
A Dynamical System Using the Voronoï Tessellation

Natalie Priebe Frank and Sean M. Hart

1. INTRODUCTION. Suppose you know the locations of post offices or cell phone satellites, and you want to know what regions they serve. Or maybe you know the locations of atoms in a crystal, and you want to know what a fundamental region looks like. There are lots of reasons you might want to make a tiling around a given discrete set of points. A natural way to do it is with “Voronoi tessellations”—so natural, in fact, that it has been rediscovered numerous times over the years.

On the other hand, if you have a tiling, you might want to decorate each tile with a few points to create or destroy symmetry. Or you might look at all the **vertices** of the tiling—points where three or more tiles meet—to extract combinatorial information. If you have a tiling, there are many ways to obtain a point set from it.

So, you can get tilings from point sets and point sets from tilings: doesn’t this give you a way to associate point sets to point sets or tilings to tilings? Once you have a map from a class of objects back to itself, you can take a dynamical systems viewpoint to analyze the situation. In this paper we are going to do exactly that, with a new dynamical system based on the vertices of Voronoi tessellations.

For those uninitiated with the Voronoi tessellation, we begin with its definition and then give the definition of our dynamical system. From there, the remainder of §1 is spent exploring the evolution of simple point sets, using these simplified examples to develop both the intuition and vocabulary needed for more interesting cases. In §2 we give a new proof of a theorem, first proved in [4], quantifying the growth in size of point sets over repeated iteration. Following that we will point out some interesting corollaries and give some estimates on the growth rate. We devote §3 to discussing what questions interest us most from the dynamical systems viewpoint.

For now, let’s turn to the definitions.

1.1. Our Dynamical System. We start with a finite point set $P \subset \mathbb{R}^2$, which we call the **generating set**, the members of which we refer to as the **generators**. (We relax the assumption that P be finite and/or planar in §3.) The **Voronoi polygon** of a point $p \in P$, denoted $V(p)$, is given by

$$V(p) = \{x \in \mathbb{R}^2 \mid \|p - x\| \leq \|p' - x\| \text{ for all } p' \in P\}.$$

Simply put, the Voronoi polygon of p contains every point in the plane that is closer to p than to any other member of P , or is equidistant between p and a nearby generator point. The **Voronoi tessellation** of P is given by

$$\Upsilon(P) = \{V(p) \mid p \in P\}.$$

A point set and its Voronoi tessellation are given in Figure 1.

Given distinct $p, p' \in P$, the sets $V(p)$ and $V(p')$ are not necessarily disjoint; in fact, the boundary of each Voronoi polygon is shared with other Voronoi polygons. If

doi:10.4169/000298910X476022

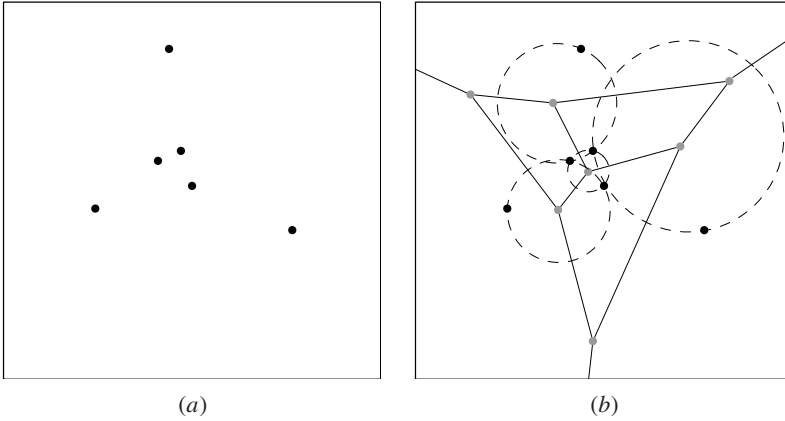


Figure 1. A point set P , its Voronoi tessellation, and $v(P)$.

$V(p) \cap V(p')$ is a line, ray, or line segment, we call it a **Voronoi edge** and denote it $e_{p,p'}$. If the intersection of three or more Voronoi tiles is a point, we call that point a **Voronoi vertex**. The sets of all Voronoi edges and vertices are denoted by $\mathcal{E}(\Upsilon(P))$ and $\mathcal{V}(\Upsilon(P))$, respectively.

It is useful to notice that $e_{p,p'}$ always lies on the perpendicular bisector of the line between p and p' . This gives a method for constructing Voronoi diagrams: for $p \in P$, sketch each perpendicular bisector between p and another member of P . Then the Voronoi polygon $V(p)$ is the intersection of all half-planes created by the perpendicular bisectors. It is also useful to notice that a Voronoi vertex is equidistant from the generator points of the tiles it is in. For proofs of these properties and a wealth of other information about Voronoi tessellations, [7] is an excellent source.

Now, the set $\mathcal{V}(\Upsilon(P))$ constitutes a point set in its own right, and so one might naturally wonder what its Voronoi tessellation looks like. And we need not stop there—the Voronoi tessellation of $\mathcal{V}(\Upsilon(P))$ will have a vertex set, too, so how does its Voronoi tessellation behave? We have the makings of a dynamical system on the set $\mathcal{P}(\mathbb{R}^2)$ of all finite point sets in the plane.

Definition 1.1. Let $P \in \mathcal{P}(\mathbb{R}^2)$. We define the **Voronoi iteration**¹ of P to be

$$v(P) = \mathcal{V}(\Upsilon(P)).$$

For $n = 2, 3, \dots$ we define the **n th Voronoi iteration** of P recursively:

$$v^n(P) = \mathcal{V}(\Upsilon(v^{n-1}(P))) = v(v^{n-1}(P)).$$

Figure 1(b) depicts the Voronoi tessellation of the six points depicted in 1(a), and so the lighter seven points—the vertices of the tessellation—constitute the Voronoi iteration of the original set. Let's begin looking at the simplest cases this dynamical system has to offer.

1.2. The Really Trivial Cases. We follow the convention of using $|P|$ to denote the cardinality of P . If $|P| = 1$, then the single Voronoi polygon is all of \mathbb{R}^2 . Since \mathbb{R}^2 is just one big tile, it has no vertices, and so $v^n(P) = \emptyset$ for all $n \geq 1$. Next, we up the

¹This is not to be confused with Lloyd's algorithm from computer graphics (see [3] and references therein), which is sometimes also called "Voronoi iteration."

ante: suppose $P = \{p_1, p_2\}$. The Voronoi tessellation then fractures the plane into two half-planes, split along the perpendicular bisector of $\overline{p_1 p_2}$, the line segment joining p_1 to p_2 . The vertex set, however, is still empty, and so $v^n(P) = \emptyset$ for all $n \geq 1$.

There are two possibilities when P has three points: either all three are collinear, or they lie on a circle. The former yields $v^n(P) = \emptyset$ for all $n \geq 1$ and in the latter, $v(P)$ is a single point. No matter how large P is, the special cases of collinearity and cocircularity always work out like this:

Proposition 1.2. *All the points in P are collinear if and only if $v^n(P) = \emptyset$ for all $n \geq 1$. All the points in P are cocircular if and only if $|v(P)| = 1$ and $v^n(P) = \emptyset$ for $n > 1$.*

The first fact follows naturally from the observation that if $p, q \in P$ are such that $e_{p,q} \in \mathcal{E}(\Upsilon(P))$, then $e_{p,q}$ lies on the perpendicular bisector of the line segment \overline{pq} . When all the points are collinear, the perpendicular bisectors are all parallel and thus do not intersect to produce new vertices. Conversely, if $v(P) = \emptyset$ then no Voronoi edges intersect, which implies all Voronoi edges are parallel, and hence P must be collinear.

To prove the cocircularity result, we introduce the concept of an **empty circle**. This is a circle whose interior does not contain any generator points and whose boundary passes through three or more generator points. Such a circle gets its name since it is “empty” of generator points. Proposition 1.2 follows immediately from the next Proposition.

Proposition 1.3. *A point $q \in \mathbb{R}^2$ is a Voronoi vertex in $\mathcal{V}(\Upsilon(P))$ if and only if it is the center of an empty circle.²*

In addition to proving Proposition 1.2, this proposition also lets us quickly find the vertex set of $\Upsilon(P)$. To do so, pick a noncollinear triple $p, q, r \in P$, and look at the unique circle passing through them. If this circle is empty, then place a vertex at its center. Once you have checked every triple, every vertex will be accounted for—and you never had to sketch an edge!

So, we know what happens when the point set is really small, or collinear, or cocircular. Let’s move on to . . .

1.3. A Slightly Less Trivial Case. Suppose $|P| = 4$. To discriminate between configurations that aren’t collinear or cocircular, we require a more sophisticated vocabulary. We say that a subset $A \subset \mathbb{R}^2$ is **convex** if, given any distinct $x, y \in A$, the line segment \overline{xy} is contained in A . (For practice, try proving that Voronoi polygons are convex!) The **convex hull** of a set P is defined to be the smallest convex set containing P . We write $CH(P)$ to mean the convex hull of P , and write $\partial CH(P)$ for its boundary.

Definition 1.4. *For $p \in P$, we say p is on the **boundary** of P and write $p \in \text{Bd}(P)$ if $p \in P \cap \partial CH(P)$. Otherwise, we say p is in the **interior** of P and write $p \in \text{Int}(P)$.*

In general it is clear that

$$|\text{Int}(P)| + |\text{Bd}(P)| = |P|. \tag{1.5}$$

²See [7, p. 61].

We say that two distinct points $p, p' \in \text{Bd}(P)$ are **neighbors on the boundary of** P if $\overline{pp'} \subset \partial CH(P)$ and no $p'' \in P - \{p, p'\}$ lies on $\overline{pp'}$. For noncollinear finite point sets, any point on the boundary has two distinct neighbors.

It is not surprising that the Voronoï tiles of boundary points would have infinite edges. We denote by $\mathcal{E}_F(\Upsilon(P))$ and $\mathcal{E}_I(\Upsilon(P))$ the sets of finite and infinite edges of $\Upsilon(P)$, respectively.

Proposition 1.6. *Assume not all $p \in P$ are collinear, and let $e_{p,p'} \in \mathcal{E}(\Upsilon(P))$. Then $e_{p,p'} \in \mathcal{E}_I(\Upsilon(P))$ if and only if p and p' are neighbors on the boundary of P . Moreover, $e_{p,p'} \in \mathcal{E}_F(\Upsilon(P))$ if and only if p and p' are not neighbors on the boundary of P .³*

This then implies

$$|\mathcal{E}_I(\Upsilon(P))| = |\text{Bd}(P)|. \tag{1.7}$$

Returning to the case of $|P| = 4$, let us assume that the points in P are neither collinear nor cocircular. We may have $|\text{Bd}(P)| = 3$ or 4 , and we are going to prove that Figure 2 represents the only possible types of Voronoï iterations.




Figure 2. Configurations satisfying (a) $|\text{Bd}(P)| = 3$ and (b) $|\text{Bd}(P)| = 4$.

First let's assume $|\text{Bd}(P)| = 3$. Since $|\text{Int}(P)| + |\text{Bd}(P)| = |P| = 4$, say $\text{Int}(P) = \{p\}$. Proposition 1.6 tells us that each of the edges of its Voronoï polygon $V(p)$ must be finite. As $|P - \{p\}| = 3$, there can be at most three edges; since $V(p)$ must be bounded, there must be exactly three edges. So $V(p)$ is a triangle, and the three infinite edges promised by Proposition 1.6 extend from its vertices. Since $|\nu(P)| = 3$, we conclude that $|\nu^2(P)| = 1$ and $\nu^n(P) = \emptyset$ for all $n \geq 3$.

Now let's assume that $|\text{Bd}(P)| = 4$ and that P is neither collinear nor cocircular. We will prove that $|\nu(P)| = 2$. By Proposition 1.3 we know that $|\nu(P)| > 1$. To prove that $|\nu(P)| \leq 2$, we use a trick that will come in handy again, during the proof of our main theorem.

For $q \in \nu(P)$, let $\rho(q)$ be the **degree** of the vertex q : the number of edges touching q . Then we have $\rho(q) \geq 3$ for each $q \in \nu(P)$. The sum of the degrees of all the vertices therefore is greater than or equal to $3 \cdot |\nu(P)|$. Counting edges instead we see

³Follows from [7, p. 59].

that each infinite edge touches exactly one vertex but each finite edge touches exactly two vertices. Thus

$$3 \cdot |v(P)| \leq \sum_{q \in v(P)} \rho(q) = |\mathcal{E}_I(\Upsilon(P))| + 2 \cdot |\mathcal{E}_F(\Upsilon(P))|.$$

Since there are only two pairs of points in P that are not neighbors, Proposition 1.6 implies $|\mathcal{E}_F(\Upsilon(P))| \leq 2$. Thus

$$3 \cdot |v(P)| \leq 4 + 2 \cdot 2 = 8,$$

and so $|v(P)| \leq 2$, proving that in this case, $|v(P)| = 2$. Moreover we can conclude that $v^n(P) = \emptyset$ for all $n > 1$.

1.4. The Rest of the Cases Are All Hard. The case $|P| = 4$ turns out to be the last case for which $v^n(P)$ is guaranteed to equal the null set for large enough n . Indeed, the case $|P| = 5$ has resisted our attempts to understand it! Figure 3 depicts configurations of larger cardinality where the point set does not iterate to a simpler configuration, which in turn enriches (and complicates) matters. We must suspend our case-by-case analysis at this point.




Figure 3. Configurations for which $|v(P)|$ is (a) equal to or (b) greater than $|P|$.

2. HOW BIG IS $v(P)$? Having played with the dynamical system for a while, we, along with our colleagues Allison Edgren and Olivia Gillham, decided to explore the question of what happens to the size of $v^n(P)$ as n goes to infinity. We had some success: a formula [4] (unpublished) that tells us exactly how many points will be in $v(P)$! We are going to share this result with you as soon as we describe one more essential piece of geometric information.

2.1. Relating Degeneracy and Cocircularity. Compare the configuration given in Figure 3(a) with those sketched in Figure 4. Each is a configuration of five points, yet each has a distinctly different iteration. We already have the vocabulary necessary to distinguish the configuration sketched in Figure 4(a) from the others: they differ with respect to the number of points on the boundary. The difference between the other two, however, is more subtle.

A meticulous reader might notice that all of the vertices in Figure 3(a) are of degree three while there is a vertex in 4(b) with degree four. A Voronoi diagram with a vertex




Figure 4. Two configurations of 5 points.

of degree greater than three is called **degenerate**; if all the vertices of the Voronoi diagram are degree three we call it **nondegenerate**. We can quantify this concept of degeneracy with the number $I_c(P)$:

Definition 2.1. Let $\{C_1, \dots, C_k\}$ be the set of all empty circles of a point set P . The number of **instances of cocircularity** of P is given by

$$I_c(P) = \sum_{i=1}^k (|C_i \cap P| - 3).$$

In a nondegenerate point set, every vertex will be of degree three, and so every empty circle will intersect exactly three points in P . In such a case, we compute $I_c(P) = 0$. When a vertex q is of degree $k > 3$, this implies that $|C_q \cap P| = k$. In this case, we say that q contributes $(k - 3)$ instances of cocircularity. This argument proves the following:

Proposition 2.2. Let $\rho(q)$ denote the degree of a vertex $q \in v(P)$. Then

$$I_c(P) = \sum_{q \in v(P)} (\rho(q) - 3).$$

As an aside, we note that another way to think about instances of cocircularity involves **Delaunay (pre)-triangulations**. These tilings arise by connecting the generators whose Voronoi polygons share an edge. If the generator set is nondegenerate, then the resulting graph is a triangulation called the Delaunay triangulation; when the generator set is degenerate, however, there exist Delaunay tiles that are not triangles. The graph is then called the Delaunay pre-triangulation, and the number of edges we need to add to obtain a triangulation is precisely the number of instances of cocircularity.

When given a degenerate configuration, one can obtain a nondegenerate point set by simply shifting any cocircular points by an arbitrarily small amount; further, when given a nondegenerate configuration, it remains nondegenerate under sufficiently small perturbations. Hence, when studying or applying Voronoi tessellations it is customary to ignore degenerate configurations: the modified nondegenerate configuration is usually “close enough” to the original, and if the points are collected from real-world instruments we can never assume that our measurements are exact anyway.

If only we could make such an assumption! In our system, computing $I_c(P)$ is usually only possible if you already know where the vertices are, and, as we shall see, incorporating degeneracy makes determining the long-term behavior much more difficult. But we cannot simply throw out these configurations; there exist situations, like that which is depicted in Figure 5, that ruin things for us. A nondegenerate point set may iterate to a degenerate one, and so assuming P is nondegenerate does not guarantee $\nu(P)$ is as well.




Figure 5. $I_c(P) = 0$ while $I_c(\nu(P)) = 1$.

2.2. The Theorem on Counting Vertices. We are ready to state and prove the theorem discovered by the second author and his colleagues in [4] (unpublished). The proof we present here, due only to the second author, is much simpler than the original.

Theorem 2.3. *If not all $p \in P$ are collinear, then*

$$|\nu(P)| = 2 \cdot |P| - |\text{Bd}(P)| - I_c(P) - 2.$$

Proof. We can use the trick from §1.3, summing the degrees of the vertices in $\nu(P)$. Since P is noncollinear, every infinite edge intersects exactly one vertex and every finite edge intersects exactly two. Hence, recalling from §1.3 that $|\mathcal{E}_I(\Upsilon(P))| = |\text{Bd}(P)|$, we get

$$\sum_{q \in \nu(P)} \rho(q) = |\mathcal{E}_I(\Upsilon(P))| + 2 \cdot |\mathcal{E}_F(\Upsilon(P))| = |\text{Bd}(P)| + 2 \cdot |\mathcal{E}_F(\Upsilon(P))|.$$

Euler's formula for finite planar graphs⁴ tells us that $V - E + F = 1$. We can apply Euler's formula to $\Upsilon(P)$ by ignoring the infinite edges and infinite Voronoi polygons. Doing this yields $|\nu(P)| - |\mathcal{E}_F(\Upsilon(P))| + |\text{Int}(P)| = 1$, and so $|\mathcal{E}_F(\Upsilon(P))| = |\nu(P)| + |\text{Int}(P)| - 1$. Thus

$$\sum_{q \in \nu(P)} \rho(q) = |\text{Bd}(P)| + 2 \cdot (|\nu(P)| + |\text{Int}(P)| - 1).$$

Recalling from §1.3 that $|P| = |\text{Bd}(P)| + |\text{Int}(P)|$, we have

$$\sum_{q \in \nu(P)} \rho(q) = 2 \cdot |\nu(P)| + 2 \cdot |P| - |\text{Bd}(P)| - 2.$$

⁴See [5, p. 5].

On the other hand, by Proposition 2.2,

$$\sum_{q \in \nu(P)} \rho(q) = 3 \cdot |\nu(P)| + I_c(P).$$

Thus

$$2 \cdot |\nu(P)| + 2 \cdot |P| - |\text{Bd}(P)| - 2 = 3 \cdot |\nu(P)| + I_c(P),$$

and so, solving for $|\nu(P)|$, we find

$$|\nu(P)| = 2 \cdot |P| - |\text{Bd}(P)| - I_c(P) - 2,$$

as desired. ■

From this and our preliminary analysis in §1, several interesting observations may be made:

Corollary 2.4.

- (i) If $|P| < 5$ then $\nu^n(P) = \emptyset$ for some $n \in \mathbb{N}$.
- (ii) If $|P| = 5$ then either $\nu^n(P) = \emptyset$ for some $n \in \mathbb{N}$ or $|\nu^n(P)| = 5$ for all $n \in \mathbb{N}$.
- (iii) If $|\nu(P)| > |P|$ then $|P| > 5$.
- (iv) If $|\text{Int}(P)| = 2$ and $I_c(P) = 0$, then $|\nu(P)| = |P|$.

Part (ii) of this corollary suggests that the case when $|P| = 5$ may be of special interest, not only because it marks a transition in the behavior of the system, but because it forms a natural subset of phase space that is closed under Voronoï iteration. We discuss this further in §3.

Theorem 2.3 alone is not enough to predict the long-term behavior of our system, since computing $|\nu^n(P)|$ requires that we can also compute $|\text{Bd}(\nu^{n-1}(P))|$ and $I_c(\nu^{n-1}(P))$. We have been able to find bounds on the size of $\nu^n(P)$, and we present one of interest next.

2.3. Bounds on the Size of $\nu^n(P)$. With no assumptions on the geometry of $\nu^n(P)$ we can give an upper bound on the size of $\nu^n(P)$, which we conjecture to be sharp.

Proposition 2.5. *For all $n > 0$, we have $|\nu^n(P)| \leq 2^n (|P| - 5) + 5$.*

Proof. By induction on n . For $n = 1$, plugging the inequalities $|\text{Bd}(P)| \geq 3$ and $I_c(P) \geq 0$ into Theorem 2.3 immediately gives us what we want. Proceeding inductively, we get

$$\begin{aligned} |\nu^n(P)| &= 2 \cdot |\nu^{n-1}(P)| - |\text{Bd}(\nu^{n-1}(P))| - I_c(\nu^{n-1}(P)) - 2 \leq 2 \cdot |\nu^{n-1}(P)| - 5 \\ &\leq 2 \cdot (2^{n-1} \cdot (|P| - 5) + 5) - 5 = 2^n \cdot (|P| - 5) + 5, \end{aligned}$$

which completes the proof. ■

We have also been able to find an upper bound on the size of the boundary of $\nu(P)$.

Theorem 2.6. *The number of points on the boundary of $\nu(P)$ does not exceed $|\text{Bd}(P)|$.*

We shall need the following lemma:

Lemma 2.7. *The interior angle for any vertex in a Voronoï polygon is strictly less than π .*

Proof. Let $p \in P$ generate the Voronoï polygon $V(p) \in \Upsilon(P)$. Any vertex of $V(p)$ lies at the intersection of two Voronoï edges $e_{p,p'}$ and $e_{p,p''}$, generated by distinct $p', p'' \in P$. From the argument given in §1.3, we deduce that $V(p)$ is convex, and hence immediately have that the interior angle made by these edges must be less than or equal to π . Seeking a contradiction, assume that this angle is equal to π . By the perpendicular bisector property from §1.2, we may obtain p' by reflecting p across $e_{p,p'}$ and p'' by reflecting p across $e_{p,p''}$. But then $p' = p''$, contradicting the assumption that p' and p'' are distinct. ■

Proof of Theorem 2.6. We have two cases: when all $q \in \nu(P)$ are collinear and when they are not. For the former, since any Voronoï tessellation is connected we must have each vertex joined to its neighbor(s). Since each vertex must at least be of degree three, there must be at least $|\nu(P)| + 2$ infinite edges, implying that $|\text{Bd}(\nu(P))| = |\nu(P)| < |\nu(P)| + 2 \leq |\mathcal{E}_I(\Upsilon(P))| = |\text{Bd}(P)|$, as needed.

To prove the result for noncollinear vertex sets, we seek to show $|\text{Bd}(\nu(P))| \leq |\mathcal{E}_I(\Upsilon(P))|$ by proving that every $q \in \text{Bd}(\nu(P))$ intersects at least one infinite Voronoï edge. This, with the equality $|\mathcal{E}_I(\Upsilon(P))| = |\text{Bd}(P)|$ from Proposition 1.6, gives us what we want. Since $q \in \text{Bd}(\nu(P))$, we may find neighbors $q', q'' \in \text{Bd}(\nu(P))$ of q on the boundary of the convex hull of $\nu(P)$. Define $H(q, q')$ to be the closed half-plane containing $\nu(P)$ that is bounded by the line through q and q' , and similarly define $H(q, q'')$; these half-planes exist by the properties of the convex hull. Thus $\nu(P)$ is contained in the intersection $H(q, q') \cap H(q, q'')$, which we denote H .

Find a point $p \in P$ such that $q \in V(p)$ and $V(p)$ has a nontrivial intersection with the complement of H , and let $e \in \mathcal{E}(\Upsilon(P))$ be an edge of $V(p)$ that touches q . If e is infinite, then we are done. If e terminates in a vertex, then by the convexity of H we have that e is contained in H . Let $e' \in \mathcal{E}(\Upsilon(P))$ be the other edge of $V(p)$ that touches q . By Lemma 2.7, the interior angle made at q must be strictly less than π . Since $V(p)$ must have a nontrivial intersection with the complement of H , we find that e' cannot be contained in H since the convexity of H would force the interior angle made at q to be greater than or equal to π . Thus it cannot terminate in a vertex, and instead is infinite, as desired. ■

Note that the size of the boundary does not steadily decrease to three in all cases. There seem to be configurations with boundaries whose sizes stay stable over many iterations.

Theorems 2.3 and 2.6 combine to give a lower bound on the size of $\nu^n(P)$ in the generic situation where there are no instances of cocircularity.

Theorem 2.8. *Let $P \in \mathcal{P}(\mathbb{R}^2)$ and $N \in \mathbb{N}$. Suppose that for all $n = 0, 1, 2, \dots, N - 1$, $I_c(\nu^n(P)) = 0$. Then $|\nu^N(P)| \geq 2^N |P| - (2^N - 1)(|\text{Bd}(P)| + 2)$.*

A point set P with $|\text{Int}(P)| > 2$ and with Voronoï iterations that never contain instances of cocircularity will have exponential growth on the order of 2^n . (If $|\text{Int}(P)| \leq 2$ then the size does not increase). Since having instances of cocircularity is a rare occurrence, we conjecture that such a point set exists. We show a typical situation in Figure 6: a few iterates of a randomly-selected point set of size 9. As Theorem 2.3

predicts, we see the size going to 13, then 21 points; if there continues not to be any cocircularity then the growth will escalate rapidly: the next four iterations contain 37, 69, 133, and then 261 points, respectively.

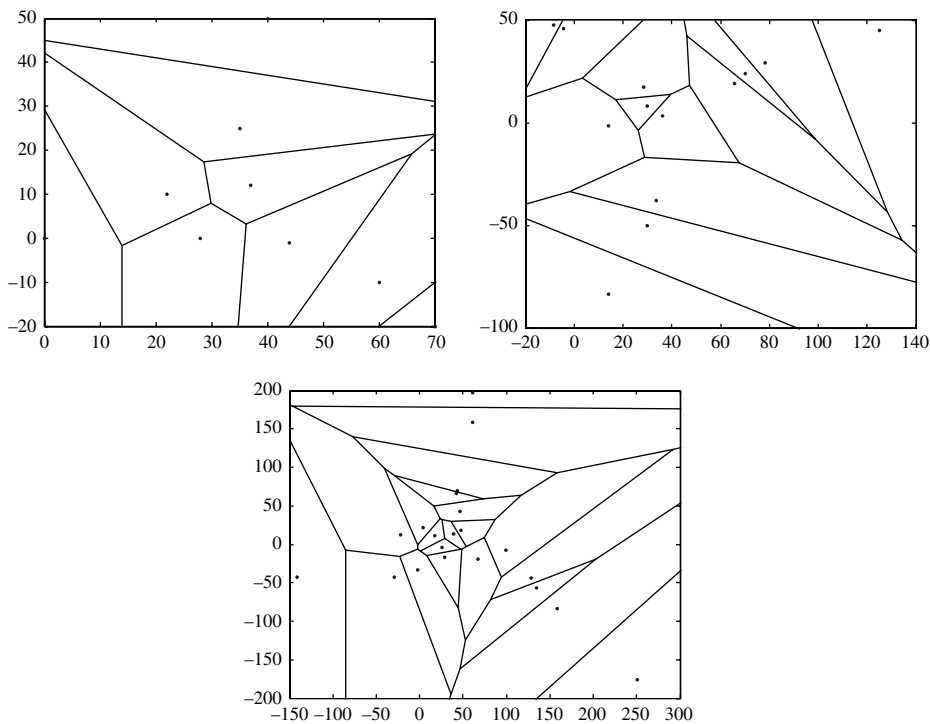


Figure 6. A few iterations of a point set with no instances of cocircularity.

3. SO WHAT'S NEXT? QUESTIONS FROM THE DYNAMICAL SYSTEMS VIEWPOINT. In dynamical systems terminology, $\mathcal{P}(\mathbb{R}^2)$ is called the **phase space** of the system, and elements of $\mathcal{P}(\mathbb{R}^2)$ are **states** of the system that evolve over time according to the map ν . The **orbit** or **trajectory** of a state P is defined to be the sequence $\{P, \nu(P), \nu^2(P), \dots\}$ ([1] and [6] are good references). So far we have only talked about one property of an orbit: the growth rate of $|\nu^n(P)|$ over time. But there are many other interesting questions we can ask about this dynamical system.

3.1. What Is the Right Topology? In order to use the machinery of modern dynamics, we ought to have a topology on $\mathcal{P}(\mathbb{R}^2)$ making Voronoi iteration continuous, at least some of the time. We need this to study major dynamical features such as *recurrence*—how orbits return to open sets over time—or topological *entropy*—a measure of the tendency of the system to become disordered. Finding the right topology is of great importance, but it has proved to be an interesting problem in its own right. We will mention some of the subtlety here.

To begin, notice that the scale of the diagrams in Figure 6 increases with each iteration. We want our topology to care about the shape of the Voronoi polygons rather than their size. Similar triangles produce similar Voronoi diagrams, so the topology ought to respect that. Let us be more precise.

Definition 3.1. A *similarity transformation* $t \in \text{Sim}(2)$ is a map of the form

$$t(\mathbf{x}) = kU\mathbf{x} + \mathbf{x}_0,$$

where $\mathbf{x}, \mathbf{x}_0 \in \mathbb{R}^2$, U is a 2×2 orthogonal matrix, and $k \in \mathbb{R}^+$. For $P, Q \subseteq \mathcal{P}(\mathbb{R}^2)$, we say P is **similar** to Q , written $P \simeq Q$, if there exists $t \in \text{Sim}(2)$ such that $t(P) = Q$.

A similarity transformation is a composition of translations, rotations, reflections, or dilatations of Euclidean space. One can show that $\text{Sim}(2)$ forms a group under the composition of functions and that the relation \simeq is an equivalence relation. Similarity transformations play nice with our dynamical system:

Theorem 3.2. For $t \in \text{Sim}(2)$, we have $t(v(P)) = v(t(P))$.

Proof. Let $q \in v(P)$. By Proposition 1.3, there is a unique empty circle C_q centered at q . Since similarity transformations preserve circles, t maps C_q to an empty circle centered at $t(q)$. Thus $t(q) \in v(t(P))$, and so $t(v(P)) \subset v(t(P))$. By similar logic, since $t(P)$ is a point set and t^{-1} is a similarity transformation, we have

$$t^{-1}(v(t(P))) \subset v(t^{-1}(t(P))) = v(P).$$

Thus $v(t(P)) \subset t(v(P))$, and so we have equality. ■

The right topology should identify similar point sets, so we aren't really looking at $\mathcal{P}(\mathbb{R}^2)$ but rather the quotient of $\mathcal{P}(\mathbb{R}^2)$ under similarity.

On the other hand, forgetting the question of similarity, we do have an intuitive idea of what it means for point sets P and Q in $\mathcal{P}(\mathbb{R}^2)$ to be "close." In fact there are already metrics, for instance the Hausdorff metric, for this. The problem is that we need our metric to respect Voronoï iteration.

Suppose we draw little ϵ -balls around all the points of P , and discover that each ball contains exactly one point of Q . If so, it is very likely that $v(P)$ and $v(Q)$ are close as well. (The obvious exception is with point sets that have nontrivial cocircularity—jiggling the points a little bit destroys the cocircularity and thus produces extra vertices in the Voronoï iteration.) This idea for a metric is promising but has a serious flaw: it can't compare sets that aren't the same cardinality. Since there are sets of the same cardinality that iterate to sets of different cardinalities, it would be nice to have the ability to consider whether those iterates are close.

Unfortunately, trying to measure distances between sets of different cardinalities opens Pandora's box. If several points of Q are inside the ϵ -ball around a point of P , we can create all sorts of bizarre patterns in $v(Q)$ that may be quite dissimilar to $v(P)$. We find ourselves unsure how to resolve these issues.

3.2. Are There Periodic Points? For $k \in \mathbb{N}$, we say P is a **period- k point** if $P \simeq v^k(P)$. As a special case, if $k = 1$ then $P \simeq v(P)$ and we call P a **fixed point**. The orbit of a periodic point repeats itself over and over again, and can be thought of as finite (modulo similarity). Periodic orbits are of substantial interest in dynamical systems theory: if periodic orbits are dense in phase space it is an indication of chaos.

Despite our attempts so far, we have not been able to find any finite periodic point set P . Curiously, it's not difficult to find an infinite set P of period 1 or 2! In Figure 7(a), the square lattice iterates to a shifted copy of itself and so has period 1; in Figure 7(b), points evenly spaced on the diagonals $y = \pm x$ iterate to points on the x - and y -axes, which then iterate back for a period of 2. In both examples one sees a high degree of cocircularity and we assume that this may be a key component in finding examples with other periods as well.

Once a fixed (or periodic) point has been located, we may ask whether it is attracting, repelling, or hyperbolic. We look at "nearby" points (again the need for a




Figure 7. Infinite configurations which are (a) period-1 and (b) period-2.

well-defined notion of distance!) and look at what their orbits do. If all nearby points come closer and closer to P , it is attracting; if they all get pushed away we call it repelling. If there is a mixture it may be hyperbolic. There's some indication that the grid in Figure 7(a) may be hyperbolic: we know it repels grids with “defects” like the one in Figure 8. However, if we shift an entire row of points up by a fixed amount, the orbit will be pulled back towards the orbit in 7(a).




Figure 8. Introducing a defect into the square lattice of 7(a).

3.3. Sensitive Dependence on Initial Conditions. This is the so-called “butterfly effect,” where a small change in an initial state can produce a large change in its orbit. Even though we lack a metric on our phase space $\mathcal{P}(\mathbb{R}^2)$, there is evidence of sensitive dependence on initial conditions. For instance, if we shift a point on the grid in Figure 7(a), it introduces a defect in the Voronoï iteration, which we see in Figure 8. As the system evolves, that defect will grow to include ever larger regions of the plane until the original grid is no longer recognizable.

One can see this effect in finite configurations as well. Consider the set depicted in Figure 4(b). Since $|\nu(P)| = 4$ we know it is doomed to iterate to the empty set. However, if we were to move any one of the four cocircular points even the slightest amount, we would see the single vertex of degree four become two vertices of degree three. This five-point configuration has a chance of iterating to five-point sets indefinitely.

3.4. Can We Go Backwards? A sensible question to ask is, what happens when we try to invert the system? Most discrete point sets in the plane are not the vertex set of a Voronoï diagram. So we must ask first when a **preimage** of a point set exists, i.e., given a set Q , does there exist a set P such that $v(P) = Q$? If there is a preimage, when is it unique? When it is not unique, how many preimages are possible, and is there a “best” one?

Some literature exists about inverting Voronoï tessellations (see, e.g., [2, 8]), and we know relatively straightforward conditions that must be satisfied for a preimage to exist. What is more subtle in our situation is that we don’t have the whole Voronoï diagram—we just have its vertices. Without the edge adjacencies the problem becomes richer. For example, in Figure 9 we depict two preimages of the lighter colored set of six points, with distinctly different Voronoï diagrams.




Figure 9. A configuration of six points with two nontrivially different preimages.

It would be nice if there was a subset of $\mathcal{P}(\mathbb{R}^2)$ that was invariant under Voronoï iteration both forward and backwards in time. We conjecture that there is a subset $\mathcal{S} \subset \mathcal{P}(\mathbb{R}^2)$ of five-point configurations with 2 in the interior and 3 on the boundary for which (1) the set is invariant under forward Voronoï iteration, and (2) there exist preimages of all orders within the set. If we could identify criteria for membership in \mathcal{S} , then we could begin to examine the orbits of \mathcal{S} for evidence of chaos.

3.5. Generalizations. In §3.2 we mentioned generalizing the Voronoï iteration to infinite point sets. However, this allows for a host of strange phenomena and so we must restrict our point sets somewhat. For example, we might require that the **derived set**—the set of all accumulation points—of P be empty. Or we might assume that P is a **Delone set**; i.e., both relatively dense and uniformly discrete. Of course for a property to have any use to us, it must be preserved under iteration. One can prove that the derived set of $v(P)$ is empty whenever the derived set of P is empty, while the property of being Delone is actually not preserved under iteration. We conclude that a natural phase space consists of all generator sets with empty derived sets.

And what if we were to no longer work on the plane? The sphere with the great circle metric immediately comes to mind. All Voronoï polygons are bounded, and so infinite edges are no longer an issue; this gives rise to a more elegant version of Theorem 2.3. Finite periodic point sets are suddenly easy to find: each Platonic solid iterates to its dual. Further, one can prove that a set of noncoplanar points will never iterate to the empty set. And this is all besides the fact that we humans live on a sphere,

so maybe the system has relevance to real-world problems such as the distribution of cell-phone satellites or the spread of diseases.

The fact that one can define the Voronoi tessellation for a set of points in any metric space means that vast generalizations are possible, whether it be to infinite sets, a higher dimension, a new metric, or a less geometric setting. Part of the appeal of this dynamical system is its intuitive definition—a definition that depends little on the size of the point sets or the space they are in. We urge interested readers to think about how the ideas in this paper play out in their favorite metric spaces!

3.6. Conclusion. We admit that we don't know much about the behavior of this dynamical system yet. We know exactly what happens for small point sets, but as soon as there are five or more points we encounter difficulties. We have a theorem [4] that measures how the point sets grow or shrink in size, but it requires geometric information that isn't always readily available. We have evidence that many point sets grow without bound, but have been unable to determine which conditions guarantee this. We've noticed evidence of sensitive dependence on initial conditions, but we don't have the machinery to measure the phenomenon. In the course of our study we've developed tools to help us with our insight on this finicky dynamical system, and we've shared some of that here. There is one thing we know for certain: plenty of problems remain. Some are simple enough to be studied by budding mathematicians, and some may be subtle enough to interest their advisors as well.

REFERENCES

1. K. T. Alligood, T. D. Sauer, and J. A. Yorke, *Chaos: An Introduction to Dynamical Systems*, Springer-Verlag, New York, 1997.
2. P. F. Ash and E. D. Bolker, Recognizing Dirichlet tessellations, *Geom. Dedicata* **19** (1985) 175–206. doi : 10.1007/BF00181470
3. Q. Du, M. Emelianenko, and L. Ju, Convergence of the Lloyd algorithm for computing centroidal Voronoi tessellations, *SIAM J. Numer. Anal.* **44** (2006) 102–119. doi : 10.1137/040617364
4. A. Edgren, O. Gillham, and S. M. Hart, On a dynamical system based on Voronoi tessellations, 2005 URSI Undergraduate Project Report (advisor: Natalie Priebe Frank), available at <http://math.vassar.edu/Faculty/Frank/NPFpublications.html>.
5. M. Henle, *A Combinatorial Introduction to Topology*, Dover, New York, 1994.
6. A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*, Encyclopedia of Mathematics and its Applications, vol. 54, Cambridge University Press, Cambridge, 1995.
7. A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd ed., John Wiley, Chichester, UK, 2000.
8. F. P. Schoenberg, T. Ferguson, and C. Li, Inverting Dirichlet tessellations, *Comput. J.* **46** (2003) 76–83. doi : 10.1093/comjnl/46.1.76

NATALIE PRIEBE FRANK received her B.S. in mathematics from Tulane University in 1991. She thought up the Voronoi dynamical system while working on her Ph.D., which she received from the University of North Carolina at Chapel Hill in 1997. She currently teaches at Vassar College, where she was delighted to recruit her coauthor for mathematical duty after his freshman year.

Department of Mathematics, Vassar College, Poughkeepsie, NY 12604
nafrank@vassar.edu

SEAN M. HART received his B.A. in mathematics from Vassar College in 2008. He is currently working on his Ph.D. at the University of Texas at Austin with the hope that one day he, too, will be struck with a great idea. He will always be grateful to his coauthor and the rest of the faculty at Vassar for nurturing his mathematical abilities.

Department of Mathematics, The University of Texas at Austin, 1 University Station, C1200,
Austin, Texas 78712
shart@math.utexas.edu