The density of unit distance-avoiding sets

Talk by Nitya Mani

Notes by Sanjana Das

September 8, 2023

§1 Introduction

Today we'll talk about the maximum density of a unit distance-avoiding subset of the plane. There has been a fairly recent breakthrough on this. Today, we'll mostly discuss the methods and results of a slightly simpler paper that comes earlier; this gets a slightly worse bound, but highlights the main ideas.

> **Question 1.1.** How dense can a subset of $\mathbb{R}^d$ be given that it doesn't have any pair of points that are a distance (exactly) 1 apart?

A related question — which we won't spend too much time discussing (but which is also important) — is the Nelson–Hadinger problem about the chromatic number of $\mathbb{R}^d$.

> **Question 1.2.** What is $\chi(\mathbb{R}^d)$, the chromatic number of the *unit distance graph* on $\mathbb{R}^d$ (the infinite graph with vertex set $\mathbb{R}^d$, where we connect points a unit distance apart)?

These are two motivating questions of this area of work.

We'll first make the notion of a unit distance graph more precise.

> **Definition 1.3.** A graph $G = (V, E)$ is a *unit distance graph* if $V \subseteq \mathbb{R}^d$ and two points are adjacent if and only if they are a unit distance apart — i.e., $E = \{(x, y) \in V^2 \mid \|x - y\| = 1\}$.

The unit distance graphs we'll see will usually be finite (i.e., we take a *finite* set of points and connect all pairs a unit distance apart) — for example, an equilateral triangle is a unit distance graph.

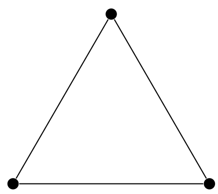

> **Definition 1.4.** A set $A \subseteq \mathbb{R}^d$ is *unit distance-free* (or 1-*avoiding*) if there do not exist any two points $x, y \in A$ with $\|x - y\| = 1$.

We always work with the L^2-distance. We'll assume that notions of density for A are well-defined — in particular, we'll assume A is measurable and periodic with respect to some lattice $\mathcal{L}$. We can do this without loss of generality — this is different from the chromatic number problem, where assuming measurability changes the answer.

We can now formulate our question more precisely.

Question 1.5. Let $m_1(\mathbb{R}^2) = \sup_A \delta(A)$ over all unit distance-free $A \subseteq \mathbb{R}^2$ (where $\delta(A)$ denotes the density of A). What can we say about $m_1(\mathbb{R}^2)$?

We'll mostly focus on trying to obtain the best *upper bound* on $m_1(\mathbb{R}^2)$ that we can.

§2 Lower bounds

Before we talk about upper bounds, we'll first give a *lower* bound on $m_1(\mathbb{R}^2)$ — this means we want to construct a reasonably dense unit distance-avoiding set A.

One construction is to tile the plane using an equilateral triangle lattice with side length 2, and place an open ball of radius $\frac{1}{2}$ at each vertex.

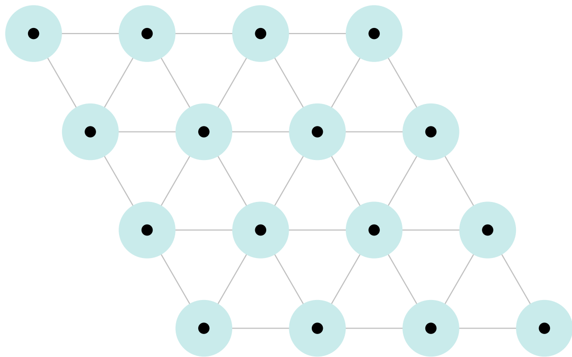

Then two points in the same ball have distance less than 1, and two points in different balls have distance greater than 1, so this gives a valid set; and the density of this set is

$$\delta(A) = \frac{\pi}{8\sqrt{3}} \approx 0.227.$$

We can actually improve this a bit. The best-known lower bound is the *Croft tortoise construction* (the name is because it looks a bit like a tortoise). The premise is similar — we again take a triangular lattice — but instead of having the distance between points on the lattice be 2, we have it be $1 + x$ for some x slightly less than 1. Then at each vertex, we put the intersection of a circle of radius $\frac{1}{2}$ and a hexagon of height x (where the *height* refers to the distance between opposite sides).

The density is maximized when $x \approx 0.965$, in which case we get density $\delta(A) \approx 0.229$.

§3 An idea for upper bounds

We'll now turn to upper bounds.

We'll first make a simple observation — fix a *specific* unit vector u, and suppose that A is a unit distance-avoiding set (throughout this talk, we'll always use A to denote a unit distance-avoiding set). Then we must have $A \cap (A+u) = \emptyset$ — if there were some point x in both A and $A+u$, then there would be a point $y \in A$ at a distance 1 from x (namely, $x-u$). This is a key observation that goes into most of the upper bounds — if we shift A by any unit distance, the resulting set won't intersect with A.

This immediately gives an upper bound of $m_1(\mathbb{R}^2) \leq \frac{1}{2}$ (by considering one specific u).

We can do better than this — suppose now that we have two unit vectors u and v forming an equilateral triangle (i.e., two unit vectors 60° apart). Then the sets A, $A+u$, and $A+v$ must *all* be disjoint; this gives a better upper bound of $m_1(\mathbb{R}^2) \leq \frac{1}{3}$.

We can generalize this a bit more — let G be *any* unit distance graph on vertices $u_1, \ldots, u_n$, and let $\alpha(G)$ denote its independence number. Then we look at the translates $A+u_1, A+u_2, \ldots, A+u_n$. Consider some (arbitrary) point $x \in \mathbb{R}^2$; then x can fall into at most $\alpha(G)$ of these translates — otherwise there would have to be some $(u_i, u_j) \in E(G)$ for which x is in both $A+u_i$ and $A+u_j$, and since $u_i \sim u_j$ we'd get two points in A a unit distance apart (namely $x-u_i$ and $x-u_j$). This tells us

$$m_1(\mathbb{R}^2) \leq \frac{\alpha(G)}{v(G)}$$

for any unit-distance graph, which gives us a new way of getting upper bounds. For example, one unit-distance graph with a small value of $\alpha(G)/v(G)$ is the *Moser spindle*, which has $\alpha(G)/v(G) = 2/7$; this gives $m_1(\mathbb{R}^2) \leq 2/7$.

In fact, we can do a bit better than this.

Definition 3.1. The *fractional chromatic number* of a graph G, denoted $\chi_f(G)$, is the minimum value of $\sum_{I \in \mathcal{I}(G)} w_I$ — where $\mathcal{I}(G)$ denotes the set of all independent sets of G — over all vectors w such that for all vertices v we have $\sum_{I \ni v} w_I \geq 1$.

Then using a similar argument, we can get $m_1(\mathbb{R}^2) \leq 1/\chi_f(G)$.

Remark 3.2. As a sidenote, we'll briefly talk about the chromatic number of the plane — we know that

$$4 \leq \chi(\mathbb{R}^2) \leq 7.$$

The lower bound comes from the Moser spindle, and the upper bound from a periodic coloring of $\mathbb{R}^2$. But if we require the color classes to be Lebesgue measurable, then the problem changes — we define $\chi_m(\mathbb{R}^2)$ as the minimum number of colors needed in this case (so that $\chi_m(\mathbb{R}^2) \geq \chi(\mathbb{R}^2)$), and in fact, we know that $\chi_m(\mathbb{R}^2) \geq 5$. We know that $m_1(\mathbb{R}^d)\chi_m(\mathbb{R}^d) \geq 1$, but we don't know if this is tight.

For a while this was where things were; then over the last 10–15 years there's been a really fruitful line of work that uses these ideas as ingredients in linear programming-based methods to get upper bounds (which we'll now discuss).

§4 Proof of upper bounds

§4.1 The autocorrelation function

We'll work with an object called the *autocorrelation function*.

Definition 4.1. Given a unit distance-avoiding set A, the *autocorrelation function* $f\colon \mathbb{R}^2 \to \mathbb{R}$ (defined with respect to A) is defined as

$$f(x) = \delta(A \cap (A-x)).$$

First, there's a few simple properties of f.

- We have $f(0) = \delta(A)$.
- If $\|x\| = 1$, then $f(x) = 0$.

Importantly, f is a *positive definite kernel* (PDK). There's a few different ways to define a PDK — for example, as a positive function with positive Fourier coefficients — but we'll use a different definition.

Definition 4.2. A function $K(x, y)$ is a *positive definite kernel* (PDK) if for all $x_1, \ldots, x_n$, we have $(K(x_i, x_j))_{i,j\in[n]} \succeq 0$ (i.e., $\sum c_i c_j K(x_i, x_j) \geq 0$ for all nonnegative $c_1, \ldots, c_n$).

Here f only takes one input, but we define $K(x, y) = f(x - y)$; this function K is a PDK.

We'll show that such an autocorrelation function f must satisfy certain properties; we'll then turn them into constraints for a linear program.

§4.2 Summary of results

First, we'll discuss the results obtained using these methods. Before, the best upper bound (due to Székely 1984) was $m_1(\mathbb{R}^2) \leq 12/43 \approx 0.2791$, using the fractional chromatic number method. Then recently, there's been a series of improvements using the linear programming approach.

- OV 2010 proved that $m_1(\mathbb{R}^2) \leq 0.2684$.
- KMOR 2016 proved that $m_1(\mathbb{R}^2) \leq 0.2588$.
- AM 2022 proved that $m_1(\mathbb{R}^2) \leq 0.2544$.

Conjecture 4.3 (Erdős) — $m_1(\mathbb{R}^2) < 1/4$.

Remark 4.4. The quantity $\frac{1}{4}$ doesn't come from thin air — it's because if we use different norms, where the unit balls are not round (e.g., L^1 or L^∞ instead of L^2), then $\frac{1}{4}$ is often the right answer. For L^2 we don't know the exact answer, and Erdős conjectured that it's strictly lower than in these cases.

Remark 4.5. This conjecture would give a different proof of the fact that $\chi_m(\mathbb{R}^2) \geq 5$.

Recent work resolved this conjecture.

Theorem 4.6 (ACMVG 2022)
We have $m_1(\mathbb{R}^2) \leq 0.247$.

All these papers use a linear programming approach centered around the autocorrelation function.

§4.3 Properties of the autocorrelation function

First, we'll write down several properties of our autocorrelation function $f(x) = \delta(A \cap (A - x))$. (We'll prove these properties later.)

(D) We have $f(0) = \delta(A)$.

(C0) For all $\|x\| = 1$, we have $f(x) = 0$.

(C1) For any finite unit distance graph G, we have

$$\sum_{x\in V(G)} f(x) \le \alpha(G)f(0).$$

(C1R) For *any* finite graph G, we have

$$\sum_{x\in V(G)} f(x) - \sum_{xy\in E(G)} f(x-y) \le \alpha(G)f(0).$$

(C2) For any finite set of points $C \subseteq \mathbb{R}^2$, we have

$$\sum_{(x,y)\in\binom{C}{2}} f(x-y) \ge |C|\, f(0) - 1.$$

(T) For any finite unit distance graph $G = (V, E)$ with $\alpha(G) \le 3$, if we define

$$\Sigma_3(G) = \sum_{(x_1,x_2,x_3)\in\binom{V}{3}} (A - x_1) \cap (A - x_2) \cap (A - x_3),$$

then we have

$$\sum_{x\in V} f(x) - 2f(0) \le \Sigma_3.$$

(T2) Under the same setup, we have

$$\Sigma_3 \le 1 - v(G)f(0) + \sum_{(x,y)\in\binom{V}{2}} f(x-y).$$

For example, (C1) is the observation we saw earlier, rephrased in terms of f, and (C1R) is a generalization of (C1).

These properties are used in the intermediate papers; the final paper uses these constraints along with a few other similar ones, and also makes slightly different choices for what C and G to use. We'll talk about where these properties come from later; first we'll talk about how we use them to get upper bounds.

§4.4 The Fourier transform

First, how do we get to linear programming? We assumed A is periodic with respect to some lattice $\mathcal{L}$. So we can think of f as a function $\mathbb{R}^2 \to \mathbb{C}$, and we can define an inner product on such functions where we average $f\overline{g}$ over bigger and bigger boxes — we define

$$\langle f, g\rangle = \lim_{T\to\infty} \frac{1}{(2T)^2} \int_{[-T,T]^2} f(x)\overline{g(x)}\, dx.$$

Now let $\mathcal{L}^*$ be the dual lattice

$$\mathcal{L}^* = \{u \in \mathbb{R}^2 \mid \langle u, v\rangle \in \mathbb{Z} \text{ for all } v \in \mathcal{L}\}.$$

For all $u \in 2\pi\mathcal{L}^*$, we define the character

$$\chi_u(x) = e^{i\langle u,x\rangle}.$$

These characters form an orthonormal system for $L^2(\mathbb{R}^2/\mathcal{L})$, which means we have a Fourier transform

$$\widehat{f}(u) = \langle f, \chi_u\rangle$$

which satisfies the usual Fourier inversion formula

$$f(x) = \sum_{u\in 2\pi\mathcal{L}^*} \widehat{f}(u)\chi_u(x) = \sum_{u\in 2\pi\mathcal{L}^*} \widehat{f}(u)e^{i\langle u,x\rangle}.$$

> **Remark 4.7.** When we said earlier that PDKs have nonnegative Fourier coefficients, these are the Fourier coefficients we meant.

We can rewrite our definition of f as

$$f(x) = \langle 1_A, 1_{A-x} \rangle.$$

By definition, we have

$$\widehat{1_{A-x}}(u) = \langle 1_A, e^{i\langle u,x \rangle} \rangle.$$

Then using Parseval, we can obtain that

$$f(x) = \sum_{u \in 2\pi \mathcal{L}^*} |\widehat{1_A}(u)|^2 e^{i\langle u,x \rangle},$$

which means these squares are the Fourier coefficients of f — i.e., $\widehat{f}(u) = |\widehat{1_A}(u)|^2$.

Next, we'll do some symmetrization. If we have some graph (on a set of points in $\mathbb{R}^2$), we can imagine rotating it; we don't want to write down a separate constraint for each rotation, so we average over all rotations — we define a function $f_\circ$ as

$$f_\circ(x) = \frac{1}{2\pi} \int_{\zeta \in S^1} f(\zeta \|x\|)$$

(so $f_\circ$ only depends on $\|x\|$ — essentially, we're averaging over all rotations). Then we can write

$$f_\circ(x) = \sum_{u \in 2\pi \mathcal{L}^*} \widehat{f}(u) J_0(\|u\| \, \|x\|),$$

where J_0 is a Bessel function. Since $J_0(\|u\| \, \|x\|)$ only depends on $\|u\|$ (and not other information about u), we can group together all u with the same magnitude to write

$$f_\circ(x) = \sum_{t \geq 0} \kappa(t) J_0(t \|x\|), \text{ where } \kappa(t) = \sum_{\|u\| = t} \widehat{f}(u).$$

(The sum over t makes sense because the values t of the form $\|u\|$ for some $u \in 2\pi\mathcal{L}^*$ are discrete.)

The important thing here is that $f(0) = f_\circ(0)$ is a linear combination of these $\kappa(t)$'s. We had a bunch of constraints that f was supposed to satisfy, and we can rewrite them as constraints on $\kappa(t)$ that hopefully look like a linear program.

§4.5 An infinite linear program

We'll first consider an infinite linear program in the variables $\kappa(t)$ (for all $t \geq 0$) and $\delta(A)$, where we want to find $\max \delta(A)$. We'll have the following conditions.

- (B) $\kappa(0) = \delta(A)^2$ — this is because $\kappa(0) = \widehat{f}(0) = |\widehat{1_A}(0)|^2 = \delta(A)^2$.
- (CP) $\kappa(t) \geq 0$ for all t — this is because $\kappa(t)$ is a sum of squares (since $\widehat{f}(u) = |\widehat{1_A}(u)|^2$).
- (CS) $\delta(A) = \sum_{t \geq 0} \kappa(t)$ — this is because $\delta(A) = f_\circ(0) = \sum_{u \in 2\pi\mathcal{L}^*} \kappa(t) J_0(0)$.
- (C0) $\sum \kappa(t) J_0(t) = 0$ — this corresponds to the condition that A is unit distance-free (since the left-hand side is precisely $f_\circ(1)$), which now corresponds to just one constraint rather than infinitely many.
- (C1R) For all finite graphs G, we have

$$\sum_{t \geq 0} \kappa(t) \left(\sum_{x \in V(G)} J_0(t \|x\|) - \sum_{xy \in E(G)} J_0(t \|x - y\|) \right) \leq \alpha(G) \delta(A).$$

  (This corresponds to our condition (C1R) on f from earlier.)

$(\cdots)$ We can translate the remaining conditions from earlier similarly.

So we now have an infinitely large linear program, and if we could solve it, then we'd get an upper bound on $m_1(\mathbb{R}^2)$ — because any actual function f must satisfy all these constraints. (If we solve the linear program, we'll get some $\kappa(t)$'s, but there's no guarantee that they'll actually correspond to some autocorrelation function; but we can get an upper bound this way.)

> **Remark 4.8.** Some of these constraints correspond to the fact that A is unit distance-avoiding. But some don't — some of these constraints are things that have to be true of *any* autocorrelation function. (In some sense, they're trying to regularize the solution to the linear program to look more like an autocorrelation function.) Adding these extra constraints is one of the things that made improvements possible.

> **Remark 4.9.** We're kind of cheating here because the first constraint has $\delta(A)^2$, which is quadratic. But this can be dealt with by dividing by $\delta(A)$ in the appropriate places.

§4.6 Discretizing the linear program

Unfortunately, this is an infinitely large linear program, so we can't solve it. But instead, we can try to discretize — rather than looking at the infinitely many possible values of t, we'll look at t which vary by ε and go up to some big value, such as ε^{-1}. However, now we have the problem that the solution to the linear program might not actually be an upper bound — we might get something too small (since now our LP is over some discretized set of variables instead — this means we're forcing all the $\kappa(t)$'s where t isn't a multiple of ε to be zero, which doesn't have to be true of the actual f).

So how do we get something that's genuinely an upper bound? The LP we're trying to solve is a maximization problem, so we're guaranteed that this maximum is actually an upper bound on the true solution, but we no longer have this guarantee when we discretize. So we'll instead consider the *dual* of our LP (which is a minimization problem). Then *any* feasible solution to the dual is an upper bound on the original. So we can discretize the dual and solve this (finite) LP, and that'll give us an upper bound for the (undiscretized) original.

> **Remark 4.10.** These ideas have been known several years before the final paper; this paper comes up with the G and C using a combination of computer search and cleverness. They use a finite list of such G and C; the KMOR paper uses a 30-point set and 5 fairly small graphs.
>
> How do people come up with such point sets and graphs? One way is by a grid search where we iterate over a bunch of possibilities, plug them in, and look at the solutions we get. Then we consider a whole host of other graphs and see which of them our candidate solution violates the worst, and then we add that graph in and try again; and we iterate.
>
> For the 3-point bounds (T1) and (T2), the graphs used are simple but surprising. The paper engineers two different graphs G_1 and G_2 with the same Σ_3, and then plug one of these graphs into (T1) and the other into (T2); the argument is very clever.

§4.7 The constraints

We'll now explain why the constraints (the properties of f described in Subsection 4.3) are true.

First, (D) and (C0) are clear. Next, (C1) essentially follows from what we discussed earlier.

Proof of (C1). If G is a unit distance graph, then for any $a \in A$ we have $x + a \in A$ for at most $\alpha(G)$ values of $x \in V(G)$ (since these points x must form an independent set). This means

$$\sum_{x \in V(G)} f(x) = \sum_{x \in V(G)} \delta(A \cap (A - x)) \leq \alpha(G) f(0) = \alpha(G)\delta(A). \qquad \square$$

We'll now look at (C1R), the generalization of (C1) where G isn't necessarily a unit distance graph.

Proof of (C1R). Let $G = (V, E)$. For each $a \in A$, define

$$g(a) = \#\{x \in V \mid a \in A - x\} - \#\{xy \in E \mid a \in (A - x) \cap (A - y)\}.$$

This essentially counts the number of times our point a is covered by the translate of a vertex x, minus the number of times it's covered by a translate corresponding to an edge. It suffices to show that $g(a) \leq \alpha(G)$ for all a; then averaging over all a gives the desired result. To do so, let

$$V' = \{x \in V \mid a \in A - x\} \text{ and } E' = \{xy \in E \mid a \in (A - x) \cap (A - y)\},$$

and consider the induced subgraph $G' = G[V']$ (which has edge set E'). First, G' must have at most $\alpha(G)$ connected components (or else we could find an independent set in G', and therefore G, of size greater than $\alpha(G)$); call these components $G_1, \ldots, G_\ell$, with $\ell \leq \alpha(G)$. This then directly gives us a bound on $|E'|$ — if G_i has v_i vertices and $e_i \geq v_i - 1$ edges, then

$$|E'| = e(G') = e_1 + \cdots + e_\ell \geq (v_1 - 1) + \cdots + (v_\ell - 1) = |V'| - \ell \geq |V'| - \alpha(G),$$

which gives $g(a) \leq \alpha(G)$, as desired. $\square$

Remark 4.11. Note that we never used the fact that A is a unit distance-avoiding set here — this is just a property of the autocorrelation function of *any* set.

We'll now verify (C2), the property of point sets C. This follows from PIE, and is again a general property of autocorrelation functions that has nothing to do with unit distance avoidance. (Finding good point sets was done by grid search in most papers.)

Proof of (C2). By inclusion-exclusion, we have $1 \geq \delta(\bigcup_{x \in C}(A - x)) \geq \sum_{x \in C} \delta(A - x) - \sum_{(x,y) \in \binom{C}{2}} \delta((A - x) \cap (A - y)) = |C|\,\delta(A) - \sum_{(x,y) \in \binom{C}{2}} f(x - y)$, which is exactly what we wanted. $\square$

We can use a similar strategy to prove (T1).

Proof of (T1). Again by inclusion-exclusion, we have $1 \geq \delta(\bigcup_{x \in V}(A - x)) = \sum_x \delta(A - x) - \sum_{x,y} \delta((A - x) \cap (A - y)) + \sum_{x,y,z} \delta((A - x) \cap (A - y) \cap (A - z))$ — note that this time we have an equality, since we can't have any intersections of four or more sets $A - x$ due to the fact that $\alpha(G) \leq 3$ (if we did have $(A - w) \cap (A - x) \cap (A - y) \cap (A - z) \neq \emptyset$, then since $\alpha(G) \leq 3$ there must be an edge among w, x, y, and z, and this would give us a unit distance in A). This gives the desired bound. $\square$

Proof of (T2). First, we have

$$\Sigma_3(G) \geq \sum_{x_1, x_2, x_3} \delta(A \cap (A - x_1) \cap (A - x_2) \cap (A - x_3))$$

(we've just thrown in an extra intersection with A here, compared to the definition of Σ_3). Let $A_x = A \cap (A - x)$. Now since $\alpha(G) \leq 3$, we have that each $a \in A$ can show up in at most three sets A_x. Some will be in exactly three, and the rest will be in at most two, and we can use this to count — we have

$$\sum_{x \in V} f(x) = \sum_{x \in V} \delta(A_x) \leq 2\delta(A) + \sum_{x_1, x_2, x_3} \delta(A \cap (A - x_1) \cap (A - x_2) \cap (A - x_3)). \qquad \square$$

§4.8 The graphs for (T1) and (T2)

As mentioned earlier, (T1) and (T2) are applied cleverly using two different graphs G_1 and G_2 with the same Σ_3 (allowing us combine the bounds). We'll get a bound of

$$\sum_{x \in V(G_2)} f(x) \leq 1 - 5f(0) + \sum_{x,y \in V(G_1)} f(x-y).$$

To construct this graph, we consider the following 8-point configuration (with edges denoting unit distances):

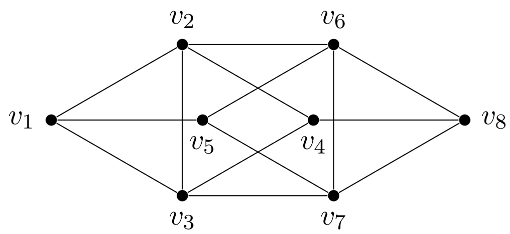

We then take G_1 and G_2 to both be graphs on 7 of these 8 vertices — G_1 consists of all the vertices except v_8, and G_2 all except v_5.

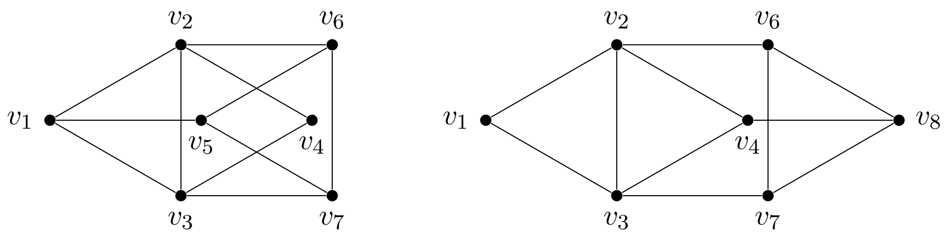

Then both graphs have the same Σ_3 (we can use the fact that A is unit distance-avoiding to cancel some of the terms), and have independence number 3.

§4.9 A final remark

There's an interesting phenomenon that most of our good constructions (e.g., the Croft construction) are kind of clumpy — we have some blobs in a bunch of places, and we might have a lot of distance-2 pairs. So we can imagine plotting a graph of $f_\circ(x)$ (as a function of $\|x\|$); we'll have a peak at $x = 0$ (where this value is $\delta(A)$), a trough of 0 at $x = 1$ (since A is unit distance-avoiding), a smaller peak at 2, a smaller trough at 3, and so on.

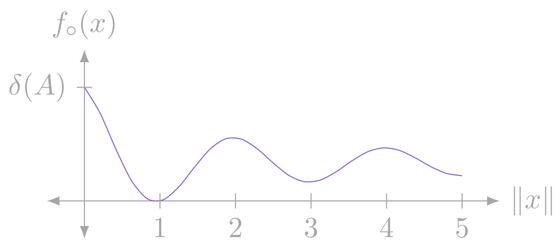

When we solve the LP, we can consider what the function $f_\circ$ we get out of it looks like. This function won't have *all* the properties that the actual function should have, but one way the authors get better bounds is by plotting this function, asking what 'looks wrong' about it, and trying to write some more constraints to fix these things.