RANDOM THOUGHTS ON RANDOM SETS

Christopher Bishop, Stony Brook

ANALYSIS AND GEOMETRY OF RANDOM SHAPES

IPAM, January 7, 2019

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

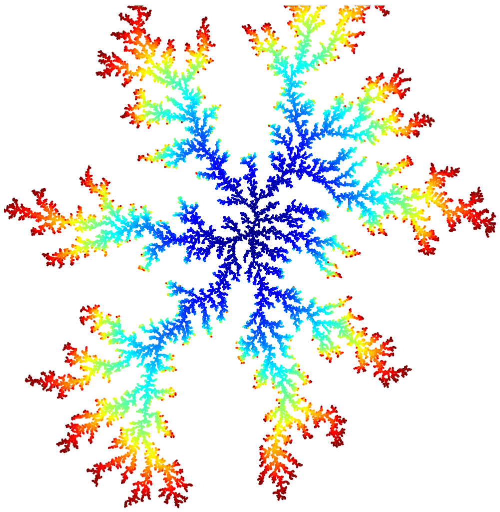

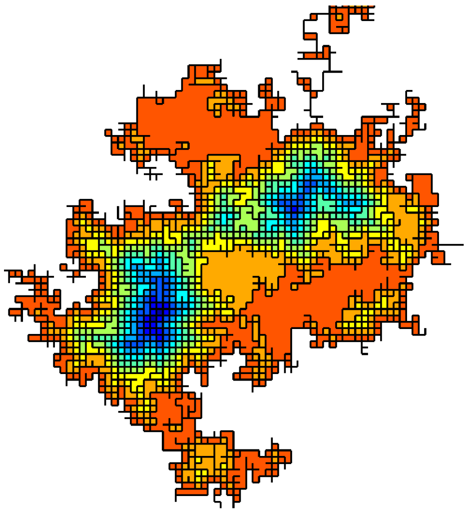

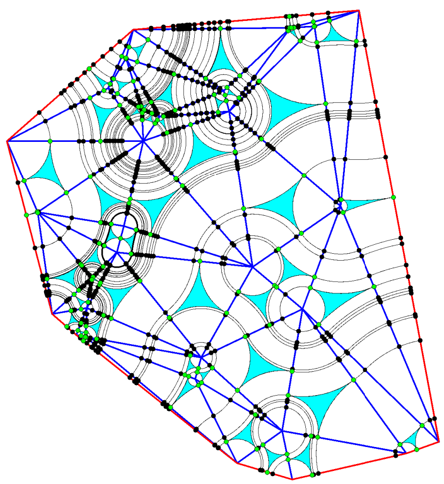

PART I: DIFFUSION LIMITED AGGREGATION (DLA)

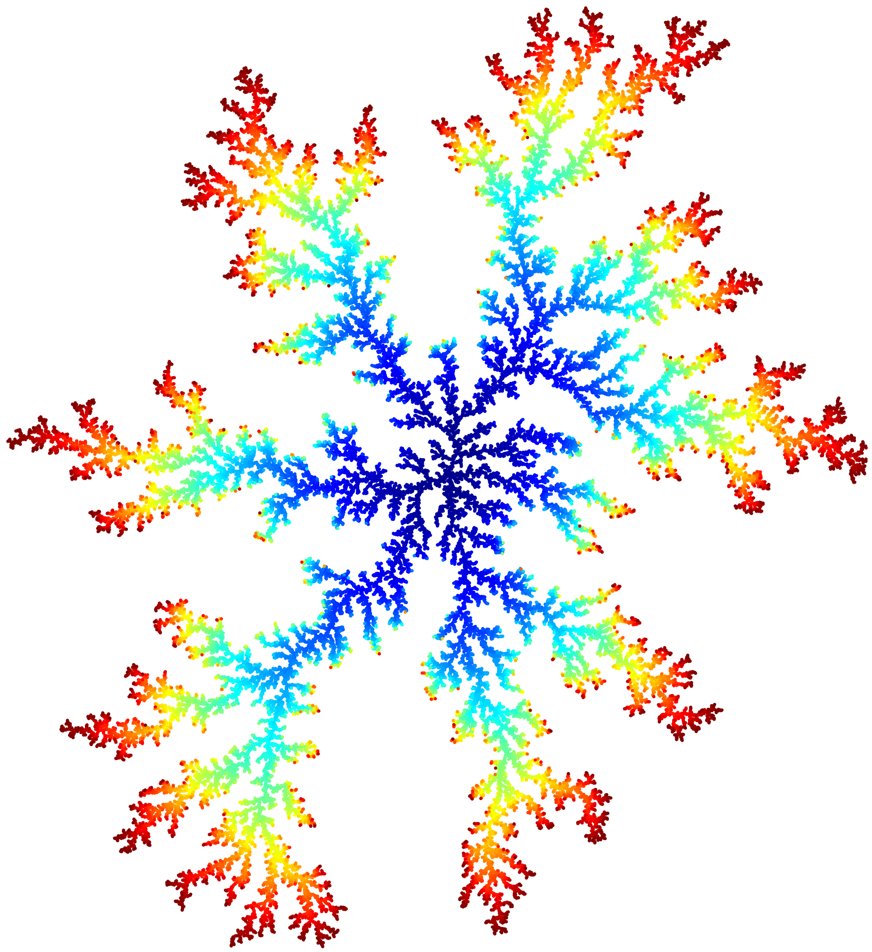

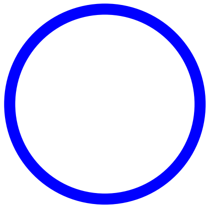

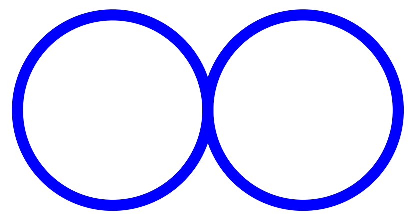

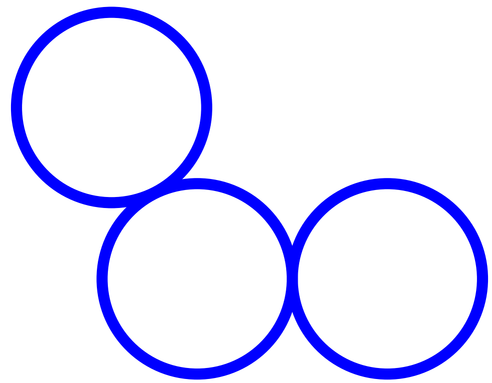

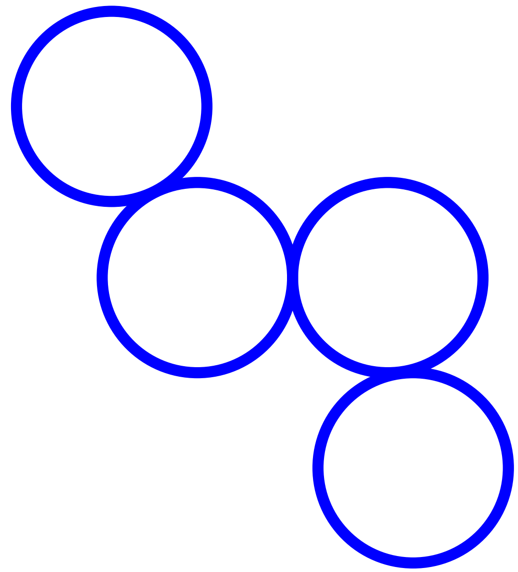

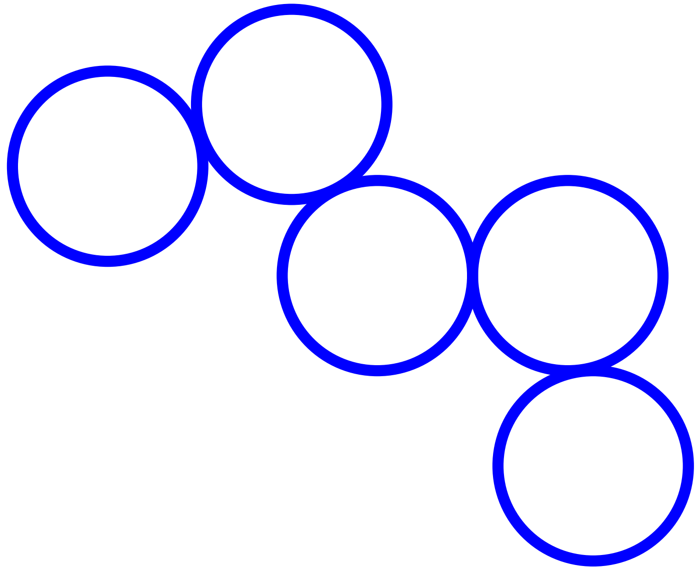

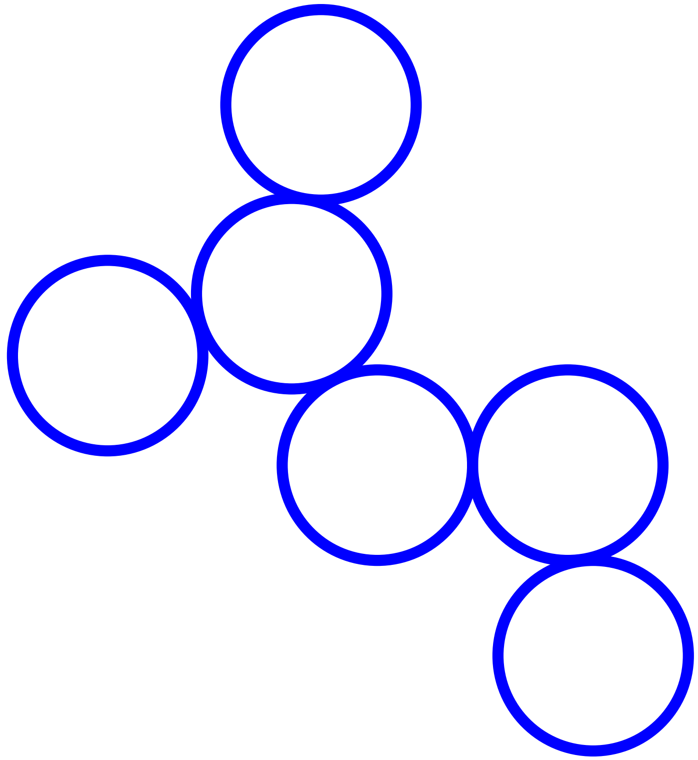

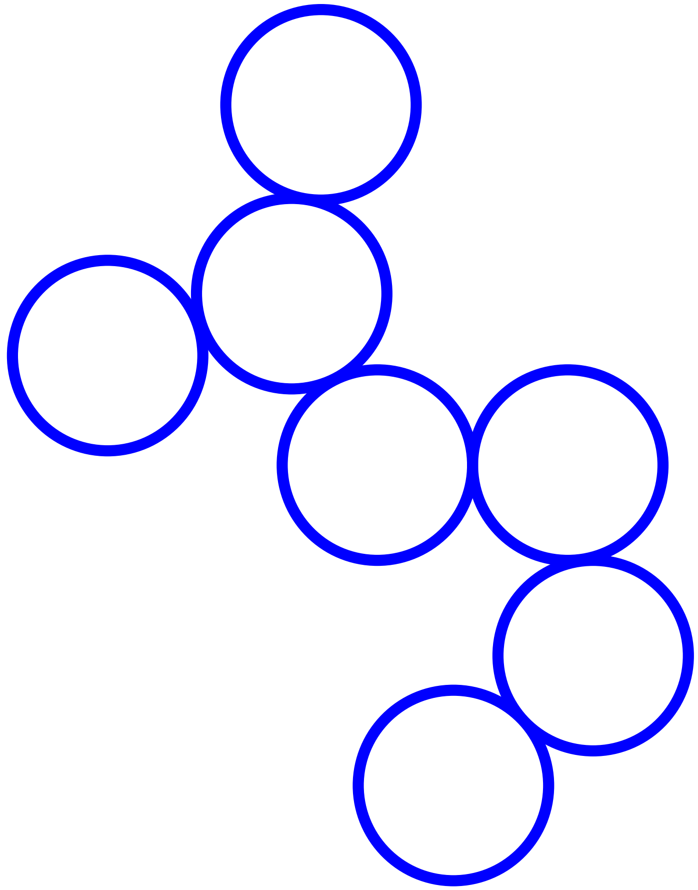

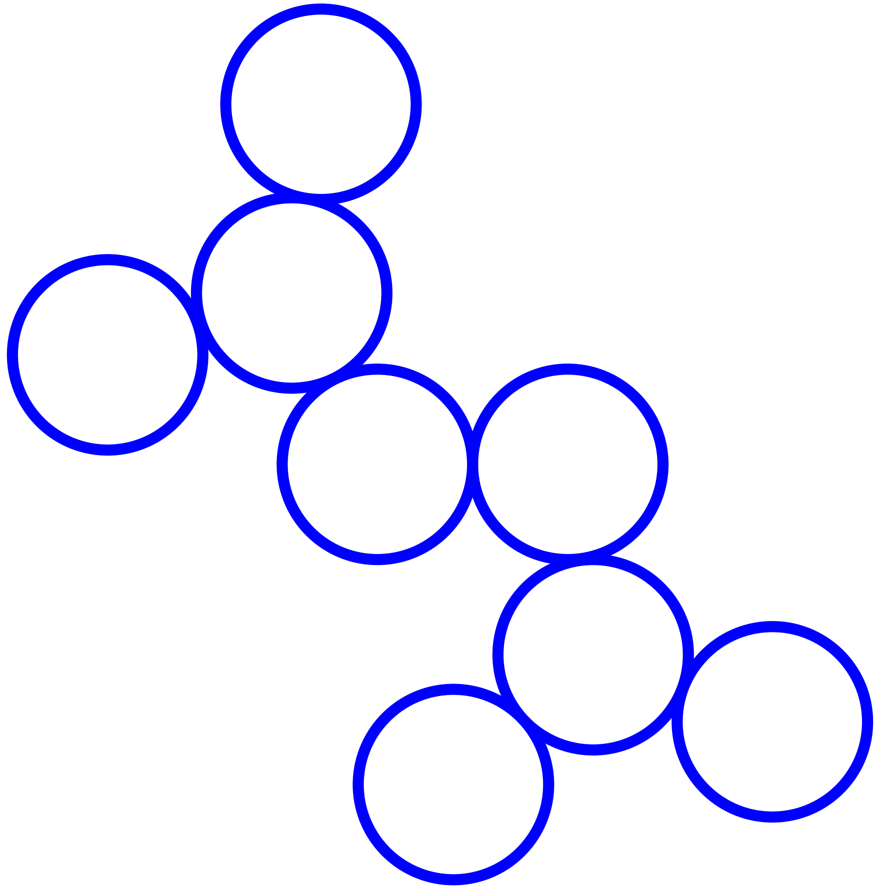

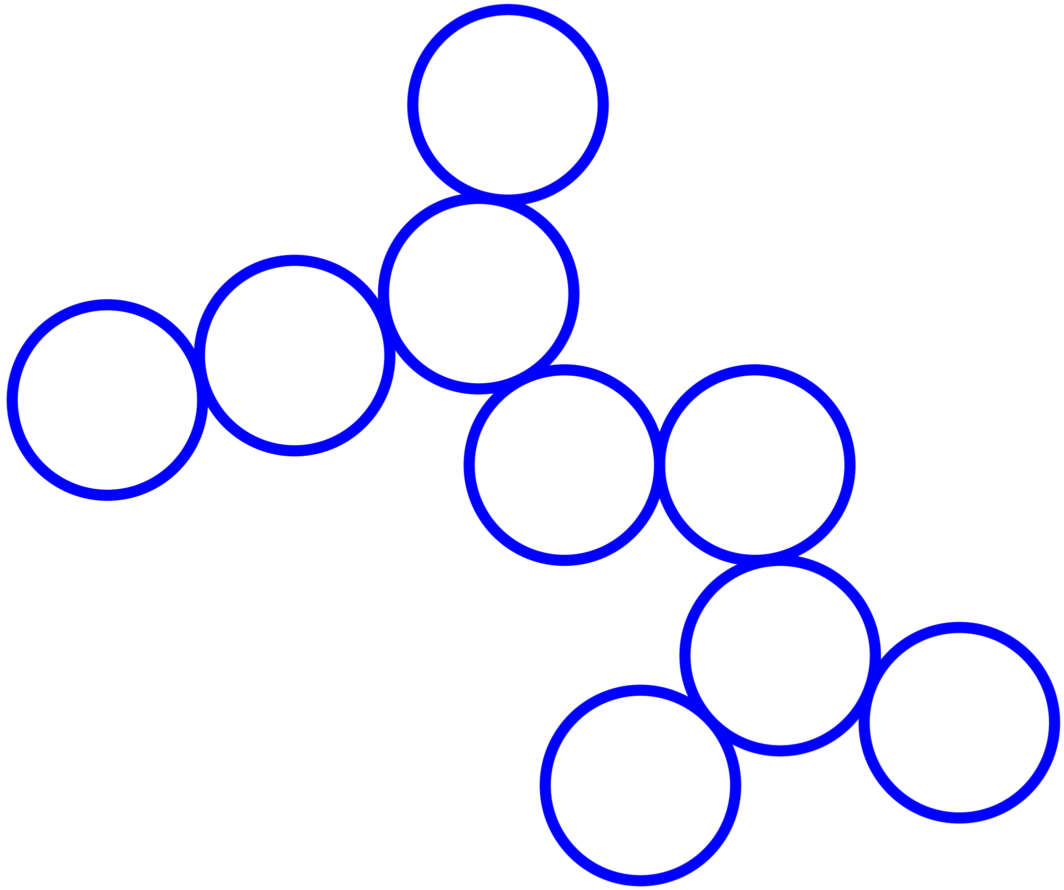

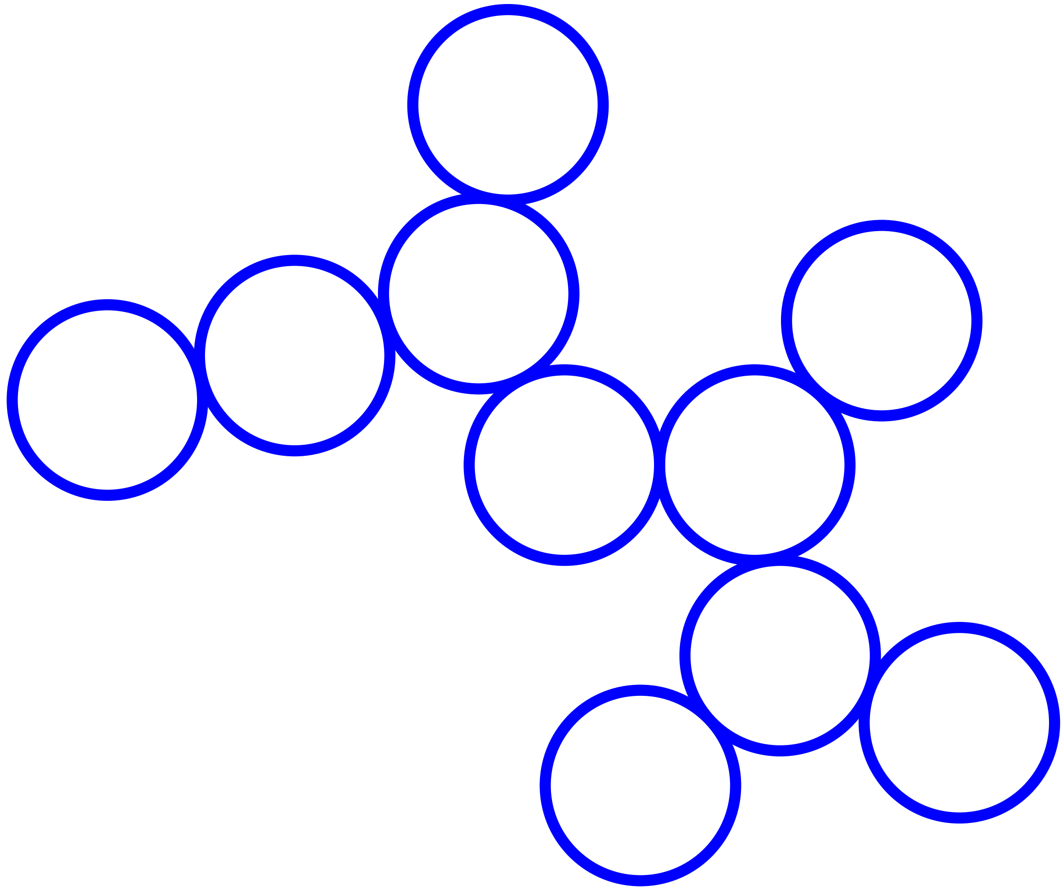

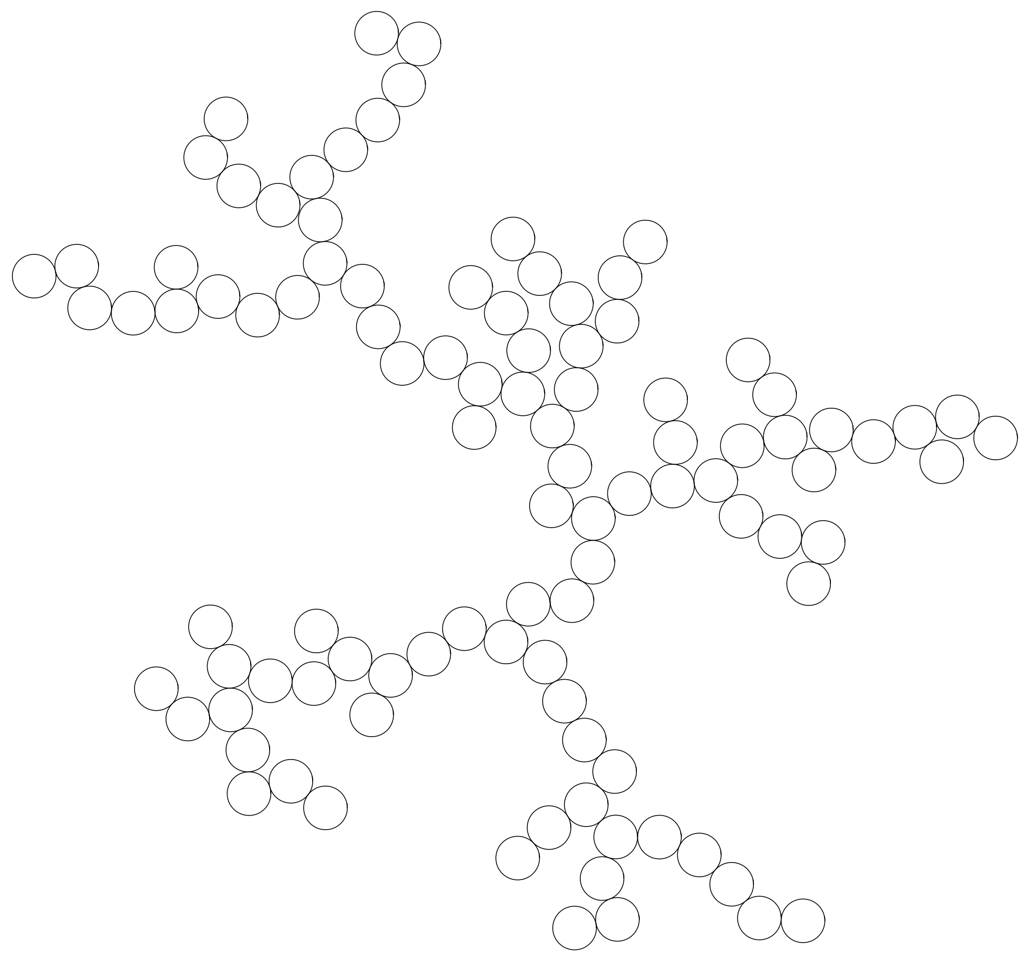

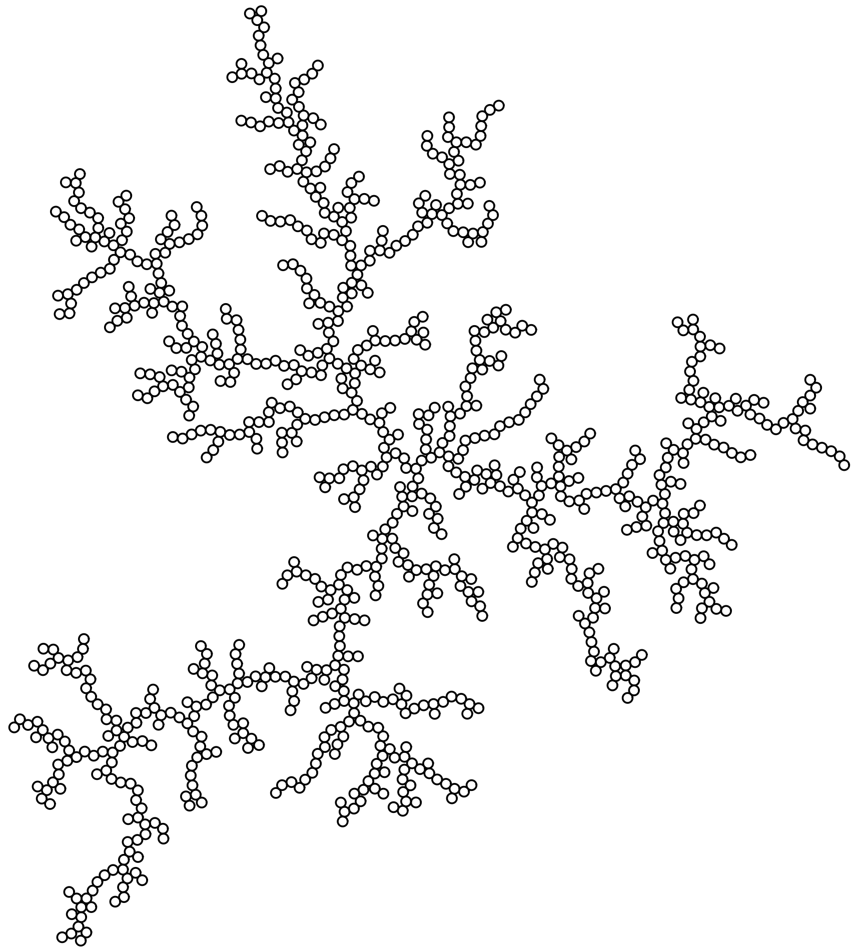

DLA, n = 10000

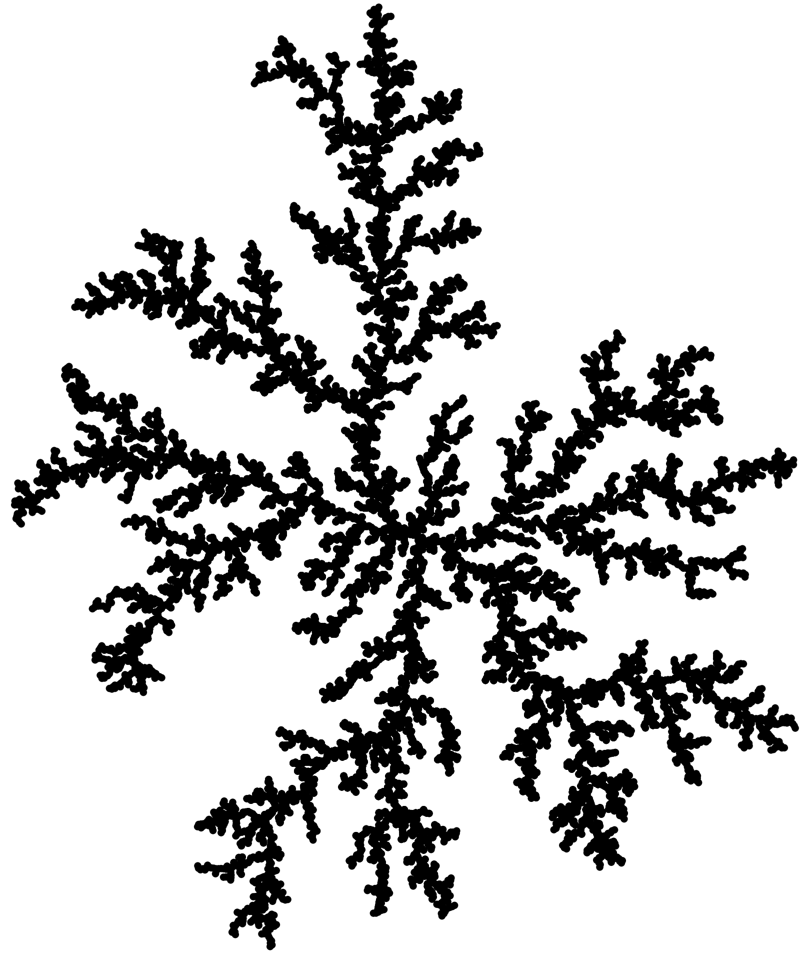

DLA, n = 10000

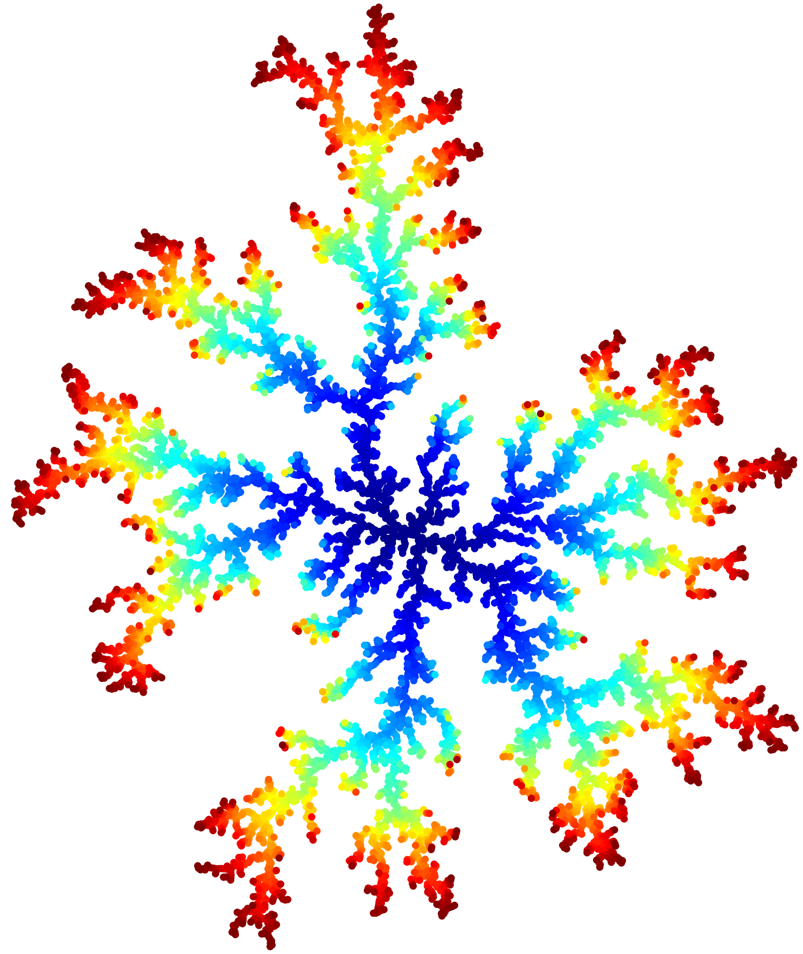

DLA, n = 100000

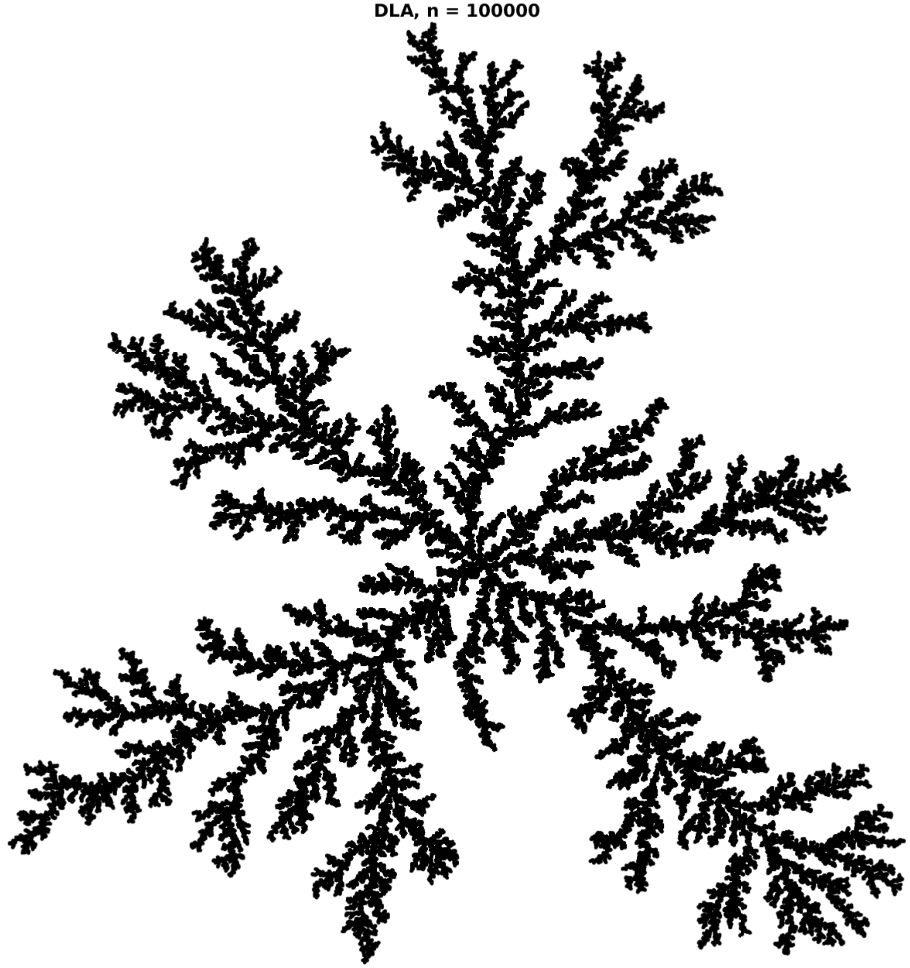

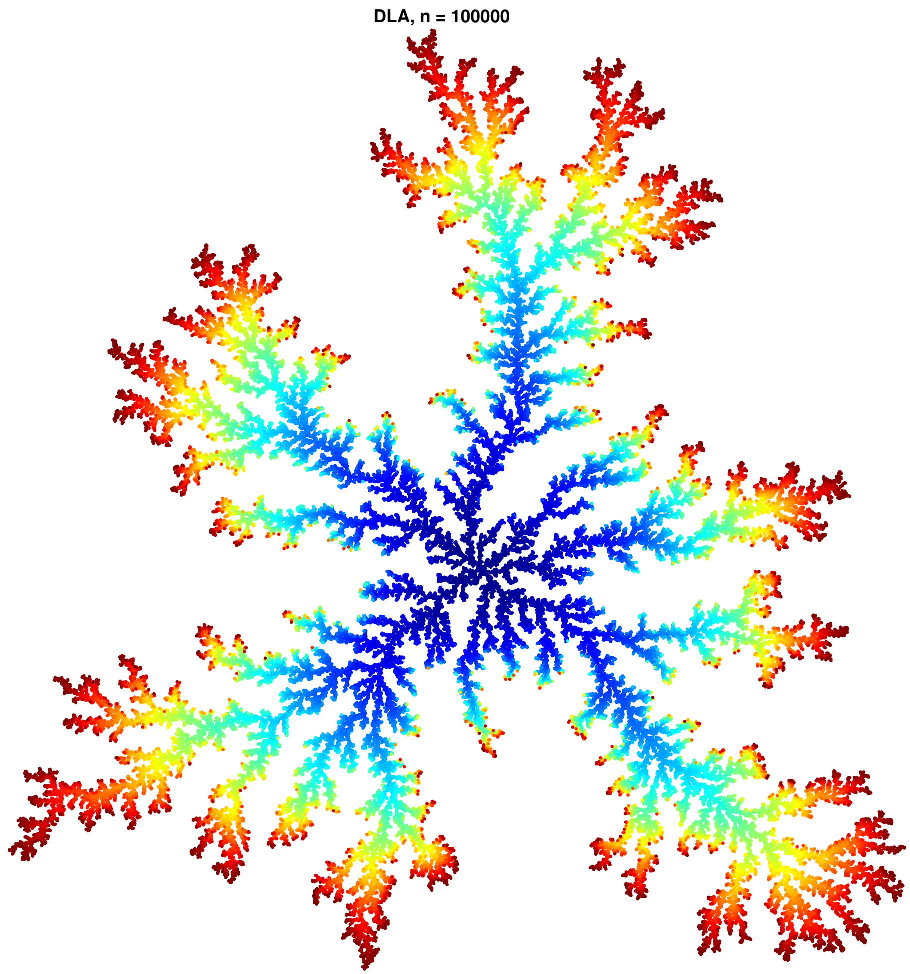

How fast does the diameter grow?

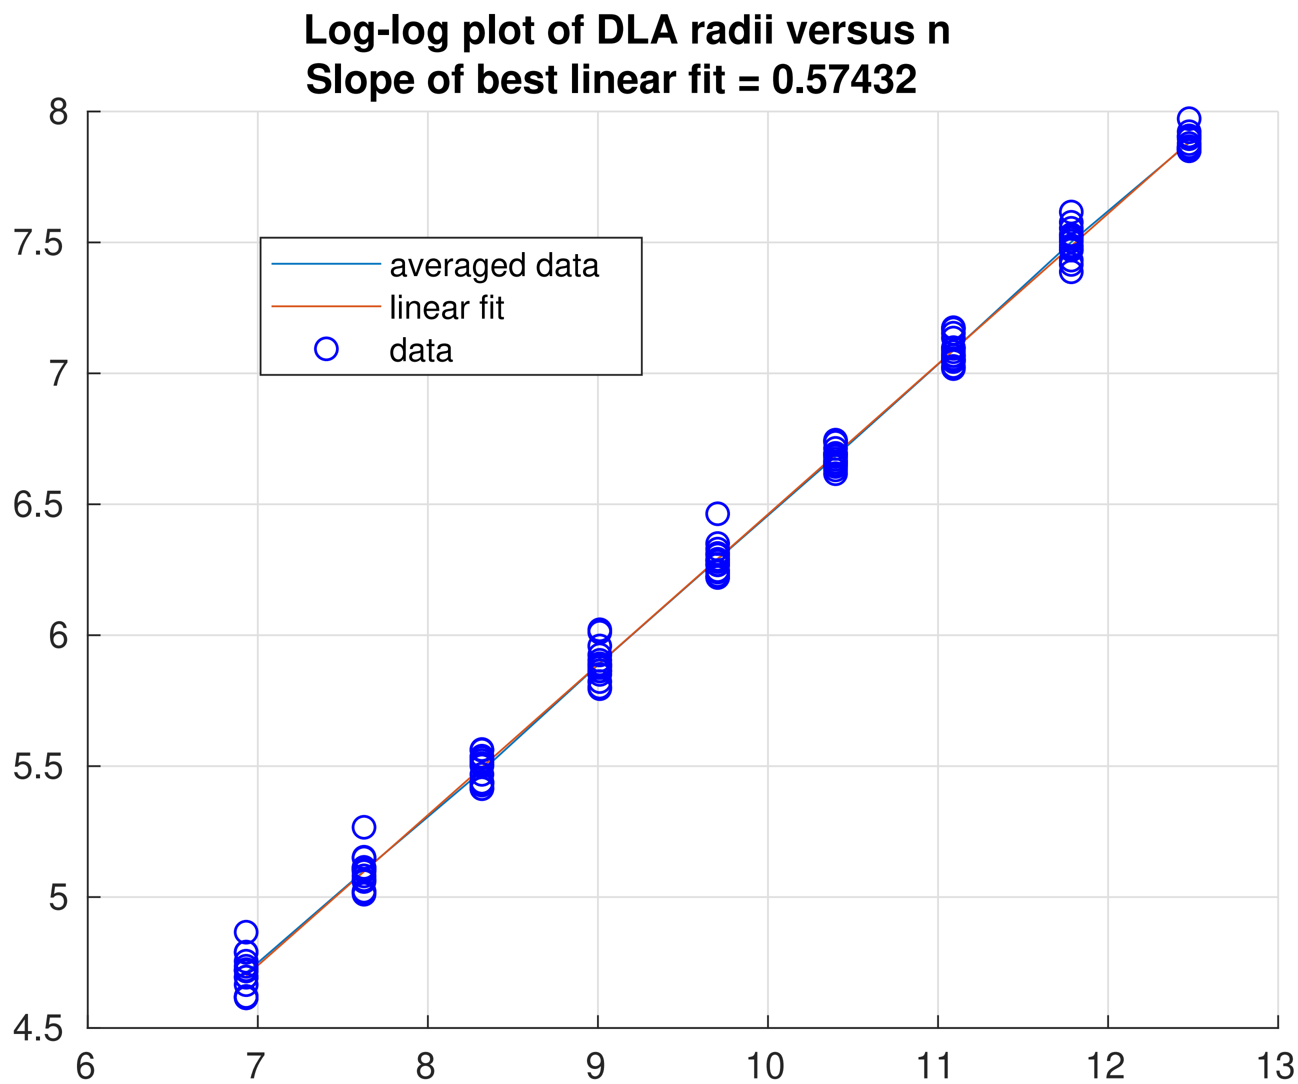

Numerical experiment for growth rate.

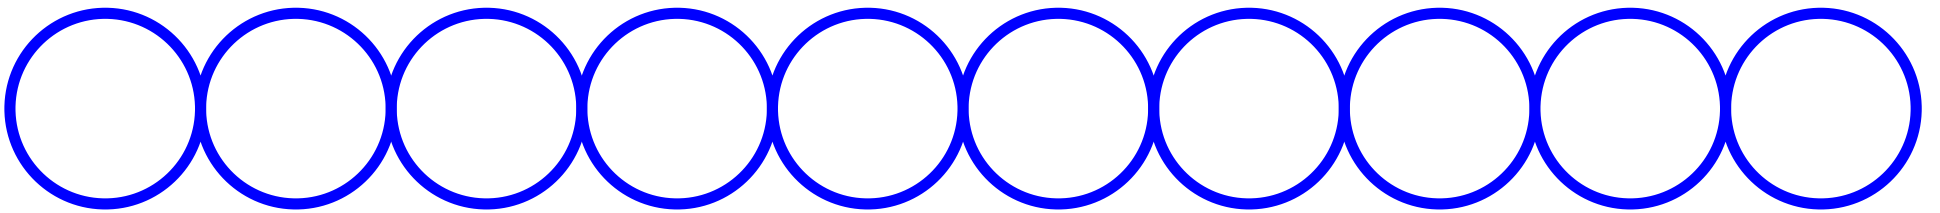

Trivial upper bound is $O(n)$.

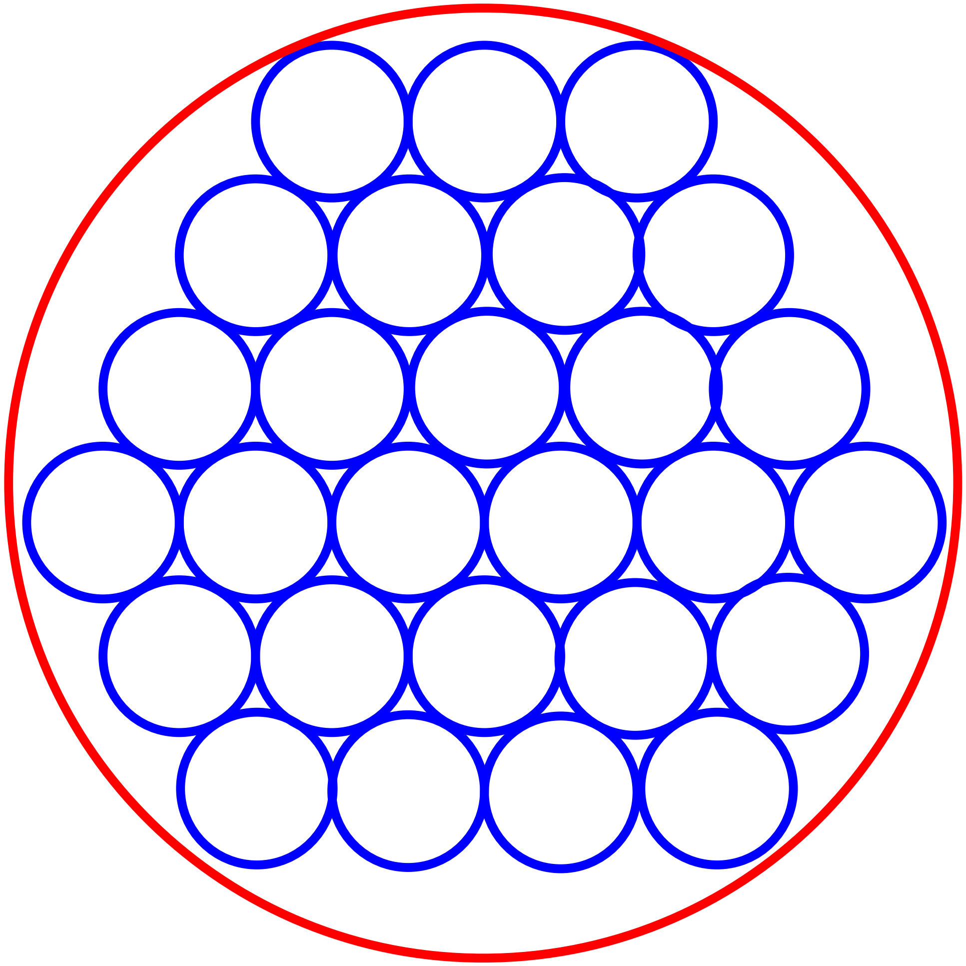

Trivial lower bound is $\Omega(\sqrt{n})$.

Theorem (Kesten): $\mathrm{diam}(\mathrm{DLA}(n)) = O(n^{2/3})$.

Theorem (Kesten): $\mathrm{diam}(\mathrm{DLA}(n)) = O(n^{2/3})$.

Equivalent: DLA takes $\gtrsim m^{3/2}$ steps to exit ball of radius m.

Sketch (following Lawler). Suppose $\beta m^{3/2}$ disks suffice to exit, $\beta > 0$.

Then cluster contains a chain of lattice points $\mathbf{z} = \{z_1, \ldots, z_k\}$ so that

$$|z_1| < m/2, \qquad |z_k| > m,$$

$$|z_j - z_{j+1}| \leq 4, j = 1, \ldots, k, \qquad m/4 \leq k \leq \beta m^{3/2}.$$

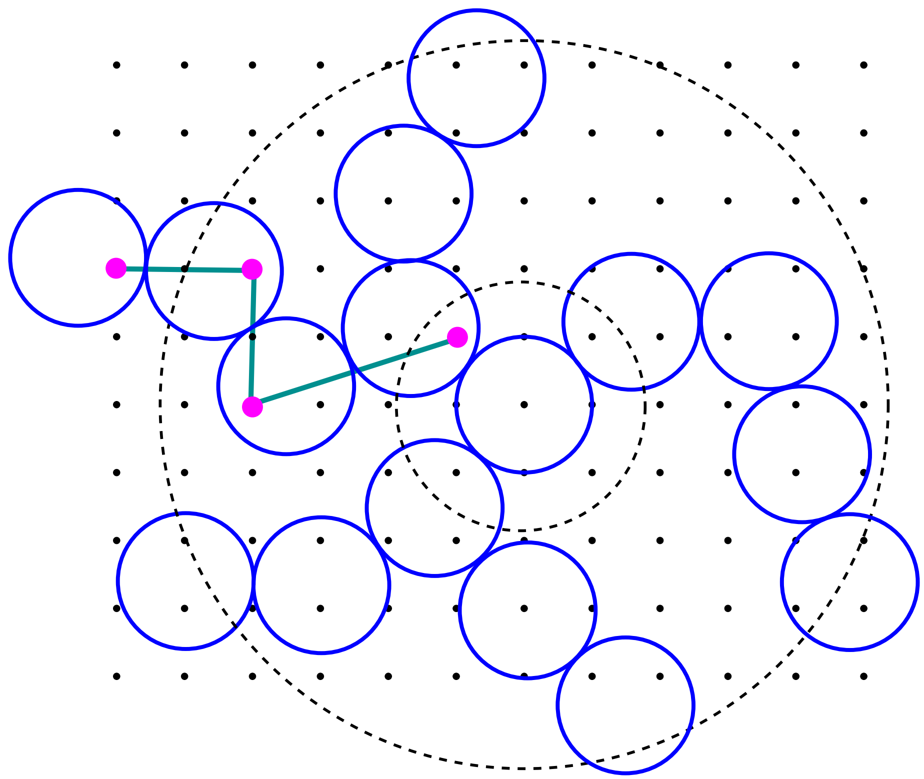

Let $W_m(\mathbf{z})$ be all clusters associated to a chain $\mathbf{z}$. Let $W_m = \cup_{\mathbf{z}} W_m(z)$.

At most $O(m^2 80^k)$ chains: $O(m^2)$ starting points and 80 choices per step.

Claim: $\text{Prob}(W_m(z)) \leq (C\beta)^k$.

Assuming claim, Kesten's theorem follows: if β small,

$$\sum_m \text{Prob}(W_m) \leq \sum_m \sum_z \text{Prob}(W_m(z)) \leq C \sum_m m^2 (80C\beta)^m < \infty.$$

By Borel-Cantelli a.s. W_m occurs only finitely often, i.e., eventually the exit time from radius m is always bigger than $\beta m^{3/2}$. QED

Proof of claim: Suppose D_j is disk covering z_j. How long do we wait between choosing D_j and D_{j+1}?

The D_{j+1} must land within distance 4 of D_j. Probability bounded using:

Beurling's thm: If $\Omega = \mathbb{C} \setminus K$, K compact and connected, $x \in K$,

$$\omega(\infty, D(x,1) \cap K, \Omega) \leq \frac{C}{\sqrt{\text{diam}(K)}}.$$

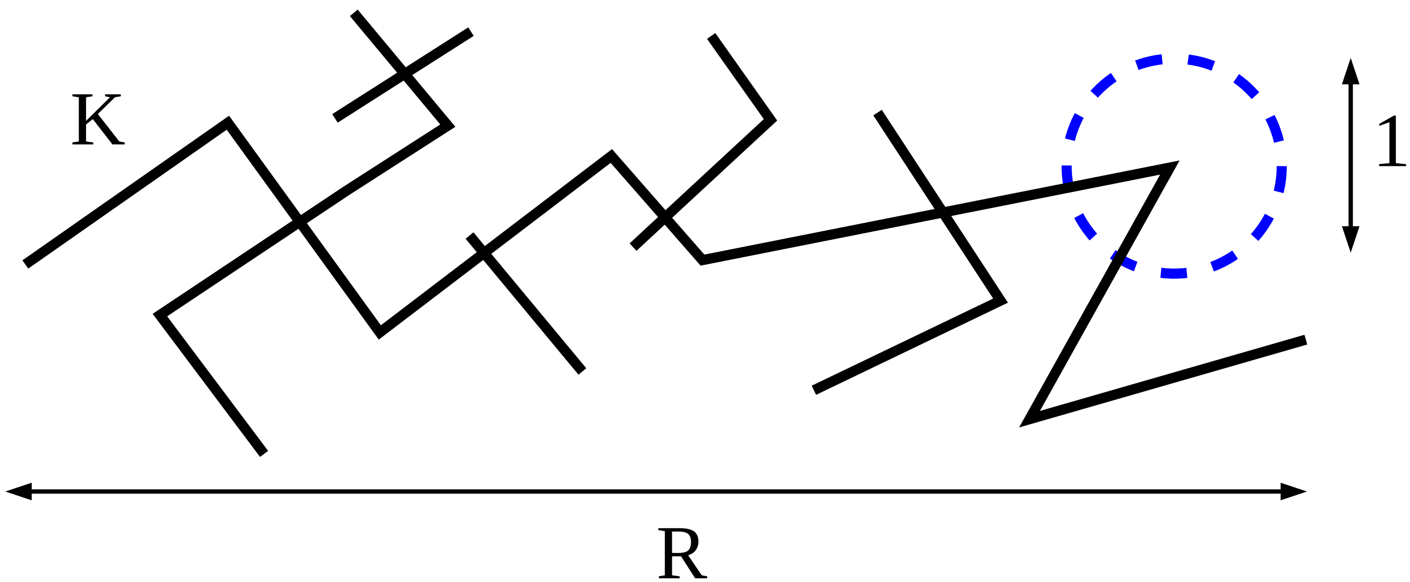

If cluster has diameter $\sim m$, probability of adding D_{j+1} at next step is

$$p \lesssim m^{-1/2}.$$

The probability of waiting t steps to add it is

$$\text{Prob}(\text{waiting} > t) \geq (1-p)^t.$$

Claim (hence Kesten's theorem) then follows from:

Lemma: If $X_1, \ldots, X_n$ are independent geometric random variables with parameter p, then

$$\text{Prob}(\sum_{k=1}^{n} X_k < \frac{an}{p}) \leq (2e^2 a)^n.$$

Beurling's theorem is sharp when K is line segment. Since clearly DLA never looks like a line segment, we should get a smaller estimate for harmonic measure, hence a longer waiting time, hence a better upper bound for the diameter of DLA.

How to make this precise?

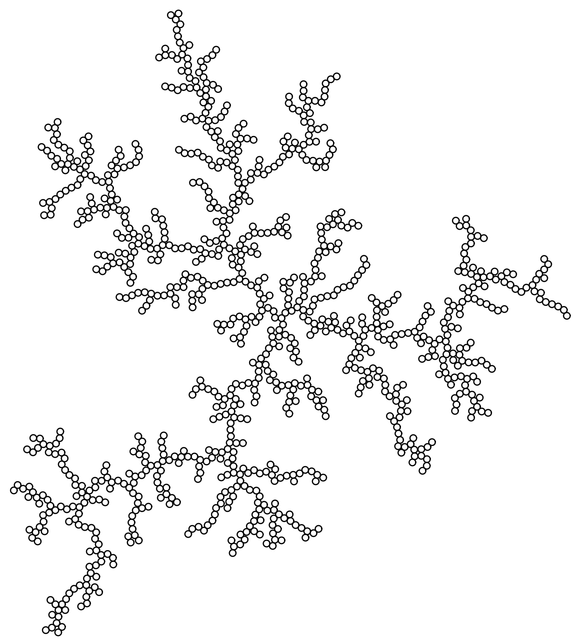

Amazingly, there is no known better lower bound than the trivial $\sqrt{n}$.

Conjecture: Almost surely,

$$\lim_{n\to\infty} \frac{\operatorname{diam}(\mathrm{DLA}(n))}{\sqrt{n}} = \infty.$$

If $\mathrm{DLA}(n)$ is roughly a disk of radius $\sqrt{n}$ then any boundary disk is hit with probability $\simeq 1/\sqrt{n}$, which gives which gives the trivial lower bound. For non-trivial lower bound, we need to show there are points that get hit with probability $\gg n^{-1/2}$.

Consider convex hull of the DLA cluster. What is the harmonic measure of the disks that touch the convex hull boundary?

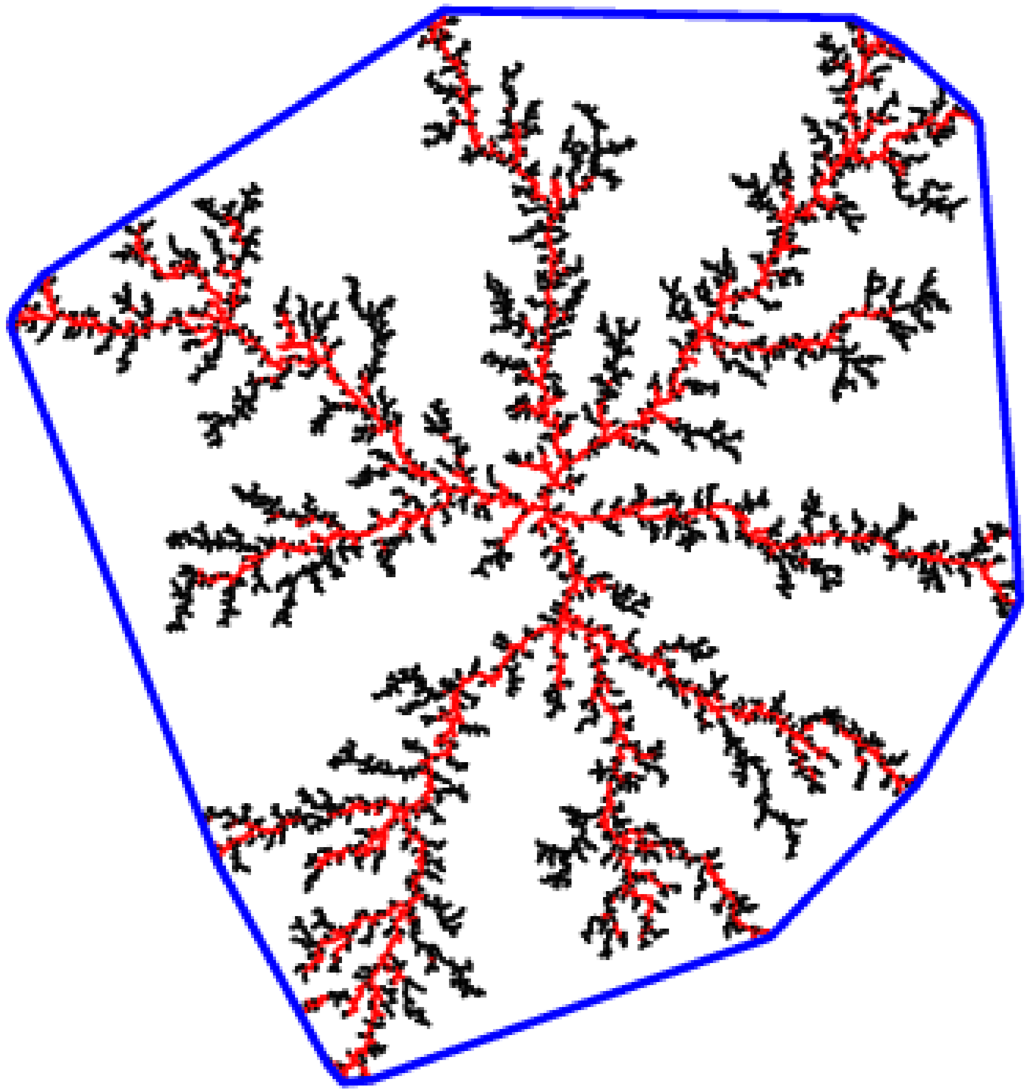

If convex hull has "sharp angles" some vertices have larger than average harmonic measure, implies faster than trivial growth.

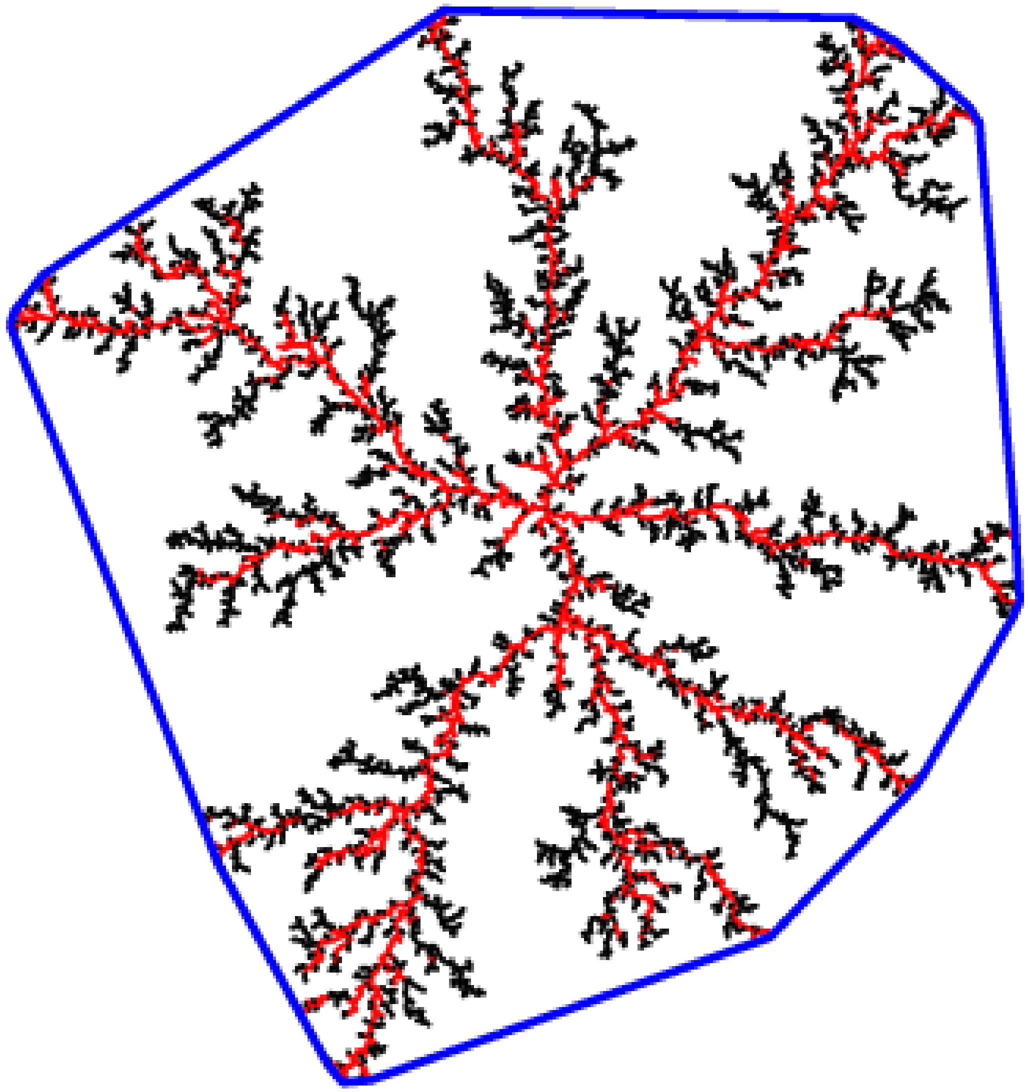

One way to have “sharp angles” is to have few vertices: if the convex hull boundary has few vertices, some of the angles should be large.

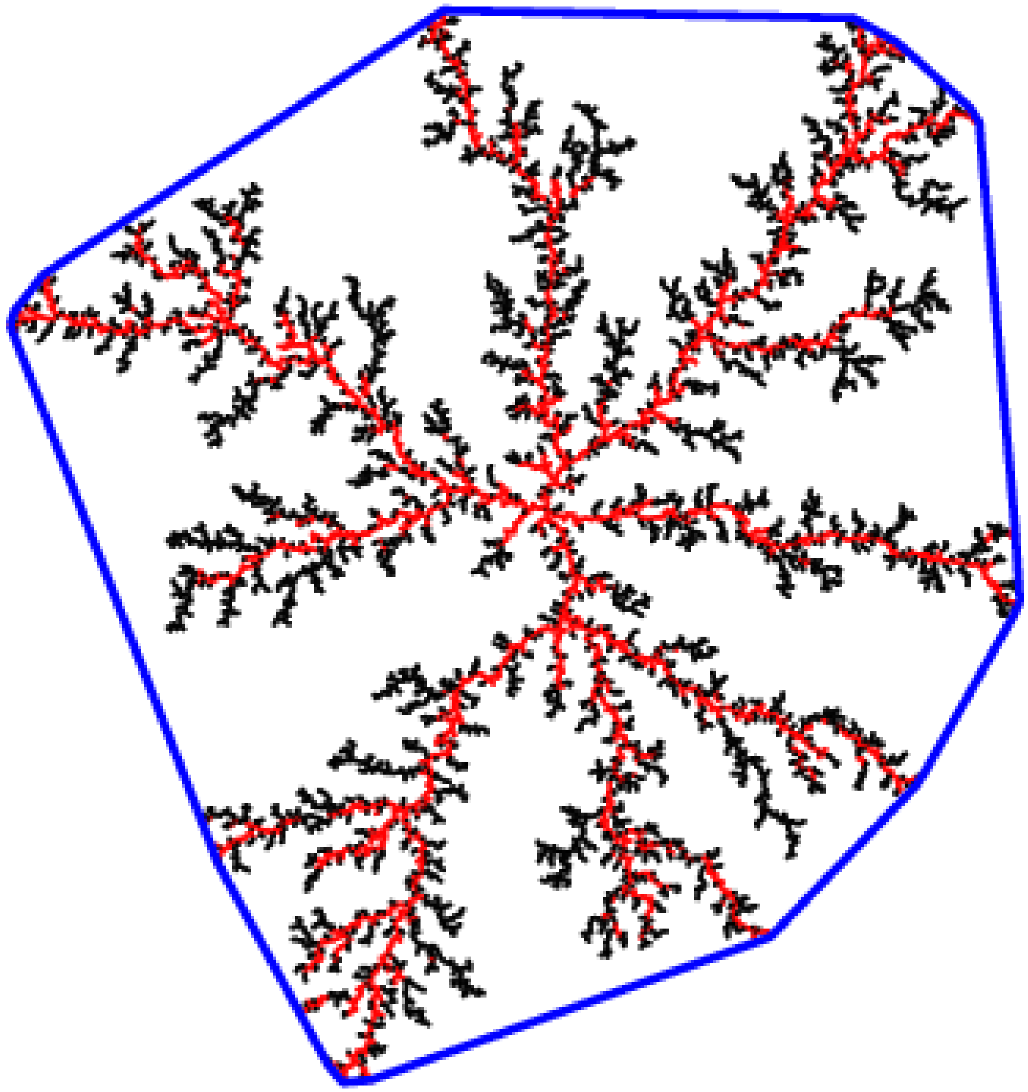

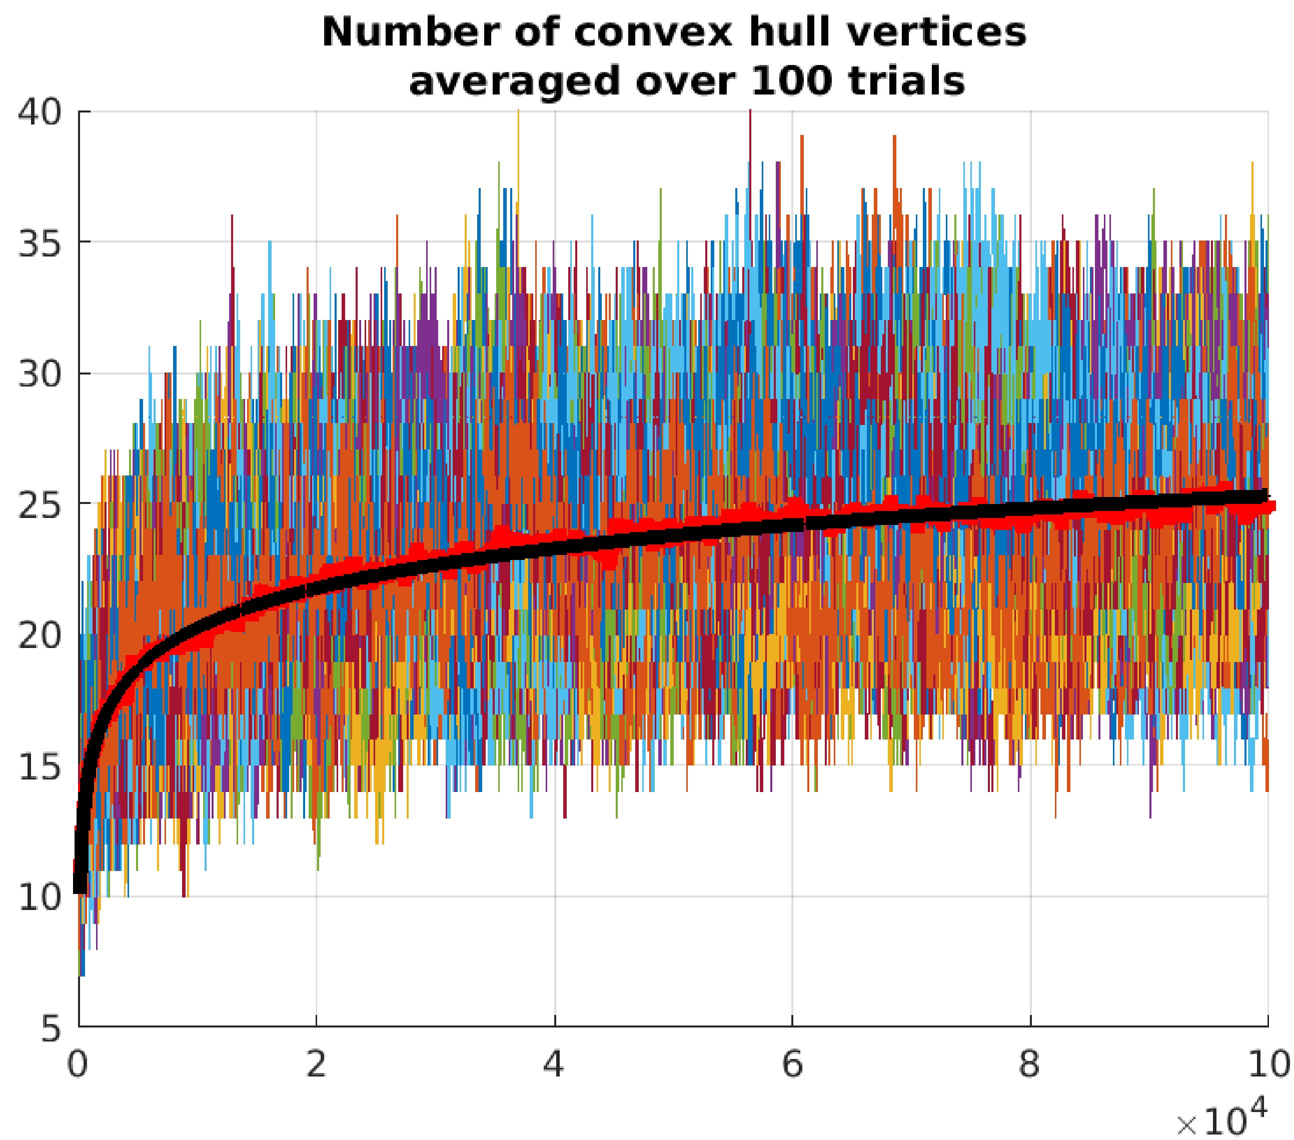

How many convex hull vertices are there at time n?

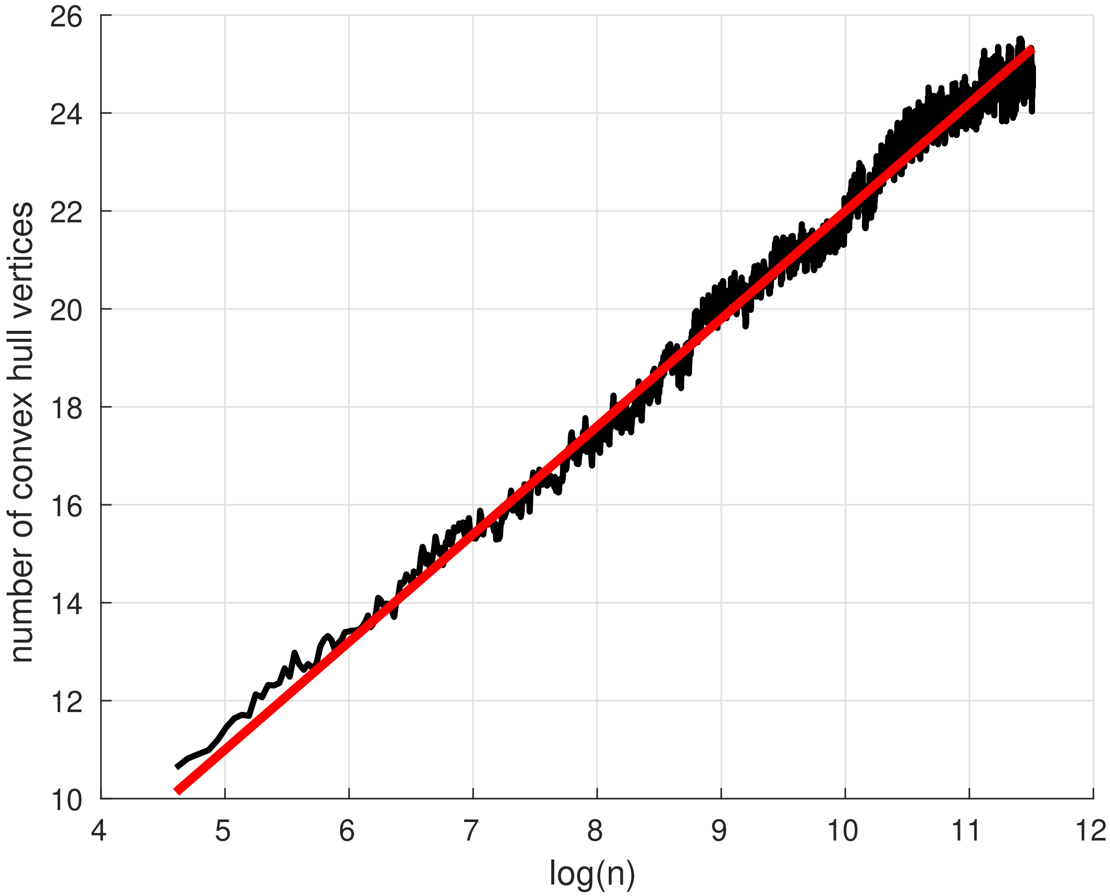

Number of convex hull vertices, averaged over 100 trials.

Plotted versus $\log(n)$, looks linear.

Numerically, $\approx (2.2) \log n$.

How to measure deviation of convex hull from disk? Gauss map? Perhaps curvature of boundary defines a measure that is not close to uniform. Then use Ahlfors distortion type estimate to show convex hull vertices are likely to get hit.

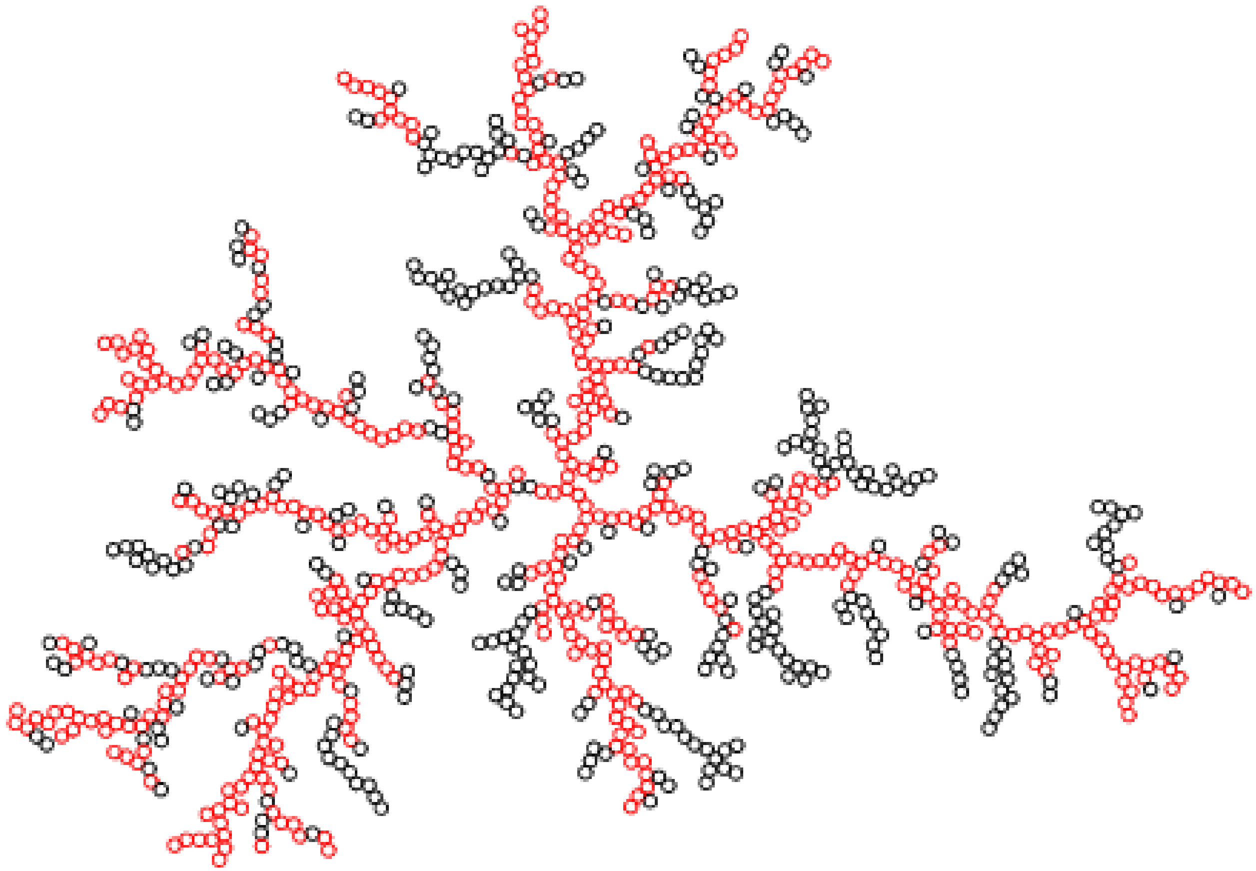

Red disks where on convex hull boundary when added.
Percentage probably tends to zero, but how fast?

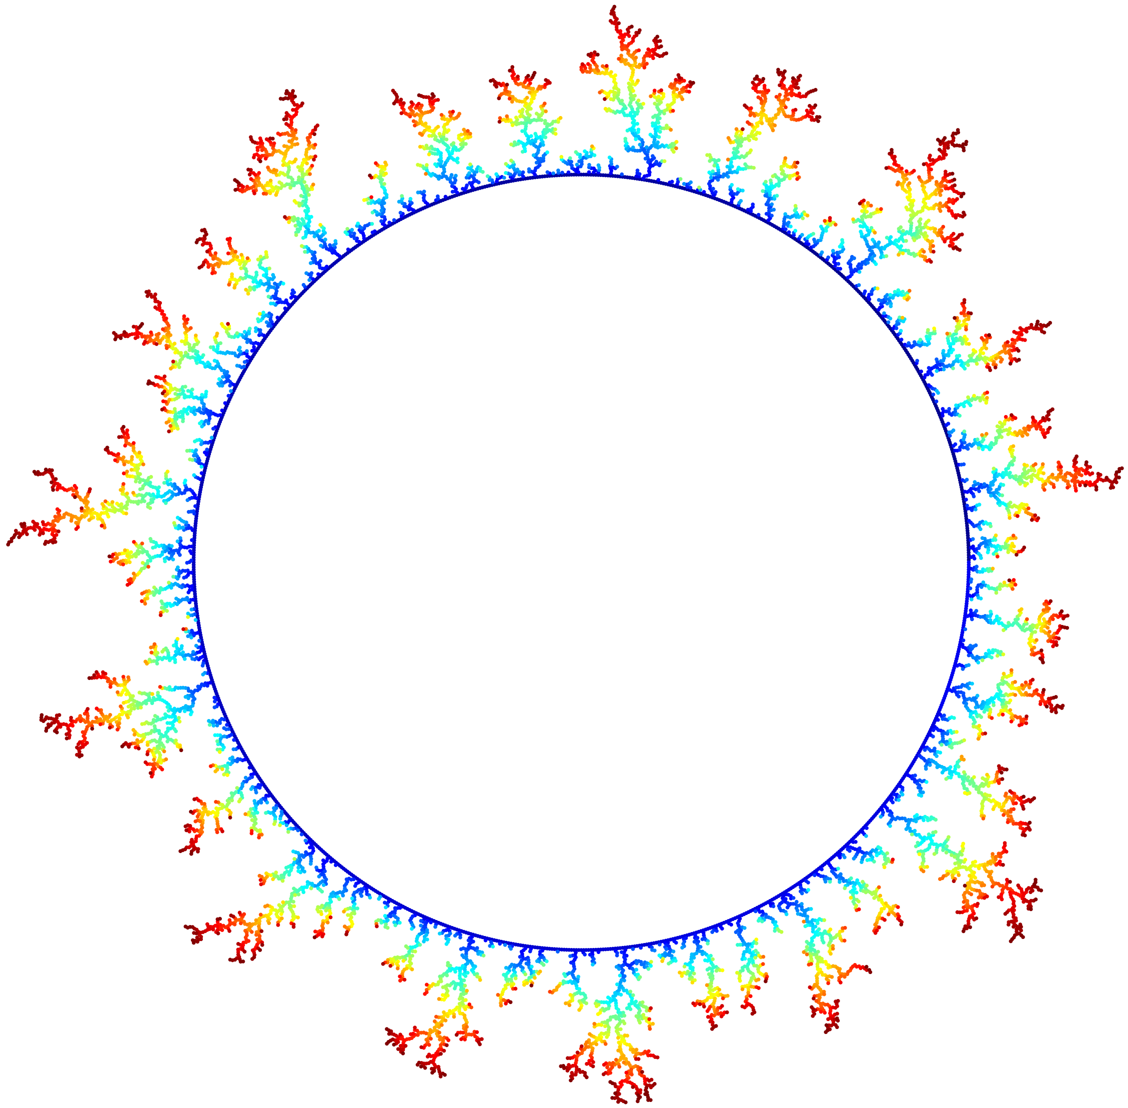

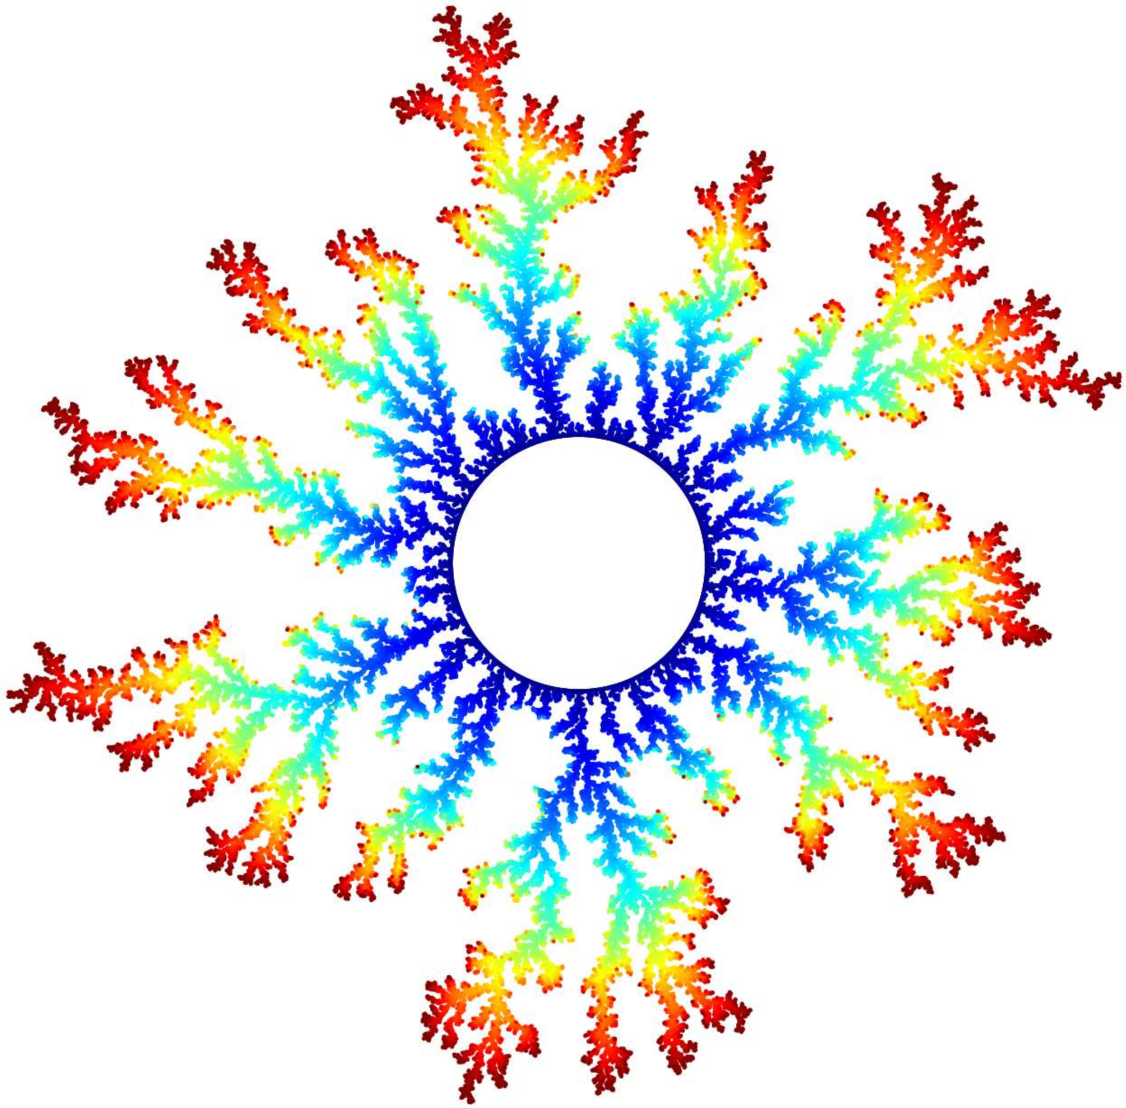

PART II: SHORT PATHS IN THE BROWNIAN TRACE

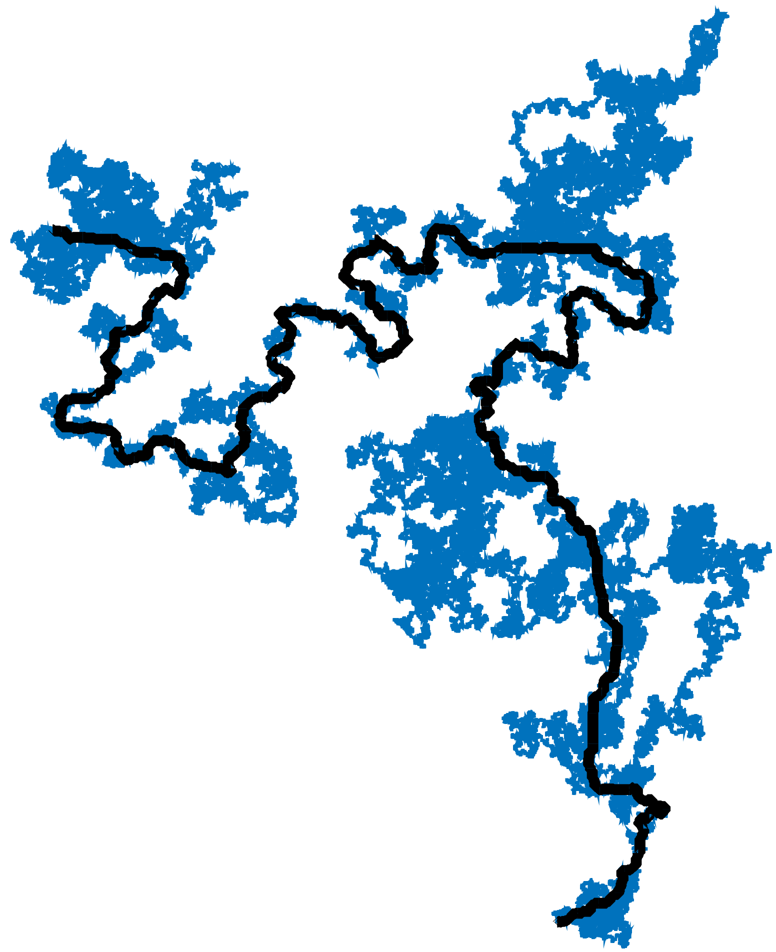

Robin Pemantle proved Brownian motion does not cover a line segment.

Question: Does it cover a rectifiable curve?

Question: Does it cover a curve of dimension $1 + \epsilon$, any $\epsilon > 0$?

Percolation dimension = minimal dimension of Jordan arc inside the set.

Frontiers have dimension $4/3$ (Lawler et. al.).

Brownian trace contains curves of dimension $5/4$ (Dapeng Zhan).

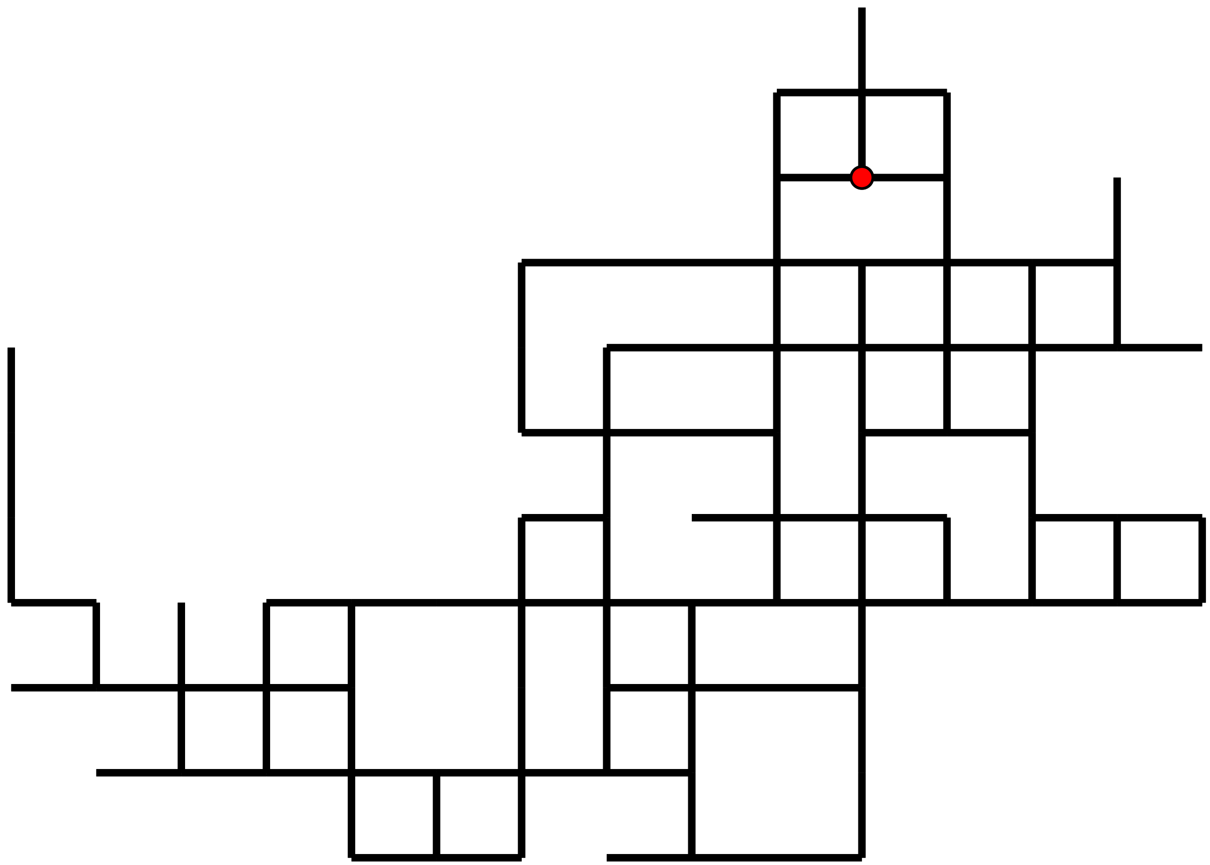

200 step random walk.

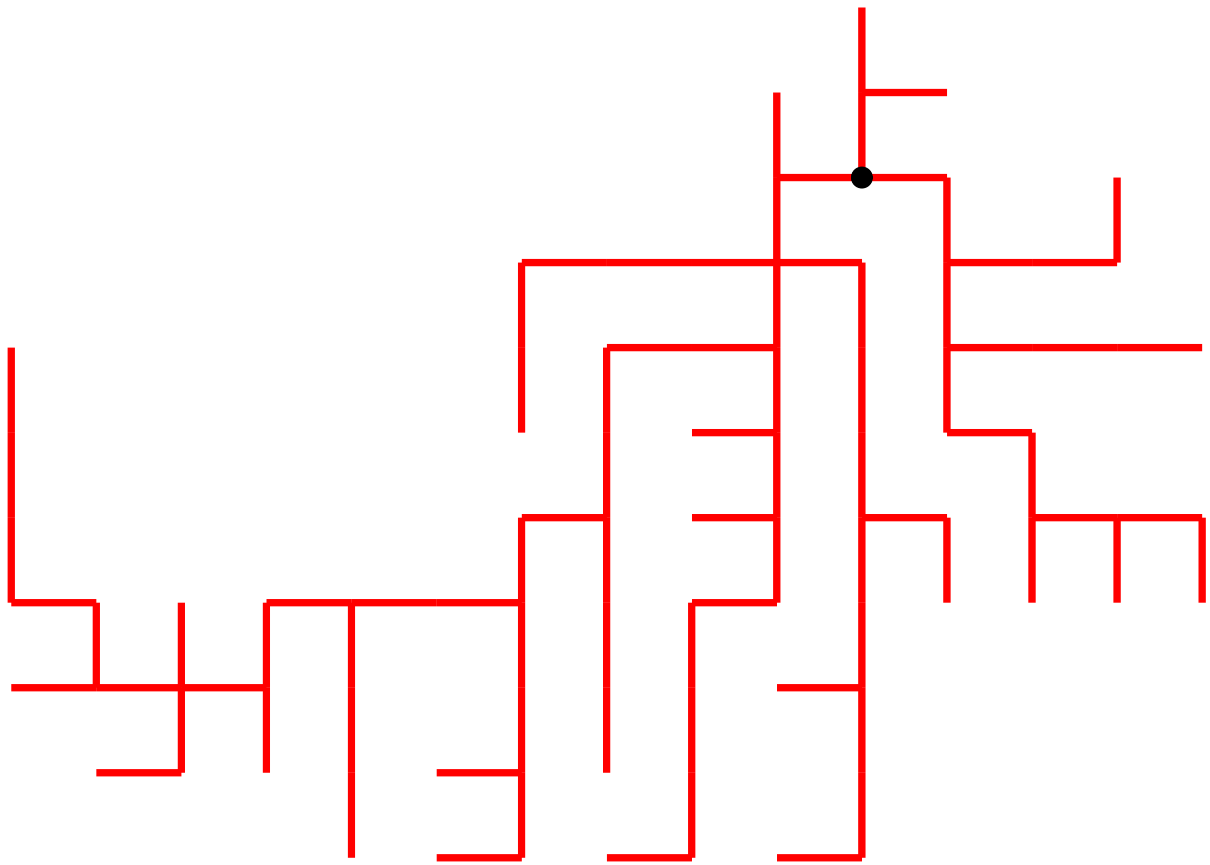

Minimal distance rooted spanning tree.

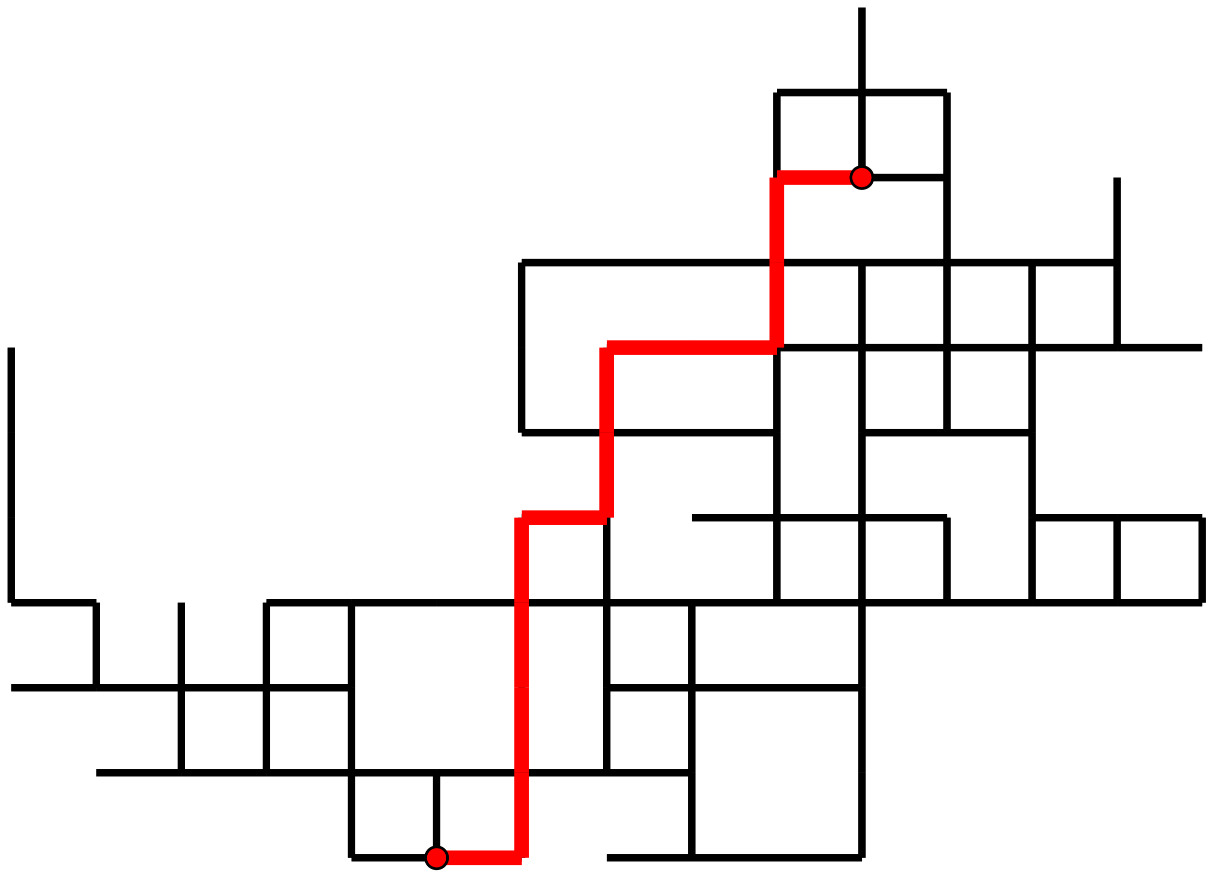

A shortest path from 0 to $\sqrt{n}/2$.

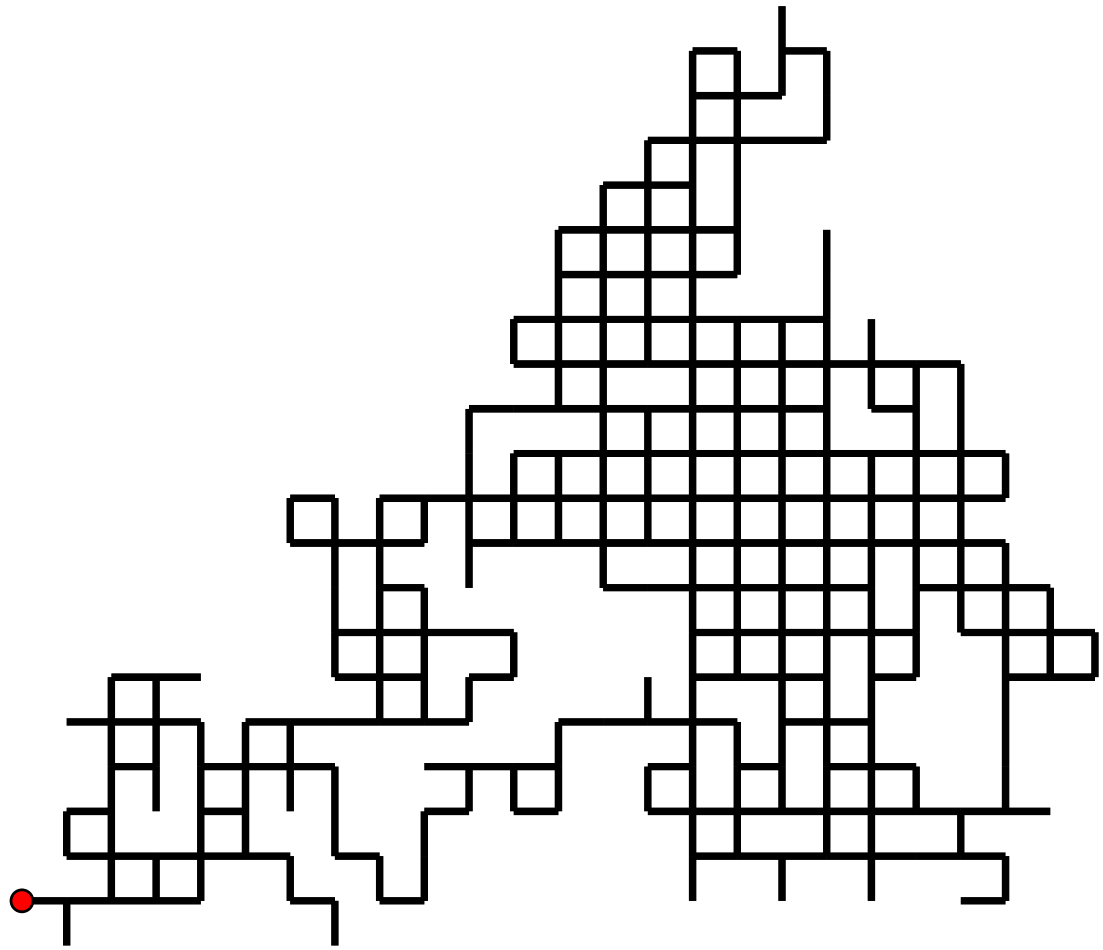

1000 step random walk.

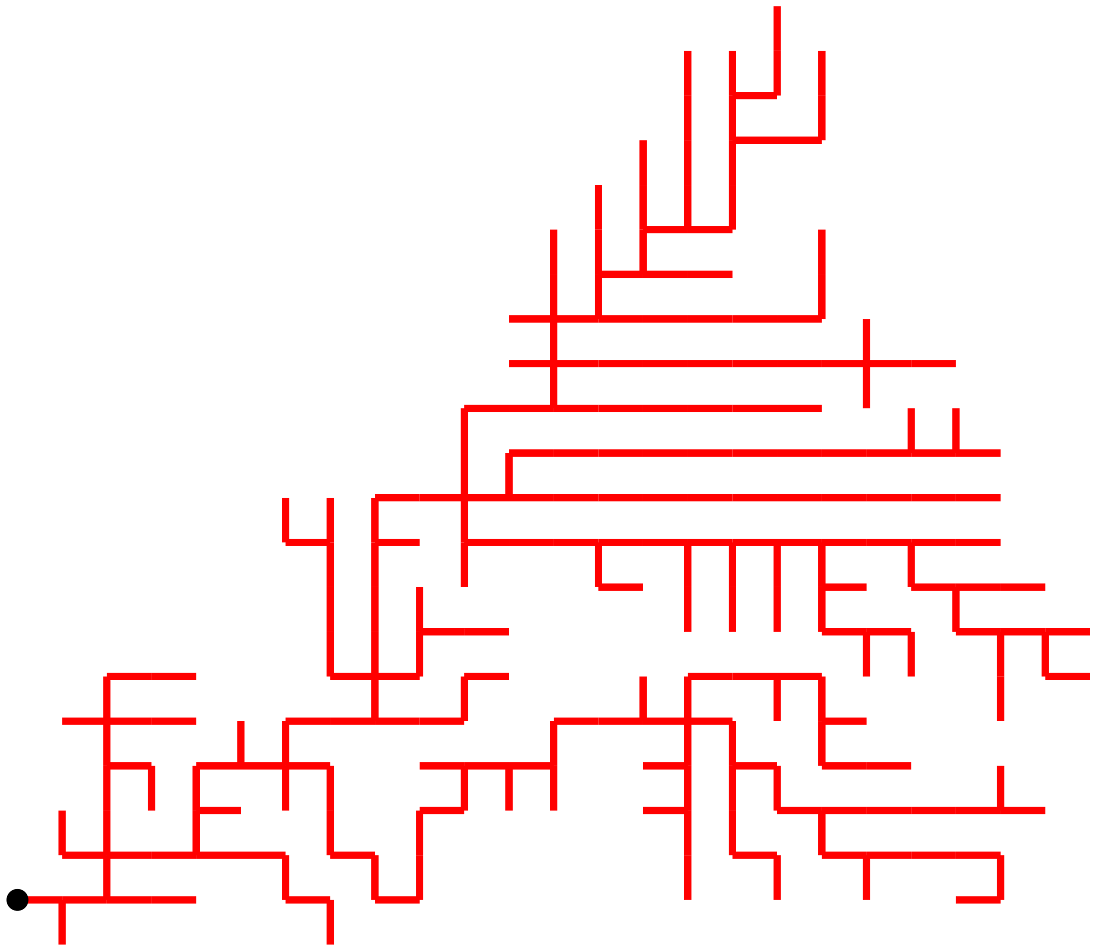

Minimal distance rooted spanning tree.

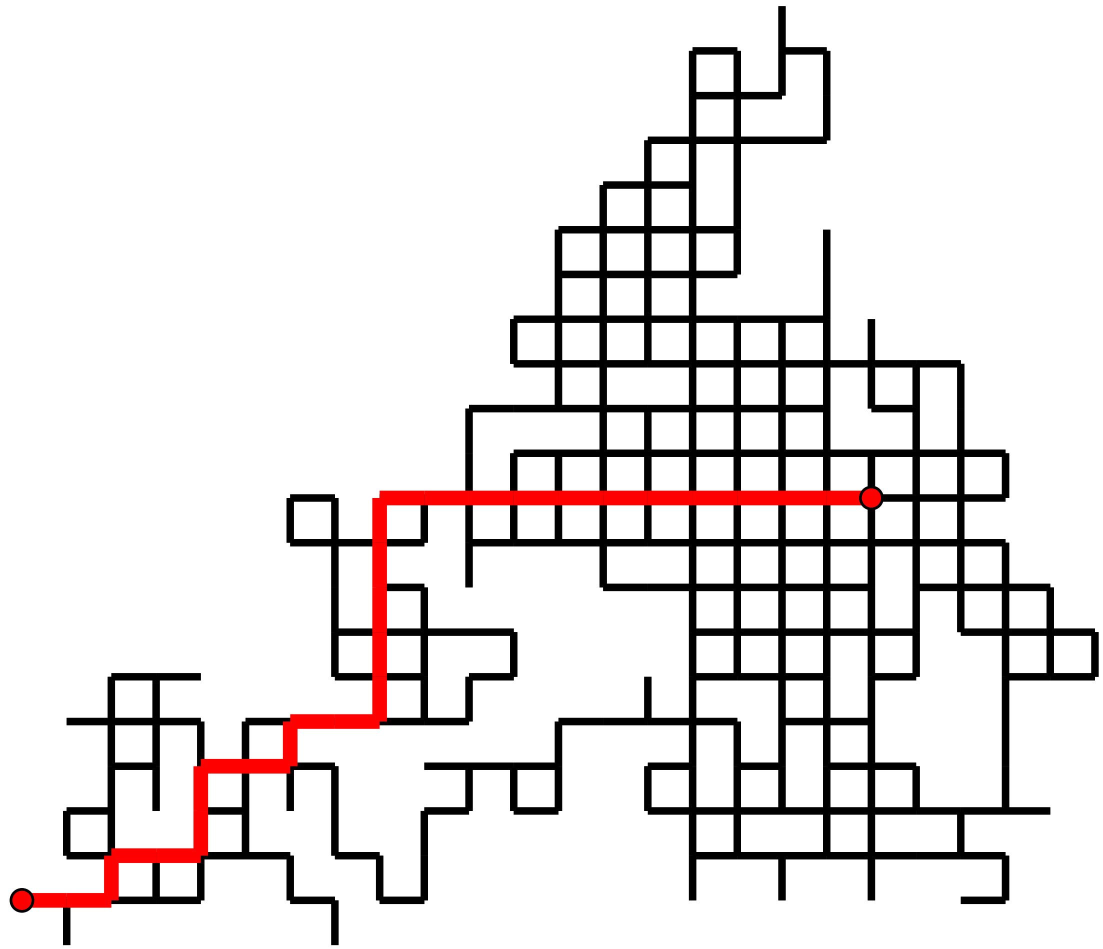

A shortest path from 0 to $\sqrt{n}/2$.

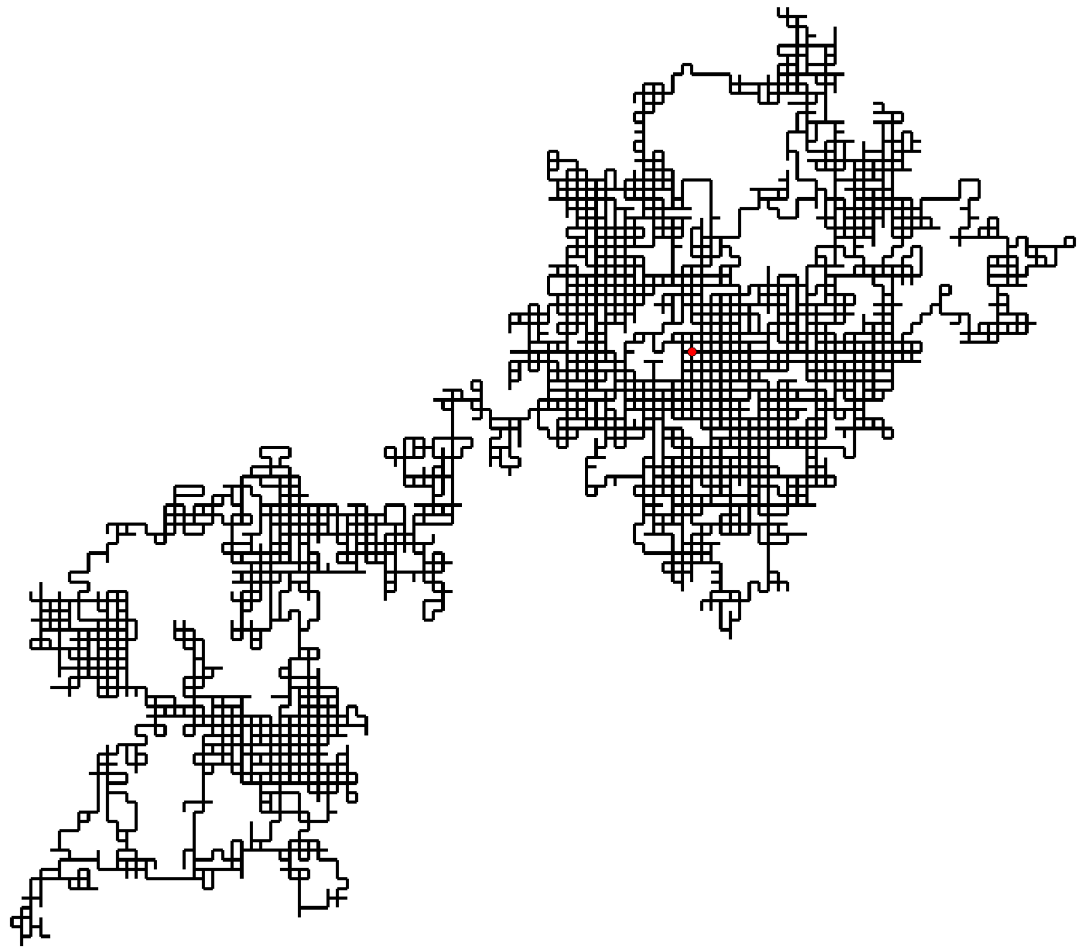

10000 step random walk.

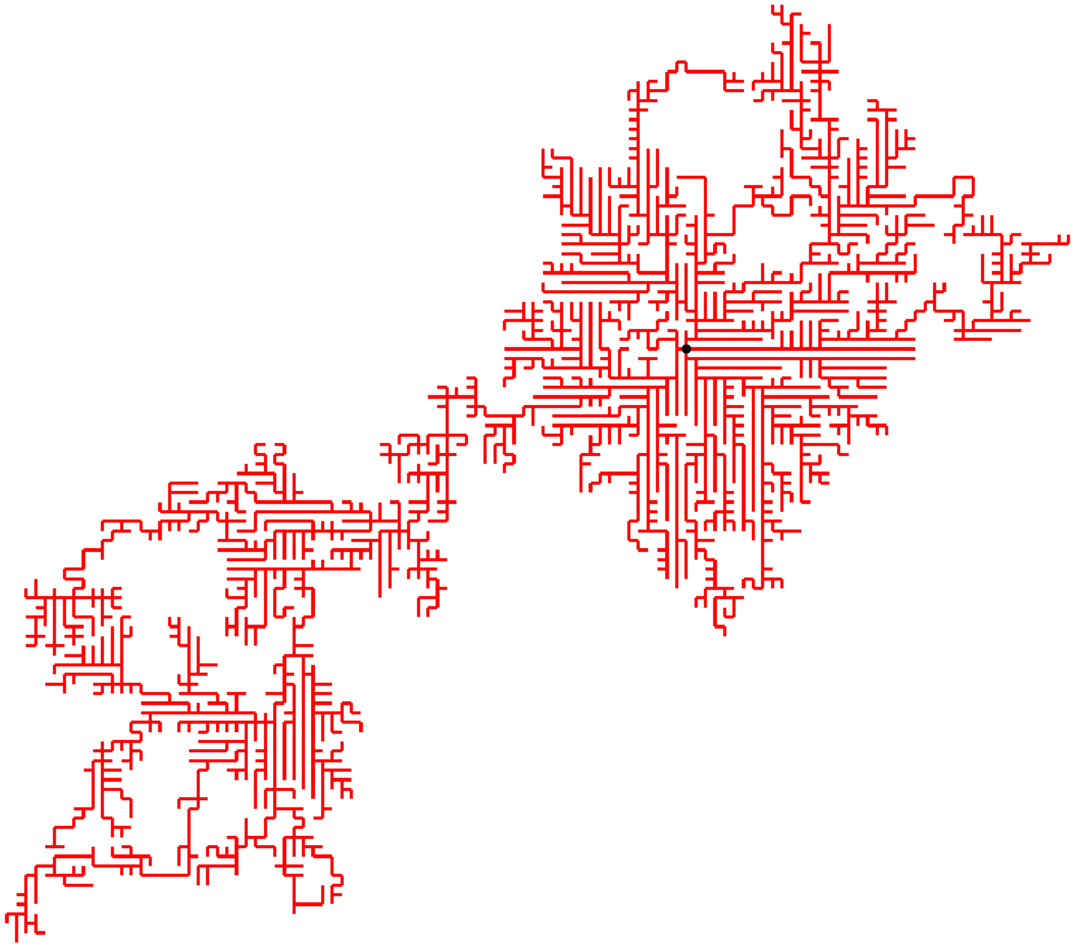

Minimal distance spanning tree, wrt to origin.

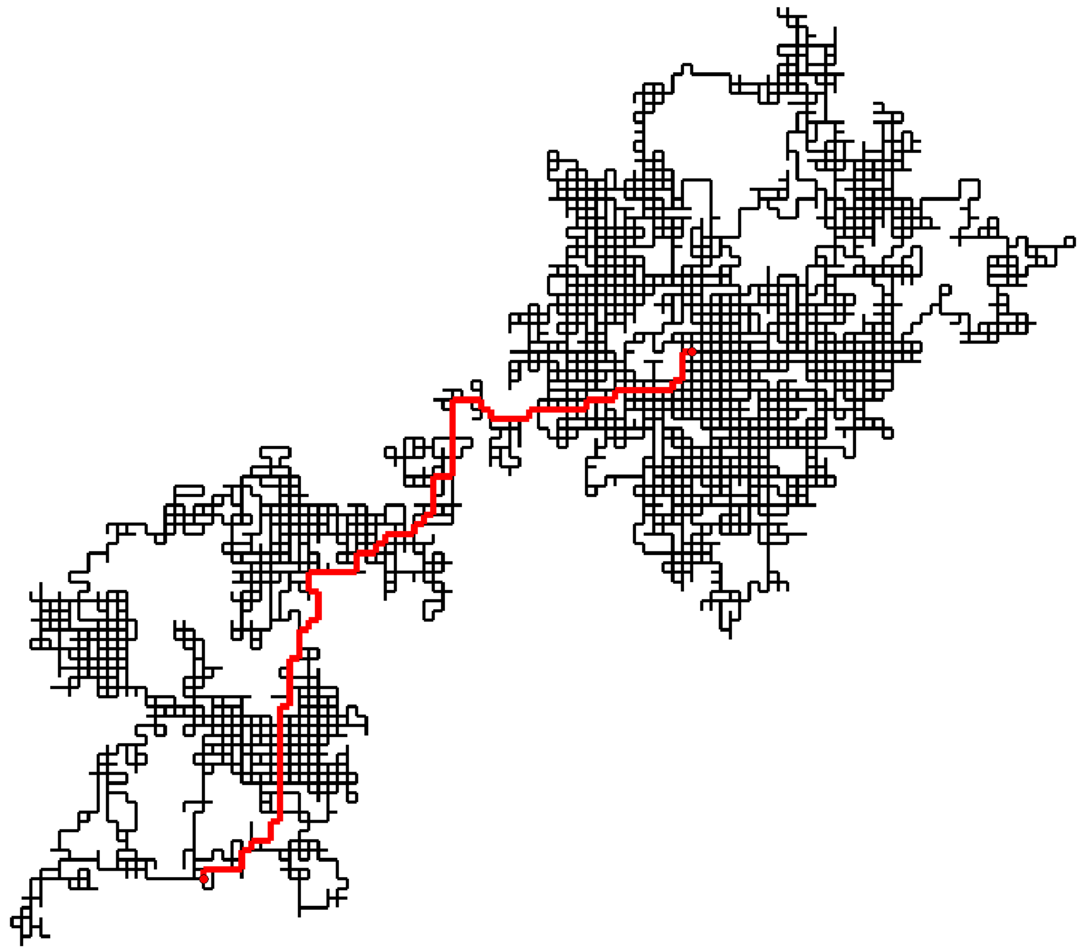

Minimal length path to distance $\sqrt{n}$.

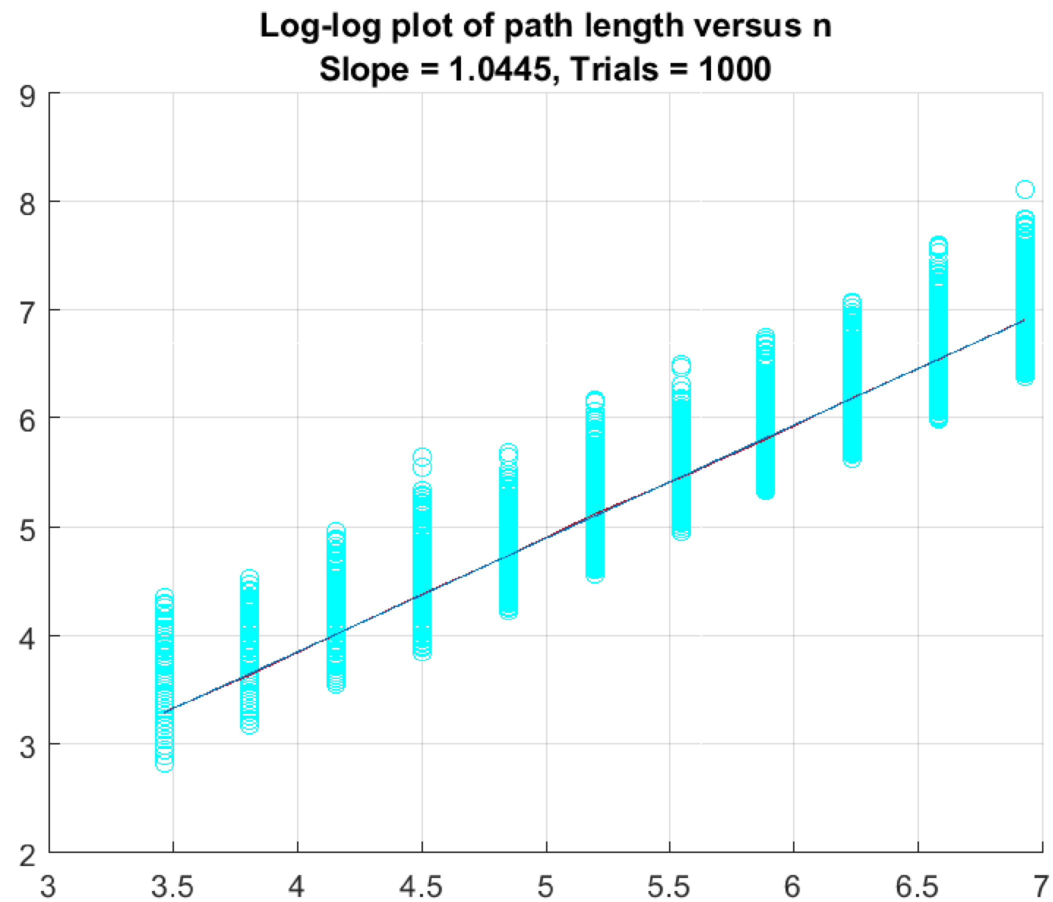

Log-log plot of graph distance 0 to $\{|z| = \frac{1}{2}\sqrt{n}\}$.

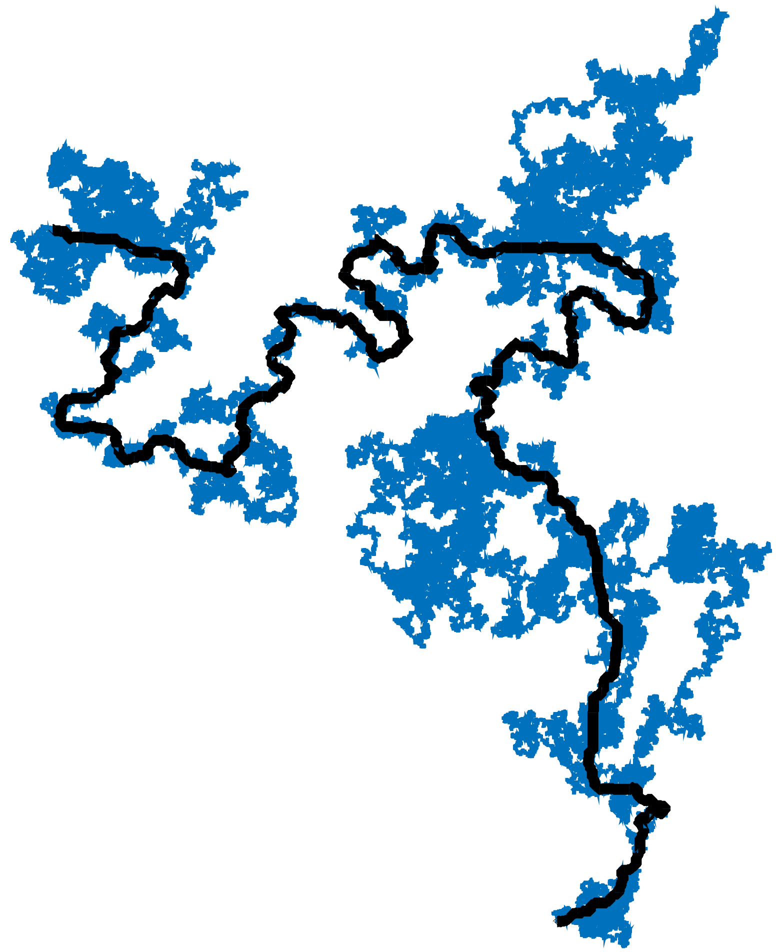

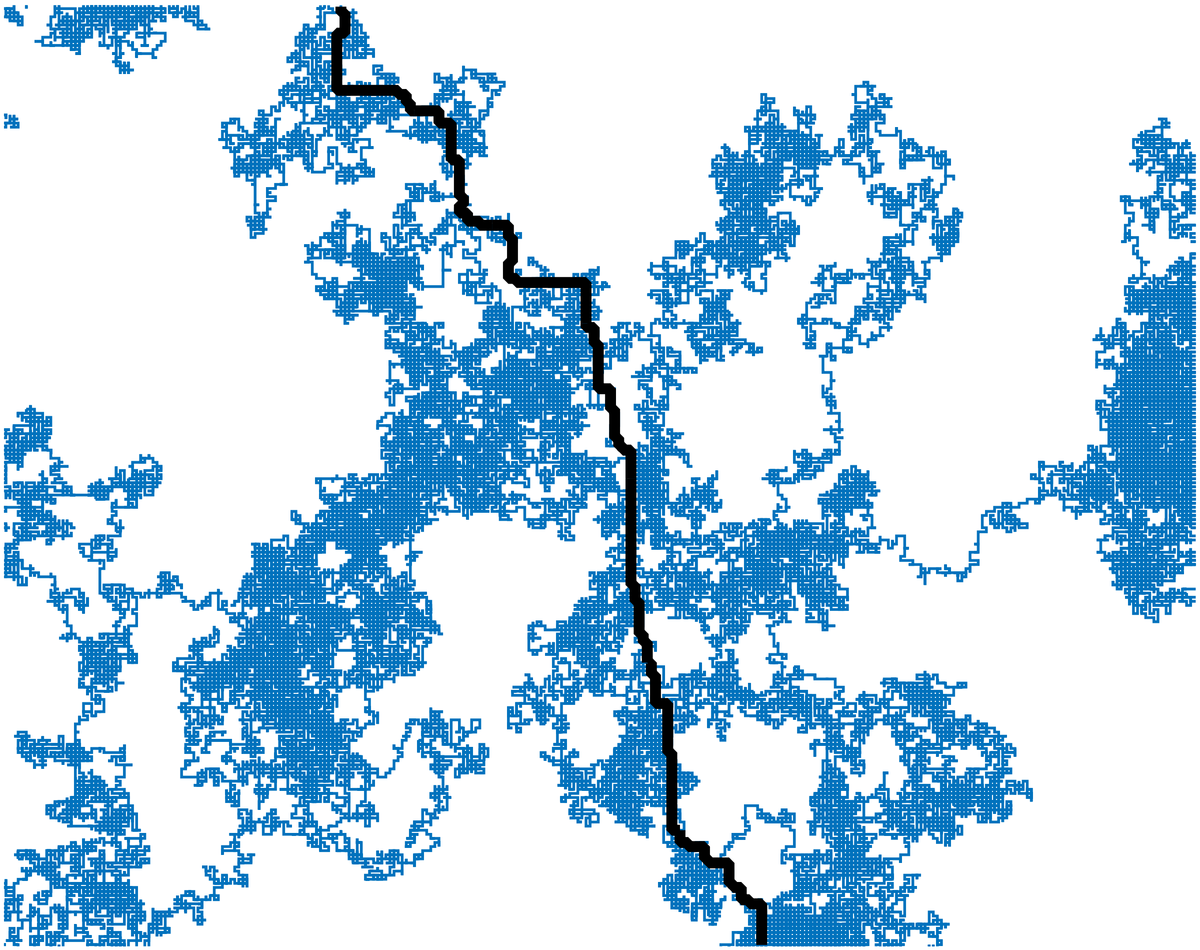

Paths can be straight where trace is dense.

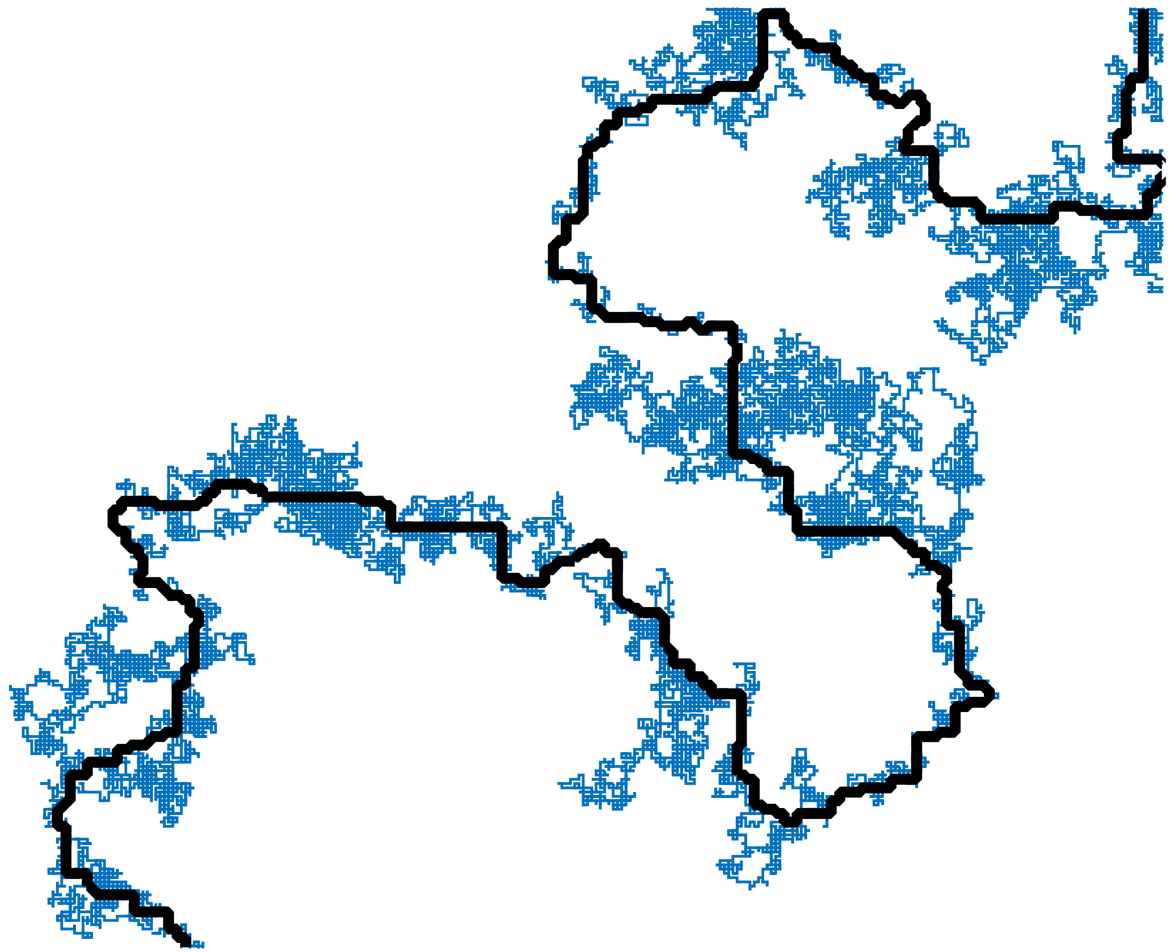

Paths wiggle more when trapped between components.

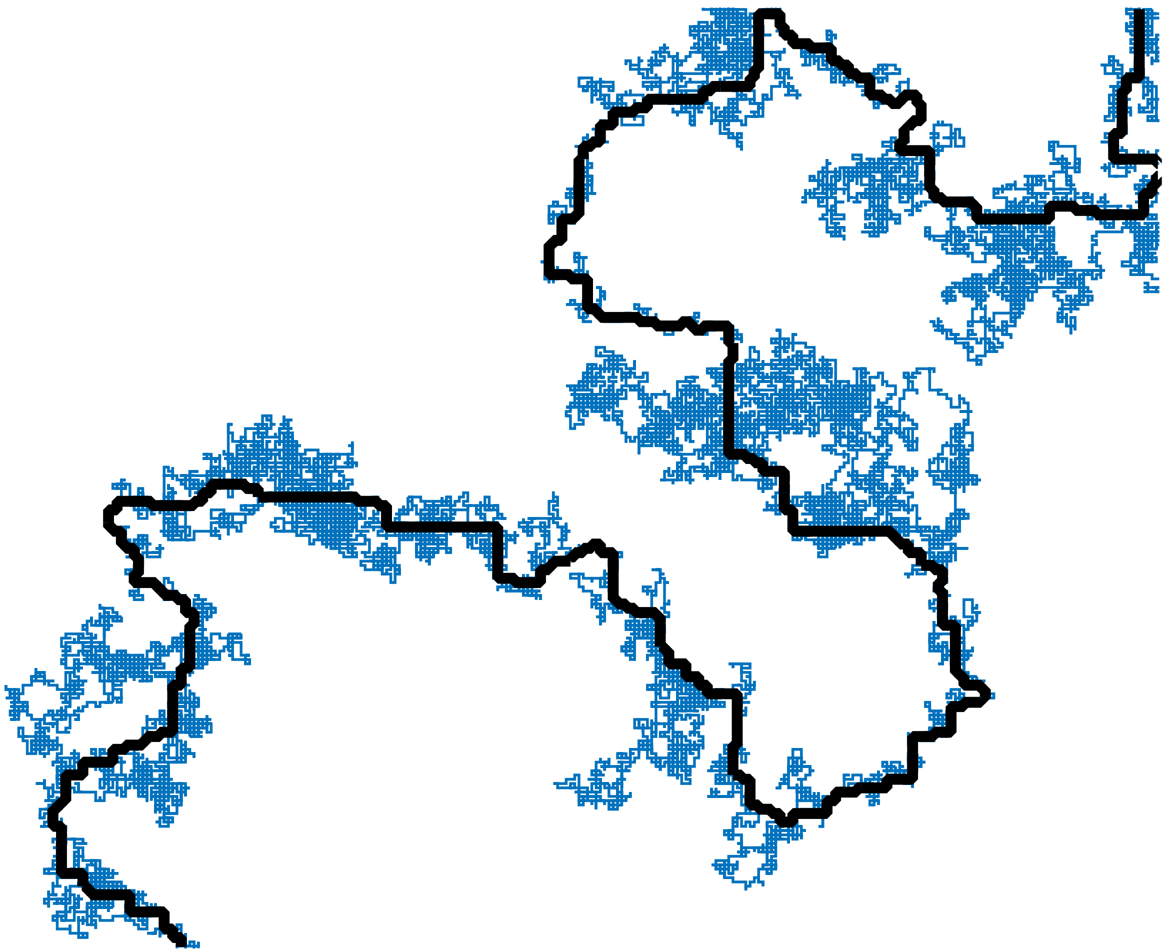

Can any two components be separated by a rectifiable curve?

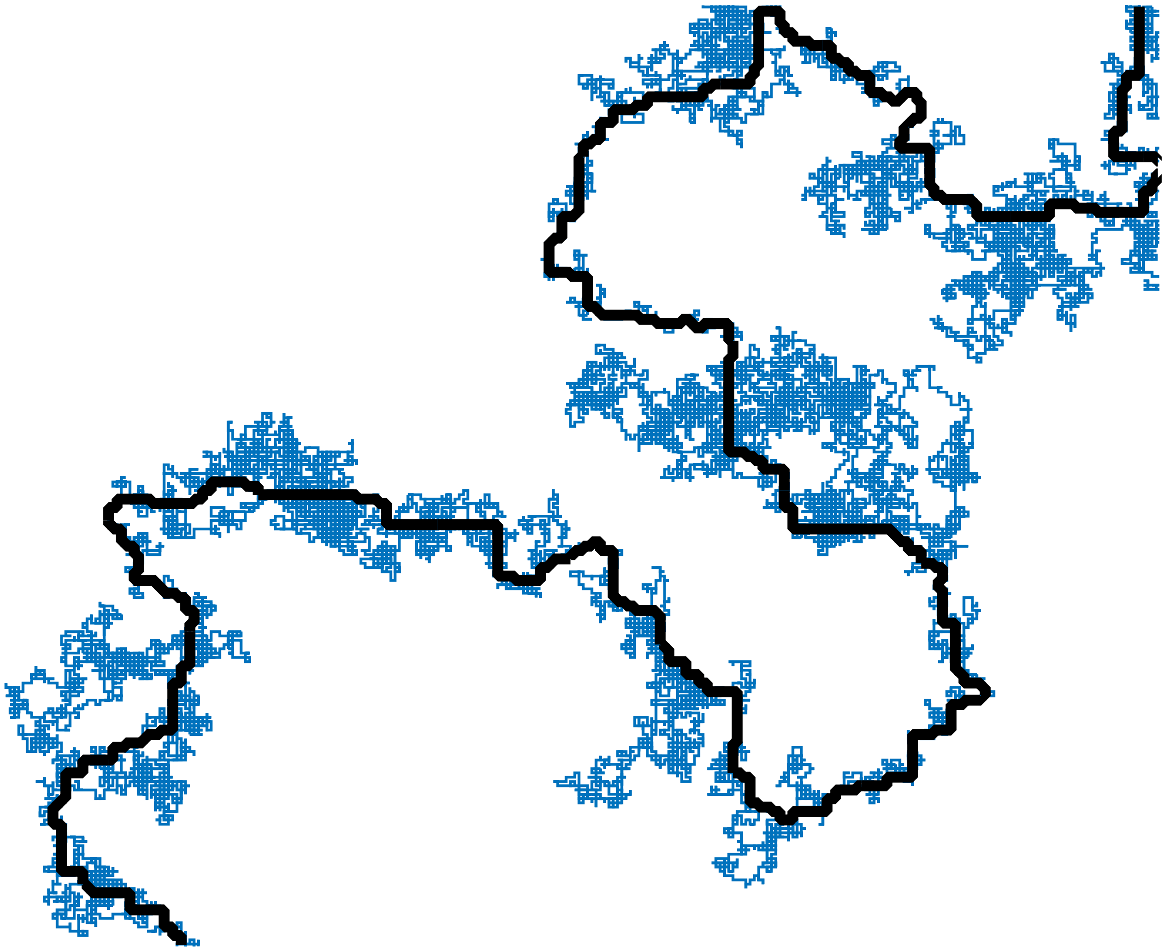

Is the intersection of two boundaries rectifiable?

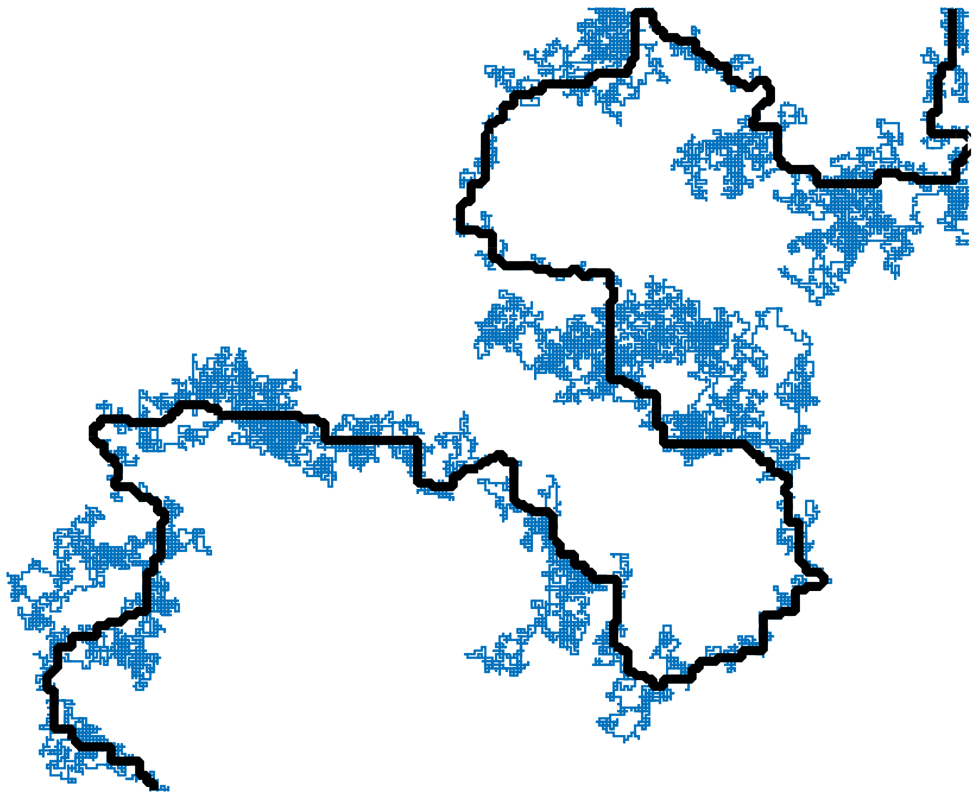

Intersection should have dimension 3/4 (= dim of cut-points).
Self-similar sets of dimension < 1 are rectifiable.
Rectifiable sets characterized by Jones' β-numbers.

For a dyadic square Q, $\beta(Q)$ measures how close E is to a line inside $3Q$. Jones TST says that E lies on a rectifiable curve if and only if

$$\sum_{Q} \beta(Q)^2 \ell(Q) < \infty,$$

where the sum is over all dyadic squares.

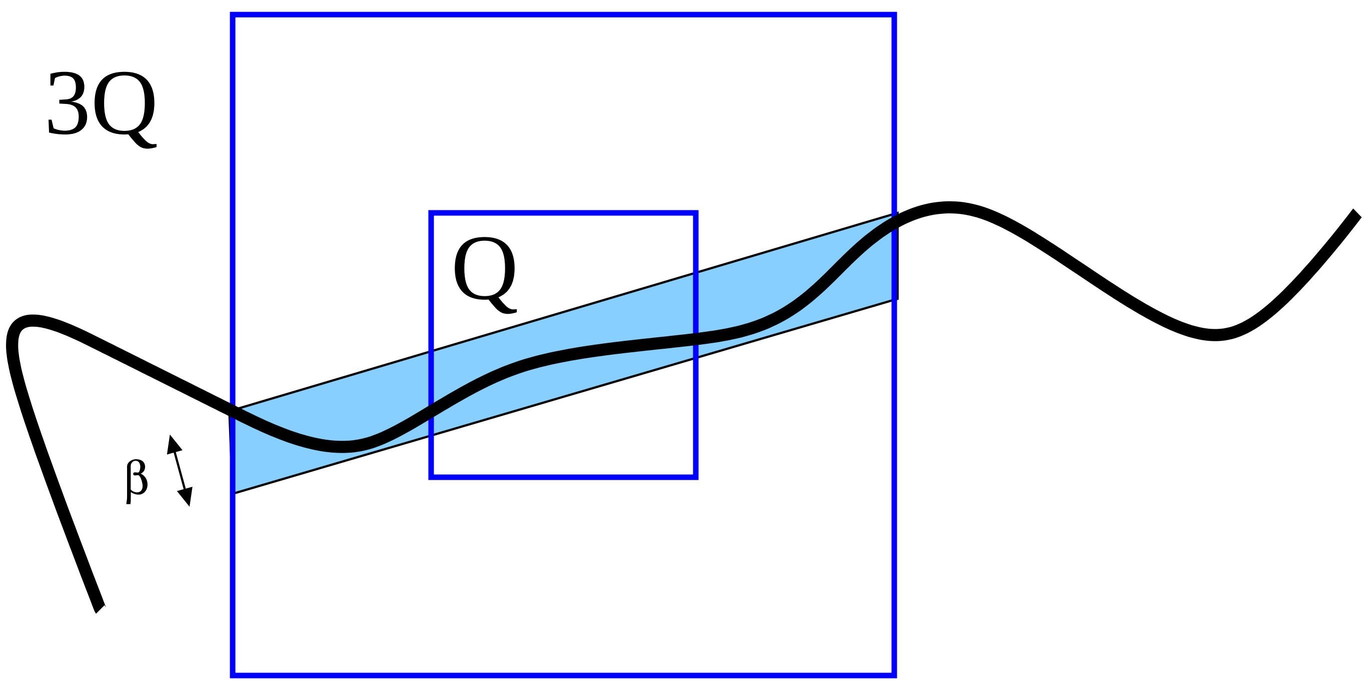

Call a square ϵ-dense for a given trace if the trace comes within ϵ of every point of the square.

Since the trace has Hausdorff dimension 2, there are many such squares at every scale. (Otherwise the trace is porus and has dimension < 2.)

Inside ϵ-dense squares, we can draw paths that are ϵ close to straight. Pemantle's proof shows that such squares cannot "line up". Can they "percolate", that is, does their union have connected components of macroscopic size (comparable to diameter of the trace)?

PART III: THE GRAPH OF COMPLEMENTARY COMPONENTS

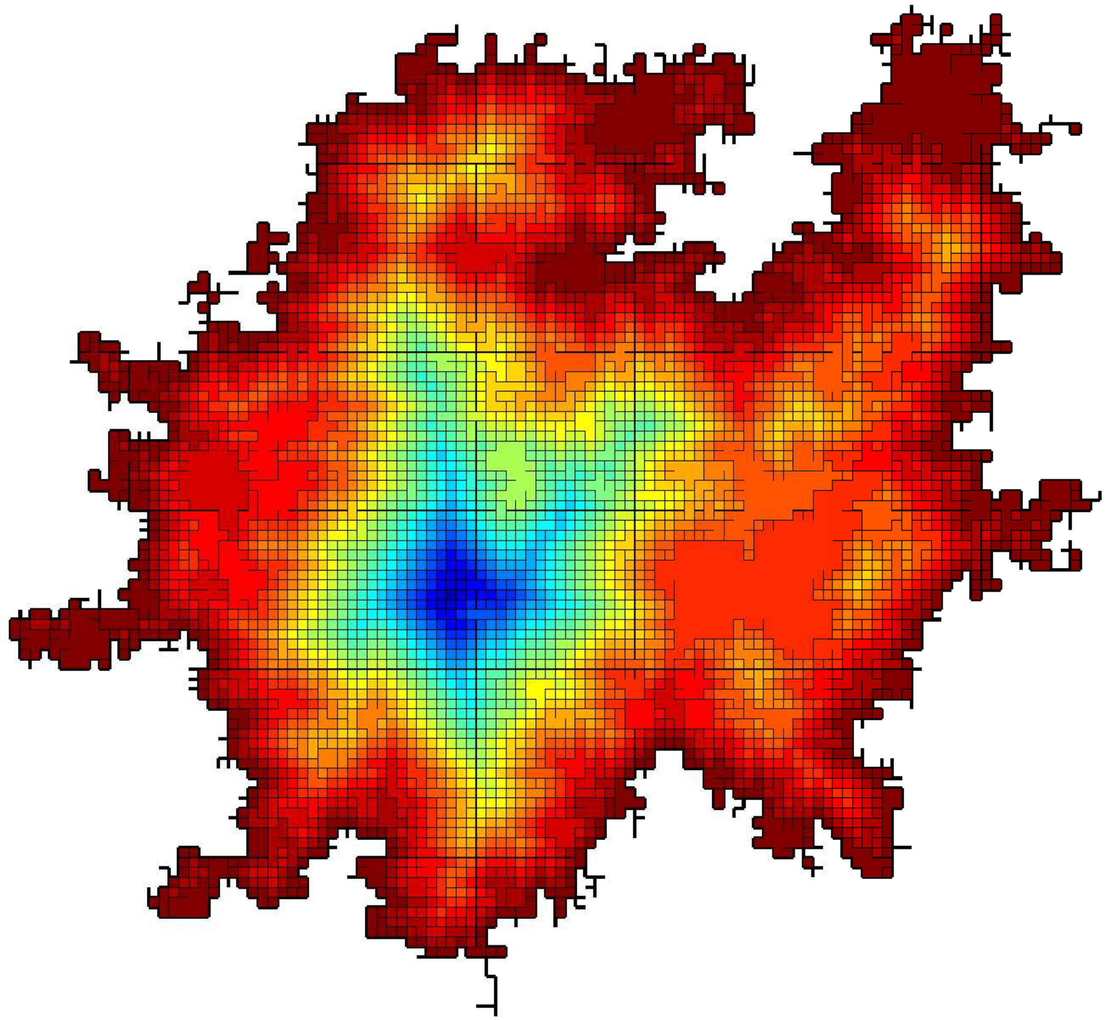

Consider the complementary components of the Brownian trace as vertices of a graph, with two being adjacent if their boundaries overlap.

Wendelin Werner conjectured this graph is connected, i.e., any two components are connected by a path hitting the trace only finitely often.

What about a path that hits the trace countably often? Hits in a set of small Hausdorff dimension?

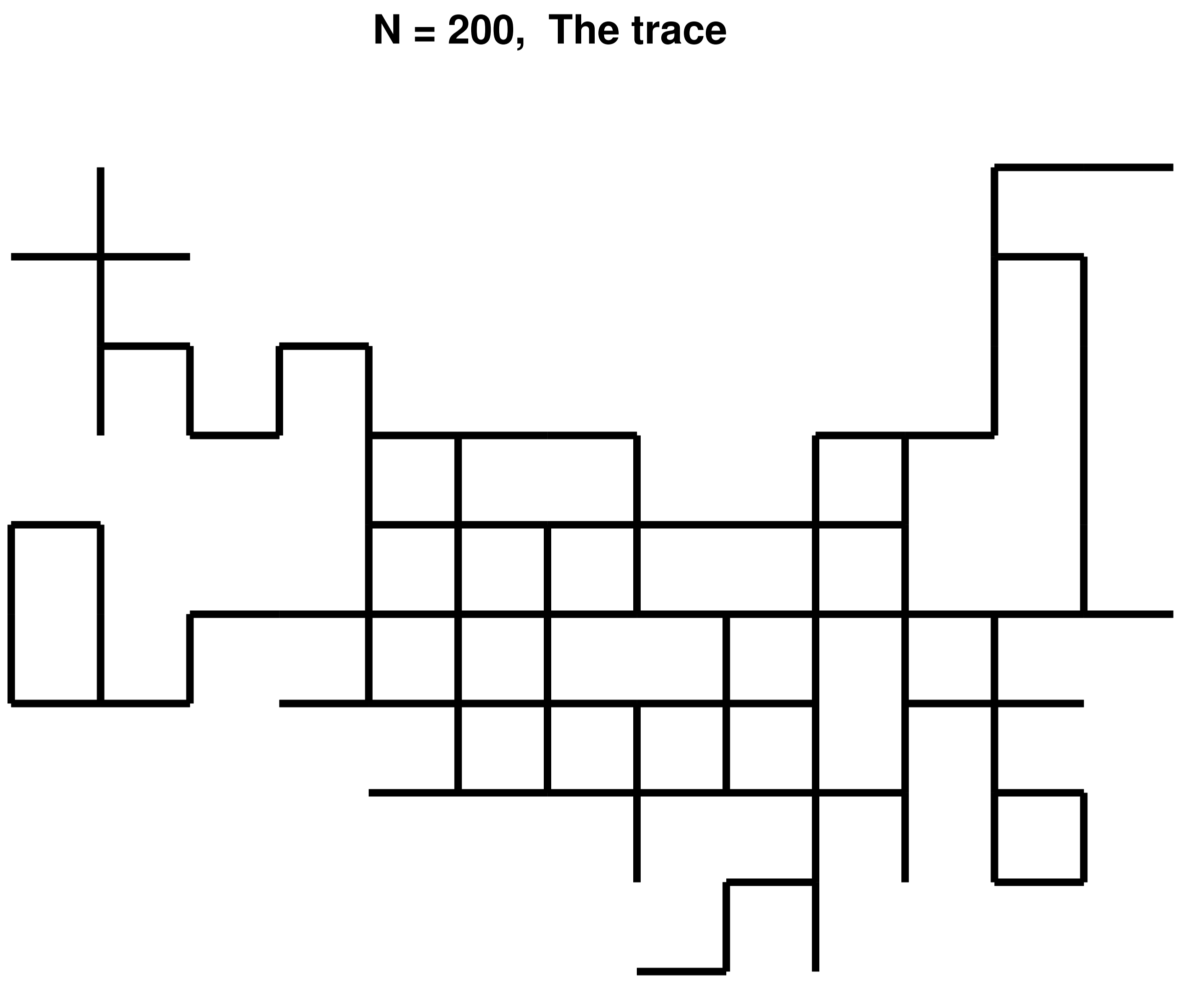

200 random steps on square grid.

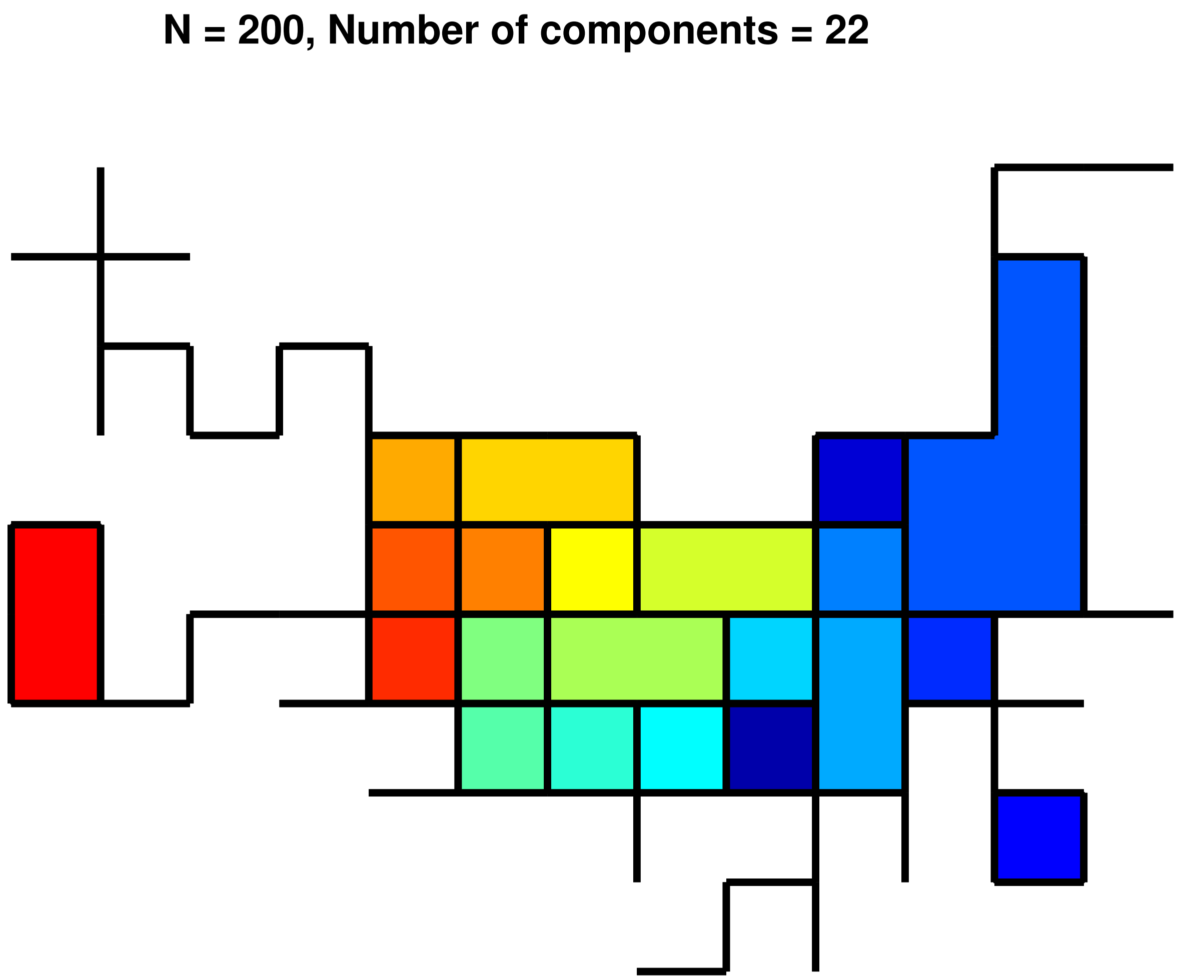

Components form a graph under edge adjacency.

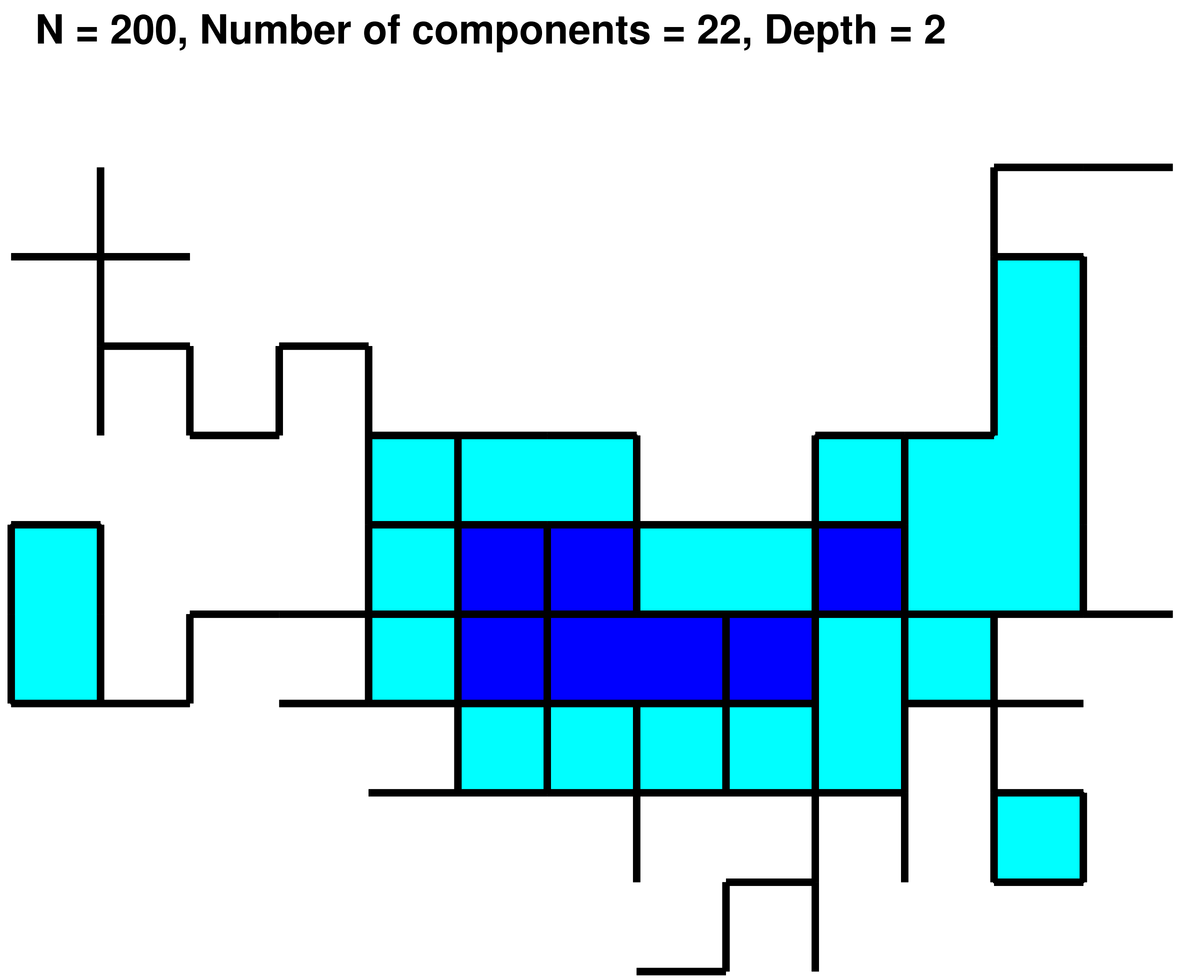

The bounded components colored by graph distance to outer component.

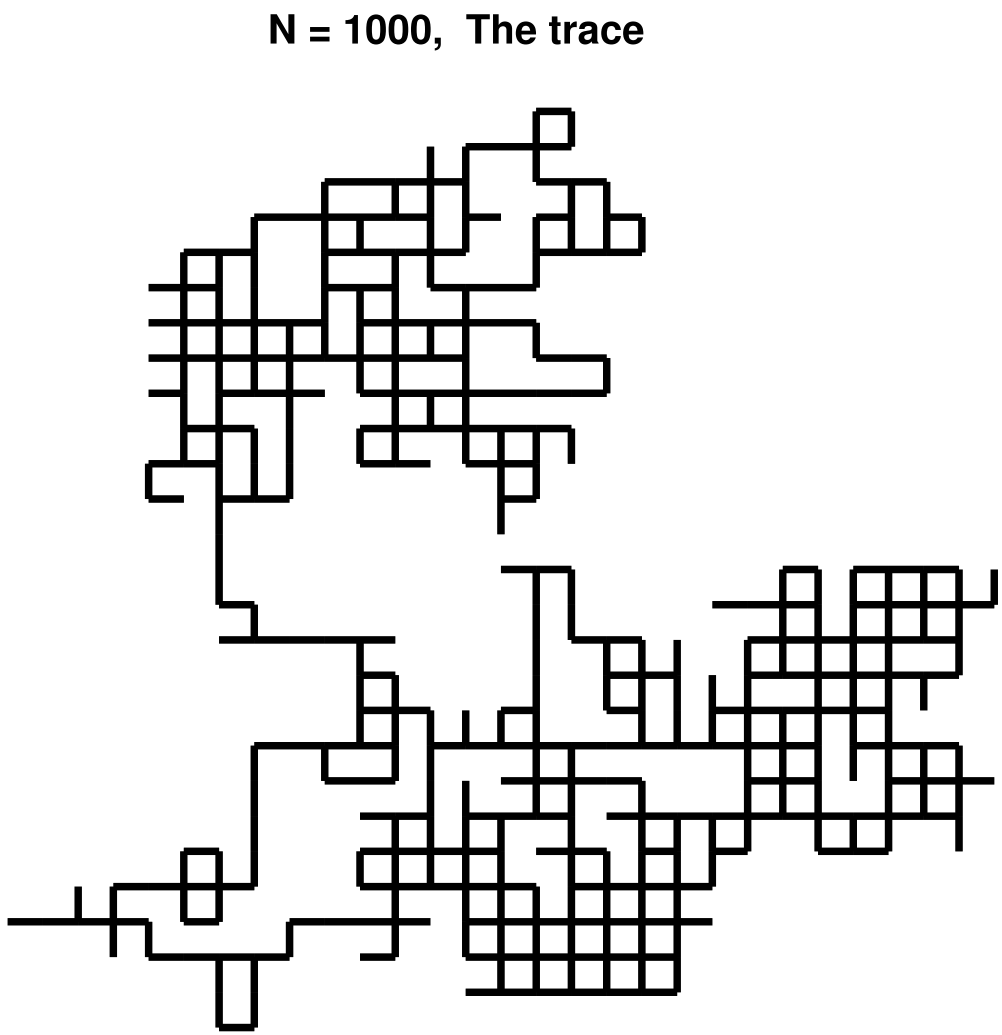

1000 random steps on square grid.

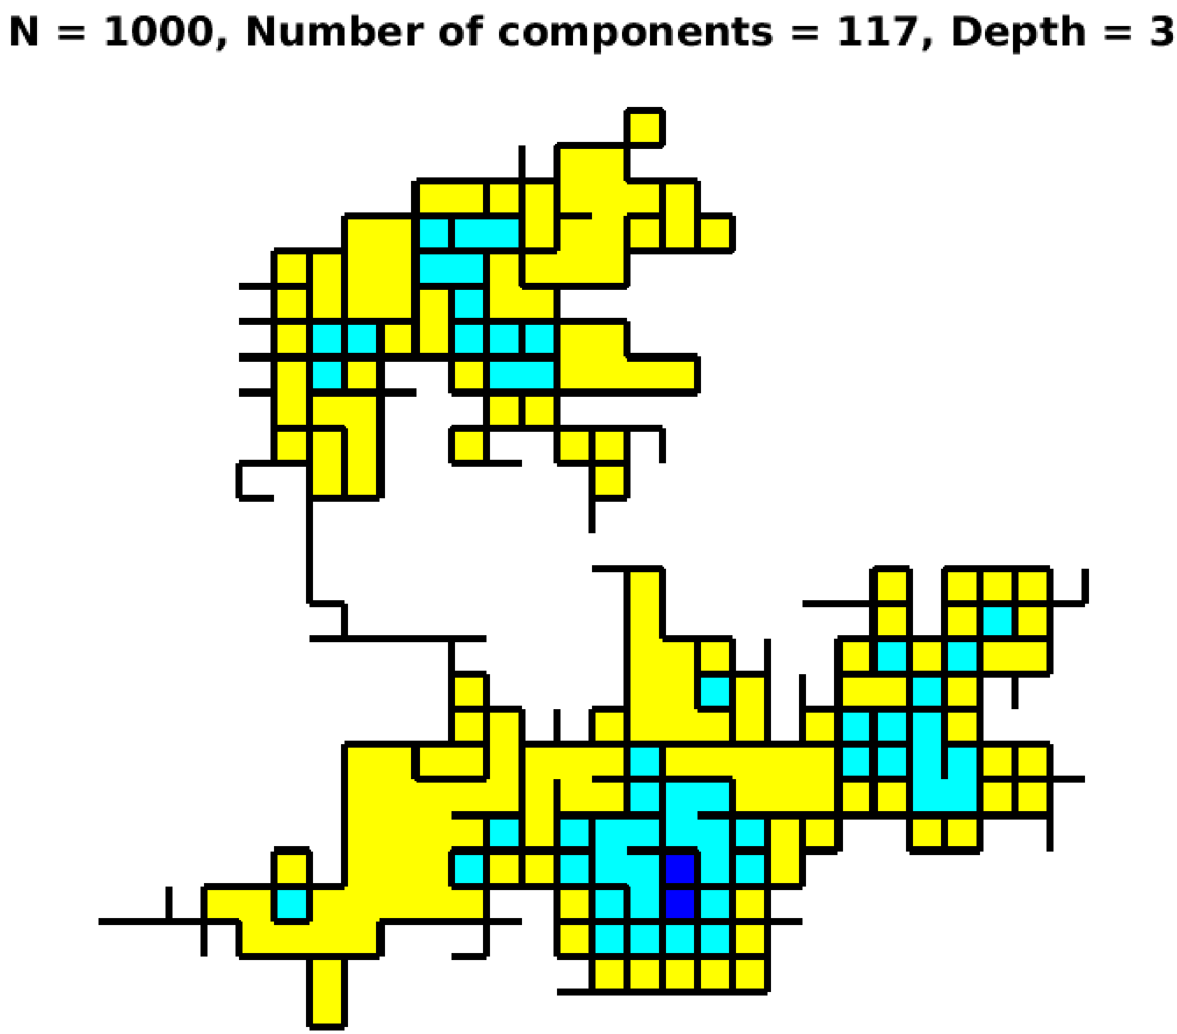

Components colored by graph distance to outer component.

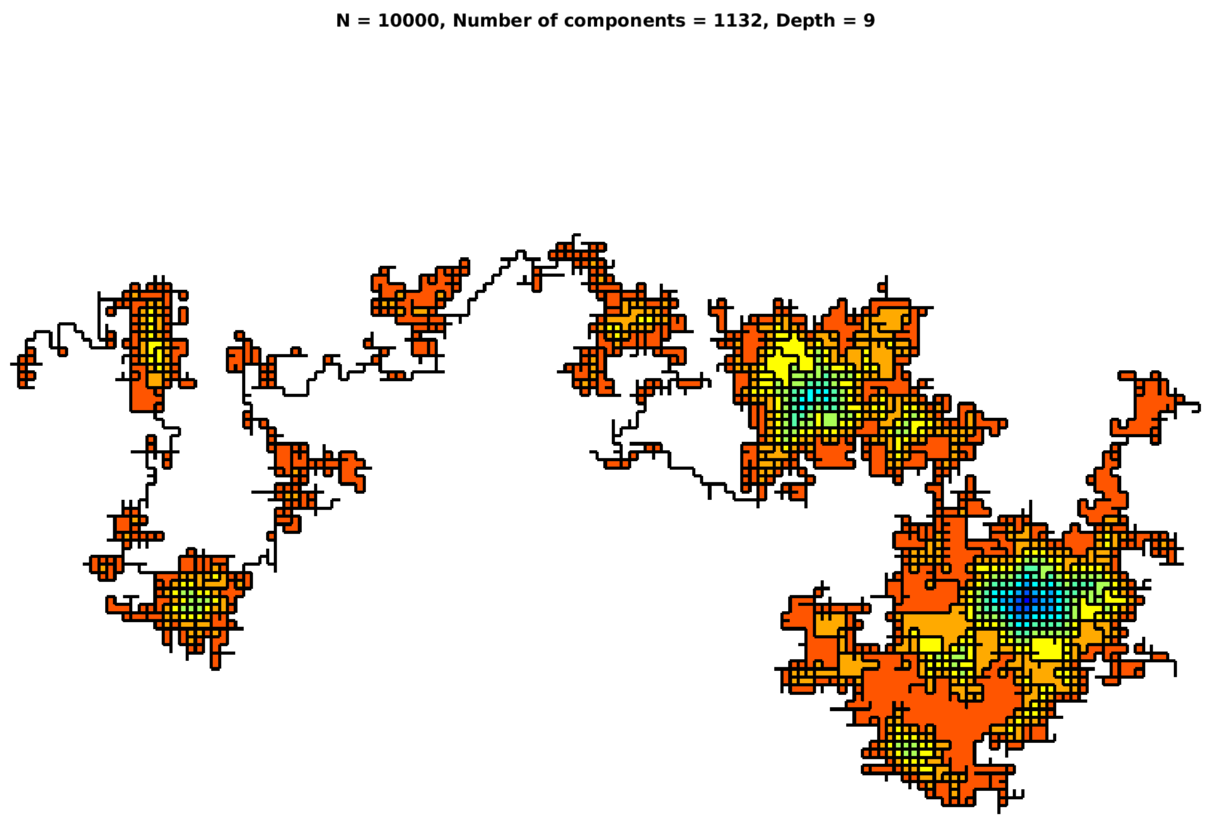

10,000 steps

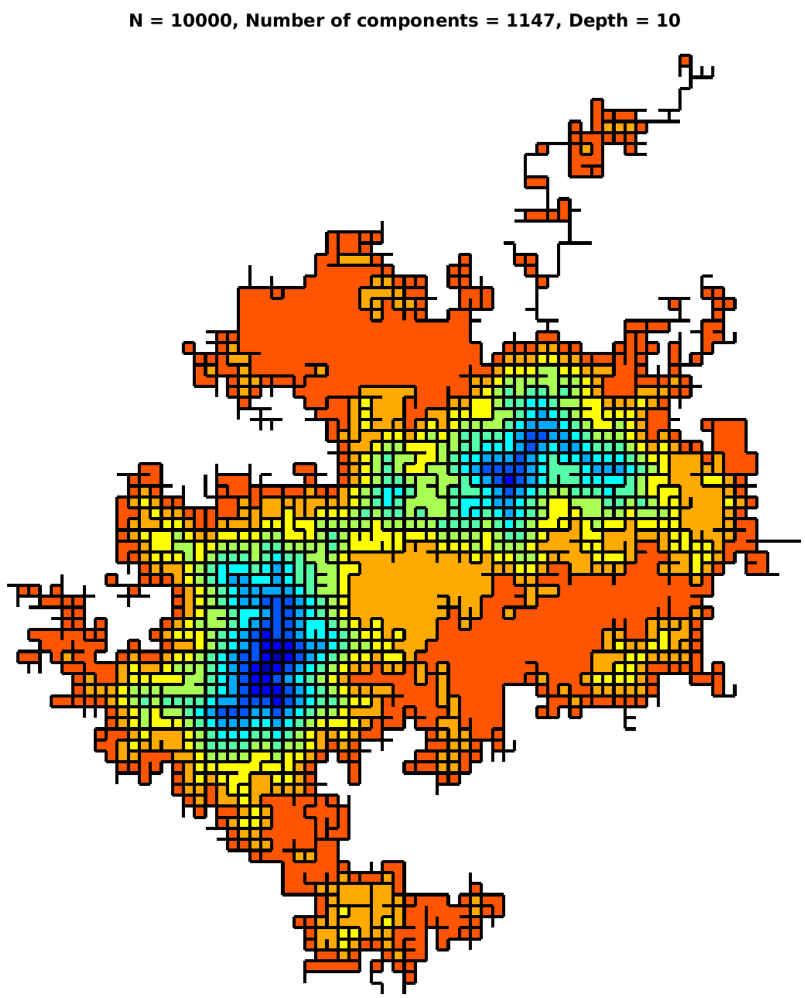

10,000 steps

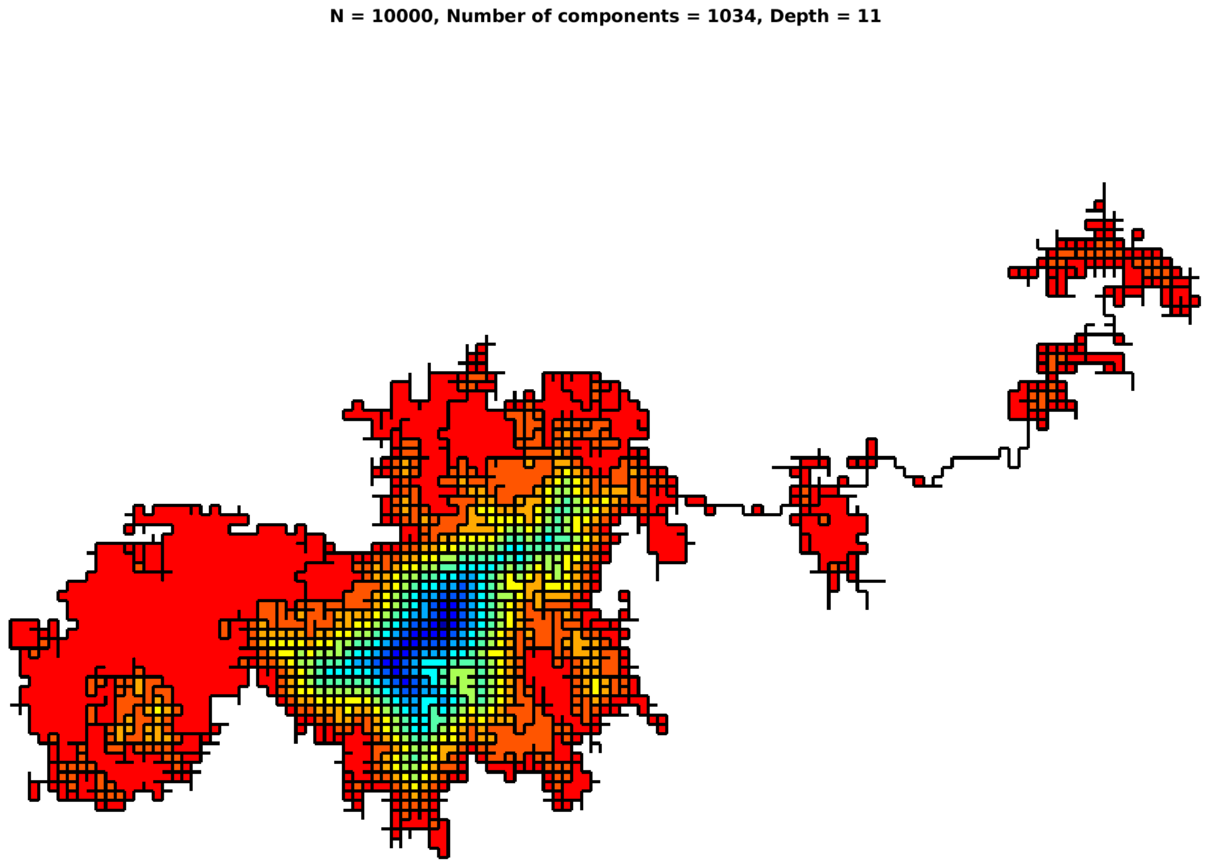

10,000 steps

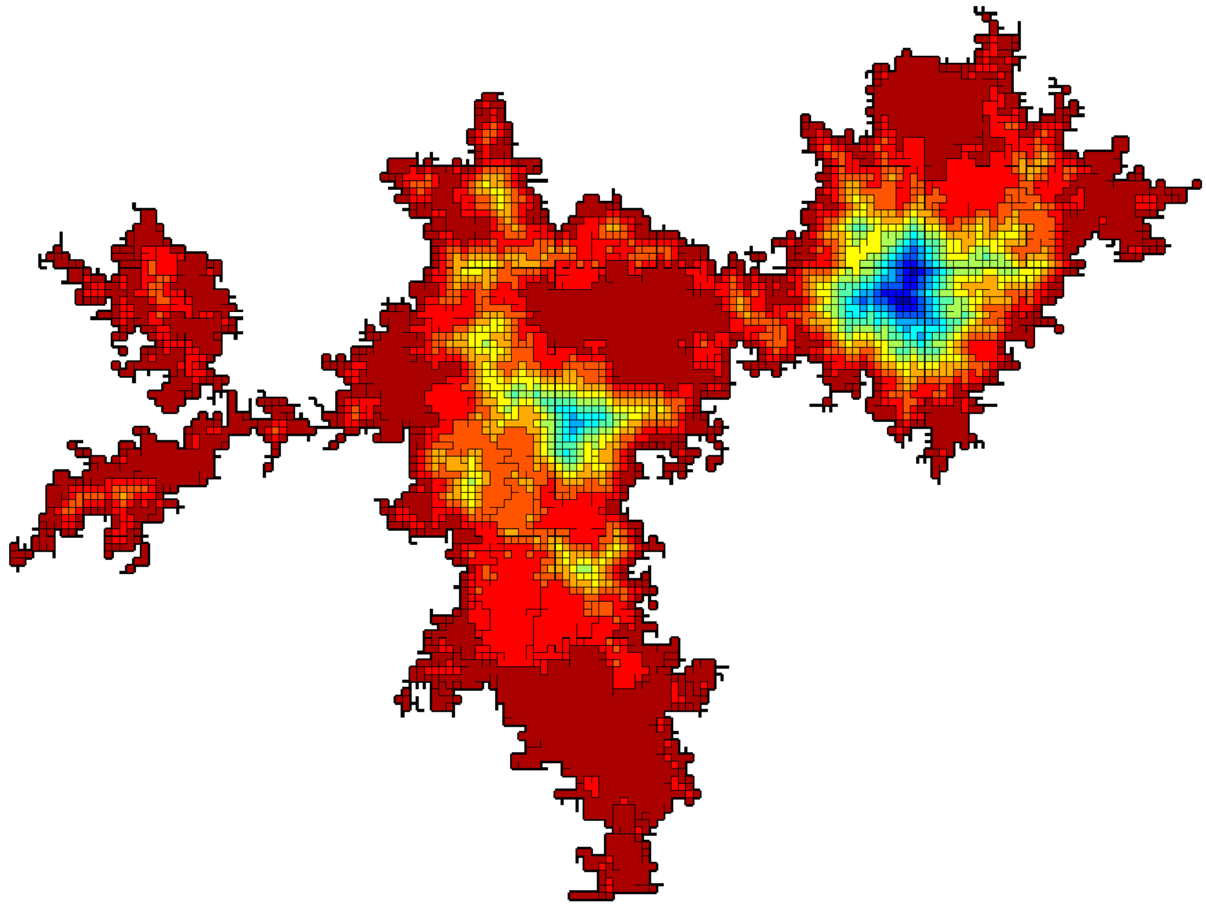

20,000 steps

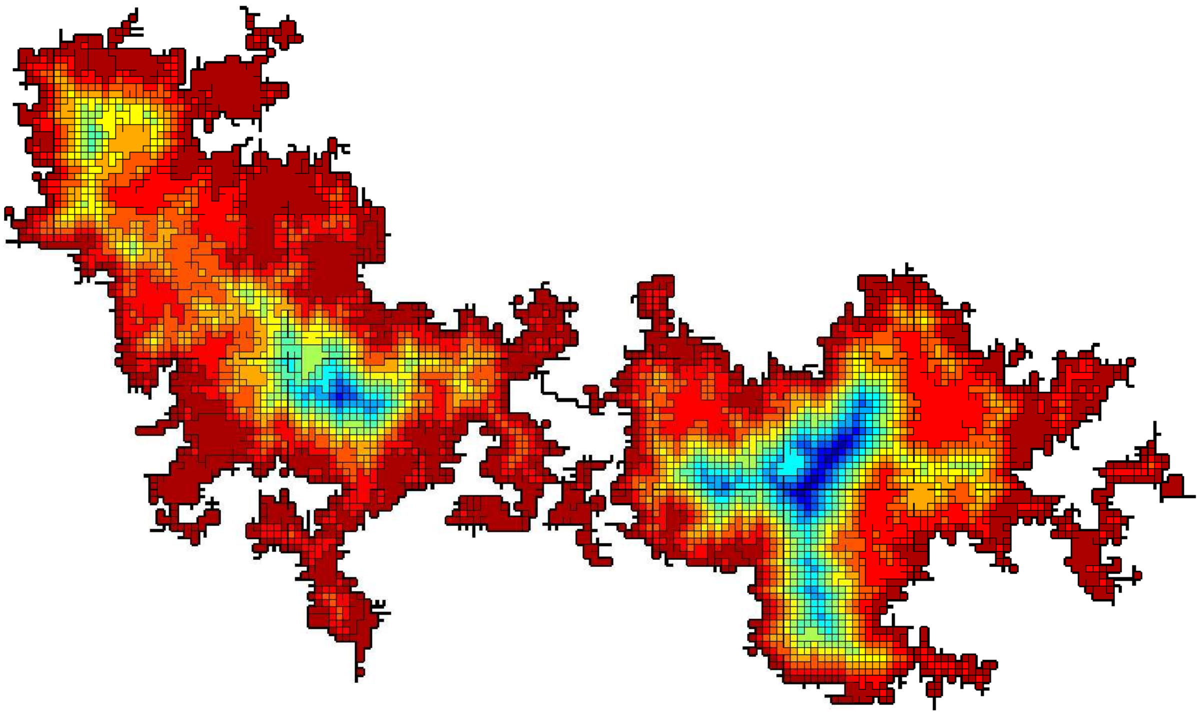

30,000 steps

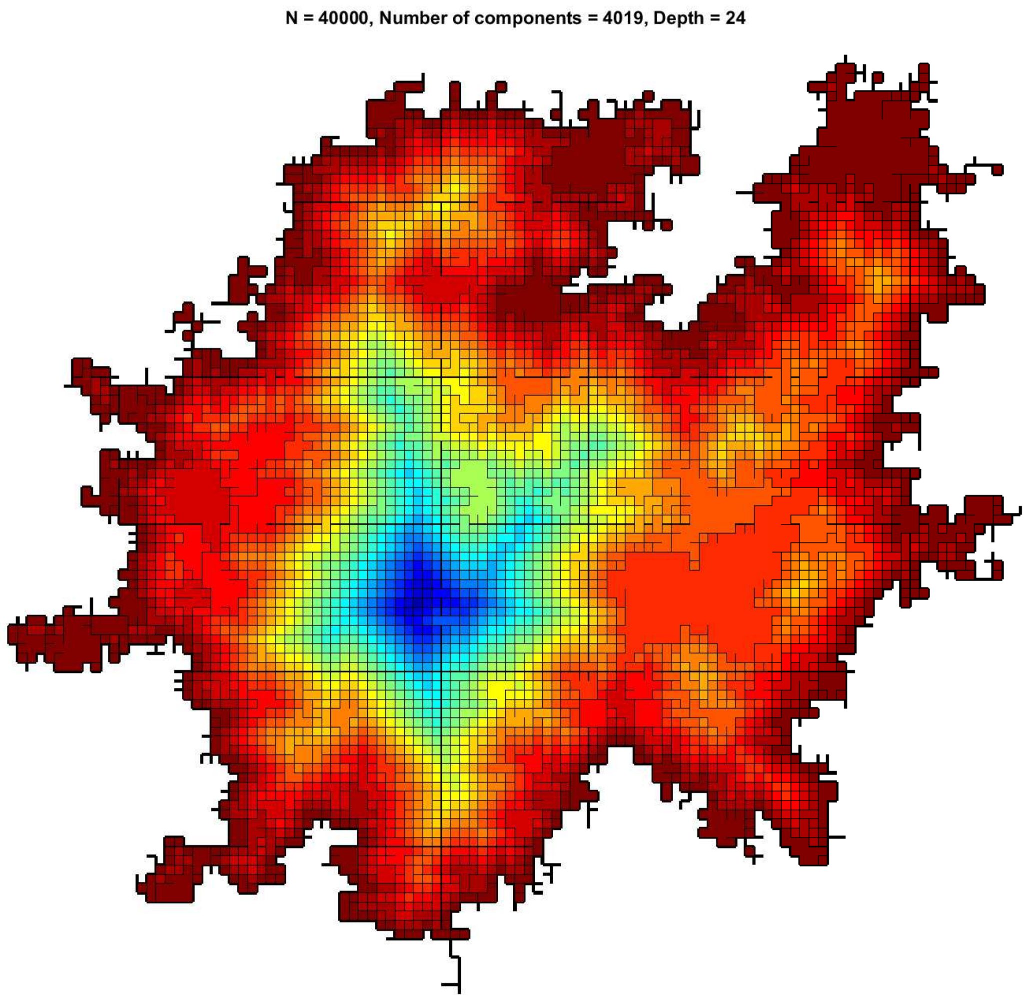

40,000 steps

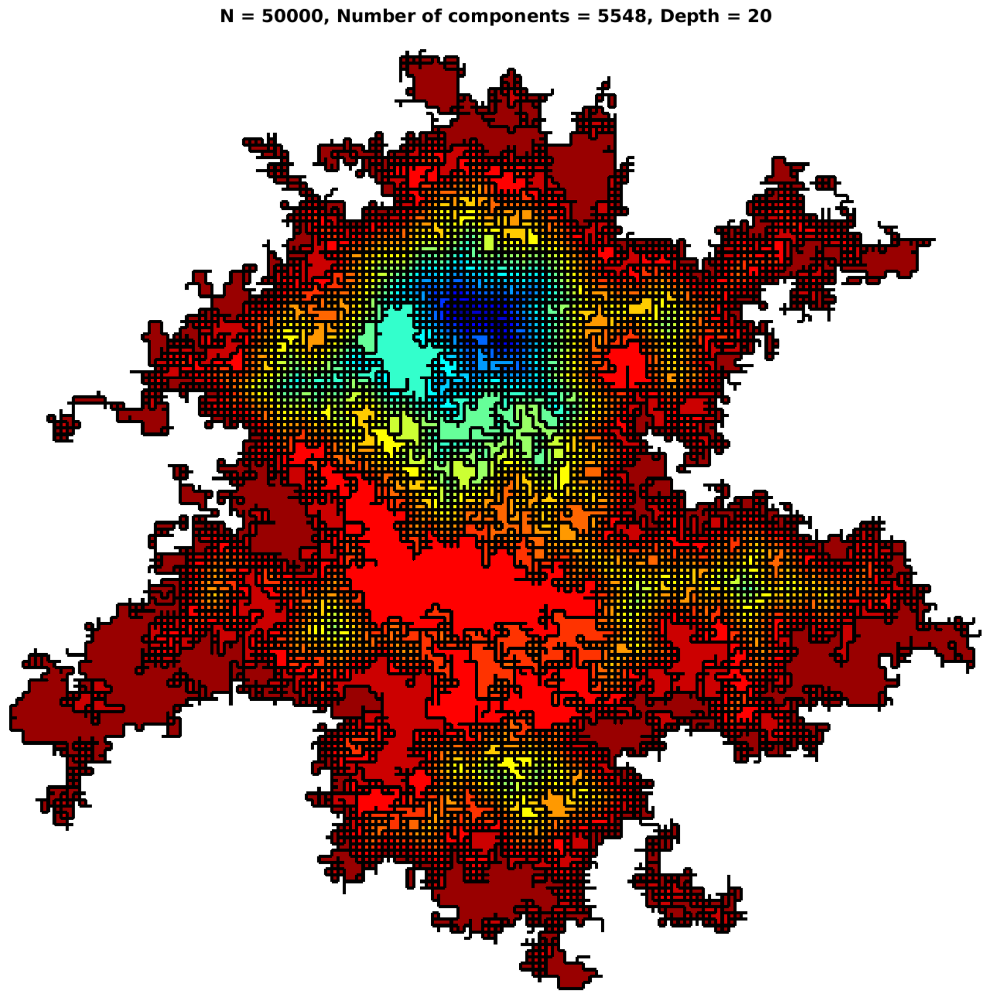

50,000 steps, 10%-component at depth 12

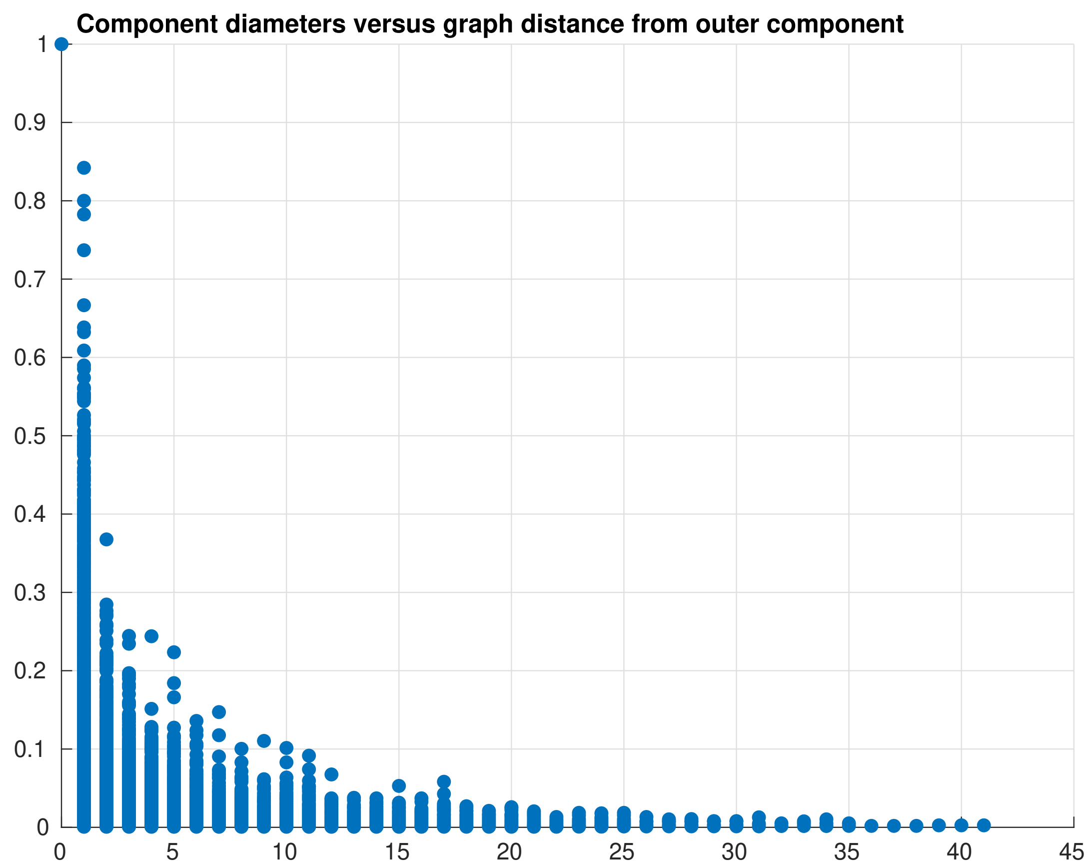

Component diameter versus graph distance from outer component.

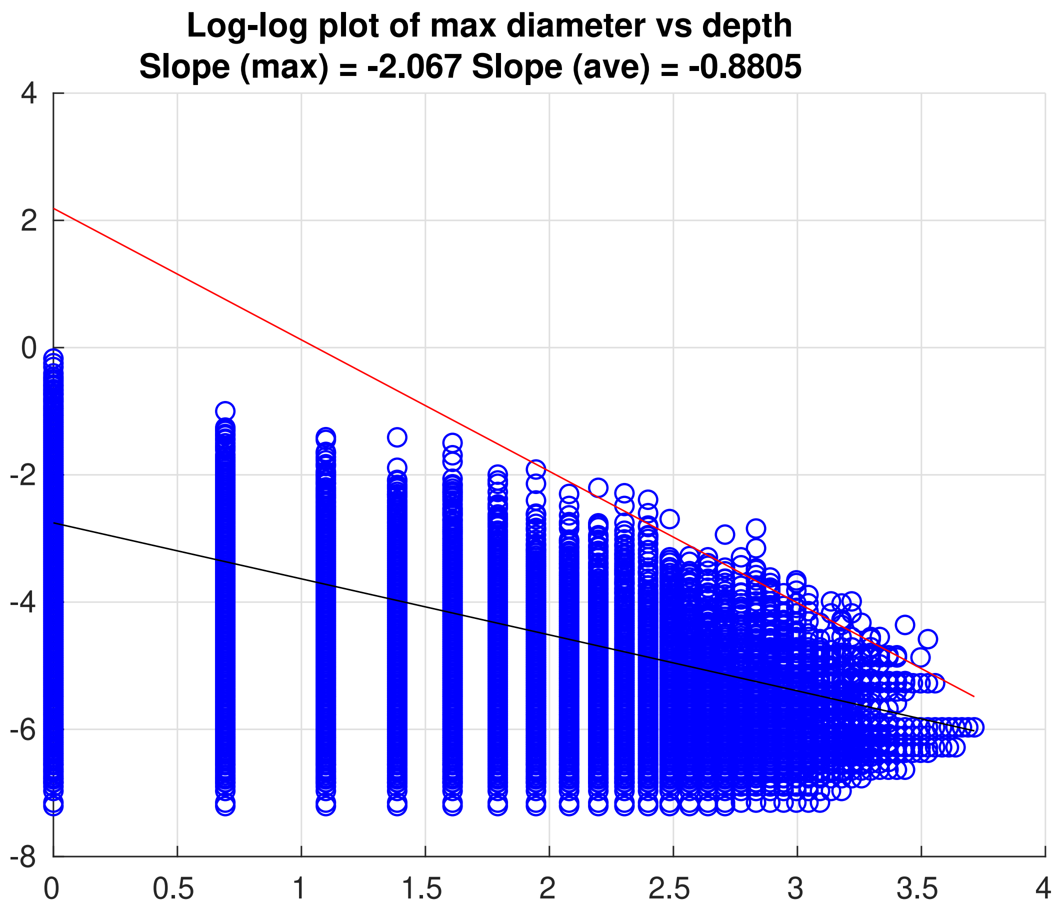

Same plot in log-log coordinates.

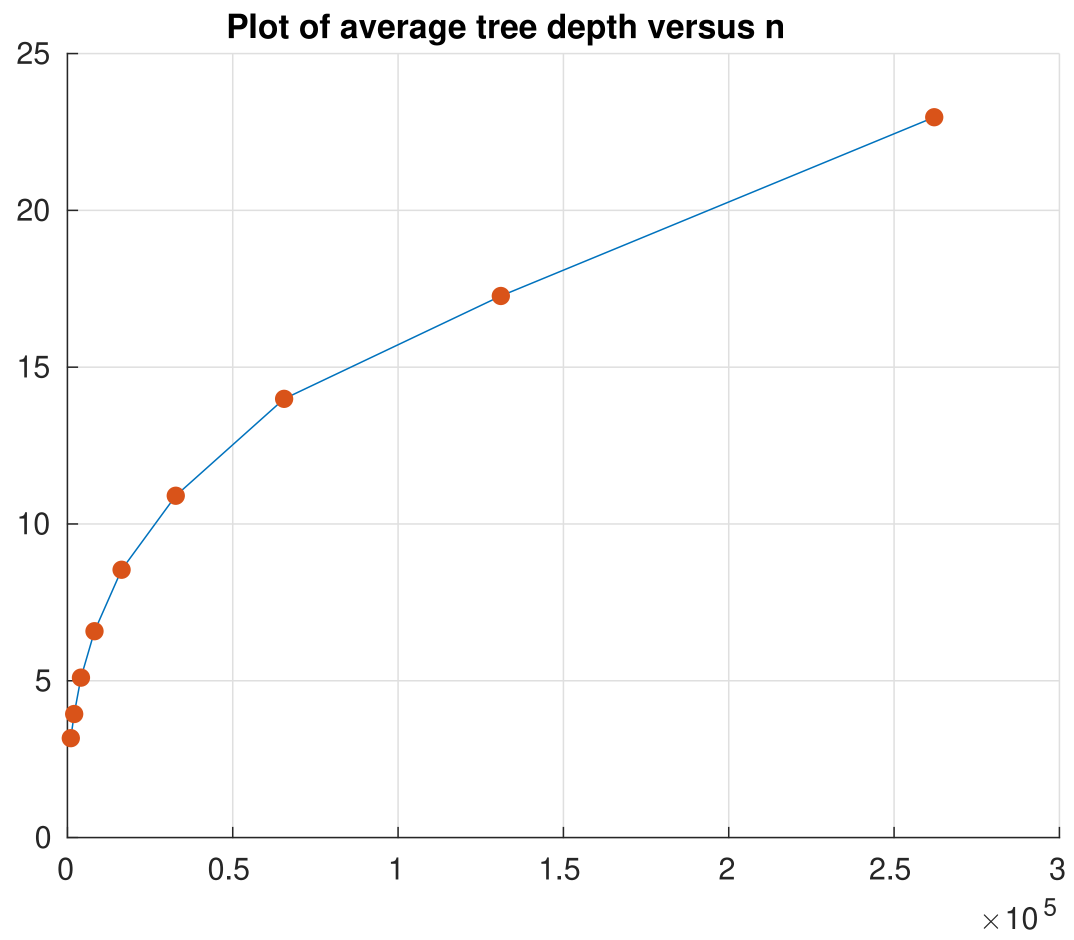

Growth of maximal graph distance (depth) to outer component.

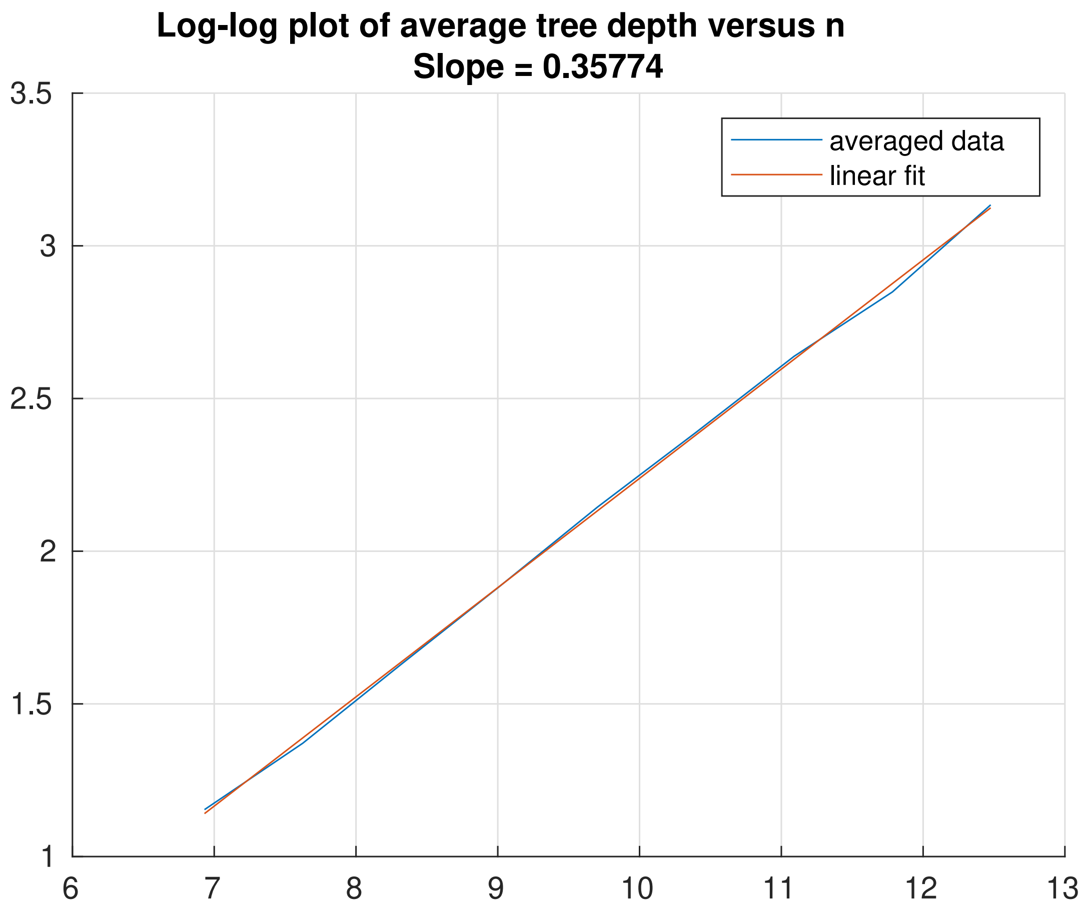

Maximum distance from outer component looks like $n^{.36}$.

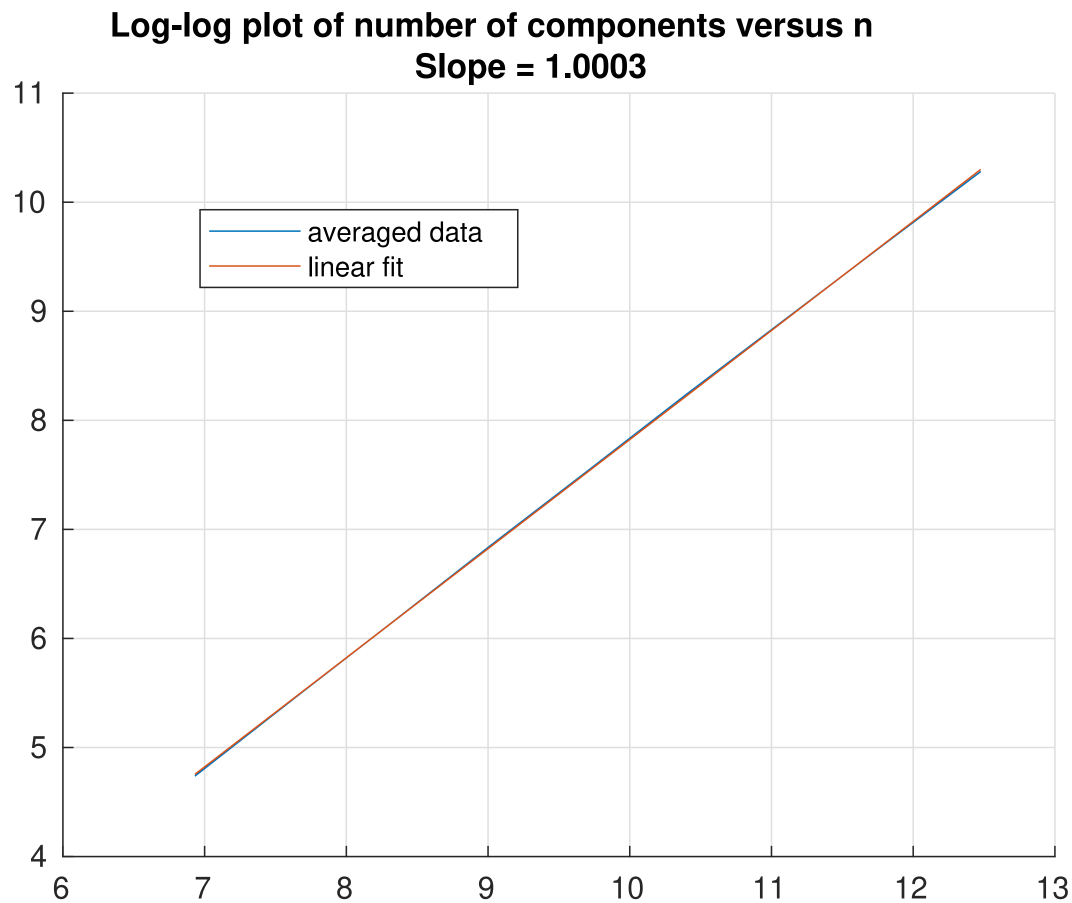

Number of components appears linear in n
Averaged over 100 trials.

How could Werner's conjecture fail?

Need large component surrounded by much smaller components. Would happen if it was surrounded by ϵ-dense squares for every ϵ.

If ϵ-dense squares "percolate", they might form macroscopic loops that cause Werner's conjecture to fail.

PART IV: TRIANGULATION FLOWS

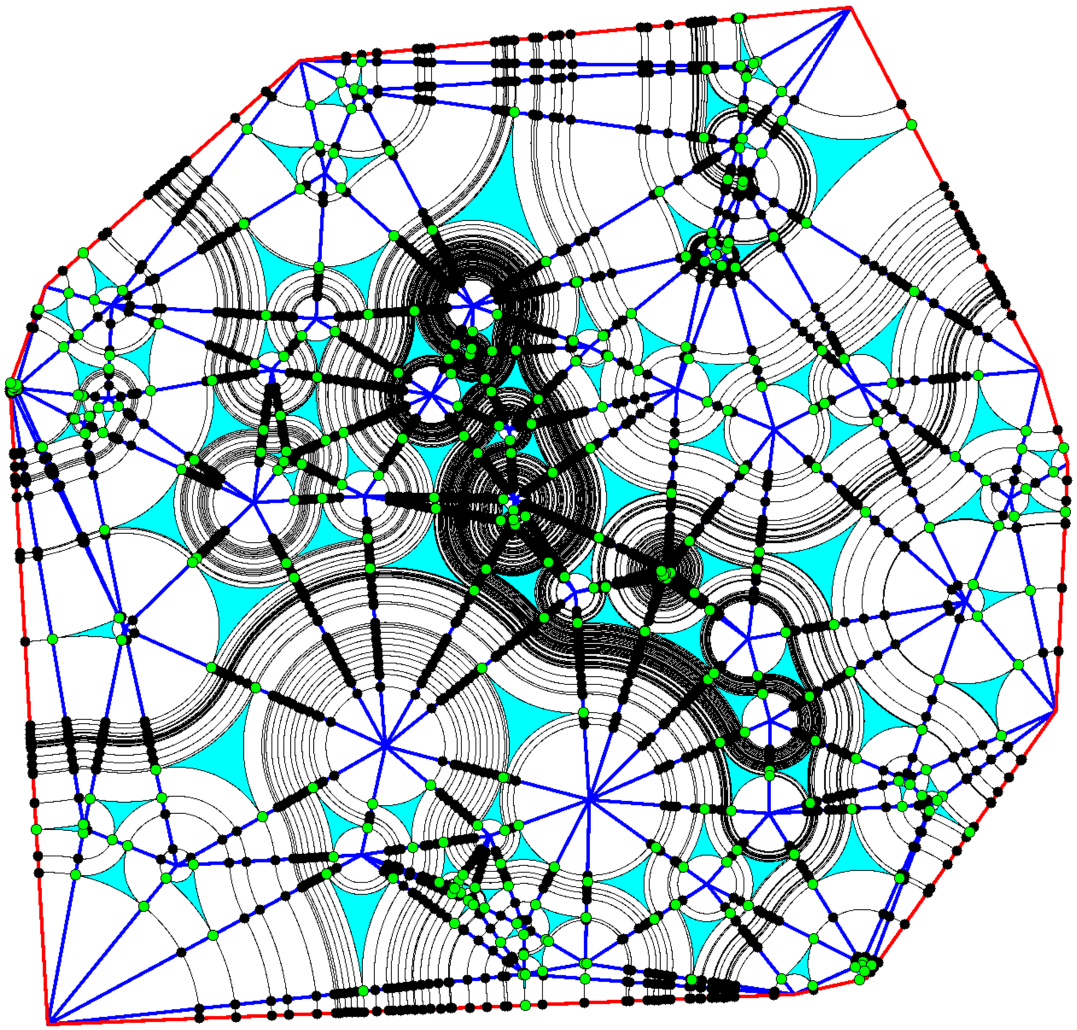

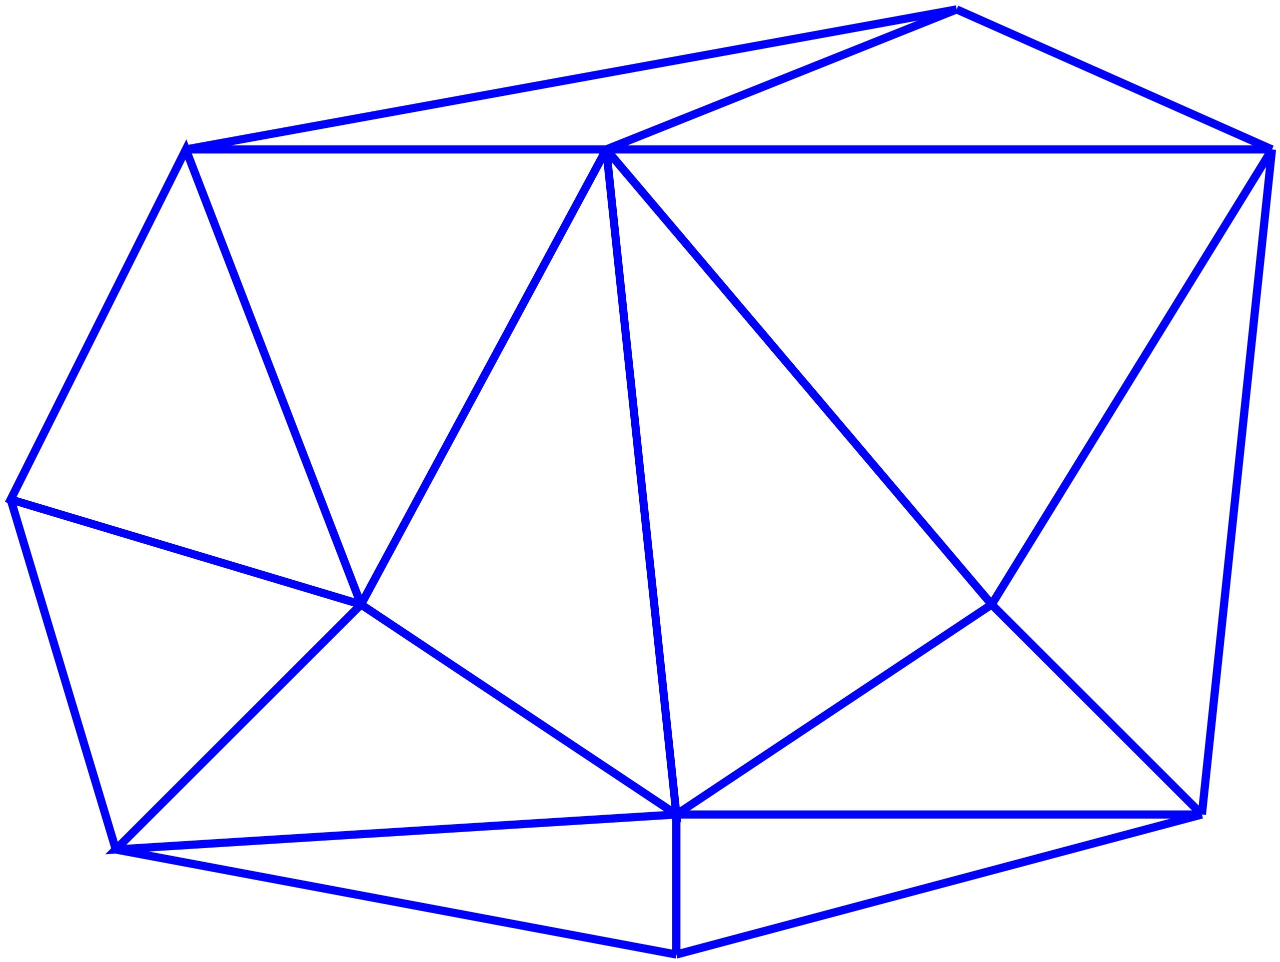

A triangulation: overlapping edges agree.

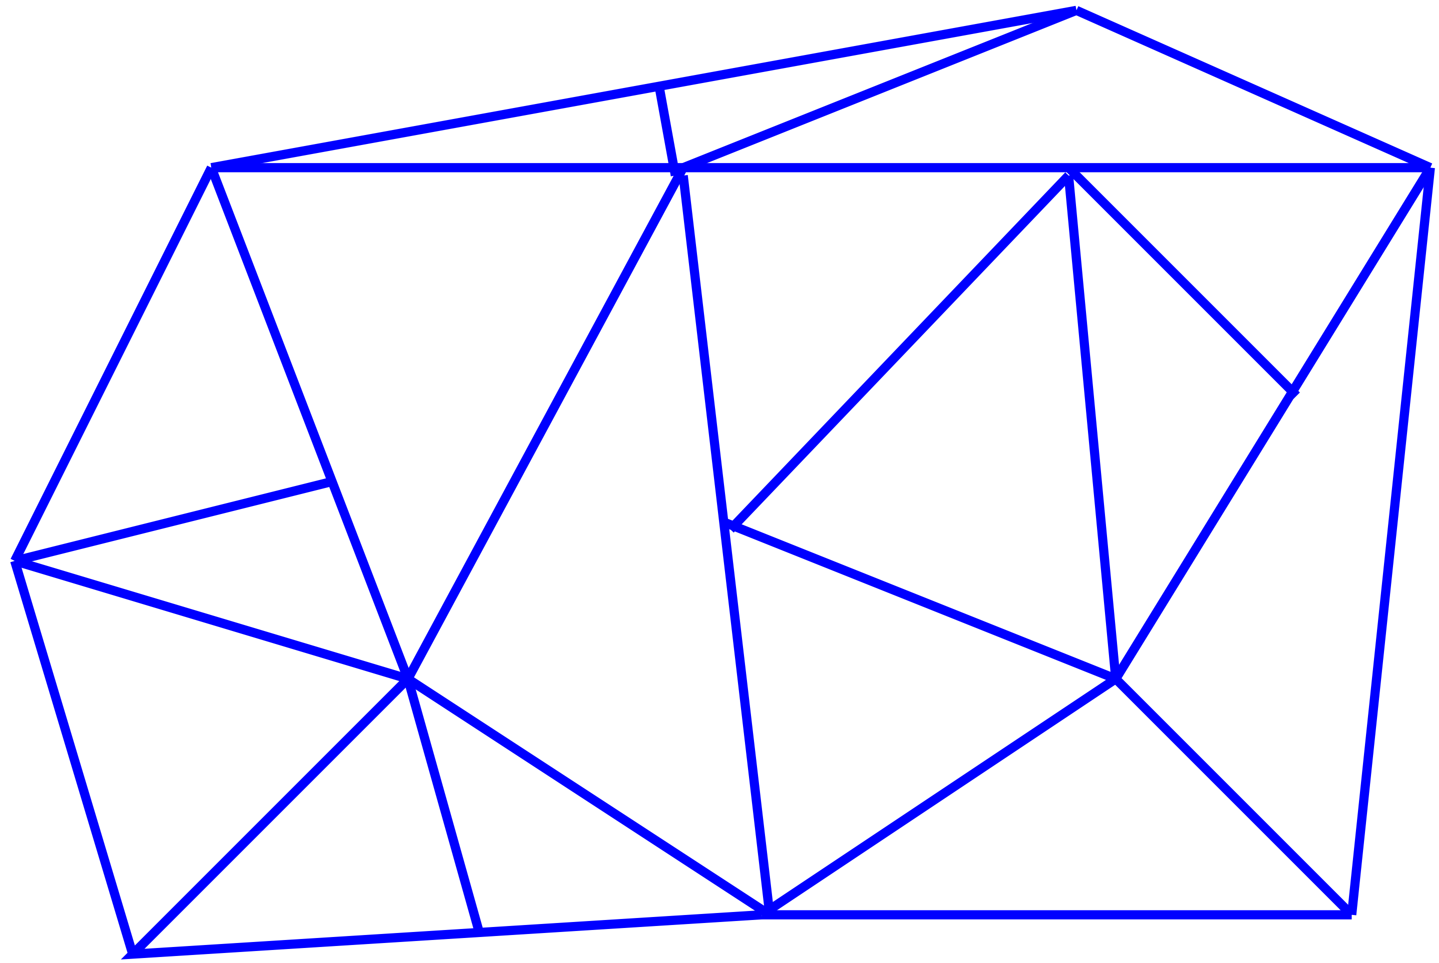

A dissection.

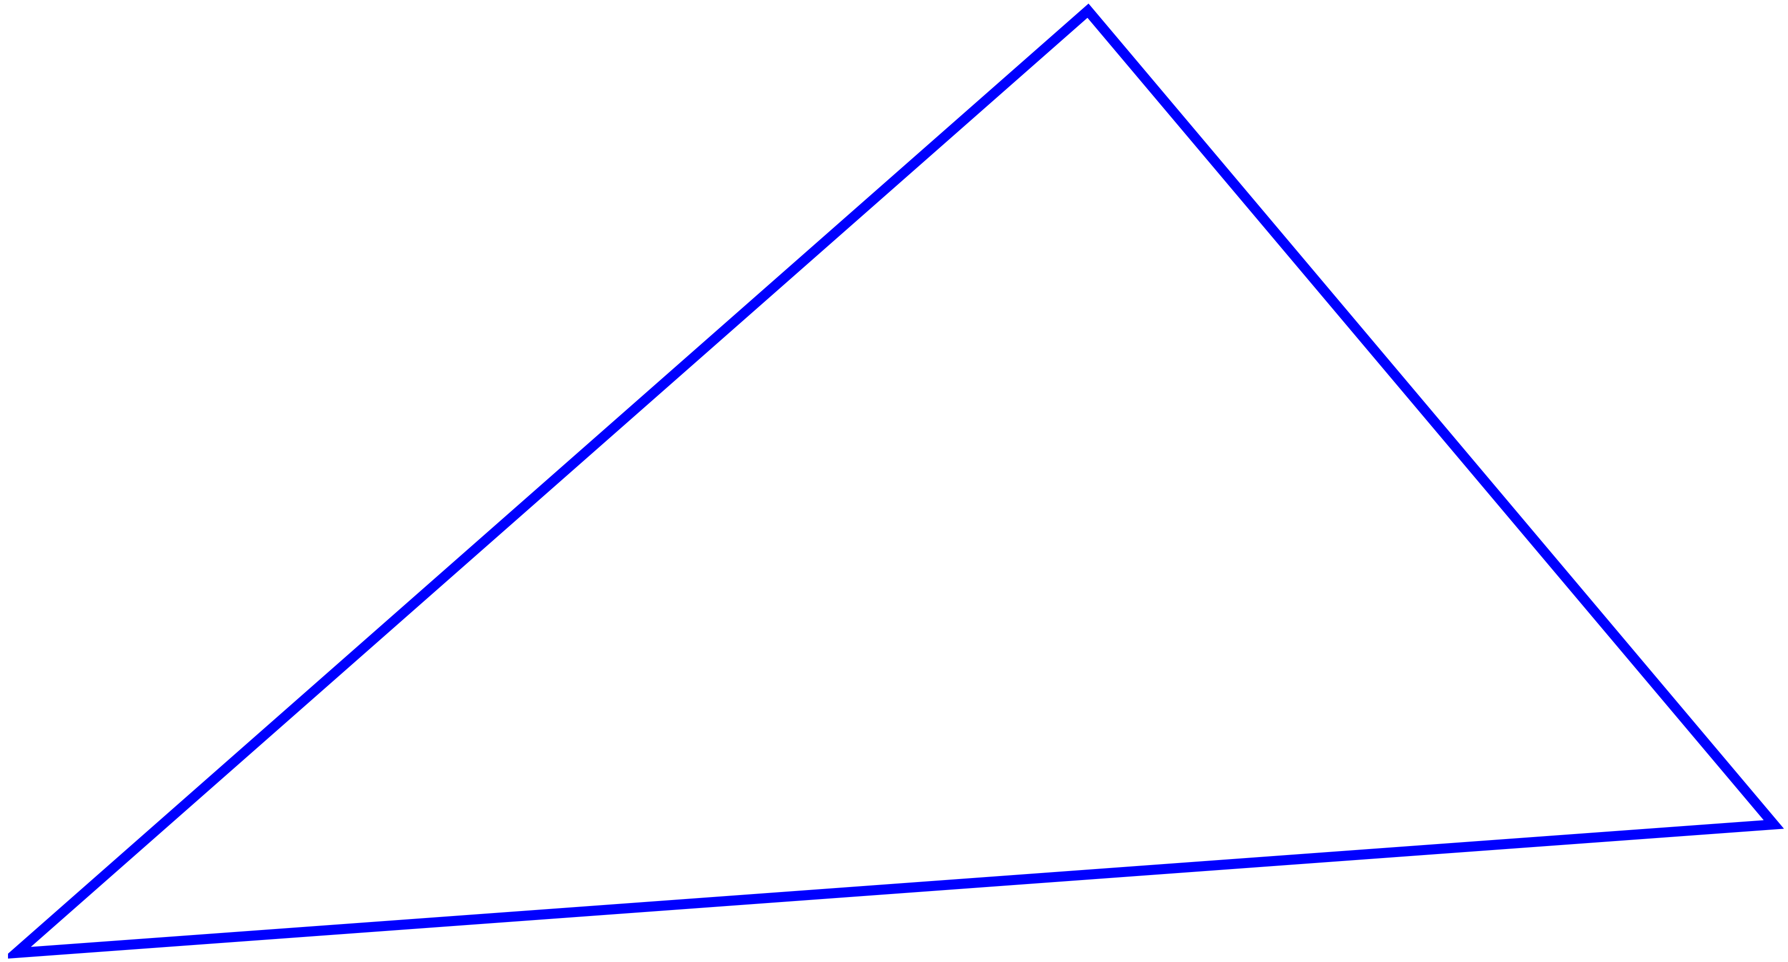

A triangle.

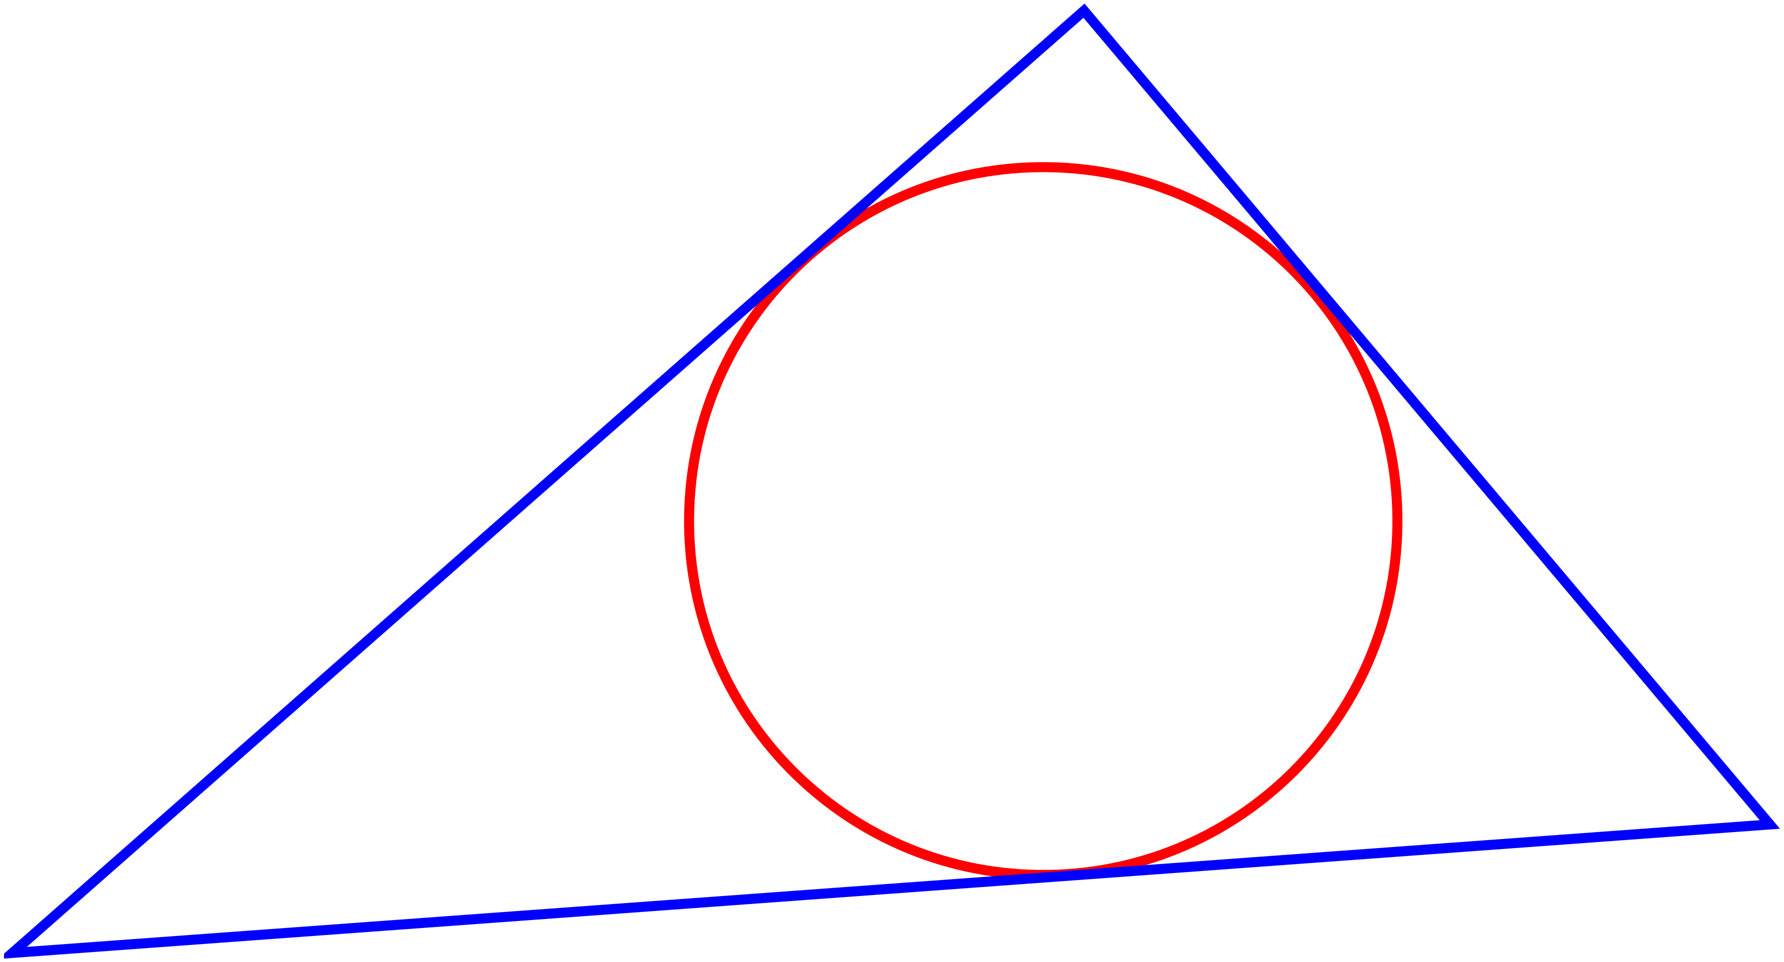

Its in-circle.

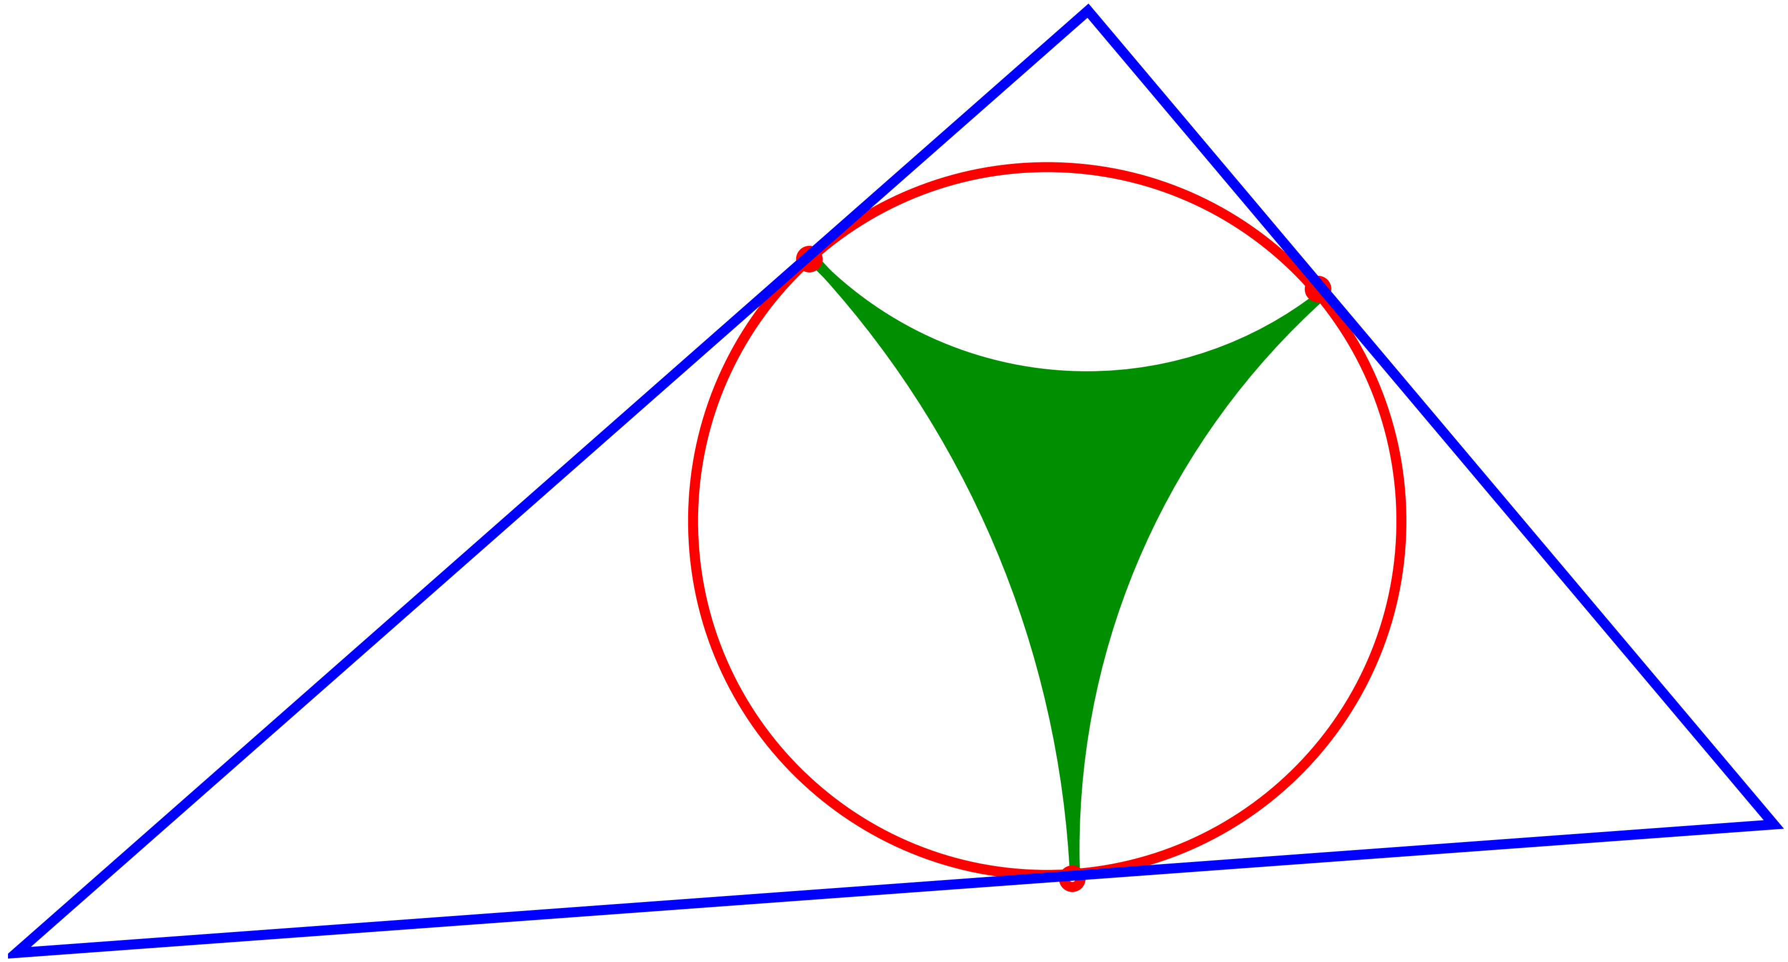

The central region and three sectors (thin version).

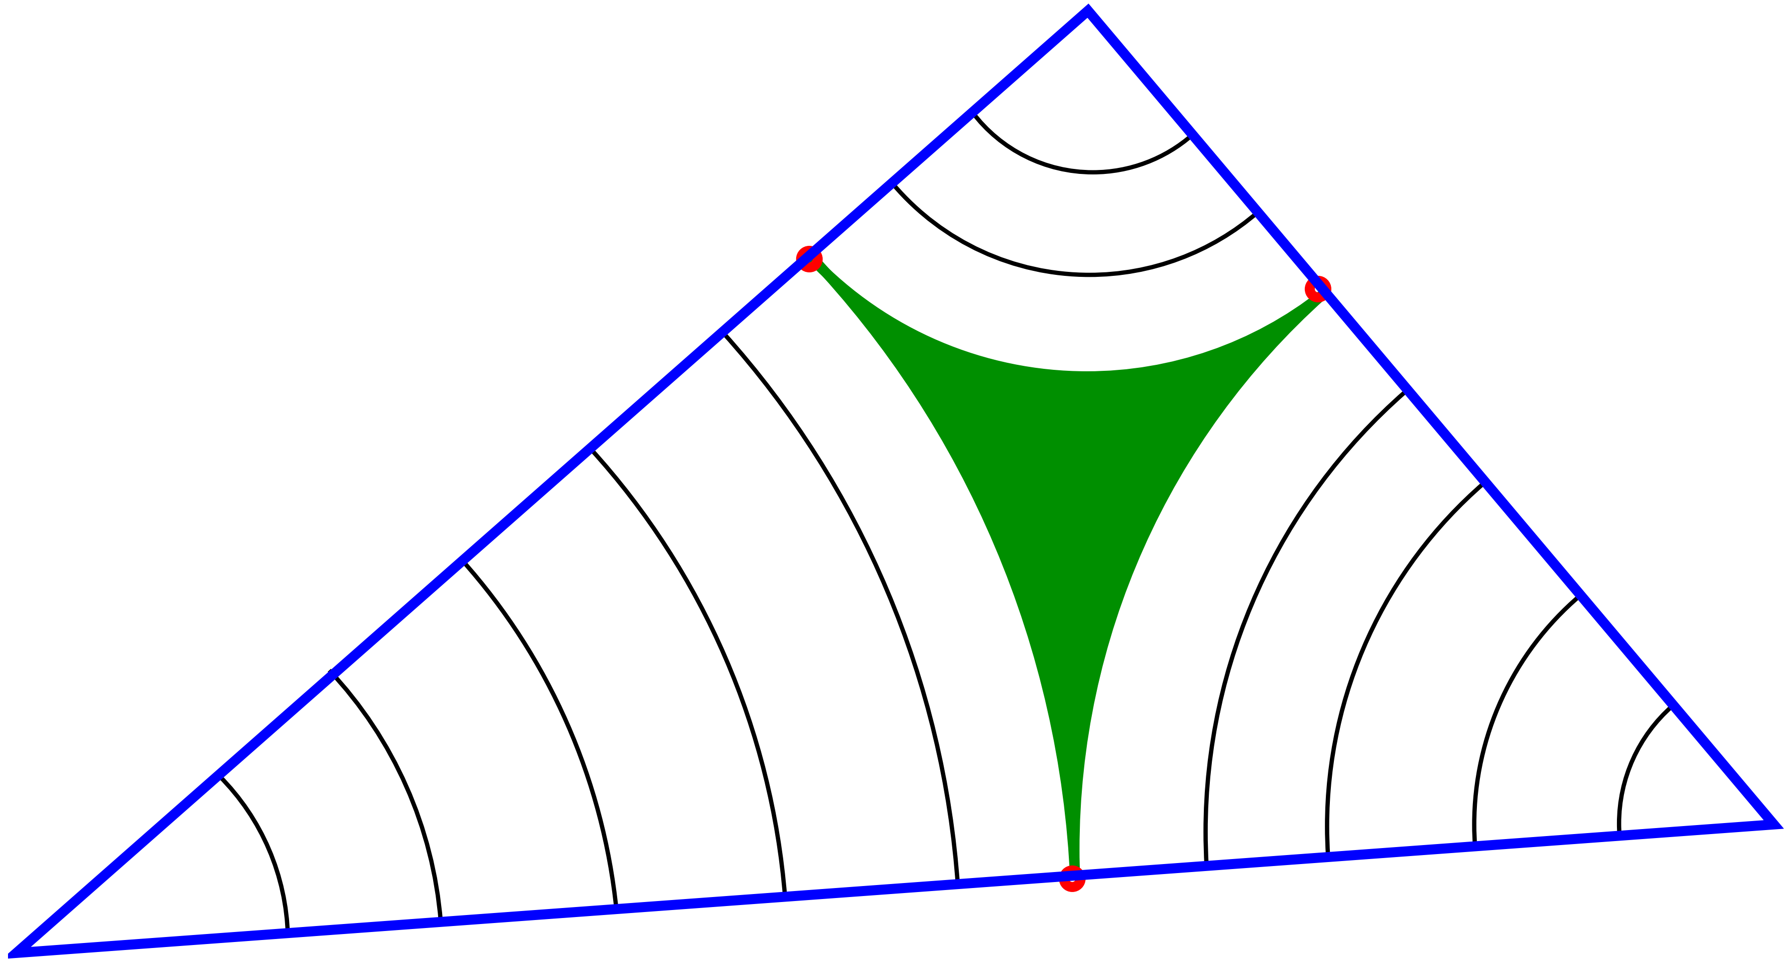

The three sectors are foliated by circular arcs.
Defines flow on a triangulation that stops at boundary or cusp point.

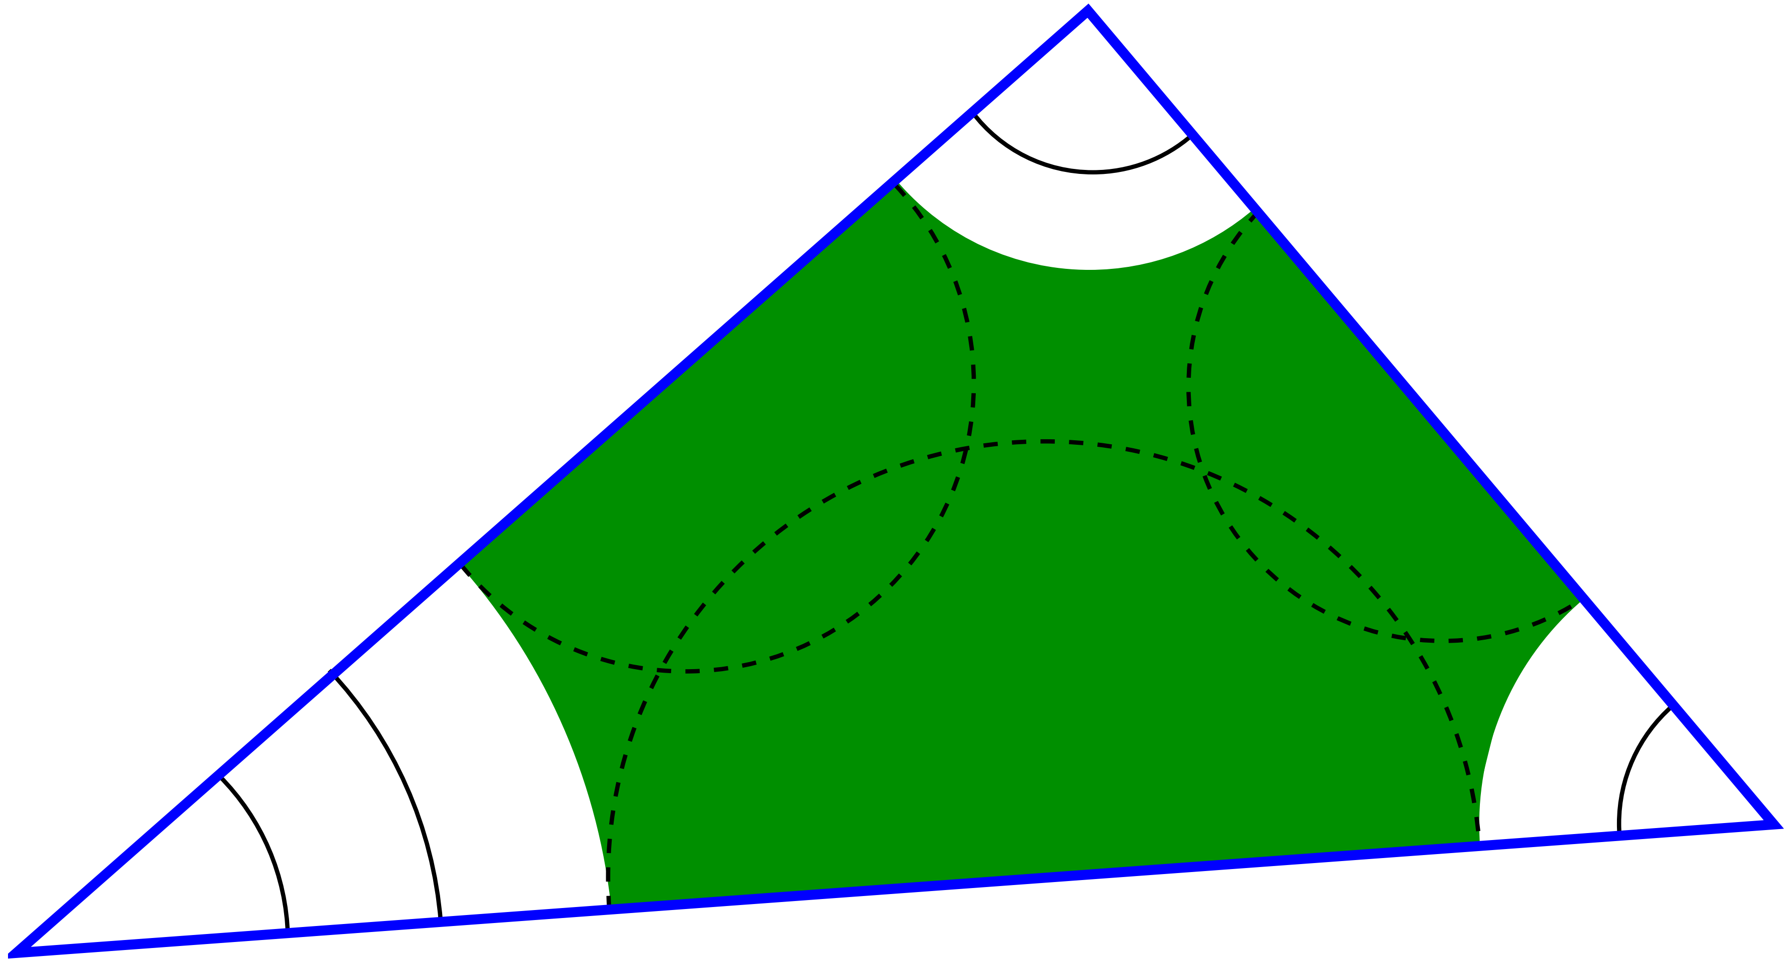

Later we will consider “thick” central regions.
Will need edges to be bases of half-disks contained inside triangle.

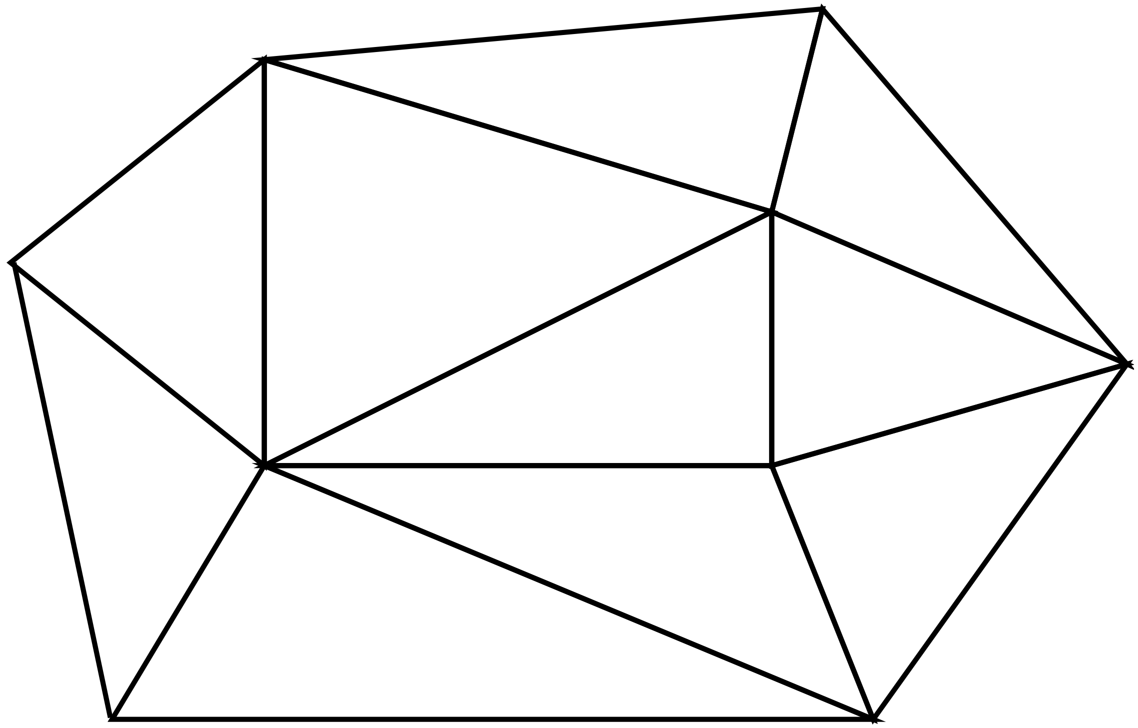

- Start with any triangulation.

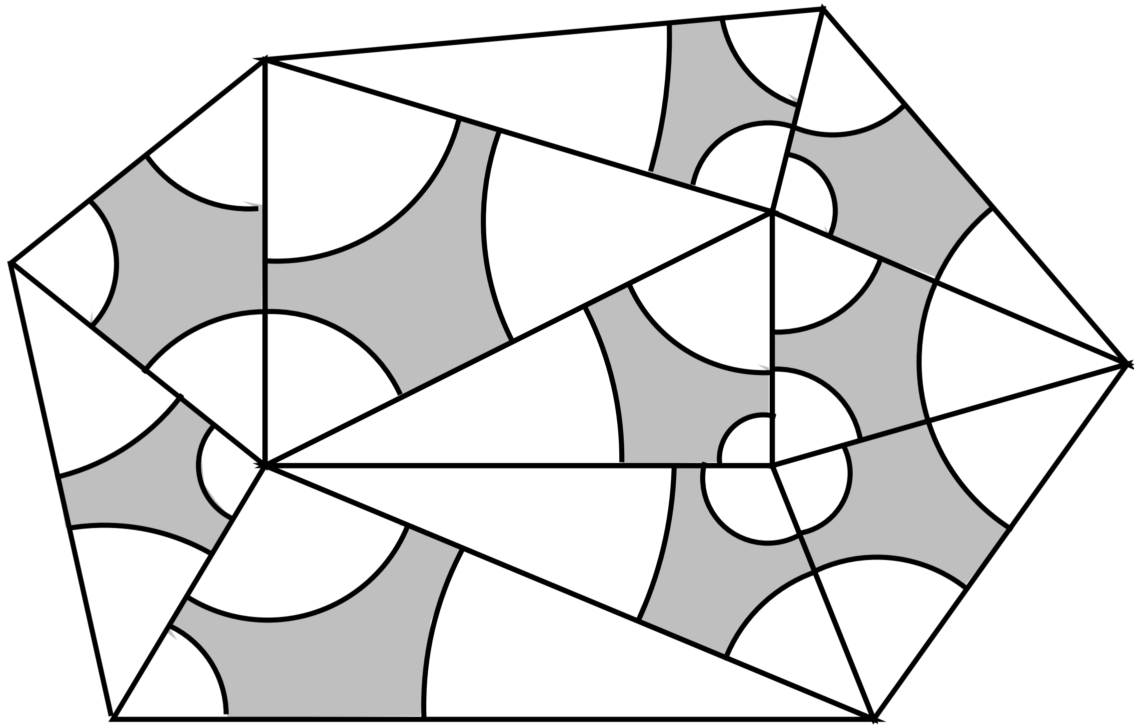

- Start with any triangulation.
- Make central parts.

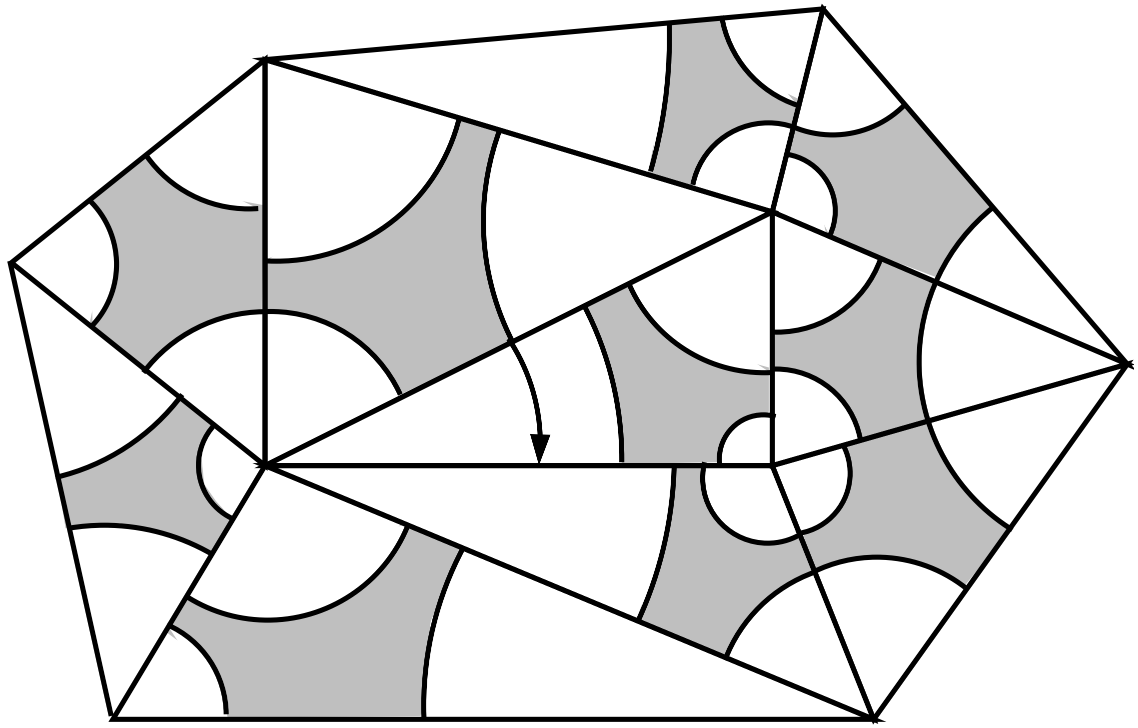

- Start with any triangulation.
- Make central parts.
- Propagate vertices until they leave thin parts.

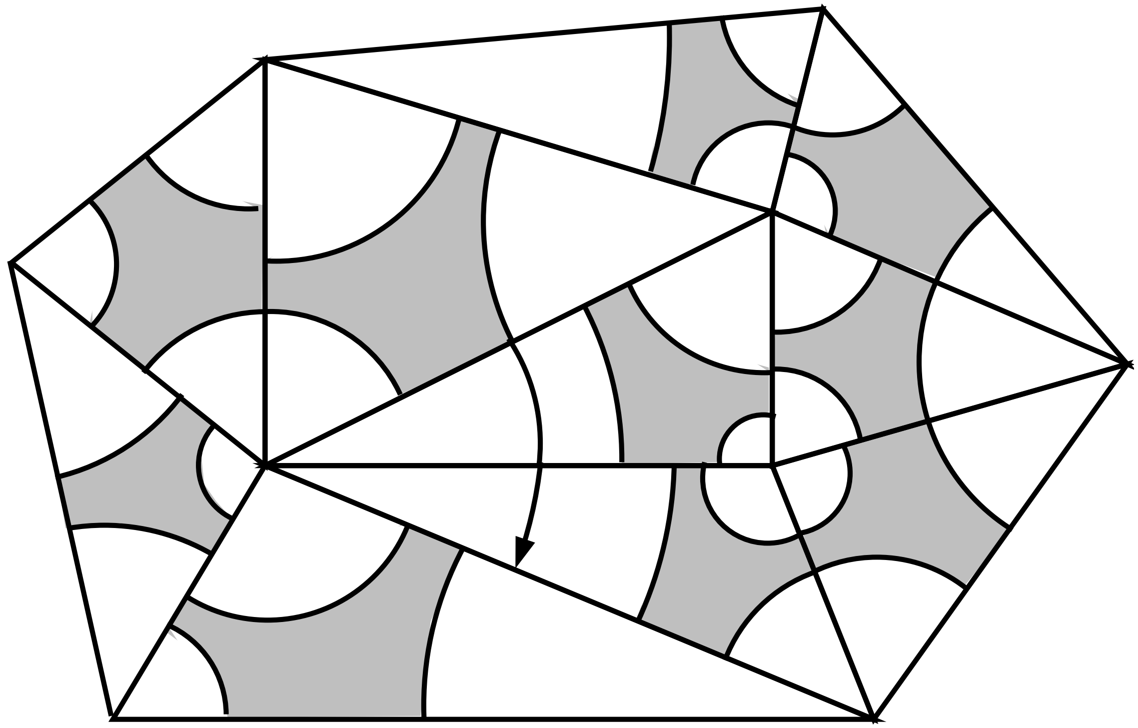

- Start with any triangulation.
- Make central parts.
- Propagate vertices until they leave thin parts.

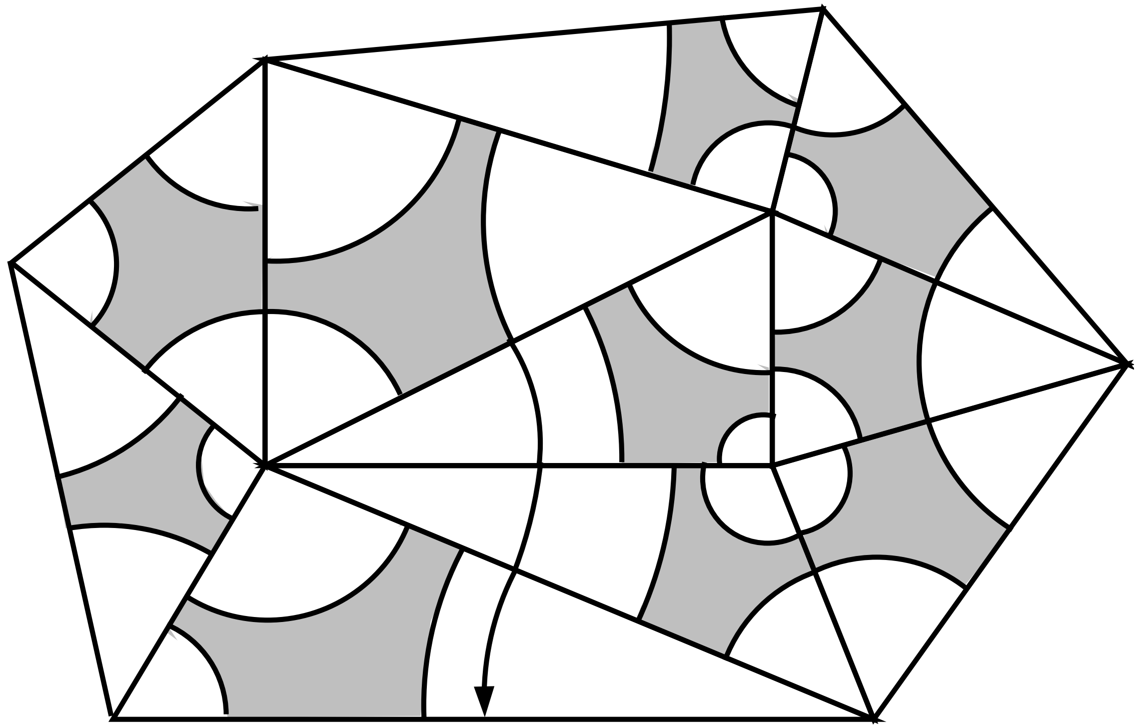

- Start with any triangulation.
- Make central parts.
- Propagate vertices until they leave thin parts.

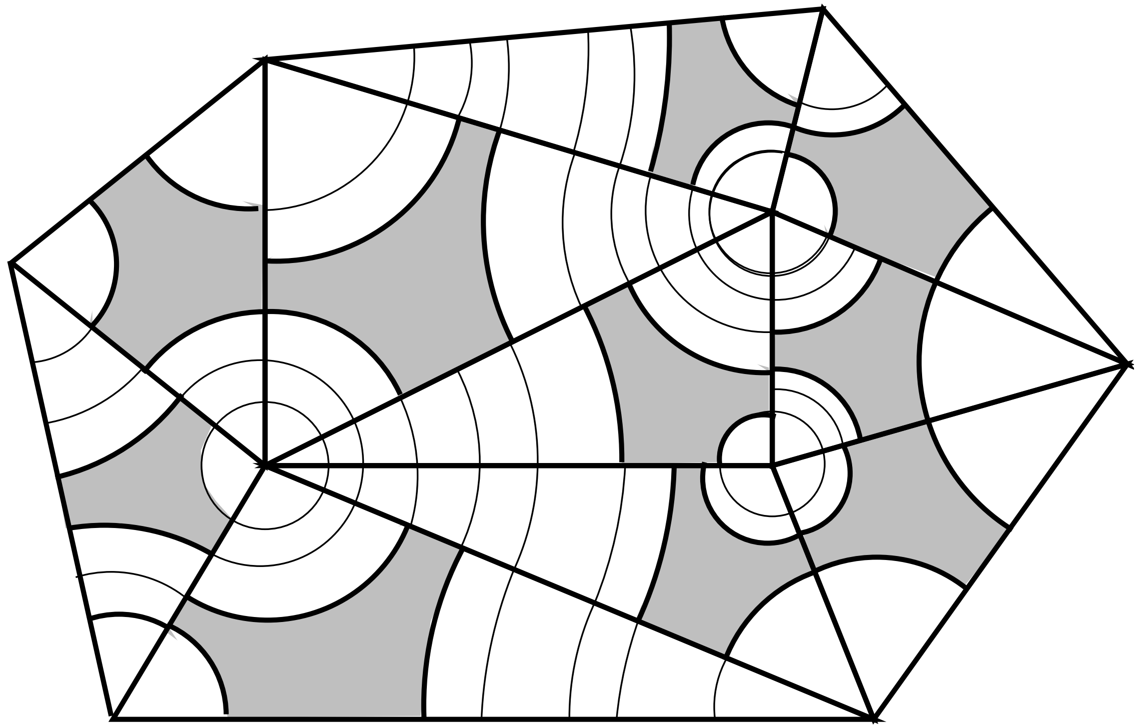

- Start with any triangulation.
- Make central parts.
- Propagate vertices until they leave thin parts.

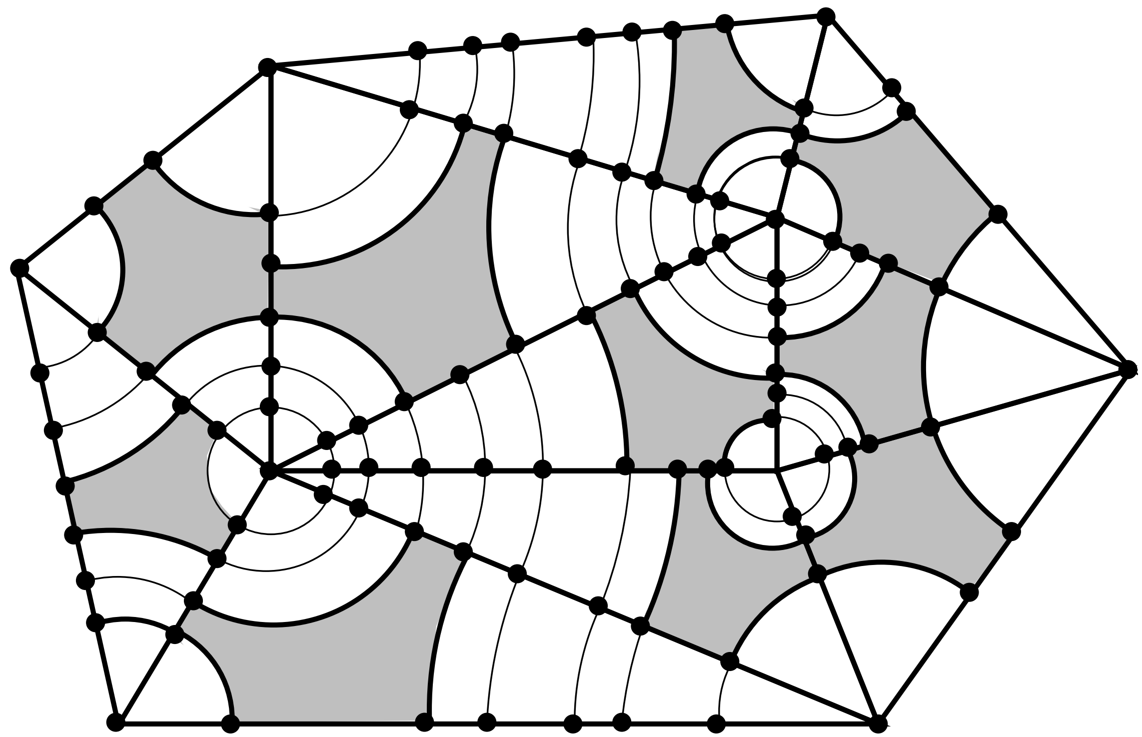

- Start with any triangulation.
- Make central parts.
- Propagate vertices until they leave thin parts.
- How many new points are created?

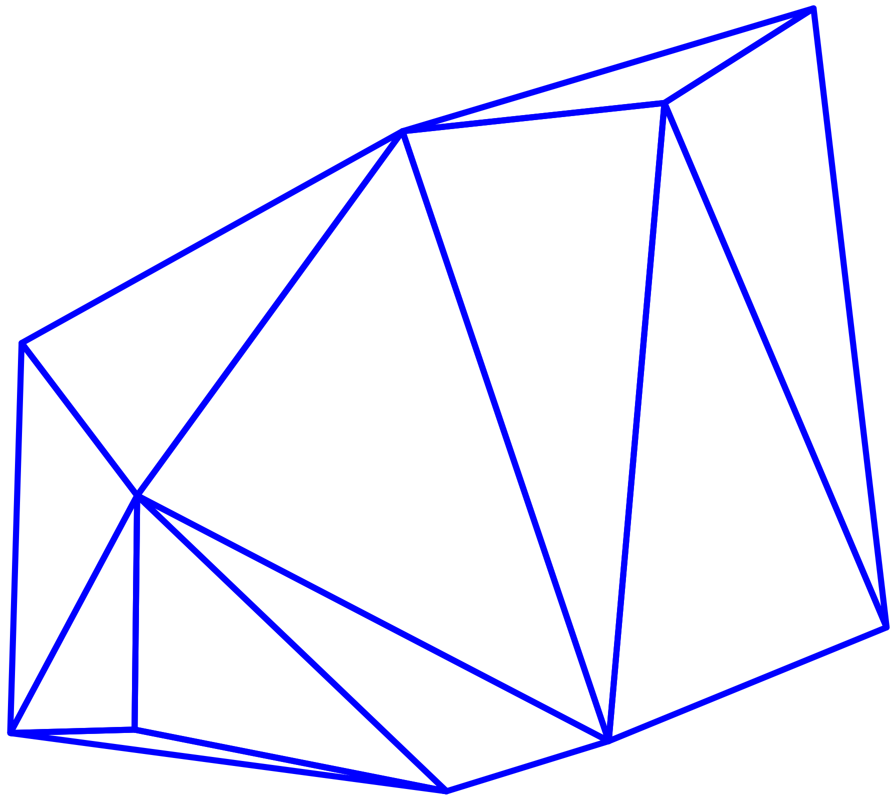

Delaunay triangulation of 10 random points,

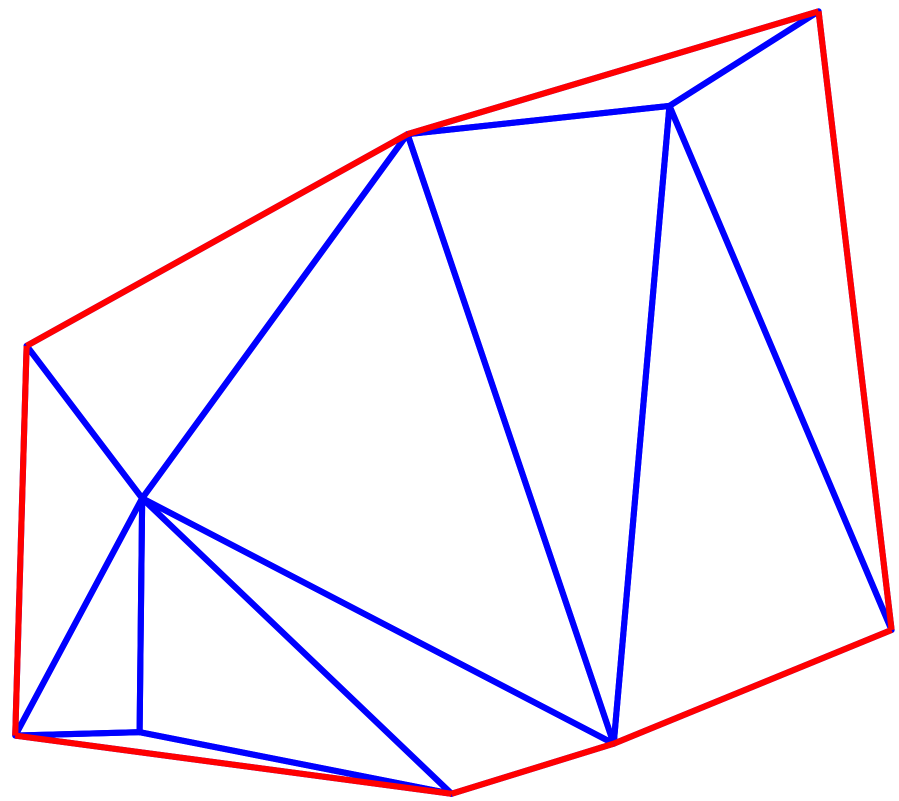

The boundary of the triangulation

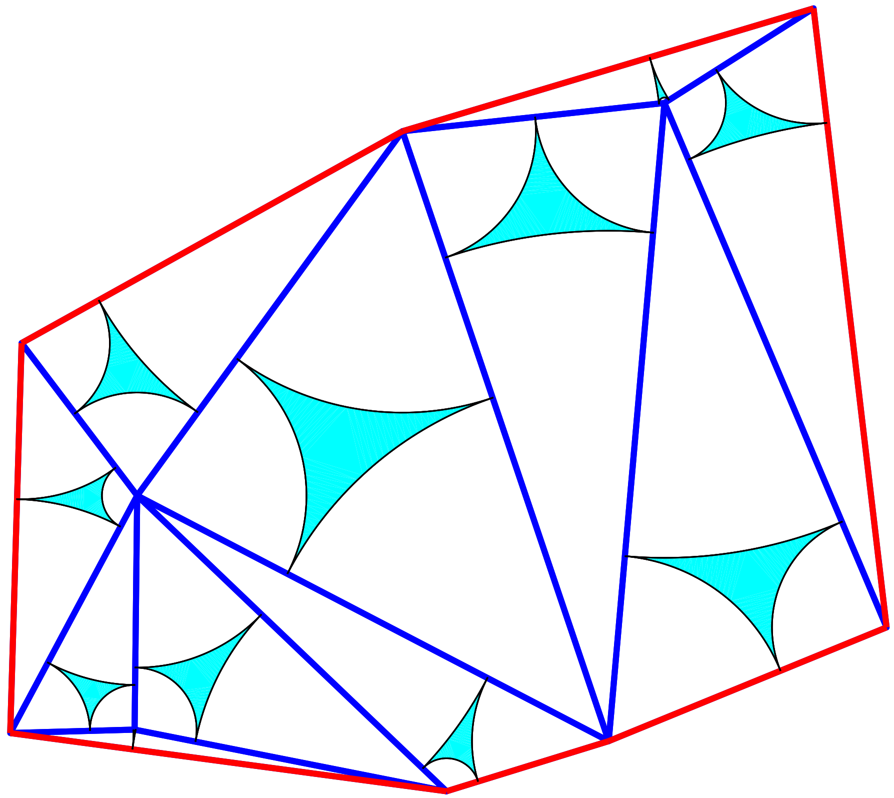

The central regions.

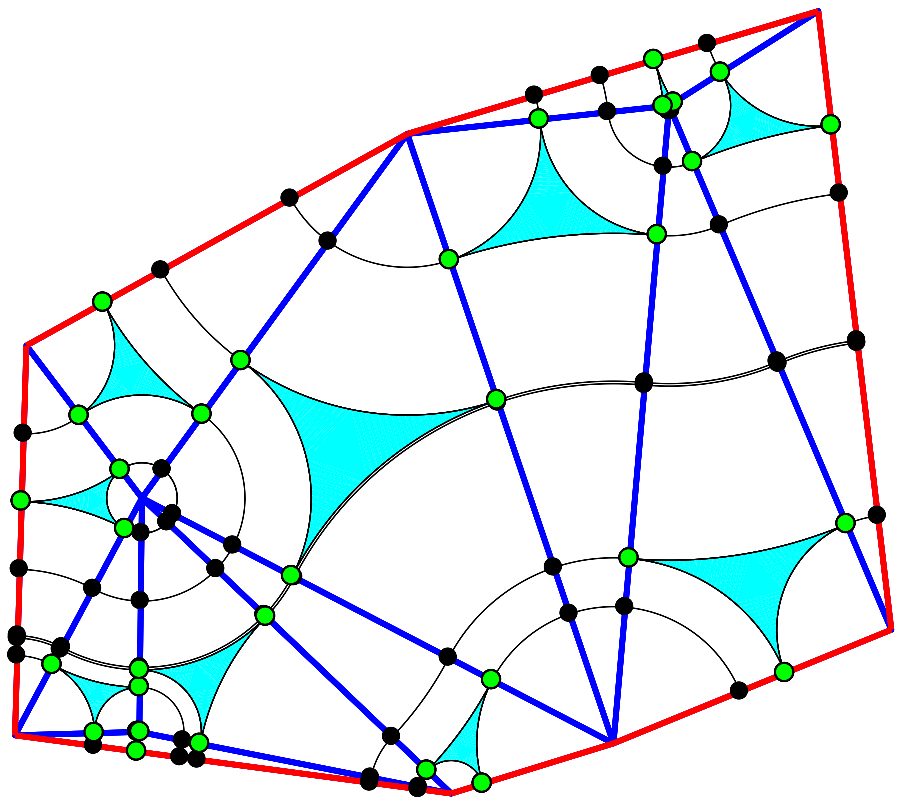

Propagation lines starting at all cusp points.

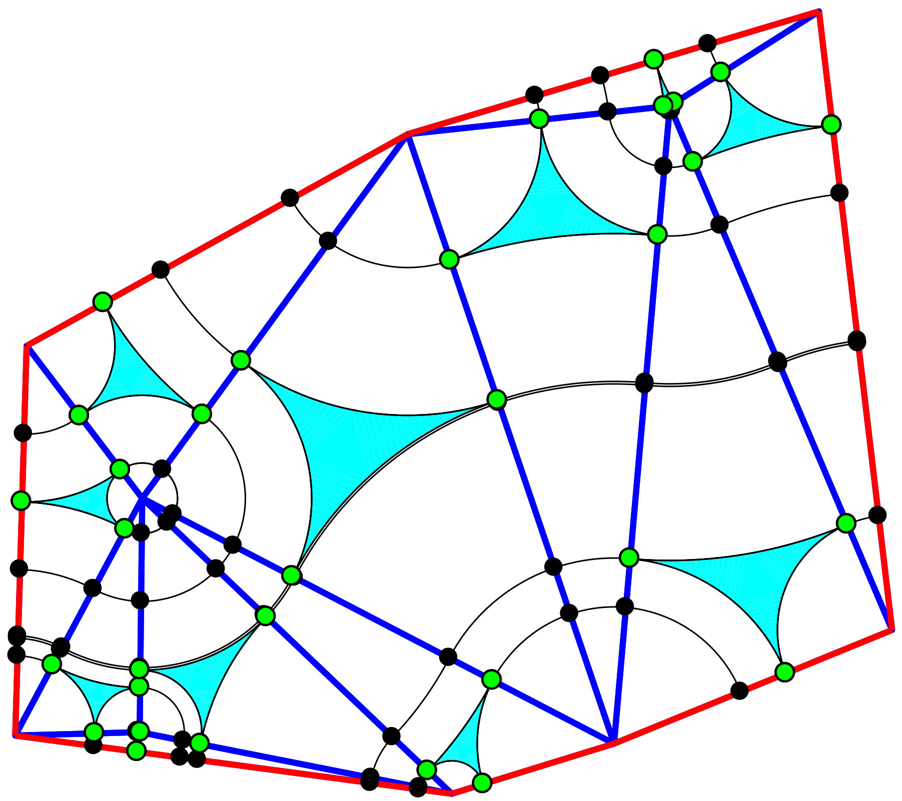

Propagation lines identify boundary points; induces tree.
Discontinuous, but piecewise length preserving.

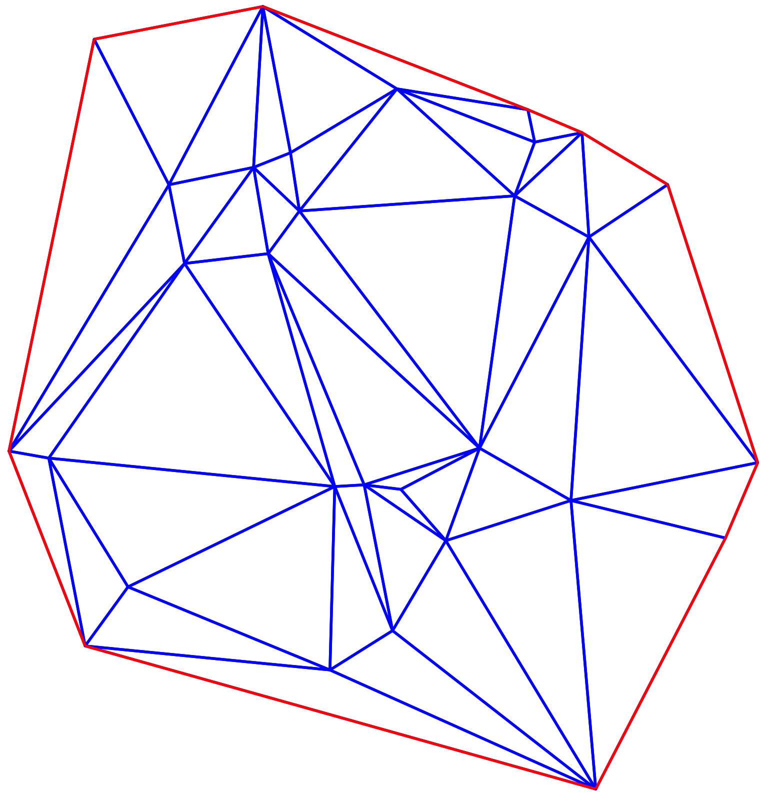

triangulation of 30 points.

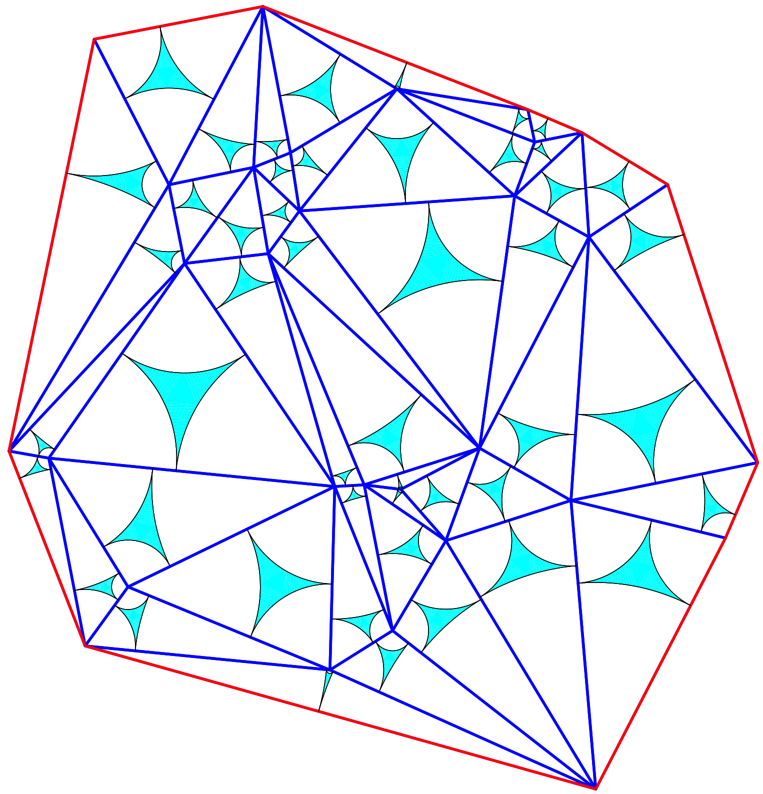

The central regions.

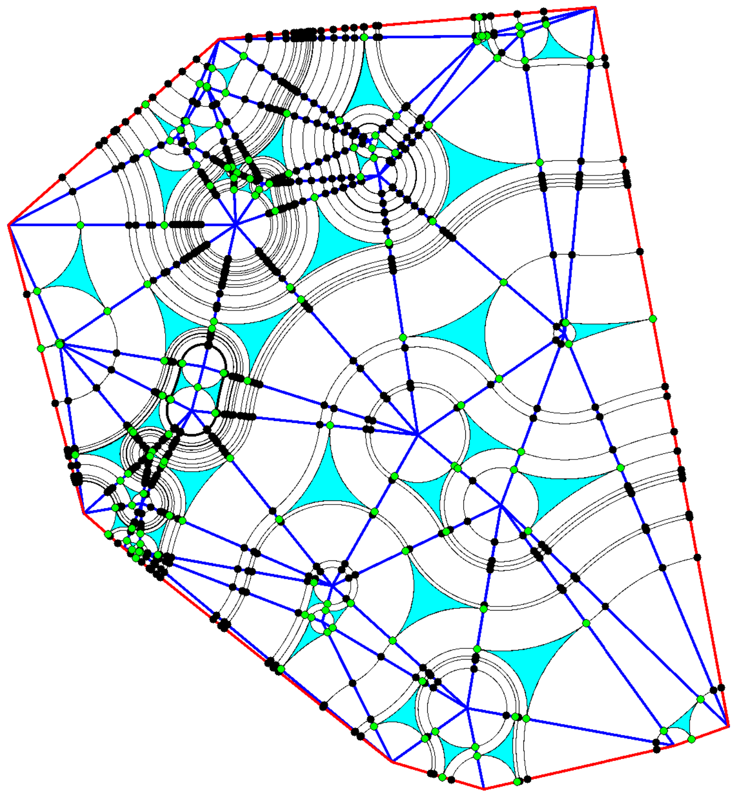

Propagation lines starting at all cusp points.

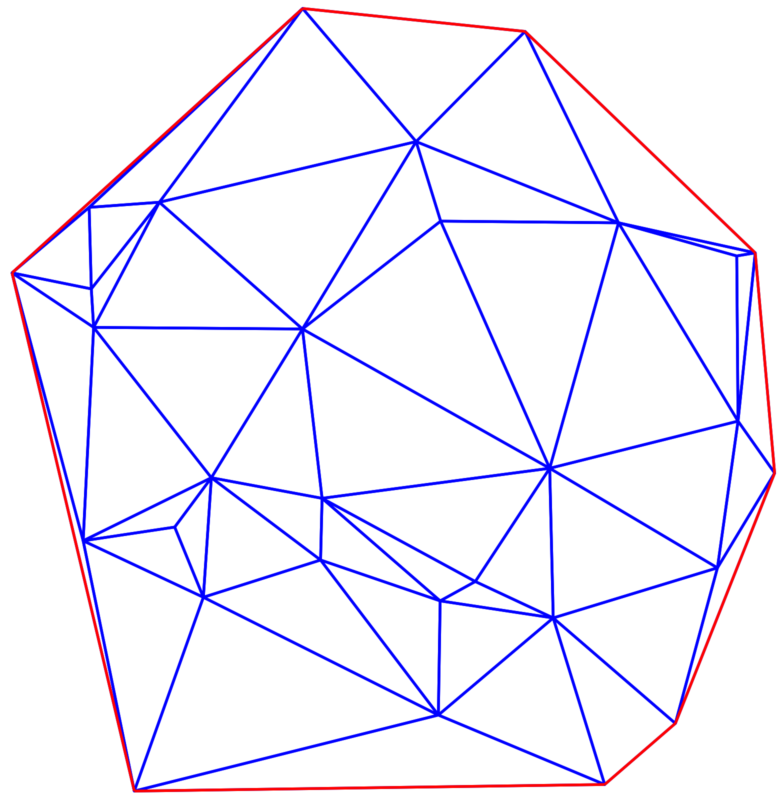

Delaunay triangulation of 30 random points in disk.

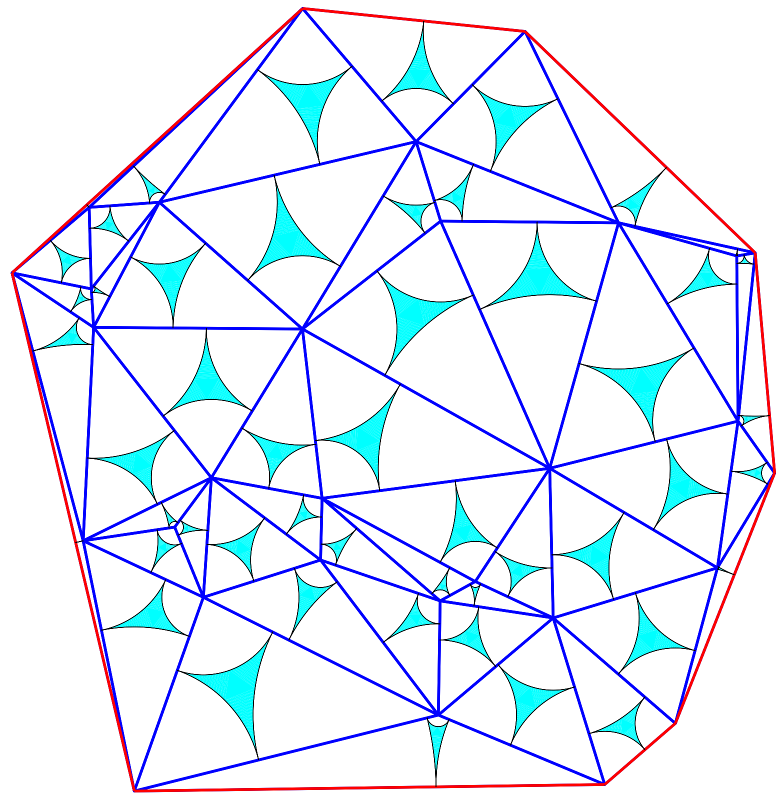

The central regions.

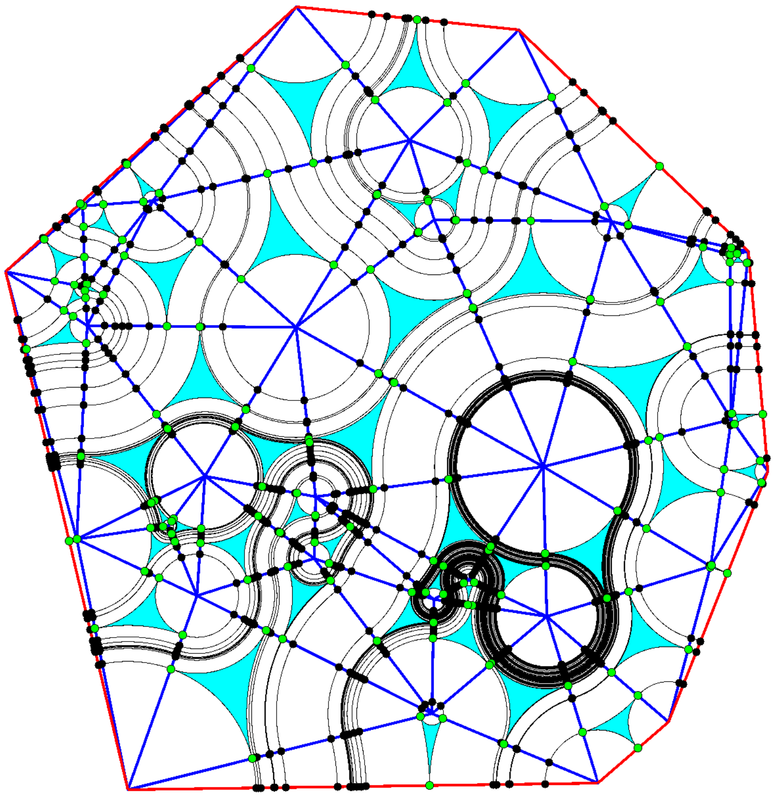

Propagation lines starting at all cusp points.

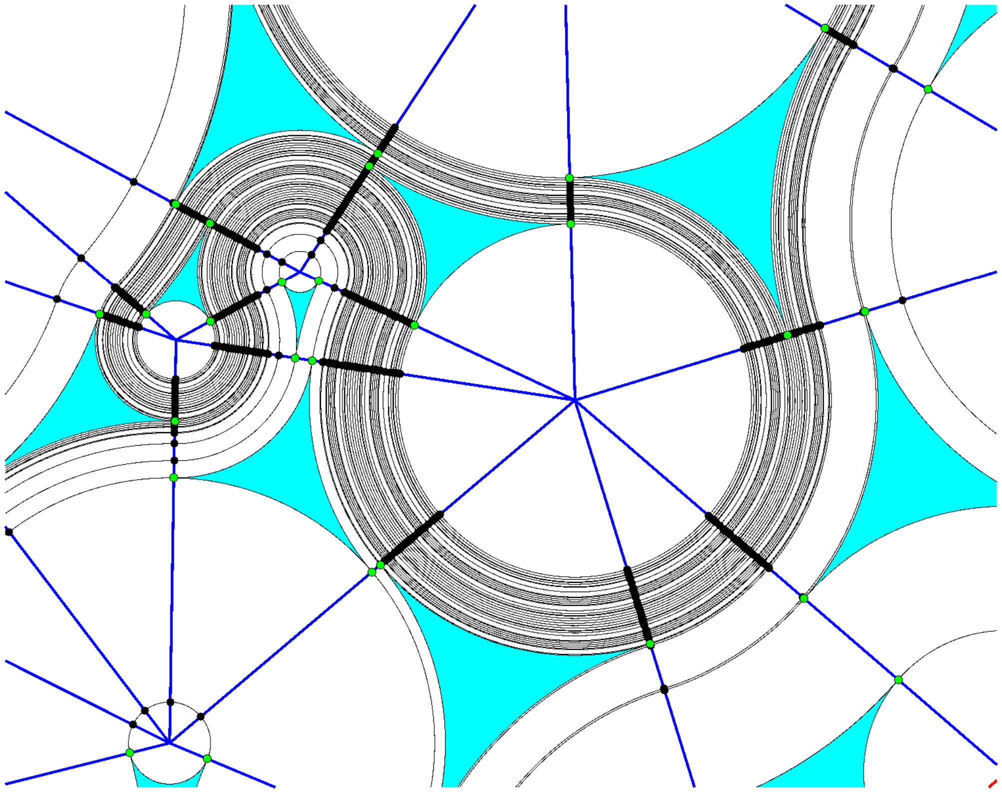

Enlargement 1.

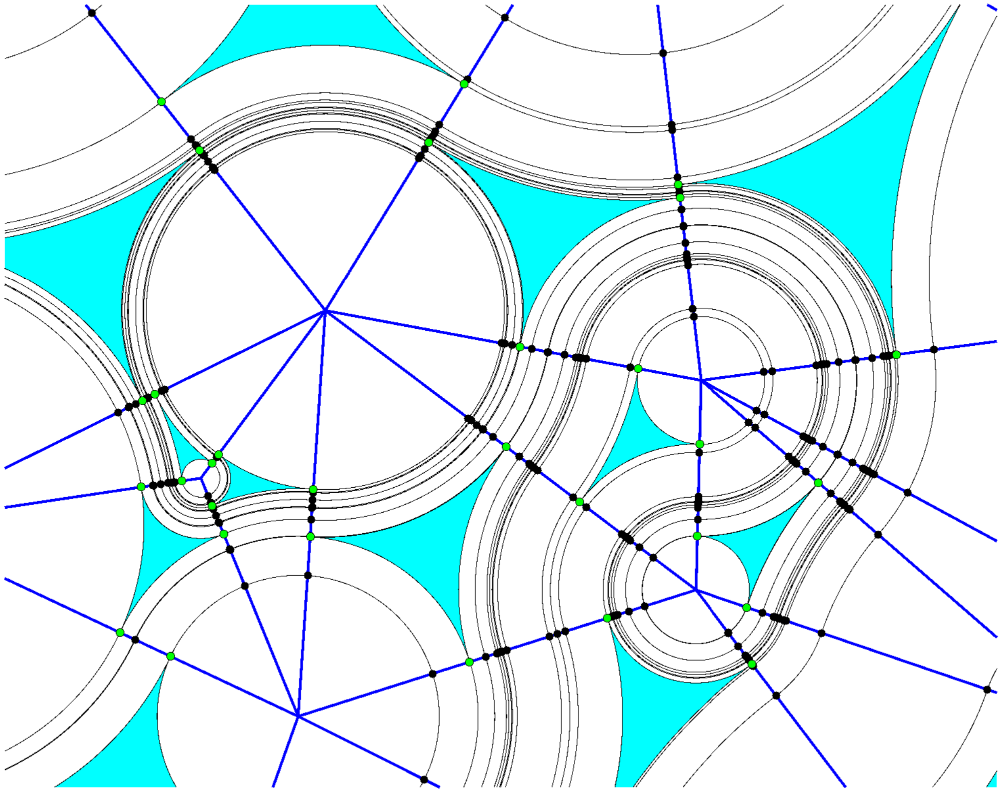

Enlargement 2.

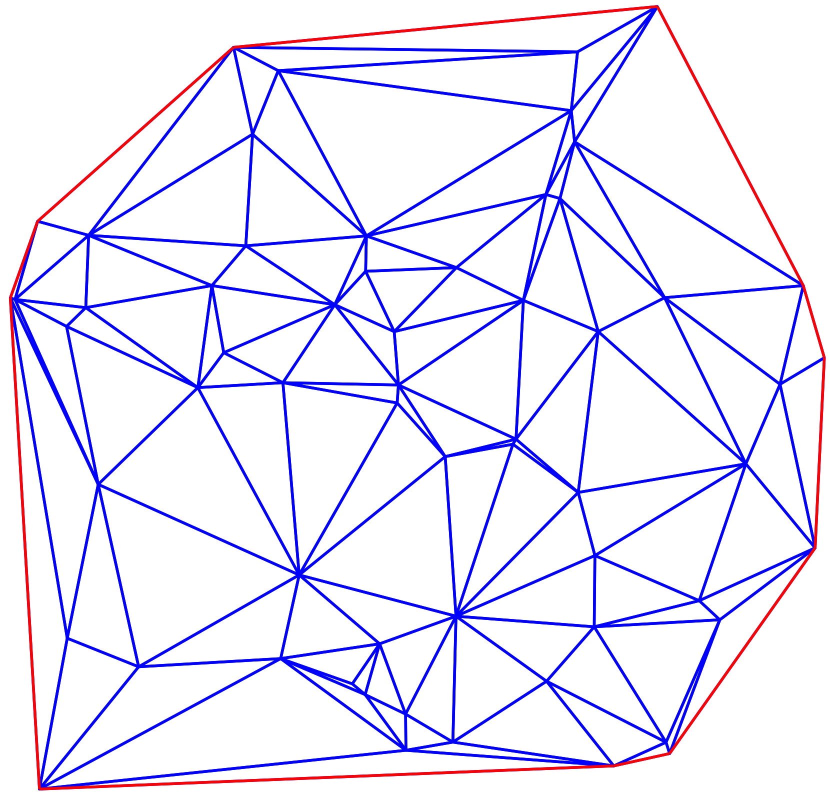

60 points

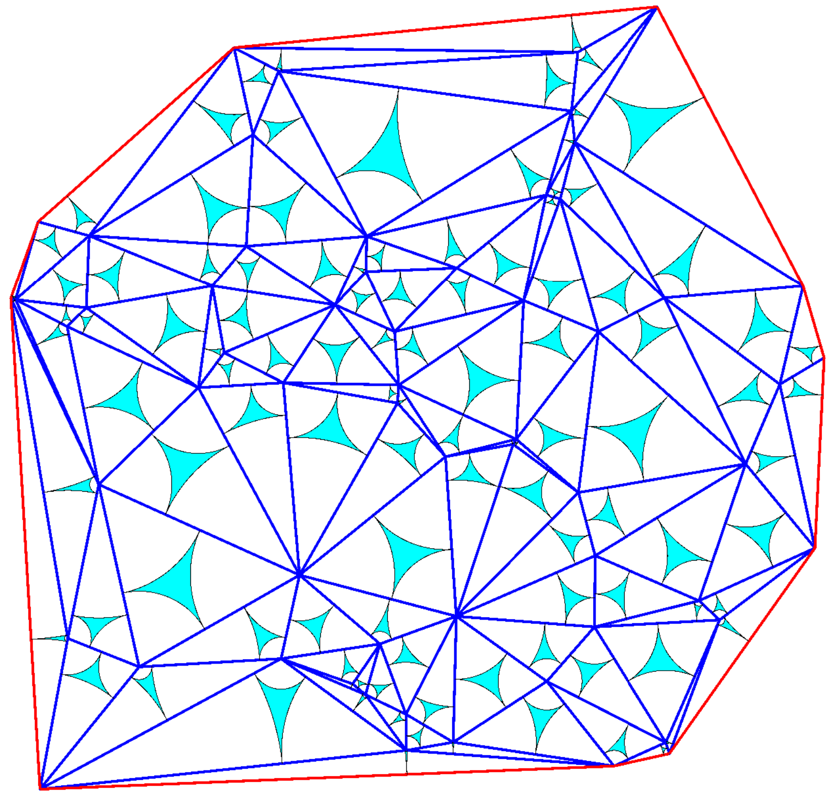

The central regions.

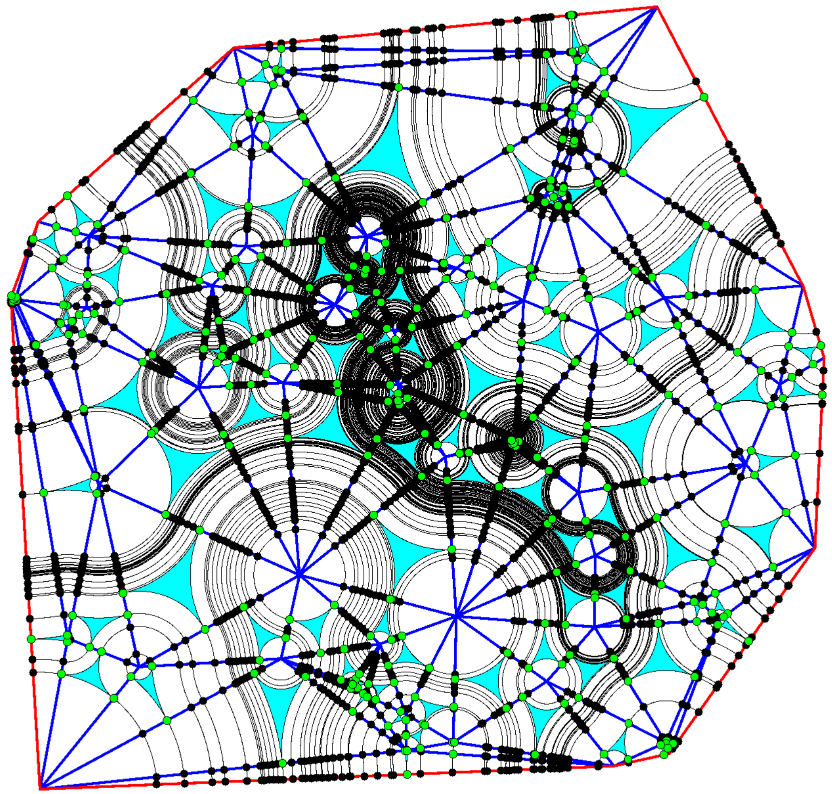

Propagation lines starting at all cusp points.

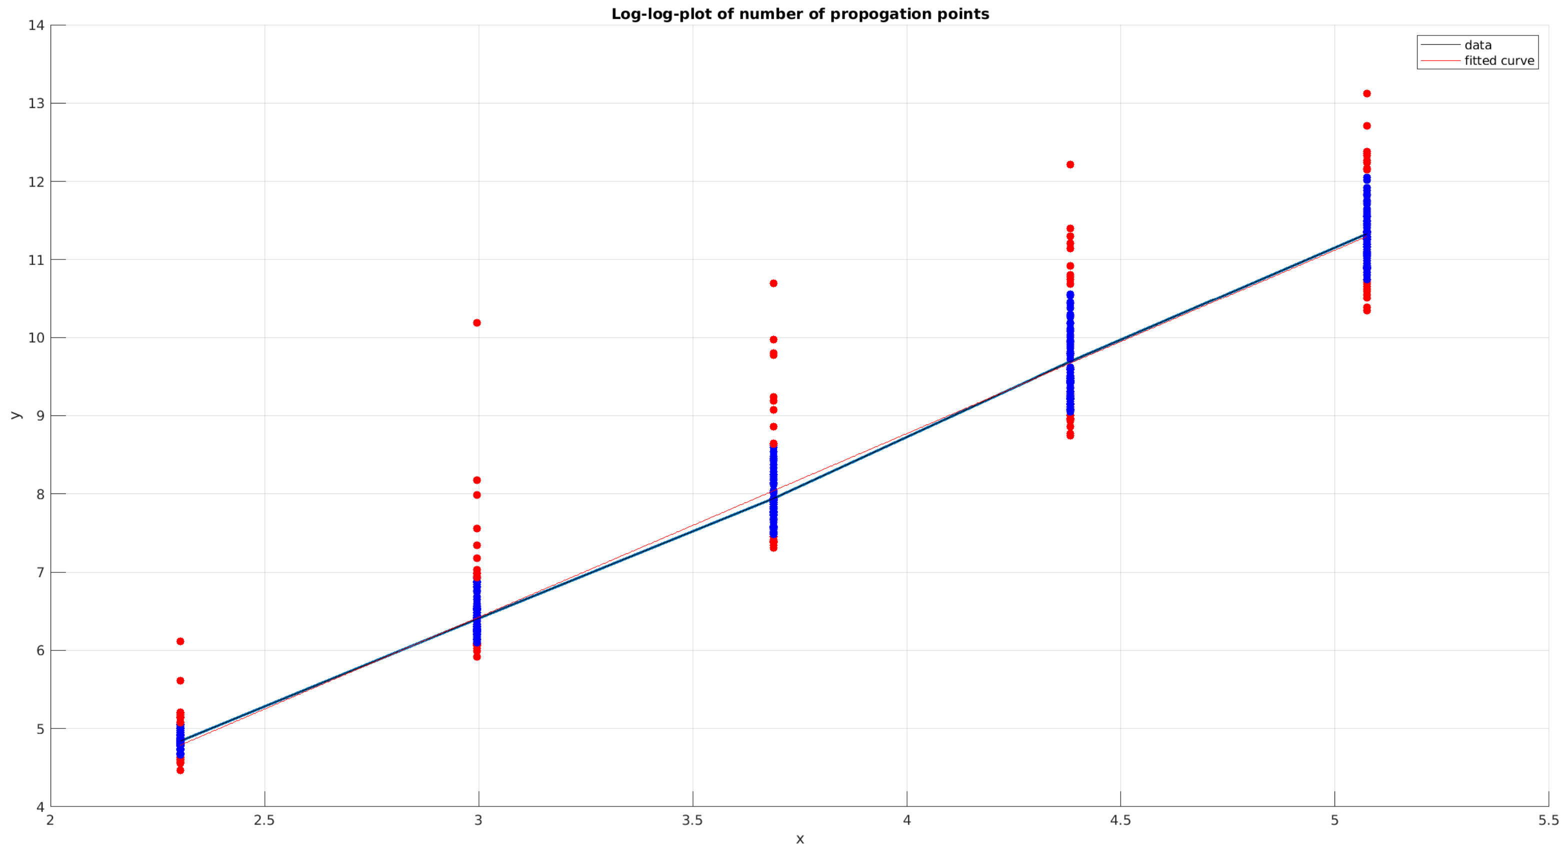

Log-log plot of number points created versus n.

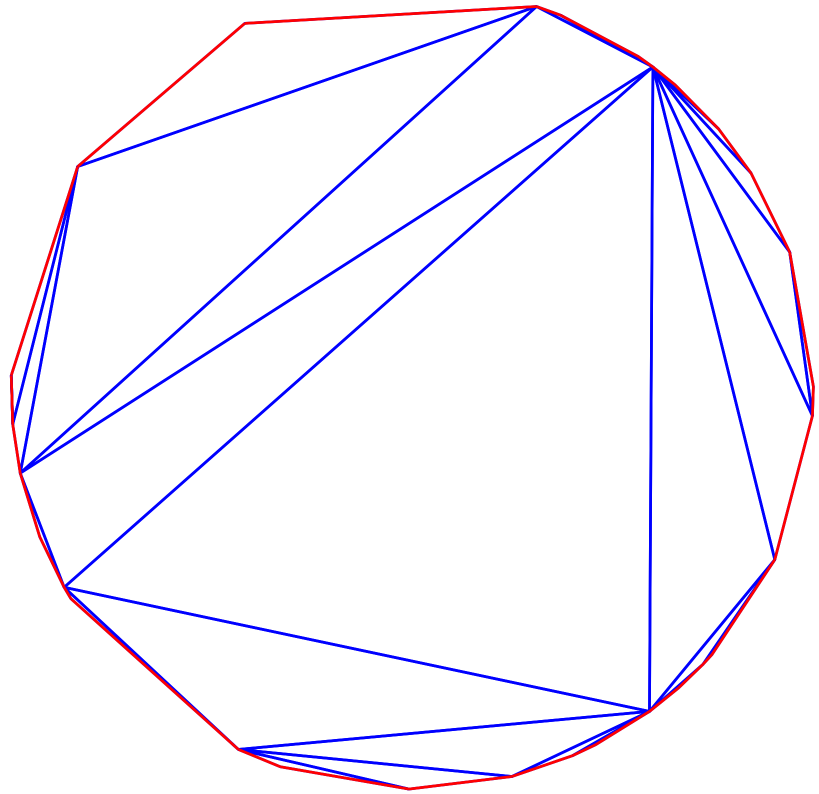

Delaunay triangulation of 30 random points on circle.

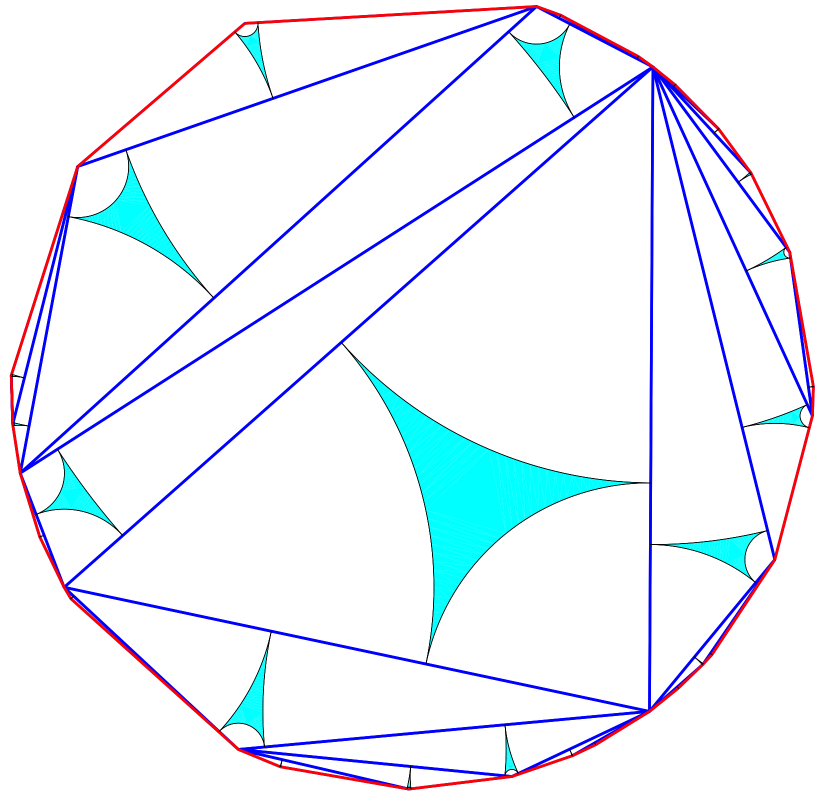

The central regions.

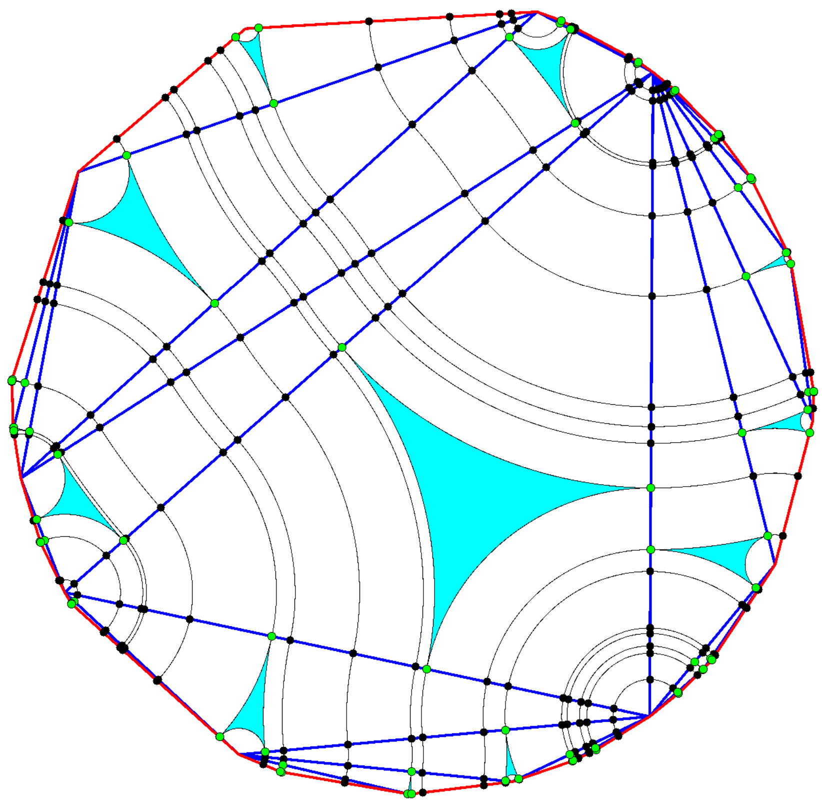

Propagation lines starting at all cusp points.

Theorem: For a triangulation of a simple polygon by diagonal, at most $O(n^2)$ points are created.

Proof: In this case, the triangles form the vertices of a tree where adjacency means sharing an edge.

Since a flow line never re-enters a triangle, it visits at most n triangles, so at most $O(n^2)$ points are generated.

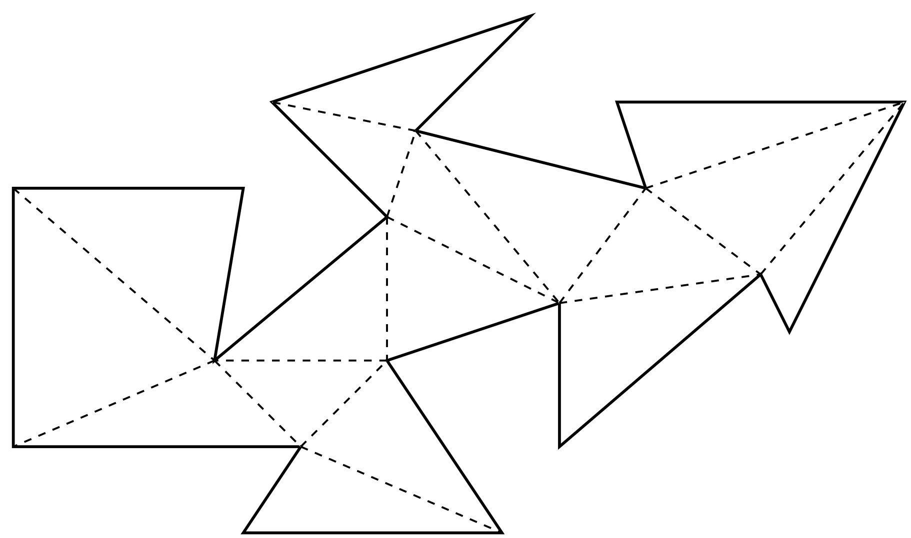

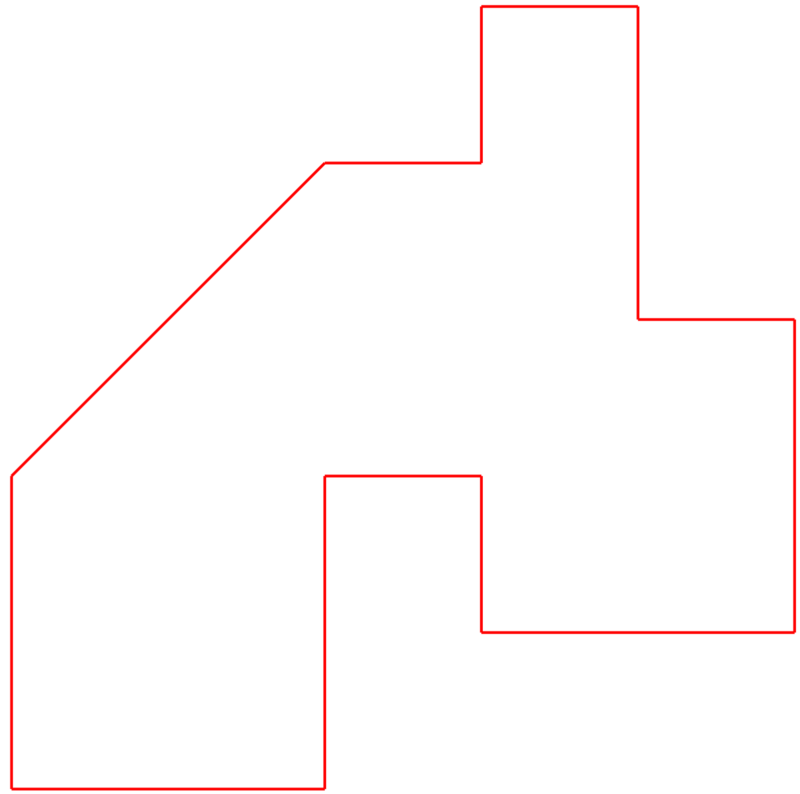

A simple polygon

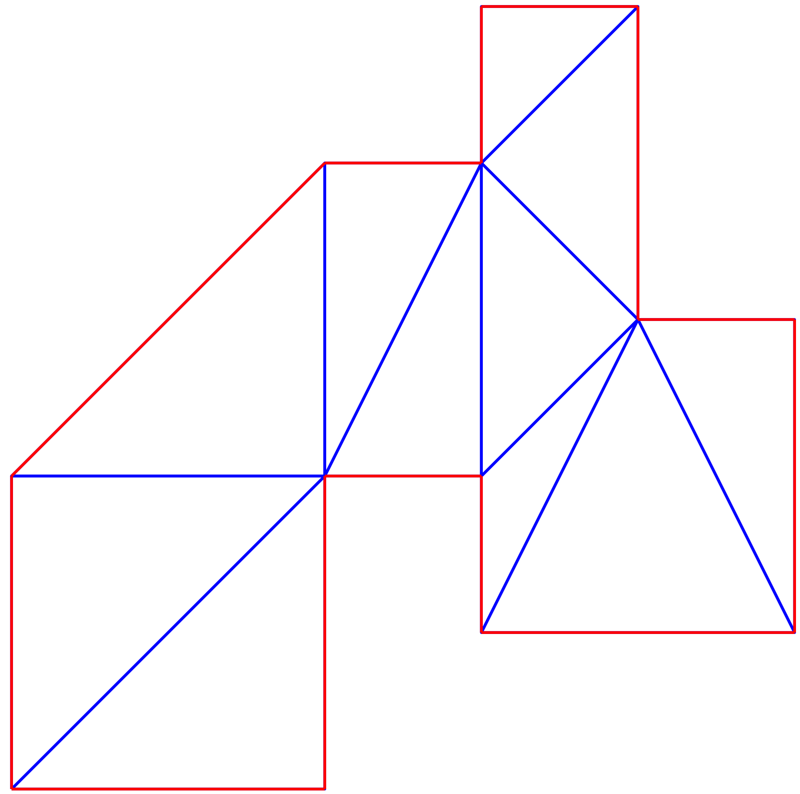

A triangulation of the polygon using diagonals.

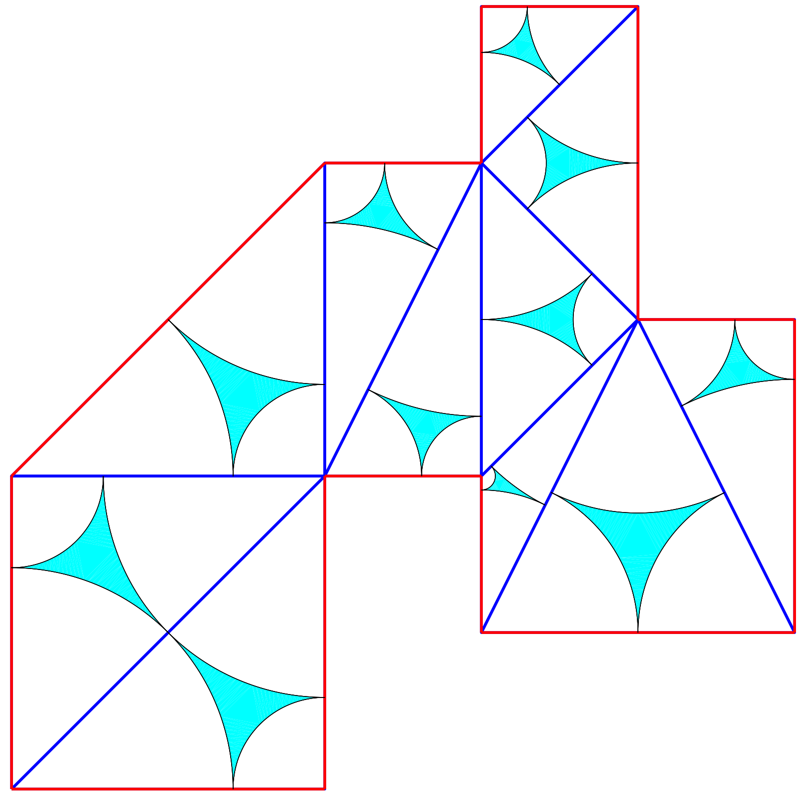

The central cusp regions.

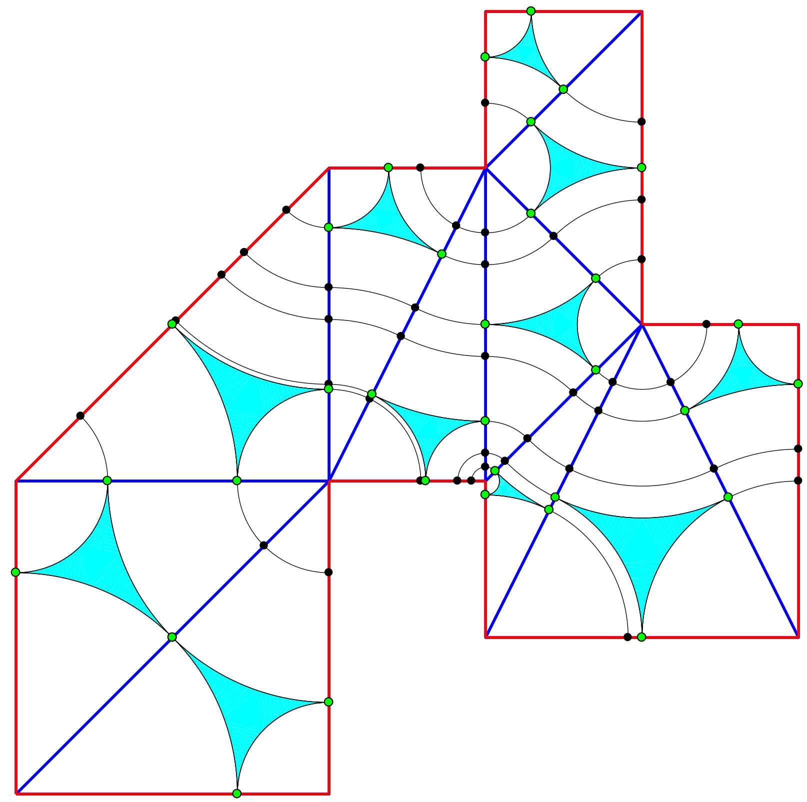

No flow line returns to a triangle, so each path creates at most n new points, for a total of $O(n^2)$.

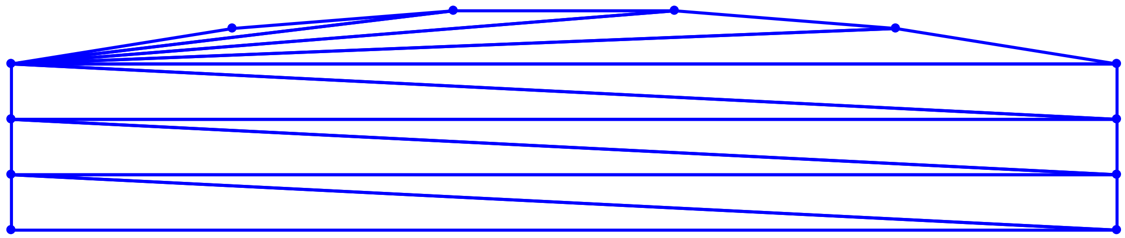

The n^2 is sharp.

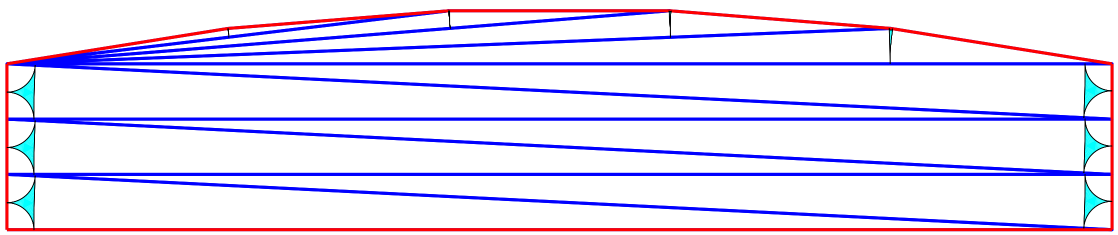

The n^2 is sharp.

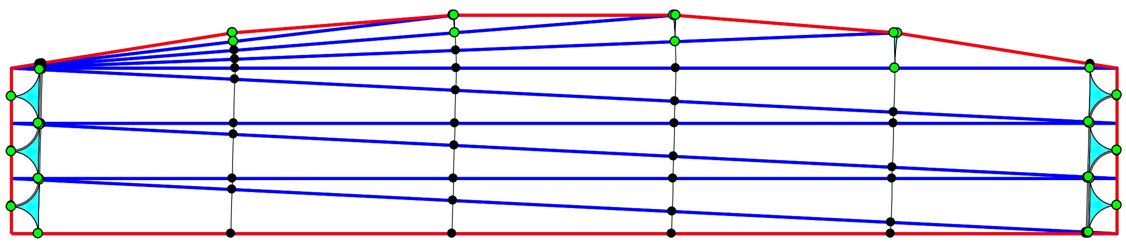

Each point "along top" creates a path that hits $\simeq n$ triangles.

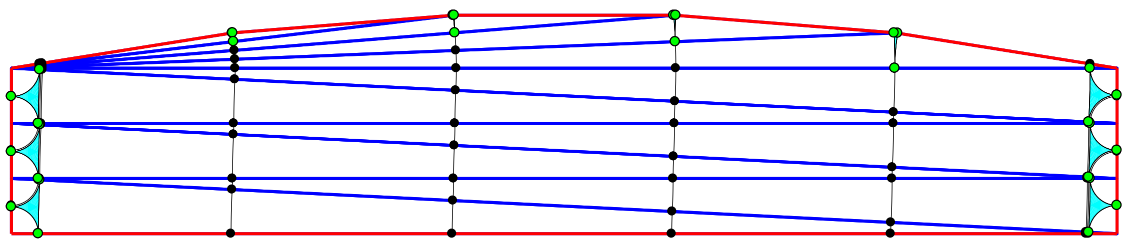

The n^2 is stable under small perturbations of vertices.

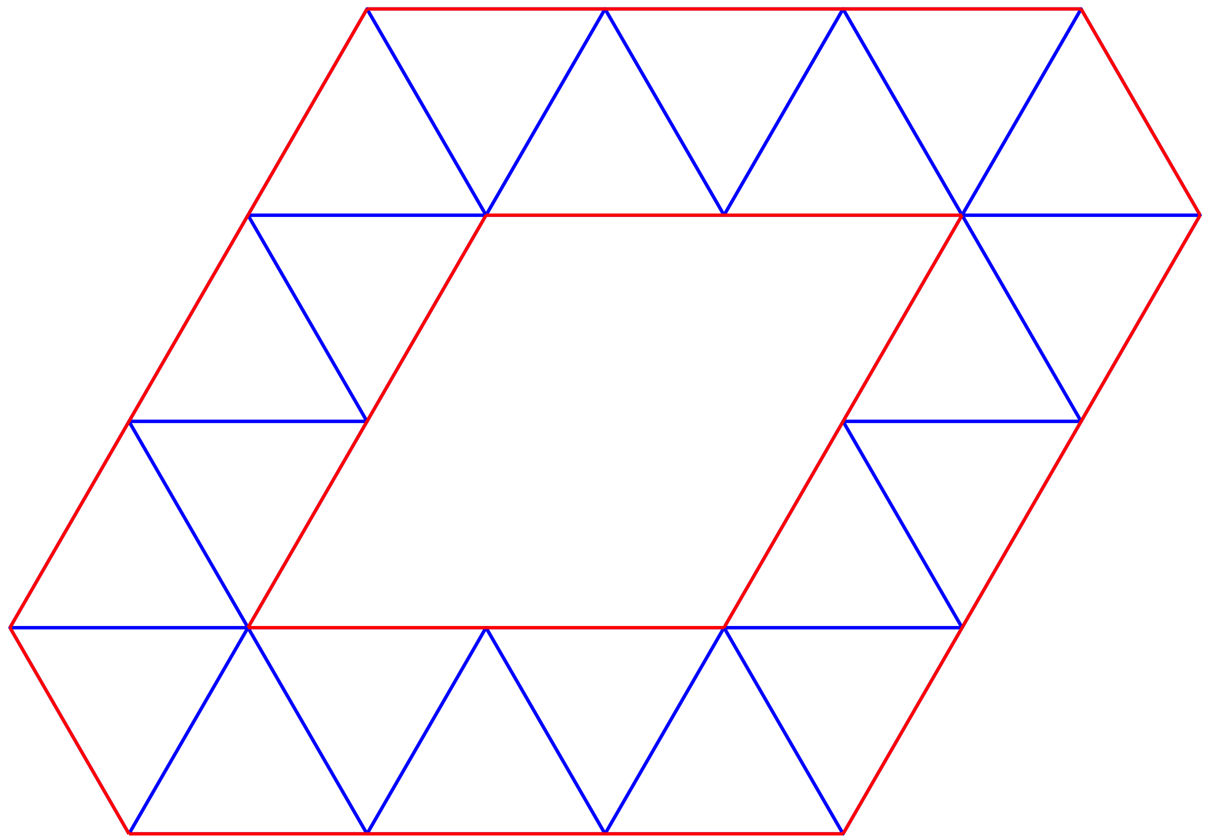

A ring of equilateral triangles.

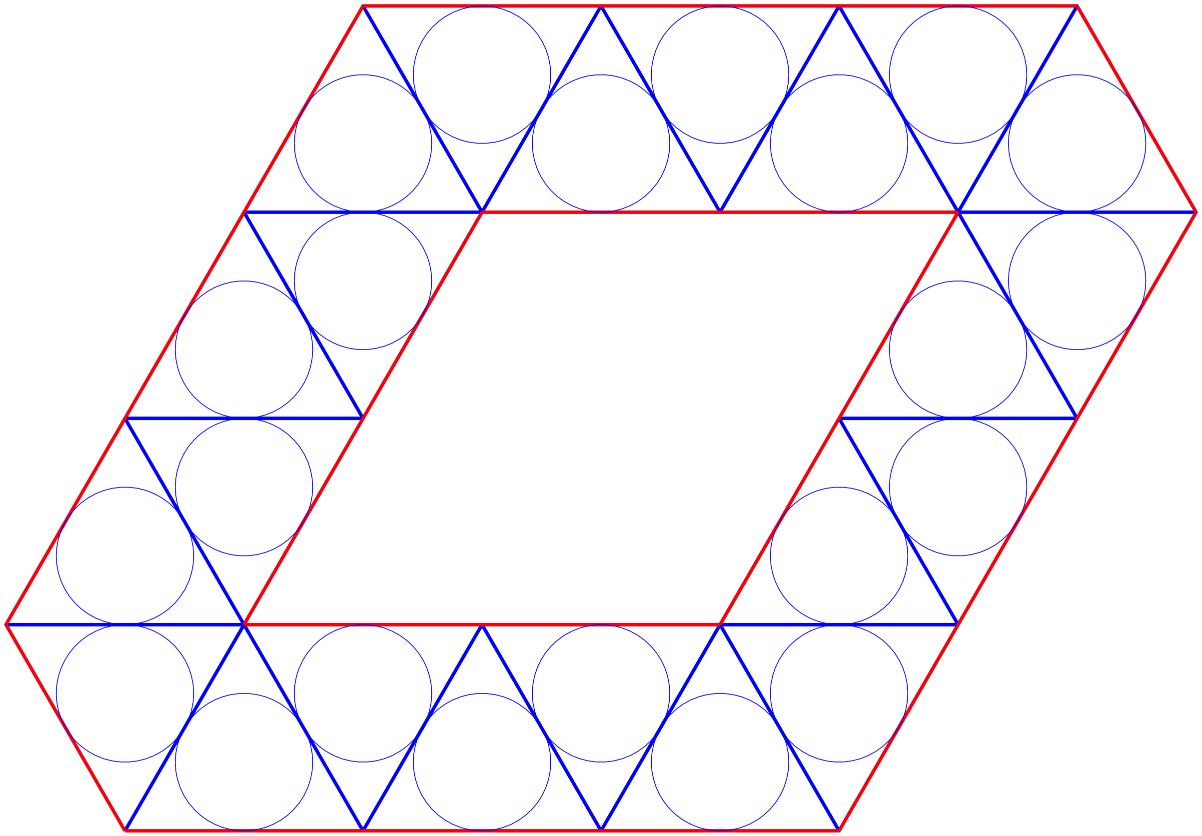

The in-circles are tangent.

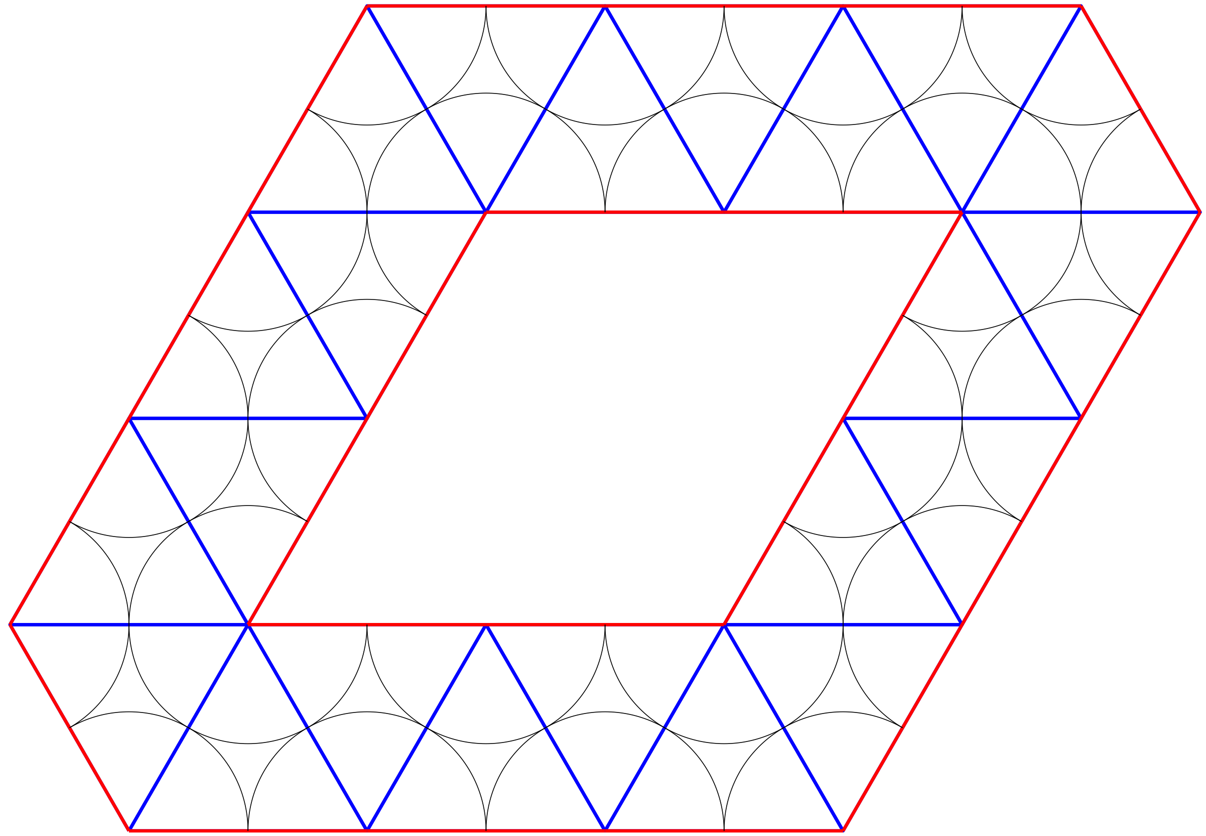

The the cusps touch; for a closed flow line.

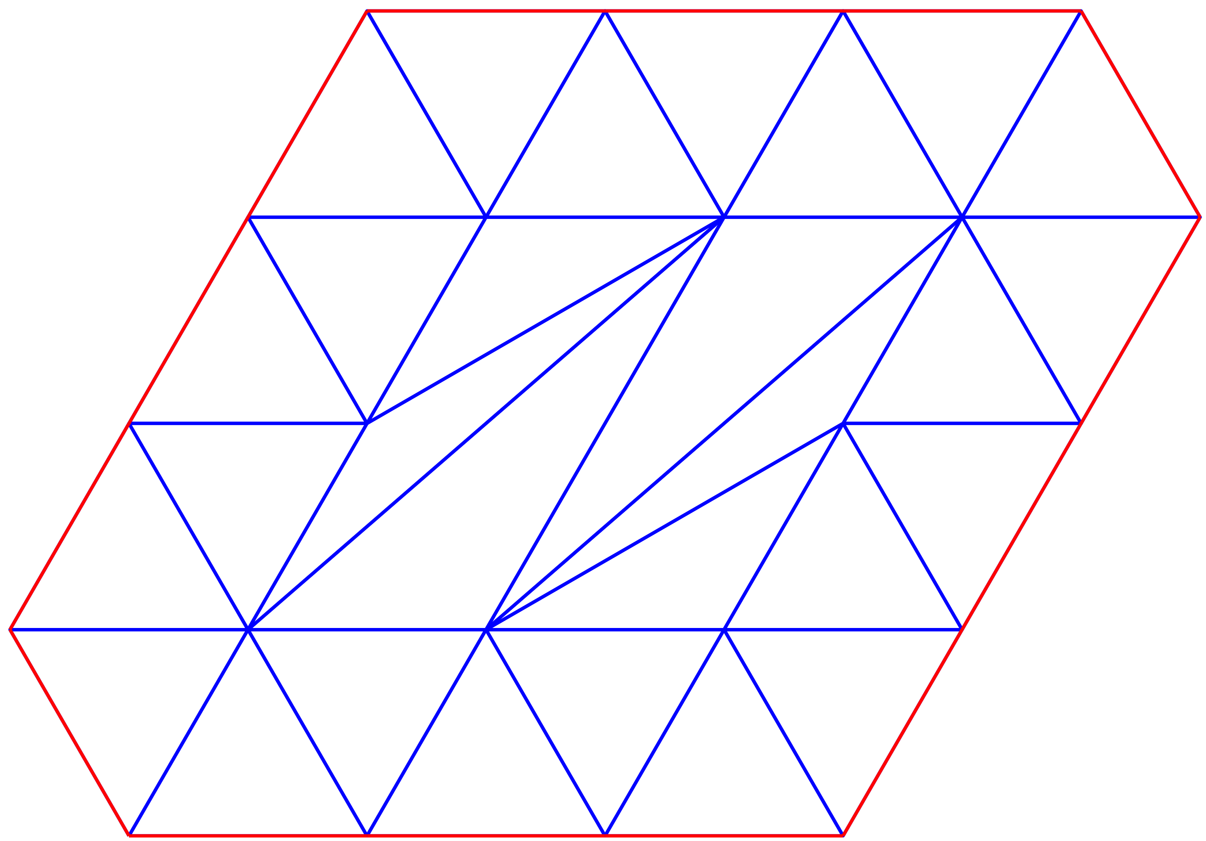

No matter how we triangulate interior, flow lines never exit.

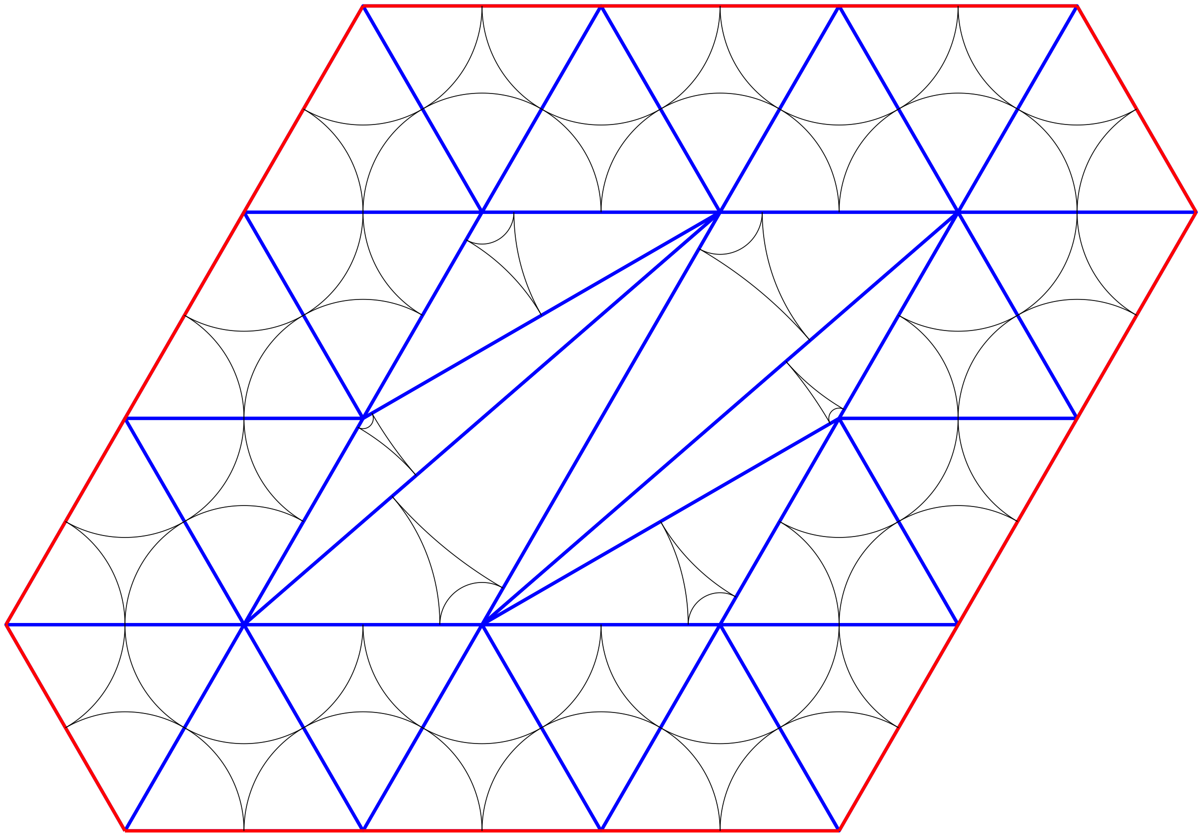

No matter how we triangulate interior, flow lines never exit.

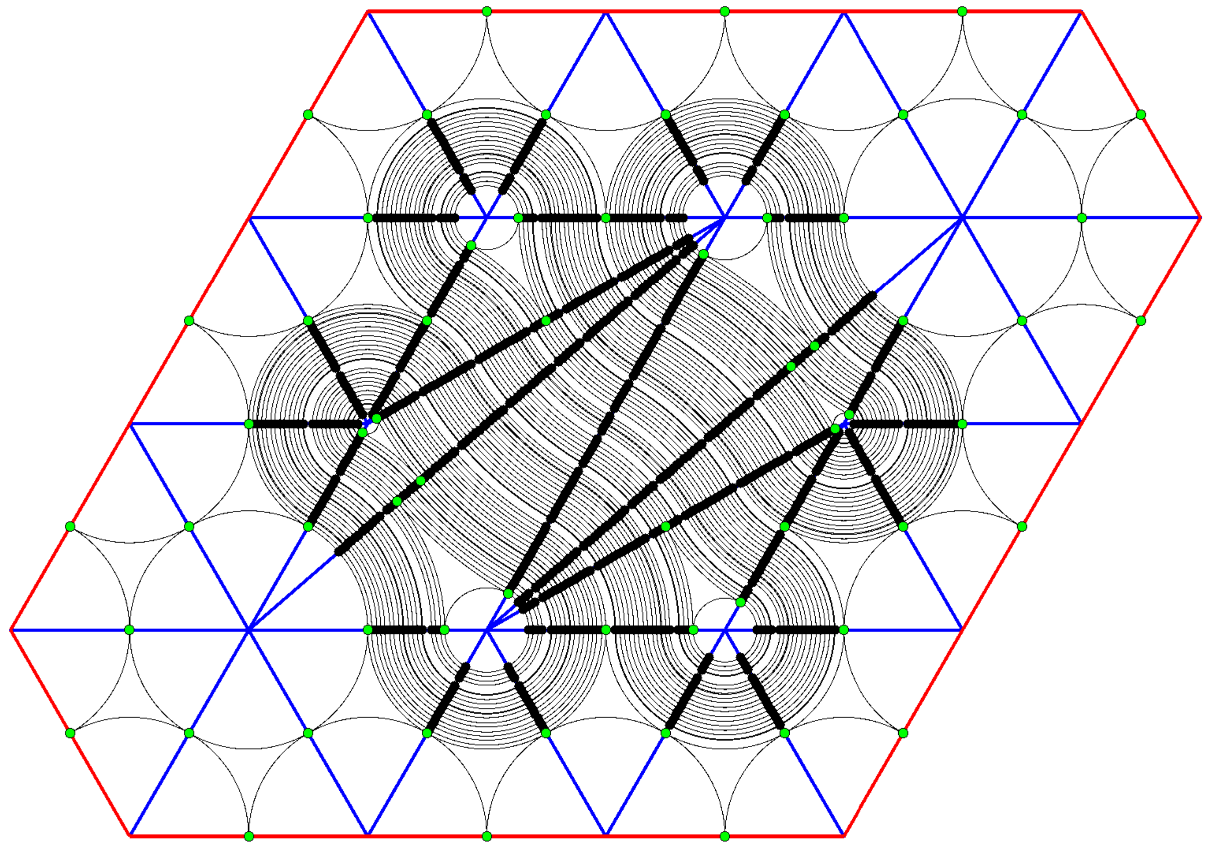

No matter how we triangulate interior, flow lines never exit.

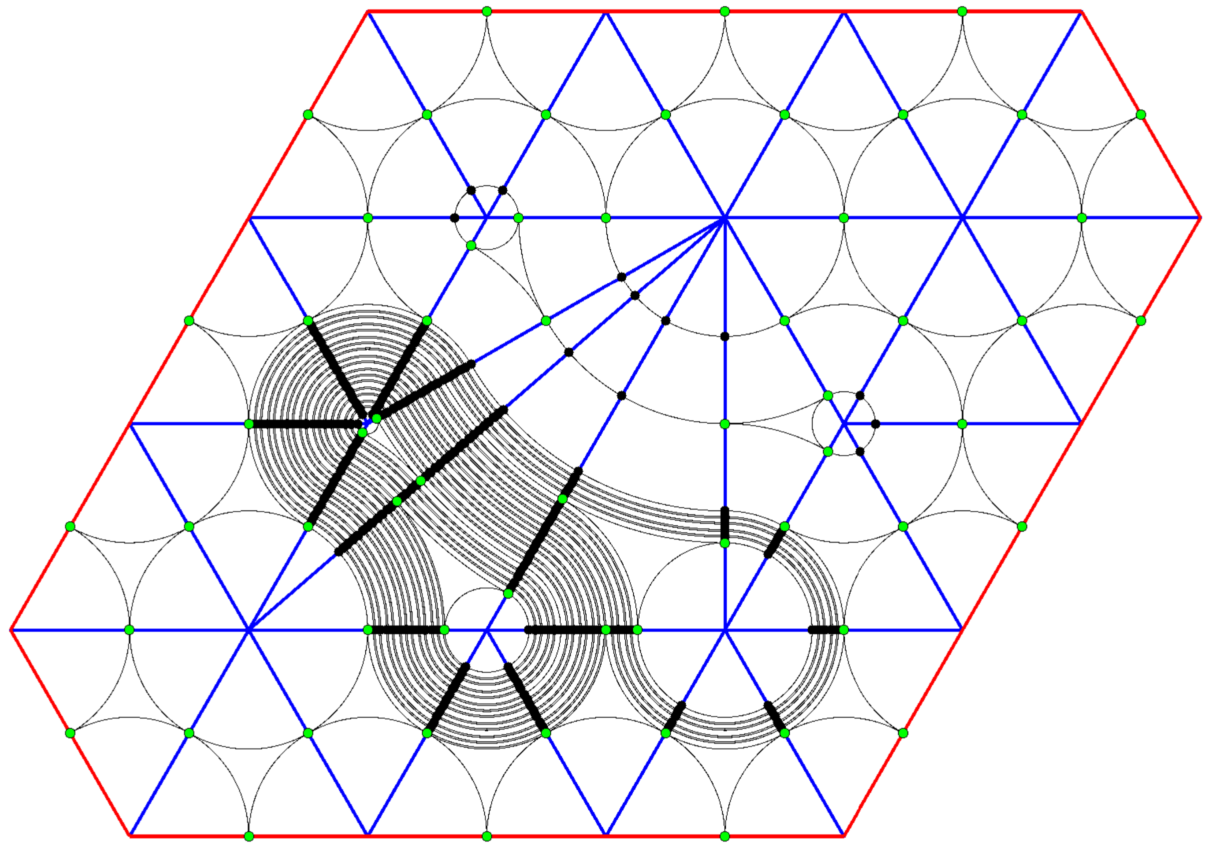

Alternative triangulation of interior: more closed orbits.

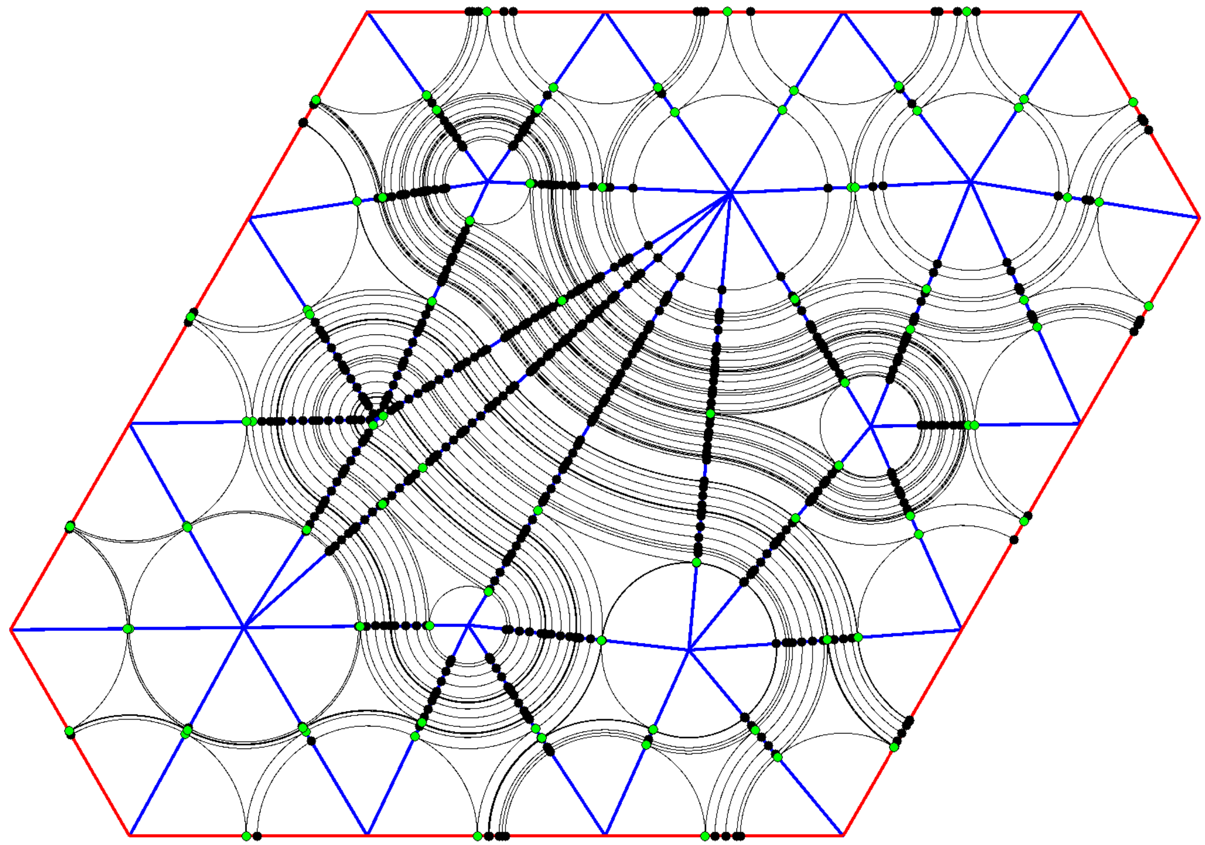

Randomly perturb some vertices; flow lines "leak" to the boundary.

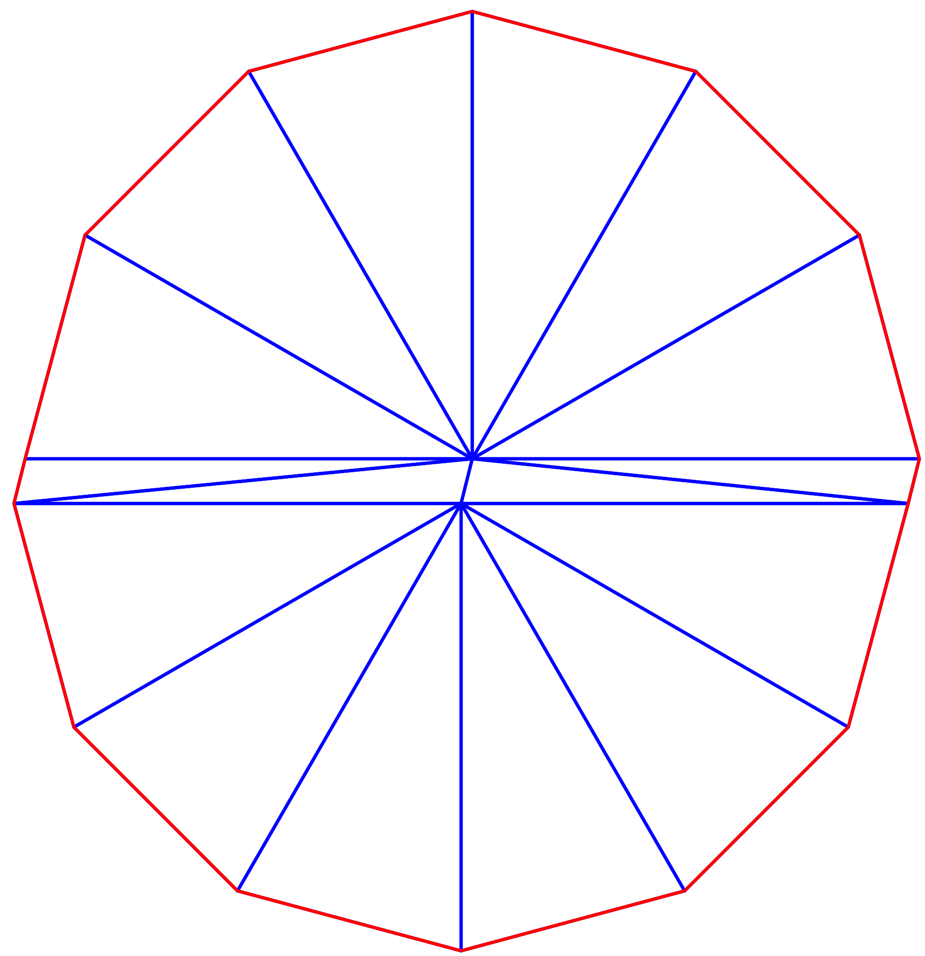

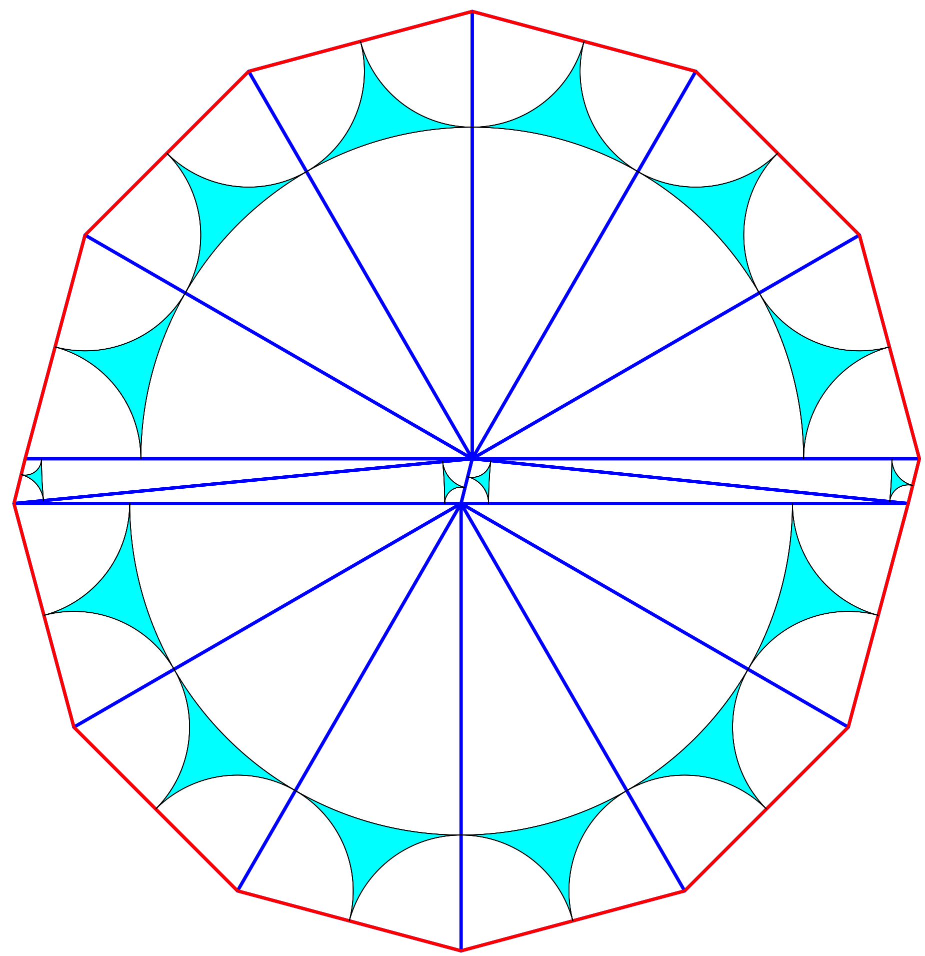

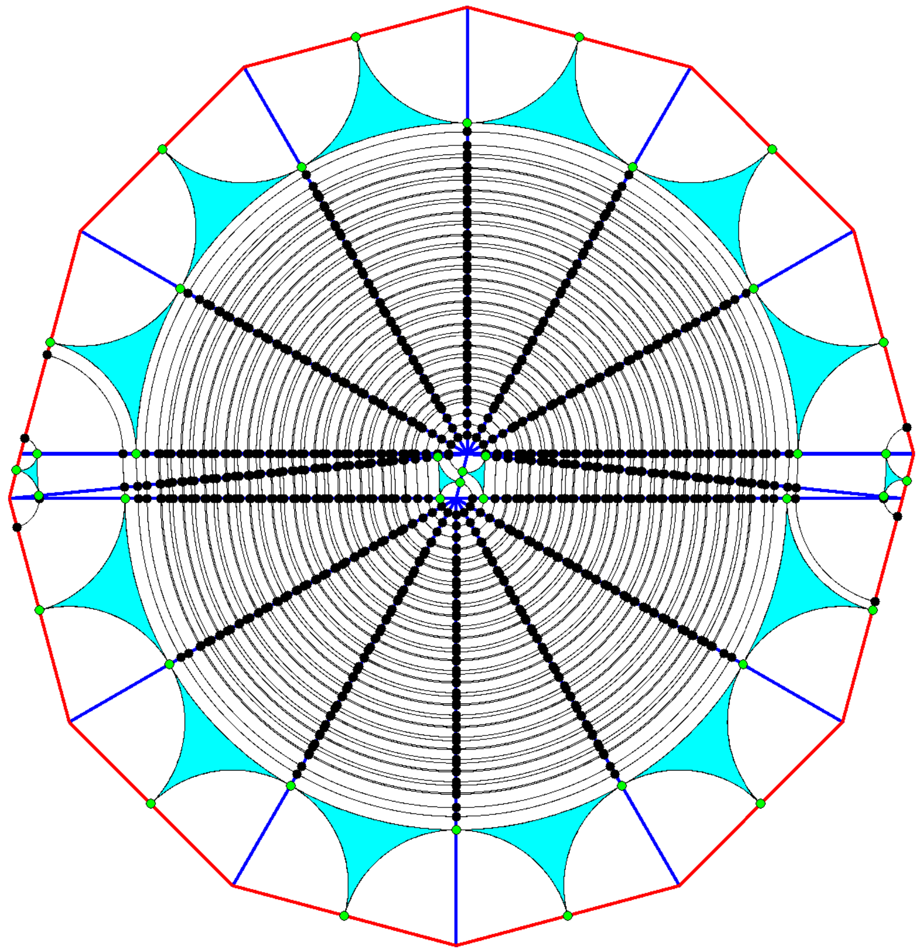

The triangulation flow arises from applications to optimal meshing and finite element methods.

A triangulation is **non-obtuse** if every angle is $\leq \pi/2$. Called a **NOT** for short. Triangulation is **acute** if every angle is $< \pi/2$.

Non-obtuse triangulations important in various applications, e.g, give better numerical methods.

For example, Vavasis showed that matrices arising from finite element method for a certain PDE have conditions numbers that grow exponentially (in number of triangles) for general triangulations, but only linearly for non-obtuse triangulations.

Other numerical methods are faster, simpler to implement, or provably correct when using non-obtuse triangulations.

Fact: every triangulation has a non-obtuse refinement (possibly using many, many more triangles).

Fact: some n-triangulations require n^2 elements in a non-obtuse refinement.

Fact: No polynomial bound is possible with angle bound $\theta < 90°$.

Question: does every triangulation have a polynomial sized non-obtuse refinement?

Fact: every triangulation has a non-obtuse refinement (possibly using many, many more triangles).

Fact: some n-triangulations require n^2 elements in a non-obtuse refinement.

Fact: No polynomial bound is possible with angle bound $\theta < 90°$.

Question: does every triangulation have a polynomial sized non-obtuse refinement?

Answer: Yes (B. 2016). Any n-triangulation has a non-obtuse refinement with $O(n^{2.5})$ elements. Gap remains between 2 and 2.5.

The segment $[v, w]$ is a **Gabriel** edge of a point set V if it is the diameter of an open disk missing V. (Special case of Delaunay condition.)

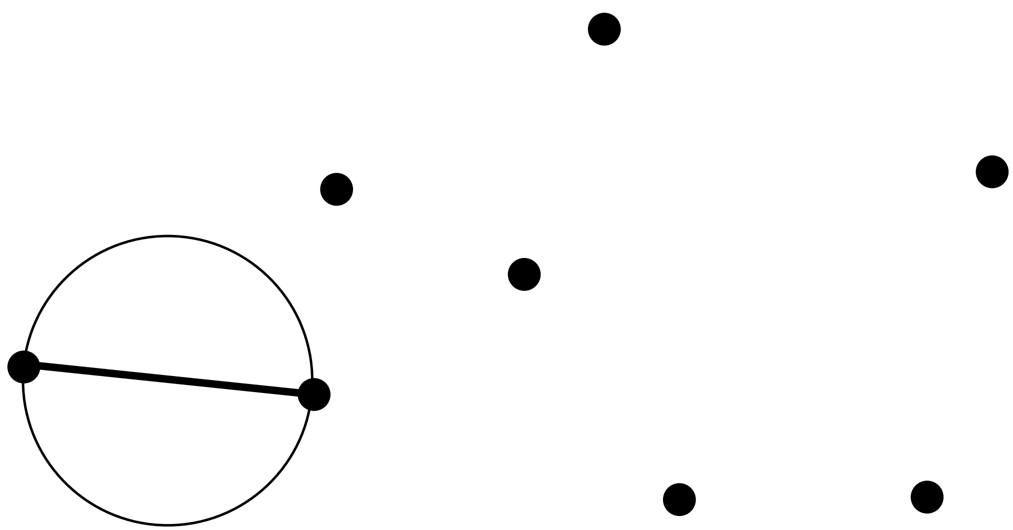

Gabriel edge.

The segment $[v, w]$ is a **Gabriel** edge of a point set V if it is the diameter of an open disk missing V. (Special case of Delaunay condition.)

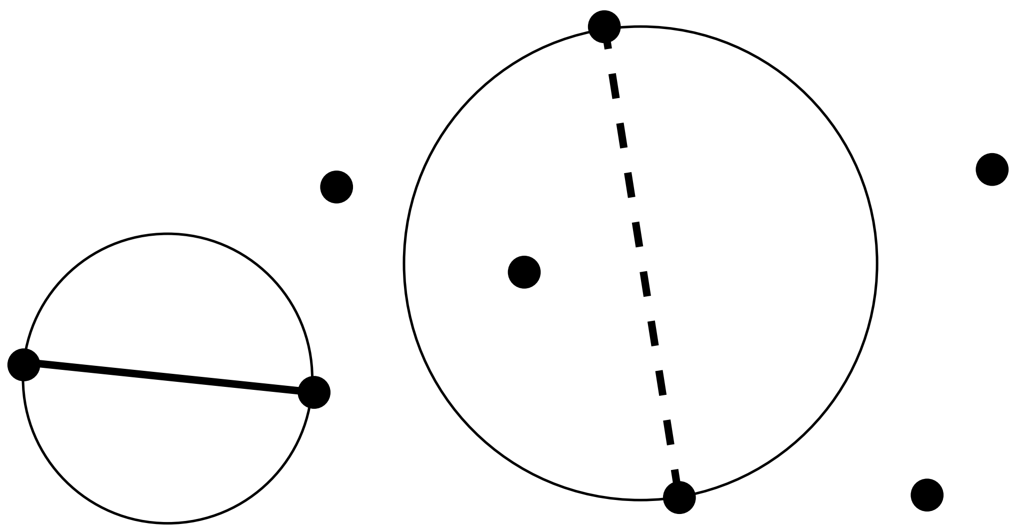

Not a Gabriel edge.

It's easy to see that every edge of a NOT is Gabriel.

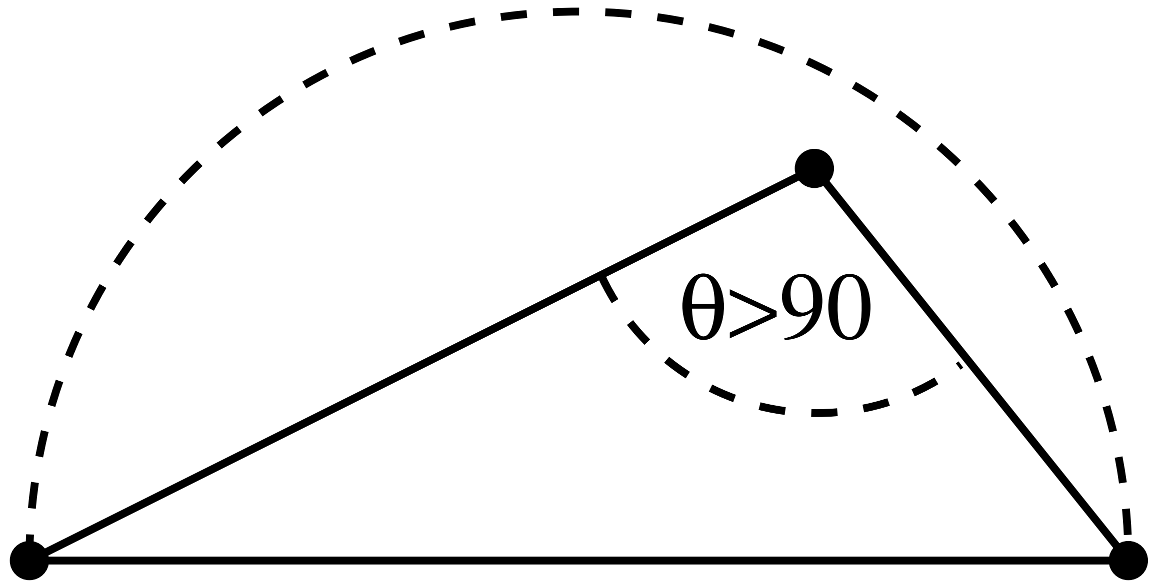

The converse is almost true in following sense.

Lemma (Bern-Mitchell-Rupert 1994): if we add N vertices to edges of a n-triangulation so every edge becomes Gabriel, then there is a non-obtuse refinement with $O(n + N)$ triangles.

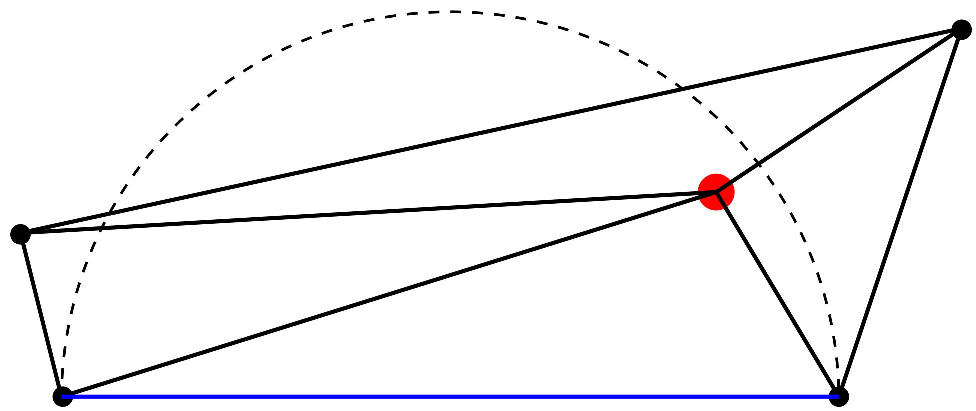

Lemma (Bern-Mitchell-Rupert 1994): if we add N vertices to edges of a n-triangulation so every edge becomes Gabriel, then there is a non-obtuse refinement with $O(n+N)$ triangles.

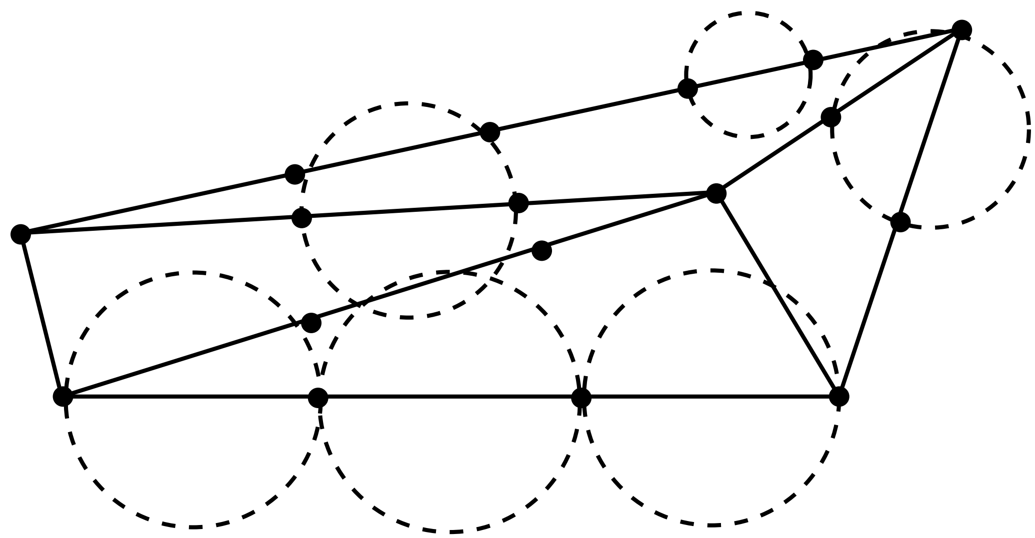

Idea: each triangle can be non-obtusely refined using only the given points on boundary. Thus refinements mesh together.

We can add such points by using triangulation flows.

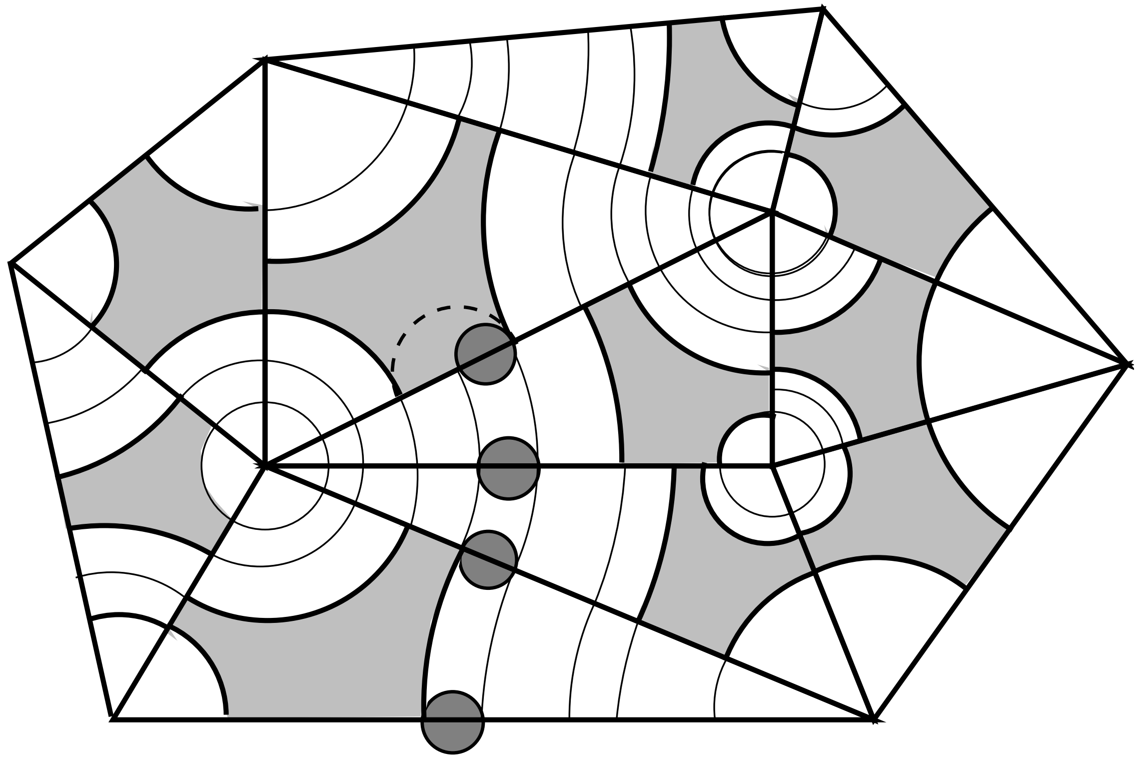

If suffices to consider "thickened" central regions. With "thick" regions, **every** flowline terminates in finitely many steps (depends on geometry).

Tube is "swept out" by fixed diameter disk. Verifies Gabriel condition: disk lies inside tube or central region or outside convex hull.

Simple-NOT-Theorem: any triangulation of a simple n-gon by diagonals can be refined to a $O(n^2)$-NOT.

Proof: We saw earlier that the triangulation flow generates at most n^2 new points. QED

Improves 1993 $O(n^4)$ bound of Marshall Bern and David Eppstein.

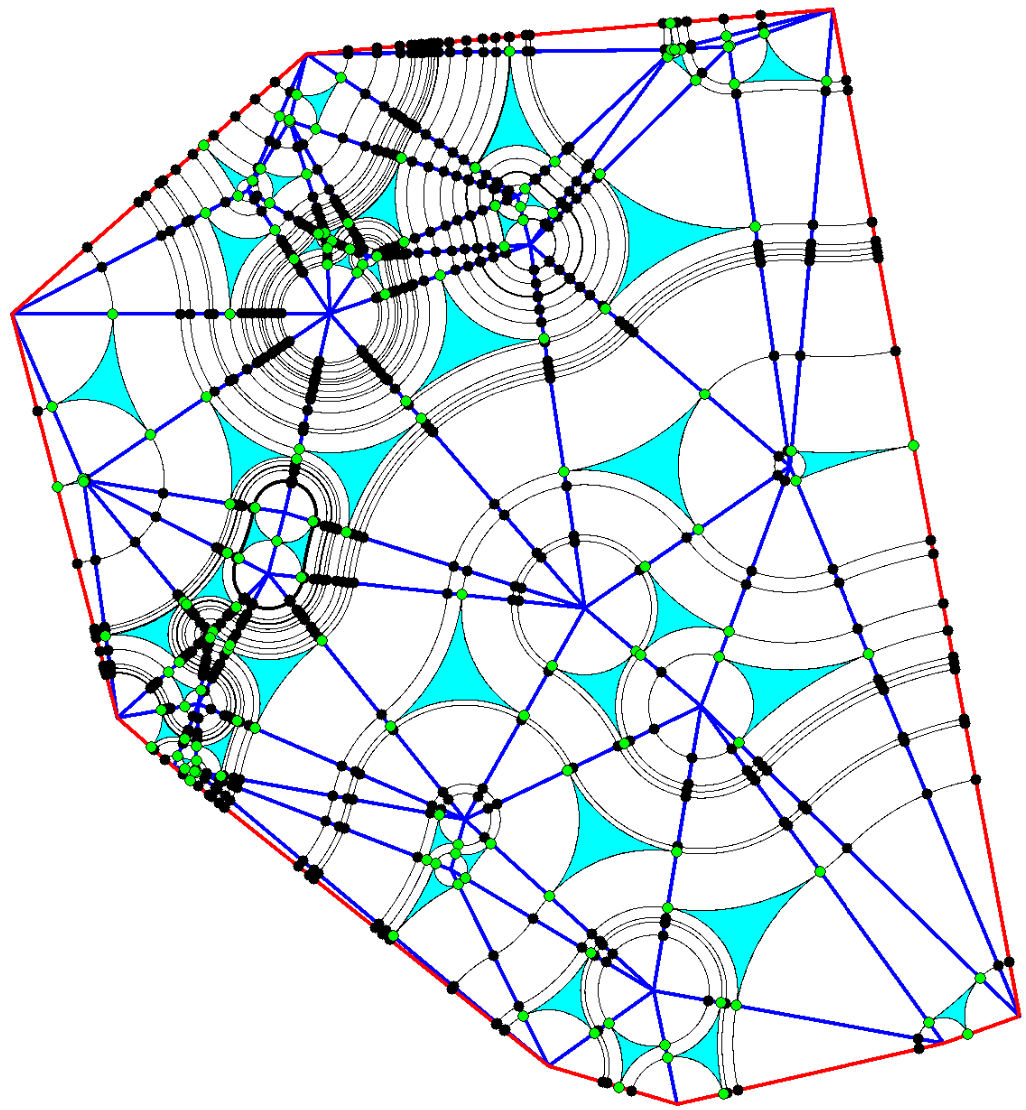

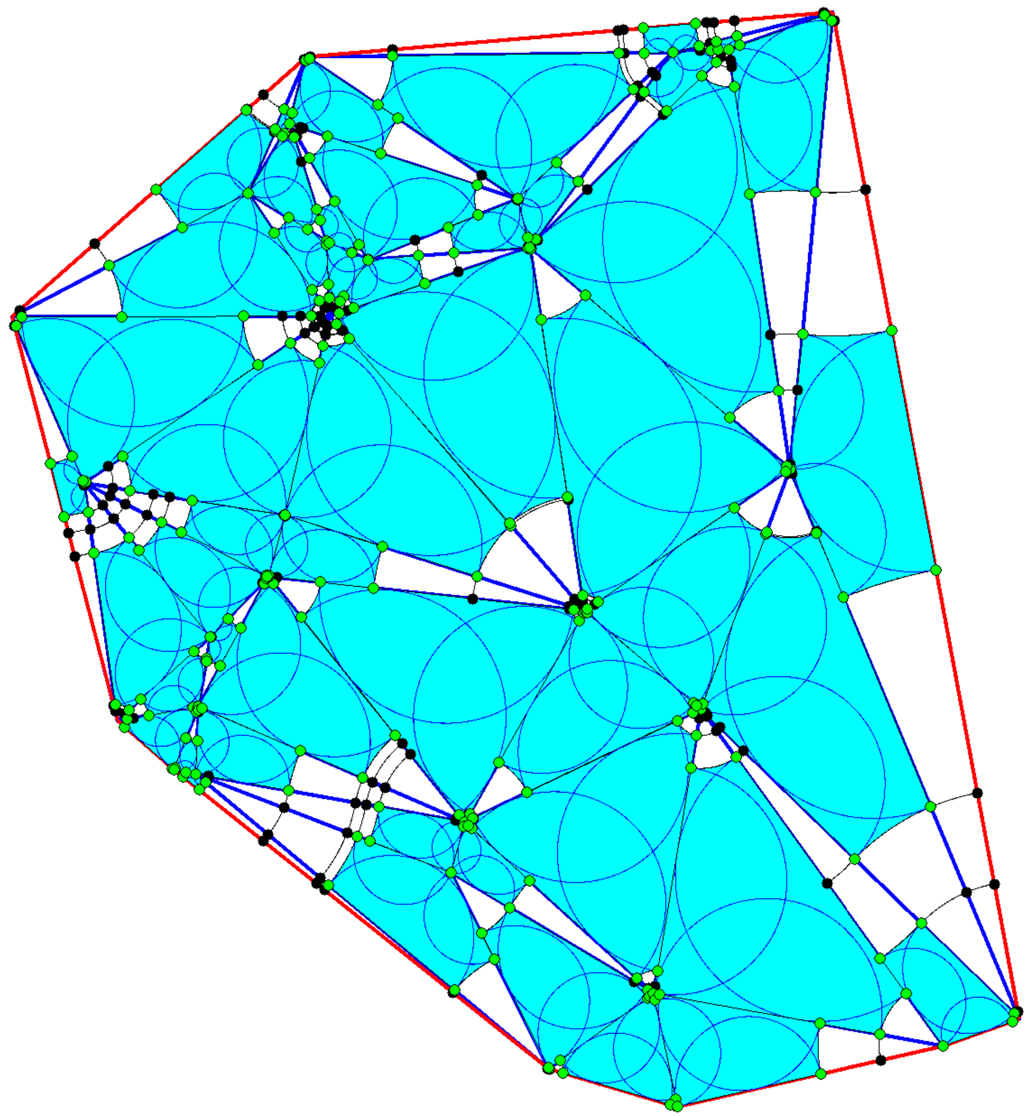

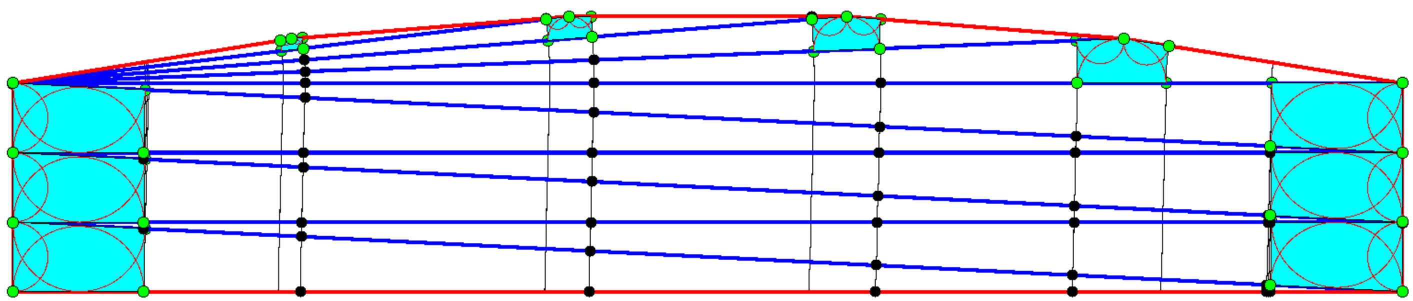

General NOT-theorem: Any triangulation with n triangles, can be refined to a NOT with $O(n^{3/2})$ using non-obtuse triangles.

Idea: we perturb the flow lines (instead of the triangulation), in order to get them to run into each other, while still giving the Gabriel condition. Requires adding about $n^{3/2}$ new lines; each visits at worst n triangles.

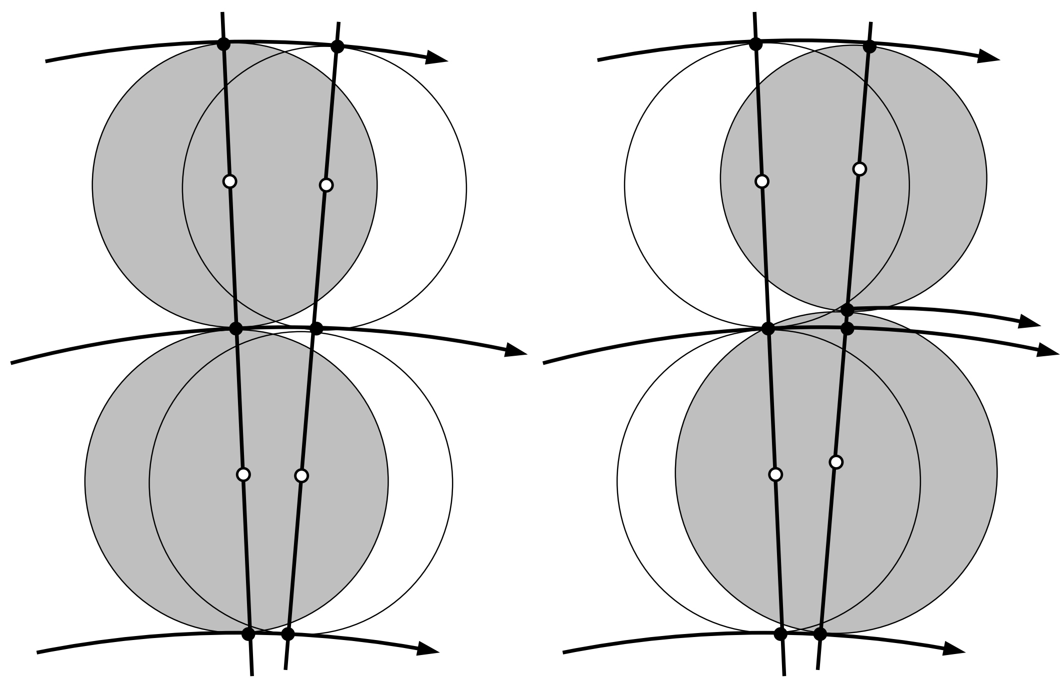

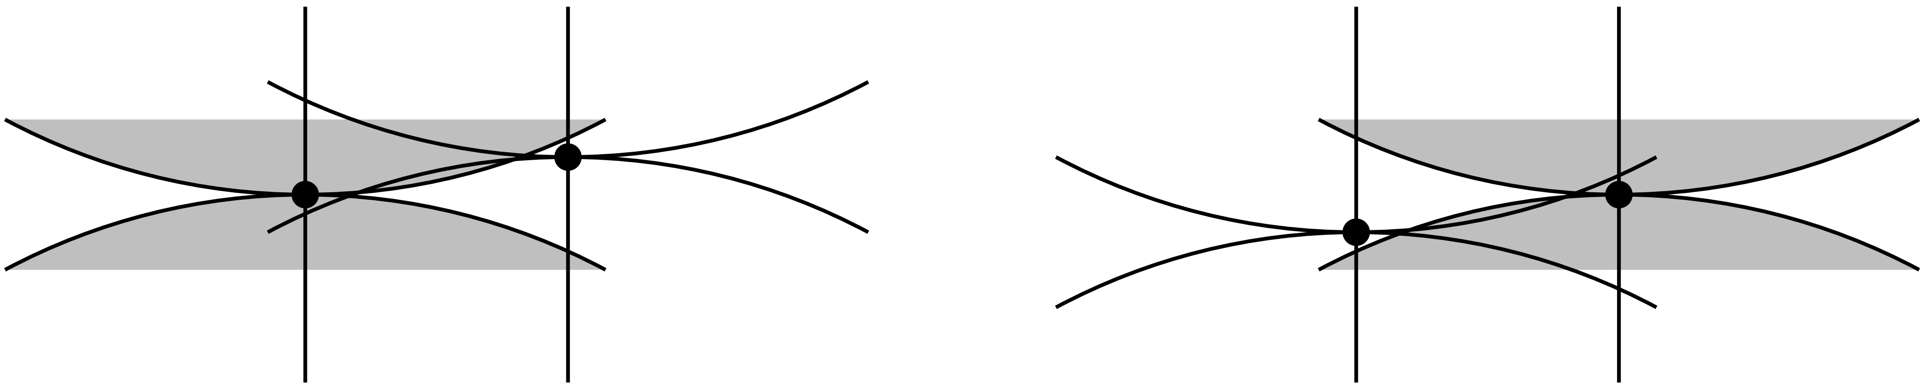

Adding lots of points is like increasing 2nd derivative of flow.

Argument is a discrete version of a "closing lemma": if we limit the amount of bending, we must estimate how many steps are needed to create closed orbits in the perturbed triangulation flow.

Can we make a precise connection to closing lemmas in surface dynamics?

What about triangulations of polyhedral surfaces?

Hyperbolic triangulations of Riemann surfaces?

What about 3 dimensions?

THANKS

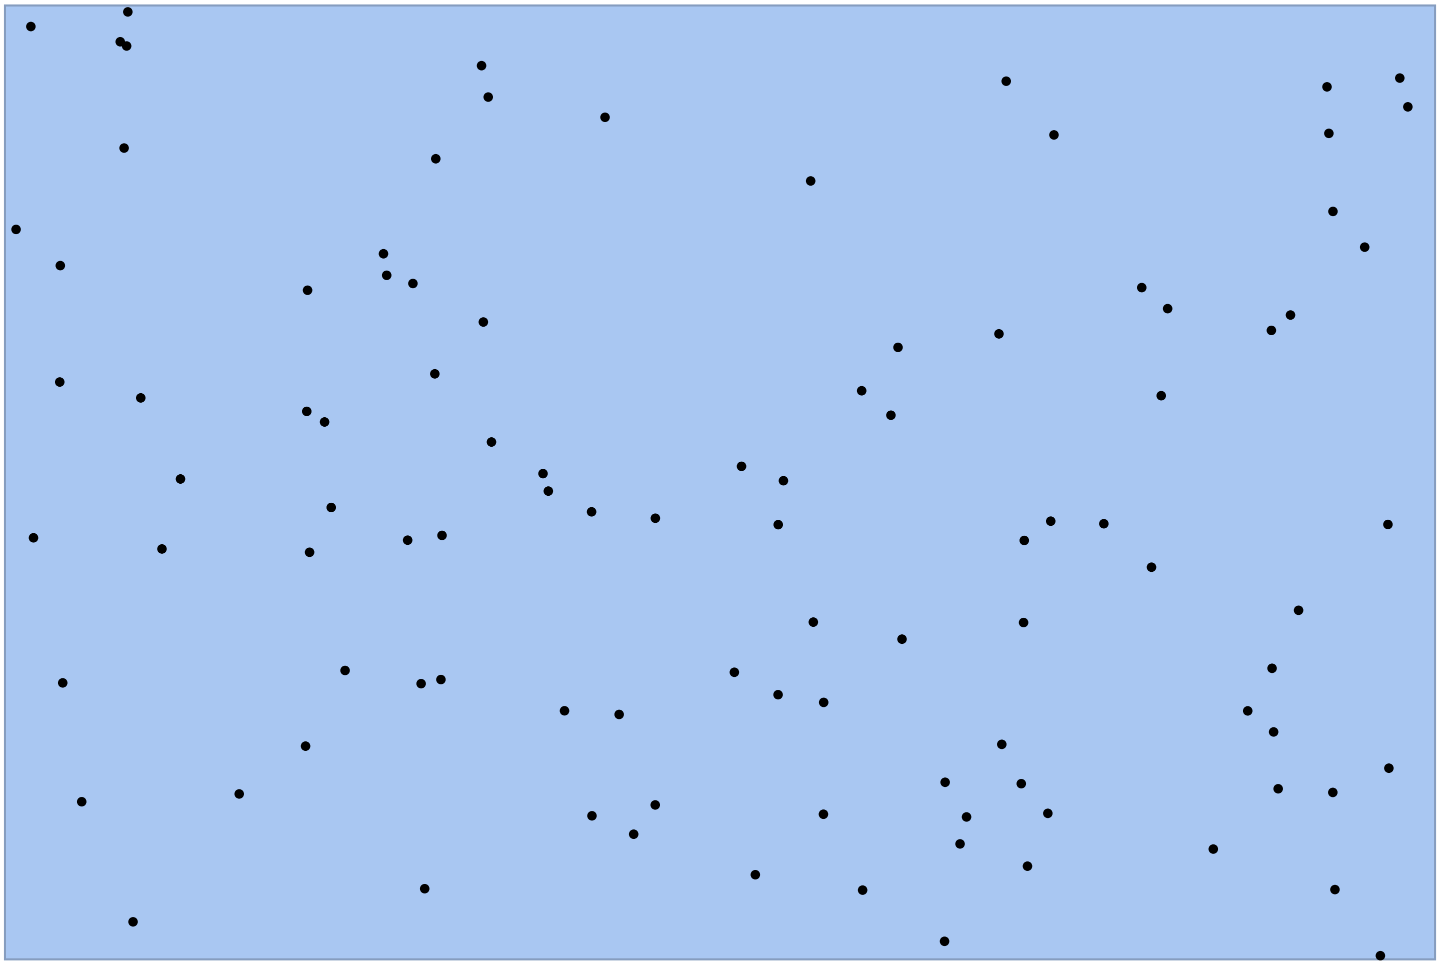

An application of the NOT theorem
Consider a finite set of points in the plane.

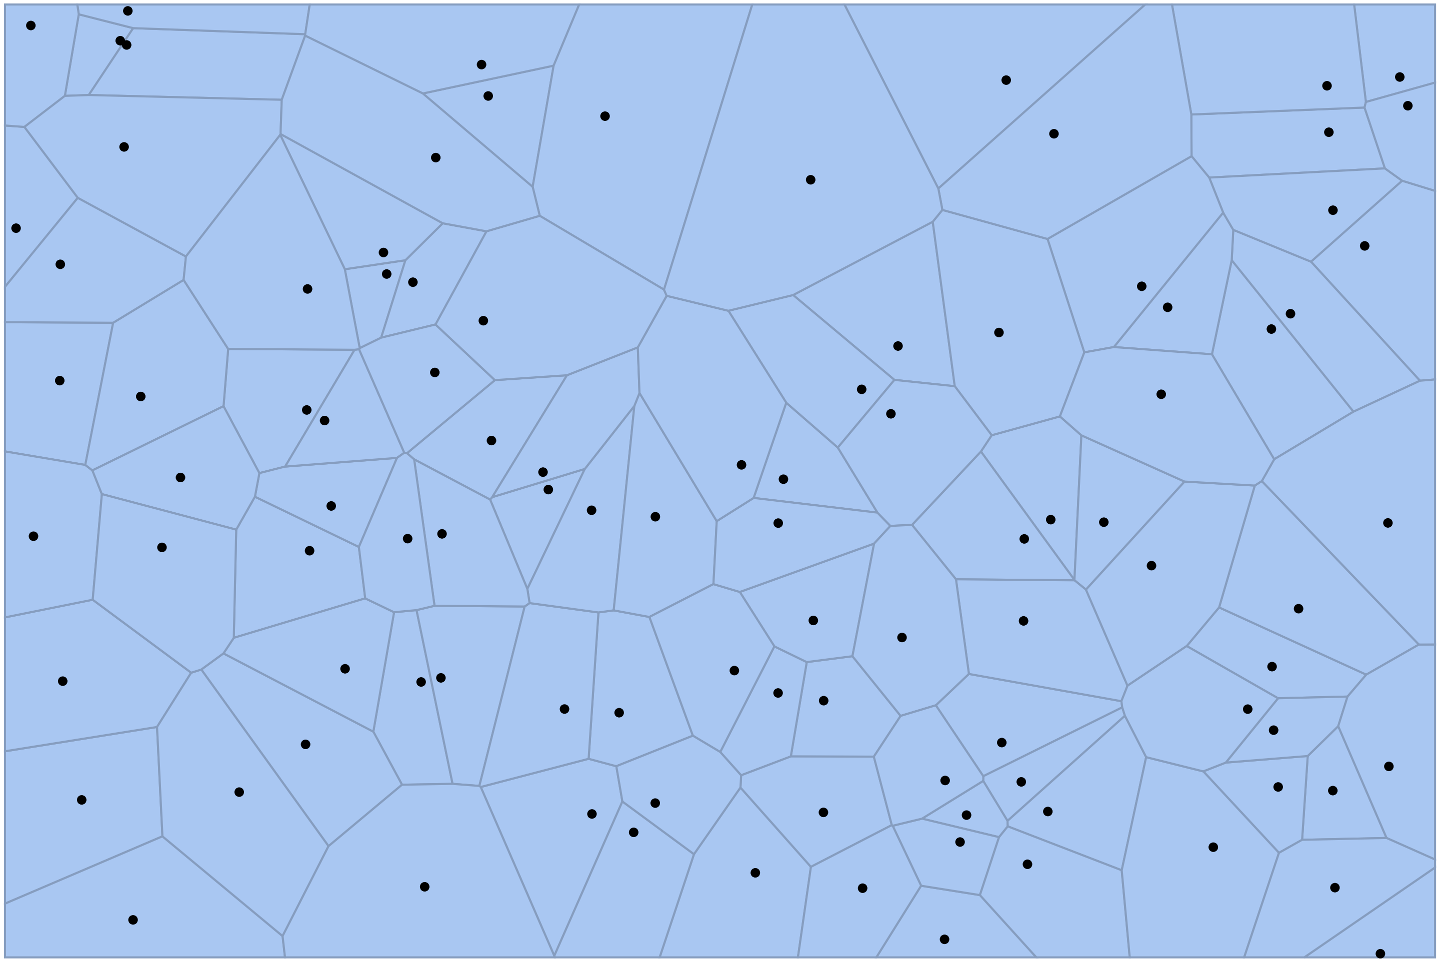

Voronoi cells (think of cell phone connecting to closest tower).

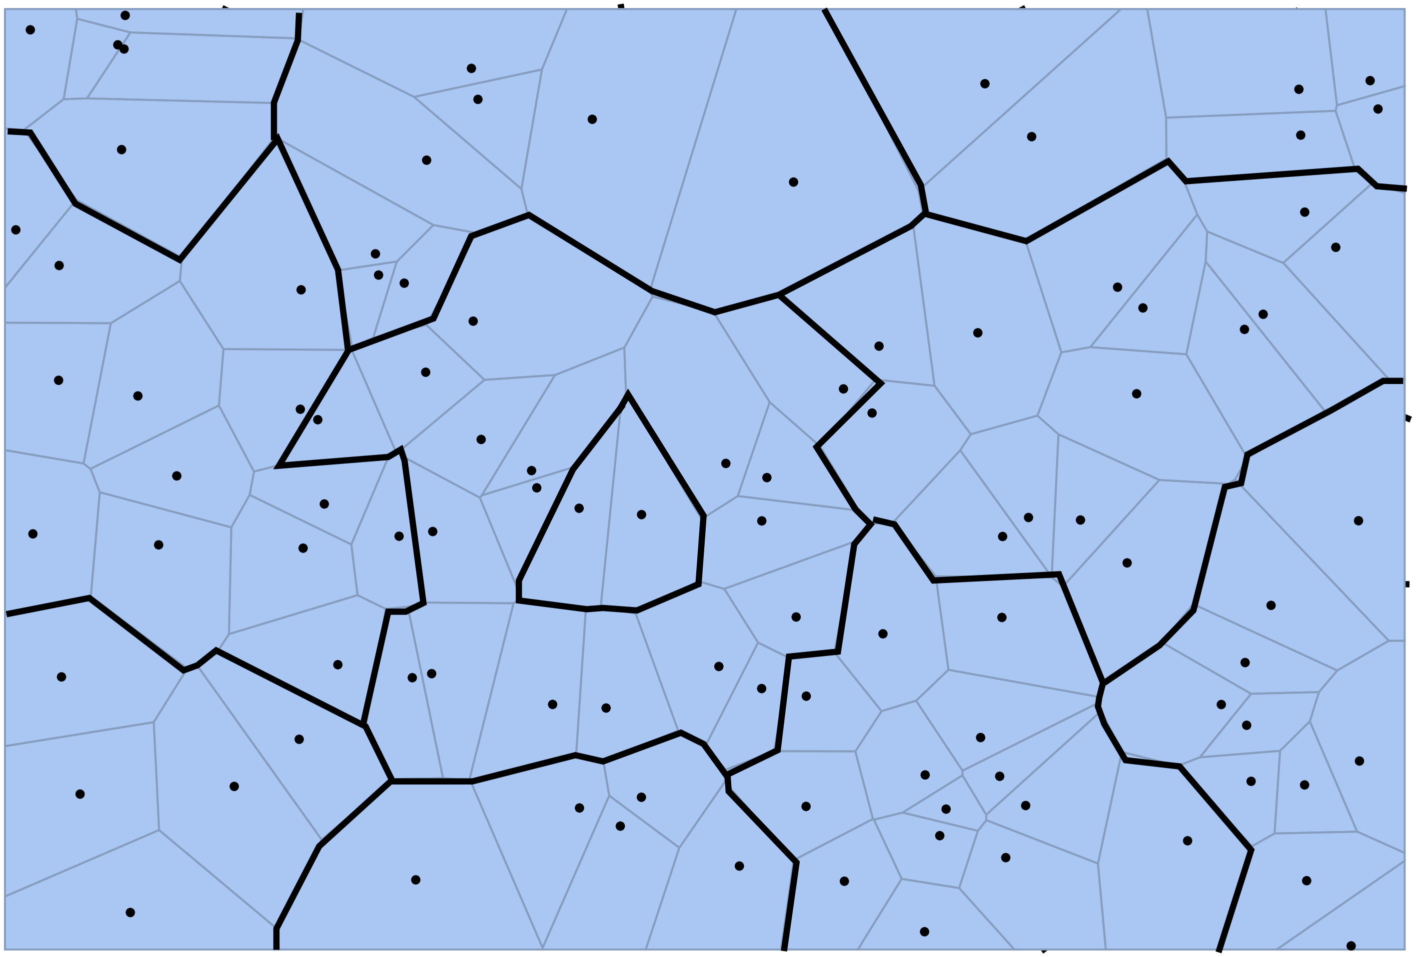

If region boundaries conform to cell boundaries, then a phone always connects to a tower in the same region.

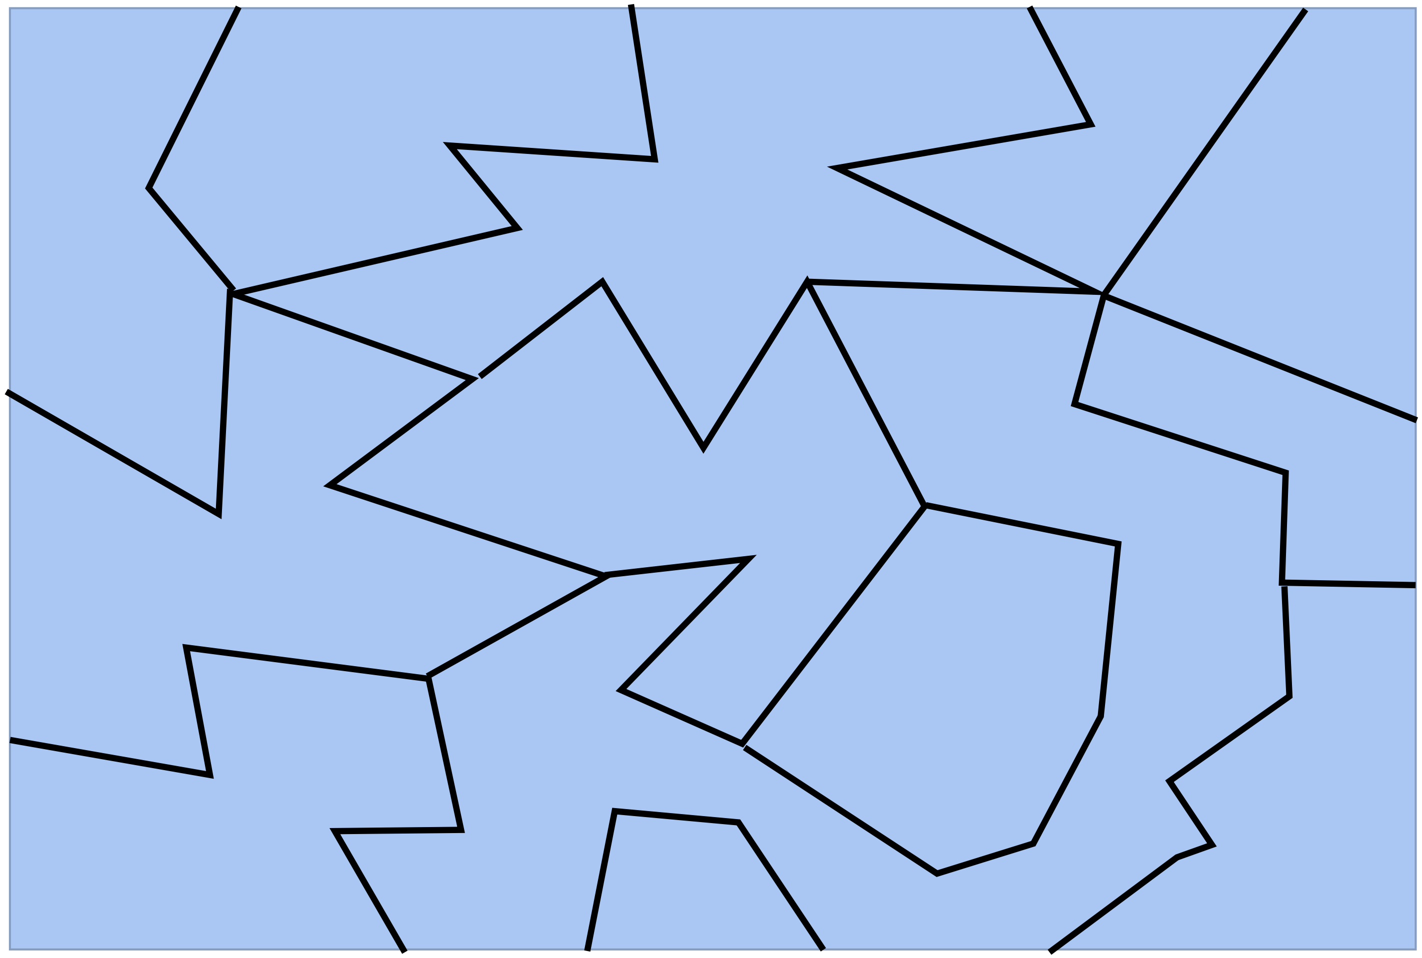

Given countries, can we place towers so this happens?

Do a polynomial number of towers suffice?

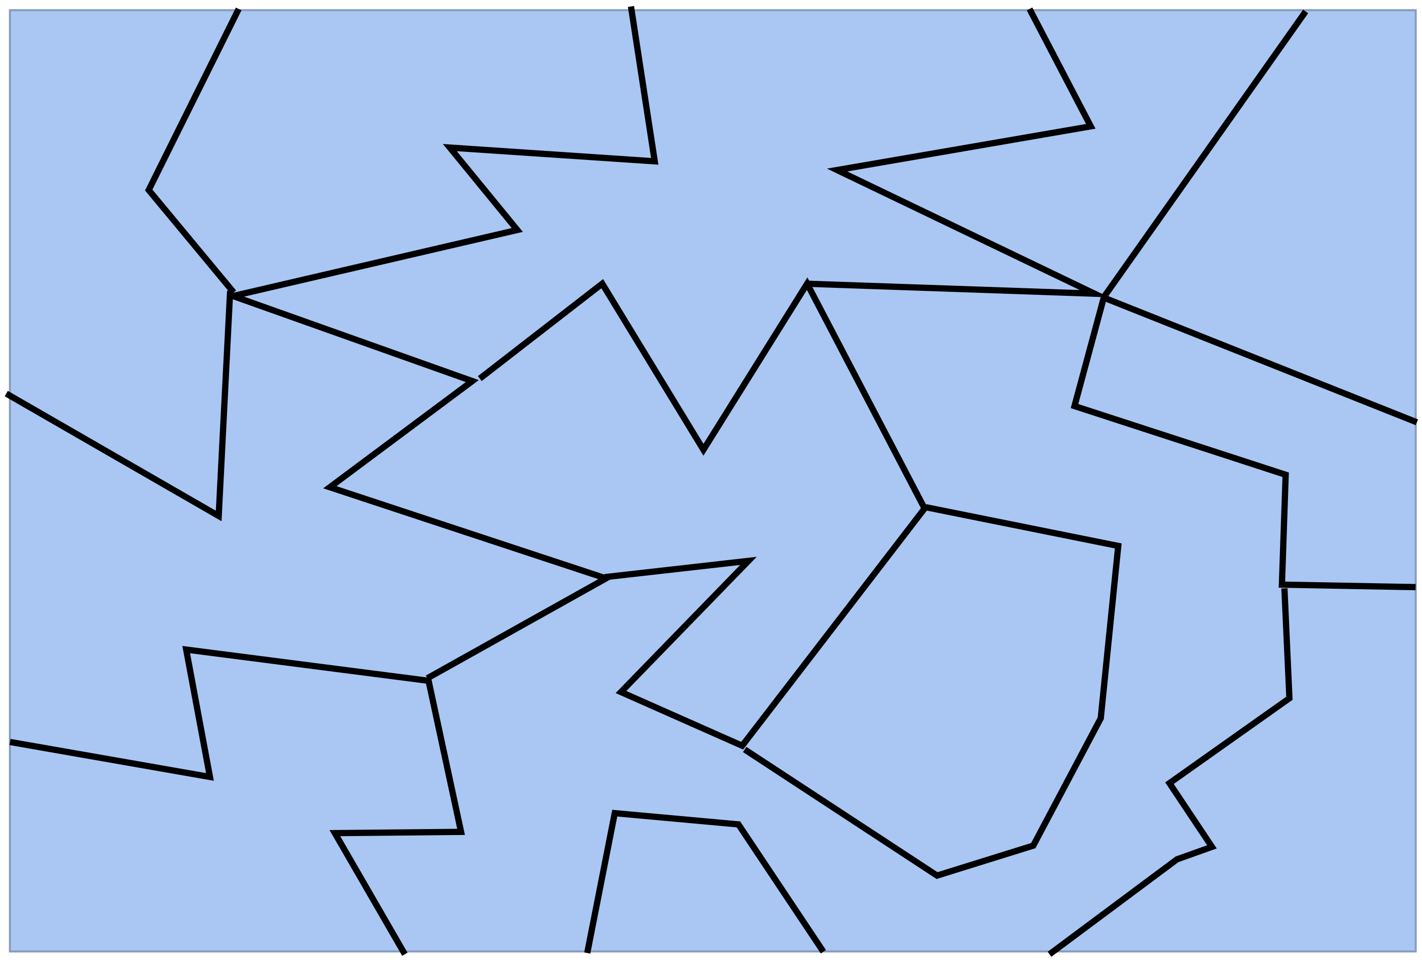

Given countries, can we place towers so this happens?

Do a polynomial number of towers suffice? Yes (B 2016)

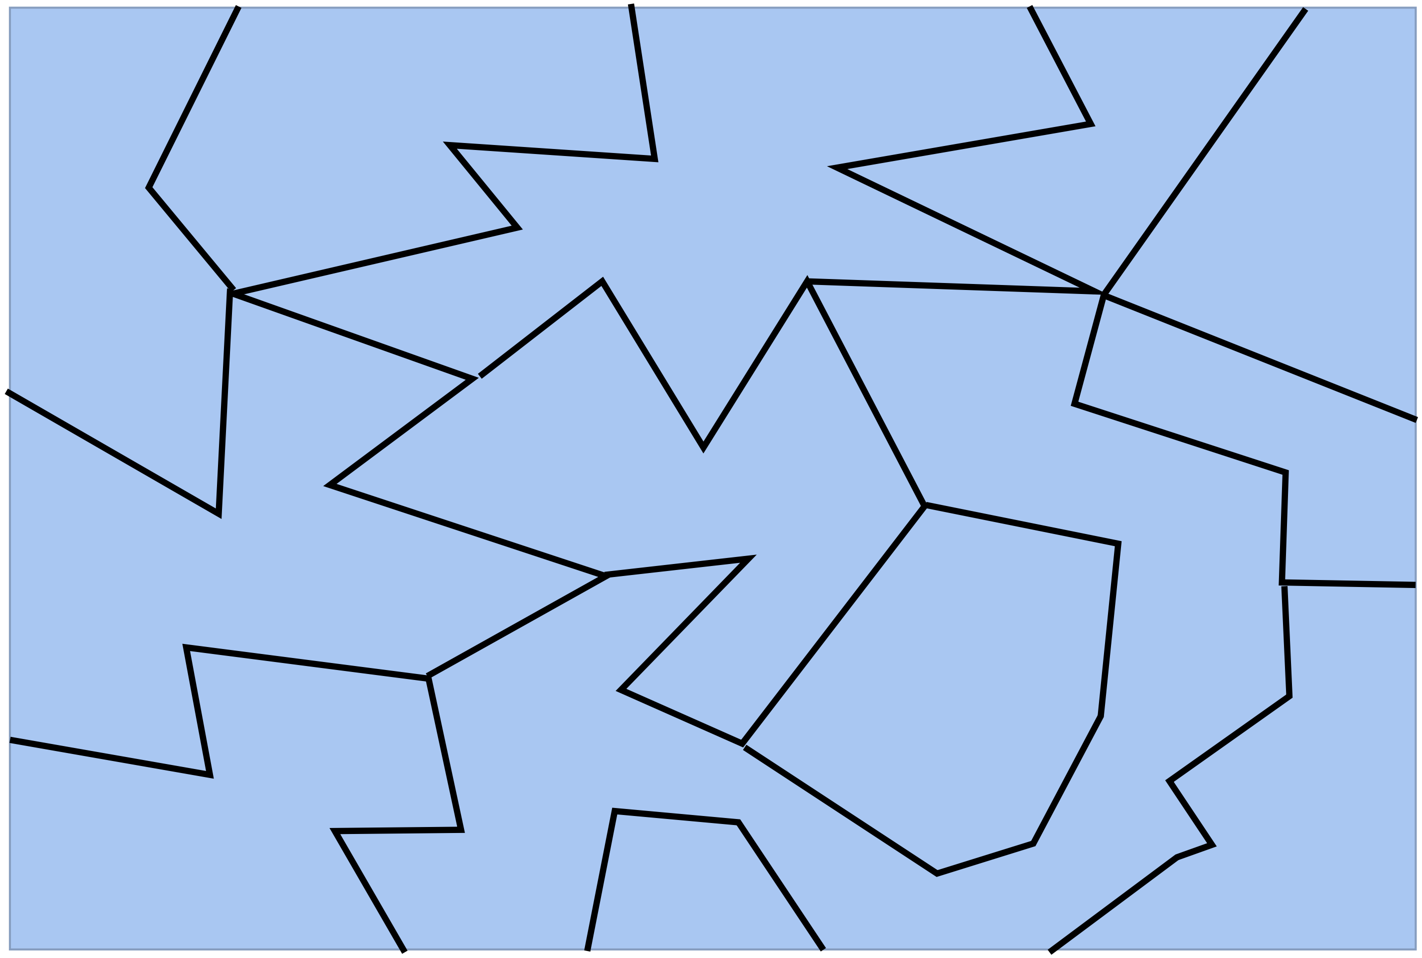

Proof: It's easy to place points explcitly if regions are all non-obtuse triangles. In general, triangulate the regions, then non-obtusely refine the triangulation.