Equilateral triangles in vector spaces over finite fields

Zachary Polansky

May 2019

Abstract

We show a subset of $\mathbb{F}_p^2$ of size $Cn^{log_7(12)}$ does not necessarily contain any equilateral triangles by giving an explicit construction of such an equilateral-free subset. We do so by providing a map between $\mathbb{F}_p^2$ and sets of points on the plane. Lastly, we examine a special case of looking at only axis-aligned triangles.

Contents

1 Introduction

Various upper bounds have been proved for the maximum size of subsets in $\mathbb{F}_p^d$ that do not contain certain configurations of points. In (2) they show a upper bound on the maximum size of a subset of $\mathbb{F}_p^d$ which does not contain every configuration of k points in general position as long as $d > \binom{k}{2}$. Specifically for triangles, this gives a bound when $d > \binom{3}{2} = 3$. A method is given in (3) to generalize to any given distance graph, as opposed to just complete graphs. They also show specifically for unit equilateral triangles, any set in $\mathbb{F}_p^d$ with size at least $24p^{\frac{3+d}{2}}$ contains some equilateral triangle. Again, this gives a nontrivial bound only when $d > 3$.

A natural question is whether there exists any bound of the form Cp^{α} with some $\alpha < d$ when d is 2 or 3, or if it is possible to find equilateral free sets with a positive percentage of the available points, or neither. For some values of p, this is trivial, as the entirety of $\mathbb{F}_p^2$ contains no equilateral triangles. (1) shows that it contains equilateral triangles iff $\sqrt{3} \in \mathbb{F}_p^2$, which is true iff $p \equiv 1, 3, \text{or} 11 \pmod{12}$. When $d = 3$, there always exists equilateral triangles, take $(1, 0, 0), (0, 1, 0), (0, 0, 1)$ for instance. Instead of looking for upper bounds, this paper will instead look at the other side, and establish the following for $d = 2$:

Theorem 1.1 There exists some C such that for any p, there's an $S \subset \mathbb{F}_p^2$ of size $|S| \geq Cn^{log_7(12)}$ which contains no three distinct points forming an equilateral triangle.

The proof will rely on the structure of $\mathbb{F}_p^2$ being similar to the normal Euclidean plane. Intuitively, many of the the operations we can do are the same. Formally, we will make use of the following bijection between the finite field and the hexagonal plane:

Theorem 1.2 There is a bijection between the points of $S \subset \mathbb{F}_p^2$ and P_p (see Definitions) for all odd $p > 3$ such a triple is an equilateral triangle in one iff it is an equilateral triangle in the other.

As such, we will look for large subsets of the hexagonal coordinate system that are equilateral-free. Lastly, because it is curious if there exists any bound of the form Cp^{α} with $\alpha < d$, we will also examine and prove the following special case:

Theorem 1.3 For any C and any $\alpha < 2$, there is some subset of P_n for large enough n with at least Cn^{α} elements, that contains no equilateral triangles with sides parallel to the coordinate axes.

This may provide evidence that there is no way to get a a bound of the form Cn^{α} on the size of equilateral free sets for $d = 2$, although trying to apply a similar proof method does not work.

2 Definitions

For any two elements x and y of $\mathbb{F}_p^d$, the d-dimensional vector space over the finite field with characteristic p prime, we define the distance between the two to be:

$$||x - y|| = (x_1 - y_1)^2 + (x_2 - y_2)^2 + ... + (x_d - y_d)^2$$

Note that this does not follow the normal intuition of a distance function as the triangle inequality does not always hold, and distances can be 0 between distinct points. Three distinct points, $x, y, z \in \mathbb{F}_p^d$ are called an **equilateral triangle** iff

$$||x - y|| = ||x - z|| = ||y - z||$$

and a subset $S \subset \mathbb{F}_p^d$ is called **equilateral − free** if it has no such triple.

We will also be making use of the hexagonal coordinate system, which consists of all points of the form

$$\{(a + \frac{b}{2}, \frac{b\sqrt{3}}{2})|a, b \in \mathbb{Z}\}$$

We can divide this into n^2 equivalency classes which we will call P_n with two points $(a_1 + \frac{b_1}{2}, \frac{b_1\sqrt{3}}{2})$ and $(a_2 + \frac{b_2}{2}, \frac{b_2\sqrt{3}}{2})$ in the same equivalency class iff $a_1 \equiv a_2 \pmod{p}$ and $b_1 \equiv b_2 \pmod{p}$. We will refer to this element of P_n by (a, b) where $a \equiv a_1 \equiv a_2 \pmod{p}$, $b \equiv b_1 \equiv b_2 \pmod{p}$, and $0 \leq a, b \leq n$. The points of P_3 are shown in Figure 1 with a coloring as an example. Note that there are not more than 3 colors along any line, and that every 3x3 parallelogram in any of the 3 orientations contains all 9 colors exactly once.

We don't define a distance function between points of P_n, but instead define a distinct triple to be an equilateral triangle if there exists a triple of points containing one from each equivalence class that is an equilateral triangle in the plane.

Figure 1: Points of the hexagonal coordinate system separated into the 3^2 equivalency classes of P_3

3 Proof of the main result

Using the bijection in theorem 1.2, our goal becomes to show there exists subsets of P_n of size Cn^α with $\alpha = log_7(12) \approx 1.277$ for some C for any n which contains no equilateral triangle.

First, looking at $n = 7$, there is a equilateral-free subset of size $7^\alpha = 12$ of P_7 as shown by the circled points in figure 2. This can be checked to not contain equilateral triangles, and has been programmatically verified as the largest such subset of P_7. The code for this is in the Appendix. The code isn't fast enough to run for larger grids, but this proof method can use a better starting grid to get a better bound, if one is found.

Next, we will show that for any k and S, if there exists an equilateral-free subset of P_k with size S, how to construct one in P_{7k} with size $12S$. The method will be to place a P_7 into each point of the P_k, and in the S originally selected points, we select the corresponding 12 points of that P_7. An example of this (with $k = 3, S = 4$) is shown in figure 3 to help with understanding. Formally, we have a set A of S pairs in P_k, and a set B of 12 pairs in P_7, we then make set

$$A' = \{(7a + b, 7c + d) | 0 \leq a, c < k, 0 \leq b, d < 7, (a, c) \in A, (b, d) \in B\}$$

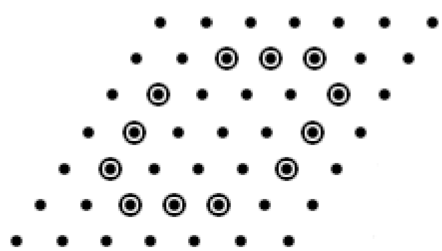

Figure 2: Maximum size equilateral-free set in P_7

Figure 3: Method for extending equilateral-free sets into larger ones

The reason A' will not contain any equilateral triangles is as follows. The points of the plane contained in it are a subset of the points defined by the 12 equivalence classes of P_7, so the only possible equilateral triangles must be within the same equivalence class in P_7. That means they must differ by multiples of 7 in axis parallel directions. Taking all such points reduces it to a copy of the original P_n, which will either be completely empty, or a copy of S, which we already know does not contain any equilateral triangles. Essentially, this is just a tensor product of the two.

With this method, we can make a subset of P_{7^n} without any equilateral triangles, which will be of size $(7^n)^\alpha = 12^n$ for any n. This doesn't prove the theorem completely though. Just because we can choose 144 points in P_{49} doesn't necessarily mean we can choose 144 points in P_{50} because we have to worry about triangles formed with wrapping around. However we can divide P_n into two sections, dividing along the main diagonal. Then, when the plane is tiled as shown in Figure 5, if we only pick points that lie under the main diagonal to be in our set, we avoid the problem of wrapping around P_n completely. This is because for any two points lying under the main diagonal, i.e. in one of the darker triangles in Figure 5, the points that complete the equilateral triangle will either be in the same darker triangle, and therefore be in the same copy of P_n, or be in a lighter triangle and therefore not be in the selected set at all. These two cases are demonstrated in the Figure as well.

By choosing a $7^a < \frac{n}{2}$ and choosing the 12^a points in the lower left P_{7^a} of P_n, we get some equilateral-free set. There is always a power of seven with $\frac{n}{14} \leq 7^a < \frac{n}{2}$ for any n, so this set will have at least $\frac{n}{14}^{log_7(12)} = 14^{-log_7(12)} n^{log_7(12)}$. □

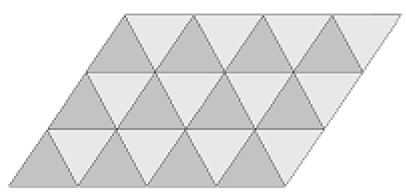

Figure 4: Division of P_n into two parts

4 Proof of Theorem 1.2

in Theorem 1.4 of (1), it is established that $\mathbb{F}^2$ contains equilateral triangles if and only if $\sqrt{3} \in \mathbb{F}$. Additionally, Lemma 4.1 in the proof of that theorem implies that if $\mathbb{F}$ does not have characteristic 2, then for any line segment (x_1, x_2), the amount of x_3s such that

$$||x_1 - x_2|| = ||x_1 - x_3|| = ||x_2 - x_3||$$

is

$$\begin{cases} 2 & \text{if 3 is a nonzero square in } \mathbb{F} \\ 1 & \text{if } 3 \equiv 0, \text{ (so if } \mathbb{F} \text{ has characteristic 3)} \\ 0 & \text{otherwise} \end{cases}$$

We will first solve to determine exactly what these x_3 are, if they exist.

Lemma 1.1 For any distinct x_1, x_2, x_3 in $\mathbb{Z}_p$ with $p > 2$, they form an equilateral triangle if and only if $x_3 = x_1 + R(x_2 - x_1)$ where

$$R = \begin{bmatrix} 1/2 & \pm\sqrt{3}/2 \\ \mp\sqrt{3}/2 & 1/2 \end{bmatrix}$$

Proof :The intuition behind R is that it represents a rotation of 60 degrees that will be applied to x_2 clockwise or counter-clockwise about x_1.
First, we can show that if $\sqrt{3} \neq 0$, the two solutions are not equal. This is because

$$\begin{bmatrix} 1/2 & \sqrt{3}/2 \\ -\sqrt{3}/2 & 1/2 \end{bmatrix} (x_2 - x_1) = \begin{bmatrix} 1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & 1/2 \end{bmatrix} (x_2 - x_1)$$

means that

$$\begin{bmatrix} 0 & \sqrt{3} \\ -\sqrt{3} & 0 \end{bmatrix} (x_2 - x_1) = 0$$

which means that either $x_2 - x_1 = 0$ which isn't allowed as they are distinct, or the matrix has nonzero kernel, which would mean its determinant is zero, so $3 = 0$, which is a contradiction.

Therefore, taking $x_3 = x_1 + R(x_2 - x_1)$ gives two distinct solutions if $\sqrt{3} \neq 0$ exists, one solution if $\sqrt{3} = 0$, and none if $\sqrt{3}$ does not exist. We know this is the total amount of solutions that will extend the segment (x_1, x_2) into an equilateral triangle, so if we show that these solutions do make an equilateral triangle, we will also have shown that no other solutions exist.

The rest is just computation. Letting $x_1 = (a, b)$, $x_2 = (c, d)$, so then $x_3 = (\frac{1}{2}(a + c) \pm \frac{\sqrt{3}}{2}(d - b), \frac{1}{2}(b + d) \mp \frac{\sqrt{3}}{2}(c - a))$, we can find the distance

$$||x_3 - x_1|| = (\frac{1}{2}(c - a) \pm \frac{\sqrt{3}}{2}(d - b))^2 + (\frac{1}{2}(d - b) \mp \frac{\sqrt{3}}{2}(c - a))^2$$

$$= (\frac{1}{4} + \frac{3}{4})(c - a)^2 + (\frac{\pm\sqrt{3}}{2} + \frac{\mp\sqrt{3}}{2})(c - a)(d - b) + (\frac{1}{4} + \frac{3}{4})(d - b)^2$$

$$= (c - a)^2 + (d - b)^2 = ||x_2 - x_1||$$

and following similarly we also get

$$||x_3 - x_2|| = (\frac{1}{4} + \frac{3}{4})(c - a)^2 + (\frac{\mp\sqrt{3}}{2} + \frac{\pm\sqrt{3}}{2})(c - a)(d - b) + (\frac{1}{4} + \frac{3}{4})(d - b)^2$$

$$= (c-a)^2 + (d-b)^2 = ||x_2 - x_1||$$

□

Now that we have this lemma, when the characteristic of $\mathbb{Z}_p$ is greater than 3, this closely matches what we expect from the plane. To get an equilateral triangle, you can start with any segment, and rotate one point about the other by 60 degrees in either direction. This intuition will give way to the bijection. Taking $\mathbb{Z}_n$ with $n > 3$ and $\sqrt{3} \in \mathbb{Z}_n$, every element of $\mathbb{Z}_n^2$ can be decomposed into $a(1,0) + b(\frac{1}{2}, \frac{\sqrt{3}}{2})$ with $0 \leq a, b < n$ (Choosing one of the two possible $\sqrt{3}$ arbitrarily but consistently). Matching up that point with $a(1,0) + b(\frac{1}{2}, \frac{\sqrt{3}}{2})$ on the plane, we get the desired correspondence. Because on the plane, x_1, x_2, x_3 also form an equilateral triangle if and only if $x_3 = x_1 + R(x_2 - x_1)$, and the bijection preserves scalar multiplication and addition, the bijection preserves equilateral triangles.

As an example, Figure 1 shows three equilateral triangles in Z_{11}^2 on the left, and the equilateral triangles that they map to in P_{11}.

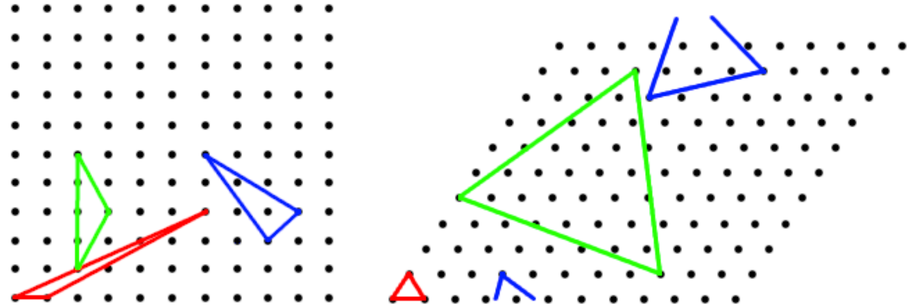

Figure 5: Three triangles in Z_{11}^2 and in P_{11}

□

5 Proof of Theorem 1.3

The proof of this theorem will rely on the fact that there exist Salem-Spencer sets of size greater than n^α for any $\alpha < 1$, as for any equilateral triangle with sides parallel to the coordinate axes, the x-coordinates of these points form an arithmetic series.

This doesn't completely solve the problem, as x-coordinates are not evenly distributed within any one n by n parallelogram of representative points in P_n, and also we have to worry about wrapping around similar to the proof of the main result. However we can employ a similar trick in which we pick only points from a more restricted shaded region, as shown in Figure 6. Formally, we choose all

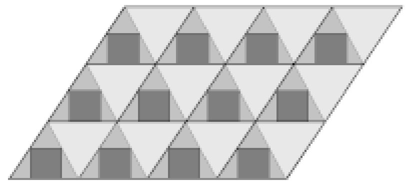

Figure 6: Darker shaded contains evenly distributed x-coordinates

$$\{(a+\frac{b}{2},\frac{b\sqrt{3}}{2})|\text{b is even}, \frac{n}{4}\le a+\frac{b}{2}<\frac{3n}{4}, 0\le b<n/2\}$$

and the corresponding equivalence classes. This represents $\lfloor\frac{n}{2}\rfloor$ choices of b, and $\lfloor\frac{n}{2}\rfloor$ choices of a for each b, so $\lfloor\frac{n}{2}\rfloor^2 \approx \frac{n}{4}$ total points, forming a square with $\lfloor\frac{n}{2}\rfloor$ possible x and y coordinates. Next, we choose a Salem-Spencer set of numbers from 0 to $\lfloor\frac{n}{2}\rfloor$ and take all points in this with an x coordinate exactly $\lfloor\frac{n}{4}\rfloor$ above a point in this set.
For any α, we can choose a n large enough such that the Salem-Spencer set will have at least $\frac{n}{2}^{\alpha-1}$ elements, and then the set we pick will have $\frac{n}{2}^{\alpha} = Cn^{\alpha}$ elements.

6 Concluding thoughts

We still haven't touched on what if $d = 3$, a similar method that was employed here to biject into space is unlikely to work, as the number of triangles for any segment in $\mathbb{F}_n^3$ depends on n, which has no parallel in space. Every slice of $\mathbb{F}_n^3$ creates a copy of $\mathbb{F}_n^2$ though scaled somehow, so there is some hope of using a solution from a lower dimension to find one for a higher dimension

If a faster program was able to find grids of larger sizes with better ratios than the 7x7 example provided, they would instantly improve the lower bound in this paper, by using the same method. If it turns out to be like the problem in Theorem 1.3, these lower bounds would not hit a stopping point, and you would have to use an argument like with Salem-Spencer sets, or perhaps Cap sets. Cap sets seem possibly related in that they also are in d dimensional fields and deal with finding large sets avoiding specific triples of points.

7 Appendix

```
#include <iostream>
#include <vector>
using namespace std;

int mx; //best solution found so far
int n; //size of the grid
```

```cpp
struct arr{ //struct used in order to be able to pass by val
	int ans[7][7];
};

void solve(arr ans, vector<pair<int, int> > marked, int x, int y, int tot){

	if(y==n){ //if we are at the end of the grid
		if(tot > mx){ //and our current solution is the best so far
			mx = tot; //update and output
			cout << "Set with size " << tot << ":" << endl;
			for(pair<int, int> p : marked){
				cout << p.first << " " << p.second << endl;
			}
			for(int i=0; i<n; i++){
				for(int j=0; j<n; j++){
					cout << ((ans.ans[i][j]==1)?"X ":"O ");
				}
				cout << endl;
			}
			cout << endl << endl;
		}
		return;
	}
	else{
		int nx = x; //get the next cell to recurse
		int ny = y;
		nx++;
		if(nx == n){
			ny++;
			nx = 0;
		}
		if(ans.ans[x][y] == -1) { //if we can't choose x, y
			solve(ans, marked, nx, ny, tot);
			return;
		}
		ans.ans[x][y] = -1;
		solve(ans, marked, nx, ny, tot); //if we don't pick x,y
		//otherwise we select x, y
		for(pair<int, int> p : marked){ //go through and mark points which now can't b
			int a = p.first;
			int b = p.second;
			ans.ans[(2*n+a+b-y)%n][(2*n+y+x-a)%n] = -1;
			ans.ans[(2*n+x+y-b)%n][(2*n+b+a-x)%n] = -1;
		}
		ans.ans[x][y] = 1;
		marked.push_back(make_pair(x, y));
		solve(ans, marked, nx, ny, tot+1);
		return;
	}

}


int main() {
	n = 7;
	arr ans;
```

```
        for(int i=0; i<n; i++){
                for(int j=0; j<n; j++){
                        ans.ans[i][j] = 0;
                }
        }
        vector<pair<int, int> > marked;
        //we can always rotate to be able to pick (0,0) and (1,0) so start with those marked
        ans.ans[0][0] = 1;
        ans.ans[1][0] = 1;
        ans.ans[0][1] = -1;
        ans.ans[1][n-1] = -1;
        marked.push_back(make_pair(0,0));
        marked.push_back(make_pair(1,0));
        solve(ans, marked, 2, 0, 2);
        return 0;
}
```

References

[1] M. Bennett, A. Iosevich, and J. Pakianathan *Three-point configurations determined by subsets of F2 via the Elekes-Sharir paradigm.* Combinatorica 34 (2014), no. 6, 689–706.

[2] D. Hart and A. Iosevich *Ubiquity of simplices in subsets of vector spaces over finite fields.* Anal. Math. 34 (2008), no. 1, 29–38.

[3] A. Iosevich and H. Parshall, *Embedding distance graphs in finite field vector spaces*, arXiv:1802.06460