A new bound and new techniques for the Erdős–Ginzburg–Ziv constant for finite abelian groups

Jared Bitz and Cory Griffith

August 19, 2017

Abstract

For a finite abelian group G, the generalized Erdős–Ginzburg–Ziv constant $\mathsf{s}_k(G)$ is the smallest m such that a sequence of m elements in G always contains a k-element subsequence which sums to zero. We produce the first exponential lower bound for $\mathsf{s}_{2p}(C_p^r)$, introduce new techniques for determining congruences between the number of zero-sum subsequences of given lengths in the groups C_p^r, and use these techniques to rederive exact values for $\mathsf{s}_{3p}(C_p^3)$ and $\mathsf{s}_{4p}(C_p^4)$.

1 Introduction

In 1961, Erdős, Ginzburg, and Ziv proved that among any $2n-1$ integers, one can always find n of them whose average is an integer [2]. Today, this result is usually phrased more abstractly in terms of finite abelian groups, which is to say, sums of finite cyclic groups. Henceforth, let p be an odd prime, and let $C_p := \mathbb{Z}/p\mathbb{Z}$. The oddness condition on p exists because in the case $p = 2$, the theory becomes trivial, but produces results distinct from that obtained with p odd. The EGZ Theorem, then, states that for any $(2p-1)$-sequence in C_p, there exists a subsequence of exactly p elements that sums to zero.

The first step in generalizing this result is defining the Erdős–Ginzburg–Ziv constant for arbitrary groups. Recall that the exponent of a finite group is the largest order among its elements.

Definition 1. *Let G be a finite abelian group, with prime exponent p. Then, the Erdős–Ginzburg–Ziv constant of G, denoted $\mathsf{s}_p(G)$, is the smallest m such that a sequence of m elements in G always contains a subsequence of p elements that sums to zero.*

It is important to note that we have defined the EGZ constant for subsequences and not subsets - we will allow repeats. Since we are concerned with sums of elements in an abelian group, the order of a sequence does not matter either. In what follows, if X is a sequence, we write $x \in X$ to mean "x appears at least once in X."

It is also easy to see why we are interested in zero-sum subsequences of lengths that are multiples of p. Indeed, showing that a sequence of a certain size must always contain a zero-sum subsequence of length k for $k \equiv a \bmod p$ and $a \neq 0$ is impossible: we can choose some element $g \in G$ of order p, and then the sequence consisting of g repeated an arbitrarily large number of times has no length k zero-sum subsequence, as $k \cdot g = a \cdot g \neq 0$. Thus, the EGZ constant $\mathsf{s}_k(G)$ cannot exist.

Of course, we can still consider sequences whose lengths are 0 mod p, and this opens the door to defining an even more general property, which is the central focus of this paper.

Definition 2. *Let G be a finite abelian group with exponent p. The k-th generalized EGZ constant, denoted $\mathsf{s}_{kp}(G)$, is the smallest m such that a sequence of m elements in G always contains a subsequence of kp elements that sums to zero.*

It may seem curious that we have been dealing only with groups of prime exponent. In most cases, and especially in the case of cyclic groups and their powers, these turn out to be the most interesting structurally, as we can synthesize results about groups with exponent p and exponent q into results about groups with exponent pq.

The EGZ problem, especially in the "smallest" case of $k = 1$, has proven to be surprisingly intractable. In the almost sixty years since the Erdős paper, the only dimensions in which s_p has been exactly determined for all p are one and two. That is, we know the precise values of $\mathsf{s}_p(G)$ when $G = C_p$ and when $G = C_p^2$. The former result of $\mathsf{s}_p(C_p) = 2p-1$ comes from the original Erdős paper, while the latter of $\mathsf{s}_p(C_p^2) = 4p-3$ is rather recent, and due to Reiher [2], [9]. Although it was Reiher who pinned down the actual value for C_p^2, Rónyai also did notable work, proving an upper bound of $4p - 2$ and in the process developing methods that have become widespread in additive combinatorics [10]. Our results largely come from attempts to extend and systematize congruence-based techniques presented by Reiher. A description of the resulting framework forms the bulk of section 4 and section 5.

More recently, Naslund has obtained exponential upper bounds on $\mathsf{s}_p(C_p^r)$ for general r by using the so-called "slice rank" method, which generalizes the standard tensor rank of a function [8]. Specifically, he shows that

$$\mathsf{s}_p(C_p^r) < 3p!(2p-1)(J(p)p)^r,$$

where $J(p)$ is a constant, depending on p, which lies between .841 and .918 and decreases as p grows. [1]

Our results are primarily in the area of higher-order generalized EGZ constants, specifically s_{2p} and s_{3p}. In this paper, we prove the following theorems.

Theorem 1. *The generalized EGZ constant $\mathsf{s}_{2p}(C_p^r)$ satisfies the following bound*

$$\mathsf{s}_{2p}(C_p^r) > \frac{p(1.25)^r}{2}$$

That is, there is a sequence of $p(1.25)^r/2$ elements in C_p^r that contains no zero-sum subsequence of length $2p$. This result, shown in section 3, is new, and in fact is the first exponential lower bound on $\mathsf{s}_{2p}(C_p^r)$ for an arbitrary dimension r. We will also derive the known exponential lower bound

$$\mathsf{s}_p(C_p^r) > (p-1)2^r$$

in section 2.

Another result is an exact value for s_{3p} in the three-dimensional case.

Theorem 2. $\mathsf{s}_{3p}(C_p^3) = 6p - 3$ *for* $p > 3$.

This comes from careful examination of systems of congruences, detailed in section 5. Both this and the following theorem are due to Kubertin [7]. Our final result is proven in section 6.

Theorem 3. $\mathsf{s}_{4p}(C_p^4) = 8p - 4$ *for* $p > 4$.

We will also use new techniques to rederive the upper bound in the known result that $5p+\frac{p-1}{2}+3 \leq \mathsf{s}_{2p}(C_p^3) \leq 6p - 3$, due to Gao and Thangadurai [4].

[1] Examining the preprint of Naslund's paper, one finds that the $2p-1$ factor in this bound is instead written as $p-1$. This difference comes from the the fact that Naslund considers sequences of unique elements only, as subsets of $\mathbb{F}_p^r$.

2 Background Theory

For the rest of the paper, unless explicitly stated otherwise, X is taken to be a sequence in C_p^r for arbitrary r.

To get a handle on $\mathsf{s}_p(C_p^r)$ for general r, we begin by restating the elementary bounds proved by Harborth [5].

Proposition 1.

$$(p-1)2^r < \mathsf{s}_p(C_p^r) < (p-1)p^r + 1.$$

Proof. For the lower bound, we will proceed by induction on the dimension. As a base case, is clear that in C_p, the sequence of $(p-1) \times 1$ and $(p-1) \times 0$ has no p-length zero-sum subsequence.

Now consider the sequence X given by taking $p-1$ of each vector in the Hamming cube $\{0,1\}^r$. We claim that no subsequence of p elements in X sums to zero. Such a subsequence must contain exactly 0 or p ones in each of the r coordinates. It is impossible that all coordinates contain p ones, since we only have $p-1$ of the vector $(1,1,\ldots,1)$, and only one choice afterwards to bring us to a p-subsequence. Thus, some coordinate must stay at 0, and we reduce to the case of the dimension one lower.

The upper bound is a pigeonhole argument. Given $(p-1)p^r$ elements of C_p^r, we either have exactly $p-1$ of each element in C_p^r or at least p of some element α. In the latter case, the p copies of α sum to 0. Otherwise, we can choose any element x and take $\frac{p-1}{2}$ copies of x and $\frac{p-1}{2}$ copies of $-x$. These, combined with a single copy of the zero vector, give a length p zero-sum subsequence. □

Remark 1. This upper bound is believed to be far larger than the actual value. Specifically, $\mathsf{s}_p(C_p^r)$ is believed to be linear in p for any fixed r, as is the case for $r = 1$ and 2. In general, for large enough k, $\mathsf{s}_{kp}(C_p^r)$ is known to be linear. Even better, it is known that

$$\mathsf{s}_{kp}(C_p^r) = (r+k)p - r \tag{1}$$

for k sufficiently large [3].

However, determining the values of k at which this behavior kicks in is an open question, which is described as determining the phase transition at r. When $r = 3$,Equation 1 holds for $k \geq 4$, as per a result of Gao and Thangadurai [4]. Theorem 2 shows that in three dimensions this behavior begins exactly at $k = 3$, since $\mathsf{s}_{3p}(C_p^3) = 6p - 3$ but $\mathsf{s}_{2p}(C_p^3) \neq 5p - 3$.

3 Lower Bounds on $\mathsf{s}_{kp}(C_p^r)$

We now describe the proof that $\mathsf{s}_{2p}(C_p^r)$ has a lower bound exponential in r. The construction is probabilistic, and uses a somewhat crude union bound to achieve its end. While more advanced probabilistic techniques such as the Lóvas Local Lemma net small gains, they do not improve the exponential constant, which is the most important element [1].

Proof of Theorem 1. Let

$$N = \frac{p(1.25-\varepsilon)^r}{2}$$

for $\varepsilon > 0$, and let $q = 1/5 + \varepsilon$. For the rest of the proof, we supress the ε, with the understanding that all quantities related to N and q are arbitrarily close to the stated values. Choose a sequence X of N

random vectors in $\{0,1\}^r$ as follows. For each v that is a term of X (hence a vector), let $v_i = 1$ with probability q and $v_i = 0$ with probability $1-q$, with each coordinate in each vector independent. Let Z be the number of zero-sum $2p$-subsequences in X. We show that $\mathbb{E}[Z] < 1$, and so there must be some possible X with no such subsequences.

Consider some arbitrary $2p$-subsequence Y of X. For Y to be zero-sum, each of the r coordinates must sum to 0 mod p, and so contains exactly 0, p, or $2p$ ones. For any coordinate $i \leq r$, let P_0 be the probability that coordinate i contains 0 ones, P_p the probability that it contains p ones, and P_{2p} the probability that it contains $2p$ ones (clearly, this is not dependent on i). We have

$$\begin{aligned} P_0 &= (1-q)^{2p} \\ P_p &= \binom{2p}{p} q^p (1-q)^p < 2^{2p} q^p (1-q)^p \\ P_{2p} &= q^{2p}, \end{aligned}$$

the values from a standard binomial distribution with probability of success q.

Now, if Q is the probability that coordinate i sums to zero, we have

$$Q = P_0 + P_p + P_{2p} < (1-q)^{2p} + 2^{2p} q^p (1-q)^p + q^{2p}.$$

For $1/5 < q < 4/5$, we see by solving for q in the inequalities $P_p > P_0$ and $P_p > P_{2p}$ that the second term dominates the first and third when p is large. So, asymptotically, we have

$$\begin{aligned} Q &< 2^{2p} \left(\frac{1}{5}\right)^p \left(\frac{4}{5}\right)^p \\ &< \left(\frac{4}{5}\right)^{2p}. \end{aligned}$$

Since Q is the probability that any one coordinate in Y sums to zero, the probability that Y as a whole is zero-sum is Q^r. Since each $2p$-subsequence of X is zero-sum with equal probability, we have

$$\begin{aligned} \mathbb{E}[Z] &= \binom{N}{2p} Q^r \\ &< \left(\frac{2N}{p}\right)^{2p} \left(\frac{4}{5}\right)^{2pr} \\ &< \left(\frac{5}{4}\right)^{2pr} \left(\frac{4}{5}\right)^{2pr} \\ &< 1, \end{aligned}$$

as desired. □

This result is notable because in addition to being the first exponential lower bound on s_{2p} for general r, it is also of the expected order of growth in p. Indeed, an order of growth larger than linear in p would be alarming, as it would indicate that the linear formulas thus observed do not generalize to higher dimensions.

As previously mentioned, one might try several strategies for raising this lower bound. In particular, the calculation of the expectation of Z really consists of performing a union bound by summing

the probabilities of each event, without regard to intersection. One might therefore attempt to use the method of alterations. This proceeds as follows. If we relax our requirements and show that $\mathbb{E}[Y] < N/2$, we can remove one element from each length $2p$ zero-sum subsequence, and be left with a new set X' with size on the same order as X, but that does not contain any length $2p$ zero-sum subsequences. As it turns out, this strategy does not result in any improvements to the exponential constant $5/4$.

The next question one might ask is how well this method generalizes to give lower bounds on s_{kp} for arbitrary k. The answer turns out to be "rather poorly." Indeed, suppose we were instead interested in kp-subsequences. If we repeat the above construction with

$$N = \frac{p(1+\delta)^r}{k}$$

and try to optimize for δ, we first must calculate the probability that a given coordinate $i \leq r$ sums to zero. We now have k component probabilities, P_0 through P_{kp}. By the same logic as above, we see

$$P_{ip} = \binom{kp}{ip} q^{ip} (1-q)^{(k-i)p}.$$

If we want the two "most skewed" terms P_0 and P_p to dominate, we set

$$(1-q)^{kp} = \binom{kp}{p} (1-q)^{(k-1)p} q^p.$$

Bounding the binomial coefficient on the right, we see

$$\begin{array}{ccccc} \left(\frac{kp}{p}\right)^p (1-q)^{(k-1)p} q^p & < & (1-q)^{kp} & < & \left(\frac{ekp}{p}\right)^p (1-q)^{(k-1)p} q^p \\ k^p q^p & < & (1-q)^p & < & (ek)^p q^p \\ kq & < & 1-q & < & ekq \\ \frac{1}{ek+1} & < & q & < & \frac{1}{k-1}. \end{array}$$

In this range, P_0 and P_p are approximately equal, and so we can approximate $Q \approx 2(1-q)^{kp}$. Then,

$$\begin{aligned} \mathbb{E}[Z] = \binom{N}{k} Q^r &< \left(\frac{kN}{p}\right)^{kp} \left[2(1-q)^{kp}\right]^r \\ &< 2^r \left[(1+\delta)(1-q)\right]^{kpr}. \end{aligned}$$

If we want $\mathbb{E}[Z] < 1$, we must have

$$\begin{aligned} 2^r \left[(1+\delta)(1-q)\right]^{kpr} &< 1 \\ (1+\delta)(1-q) &< 2^{\frac{-1}{kp}} \\ \delta &< \frac{2^{\frac{-1}{kp}}}{1-q} - 1. \end{aligned}$$

For q in the range determined above, the permissible values of δ approach zero as k grows, since the numerator and denominator of the fraction both go to one. Therefore, this gives us nothing if we seek a general bound for $\mathsf{s}_{kp}(C_p^r)$ for completely general k and p. However, if we are interested in specific values of k, i.e. s_{3p}, one can still attempt to optimize q in order to achieve an exponential lower bound - it just becomes less effective for large k.

4 Algebraic Lemmas

For the rest of the paper, assume that $p > r$ is a prime. Further, the case $p = 2$ is in general trivial. We therefore also assume that p is odd.

In order to prove Theorem 2, we will first need to develop several lemmas that provide relations and congruences between the number of zero-sum subsequences in a given X. The method of examination of these relations is entirely different from the approaches previously in this paper. We will therefore need some new notation.

Definition 3. *Let X be a sequence of elements of C_p^r. Then if Y is a subsequence of X, we write $(kp \mid Y)$ to indicate the number of zero-sum subsequences of Y with length kp.*

We also β_k to indicate $(kp \mid X)$, with $\beta_0 = 1$. We will only consider β_k to be defined if $k \geq 0$ is small enough so that $kp \leq |X|$. Finally, let $\vec{\beta}$ refer to the vector of all β_i, meaning $\vec{\beta} := (\beta_0, \beta_1, \ldots, \beta_\ell)$.

We will also frequently omit the modulus in expressions of the form $x \equiv y \pmod{p}$, letting it be clear from context.

The first result we will need is the Chevalley-Warning Theorem. This is a well-known, purely algebraic theorem about common zeros of polynomials over finite fields.

Lemma 1 (Chevalley-Warning Theorem)**.** *Let $\mathbb{F}$ be a finite field, and $f_1, \ldots, f_r \in \mathbb{F}[X_1, \ldots, X_n]$, where*

$$n > \sum_{i=1}^{r} \deg(f_i),$$

Then, the number of common solutions $(a_1, \ldots, a_n) \in \mathbb{F}^n$ to all of the f_i is divisible by the characteristic of $\mathbb{F}$.

Proof. See [9]. □

Our strategy, then, will be to design a collection of polynomials such that their zeros give us information about the number of zero-sum subsequences of a given subsequence. Our first step is the following lemma:

Lemma 2. *Pick s, k so that $kp - s > (r+1)(p-1)$. Let X be a sequence in C_p^r, with $|X| = kp - s$. Then*

$$\sum_i (-1)^i \beta_i \equiv 0 \pmod{p}$$

The result is a generalization of a construction used by Reiher to arbitrary dimensions [9].

Proof. Let the sequence X be given by $\{(a_{1i}, \ldots, a_{ri})\}$, and let $n = |X|$. In other words, we fix an (arbitrary) ordering on the elements of X so that we may use them in sums. Then, a_{ji} is the j-th

coordinate of the i-th vector in X. Now, define a collection of $r+1$ polynomials over $\mathbb{F}_p$ by

$$\begin{aligned} f_0(x_1,\ldots,x_n) &= \sum_{i=1}^{n} x_i^{p-1} \\ f_1(x_1,\ldots,x_n) &= \sum_{i=1}^{n} a_{1i} x_i^{p-1} \\ &\vdots \\ f_r(x_1,\ldots,x_r) &= \sum_{i=1}^{n} a_{ri} x_i^{p-1}. \end{aligned}$$

We can think of the x_i as indicating whether or not to choose a specific element in X as part of a particular subsequence. In other words, we can associate each subsequence Y of X to the set of $(x_1,\ldots,x_n)$ such that $x_i \neq 0$ if and only if the i-th element of X is in Y. Then, f_0 is zero when $|Y| \equiv 0 \bmod p$. For $i > 0$, f_i is zero when the the corresponding coordinate in Y sums to zero across the chosen elements. So, the common zeros of the f_i exactly encode the zero-sum subsequences of X whose lengths are multiples of p.

Now, suppose we have a zero-sum subsequence of length lp. To count the common zeroes of the f_i that this subsequence induces, we note that exactly lp of the x_i must be non-zero, and that we can choose for each of these any element of $\mathbb{F}_p^\times$. So, the total number of zeros is $(p-1)^{lp} \equiv (-1)^{lp}$. When we sum across all possible lengths of subsequences (including the trivial subsequence of length zero), we see by the Chevalley-Warning Theorem that

$$\sum_i (-1)^i \beta_i \equiv 0 \mod p$$

as desired. □

This single congruence can, somewhat surprisingly, be used to generate a family of further independent congruences on the $(lp \mid X)$ that we encode as a matrix over $\mathbb{F}_p$.

For our next lemma, recall the conventions that $\binom{x}{y} = 0$ if $x < y$, and that $\binom{0}{0} = 1$.

Lemma 3. *Pick s with $0 \leq s < r$. Consider a sequence X in C_p^r with $|X| = kp - r + s$ for some $k \geq r+1$. Let A be the $(k-r) \times k$ matrix*

$$\begin{pmatrix} \binom{k-1}{k-1} & -\binom{k-2}{k-2} & \cdots & (-1)^{k-1}\binom{0}{0} \\ \binom{k-1}{k-2} & -\binom{k-2}{k-3} & \cdots & (-1)^{k-1}\binom{0}{1} \\ \vdots & \vdots & \ddots & \vdots \\ \binom{k-1}{k-r-1} & -\binom{k-2}{k-r-2} & \cdots & (-1)^{k-1}\binom{0}{k-r-1} \end{pmatrix}$$

Then $A\vec{\beta} \equiv \vec{0}$, when the entries are taken modulo p.

As an example, consider a sequence of $6p-3$ elements in C_p^3. The matrix encoding the relations of

the above lemma is

$$\begin{pmatrix} 1 & -1 & 1 & -1 & 1 & -1 \\ 5 & -4 & 3 & -2 & 1 & 0 \\ 10 & -6 & 3 & -1 & 0 & 0 \end{pmatrix} \begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \\ \beta_5 \end{pmatrix} \equiv \vec{0}$$

Notice that the first row encodes the alternating sum from the previous lemma.

Definition 4. *A matrix determined as in the lemma for a sequence X is called the relation matrix for X.*

Note that the relation matrix for X is only dependent on $|X|$, and not the actual elements it contains. Relation matrices are always truncated versions of Pascal's triangle, with an alternating sign on each column. [2]

Proof of Lemma 3. For each $r+1 \leq n \leq k$, consider some arbitrary subsequence $Y \subset X$ with $|Y| = np - r + s$. As long as $n \geq r+1$, Lemma 2 still holds, and

$$1 - (p \mid Y) + (2p \mid Y) - \cdots + (-1)^{n-1}((n-1)p \mid Y) \equiv 0$$

If we sum across all possible Y of size $np - r + s$, we see that

$$\sum_{|Y|=np-r+s} \left[1 - (p \mid Y) + (2p \mid Y) - \cdots + (-1)^{n-1}((n-1)p \mid Y)\right] \equiv 0$$

Now, each lp-length zero-sum subsequence is counted multiple times in this sum. If we know exactly how many times, at least modulo p, this will give us a new equation

$$\sum_i c_i \beta_i \equiv 0 \tag{2}$$

To calculate c_l, consider how many of our Y contain any particular length lp zero-sum subsequence. We need a total of $np - r + s$ elements, and have already fixed lp. Thus, we can choose the remaining $(n-l)p - r + s$ from among $|X| - lp = (k-l)p - r + s$ elements. So, each such subsequence is counted

$$\binom{(k-l)p - r + s}{(n-l)p - r + s}$$

times. We wish to evaluate this modulo p. Writing it out explicitly and gathering the p factors gives

$$\begin{aligned} \binom{(k-l)p - r + s}{(n-l)p - r + s} &\equiv \frac{[(k-l)p - r + s]!}{[(k-n)p]![(n-l)p - r + s]!} \\ &\equiv \frac{(k-l-1)! \cdot p^{k-l-1} \cdot (p - r + s)! \cdot ((p-1)!)^{k-l-1}}{[(k-n)! \cdot p^{k-n} \cdot ((p-1)!)^{k-n}][(n-l-1)! \cdot p^{n-l-1} \cdot (p-r+s)! \cdot ((p-1)!)^{n-l-1}]} \\ &\equiv \frac{(k-l-1)!}{(k-n)!(n-l-1)!} \\ &\equiv \binom{k-l-1}{n-l-1} \end{aligned}$$

[2] A program to generate relation matrices has been included in Appendix A

So, Equation 2 is calculated to be

$$\binom{k-1}{n-1}\beta_0 - \binom{k-2}{n-2}\beta_1 + \cdots + (-1)^{n-1}\beta_n \equiv 0$$

which gives us precisely the desired relationship on a row of the relation matrix of X. □

By summing over appropriately sized subsets of X, we were able to derive relations involving fewer and fewer of the terms $(lp \mid X)$. We will now use these relations to prove a rather extreme restriction on how many of these terms can take on a value of 0 mod p.

5 Upper bounds on $\mathsf{s}_{kp}(C_p^3)$

Having developed the results of section 4, there are a number of strategies one can pursue using the following idea: if X is a sufficiently long sequence and we would like to show that some β_i is non-zero, we may set it to zero and use our relations to produce some oddity from which we can derive a contradiction. The following result and its corollary subsumed all our previous attempts, and produced a quick proof of Theorem 2.

Lemma 4. *Let A be a relation matrix. Then any $(k-r)\times(k-r)$ submatrix of A is invertible.*

Proof. This is an immediate consequence of the first lemma of [6]. □

Lemma 5 (Independence Lemma)**.** *Let X be a sequence in C_p^r, of length $kp-r+s$, with $k \geq r+1$ and $0 \leq s < r$. Then at most $r-1$ of the β_i satisfy $\beta_i \equiv 0$.*

Proof. Suppose r of the β_i satisfy $\beta_i \equiv 0$, say $\beta_{t_1}, \ldots, \beta_{t_r}$. Form the relation matrix A of X. Then by removing the t_i-th columns, we obtain a $(k-r)\times(k-r)$ matrix A'. Letting $\vec{\beta}'$ be $\vec{\beta}$ with the β_{t_i} removed, we have $A'\vec{\beta}' \equiv \vec{0}$. By Lemma 4, we know that A' is invertible. Thus $\vec{\beta}' \equiv \vec{0}$, but by definition $\beta_0 = 1$, so that β_0 is in $\vec{\beta}'$, yielding a contradiction. □

A brief expository note may be helpful here. One of the difficulties of examining the EGZ constant is that one can obtain no useful information about the overlap between two sequences unless they are completely disjoint or one is contained in the other. For example, suppose we wish to determine information from the existence of particular zero-sum subsequences of length p and $3p$ in a sequence X. If they are disjoint, we obtain a length $4p$ zero-sum subsequence. If one is contained in the other, we obtain a length $2p$ zero-sum subsequence. In both these cases, we may have significantly increased our understanding of X. On the other hand, if the sequences have even a single element in common, there is no evident information to be gained.

This example suggests that the statements about X that will most easily build up into other statements are those concerning the existence of disjoint zero-sum subsequences, especially those contained in zero-sum subsequences we already have information about. The following corollary allows us to use the existence of long zero-sum subsequences to obtain statements about the existence and disjointness of strictly contained zero-sum subsequences.

Definition 5. *Given a zero-sum subsequence $Z \subseteq X$ of length a multiple of p, we say Z splits if Z strictly contains a zero-sum subsequence W that also has length a multiple of p.*

Evidently, the sum across W and $Z\backslash W$ both come to zero. The lemma may then be interpreted as a description of exactly which lengths of subsequences can be "atomic" in each dimension, and tells us that this is independent of p.

Lemma 6 (Splitting Lemma)**.** *Let X be a zero-sum sequence of length $kp > (r+1)(p-1)$. Then X splits.*

Proof. Remove any vector from X to obtain a sequence to which we can apply Lemma 5. □

Theorem 2 now follows rapidly.

Proof. For the lower bound, it is clear that the sequence consisting of $3p-1$ copies of the zero vector and $p-1$ copies of each standard basis vector e_1, e_2, e_3 contains no zero-sum subsequence of length $3p$. For the upper bound, fix a sequence X with $|X| = 6p-3$, and suppose that $\beta_3 = 0$. By the splitting lemma we know any zero-sum subsequence of length $5p$ must split, and since $\beta_3 = 0$, any such splitting must involve a length p zero-sum subsequence. Thus if $\beta_1 = 0$, we have $\beta_5 = 0$, which contradicts the independence lemma, so that we can pick a zero-sum subsequence Y of length p.

Note that $X\backslash Y$ has length $5p-3$, so we may apply the independence lemma. Certainly $(3p \mid X\backslash Y) = 0$, so at most one of $(p \mid X\backslash Y), (2p \mid X\backslash Y), (4p \mid X\backslash Y)$ is zero. Any zero-sum subsequence of length $4p$ must split, and by assumption any such splitting must be into pairs of zero-sum subsequences of length $2p$. Thus $(2p|X\backslash Y) = 0$ implies $(4p \mid X\backslash Y) = 0$, which is a contradiction. Letting $Z \subset X\backslash Y$ be any zero-sum subsequence with $|Z| = 2p$, putting Z and Y together yields a length $3p$ zero-sum subsequence. □

To further demonstrate the simplicity and clarity these results make possible, we will also rederive a known upper bound from [4].

Theorem 4. $\mathsf{s}_{2p}(C_p^3) \leq 6p-3$.

Proof. Pick X with $|X| = 6p-3$, and suppose $\beta_2 = 0$. Then by the splitting lemma $\beta_4 = 0$, so there must be a length p zero-sum subsequence $Y \subset X$. Note that $|X\backslash Y| = 5p-3$, and that $(2p \mid X\backslash Y) = (4p \mid X\backslash Y) = 0$, so by the independence lemma we can pick a zero-sum subsequence $Z \subset X\backslash Y$ with $|Z| = p$. Combining Z with Y then gives a contradiction. □

6 Avenues for Further Research

At this point in the paper, we reach the boundary of speculation and conjecture. The framework we have developed seems potentially suited to proving upper bounds on "diagonals". To illustrate why, consider the following new result.

Proof of Theorem 3. This is a lower bound, again by the construction using $4p-1$ copies of the zero vector and $p-1$ of each standard basis vector. We now show that it is also an upper bound. Let $|X| = 8p-4$, and suppose $\beta_4 = 0$. By the independence lemma at most two of $\beta_1, \ldots, \beta_7$ are zero. Further, any zero-sum subsequence of length $7p$ splits.

If $\beta_2 > 0$, pick a zero-sum subsequence Y of length $2p$. Then $|X\backslash Y| = 6p-4$, so applying the independence lemma gives that at most two of $(p \mid X\backslash Y), (2p \mid X\backslash Y), (3p \mid X\backslash Y), (5p \mid X\backslash Y)$ are zero. If $X\backslash Y$ has a zero-sum subsequence of length $2p$, we may combine it with Y to obtain a length $4p$ zero-sum subsequence, giving a contradiction. Thus $(2p \mid X\backslash Y) = 0$. This implies any length $5p$

zero-sum subsequence in $X\backslash Y$ splits into zero-sum subsequences with lengths $3p, p, p$, but combining the two p length subsequences contradicts $(2p \mid X\backslash Y) = 0$. Thus $(5p \mid X\backslash Y) = 0$, so $X\backslash Y$ contains some zero-sum subsequence Q with $|Q| = p$.

Set $W = X\backslash(Y \cup Q)$, so that $|W| = 5p - 4$. At most one of $(p \mid W), (3p \mid W)$ are zero. If we have a length p zero-sum subsequence in W, combining it with Q gives a length $2p$ zero-sum subsequence in $X\backslash Y$, which is a contradiction, but if we have a length $3p$ zero-sum subsequence, combining it with Q gives a length $4p$ zero-sum subsequence, which is also a contradiction.

Our starting assumption here was that $\beta_2 > 0$, so we now know that $\beta_2 = 0$. The splitting lemma immediately gives us that $\beta_5 = 0$, so the remaining β_i are non-zero. In particular, there exists a length $7p$ zero-sum subsequence. This splits, but the only possible splitting is into subsequences of lengths $3p, 3p$ and p. Combining the last two gives a length $4p$ zero-sum subsequence, which is a contradiction. □

Note that here, in Theorem 2, and in the primordial result $\mathsf{s}_p(C_p) = 2p - 1$, as well as for other values not mentioned in this paper, we have that $\mathsf{s}_{np}(C_p^n) = 2np - n$ (although our methods require $p > n$). In fact this is known as a lower bound generally [4]. When holding p fixed and laying out values of $\mathsf{s}_{kp}(C_p^r)$ in an array according to k and r, these are diagonal entries.

We have reason to suspect that for a fixed r, the least k for which $\mathsf{s}_{kp}(C_p^r)$ is amenable to our methods will be either r or not much less than r. The idea is as follows.

Definition 6. *A zero-sum subsequence Y of X is atomic if it is of such a length that it does not necessarily split, and $|Y|$ is a multiple of p. We also call lengths atomic if subsequences of that length are atomic, and β_i atomic if they correspond to atomic lengths.*

Given a sequence X in C_p^r, we know that $\beta_1, \ldots, \beta_r$ are atomic. The hypothesis that X contains a kp length zero-sum subsequence is checked by setting $\beta_k = 0$ and searching for a contradiction. If a zero-sum subsequence of X splits, certainly it can be split entirely into atomic subsequences Y_1, Y_2, Y_ℓ. However, no combination of these subsequences can have total length summing to kp. Thus, the β_m that allows sequences to split in the fewest ways when set to zero should intuitively be the largest atomic β_i, which is to say, β_k. In other words, the least k for which our methods reach their peak applicability when attempting to determine $\mathsf{s}_{kp}(C_p^r)$ should just be $k = r$.

It is widely believed that the phase transition of 1 occurs precisely at the diagonal. If the pattern continues, and can be proven by examination of congruences, the result would provide a strong indication that this conjecture is correct.

A Relation Matrix Generator

```
#include <iostream>

using namespace std;

int main(int argc, char* argv[]) {
    int k;
    try{
        k = stoi(argv[1]);
    } catch (...) {
```

```
    cout << "Error: argument must be an integer" << endl;
    return 0;
}
cout << "Calculating for |X| = " << k << "p - 1." << endl << endl;

int numRows = k - 3;
int numCols = k;
int relationMatrix[numRows][numCols];

//Construct Pascal's triangle
for(int row = 0; row < numRows; row++) {
    for(int col = numCols - 1; col >= 0; col--) { //Build from right
        if(row == 0 || row == numCols - 1 - col){ //First row is ones
            relationMatrix[row][col] = 1;
        } else if (col > numCols - row - 1){ //Zeros in the bottom right
            relationMatrix[row][col] = 0;
        } else { //Sum of the 'above' elements otherwise
            relationMatrix[row][col] = abs(relationMatrix[row][col + 1])
                + abs(relationMatrix[row - 1][col + 1]);
        }
        relationMatrix[row][col] *= col % 2 == 0 ? 1 : -1;
    }
}

if(argc == 2) {
    //Row reduction - easy because the matrix is upper triangular
    for(int row = 0; row < numRows; row++) {
        int pivot = numCols - row - 1; //Location of the 1 to keep
        for(int col = pivot - 1; col > 2; col--) { //Always two free variables
            int curVal = relationMatrix[row][col];
            int multiplier = curVal * (col % 2 == 0 ? 1 : -1);
            for(int i = 0; i < pivot; i++) {
                relationMatrix[row][i] -=  multiplier *
                    relationMatrix[numCols - 1 - col][i];
            }
        }
        if(relationMatrix[row][pivot] == -1) { //Make the pivot always +1
            for(int col = 0; col <= pivot; col++)
                relationMatrix[row][col] *= -1;
        }
    }
}

//Print matrix
for(int col = 0; col < numCols; col++) { //Header
    cout << "b_" << col << "\t";
```

```
    }
    cout << endl;
    for(int col = 0; col < numCols; col++) { //Separator
        cout << "--------";
    }
    cout << endl;
    for(int row = 0; row < numRows; row++) { //Values
        for(int col = 0; col < numCols; col++) {
            cout << relationMatrix[row][col] << "\t";
        }
        cout << endl;
    }
    cout << endl;
}
```

References

[1] Noga Alon. *The probabilistic method*. Wiley, Hoboken, N.J, 2008.

[2] P Erdős, A Ginzburg, and A Ziv. Theorem in the additive number theory. *Bull. Research Council Israel*, 10:41–43, 1961.

[3] W. Gao, D. Han, J. Peng, and F. Sun. On zero-sum subsequences of length k exp (G). *J. Combin. Theory*, 125(August):240–253, 2014.

[4] Weidong Gao and Ravindranathan Thangadurai. On zero-sum sequences of prescribed length. *Aequationes Mathematicae*, 72:201–212, 2006.

[5] Heiko Harborth. Ein Extremalproblem fur Gitterpunkte. *J. Reine Angew. Math.*, pages 356–360, 1973.

[6] Michael Hua, Samuel Freedman, M Hua, D Capodilupo, S B Damelin, S Freedman, J Sun, and M Yu. The Truncated , Upper-Triangle Pascal Matrix Linear Independence and Results in Matroid Theory. (August):1–4, 2015.

[7] Silke Kubertin. Zero-sums of length kq in $\mathbb{Z}_q^d$. *Acta Arithmetica*, 2005.

[8] Eric Naslund. Exponential bounds for the Erdős Ginzburg Ziv constant. 2017.

[9] Christian Reiher. On Kemnitz' conjecture concerning lattice-points in the plane. In *Ramanujan Journal*, 2007.

[10] Lajos Rónyai. On a Conjecture of Kemnitz. *Combinatorica*, 20(4):569–573, apr 2000.