

UPPER TAIL LARGE DEVIATIONS FOR ARITHMETIC PROGRESSIONS IN A RANDOM SET

BHASWAR B. BHATTACHARYA, SHIRSHENDU GANGULY, XUANCHENG SHAO, AND YUFEI ZHAO

ABSTRACT. Let X_k denote the number of k -term arithmetic progressions in a random subset of $\mathbb{Z}/N\mathbb{Z}$ or $\{1, \dots, N\}$ where every element is included independently with probability p . We determine the asymptotics of $\log \mathbb{P}(X_k \geq (1 + \delta)\mathbb{E}X_k)$ (also known as the large deviation rate) where $p \rightarrow 0$ with $p \geq N^{-c_k}$ for some constant $c_k > 0$, which answers a question of Chatterjee and Dembo. The proofs rely on the recent nonlinear large deviation principle of Eldan, which improved on earlier results of Chatterjee and Dembo. Our results complement those of Warnke, who used completely different methods to estimate, for the full range of p , the large deviation rate up to a constant factor.

CONTENTS

1.	Introduction	1
2.	Statements of results	3
3.	Gaussian width and non-linear large deviations	7
4.	Bounds on Gaussian width	11
5.	Variational problem at the macroscopic scale	16
6.	Variational problem at the microscopic scale	21
7.	Replica symmetry	25
8.	Maximizing the number of k -APs	27
	Appendix A. Proofs of further bounds on Gaussian widths	31
	References	32

1. INTRODUCTION

Let X_k denote the number of k -term arithmetic progressions (k -AP) in the random set Ω_p , where Ω is taken to be either $\mathbb{Z}/N\mathbb{Z}$ or $[N] := \{1, \dots, N\}$ throughout this paper, and Ω_p denotes the random subset of Ω where every element is included independently with probability p . The *upper tail problem* asks to estimate the probability that X_k significantly exceeds its expectation. This problem has received some interest over the years [8, 25, 31]. More generally, the problem of computing tail probabilities of a sum of weakly dependent random variables has a long and interesting history [2, 6, 8, 9, 10, 13, 14, 22, 23, 24, 25, 26, 28, 27, 30, 31]. There has been many exciting recent developments, particularly in the setting of random graphs [6, 8, 14], where one is interested in the concentration of the number of triangles in

SG was supported by a Miller Research Fellowship.

XS was supported by a Glasstone Research Fellowship.

YZ was supported by an Esmée Fairbairn Junior Research Fellowship at New College, Oxford and NSF Award DMS-1362326.

an Erdős–Rényi random graph $\mathcal{G}(N, p)$. See Chatterjee’s recent survey [7] and the references therein for an introduction to recent developments on large deviations in random graphs.

The recent work of Warnke [31] settles the question of the asymptotic order of $\mathbb{P}(X_k \geq (1 + \delta)\mathbb{E}X_k)$ when $\Omega = [N]$. Warnke [31] shows that for fixed $\delta > 0$ and $k \geq 3$, there exists constants $c, C > 0$ (depending only on k) such that

$$p^{C\sqrt{\delta}Np^{k/2}} \leq \mathbb{P}(X_k \geq (1 + \delta)\mathbb{E}X_k) \leq p^{c\sqrt{\delta}Np^{k/2}}, \quad (1.1)$$

as long as $p = p_N \geq (\log N/N)^{1/(k-1)}$ and p is bounded away from 1. Prior to Warnke’s work, the best upper bound [25] was $\mathbb{P}(X_k \geq (1 + \delta)\mathbb{E}X_k) \leq e^{-c_\delta p^{k/2}N}$, for some constant $c_\delta > 0$ depending on k and δ . However, the natural question of precise asymptotics still remained open.

The main result of this paper shows for every $k \geq 3$, fixed $\delta > 0$, and $\Omega = [N]$, if $p = p_N \geq N^{-\frac{1}{6k(k-1)}} \log N$ and $p \rightarrow 0$, then, as $N \rightarrow \infty$,

$$\mathbb{P}(X_k \geq (1 + \delta)\mathbb{E}X_k) = p^{(1+o(1))\sqrt{\delta}Np^{k/2}}. \quad (1.2)$$

The lower bound to the probability can be seen by forcing an interval of length $(1 + o(1))\sqrt{\delta}p^{k/2}N$ to be present in Ω , so that it generates the extra $\delta\mathbb{E}X_k$ many k -APs as desired. For the special case of $k = 3$, which was also treated in [8], methods in [15] combined with Fourier analysis allow us to take $p \geq N^{-1/18} \log N$, improving on the $p \geq N^{-1/162}(\log N)^{34/162}$ hypothesis in [8]. For $k \geq 4$, (1.2) is the first large deviation result for k -AP counts allowing p to decay as N^{-c} , thereby answering a question posed by Chatterjee and Dembo [8, Section 1.8] and improving on Warnke’s result in the appropriate regime.¹

The proofs rely on the powerful nonlinear large deviation principle (LDP) developed by Chatterjee and Dembo [8], which was recently improved by Eldan [15] using different methods, namely stochastic control theory. These LDPs reduce the determination of the large deviation rate (i.e., asymptotics of log-probability) in many combinatorial problems to a natural variational problem involving entropies. For the problem of upper tails of subgraph counts in a sparse random graph, the corresponding variational problem was recently solved [2, 28]. For arithmetic progressions, Chatterjee and Dembo were able to verify the hypotheses of their LDP for 3-term arithmetic progressions but not longer ones [8, Section 1.5]. More recently, Eldan [15] provided a different, but related, set of hypotheses for his LDP, involving the supremum of an associated Gaussian process (see Theorem 3.1). We prove the necessary bounds to apply Eldan’s LDP for arithmetic progressions of arbitrary fixed length. We remark that similar arguments can also be used to verify Chatterjee and Dembo’s LDP hypotheses for arithmetic progressions of any fixed length. Some interesting open problems in additive combinatorics arise in the analysis of this Gaussian process (see Section 4).

After establishing the LDP, we solve the corresponding variational problem. Here, there are two regimes, below, depending on how δ decays to zero compared to p .

- (1) In the *macroscopic* (large δ) regime, where $\delta^{-3}p^{k-2}(\log(1/p))^2 \rightarrow 0$, the solution of the variational problem reduces to the extremal problem of maximizing the number of k -APs in a set of given size, which was solved by Green and Sisask [18] for 3-APs, and extended to k -APs in Theorem 2.4. The solution to this extremal problem is

¹In a previous arXiv version of this paper, we proved the result with p decaying extremely slowly, as our previous proof depended on the heavy-powered inverse theorem for Gowers uniformity norms due to Green, Tao, and Ziegler [20].

attained by an interval. The case of fixed $\delta > 0$, namely the asymptotic (1.2), belongs to this regime.

- (2) In the *microscopic* (small δ) regime, where $\delta^{-3}p^{k-2}(\log(1/p))^2 \rightarrow \infty$, the variational problem exhibits a rather different qualitative behavior compared to the previous case. We show that

$$\mathbb{P}(X_k \geq (1 + \delta)\mathbb{E}X_k) = e^{-(c+o(1))\delta^2 Np},$$

for some explicit constant $c > 0$ depending on k and Ω .

Note that in both these regimes above we require $p = p_N \rightarrow 0$. If both p and δ are fixed, the situation is quite different, and we report some partial results for this setting in Section 7.

It is worth comparing these results for arithmetic progressions to the corresponding results for triangles in a random graph. Let X_{K_3} denote the number of copies of K_3 in $\mathcal{G}(N, p)$. It was shown in [28] (again relying on [8], which was recently improved in [15]) that for $p = p_N \rightarrow 0$ with $p \geq N^{-1/18} \log N$ and fixed $\delta > 0$, one has

$$\mathbb{P}(X_{K_3} \geq (1 + \delta)\mathbb{E}X_{K_3}) = p^{(1+o(1)) \min\{\delta^{2/3}/2, \delta/3\} N^2 p^2}.$$

This result was extended in [2], which determined the upper tail large deviation rate of X_H for every graph H . The extra complexity in the above expression, as compared to (1.2), arises from the dichotomy of methods of generating many extra triangles: we can either force a clique to be present, or force a small subset of vertices to be connected to all other vertices.

2. STATEMENTS OF RESULTS

2.1. Notation. We recall some standard asymptotic notations. For two nonnegative sequences $(f_n)_{n \geq 1}$ and $(g_n)_{n \geq 1}$, $f_n \lesssim g_n$ means $f_n = O(g_n)$; $f_n \sim g_n$ means $f_n = (1 + o(1))g_n$; and $f_n \asymp g_n$ means $f_n = \Theta(g_n)$, i.e., $f_n \lesssim g_n \lesssim f_n$. Subscripts in the above notation, for example, $O_\square(\cdot)$, $\lesssim_\square$, $\asymp_\square$ denote that the hidden constants may depend on the subscripted parameters. We always treat k (as in k -AP) as a constant, and the dependence of the hidden constants on k is always implicitly assumed and may be suppressed in the asymptotic notation.

For any set A in some ambient abelian group (in this paper the ambient group will always be either $\mathbb{Z}$ or $\mathbb{Z}/N\mathbb{Z}$) and $k \geq 3$, we write $T_k(A)$ to denote the number of pairs (a, b) of elements in the ambient group such that $a, a + b, a + 2b, \dots, a + (k - 1)b \in A$. Note that every non-trivial k -AP is counted twice, and every trivial k -AP (i.e., $b = 0$) is counted once. It will be convenient to state our results in terms of $T_k(A)$.

For $p \in (0, 1)$ and a subset Ω in the ambient group, denote by $\Omega_p \subset \Omega$ the random set obtained by independently including each element in Ω with probability p . Throughout the paper, we will consider two settings:

- (1) $\Omega = [N] = \{1, \dots, N\}$, and
- (2) $\Omega = \mathbb{Z}/N\mathbb{Z}$.

The ambient group is $\mathbb{Z}$ and $\mathbb{Z}/N\mathbb{Z}$ in the two respective settings. Note that, as long as $p^{k-1}N \rightarrow \infty$, it is easy to see that $\mathbb{E}T_k(\Omega_p) = p^k(T_k(\Omega) - |\Omega|) + p|\Omega| \sim p^k T_k(\Omega)$. In this paper we are interested in the *upper tail probability*, $\mathbb{P}(T_k(\Omega_p) \geq (1 + \delta)\mathbb{E}T_k(\Omega_p))$, when $p = p_N \rightarrow 0$, as $N \rightarrow \infty$ ($\delta = \delta_N$ may also depend on N).

The relative entropy function with respect to Bernoulli(p) is denoted

$$I_p(x) := x \log \frac{x}{p} + (1 - x) \log \frac{1 - x}{1 - p}.$$

Finally, denote the weighted k -AP count of a function $f : \Omega \rightarrow \mathbb{R}$ by

$$T_k(f) := \sum_{a,b} f(a)f(a+b) \cdots f(a+(k-1)b). \quad (2.1)$$

Here a and b each range over all elements of the ambient group (either $\mathbb{Z}$ or $\mathbb{Z}/N\mathbb{Z}$) such that $\{a, a+b, \dots, a+(k-1)b\} \subset \Omega$. By convention, when $\Omega = [N]$, we set $f(x) = 0$ for all $x \notin \Omega$. Note that $T_k(A) = T_k(\mathbf{1}_A)$, where $\mathbf{1}_A$ is the indicator function of A .

2.2. Large deviation principle. Let us write

$$\phi_p^{(k,\Omega)}(\delta) := \inf_{f: \Omega \rightarrow [0,1]} \left\{ \sum_{a \in \Omega} I_p(f(a)) : T_k(f) \geq (1+\delta)p^k T_k(\Omega) \right\} \quad (2.2)$$

for the natural large deviations variational problem for upper tails of k -AP counts. We will establish in Section 3.1 the following LDP for k -APs, via Eldan's LDP [15].

Theorem 2.1. *Fix $k \geq 3$. Let $\Omega = [N]$ or $\mathbb{Z}/N\mathbb{Z}$. Let $p = p_N$ be bounded away from 1, and $\delta = \delta_N > 0$ with $\delta = O(1)$ such that*

$$\min\{\delta p^k, \delta^2 p\} \geq N^{-\frac{1}{6(k-1)}} \log N. \quad (2.3)$$

Then, as $N \rightarrow \infty$,

$$-\log \mathbb{P}(T_k(\Omega_p) \geq (1+\delta)\mathbb{E}T_k(\Omega_p)) = (1+o(1))\phi_p^{(k,\Omega)}(\delta + o(\delta)). \quad (2.4)$$

Furthermore, for $k = 3$, the right-hand side of (2.3) can be relaxed to $N^{-1/6}(\log N)^{7/6}$; for $k = 4$, it can be relaxed to $N^{-1/12}(\log N)^{13/12}$.

Remark. For fixed $\delta > 0$, the theorem requires p to decay slower than $N^{-\frac{1}{6k(k-1)}}$. Very recently, Briët and Gopi [5] improved the exponent from $\frac{1}{6k(k-1)}$ to $\frac{1}{6k\lceil(k-1)/2\rceil}$, which improves our result for all $k \geq 5$. It remains an open problem to extend the range of validity of p . In comparison, Warnke's asymptotics (1.1) on the order of log-probability holds for all $p \geq (\log N/N)^{1/(k-1)}$.

In the above theorem, δ is allowed to decay as a function of N , and there is a qualitative change in the behavior of $\phi_p^{(k,\Omega)}(\delta)$ depending on how quickly δ decays compared to p . Drawing a parallel from statistical physics², we refer to the two regimes by:

- *Macroscopic scale:* when δ is “large”, namely when $\delta^{-3}p^{k-2}(\log(1/p))^2 \rightarrow 0$; and
- *Microscopic scale:* when δ is “small”, namely when $\delta^{-3}p^{k-2}(\log(1/p))^2 \rightarrow \infty$.

2.3. Macroscopic scale. In the macroscopic scale, $\delta^{-3}p^{k-2}(\log(1/p))^2 \rightarrow 0$. Here, it might be helpful to think of δ as a constant or tending to zero “slowly” compared to p . We will establish, in Section 5, the following asymptotic solution to the variational problem (2.2) in this regime.

²The terms macroscopic/microscopic appear in many contexts in the statistical mechanics literature, and is generically used to describe large/small scale behaviors, respectively. However, the exact definitions vary depending on the problem in question.

Theorem 2.2. *Fix $k \geq 3$. Let $\Omega = [N]$ or $\mathbb{Z}/N\mathbb{Z}$, and in the latter case assume that N is prime. Let $p = p_N \rightarrow 0$ and $\delta = \delta_N > 0$ be such that $\delta = O(1)$, $\delta p^k N^2 \rightarrow \infty$, and $\delta^{-3} p^{k-2} (\log(1/p))^2 \rightarrow 0$. Then, as $N \rightarrow \infty$,*

$$\begin{aligned} \phi_p^{(k,\Omega)}(\delta) &= (1 + o(1)) \sqrt{(k-1)\delta p^k T_k(\Omega) \log(1/p)} \\ &= \begin{cases} (1 + o(1)) \sqrt{\delta p^{k/2} N \log(1/p)} & \text{if } \Omega = [N], \\ (1 + o(1)) \sqrt{(k-1)\delta p^{k/2} N \log(1/p)} & \text{if } \Omega = \mathbb{Z}/N\mathbb{Z} \text{ with prime } N. \end{cases} \end{aligned}$$

Remark. In the case of constant δ , say, while the above solution to the variation problem needs only $p^k N^2 \rightarrow \infty$, a stronger condition $p^{k-1} N \rightarrow \infty$ (implying $\mathbb{E}T_k(\Omega_p) \sim p^k T_k(\Omega)$) is necessary even just for the concentration of the random variable $T_k(\Omega_p)$.

We prove Theorem 2.2 by first reducing the variational problem to an extremal problem in additive combinatorics, namely that of determining the size of the smallest subset of Ω with a given number of k -APs, or equivalently, the maximum number of k -APs in a subset of Ω of a given size.

Proposition 2.3. *Under the same hypotheses as Theorem 2.2 (except that N is not required to be prime in the case $\Omega = \mathbb{Z}/N\mathbb{Z}$), as $N \rightarrow \infty$,*

$$\phi_p^{(k,\Omega)}(\delta) = (1 + o(1)) \log(1/p) \cdot \min_{S \subset \Omega} \left\{ |S| : T_k(S) \geq \delta p^k T_k(\Omega) \right\}. \quad (2.5)$$

The number of k -APs in a set of given (sufficiently small) size is always maximized by an interval, as stated precisely below. The theorem below was proved for 3-APs by Green and Sisask [18] and extended to all k -APs in Section 8.

Theorem 2.4. *Fix a positive integer $k \geq 3$. There exists some constant $c_k > 0$ such that the following statement holds. Let $A \subset \mathbb{Z}$ be a subset with $|A| = n$, or $A \subset \mathbb{Z}/N\mathbb{Z}$ with N prime and $|A| = n \leq c_k N$. Then $T_k(A) \leq T_k([n])$.*

After some algebra one easily obtains

$$T_k([n]) = \frac{n^2}{k-1} - \frac{r^2}{k-1} + r,$$

where $r \in \{1, 2, \dots, k-1\}$ is chosen such that $n \equiv r \pmod{k-1}$. In particular we have

$$\frac{n^2}{k-1} \leq T_k([n]) \leq \frac{n^2}{k-1} + \frac{1}{4}(k-1). \quad (2.6)$$

From this formula we can easily deduce Theorem 2.2 from Proposition 2.3.

Combining with the large deviation principle Theorem 2.1, we obtain the following corollary on the large deviation rate for upper tails of k -AP counts.

Corollary 2.5. *Fix $k \geq 3$. Let $\Omega = [N]$ or $\mathbb{Z}/N\mathbb{Z}$, and in the latter case assume that N is prime. Let $p = p_N \rightarrow 0$ and $\delta = \delta_N > 0$ be such that $\delta = O(1)$, $\delta^{-3} p^{k-2} (\log(1/p))^2 \rightarrow 0$, and*

$$\delta p^k \geq N^{-\frac{1}{6(k-1)}} \log N \quad (2.7)$$

(for $k = 3$, the right-hand side can be relaxed to $N^{-1/6} (\log N)^{7/6}$; for $k = 4$, it can be relaxed to $N^{-1/12} (\log N)^{13/12}$). Then, as $N \rightarrow \infty$, the random variable $X_k = T_k(\Omega_p)$ satisfies

$$\frac{-\log \mathbb{P}(X_k \geq (1 + \delta) \mathbb{E}X_k)}{\sqrt{\delta p^{k/2} N \log(1/p)}} = \begin{cases} 1 + o(1) & \text{if } \Omega = [N], \\ \sqrt{k-1} + o(1) & \text{if } \Omega = \mathbb{Z}/N\mathbb{Z} \text{ with prime } N. \end{cases}$$

Remark. When $\Omega = \mathbb{Z}/N\mathbb{Z}$, it is standard to assume $N \rightarrow \infty$ along the primes to avoid torsion issues. Without the primality assumption, the leading constant in the large deviations rate function could depend on the subsequence along which N goes to infinity (a similar issue was discussed in [12]). For example, when $N = N_1 N_2$, the maximum number of k -APs in $\mathbb{Z}/N\mathbb{Z}$ in a set of size N_1 is given by $A = N_2 \mathbb{Z}/N\mathbb{Z}$, which has $T_k(A) = N_1^2$, more than $T_k([N_1]) = N_1^2/(k-1) + O(1)$. Thus, as a consequence of Proposition 2.3, we could have a different constant in Corollary 2.5 for $\Omega = \mathbb{Z}/N\mathbb{Z}$ if $N \rightarrow \infty$ along some sequence other than the primes.

2.4. Microscopic scale. In the microscopic scale, $\delta^{-3} p^{k-2} (\log(1/p))^2 \rightarrow \infty$. Here $\delta = \delta_N$ is thought as tending to zero relatively quickly compared to $p = p_N$. We will establish, in Section 6, the following asymptotic solution to the variational problem (2.2) in this regime.

Theorem 2.6. *Fix $k \geq 3$. Let $\Omega = \mathbb{Z}/N\mathbb{Z}$ or $[N]$. Let $p = p_N \in (0, 1)$ and $\delta = \delta_N > 0$ be such that $p \rightarrow 0$ and $\delta^{-3} p^{k-2} (\log(1/p))^2 \rightarrow \infty$. Then, as $N \rightarrow \infty$,*

$$\phi_p^{(k, \Omega)}(\delta) = \begin{cases} \left(\frac{1}{2k^2} + o(1) \right) \delta^2 N p & \text{if } \Omega = \mathbb{Z}/N\mathbb{Z}, \\ \left(\frac{1}{2\gamma_k} + o(1) \right) \delta^2 N p & \text{if } \Omega = [N], \end{cases}$$

where

$$\gamma_k = \frac{4}{3} \left(k + \sum_{0 \leq i < j < k} \frac{(k-1)^2 - i^2 - (k-1-j)^2}{(k-1-i)j} \right). \quad (2.8)$$

Remark. The first few values of γ_k are $\gamma_3 = 28/3$, $\gamma_4 = 17$, and $\gamma_5 = 718/27$. We are not aware of a closed-form expression for γ_k . However, one always has $\gamma_k \geq k^2$, and asymptotically $\lim_{k \rightarrow \infty} \gamma_k/k^2 = (30 - 2\pi^2)/9 \approx 1.14$.

Combining with the large deviation principle Theorem 2.1, we obtain the following corollary.

Corollary 2.7. *Fix $k \geq 3$. Let $\Omega = [N]$ or $\mathbb{Z}/N\mathbb{Z}$, and in the latter case assume that N is prime. Let $p = p_N \rightarrow 0$ and $\delta = \delta_N > 0$ be such that $\delta = O(1)$, $\delta^{-3} p^{k-2} (\log(1/p))^2 \rightarrow \infty$, and*

$$\min\{\delta p^k, \delta^2 p\} \geq N^{-\frac{1}{6(k-1)}} \log N$$

(for $k = 3$, the right-hand side can be relaxed to $N^{-1/6} (\log N)^{7/6}$; for $k = 4$, it can be relaxed to $N^{-1/12} (\log N)^{13/12}$). Then, as $N \rightarrow \infty$, the random variable $X_k = T_k(\Omega_p)$ satisfies

$$\frac{-\log \mathbb{P}(X_k \geq (1 + \delta) \mathbb{E} X_k)}{\delta^2 N p} = \begin{cases} 1/(2k^2) + o(1) & \text{if } \Omega = \mathbb{Z}/N\mathbb{Z}, \\ 1/(2\gamma_k) + o(1) & \text{if } \Omega = [N], \end{cases}$$

where γ_k is defined in (2.8).

The upper bound on $\phi_p^{(k, \Omega)}$ in Theorem 2.6 for $\Omega = \mathbb{Z}/N\mathbb{Z}$ is obtained by taking the constant function on Ω with value $p(1 + \delta)^{1/k}$, which turns out to be tight asymptotically to the first order. This behavior, where the solution to the variational problem is obtained by a constant function, at least asymptotically, suggests that the reason for many k -APs in the microscopic scale is a uniform boost in the density of the set, and such phenomena are referred to in the literature as *replica symmetry* [10]. (Admittedly we are somewhat abusing terminology here, as replica symmetry in previous works [10, 27, 32] on random graphs refer to setting of constant p and δ). In Section 7 we record some partial results on

replica symmetry for constant p and δ for k -APs. On the contrary, in the macroscopic scale, many k -APs are created by a smaller set arranged in a special structure, for example, an interval, and this is referred as *replica symmetry breaking*.

When $\Omega = [N]$, in the microscopic scale, the asymptotically optimal solution to the variational problem turns out not to be a constant function, but rather, a function that assigns each $a \in [N]$ to a number proportional to the number of k -APs in $[N]$ containing a . This is due to the asymmetry of the elements in $[N]$, as the elements in the middle bulk are contained in more k -APs than those in the fringe. Even though the constant function does not asymptotically minimize the variational problem in this setting, the solution nevertheless exhibits some features of replica symmetry (by analogy to the $\mathbb{Z}/N\mathbb{Z}$ setting). We find this new phenomenon interesting, as we are not aware of analogous results in the random graph setting.

3. GAUSSIAN WIDTH AND NON-LINEAR LARGE DEVIATIONS

In this section we apply Eldan's non-linear large deviation principle [15] to k -AP upper tails, reducing the large deviation rate problem to a variational problem. The proof relies on bounding the Gaussian width of a set of gradients, which will be done in Section 4.

3.1. Eldan's LDP. We start with a short discussion of Eldan's [15] result (adapted to our setting). For any $K \subset \mathbb{R}^N$ define the *Gaussian width* of K by

$$\text{GW}(K) := \mathbb{E} \left[\sup_{x \in K} \langle x, Z \rangle \right]$$

where the expectation is taken over $Z \sim \text{Normal}(0, I_N)$, a standard Gaussian random vector in $\mathbb{R}^N$.

For any function $F: \{0, 1\}^N \rightarrow \mathbb{R}$, define its discrete derivatives by

$$\text{DiscDerv}_i F(y) = F(y_1, \dots, y_{i-1}, 1, y_{i+1}, \dots, y_N) - F(y_1, \dots, y_{i-1}, 0, y_{i+1}, \dots, y_N)$$

for any $i \in [N]$ and $y = (y_1, y_2, \dots, y_N) \in \{0, 1\}^N$, and its discrete gradient by

$$\text{DiscGrad } F(y) = (\text{DiscDerv}_1 F(y), \dots, \text{DiscDerv}_N F(y)).$$

A key quantity is the Gaussian width of the set of all discrete gradients of F , which we denote by

$$\text{DiscGW}(F) := \text{GW} \left(\{ \text{DiscGrad } F(y) : y \in \{0, 1\}^N \} \cup \{0\} \right). \quad (3.1)$$

Define the *discrete Lipschitz constant* of F by

$$\text{DLip}(f) = \max_{i \in [N], y \in \{0, 1\}^N} \text{DiscDerv}_i f(y).$$

Improving an earlier result of Chatterjee and Dembo [8], Eldan [15] proved a large deviation principle for general non-linear functions $F: \{0, 1\}^N \rightarrow \mathbb{R}$ satisfying certain hypotheses on its set of discrete gradients. The large deviation rate is given in terms of the natural variational problem:

$$\varphi_p^F(t) := \inf_{y \in [0, 1]^N} \left\{ \sum_{i=1}^N I_p(y_i) : \mathbb{E} F(Y) \geq tN \right\}, \quad (3.2)$$

where the expectation is taken with respect to a random vector $Y = (Y_1, Y_2, \dots, Y_N)$ with $Y_i \sim \text{Bernoulli}(y_i)$ independently for every $i \in [N]$.

Theorem 3.1 (Eldan [15]). *Let $X = (X_1, X_2, \dots, X_N) \in \{0, 1\}^N$ be a random vector with i.i.d. $X_i \sim \text{Bernoulli}(p)$. Given a function $F: \{0, 1\}^N \rightarrow \mathbb{R}$, for every $t, \varepsilon \in \mathbb{R}$ with $0 < \varepsilon < \varphi_p^F(t - \varepsilon)/N$, we have*

$$\log \mathbb{P}(F(X) \geq tN) \leq -\varphi_p^F(t - \varepsilon) \left(1 - \frac{6L(\log N)^{1/6}}{N^{1/3}}\right)$$

with

$$L = \frac{1}{\varepsilon} \left(2\text{DLip}(F) + \frac{1}{\varepsilon\sqrt{N}}\text{DLip}(F) + |\log(p(1-p))|\right)^{2/3} \left(\text{DiscGW}(F) + \frac{1}{\varepsilon}\text{DLip}(F)^2\right)^{1/3}. \quad (3.3)$$

Moreover, whenever the assumption $\frac{2}{N\varepsilon^2}\text{DLip}(F)^2 \leq \frac{1}{2}$ holds, the following lower bound holds:

$$\log \mathbb{P}(F(X) \geq (t - \varepsilon)N) \geq -\varphi_p^F(t) \left(1 + \frac{2}{N\varepsilon^2}\text{DLip}(F)^2\right) - \log 10.$$

Denote the usual gradient and partial derivatives of $F: \mathbb{R}^N \rightarrow \mathbb{R}$ by

$$\nabla F := (\partial_1 F, \dots, \partial_N F).$$

Define

$$\text{GW}(F) := \text{GW}(\{\nabla F(y) : y \in \{0, 1\}^N\} \cup \{0\}) \quad (3.4)$$

for the continuous analog of $\text{DiscGW}(F)$ from (3.1). These two quantities differ only negligibly in our applications.

Lemma 3.2. *For any twice-differentiable $F: \mathbb{R}^N \rightarrow \mathbb{R}$, we have*

$$|\text{DiscGW}(F) - \text{GW}(F)| \lesssim \sum_{i=1}^N \sup_{x \in [0, 1]^N} |\partial_{ii} F(x)|,$$

where $\partial_{ij} F = \partial^2 F / \partial x_i \partial x_j$ denotes the (i, j) -th partial derivative of F and $\text{DiscGW}(F)$ is defined by considering the restriction of F to $\{0, 1\}^N$.

Proof. Applying the intermediate value theorem (twice), we have

$$|\text{DiscDerv}_i F(y) - \partial_i F(y)| \leq \sup_{x \in [0, 1]^N} |\partial_{ii} F(x)|$$

for any $y \in \{0, 1\}^N$. Thus for any $Z = (Z_1, \dots, Z_N) \in \mathbb{R}^N$ and any $y \in \{0, 1\}^N$,

$$|\langle \text{DiscGrad } F(y), Z \rangle - \langle \nabla F(y), Z \rangle| \leq \sum_i \sup_{x \in [0, 1]^N} |\partial_{ii} F(x)| |Z_i|.$$

The result follows by first taking the supremum over y , and then taking an expectation over $Z \sim \text{Normal}(0, I_N)$ and using $\mathbb{E}|Z_i| = O(1)$. $\square$

3.2. LDP for k -AP. Now, we apply Theorem 3.1 to derive a large deviation principle for k -AP counts, conditioned on bounds for the Gaussian width of the gradients of the k -AP counting function.

By viewing points in $\mathbb{R}^\Omega$ as functions $\Omega \rightarrow \mathbb{R}$, the previously defined k -AP functional T_k can be viewed as a function on $\mathbb{R}^\Omega$ by

$$T_k(y) = \sum_{a, b} y_a y_{a+b} y_{a+2b} \dots y_{a+(k-1)b}, \quad y \in \mathbb{R}^\Omega.$$

In the case $\Omega = [N]$, the indices a and b both range over $\mathbb{Z}$, and we set $y_a = 0$ if $a \notin [N]$. In the case $\Omega = \mathbb{Z}/N\mathbb{Z}$, the indices a and b both range over $\mathbb{Z}/N\mathbb{Z}$ and the indices of y are taken mod N . Recall the definition (2.2) of the variational problem for upper tails of k -AP counts, reproduced here:

$$\phi_p^{(k,\Omega)}(\delta) := \inf_{y \in [0,1]^\Omega} \left\{ \sum_{a \in \Omega} I_p(y_a) : T_k(y) \geq (1 + \delta)p^k T_k(\Omega) \right\}. \quad (3.5)$$

We will apply Theorem 3.1 to the function $F = T_k/N$. The relevant Gaussian width is

$$\text{GW}(T_k/N) = \text{GW} \left(\left\{ \frac{1}{N} \nabla T_k(y) : y \in \{0, 1\}^\Omega \right\} \right).$$

Note that we do not need to include the origin in the definition since $\nabla T_k(0) = 0$. We have the “trivial” bounds:

$$\sqrt{N} \lesssim \text{GW}(T_k/N) \lesssim N. \quad (3.6)$$

The lower bound comes from considering the constant vector $y = (1, \dots, 1)$, and the upper bound comes from noting that $\frac{1}{N} \nabla T_k(y)$ is coordinatewise $O(1)$ for all $y \in \{0, 1\}^\Omega$.

The main result of this section is the following proposition, showing that any power-saving improvement to the trivial upper bound to $\text{GW}(T_k/N)$ leads to a large deviation principle allowing the probability p to decay as N^{-c} . Combining it with bounds on the Gaussian width (to be proved in the next section) gives Theorem 2.1.

Proposition 3.3. *Fix $k \geq 3$. Suppose we have real constants σ, τ such that*

$$\text{GW}(T_k/N) = O(N^{1-\sigma}(\log N)^\tau). \quad (3.7)$$

Let $p = p_N$ be bounded away from 1, and $\delta = \delta_N > 0$ be such that $\delta = O(1)$ and

$$N^{-\sigma/3}(\log N)^{\tau/3+1} \lesssim \min\{\delta p^k, \phi_p^{(k,\Omega)}(\delta/2)/N\}. \quad (3.8)$$

Then

$$-\log \mathbb{P}(T_k(\Omega_p) \geq (1 + \delta)\mathbb{E}T_k(\Omega_p)) = (1 + o(1))\phi_p^{(k,\Omega)}(\delta + o(\delta)).$$

As long as one can prove an estimate (3.7) on Gaussian width with $\sigma > 0$, we can allow p to decay as N^{-c} for some constant $c > 0$. In Theorem 4.1, we show that one can take $\sigma = 1/(2(k-1))$ and $\tau = 0$ (with better bounds for $k = 3, 4$). From the asymptotic solutions to the variational problems (Theorem 2.2 and 2.6), and noting that $\phi_p^{(\Omega,k)}(\delta)$ is monotonic in δ , we have, as long as $\delta p^k N^2 \rightarrow \infty$,

$$\phi_p^{(\Omega,k)}(\delta/2)/N \asymp \min\{\sqrt{\delta} p^{k/2} \log(1/p), \delta^2 p\}.$$

Combining these asymptotics, we see that hypothesis (3.8) translates into hypothesis (2.3) in Theorem 2.1, and hence Proposition 3.3 implies Theorem 2.1.

In the rest of this section, we prove Proposition 3.3. We first prove some easy estimates on the various quantities that appear in Theorem 3.1. Recall the definitions of $\text{DiscGW}(F)$ and $\text{GW}(F)$ from (3.1) and (3.4).

Lemma 3.4. *For any $k \geq 3$, we have*

$$\text{DLip}(T_k/N) = O(1)$$

and

$$\text{DiscGW}(T_k/N) = \text{GW}(T_k/N) + O(1).$$

Proof. The first claim follows from noting that in $T_k(y) = \sum_{a,b} y_a y_{a+b} \cdots y_{a+(k-1)b}$, every variable y_a appears in $O(N)$ terms.

The second claim follows from Lemma 3.2, as $\partial_{aa} T_k(y)$ is uniformly bounded for all $y \in [0, 1]^\Omega$ and all $a \in \Omega$. $\square$

Proof of Proposition 3.3. We apply Theorem 3.1 for $F = T_k/N$. Set

$$\varepsilon = N^{-1/3} (\log N)^{11/12} \text{GW}(T_k/N)^{1/3} \lesssim N^{-\sigma/3} (\log N)^{\tau/3+11/12}.$$

By (3.8),

$$\varepsilon = o(\delta p^k) \quad \text{and} \quad \varepsilon = o(\phi_p^{(k,\Omega)}(\delta/2)/N). \quad (3.9)$$

Note that $\sigma \leq 1/2$ due to the lower bound $\text{GW}(T_k/N) \gtrsim \sqrt{N}$ in (3.6). So in particular, $N\varepsilon^2 \rightarrow \infty$. Also, $\log(1/p) = O(\log N)$ by (3.8).

Recall L from (3.3). Using Lemma 3.4 and earlier estimates, we have

$$L \lesssim \varepsilon^{-1} (\log N)^{2/3} \text{GW}(T_k/N)^{1/3}.$$

Thus, as $N \rightarrow \infty$,

$$LN^{-1/3} (\log N)^{1/6} \lesssim \varepsilon^{-1} N^{-1/3} (\log N)^{5/6} \text{GW}(T_k/N)^{1/3} = (\log N)^{-1/12} \rightarrow 0.$$

Let $Y = (Y_1, \dots, Y_N)$ be a random vector with $Y_i \sim \text{Bernoulli}(y_i)$ independently for all $i \in [N]$. Then

$$\mathbb{E}T_k(Y) = T_k(y) + O(N) = T_k(y) + o(\delta p^k N^2).$$

The discrepancy $O(N)$ comes from terms in $T_k(y)$ where some y_i may appear more than once. Setting

$$t = (1 + \delta)p^k T_k(\Omega)/N^2,$$

we see that

$$\mathbb{E}T_k(Y)/N \geq (t + o(\delta p^k))N \quad \text{is equivalent to} \quad T_k(y) \geq (1 + \delta + o(\delta))p^k T_k(\Omega).$$

Comparing the definition of $\varphi_p^F(t)$ from (3.2) for $F = T_k/N$ and $\phi_p^{(k,\Omega)}(\delta)$ from (3.5) (and noting that $\varphi_p^F(t)$ is non-decreasing in t and $\phi_p^{(k,\Omega)}(\delta)$ is non-decreasing in δ), we obtain

$$\varphi_p^{T_k/N}(t \pm \varepsilon) = \phi_p^{(k,\Omega)}(\delta + o(\delta)).$$

The hypothesis $0 < \varepsilon < \frac{1}{N}\varphi_p^{T_k/N}(t - \varepsilon)$ in Theorem 3.1 is also satisfied due to (3.9).

Applying Theorem 3.1, we obtain the upper bound to the log-probability

$$\log \mathbb{P}(T_k(\Omega_p) \geq (1 + \delta)p^k T_k(\Omega)) \leq -(1 - o(1))\varphi_p^{T_k/N}(t - \varepsilon) \sim -\phi_p^{(k,\Omega)}(\delta - o(\delta)),$$

as well as the lower bound (changing t to $t + \varepsilon$ when applying Theorem 3.1),

$$\log \mathbb{P}(T_k(\Omega_p) \geq (1 + \delta)p^k T_k(\Omega)) \geq -(1 - o(1))\varphi_p^{T_k/N}(t + \varepsilon) - O(1) \sim -\phi_p^{(k,\Omega)}(\delta + o(\delta)).$$

Combining the upper and lower bounds, and recalling that $\mathbb{E}T_k(\Omega_p) \sim p^k T_k(\Omega)$, the result follows. $\square$

4. BOUNDS ON GAUSSIAN WIDTH

In this section, we establish bounds on the Gaussian width of the set of gradients of T_k . These bounds can be used in Proposition 3.3 from the previous section to deduce Theorem 2.1 on the LDP for APs.

Our main result of this section is stated below. Recall from Section 3.1 that

$$\begin{aligned} \text{GW}(T_k/N) &= \text{GW}(\{\nabla T_k(y)/N : y \in \{0, 1\}^\Omega\}) \\ &= \mathbb{E}_{Z \sim \text{Normal}(0, I_N)} \sup_{y \in \{0, 1\}^\Omega} \left\langle \frac{1}{N} \nabla T_k(y), Z \right\rangle. \end{aligned} \quad (4.1)$$

Theorem 4.1. *For any fixed $k \geq 3$,*

$$\text{GW}(T_k/N) = O(N^{1 - \frac{1}{2(k-1)}}). \quad (4.2)$$

Furthermore, for $k = 3$, the bound can be tightened to

$$\text{GW}(T_3/N) = \Theta(\sqrt{N \log N}).$$

For $k = 4$, the bound can be improved to

$$\text{GW}(T_4/N) = O(N^{3/4}(\log N)^{1/4}).$$

Remark. After a preprint of our paper had appeared, Briët and Gopi [5] improved the bound to $\text{GW}(T_k/N) = O(N^{1 - \frac{1}{2\lceil (k-1)/2 \rceil}}(\log N)^{1/2})$ for all $k \geq 5$.

We have an easy lower bound $\text{GW}(T_k/N) \gtrsim \sqrt{N}$ deduced by taking the constant vector $y = (1, \dots, 1)$ in (4.1). We conjecture that it is essentially tight.

Conjecture 4.2. *For any fixed $k \geq 4$,*

$$\text{GW}(T_k/N) = \sqrt{N}(\log N)^{O(1)}.$$

For the proof of Theorem 4.1, we go back to viewing T_k as an operator on functions $\Omega \rightarrow \mathbb{R}$ (as opposed to a function on points in $\mathbb{R}^\Omega$). We define a multilinear version of T_k by setting, for $f_0, \dots, f_{k-1}: \Omega \rightarrow \mathbb{R}$,

$$T_k(f_0, \dots, f_{k-1}) := \sum_{a, b} f_0(a) f_1(a+b) \cdots f_{k-1}(a + (k-1)b),$$

so that for $f: \Omega \rightarrow \mathbb{R}$,

$$T_k(f) = T_k(\underbrace{f, f, \dots, f}_{k \text{ times}}).$$

We identify points in $\mathbb{C}^\Omega$ with functions $\Omega \rightarrow \mathbb{C}$ and maintain the notation $\langle \cdot, \cdot \rangle$ for inner products, so that for $f, h: \Omega \rightarrow \mathbb{C}$,

$$\langle f, h \rangle := \sum_{a \in \Omega} f(a) \overline{h(a)}.$$

The gradient ∇T_k of T_k maps a function $f: \Omega \rightarrow \mathbb{R}$ to the function $\nabla T_k(f): \Omega \rightarrow \mathbb{R}$ defined by

$$\nabla T_k(f)(a) = \sum_b \sum_{i=0}^{k-1} \prod_{\substack{0 \leq j \leq k-1 \\ j \neq i}} f(a + (j-i)b), \quad a \in \Omega.$$

Since T_k is multi-linear, for any $f, h: \Omega \rightarrow \mathbb{R}$,

$$\langle \nabla T_k(f), h \rangle = \sum_{j=0}^{k-1} T_k(\underbrace{f, f, \dots, f}_j \text{ times}, h, \underbrace{f, f, \dots, f}_{k-1-j} \text{ times}). \quad (4.3)$$

4.1. 3-APs and Fourier analysis. Here, we prove the claim in Proposition 4.7 for $k = 3$ using Fourier analysis. It will be easier to work in the setting $\Omega = \mathbb{Z}/N\mathbb{Z}$. The corresponding bounds for $\Omega = [N] \subset \mathbb{Z}$ can be easily derived by embedding $[N]$ in $\mathbb{Z}/N'\mathbb{Z}$ for some larger $N' \in [2N, 3N]$ so that 3-APs in $[N] \subset \mathbb{Z}/N'\mathbb{Z}$ do not wrap-around zero in the cyclic group.

Given $f: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{C}$, define its discrete Fourier transform by

$$\widehat{f}(r) = \frac{1}{N} \sum_{a \in \mathbb{Z}/N\mathbb{Z}} f(a) \omega^{-ar}, \quad r \in \mathbb{Z}/N\mathbb{Z},$$

where $\omega = e^{2\pi i/N}$. The inverse transform is given by

$$f(a) = \sum_{r \in \mathbb{Z}/N\mathbb{Z}} \widehat{f}(r) \omega^{ar}.$$

The following standard identity relates T_3 with the Fourier transform.

Lemma 4.3. *For $f, g, h: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{R}$,*

$$\frac{1}{N^2} T_3(f, g, h) = \sum_{r \in \mathbb{Z}/N\mathbb{Z}} \widehat{f}(r) \widehat{g}(-2r) \widehat{h}(r).$$

Proof. Expanding the left-hand side using the inverse transform, we have

$$\begin{aligned} T_3(f, g, h) &= \sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} f(a) g(a+b) h(a+2b) \\ &= \sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} \sum_{r, s, t \in \mathbb{Z}/N\mathbb{Z}} \widehat{f}(r) \widehat{g}(s) \widehat{h}(t) \omega^{ar+(a+b)s+(a+2b)t} \\ &= N^2 \sum_{r \in \mathbb{Z}/N\mathbb{Z}} \widehat{f}(r) \widehat{g}(-2r) \widehat{h}(r), \end{aligned}$$

where the final step follows from noting that

$$\sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} \omega^{ar+(a+b)s+(a+2b)t} = \begin{cases} N^2 & \text{if } r + s + t = 0 \text{ and } s + 2t = 0, \\ 0 & \text{otherwise.} \end{cases}$$

□

The above identity leads to the following bound, showing that T_3 is controlled by the Fourier transform of its inputs.

Lemma 4.4. *Let $f_0, f_1, f_2: \mathbb{Z}/N\mathbb{Z} \rightarrow [-1, 1]$. For each $i = 0, 1, 2$,*

$$\frac{1}{N^2} |T_3(f_0, f_1, f_2)| \lesssim \|f_i\|_\infty =: \max_r |\widehat{f}_i(r)|.$$

Remark. If N is odd, then the $\lesssim$ can be replaced by $\leq$.

Proof. Using Lemma 4.3 and the Cauchy–Schwarz inequality, we have

$$\begin{aligned}
\frac{1}{N^2} |T_3(f_0, f_1, f_2)| &\leq \|\widehat{f}_0\|_\infty \sum_{r \in \mathbb{Z}/N\mathbb{Z}} |\widehat{f}_1(-2r)| |\widehat{f}_2(r)| \\
&\leq \|\widehat{f}_0\|_\infty \left(\sum_r |\widehat{f}_1(-2r)|^2 \right)^{1/2} \left(\sum_r |\widehat{f}_2(r)|^2 \right)^{1/2} \\
&\lesssim \|\widehat{f}_0\|_\infty \left(\sum_r |\widehat{f}_1(r)|^2 \right)^{1/2} \left(\sum_r |\widehat{f}_2(r)|^2 \right)^{1/2} \\
&\leq \|\widehat{f}_0\|_\infty.
\end{aligned}$$

The final step follows from Parseval’s identity: $\sum_r |\widehat{f}_j(r)|^2 = \frac{1}{N} \sum_a |f_j(a)|^2 \leq 1$. The proofs for $i = 1, 2$ are analogous. $\square$

Proof of Theorem 4.1 for $k = 3$. For any $g: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{R}$, we have

$$\langle \nabla T_3(f), g \rangle = T_3(g, f, f) + T_3(f, g, f) + T_3(f, f, g),$$

so by Lemma 4.4,

$$\frac{1}{N} |\langle \nabla T_3(f), g \rangle| \lesssim N \max_r |\widehat{g}(r)|.$$

Now let g be a random function taking i.i.d. standard normal values. Then each $\widehat{g}(r)$ is a normally distributed complex number with $\mathbb{E}[|\widehat{g}(r)|^2] = 1/N$, since the Fourier transform is a unitary operator. Standard results about the supremum of Gaussian processes, e.g., Lemma 4.5 below, then gives $\mathbb{E} \sup_r |\widehat{g}(r)| \lesssim \sqrt{(\log N)/N}$. Thus

$$\text{GW}(T_3/N) = \mathbb{E} \sup_{f: \mathbb{Z}/N\mathbb{Z} \rightarrow \{0,1\}} \frac{1}{N} \langle \nabla T_3(f), g \rangle \lesssim N \mathbb{E} \sup_r |\widehat{g}(r)| \lesssim \sqrt{N \log N}.$$

The matching lower bound is proved in Appendix A. $\square$

See Appendix A for proof of the bound $\text{GW}(T_4/N) = O(N^{3/4}(\log N)^{1/4})$ in Theorem 4.1, which extends the above Fourier analytic technique.

4.2. k -APs and the Chinese remainder theorem. Our strategy for proving Theorem 4.1 is to show that the set of gradients $\nabla T_k(f)/N$ over all $f: \Omega \rightarrow \{0,1\}$ is contained in the convex hull of a small set of bounded functions. We start with a standard bound on Gaussian width of sets. The proof is included for completeness.

Lemma 4.5 (Small sets have small Gaussian widths). *If $S \subset [-1, 1]^N$, then*

$$\text{GW}(S) = O(\sqrt{N \log |S|}).$$

Proof. We may assume that $|S| \geq 2$. We have the following standard tail bound for $Z \sim \text{Normal}(0, \sigma^2)$: $\mathbb{P}(Z > t) \leq e^{-t^2/(2\sigma^2)}$. With $Z \sim \text{Normal}(0, I_N)$ a standard Gaussian vector in $\mathbb{R}^N$, one has $\langle y, Z \rangle \sim \text{Normal}(0, |y|^2)$ for every $y \in \mathbb{R}^N$. Thus $\mathbb{P}(\langle y, Z \rangle > t) \leq e^{-t^2/(2|y|^2)} \leq e^{-t^2/(2N)}$ for every $y \in [-1, 1]^N$. So for any $t > 0$, we have by the union bound

$$\mathbb{P}(\sup_{y \in S} \langle y, Z \rangle > t) \leq |S| e^{-t^2/(2N)}.$$

Set $u = 10\sqrt{N \log |S|}$. Then

$$\begin{aligned} \text{GW}(S) &= \mathbb{E}_{\mathbb{Z}} \sup_{y \in S} \langle y, Z \rangle \leq u + \int_u^\infty \mathbb{P}(\sup_{y \in S} \langle y, Z \rangle > t) dt \\ &\leq u + \int_u^\infty |S| e^{-t^2/(2N)} dt = O(\sqrt{N \log |S|}). \quad \square \end{aligned}$$

Corollary 4.6. *If $\{\nabla T_k(f)/N : f : \Omega \rightarrow \{0, 1\}\}$ is contained in the convex hull of a set S of uniformly bounded functions³ on Ω , then $\text{GW}(T_k/N) = O(\sqrt{N \log |S|})$.*

We will prove the following bounds on the size of S , which imply the bound $\text{GW}(T_k/N) = O(N^{1-\frac{1}{2(k-1)}})$ in Theorem 4.1.

Proposition 4.7. *Fix $k \geq 3$. The set $\{\nabla T_k(f)/N : f : \Omega \rightarrow \{0, 1\}\}$ is contained in the convex hull of a set S of uniformly bounded functions on Ω , where $|S| = \exp(O(N^{1-1/(k-1)}))$.*

We can conjecture that the bound on $|S|$ can be improved to $e^{N^{o(1)}}$, or perhaps even the following stronger bound, which would imply Conjecture 4.2.

Conjecture 4.8. *In Proposition 4.7, for every fixed $k \geq 4$, one can have $|S| = e^{(\log N)^{O(1)}}$.*

Remark. For $k = 3$, one can have $|S| = O(N)$ by considering the Fourier basis and using Lemma 4.4. For the first open case $k = 4$, intuition from the theory of *higher-order Fourier analysis* (e.g., [29]) suggests that perhaps it suffices to take the set of “quadratic characters”, i.e., functions of the form $a \mapsto e^{2\pi i q(a)}$ where $q(a)$ behaves quadratically (very loosely speaking). More generally, perhaps we can take a set of *nilsequences* [19]. However, this method is currently incapable of proving the conjecture due to poor quantitative dependencies in the theory of higher order Fourier analysis. See [16, Section 4.1] for the analogous problem in the language of ergodic averages.

For the proof of Proposition 4.7, it will be easier to work in the setting $\Omega = [N] \subset \mathbb{Z}$. The proof can be easily adapted to work for $\Omega = \mathbb{Z}/N\mathbb{Z}$ by chopping $\mathbb{Z}/N\mathbb{Z}$ into a bounded number of intervals and analyzing their contributions separately.

The argument we present here is essentially the same as that appearing in [17, Section 4], which in turn is motivated by the random sampling idea in [11]. Roughly speaking, we can obtain a good estimate of $\nabla T_k(f)$ by knowing f only on $A \subset [N]$, where A is a random subset of slightly more than $N^{1-1/(k-1)}$ elements. Thus any gradient $\nabla T_k(f)$ can be well approximated by one of $2^{|A|}$ many possible functions.⁴ Actually one can achieve a similar effect (more simply and with better bounds) deterministically by partitioning $[N]$ modulo $k-1$ distinct primes of sizes roughly $N^{1/(k-1)}$ each, which is how we shall proceed. In this approach, the Chinese remainder theorem play a role similar to independence for random variables.

Proof of Proposition 4.7. We may assume that N is large. To analyze $\nabla T_k(f)$, we make the following definition. For $c_1, \dots, c_{k-1} \in \mathbb{Z}$, let $\Phi(c_1, \dots, c_{k-1})$ be the set of all functions

³Here, and in the sequel, uniformly bounded functions refer to those functions f with $|f(a)| = O(1)$ point-wise.

⁴This random sampling argument can also be used in conjunction with Chatterjee and Dembo’s results [8] to obtain a large deviation principle for k -APs, allowing the probability p to decay as N^{-c_k} for some constant $c_k > 0$, though with a smaller c_k than what is shown here.

$F: [N] \rightarrow \mathbb{R}$ of the form

$$F(a) = \frac{1}{N} \sum_{b \in \mathbb{Z}} f(a + c_1 b) f(a + c_2 b) \dots f(a + c_{k-1} b), \quad a \in [N] \quad (4.4)$$

for some $f: [N] \rightarrow \{0, 1\}$ (set $f(a) = 0$ if $a \notin [N]$). It suffices to show that $\Phi(c_1, \dots, c_{k-1})$ is contained in the convex hull of $\exp(O(N^{1-1/(k-1)}))$ many bounded functions, whenever $\{c_1, \dots, c_{k-1}\}$ is any of

$$\{1, 2, \dots, k-1\}, \{-1, 1, \dots, k-2\}, \dots, \{-(k-1), \dots, -2, -1\}.$$

For clarity, we assume that $(c_1, \dots, c_k) = (1, \dots, k-1)$, as the other cases are similar.

Pick $k-1$ distinct primes $q_1, \dots, q_{k-1} \in [(N/2)^{\frac{1}{k-1}}, N^{\frac{1}{k-1}}]$ (such a choice always exist when N is large due to bounds on large gaps between primes [1]). We have $q_1 q_2 \dots q_{k-1} \in [N/2, N]$. Write

$$f = \sum_{r_i \in \mathbb{Z}/q_i \mathbb{Z}} f_{r_i + q_i \mathbb{Z}} \quad (4.5)$$

where

$$f_{r_i + q_i \mathbb{Z}}(a) = \begin{cases} f(a) & \text{if } a \equiv r_i \pmod{q_i}, \\ 0 & \text{if } a \not\equiv r_i \pmod{q_i}, \end{cases}$$

is the restriction of f to the residue class $r_i \pmod{q_i}$. By using (4.5) to partition the i -th factor $f(a + ib)$ in each term in (4.4), we obtain

$$F(a) = \frac{1}{N} \sum_{b \in \mathbb{Z}} \prod_{i=1}^{k-1} \sum_{r_i \in \mathbb{Z}/q_i \mathbb{Z}} f_{r_i + q_i \mathbb{Z}}(a + ib) = \frac{1}{q_1 \dots q_k} \sum_{r_1, \dots, r_{k-1}} v_{r_1, \dots, r_{k-1}; f}(a)$$

where the sum is taken over $r_1 \in \mathbb{Z}/q_1 \mathbb{Z}, \dots, r_{k-1} \in \mathbb{Z}/q_{k-1} \mathbb{Z}$, and

$$v_{r_1, \dots, r_{k-1}; f}(a) := \frac{q_1 q_2 \dots q_{k-1}}{N} \sum_{b \in \mathbb{Z}} \prod_{i=1}^{k-1} f_{r_i + q_i \mathbb{Z}}(a + ib). \quad (4.6)$$

We see that F lies in the convex hull of

$$V = \{v_{r_1, \dots, r_{k-1}; f} : f \text{ is } \{0, 1\}\text{-valued}, r_1 \in \mathbb{Z}/q_1 \mathbb{Z}, \dots, r_{k-1} \in \mathbb{Z}/q_{k-1} \mathbb{Z}\}.$$

For fixed i and r_i , the number of possibilities for $f_{r_i + q_i \mathbb{Z}}$ is at most $2^{N/q_i+1}$ as f ranges over $\{0, 1\}$ -valued functions. It follows that

$$|V| \leq q_1 q_2 \dots q_{k-1} \prod_{i=1}^{k-1} 2^{N/q_i+1} = 2^{O(N^{1-1/(k-1)})},$$

by our choice of q_i . It remains to show that all functions in V are bounded. Indeed, for any fixed $a \in [N]$, the simultaneous congruence conditions $a + ib \equiv r_i \pmod{q_i}$, for $1 \leq i \leq k-1$, in the unknown $b \in [-N, N]$ have $O(N/(q_1 q_2 \dots q_{k-1}) + 1)$ solutions by the Chinese Remainder Theorem. Thus the number of nonzero summands in (4.6) is $O(N/(q_1 q_2 \dots q_{k-1}) + 1)$, and it follows that

$$|v_{r_1, \dots, r_{k-1}; f}(a)| \lesssim 1 + \frac{q_1 q_2 \dots q_{k-1}}{N} \lesssim 1,$$

as desired. This completes the proof. $\square$

This completes the proof of Theorem 4.1, which bounds $\text{GW}(T_k/N)$, other than the matching lower bound $\text{GW}(T_3/N) \gtrsim \sqrt{N \log N}$ and the improvement $\text{GW}(T_4/N) \lesssim N^{3/4}(\log N)^{1/4}$, whose proofs can be found in Appendix A.

5. VARIATIONAL PROBLEM AT THE MACROSCOPIC SCALE

The goal of this section is to prove Proposition 2.3, which reduces the entropic variational problem in the macroscopic scale, i.e., $\delta^{-3}p^{k-2}(\log(1/p))^2 \rightarrow 0$, to a corresponding extremal problem in additive combinatorics:

$$\phi_p^{(k,\Omega)}(\delta) = (1 + o(1)) \log(1/p) \cdot \min_{S \subset \Omega} \left\{ |S| : T_k(S) \geq \delta p^k T_k(\Omega) \right\} \quad (5.1)$$

(provided that $\delta = O(1)$ and $\delta p^k N^2 \rightarrow \infty$). Let us provide an overview of the proof strategy. The upper bound on $\phi_p^{(k,\Omega)}(\delta)$ follows by considering a function which takes value 1 on an interval and p elsewhere (Section 5.1). For the lower bound, suppose $T_k(f) \geq (1 + \delta)p^k T_k(\Omega)$. Write $f(a) = p + g(a)$. Let g' denote a function obtained from g by changing each $g(a)$ to 0 if $g(a)$ is already sufficiently close to zero (the exact threshold will be specified in the proof). It will be shown that $T_k(f) \approx p^k T_k(\Omega) + T_k(g')$, so that $T_k(g') \geq (1 - o(1))\delta p^k T_k(\Omega)$. For now, it is fine to pretend that g' is an indicator function of a set, so that we have a lower bound on the number of k -APs of the set. We will prove an extremal result on maximizing the number of k -APs of a set of given size. This will imply a lower bound on $\sum_{a \in \Omega} g'(a)$, thereby giving the desired lower bound on $\sum_a I_p(f(a)) \approx \sum_a g'(a) \log(1/p)$.

5.1. Proof of upper bound in Proposition 2.3. We begin by noting that the right-hand side expression of (5.1) does not change if δ is replaced by some $\delta' = \delta + o(\delta)$. This is equivalent to the fact that the maximum number of k -APs in a subset of Ω of size n does not change significantly if n is changed to $n + o(n)$. This is clear from the exact formula in Theorem 2.4 and (2.6) whenever the hypothesis of the theorem applies, or otherwise in general from the easy lemma below (applied with $n \rightarrow \infty$, $n \leq N$, and $s = o(n)$, so that $M_n \gtrsim n^2$ and $M_{n+o(n)} \sim M_n$).

Lemma 5.1. *Let $k \geq 3$ and $\Omega = \mathbb{Z}/N\mathbb{Z}$ or $[N]$. Let $M_n = \max_{A \subset \Omega: |A| \leq n} T_k(A)$. Then*

$$M_n \leq M_{n+s} \leq M_n + ks(n+s)$$

for all $n, s \geq 0$.

Proof. It is easy to see that $T_k(A \cup \{a\}) \leq T_k(A) + k|A| + 1$ by counting the number of new k -APs that are formed with the addition of a new element a to A . Thus $M_{n+1} \leq M_n + kn + 1$. The lemma follows by iterating this bound. $\square$

It is not too hard to prove that $\phi_p^{(k,\Omega)}(\delta)$ is at most the right-hand side quantity in (5.1). Take any $S \subset \Omega$ with $T_k(S) \geq (1 - p^k)^{-1} \delta p^k T_k(\Omega)$ (here we are implicitly changing the δ in the right-hand side of (5.1) to $\delta' = (1 - p^k)^{-1} \delta = \delta + o(\delta)$), and let f in the variational problem (2.2) be the function

$$f(a) = \begin{cases} 1 & \text{if } a \in S, \\ p & \text{if } a \notin S. \end{cases} \quad (5.2)$$

So that

$$T_k(f) \geq (1 - p^k) T_k(S) + p^k T_k(\Omega) \geq (1 + \delta) p^k T_k(\Omega),$$

and

$$\sum_{a \in \Omega} I_p(f(a)) = |S| \log(1/p).$$

This proves the upper bound in Proposition 2.3.

5.2. Proof of lower bound in Proposition 2.3. To begin with note that taking $S \subset \Omega$ to be an interval of size $\lceil \sqrt{k\delta' T_k(\Omega)} p^{k/2} \rceil$, where $\delta' = (1 - p^k)^{-1} \delta = \delta + o(\delta)$, ensures $T_k(S) \geq (1 - p^k)^{-1} \delta p^k T_k(\Omega)$. Then taking f as in (5.2) gives $T_k(f) \geq (1 + \delta) p^k T_k(\Omega)$ and $\sum_{a \in \Omega} I_p(f(a)) = \sqrt{k\delta' T_k(\Omega)} p^{k/2} \log(1/p) \leq 2\sqrt{k\delta} N p^{k/2} \log(1/p)$, for N large enough. Therefore, to show that the left-hand side of (2.5) is greater than its right-hand side, it suffices to restrict our attention to functions $f: \Omega \rightarrow [p, 1]$ that satisfy

$$\sum_{a \in \Omega} I_p(f(a)) \leq 2\sqrt{k\delta} N p^{k/2} \log(1/p) \quad \text{and} \quad T_k(f) \geq (1 + \delta) p^k T_k(\Omega) \quad (5.3)$$

(note that we can restrict the range of f to $[p, 1]$ since $I_p(\cdot)$ is decreasing in $[0, p]$). We will show that for such f ,

$$\sum_{a \in \Omega} I_p(f(a)) \geq (1 + o(1)) \log(1/p) \cdot \min_{S \subset \Omega} \left\{ |S| : T_k(S) \geq \delta p^k T_k(\Omega) \right\}.$$

We recall a useful asymptotic estimate of $I_p(\cdot)$ from [28, Lemma 3.3].

Lemma 5.2. *Let $p \rightarrow 0$, and $x = x(p) \in [0, 1 - p]$. If $x = o(p)$, then $I_p(p + x) \sim x^2/(2p)$. If $x/p \rightarrow \infty$, then $I_p(p + x) \sim x \log(x/p)$.*

Let $f(a) = p + g(a)$, where $g: \Omega \rightarrow [0, 1 - p]$. (When $\Omega = [N]$, we set both f and g to be zero outside $[N]$.) Note the following bounds on g : by convexity of $I_p(\cdot)$, (5.3), and Lemma 5.2,

$$\begin{aligned} I_p\left(p + \frac{1}{N} \sum_{a \in \Omega} g(a)\right) &= I_p\left(\frac{1}{N} \sum_{a \in \Omega} f(a)\right) \leq \frac{1}{N} \sum_{a \in \Omega} I_p(f(a)) \\ &\leq \sqrt{k\delta} p^{k/2} \log(1/p) \sim I_p\left(p + 2(k\delta)^{1/4} p^{(k+2)/4} \sqrt{\log(1/p)}\right), \end{aligned}$$

since $\delta^{1/4} p^{(k+2)/4} \sqrt{\log(1/p)} = o(p)$ whenever $\delta = O(1)$. Moreover, since $I_p(p + x)$ is increasing for $x \in [0, p - 1]$,

$$\frac{1}{N} \sum_{a \in \Omega} g(a) \lesssim \delta^{1/4} p^{(k+2)/4} \sqrt{\log(1/p)}. \quad (5.4)$$

The proof of the lower bound on $\sum_{a \in \Omega} I_p(f(a))$ proceeds via two-step thresholding on the function g . At each step, we choose some threshold τ and decompose g into its small and large components:

$$g = g_{\leq \tau} + g_{> \tau}.$$

Here $g_{\leq \tau}(a) := g(a) \mathbf{1}\{g(a) \leq \tau\}$, and $g_{> \tau} := g(a) \mathbf{1}\{g(a) > \tau\}$.

- (1) (*Thresholding*) First, we perform the decomposition with $\tau = p^{3/4}$ and show that the contribution to $T_k(f)$ from the small component $g_{\leq \tau}$ is negligible.
- (2) (*Bootstrapping*) Next, we bootstrap the argument in the first step and take a higher threshold $\tau = p^{o(1)}$.

The following lemma will be useful later.

Lemma 5.3. *Let $\Omega = \mathbb{Z}/N\mathbb{Z}$ or $[N]$. Let $x, y \in \{0, 1, \dots, k-1\}$ with $x \neq y$. For any $f: \Omega \rightarrow [0, 1]$, one has*

$$\sum_{a,b} f(a+xb)f(a+yb) \leq (k-1) \left(\sum_{a \in \Omega} f(a) \right)^2,$$

where the sum on the left-hand side is taken over all pairs of elements (a, b) in the ambient group such that $\{a+xb, a+yb\} \subset \Omega$.

Proof. The lemma follows from observing that after expanding the right-hand side, for any $c, d \in \Omega$, there are at most $k-1$ pairs of elements a, b in the ambient group such that $a+xb = c$ and $a+yb = d$. Indeed, subtracting the two equations gives $(x-y)b = c-d$. Since $0 < |x-y| \leq k-1$, the number of solutions for b is at most 1 when $\Omega = [N]$ and at most $(k-1)$ when $\Omega = \mathbb{Z}/N\mathbb{Z}$. $\square$

5.2.1. *The thresholding step.* In this section we formalize step (1) above. Recall that f is as in (5.3), and $f = p + g$, where $g: \Omega \rightarrow [0, 1]$ satisfies $0 \leq g \leq 1-p$. From (5.3) we know that

$$T_k(p + g_{\leq \tau} + g_{> \tau}) = T_k(f) \geq (1 + \delta)p^k T_k(\Omega).$$

The expression $T_k(p + g_{\leq \tau} + g_{> \tau})$, when written out as a sum, expands into a number of components. The following lemma shows that, with an appropriate choice of τ , the only non-negligible contributions are $T_k(p) = p^k T_k(\Omega)$ and $T_k(g_{> \tau})$.

Lemma 5.4. *Assume that $\delta^{-3}p^{k-2} \log^2(1/p) \rightarrow 0$, and $f = p + g: \Omega \rightarrow [p, 1]$ satisfies (5.3). Let $\tau = p^{3/4}$. Then*

$$T_k(g_{> \tau}) \geq (1 - o(1))\delta p^k T_k(\Omega) \quad (5.5)$$

and

$$\sum_{a \in \Omega} g_{> \tau}(a) \lesssim \sqrt{\delta} N p^{k/2}. \quad (5.6)$$

Remark. In the above lemma, one may take $\tau = p^s$ for any fixed $2/3 < s < 1$.

Proof. From Lemma 5.2 we have $I_p(p + g_{> \tau}(x)) \asymp g_{> \tau}(x) \log(1/p)$. Thus

$$\sum_{a \in \Omega} g_{> \tau}(a) \log(1/p) \asymp \sum_{a \in \Omega} I_p(p + g_{> \tau}(a)) \leq \sum_{a \in \Omega} I_p(f(a)) \lesssim \sqrt{\delta} N p^{k/2} \log(1/p),$$

where the final step uses (5.3). This gives us (5.6).

By expanding, we have

$$T_k(p + g_{\leq \tau} + g_{> \tau}) = \sum_{X,Y,Z} T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau}), \quad (5.7)$$

where the sum is over all ordered partitions (X, Y, Z) of the set $\{0, 1, \dots, k-1\}$, and

$$T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau}) := \sum_{a,b} p^{|X|} \prod_{y \in Y} g_{\leq \tau}(a+yb) \prod_{z \in Z} g_{> \tau}(a+zb). \quad (5.8)$$

Here the sum is taken over all pairs of elements (a, b) in the ambient group such that $\{a, a+b, \dots, a+(k-1)b\} \subset \Omega$. We say that $T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau})$ *contributes negligibly* to the sum (2.1) (or *negligible* for short) if $T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau}) = o(\delta N^2 p^k)$. We will show that all terms except for $(X, Y, Z) = (\{0, \dots, k-1\}, \emptyset, \emptyset)$ and $(\emptyset, \emptyset, \{0, \dots, k-1\})$ are negligible, i.e., the only non-negligible terms are $p^k T_k(\Omega)$ and $T_k(g_{> \tau})$. This would prove (5.5), due to the assumption $T_k(f) \geq (1 + \delta)p^k T_k(\Omega)$ from (5.3).

First, if $|Z| \geq 2$, then by Lemma 5.3 and (5.6),

$$T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau}) \leq (k-1)\tau^{|X \cup Y|} \left(\sum_{a \in \Omega} g_{> \tau}(a) \right)^2 \lesssim \tau^{|X \cup Y|} \delta N^2 p^k.$$

Therefore, if $|Z| \geq 2$, then the contribution from $T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau})$ is negligible unless $|X \cup Y| = 0$, i.e., $(X, Y, Z) = (\emptyset, \emptyset, \{0, \dots, k-1\})$.

Next, if $|X| = k-1$ and $|Y \cup Z| = 1$, then by (5.4),

$$T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau}) \leq kNp^{k-1} \sum_{a \in \Omega} g(a) \lesssim \delta^{1/4} N^2 p^{(5k-2)/4} \sqrt{\log(1/p)} = o(\delta N^2 p^k),$$

where in the final step we use the macroscopic scale assumption $\delta^{-3} p^{k-2} (\log(1/p))^2 \rightarrow 0$. Therefore these terms are negligible.

Finally, if $|Z| \leq 1$ and $|Y \cup Z| \geq 2$, then by (5.4) and Lemma 5.3,

$$T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau}) \leq (k-1)\tau^{k-2} \left(\sum_{a \in \Omega} g(a) \right)^2 \lesssim \tau^{k-2} N^2 \delta^{1/2} p^{(k+2)/2} \log(1/p) = o(\delta N^2 p^k),$$

where the last step holds due to $\tau = o(p^{2/3} (\log(1/p))^{-2/(3k-6)})$ and the macroscopic scale assumption on δ .

It follows from the above analysis that the only non-negligible contributions to the sum (5.7) are $(X, Y, Z) = (\{0, \dots, k-1\}, \emptyset, \emptyset)$ and $(\emptyset, \emptyset, \{0, \dots, k-1\})$, so that

$$T_k(f) = p^k T_k(\Omega) + T_k(g_{> \tau}) + o(\delta p^k N^2).$$

By the assumption $T_k(f) \geq (1+\delta)p^k T_k(\Omega)$ from (5.3), we obtain $T_k(g_{> \tau}) \geq (1-o(1))\delta p^k T_k(\Omega)$ as desired. $\square$

5.2.2. The bootstrapping step. Now, we strengthen Lemma 6.3 by replacing τ with any $\tau = o(1)$.

Lemma 5.5. *Assume that $\delta^{-3} p^{k-2} \log^2(1/p) \rightarrow 0$, and $f = p + g: \Omega \rightarrow [p, 1]$ satisfies (5.3). For any $\tau = o(1)$, we have*

$$T_k(g_{> \tau}) \geq (1-o(1))\delta p^k T_k(\Omega).$$

Proof. Observe that

$$T_k(g_{> p^{3/4}}) - T_k(g_{> \tau}) \leq k^2 \tau \left(\sum_{a \in \Omega} g_{> p^{3/4}}(a) \right)^2 \lesssim \tau \delta p^k N^2 = o(\delta p^k N^2). \quad (5.9)$$

To prove the first inequality, note that the left-hand side can be bounded above by a sum of k terms, where the j -th term (for $1 \leq j \leq k$) is itself the following sum

$$\sum_{a, b \in \Omega} g_1(a) g_2(a+b) \cdots g_k(a+(k-1)b)$$

where $g_j = g_{\leq \tau}$ and all other g_i 's are set to $g_{> p^{3/4}}$, and this sum can be bounded by $\tau(k-1)(\sum_{a \in \Omega} g_{> p^{3/4}})^2$ using Lemma 5.3. The second part of (5.9) follows from (5.6). By (5.5) we have $T_k(g_{> p^{3/4}}) \geq (1-o(1))\delta p^k T_k(\Omega)$, and thus by (5.9) we have $T_k(g_{> \tau}) \geq (1-o(1))\delta p^k T_k(\Omega)$ as well. $\square$

5.2.3. *Completing the proof of Proposition 2.3.* Assume that $f: \Omega \rightarrow [p, 1]$ satisfies (5.3). Let $f = p + g$, and choose some threshold τ satisfying $\tau = p^{o(1)} \rightarrow 0$ (e.g., $\tau = 1/\log(1/p)$ for concreteness). By Lemma 5.5, one has

$$T_k(g_{>\tau}) \geq (1 - o(1))\delta p^k T_k(\Omega).$$

This implies, by Lemma 5.6 below (used in the contrapositive in conjunction with Lemma 5.1),

$$\sum_{a \in \Omega} g_{>\tau}(a) \geq (1 + o(1)) \min_{S \subset \Omega} \left\{ |S| : T_k(S) \geq \delta p^k T_k(\Omega) \right\}.$$

By Lemma 5.2, $I_p(p + g_{>\tau}(a)) \sim g_{>\tau}(x) \log(1/p)$ since $\tau = p^{o(1)}$. Therefore,

$$\begin{aligned} \sum_{a \in \Omega} I_p(f(a)) &\geq \sum_{a \in \Omega} I_p(p + g_{>\tau}(a)) = (1 + o(1)) \log(1/p) \sum_{a \in \Omega} g_{>\tau}(a) \\ &\geq (1 + o(1)) \log(1/p) \min_{S \subset \Omega} \left\{ |S| : T_k(S) \geq \delta p^k T_k(\Omega) \right\}, \end{aligned}$$

thereby completing the proof of Proposition 2.3, modulo the following lemma, which says that the problem of maximizing the number of k -APs in a set remains roughly unchanged even if we allow the elements to be weighted.

Lemma 5.6. *Let $f: \Omega \rightarrow [0, 1]$ be such that $\sum_{a \in \Omega} f(a) = m$. Then*

$$T_k(f) \leq (1 + o(1)) \max_{A \subset \Omega: |A| \leq m} T_k(A) \quad (5.10)$$

provided that $m \rightarrow \infty$ as $N \rightarrow \infty$.

Proof. Let $M_s = \max_{A \subset \Omega: |A| \leq s} T_k(A)$. Let Ω_f be a random subset of Ω chosen by including element $a \in \Omega$ with probability $f(a)$ independently for all $a \in \Omega$. Note that $\mathbb{E}[|\Omega_f|] = m$, and for $a, b \in \Omega$,

$$f(a)f(a+b) \cdots f(a+(k-1)b) \leq \mathbb{P}(a, a+b, \dots, a+(k-1)b \in A).$$

(it is always an equality when the elements of the k -AP are distinct). This implies

$$T_k(f) \leq \mathbb{E} T_k(\Omega_f) \leq \sum_{s \geq 1} \mathbb{P}(|\Omega_f| = s) M_s.$$

For $s \leq m + m^{2/3}$, we bound $M_s \leq M_{m+m^{2/3}} = (1 + o(1))M_m$. And for $s > m + m^{2/3}$ we use the trivial bound $M_s \leq s^2$. Thus

$$\begin{aligned} T_k(f) &\leq (1 + o(1))M_m + 4m^2 \mathbb{P}(m + m^{2/3} < |A| \leq 2m) + \sum_{s > 2m} \mathbb{P}(|A| = s) s^2 \\ &\leq (1 + o(1))M_m + 4m^2 \mathbb{P}(|A| > m + m^{2/3}) + \sum_{s > 2m} \mathbb{P}(|A| \geq s) s^2 \\ &\leq (1 + o(1))M_m + 4m^2 e^{-m^{1/3}/3} + \sum_{s > 2m} s^2 e^{-s/6} \\ &= (1 + o(1))M_m, \end{aligned}$$

where the penultimate step uses Chernoff bound in the following form: if X is a sum of independent indicator random variables, and $\mathbb{E}X = \mu$, then for any $\delta > 0$, $\mathbb{P}(X \geq (1 + \delta)\mu) \leq e^{-\min\{\delta^2, \delta\}\mu/3}$. $\square$

This completes the proof of Proposition 2.3. Combining it with Theorem 2.4, which will be proved in Section 8, yields Theorem 2.2, the asymptotic solution to the variational problem in the macroscopic scale.

6. VARIATIONAL PROBLEM AT THE MICROSCOPIC SCALE

In this section, we prove Theorem 2.6, which solves the variational problem in the microscopic scale, i.e., $\delta^{-3}p^{k-2}\log^2(1/p) \rightarrow \infty$. The following theorem unifies the settings $\Omega = [N]$ and $\mathbb{Z}/N\mathbb{Z}$.

Theorem 6.1. *Fix $k \geq 3$. Let $\Omega = [N]$ or $\mathbb{Z}/N\mathbb{Z}$. Then for $p = p_N \rightarrow 0$ and $\delta = \delta_N > 0$ such that $\delta^{-3}p^{k-2}(\log(1/p))^2 \rightarrow \infty$, we have*

$$\phi_p^{(k,\Omega)}(\delta) = (1 + o(1)) \frac{\delta^2 T_k(\Omega)^2 p}{2 \sum_{a \in \Omega} \nu_a^2}, \quad (6.1)$$

where ν_a is the number of k -APs in Ω containing $a \in \Omega$ (the constant k -AP $a, a, \dots, a$ is counted k times), i.e., ν_a is the number of triples (x, y, j) where x, y are elements in the ambient group, $j \in \{0, 1, \dots, k-1\}$, so that $a = x + jy$ and $\{x, x+y, \dots, x+(k-1)y\} \subset \Omega$.

In Section 6.1, we prove Theorem 6.1, following similar thresholding strategy to the macroscopic setting. In Section 6.2, we then compute the rate formulae in Theorem 2.6.

6.1. Proof of Theorem 6.1. We begin by showing that $\phi_p^{(k,\Omega)}(\delta)$ is at most the right-hand side of (6.1). The claim follows from an explicit construction of a function f in the variational problem (2.2), as given by the following lemma.

Lemma 6.2. *Let $\Omega = [N]$ or $\mathbb{Z}/N\mathbb{Z}$, and let ν_a be defined as in Theorem 6.1. Define $g: \Omega \rightarrow [0, 1]$ by*

$$g(a) = \frac{\delta T_k(\Omega) \nu_a}{\sum_{a \in \Omega} \nu_a^2} p.$$

Then $f = p + g$ satisfies

$$I_p(f(a)) \sim \frac{\delta^2 T_k(\Omega)^2 p}{2 \sum_{a \in \Omega} \nu_a^2} \quad \text{and} \quad T_k(f) \geq (1 + \delta) p^k T_k(\Omega). \quad (6.2)$$

As a consequence, $\phi_p^{(k,\Omega)}(\delta) \leq (1 + o(1)) \frac{\delta^2 T_k(\Omega)^2 p}{2 \sum_{a \in \Omega} \nu_a^2}$.

Proof. We have $\nu_a \asymp N$, $\sum_{a \in \Omega} \nu_a^2 \asymp N^3$, so $g(a) \asymp \delta p = o(p)$, and thus by Lemma 5.2,

$$\sum_{a \in \Omega} I_p(f(a)) \sim \sum_{a \in \Omega} \frac{g(a)^2}{2p} \sim \frac{\delta^2 T_k(\Omega)^2 p}{2 \sum_{a \in \Omega} \nu_a^2}.$$

Next, by expanding we have

$$\begin{aligned} T_k(f) &= T_k(p + g) \geq p^k T_k(\Omega) + p^{k-1} \sum_{x,y} \sum_{j=0}^{k-1} g(x + jy) \\ &= p^k T_k(\Omega) + p^{k-1} \sum_{a \in \Omega} g(a) \nu_a = (1 + \delta) p^k T_k(\Omega), \end{aligned}$$

where (x, y) ranges over all pairs of elements in the ambient group such that $\{x, x+y, \dots, x+(k-1)y\} \subset \Omega$. This proves (6.2) and the upper bound on $\phi_p^{(k,\Omega)}(\delta)$ in Theorem 6.1. $\square$

Now we prove the lower bound on $\phi_p^{(k,\Omega)}(\delta)$. To begin with, using $\sum_{a \in \Omega} \nu_a = kT_k(\Omega)$ and the Cauchy–Schwarz inequality, we have

$$\phi_p^{(k,\Omega)}(\delta) \leq (1 + o(1)) \frac{\delta^2 T_k(\Omega)^2 p}{2 \sum_{a \in \Omega} \nu_a^2} \leq (1 + o(1)) \frac{\delta^2 p N}{2k^2}. \quad (6.3)$$

Therefore, to prove the lower bound on $\phi_p^{(k,\Omega)}(\delta)$ in (6.1), we can restrict our attention to functions $f = p + g: \Omega \rightarrow [p, 1]$ satisfying

$$\sum_{a \in \Omega} I_p(f(a)) \leq \delta^2 N p \quad \text{and} \quad T_k(f) \geq (1 + \delta) p^k T_k(\Omega). \quad (6.4)$$

Then by convexity, (6.4), and Lemma 5.2,

$$I_p \left(p + \frac{1}{N} \sum_{a \in \Omega} g(a) \right) \leq \frac{1}{N} \sum_{a \in \Omega} I_p(p + g(a)) \leq \delta^2 p \sim I_p(p + \sqrt{2} \delta p).$$

Since $I_p(p + x)$ is increasing for $x \in [0, 1 - p]$, we have

$$\sum_{a \in \Omega} g(a) \lesssim \delta p N. \quad (6.5)$$

The following lemma gives a lower bound on the weighted average of any function g satisfying $f = p + g: \Omega \rightarrow [p, 1]$ with $T_k(f) \geq (1 + \delta) p^k T_k(\Omega)$, in the microscopic regime.

Lemma 6.3. *Assume that $\delta^{-3} p^{k-2} \log^2(1/p) \rightarrow \infty$, and $f = p + g: \Omega \rightarrow [p, 1]$ with $T_k(f) \geq (1 + \delta) p^k T_k(\Omega)$. Then*

$$\sum_{a \in \Omega} g(a) \nu_a \geq (1 - o(1)) \delta p T_k(\Omega). \quad (6.6)$$

Proof. Set threshold $\tau = p^{3/4}$. As in Section 5.2.1, write $g = g_{\leq \tau} + g_{> \tau}$. As in Lemma 6.3, we have by Lemma 5.2 and (6.4),

$$\sum_{a \in \Omega} g_{> \tau}(a) \log(1/p) \asymp \sum_{a \in \Omega} I_p(p + g_{> \tau}(a)) \leq \sum_{a \in \Omega} I_p(f(a)) \leq \delta^2 N p, \quad (6.7)$$

thereby gaining an extra $\log(1/p)$ factor compared to (6.5).

By expanding, we have

$$T_k(p + g_{\leq \tau} + g_{> \tau}) = \sum_{X, Y, Z} T_{X, Y, Z}(p, g_{\leq \tau}, g_{> \tau}), \quad (6.8)$$

where $T_{X, Y, Z}(p, g_{\leq \tau}, g_{> \tau})$ is the same as earlier (5.8). We say that $T_{X, Y, Z}(p, g_{\leq \tau}, g_{> \tau})$ contributes negligibly to the sum (6.8) if $T_{X, Y, Z}(p, g_{\leq \tau}, g_{> \tau}) = o(\delta N^2 p^k)$. We will show that unless $|X| = k - 1$ or k , the term contributes negligibly.

Indeed, as in Section 5.2.1, if $|Z| \geq 2$, then by Lemma 5.3 and (6.7),

$$T_{X, Y, Z}(p, g_{\leq \tau}, g_{> \tau}) \leq (k - 1) \left(\sum_{a \in \Omega} g_{> \tau}(a) \right)^2 \lesssim \left(\frac{\delta^2 N p}{\log(1/p)} \right)^2 = o(\delta N^2 p^k),$$

where the final step uses the microscopic scale hypothesis $\delta^{-3} p^{k-2} \log^2(1/p) \rightarrow \infty$.

If $|Z| \leq 1$ and $|Y \cup Z| \geq 2$, then by Lemma 5.3 and (6.5),

$$T_{X,Y,Z}(p, g_{\leq \tau}, g_{> \tau}) \leq (k-1)\tau^{k-2} \left(\sum_{a \in \Omega} g(a) \right)^2 \lesssim \tau^{k-2} \delta^2 p^2 N^2 = o(\delta N^2 p^k),$$

where the last step holds due to $\tau = o(p^{2/3}(\log(1/p))^{-2/(3k-6)})$ and the microscopic scale hypothesis on δ .

This shows that the non-negligible contributions are those terms with $|X| = k-1$ or k . Hence

$$\begin{aligned} T_k(f) &= p^k T_k(\Omega) + p^{k-1} \sum_{x,y} \sum_{j=0}^{k-1} g(x+jy) + o(\delta p^k N^2) \\ &= p^k T_k(\Omega) + p^{k-1} \sum_{a \in \Omega} g(a) \nu_a + o(\delta p^k N^2), \end{aligned}$$

where the first sum runs over all pairs (x, y) of elements in the ambient group such that $\{x, x+y, \dots, x+(k-1)y\} \subset \Omega$. Since $T_k(f) \geq (1+\delta)p^k T_k(\Omega)$, we have $\sum_{a \in \Omega} g(a) \nu_a \geq (1-o(1))\delta p T_k(\Omega)$, as required. $\square$

In the final step of the proof of Theorem 6.1, we convert the lower bound on the weighted sum of g from the above lemma to a lower bound on the entropy $\sum_a I_p(p+g(a))$. We consider the two cases $\Omega = \mathbb{Z}/N\mathbb{Z}$ and $[N]$ separately.

Case 1: $\Omega = \mathbb{Z}/N\mathbb{Z}$. In this case, we have $\nu_a = kN$ for all $a \in \Omega$ and $T_k(\Omega) = N^2$. By convexity of $I_p(\cdot)$, (6.6), and Lemma 5.2,

$$\begin{aligned} \sum_{a \in \Omega} I_p(p+g(a)) &\geq N I_p\left(p + \frac{1}{N} \sum_{a \in \Omega} g(a)\right) \\ &\geq N I_p\left(p + (1-o(1))\frac{\delta p}{k}\right) \sim \frac{\delta^2 N p}{2k^2}. \end{aligned}$$

This combined with Lemma 6.2, proves Theorem 6.1, when $\Omega = \mathbb{Z}/N\mathbb{Z}$.

Case 2: $\Omega = [N]$. In this case, the quantities ν_a are unequal, and the solution requires an extra step. We use the estimate⁵

$$I_p(p+x) \geq (1+o(1))p^{0.3}x \quad \text{uniformly for all } x \in [p^{1.1}, 1-p]. \quad (6.9)$$

It follows by (6.9) and (6.4) that

$$\begin{aligned} \sum_{a \in \Omega} g_{> p^{1.1}}(a) &\leq (1+o(1))p^{-0.3} \sum_{a \in \Omega} I_p(p+g_{> p^{1.1}}(a)) \\ &\leq (1+o(1))p^{-0.3} \sum_{a \in \Omega} I_p(f(a)) \leq (1+o(1))\delta^2 p^{0.7} N = o(\delta p N) \end{aligned}$$

where the final step, $\delta = o(p^{0.3})$, is due to the microscopic scale hypothesis. Since $\nu_a \asymp N$ for all $a \in \Omega$, the above estimate along with (6.6) gives

$$\sum_{a \in \Omega} g_{\leq p^{1.1}}(a) \nu_a \geq (1-o(1))\delta p T_k(\Omega).$$

⁵*Proof of estimate:* by Lemma 5.2, for $x \in [p^{1.1}, \frac{1}{2}p^{0.9}]$ we have $I_p(p+x) \geq I_p(p+p^{1.1}) \sim \frac{1}{2}p^{1.2} \geq p^{0.3}x$, and for $x \geq \frac{1}{2}p^{0.9}$ we have $I_p(p+x) \sim x \log(x/p) \geq (1+o(1))p^{0.3}x$.

By the Cauchy–Schwarz inequality, we have

$$\sum_{a \in \Omega} g_{\leq p^{1.1}}(a)^2 \geq (1 - o(1)) \frac{\delta^2 p^2 T_k(\Omega)^2}{\sum_{a \in \Omega} \nu_a^2}.$$

It follows by Lemma 5.2 and the above estimate that

$$\sum_{a \in \Omega} I_p(f(a)) \geq \sum_{a \in \Omega} I_p(p + g_{\leq p^{1.1}}(a)) \sim \sum_{a \in \Omega} \frac{g_{\leq p^{1.1}}(a)^2}{2p} \geq (1 + o(1)) \frac{\delta^2 T_k(\Omega)^2}{2 \sum_{a \in \Omega} \nu_a^2},$$

which combined with Lemma 6.2 completes the proof of Theorem 6.1.

6.2. Microscopic rate function. In the case $\Omega = \mathbb{Z}/N\mathbb{Z}$, by symmetry, $\nu_s = kN$ for all $s \in \Omega$, and $T_k(\Omega) = N^2$. Thus Theorem 2.6 for $\Omega = \mathbb{Z}/N\mathbb{Z}$ follows from Theorem 6.1.

When $\Omega = [N] \subset \mathbb{Z}$, the derivation of the formula in Theorem 2.6 is routine though a bit more involved. For each $s \in [N]$ and $0 \leq j < k$, let $\nu_{s,j}$ denote the number of pairs of $a, b \in \mathbb{Z}$ such that $a, a + (k-1)b \in [N]$ and $a + jb = s$, i.e., the number of k -APs (allowing zero or negative common difference) contained in $[N]$ and whose $(j+1)$ -th term is s . It is easy to see that $\nu_{s,j}$ equals to the number of $b \in \mathbb{Z}$ satisfying $1 \leq s + (k-1-j)b \leq N$ and $1 \leq s - jb \leq N$, so that

$$\nu_{s,j} = \min \left\{ \left\lfloor \frac{s-1}{j} \right\rfloor, \left\lfloor \frac{N-s}{k-1-j} \right\rfloor \right\} + \min \left\{ \left\lfloor \frac{s-1}{k-1-j} \right\rfloor, \left\lfloor \frac{N-s}{j} \right\rfloor \right\} + 1.$$

Thus, for each $s \in [N]$,

$$\nu_s = \sum_{j=0}^{k-1} \nu_{s,j} = k + 2 \sum_{j=0}^{k-1} \min \left\{ \left\lfloor \frac{s-1}{j} \right\rfloor, \left\lfloor \frac{N-s}{k-1-j} \right\rfloor \right\}.$$

By Riemann sum, we have

$$\lim_{N \rightarrow \infty} N^{-3} \sum_{s=1}^N \nu_s^2 = 4 \sum_{i,j=0}^{k-1} \beta_{i,j},$$

where

$$\beta_{i,j} = \int_0^1 \min \left\{ \frac{x}{i}, \frac{1-x}{k-1-i} \right\} \min \left\{ \frac{x}{j}, \frac{1-x}{k-1-j} \right\} dx.$$

Observing that $\min\{x/i, (1-x)(k-1-i)\}$ is piecewise-linear with the kink at $x = i/(k-1)$, we can compute the above integral: for all $0 \leq i \leq j \leq k-1$,

$$\begin{aligned} \beta_{ij} &= \int_0^{\frac{i}{k-1}} \frac{x^2}{ij} dx + \int_{\frac{i}{k-1}}^{\frac{j}{k-1}} \frac{(1-x)x}{(k-1-i)j} dx + \int_{\frac{j}{k-1}}^1 \frac{(1-x)^2}{(k-1-i)(k-1-j)} dx \\ &= \frac{(k-1)^2 - i^2 - (k-1-j)^2}{6(k-1)^2(k-1-i)j}. \end{aligned}$$

In particular, for all $0 \leq i \leq k-1$,

$$\beta_{i,i} = \frac{1}{3(k-1)^2}.$$

Therefore,

$$\lim_{N \rightarrow \infty} N^{-3} \sum_{s=1}^N \nu_s^2 = 4 \sum_{i,j=0}^{k-1} \beta_{i,j} = 4 \sum_{i=0}^{k-1} \beta_{i,i} + 8 \sum_{0 \leq i < j < k} \beta_{i,j} = \frac{\gamma_k}{(k-1)^2},$$

where

$$\gamma_k = \frac{4}{3} \left(k + \sum_{0 \leq i < j < k} \frac{(k-1)^2 - i^2 - (k-1-j)^2}{(k-1-i)j} \right).$$

Since $T_k([N]) \sim N^2/(k-1)$ for fixed k as $N \rightarrow \infty$, Theorem 2.6 for $\Omega = [N]$ follows from Theorem 6.1.

As for the remark following Theorem 2.6, we always have $\gamma_k \geq k^2$ due to (6.3). The first few values of γ_k are

$$\gamma_3 = \frac{28}{3}, \quad \gamma_4 = 17, \quad \gamma_5 = \frac{718}{27}.$$

The asymptotic dependence of γ_k on k can be computed via a Riemann integral⁶ (note that the integrand takes value in $[0, 2]$ in the given domain):

$$\lim_{k \rightarrow 0} \frac{\gamma_k}{k^2} = \frac{4}{3} \int_{0 \leq x \leq y \leq 1} \frac{1 - x^2 - (1-y)^2}{(1-x)y} dx dy = \frac{40 - 2\pi^2}{9} \approx 1.14.$$

7. REPLICA SYMMETRY

In this section, we record a partial result on exact replica symmetry for constant values of p and δ in the case of $\Omega = \mathbb{Z}/N\mathbb{Z}$, analogous to results about dense random graphs in [10, 27]. Unlike previous sections, where we solve the variational problem asymptotically as $p \rightarrow 0$, the following theorem gives *exact* replica symmetry, i.e., we give sufficient conditions on constants p and δ so that the constant function uniquely minimizes the variational problem. Unlike the results in [27], we do not know if the following theorem gives the full replica symmetry phase (it probably does not). The proof is nearly identical to the one in [27], the only difference being the Hölder-like inequality (Lemma 7.2) for k -APs.

Using arguments very similar to [32], one can also prove regions of replica symmetry in the lower tail. Details are omitted.

Theorem 7.1. *Take any $0 < p \leq q < 1$, positive integer $k \geq 3$ and prime N . Suppose that $(q^{k/2}, I_p(q))$ lies on the convex minorant of the function $x \mapsto I_p(x^{2/k})$. Then the constant function $f \equiv q$ is the unique minimizer to the variational problem (2.2) with $\Omega = \mathbb{Z}/N\mathbb{Z}$ and $(1+\delta)p^k = q^k$, so that $\phi_p^{(k, \mathbb{Z}/N\mathbb{Z})}(\delta) = NI_p(q)$.*

Remark. The hypothesis that N is prime is mainly for convenience, and it is likely unnecessary here. For example, the proof shows that when k is even, there is no requirement on N , and when k is odd, $\gcd(k-1, N) = 1$ suffices.

Proof. In the variational problem (2.2), suppose that $f: \mathbb{Z}/N\mathbb{Z} \rightarrow [0, 1]$ satisfies $T_k(f) \geq (1+\delta)p^k N^2 = q^k N^2$. By Lemma 7.2 below, we have

$$\sum_{a \in \mathbb{Z}/N\mathbb{Z}} f(a)^{k/2} \geq q^{k/2} N. \quad (7.1)$$

Let $J(x) = I_p(x^{2/k})$, and $\ell(x) = J'(q^{k/2})(x - q^{k/2}) + I_p(q)$. Then ℓ is the tangent line to $J(x)$ at $x = q^{k/2}$, and by the convex minorant condition, we have $J(x) \geq \ell(x)$ for all $x \in [0, 1]$.

⁶The integral was computed using MATHEMATICA.

Since $I_p(x)$ is increasing in $x \in [p, 1]$, we have $J'(q^{k/2}) \geq 0$. It follows that

$$\begin{aligned}
\sum_{a \in \Omega} I_p(f(a)) &= \sum_{a \in \Omega} J(f(a)^{k/2}) \\
&\geq \sum_{a \in \Omega} \ell(f(a)^{k/2}) \\
&= J'(q^{k/2}) \left(\sum_{a \in \Omega} f(a)^{k/2} - q^{k/2} N \right) + I_p(q) \\
&\geq I_p(q)
\end{aligned}$$

by (7.1). Equality occurs if and only if $f(a) = q$ for all $a \in \mathbb{Z}/N\mathbb{Z}$. □

Lemma 7.2. *Let $k \geq 3$ and N a prime. For any $f: \mathbb{Z}/N\mathbb{Z} \rightarrow [0, \infty)$, one has*

$$T_k(f) \leq \left(\sum_{a \in \mathbb{Z}/N\mathbb{Z}} f(a)^{k/2} \right)^2.$$

Proof. Define

$$h_{i,j}(a, b) = \sqrt{f(a + ib)f(a + jb)}.$$

By Hölder's inequality, one has

$$\begin{aligned}
T_k(f) &= \sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} f(a)f(a+b) \cdots f(a+(k-1)b) \\
&= \sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} h_{0,1}(a, b)h_{1,2}(a, b) \cdots h_{k-2,k-1}(a, b)h_{k-1,0}(a, b) \\
&\leq \prod_{(i,j) \in \{(0,1), (1,2), \dots, (k-2,k-1), (k-1,0)\}} \left(\sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} h_{i,j}(a, b)^k \right)^{1/k} \\
&= \left(\sum_{a \in \mathbb{Z}/N\mathbb{Z}} f(a)^{k/2} \right)^2,
\end{aligned}$$

where the last step is due to

$$\begin{aligned}
\sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} h_{ij}(a, b)^k &= \sum_{a, b \in \mathbb{Z}/N\mathbb{Z}} f(a + ib)^{k/2} f(a + jb)^{k/2} \\
&= \left(\sum_{a \in \mathbb{Z}/N\mathbb{Z}} f(a)^{k/2} \right)^2 \quad \text{if } \gcd(i - j, N) = 1.
\end{aligned}$$

Note that when k is even, we can instead write

$$f(a)f(a+b) \cdots f(a+(k-1)b) = h_{0,1}(a, b)^2 h_{2,3}(a, b)^2 \cdots h_{k-2,k-1}(a, b)^2$$

and remove the need for primality hypothesis on N . □

8. MAXIMIZING THE NUMBER OF k -APS

In this section we prove Theorem 2.4, repeated below for convenience.

Theorem 2.4. *Fix a positive integer $k \geq 3$. There exists some constant $c_k > 0$ such that the following statement holds. Let $A \subset \mathbb{Z}$ be a subset with $|A| = n$, or $A \subset \mathbb{Z}/N\mathbb{Z}$ with N prime and $|A| = n \leq c_k N$. Then $T_k(A) \leq T_k([n])$.*

In Section 8.1, we will prove Theorem 2.4 when $A \subset \mathbb{Z}$ using a simple combinatorial argument. Unfortunately this proof does not extend to the case $A \subset \mathbb{Z}/N\mathbb{Z}$, due to the lack of a natural ordering in $\mathbb{Z}/N\mathbb{Z}$. Following the idea in [18], we will attempt to replace the original set $A \subset \mathbb{Z}/N\mathbb{Z}$ by a Freiman model $\tilde{A} \subset \mathbb{Z}$ (so that in particular $|A| = |\tilde{A}|$ and $T_k(A) = T_k(\tilde{A})$). This technique, called *rectification*, was first investigated in [3]. The following lemma gives a simple example of rectification:

Lemma 8.1. *Let N be a positive integer. Let $\tilde{A} \subset \mathbb{Z} \cap (-N/4, N/4)$ be a subset, and let $A \subset \mathbb{Z}/N\mathbb{Z}$ be the image of $\tilde{A}$ under the natural projection $\mathbb{Z} \rightarrow \mathbb{Z}/N\mathbb{Z}$. Then $T_k(A) = T_k(\tilde{A})$.*

Proof. Let $\pi: \tilde{A} \rightarrow A$ be the natural projection map. We need to show that both π and π^{-1} preserve k -APS. It suffices to show that π is a *Freiman isomorphism*, in the sense that for any $a_1, a_2, a_3, a_4 \in \tilde{A}$, $a_1 + a_2 = a_3 + a_4$ if and only if $\pi(a_1) + \pi(a_2) = \pi(a_3) + \pi(a_4)$. The only if direction is clear. The if direction follows from the fact that

$$-N < a_1 + a_2 - a_3 - a_4 < N$$

whenever $a_1, a_2, a_3, a_4 \in \tilde{A}$. □

A more sophisticated rectification lemma is given in [18, Theorem 4.1], which allows us to prove Theorem 2.4 when the set A has small doubling, in the sense that $|A + A|/|A|$ is small (see Lemma 8.3 below). After stating some preparatory lemmas in Section 8.2, we will then prove the general case of Theorem 2.4 in Section 8.3 using a structural decomposition theorem [18, Proposition 3.2], which allows us to deduce that if $T_k(A)$ is close to maximal then A must have small doubling.

8.1. Proof of Theorem 2.4 when $A \subset \mathbb{Z}$. In this subsection we prove the case when $A \subset \mathbb{Z}$ by induction on k .⁷ The statement is trivial when $k = 2$ since $T_k(A) = n^2$ always. Now let $k \geq 3$, and assume that the statement has been proved for smaller values of k . It is convenient to count the number of nontrivial increasing k -APS in A :

$$T'_k(A) = \#\{(a, b) : b > 0, a, a + b, \dots, a + (k - 1)b \in A\}.$$

Clearly $T_k(A) = 2T'_k(A) + n$. Thus it suffices to show that $T'_k(A) \leq T'_k([n])$. Arrange the elements in A in increasing order:

$$A = \{a_0 < a_1 < \dots < a_{n-1}\}.$$

Choose $m \in \mathbb{Z}$ such that

$$\frac{k-2}{k-1}(n-1) \leq m < \frac{k-2}{k-1}n + 1.$$

There are two types of k -APS counted in $T'_k(A)$: those whose second largest element is at least a_m , and those whose second largest element is smaller than a_m . If the second largest

⁷We are grateful to Anton Bankevich for suggesting the proof in this subsection.

element of a k -AP in A is a_i for some $m \leq i < n$, then there are at most $n - 1 - i$ possibilities for its largest element. Thus the number of k -APs in A of the first type is at most

$$\sum_{m \leq i < n} (n - 1 - i).$$

For k -APs in A of the second type, their first $k - 1$ terms form a $(k - 1)$ -AP in $\{a_0, \dots, a_{m-1}\}$. Thus the number of k -APs in A of the second type is at most

$$T'_{k-1}(\{a_0, \dots, a_{m-1}\}) \leq T'_{k-1}(\{0, 1, \dots, m - 1\})$$

by induction hypothesis. It follows that

$$T'_k(A) \leq \sum_{m \leq i < n} (n - 1 - i) + T'_{k-1}(\{0, 1, \dots, m - 1\}).$$

To conclude the proof, we claim that the first term on the right hand side above is equal to the number of k -APs in $\{0, 1, \dots, n - 1\}$ of the first type:

$$\sum_{m \leq i < n} (n - 1 - i) = \#\{(a, b) : b > 0, a \geq 0, a + (k - 1)b < n, a + (k - 2)b \geq m\}, \quad (8.1)$$

and the second term is equal to the number of k -APs in $\{0, 1, \dots, n - 1\}$ of the second type:

$$T'_{k-1}(\{0, 1, \dots, m - 1\}) = \#\{(a, b) : b > 0, a \geq 0, a + (k - 1)b < n, a + (k - 2)b < m\}. \quad (8.2)$$

To prove (8.1), it suffices to show that for any $m \leq i < n$ we have

$$n - 1 - i = \#\{(a, b) : b > 0, a \geq 0, a + (k - 1)b < n, a + (k - 2)b = i\}.$$

This follows from the fact that any choice of the value of $j = a + (k - 1)b$ from $\{i + 1, \dots, n - 1\}$ uniquely determines an admissible (a, b) since

$$i \geq (k - 2)(j - i)$$

by our choice of m . To prove (8.2), note that any (a, b) with $b > 0$, $a \geq 0$, and $a + (k - 2)b < m$ automatically satisfies $a + (k - 1)b < n$ since

$$a + (k - 1)b = \frac{k - 1}{k - 2} [a + (k - 2)b] - \frac{a}{k - 2} \leq \frac{k - 1}{k - 2} (m - 1) < n$$

by our choice of m .

8.2. Proof of Theorem 2.4 when $A \subset \mathbb{Z}/N\mathbb{Z}$: preparations. In this subsection, we collect a few lemmas that will allow us to reduce the problem in $\mathbb{Z}/N\mathbb{Z}$ to the simpler one in $\mathbb{Z}$. From now on we fix some large prime N and work in $\mathbb{Z}/N\mathbb{Z}$. To begin with, we show that k -AP counts are controlled by additive energy. For two subsets $A, B \subset \mathbb{Z}/N\mathbb{Z}$, the additive energy $E(A, B)$ is defined by

$$E(A, B) = \#\{(a, a', b, b') \in A \times A \times B \times B : a + b = a' + b'\},$$

and note the trivial bound

$$E(A, B) \leq \min(|A|^2|B|, |A||B|^2) \leq |A|^{3/2}|B|^{3/2}.$$

We will also consider additive energy of dilates $\ell \cdot A$ for a positive integer ℓ , defined by

$$\ell \cdot A = \{\ell a : a \in A\}.$$

For subsets $A_1, \dots, A_k \subset \mathbb{Z}/N\mathbb{Z}$, define the asymmetric k -AP count by

$$T_k(A_1, \dots, A_k) = \#\{(a, b) : a \in A_1, a + b \in A_2, \dots, a + (k - 1)b \in A_k\}.$$

Clearly if $A_1 = \dots = A_k = A$ then $T_k(A, \dots, A) = T_k(A)$.

Lemma 8.2. *Let $A_1, \dots, A_k \subset \mathbb{Z}/N\mathbb{Z}$ be subsets, and let $n = \max(|A_1|, \dots, |A_k|)$. If*

$$\max_{\ell, \ell' \in \{1, 2\}} E(\ell \cdot A_i, \ell' \cdot A_{i+1}) \leq \varepsilon n^3$$

for some $\varepsilon \in (0, 1)$ and $1 \leq i < k$, then $T_k(A_1, \dots, A_k) \leq \varepsilon^{1/6} n^2$.

Proof. Note that $T_k(A_1, \dots, A_k)$ is trivially bounded by either $T_3(A_{i-1}, A_i, A_{i+1})$ or $T_3(A_i, A_{i+1}, A_{i+2})$. The conclusion then follows immediately from [18, Lemma 4.2]. $\square$

The following lemma shows that Theorem 2.4 holds if A has small doubling.

Lemma 8.3. *Let $K \geq 1$. The following statement holds if c is sufficiently small in terms of K . Let $A \subset \mathbb{Z}/N\mathbb{Z}$ be a subset with $|A| = n \leq cN$ and $|A + A| \leq K|A|$. Then $T_k(A) \leq T_k([n])$.*

Proof. By [18, Theorem 4.1], there exists a dilate $d \cdot A$ of A (for some $d \in (\mathbb{Z}/N\mathbb{Z})^*$) that is contained in an interval of length at most $N/2k$. After dilation and translation we may assume that $A \subset [1, N/2k]$. Now that any k -AP in A (as a subset of $\mathbb{Z}/N\mathbb{Z}$) is also a k -AP in $\mathbb{Z}$, we may apply the $A \subset \mathbb{Z}$ case of Theorem 2.4 to conclude that $T_k(A) \leq T_k([n])$, as desired. $\square$

8.3. Proof of Theorem 2.4 when $A \subset \mathbb{Z}/N\mathbb{Z}$. We continue to work in $\mathbb{Z}/N\mathbb{Z}$. In view of Lemma 8.3, we may assume that n is sufficiently large in terms of k . We first establish a rough structure theorem for sets with close to maximal number of k -APs.

Lemma 8.4. *Let $c > 0$ be sufficiently small. Let $A \subset \mathbb{Z}/N\mathbb{Z}$ be a subset with $|A| = n \leq cN$, and suppose that n is sufficiently large in terms of k . If*

$$T_k(A) \geq \frac{n^2}{k-1} \left(1 - \frac{1}{100k^2}\right),$$

then there exists a dilate $d \cdot A$ of A for some $d \in (\mathbb{Z}/N\mathbb{Z})^$, such that all but at most $n/10k^2$ elements in $d \cdot A$ lies in an interval of length at most $N/100k$.*

Proof. Choose a sufficiently small $\varepsilon > 0$, and then choose some $\varepsilon' > 0$ sufficiently small in terms of ε . Apply [18, Proposition 3.2] to obtain a structural decomposition $A = A_1 \cup \dots \cup A_m \cup A_0$ into disjoint subsets satisfying the following properties:

- (1) $|A_i| \gtrsim_\varepsilon |A|$ for each $1 \leq i \leq m$, so that $m \lesssim_\varepsilon 1$;
- (2) $|A_i + A_i| \lesssim_{\varepsilon, \varepsilon'} |A_i|$ for each $1 \leq i \leq m$;
- (3) $E(\ell \cdot A_i, \ell' \cdot A_j) \leq \varepsilon' |A_i|^{3/2} |A_j|^{3/2}$ whenever $1 \leq i < j \leq m$ and $\ell, \ell' \in \{1, 2\}$;
- (4) $E(\ell \cdot A_0, \ell' \cdot A) \leq \varepsilon |A|^3$ whenever $\ell, \ell' \in \{1, 2\}$.

To estimate $T_k(A)$, we write it as the sum of m^k terms of the form $T_k(A_{i_1}, \dots, A_{i_k})$ for some $1 \leq i_1, \dots, i_k \leq m$, and the k terms $T_k(A_0, A, \dots, A), T_k(A, A_0, \dots, A), \dots, T_k(A, A, \dots, A_0)$. To estimate $T_k(A_{i_1}, \dots, A_{i_k})$ when $i_1, \dots, i_k$ are not all the same, we use Lemma 8.2 and property (3) to obtain

$$T_k(A_{i_1}, \dots, A_{i_k}) \leq \varepsilon^{1/6} n^2.$$

Thus the total contributions from these terms are bounded by $\varepsilon^{1/6} m^k n^2$. Moreover, by Lemma 8.2 and property (4) we also have

$$T_k(A_0, A, \dots, A) \leq \varepsilon^{1/6} n^2,$$

and similarly for the other $k - 1$ terms involving A_0 . Thus we have shown that

$$T_k(A) \leq \sum_{i=1}^m T_k(A_i) + (\varepsilon'^{1/6} m^k + \varepsilon^{1/6} k) n^2 \leq \sum_{i=1}^k T_k(A_i) + \frac{n^2}{100k^3},$$

if $\varepsilon, \varepsilon'$ are small enough (recall that m depends only on ε). If c is sufficiently small in terms of $\varepsilon, \varepsilon'$, then Lemma 8.3 can be applied to each A_i in view of property (2) to get

$$T_k(A) \leq \frac{1}{k-1} \sum_{i=1}^m |A_i|^2 + \frac{1}{4}(k-1)m + \frac{n^2}{100k^3} \leq \frac{n}{k-1} \max_{1 \leq i \leq m} |A_i| + \frac{1}{4}(k-1)m + \frac{n^2}{100k^3}.$$

Combined with the lower bound for $T_k(A)$ in the hypothesis, this implies that $|A_i| \geq (1 - 1/10k^2)n$ for some $1 \leq i \leq k$. By [18, Theorem 4.1], after dilation we may assume that A_i is contained in an interval of length at most $N/100k$. $\square$

We are now ready to prove Theorem 2.4. In view of Lemma 8.4, after a suitable dilation and translation we may assume that the set $A_0 = A \cap [-N/100k, N/100k]$ has size $|A_0| \geq (1 - 1/10k^2)|A|$. Let

$$A_1 = A \cap ([-0.1N, 0.1N] \cup [0.4N, 0.6N]).$$

Write $n_1 = |A_1|$ and $n_2 = |A \setminus A_1| = n - n_1$. Note that for any k -AP in A , if at least two of its first three terms lie in A_0 , then it is entirely contained in A_1 . Thus any k -AP in A that is not entirely contained in A_1 must have at least one term outside A_1 and at least two out of the first three terms outside A_0 , and the number of these k -APs is then bounded by $3kn_2|A \setminus A_0| \leq nn_2/2k$. It follows that

$$T_k(A) \leq T_k(A_1) + \frac{nn_2}{2k}.$$

Note that $2 \cdot A_1 \subset [-0.2N, 0.2N]$. Thus from Lemma 8.1 and the integer case of Theorem 2.4 proved in Section 8.1, we obtain

$$T_k(A_1) = T_k(2 \cdot A_1) \leq T_k([n_1]). \quad (8.3)$$

Using (2.6) we arrive at

$$T_k(A) \leq \frac{n_1^2}{k-1} + \frac{1}{4}(k-1) + \frac{nn_2}{2k}.$$

If $n_2 > 0$, then

$$\frac{1}{4}(k-1) + \frac{nn_2}{2k} \leq \frac{nn_2}{k} \leq \frac{n^2 - n_1^2}{k-1},$$

provided that $n \geq k(k-1)/2$ (which we may assume). Thus in this case we have

$$T_k(A) \leq \frac{n^2}{k-1} \leq T_k([n])$$

as desired. If $n_2 = 0$, then $A = A_1$ and the desired conclusion follows from (8.3).

APPENDIX A. PROOFS OF FURTHER BOUNDS ON GAUSSIAN WIDTHS

A.1. Matching lower bound for 3-APs. Here we prove the matching lower bound

$$\text{GW}(T_3/N) \gtrsim \sqrt{N \log N}$$

in Theorem 4.4. As in Section 4.1 we only work out the $\Omega = \mathbb{Z}/N\mathbb{Z}$ setting here. The $\Omega = [N] \subset \mathbb{Z}$ setting can be deduced similarly by embedding $[N]$ in a larger cyclic group to avoid APs from wrapping around zero.

First we show that if $h: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{R}$ is a random function with $h(a) \sim \text{Normal}(0, 1)$ i.i.d. for $a \in \mathbb{Z}/N\mathbb{Z}$, then, with probability at least $1/2$, we can find some $f: \mathbb{Z}/N\mathbb{Z} \rightarrow [0, 1]$ such that

$$\langle \frac{1}{N} \nabla T_3(f), h \rangle \gtrsim \sqrt{N \log N}. \quad (\text{A.1})$$

Indeed, the real components of $\widehat{h}(2s)$ for integers $0 \leq s < N/4$ are independent Gaussians with variance $\Theta(1/N)$ (since the Fourier transform is orthogonal), so with probability at least 0.9 there is some integer $0 < s < N/4$ such that $\Re \widehat{h}(2s) \gtrsim \sqrt{(\log N)/N}$ and furthermore $\widehat{h}(0) = O(1/\sqrt{N})$. Then, setting $f(a) = (1 + \cos(2\pi sa/N))/2$ so that $\widehat{f}(0) = 1/2$, $\widehat{f}(\pm s) = 1/4$ and $\widehat{f}(r) = 0$ for all $r \notin \{0, s, -s\}$, we obtain, by Lemma 4.3,

$$\frac{1}{N^2} T_3(f, h, f) = \widehat{f}(0)^2 \widehat{h}(0) + \widehat{f}(s)^2 \widehat{h}(-2s) + \widehat{f}(-s)^2 \widehat{h}(2s) = \frac{1}{4} \widehat{h}(0) + \frac{1}{8} \Re \widehat{h}(2s) \gtrsim \sqrt{\frac{\log N}{N}},$$

and

$$\frac{1}{N^2} T_3(h, f, f) = \frac{1}{N^2} T_3(f, f, h) = \widehat{f}(0)^2 \widehat{h}(0) = O(N^{-1/2})$$

Thus (A.1) holds since $\langle \nabla T_3(f), h \rangle = T_3(f, f, h) + T_3(f, h, f) + T_3(h, f, f)$.

Finally, one can convert $f: \mathbb{Z}/N\mathbb{Z} \rightarrow [0, 1]$ to a $\{0, 1\}$ -valued function by changing the value of f at $a \in \mathbb{Z}/N\mathbb{Z}$ to 1 with probability $f(a)$ and 0 with probability $1 - f(a)$. A routine probabilistic argument shows that (A.1) holds for some $f: \mathbb{Z}/N\mathbb{Z} \rightarrow \{0, 1\}$, thereby proving $\text{GW}(T_3/N) \gtrsim \sqrt{N \log N}$.

A.2. Improved upper bound for 4-AP. Here we prove the final claim in Theorem 4.1 that

$$\text{GW}(T_4/N) \lesssim N^{3/4} (\log N)^{1/4}. \quad (\text{A.2})$$

As earlier we only discuss the $\Omega = \mathbb{Z}/N\mathbb{Z}$ setting here. We have

$$\langle \nabla T_4(f), h \rangle = T_4(h, f, f, f) + T_4(f, h, f, f) + T_4(f, f, h, f) + T_4(f, f, f, h). \quad (\text{A.3})$$

Define the following multiplicative analogue of the finite difference

$$\Delta_s h(x) := f(x) \overline{h(x+s)}$$

(since we will be working with real-valued functions, one can ignore the conjugation). The Fourier transform of $\Delta_s h$ controls T_4 , as shown by the following lemma, which plays a similar role to Lemma 4.4 for T_3 .

Lemma A.1. *For any $f: \mathbb{Z}/N\mathbb{Z} \rightarrow [-1, 1]$ and any $h: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{R}$,*

$$\frac{1}{N^2} |T_4(h, f, f, f)| \lesssim \left(\frac{1}{N} \sum_{s \in \mathbb{Z}/N\mathbb{Z}} \|\widehat{\Delta_s h}\|_\infty \right)^{1/2}.$$

The same holds with the left-hand side replaced by $T(f, h, f, f)$, $T(f, f, h, f)$, or $T(f, f, f, h)$.

Proof. We prove the inequality for $T(h, f, f, f)$ as the other cases are analogous. We have

$$\begin{aligned}
|T_4(h, f, f, f)|^2 &= \left| \sum_a \left(\sum_b h(a-3b)f(a-2b)f(a-b) \right) f(a) \right|^2 \\
&\leq N \sum_a \left(\sum_b h(a-3b)f(a-2b)f(a-b) \right)^2 \quad [\text{Cauchy-Schwarz}] \\
&= N \sum_{a,b,b'} h(a-3b)h(a-3b')f(a-2b)f(a-2b')f(a-b)f(a-b') \\
&= N \sum_{a,b,s} \Delta_{3s}h(a-3b)\Delta_{2s}f(a-2b)\Delta_s f(a-b) \quad [\text{by setting } s = b - b'] \\
&= N \sum_s T_3(\Delta_{3s}h, \Delta_{2s}f, \Delta_s f) \\
&\lesssim N^3 \sum_s \|\widehat{\Delta_{3s}h}\|_\infty \quad [\text{by Lemma 4.4 on } T_3] \\
&\lesssim N^3 \sum_s \|\widehat{\Delta_s h}\|_\infty.
\end{aligned}$$

The lemma follows by rearranging. $\square$

Let $h: \mathbb{Z}/N\mathbb{Z} \rightarrow \mathbb{R}$ be a random function with independent standard Gaussian values. Noting that

$$\widehat{\Delta_s h}(r) = \frac{1}{N} \sum_{a \in \mathbb{Z}/N\mathbb{Z}} h(a)h(a+s)\omega^{-ar}$$

is a quadratic form of Gaussians, and using tail bounds for such random variables [21] (also [4, Example 2.12]) we find that for every $s \neq 0$,

$$\|\widehat{\Delta_s h}\|_\infty = O\left(\sqrt{\frac{\log N}{N}}\right) \quad \text{with probability } 1 - O(N^{-10}). \quad (\text{A.4})$$

We always have $\|\widehat{\Delta_s h}\|_\infty \leq 1$ for all s . Thus

$$\begin{aligned}
\text{GW}(T_4/N) &= \mathbb{E}_h \sup_{f: \mathbb{Z}/N\mathbb{Z} \rightarrow \{0,1\}} \langle \nabla T_4(f)/N, h \rangle \\
&\lesssim \mathbb{E}_h \sqrt{N} \left(\sum_{s \in \mathbb{Z}/N\mathbb{Z}} \|\widehat{\Delta_s h}\|_\infty \right)^{1/2} \quad [\text{by (A.3) and Lem. A.1}] \\
&\lesssim N^{3/4} (\log N)^{1/4}. \quad [\text{by (A.4)}]
\end{aligned}$$

Acknowledgments. We thank Ben Green and Freddie Manners for helpful discussions. We also thank the anonymous referee for helpful comments that improved the exposition of the paper.

REFERENCES

- [1] R. C. Baker, G. Harman, and J. Pintz. The difference between consecutive primes. II. *Proc. London Math. Soc. (3)*, 83(3):532–562, 2001. $\uparrow 15$
- [2] B. B. Bhattacharya, S. Ganguly, E. Lubetzky, and Y. Zhao. Upper tails and independence polynomials in random graphs. *Adv. Math.*, 319:313–347, 2017. $\uparrow 1, \uparrow 2, \uparrow 3$

- [3] Y. F. Bilu, V. F. Lev, and I. Z. Ruzsa. Rectification principles in additive number theory. *Discrete Comput. Geom.*, 19(3, Special Issue):343–353, 1998. Dedicated to the memory of Paul Erdős. ↑27
- [4] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013. A nonasymptotic theory of independence, With a foreword by Michel Ledoux. ↑32
- [5] J. Briët and S. Gopi. Gaussian width bounds with applications to arithmetic progressions in random settings. arXiv:1711.05624. ↑4, ↑11
- [6] S. Chatterjee. The missing log in large deviations for triangle counts. *Random Structures Algorithms*, 40(4):437–451, 2012. ↑1
- [7] S. Chatterjee. An introduction to large deviations for random graphs. *Bull. Amer. Math. Soc.*, to appear. ↑2
- [8] S. Chatterjee and A. Dembo. Nonlinear large deviations. *Adv. Math.*, 299:396–450, 2016. ↑1, ↑2, ↑3, ↑7, ↑14
- [9] S. Chatterjee and P. S. Dey. Applications of Stein’s method for concentration inequalities. *Ann. Probab.*, 38(6):2443–2485, 2010. ↑1
- [10] S. Chatterjee and S. R. S. Varadhan. The large deviation principle for the Erdős-Rényi random graph. *European J. Combin.*, 32(7):1000–1017, 2011. ↑1, ↑6, ↑25
- [11] D. Conlon and W. T. Gowers. Combinatorial theorems in sparse random sets. *Ann. of Math. (2)*, 184(2):367–454, 2016. ↑14
- [12] E. Croot. The minimal number of three-term arithmetic progressions modulo a prime converges to a limit. *Canad. Math. Bull.*, 51:47–56, 2008. ↑6
- [13] B. Demarco and J. Kahn. Tight upper tail bounds for cliques. *Random Structures Algorithms*, 41(4):469–487, 2012. ↑1
- [14] B. DeMarco and J. Kahn. Upper tails for triangles. *Random Structures Algorithms*, 40(4):452–459, 2012. ↑1
- [15] R. Eldan. Gaussian-width gradient complexity, reverse log-Sobolev inequalities and nonlinear large deviations. arXiv:1612.04346. ↑2, ↑3, ↑4, ↑7, ↑8
- [16] N. Frantzikinakis. Some open problems on multiple ergodic averages. *arXiv preprint arXiv:1103.3808*, 2016. ↑14
- [17] N. Frantzikinakis, E. Lesigne, and M. Wierdl. Random differences in Szemerédi’s theorem and related results. *J. Anal. Math.*, 130:91–133, 2016. ↑14
- [18] B. Green and O. Sisask. On the maximal number of 3-term arithmetic progressions in subsets of $\mathbb{Z}/p\mathbb{Z}$. *Bull. Lond. Math. Soc.*, 40(6):945–955, 2008. ↑2, ↑5, ↑27, ↑29, ↑30
- [19] B. Green and T. Tao. The quantitative behaviour of polynomial orbits on nilmanifolds. *Ann. of Math. (2)*, 175(2):465–540, 2012. ↑14
- [20] B. Green, T. Tao, and T. Ziegler. An inverse theorem for the Gowers $U^{s+1}[N]$ -norm. *Ann. of Math. (2)*, 176(2):1231–1372, 2012. ↑2
- [21] D. L. Hanson and F. T. Wright. A bound on tail probabilities for quadratic forms in independent random variables. *Ann. Math. Statist.*, 42:1079–1083, 1971. ↑32
- [22] S. Janson, K. Oleszkiewicz, and A. Ruciński. Upper tails for subgraph counts in random graphs. *Israel J. Math.*, 142:61–92, 2004. ↑1
- [23] S. Janson and A. Ruciński. The infamous upper tail. *Random Structures Algorithms*, 20(3):317–342, 2002. ↑1
- [24] S. Janson and A. Ruciński. The deletion method for upper tail estimates. *Combinatorica*, 24(4):615–640, 2004. ↑1
- [25] S. Janson and A. Ruciński. Upper tails for counting objects in randomly induced subhypergraphs and rooted random graphs. *Ark. Mat.*, 49(1):79–96, 2011. ↑1, ↑2
- [26] J. H. Kim and V. H. Vu. Divide and conquer martingales and the number of triangles in a random graph. *Random Structures Algorithms*, 24(2):166–174, 2004. ↑1
- [27] E. Lubetzky and Y. Zhao. On replica symmetry of large deviations in random graphs. *Random Structures Algorithms*, 47(1):109–146, 2015. ↑1, ↑6, ↑25
- [28] E. Lubetzky and Y. Zhao. On the variational problem for upper tails in sparse random graphs. *Random Structures Algorithms*, 50:420–436, 2017. ↑1, ↑2, ↑3, ↑17
- [29] T. Tao. *Higher order Fourier analysis*, volume 142 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2012. ↑14

- [30] V. H. Vu. A large deviation result on the number of small subgraphs of a random graph. *Combin. Probab. Comput.*, 10(1):79–94, 2001. $\uparrow 1$
- [31] L. Warnke. Upper tails for arithmetic progressions in random subsets. *Israel J. Math.*, 221:317–365, 2017. $\uparrow 1, \uparrow 2$
- [32] Y. Zhao. On the lower tail variational problem for random graphs. *Combin. Probab. Comput.*, 26:301–320, 2017. $\uparrow 6, \uparrow 25$

DEPARTMENT OF STATISTICS, UNIVERSITY OF PENNSYLVANIA, PHILADELPHIA, PA 19104, USA
E-mail address: `bhaswar@wharton.upenn.edu`

DEPARTMENT OF STATISTICS, UC BERKELEY, CALIFORNIA, CA 94720, USA
E-mail address: `sganguly@berkeley.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF KENTUCKY, LEXINGTON, KY 40506, USA
E-mail address: `xuancheng.shao@uky.edu`

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA 02139, USA
E-mail address: `yufeiz@mit.edu`