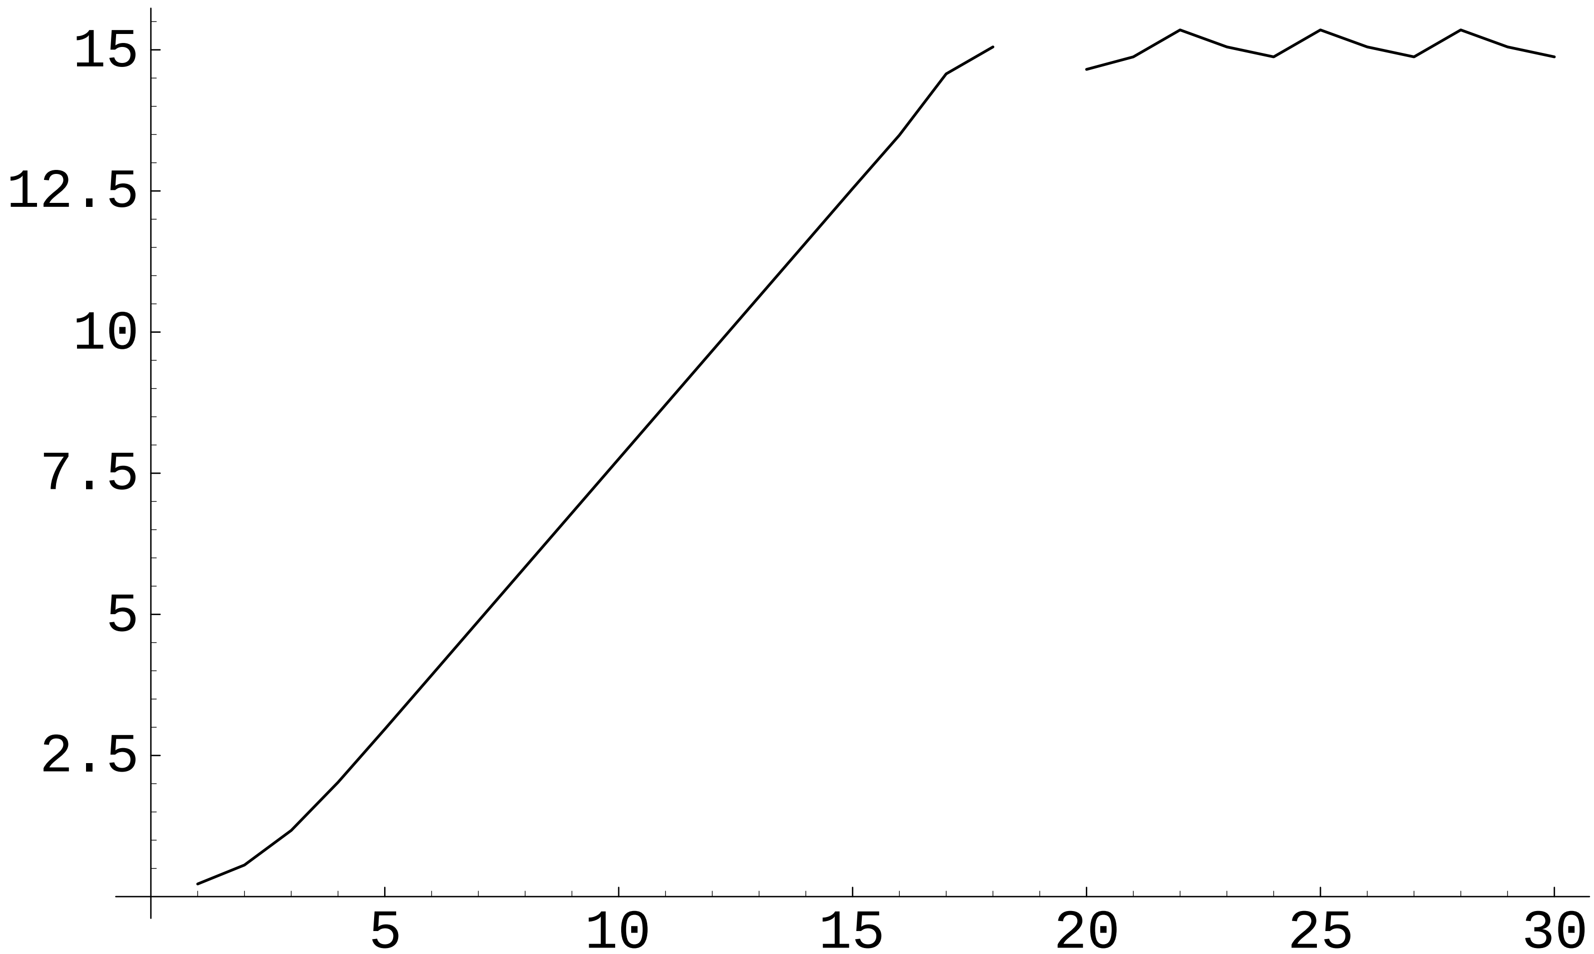

30 iterations of Davis' method - QC error.
We hit machine error $\approx 10^{-15}$.

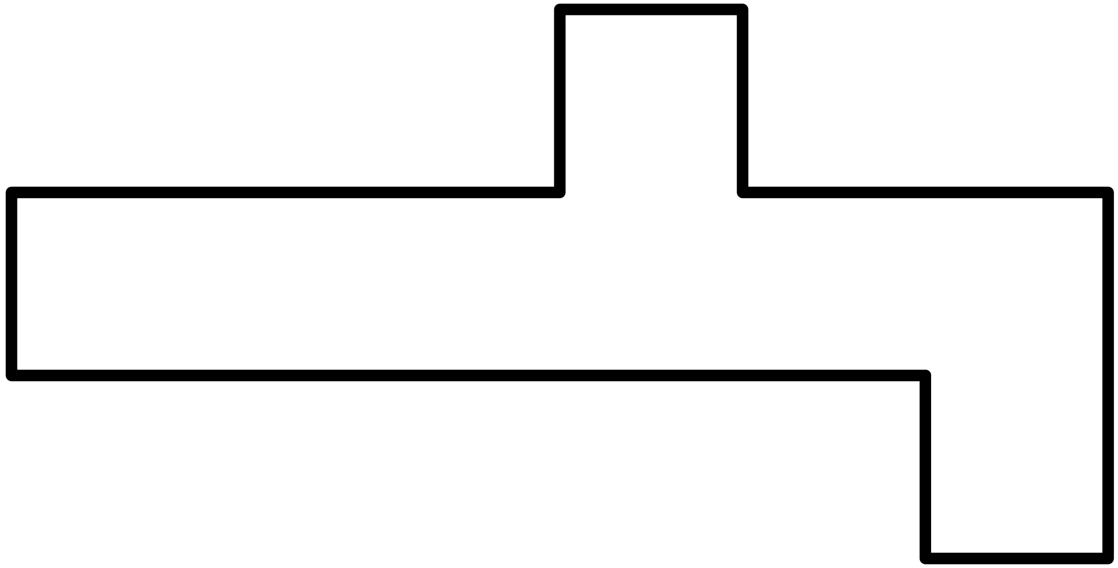

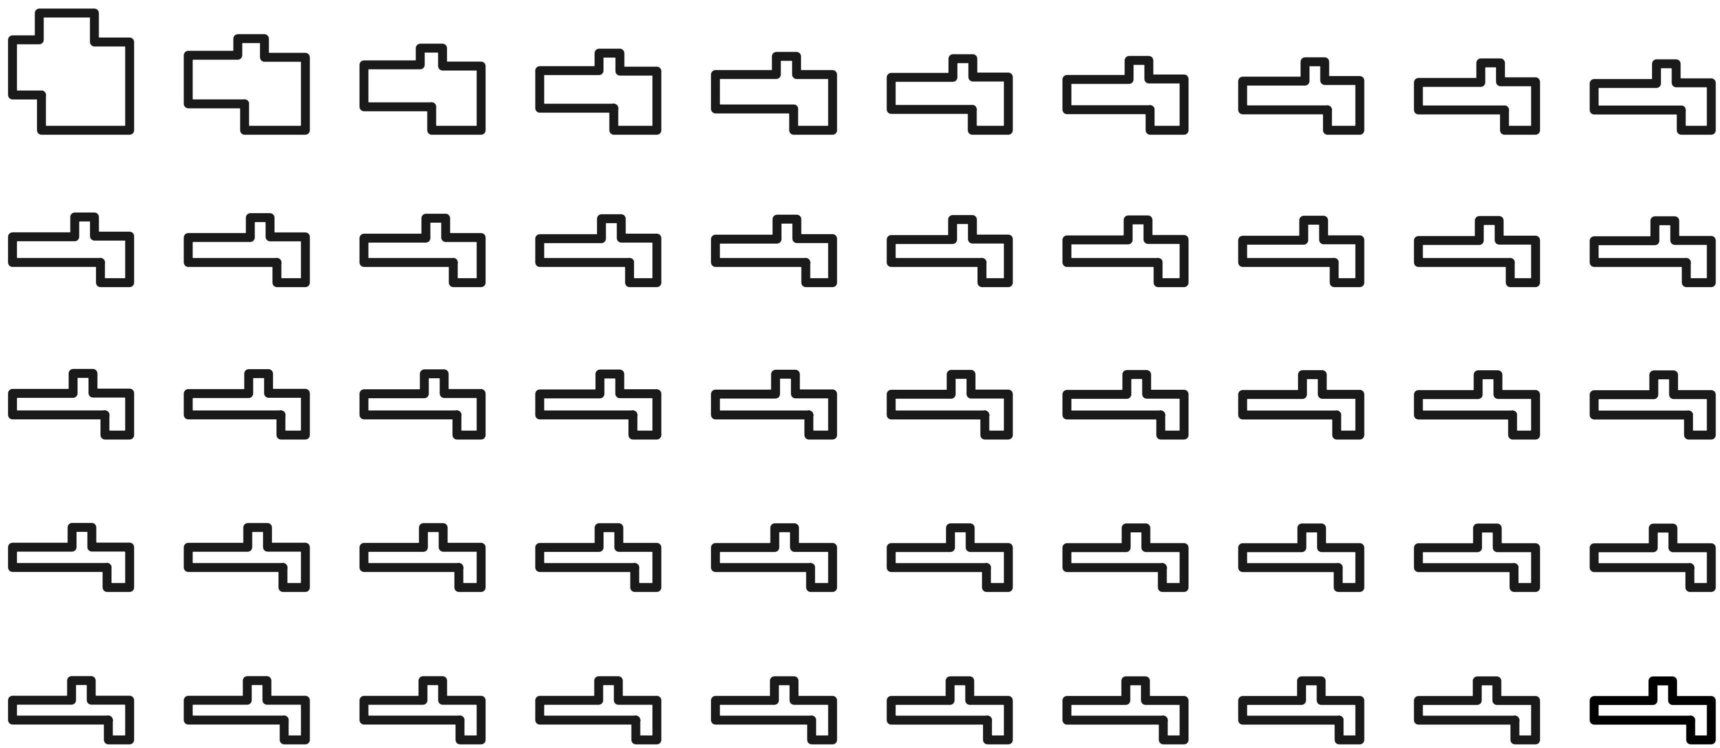

50 iterations of Davis' method.

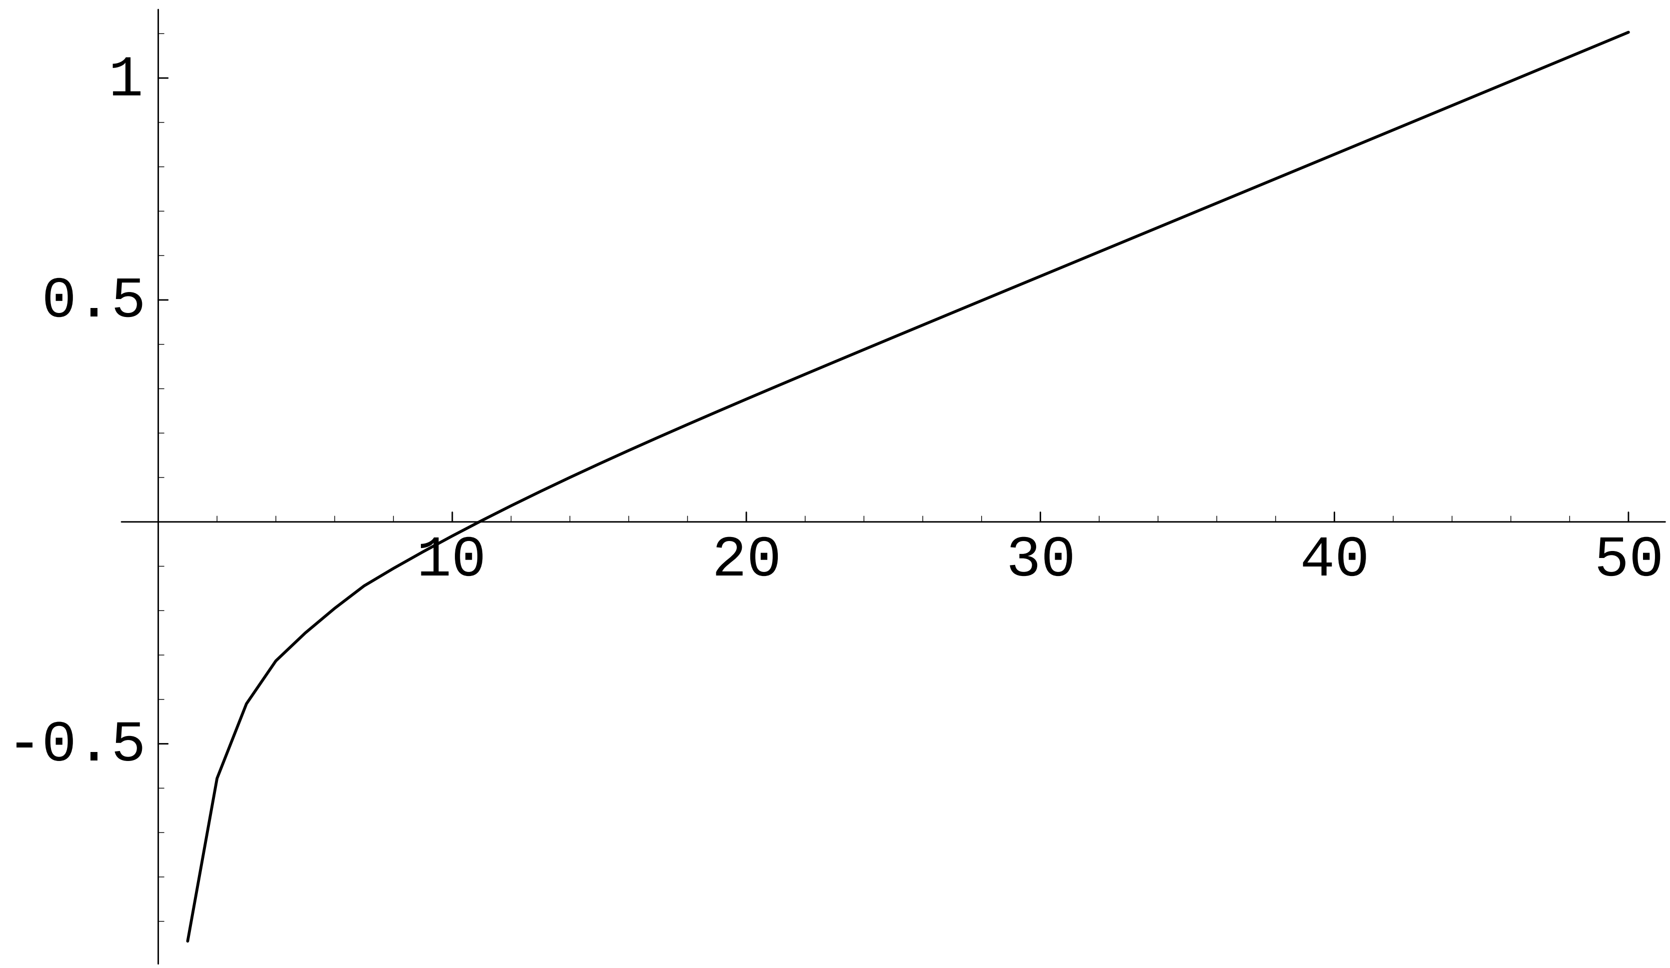

QC error for 50 iterations of Davis' method.

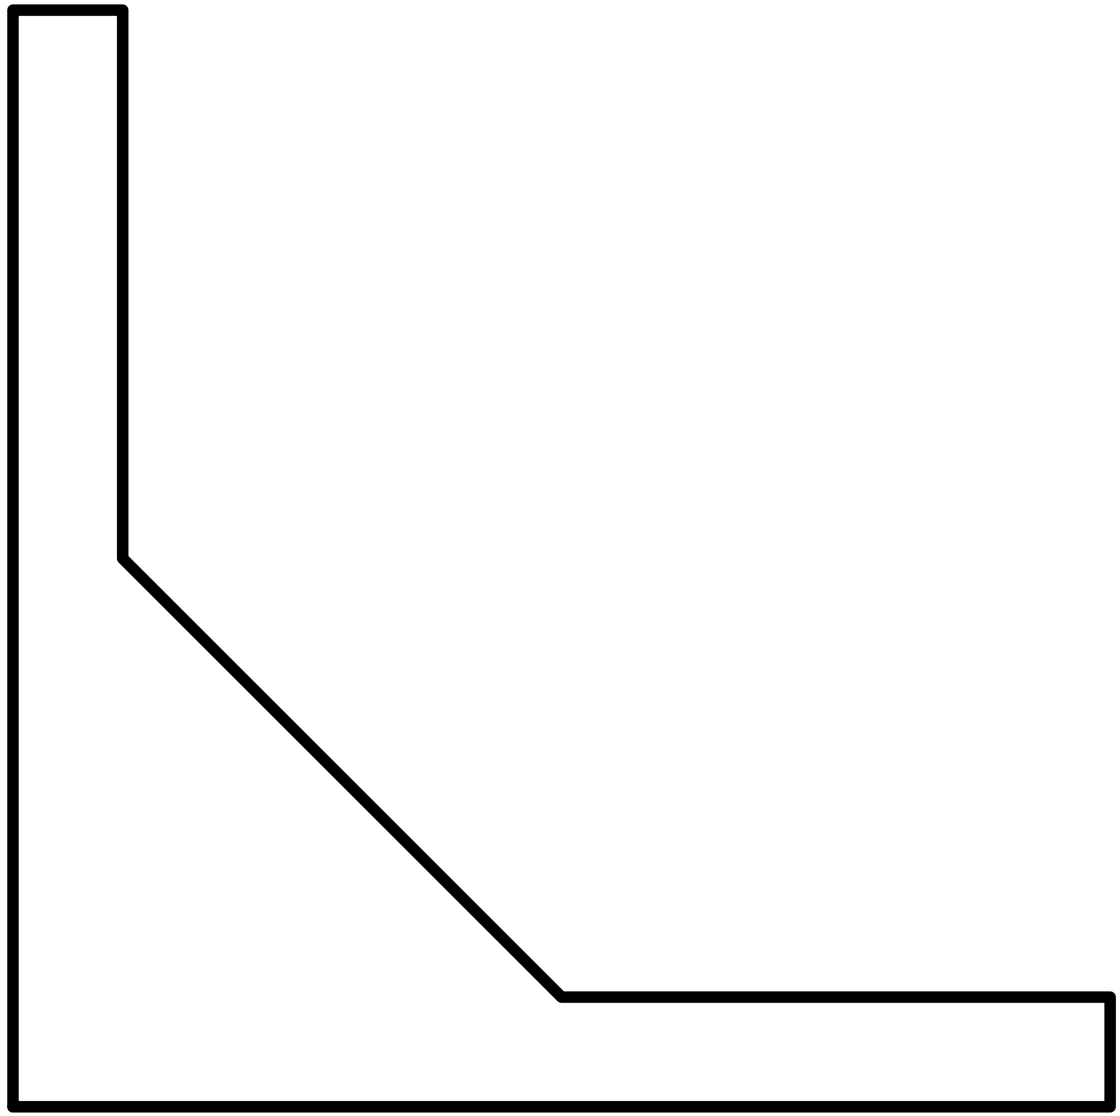

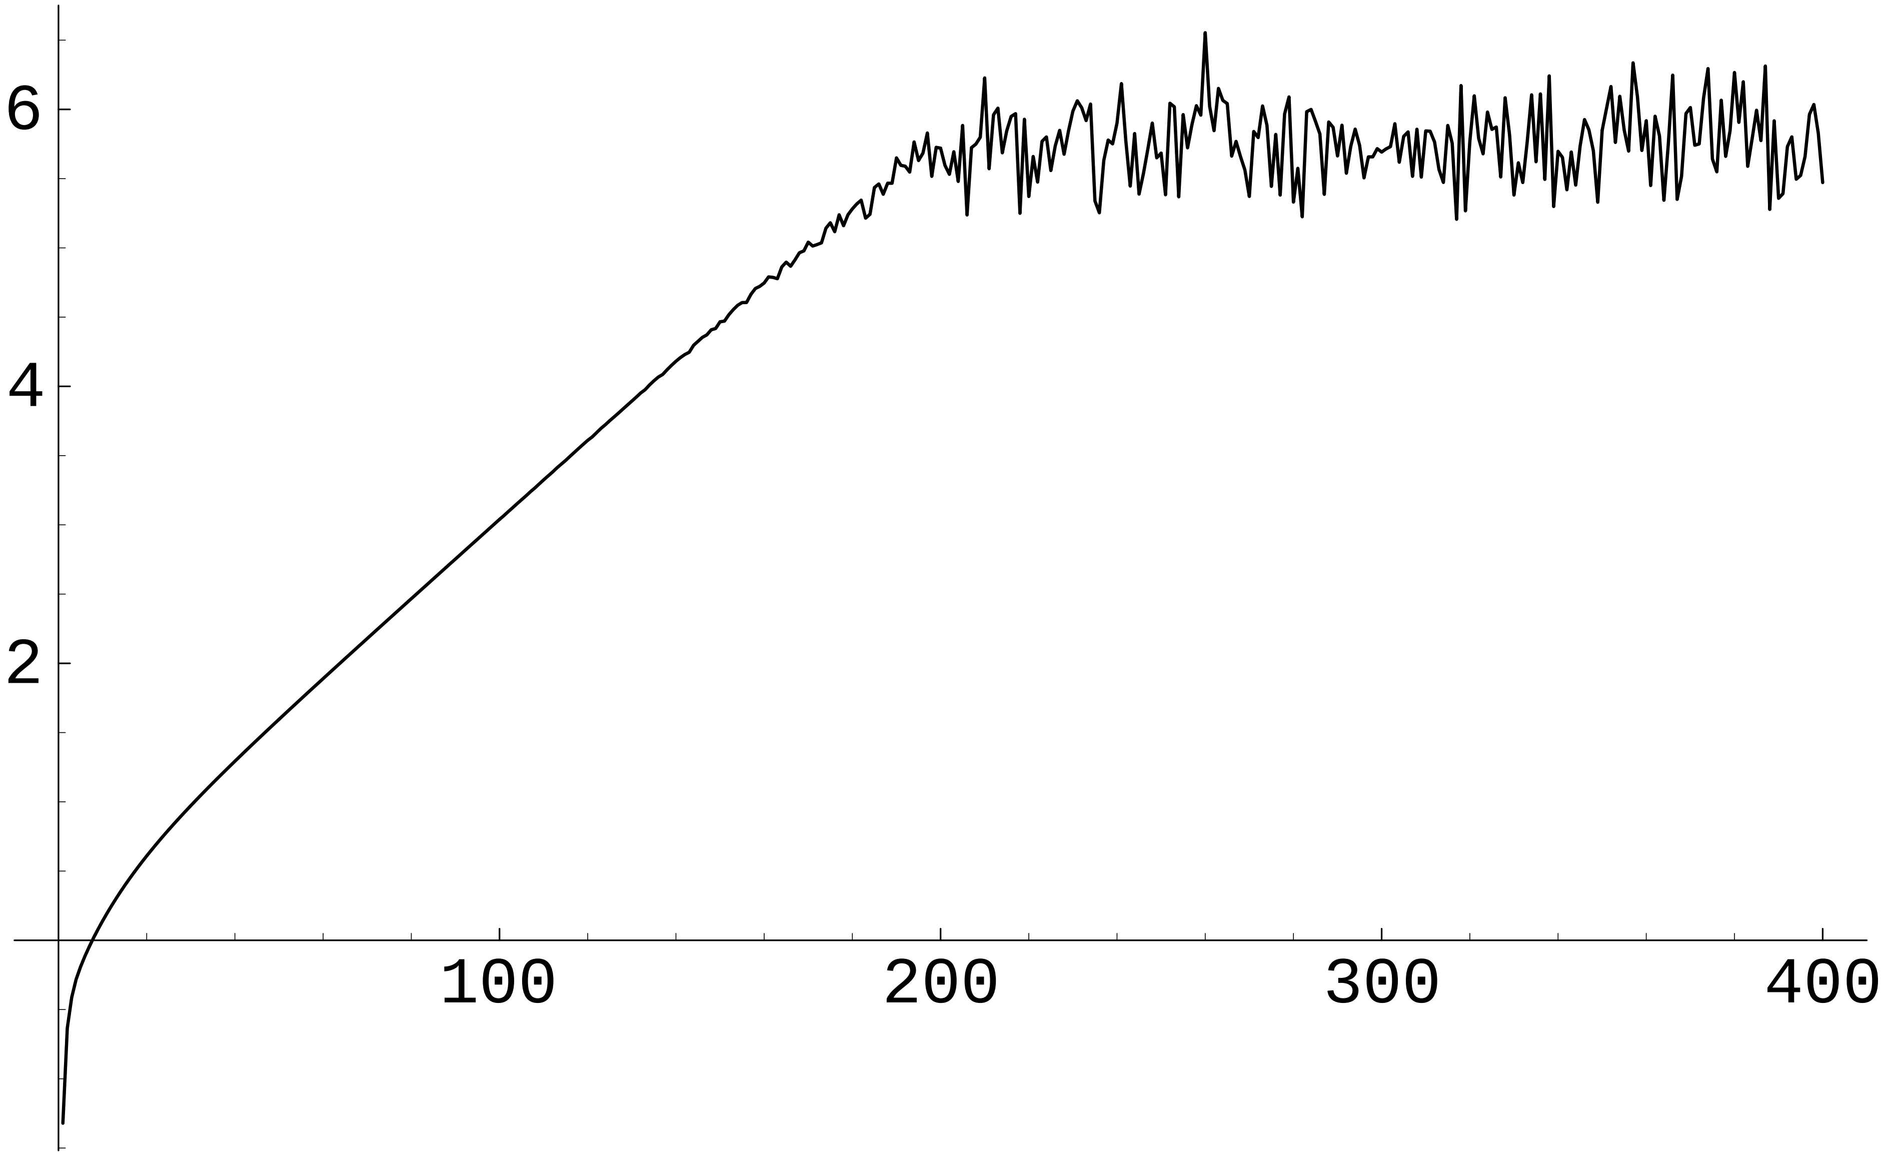

QC error for 400 iterations of Davis' method.

Cross Ratios and Delaunay Triangulations (Driscoll and Vavasis)

Identifies n-tuples on T with R^{n-3} using cross ratios.

Uses standard high-dimensional root finding. Different versions based in different approximations of the derivative matrix:

- Full CRDT: discrete approximation (samples all directions)
- Simple CRDT: assume derivative = Identity
- Shortcut CRDT: use Broyden updates

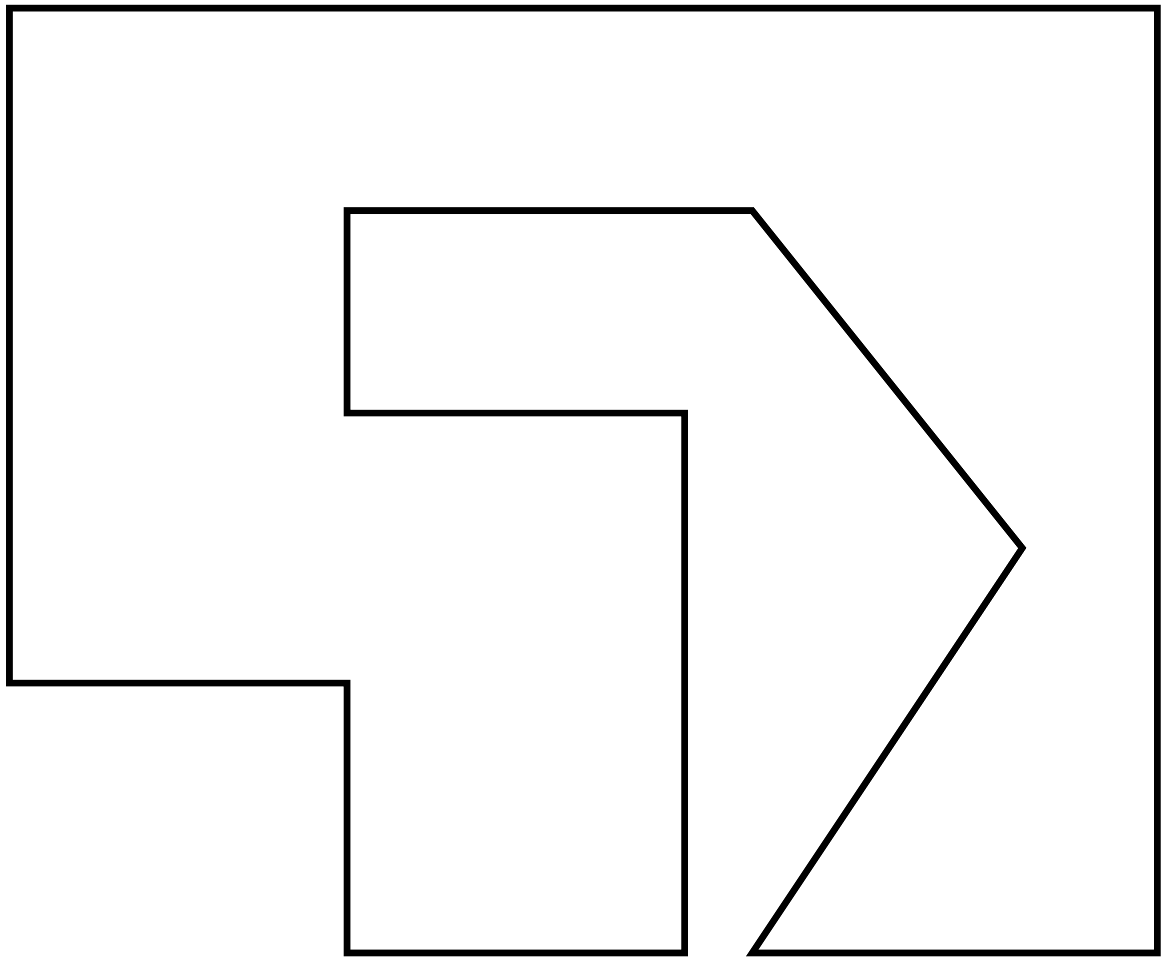

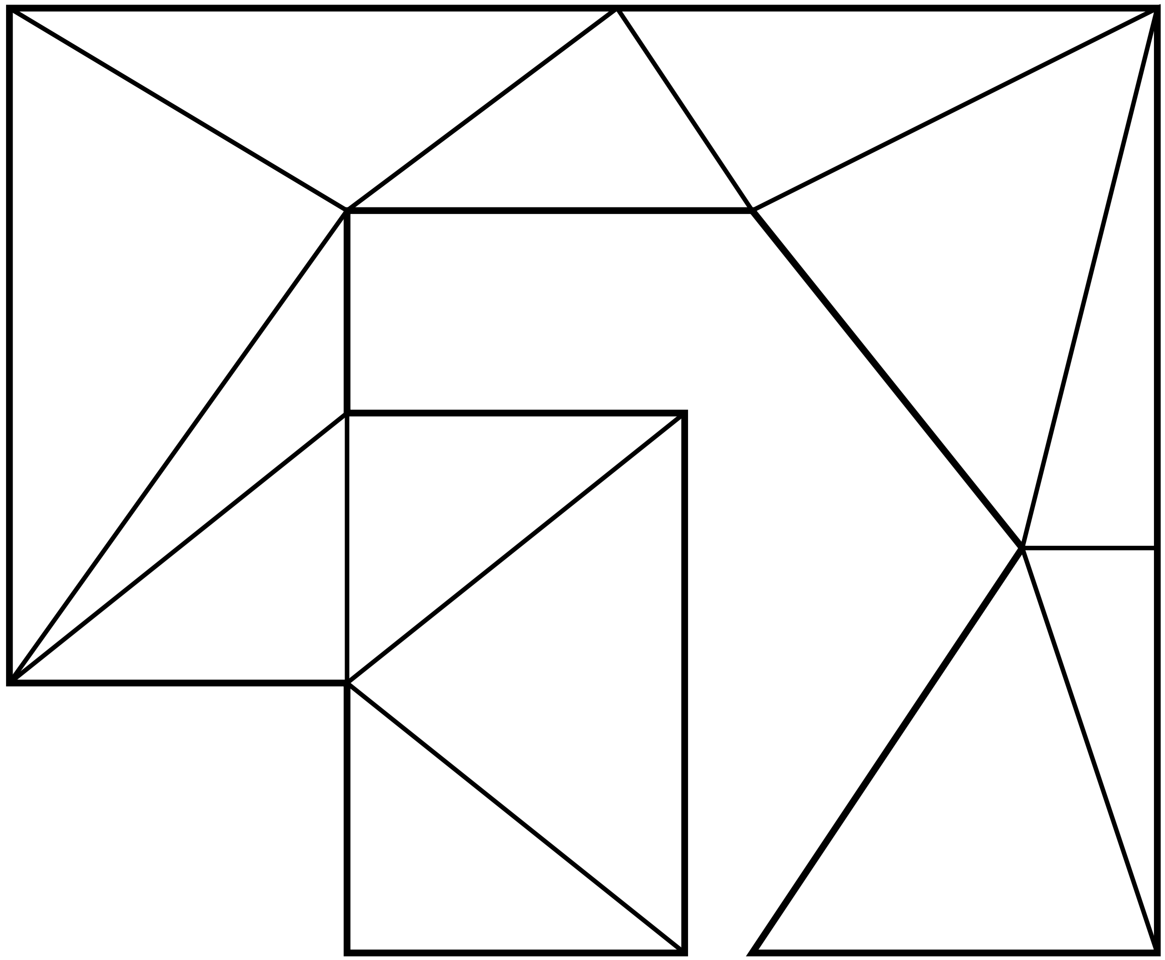

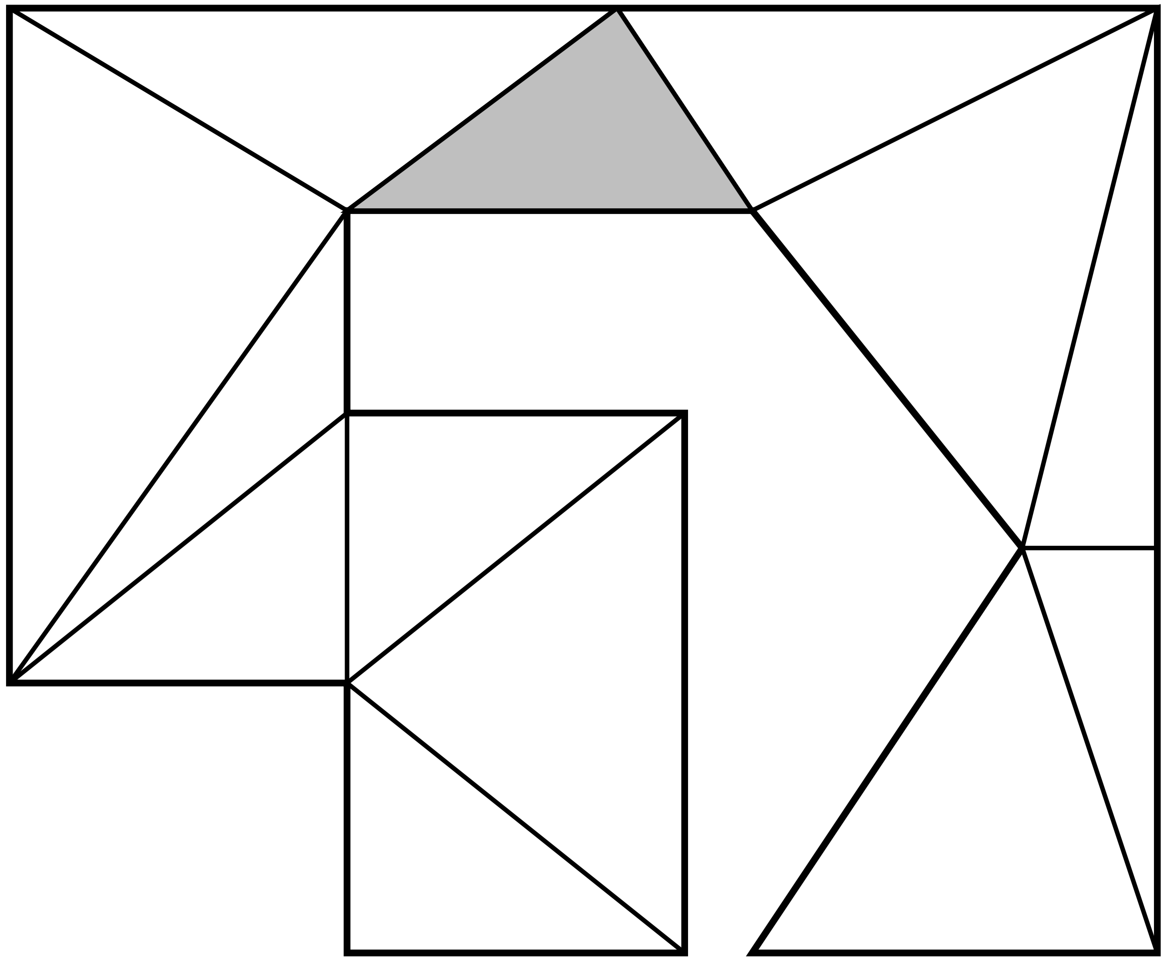

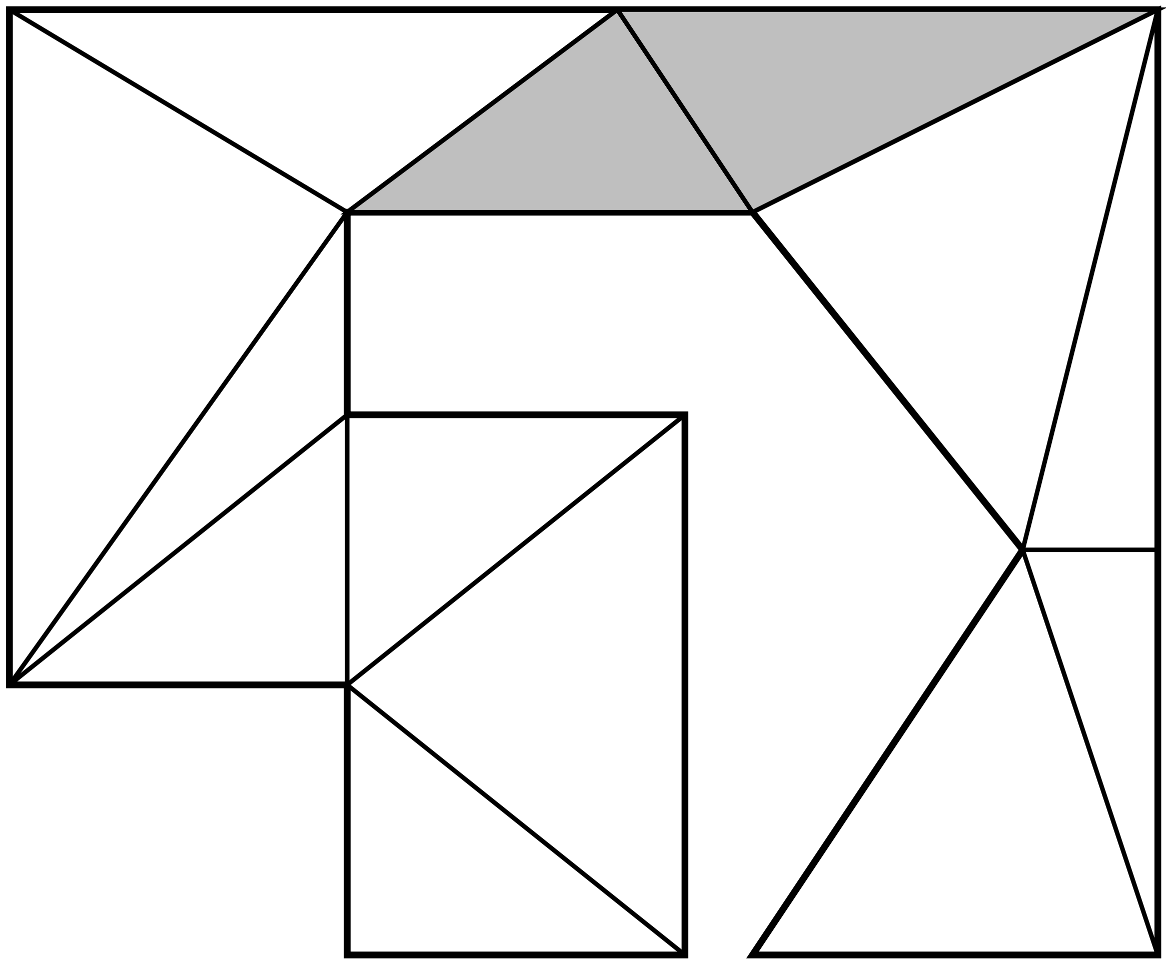

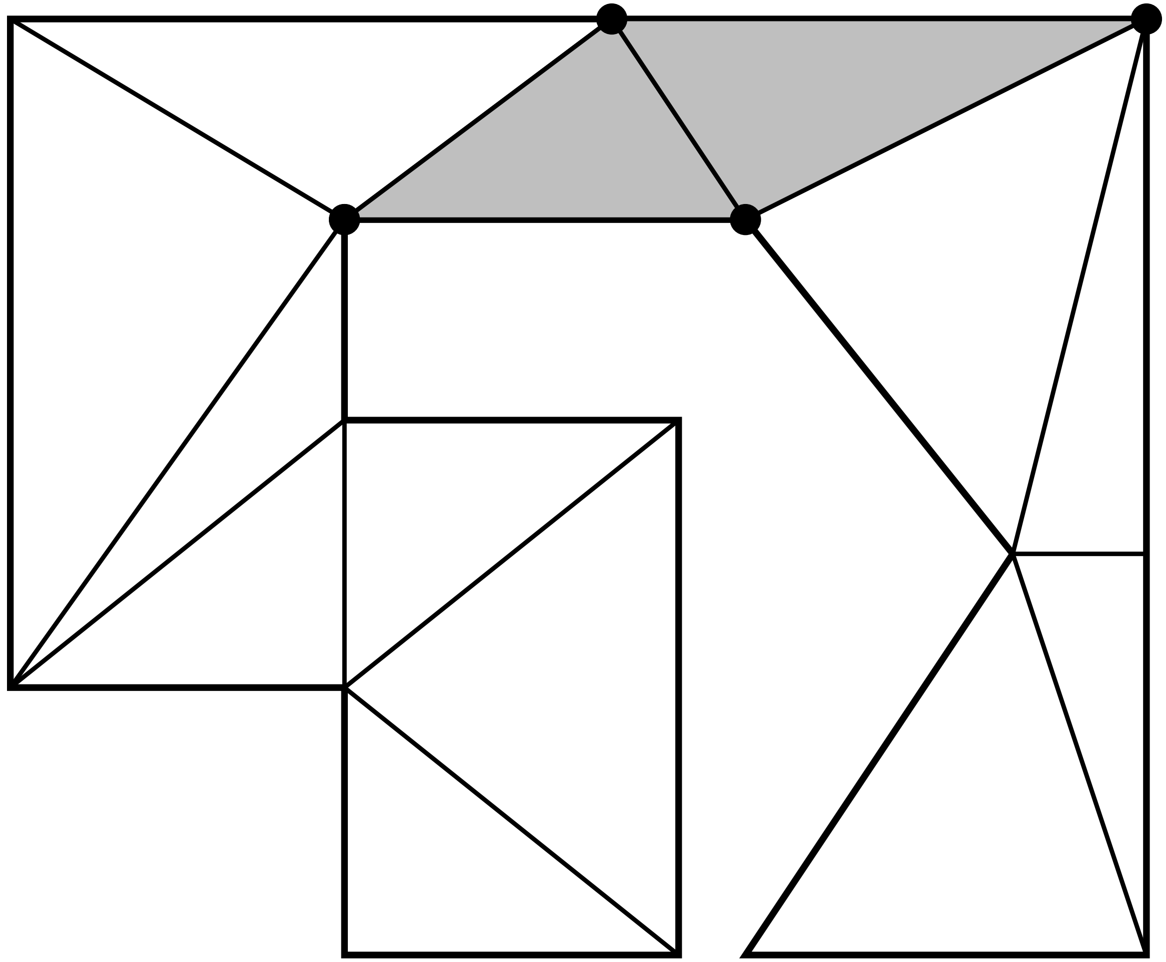

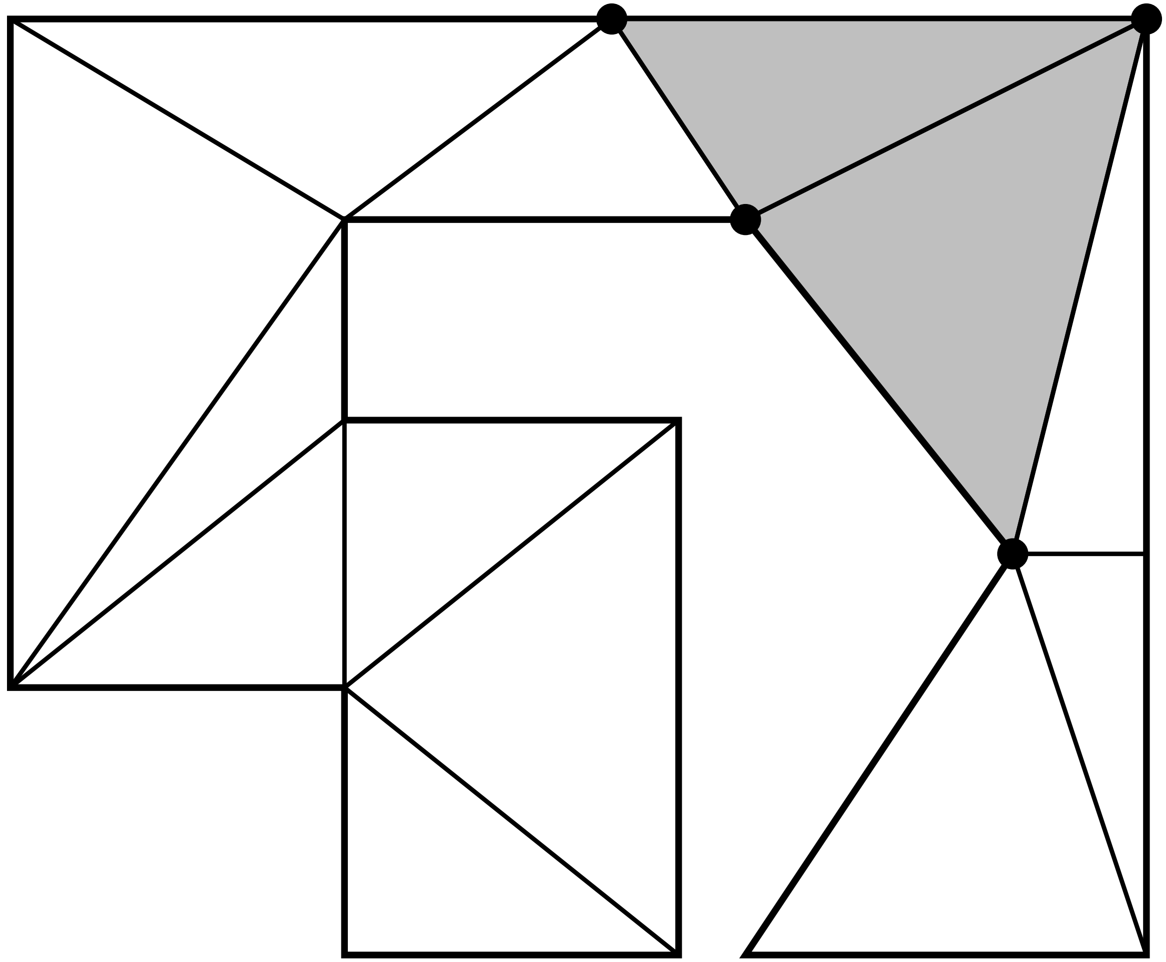

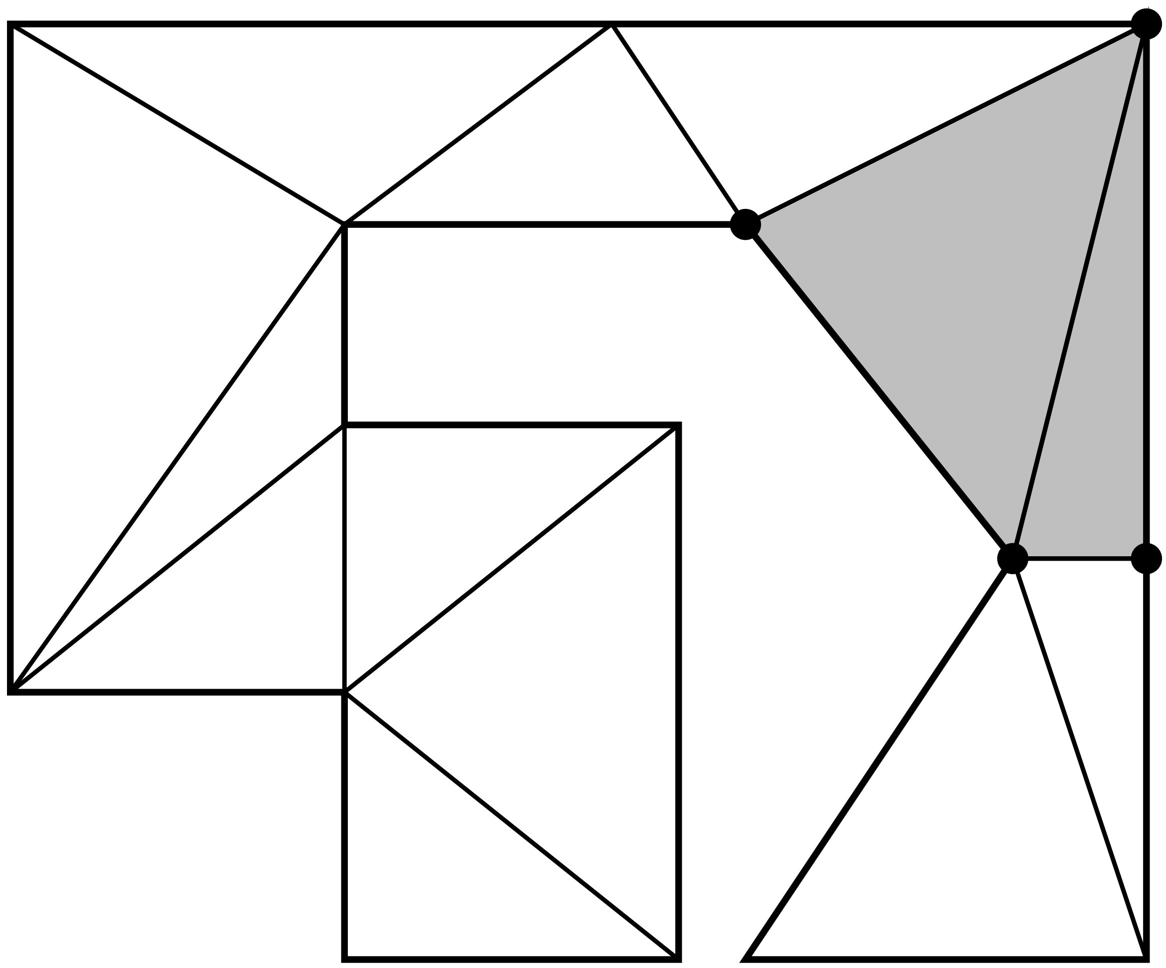

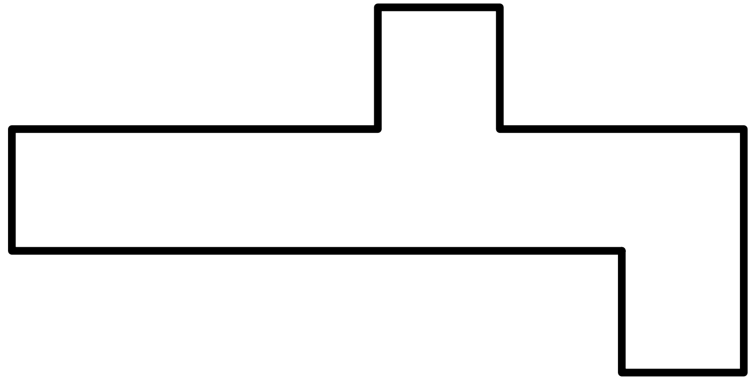

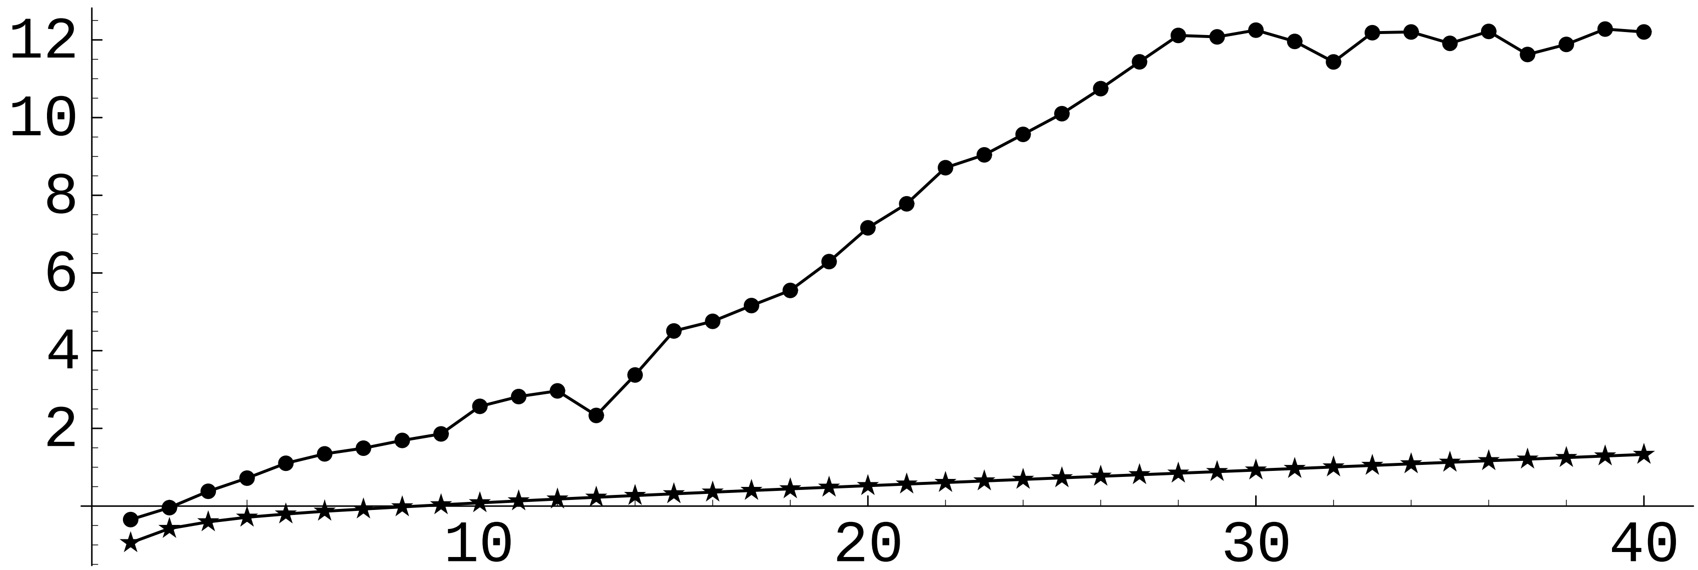

Comparison of Davis and CRDT

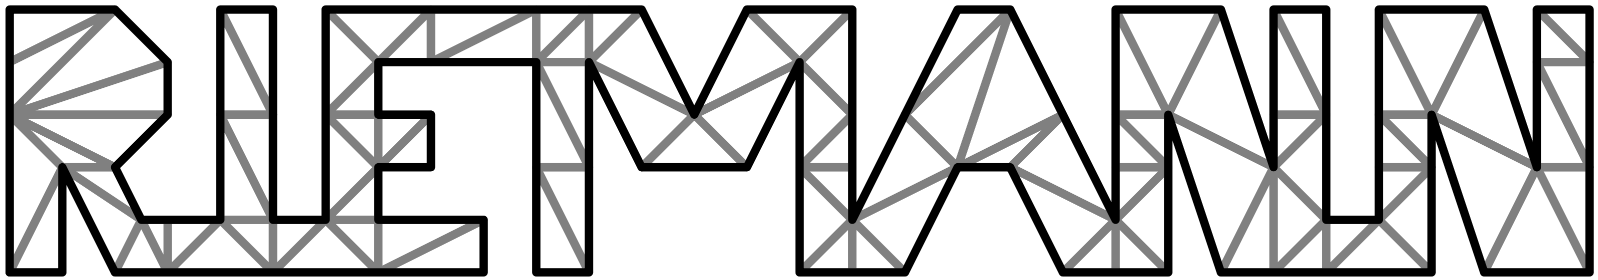

A 98-gon and its Delaunay triangulation.

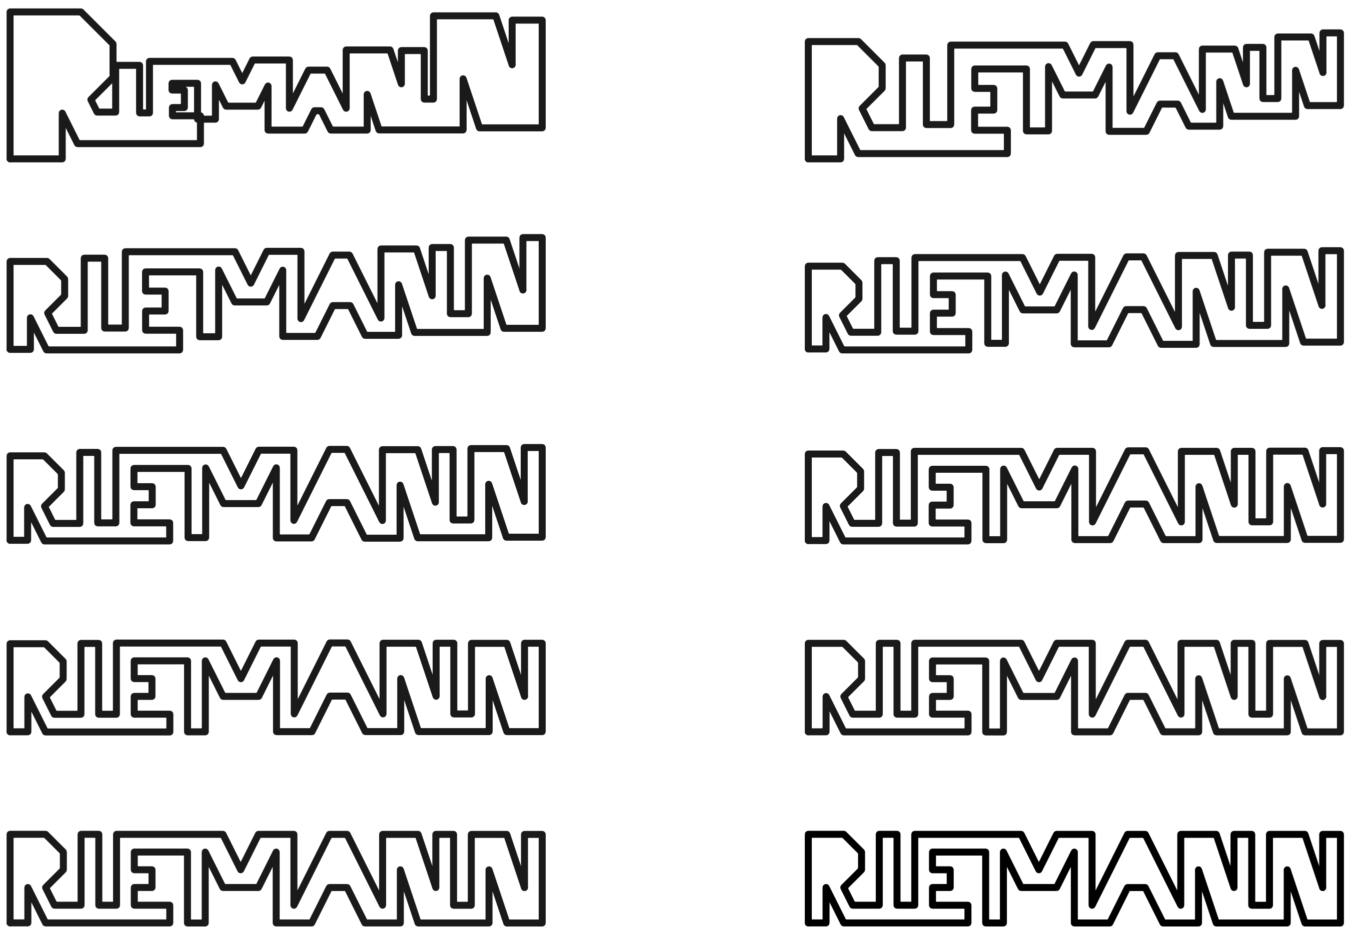

10 iterations of shortcut CRDT applied to a 98-gon.

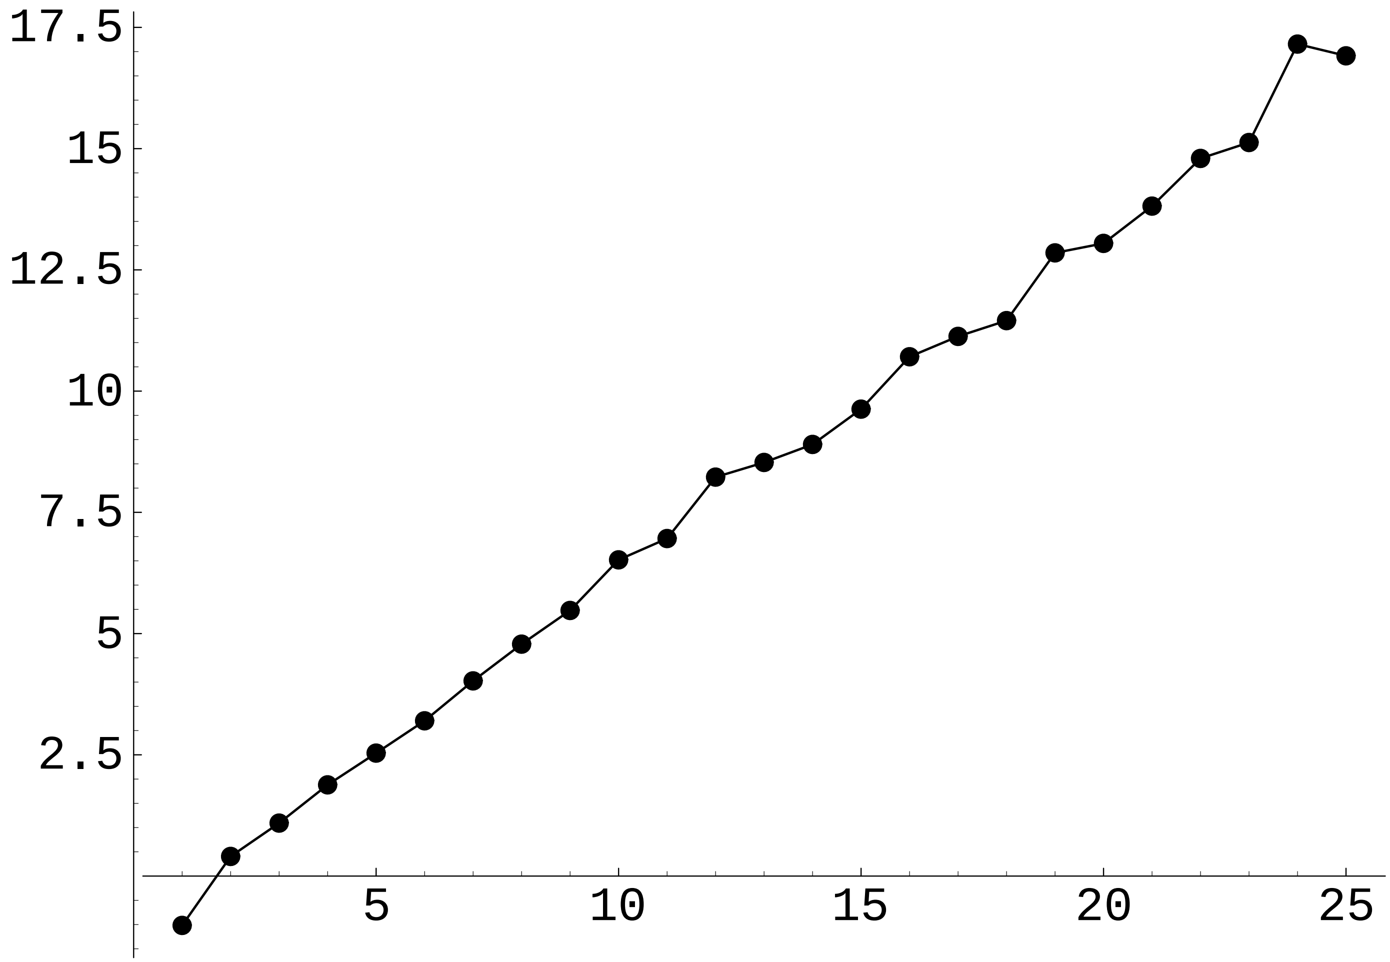

QC error of shortcut CRDT applied to the 98-gon.

Davis, CRDT give conformal maps onto approximate domains.

Instead, guess "approximately conformal" map onto correct polygon.

Iteration makes the map more conformal.

Theorem: Can compute ϵ-conformal map onto n-gon in time $C_\epsilon \cdot n$.

Theorem: Can compute ϵ-conformal map onto n-gon in time $C_\epsilon \cdot n$.

ϵ-conformal $= 1 + \epsilon$ quasiconformal.

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

Data held as $O(n)$ Laurent series of length $p = \log \frac{1}{\epsilon}$.

Bottleneck is doing $O(1)$ FFTs per vertex of polygon.

Gives $O(\epsilon)$ approximation of SC-parameters.

"Scale invariant approximation" = still ϵ-good after rescaling by Möbius transformations = relative gap sizes are ϵ-accurate.

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.

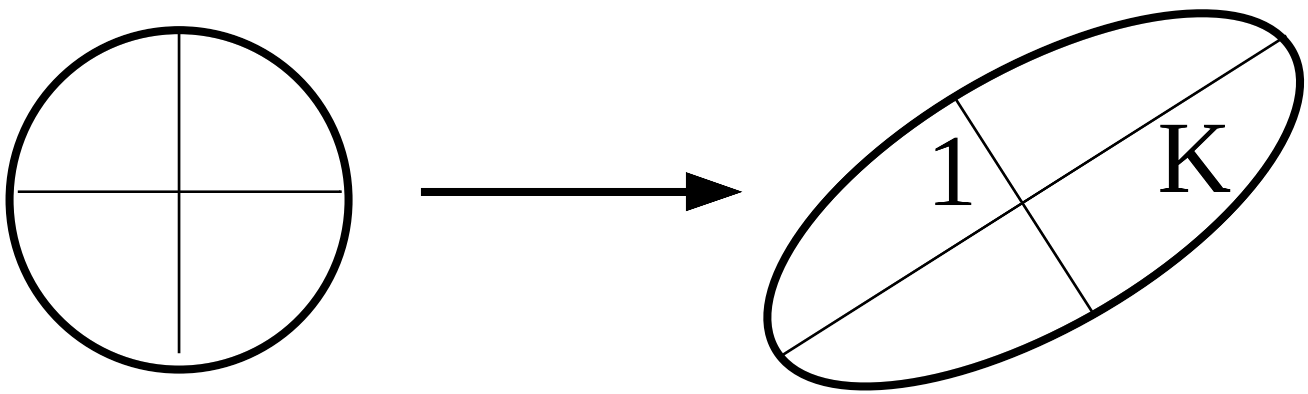

1-QC = conformal

For n-tuples z, w on unit circle,

$$d_{QC}(z, w) = \inf \log K \text{ s.t. } \exists\ K\text{-QC map } h(z) = w.$$

QC approximation implies "scale invariant" Euclidean approximation.

QC maps are differentiable almost everywhere.

$$f_z = f_x - i f_y \qquad f_{\overline{z}} = f_x + i f_y.$$

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

μ encodes amount and direction of stretching.

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

Homeomorphic solutions of Beltrami equation

$$f_z = \mu f_{\overline{z}}$$

where $|\mu| < k < 1$.

Measurable Riemann Mapping Theorem:

Given QC $g : D \to \Omega$, there is $f : D \to D$ so that $g \circ f$ is conformal.

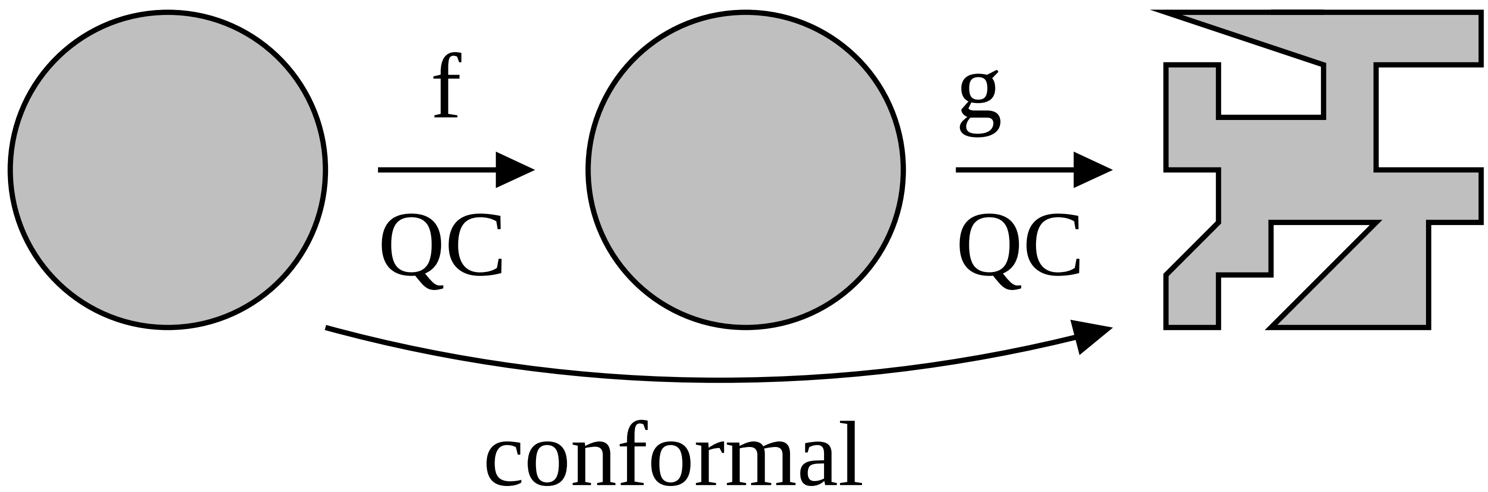

Measurable Riemann Mapping Theorem:
Given QC $g : D \to \Omega$, there is $f : D \to D$ so that $g \circ f$ is conformal.

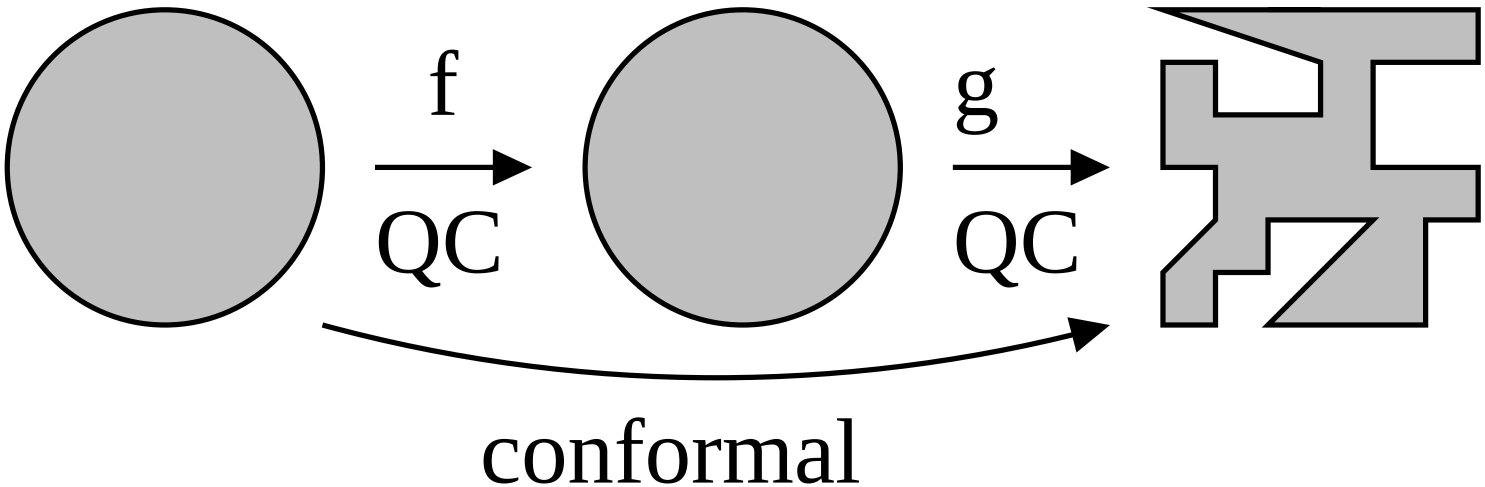

Need to solve Beltrami equation $f_z = \mu_g f_{\overline{z}}$.

Explicit exact solution by power series of singular integral operators.

Linearization can be solved by convolution with $1/z$.

Iteration converges quadratically to solution (Newton's method).

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

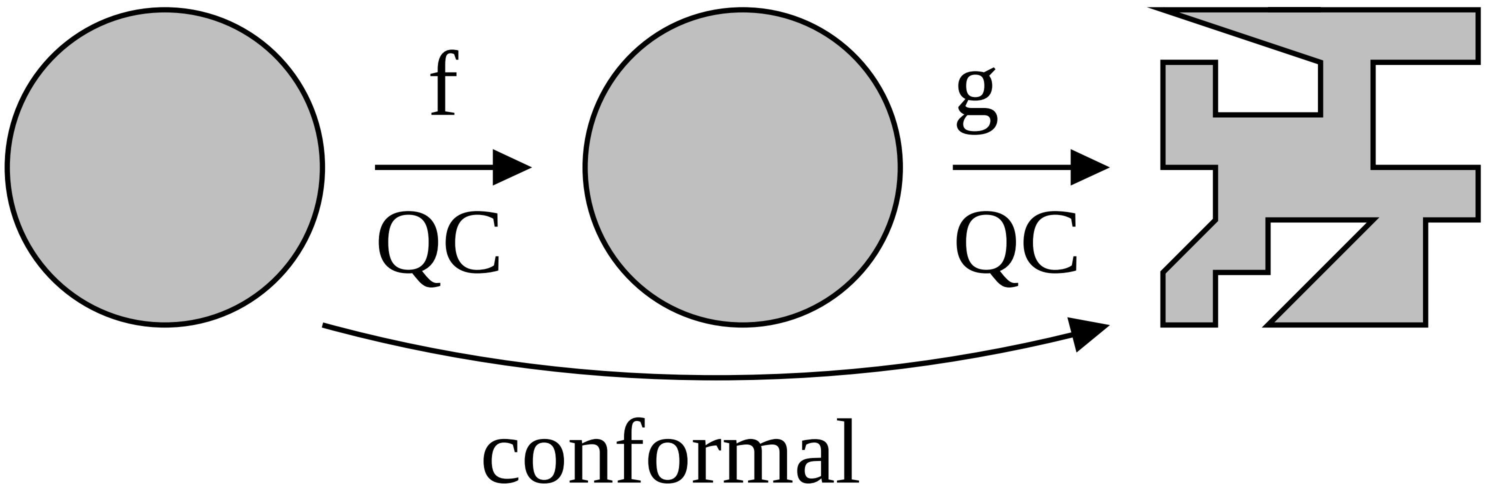

- Find initial QC map g to polygon.
- Solve a PDE for $f : D \to D$.

I will only address finding a good g.

Want g to be **fast** to compute and guaranteed **close** to correct answer.

Want g to be **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.

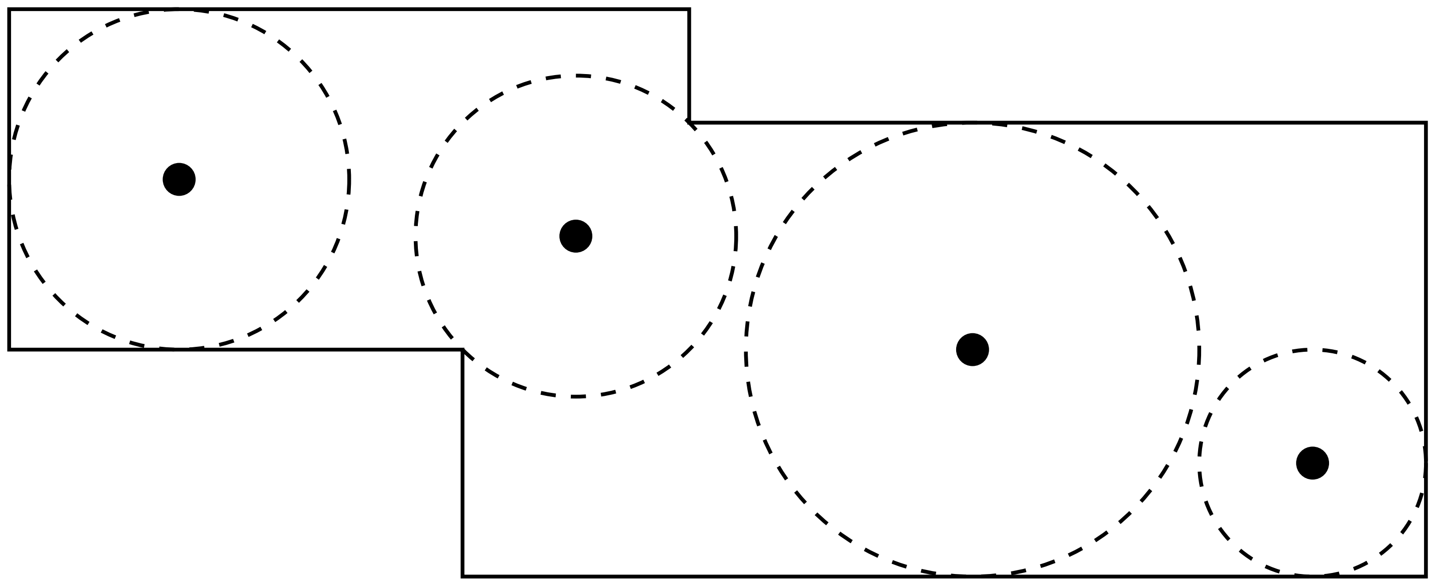

Medial axis:
centers of disks that hit boundary in at least two points.

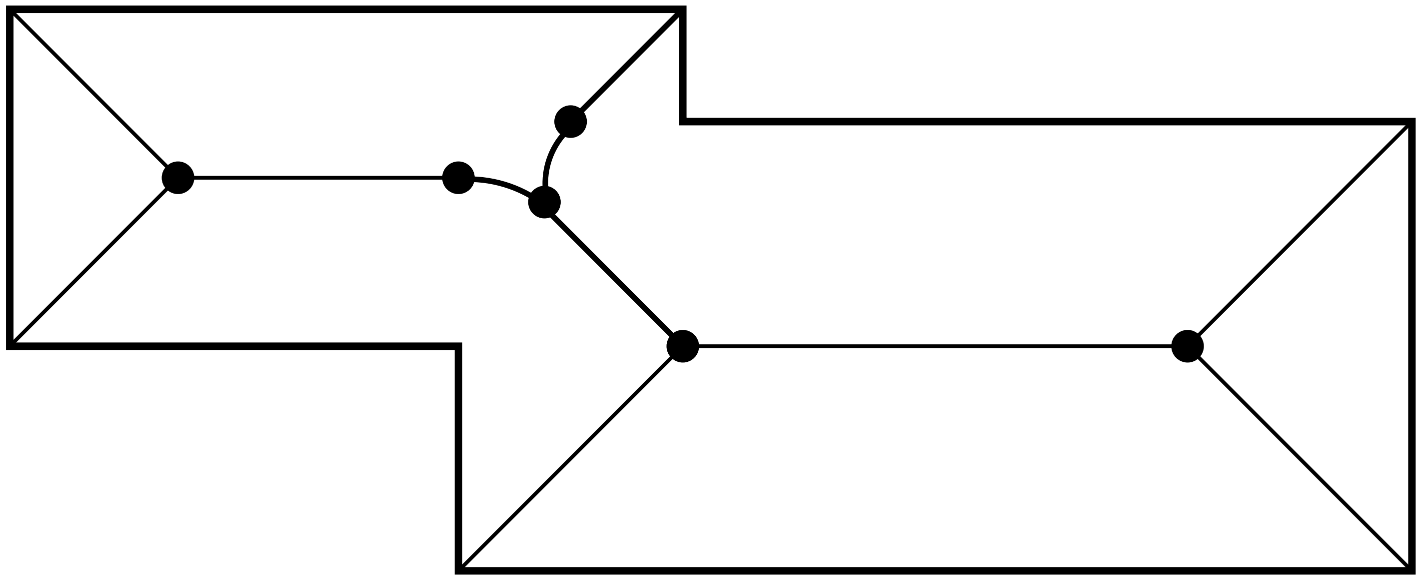

Medial axis of a polygon is a finite tree.

Computable in O(n), Chin-Snoeyink-Wang (1999).

Related to Voronoi diagrams: divides polygon according to nearest edge.

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

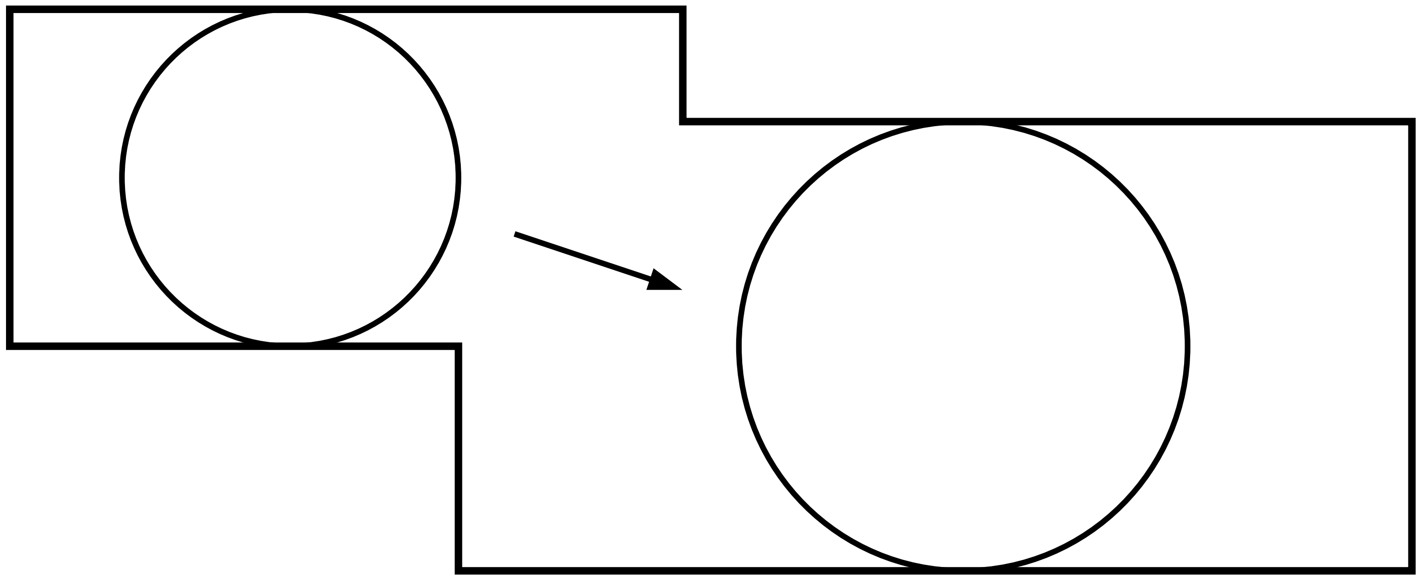

Claim: there is a "natural" choice of conformal map between any two medial axis disks.

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

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

Conformally maps disks to disks (or half-planes).

Form a group under composition.

Uniquely determined by images of 3 distinct points.

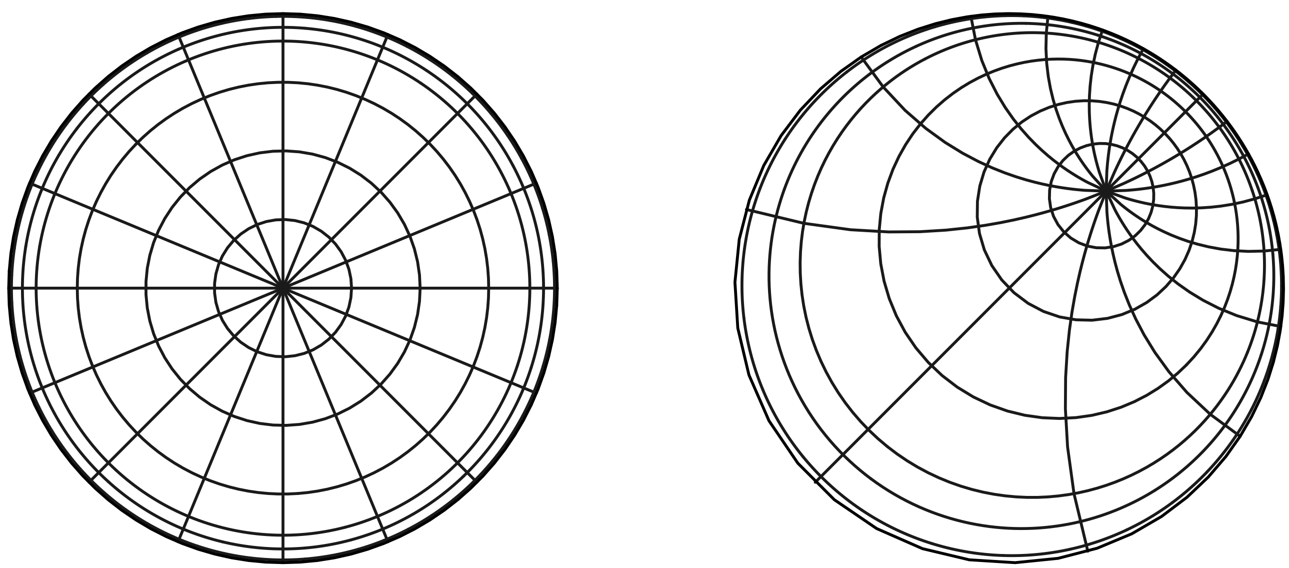

Intersecting circles:

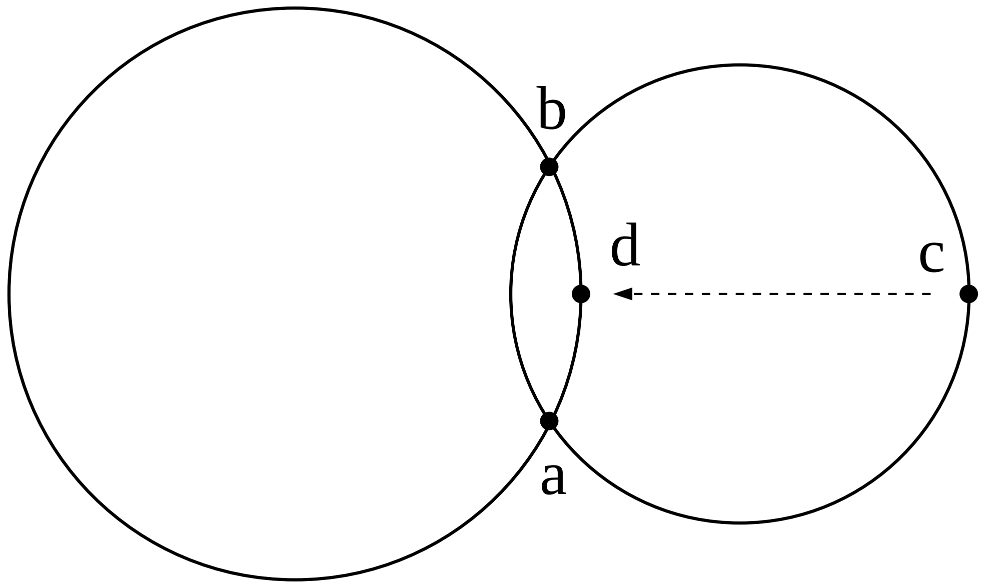

Fix intersection points a, b and map $c \to d$ as shown.

Determines unique Möbius map between disks.

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

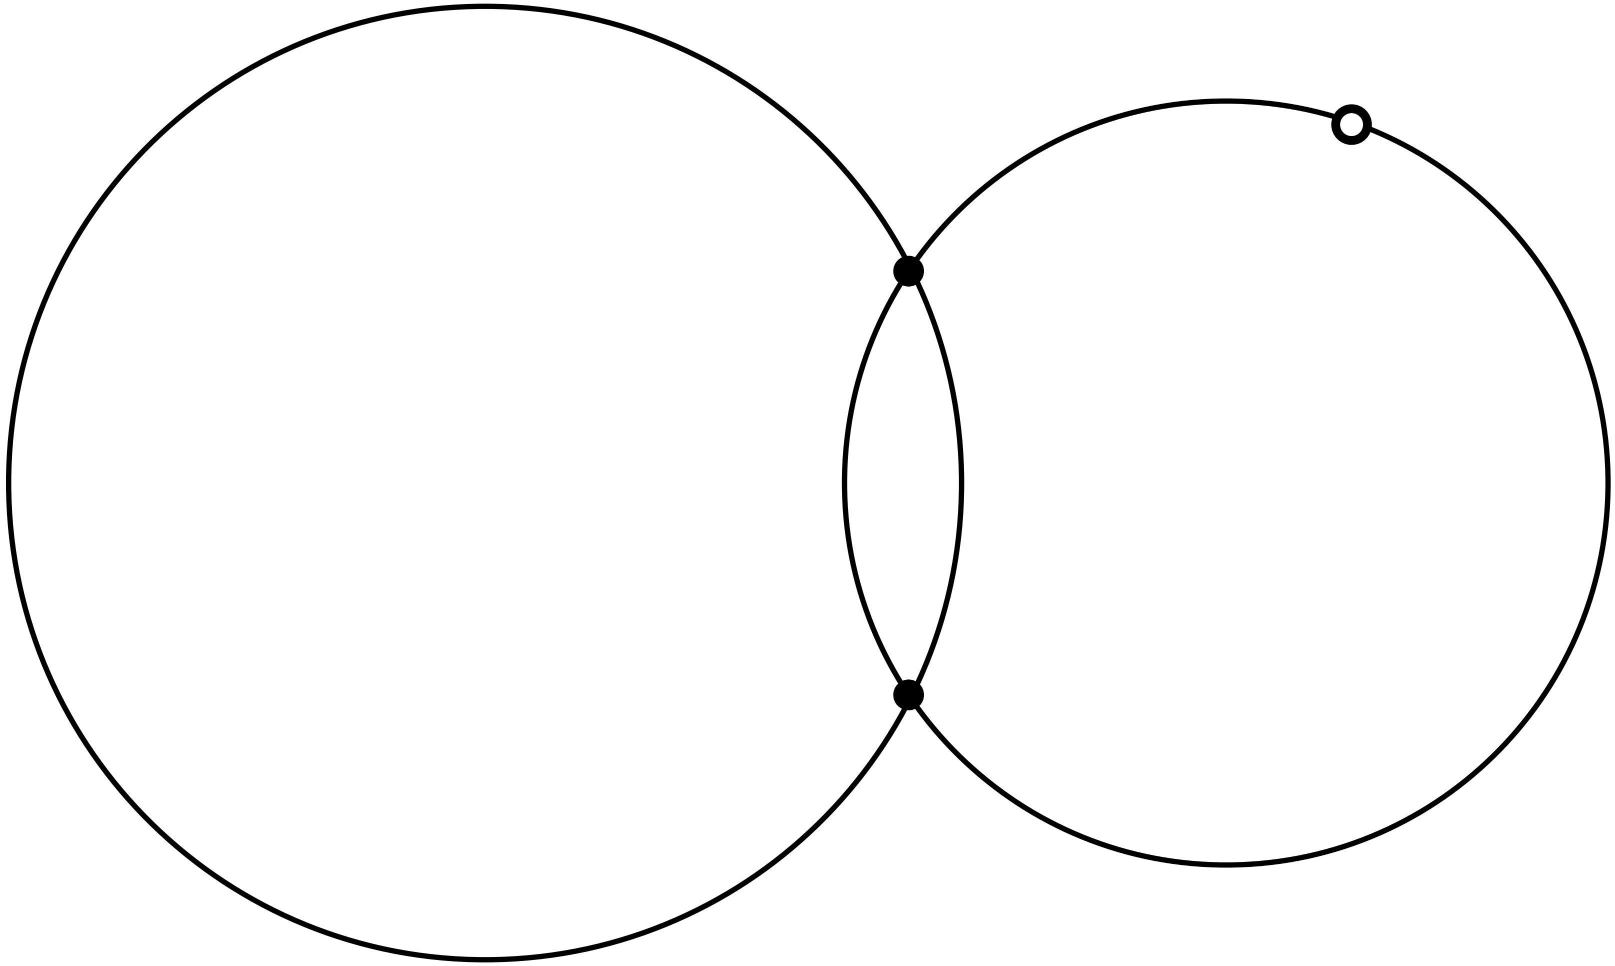

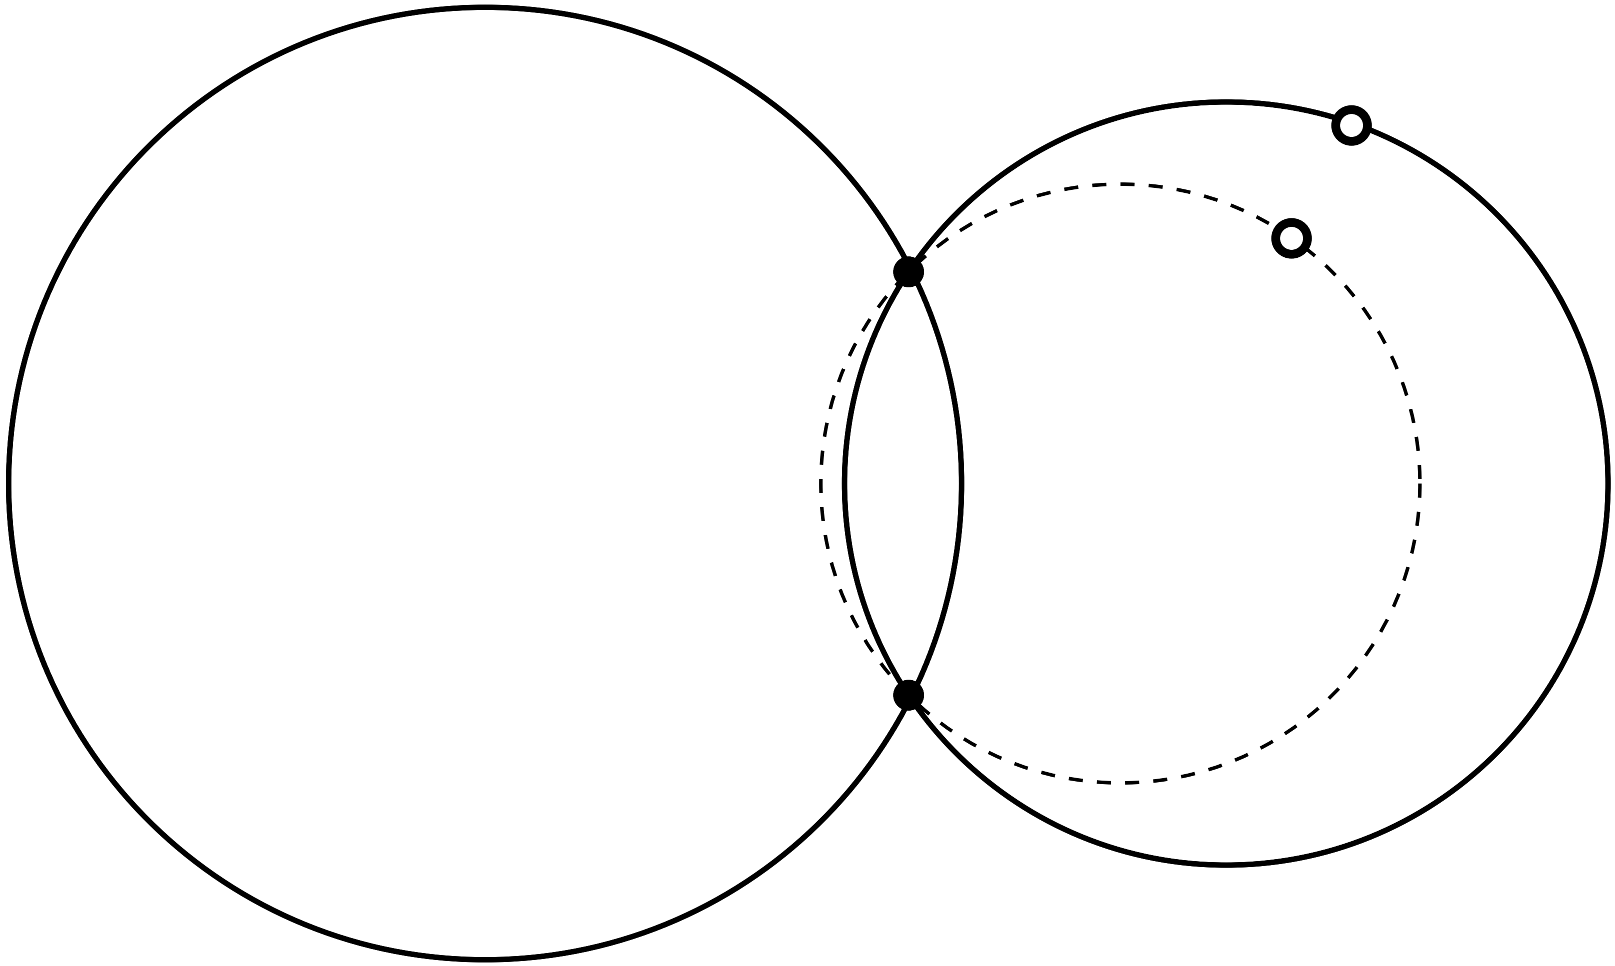

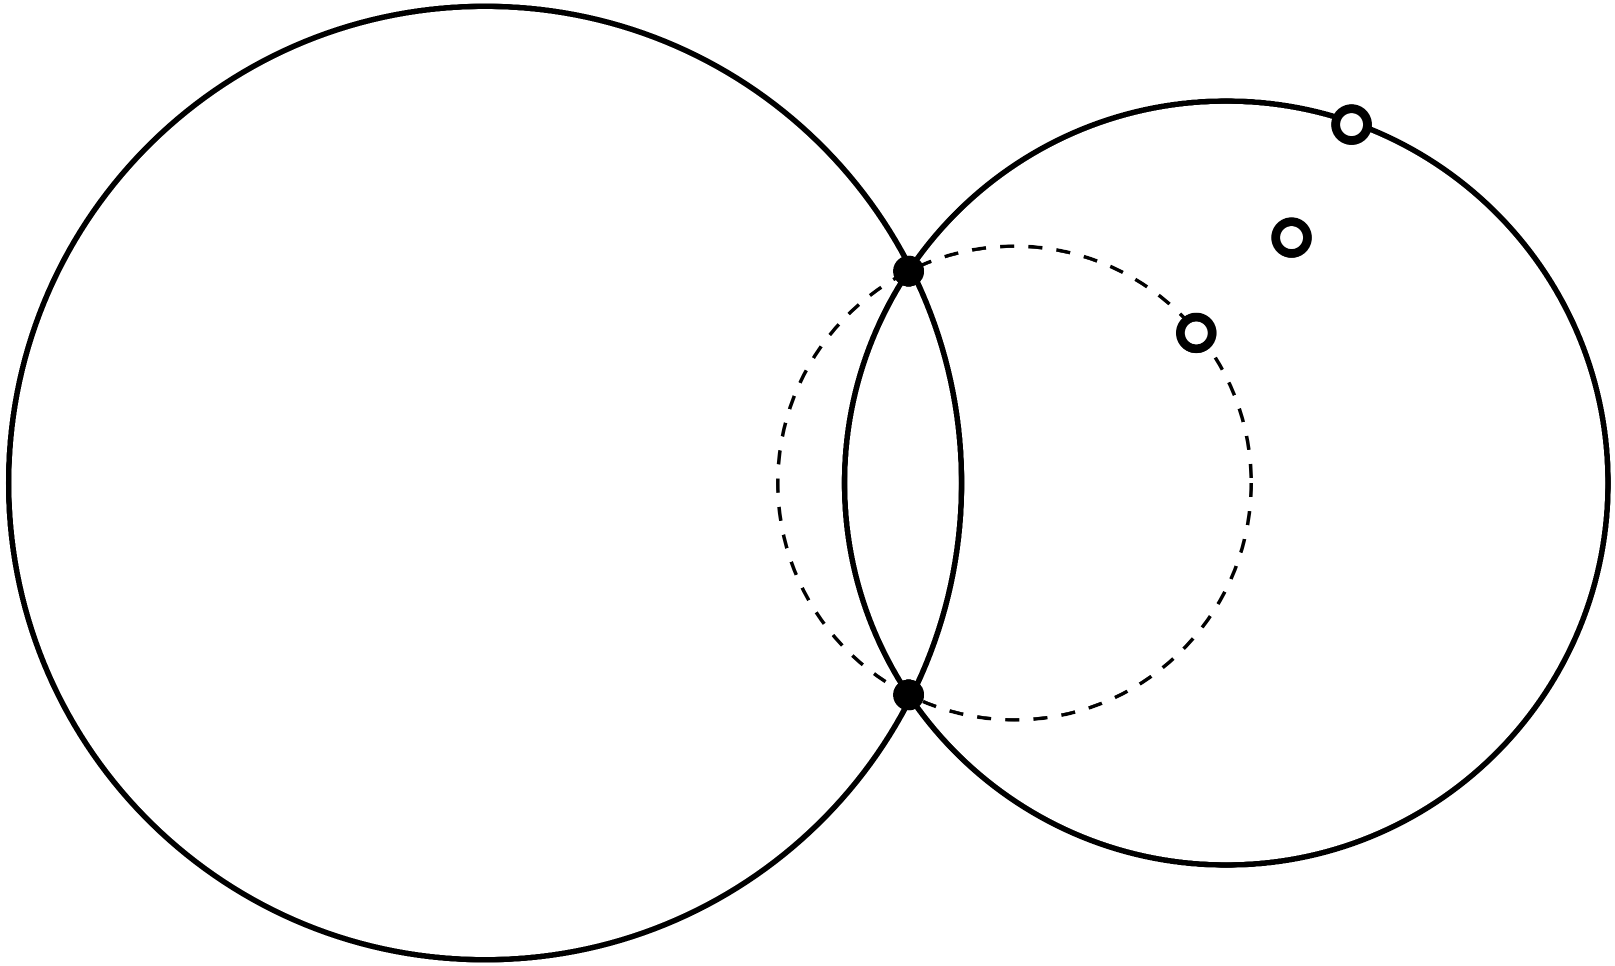

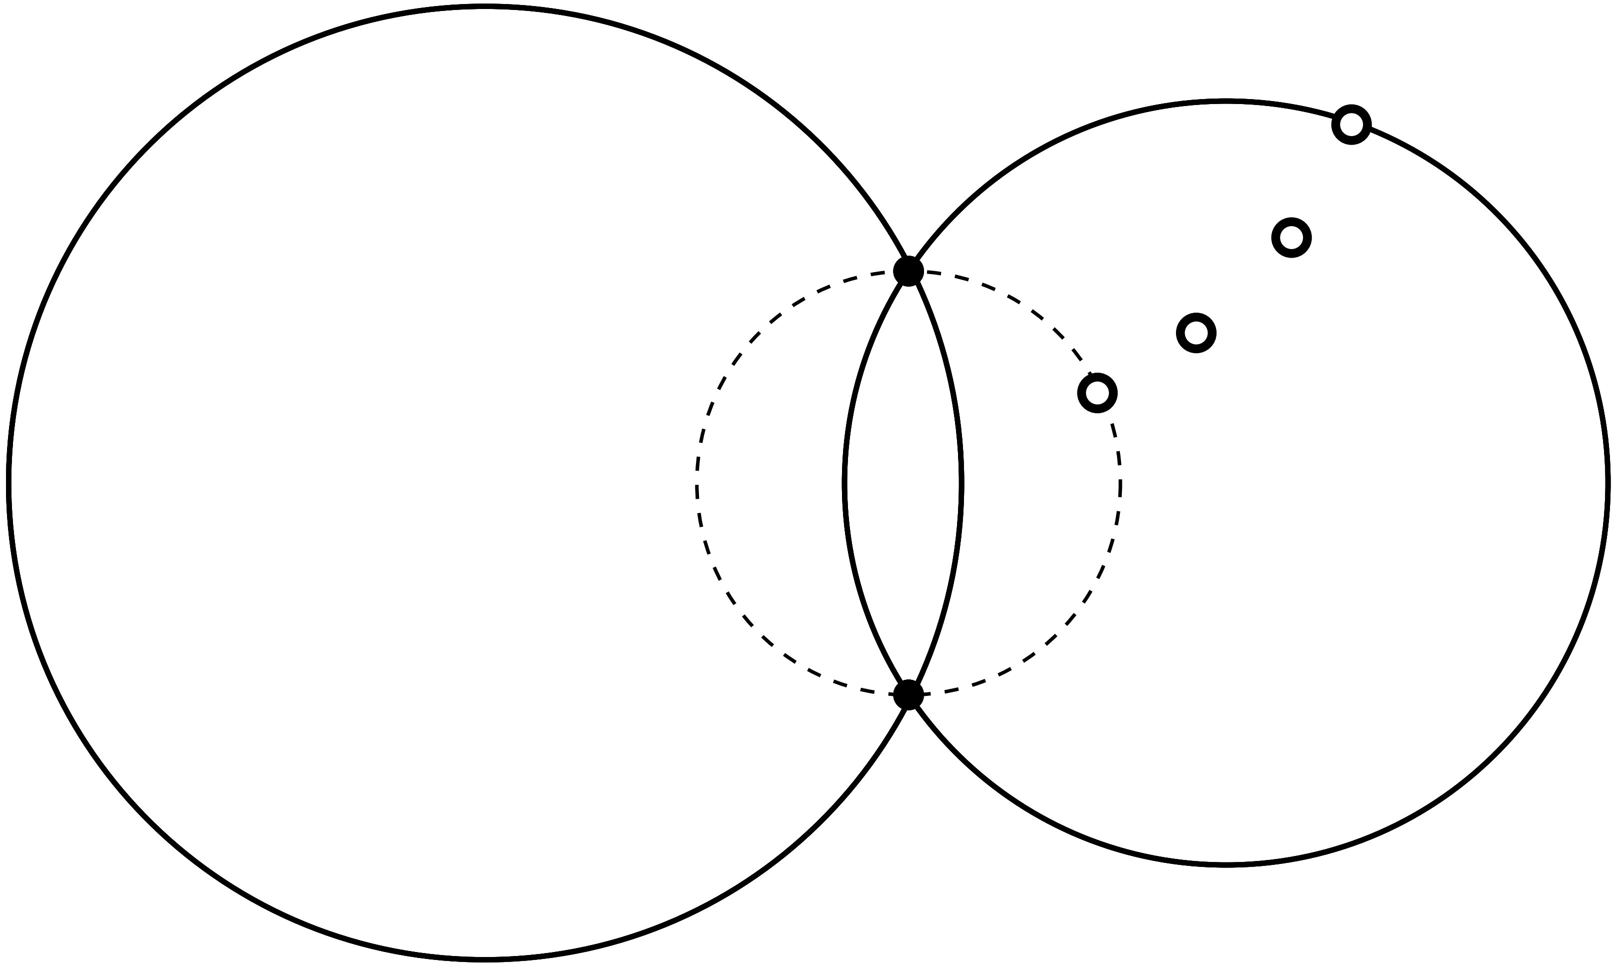

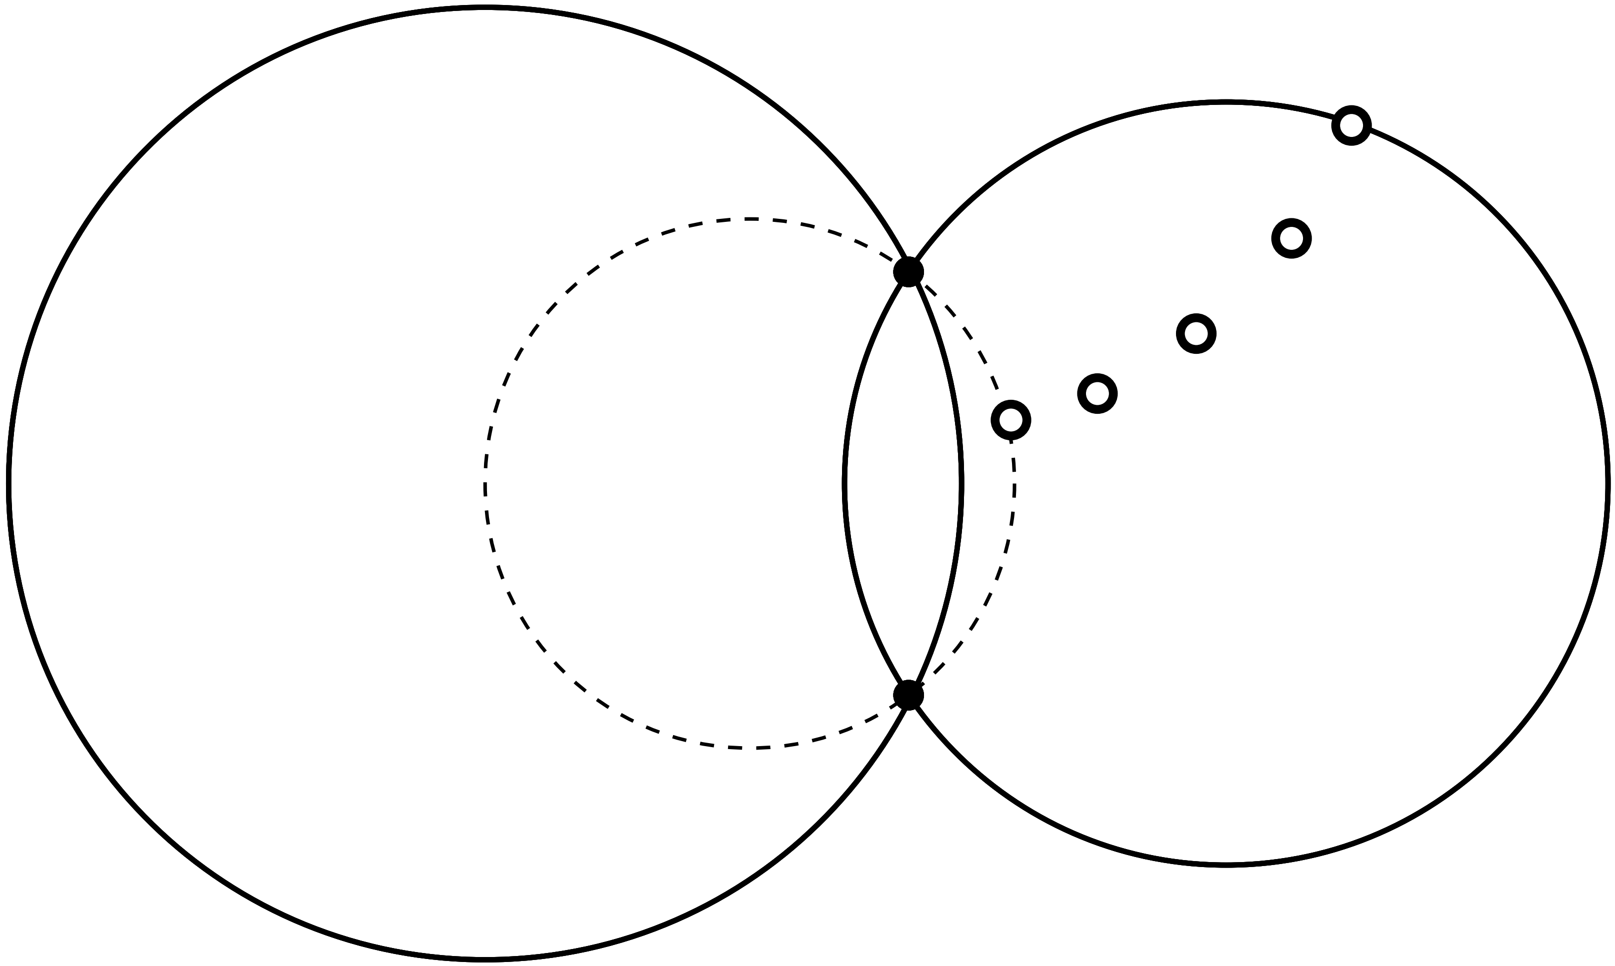

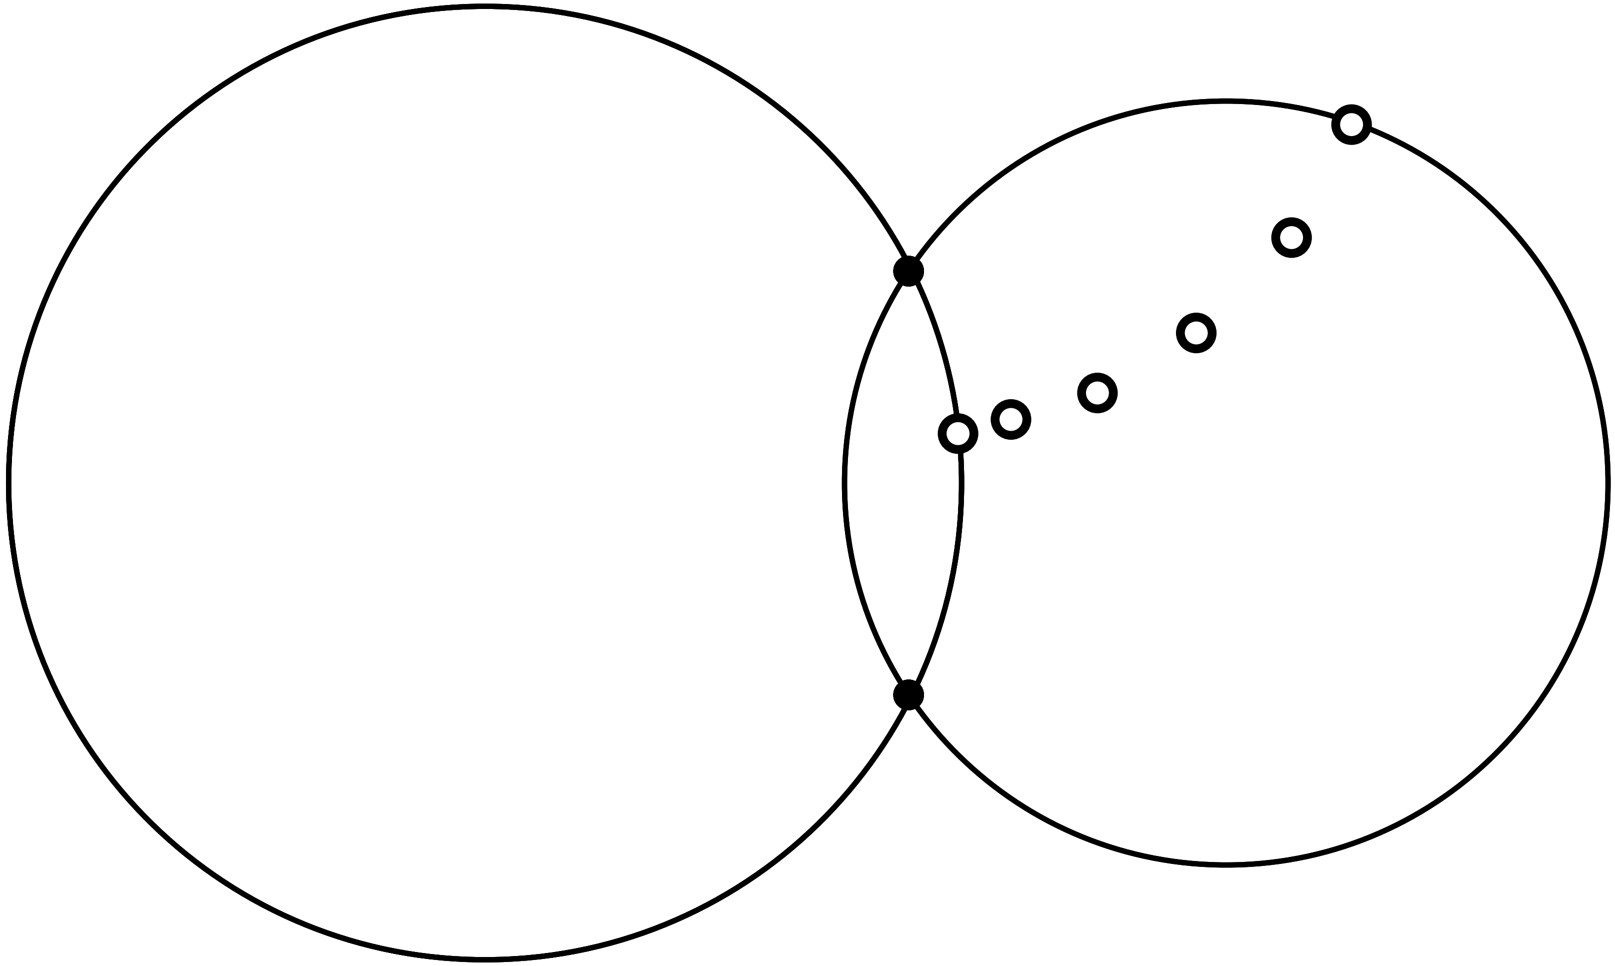

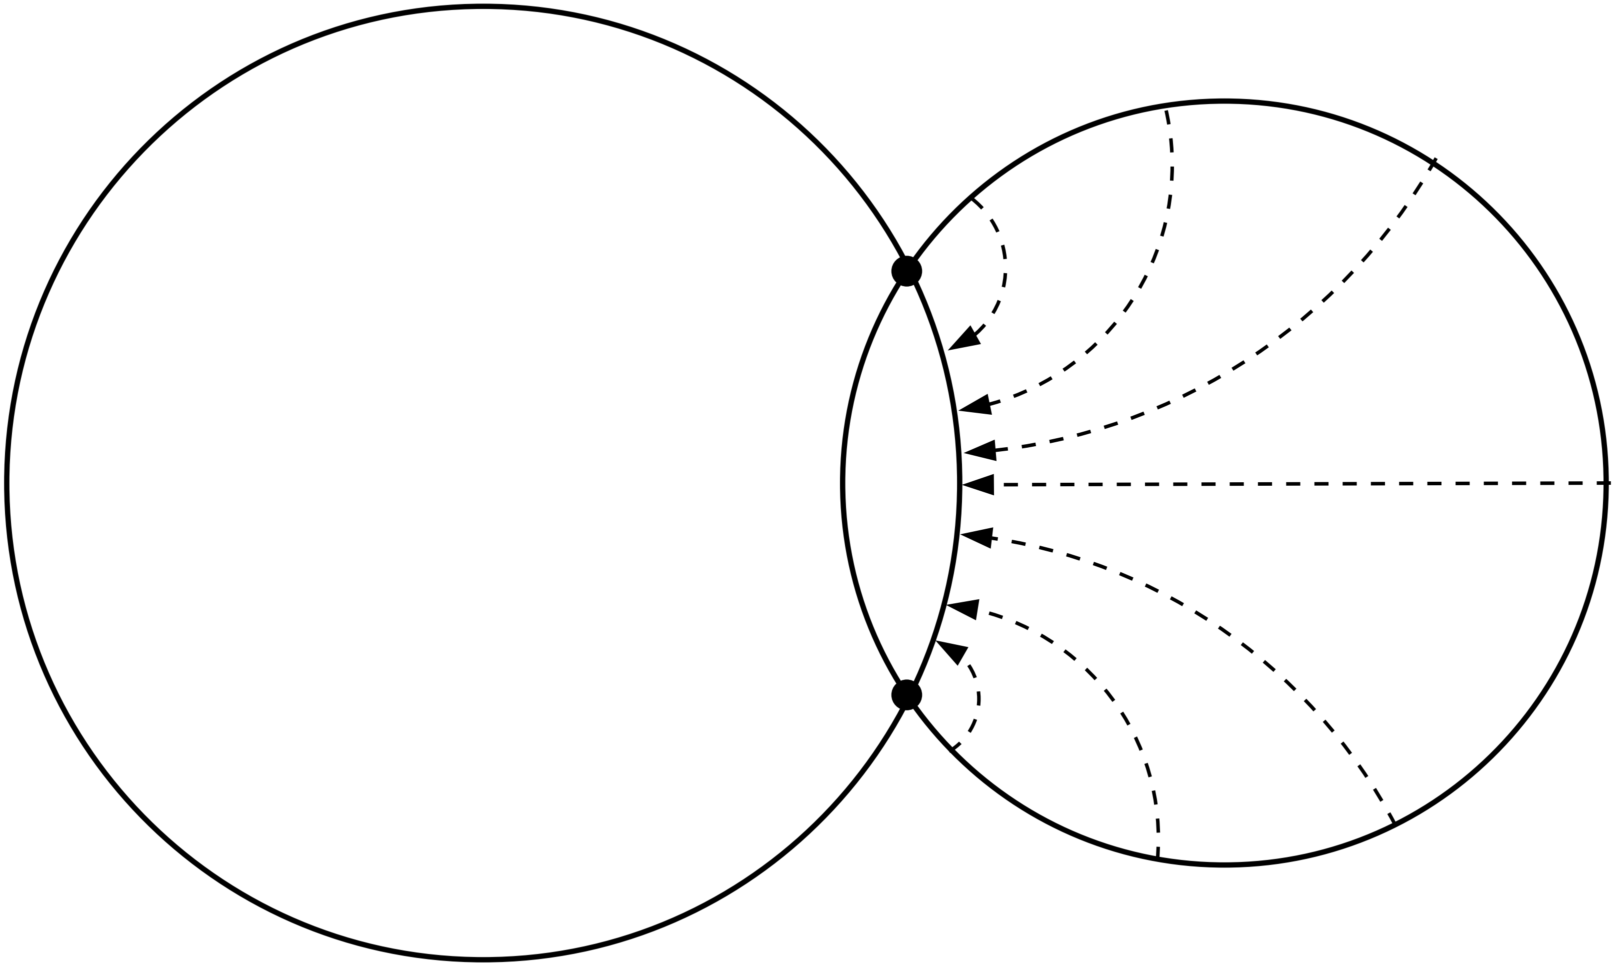

Points follow circular paths, perpendicular to boundary.

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

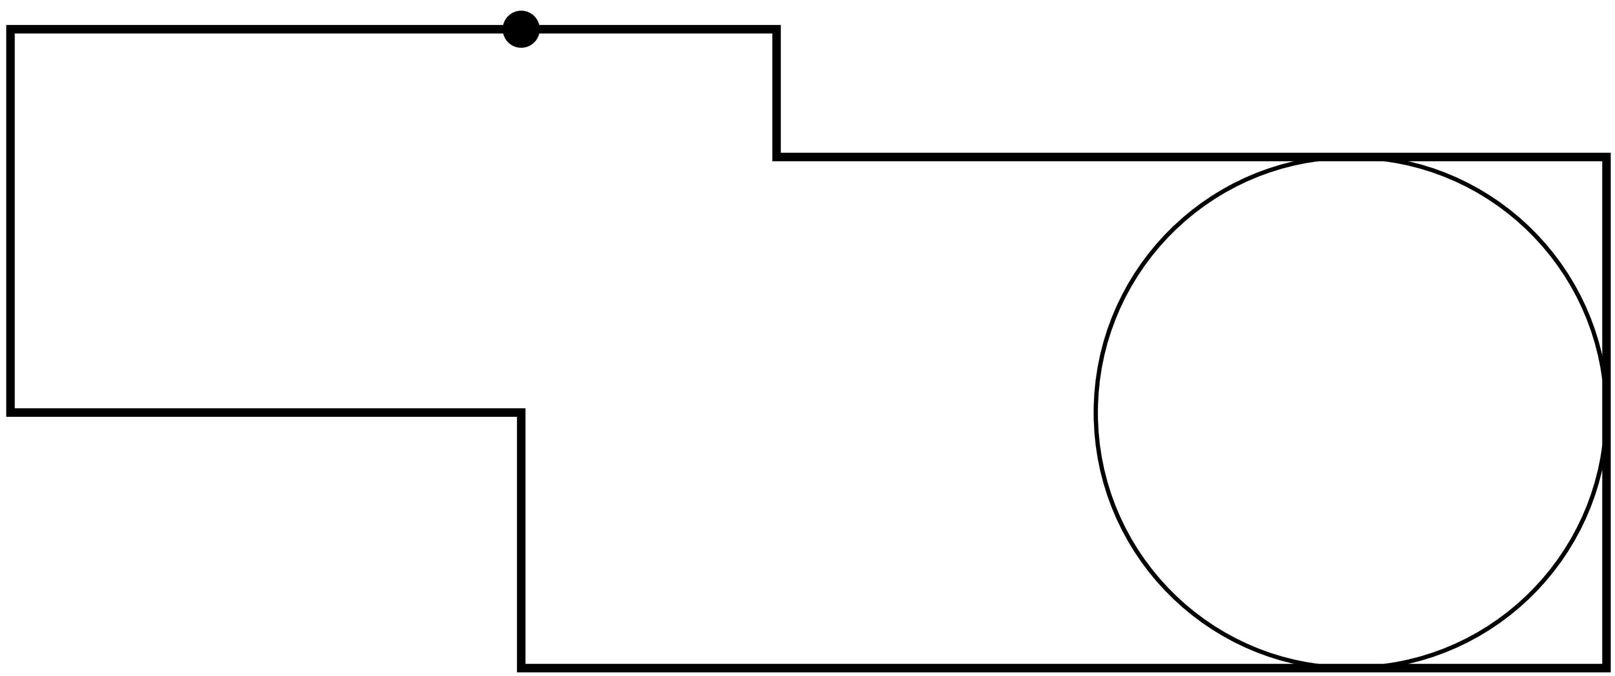

- Fix a “root” MA disk D.

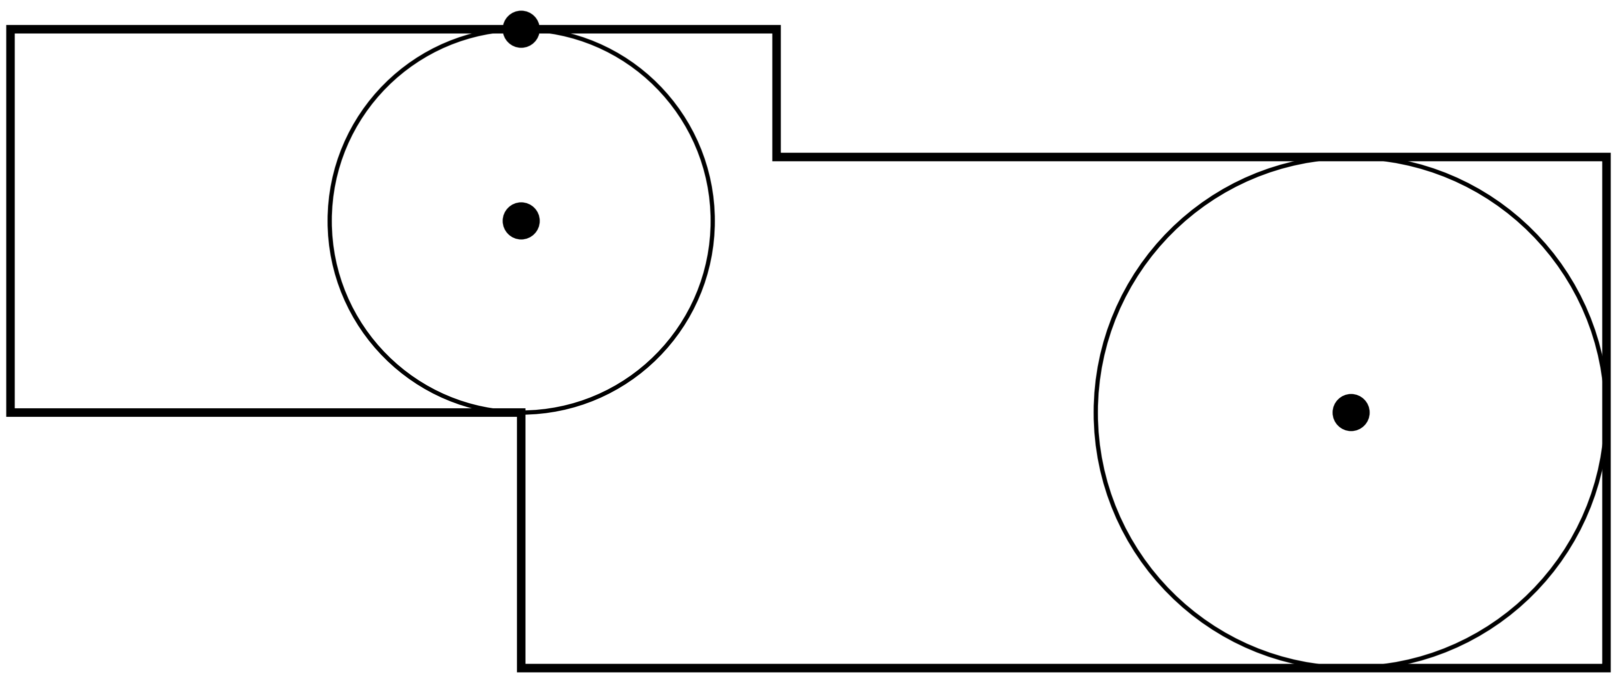

- For any $z \in P$, take MA disk D_z touching z.

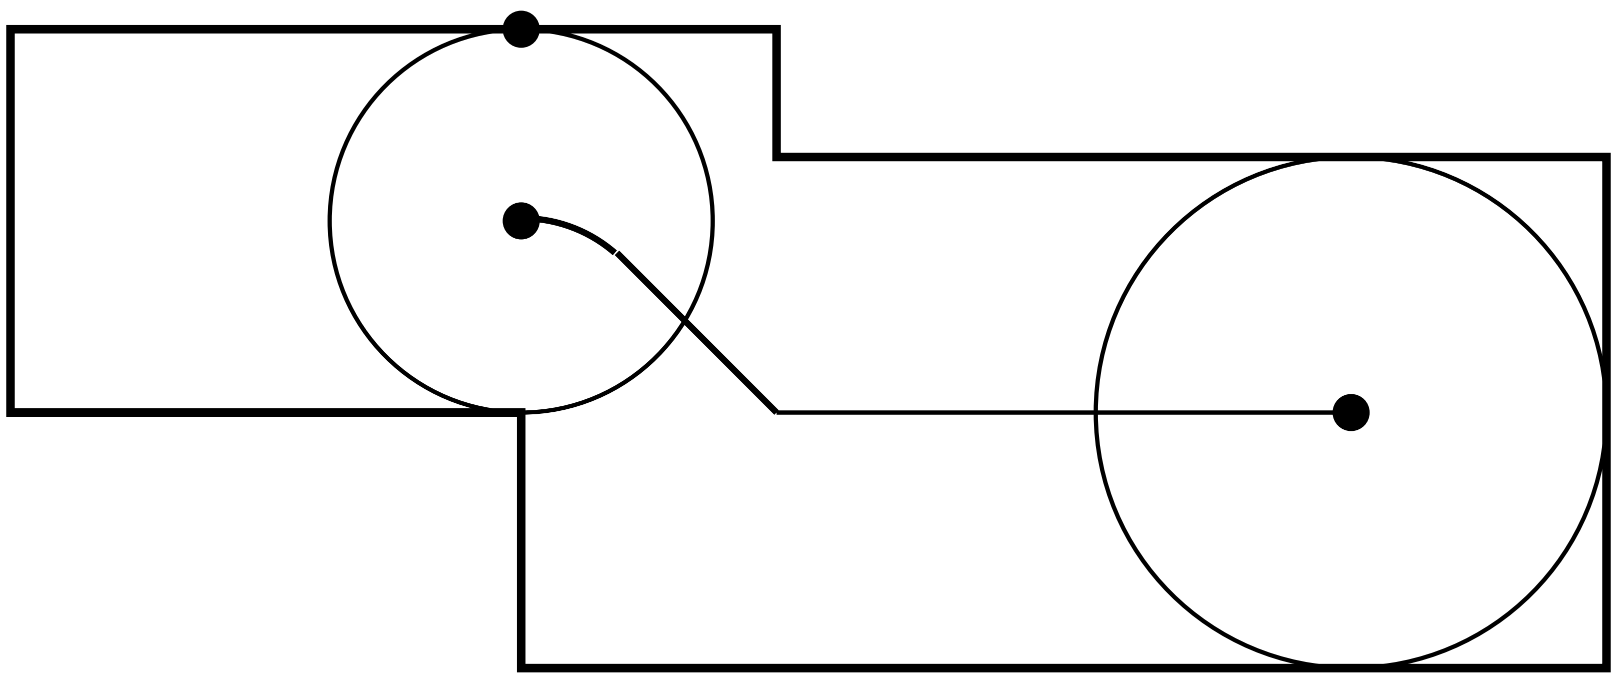

- Connect D_z to D on MA.

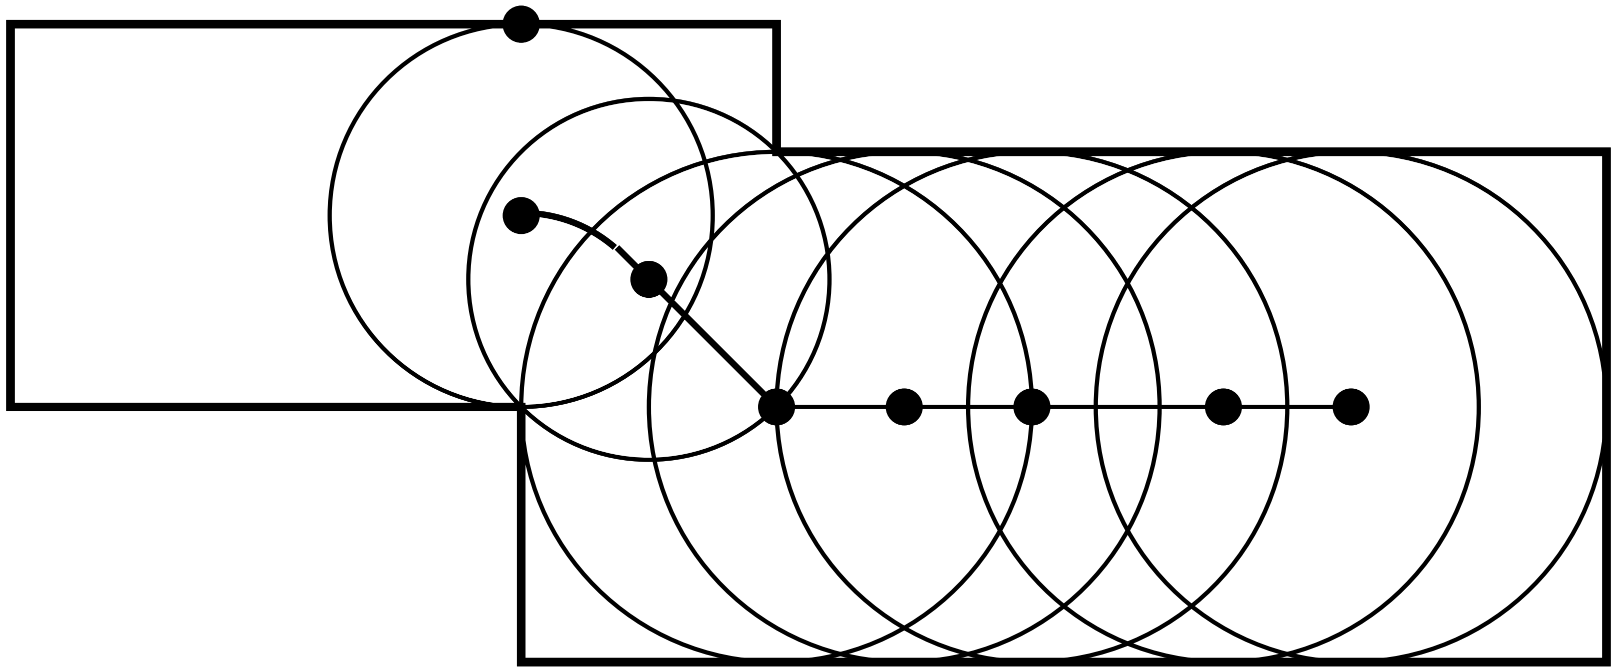

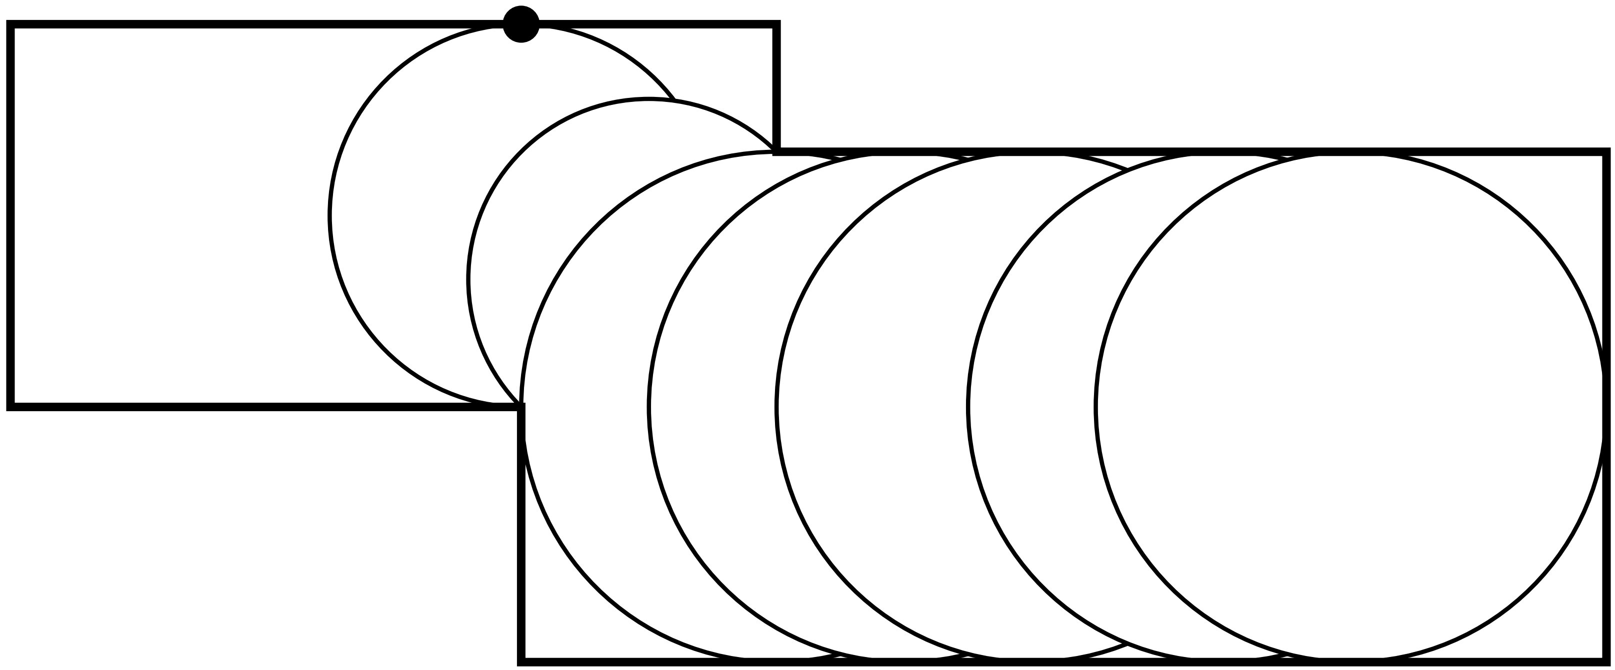

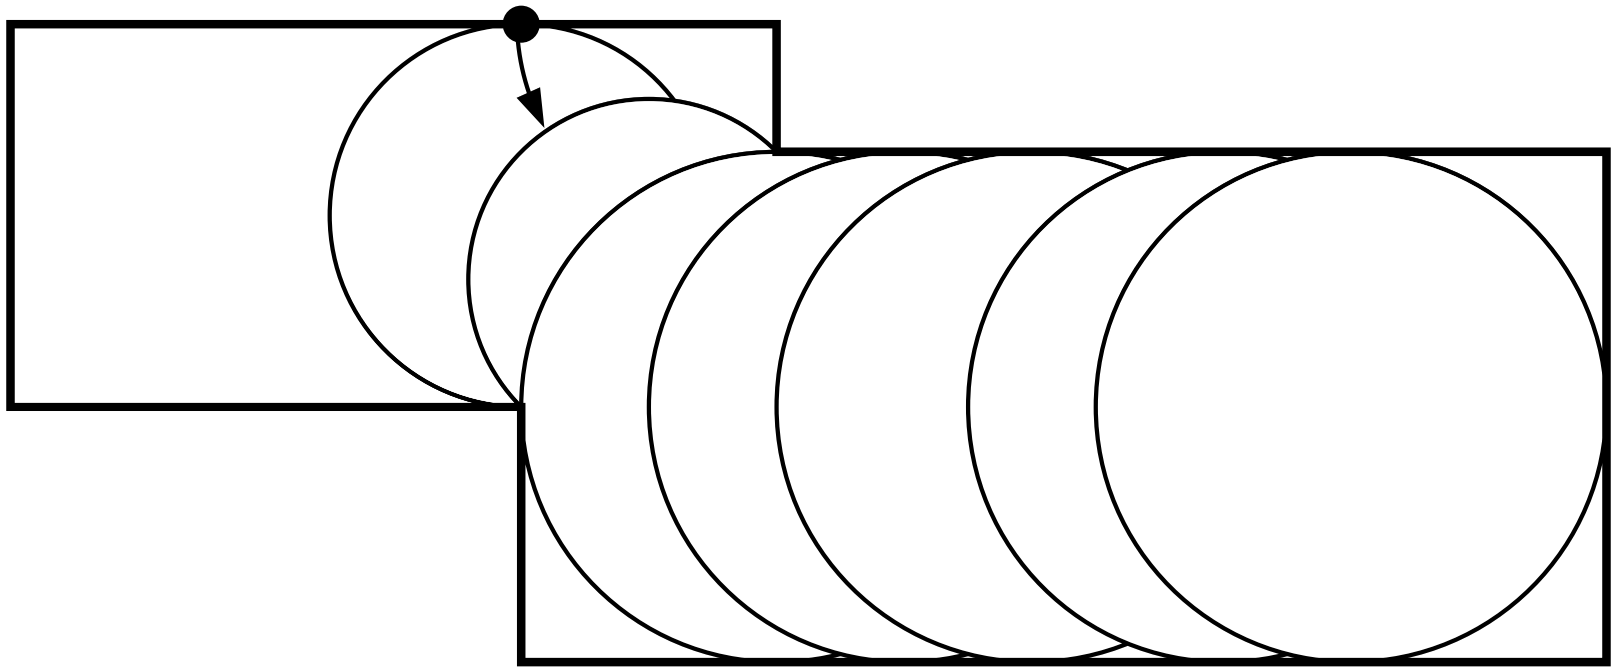

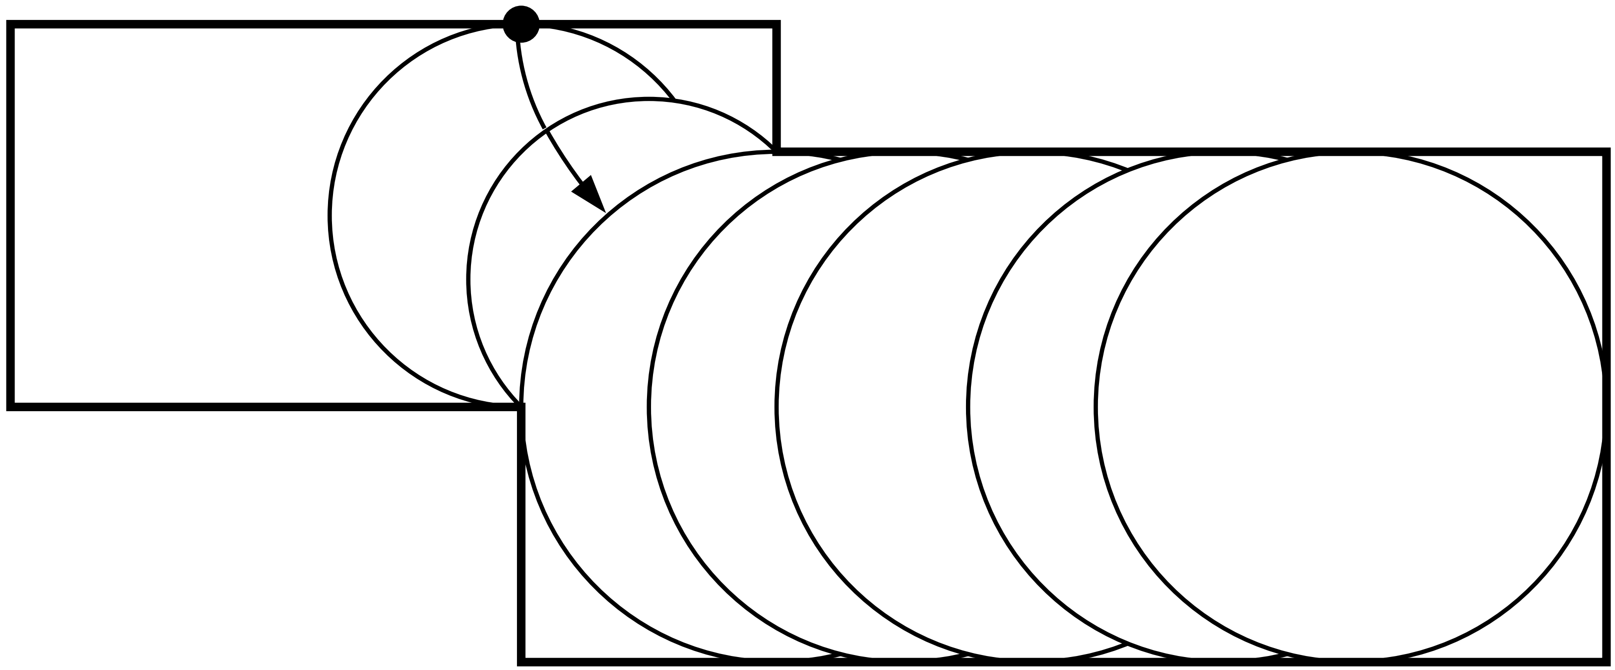

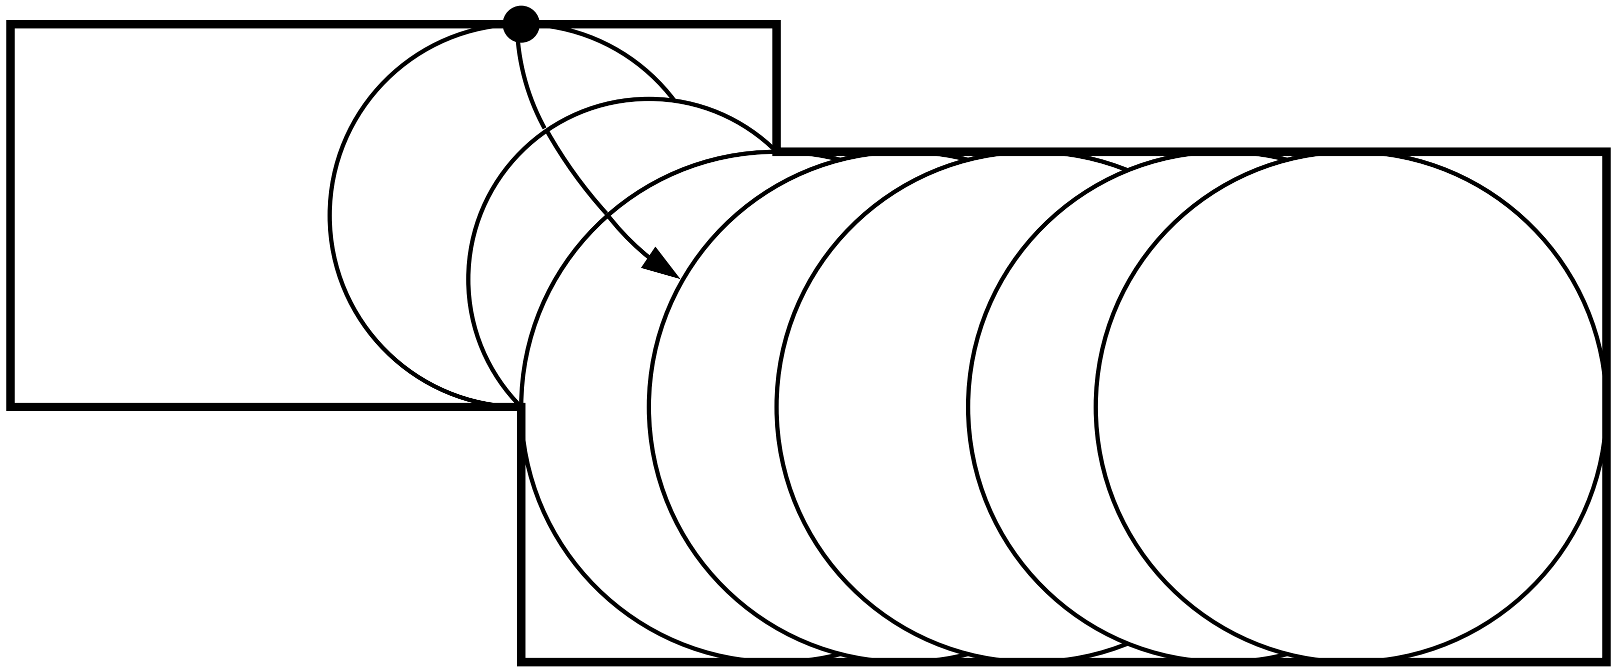

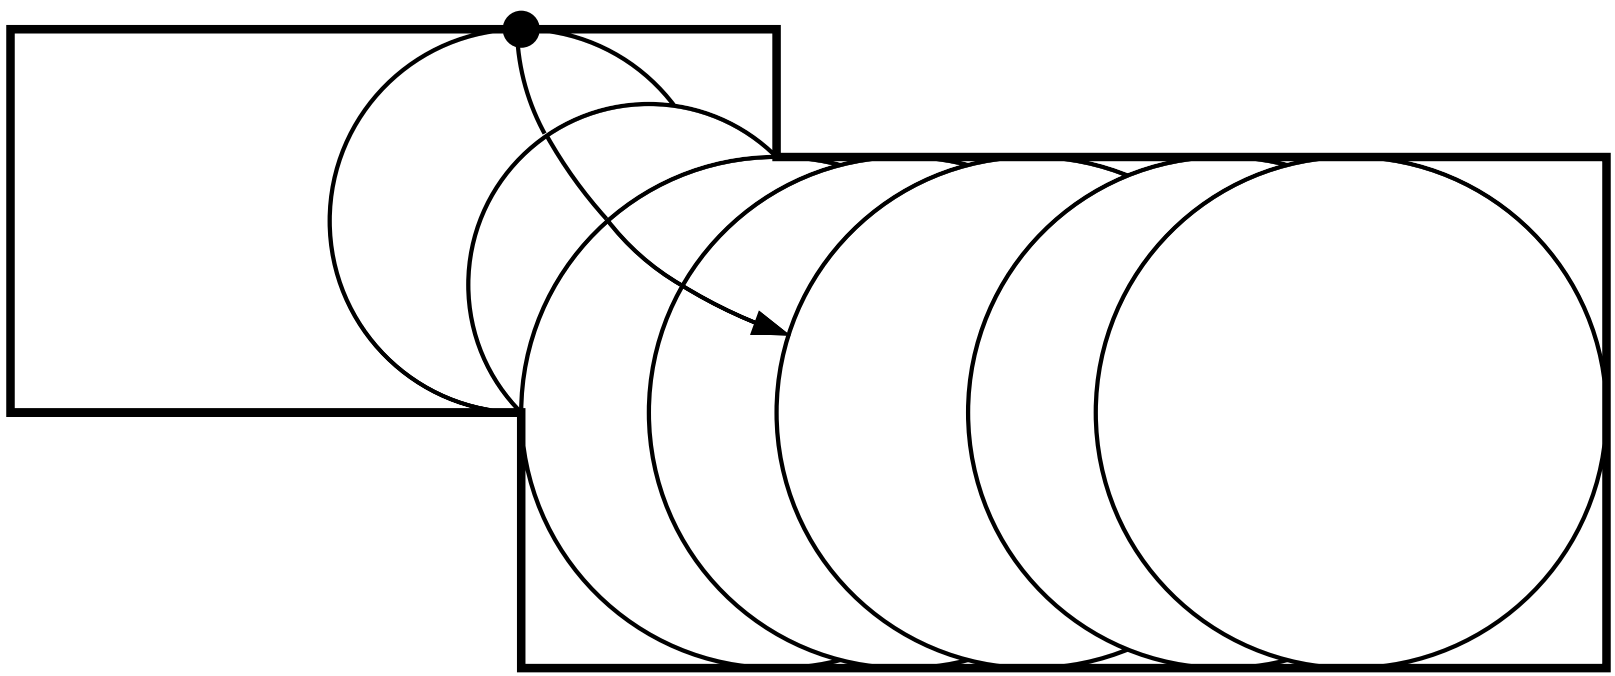

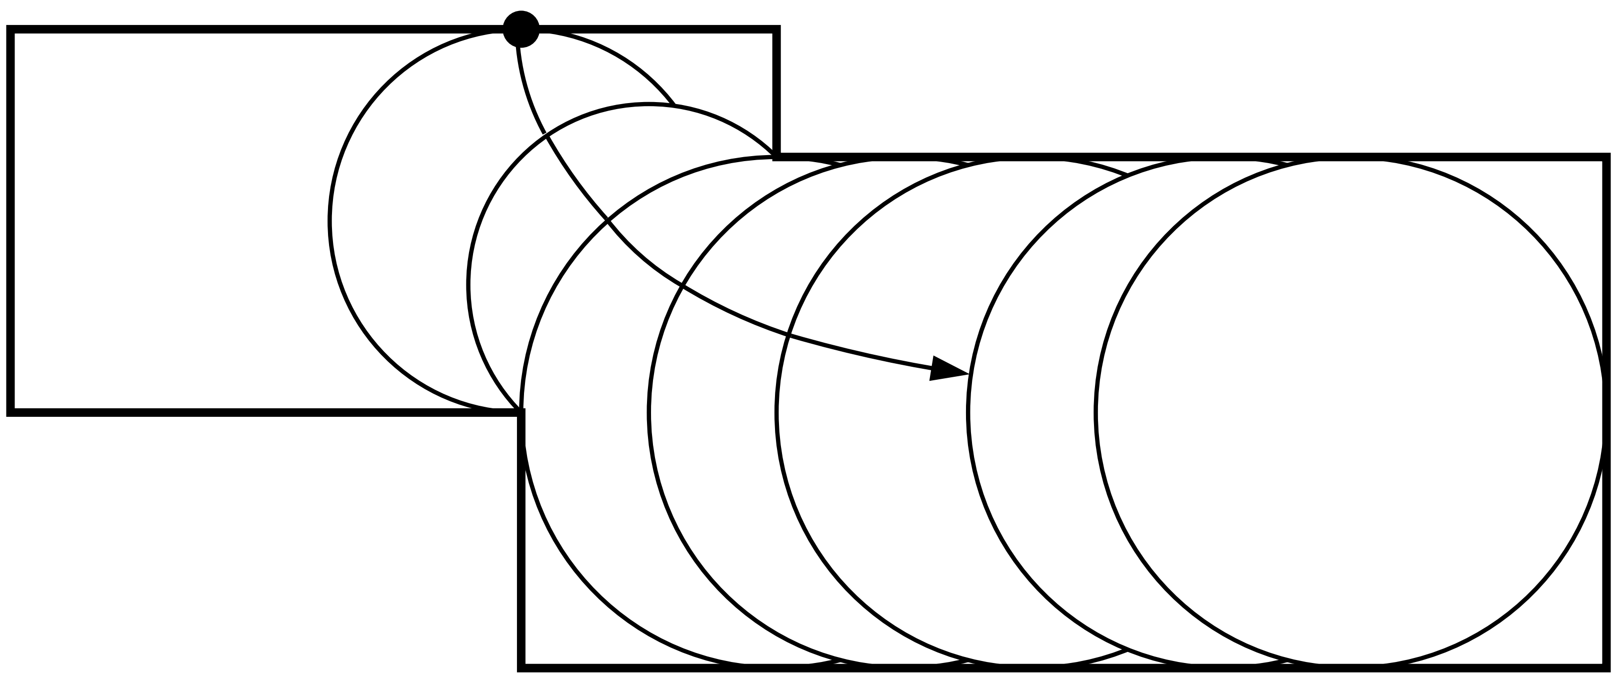

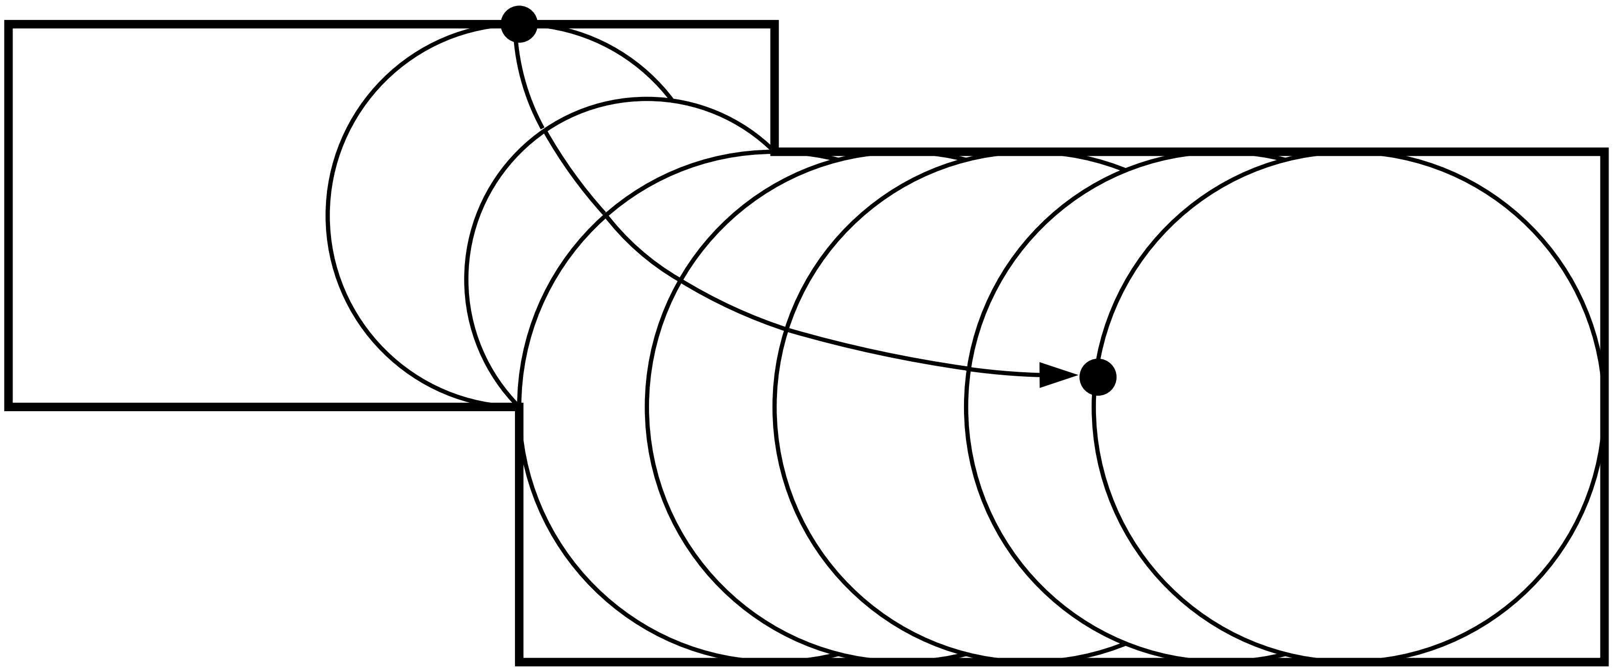

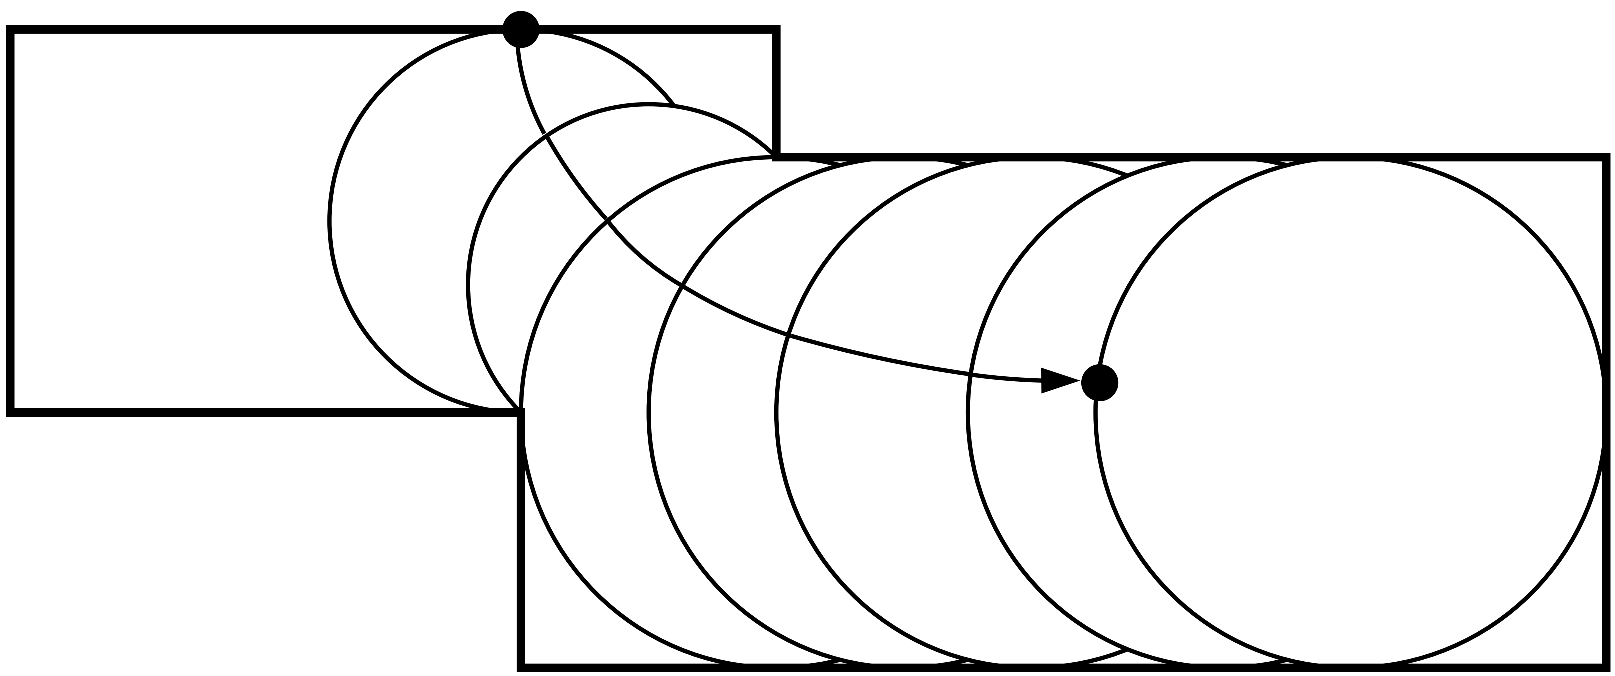

We discretize only to draw picture.

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

Similar flow for any simply connected domain.

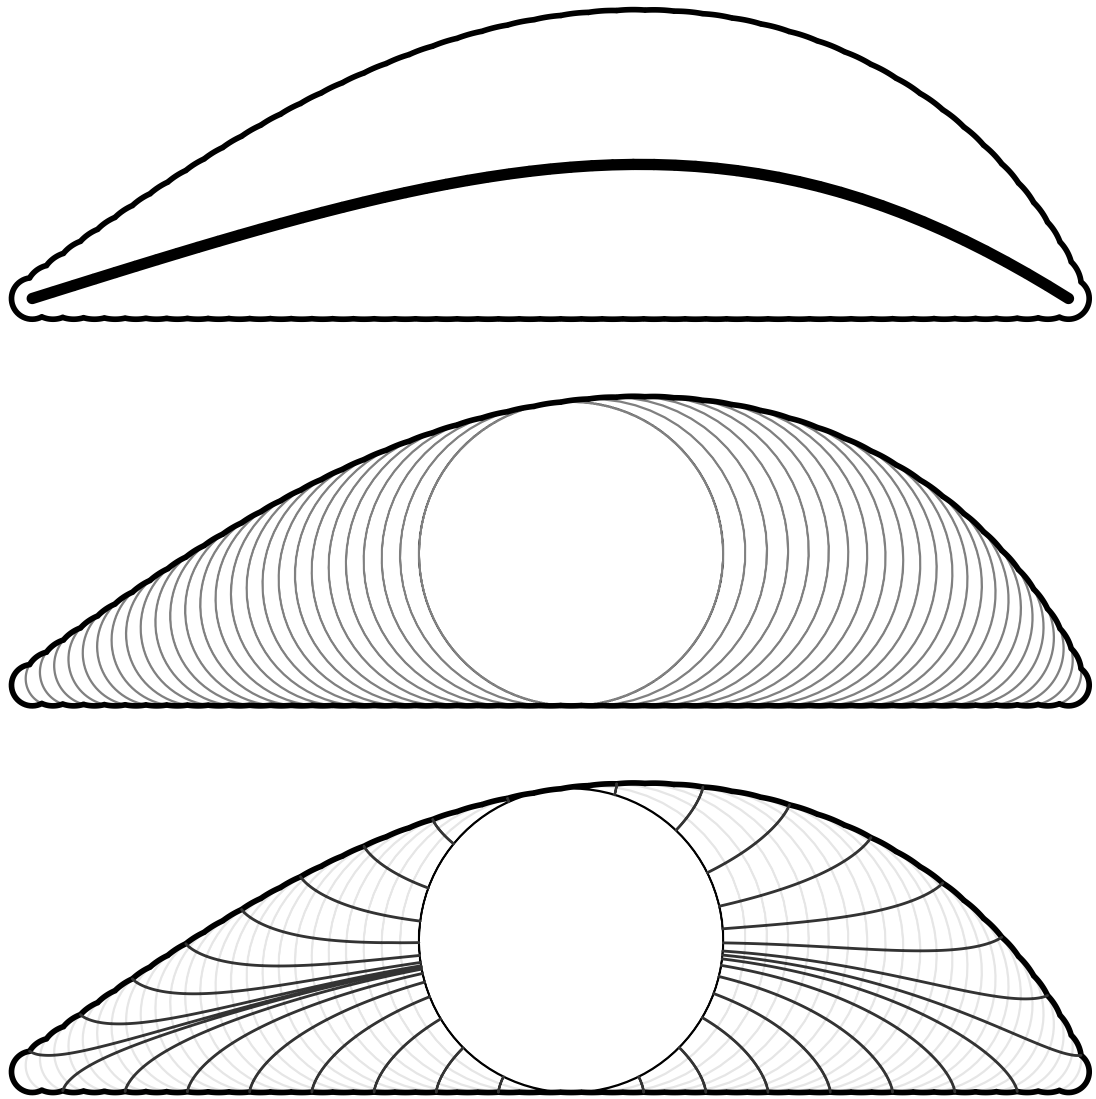

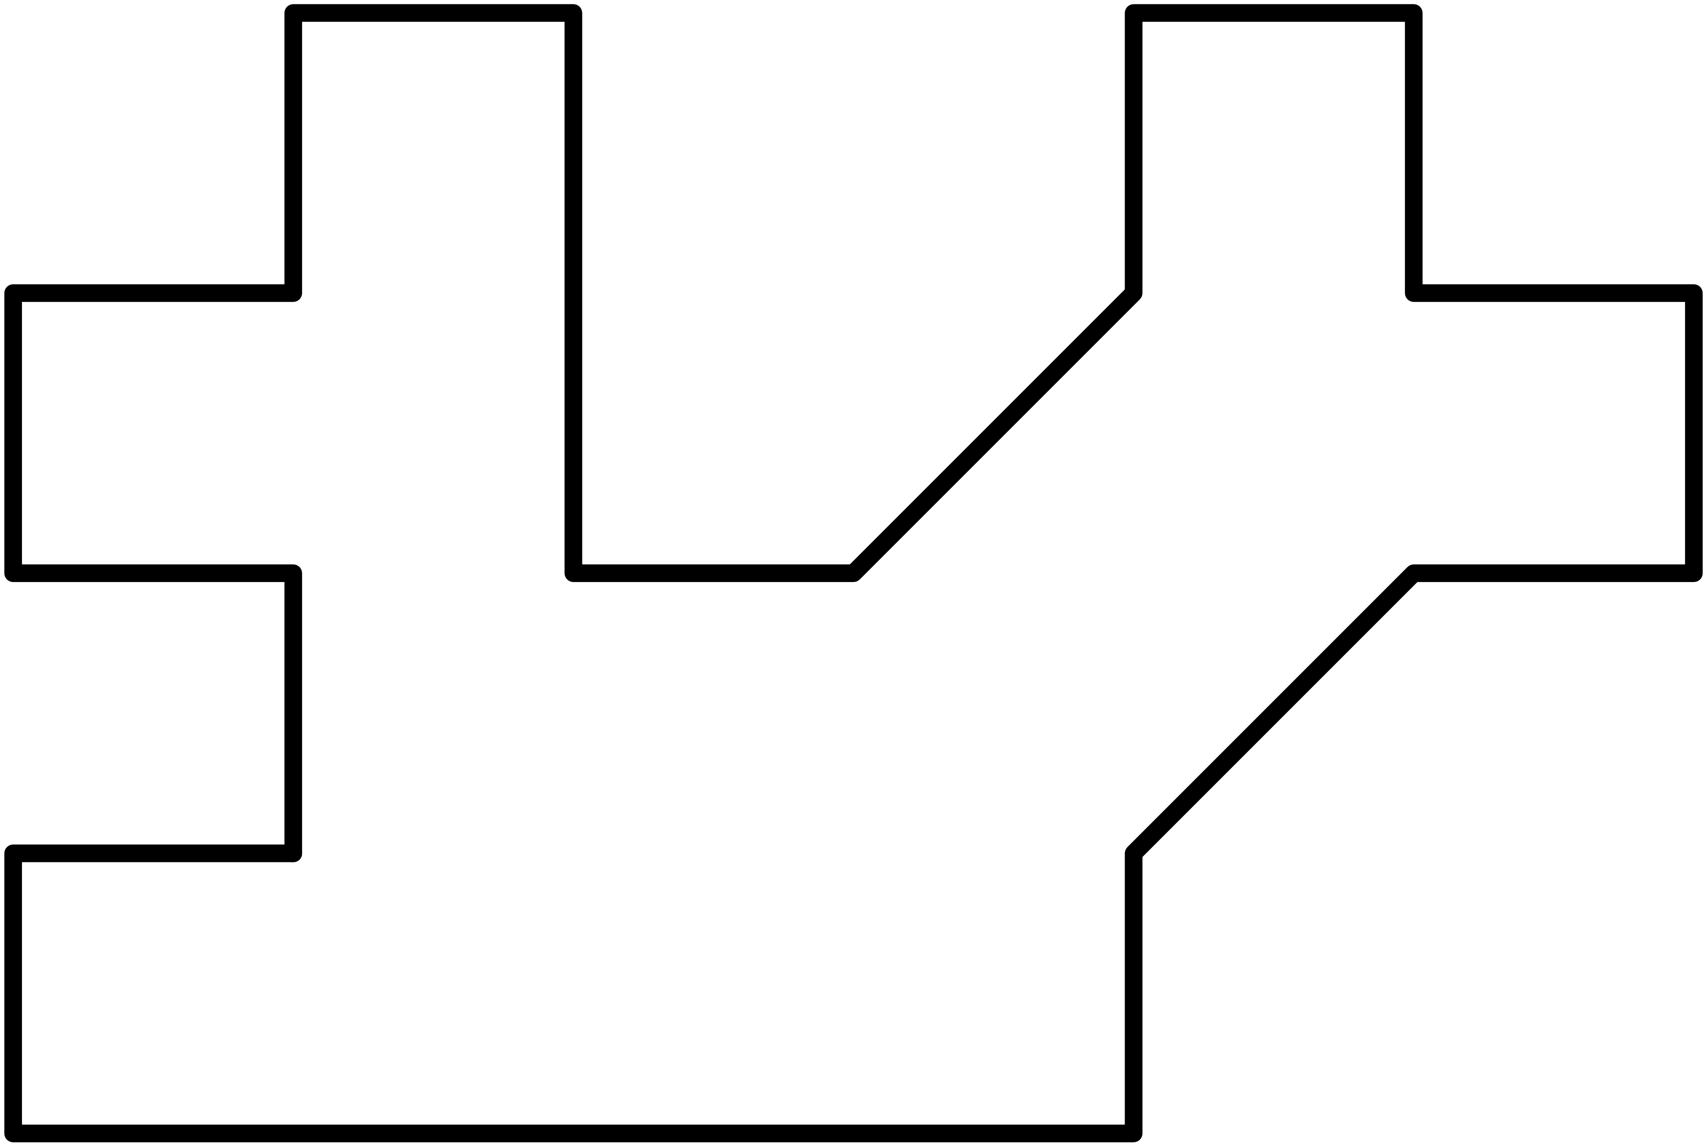

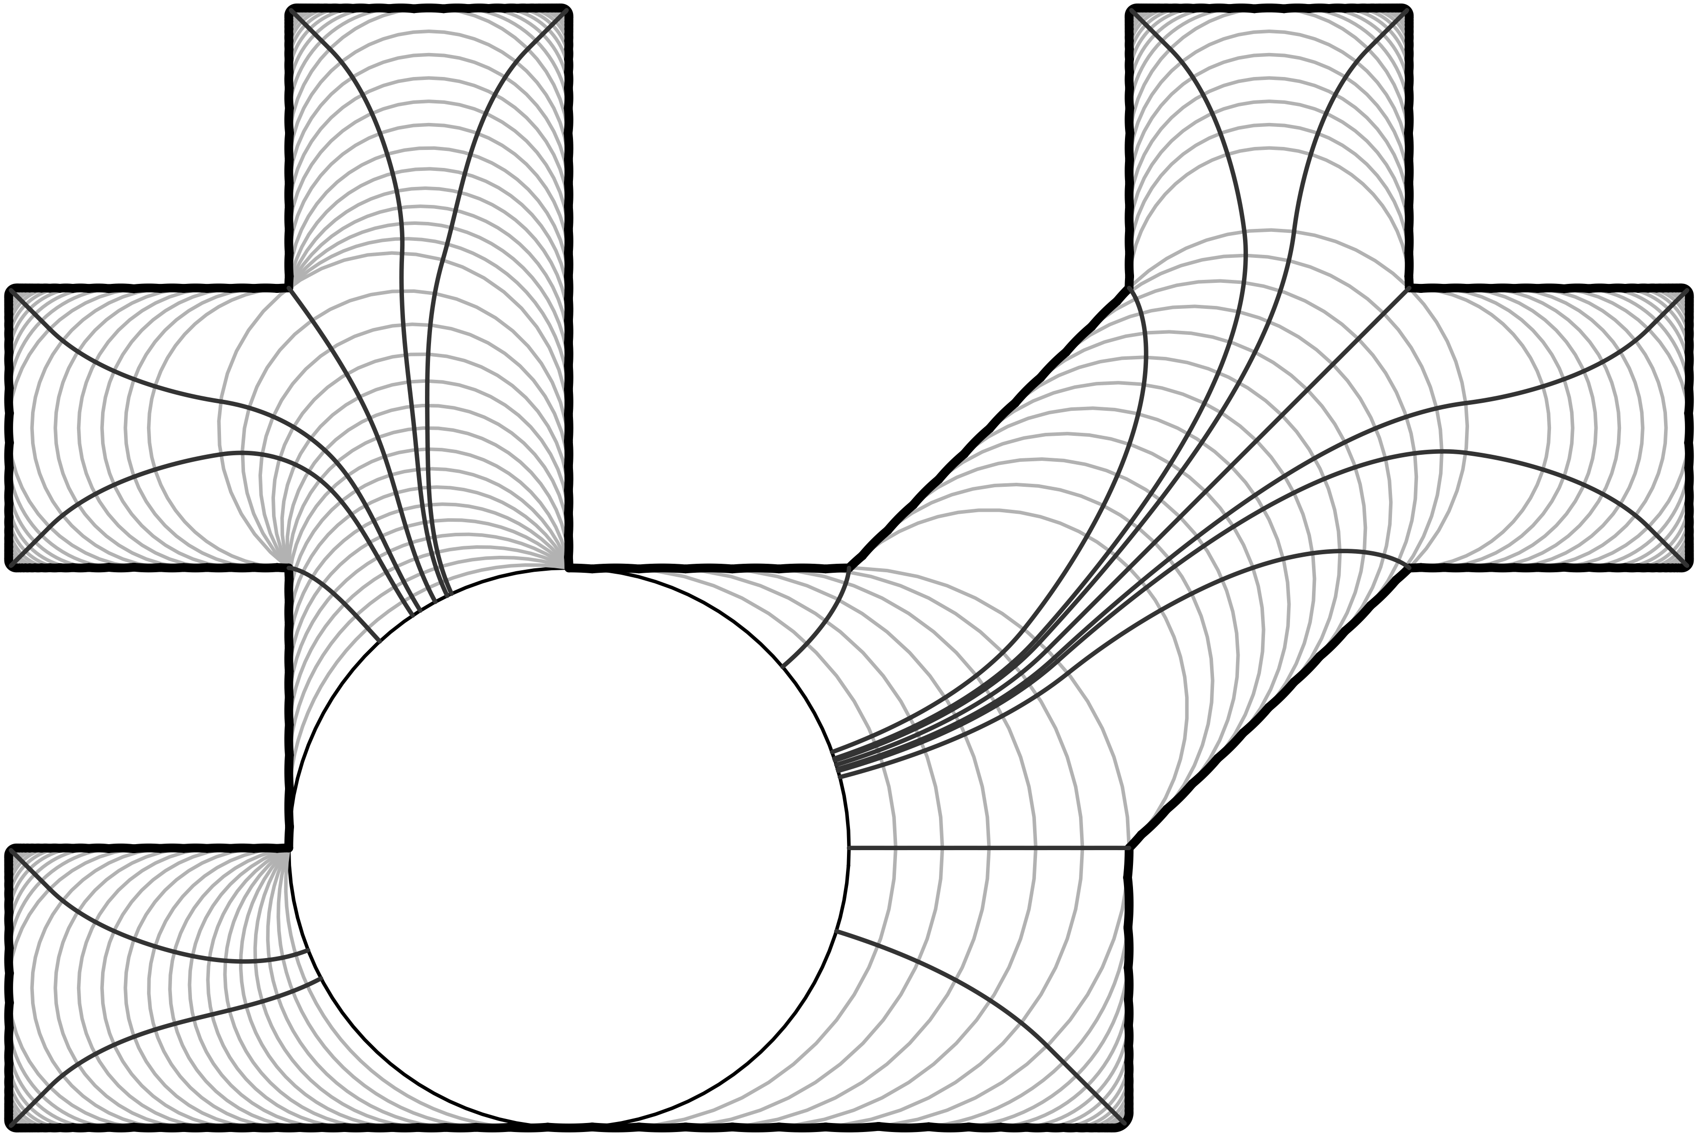

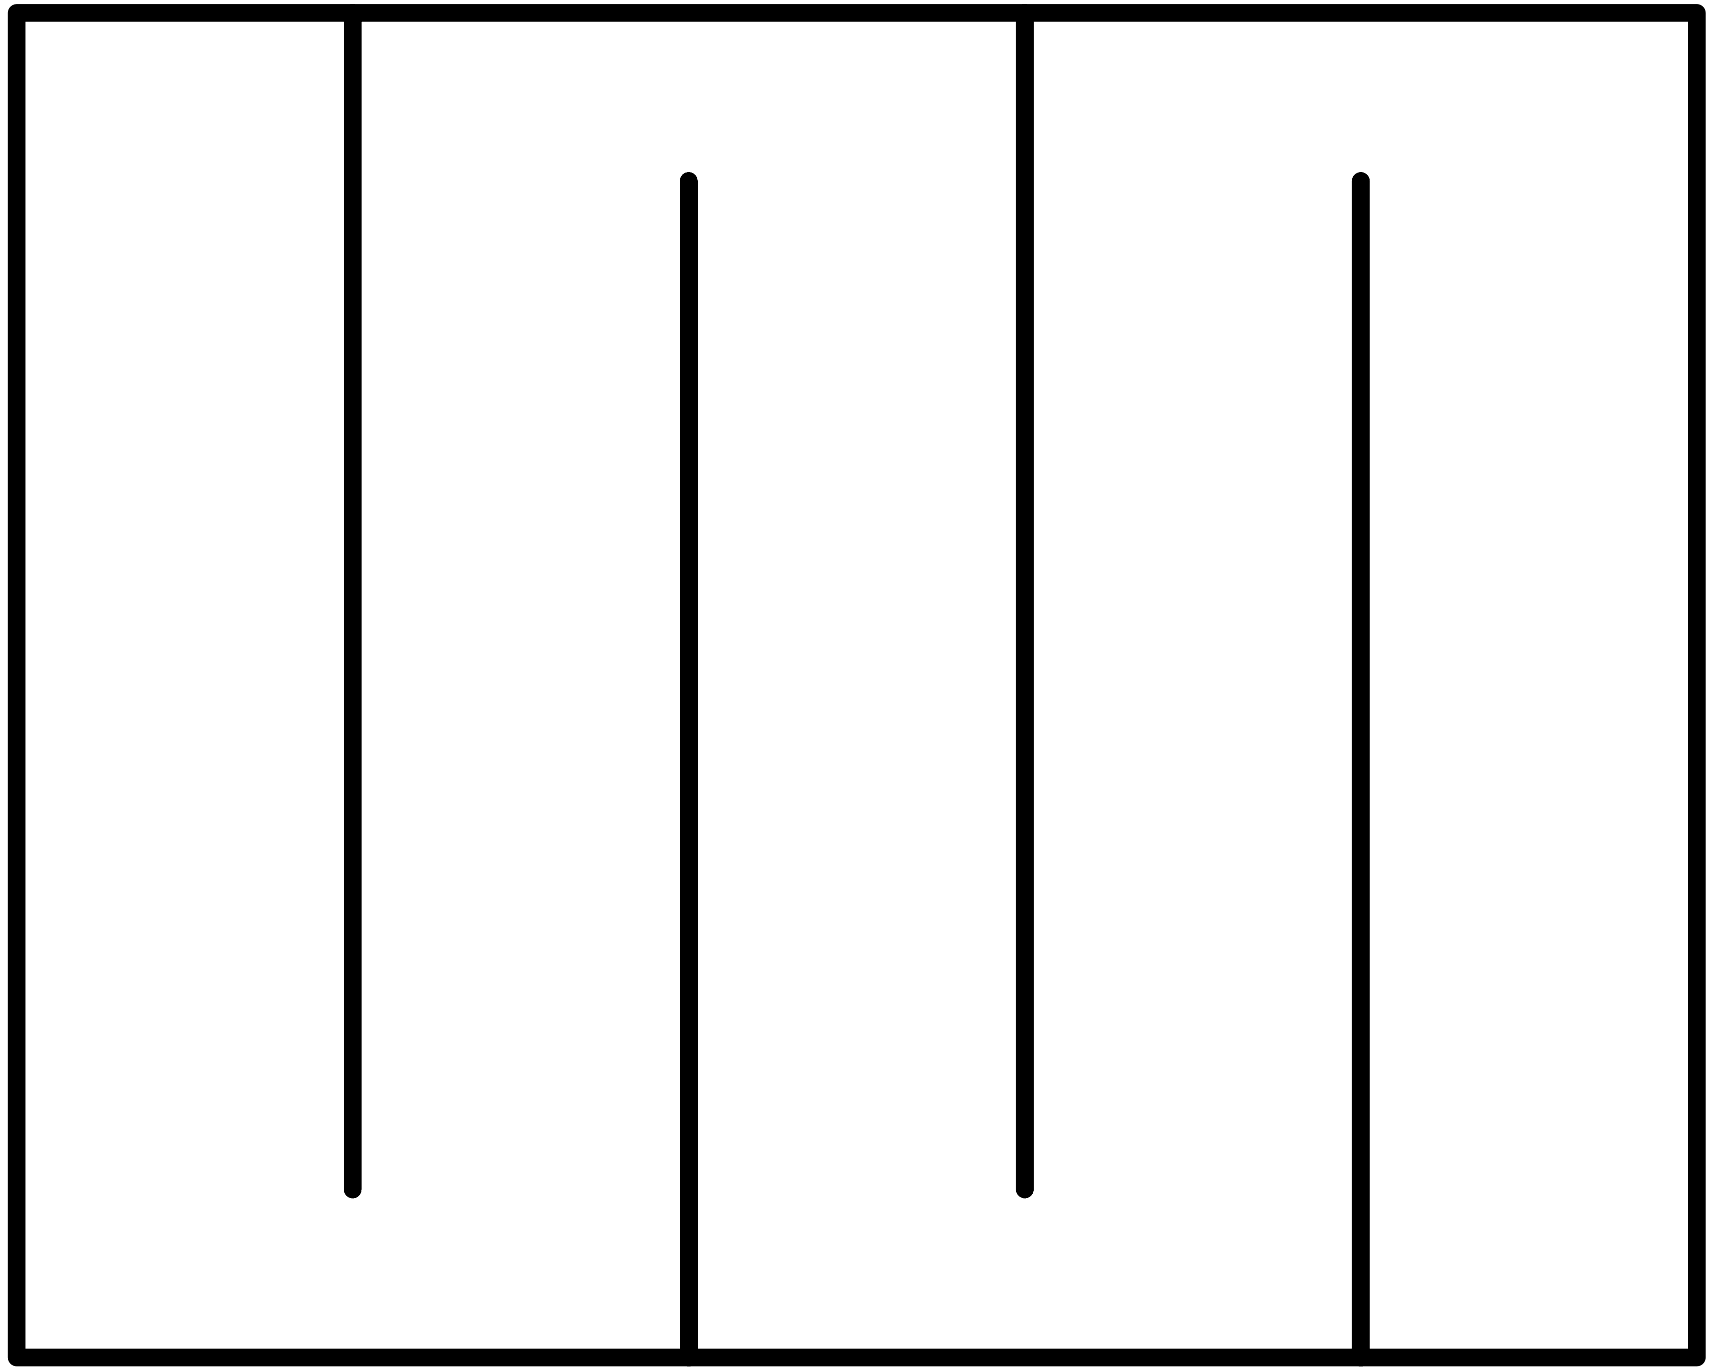

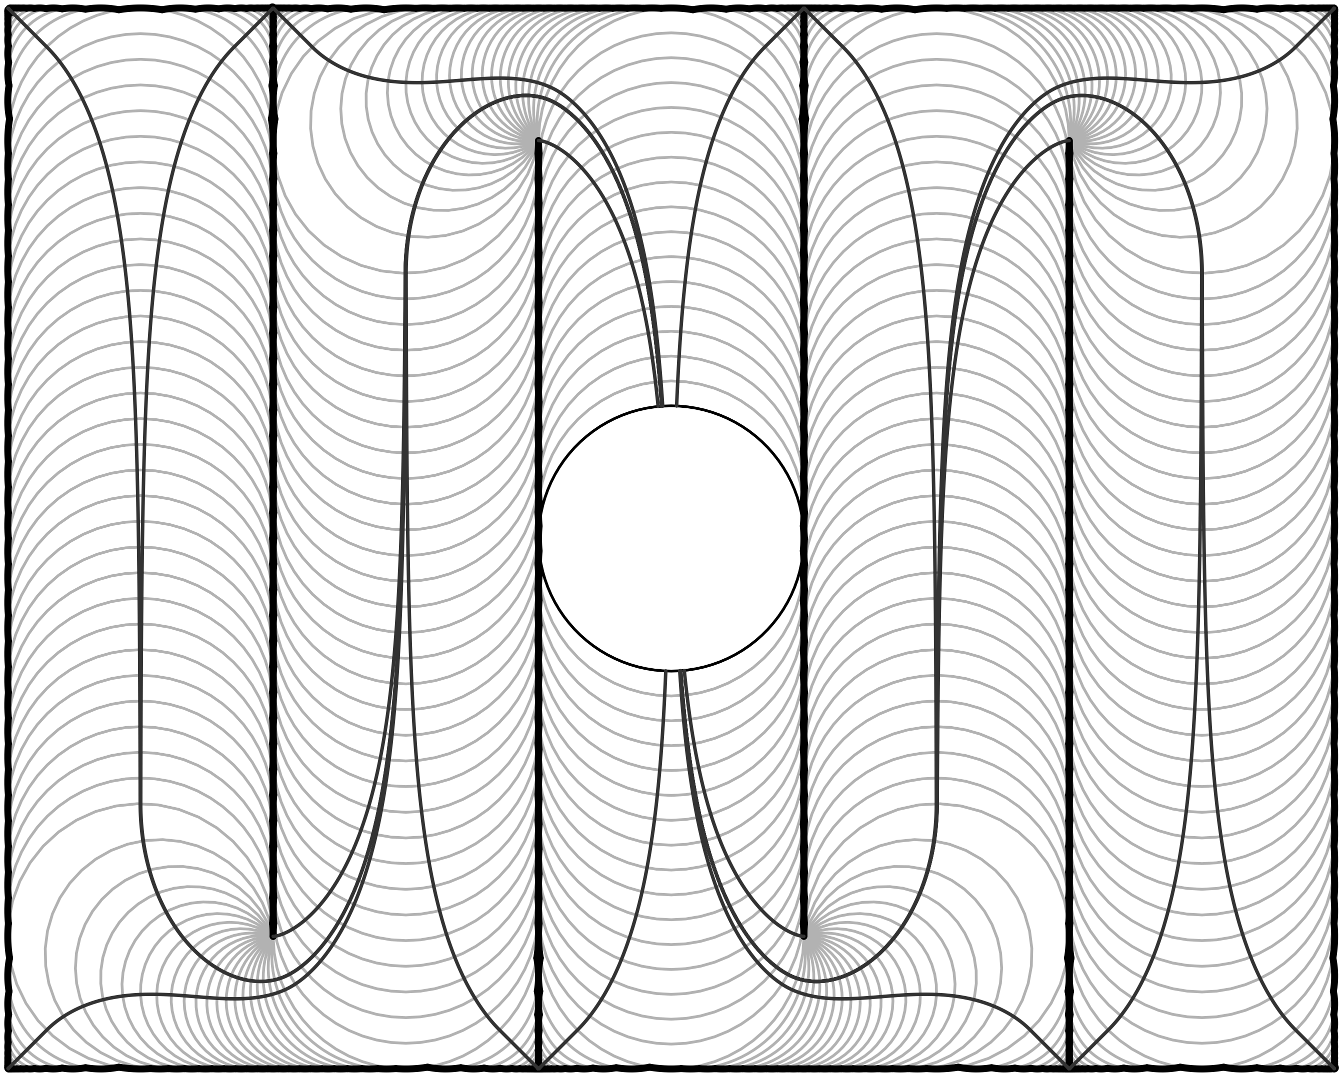

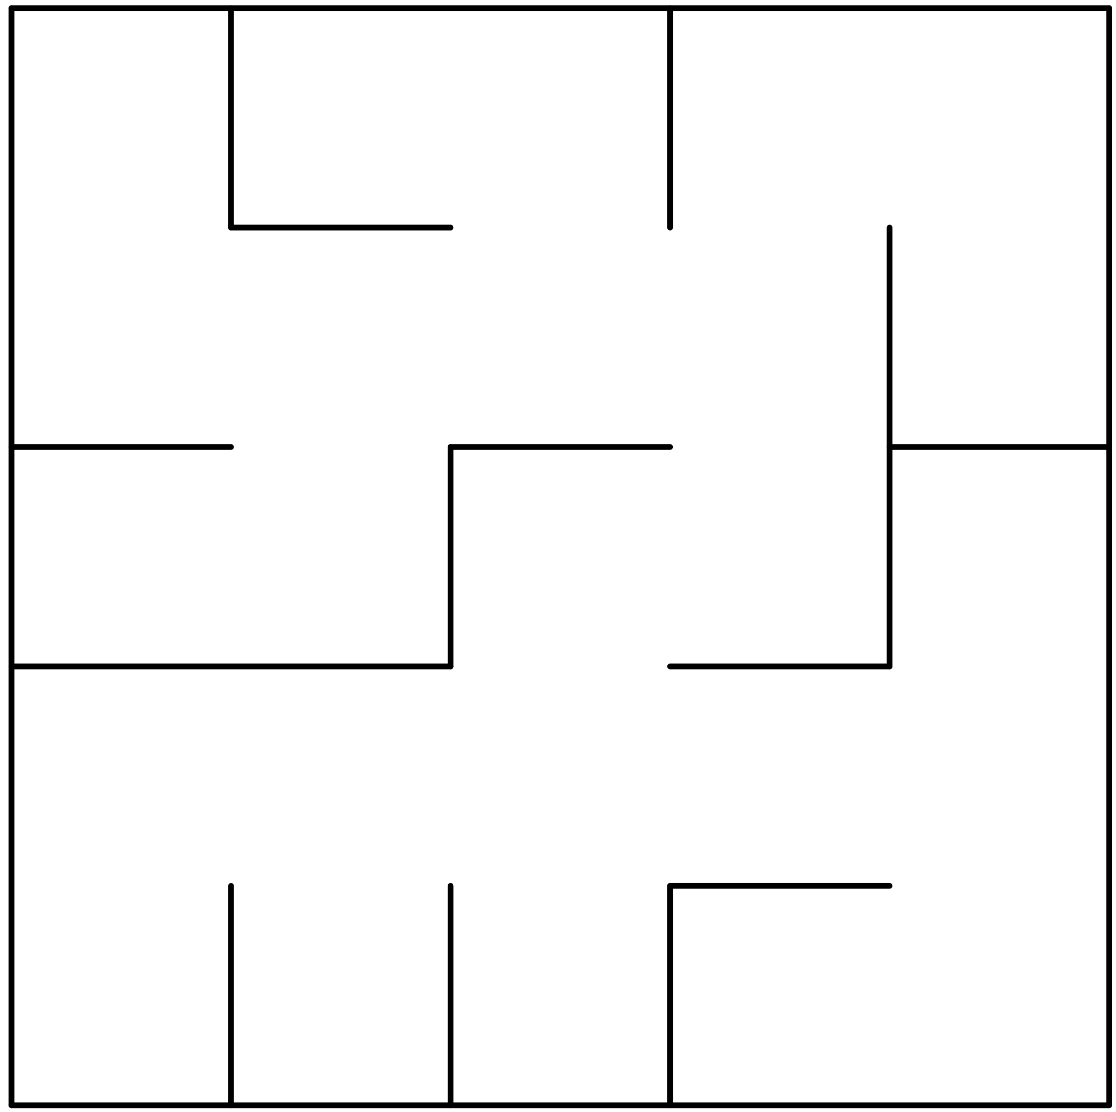

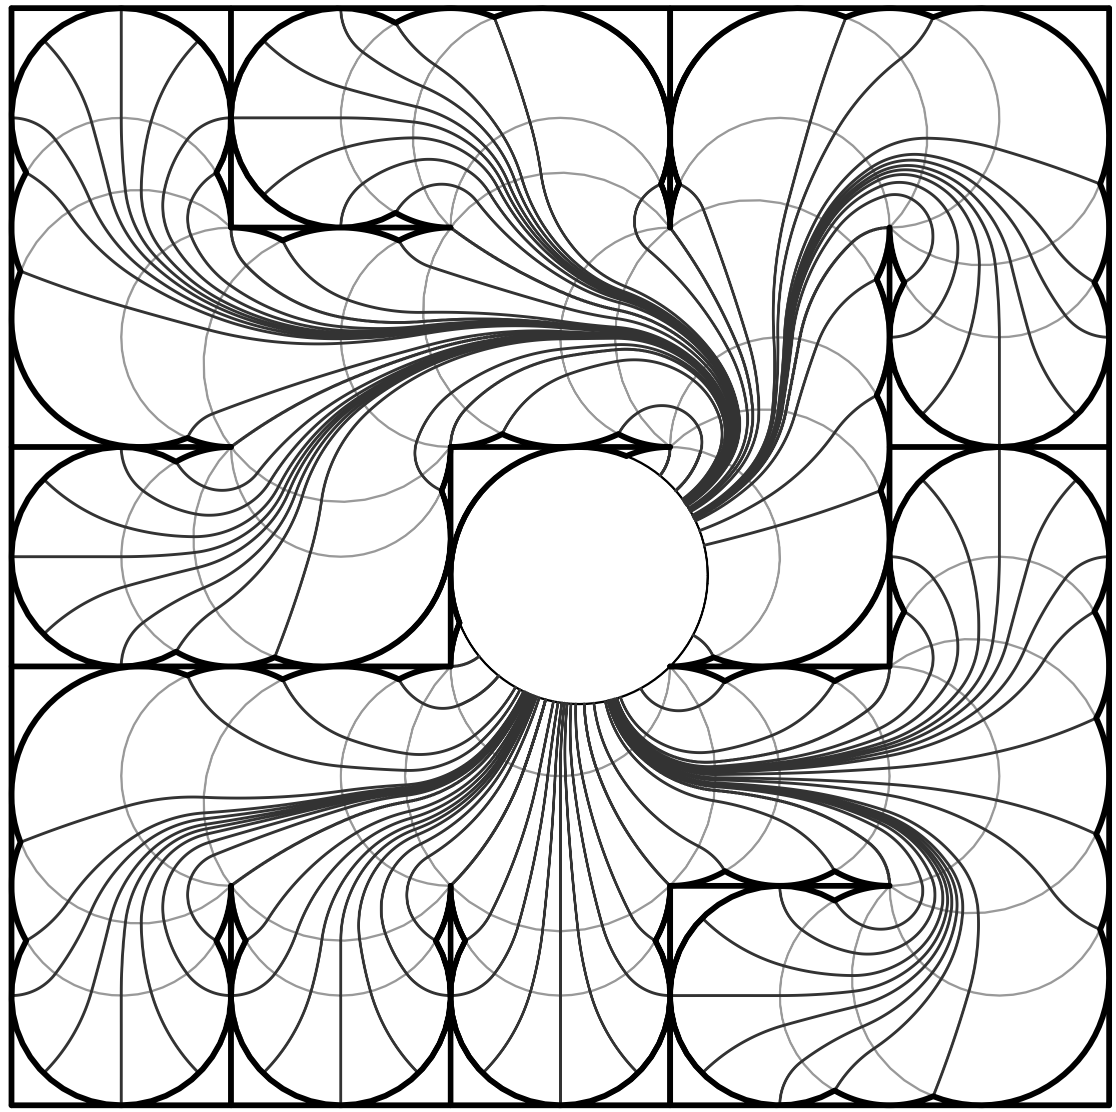

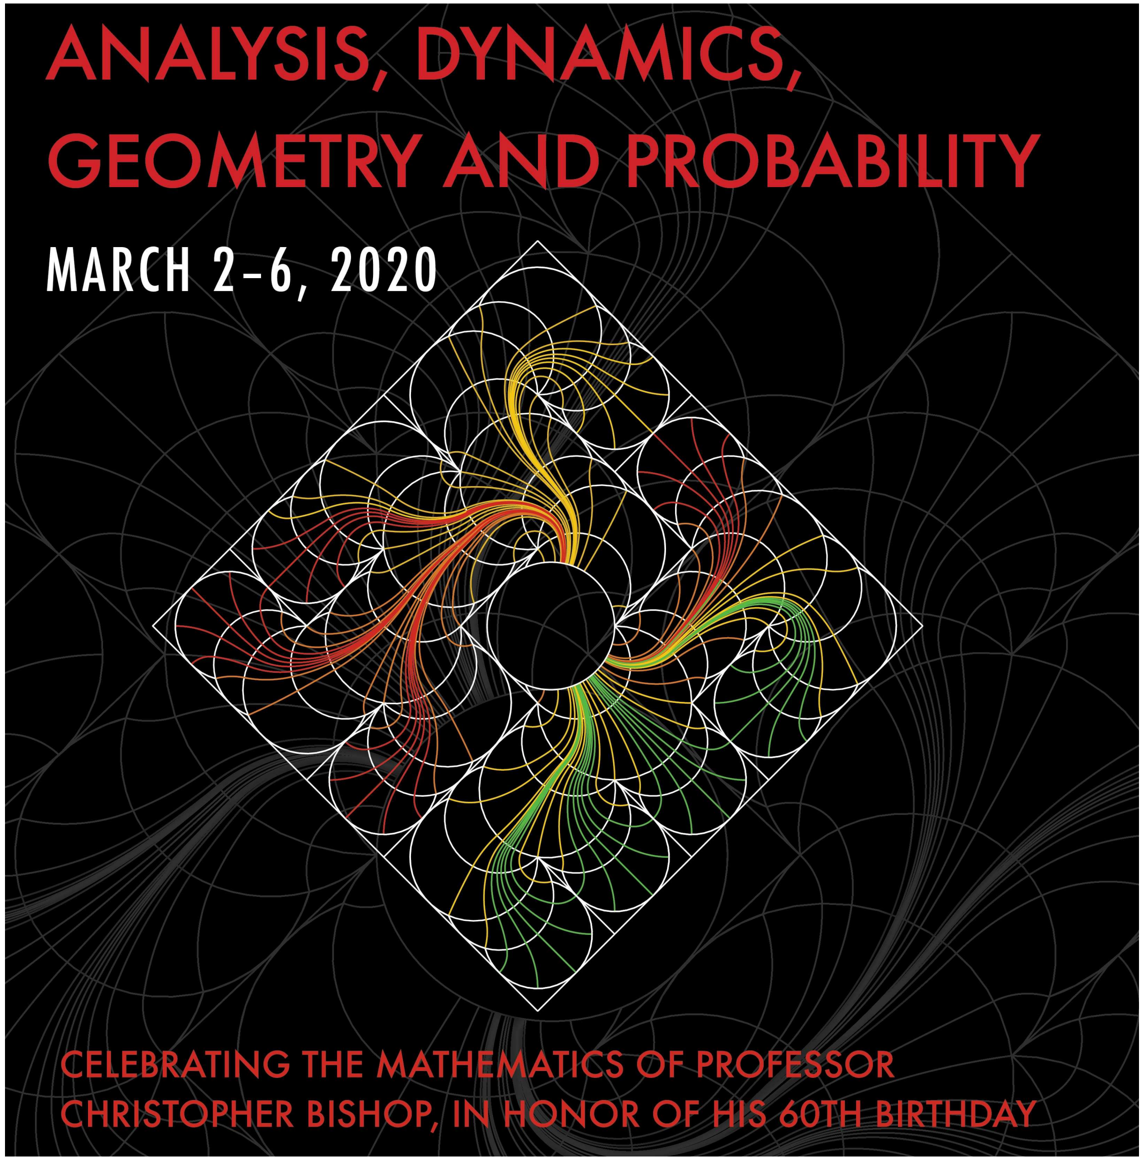
ANALYSIS, DYNAMICS,
GEOMETRY AND PROBABILITY
MARCH 2–6, 2020
CELEBRATING THE MATHEMATICS OF PROFESSOR
CHRISTOPHER BISHOP, IN HONOR OF HIS 60TH BIRTHDAY

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

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

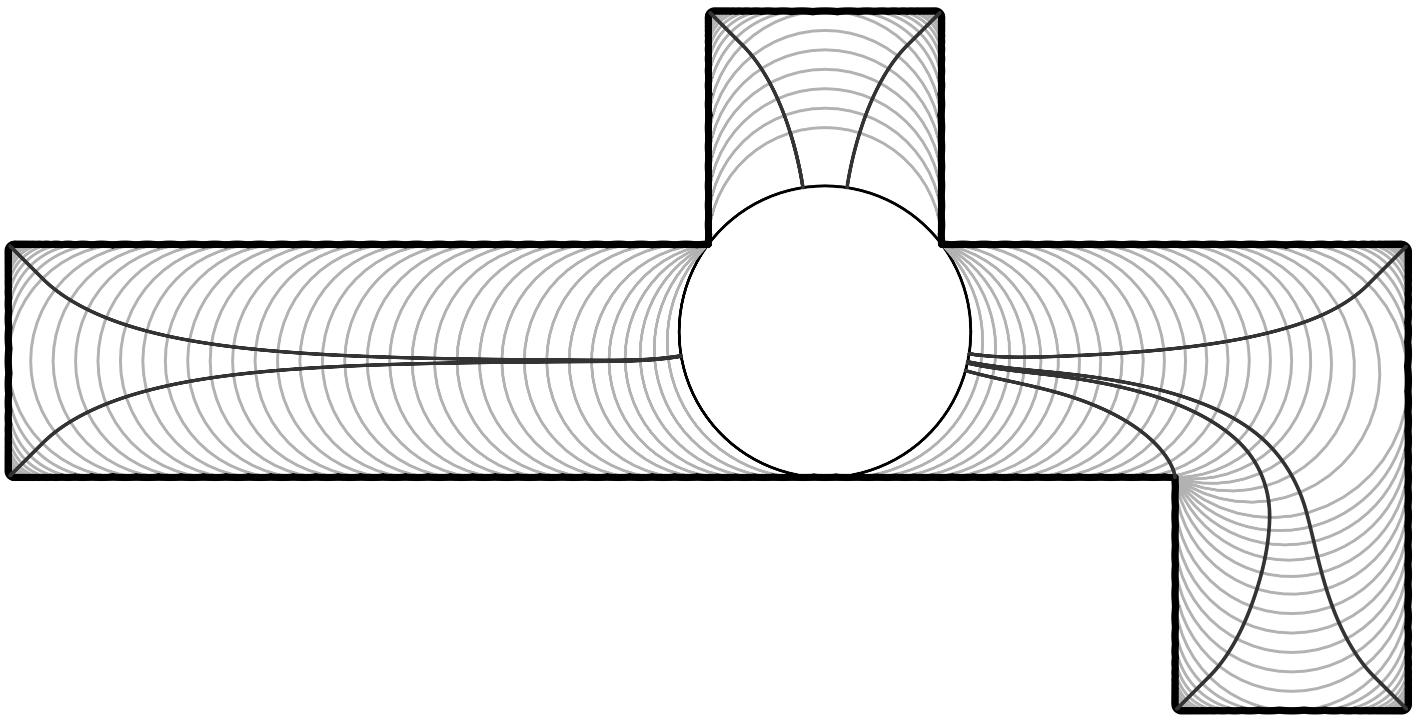

How close is medial axis map to conformal map?

How close is medial axis map to conformal map?

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

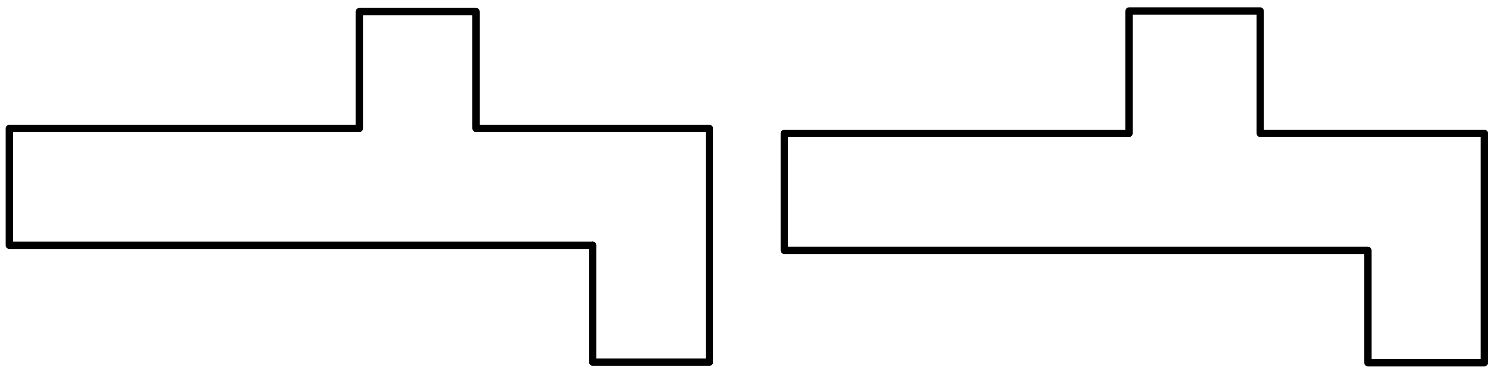

Target Polygon   MA Parameters

How close is medial axis map to conformal map?

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

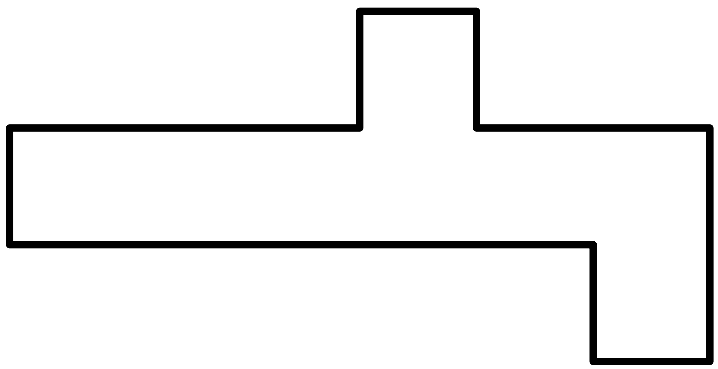

Target Polygon

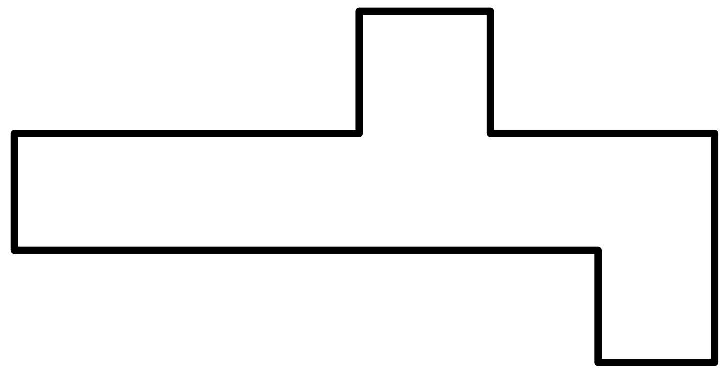

MA Parameters

Looks pretty close. How do we measure error?

How close is medial axis map to conformal map?

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

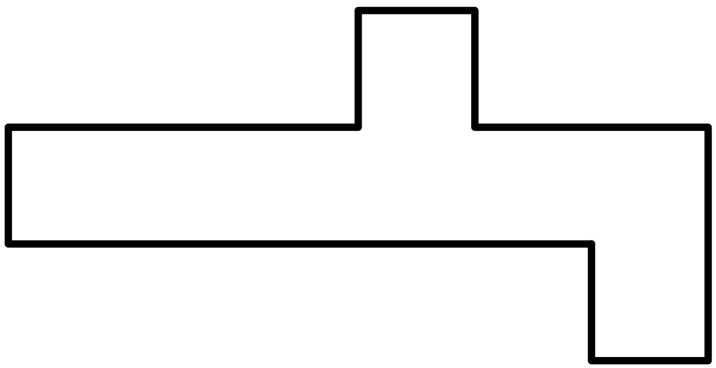

Target Polygon

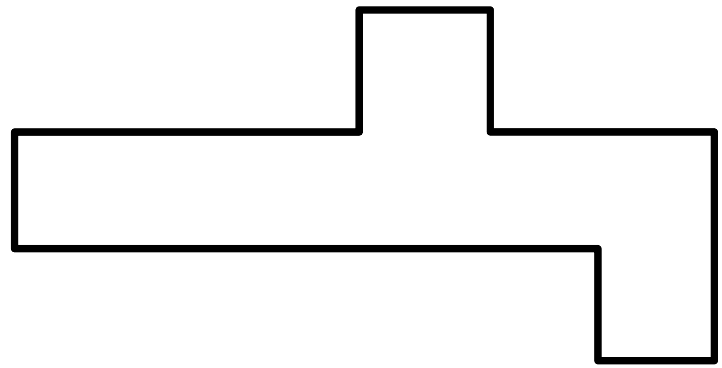

MA Parameters

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.

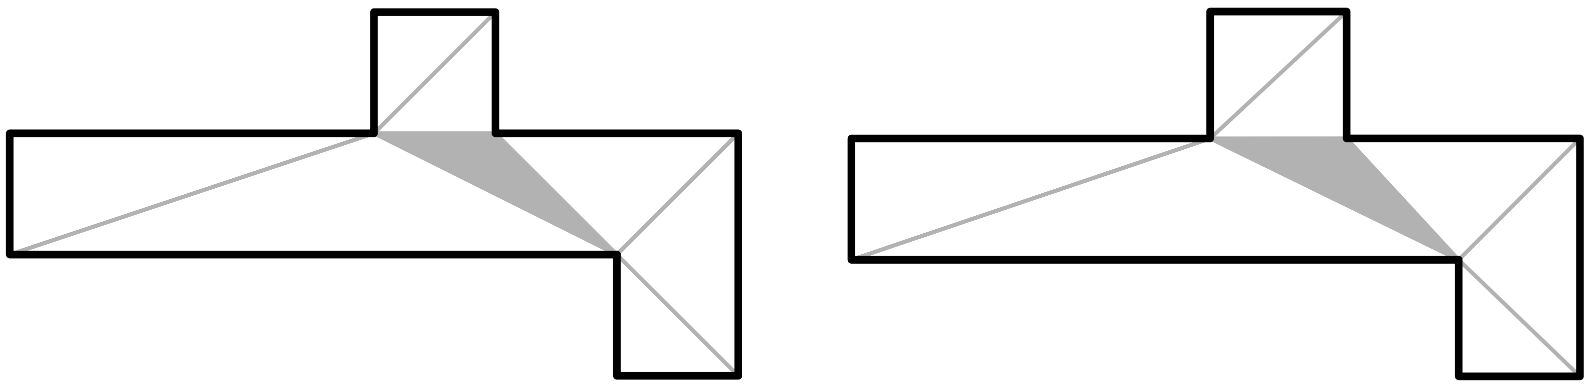

The most distorted triangle is shaded. Here $K = 1.24$.

Upper bound for QC-distance:

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

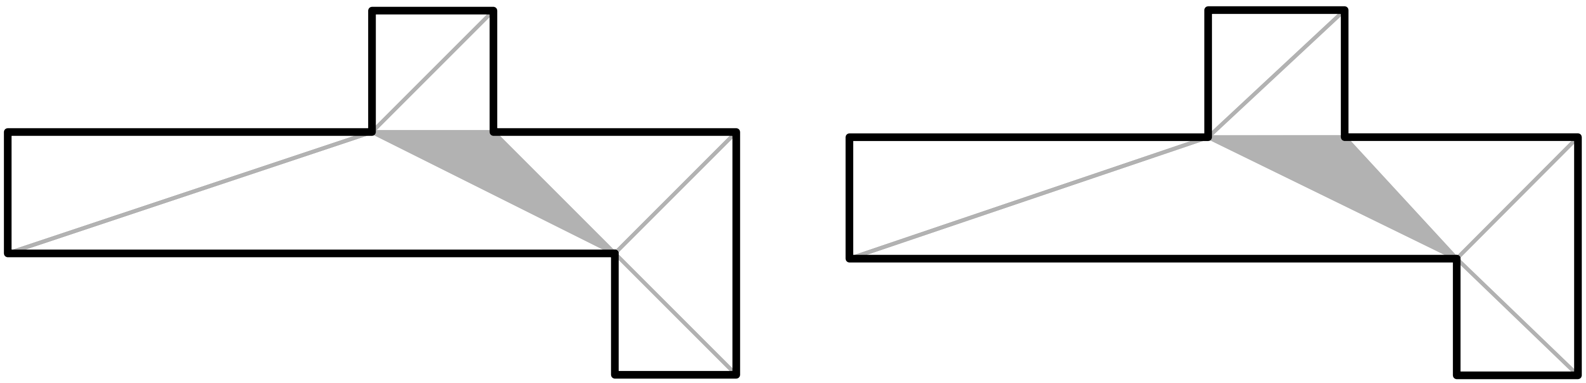

The most distorted triangle is shaded. Here $K = 1.24$.

Bounds QC-error of guessed parameters without needing to know the true SC-parameters.

Theorem: Medial axis map gives QC-error < 8.

Theorem: Medial axis map gives QC-error < 8.

Why is this theorem true?

Theorem: Medial axis map gives QC-error < 8.

Why is this theorem true?

Short answer: hyperbolic 3-manifold theory

Theorem: Medial axis map gives QC-error < 8.

Why is this theorem true?

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

Theorem: Medial axis map gives QC-error < 8.

Why is this theorem true?

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

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

The **dome** of Ω is upper envelope of all hemispheres with base in Ω.

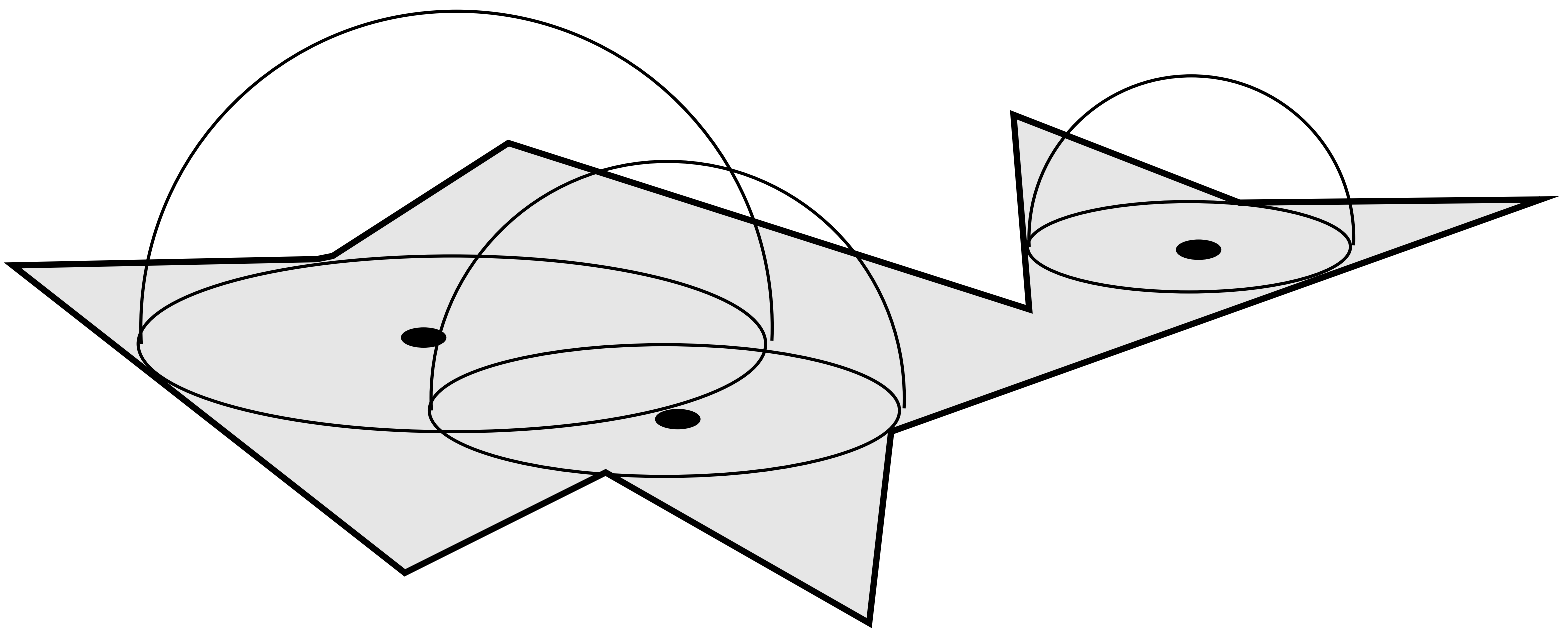

Suffices to take medial axis disks (= maximal disks).

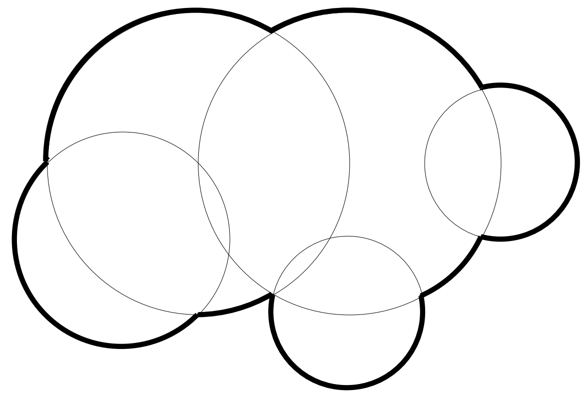

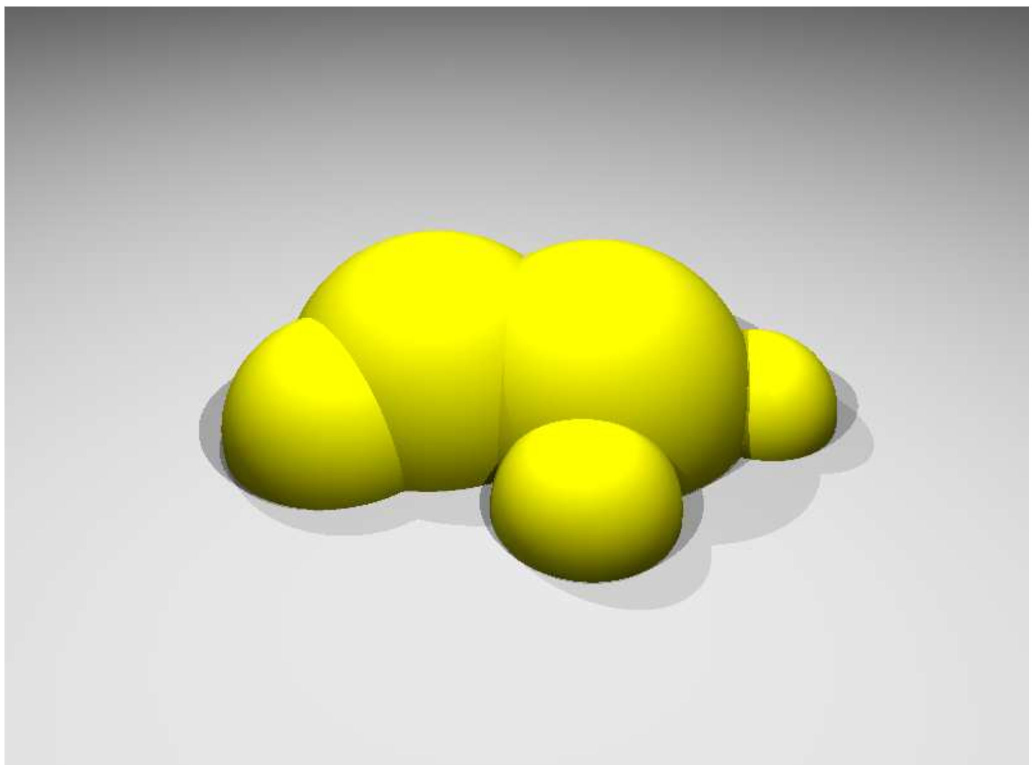

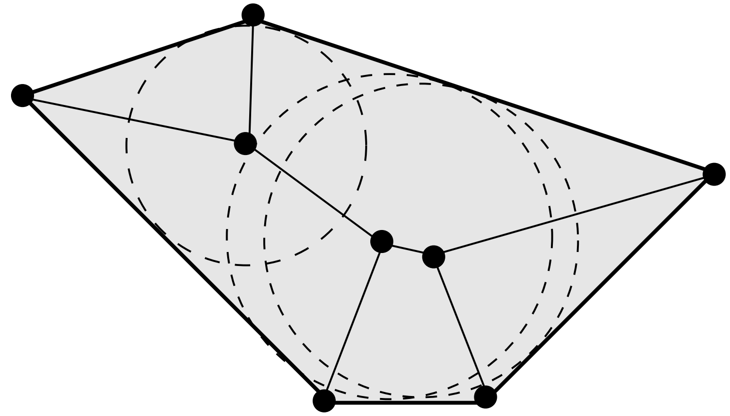

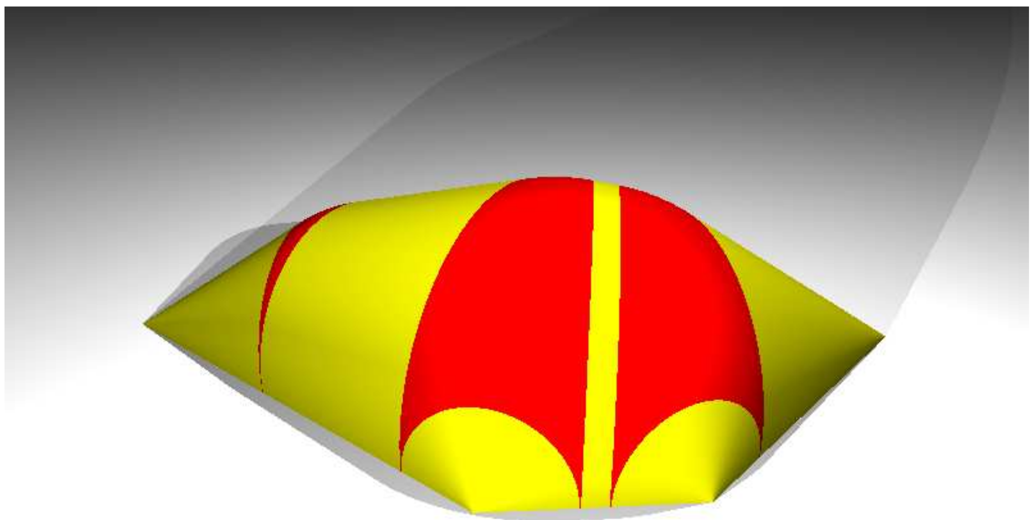

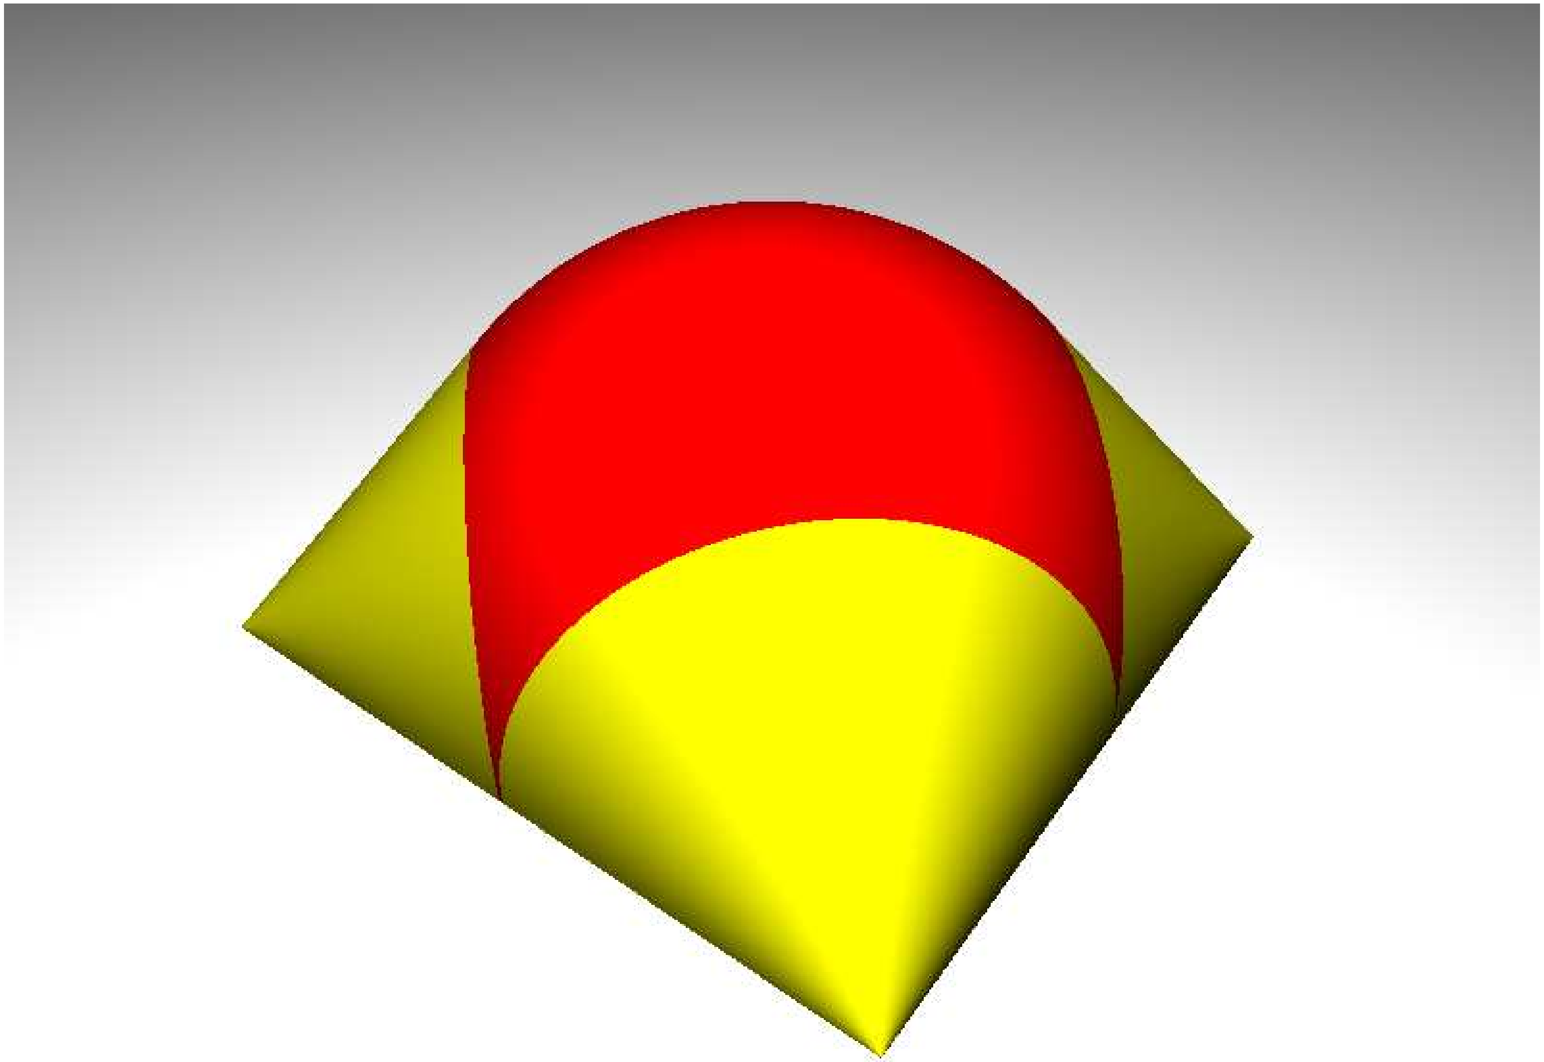

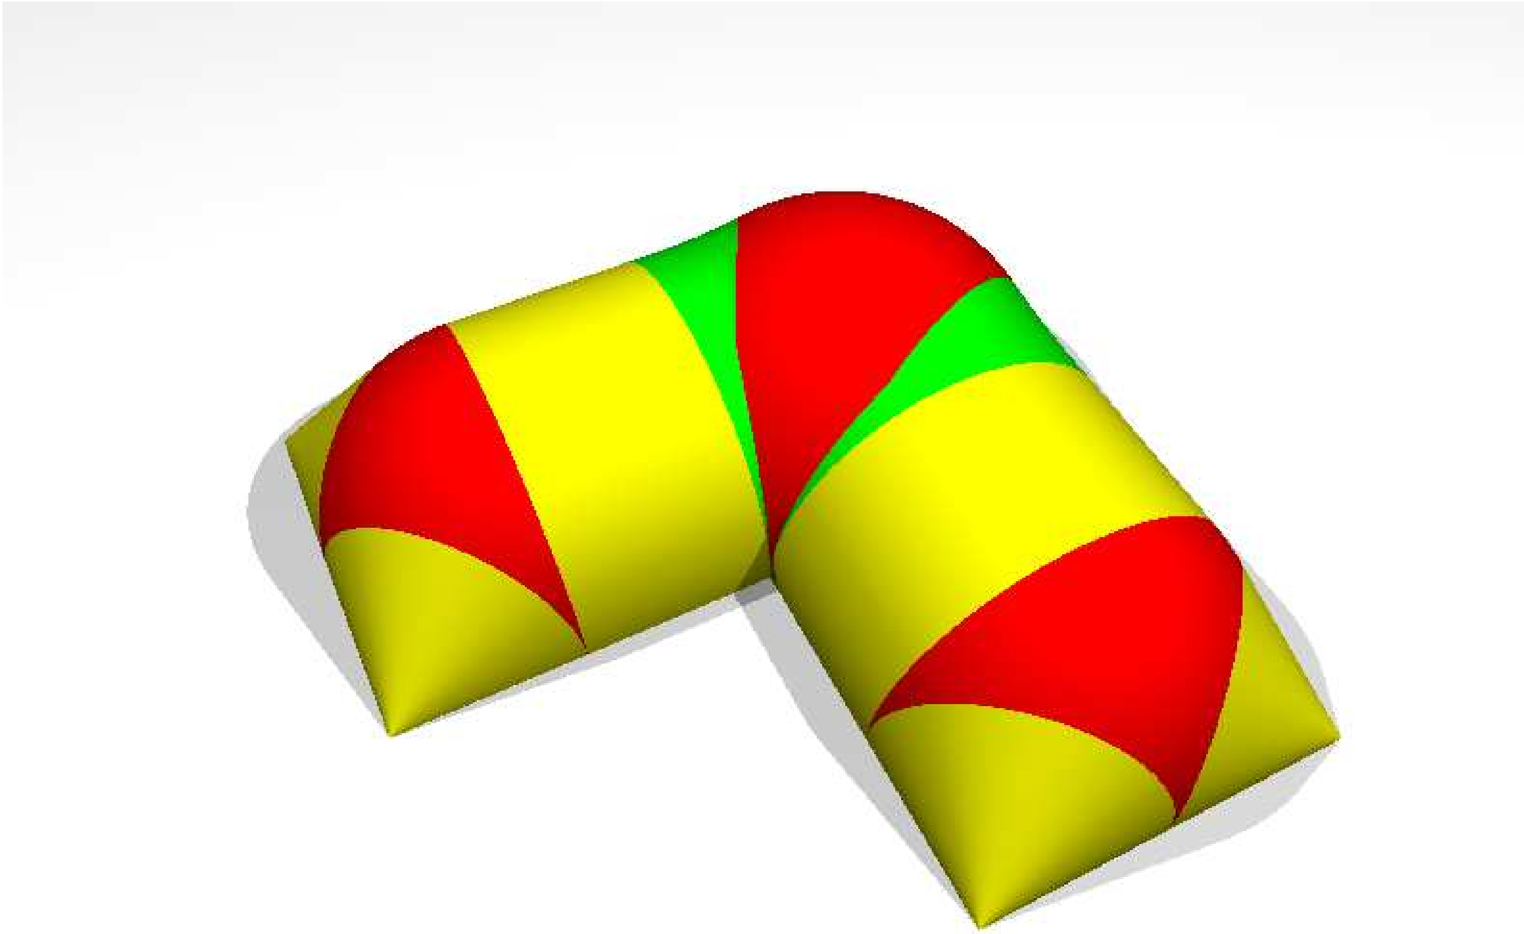

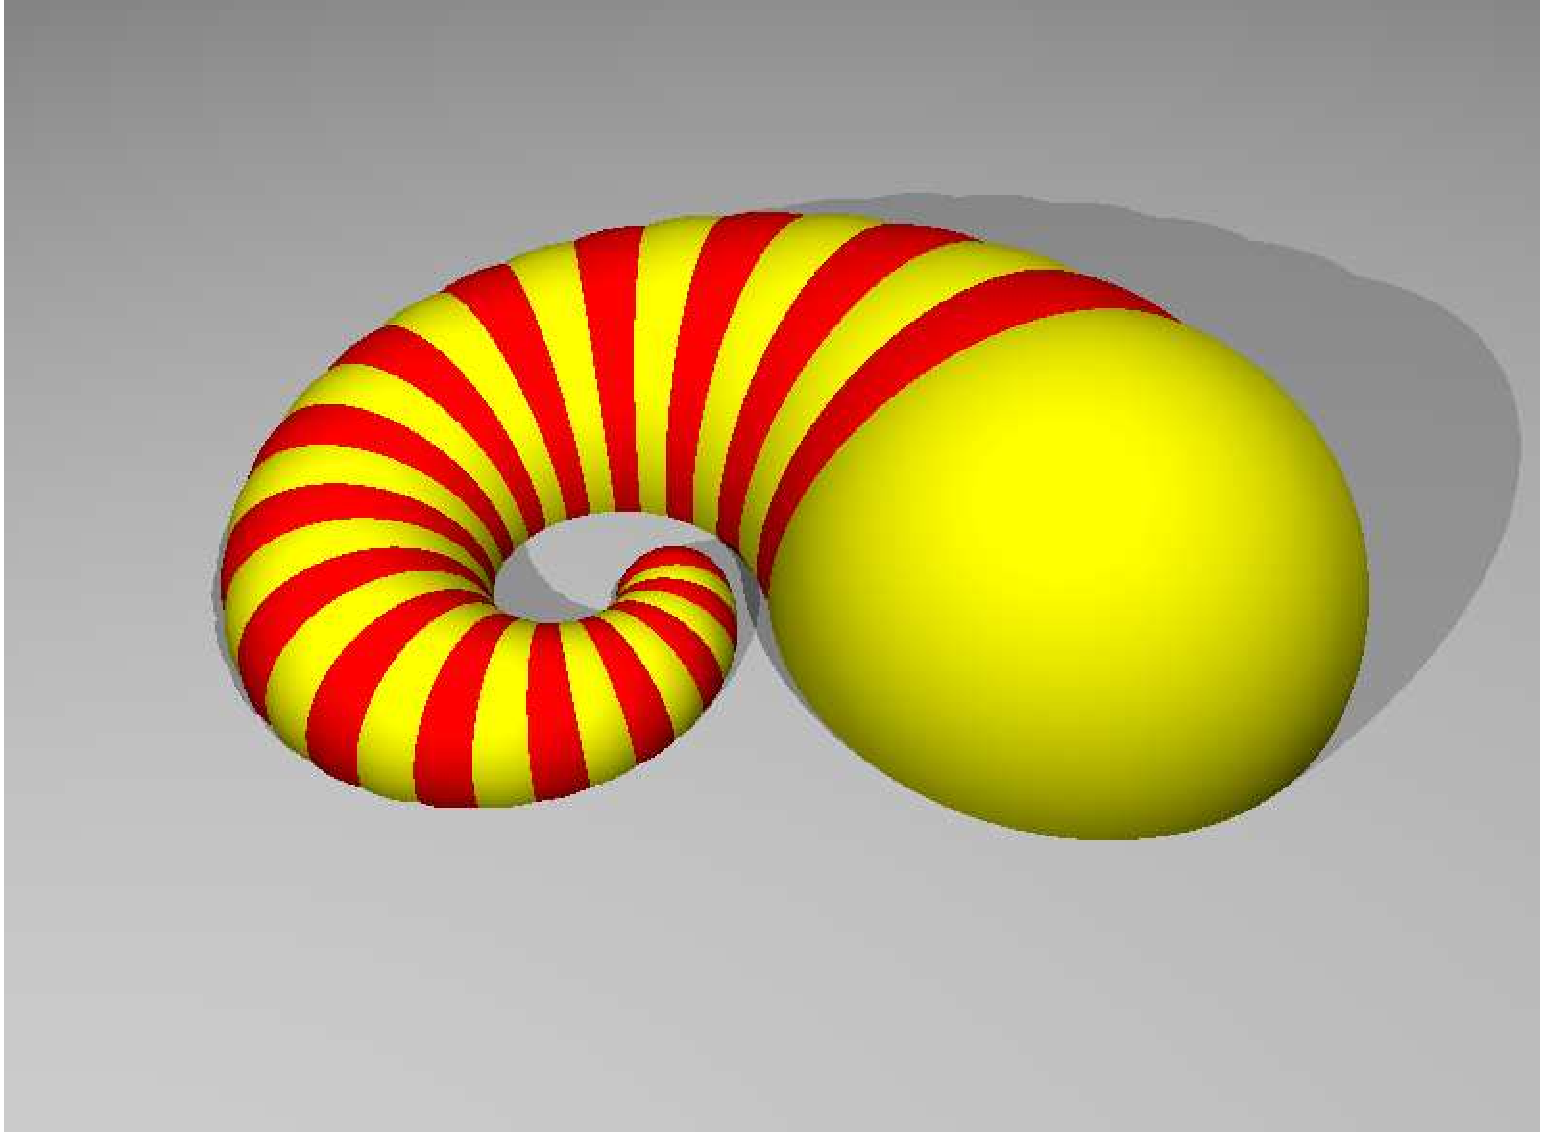

Every dome has conformal map to disk by “flattening”.

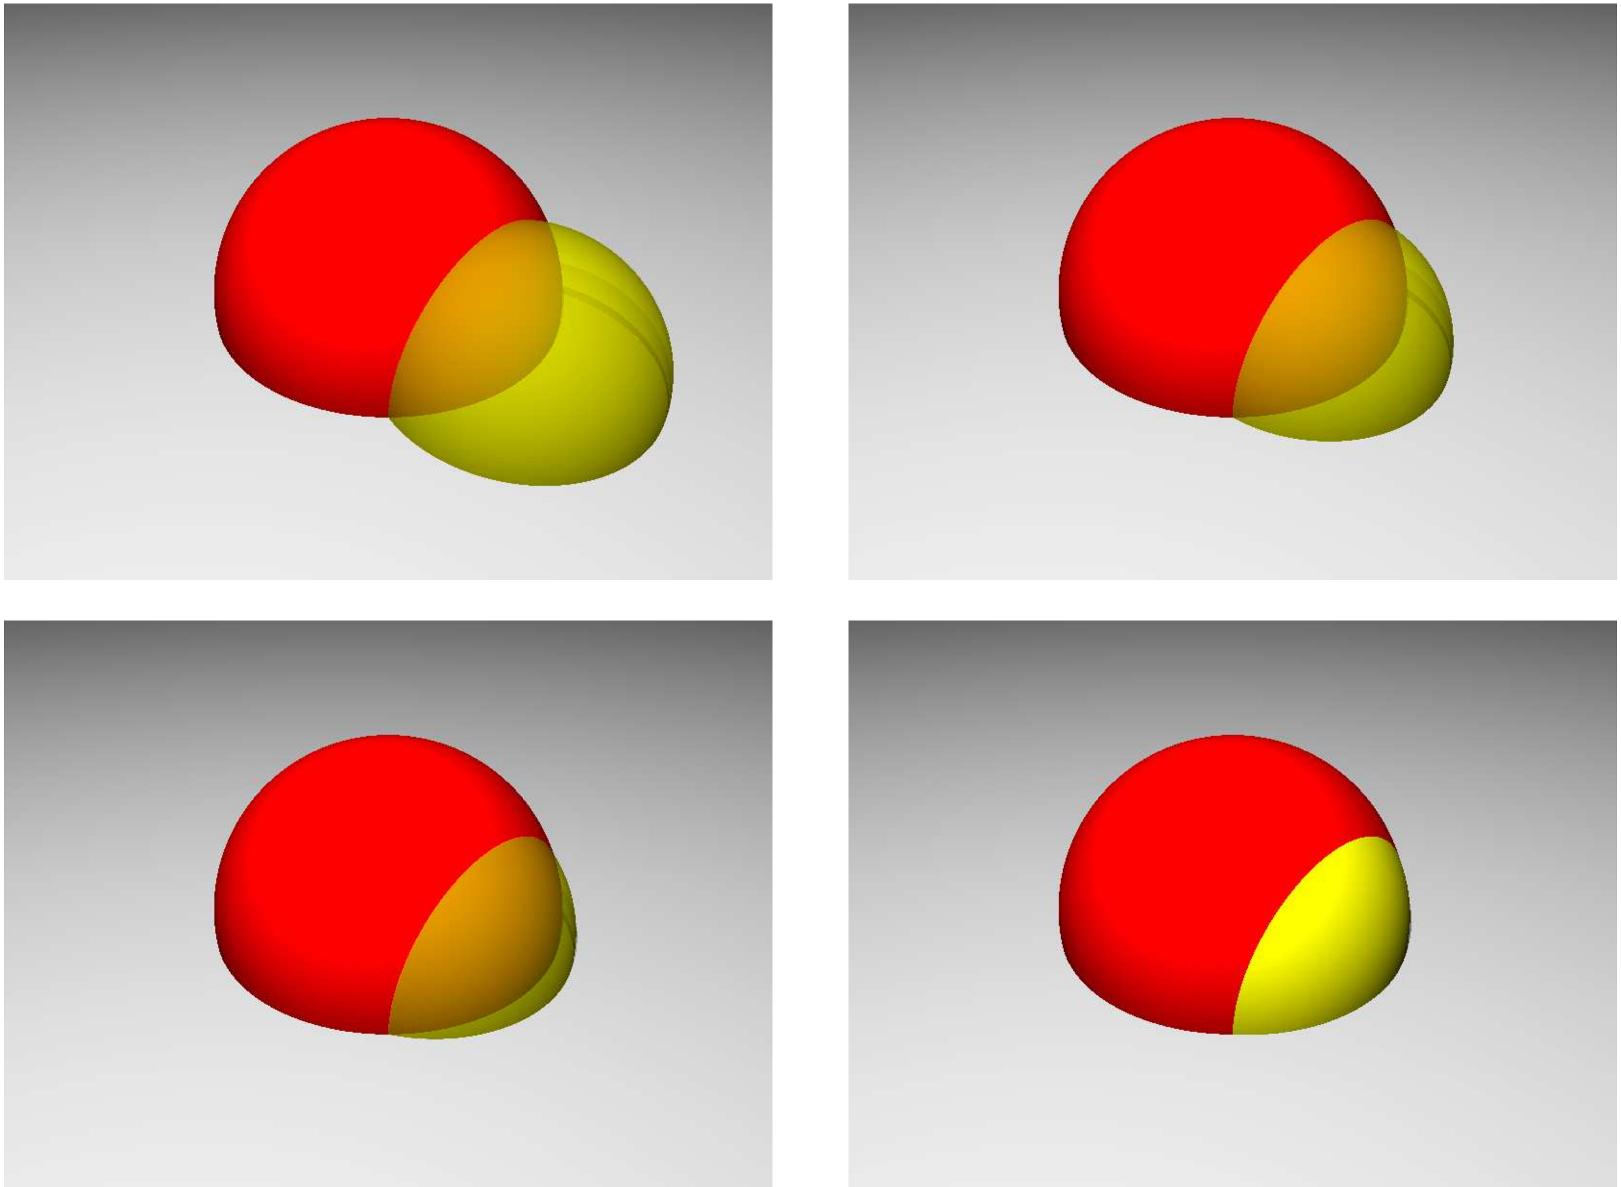

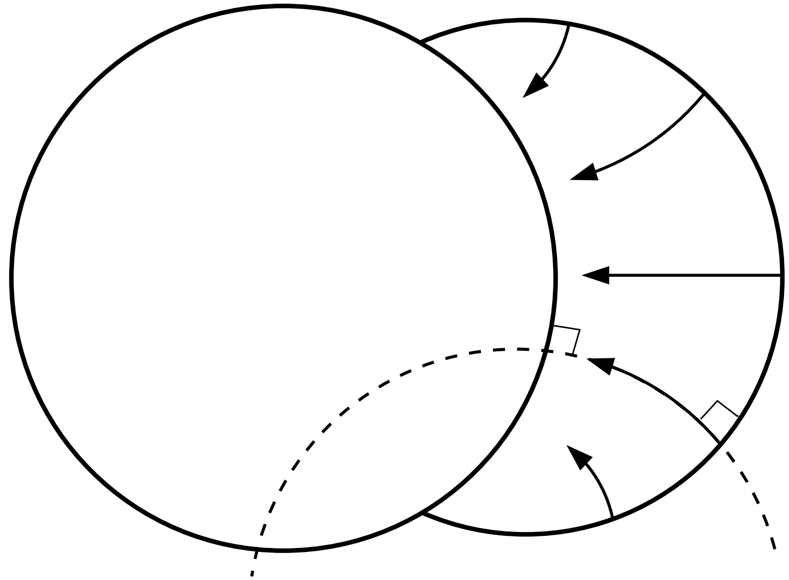

Medial axis map = boundary of flattening map (iota)

= boundary of conformal map of dome to hemisphere

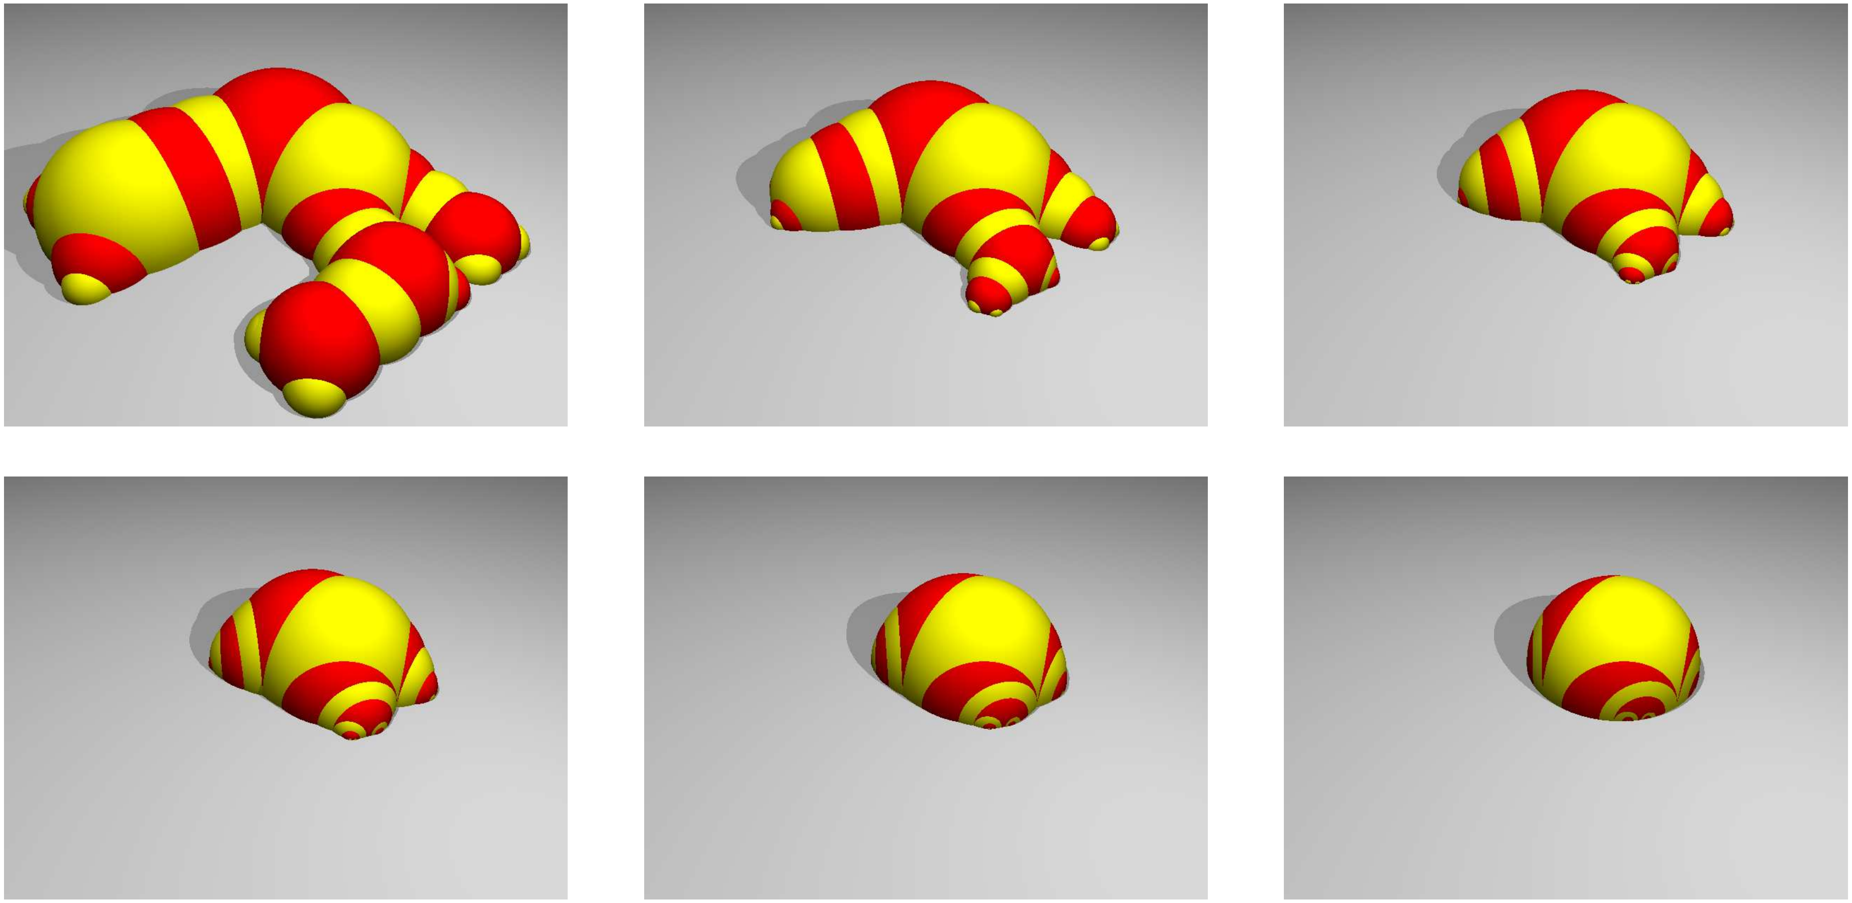

Map dome to hemisphere by makings sides flush.

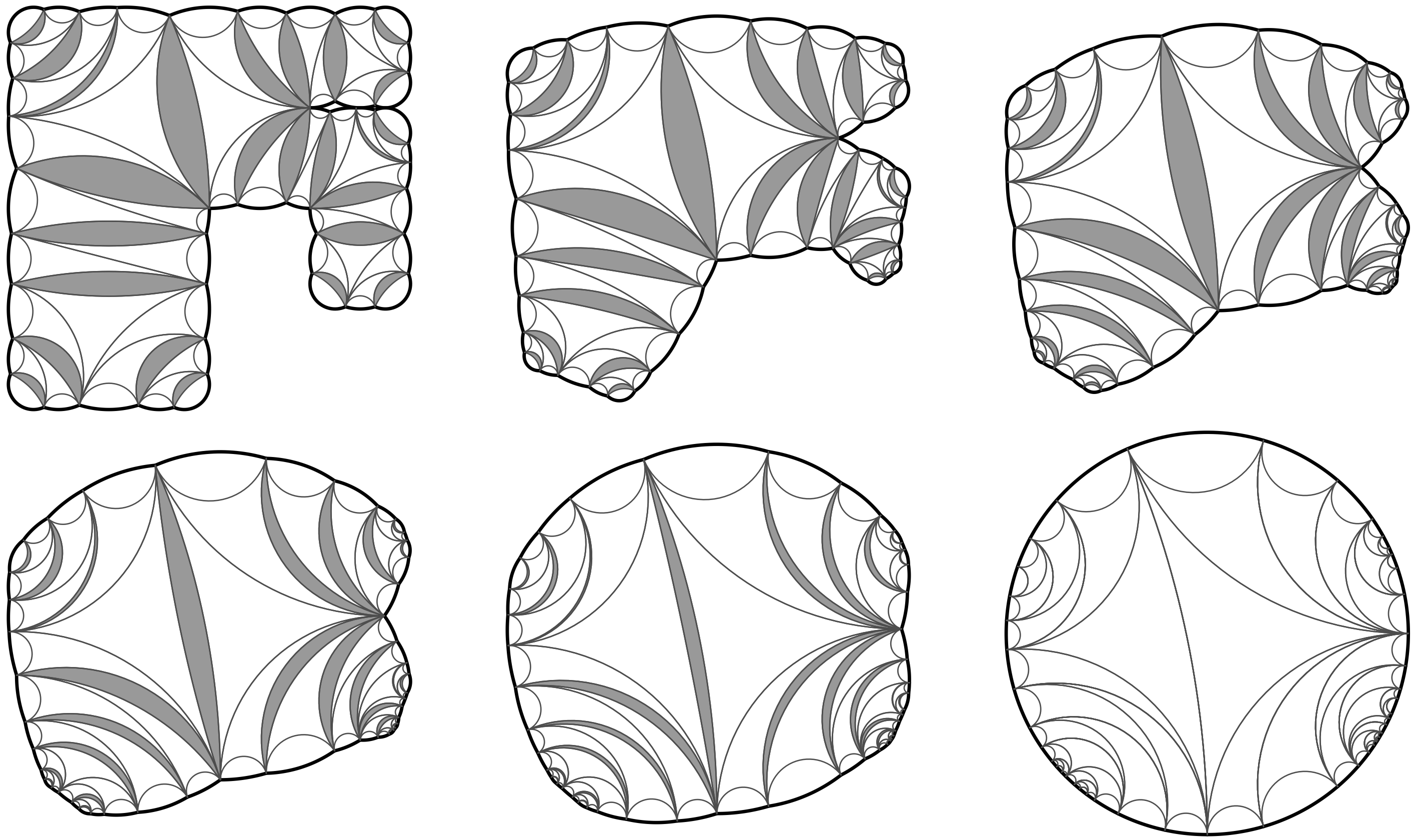

Planar version of previous figure.

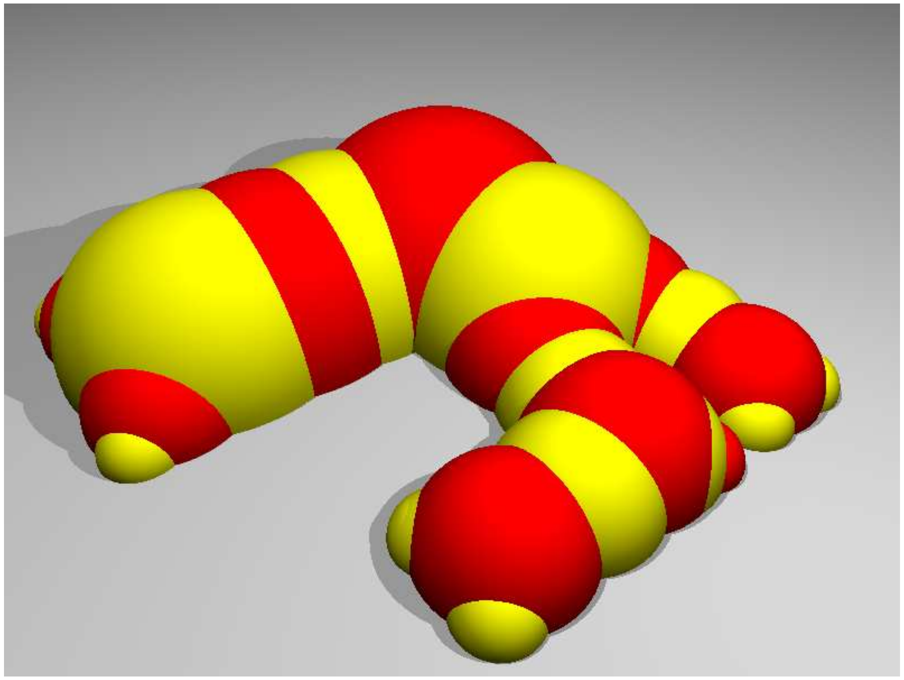 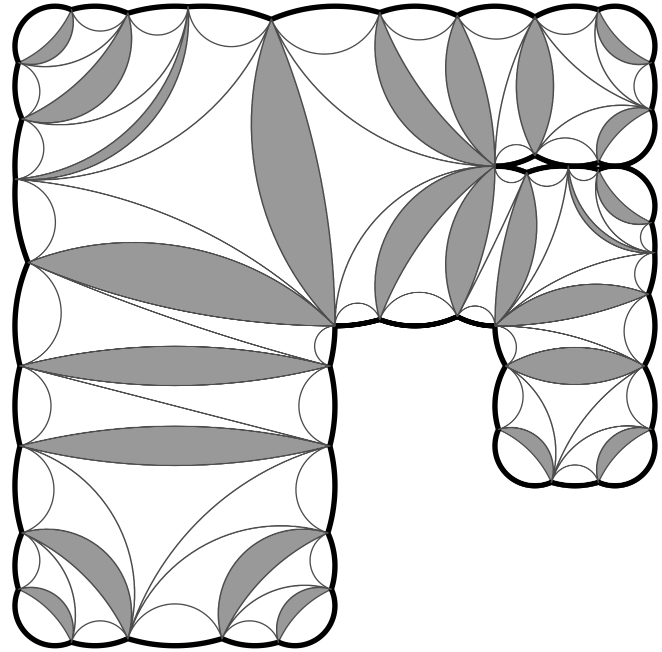

How are these connected?

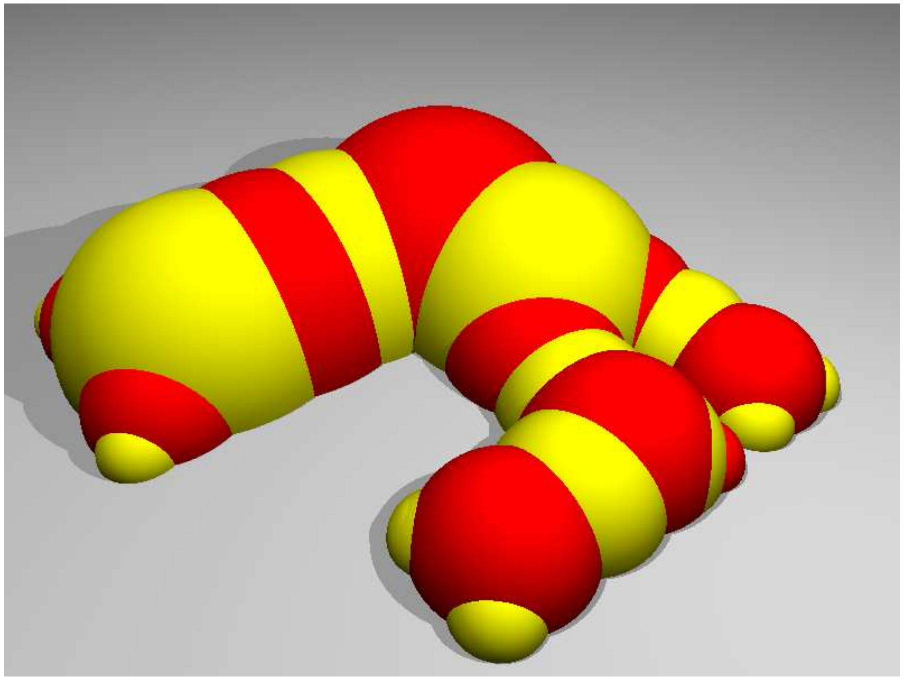 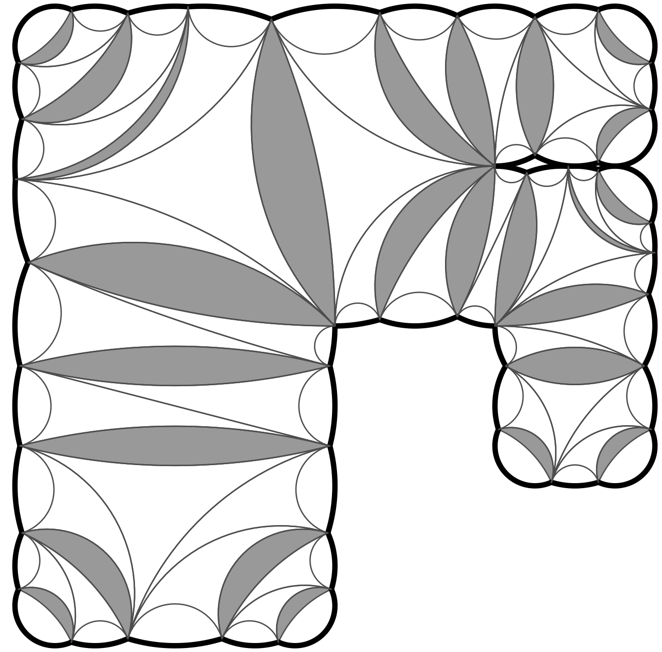

How are these connected? By projection.

Nearest point map in R^n is Lipschitz.

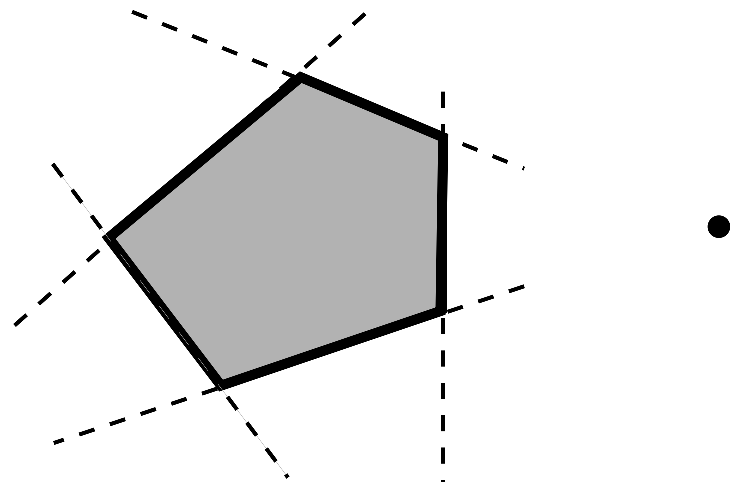

Nearest point map in R^n is Lipschitz.

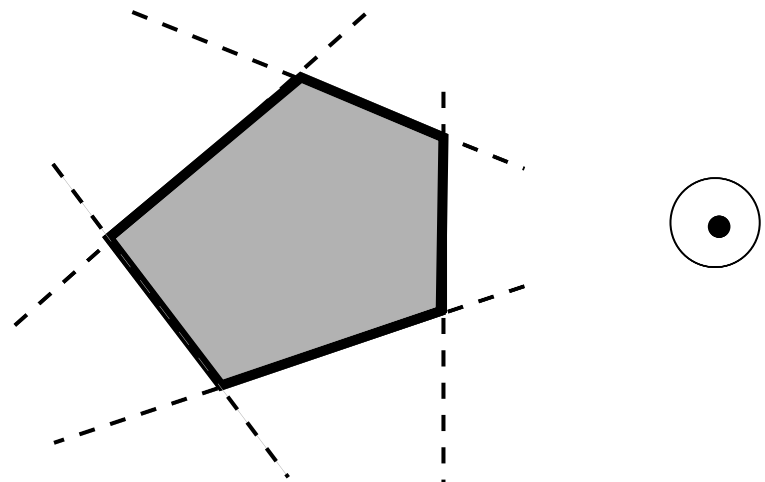

Nearest point map in R^n is Lipschitz.

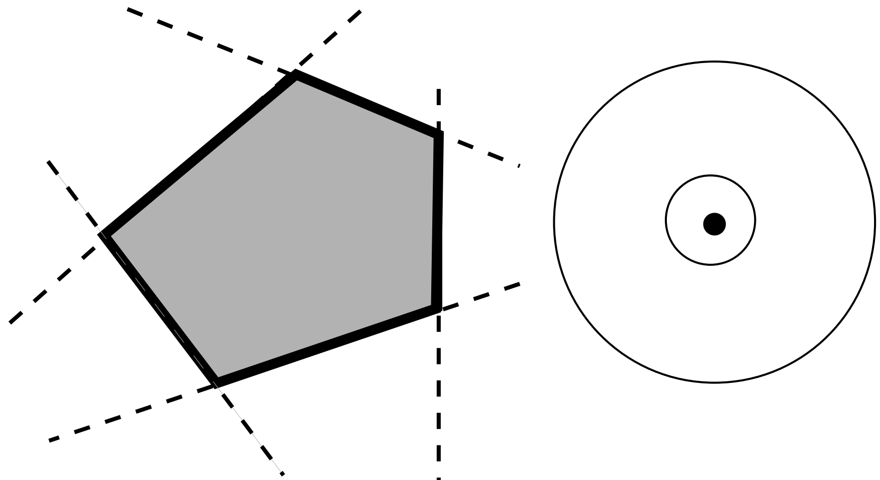

Nearest point map in R^n is Lipschitz.

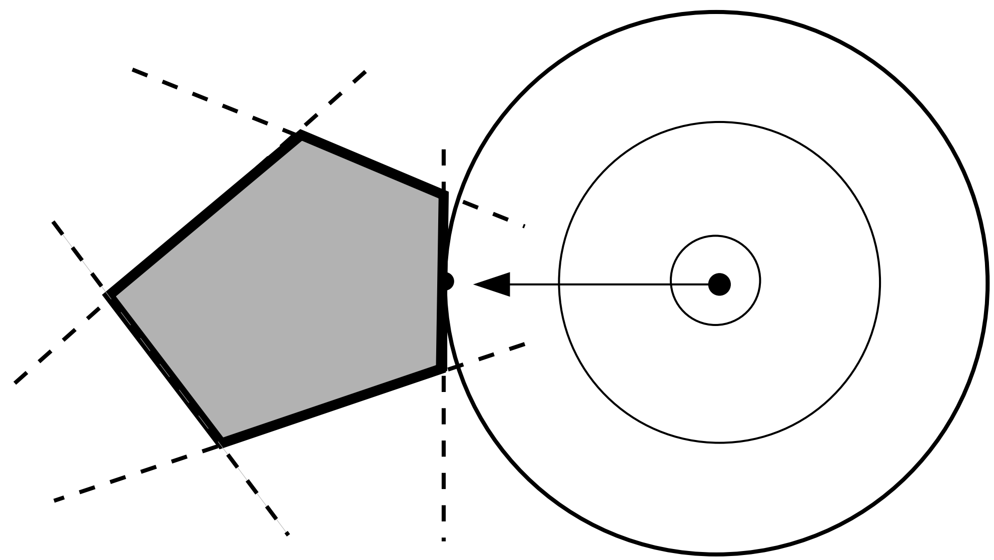

Nearest point map in R^n is Lipschitz.

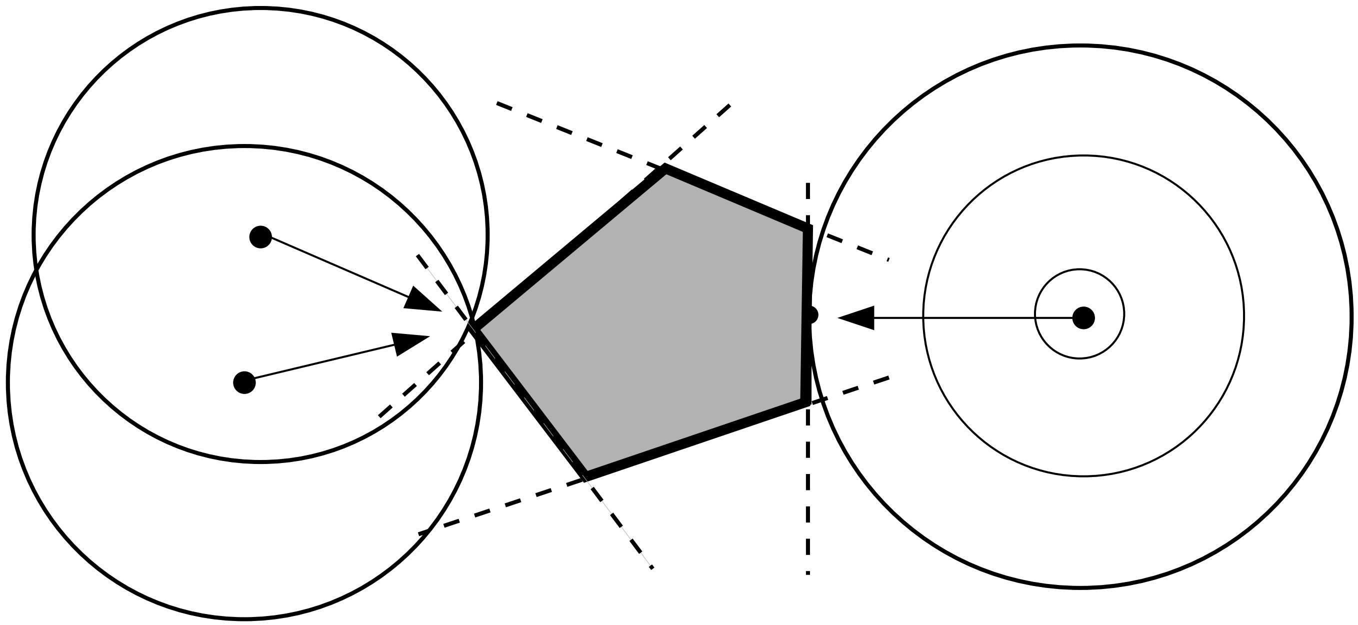

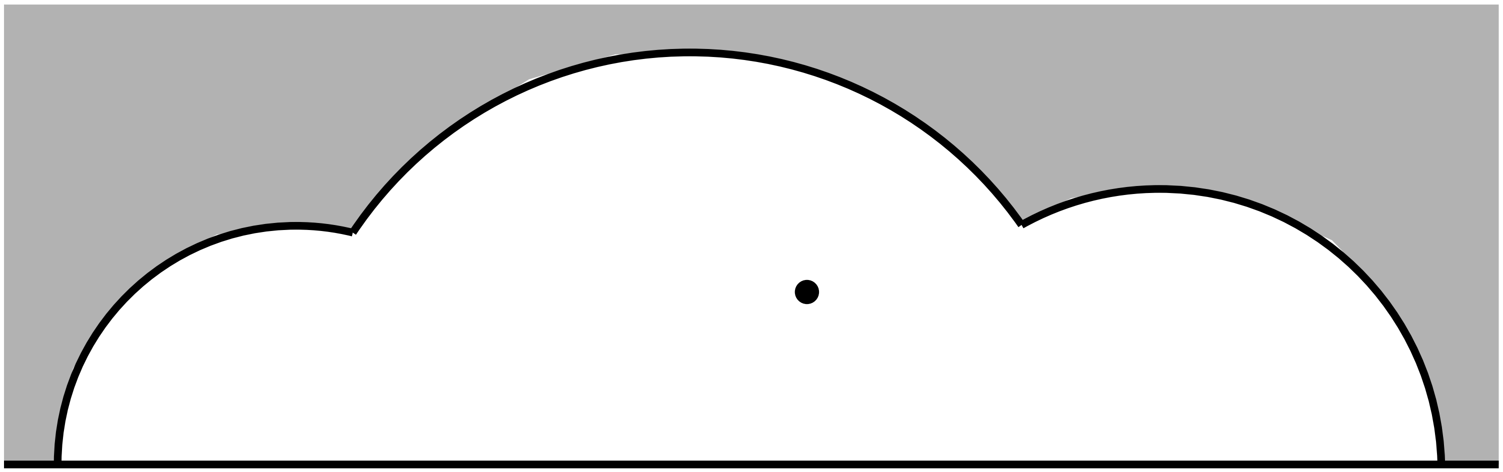

Region below dome is union of hemispheres

Hemispheres = hyperbolic half-spaces.

Region above dome is hyperbolically convex.

Consider nearest point retraction onto this convex set.

Hyperbolic metric on disk given by

$$d\rho = \frac{ds}{1-|z|^2} \simeq \frac{ds}{\mathrm{dist}(z, \partial D)}.$$

- Geodesics are circles perpendicular to boundary.
- Shaded region is hyperbolically convex.
- Metric transfers via conformal maps to other domains.
- For simply connected regions $d\rho \simeq ds/\mathrm{dist}(z, \partial\Omega)$.

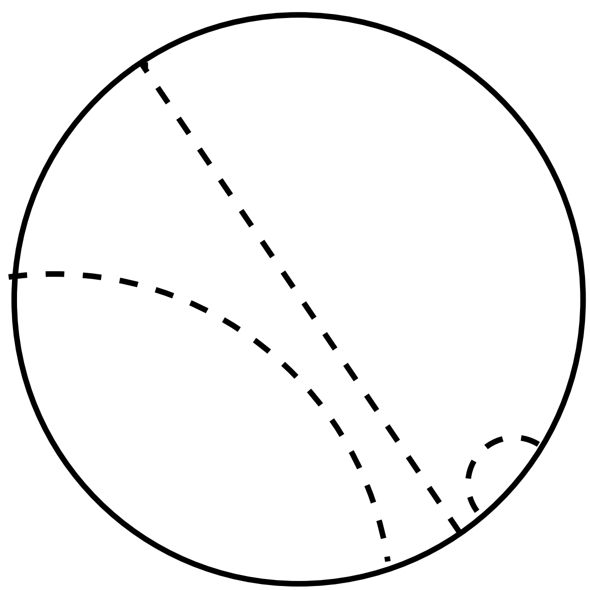
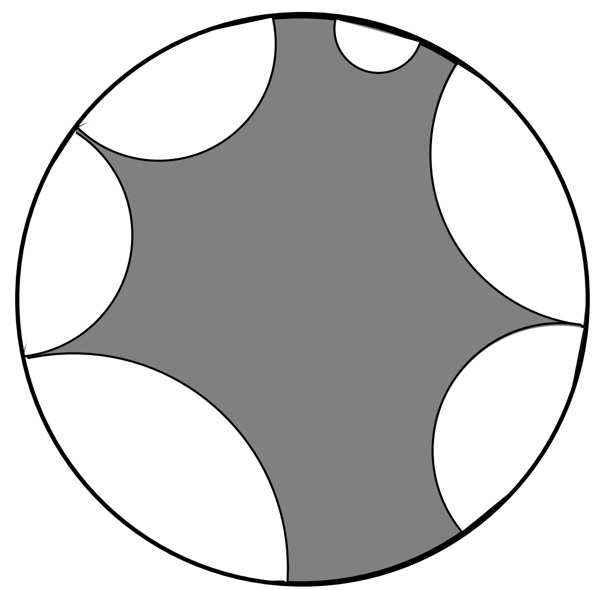

In the upper half-space $R^3_+ = \{(x, y, t) : t > 0\}$, metric is $d\rho = ds/2t$.

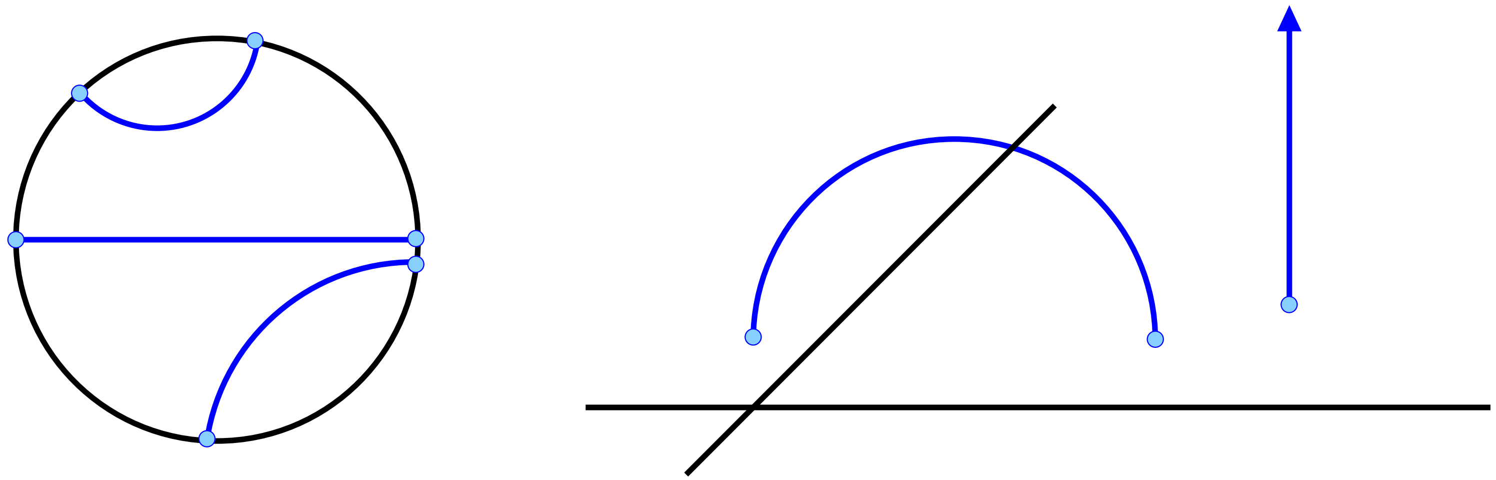

Geodesics in R^3_+ are vertical rays or semi-circles perpendicular to R^2.

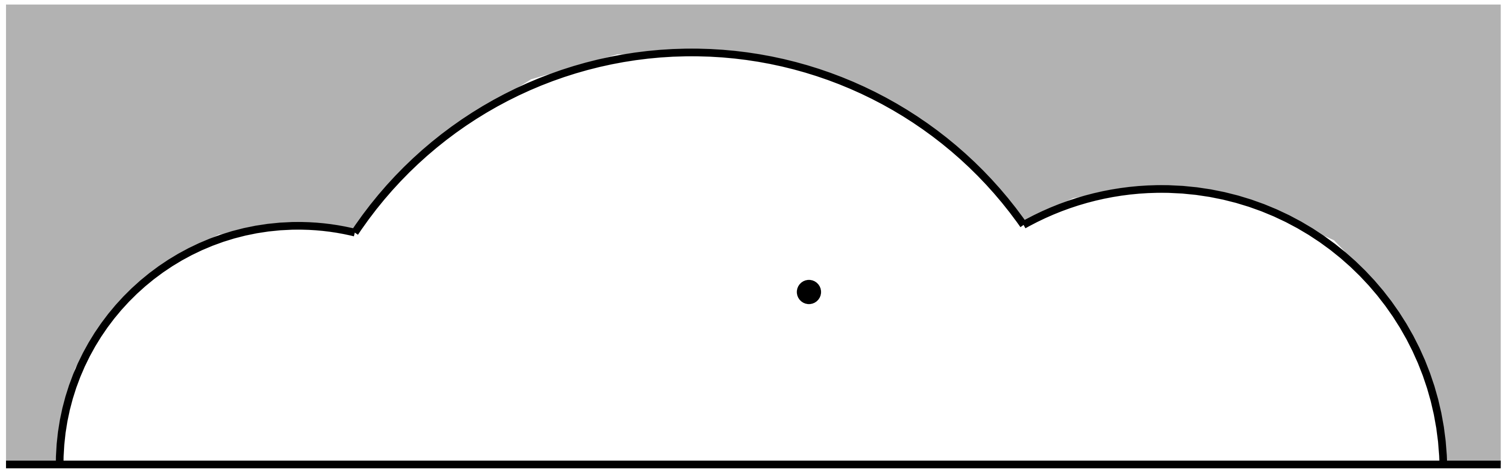

Region below dome is union of hemispheres

Hemispheres = hyperbolic half-spaces.

Region above dome is hyperbolically convex.

Consider nearest point retraction onto this convex set.

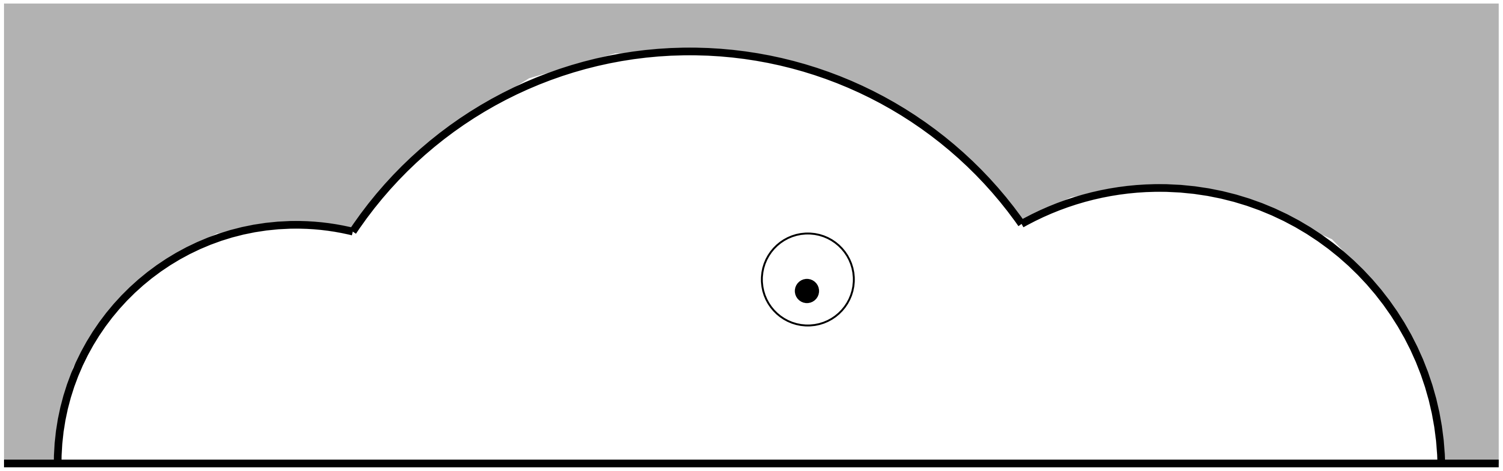

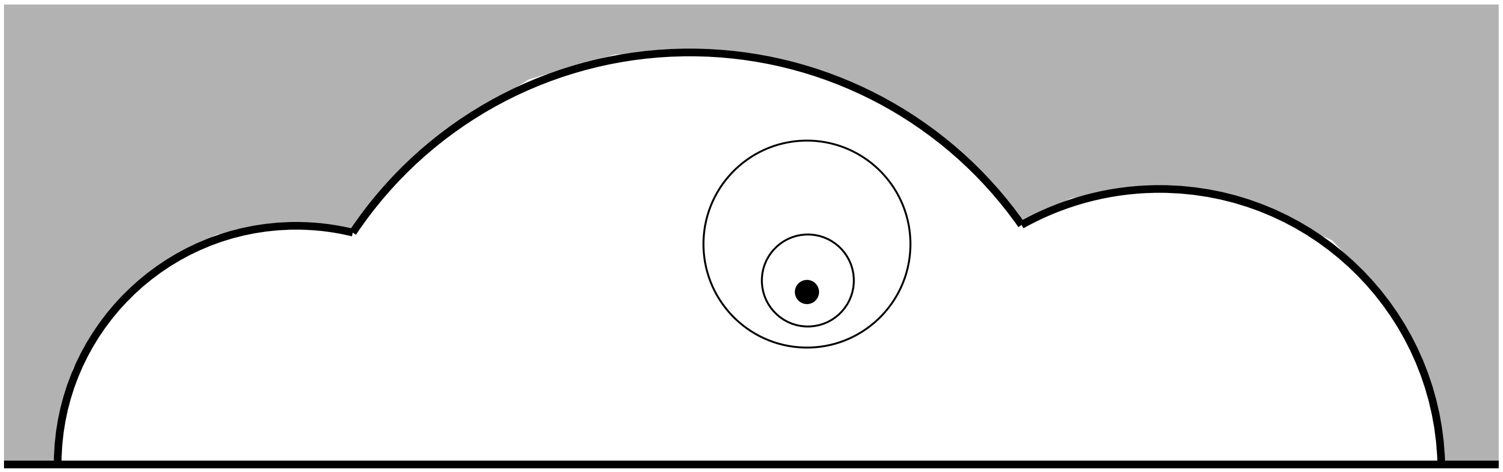

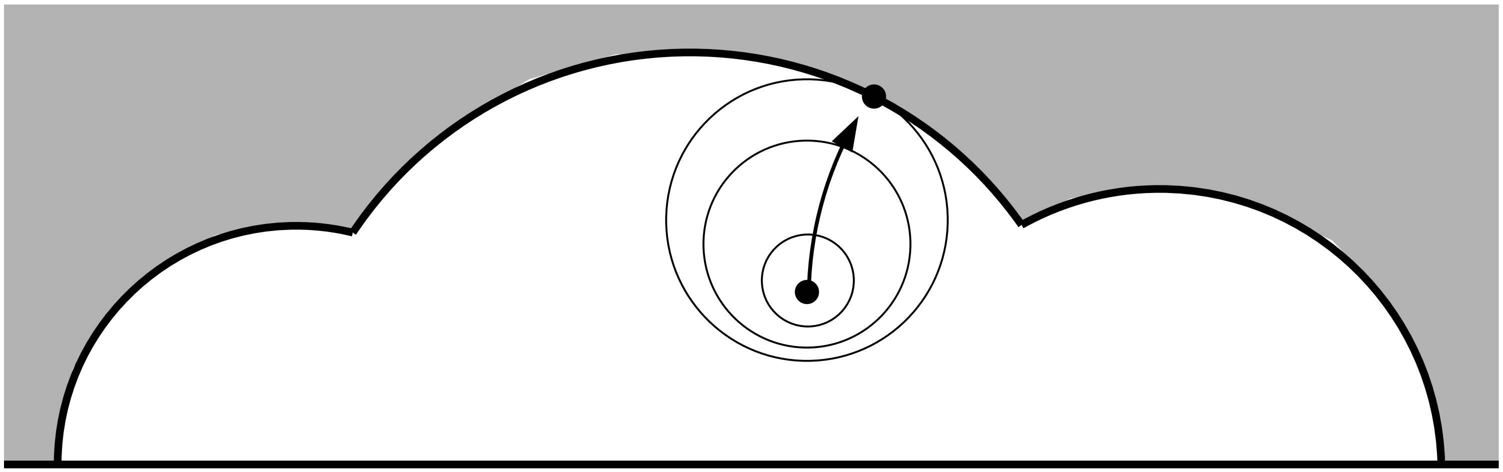

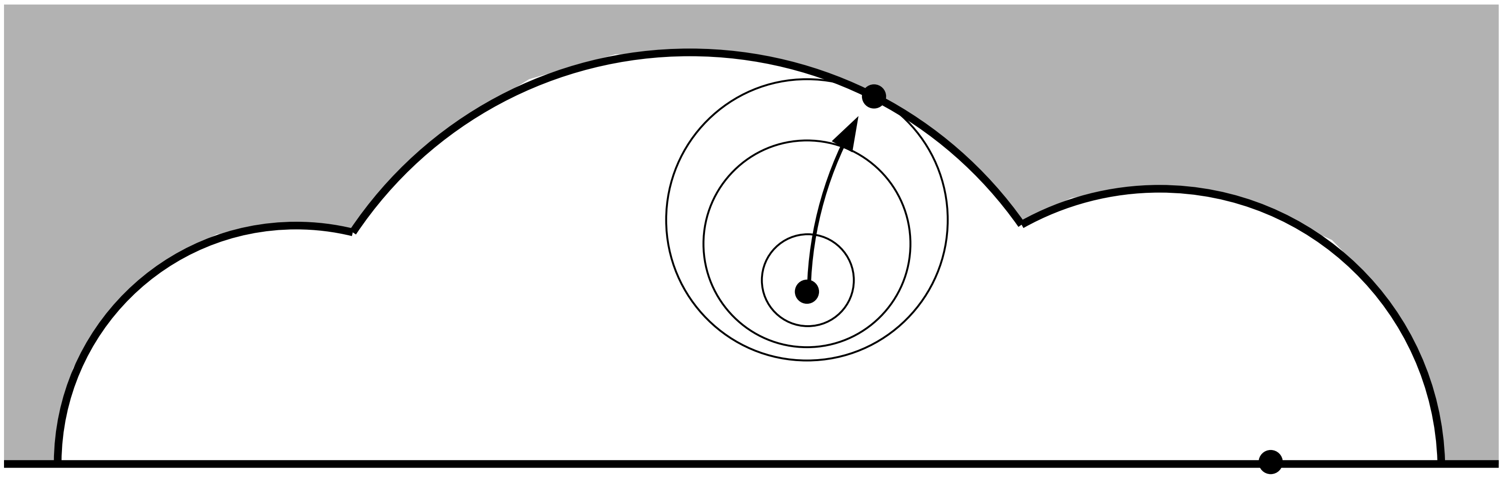

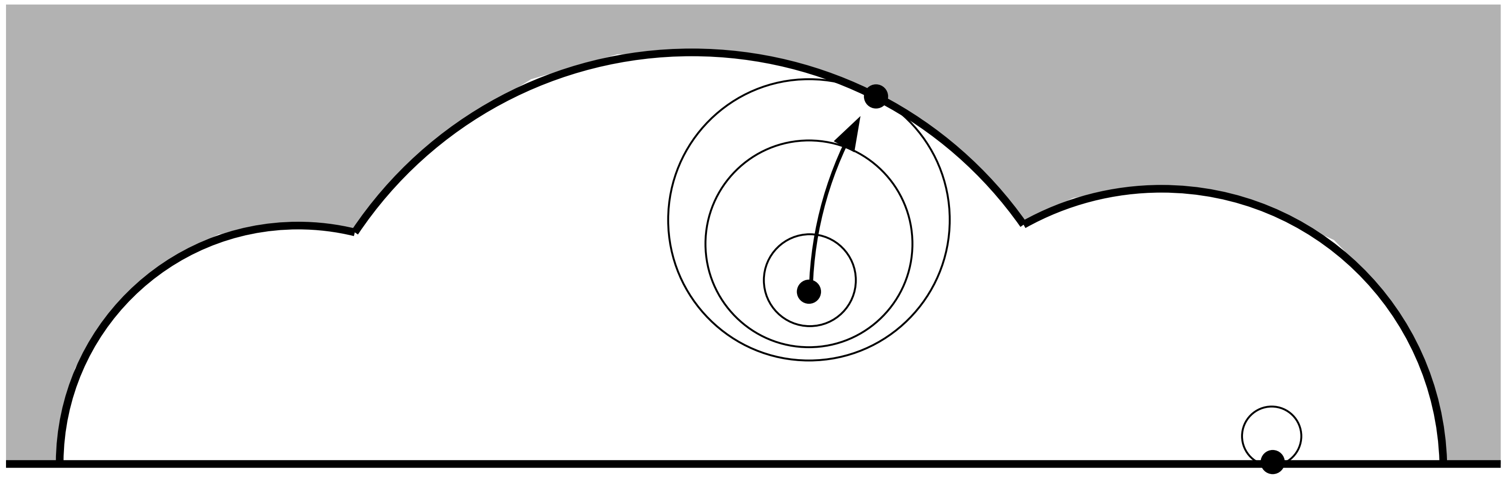

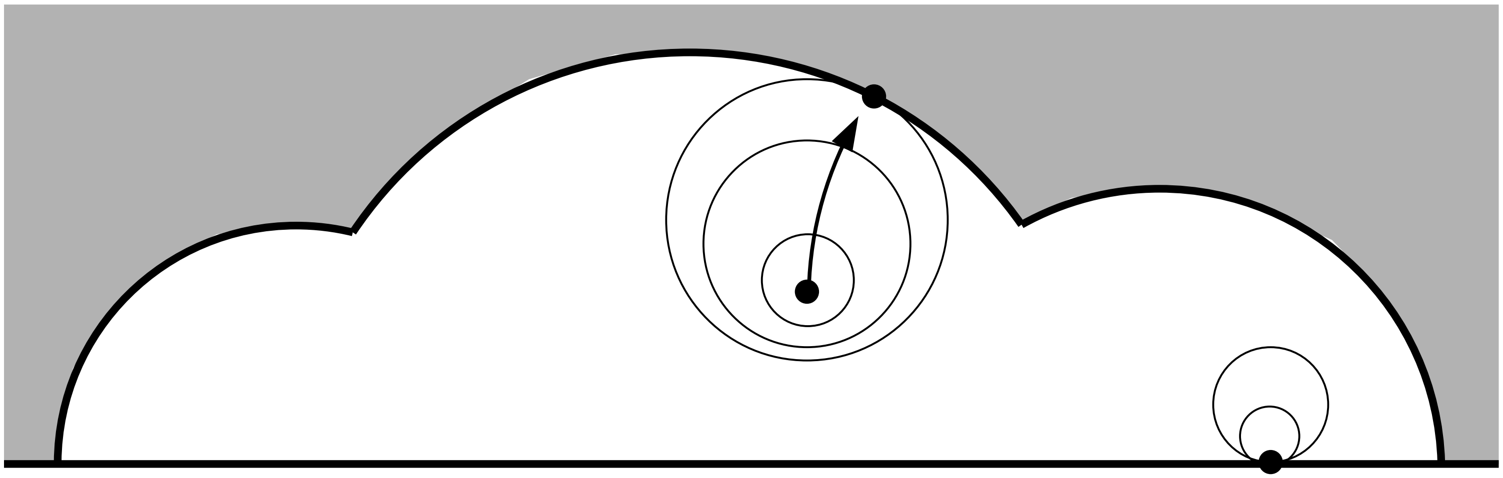

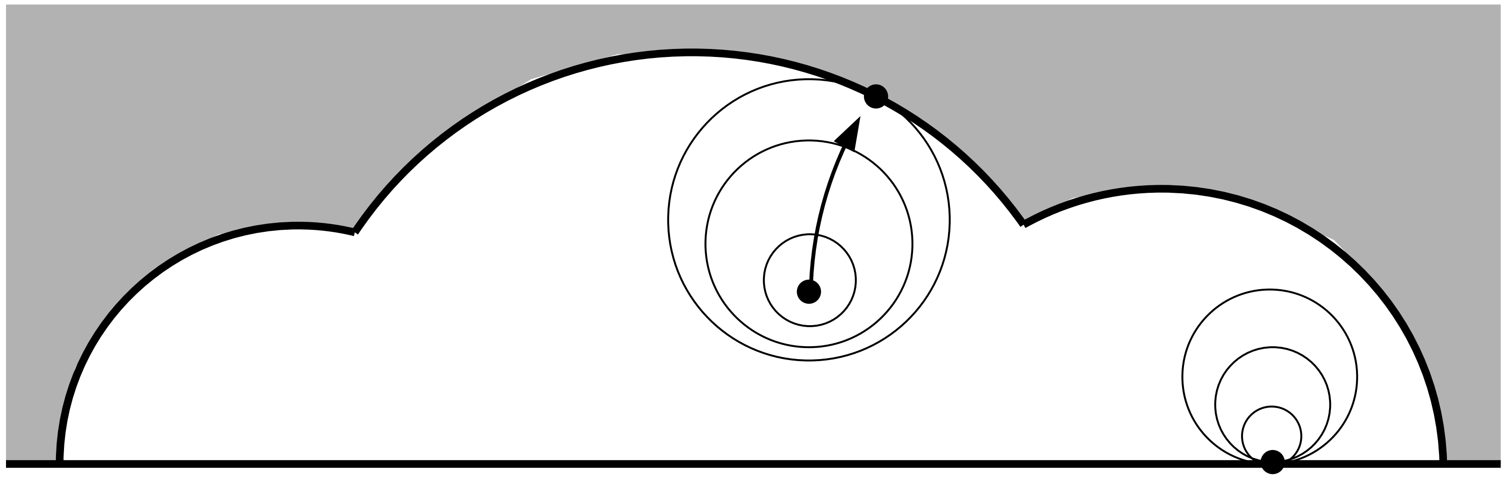

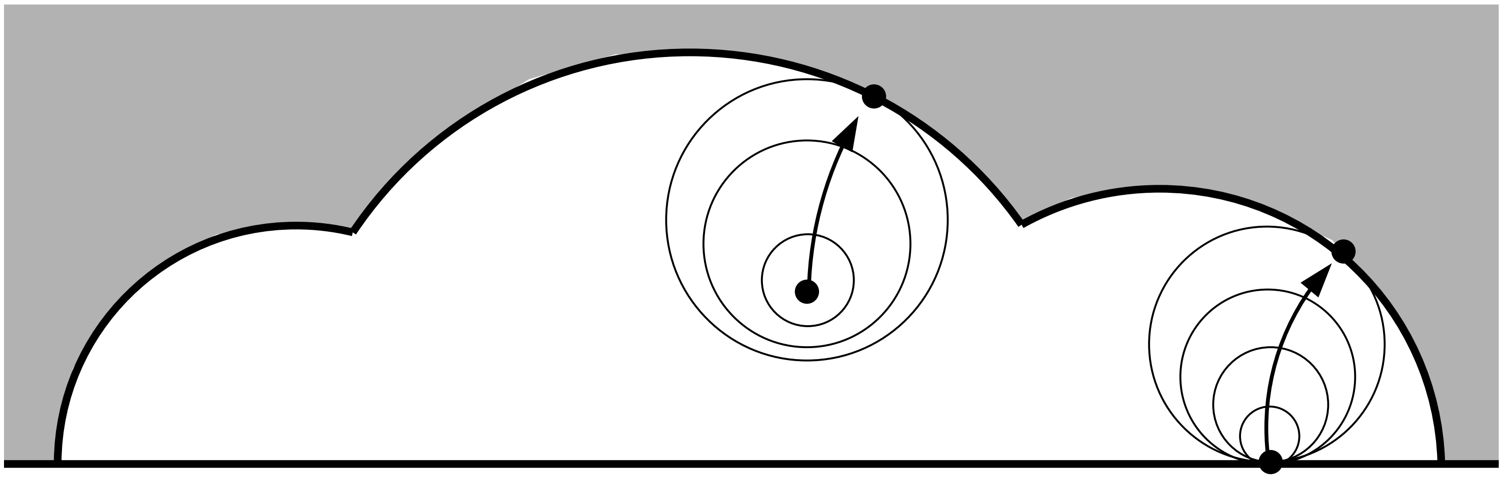

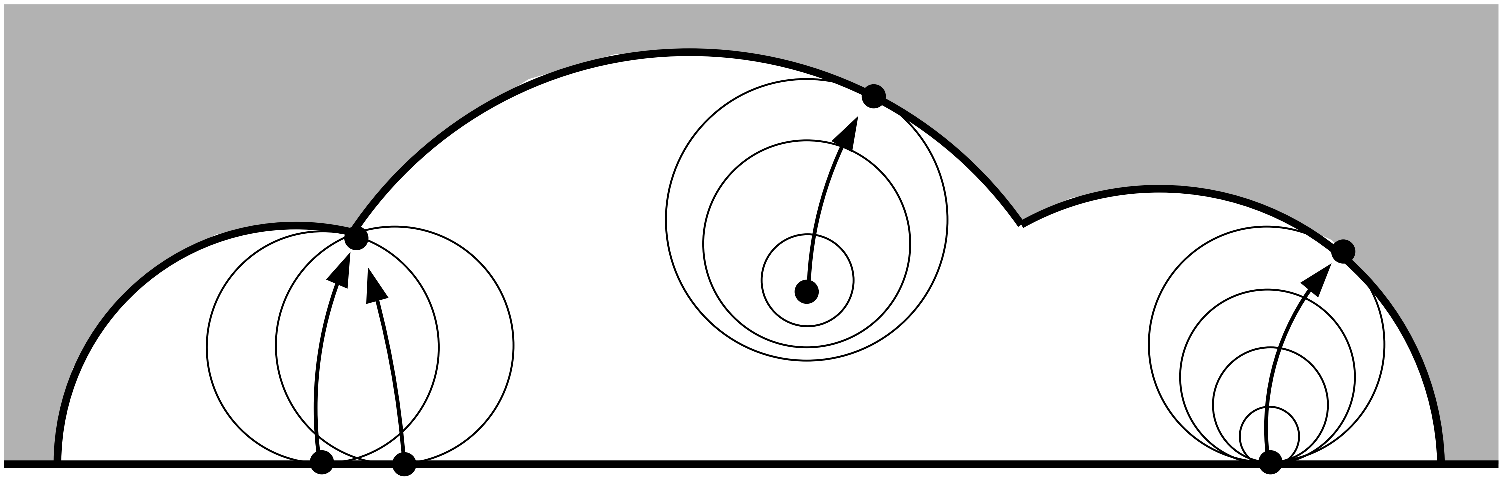

Need not be a homeomorphism, but . . .

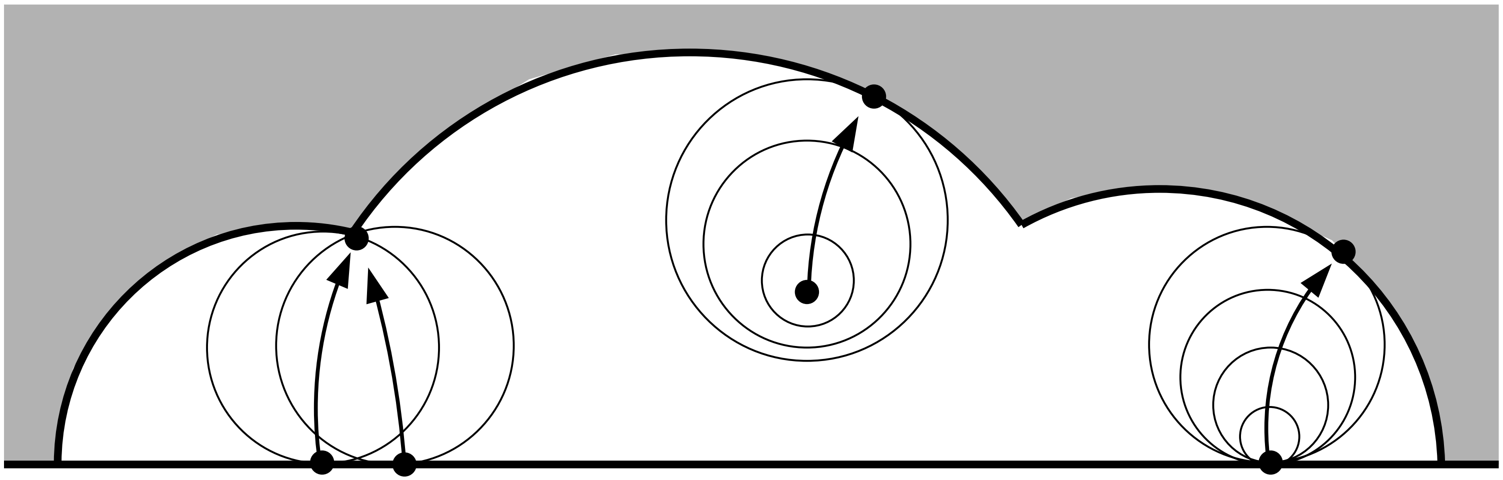

Need not be a homeomorphism, but it is a **quasi-isometry**

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

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

Metrics are hyperbolic metrics on Ω and S.

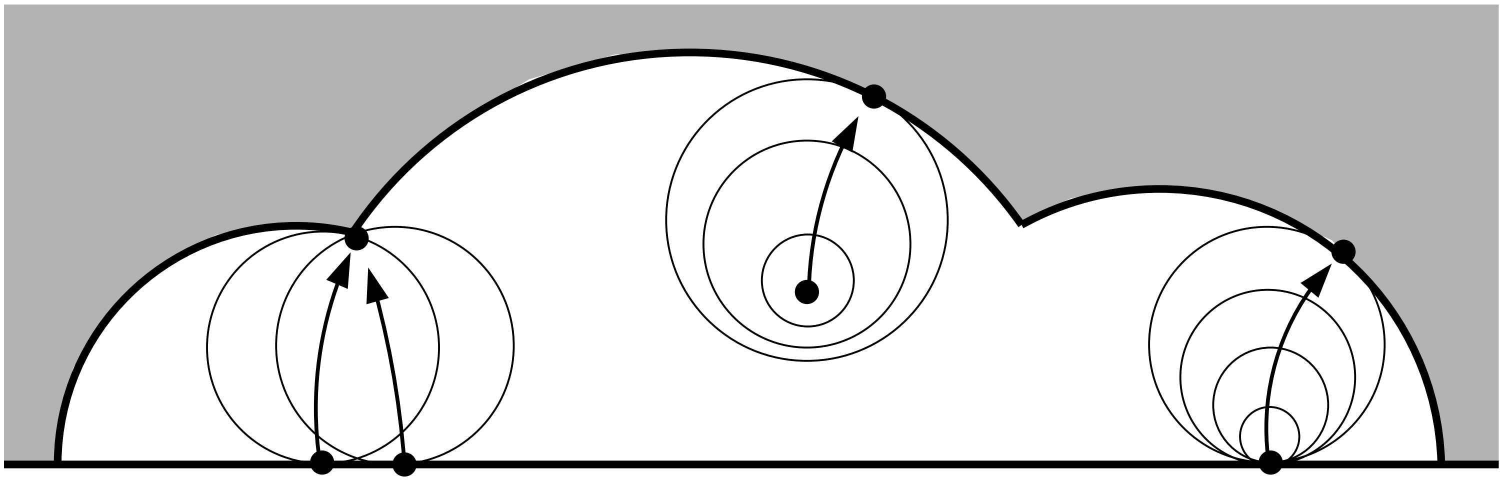

"Smoothing" gives K-QC map fixing boundary points.

Sullivan's theorem: K is independent of domain.

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

Best value unknown, $2.1 < K < 7.82$.

  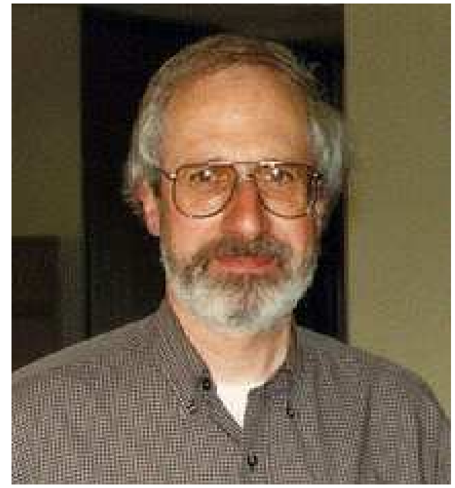

Dennis Sullivan　David Epstein　Al Marden

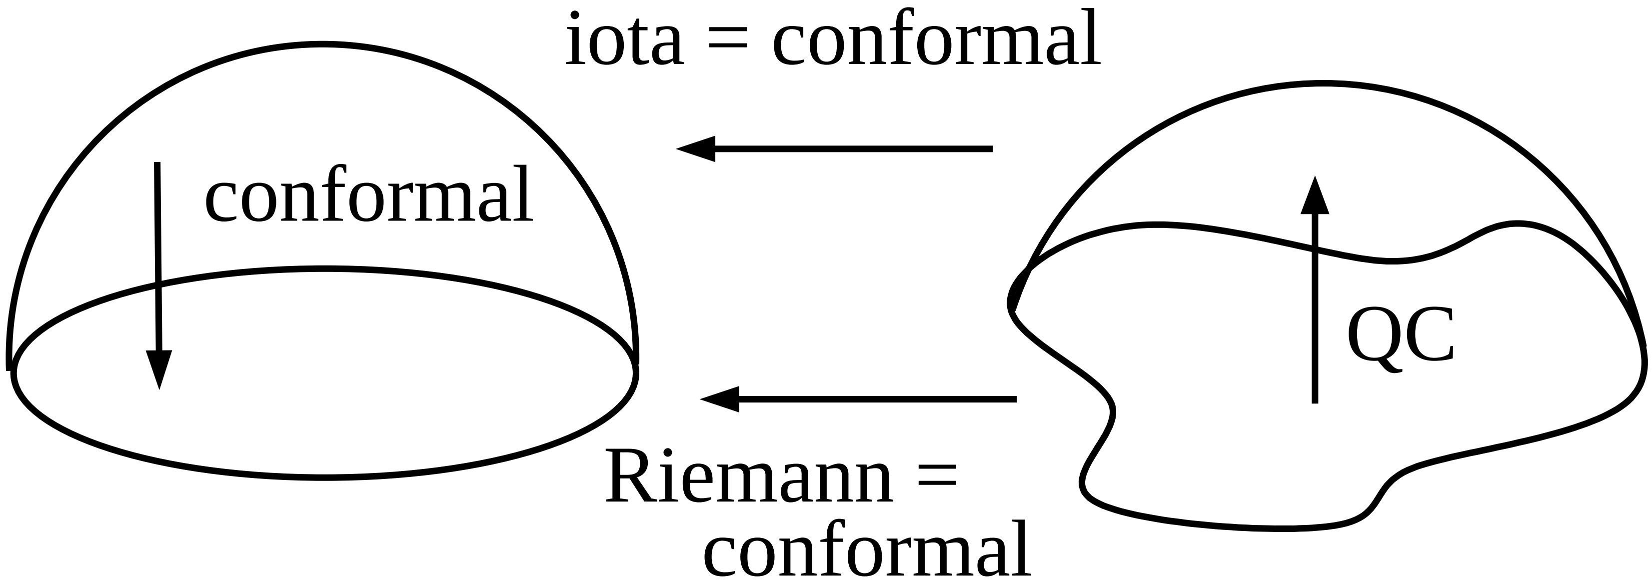

Iota = conformal from dome to disk.

Medial axis flow = boundary values of iota

Riemann map = conformal from base to disk

Sullivan's thm $\Rightarrow$ MA-map has K-QC extension

Sketch of proof that R is quasi-isometry

One direction: R is Lipschitz.

Other direction: R^{-1} is Lipschitz at distances ≥ 1.

Fact 1: If $z \in \Omega$, $\infty \notin \Omega$,

$$r \simeq \mathrm{dist}(z, \partial\Omega) \simeq \mathrm{dist}(R(z), R^2) \simeq |z - R(z)|.$$

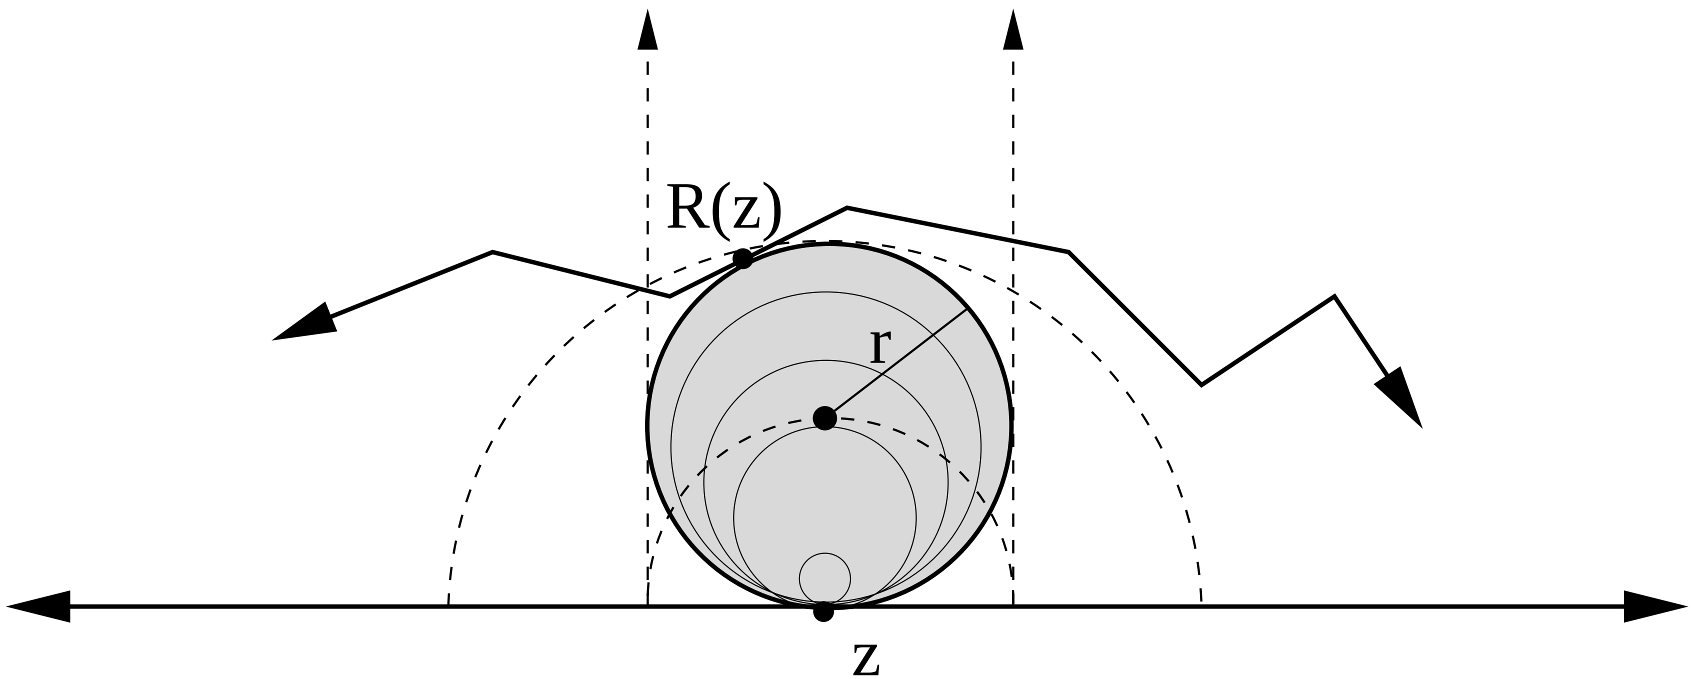

Fact 2: R is Lipschitz.

- Ω simply connected $\Rightarrow d\rho \simeq |dz|/\text{dist}(z, \partial\Omega)$.

- $z \in D \subset \Omega$ and $R(z) \in \text{Dome(D)} \Rightarrow z$ in hyperbolic convex hull of $\partial\Omega \cap \partial D$ in D.

$$\Rightarrow \text{dist}(z, \partial\Omega)/\sqrt{2} \leq \text{dist}(z, \partial D) \leq \text{dist}(z, \partial\Omega)$$

$$\Rightarrow \quad \rho_\Omega(z) \simeq \rho_D(z) = \rho_{\text{Dome}}(R(z)).$$

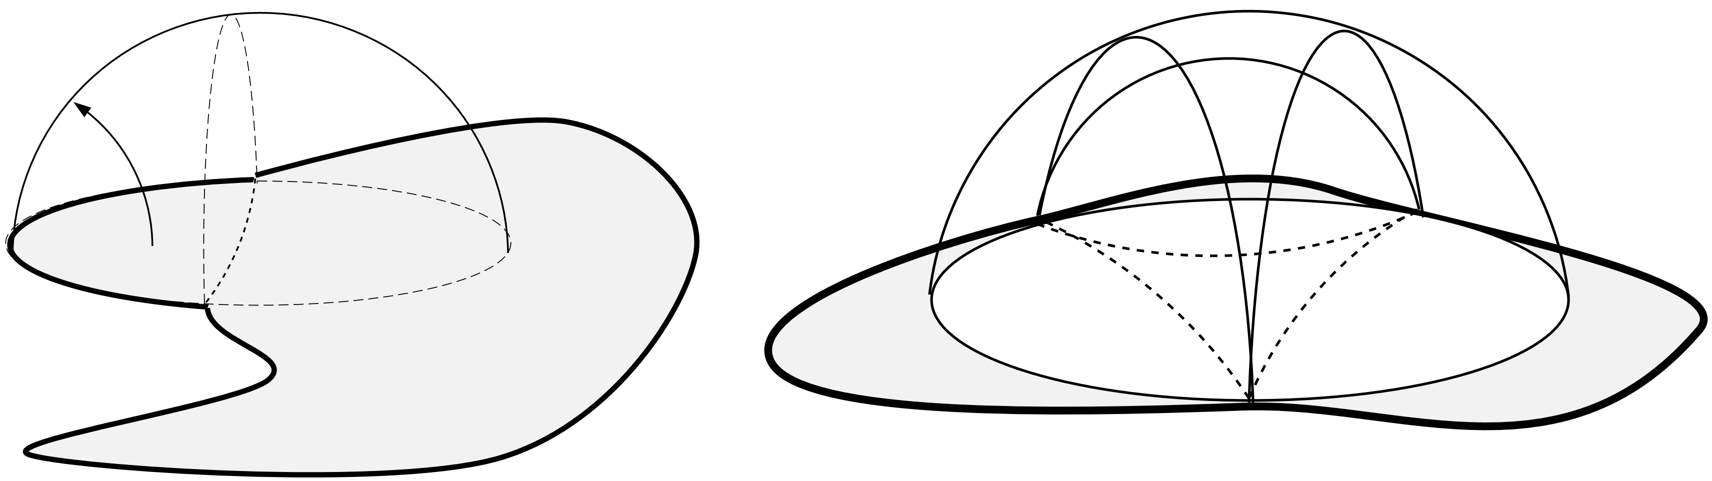

Fact 3: $\rho_S(R(z), R(w)) \leq 1 \Rightarrow \rho_\Omega(z, w) \leq C$.

Suppose $\text{dist}(R(z), R^2) = r$.

Suppose γ is geodesic on dome from $R(z)$ to $R(w)$.

$$\Rightarrow \qquad \text{dist}(\gamma, R^2) \simeq r$$
$$\Rightarrow \qquad \text{dist}(R^{-1}(\gamma), \partial\Omega) \simeq r, \qquad R^{-1}(\gamma) \subset D(z, Cr)$$
$$\Rightarrow \qquad \rho_\Omega(z, w) \leq C$$

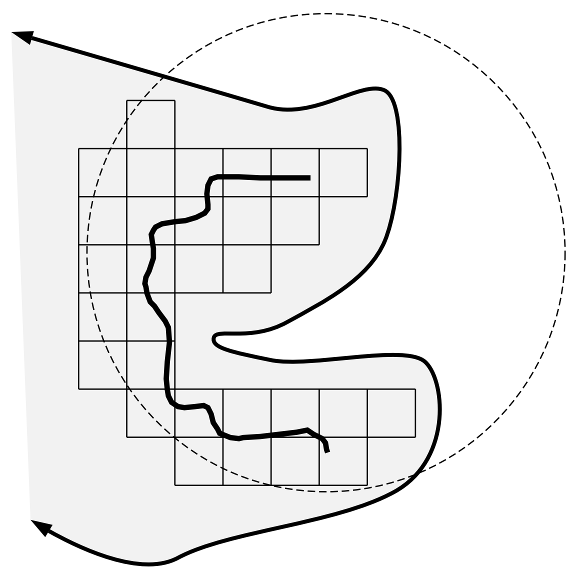

Moreover, $g = \iota \circ \sigma : \Omega \to D$ is locally Euclidean Lipschitz.

$$|g'(z)| \simeq \frac{\mathrm{dist}(g(z), \partial D)}{\mathrm{dist}(z, \partial\Omega)}.$$

Use Fact 1

$$\begin{aligned}
\mathrm{dist}(z, \partial\Omega) &\simeq \mathrm{dist}(R(z), R^2) \\
&\simeq \exp(-\rho_{R^3_+}(R(z), z_0)) \\
&\gtrsim \exp(-\rho_S(R(z), z_0)) \\
&= \exp(-\rho_D(g(z), 0)) \\
&\simeq \mathrm{dist}(g(z), \partial D)
\end{aligned}$$

The factorization theorem: Riemann map $f = h \circ g$ where

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

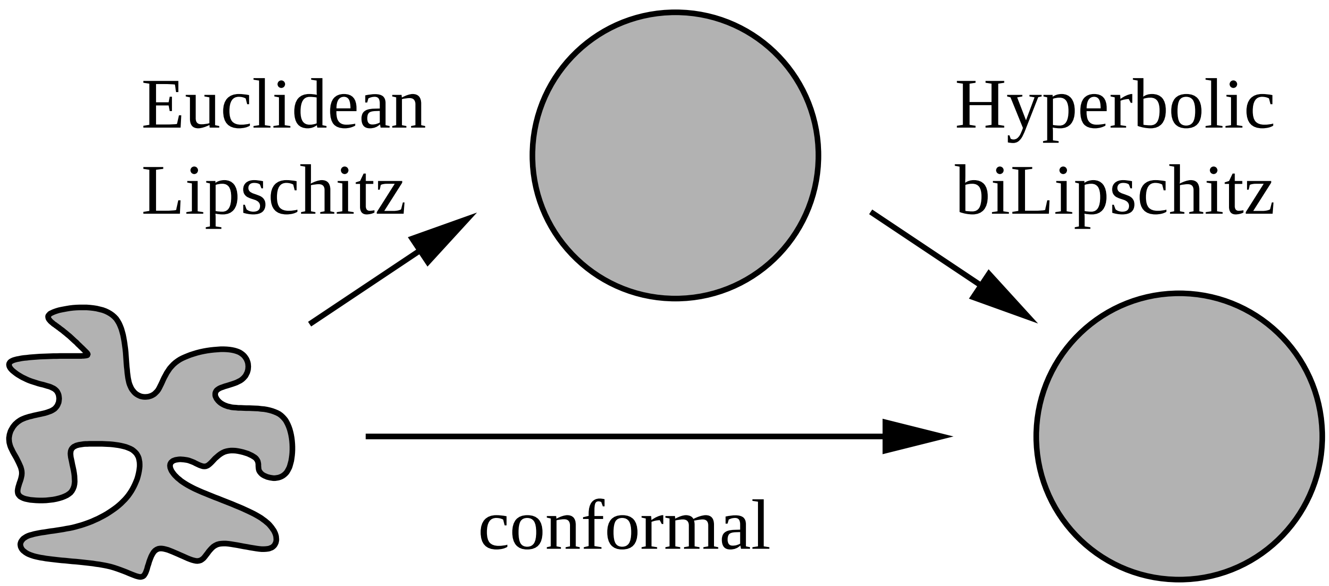

Cor: Any s.c. domain can be mapped 1-1, onto a disk by a contraction for the internal path metric.

Crescents in base can map to geodesics on surface.

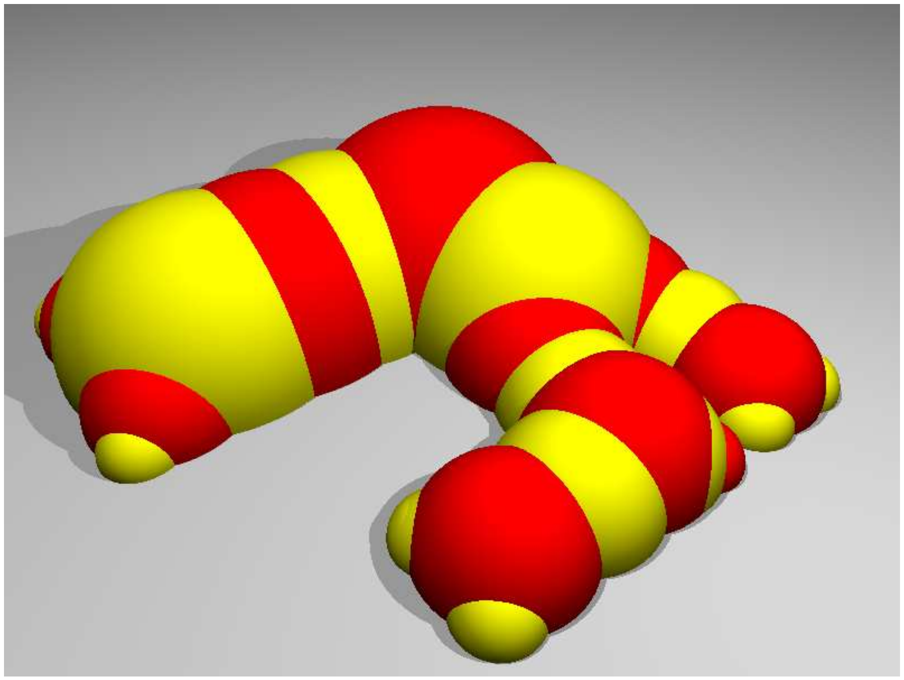

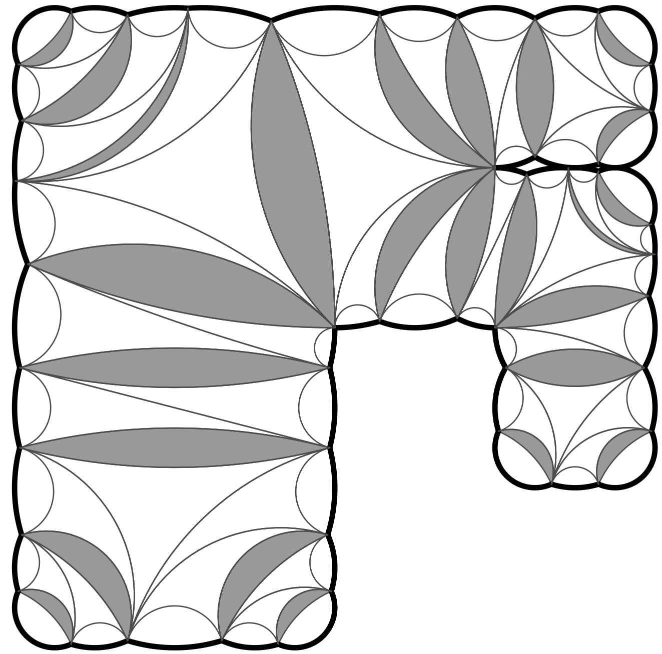

Gray collapses to bending lines, “width = angle”.

White maps isometrically to dome.

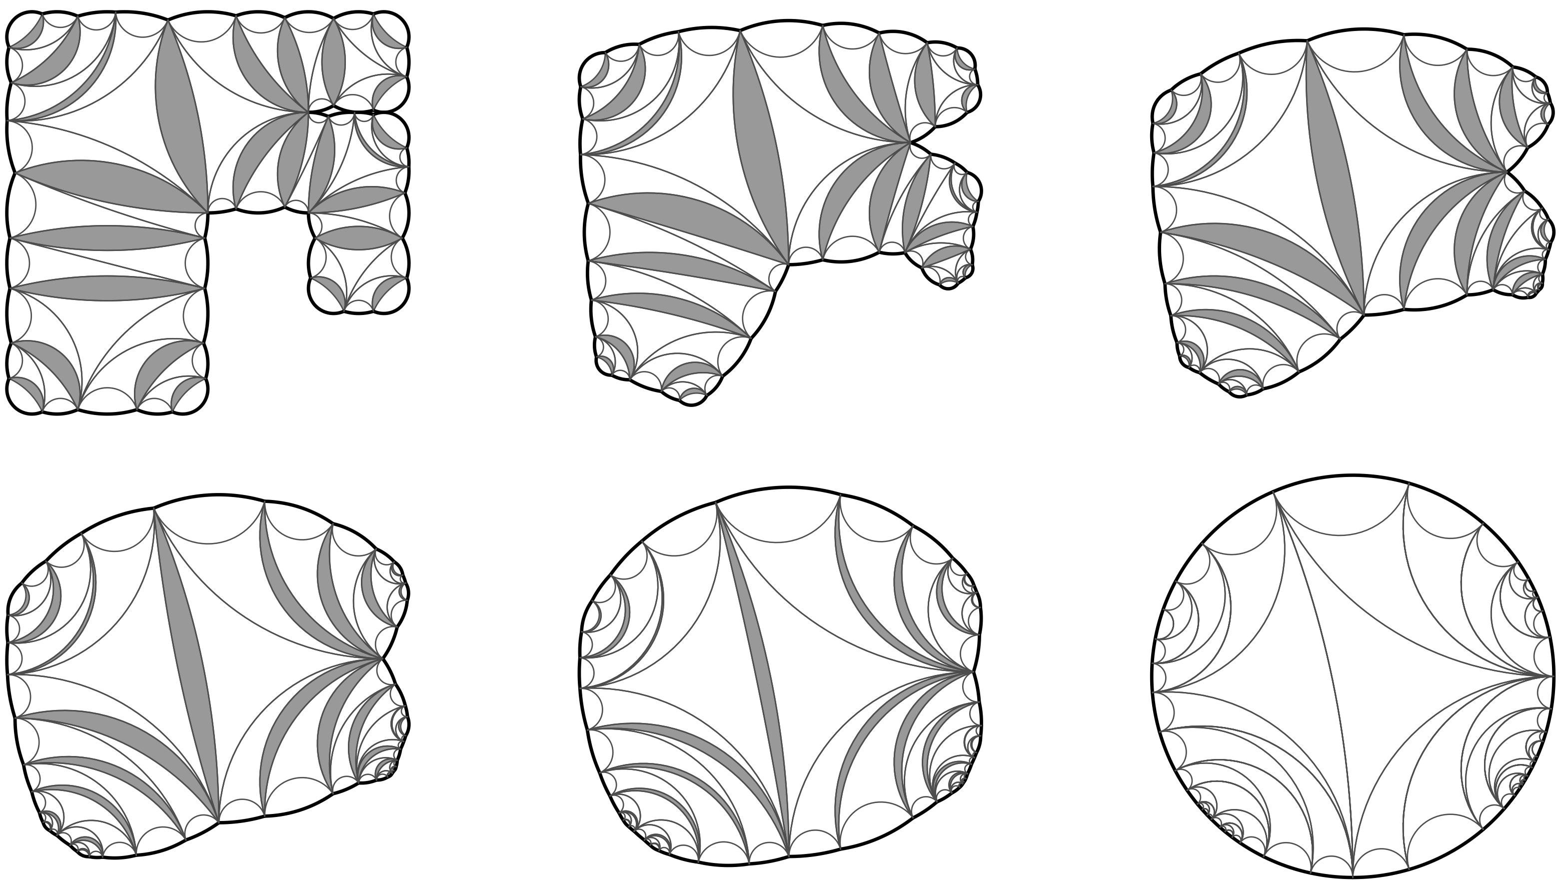

Angle scaling family - crescent angles decrease

"Morphs" region to disk. Intermediate regions need not be planar.

Gives continuation method to apply Newton's method.

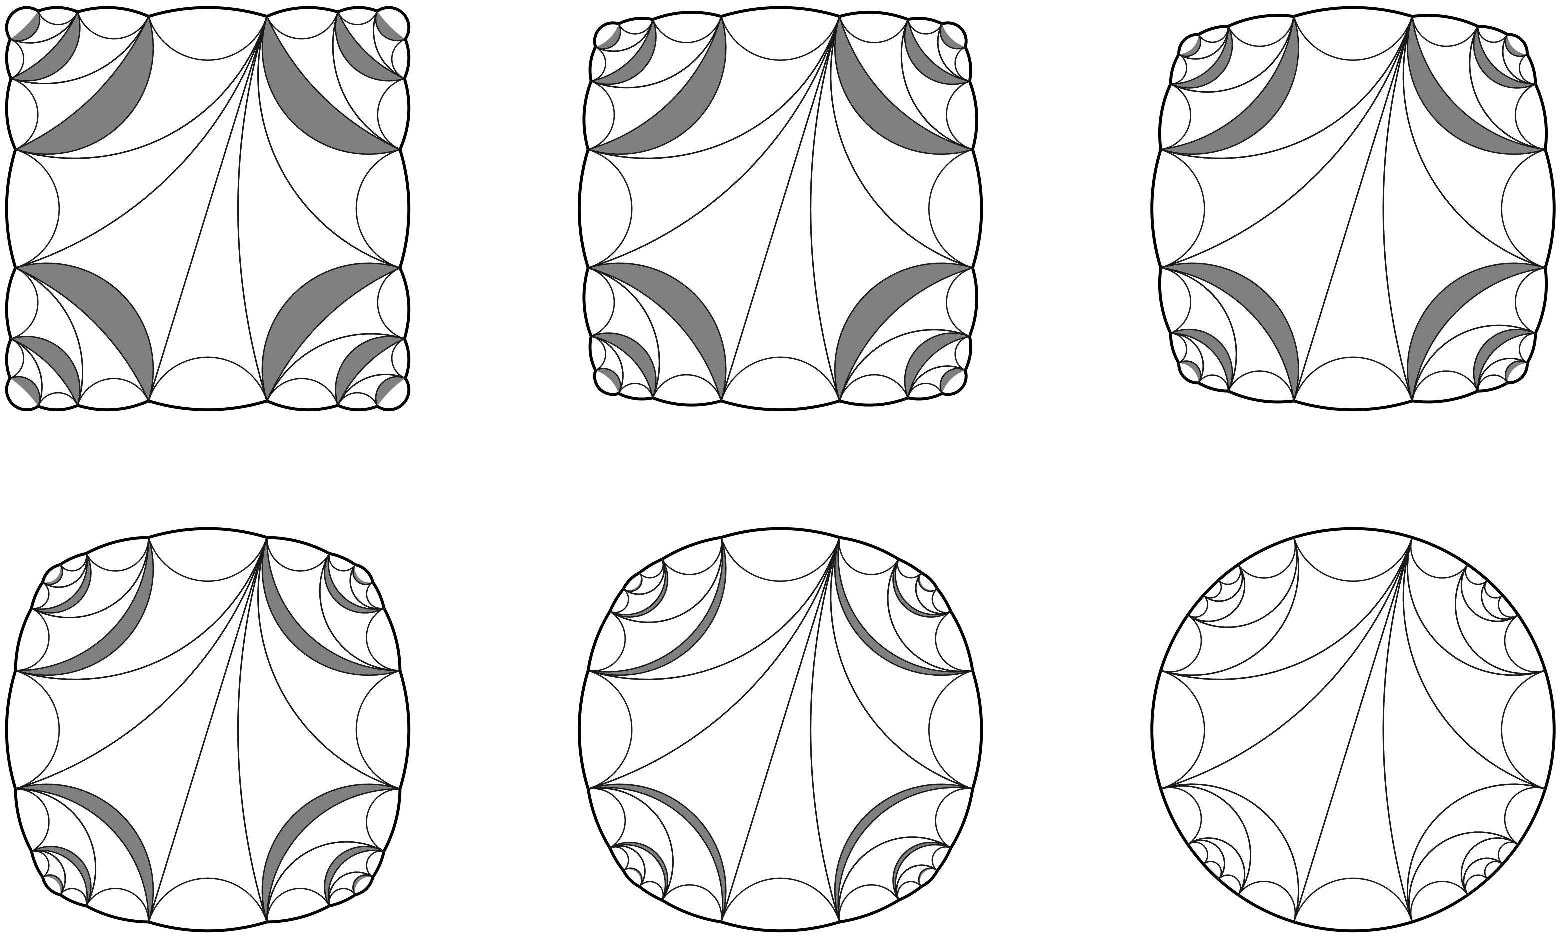

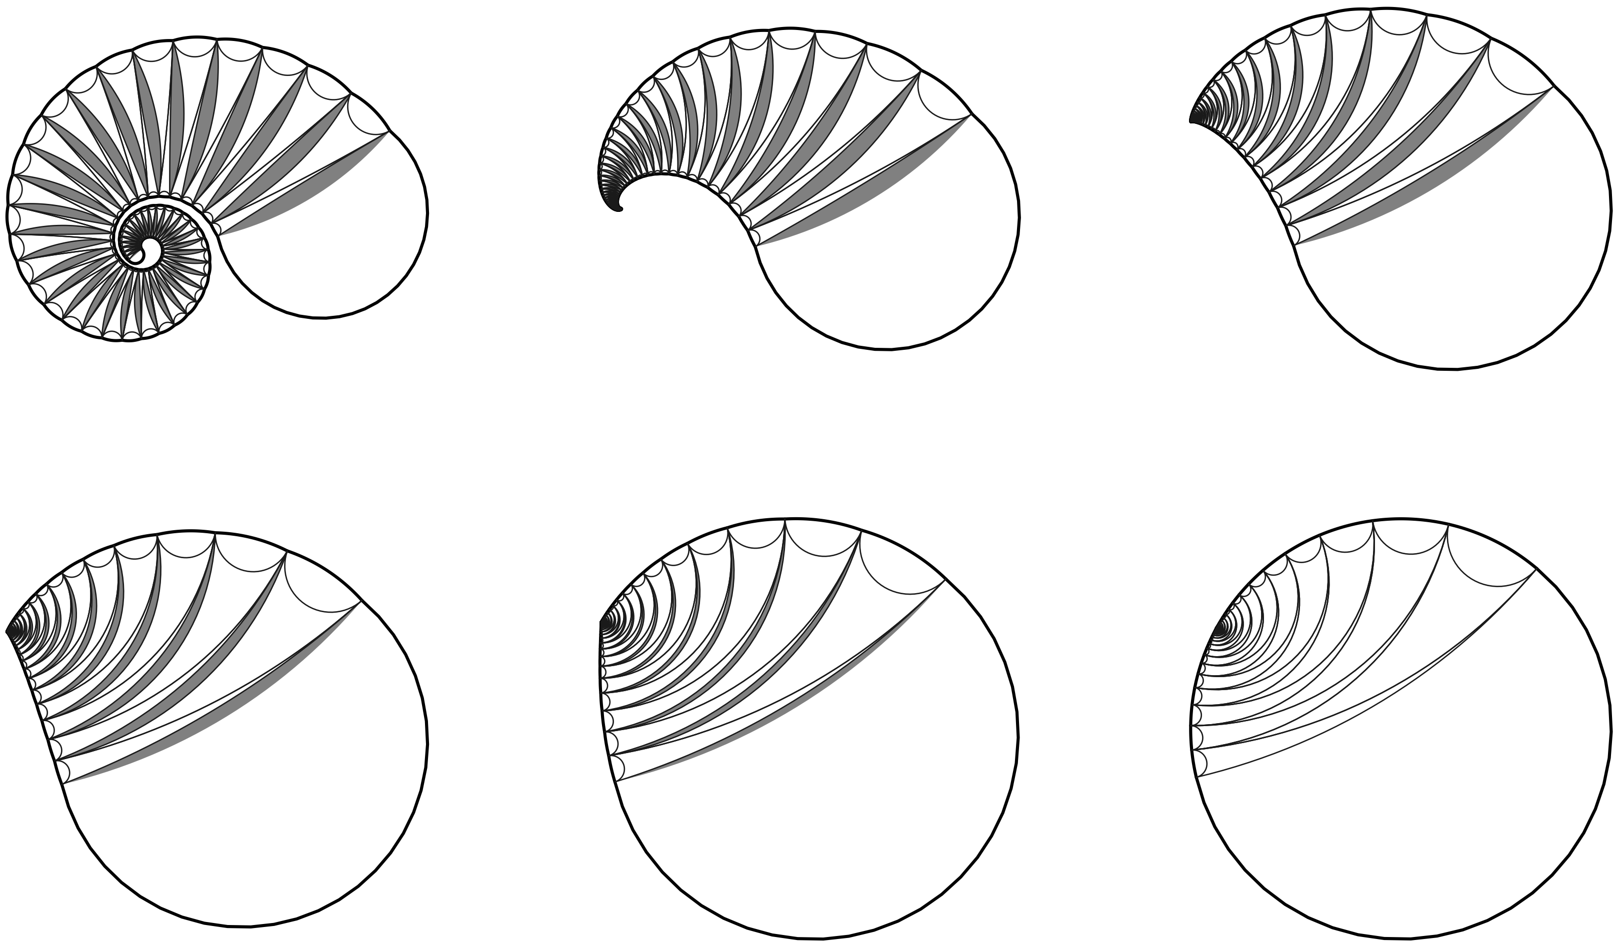

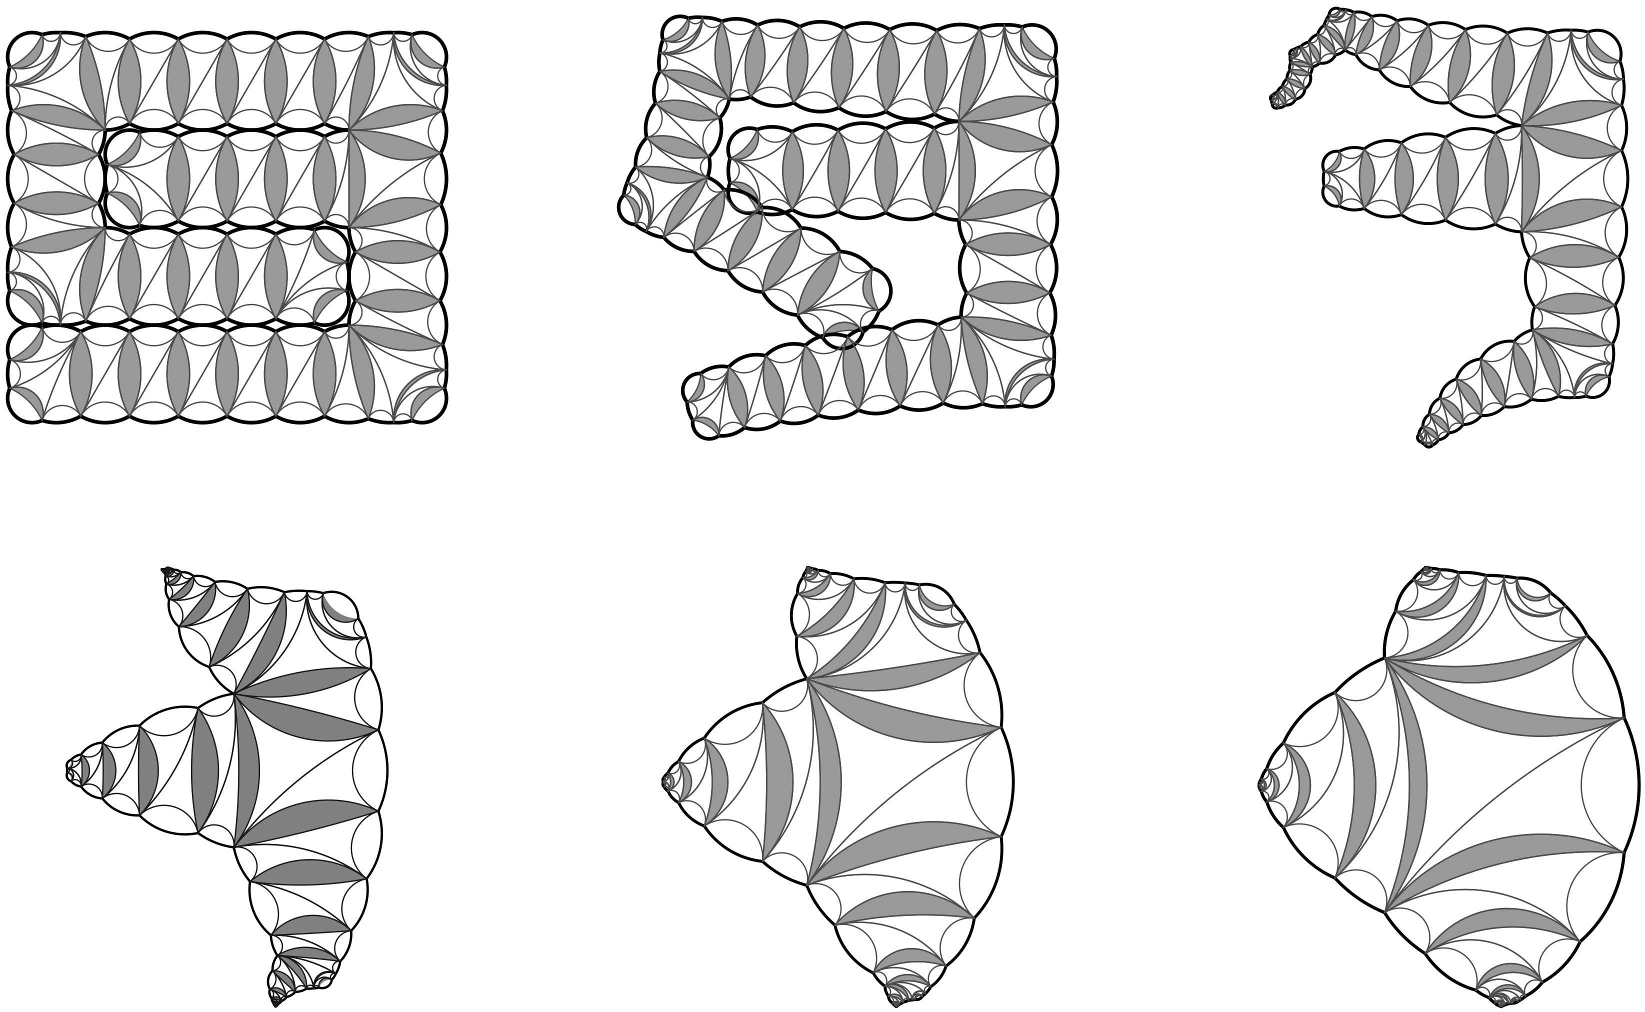

Application to computational geometry: Quadrilateral meshes

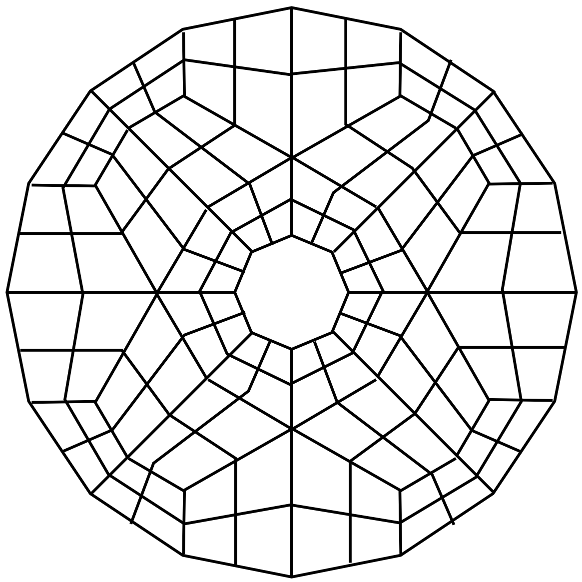

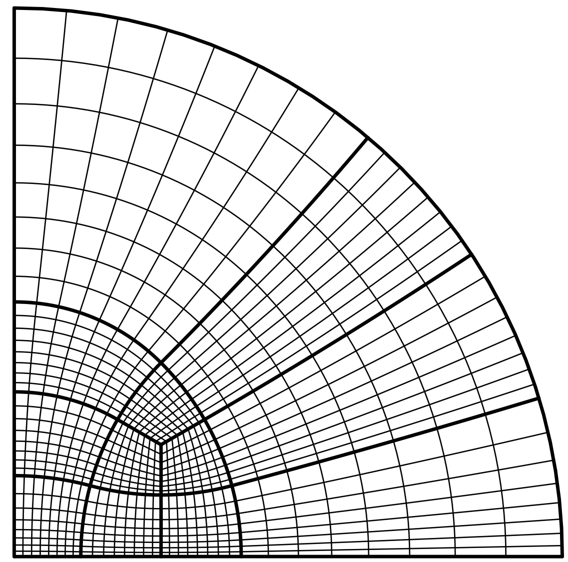

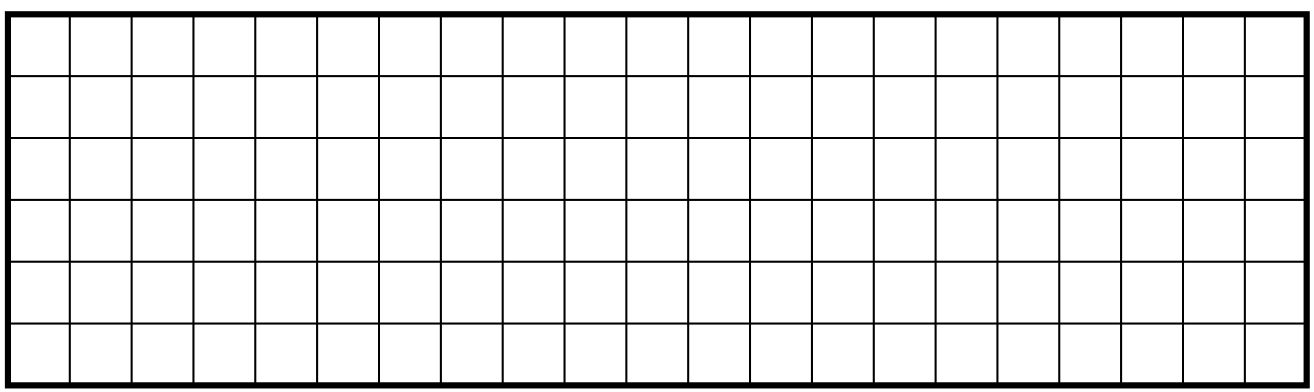

Marshall Bern and David Eppstein, 2000.

- n-gons have $O(n)$ quad-mesh with angles $\leq 120^\circ$.
- $O(n \log n)$ work.
- Regular hexagon (and Euler's formula) shows 120° is sharp.

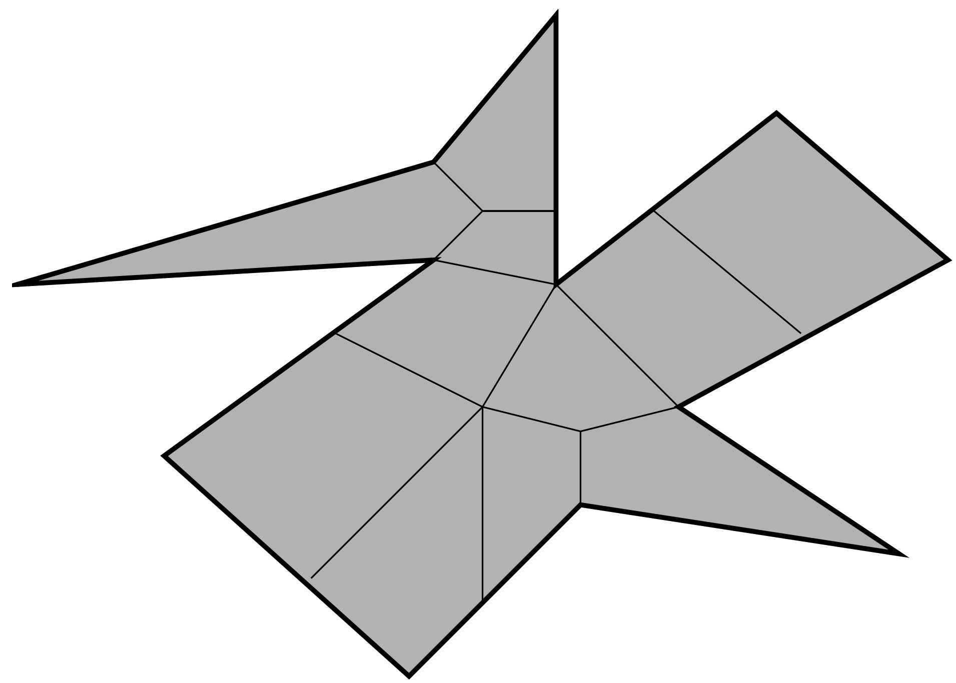

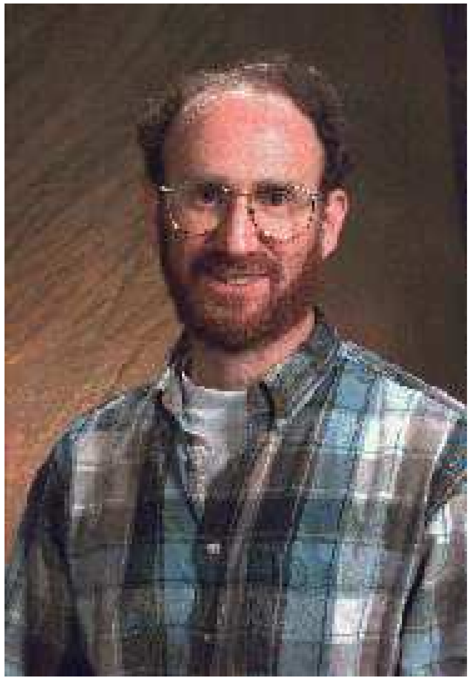

Marshall Bern

David Eppstein

David Epstein

David Eppstein

$P =$ hyperbolic geometry, University of Warwick

$P^2 =$ computational geometry, UC Irvine

Bern asked: can we bound angles from below?

Bern asked: 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.

Bern asked: 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.

Proof uses conformal mapping, plus an idea from hyperbolic manifolds: **thick/thin decompositions**.

Idea of proof: transfer mesh from disk

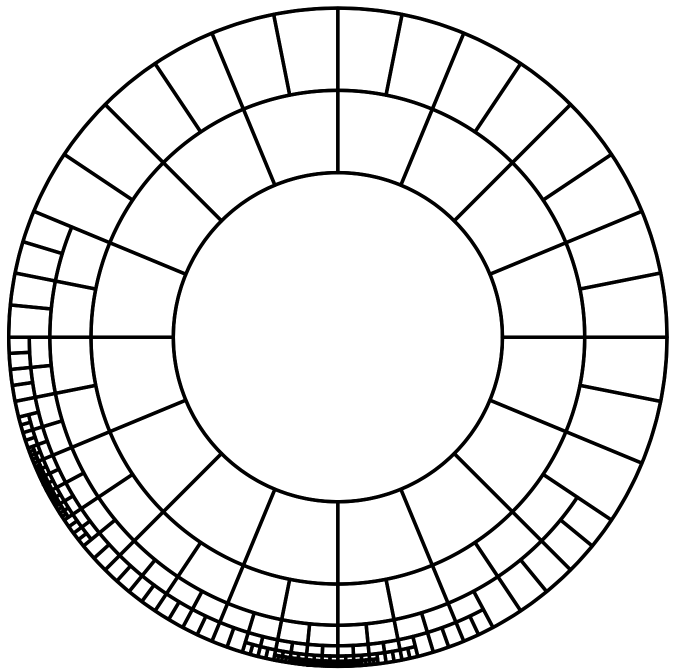

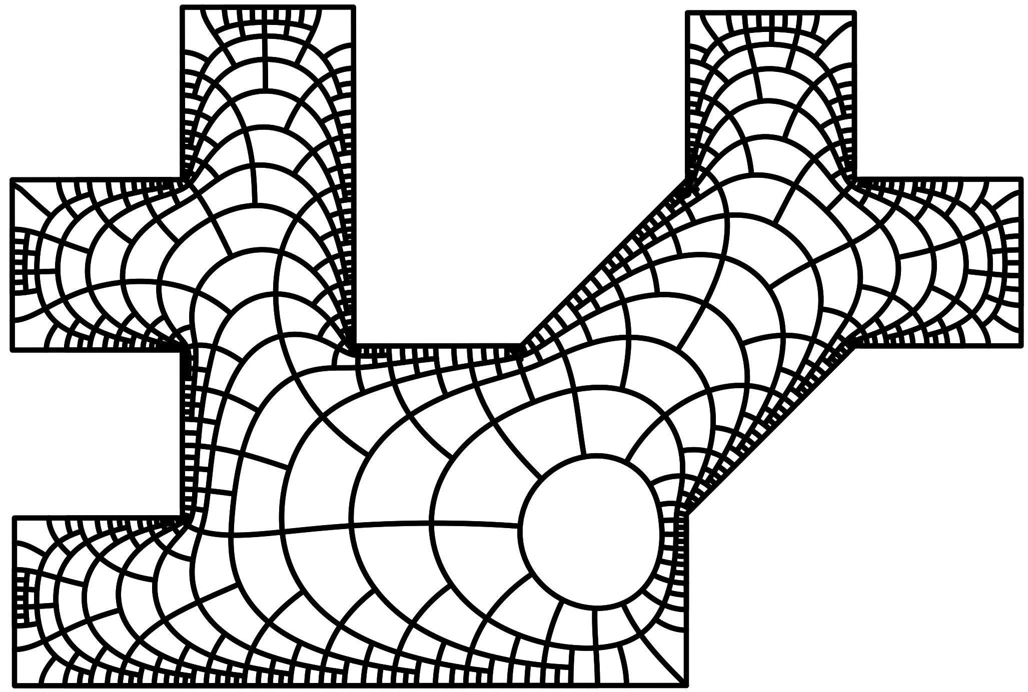

Idea of proof: transfer mesh from disk

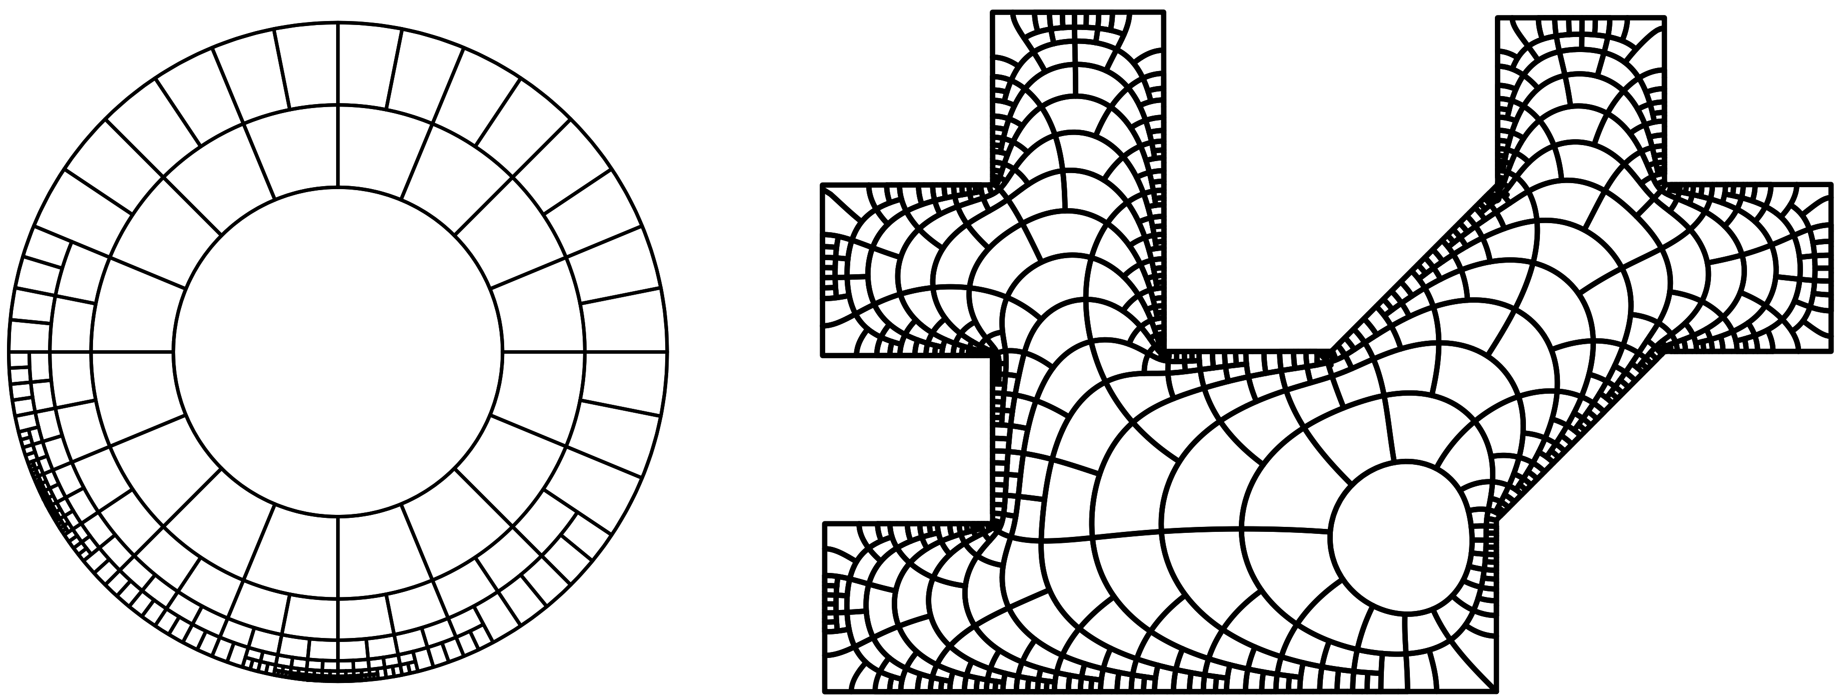

But number of boxes depends on geometry (not just n).

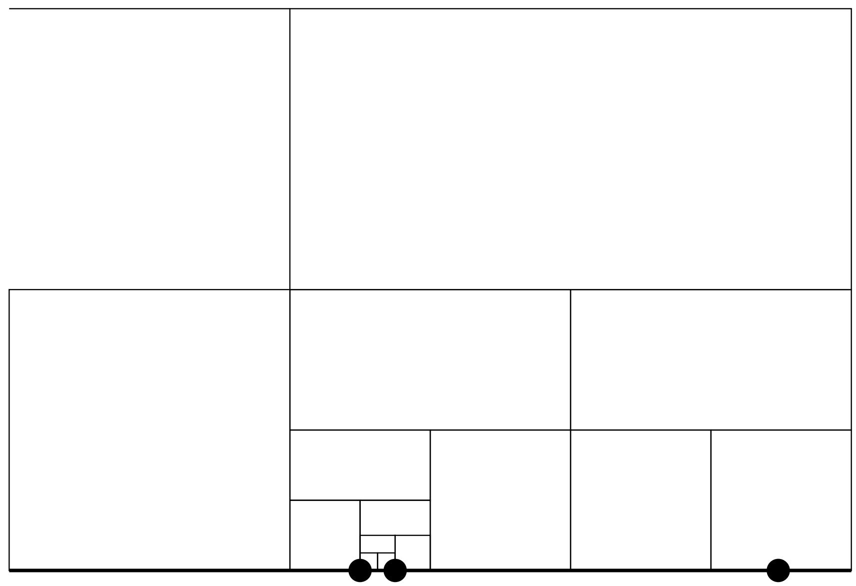

Obvious idea: subdivide disk into boxes until preimages are separated.

Gives too many boxes ($\gg n$) if pre-vertices cluster tightly.

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

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

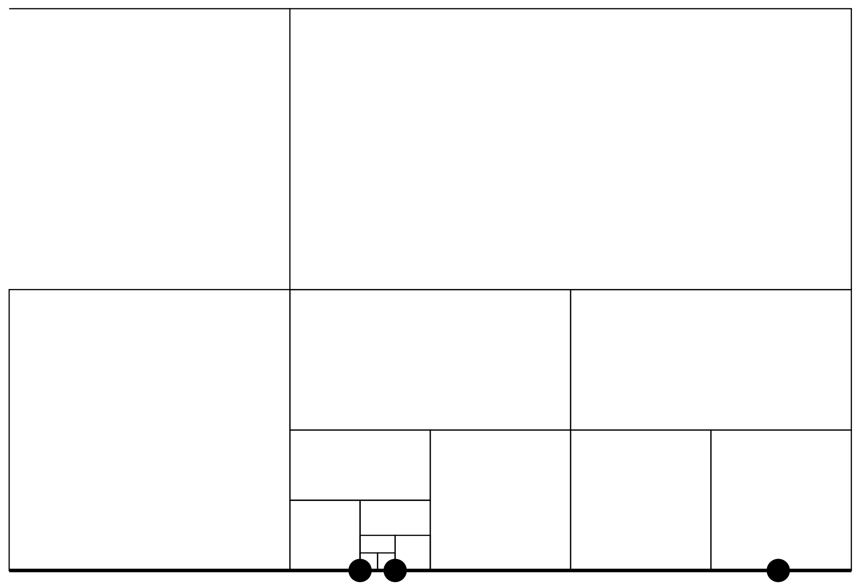

Subdivide until boxes separate points.

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.

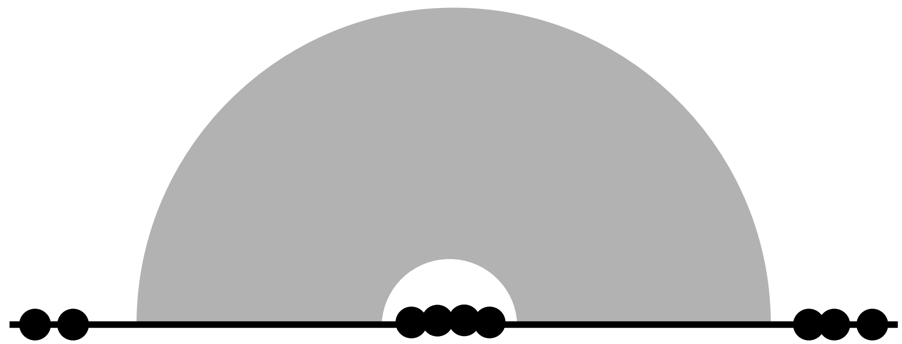

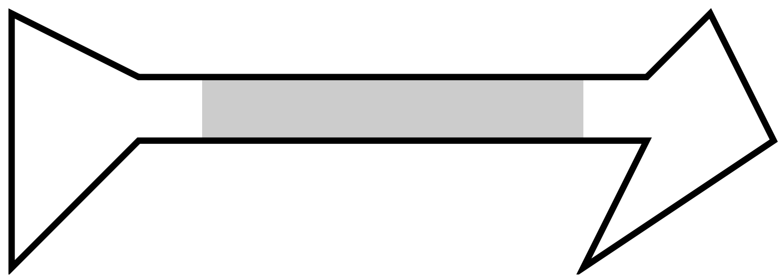

Maps to a **hyperbolic thin piece** dividing polygon.

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

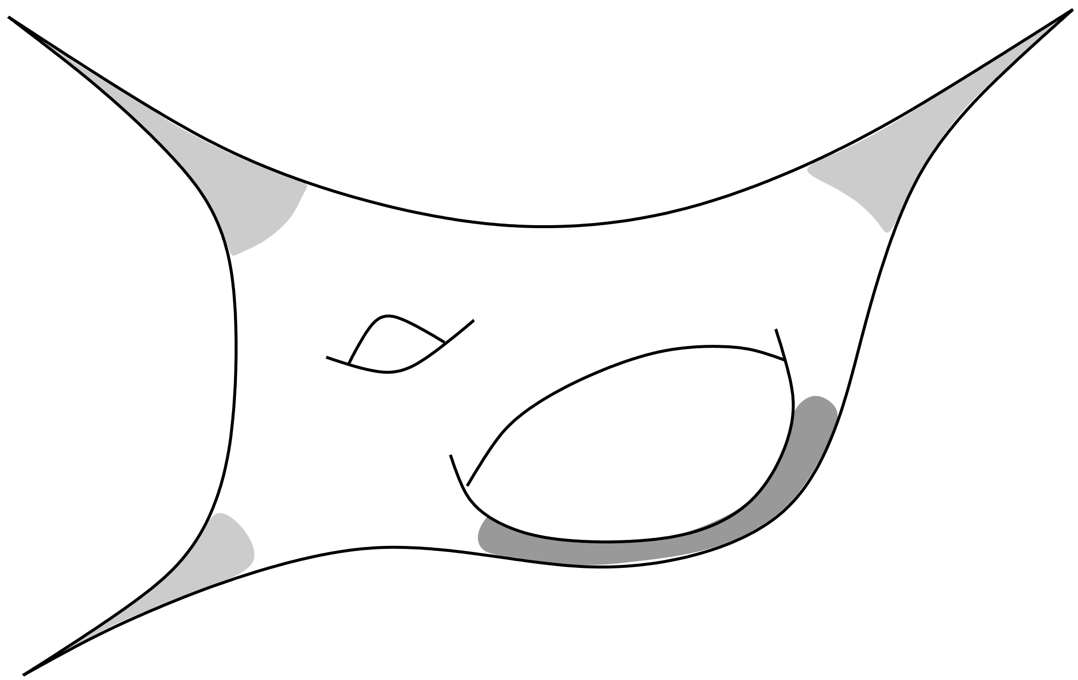

parabolic = puncture,   hyperbolic = handle

Thick and Thin parts of a polygon

Thin part is union of short curves between edges.

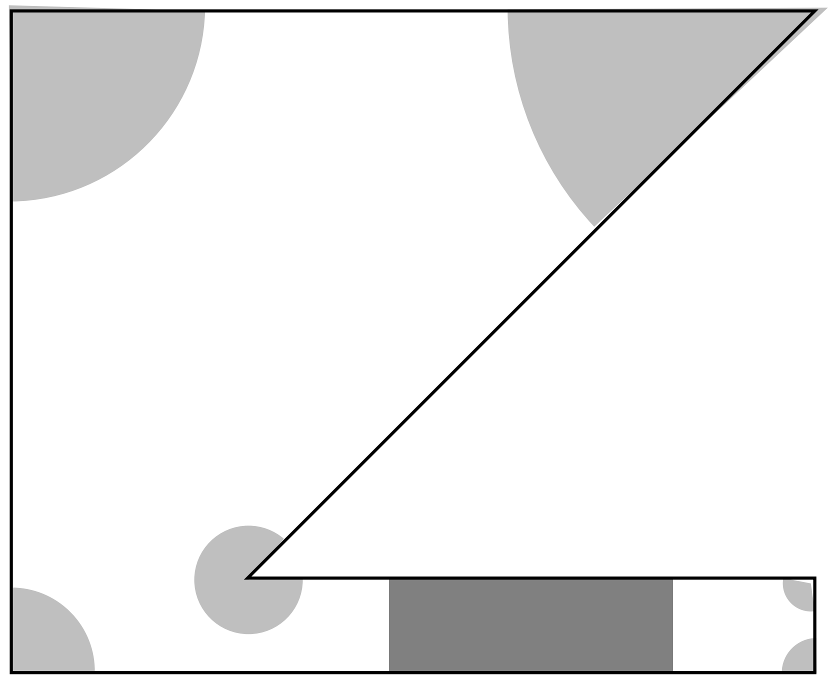

Parabolic = adjacent,   Hyperbolic = non-adjacent

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

Thick parts = remaining components (white)

More examples of hyperbolic thin parts.

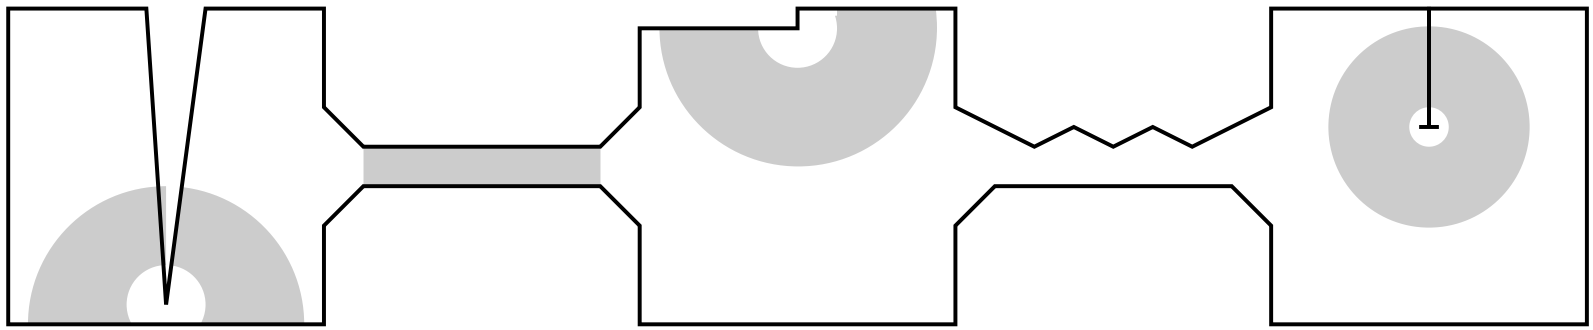

Inside thick regions (white) conformal pre-vertices are well separated on circle (no clusters).

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".

- Conformally map mesh on disk to thick parts.

Thin parts are meshed by explicit construction (easy).

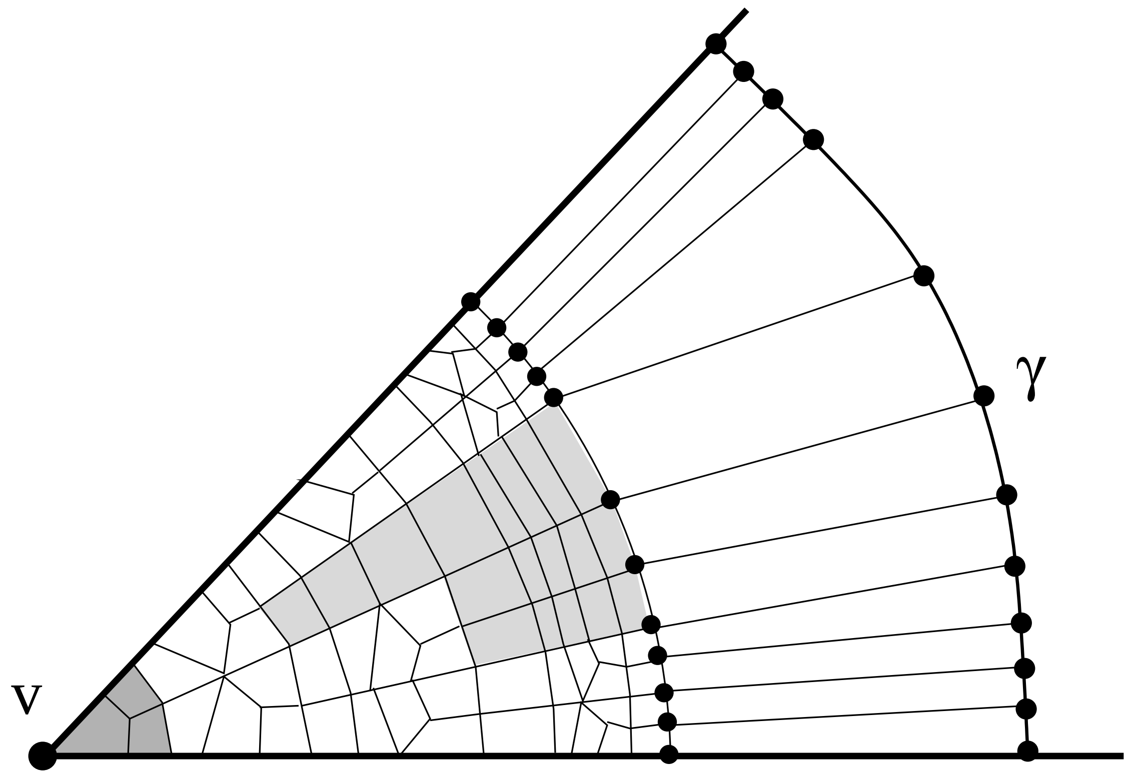

Thick parts: transfer mesh from disk

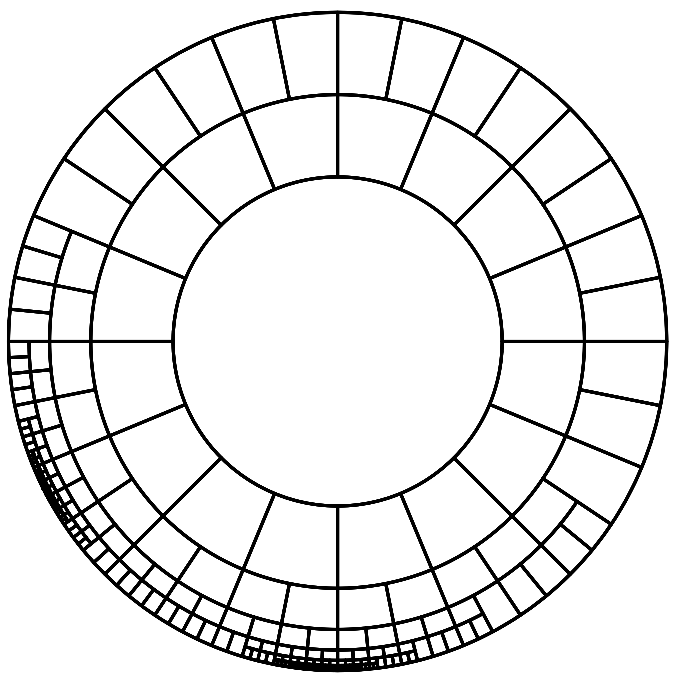

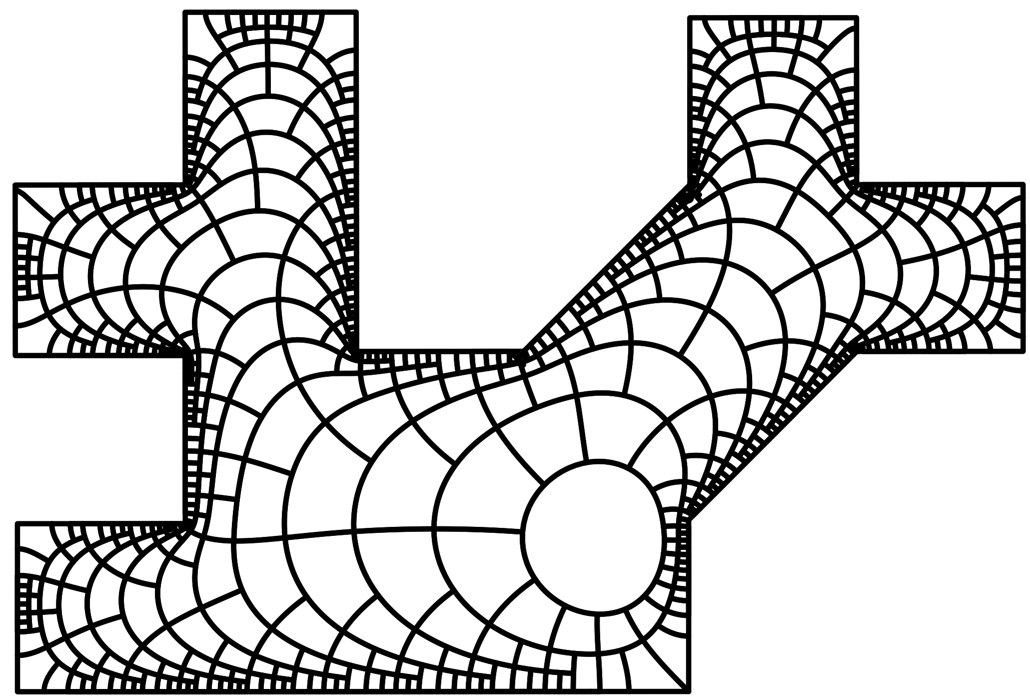

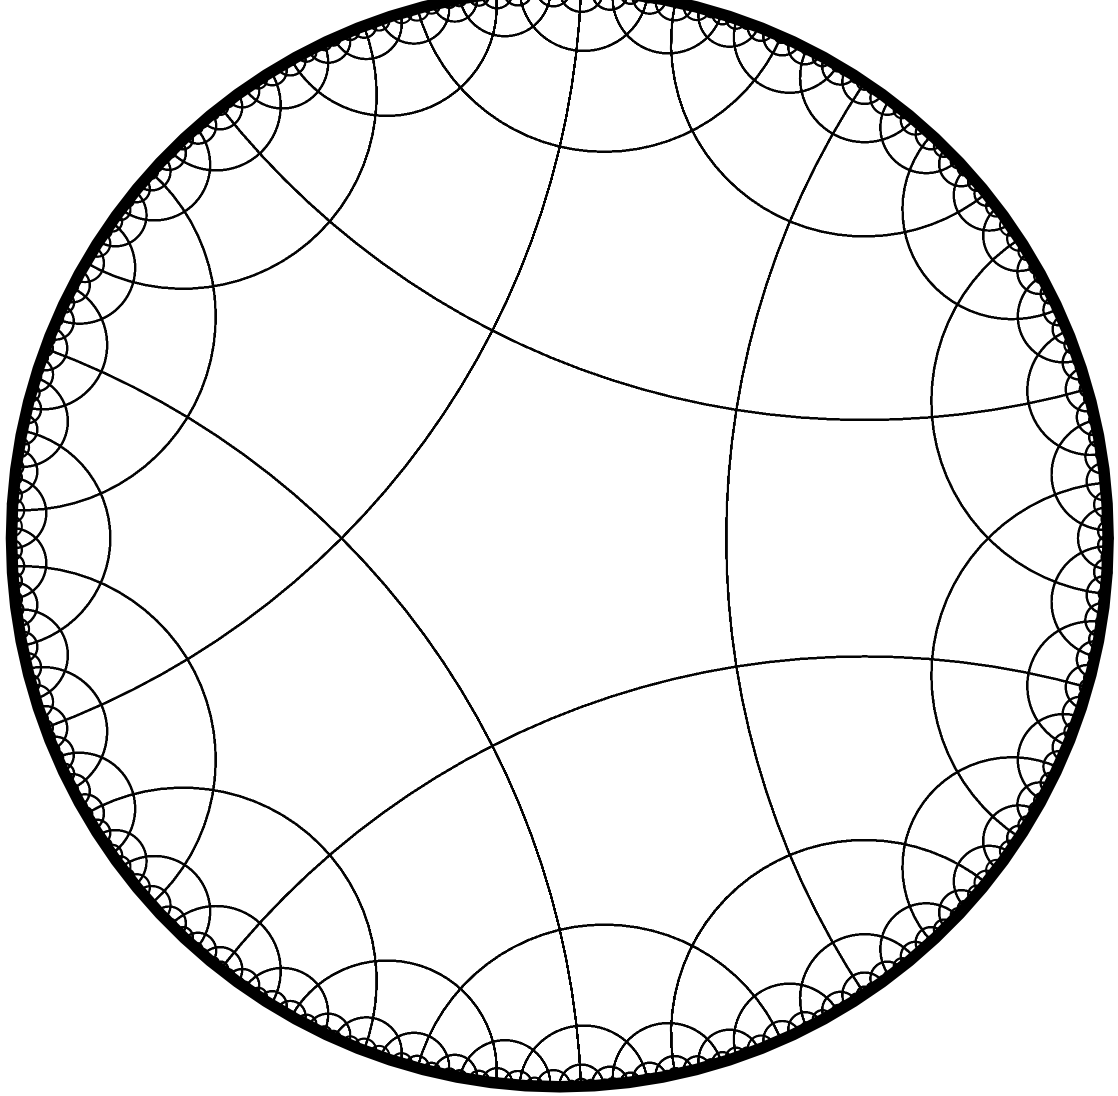

Conformal map from polygon to disk takes thick and thin parts to disk.

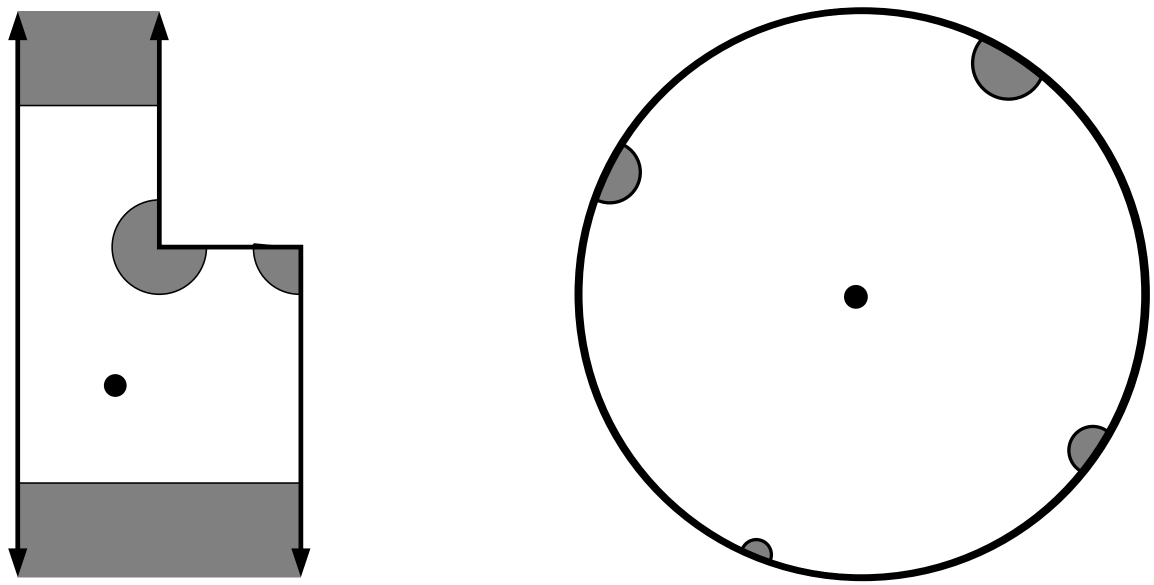

Conformal map from polygon to disk takes thick and thin parts to disk.

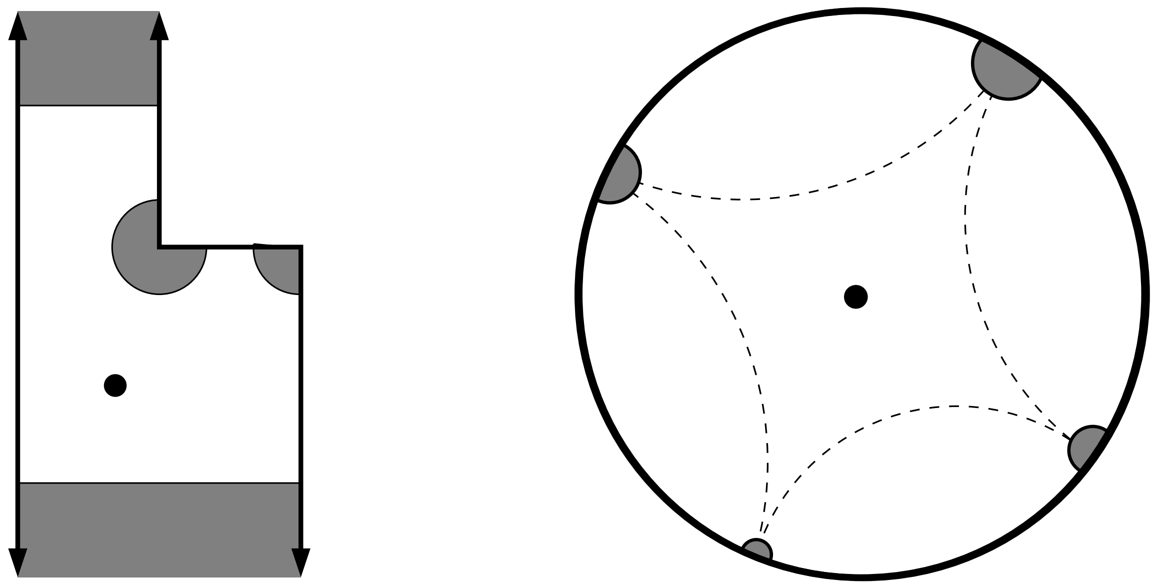

Draw (hyperbolic) convex hull of thin regions.

Conformal map from polygon to disk takes thick and thin parts to disk.

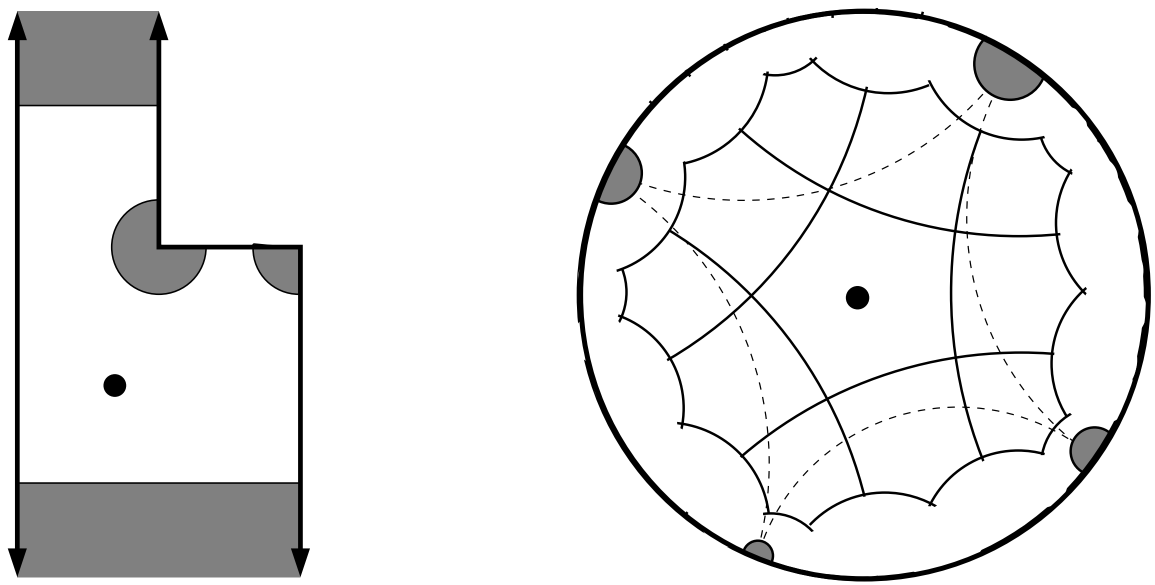

Draw (hyperbolic) convex hull of thin regions.

Take pentagons from tesselation hitting convex hull but missing thin parts.

Conformal map from polygon to disk takes thick and thin parts to disk.

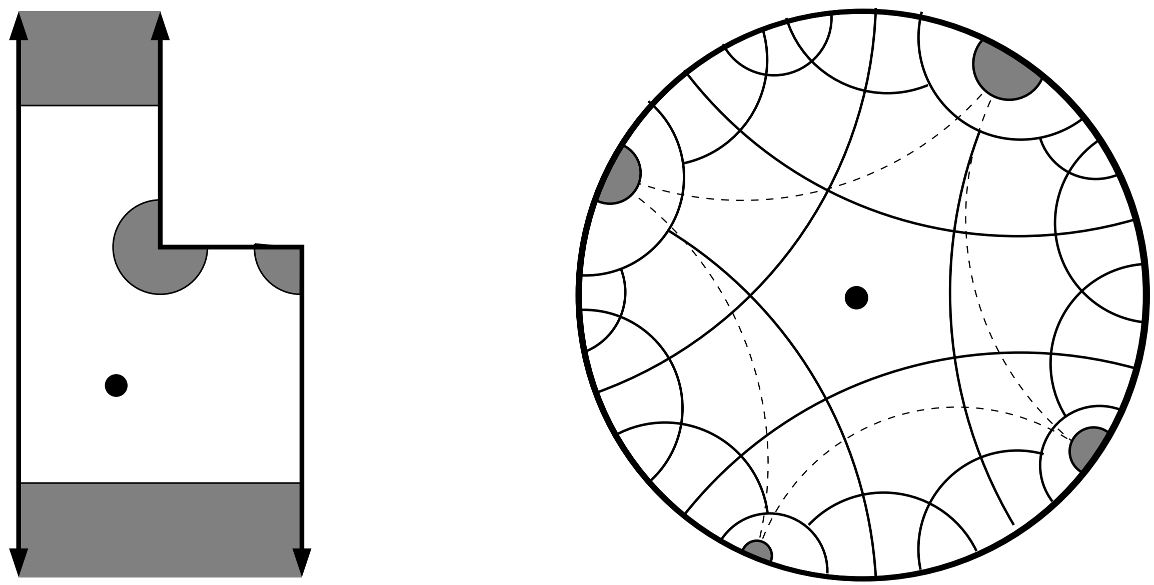

Draw (hyperbolic) convex hull of thin regions.

Take pentagons from tesselation hitting convex hull but missing thin parts.
Extend pentagon edges to boundary.

Conformal map from polygon to disk takes thick and thin parts to disk.

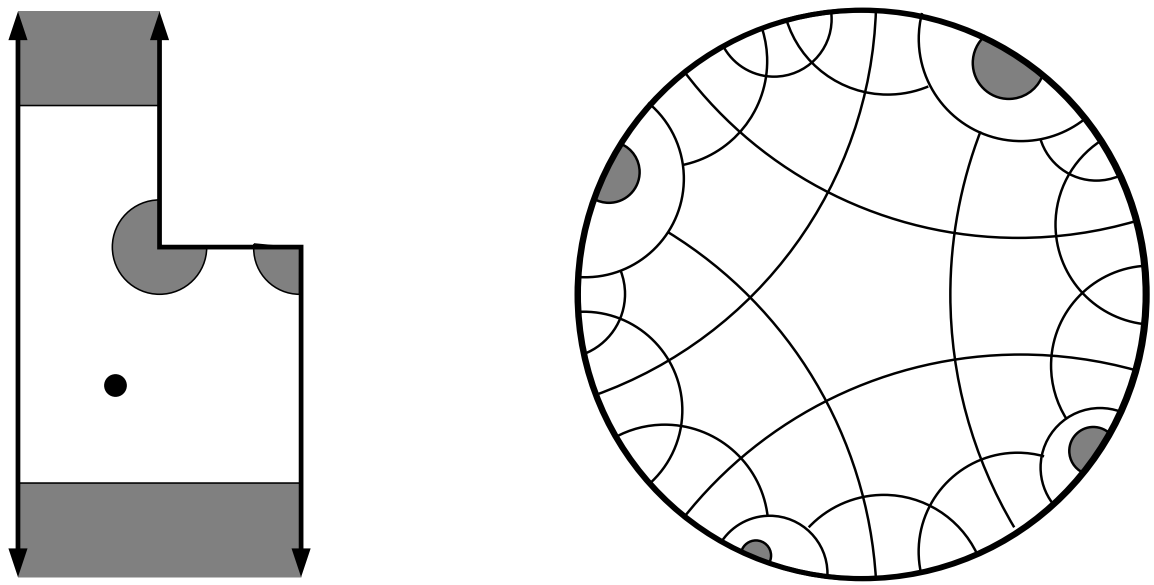

Draw (hyperbolic) convex hull of thin regions.

Take pentagons from tesselation hitting convex hull but missing thin parts.
Extend pentagon edges to boundary.

Analog of Whitney or quadtreee construction.

Conformal map from polygon to disk takes thick and thin parts to disk.

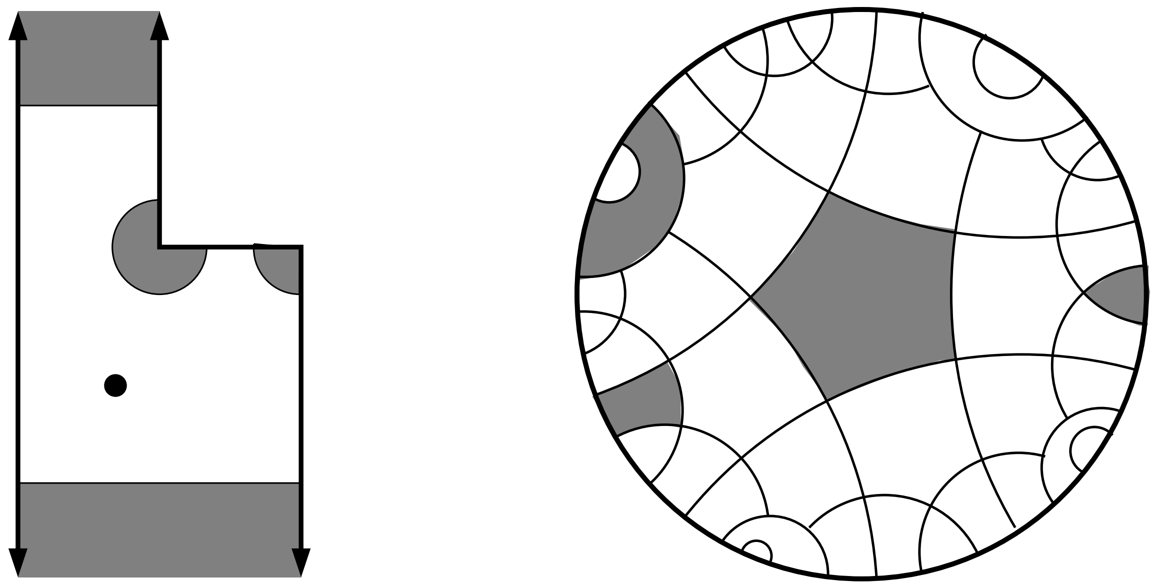

Draw (hyperbolic) convex hull of thin regions.

Take pentagons from tesselation hitting convex hull but missing thin parts.
Extend pentagon edges to boundary.

Pentagons, quadrilaterals, triangles and half-annuli.

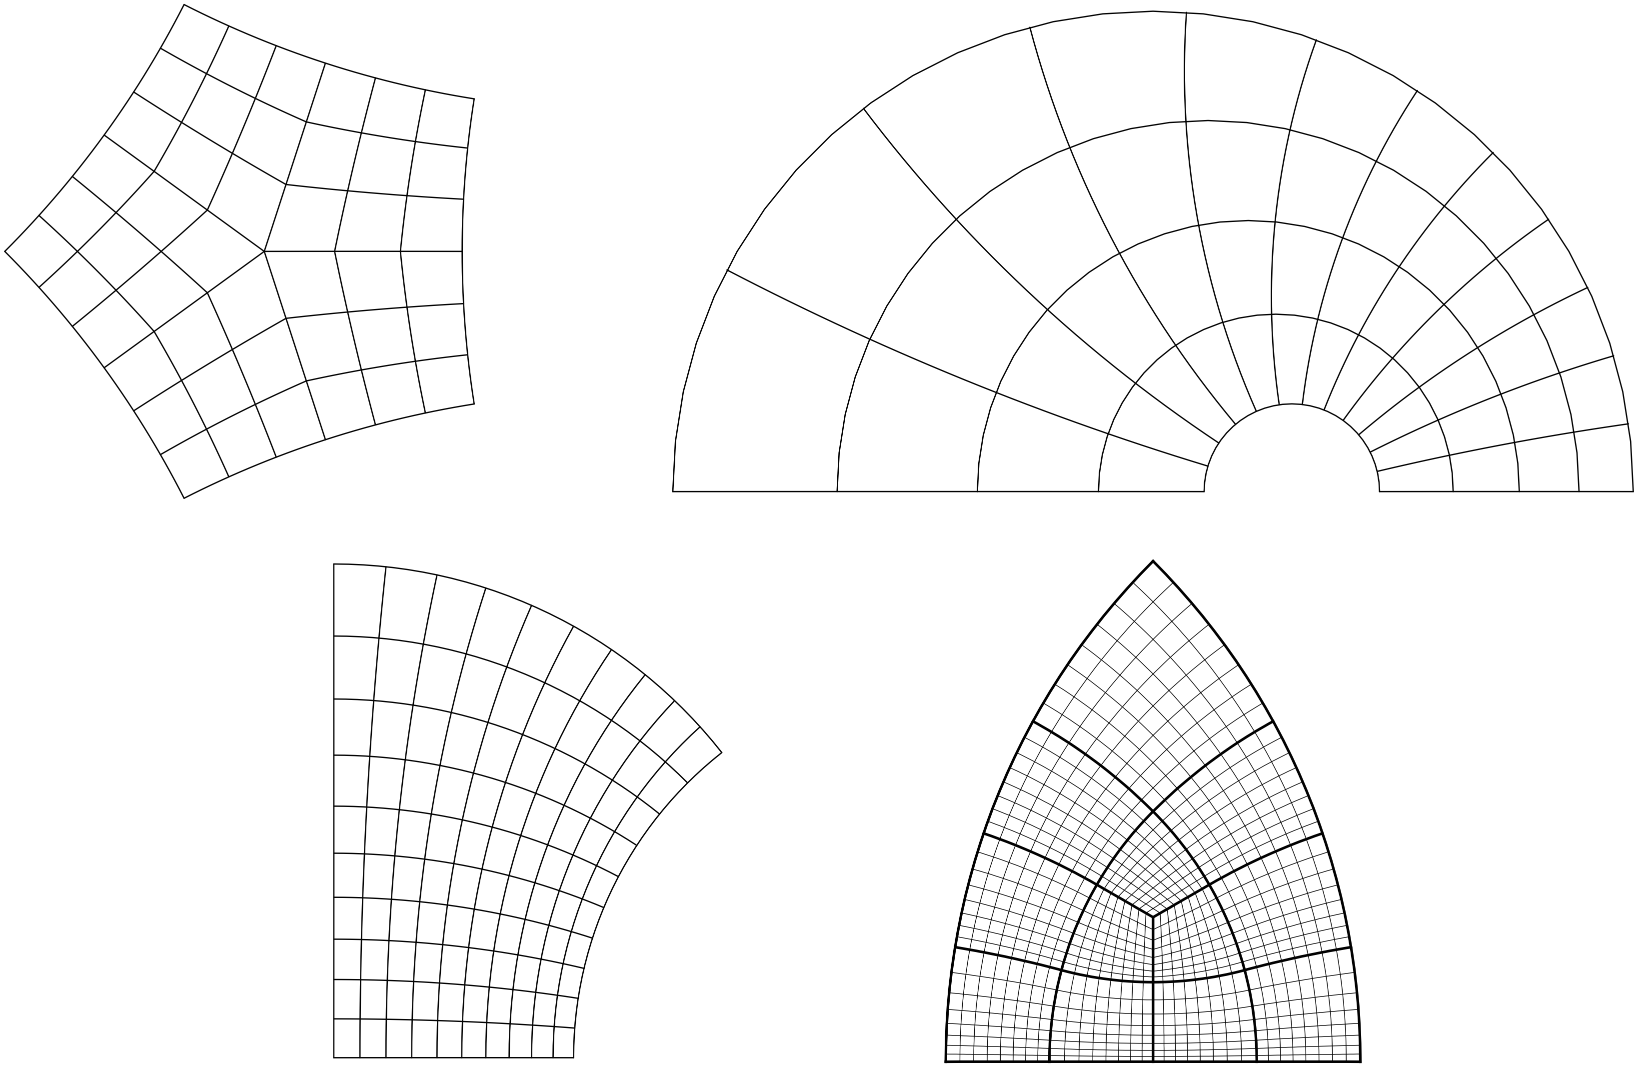

Meshes designed to match along common edges.

So far, we have been meshing simple polygons.

Harder case: mesh PSLGs.

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

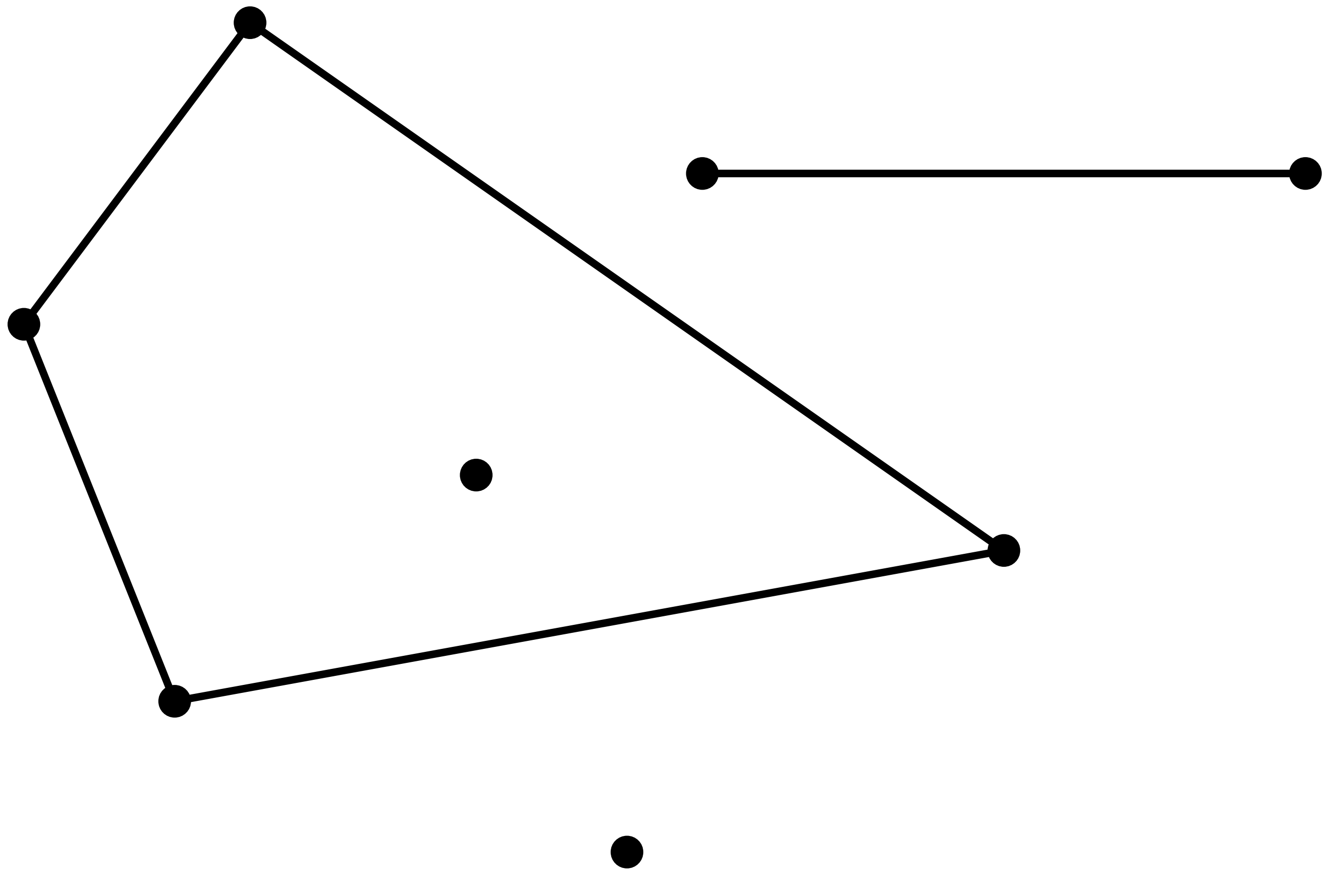

A **Planar Straight Line Graph** (PSLG) is a finite point set plus a set of disjoint edges between them.

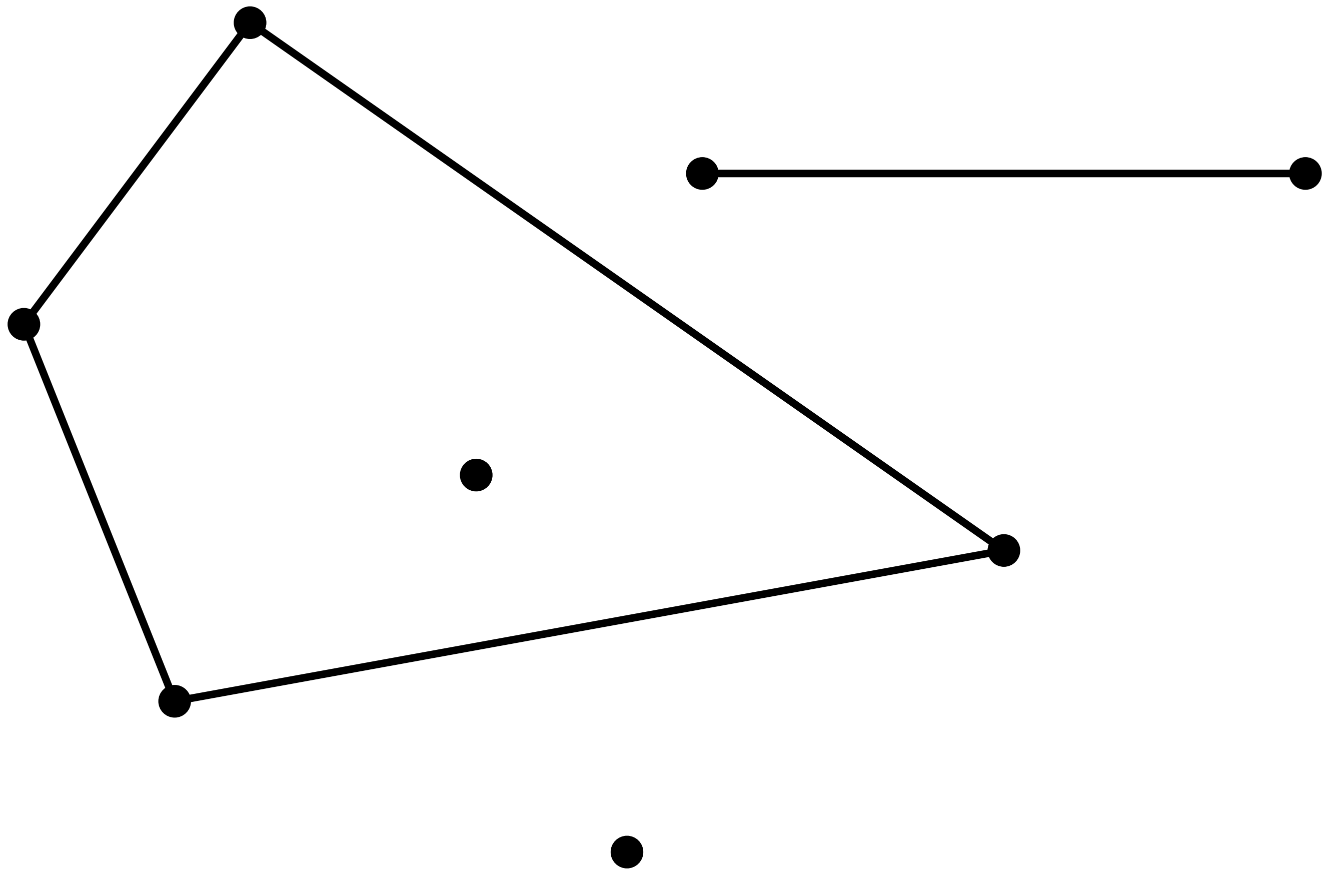

Size = number of vertices = n.

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

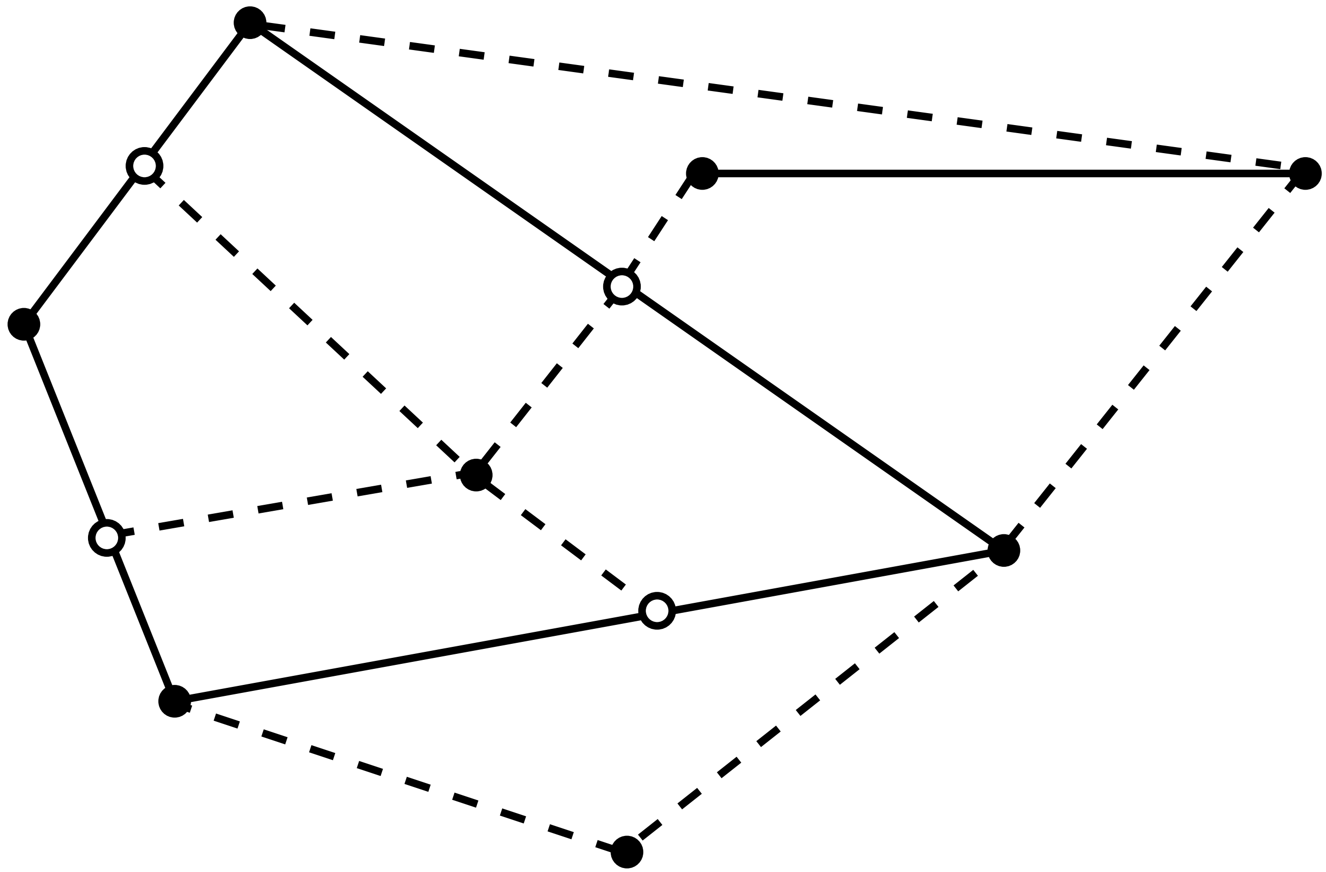

Mesh convex hull conforming to PSLG.

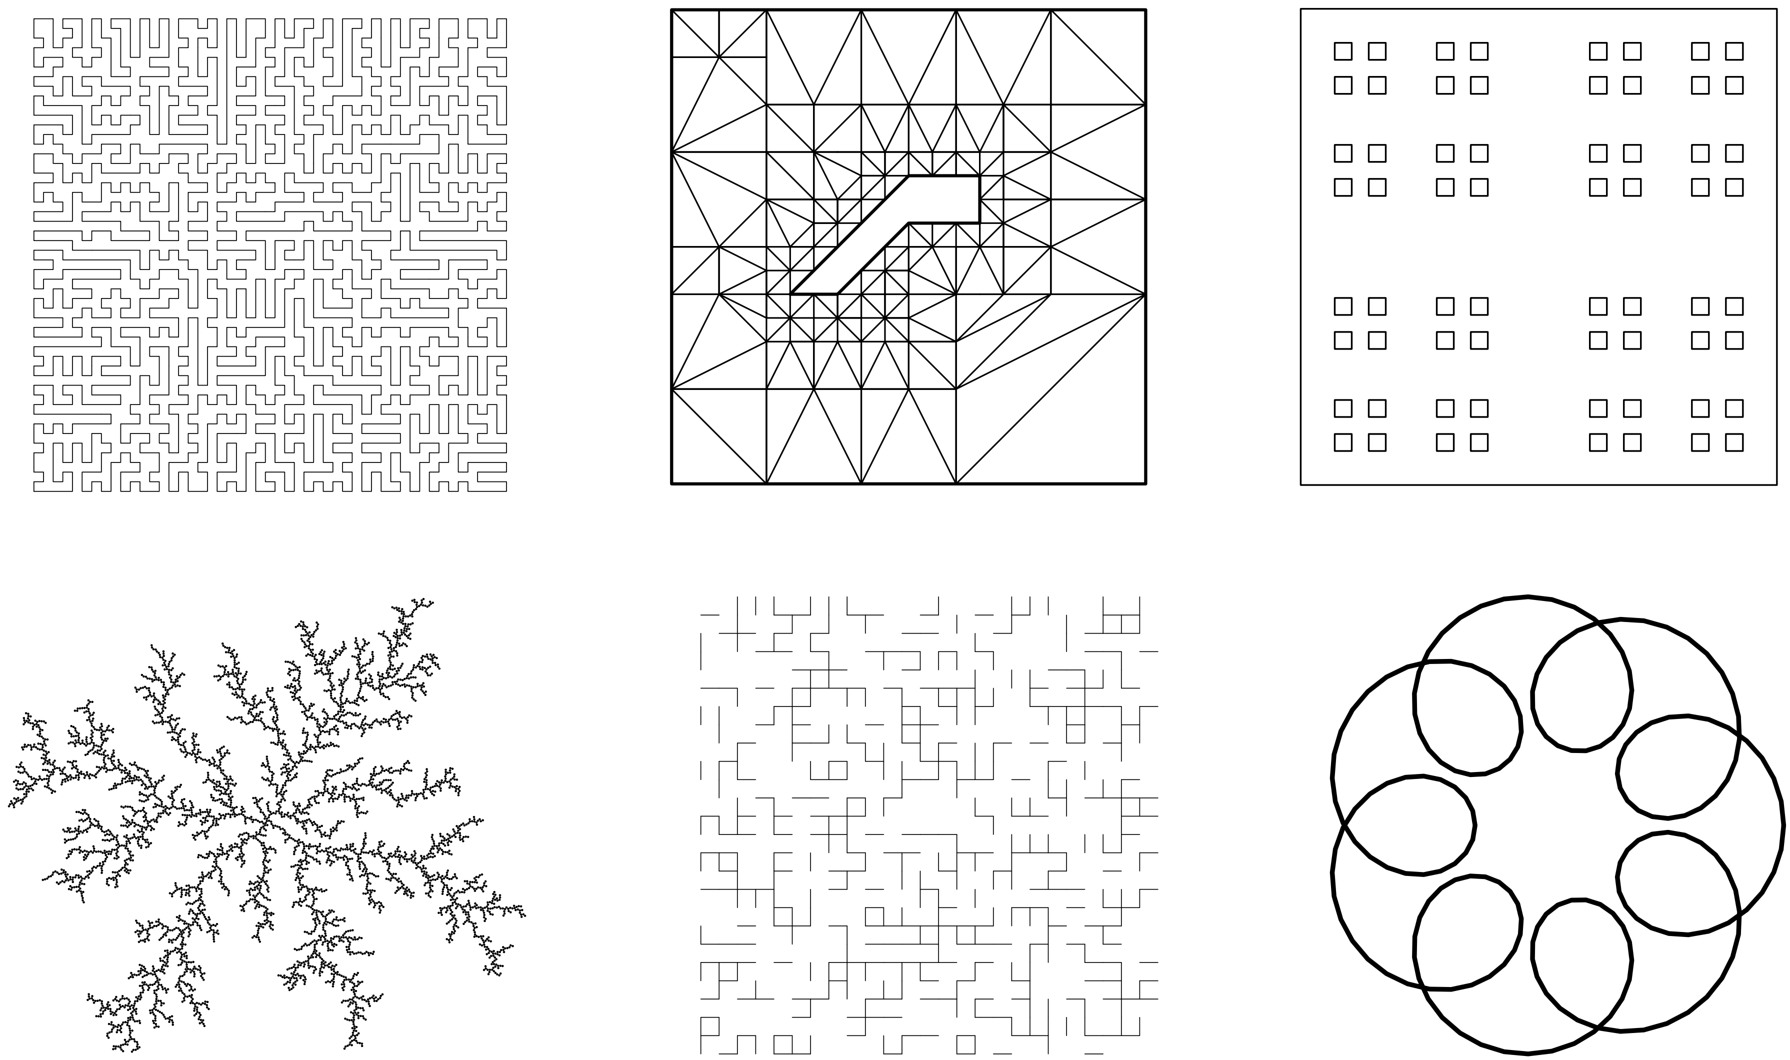

More PSLGs

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

Angles and complexity sharp.

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.

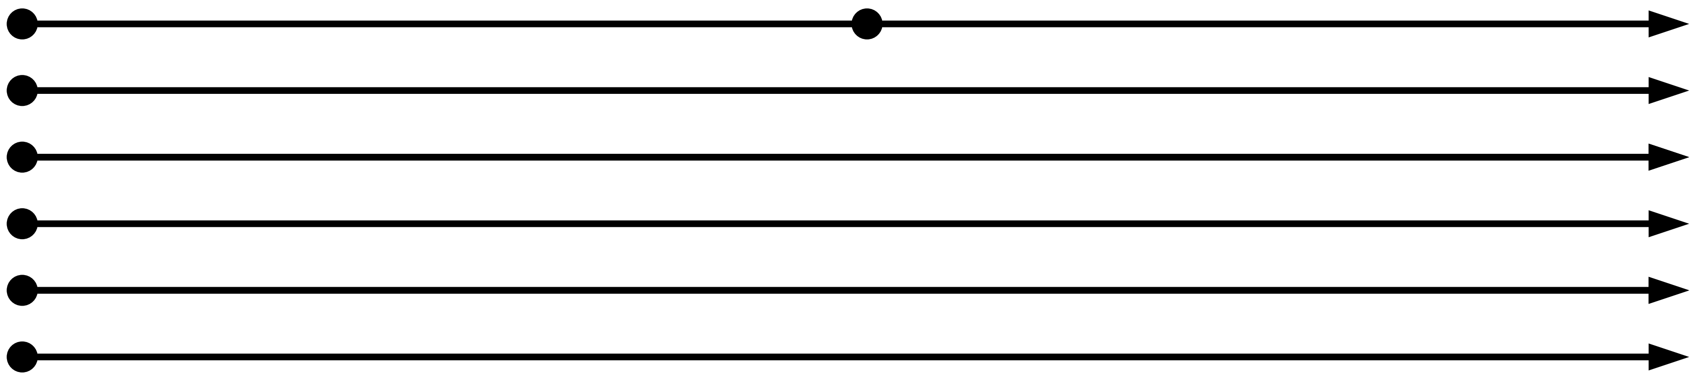

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

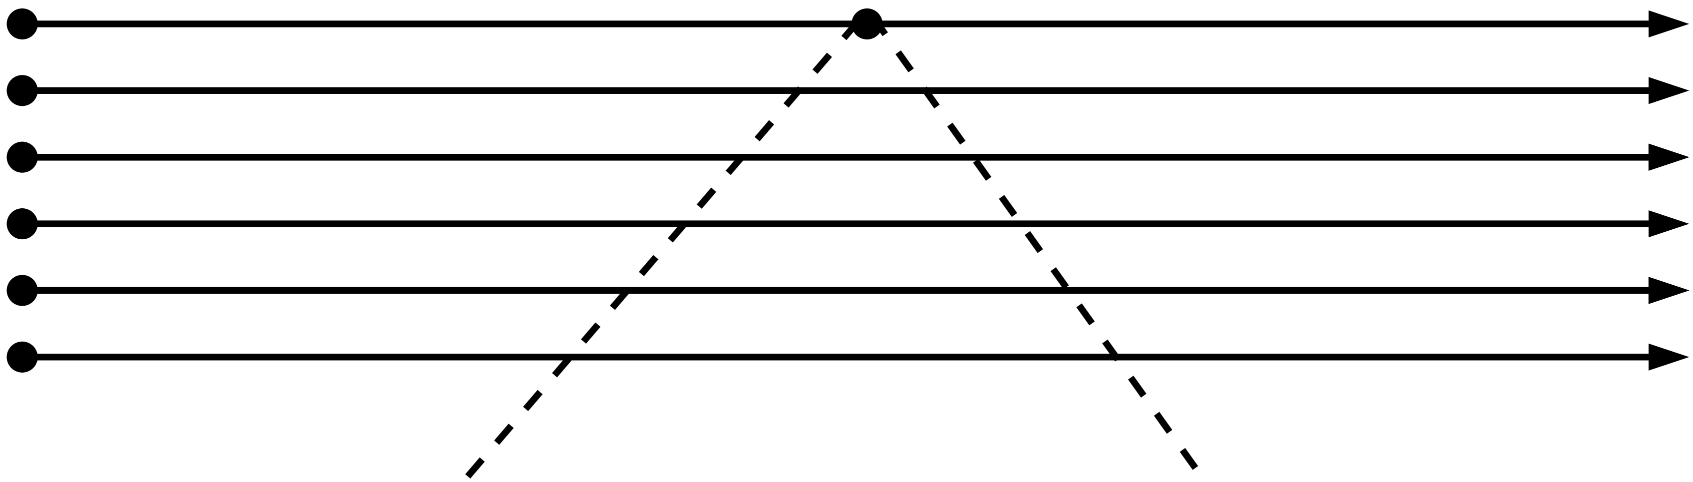

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

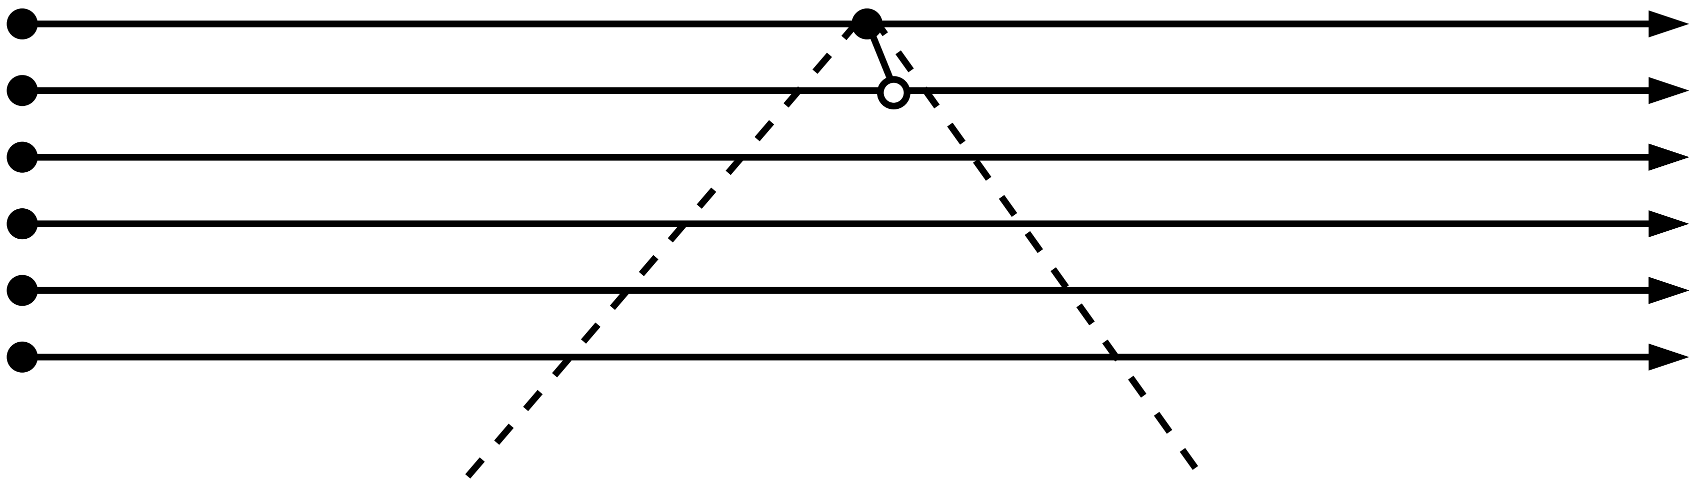

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

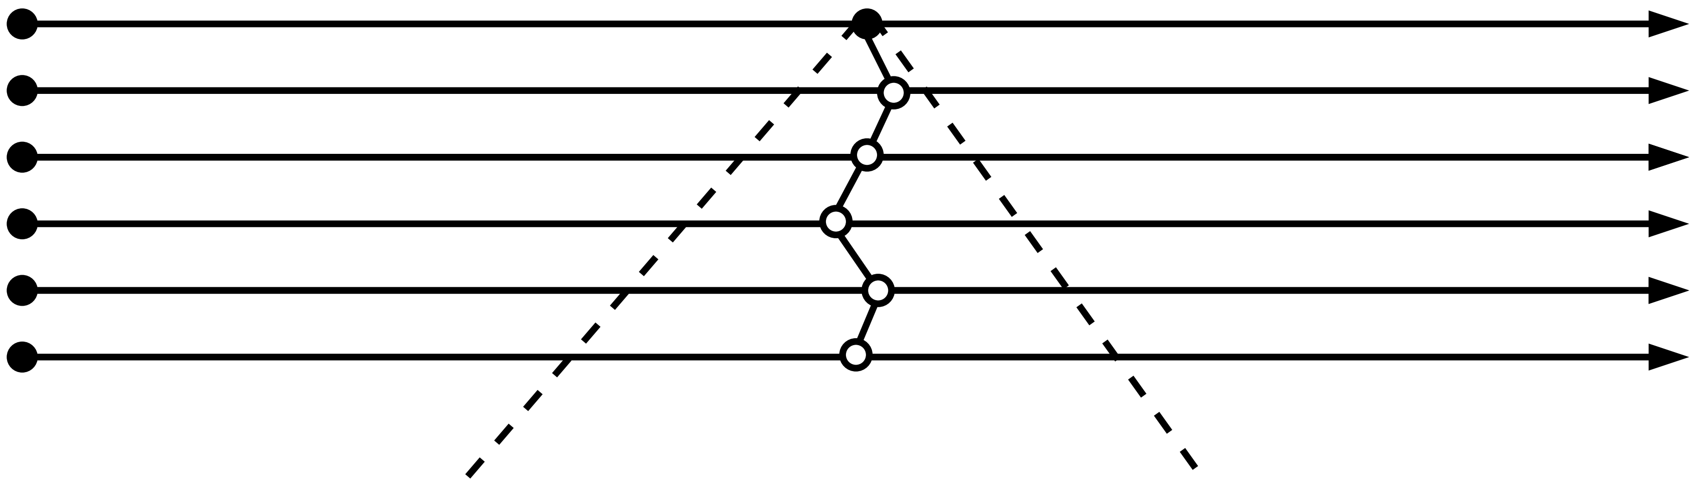

The $O(n^2)$ is sharp because any mesh with maximum angle $\leq \theta < 180^\circ$ sometimes requires n^2 elements.

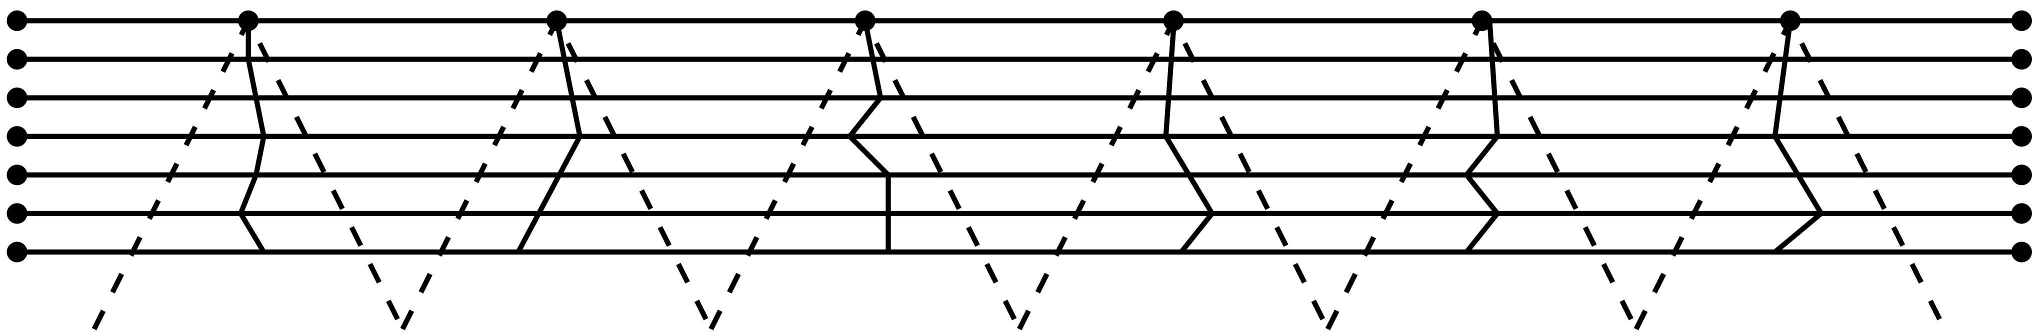

Theorem: Every PSLG of size n has $O(n^2)$ quad-mesh with all angles $\leq 120°$ and all new angles $\geq 60°$.

I won't discuss proof of this here. Uses results several ideas:

- Thick/thin decomposition

- Sinks = a polygon that can be quad-meshed with angles between 60° and 120° using a given (even) set of boundary vertices, without adding any extra boundary vertices.
  Square is not a sink, regular octagon is.
  For sinks, every "even boundary meshing problem" has a solution.
  Sinks stop propagation of vertices as in previous example.

- optimal triangulations of PSLGs

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

Cor (of proof): All vertex degrees are 4 with $O(n)$ exceptions.

Cor: Mesh is a disjoint union of $O(n)$ combinatorial rectangular grids.

Cor: All angles are in $[89^\circ, 91^\circ]$ with $O(n)$ exceptions.

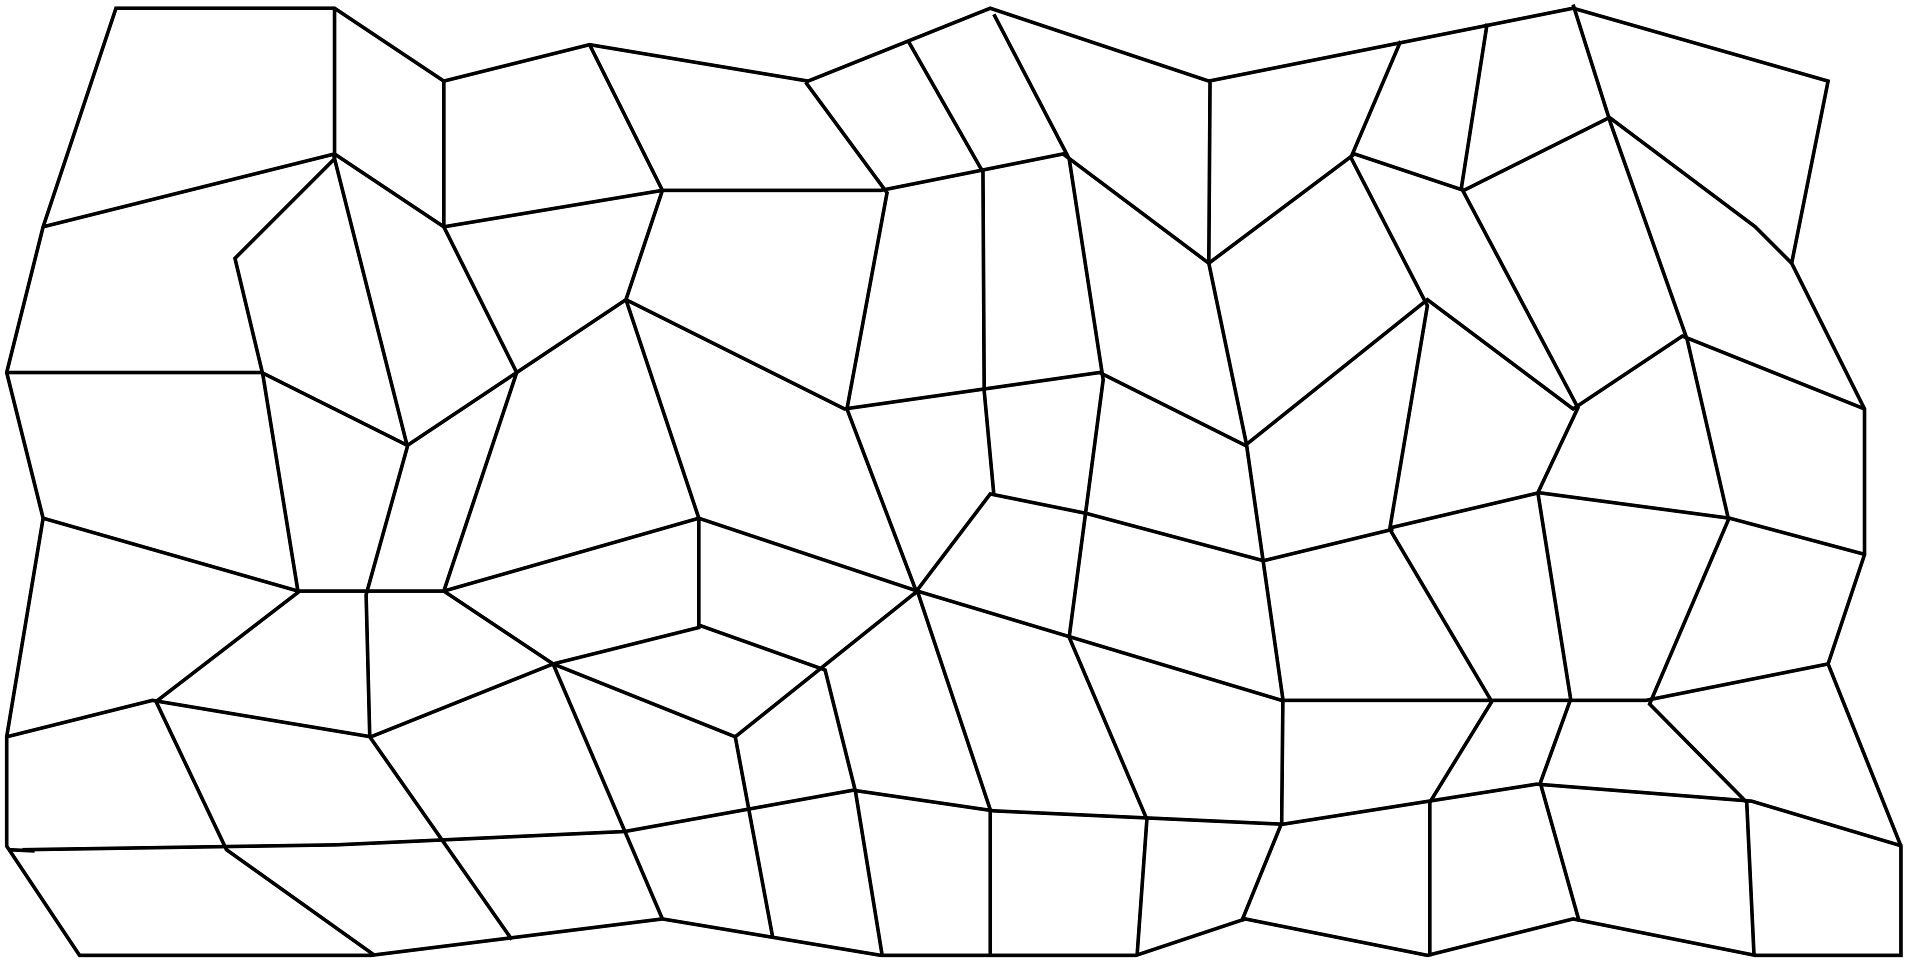

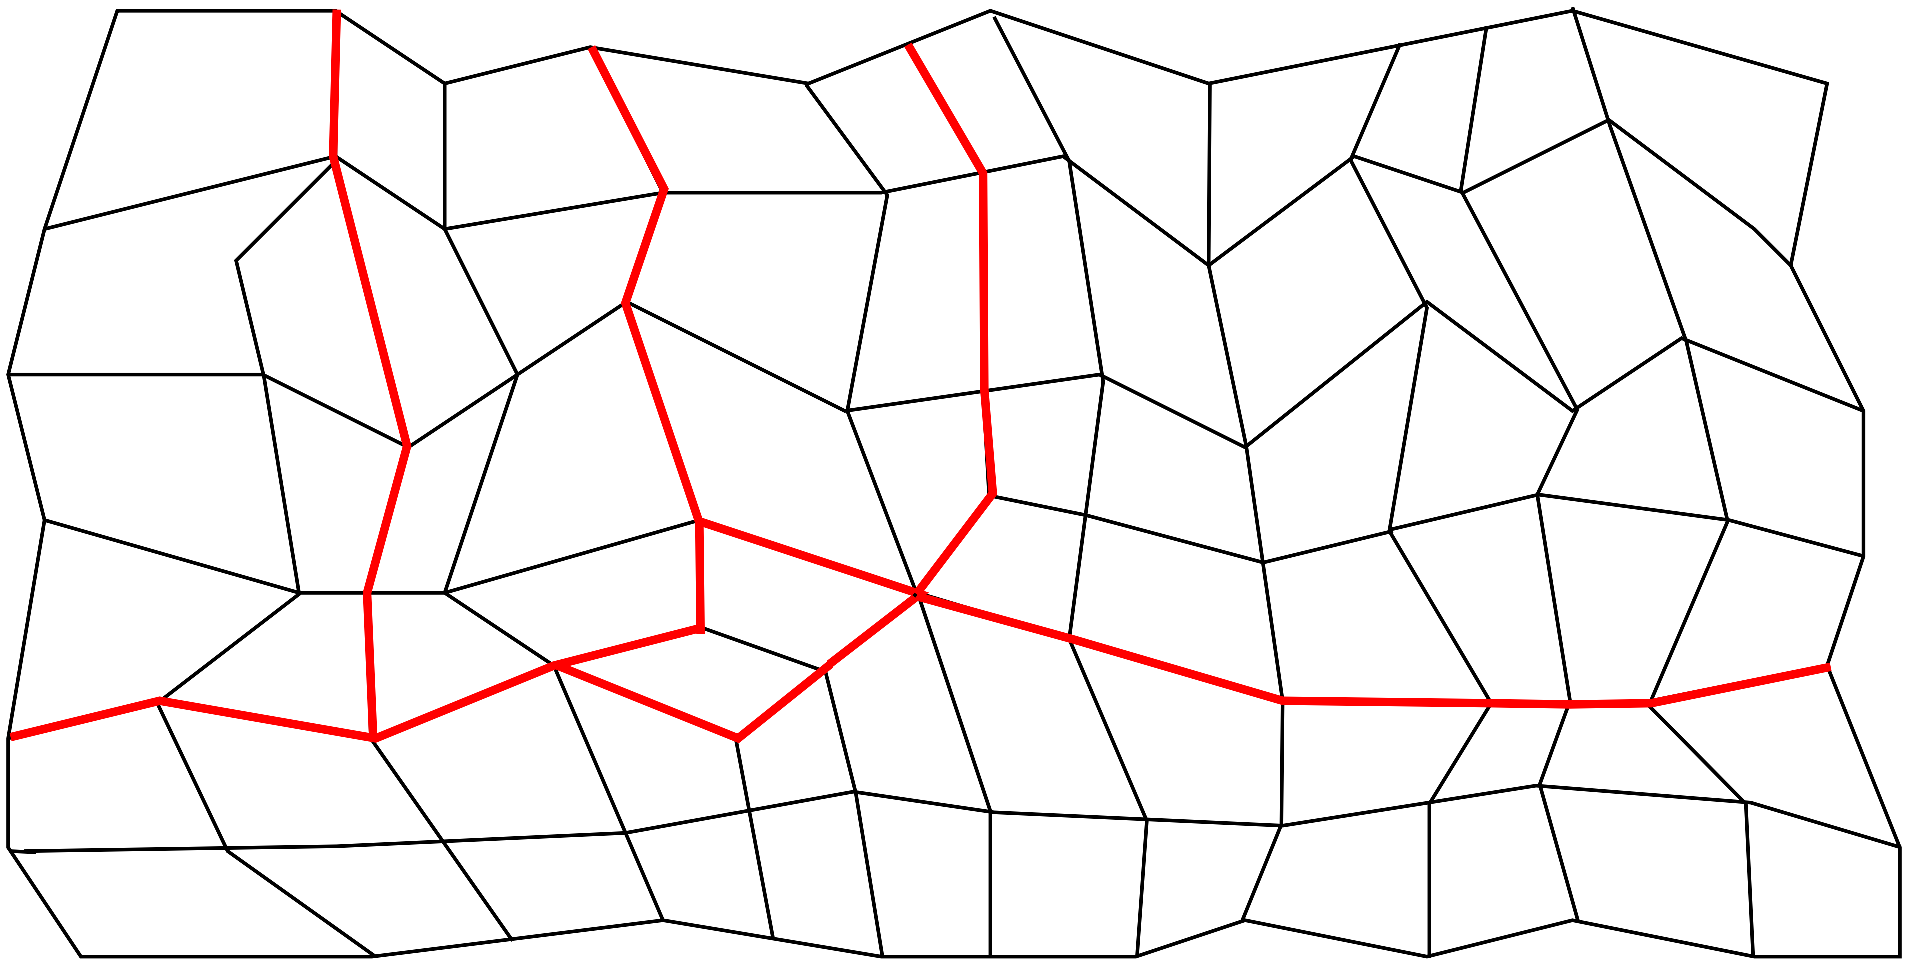

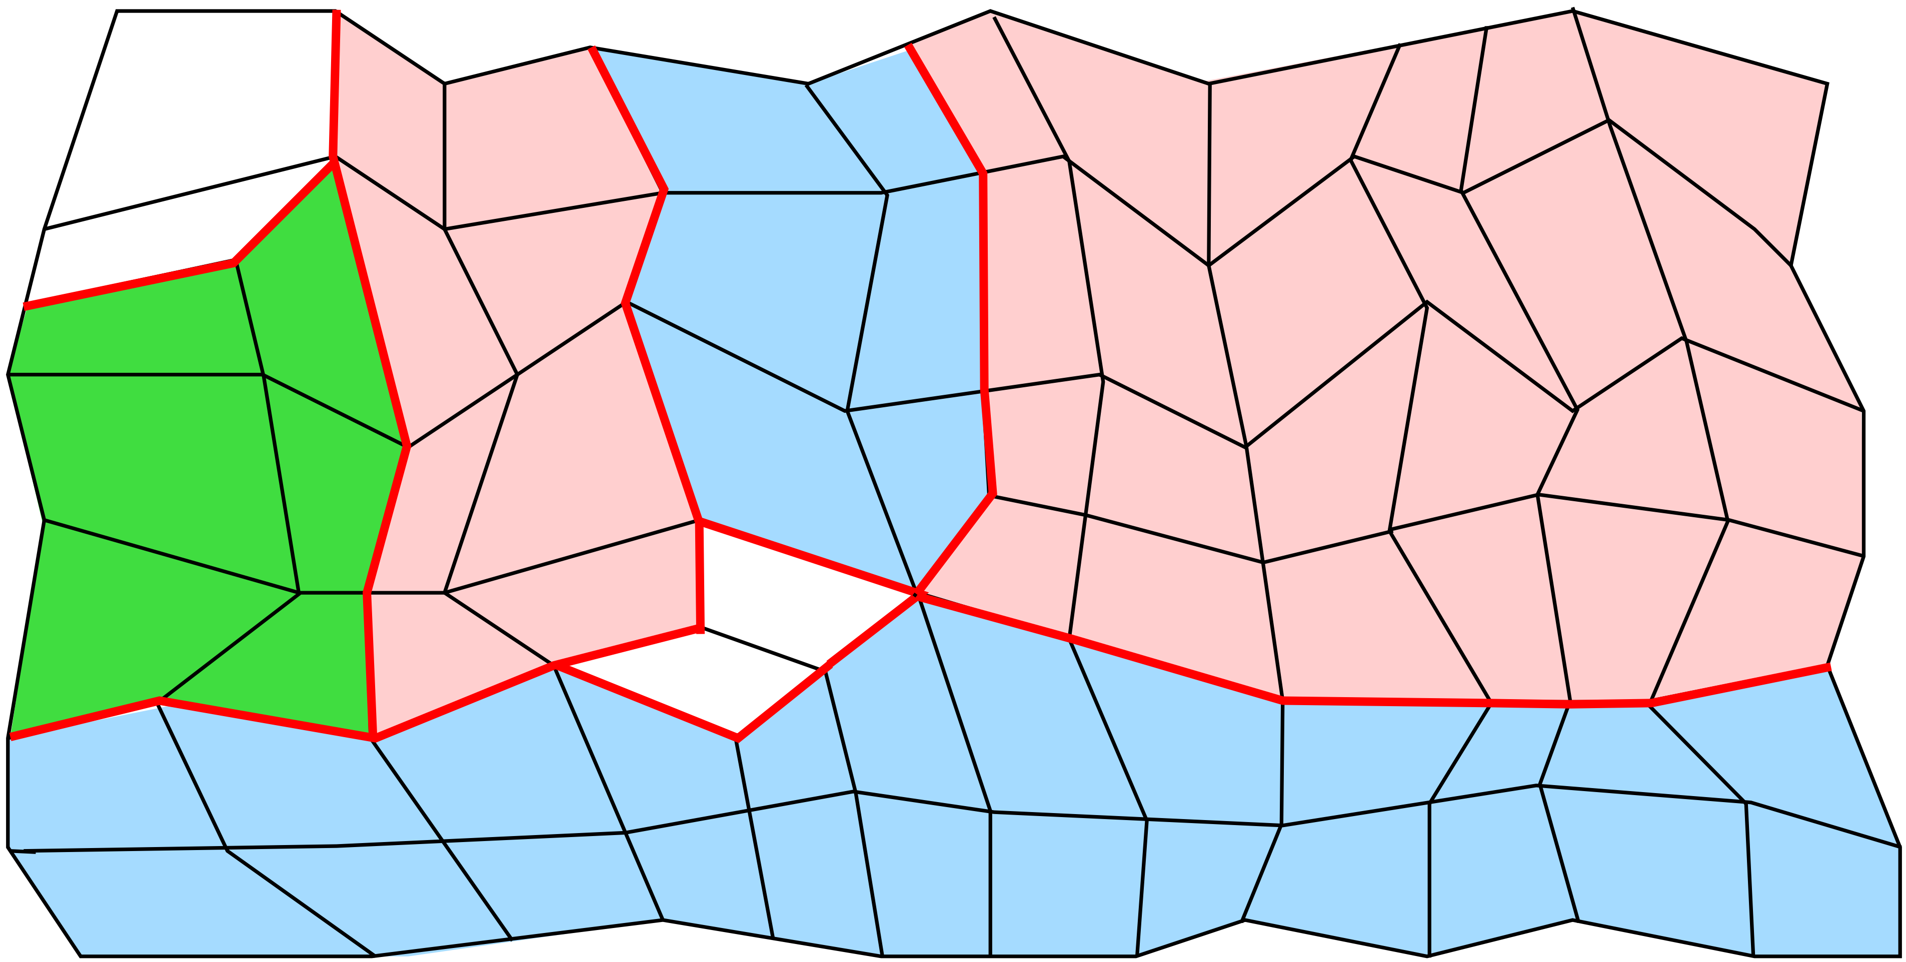

Theorem: Every PSLG of size n has $O(n^2)$ quad-mesh with all angles $\leq 120°$ and all new angles $\geq 60°$.

Every quad-mesh gives a triangulation: add diagonals.

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

Every quad-mesh gives a triangulation: add diagonals.

Corollary: Every PSLG has a $O(n^2)$ triangulation, all angles $\leq 120°$.

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

Every quad-mesh gives a triangulation: add diagonals.

Corollary: Every PSLG has a $O(n^2)$ triangulation, all angles $\leq 120^\circ$.

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?

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

Every quad-mesh gives a triangulation: add diagonals.

Corollary: Every PSLG has a $O(n^2)$ triangulation, all angles $\leq 120^\circ$.

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**

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

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

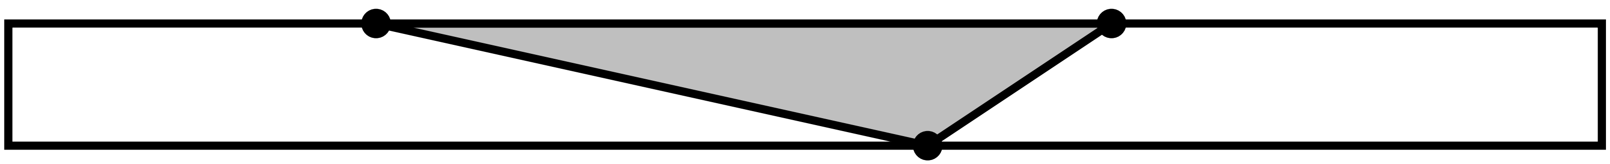

So, bounded complexity $\Rightarrow$ no lower angle bound.

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

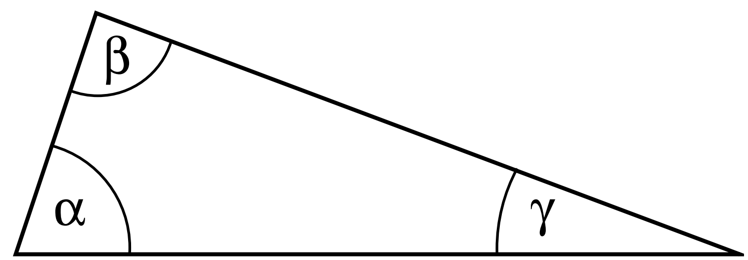

$$\alpha, \beta < (90^\circ - \epsilon) \Rightarrow \gamma = 180^\circ - \alpha - \beta \geq 2\epsilon.$$

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

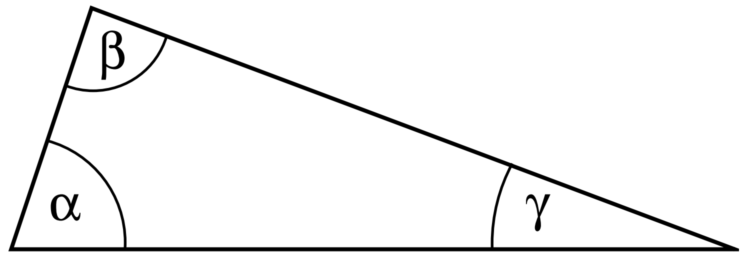

$$\alpha, \beta < (90^\circ - \epsilon) \Rightarrow \gamma = 180^\circ - \alpha - \beta \geq 2\epsilon.$$

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

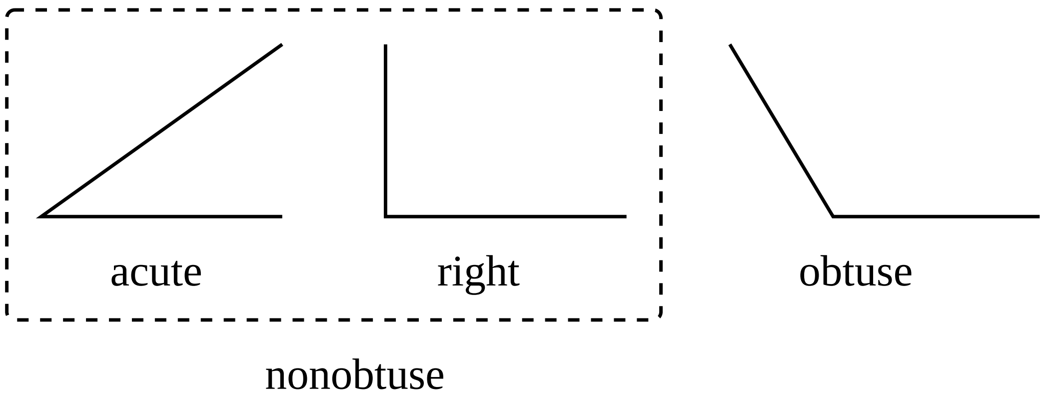

Is NOT="non-obtuse triangulation" possible?

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

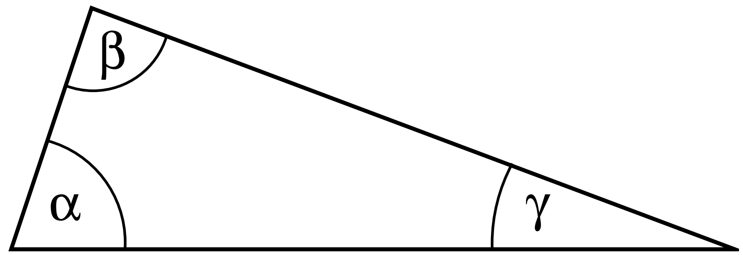

$$\alpha, \beta < (90^\circ - \epsilon) \Rightarrow \gamma = 180^\circ - \alpha - \beta \geq 2\epsilon.$$

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

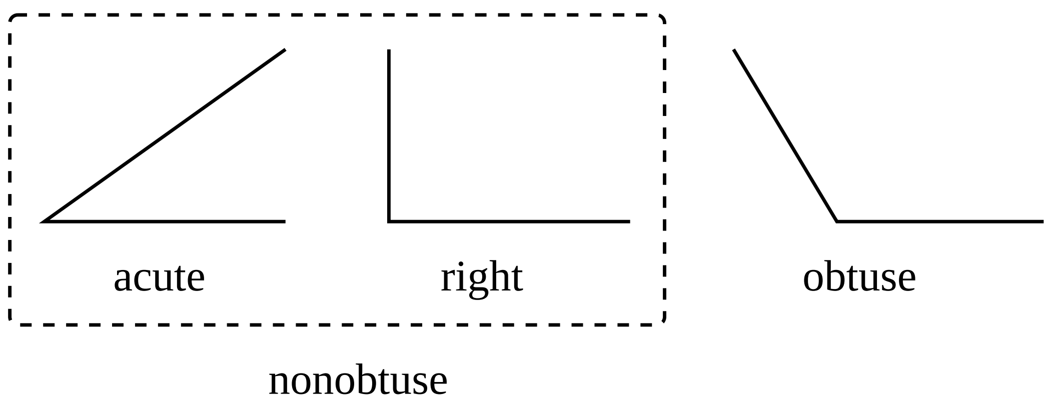

Is NOT="non-obtuse triangulation" possible? **Yes**

Brief history of NOTs:

- Always possible (no bound): 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
- PSLG's: exist, n^2 lower bound

Do a polynomial number of triangles suffice?

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)

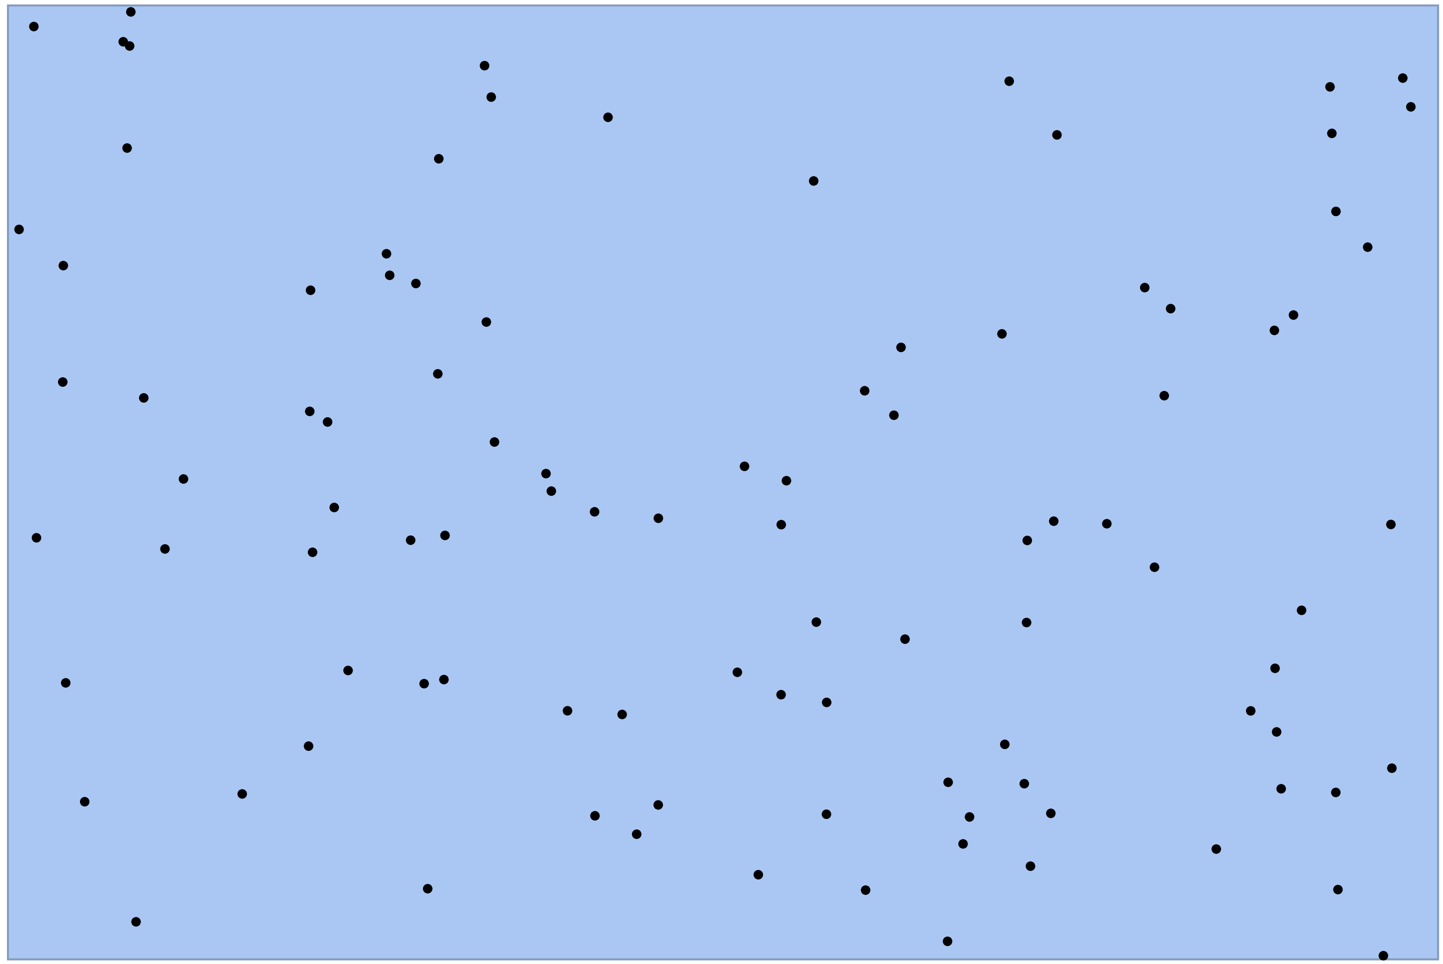

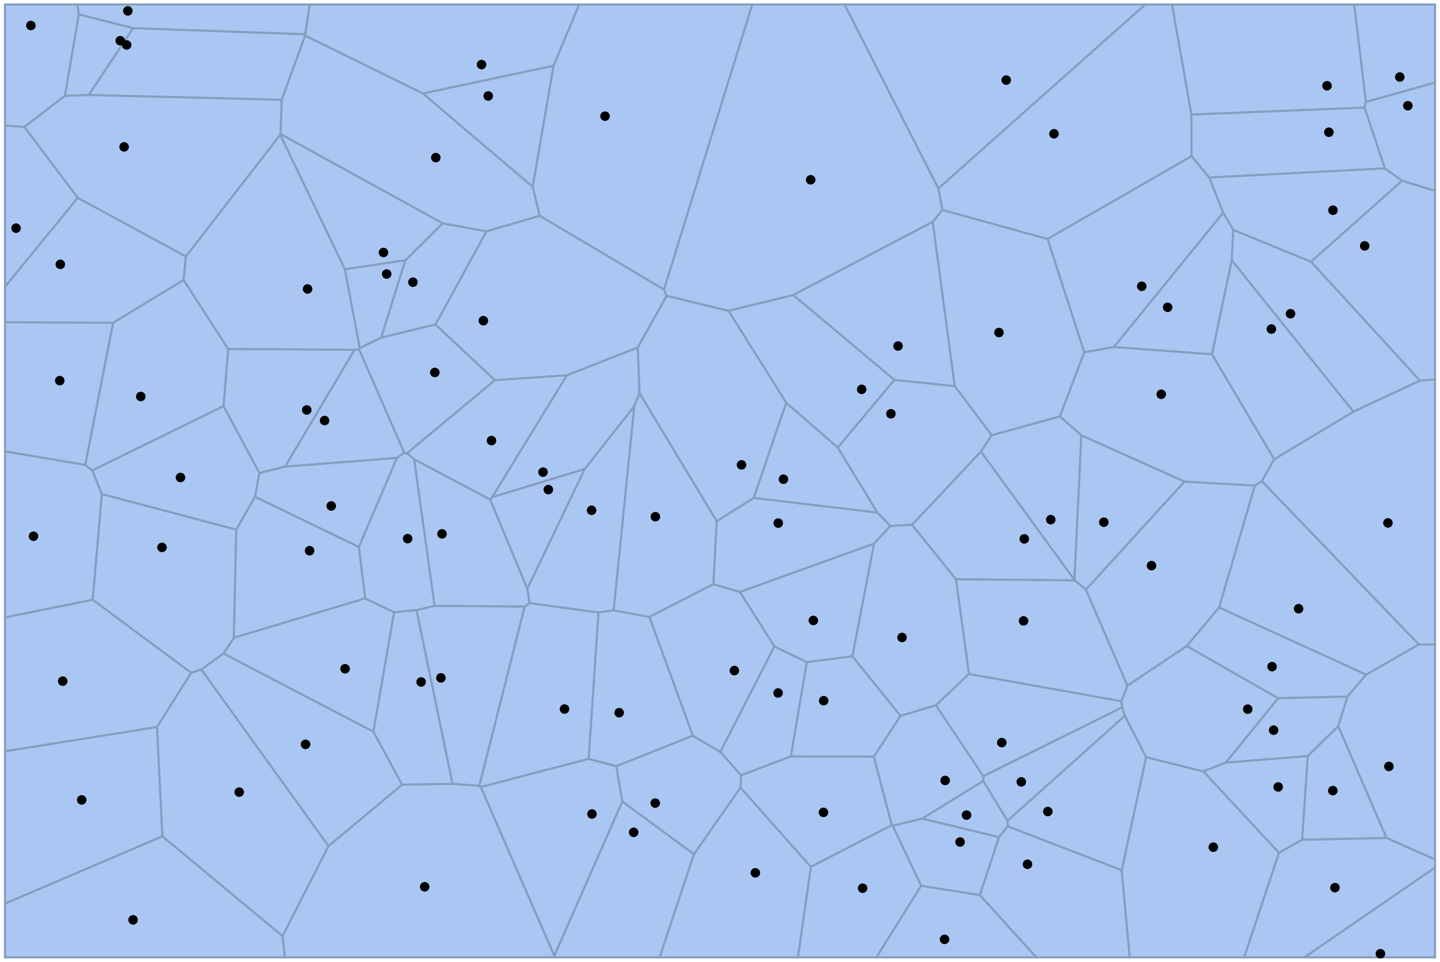

Voronoi cells (think of cell phone connecting to closest tower).

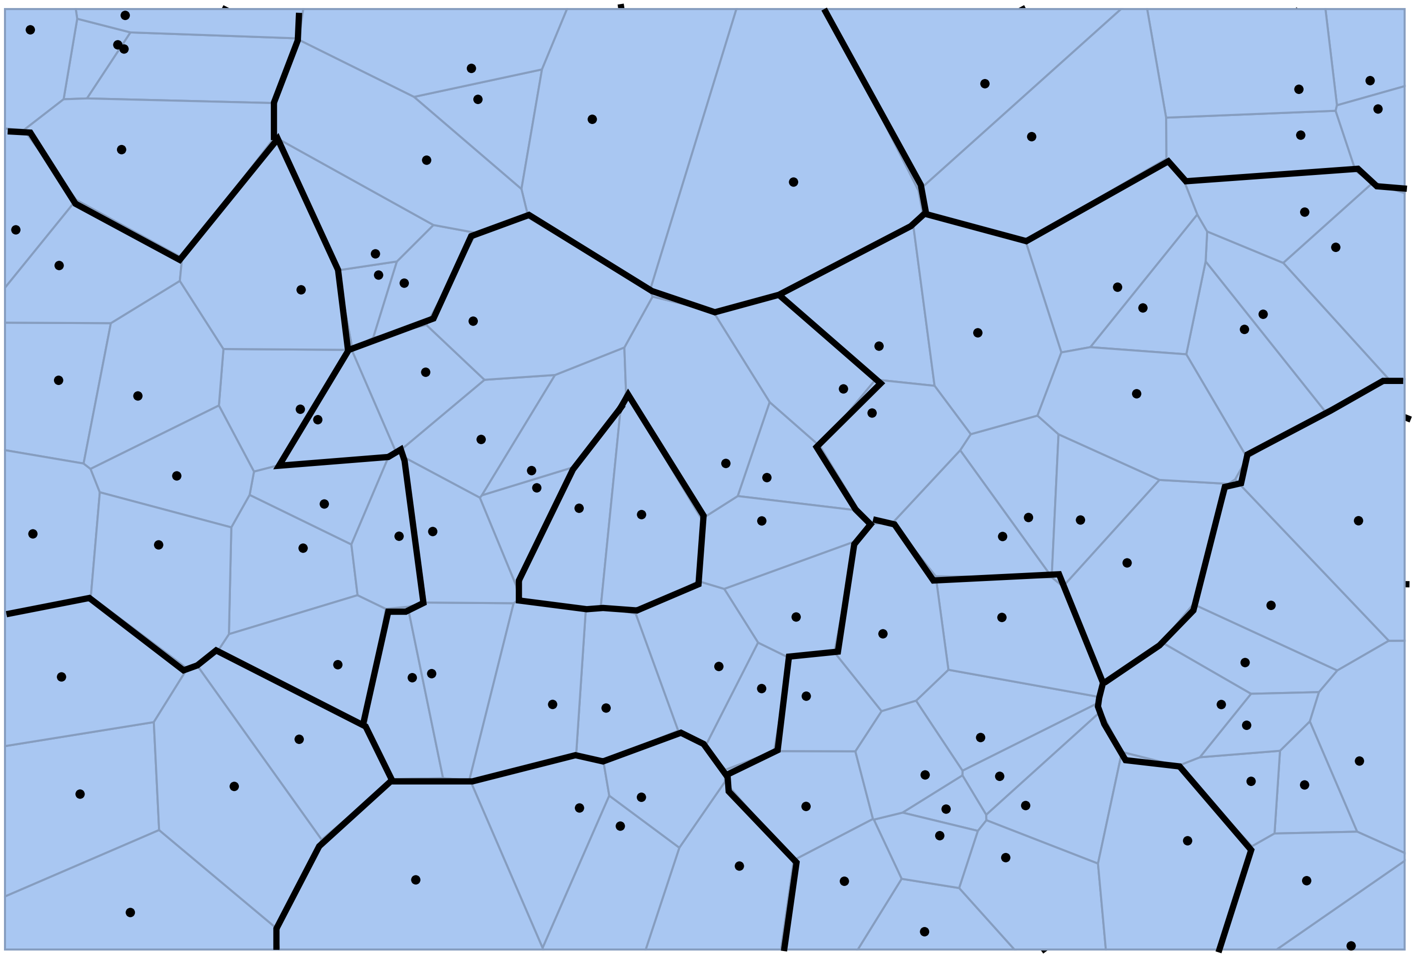

If country boundaries conform to cell boundaries, then a phone always connects to a tower in the same country.

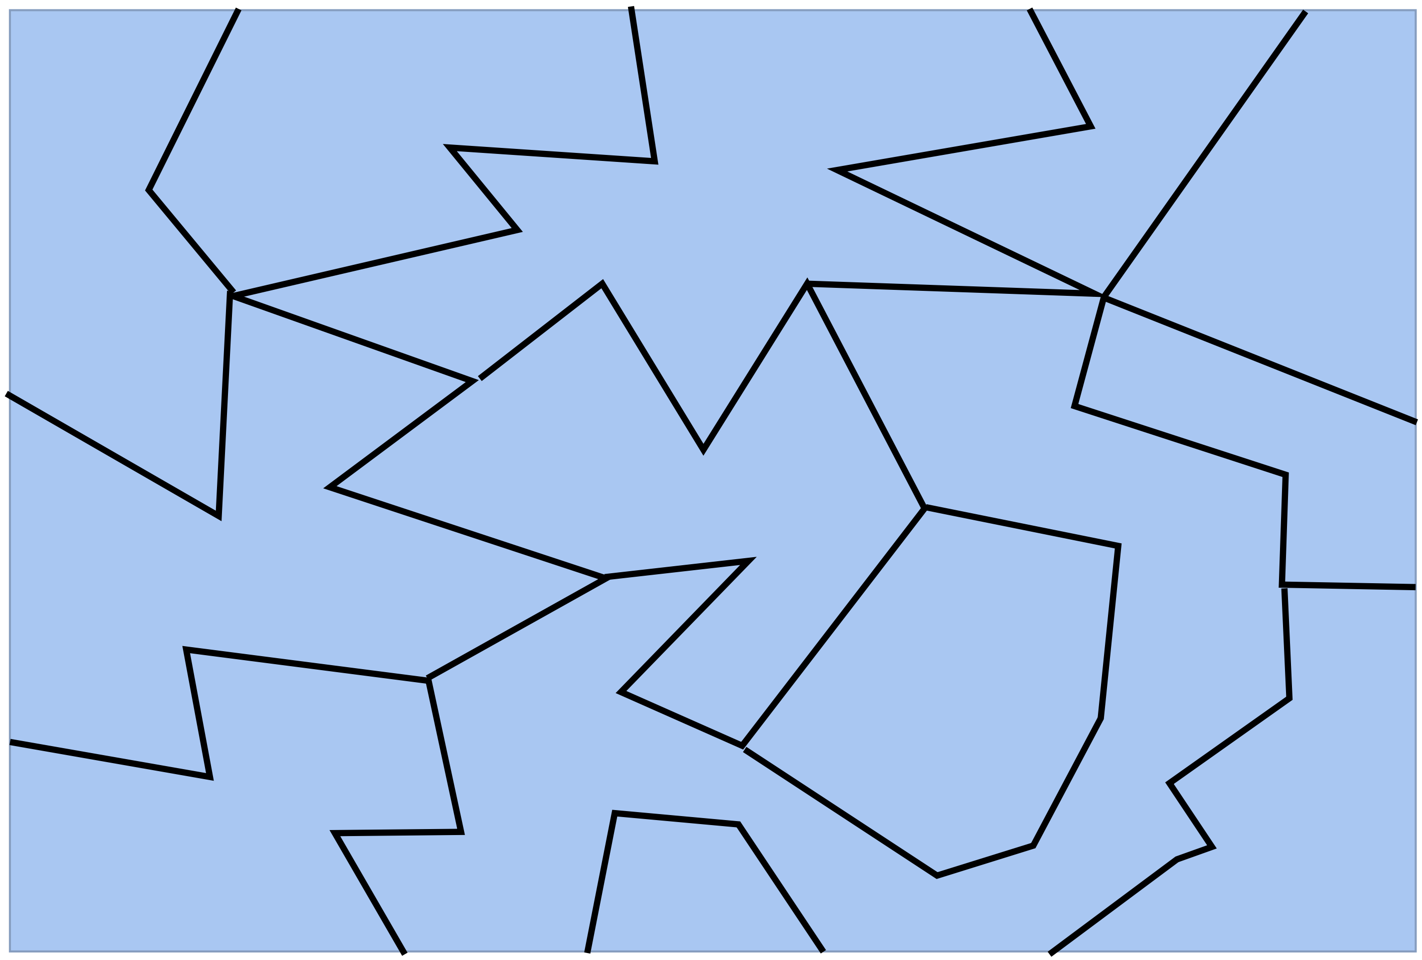

Given polygonal countries, can we place towers so this happens?

Do a polynomial number of towers suffice?

Salzberg, Delcher, Heath, Kasif, 1995, *Best-case results for nearest-neighbor learning.*

Given polygon Γ find sets I, O so that

$$\mathrm{int}(\Gamma) = \{z : \mathrm{dist}(z, I) < \mathrm{dist}(z, O)\},$$

(Voronoi diagram of $I \cup O$ covers Γ.)

Easy for nonobtuse triangles. Reduces problem to finding a NOT.

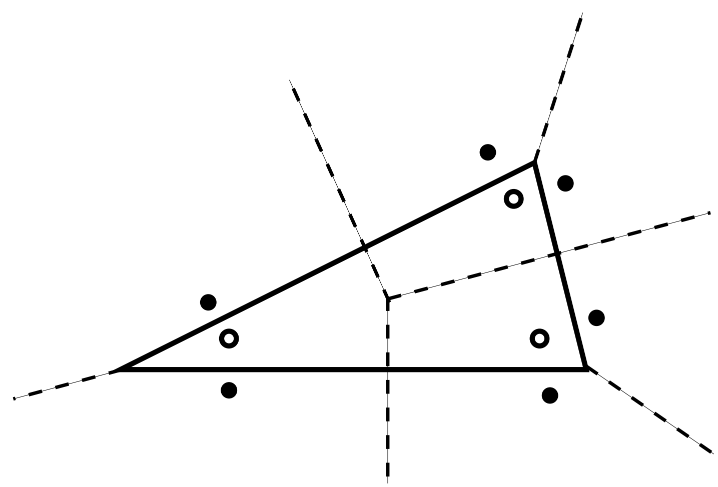

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

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.

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

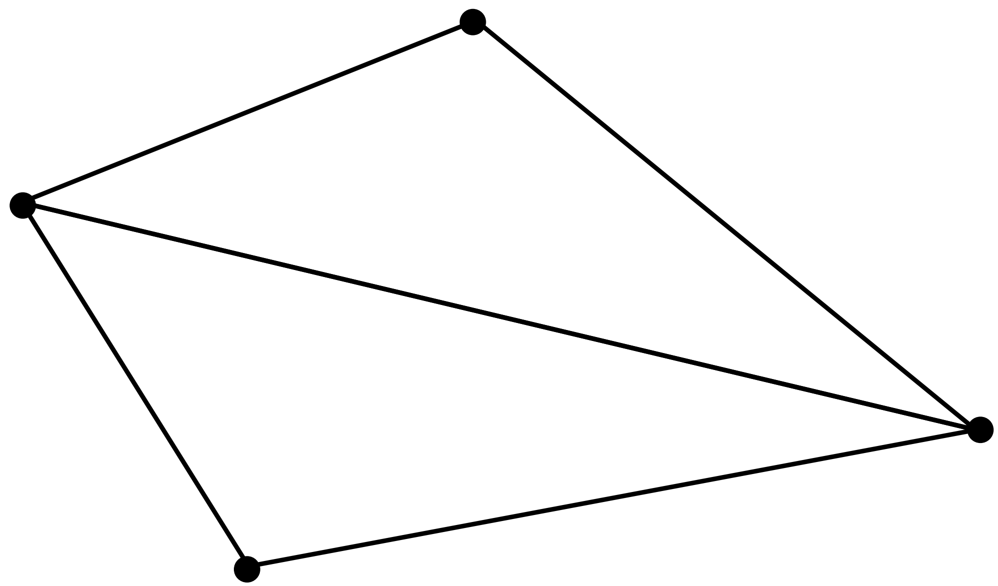

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

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

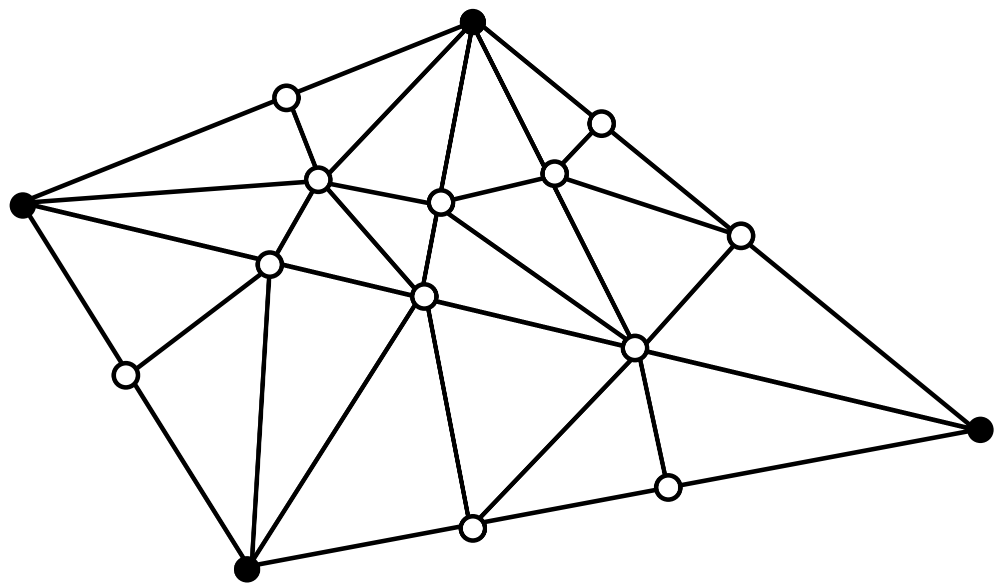

NOT-Theorem: Every PSLG has an $O(n^{2.5})$-NOT.

Cor: Any PSLG has a conforming Delaunay triangulation of size $O(n^{2.5})$.
(see Edelsbrunner, Tan 1993 for $O(n^3)$)

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

Cor: Any PSLG has a conforming Delaunay triangulation of size $O(n^{2.5})$. (see Edelsbrunner, Tan 1993 for $O(n^3)$)

Cor: Every PSLG has a triangulation with all angles $\leq 90° + \epsilon$ and $O(n^2/\epsilon^2)$ elements.

Follows from proof of NOT-theorem.

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

Cor: Any PSLG has a conforming Delaunay triangulation of size $O(n^{2.5})$. (see Edelsbrunner, Tan 1993 for $O(n^3)$)

Cor: Every PSLG has a triangulation with all angles $\leq 90° + \epsilon$ and $O(n^2/\epsilon^2)$ elements.

Cor: 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.

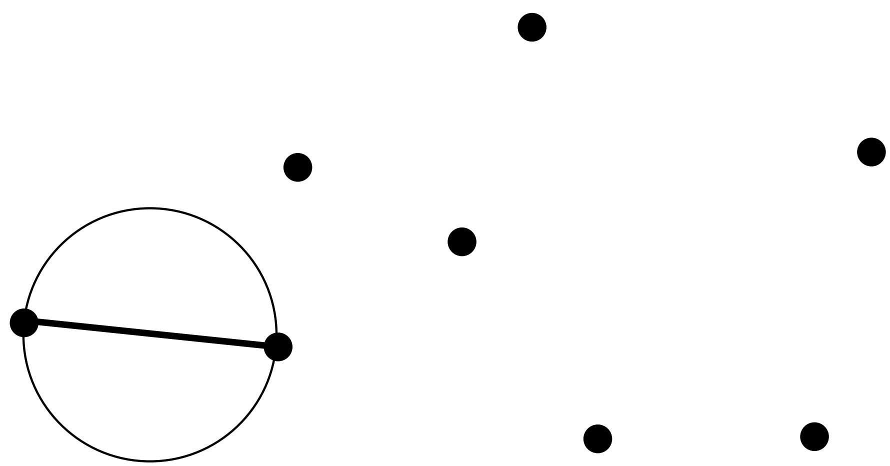

Gabriel edge.

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

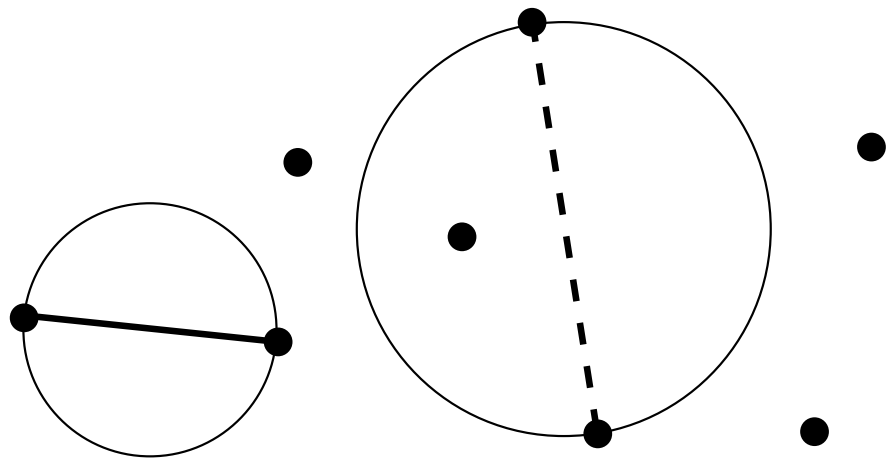

Not a Gabriel edge.

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

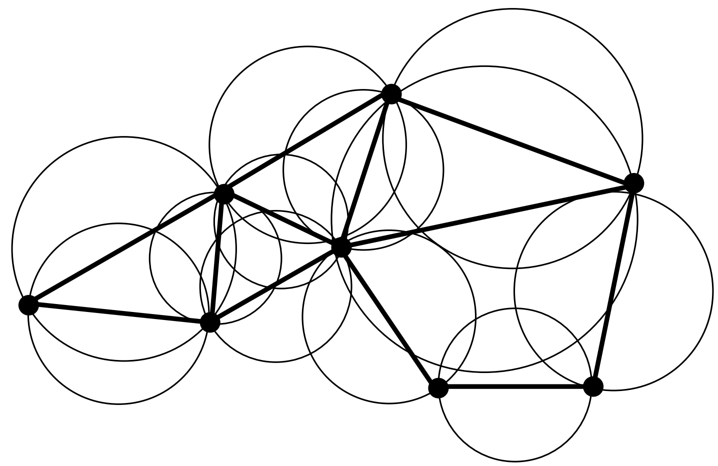

Gabriel graph contains the minimal spanning tree.

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

Gabriel edge is a special case of a **Delaunay** edge: $[v, w]$ is a **chord** of an open disk not hitting V.

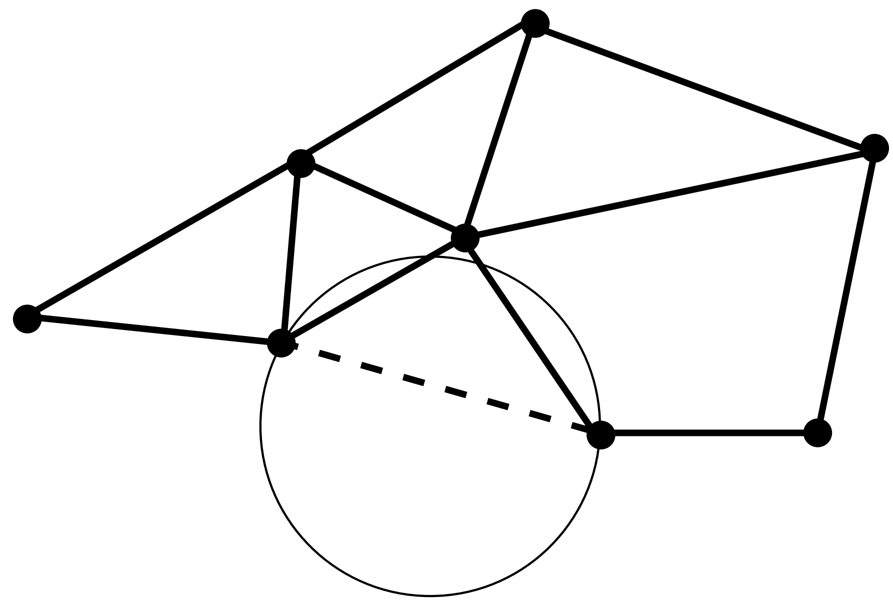

Gabriel edge is a special case of a **Delaunay** edge: $[v, w]$ is a **chord** of an open disk not hitting V.

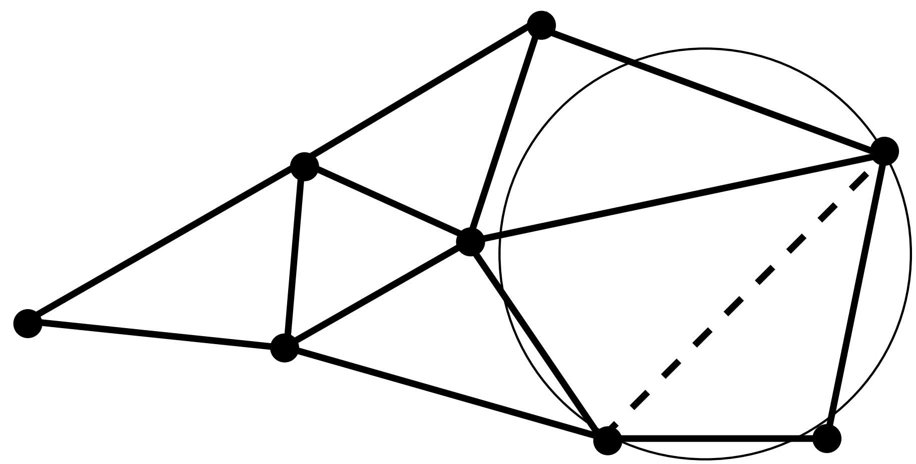

Gabriel edge is a special case of a **Delaunay** edge: $[v, w]$ is a **chord** of an open disk not hitting V.

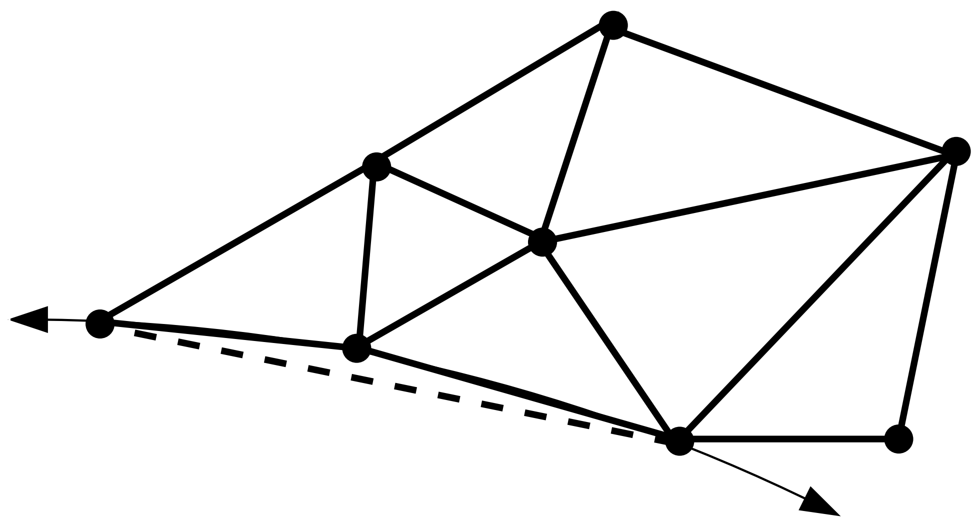

Gabriel edge is a special case of a **Delaunay** edge: $[v, w]$ is a **chord** of an open disk not hitting V.

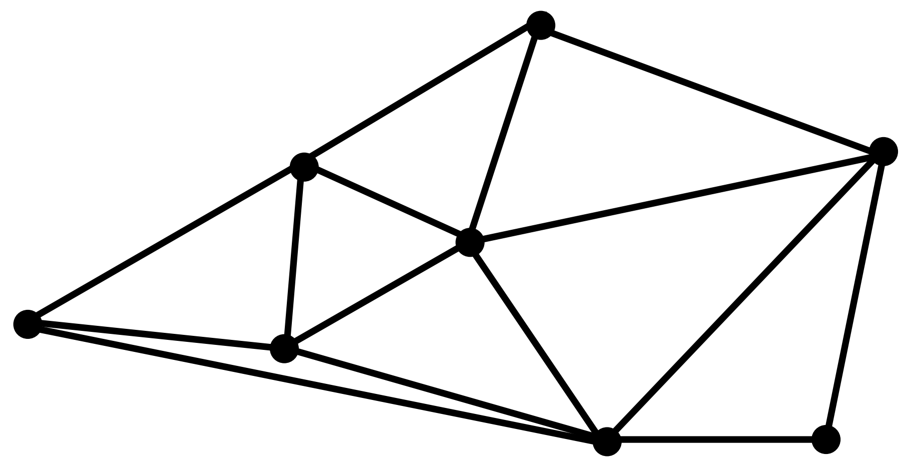

Delaunay edges triangulate.

DT minimizes the maximum angle, among all triangulations.

Thus, if a point set has a NOT, then DT = NOT.

Known: If we can add N points to a triangulation of n points so that every edge becomes Gabriel, then there is a $O(n + N)$ NOT.

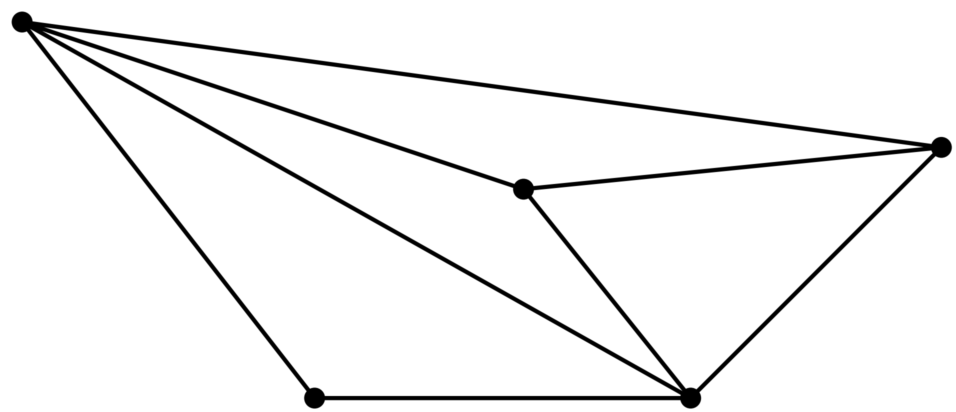

Follows from work of Bern, Mitchell, Ruppert (1994).

Known: If we can add N points to a triangulation of n points so that every edge becomes Gabriel, then there is a $O(n + N)$ NOT.

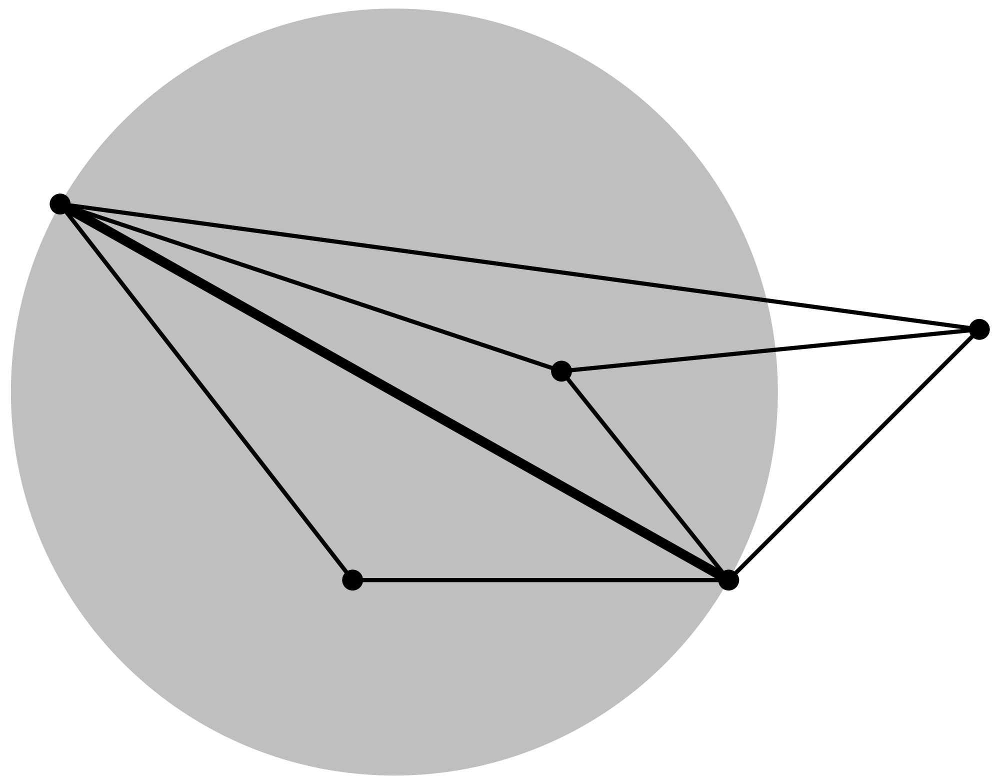

Known: If we can add N points to a triangulation of n points so that every edge becomes Gabriel, then there is a $O(n+N)$ NOT.

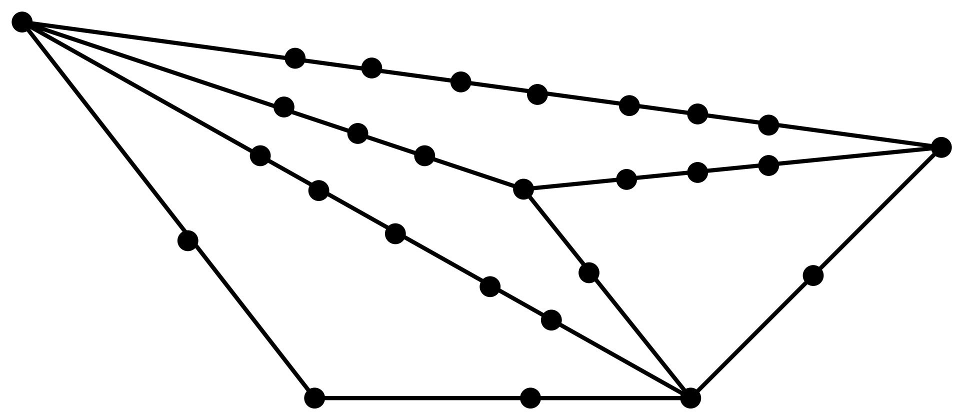

Known: If we can add N points to a triangulation of n points so that every edge becomes Gabriel, then there is a $O(n+N)$ NOT.

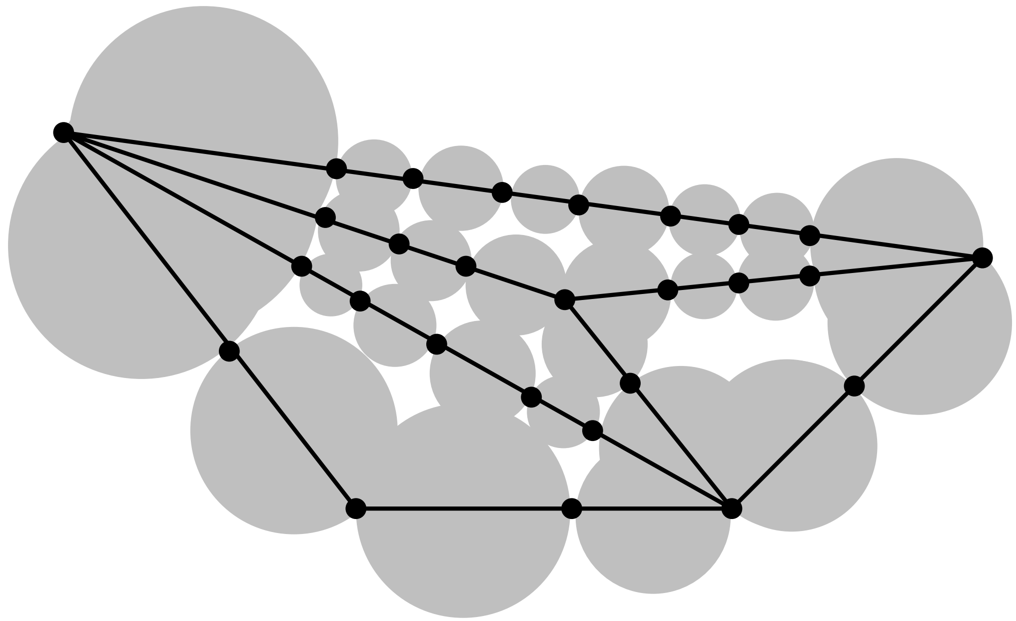

Theorem: Given a triangulation, we can add $O(n^{2.5})$ points to the edges so that every new edge is Gabriel.

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.

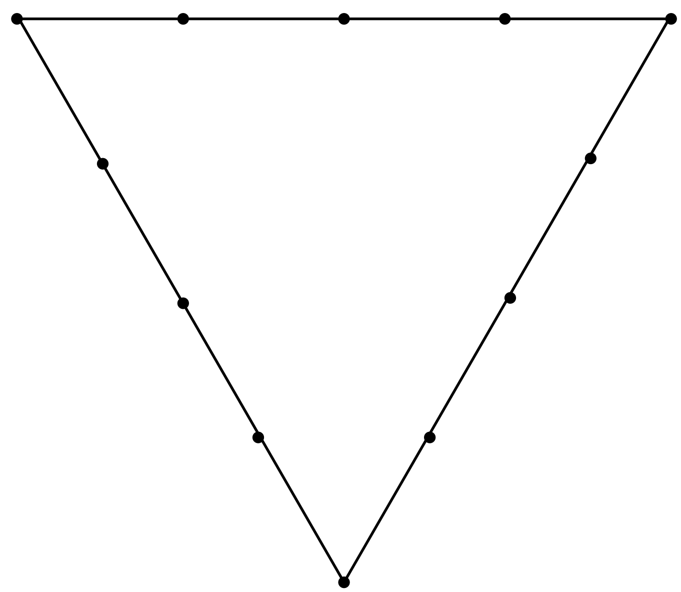

Gabriel edges $\Rightarrow$ non-obtuse triangulation

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.

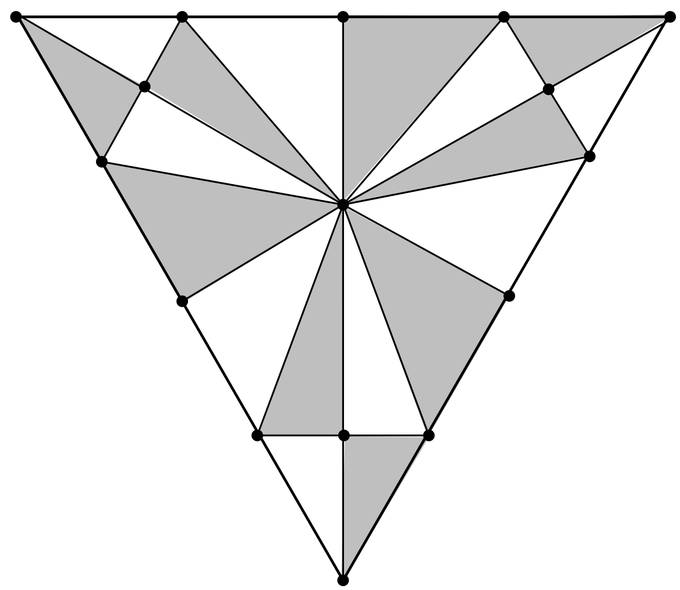

Gabriel edges $\Rightarrow$ non-obtuse triangulation

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. Done.

Proof of BMR:

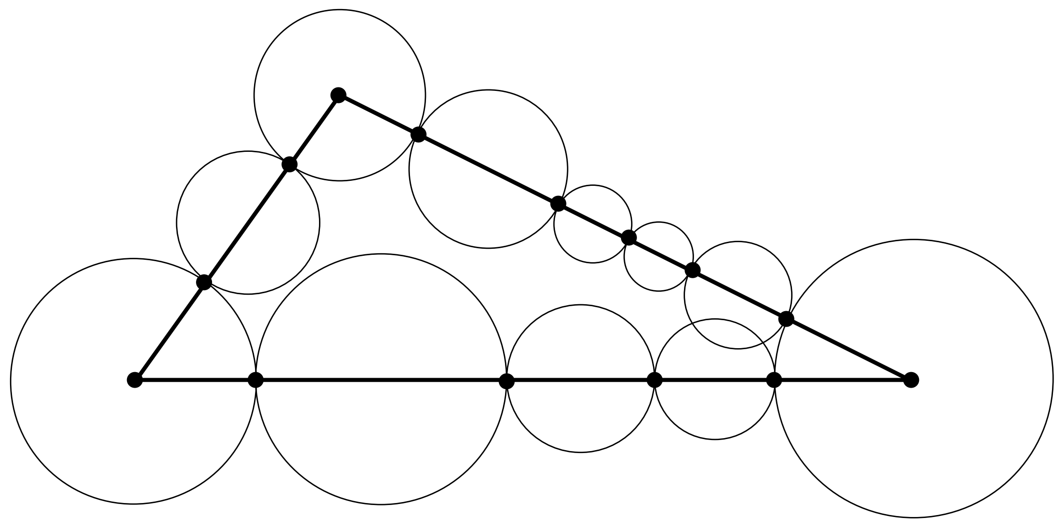

Proof of BMR:

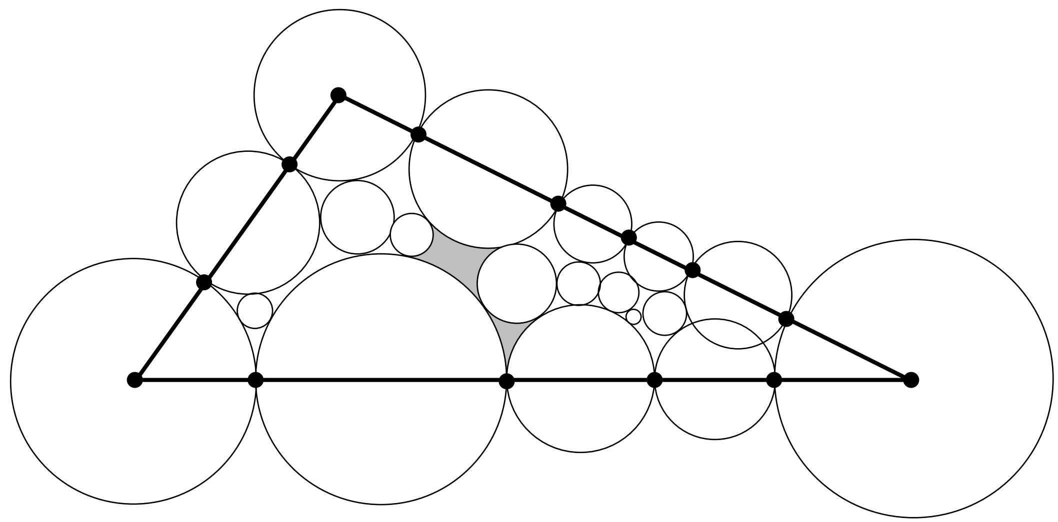

• Pack with disjoint disks so only 3-sided and 4-sided regions remain.

Proof of BMR:

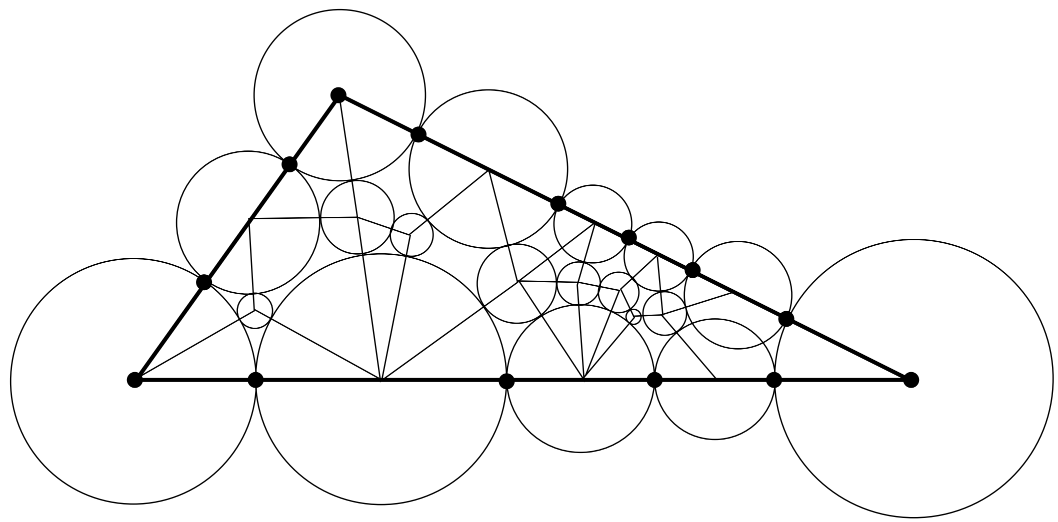

- Pack with disjoint disks so only 3-sided and 4-sided regions remain.
- Connect centers.

Proof of BMR:

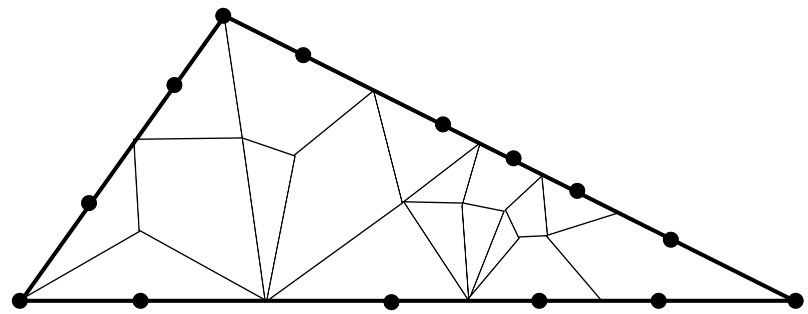

- Pack with disjoint disks so only 3-sided and 4-sided regions remain.
- Connect centers.
- Divides triangle into triangles and quadrilaterals.

Proof of BMR:

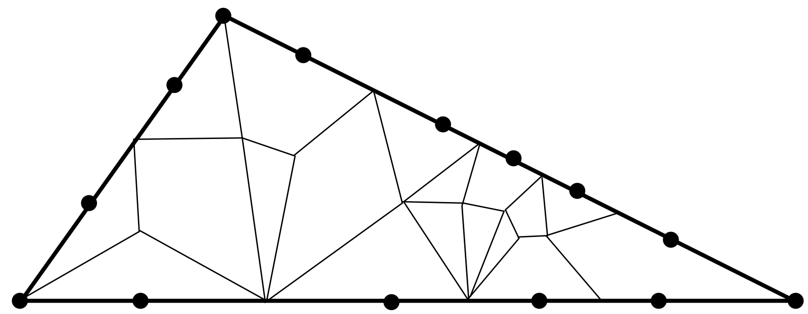

- Pack with disjoint disks so only 3-sided and 4-sided regions remain.
- Connect centers.
- Divides triangle into triangles and quadrilaterals.

Nonobtusely triangulate each region without new boundary vertices. Several cases.

Proof of BMR:

First case: decompose 3-region into right triangles

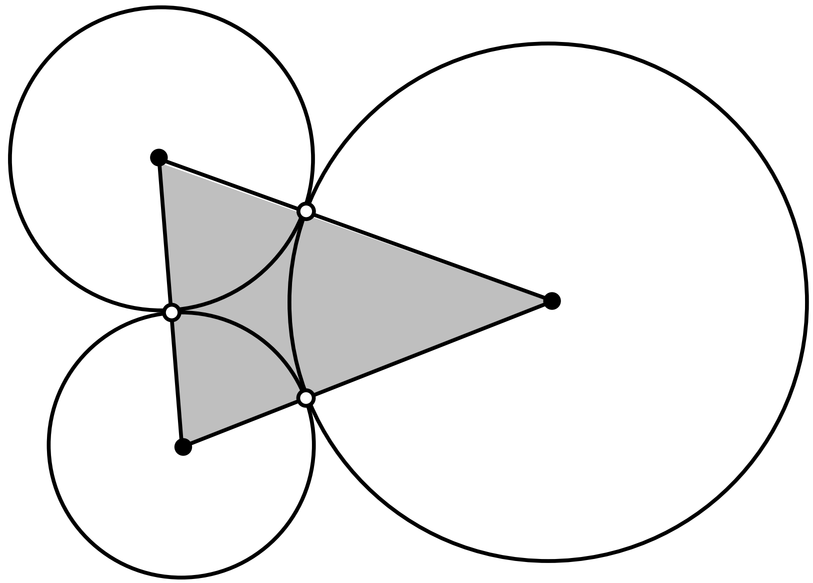

Proof of BMR:

First case: decompose 3-region into right triangles

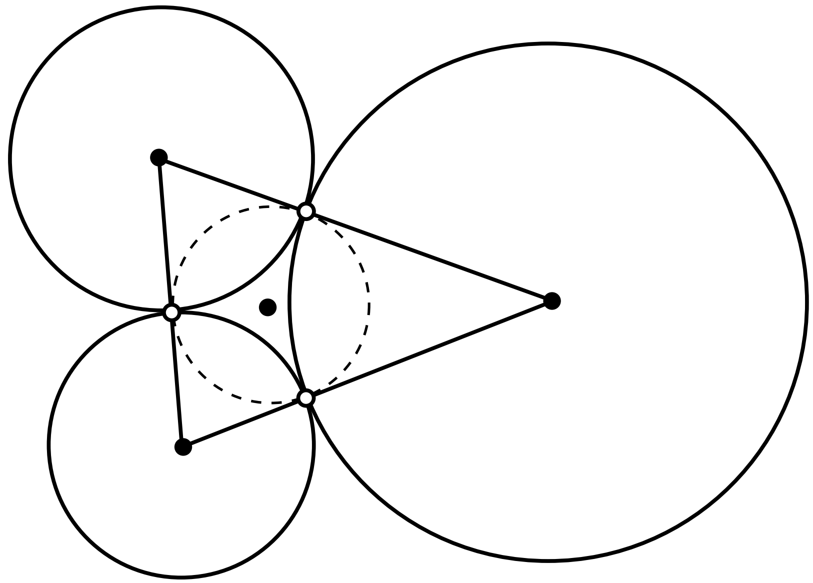

- Add center of circle through the three tangent points.

Proof of BMR:

First case: decompose 3-region into right triangles

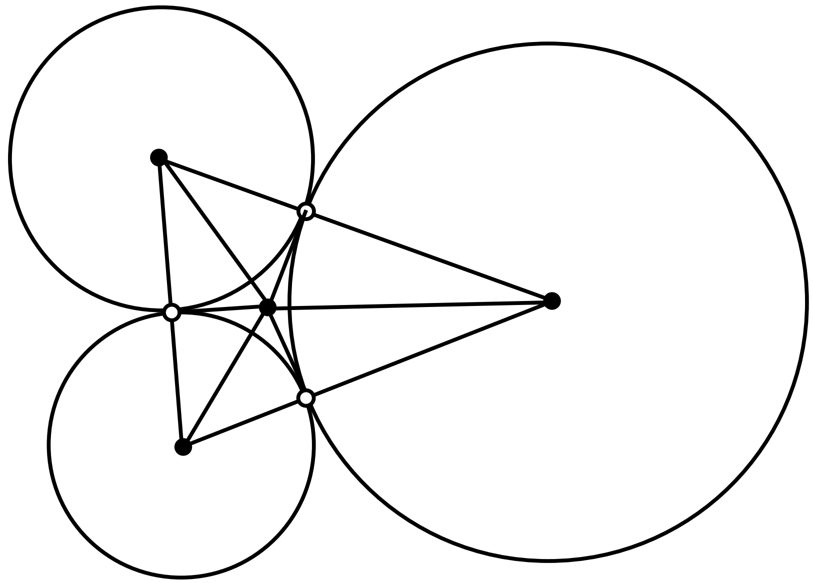

- Add center of circle through the three tangent points.
- Connect center to tangent points and centers of circles

Proof of BMR:

First case: decompose 3-region into right triangles

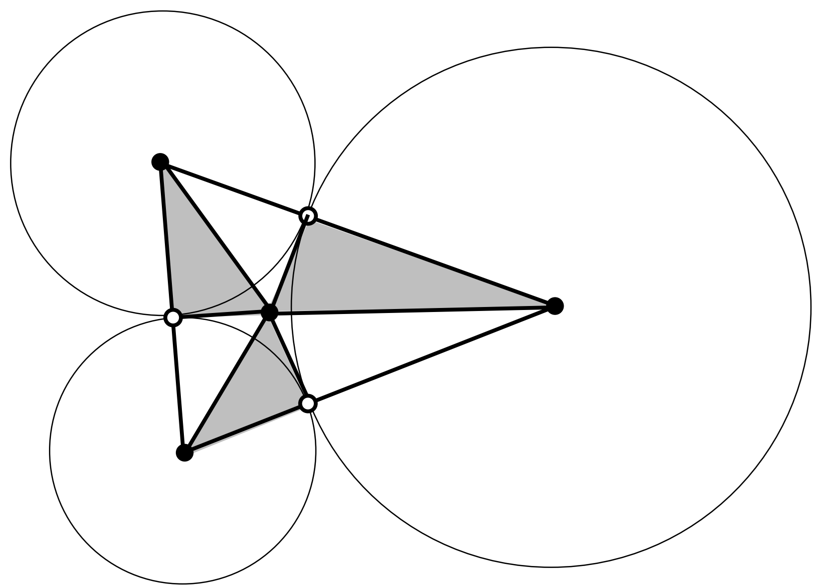

- Add center of circle through the three tangent points.
- Connect center to tangent points and centers of circles

Proof of BMR:

The 4-regions are similar (but several cases arise).

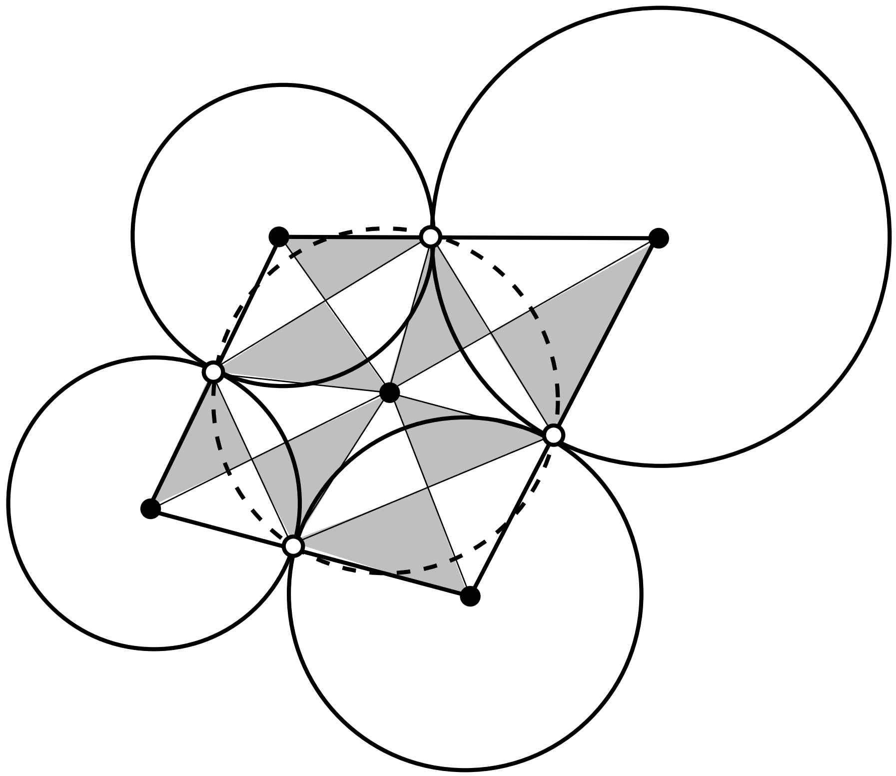

So Gabriel covering implies a nonobtuse triangulation.

Proof of BMR:

The 4-regions are similar (but several cases arise).

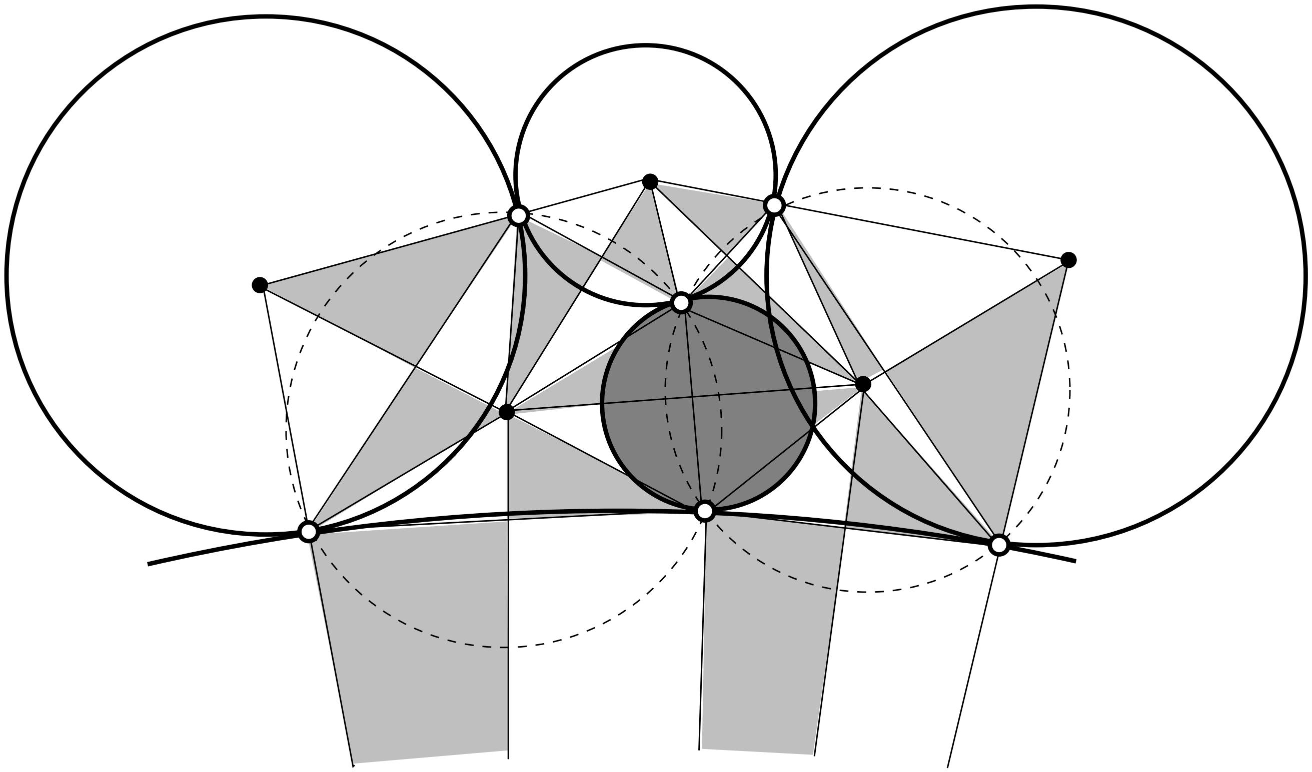

So Gabriel covering implies a nonobtuse triangulation.

Proof of BMR:

The 4-regions are similar (but several cases arise).

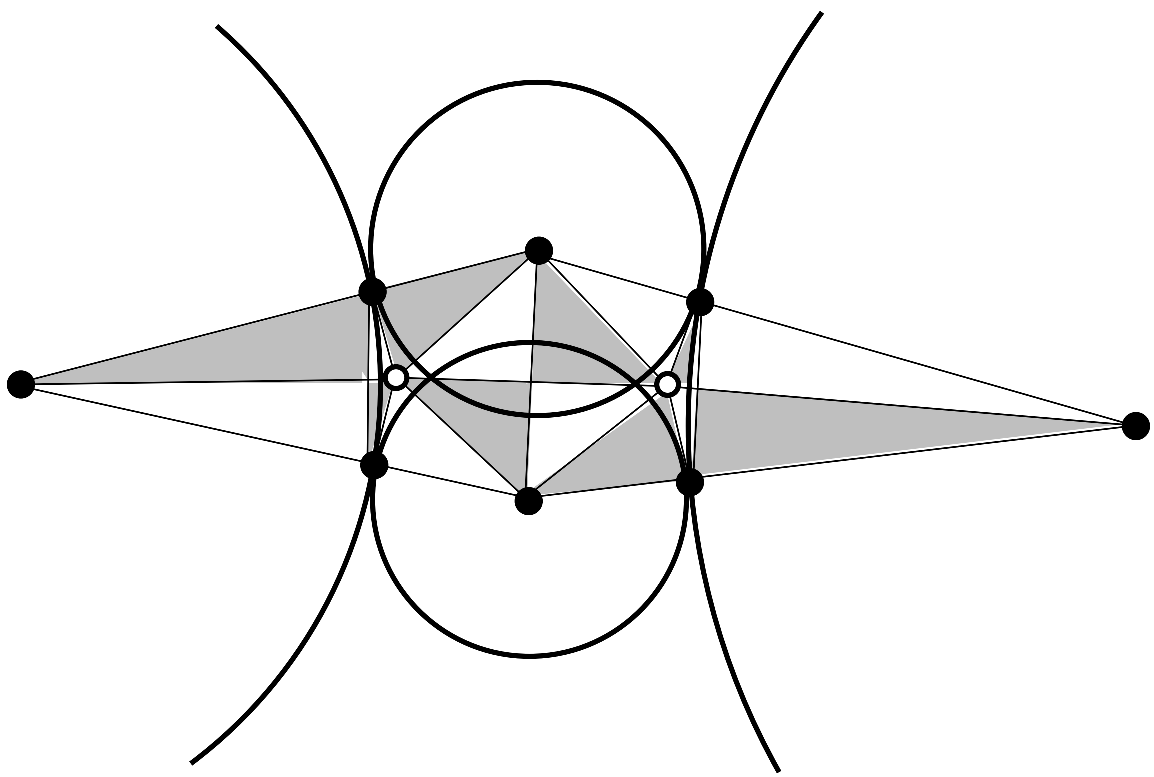

So Gabriel covering implies a nonobtuse triangulation.

Construct Gabriel points:

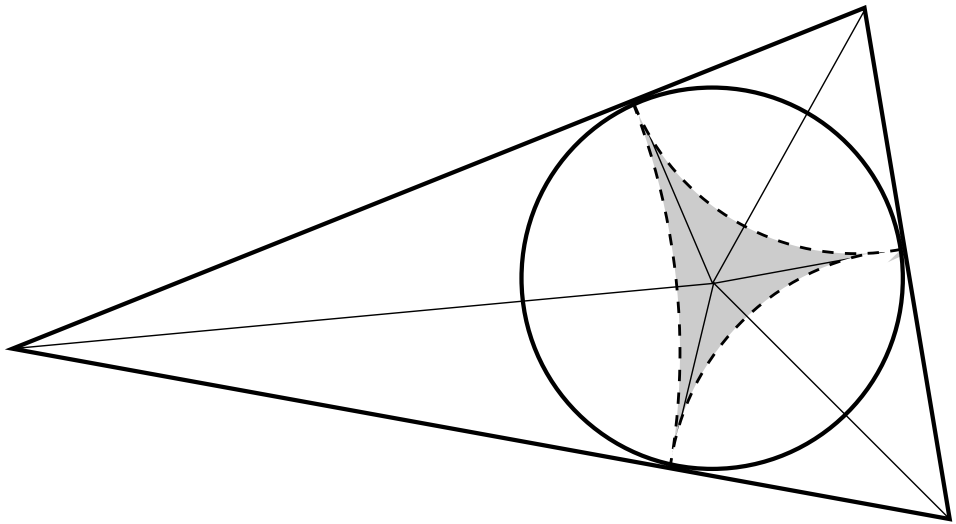

Break every triangle into thick and thin parts.

Thin parts = corners,   Thick part = central region

Construct Gabriel points:

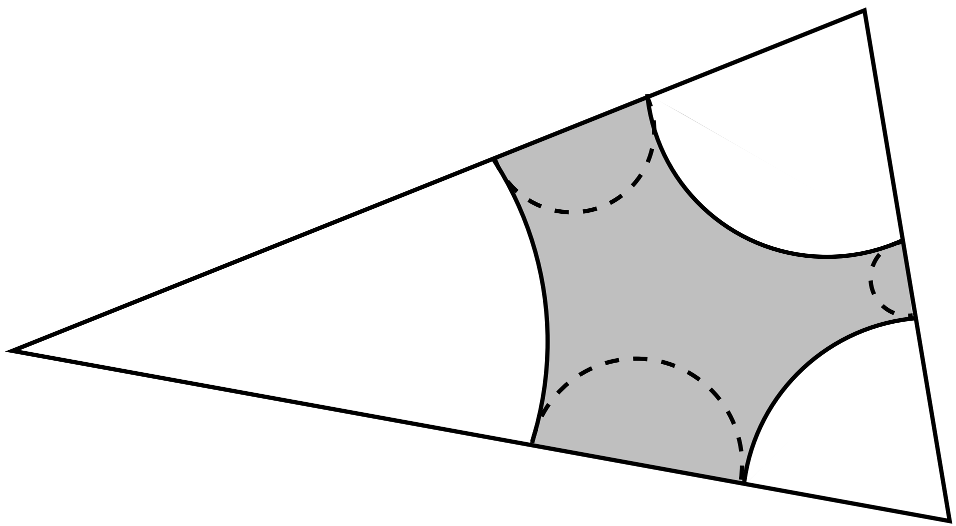

Divide triangle into thick and thin parts.

Thick sides are base of half-disk inside triangle.

Thick version prevents infinite propagation.

Construct Gabriel points:

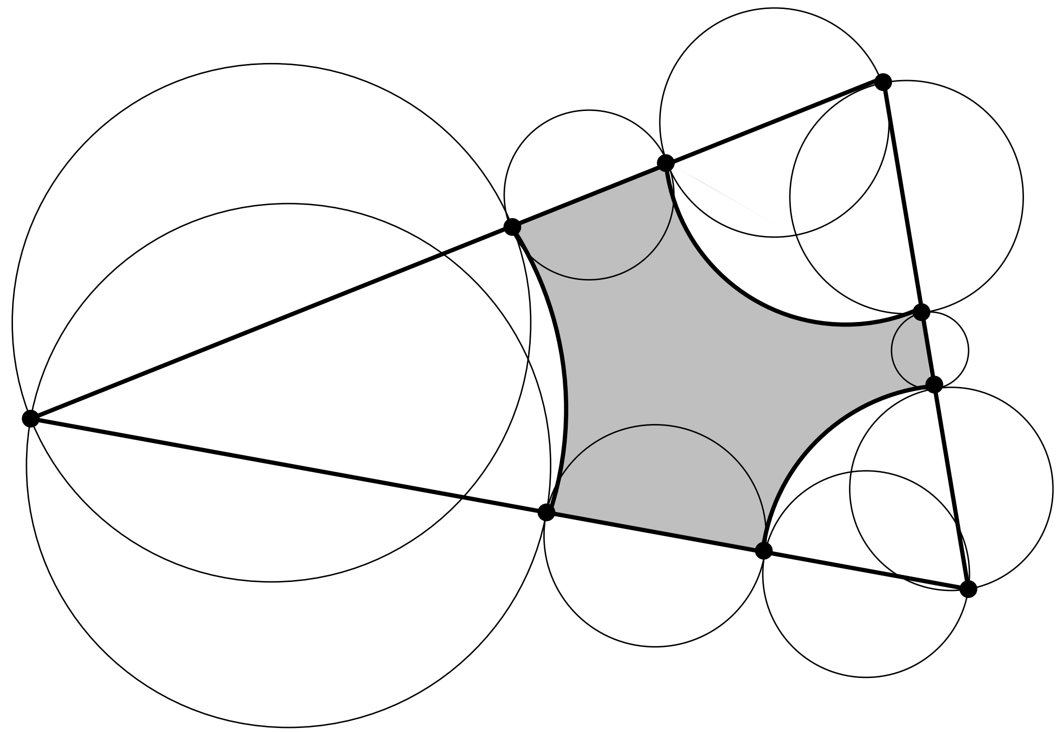

Then vertices of thick part give Gabriel edges.

Construct Gabriel points:

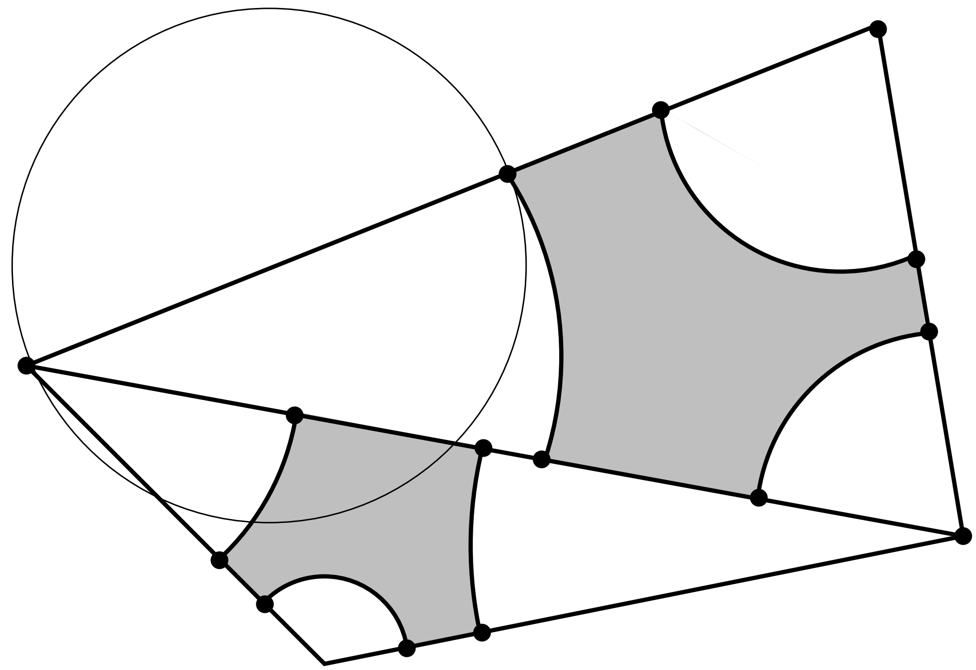

But, adjacent triangle can make Gabriel condition fail.

Construct Gabriel points:

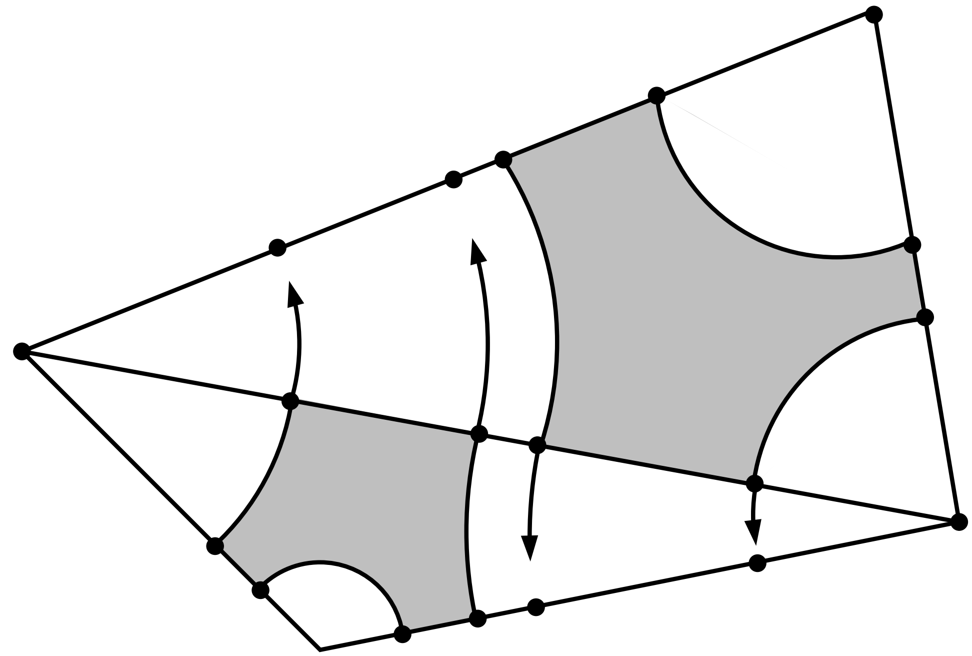

Idea: "Push" vertices across the thin parts.

Construct Gabriel points:

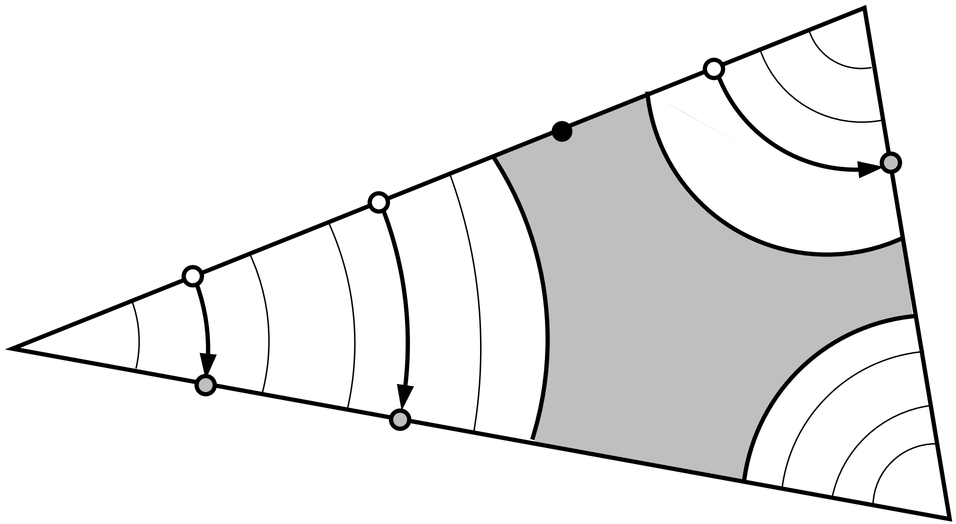

Thin parts foliated by circles centered at vertices.

Push vertices along foliation paths.

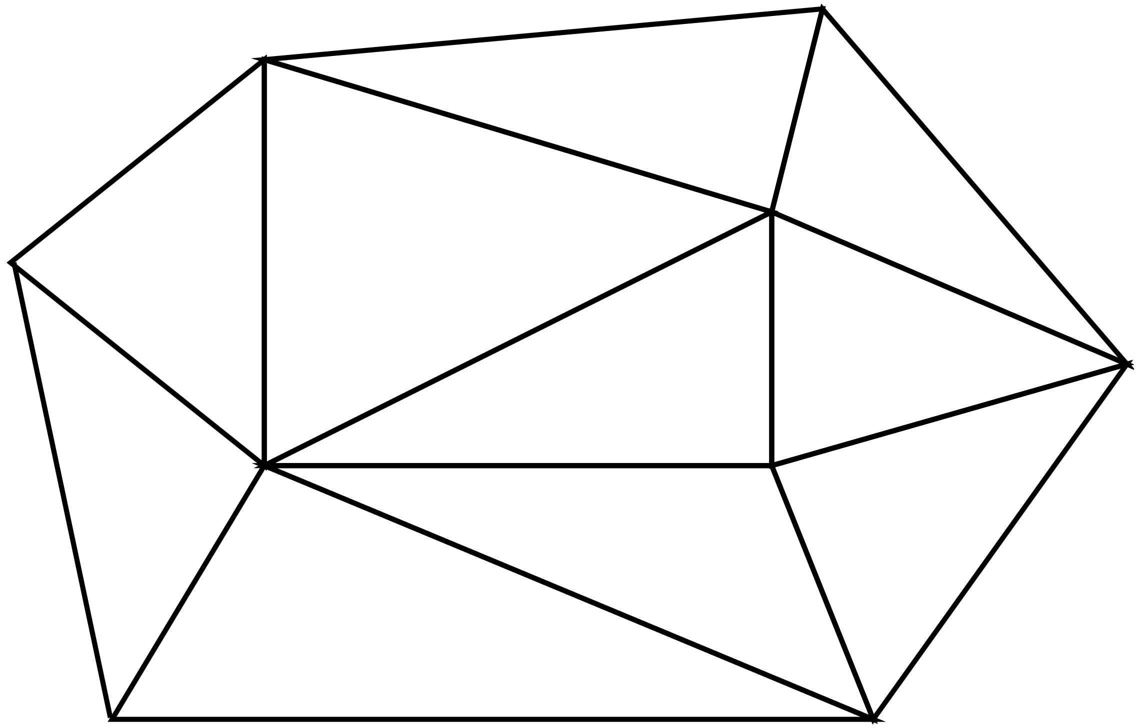

- Start with any triangulation.

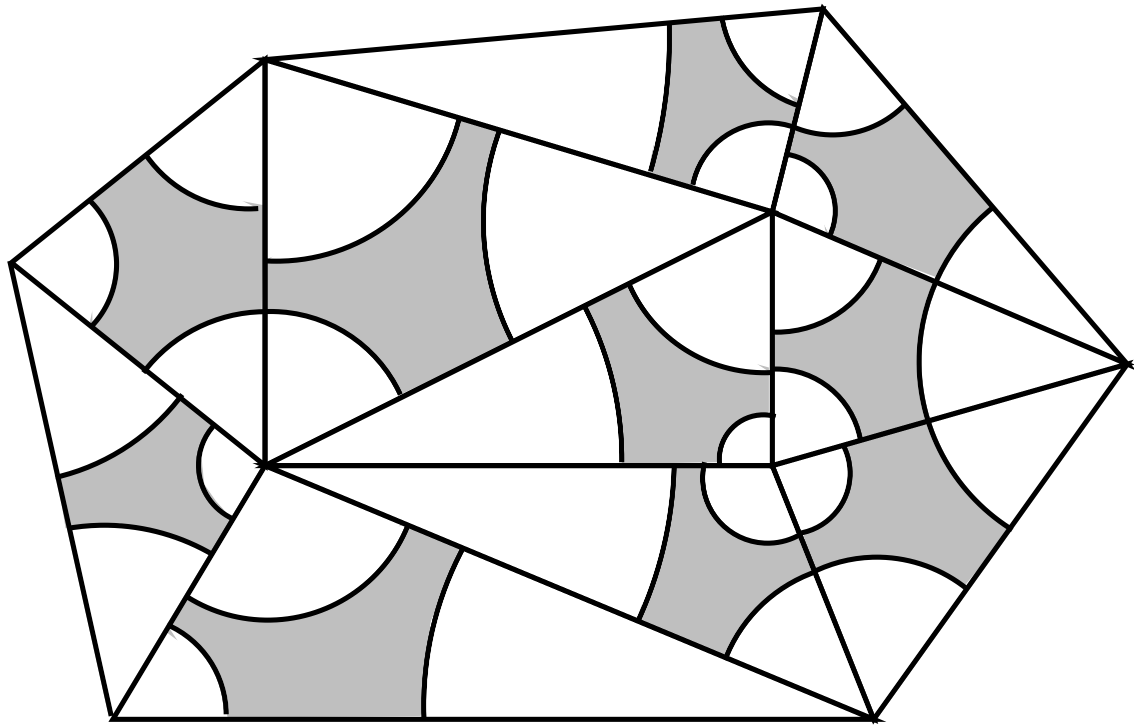

- Start with any triangulation.
- Make thick/thin parts.

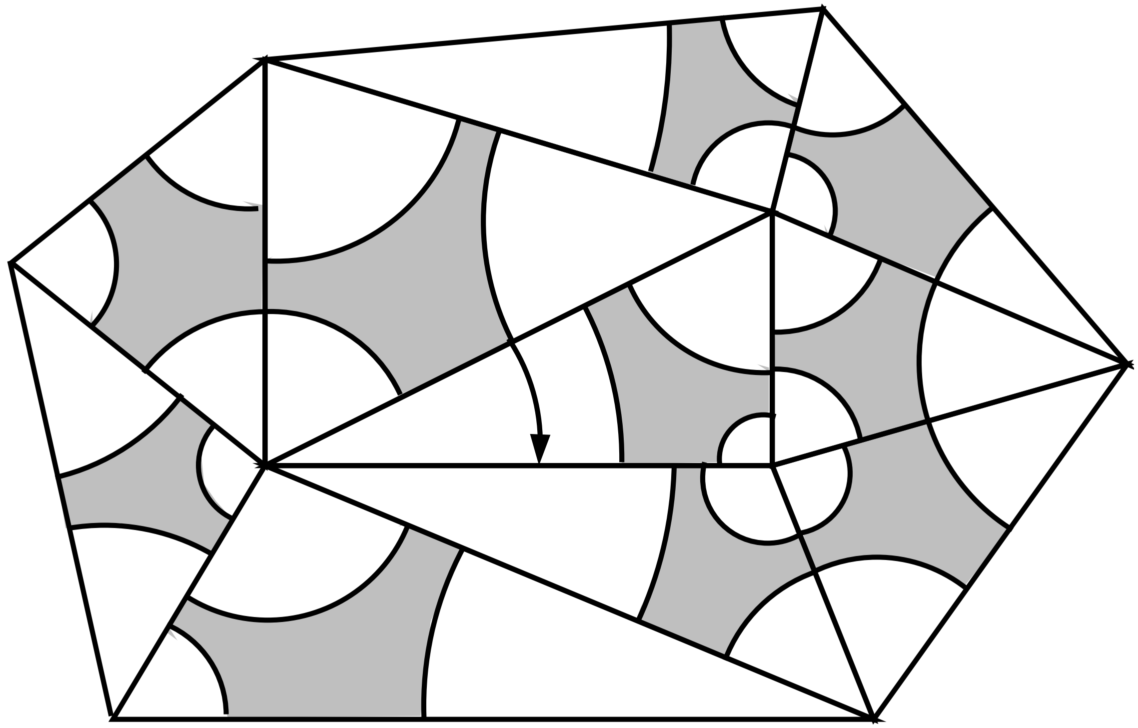

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

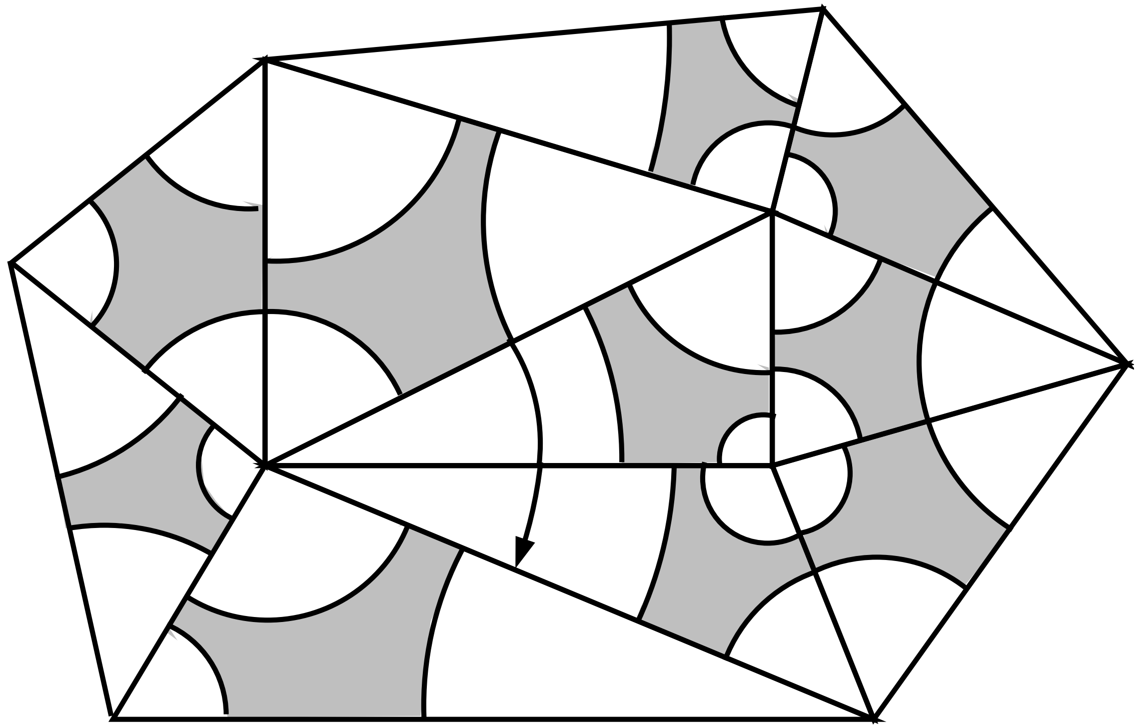

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

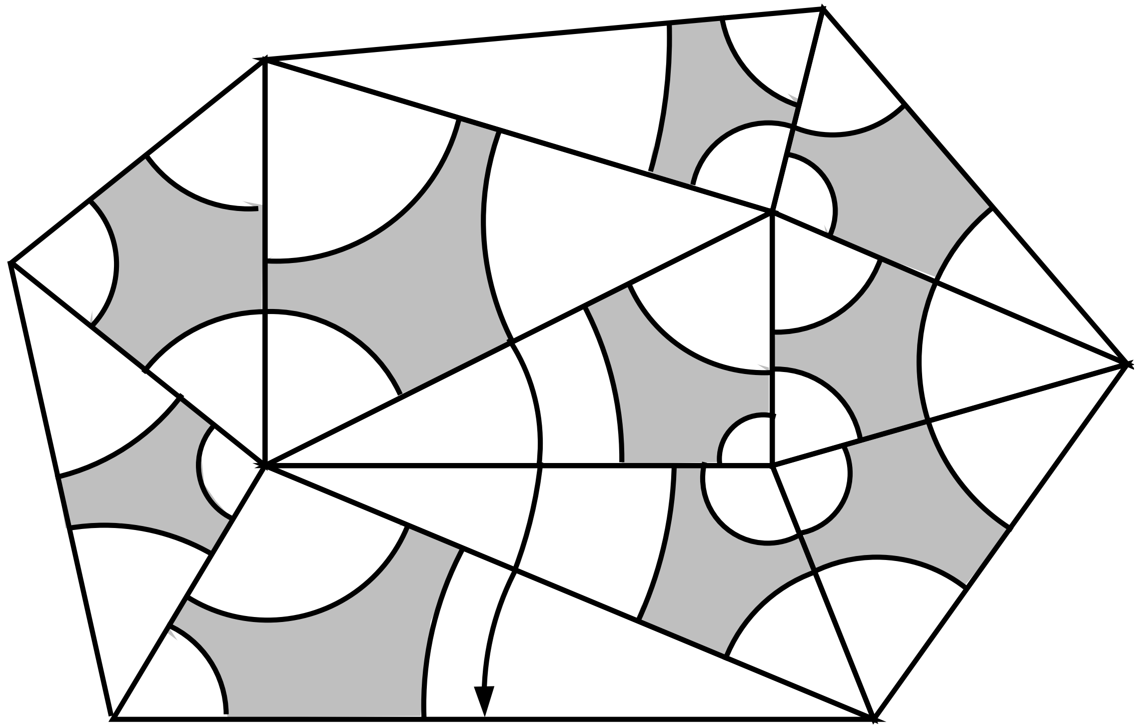

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

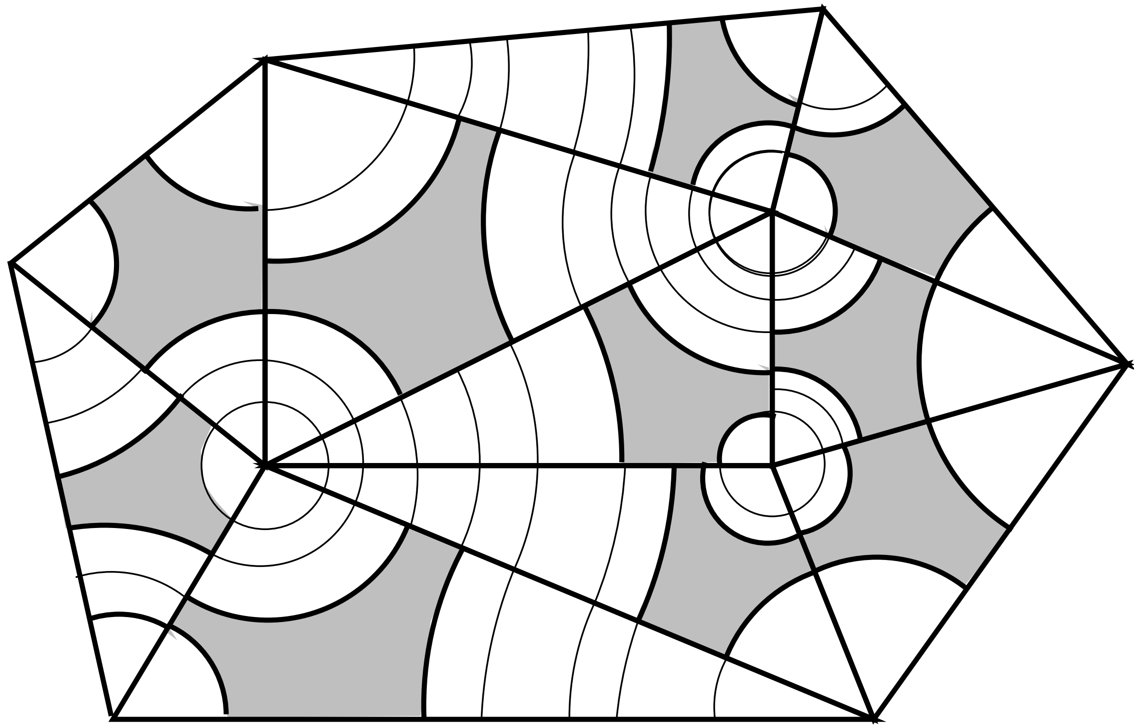

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.

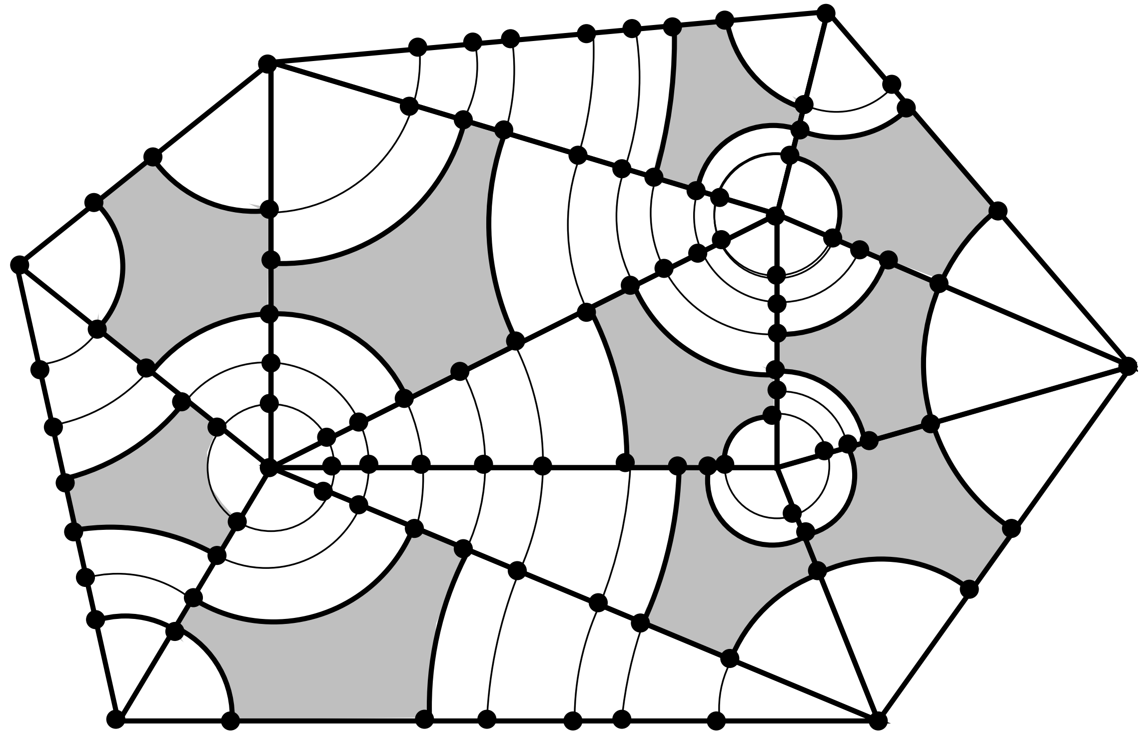

- Start with any triangulation.
- Make thick/thin parts.
- Propagate vertices until they leave thin parts.
- Intersections satisfy Gabriel condition. Why?

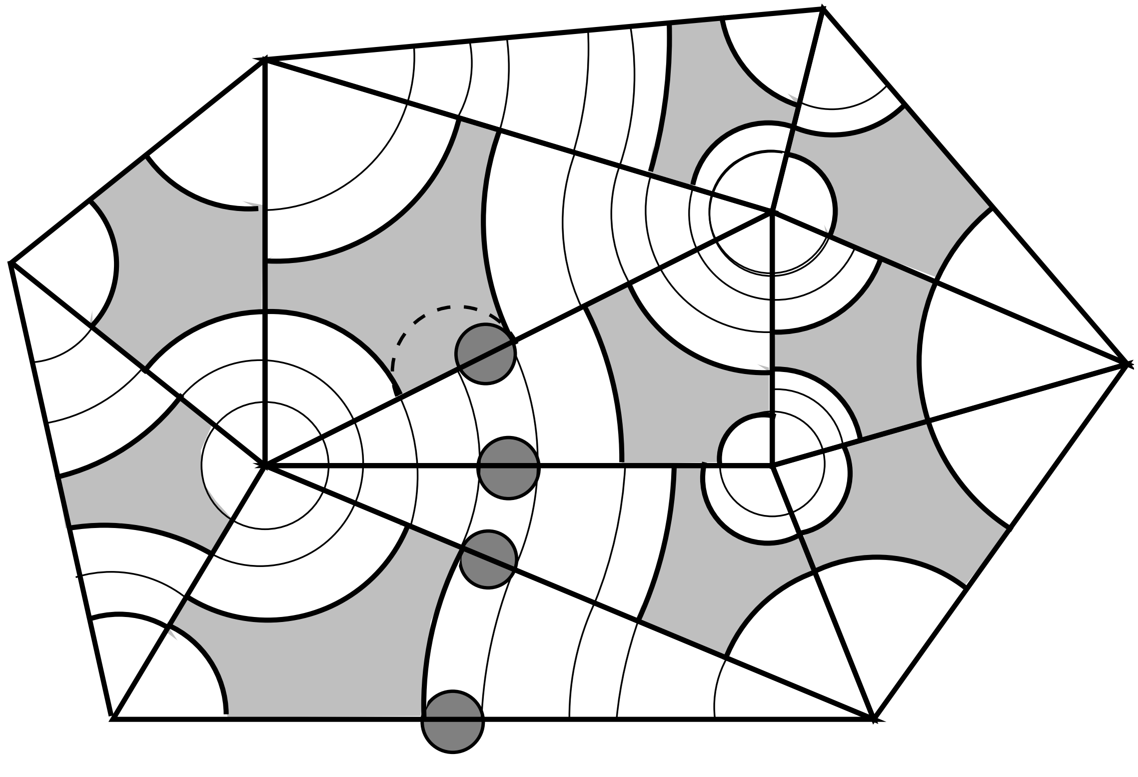

Tube is “swept out” by fixed diameter disk.

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

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

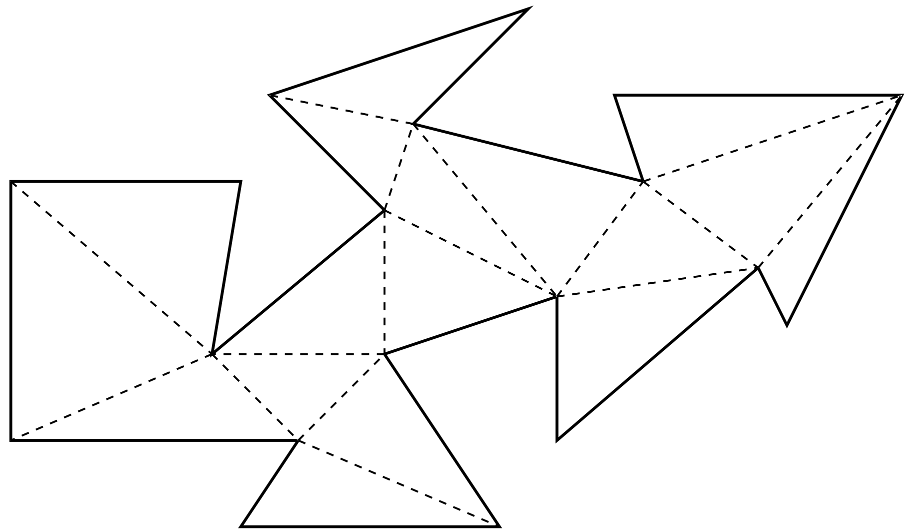

In triangulation of a n-gon, adjacent triangles form tree.

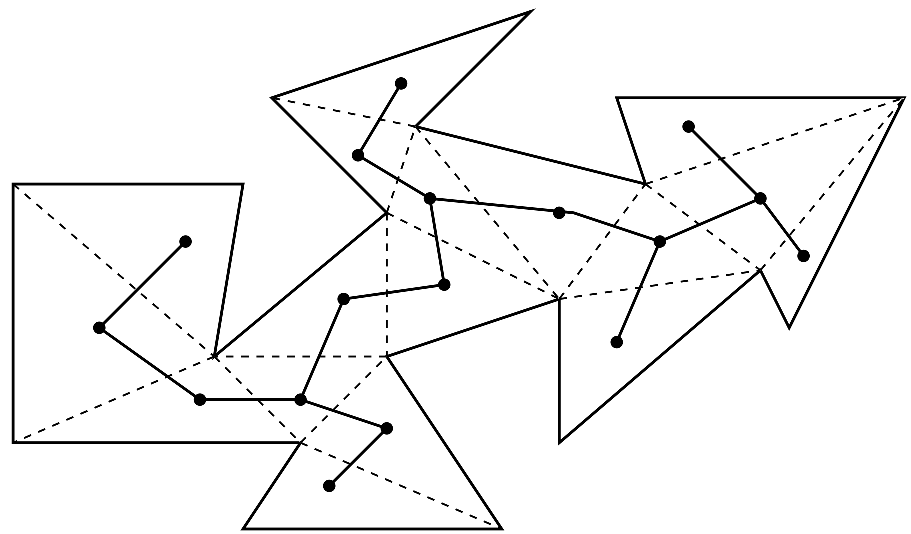

In triangulation of a n-gon, adjacent triangles form tree.

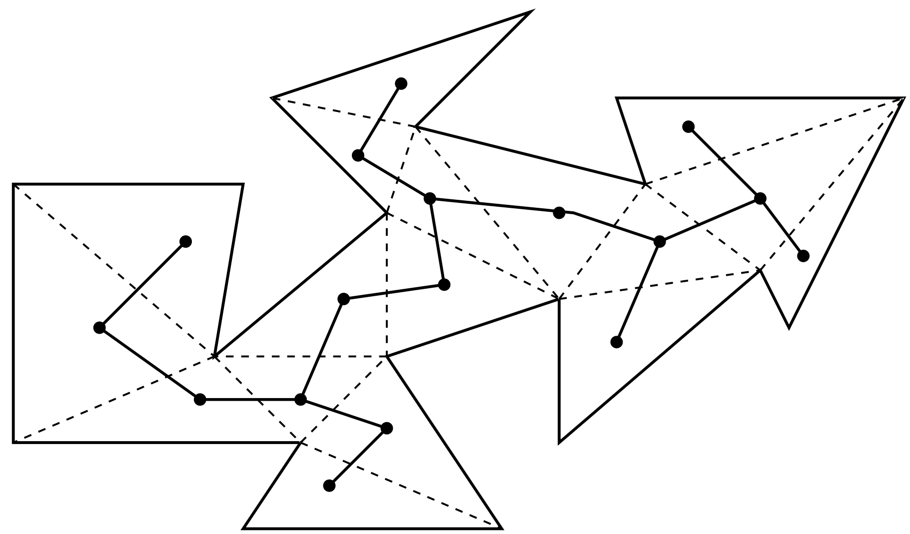

Hence foliation paths never revisit a triangle.

$O(n)$ starting points, so $O(n^2)$ points are created.

Thm: Triangulation of a n-gon has a $O(n^2)$ NOT.

Improves $O(n^4)$ by Bern and Eppstein (1992).

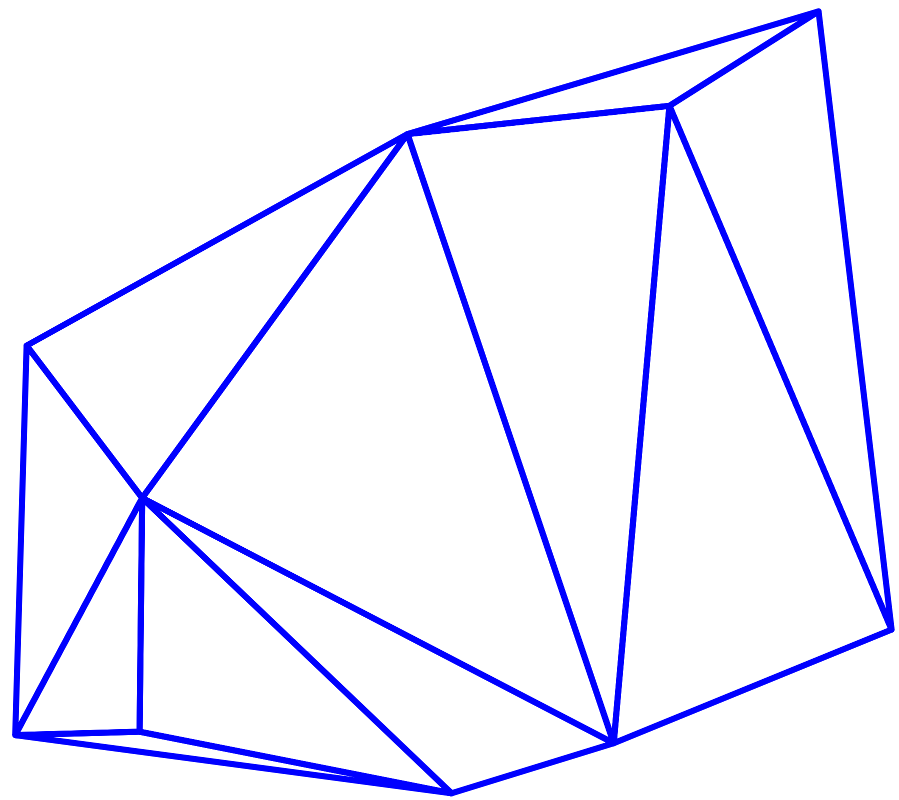

Delaunay triangulation of 10 random points,

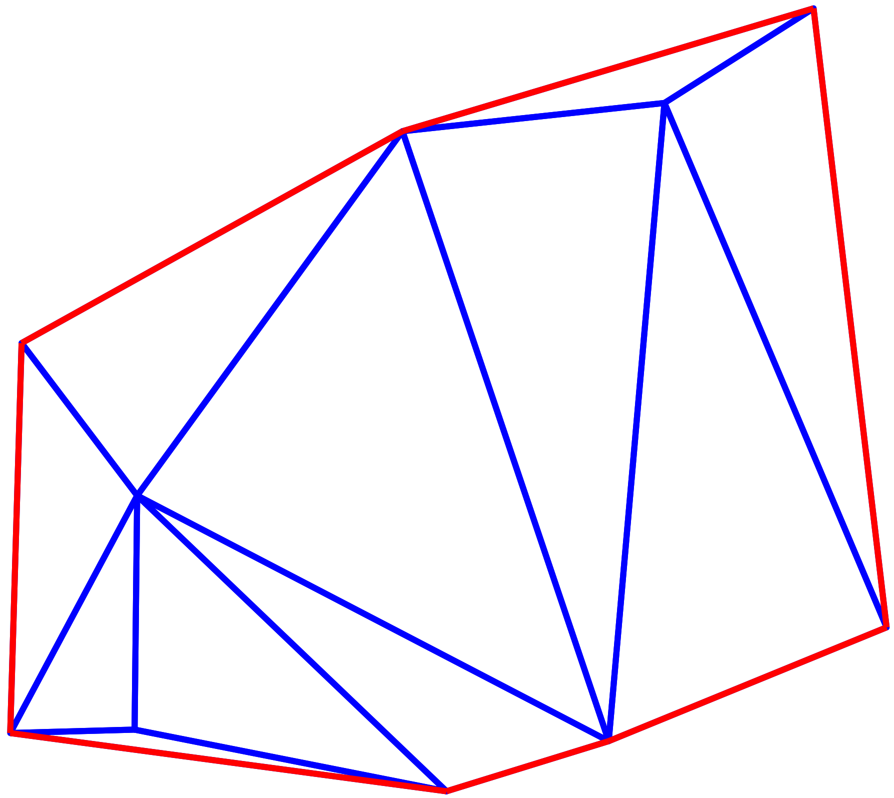

The boundary of the triangulation

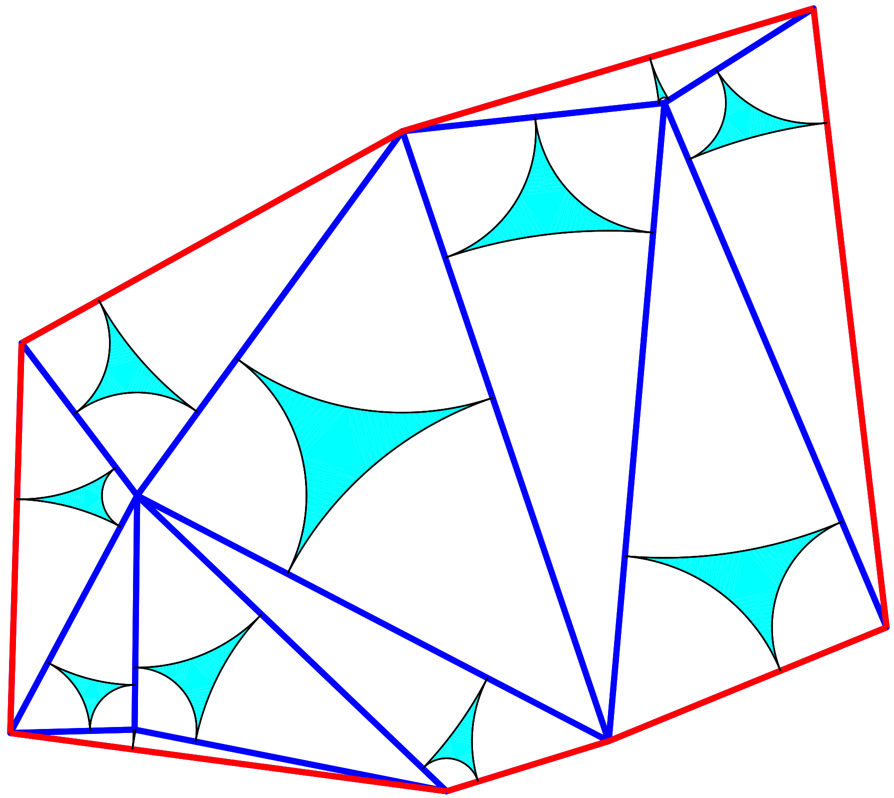

The central regions.

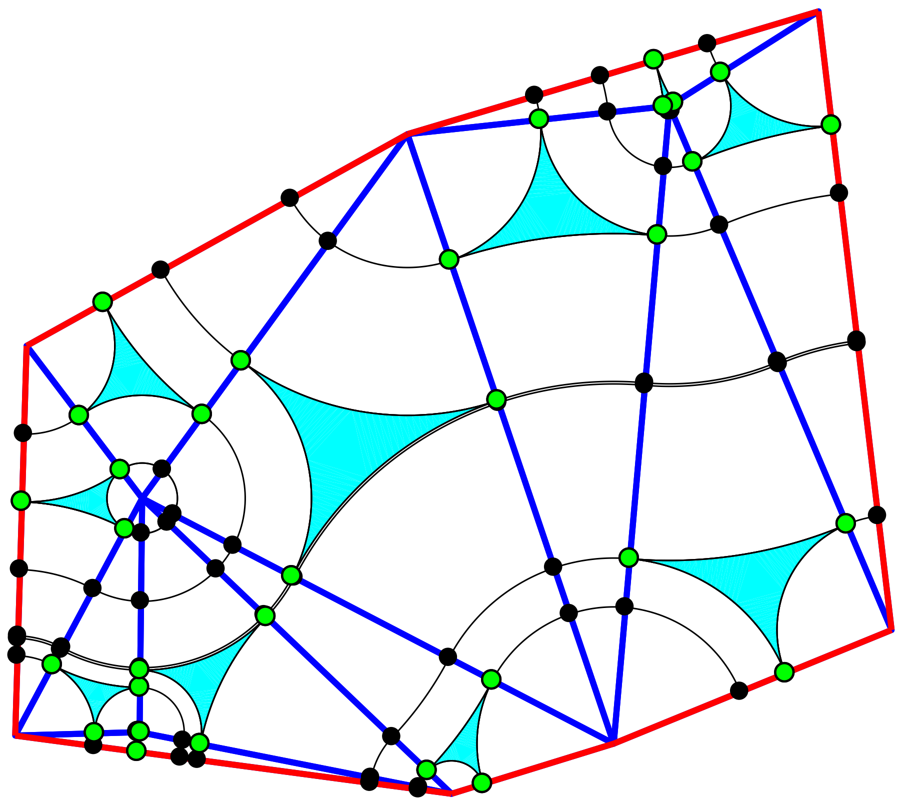

Propagation lines starting at all cusp points.

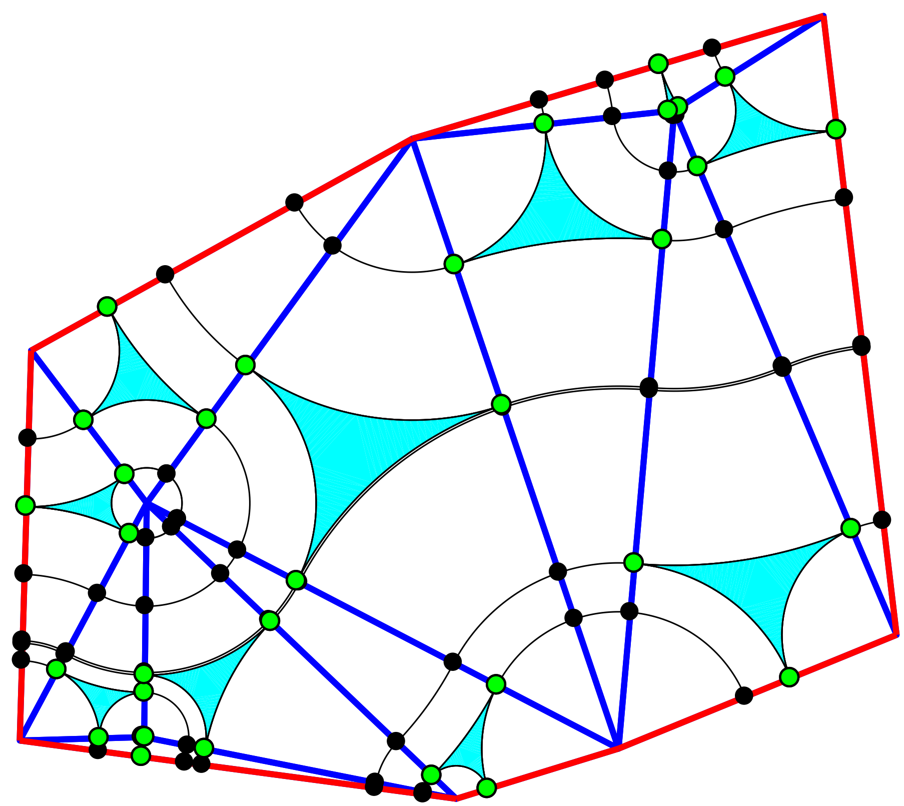

Propagation lines identify boundary points; induces tree.
Discontinuous, but piecewise length preserving.

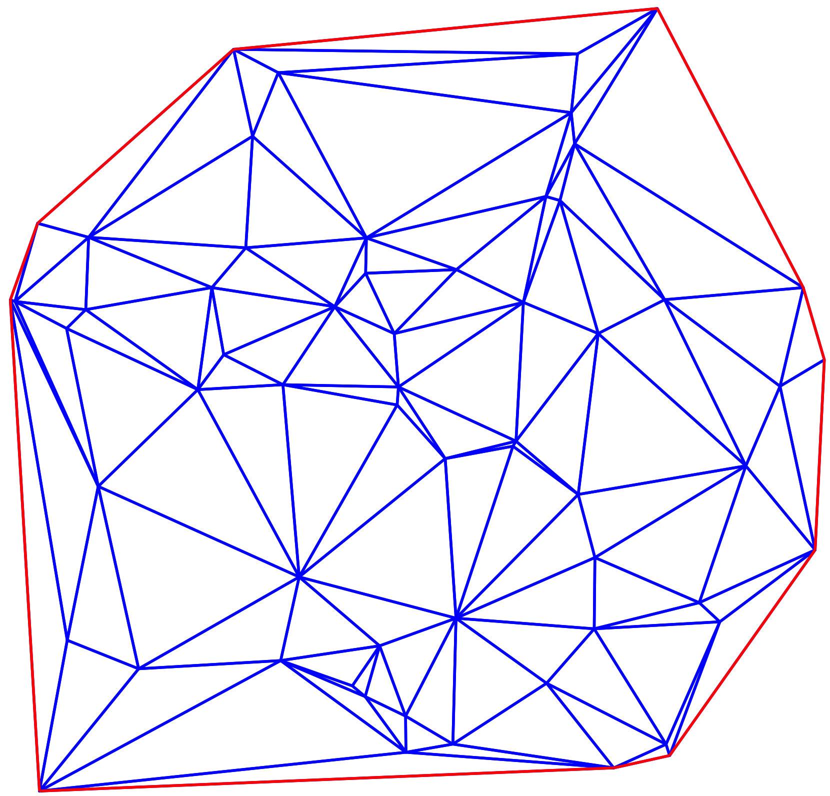

60 points

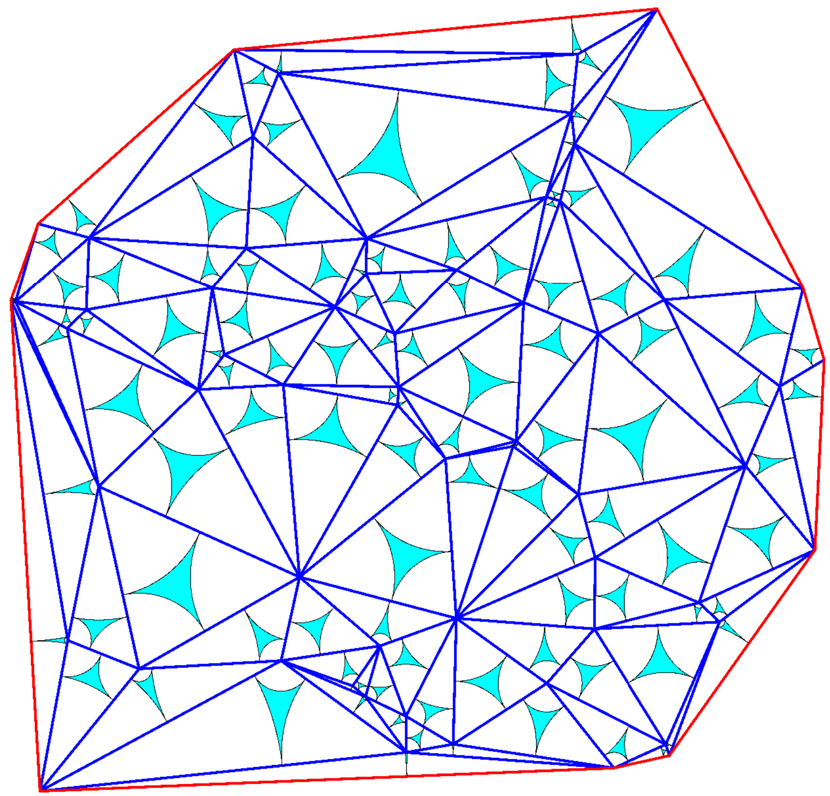

The central regions.

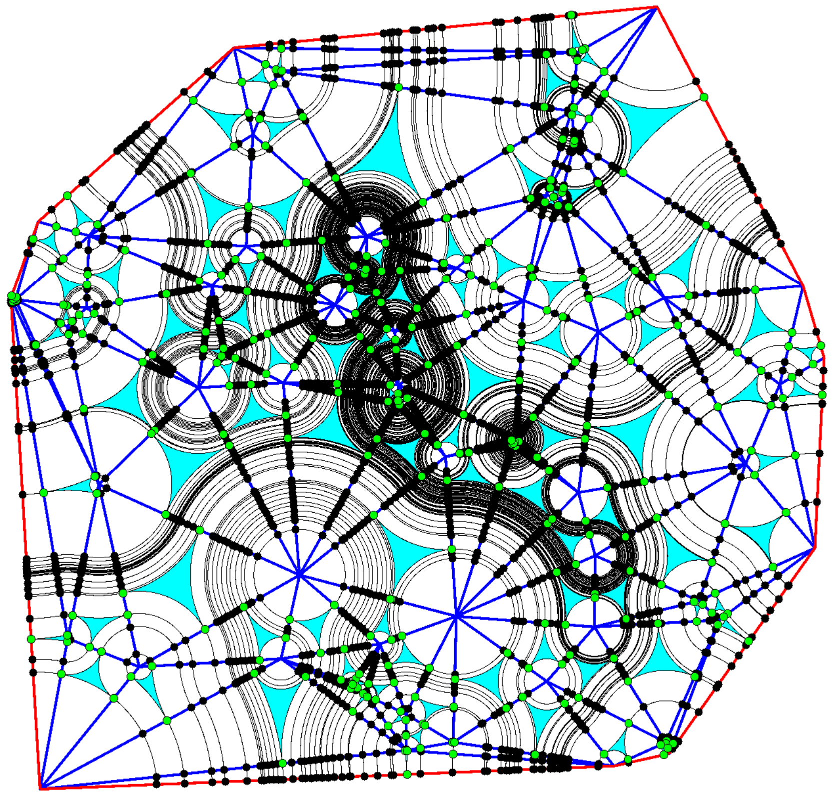

Propagation lines starting at all cusp points.

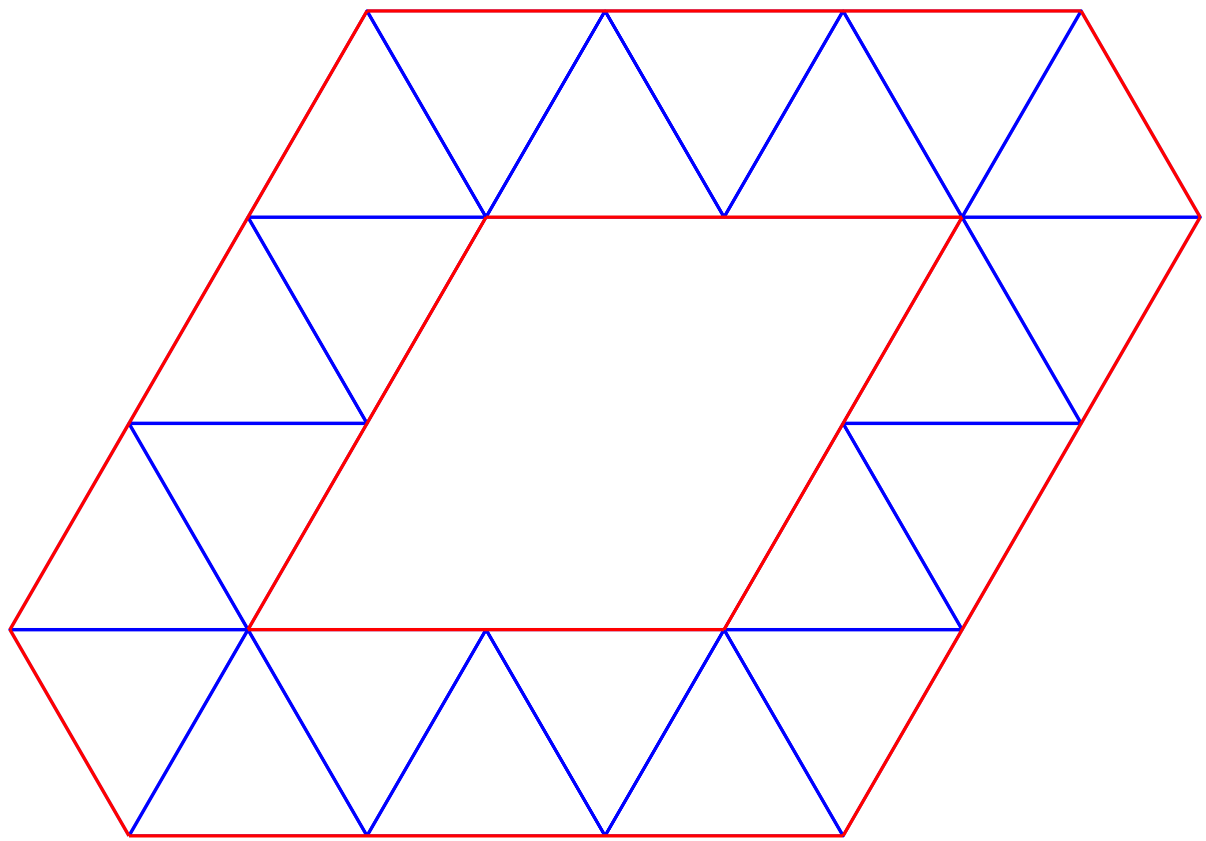

A ring of equilateral triangles.

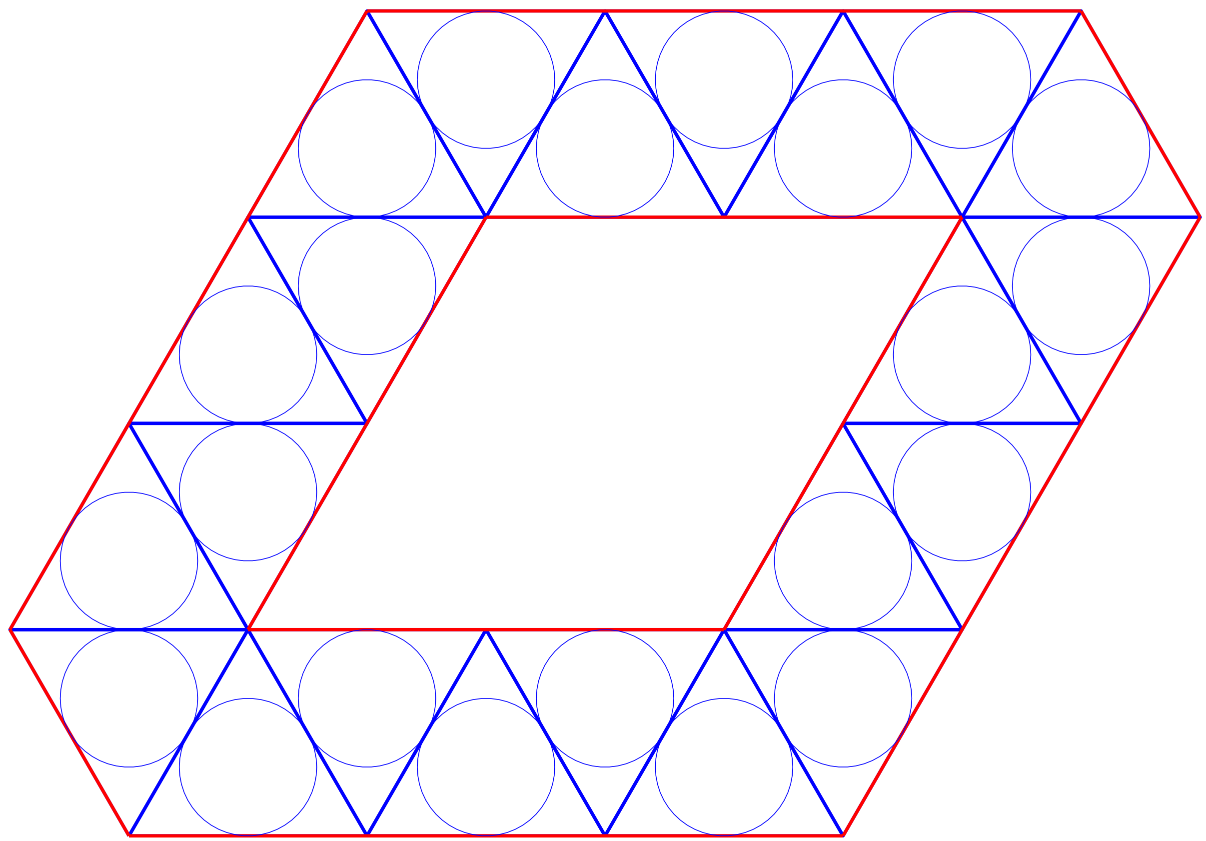

The in-circles are tangent.

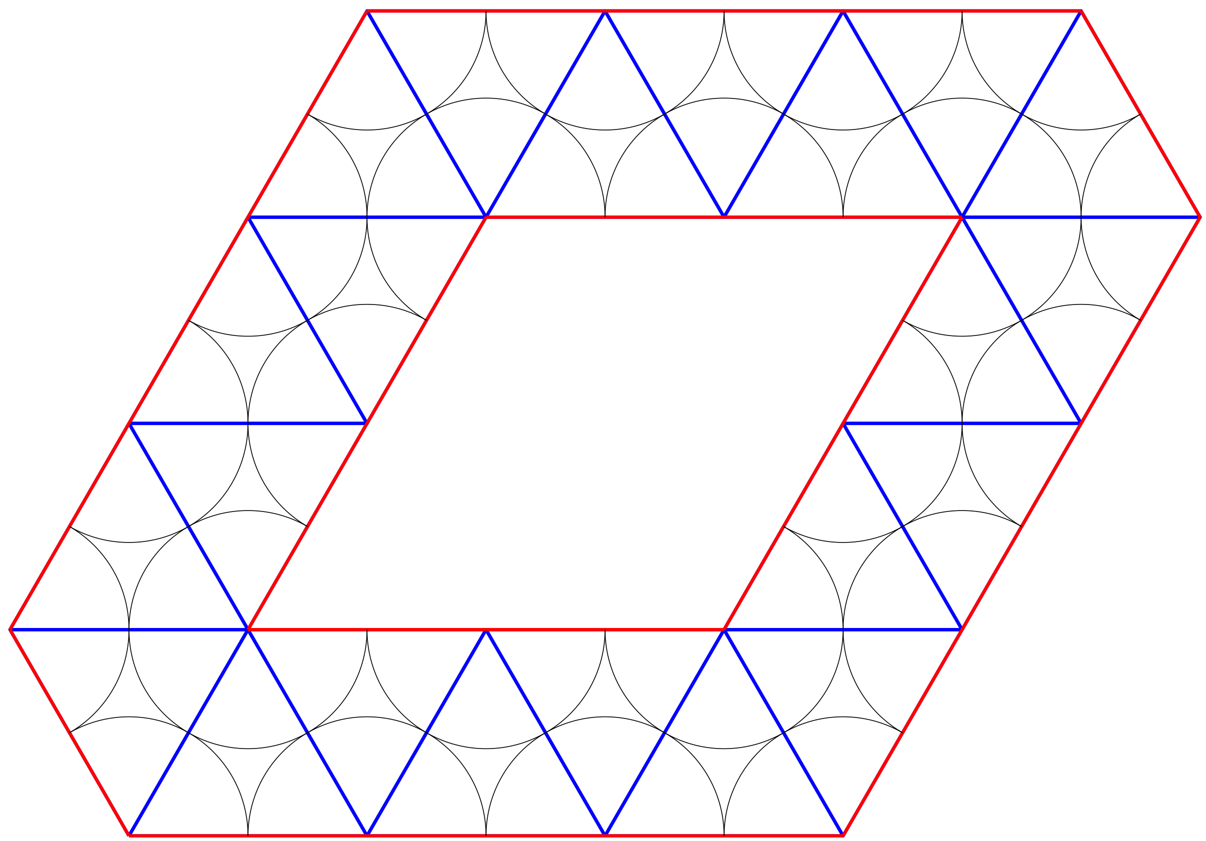

The the cusps touch; for a closed flow line.

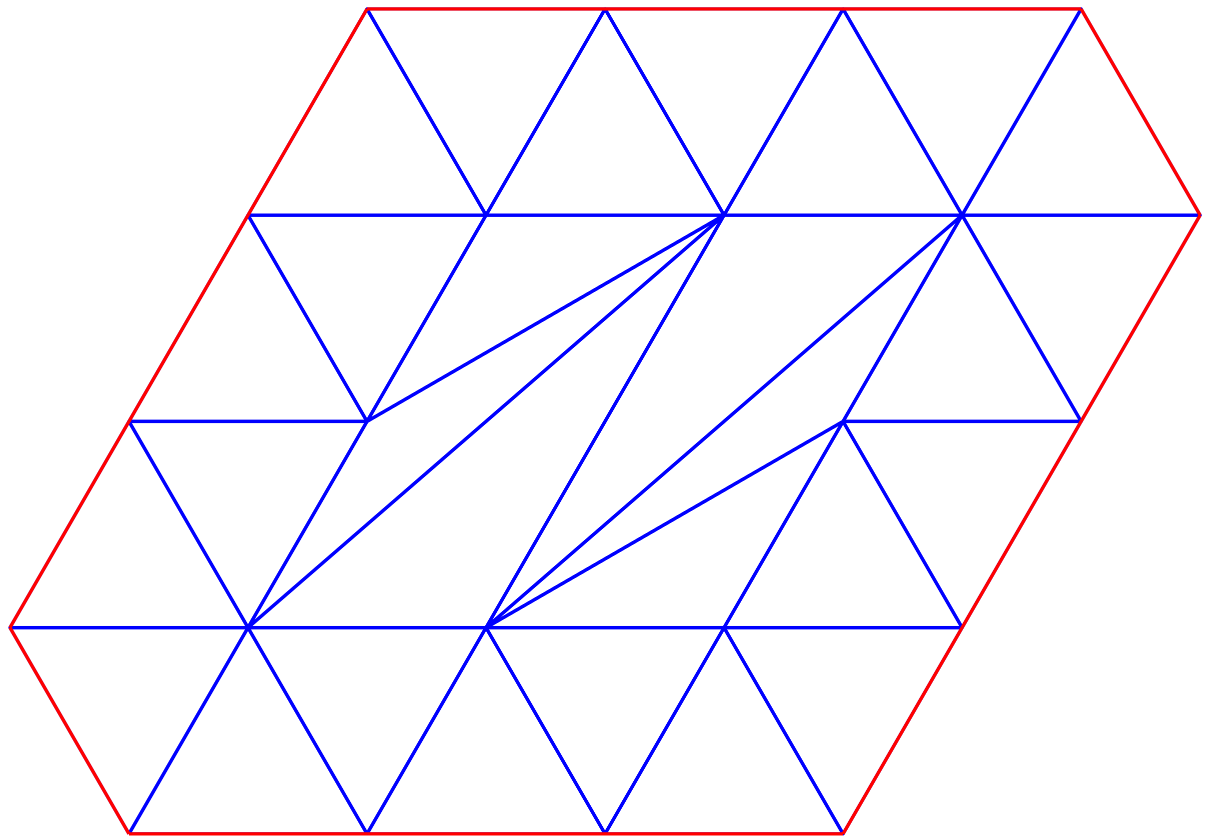

No matter how we triangulate interior, flow lines never exit.

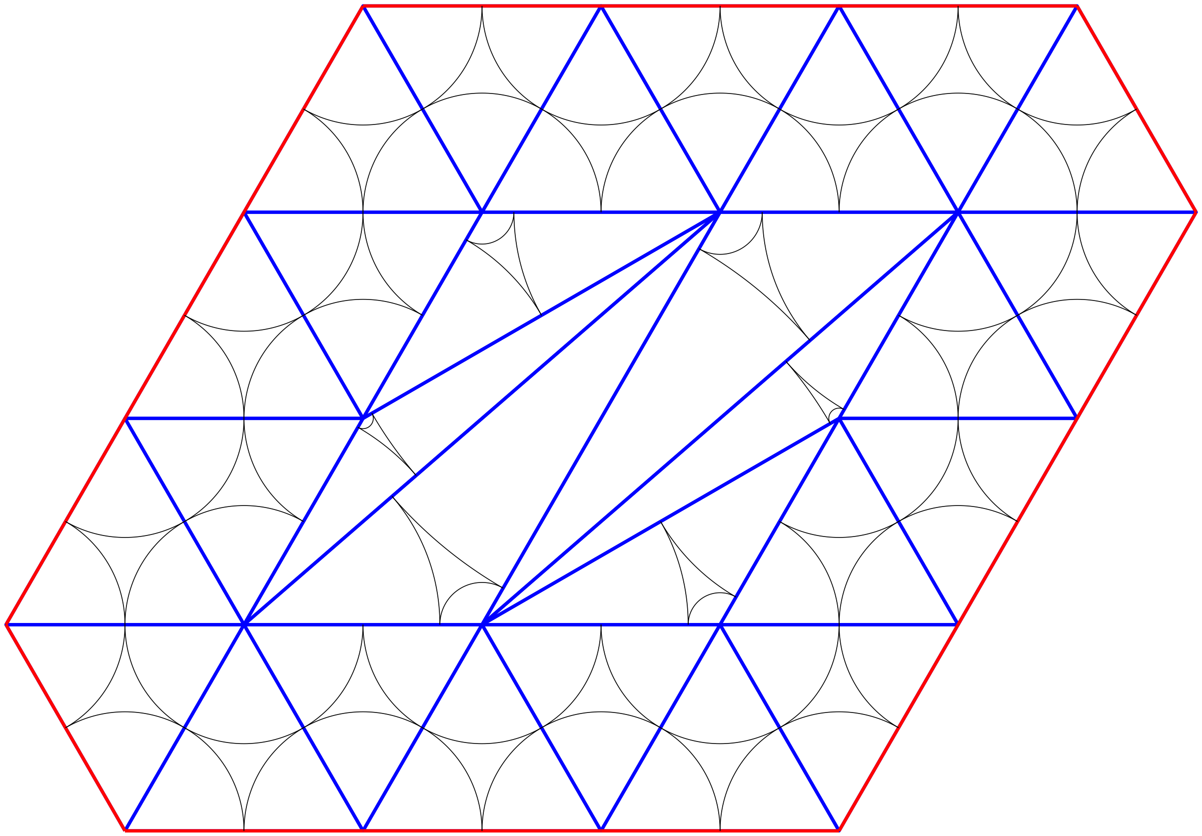

No matter how we triangulate interior, flow lines never exit.

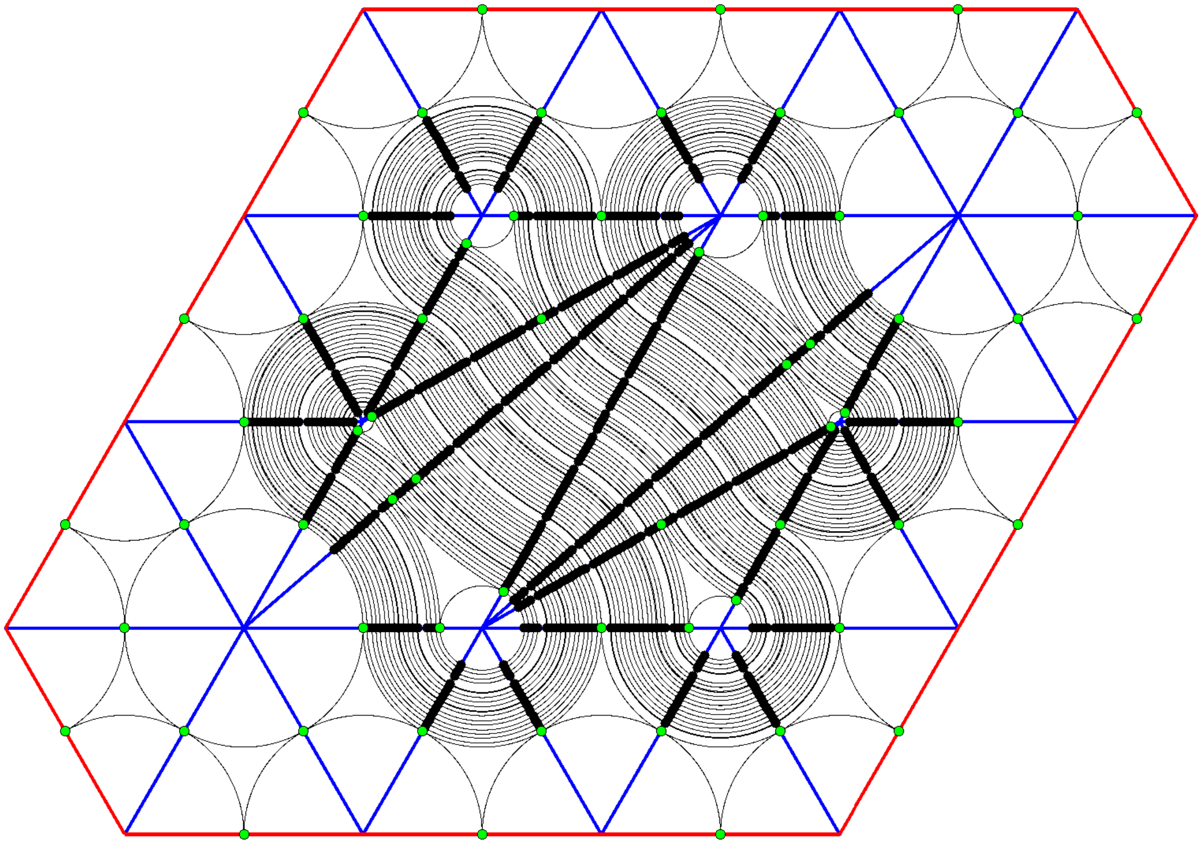

No matter how we triangulate interior, flow lines never exit.

Idea: perturb paths to terminate sooner.

If a path returns to same thin edge at least 3 times it has a sub-path that looks like one of these:

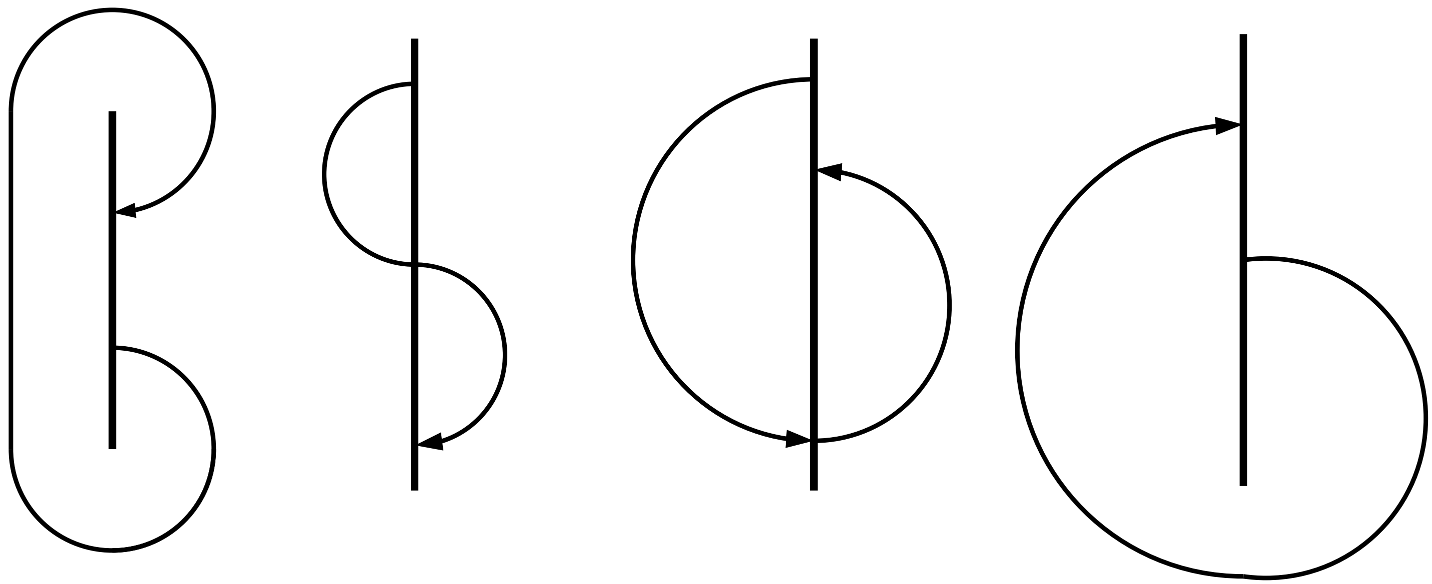

C-curve, S-curve, G-curves

Return region consists of paths "parallel" to one of these.

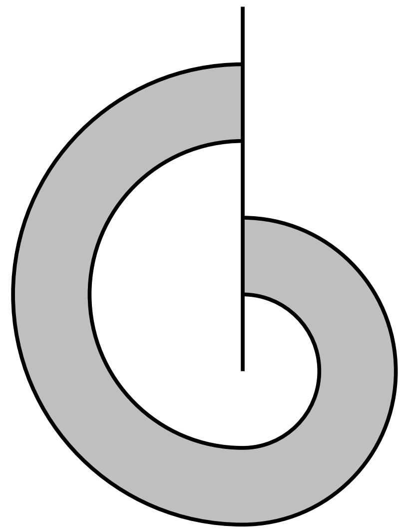

Return region consists of paths "parallel" to one of these.

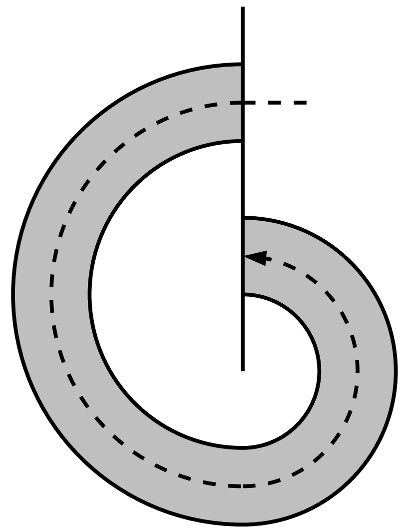

There are $O(n)$ return regions and every propagation path enters one after crossing at most $O(n)$ thin parts.

Return region consists of paths "parallel" to one of these.

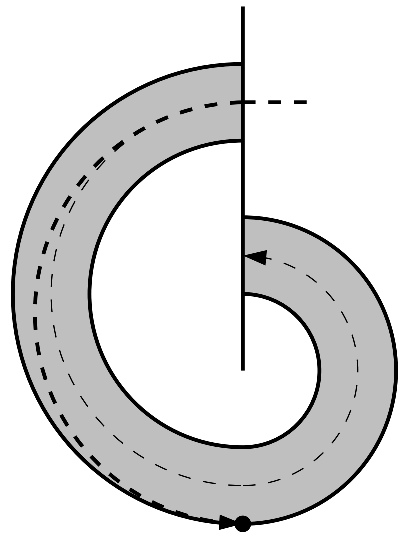

There are $O(n)$ return regions and every propagation path enters one after crossing at most $O(n)$ thin parts.

IDEA: bend paths to terminate before they exit.

Gives $O(n^2)$ if it works.

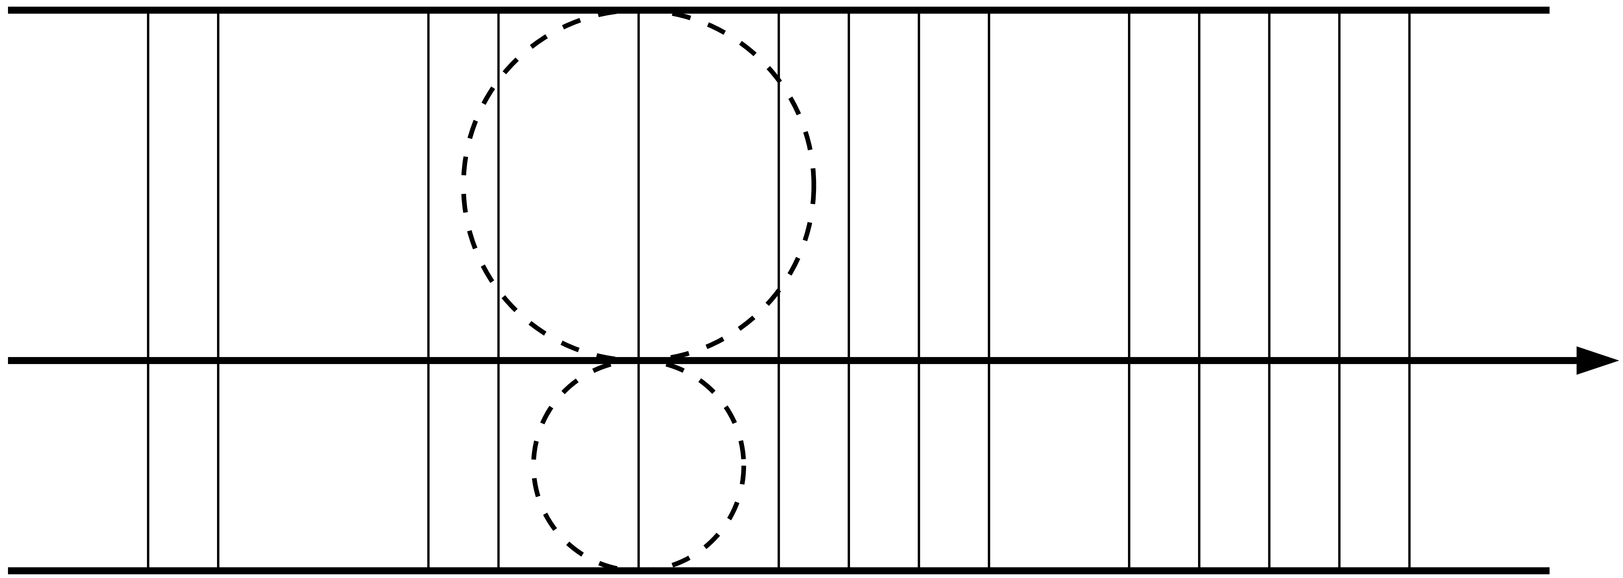

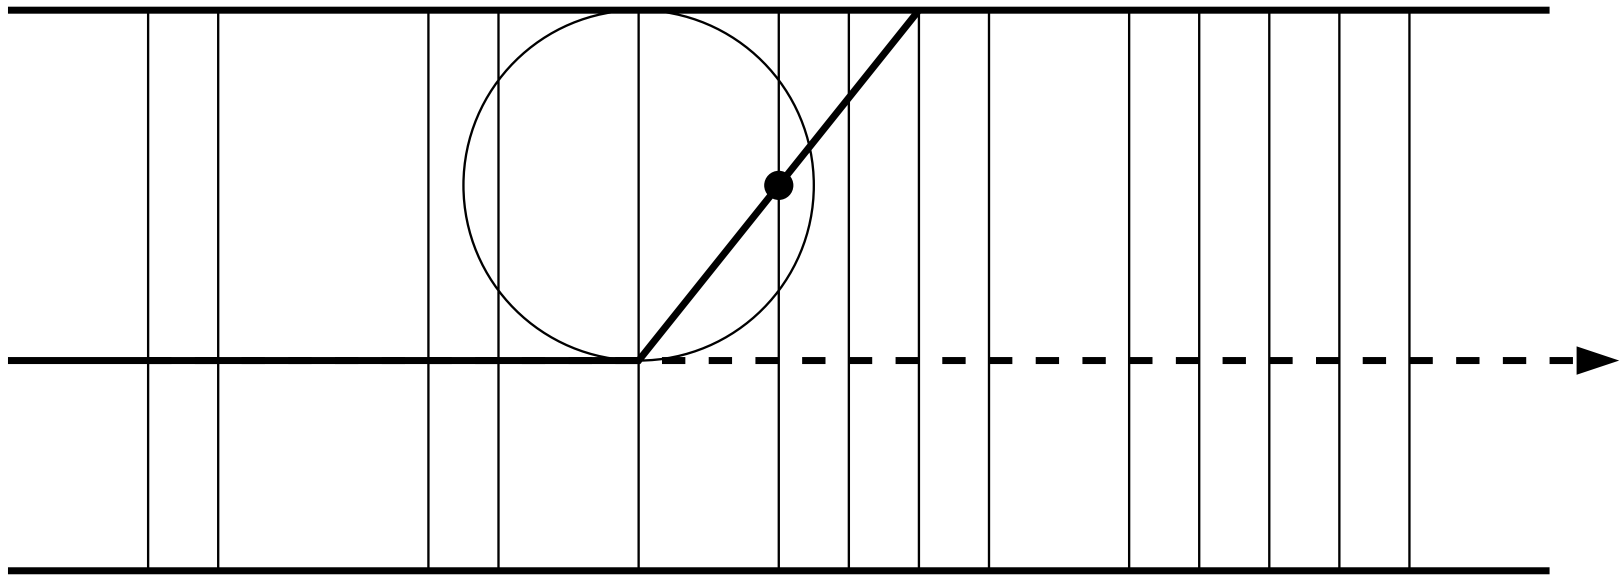

If path bends too fast, Gabriel condition can fail.

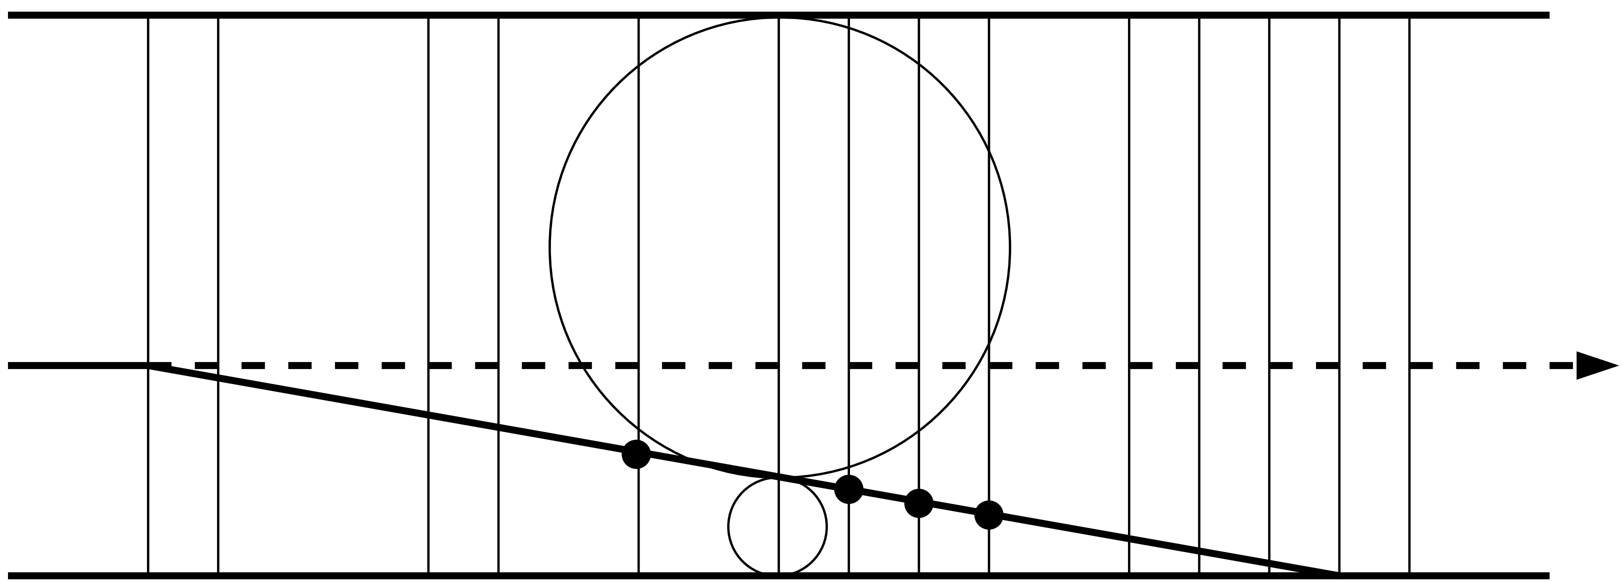

Bend slowly enough to satisfy Gabriel condition.

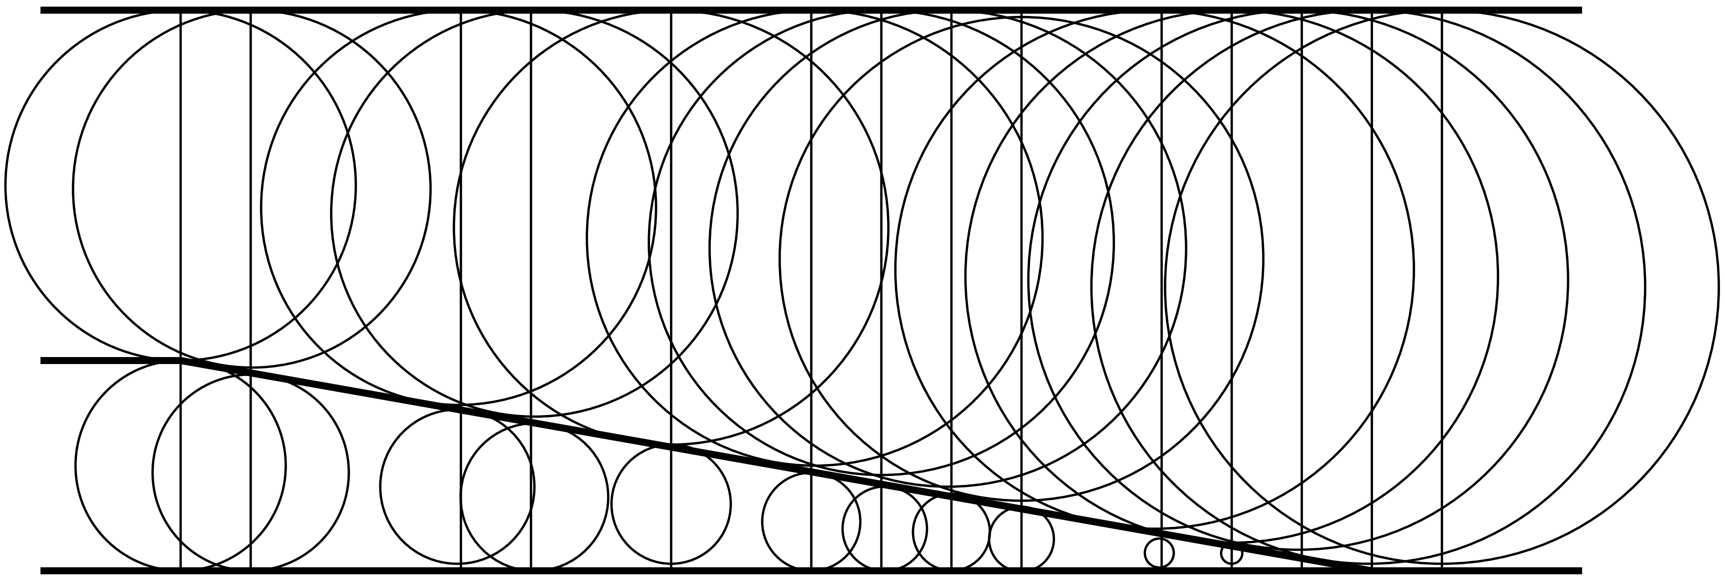

Bend slowly enough to satisfy Gabriel condition.

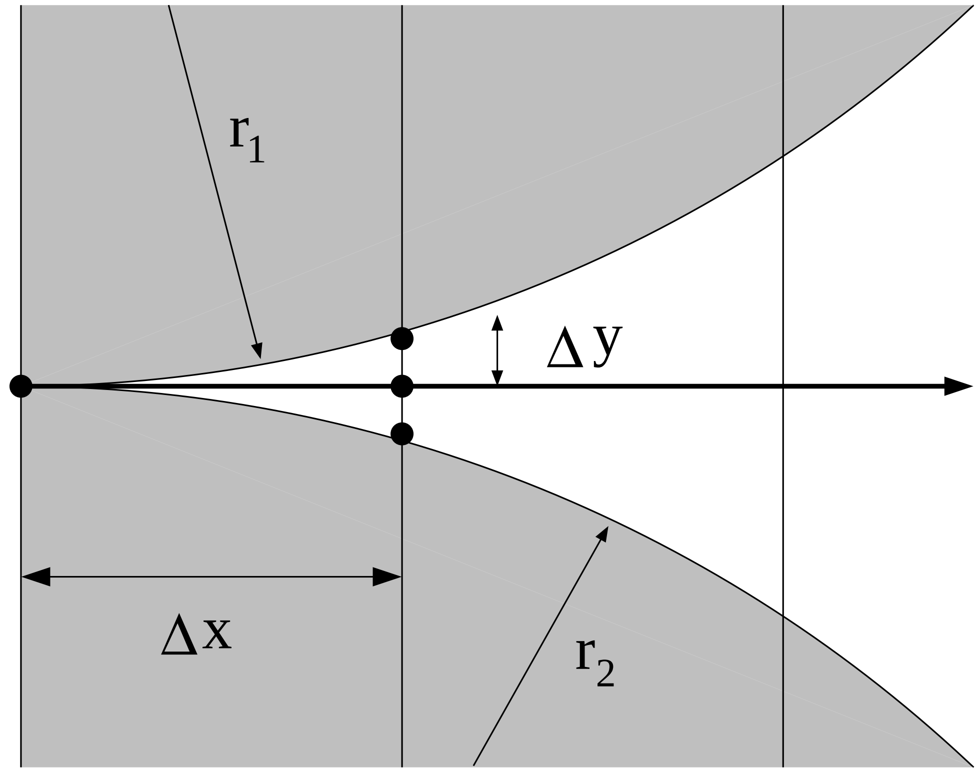

$$\Delta y \approx (\Delta x/r)^2 r = (\Delta x)^2/r.$$

$$r = \max(r_1, r_2).$$

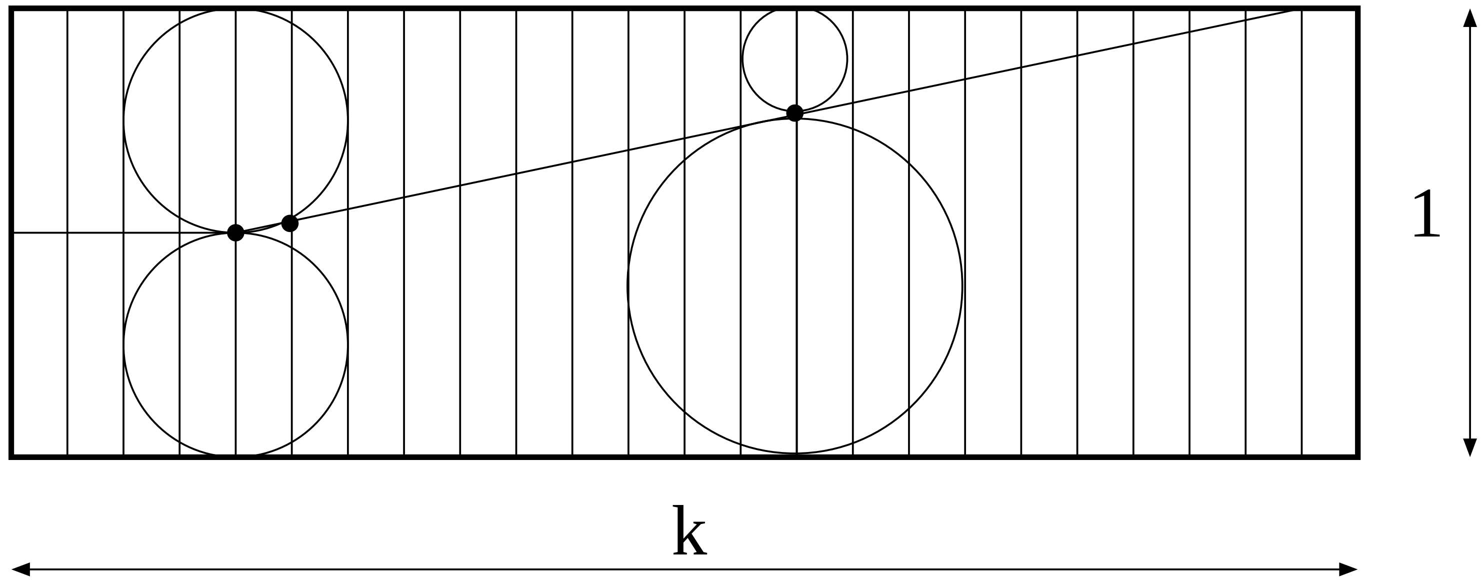

$k \times 1$ region crossing n (equally spaced) thin parts,

$r \approx 1, \quad \Delta x \approx k/n, \quad \Rightarrow \quad \Delta y \approx k^2/n^2$

Need $1 \leq \sum \Delta y = n\Delta y = k^2/n$.

Bent path hits side of region if $k \gg \sqrt{n}$.

Easy case: return region length > width.

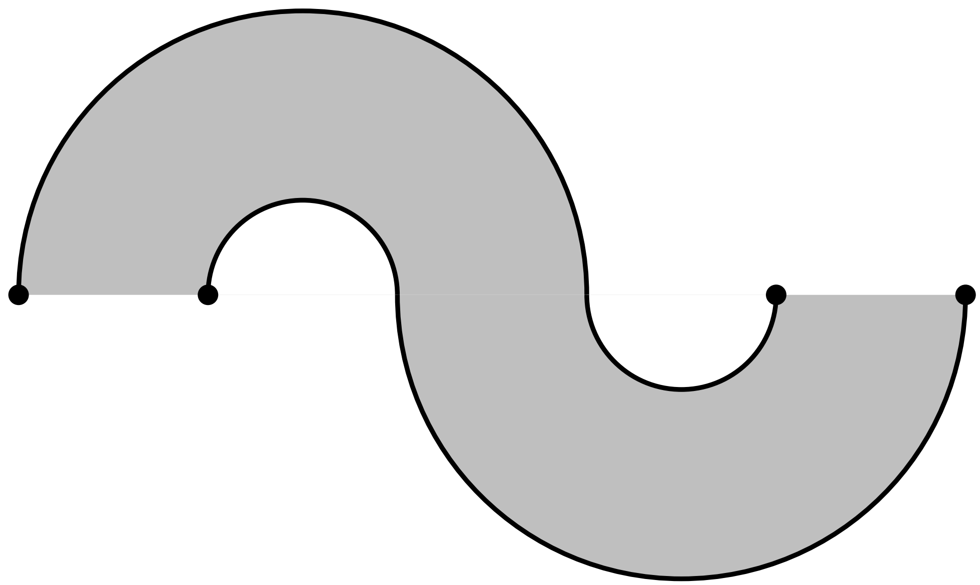

- Show there are $O(n)$ return regions.

Easy case: return region length > width.

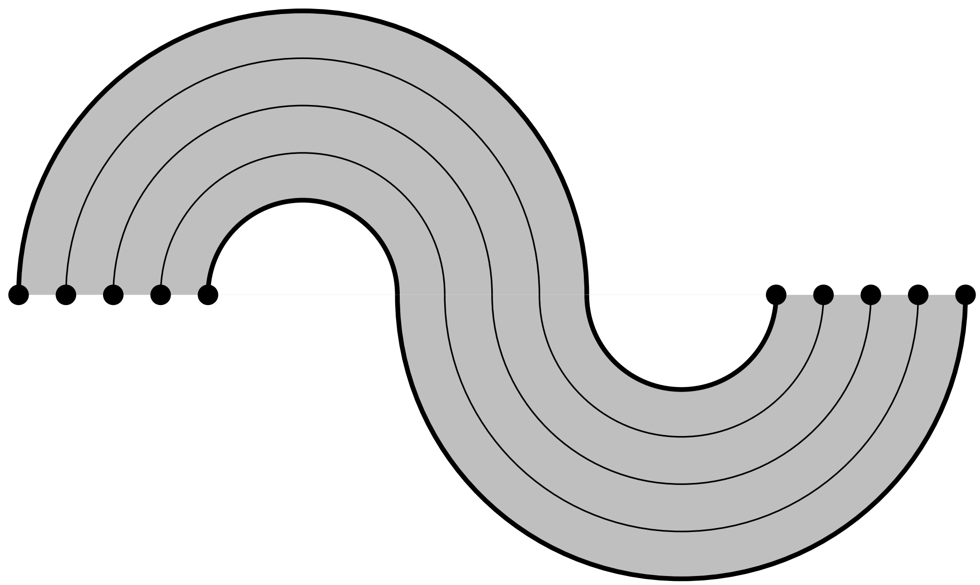

- Show there are $O(n)$ return regions.
- Divide each region into $O(\sqrt{n})$ long parallel tubes.

Easy case: return region length > width.

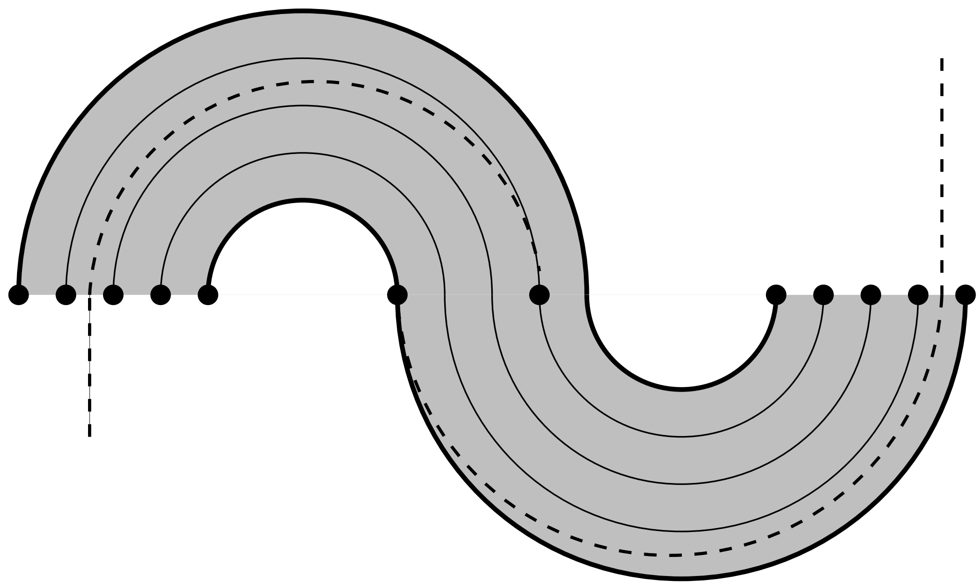

- Show there are $O(n)$ return regions.
- Divide each region into $O(\sqrt{n})$ long parallel tubes.
- Entering paths can be bent and terminated.
  Total vertices created $= O(n^2)$, but ...

Easy case: return region length $>$ width.

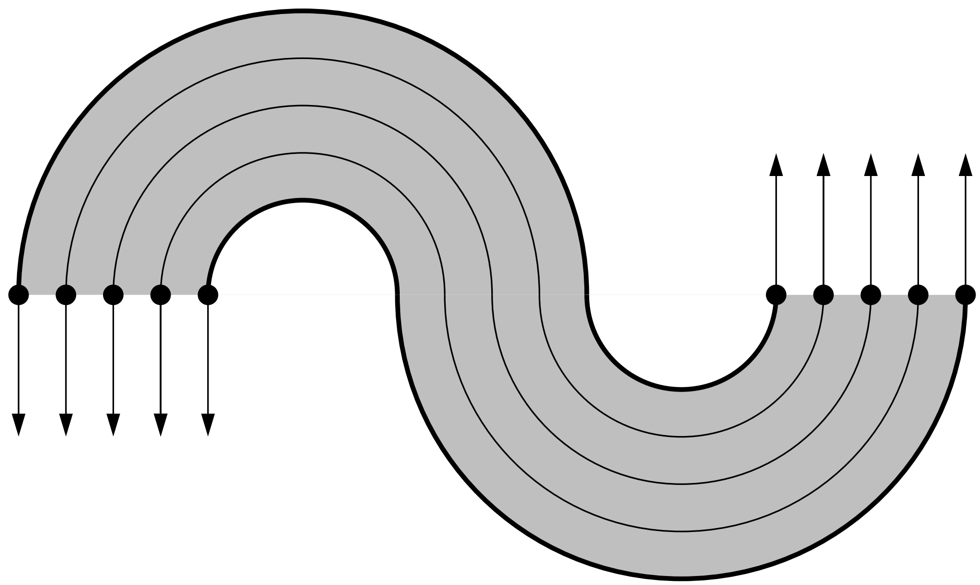

- Show there are $O(n)$ return regions.
- Divide each region into $O(\sqrt{n})$ long parallel tubes.
- Entering paths can be bent and terminated.
  Total vertices created $= O(n^2)$, but ...
- Each region has $O(\sqrt{n})$ new vertices to propagate.
  Vertices created is $O(\sqrt{n} \cdot n^2) = O(n^{2.5})$.

Hard case is spirals:

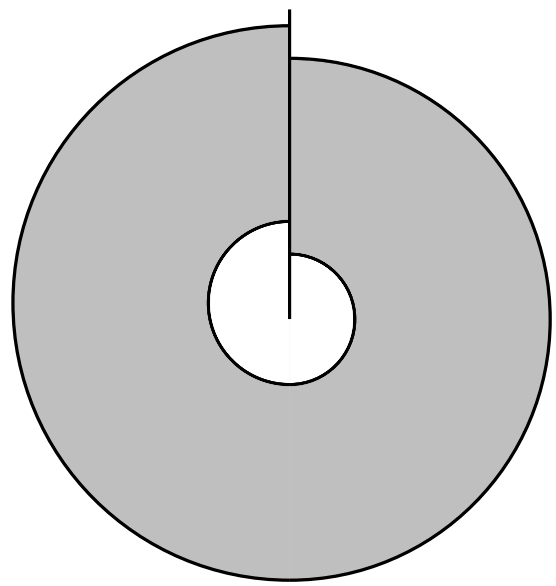

Hard case is spirals:

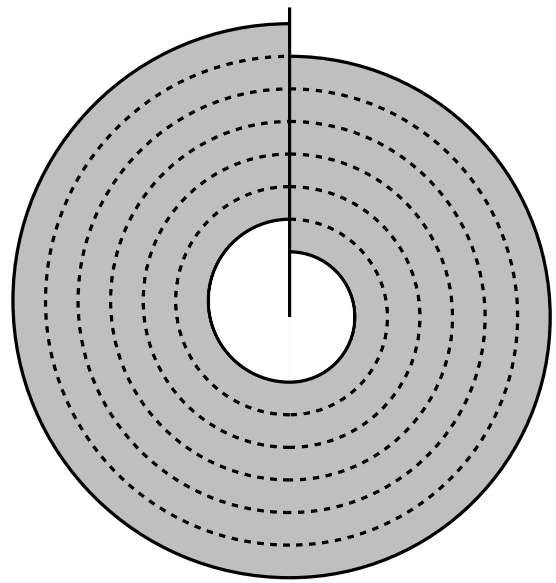

Curves may spiral arbitrarily often.

No curve can be allowed to pass all the way through the spiral. We stop them in a multi-stage construction.

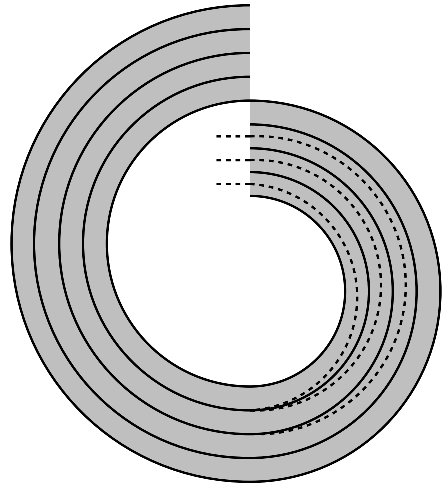

- Start with $\sqrt{n}$ parallel tubes at entrance of spiral.
  Terminate entering paths (1 spiral).

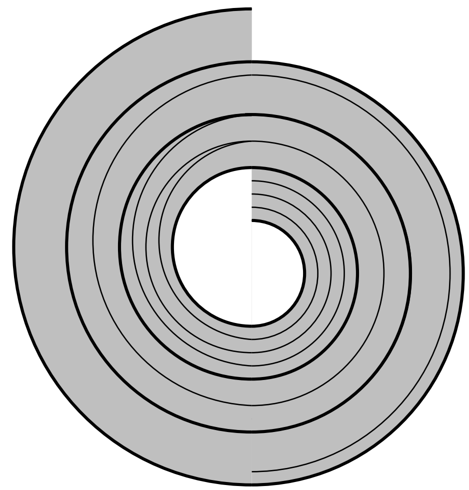

- Start with $\sqrt{n}$ parallel tubes at entrance of spiral. Terminate entering paths (1 spiral).
- Merge $\sqrt{n}$ tubes to single tube ($n^{1/3}$ spirals). (get longer as we move out.)

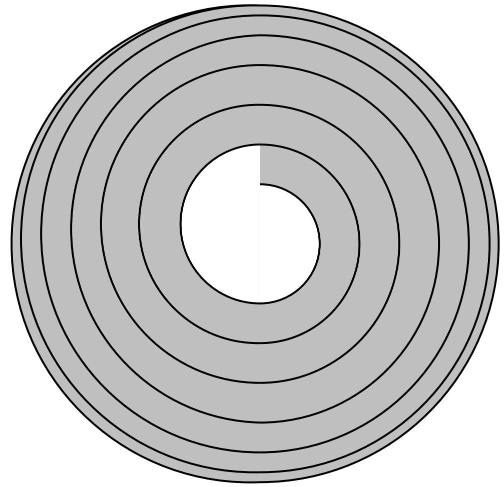

- Start with $\sqrt{n}$ parallel tubes at entrance of spiral. Terminate entering paths (1 spiral).
- Merge $\sqrt{n}$ tubes to single tube ($n^{1/3}$ spirals). (spirals get longer as we move out.)
- Make tube edge self-intersect ($n^{1/2}$ spirals)

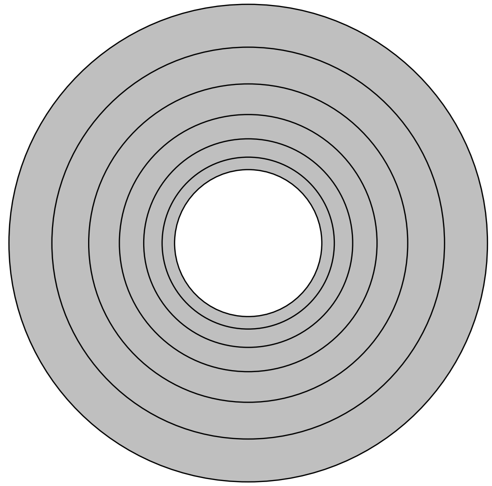

- Start with $\sqrt{n}$ parallel tubes at entrance of spiral. Terminate entering paths (1 spiral).
- Merge $\sqrt{n}$ tubes to single tube ($n^{1/3}$ spirals). (spirals get longer as we move out.)
- Make tube edge self-intersect ($n^{1/2}$ spirals)
- Loops with increasing gaps ($n^{1/2}$ loops, n spirals)
- Beyond radius n spiral is empty.

Careful estimates needed to get $O(n^{2.5})$ vertices.

Almost Nonobtuse Triangulation: Replace cusps by cones of angle ϵ. Same construction in thick parts.

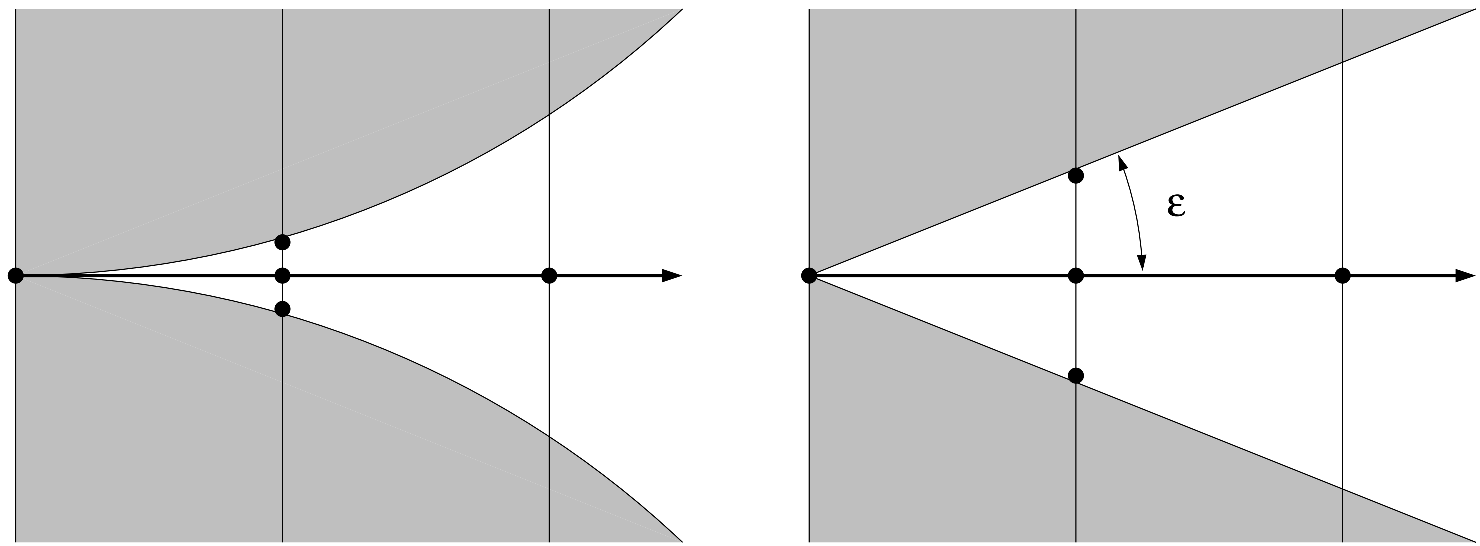

Paths can be terminated inside a $1 \times \frac{1}{\epsilon}$ tube.

Thm: Uses angles $\leq 90^\circ + \epsilon$ and $O(\frac{n^2}{\epsilon^2})$ triangles.

Summary:

Delaunay triangulation (DT) is easy to compute but is non-conforming and allows large angles.

We can add $O(n^{1.5})$ new points and perturb foliation paths to end quickly. Gives $O(n^{2.5})$ Gabriel edges.

From this we get a $O(n^{2.5})$ conforming NOT-DT.

DT or NOT-DT, that is the question.
Whether 'tis nobler in the mind to suffer
the slings and arrows of obtuse angles,
or take arms against a sea of paths,
and by perturbing end them?

Questions:

- Implementable?
- Average versus worst case bounds?
- Improve 2.5 to 2 in NOT theorem?
- Relation to closing lemma for vector fields?
- 3-D meshes? The eightfold way? Ricci flow?
- Polynomial meshes using acute tetrahedra? Even cube is hard.
- 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?

More Questions:

- Does Davis' method converge?
- Does CRDT converge?
- Classify sinks.
- Do NOTs and Delaunay triangulations have same complexity?
- Are rectangular grid sub-meshes good for anything?