

The Average-Case Complexity of Counting Cliques in Erdős-Rényi Hypergraphs

Enric Boix-Adserà* Matthew Brennan† Guy Bresler‡

October 15, 2020

Abstract

We consider the problem of counting k -cliques in s -uniform Erdős-Rényi hypergraphs $G(n, c, s)$ with edge density c and show that its fine-grained average-case complexity can be based on its worst-case complexity. We give a worst-case to average-case reduction for counting k -cliques on worst-case hypergraphs given a blackbox solving the problem on $G(n, c, s)$ with low error probability. Our approach is closely related to [Goldreich and Rothblum, FOCS18], which showed a worst-case to average-case reduction for counting cliques for an efficiently-sampleable distribution on graphs.

Our reduction has the following implications:

- *Dense Erdős-Rényi graphs and hypergraphs:* Counting k -cliques on $G(n, c, s)$ with k and c constant matches its worst-case time complexity up to a $\text{polylog}(n)$ factor. Assuming randomized ETH, it takes $n^{\Omega(k)}$ time to count k -cliques in $G(n, c, s)$ if k and c are constant.
- *Sparse Erdős-Rényi graphs and hypergraphs:* When $c = \Theta(n^{-\alpha})$, we give several algorithms exploiting the sparsity of $G(n, c, s)$ that are faster than the best known worst-case algorithms. Complementing this, based on a fine-grained worst-case assumption, our reduction implies a different average-case phase diagram for each fixed α depicting a tradeoff between a runtime lower bound and k . Surprisingly, in the hypergraph case ($s \geq 3$), these lower bounds are tight against our algorithms exactly when c is above the Erdős-Rényi k -clique percolation threshold.

Our reduction is the first worst-case to average-case reduction for a problem over Erdős-Rényi hypergraphs and is the first mapping from a worst-case problem to an average-case problem with a different fine-grained complexity that we are aware of. We also give a variant of our worst-case to average-case reduction for computing the parity of the k -clique count that requires a milder assumption on the error probability of the blackbox solving the problem on $G(n, c, s)$.

*Massachusetts Institute of Technology. Department of EECS. Email: eboix@mit.edu.

†Massachusetts Institute of Technology. Department of EECS. Email: brennanm@mit.edu.

‡Massachusetts Institute of Technology. Department of EECS. Email: guy@mit.edu.

1 Introduction

We consider the average-case complexity of counting k -cliques in s -uniform Erdős-Rényi hypergraphs $G(n, c, s)$, where every s -subset of the n vertices is a hyperedge independently with probability c . Our main result is a reduction for counting k -cliques on worst-case hypergraphs given a blackbox algorithm solving the problem on $G(n, c, s)$ with low error probability. Our approach is closely related to the recent work [GR18], which showed a worst-case to average-case reduction for counting cliques for a particular efficiently-samplable distribution on graphs. Our reduction yields two different sets of average-case lower bounds for counting k -cliques in graphs sampled from the natural distribution $G(n, c, s)$ in the dense and sparse cases of $c = \Theta(1)$ and $c = \Theta(n^{-\alpha})$, with tradeoffs between runtime and c . We also show that these average-case lower bounds often match algorithmic upper bounds.

The complexity of clique problems on Erdős-Rényi random graphs has become a central topic in average-case complexity, discrete probability and high-dimensional statistics. A body of work has analyzed algorithms for finding large cliques in Erdős-Rényi graphs¹ [Kuc95; AKS98; FK00; McS01; FR10; AV11; DGP14; DM15; CX16], and hardness results have been shown for greedy algorithms [Kar76; GM75; Jer92; McD84; Pit82], local algorithms [GS14; CE15; RV17], query models [Fei+18], bounded-depth circuits [Ros08], monotone circuits [Ros10], low-degree sum of squares (SOS) relaxations [Bar+16], statistical query algorithms [Fel+13], and resolution [Ats+18]. The hardness of clique problems on Erdős-Rényi graphs has been used as an average-case assumption in cryptography [JP00] and to show information-computation gaps in a variety of statistical problems [BR13; KZ14; Che15; HWX15; MW15; BBH18; BBH19; BB19].

All of the above lower bounds for clique problems on Erdős-Rényi random graphs are against restricted classes of algorithms. One reason for this is that there are general obstacles to basing average-case complexity on worst-case complexity. For example, natural approaches to polynomial-time worst-case to average-case reductions for NP-complete problems fail unless $\text{coNP} \subseteq \text{NP/poly}$ [FF93; BT06b; BT06a]. The objective of this work is to show that this worst-case characterization of average-case complexity is possible in a fine-grained sense for the natural problem of counting k -cliques in s -uniform Erdős-Rényi hypergraphs $G(n, c, s)$ with edge density c .

A motivating recent work by Goldreich and Rothblum [GR18] also considered worst-case to average-case reductions for k -clique counting. They provided such a reduction mapping to an efficiently samplable distribution on graphs with a high min-entropy of $\tilde{\Omega}(n^2)$. In contrast to [GR18], our objectives are to: (1) map precisely to the natural distribution $G(n, c, s)$ for different edge densities c , including $c = \Theta(1)$ and the sparse case $c = \Theta(n^{-\alpha})$; and (2) to characterize the tradeoff between the time-complexity of counting k -cliques in $G(n, c, s)$ and the sparsity parameter α . Achieving this requires new ingredients for the self-reducibility of counting k -cliques as a low-degree polynomial and a tight analysis of random biased binary expansions over \mathbb{F}_p with finite Fourier analysis.

However, our techniques also come at the cost of requiring a low error probability ($1/\text{polylog}(n)$ in the dense case and $1/\text{poly}(n)$ in the sparse case) for the average-case blackbox solving k -clique counting on $G(n, c, s)$. This is in contrast to [GR18], where a very high error probability of $1 - 1/\text{polylog}(n)$ is tolerated. It remains an interesting open problem to extend our results for $G(n, c, s)$ to tolerate higher error blackboxes. This error tolerance and open problem are discussed further in Sections 2.2 and 6, and how our techniques relate to those in [GR18] is discussed in Sections 1.2 and 3. As a step towards increasing the allowed blackbox error, we also give a variant of our reduction for computing the *parity* of the k -clique count that only requires a constant bound on the error probability (for each fixed k) of the blackbox algorithm solving the problem on $G(n, c, s)$

¹In both ordinary Erdős-Rényi graphs and the planted clique model.

when $c = 1/2$. We now give an overview of our contributions.

1.1 Overview of Main Results

We provide two complementary main results on the fine-grained average-case complexity of counting k -cliques in $G(n, c, s)$. The precise formulations of the problems we consider are in Section 2.1.

Worst-case to average-case reduction. We give a worst-case to average-case reduction from counting k -cliques in worst-case s -uniform hypergraphs to counting k -cliques in hypergraphs drawn from $G(n, c, s)$. The key guarantees of this reduction are summarized in the following simplified version of our main theorem.

Theorem 1.1 (Simplified Main Result). *If $2 \leq s \leq k$ are constant integers and $c = c(n)$ satisfies $0 < c \leq 1 - \Omega(1)$, then there is a parameter $\Upsilon_{\#} = c^{-\binom{k}{s}}(\log n)^{O(1)}$ such that the following holds. If there is a randomized algorithm counting k -cliques in time $O(n^t)$ with error probability less than $1/\Upsilon_{\#}$ on hypergraphs drawn from $G(n, c, s)$, then there is a randomized algorithm counting k -cliques on worst-case s -uniform hypergraphs with error probability less than $1/3$ running in time $O(\Upsilon_{\#} \cdot n^{\max\{t, s\}})$.*

We discuss the necessity of the error tolerance and the multiplicative slowdown in our worst-case to average-case reduction in Section 2.2. This result has a number of consequences for basing the average-case fine-grained complexity of k -clique counting over Erdős-Rényi hypergraphs on its worst-case complexity, which we now overview.

Counting k -cliques in worst-case hypergraphs is known to take $n^{\Omega(k)}$ time for randomized algorithms assuming the randomized Exponential Time Hypothesis (rETH)² if k does not grow with n [Che+06; Cal+08]. The best known worst-case algorithms up to subpolynomial factors are the $O(n^{\omega \lceil k/3 \rceil})$ time algorithm of [NP85] in the graph case of $s = 2$ and exhaustive $O(n^k)$ time search on worst-case hypergraphs with $s \geq 3$. Here, $\omega \leq 2.373$ denotes the matrix multiplication constant. Our reduction is the first worst-case to average-case reduction to Erdős-Rényi hypergraphs. It has different implications for the cases of dense and sparse hypergraphs because of the factor $\Upsilon_{\#}$, as described next.

1. *Dense Erdős-Rényi graphs and hypergraphs.* When k and c are constant, our reduction constructs an efficient k -clique counting algorithm that succeeds on a worst-case input hypergraph with high probability, using $\text{polylog}(n)$ queries to an average-case oracle that correctly counts k -cliques on a $1 - 1/\text{polylog}(n)$ fraction of Erdős-Rényi hypergraphs drawn from $G(n, c, s)$. This essentially shows that k -clique counting in the worst-case matches that on dense Erdős-Rényi hypergraphs. More precisely, k -clique counting on $G(n, c, s)$ with k, c and s constant must take $\tilde{\Omega}(n^{\omega \lceil k/3 \rceil})$ time when $s = 2$ and $\tilde{\Omega}(n^k)$ time when $s \geq 3$, unless there are faster worst-case algorithms. Furthermore, our reduction shows that it is rETH-hard to count k -cliques in $n^{o(k)}$ time on $G(n, c, s)$ with k, c and s constant.
2. *Sparse Erdős-Rényi graphs and hypergraphs.* Our reduction also applies with a different multiplicative slowdown and error tolerance to the sparse case of $c = \Theta(n^{-\alpha})$, where the fine-grained complexity of k -clique counting on $G(n, c, s)$ is very different than on worst-case inputs. Our reduction implies fine-grained lower bounds of $\tilde{\Omega}(n^{\omega \lceil k/3 \rceil - \alpha \binom{k}{2}})$ when $s = 2$ and $\tilde{\Omega}(n^{k - \alpha \binom{k}{s}})$ when $s \geq 3$ for inputs drawn from $G(n, c, s)$, unless there are faster worst-case algorithms.

²rETH asserts that any randomized algorithm takes at least 2^{cn} time to solve 3-SAT in the worst-case, for some constant $c > 0$.

We remark that in the hypergraph case of $s \geq 3$, this lower bound matches the expectation of the quantity being counted, the number of k -cliques in $G(n, c, s)$, up to $\text{polylog}(n)$ factors.³

Precise statements of our results can be found in Section 2.2. For simplicity, our results should be interpreted as applying to algorithms that succeed with probability $1 - (\log n)^{-\omega(1)}$ in the dense case and $1 - n^{-\omega(1)}$ in the sparse case.

We also give a second worst-case to average-case reduction for computing the parity of the number of k -cliques which has a weaker requirement of $1 - \Theta_{k,s}(1)$ on the error probability for the blackbox solving the problem on $G(n, c, s)$ in the dense case of $c = 1/2$. We provide an overview of our multi-step worst-case to average-case reduction in Section 1.2. The steps are described in detail in Section 3.

Algorithms for k -clique counting on $G(n, c, s)$. We also analyze several natural algorithms for counting k -cliques in sparse Erdős-Rényi hypergraphs. These include an extension of the natural greedy algorithm mentioned previously from k -CLIQUE to counting k -cliques, a modification to this algorithm using the matrix multiplication step of [NP85] and an iterative algorithm achieving nearly identical guarantees. These algorithms count k -cliques in $G(n, c, s)$ when $c = \Theta(n^{-\alpha})$ with several different runtimes, the best of which are as follows:

- $\tilde{O}\left(n^{k+1-\alpha\binom{k}{s}}\right)$ if $s \geq 3$ and $k < \tau + 1$;
- $\tilde{O}\left(n^{\tau+2-\alpha\binom{\tau+1}{s}}\right)$ if $s \geq 3$ and $\tau + 1 \leq k \leq \kappa + 1$; and
- $\tilde{O}\left(n^{\omega\lceil k/3\rceil + \omega - \omega\alpha\binom{\lceil k/3\rceil}{2}}\right)$ if $s = 2$ and $k \leq \kappa + 1$.

Here, τ and κ are the largest positive integers satisfying that $\alpha\binom{\tau}{s-1} < 1$ and $\alpha\binom{\kappa}{s-1} < s$. The thresholds κ and τ have natural interpretations as roughly the clique number and most frequent clique size in the graph $G(n, c, s)$, respectively. Throughout, we restrict our attention to k with $k \leq \kappa + 1$ since the probability that the largest clique in G has size $\omega(G) > \kappa + 1$ is $1/\text{poly}(n)$.

The threshold $\tau + 1$ also has a natural interpretation as the k -clique percolation threshold [DPV05; PDV07; DGM08], which roughly corresponds to the largest value of k at which a local search algorithm can explore all the cliques in the hypergraph starting from any given clique. In particular, in the graph case of $s = 2$, $\tau + 1$ is the largest integer k such that $\alpha < \frac{1}{k-1}$, which is exactly the k -clique percolation threshold described below. Given a hypergraph G , define two k -cliques of G to be adjacent if they share $(k-1)$ of their k vertices. This induces a hypergraph G_k on the set of k -cliques. For graphs G drawn from $G(n, c)$, [DPV05] introduced the k -clique percolation threshold of $c = \frac{1}{k-1} \cdot n^{-\frac{1}{k-1}}$, above which a giant component emerges in G_k . This threshold and extensions were rigorously established in [BR09]. Following the same heuristic as in [DPV05], our threshold $\tau + 1$ is a natural extension of the k -clique percolation threshold to the hypergraph case of $s \geq 3$.

Comparing our upper and lower bounds. A comparison of our algorithmic guarantees and average-case lower bounds based on the best known worst-case algorithms for counting k -cliques is shown in Figure 1.

³For the sub-class of algorithms that enumerate k -cliques one by one, the k -clique count is a trivial lower bound on the runtime. Our general lower bound matches this heuristic lower bound.

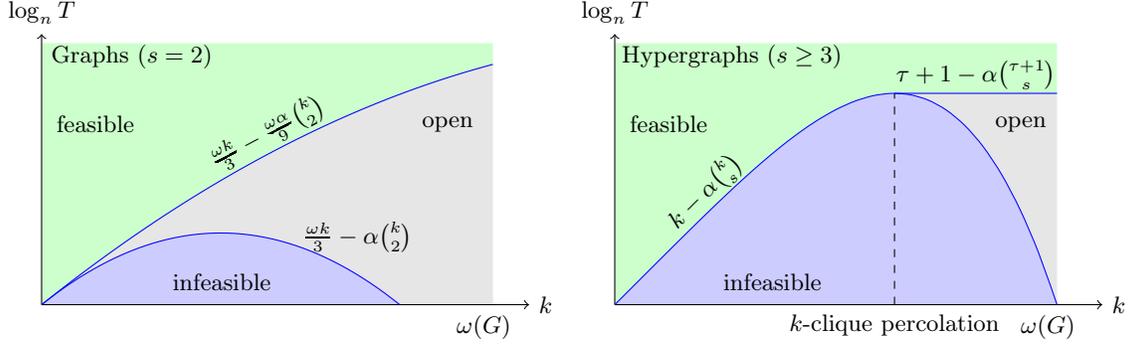


Figure 1: Comparison of our algorithms and average-case lower bounds for counting k -cliques in sparse Erdős-Rényi Hypergraphs $G(n, c, s)$ with $c = \Theta(n^{-\alpha})$. Green denotes runtimes T feasible for each k , blue denotes T infeasible given that the best known worst-case algorithms are optimal and gray denotes T for which the complexity of counting k -cliques is open after this work. The left plot shows the graph case of $s = 2$ and the right plot shows the hypergraph case of $s \geq 3$. For simplicity, all quantities shown are up to constant $O_{k,\alpha}(1)$ additive error.

1. *Graph Case ($s = 2$).* In the graph case, our lower and upper bounds have the same form and show that the exponent in the optimal running time is $\frac{\omega k}{3} - C\alpha \binom{k}{2} + O_{k,\alpha}(1)$ where $\frac{\omega}{9} \leq C \leq 1$ as long as $k \leq \kappa + 1 = 2\alpha^{-1} + 1$. As shown in Figure 1, our upper and lower bounds approach each other for k small relative to $\kappa + 1$.
2. *Hypergraph Case ($s \geq 3$).* In the hypergraph case of $s \geq 3$, the exponents in our lower and upper bounds are nearly identical at $k - \alpha \binom{k}{s} + O_{k,\alpha}(1)$ up to the k -clique percolation threshold. After this threshold, our lower bounds slowly deteriorate relative to our algorithms until they become trivial at the clique number of G by $k = \kappa + 1$.

Because we consider sparse Erdős-Rényi hypergraphs, for each n, k , and s we actually have an entire family of problems parametrized by the edge probability c and the behavior changes as a function of c ; this is the first worst-to-average-case hardness result we are aware of for which the complexity of the same problem over worst-case versus average-case inputs is completely different and can be sharply characterized over the whole range of c starting from the same assumption. It is surprising that our worst-case to average-case reduction techniques – which range from the self-reducibility of polynomials to random binary expansions – together yield tight lower bounds matching our algorithms in the hypergraph case.

Two interesting problems left open by our work are to show average-case lower bounds with an improved constant C in the graph case and to show tight average-case lower bounds beyond the k -clique percolation threshold in the case $s \geq 3$. These, other open problems and some extensions of our methods are discussed in Section 6.

1.2 Overview of Reduction Techniques

For clarity of exposition, in this section we will restrict our discussion to the graph case $s = 2$, as well as the case of constant k . A key step of our worst-case to average-case reduction uses the random self-reducibility of multivariate low-degree polynomials – i.e., evaluating a polynomial on any worst-case input can be efficiently reduced to evaluating it on several random inputs. This result follows from a line of work [Lip89; FF93; Gem+91; GS92] that provides a method to efficiently compute a polynomial $P : \mathbb{F}^N \rightarrow \mathbb{F}$ of degree $d \leq |\mathbb{F}|/20$ on any worst-case input $x \in \mathbb{F}^N$, given an

oracle $\tilde{P} : \mathbb{F}^N \rightarrow \mathbb{F}$ that agrees with P on a $\frac{1}{2} + \frac{1}{\text{poly}(N)}$ fraction of inputs. Thus, for any low-degree polynomial over a large enough finite field, evaluating the polynomial on a random element in the finite field is roughly as hard as evaluating the polynomial on any adversarially chosen input.

Random self-reducibility for counting k -cliques. With the random self-reducibility of polynomials in mind, a natural approach is to express the number of k -cliques in a graph as a low-degree polynomial of the $n \times n$ adjacency matrix A

$$P(A) = \sum_{\substack{S \subset [n] \\ |S|=k}} \left(\prod_{i < j \in S} A_{ij} \right).$$

This polynomial has been used in a number of papers, including by Goldreich and Rothblum [GR18] to construct a distribution on dense graphs for which counting k -cliques is provably hard on average. However, their techniques are primarily focused on the error probability requirement for the average-case blackbox. As a result, the distribution they obtain is far from Erdős-Rényi and their approach does not yield tight bounds for sparse graphs.

The significant obstacle that arises in applying the random self-reducibility of P is that one needs to work over a large enough finite field \mathbb{F}_p , so evaluating P on worst-case graph inputs in $\{0, 1\}^{\binom{n}{2}}$ only reduces to evaluating P on uniformly random inputs in $\mathbb{F}_p^{\binom{n}{2}}$. In order to further reduce to evaluating P on graphs, given a random input $A \in \mathbb{F}_p^{\binom{n}{2}}$ [GR18] uses several gadgets (including replacing vertices by independent sets and taking disjoint unions of graphs) in order to create a larger unweighted random graph A' whose k -clique count is equal to $k! \cdot P(A) \pmod{p}$ for appropriate p . However, any nontrivial gadget-based reduction seems to have little hope of arriving at something close to the Erdős-Rényi distribution, because gadgets inherently create non-uniform structure.

Reducing to k -partite graphs. We instead consider a different polynomial for graphs on nk vertices with $nk \times nk$ adjacency matrix A ,

$$P'(A) = \sum_{v_1 \in [n]} \sum_{v_2 \in [2n] \setminus [n]} \cdots \sum_{v_k \in [kn] \setminus [(k-1)n]} \left(\prod_{1 \leq i < j \leq k} A_{v_i v_j} \right).$$

The polynomial P' correctly counts the number of k -cliques if A is k -partite with vertex k -partition $[n] \sqcup ([2n] \setminus [n]) \sqcup \cdots \sqcup ([kn] \setminus [(k-1)n])$. We first reduce clique-counting in the worst case to computing P' in the worst case; this is a simple step, because it is a purely worst-case reduction. Next, we construct a recursive counting procedure that reduces evaluating P' on Erdős-Rényi graphs to counting k -cliques in Erdős-Rényi graphs. Therefore, it suffices to prove that if evaluating P' is hard in the worst case, then evaluating P' on Erdős-Rényi graphs is also hard.

Applying the Chinese Remainder theorem as well as the random self-reducibility of polynomials, computing P' on worst-case inputs in $\{0, 1\}^{\binom{nk}{2}}$ reduces to computing P' on several uniformly random inputs in $\mathbb{F}_p^{\binom{nk}{2}}$, for several different primes p each on the order of $\Theta(\log n)$. The main question is: how can one evaluate P' on inputs $X \sim \text{Unif}[\mathbb{F}_p^{\binom{nk}{2}}]$ using an algorithm that evaluates P' on $G(n, c, 2)$ Erdős-Rényi graphs (i.e., inputs $Y \sim \text{Ber}(c)^{\otimes \binom{nk}{2}}$)?

Eliminating weights with random sparse binary expansions. We solve this by decomposing the random weighted graph $X \sim \text{Unif}[\mathbb{F}_p^{\binom{nk}{2}}]$ into a weighted sum of graphs $Y^{(0)}, \dots, Y^{(t)} \in \{0, 1\}^{\binom{nk}{2}}$ such that each $Y^{(i)}$ is close to Erdős-Rényi $G(n, c, 2)$. Specifically, this additive decomposition satisfies $X \equiv \sum_{i=0}^t 2^i Y^{(i)} \pmod{p}$, i.e., that we can write X as a binary expansion modulo p of Erdős-Rényi graphs. Importantly, in Section 4 we derive near-optimal bounds on t and prove that we can take t to be quite small, growing only as $\text{poly}(c^{-1}(1-c)^{-1} \log(p))$. This technique seems likely to have applications elsewhere. For the unbiased case of $c = 1/2$, a version of this binary expansions technique appeared previously in [GR17].

Now, using the binary expansion decomposition of X , we algebraically manipulate P' as follows:

$$\begin{aligned} P'(X) &= \sum_{v_1 \in [n]} \sum_{v_2 \in [2n] \setminus [n]} \cdots \sum_{v_k \in [kn] \setminus [(k-1)n]} \prod_{1 \leq i < j \leq k} \left(\sum_{l \in \{0, \dots, t\}} 2^l \cdot Y_{v_i v_j}^{(l)} \right) \\ &= \sum_{f \in \{0, \dots, t\}^{\binom{k}{2}}} \left(\prod_{1 \leq i < j \leq k} 2^{f_{ij}} \right) \\ &\quad \times \left(\sum_{v_1 \in [n]} \sum_{v_2 \in [2n] \setminus [n]} \cdots \sum_{v_k \in [kn] \setminus [(k-1)n]} \prod_{1 \leq i < j \leq k} Y_{v_i v_j}^{(f_{ij})} \right) \\ &= \sum_{f \in \{0, \dots, t\}^{\binom{k}{2}}} \left(\prod_{1 \leq i < j \leq k} 2^{f_{ij}} \right) P'(Y^{(f)}). \end{aligned}$$

Here $Y^{(f)}$ is the nk -vertex graph with entries given by $Y_{ab}^{(f_{\bar{a}\bar{b}})}$ for $1 \leq a < b \leq nk$, where $\bar{a} = \lceil a/n \rceil$ and $\bar{b} = \lceil b/n \rceil$. We thus reduce the computation of $P'(X)$ to the computation of a weighted sum of $\text{poly}(c^{-1}(1-c)^{-1} \log(n))^{\binom{k}{2}}$ different evaluations of P' at graphs close in total variation to $G(n, c, 2)$. This concludes our reduction.⁴

We remark that an important difference between our reduction and the reduction in [GR18] is the number of and structure of the calls to the average-case blackbox. Our reduction requires many successful calls to the blackbox in order to obtain a single correct evaluation of the polynomial $P'(A)$, which is where our low error probability requirement comes from. The gadgets in [GR18] are specifically designed to only require a single successful call to obtain a single correct evaluation of $P(A)$. Thus even given a blackbox with a constant error probability, the Berkelamp-Welch algorithm can recover $P(A)$ in the case of [GR18].

We also give a different worst-case to average-case reduction for determining the parity of the number of k -cliques in Erdős-Rényi hypergraphs, as discussed in Sections 2.2 and 3.

1.3 Related Work on Worst-Case to Average-Case Reductions

The random self-reducibility of low-degree polynomials serves as the basis for several worst-case to average-case reductions found in the literature. One of the first applications of this method was

⁴If we had instead worked with P , then this argument would fail. The argument uses the k -partiteness structure of P' as follows: for every pair of vertices $a, b \in [nk]$ and $f \in \{0, \dots, t\}^{\binom{k}{2}}$, the term $Y_{ab}^{(f_{ij})}$ appearing in the sum is uniquely determined by $a \in [ik] \setminus [(i-1)k]$ and $b \in [jk] \setminus [(j-1)k]$. So given f we can define a graph $Y^{(f)}$ uniquely. On the other hand, running the same argument with the polynomial P , the term $Y_{ab}^{(f_{ij})}$ for many different i, j would appear in the sum, and there is no way to uniquely define a graph $Y^{(f)}$.

to prove that the permanent is hard to evaluate on random inputs, even with polynomially-small probability of success, unless $\mathsf{P}^{\#\mathsf{P}} = \mathsf{BPP}$ [Sud97; CPS99]. (Under the slightly stronger assumption that $\mathsf{P}^{\#\mathsf{P}} \neq \mathsf{AM}$, and with different techniques, [FL92] proved that computing the permanent on large finite fields is hard even with exponentially small success probability.) Recently, [Bal+17] used the polynomial random self-reducibility result in the fine-grained setting in order to construct polynomials that are hard to evaluate on most inputs, assuming fine-grained hardness conjectures for problems such as 3-SUM, ORTHOGONAL-VECTORS, and/or ALL-PAIRS-SHORTEST-PATHS. The random self-reducibility of polynomials was also used by Gamarnik [Gam18] in order to prove that exactly computing the partition function of the Sherrington-Kirkpatrick model in statistical physics is hard on average.

If a problem is random self-reducible, then random instances of the problem are essentially as hard as worst-case instances, and therefore one may generate a hard instance of the problem by simply generating a random instance. Because of this, random self-reducibility plays an important role in cryptography: it allows one to base cryptographic security on random instances of a problem, which can generally be generated efficiently. A prominent example of a random-self reducible problem with applications to cryptography is the problem of finding a short vector in a lattice. In a seminal paper, Ajtai [Ajt96] gave a worst-case to average-case reduction for this short-vector problem. His ideas were subsequently applied to prove the average-case hardness of the Learning with Errors (LWE) problem, which underlies lattice cryptography [Ajt96; Reg09]. A good survey covering worst-case to average-case reductions in lattice cryptography is [Reg].

There are known restrictions on problems that are self-reducible. For example, non-adaptive worst-case to average-case reductions for NP-complete problems fail unless $\mathsf{coNP} \subseteq \mathsf{NP}/\text{poly}$ [FF93; BT06b; BT06a].

1.4 Notation and Preliminaries

A s -uniform hypergraph $G = (V(G), E(G))$ consists of a vertex set $V(G)$ and a hyperedge set $E(G) \subseteq \binom{V(G)}{s}$. A k -clique C in G is a subset of vertices $C \subset V(G)$ of size $|C| = k$ such that all of the possible hyperedges between the vertices are present in the hypergraph: $\binom{C}{s} \subseteq E(G)$. We write $\text{cl}_k(G)$ to denote the set of k -cliques of the hypergraph G . One samples from the Erdős-Rényi distribution $G(n, c, s)$ by independently including each of the $\binom{n}{s}$ hyperedges with probability c .

We denote the law of a random variable X by $\mathcal{L}(X)$. We use $T(A, n)$ to denote the worst-case run-time of an algorithm A on inputs of size parametrized by n . We work in the Word RAM model of computation, where the words have $O(\log n)$ bits. All algorithms in this paper are randomized, and each (possibly biased) coin flip incurs constant computational cost.

2 Problem Formulations and Average-Case Lower Bounds

2.1 Clique Problems and Worst-Case Fine-Grained Conjectures

In this section, we formally define the problems we consider and the worst-case fine-grained complexity conjectures off of which our average-case lower bounds are based. We focus on the following computational problems.

Definition 2.1. $\#(k, s)$ -CLIQUE denotes the problem of counting the number of k -cliques in an s -uniform hypergraph G .

Definition 2.2. $\text{PARITY-}(k, s)$ -CLIQUE denotes the problem of counting the number of k -cliques up to parity in an s -uniform hypergraph G .

Definition 2.3. $\text{DECIDE-}(k, s)\text{-CLIQUE}$ denotes the problem of deciding whether or not an s -uniform hypergraph G contains a k -clique.

Both $\#(k, s)\text{-CLIQUE}$ and $\text{DECIDE-}(k, s)\text{-CLIQUE}$ are fundamental problems that have long been studied in computational complexity theory and are conjectured to be computationally hard in the worst-case setting. When k is allowed to be an unbounded input to the problem, $\text{DECIDE-}(k, s)\text{-CLIQUE}$ is known to be NP-complete [Kar72] and $\#(k, s)\text{-CLIQUE}$ is known to be $\#\text{P}$ -complete [Val79]. In this work, we consider the fine-grained complexity of these problems, where k either can be viewed as a constant or a very slow-growing parameter compared to the number n of vertices of the hypergraph. In this context, $\text{PARITY-}(k, s)\text{-CLIQUE}$ can be interpreted as an intermediate problem between the other two clique problems that we consider. The worst-case reduction from $\text{PARITY-}(k, s)\text{-CLIQUE}$ to $\#(k, s)\text{-CLIQUE}$ is immediate. As we show in Appendix A, in the worst-case setting, $\text{DECIDE-}(k, s)\text{-CLIQUE}$ also reduces to $\text{PARITY-}(k, s)\text{-CLIQUE}$ with a multiplicative overhead of $O(k2^k)$ time.

When k is a constant, the trivial brute-force search algorithms for these problems are efficient in the sense that they take polynomial time. However, these algorithms do not remain efficient under the lens of fine-grained complexity since brute-force search requires $\Theta(n^k)$ time, which can grow significantly as k grows. In the hypergraph case of $s \geq 3$, no algorithm taking time $O(n^{k-\epsilon})$ on any of these problems is known, including for $\text{DECIDE-}(k, s)\text{-CLIQUE}$ [Yus06]. In the graph case of $s = 2$, the fastest known algorithms for all of these problems take $\Theta(n^{\omega \lceil k/3 \rceil})$ time, where $2 \leq \omega < 2.4$ is the fast matrix multiplication constant [IR78; NP85]. Since this is the state of the art, one may conjecture that $\text{DECIDE-}(k, s)\text{-CLIQUE}$ and $\#(k, s)\text{-CLIQUE}$ take $n^{\Omega(k)}$ time in the worst case.

Supporting this conjecture, Razborov [Raz85] proves that monotone circuits require $\tilde{\Omega}(n^k)$ operations to solve $\text{DECIDE-}(k, 2)\text{-CLIQUE}$ in the case of constant k . Monotone circuit lower bounds are also known in the case when $k = k(n)$ grows with n [AB87; AM05]. In [DF95], $\text{DECIDE-}(k, 2)\text{-CLIQUE}$ is shown to be $\text{W}[1]$ -hard. In other words, this shows that if $\text{DECIDE-}(k, 2)\text{-CLIQUE}$ is fixed-parameter tractable – admits an algorithm taking time $f(k) \cdot \text{poly}(n)$ – then any algorithm in the parametrized complexity class $\text{W}[1]$ is also fixed-parameter-tractable. This provides further evidence that $\text{DECIDE-}(k, 2)\text{-CLIQUE}$ is intractable for large k . Finally, [Che+06] shows that solving $\text{DECIDE-}(k, 2)\text{-CLIQUE}$ in $n^{o(k)}$ time is ETH-hard for constant k^5 . We therefore conjecture that the k -clique problems take $n^{\Omega(k)}$ time on worst-case inputs when k is constant, as formalized below.

Conjecture 2.4 (Worst-case hardness of $\#(k, s)\text{-CLIQUE}$). *Let k be constant. Any randomized algorithm A for $\#(k, s)\text{-CLIQUE}$ with error probability less than $1/3$ takes time at least $n^{\Omega(k)}$ in the worst case for hypergraphs on n vertices.*

Conjecture 2.5 (Worst-case hardness of $\text{PARITY-}(k, s)\text{-CLIQUE}$). *Let k be constant. Any randomized algorithm A for $\text{PARITY-}(k, s)\text{-CLIQUE}$ with error probability less than $1/3$ takes time at least $n^{\Omega(k)}$ in the worst case for hypergraphs on n vertices.*

Conjecture 2.6 (Worst-case hardness of $\text{DECIDE-}(k, s)\text{-CLIQUE}$). *Let k be constant. Any randomized algorithm A for $\text{DECIDE-}(k, s)\text{-CLIQUE}$ with error probability less than $1/3$ takes time at least $n^{\Omega(k)}$ in the worst case for hypergraphs on n vertices.*

The conjectures are listed in order of increasing strength. Since Conjecture 2.6 is implied by rETH, they all follow from rETH. We also formulate a stronger version of the clique-counting

⁵These hardness results also apply to $\text{DECIDE-}(k, s)\text{-CLIQUE}$ for $s \geq 3$ since there is a reduction from $\text{DECIDE-}(k, 2)\text{-CLIQUE}$ to $\text{DECIDE-}(k, s)\text{-CLIQUE}$ in n^s time. The reduction proceeds by starting with a graph G and constructing an s -uniform hypergraph G' that contains a s -hyperedge for every s -clique in G . The k -cliques of G and G' are in bijection. This construction also reduces $\#(k, 2)\text{-CLIQUE}$ to $\#(k, s)\text{-CLIQUE}$.

hardness conjecture, which asserts that the current best known algorithms for k -clique counting are optimal.

Conjecture 2.7 (Strong worst-case hardness of $\#(k, s)$ -CLIQUE). *Let k be constant. Any randomized algorithm A for $\#(k, s)$ -CLIQUE with error probability less than $1/3$ takes time $\tilde{\Omega}(n^{\omega\lceil k/3\rceil})$ in the worst case if $s = 2$ and $\tilde{\Omega}(n^k)$ in the worst case if $s \geq 3$.*

2.2 Average-Case Lower Bounds for Counting k -Cliques in $G(n, c, s)$

Our first main result is a worst-case to average-case reduction solving either $\#(k, s)$ -CLIQUE or PARITY- (k, s) -CLIQUE on worst-case hypergraphs given a blackbox solving the problem on *most* Erdős-Rényi hypergraphs drawn from $G(n, c, s)$. We discuss this error tolerance over sampling Erdős-Rényi hypergraphs as well as the multiplicative overhead in our reduction below. These results show that solving the k -clique problems on Erdős-Rényi hypergraphs $G(n, c, s)$ is as hard as solving them on worst-case hypergraphs, for certain choices of k, c and s . Therefore the worst-case hardness assumptions, Conjectures 2.4, 2.5 and 2.7, imply average-case hardness on Erdős-Rényi hypergraphs for $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE.

Theorem 2.8 (Worst-case to average-case reduction for $\#(k, s)$ -CLIQUE). *There is an absolute constant $C > 0$ such that if we define*

$$\Upsilon_{\#}(n, c, s, k) \triangleq (C(c^{-1}(1-c)^{-1})(s \log k + s \log \log n)(\log n))^{\binom{k}{s}}$$

then the following statement holds. Let A be a randomized algorithm for $\#(k, s)$ -CLIQUE with error probability less than $1/\Upsilon_{\#}$ on hypergraphs drawn from $G(n, c, s)$. Then there exists an algorithm B for $\#(k, s)$ -CLIQUE that has error probability less than $1/3$ on any hypergraph, such that

$$T(B, n) \leq (\log n) \cdot \Upsilon_{\#} \cdot (T(A, nk) + (nk)^s),$$

where $T(\mathcal{A}, \ell)$ denotes the runtime of algorithm \mathcal{A} on ℓ -vertex hypergraphs.

For PARITY- (k, s) -CLIQUE we also give an alternative reduction with an improved reduction time and error tolerance in the dense case when $c = 1/2$.

Theorem 2.9 (Worst-case to average-case reduction for PARITY- (k, s) -CLIQUE). *We have that:*

1. *There is an absolute constant $C > 0$ such that if we define*

$$\Upsilon_{P,1}(n, c, s, k) \triangleq \left(C(c^{-1}(1-c)^{-1})(s \log k) \left(s \log n + \binom{k}{s} \log \log \binom{k}{s} \right) \right)^{\binom{k}{s}}$$

then the following statement holds. Let A be a randomized algorithm for PARITY- (k, s) -CLIQUE with error probability less than $1/\Upsilon_{P,1}$ on hypergraphs drawn from $G(n, c, s)$. Then there exists an algorithm B for PARITY- (k, s) -CLIQUE that has error probability less than $1/3$ on any hypergraph, such that

$$T(B, n) \leq \Upsilon_{P,1} \cdot (T(A, nk) + (nk)^s)$$

2. *There is an absolute constant $C > 0$ such that if we define*

$$\Upsilon_{P,2}(s, k) \triangleq (Cs \log k)^{\binom{k}{s}}$$

then the following statement holds. Let A be a randomized algorithm for PARITY- (k, s) -CLIQUE with error probability less than $1/\Upsilon_{P,2}$ on hypergraphs drawn from $G(n, 1/2, s)$. Then there exists an algorithm B for PARITY- (k, s) -CLIQUE that has error probability less than $1/3$ on any hypergraph, such that

$$T(B, n) \leq \Upsilon_{P,2} \cdot (T(A, nk) + (nk)^s)$$

Our worst-case to average-case reductions yield the following fine-grained average-case lower bounds for k -clique counting and parity on Erdős-Rényi hypergraphs based on Conjectures 2.4 and 2.7. We separate these lower bounds into the two cases of dense and sparse Erdős-Rényi hypergraphs. We remark that, for all constants k , an error probability of less than $(\log n)^{-\omega(1)}$ suffices in the dense case and error probability less than $n^{-\omega(1)}$ suffices in the sparse case.

Corollary 2.10 (Average-case hardness of $\#(k, s)$ -CLIQUE on dense $G(n, c, s)$). *If $k, c, \epsilon > 0$ are constant, then we have that*

1. *Assuming Conjecture 2.4, then any algorithm A for $\#(k, s)$ -CLIQUE that has error probability less than $(\log n)^{-\binom{k}{s}-\epsilon}$ on Erdős-Rényi hypergraphs drawn from $G(n, c, s)$ must have runtime at least $T(A, n) \geq n^{\Omega(k)}$.*
2. *Assuming Conjecture 2.7, then any algorithm A for $\#(k, s)$ -CLIQUE that has error probability less than $(\log n)^{-\binom{k}{s}-\epsilon}$ on Erdős-Rényi hypergraphs drawn from $G(n, c, s)$ must have runtime at least $T(A, n) \geq \tilde{\Omega}(n^{\omega\lceil k/3 \rceil})$ if $s = 2$ and $T(A, n) \geq \tilde{\Omega}(n^k)$ if $s \geq 3$.*

Corollary 2.11 (Average-case hardness of $\#(k, s)$ -CLIQUE on sparse $G(n, c, s)$). *Let $k, \alpha, \epsilon > 0$ be constants and $c = \Theta(n^{-\alpha})$. Assuming Conjecture 2.7, then any algorithm A for $\#(k, s)$ -CLIQUE that has error probability less than $n^{-\alpha\binom{k}{s}-\epsilon}$ on Erdős-Rényi hypergraphs drawn from $G(n, c, s)$ must have runtime at least $T(A, n) \geq \tilde{\Omega}(n^{\omega\lceil k/3 \rceil - \alpha\binom{k}{s}})$ if $s = 2$ and $T(A, n) \geq \tilde{\Omega}(n^{k - \alpha\binom{k}{s}})$ if $s \geq 3$.*

We remark that Conjecture 2.4 implies there is a constant $C > 0$ such that a version of Corollary 2.11 holds with the weaker conclusion $T(A, n) \geq n^{\Omega(k)}$ holds for all $\alpha \leq C\binom{k}{s}$. For PARITY- (k, s) -CLIQUE, we consider here the implications of Theorem 2.9 only for $c = 1/2$, since this is the setting in which we obtain substantially different lower bounds than for $\#(k, s)$ -CLIQUE. As shown, an error probability of $o(1)$ on $G(n, 1/2, s)$ hypergraphs suffices for our reduction to succeed.

Corollary 2.12 (Average-case hardness of PARITY- (k, s) -CLIQUE on $G(n, 1/2, s)$). *Let k be constant. Assuming Conjecture 2.5, there is a small enough constant $\epsilon \triangleq \epsilon(k, s)$ such that if any algorithm A for PARITY- (k, s) -CLIQUE has error less than ϵ on $G(n, 1/2, s)$ then A must have runtime at least $T(A, n) \geq n^{\Omega(k)}$.*

We remark on one subtlety of our setup in the sparse case. Especially in our algorithms section, we generally restrict our attention to $c = \Theta(n^{-\alpha})$ satisfying $\alpha \leq k\binom{k}{s}^{-1} = s\binom{k}{s-1}^{-1}$, which is necessary for the expected number of k -cliques in $G(n, c, s)$ to not tend to zero. However, even when this expectation is decaying, the problem $\#(k, s)$ -CLIQUE as we formulate it is still nontrivial. The simple algorithm that always outputs zero fails with a polynomially small probability that does not appear to meet the $1/\Upsilon_{\#}$ requirement in our worst-case to average-case reduction. A simple analysis of this error probability can be found in Lemma 5.1. Note that even when $\alpha > s\binom{k}{s-1}^{-1}$, GREEDY-RANDOM-SAMPLING and its derivative algorithms in Section 5 still has guarantees and succeeds with probability $1 - n^{-\omega(1)}$. We now discuss the multiplicative overhead and error tolerance in our worst-case to average-case reduction for $\#(k, s)$ -CLIQUE.

Discussion of the Multiplicative Slowdown $\Upsilon_{\#}$ In the sparse case of $c = \Theta(n^{-\alpha})$, our algorithmic upper bounds in Section 5 imply lower bounds on the multiplicative overhead factor $\Upsilon_{\#}$ in Theorem 2.8. In the hypergraph case of $s \geq 3$ and below the k -clique percolation threshold, it must follow that the overhead is at least $\Upsilon_{\#} = \tilde{\Omega}\left(n^{\alpha \binom{k}{s}}\right) = \tilde{\Omega}\left(c^{-\binom{k}{s}}\right)$. Otherwise, our algorithms combined with our worst-case to average-case reduction would contradict Conjecture 2.7. Up to polylog(n) factors, this exactly matches the $\Upsilon_{\#}$ from our reduction. In the graph case of $s = 2$, it similarly must follow that the overhead is at least $\Upsilon_{\#} = \tilde{\Omega}\left(n^{\frac{\omega\alpha}{9} \binom{k}{s}}\right) = \tilde{\Omega}\left(c^{-\frac{\omega}{9} \binom{k}{s}}\right)$ to not contradict Conjecture 2.7. This matches the $\Upsilon_{\#}$ from our reduction up to a constant factor in the exponent.

Discussion of the Error Tolerance $1/\Upsilon_{\#}$ Notice that our worst-case to average-case reductions in Theorems 2.8 and 2.9 require that the error of the average-case blackbox on Erdős-Rényi hypergraphs go to zero as k goes to infinity. This error tolerance requirement is unavoidable. When $k = \omega(\log n)$ in the dense Erdős-Rényi graph case of $G(n, 1/2)$, there is a k -clique with at most $\binom{n}{k} 2^{-\binom{k}{2}} = o(1)$ probability by a union bound on k -subsets of vertices. So in this regime clique-counting on $G(n, 1/2)$ with constant error probability is not hard: the algorithm that always outputs zero achieves $o(1)$ average-case error.

If $k \triangleq 3 \log_2 n$, then the probability of a k -clique on $G(n, 1/2)$ is less than $\binom{n}{k} 2^{-\binom{k}{2}} \leq 2^{-k^2/6}$. So average-case k -clique counting is not hard with error more than $2^{-k^2/6}$. On the other hand, our $\#(k, 2)$ -CLIQUE reduction works with average-case error less than $1/\Upsilon_{\#} = 2^{-\Omega(k^2 \log \log n)}$. And our PARITY- $(k, 2)$ -CLIQUE reduction is more lenient, requiring error only less than $2^{-\Omega(k^2 \log \log \log n)}$. Thus, the error bounds required by our reductions are quite close to the $2^{-k^2/6}$ error bound that is absolutely necessary for any reduction in this regime.

In the regime where $k = O(1)$ is constant and on $G(n, 1/2)$, our PARITY- $(k, 2)$ -CLIQUE reduction only requires a small constant probability of error and our $\#(k, 2)$ -CLIQUE reduction requires less than a $1/\text{polylog}(n)$ probability of error. We leave it as an intriguing open problem whether the error tolerance of our reductions can be improved in this regime.

Finally, we remark that the error tolerance of the reduction must depend on c . The probability that a $G(n, c)$ graph contains a k -clique is less than $(nc^{(k-1)/2})^k$. For example, if $c = 1/n$ then the probability that there exists a k -clique is less than $n^{-\Omega(k^2)}$. As a result, no worst-case to average-case reduction can tolerate average-case error more than $n^{-O(k^2)}$ on $G(n, 1/n)$ graphs. And therefore our reductions for $\#(k, 2)$ -CLIQUE and for PARITY- $(k, 2)$ -CLIQUE are close to optimal when $c = 1/n$, because our error tolerance scales as $n^{-O(k^2 \log \log n)}$.

3 Worst-Case to Average-Case Reduction for $G(n, c, s)$

In this section, we give our main worst-case to average-case reduction that transforms a blackbox solving $\#(k, s)$ -CLIQUE on $G(n, c, s)$ into a blackbox solving $\#(k, s)$ -CLIQUE on a worst-case input hypergraph. This also yields a worst-case to average-case reduction for PARITY- (k, s) -CLIQUE and proves Theorems 2.8 and 2.9. The reduction involves the following five main steps, the details of which are in Sections 3.1 to 3.5.

1. Reduce $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE on general worst-case hypergraphs to the worst-case problems with inputs that are k -partite hypergraphs with k parts of equal size.
2. Reduce the worst-case problem on k -partite hypergraphs to the problem of computing a low-degree polynomial $P_{n,k,s}$ on $N \triangleq N(n, k, s)$ variables over a small finite field \mathbb{F} .

3. Reduce the problem of computing $P_{n,k,s}$ on worst-case inputs to computing $P_{n,k,s}$ on random inputs in \mathbb{F}^N .
4. Reduce the problem of computing $P_{n,k,s}$ on random inputs in \mathbb{F}^N to computing $P_{n,k,s}$ on random inputs in $\{0, 1\}^N$. This corresponds to $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE on k -partite Erdős-Rényi hypergraphs.
5. Reduce the resulting average-case variants of $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE on k -partite Erdős-Rényi hypergraphs to non- k -partite Erdős-Rényi hypergraphs.

These steps are combined in Section 3.6 to complete the proofs of Theorems 2.8 and 2.9. Before proceeding to our worst-case to average-case reduction, we establish some definitions and notation, and also give pseudocode for the counting reduction in Figure 2 – the parity reduction is similar.

The intermediate steps of our reduction crucially make use of k -partite hypergraphs with k parts of equal size, defined below.

Definition 3.1 (*k*-Partite Hypergraphs). *Given a s -uniform hypergraph G on nk vertices with vertex set $V(G) = [n] \times [k]$, define the vertex labelling*

$$L : (i, j) \in [n] \times [k] \mapsto j \in [k]$$

If for all $e = \{u_1, \dots, u_s\} \in E(G)$, the labels $L(u_1), L(u_2), \dots, L(u_s)$ are distinct, then we say that G is k -partite with k parts of equal size n .

In our reduction, it suffices to consider only k -partite hypergraphs with k parts of equal size. For ease of notation, our k -partite hypergraphs will always have nk vertices and vertex set $[n] \times [k]$. In particular, the edge set of a k -partite s -uniform hypergraph is an arbitrary subset of

$$E(G) \subseteq \{\{u_1, \dots, u_s\} \subset V(G) : L(u_1), \dots, L(u_s) \text{ are distinct}\}$$

Taking edge indicators yields that the k -partite hypergraphs on nk vertices we consider are in bijection with $\{0, 1\}^N$, where $N \triangleq N(n, k, s) = \binom{k}{s} n^s$ is this size of this set of permitted hyperedges. Thus we will refer to elements $x \in \{0, 1\}^N$ and k -partite s -uniform hypergraphs on nk vertices interchangeably. This definition also extends to Erdős-Rényi hypergraphs.

Definition 3.2 (*k*-Partite Erdős-Rényi Hypergraphs). *The k -partite s -uniform Erdős-Rényi hypergraph $G(nk, c, s, k)$ is a distribution over hypergraphs on nk vertices with vertex set $V(G) = [n] \times [k]$. A sample from $G(nk, c, s, k)$ is obtained by independently including hyperedge each $e = \{u_1, \dots, u_s\} \in E(G)$ with probability c for all e with $L(u_1), L(u_2), \dots, L(u_s)$ distinct.*

Viewing the hypergraphs as elements of $G(nk, c, s, k)$ as a distribution on $\{0, 1\}^N$, it follows that $G(nk, c, s, k)$ corresponds to the product distribution $\text{Ber}(c)^{\otimes N}$.

3.1 Worst-Case Reduction to k -Partite Hypergraphs

In the following lemma, we prove that the worst-case complexity of $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE are nearly unaffected when we restrict the inputs to be worst-case k -partite hypergraphs. This step is important, because the special structure of k -partite hypergraphs will simplify future steps in our reduction.

Lemma 3.3. *Let A be an algorithm for $\#(k, s)$ -CLIQUE, such that A has error probability less than $1/3$ for any k -partite hypergraph G on nk vertices. Then, there is an algorithm B for $\#(k, s)$ -CLIQUE with error probability less than $1/3$ on any hypergraph G satisfying that $T(B, n) \leq T(A, n) + O(k^s n^s)$. Furthermore, the same result holds for PARITY- (k, s) -CLIQUE in place of $\#(k, s)$ -CLIQUE.*

Algorithm TO-ER- $\#(G, k, A, c)$

Inputs: s -uniform hypergraph G with vertex set $[n]$, parameters k, c , algorithm A for $\#(k, s)$ -CLIQUE on Erdős-Rényi hypergraphs with density c .

1. Construct an s -uniform hypergraph G' on vertex set $[n] \times [k]$ by defining

$$E(G') = \left\{ \{(v_1, t_1), (v_2, t_2), \dots, (v_s, t_s)\} \right. \\ \left. : \{v_1, \dots, v_s\} \in E(G) \text{ and } \begin{matrix} 1 \leq v_1 < v_2 < \dots < v_s \leq n \\ 1 \leq t_1 < t_2 < \dots < t_s \leq k \end{matrix} \right\}.$$

Since G' is k -partite, view it as an indicator vector of edges $G' \in \{0, 1\}^N$ for $N := N(n, k, s) = \binom{k}{s} n^s$.

2. Find the first T primes $12 \binom{k}{s} < p_1 < \dots < p_T$ such that $\prod_{i=1}^T p_i > n^k$.
3. Define $L : (a, b) \in [n] \times [k] \mapsto b \in [k]$, and

$$P_{n,k,s}(x) = \sum_{\substack{\{u_1, \dots, u_k\} \in V(G') \\ L(u_i) = i \ \forall i}} \prod_{\substack{S \subseteq [k] \\ |S| = s}} x_{u_S}$$

For each $1 \leq t \leq T$, compute $P_{n,k,s}(G') \pmod{p_t}$, as follows:

- (1) Use the procedure of [GS92] in order to reduce the computation of $P_{n,k,s}(G') \pmod{p_t}$ to the computation of $P_{n,k,s}$ on $M = 12 \binom{k}{s}$ distinct inputs $x_1, \dots, x_M \sim \text{Unif}[\mathbb{F}_{p_t}^N]$.
- (2) For each $1 \leq m \leq M$, compute $P_{n,k,s}(x_m) \pmod{p_t}$ as follows:
 - (i) Use the rejection sampling procedure of Lemma 4.4 in order to sample $(Y^{(0)}, \dots, Y^{(B)})$ close to $(\text{Ber}(c)^{\otimes N})^{\otimes B}$ in total variation distance, such that $x_m \equiv \sum_{b=0}^B 2^b \cdot Y^{(b)} \pmod{p_t}$. It suffices to take $B = \Theta(c^{-1}(1-c)^{-1}s(\log n)(\log p_t))$.
 - (ii) For each function $a : \binom{[k]}{s} \rightarrow \{0, \dots, B\}$, define $Y_S^{(a)} = Y^{a(L(S))}$ for all $S \in [N] \subset \binom{[n]}{s}$. Note that for each a , the corresponding $Y^{(a)}$ is approximately distributed as $\text{Ber}(c)^{\otimes N}$. Use algorithm A and the recursive counting procedure of Lemma 3.10 in order to compute $P_{n,k,s}(Y^{(a)})$ for each a .
 - (iii) Set $P_{n,k,s}(G') \leftarrow \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, B\}} 2^{|a|_1} \cdot P_{n,k,s}(Y^{(a)})$.
4. Since $0 \leq P_{n,k,s}(G') \leq n^k$, use Chinese remaindering and the computations of $P_{n,k,s}(G') \pmod{p_i}$ in order to calculate and output $P_{n,k,s}(G')$.

Figure 2: Reduction TO-ER- $\#$ for showing computational lower bounds for average-case $\#(k, s)$ -CLIQUE on Erdős-Rényi $G(n, c, s)$ hypergraphs based on the worst-case hardness of $\#(k, s)$ -CLIQUE.

Proof. Let G be an s -uniform hypergraph on n vertices. Construct the s -uniform hypergraph G'

on the vertex set $V(G') = [n] \times [k]$ with edge set

$$E(G') = \left\{ \{(v_1, t_1), (v_2, t_2), \dots, (v_s, t_s)\} : \{v_1, \dots, v_s\} \in E(G) \text{ and } \begin{matrix} 1 \leq v_1 < v_2 < \dots < v_s \leq n \\ 1 \leq t_1 < t_2 < \dots < t_s \leq k \end{matrix} \right\}$$

The hypergraph G' can be constructed in $O(k^s n^s)$ time. Note that G' is k -partite with the vertex partition $L : (i, j) \in [n] \times [k] \mapsto j \in [k]$. There is also a bijective correspondence between k -cliques in G' and k -cliques in G given by

$$\{v_1, v_2, \dots, v_k\} \mapsto \{(v_1, 1), (v_2, 2), \dots, (v_k, k)\}$$

where $v_1 < v_2 < \dots < v_k$. Thus, the k -partite s -uniform hypergraph G' on nk vertices has exactly the same number of k -cliques as G . It suffices to run A on G' and to return its output. \square

A corollary to Lemma 3.3 is that any worst-case hardness for $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE on general s -uniform hypergraphs immediately transfers to the k -partite case. For instance, the lower bounds of Conjectures 2.4, 2.5, and 2.7 imply corresponding lower bounds in the k -partite case. Going forward in our worst-case to average-case reduction, we may restrict our attention to k -partite hypergraphs without loss of generality.

3.2 Counting k -Cliques as a Low-Degree Polynomial

A key step in our worst-case to average-case reduction is to express the number of k -cliques as a low-degree polynomial in the adjacency matrix. As mentioned in the introduction, a similar step – but without the k -partiteness constraint – appears in the worst-case to average-case reduction of Goldreich and Rothblum [GR18].

Let $\mathcal{E} \subset \binom{V(G)}{s}$ be the set of possible hyperedges that respect the k -partition: i.e., $\mathcal{E} = \{A \in \binom{V(G)}{s} : |L(A)| = s\}$. Let $N \triangleq N(n, k, s) = |\mathcal{E}|$ and identify \mathcal{E} with $[N]$ through a bijection $\pi : [N] \rightarrow \mathcal{E}$. To simplify the notation, we will omit the map π in the proof, and simply treat $[N]$ and \mathcal{E} as the same set. Thus, each $x \in \{0, 1\}^N$ corresponds to a k -partite hypergraph where x_A is the indicator that $A \in \mathcal{E}$ is an edge in the hypergraph. The number of k -cliques of a k -partite hypergraph $x \in \{0, 1\}^N$ is a degree- D polynomial $P_{n,k,s} : \{0, 1\}^N \rightarrow \mathbb{Z}$ where $D \triangleq D(k, s) = \binom{k}{s}$:

$$P_{n,k,s}(x) = \sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall i L(u_i) = i}} \prod_{\substack{S \subset [k] \\ |S| = s}} x_{u_S} \quad (1)$$

For any finite field \mathbb{F} , this equation defines $P_{n,k,s}$ as a polynomial over that finite field. For clarity, we write this polynomial over \mathbb{F} as $P_{n,k,s,\mathbb{F}} : \mathbb{F}^N \rightarrow \mathbb{F}$. Observe that for any hypergraph $x \in \{0, 1\}^N$, we have that

$$P_{n,k,s,\mathbb{F}}(x) = P_{n,k,s}(x) \pmod{\text{char}(\mathbb{F})}$$

where $\text{char}(\mathbb{F})$ is the characteristic of the finite field. We now reduce computing $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE on a k -partite hypergraph $x \in \{0, 1\}^N$ to computing $P_{n,k,s,\mathbb{F}}(x)$ for appropriate finite fields \mathbb{F} . This is formalized in the following two propositions.

Proposition 3.4. *Let $x \in \{0, 1\}^N$ denote a s -uniform hypergraph that is k -partite with vertex labelling L . Let p_1, p_2, \dots, p_t be t distinct primes, such that $\prod_i p_i > n^k$. Then, solving $\#(k, s)$ -CLIQUE reduces to computing $P_{n,k,s,\mathbb{F}_{p_i}}(x)$ for all $i \in [t]$, plus $O(k \log n)$ additive computational overhead. Moreover, computing $P_{n,k,s,\mathbb{F}_{p_i}}(x)$ for all $i \in [t]$ reduces to computing $\#(k, s)$ -CLIQUE, plus $O(tk \log n)$ computational overhead.*

Proof. Note that $P_{n,k,s}(x) \leq n^k$ since there are at most n^k cliques in the hypergraph. So the claim follows from the Chinese Remainder Theorem and the fact that for any $i \in [t]$, it holds that $P_{n,k,s,\mathbb{F}_{p_i}}(x) \equiv P_{n,k,s}(x) \pmod{p_i}$. \square

Proposition 3.5. *Let \mathbb{F} be a finite field of characteristic 2. Let $x \in \{0,1\}^N$ be a s -uniform hypergraph that is k -partite with vertex labelling L . Then solving PARITY- (k,s) -CLIQUE for x is equivalent to computing $P_{n,k,s,\mathbb{F}}(x)$.*

Proof. This is immediate from $P_{n,k,s,\mathbb{F}}(x) \equiv P_{n,k,s}(x) \pmod{\text{char}(\mathbb{F})}$. \square

3.3 Random Self-Reducibility: Reducing to Random Inputs in \mathbb{F}^N

Expressing the number and parity of cliques as low-degree polynomials allows us to perform a key step in the reduction: because polynomials over finite fields are random self-reducible, we can reduce computing $P_{n,k,s,\mathbb{F}}$ on worst-case inputs to computing $P_{n,k,s,\mathbb{F}}$ on several uniformly random inputs in \mathbb{F}^N .

The following well-known lemma states the random self-reducibility of low-degree polynomials. The lemma first appeared in [GS92]. We follow the proof of [Bal+17] in order to present the lemma with explicit guarantees on the running time of the reduction.

Lemma 3.6 (Theorem 4 of [GS92]). *Let \mathbb{F} be a finite field with $|\mathbb{F}| = q$ elements. Let $N, D > 0$. Suppose $9 < D < q/12$. Let $f : \mathbb{F}^N \rightarrow \mathbb{F}$ be a polynomial of degree at most D . If there is an algorithm A running in time $T(A, N)$ such that*

$$\mathbb{P}_{x \sim \text{Unif}[\mathbb{F}^N]}[A(x) = f(x)] > 2/3,$$

then there is an algorithm B running in time $O((N + D)D^2 \log^2 q + T(A, N) \cdot D)$ such that for any $x \in \mathbb{F}^N$, it holds that $\mathbb{P}[B(x) = f(x)] > 2/3$.

For completeness, we provide a proof of this lemma in Appendix B. Lemma 3.6 implies that if we can efficiently compute $P_{n,k,s,\mathbb{F}}$ on at least a $2/3$ fraction of randomly chosen inputs in \mathbb{F}^N , then we can efficiently compute the polynomial $P_{n,k,s,\mathbb{F}}$ over a worst-case input in \mathbb{F}^N .

3.4 Reduction to Evaluating the Polynomial on $G(nk, c, s, k)$

So far, we have reduced worst-case clique-counting over unweighted hypergraphs to the average-case problem of computing $P_{n,k,s,\mathbb{F}}$ over k -partite hypergraphs with random edge weights in \mathbb{F} . It remains to reduce from computing $P_{n,k,s,\mathbb{F}}$ on inputs $x \sim \text{Unif}[\mathbb{F}^N]$ to random hypergraphs, which correspond to $x \sim \text{Unif}[\{0,1\}^N]$. Since $\{0,1\}^N$ is an exponentially small subset of \mathbb{F}^N if $|\mathbb{F}| > 2$, the random weighted and unweighted hypergraph problems are very different. In this section, we carry out this reduction using two different arguments for PARITY- (k,s) -CLIQUE and $\#(k,s)$ -CLIQUE. The latter reduction is based on the total variation convergence of random binary expansion modulo p to $\text{Unif}[\mathbb{F}_p]$ and related algorithmic corollaries from Section 4.

We first present the reduction that will be applied in the case of PARITY- (k,s) -CLIQUE. Recall $D = \binom{k}{s}$ is the degree of $P_{n,k,s}$. The following lemma will be used only for the PARITY- (k,s) -CLIQUE case:

Lemma 3.7. *Let p be prime and $t \geq 1$. Suppose A is an algorithm that computes $P_{n,k,s,\mathbb{F}_p}(y)$ with error probability less than $\delta \triangleq \delta(n)$ for $y \sim \text{Unif}[\mathbb{F}_p^N]$ in time $T(A, n)$. Then there is an algorithm B that computes $P_{n,k,s,\mathbb{F}_{p^t}}(x)$ with error probability less than $t^D \cdot \delta$ for $x \sim \text{Unif}[\mathbb{F}_{p^t}^N]$ in time $T(B, n) = O(Nt^4(\log p)^3 + t^D \cdot T(A, n))$.*

Proof. We give a reduction computing $P_{n,k,s,\mathbb{F}_{p^t}}(x)$ where $x \sim \text{Unif}[\mathbb{F}_{p^t}^N]$ given blackbox access to A . Let β be such that $\beta, \beta^p, \beta^{p^2}, \dots, \beta^{p^{t-1}} \in \mathbb{F}_{p^t}$ forms a normal basis for \mathbb{F}_{p^t} over \mathbb{F}_p . Now for each $i \in [N]$, compute the basis expansion

$$x_i = x_i^{(0)}\beta + x_i^{(1)}\beta^p + \dots + x_i^{(t-1)}\beta^{p^{t-1}}.$$

One can find a generator for a normal basis $\beta \in \mathbb{F}_{p^t}$ in time $O((t^2 + \log p)(t \log p)^2)$ by Bach et al. [BDS93]. Computing $x^{(0)}, \dots, x^{(t-1)}$ then takes time $O(Nt^3(\log p)^3)$ because N applications of Gaussian elimination each take at most $O(t^3)$ operations over \mathbb{F}_p .⁶ Note that since x is uniformly distributed and $\beta, \beta^p, \dots, \beta^{p^{t-1}}$ form a basis, it follows that $x^{(0)}, x^{(1)}, \dots, x^{(t-1)}$ are distributed i.i.d according to $\text{Unif}[\mathbb{F}_p^N]$.

Given a coloring of the hyperedges $b : [N] \rightarrow \{0, 1, \dots, t-1\}$, define $x^{(b)} \in \mathbb{F}_p^N$ as $x_i^{(b)} = x_i^{(b(i))}$ for all $i \in [N]$. Observe that for any fixed coloring b , the vector $x^{(b)}$ is uniform in \mathbb{F}_p^N .

In our proof, for every map $a : \binom{[k]}{s} \rightarrow \{0, 1, \dots, t-1\}$, we construct a coloring $a \circ L : [N] \rightarrow \{0, \dots, t-1\}$ of the hyperedges $[N]$ using the k -partiteness of the hypergraph. Given a hyperedge $W = \{w_1, \dots, w_s\} \in \mathcal{E} = [N]$, we have that $L(W) \in \binom{[k]}{s}$ by the k -partiteness of the hypergraph, and hence the color $(a \circ L)(W) \triangleq a(L(W))$ is well-defined. As above, for any fixed a , the vector $x^{(a \circ L)}$ is uniform in \mathbb{F}_p^N .

We now manipulate $P_{n,k,s,\mathbb{F}_{p^t}}$. First we write each entry x_{u_S} in the normal basis, and then we redistribute terms to write $P_{n,k,s,\mathbb{F}_{p^t}}$ as a weighted sum of clique-counts modulo p :

$$\begin{aligned} P_{n,k,s,\mathbb{F}_{p^t}}(x) &= \sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall j L(u_j)=j}} \prod_{S \in \binom{[k]}{s}} x_{u_S} \\ &= \sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall j L(u_j)=j}} \prod_{S \in \binom{[k]}{s}} \left(\sum_{i=0}^{t-1} x_{u_S}^{(i)} \beta^{p^i} \right) \\ &= \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, t-1\}} \left(\sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall i L(u_i)=i}} \prod_{S \in \binom{[k]}{s}} \left(x_{u_S}^{(a(S))} \beta^{p^{a(S)}} \right) \right) \\ &= \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, t-1\}} \left(\prod_{S \in \binom{[k]}{s}} \beta^{p^{a(S)}} \right) \left(\sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall i L(u_i)=i}} \prod_{S \in \binom{[k]}{s}} x_{u_S}^{(a(S))} \right) \\ &= \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, t-1\}} \left(\prod_{S \in \binom{[k]}{s}} \beta^{p^{a(S)}} \right) P_{n,k,s,\mathbb{F}_p}(x^{(a \circ L)}) \end{aligned}$$

Since $x^{(a \circ L)} \sim \text{Unif}[\mathbb{F}_p^N]$ for each fixed map a , computing $P_{n,k,s,\mathbb{F}_{p^t}}(x)$ reduces to evaluating P_{n,k,s,\mathbb{F}_p} on t^D uniformly random inputs in \mathbb{F}_p^N and outputting a weighted sum of the evaluations. The error probability is bounded by a union bound. \square

⁶For a good survey on normal bases, we recommend [Gao93].

We now give the reduction to evaluating $P_{n,k,s}$ on random hypergraphs drawn from $G(nk, c, s, k)$ in the case of $\#(k, s)$ -CLIQUE. One of the main lemmas driving the reduction is the following.

Lemma 3.8. *There is an absolute constant K such that the following holds. For any $\epsilon > 0$, $c \in (0, 1)$, prime $p > 2$, and $t \geq K \cdot c^{-1}(1 - c)^{-1} \log(p/\epsilon^2) \log p$, there is an $O(pt \log(p/\epsilon))$ -time algorithm that, given $x \in [p]$, samples random variables $X_0, \dots, X_{t-1} \in \{0, 1\}$ satisfying $\sum_{i=0}^{t-1} 2^i \cdot X_i \equiv x \pmod{p}$ almost surely. Moreover, if x is chosen uniformly at random from $[p]$ then $d_{TV}(\mathcal{L}(X_1, \dots, X_t), \text{Ber}(c)^{\otimes t}) \leq \epsilon$.*

The proof of Lemma 3.8 is deferred to Section 4. It is a central ingredient in the $\#(k, s)$ -CLIQUE reduction and will be used through the following lemma.

Lemma 3.9. *Let p be prime and let $c = c(n), \gamma = \gamma(n) \in (0, 1)$. Suppose that A is an algorithm that computes $P_{n,k,s,\mathbb{F}_p}(y)$ with error probability less than $\delta \triangleq \delta(n)$ when $y \in \{0, 1\}^N$ is drawn from $G(nk, c, s, k)$. Then, for some $t = O(c^{-1}(1 - c)^{-1} \log(Np/\gamma) \log p)$, there is an algorithm B that evaluates $P_{n,k,s,\mathbb{F}_p}(x)$ with error probability at most $\gamma + t^D \cdot \delta$ when $x \sim \text{Unif}[\mathbb{F}_p^N]$ in time upper bounded by $T(B, n) = O(Npt \log(Np/\gamma) + t^D \cdot T(A, n))$.*

Proof. We give a reduction computing $P_{n,k,s,\mathbb{F}_p}(x)$ where $x \sim \text{Unif}[\mathbb{F}_p^N]$ given blackbox access to A . We first handle the case in which $p > 2$. For each $j \in [N]$, apply the algorithm from Lemma 3.8 to sample $x_j^{(0)}, x_j^{(1)}, \dots, x_j^{(t-1)} \in \{0, 1\}$ satisfying

$$d_{TV} \left(\mathcal{L}(x_j^{(0)}, \dots, x_j^{(t-1)}), \text{Ber}(c)^{\otimes t} \right) \leq \epsilon \triangleq \gamma/N \quad \text{and} \quad \sum_{i=0}^{t-1} 2^i x_j^{(i)} \equiv x_j \pmod{p}$$

By Lemma 3.8, we may choose $t = O(c^{-1}(1 - c)^{-1} \log(Np/\gamma) \log p)$ and this sampling can be carried out in $O(Npt \log(Np/\gamma))$ time. By the total variation bound, for each j we may couple $(x_j^{(0)}, \dots, x_j^{(t-1)})$ with $(Z_j^{(0)}, \dots, Z_j^{(t-1)}) \sim \text{Ber}(c)^{\otimes t}$, so that $\mathbb{P}[x_j^{(i)} = Z_j^{(i)} \forall i, j] \geq 1 - \gamma$. Moreover, $x_j^{(i)}$ is independent of $x_l^{(k)}$ whenever $j \neq l$, so we may choose the Z so that $Z_j^{(i)}$ is independent of $Z_l^{(k)}$ whenever $j \neq l$.

As in the proof of Lemma 3.7, given any coloring $b : [N] \rightarrow \{0, \dots, t-1\}$, we define $Z^{(b)} \in \{0, 1\}^N$ by $Z_j^{(b)} = Z_j^{(b(j))}$, for all $j \in [N]$. We also note that for any fixed b , the entries $Z_1^{(b)}, \dots, Z_N^{(b)}$ are independent and distributed as $\text{Ber}(c)$. Therefore,

$$Z^{(b)} \sim G(nk, c, s, k)$$

Now compute the following quantity, writing each entry of Z as a binary expansion and redistributing terms similarly to the calculations in Lemma 3.7. We are working in \mathbb{F}_p so the equalities hold modulo p :

$$\begin{aligned} \tilde{P}_{n,k,s,\mathbb{F}_p}(Z) &\triangleq \sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall j \ L(u_j)=j}} \prod_{S \in \binom{[k]}{s}} \left(\sum_{i=0}^{t-1} 2^i \cdot Z_{u_S}^{(i)} \right) \\ &= \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, t-1\}} \left(\sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall i \ L(u_i)=i}} \prod_{S \in \binom{[k]}{s}} \left(2^{a(S)} \cdot Z_{u_S}^{(a(S))} \right) \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, t-1\}} \left(\prod_{S \in \binom{[k]}{s}} 2^{a(S)} \right) \left(\sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall i L(u_i) = i}} \prod_{S \in \binom{[k]}{s}} Z_{u_S}^{(a(S))} \right) \\
&= \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, t-1\}} \left(\prod_{S \in \binom{[k]}{s}} 2^{a(S)} \right) P_{n, k, s, \mathbb{F}_p}(Z^{(a \circ L)}).
\end{aligned}$$

We may use algorithm A to evaluate the t^D values of $P_{n, k, s, \mathbb{F}_p}(Z^{(a \circ L)})$, with probability $< t^D \cdot \delta$ of any error (by a union bound). Computing $\tilde{P}_{n, k, s, \mathbb{F}_p}(Z)$ reduces to computing a weighted sum over the t^D evaluations. Conditioned on the event that $x_j^{(i)} = Z_j^{(i)} \forall i, j$, then $P_{n, k, s, \mathbb{F}_p}(x) = \tilde{P}_{n, k, s, \mathbb{F}_p}(Z)$, because

$$\begin{aligned}
P_{n, k, s, \mathbb{F}_p}(x) &= \sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall j L(u_j) = j}} \prod_{S \in \binom{[k]}{s}} x_{u_S} \\
&= \sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall j L(u_j) = j}} \prod_{S \in \binom{[k]}{s}} \left(\sum_{i=0}^{t-1} 2^i \cdot x_{u_S}^{(i)} \right) \\
&= \sum_{\substack{\{u_1, \dots, u_k\} \subset V(G) \\ \forall j L(u_j) = j}} \prod_{S \in \binom{[k]}{s}} \left(\sum_{i=0}^{t-1} 2^i \cdot Z_{u_S}^{(i)} \right) = \tilde{P}_{n, k, s, \mathbb{F}_p}(Z).
\end{aligned}$$

Since $\mathbb{P}[x_j^{(i)} = Z_j^{(i)} \forall i, j] \geq 1 - \gamma$, by a union bound with the error in calculation we have computed $P_{n, k, s, \mathbb{F}_p}(x)$ with probability of error $\leq \gamma + t^D \cdot \delta$. The claim follows for the case $p > 2$.

If $p = 2$, then the proof is almost identical, except that since $2 \equiv 0 \pmod{2}$, we may no longer use the result on random binary expansions of Lemma 3.8. In this case, for each $j \in [N]$ we sample $x_j^{(0)}, \dots, x_j^{(t-1)} \in \{0, 1\}^N$ such that each $d_{\text{TV}}(\mathcal{L}(x_j^{(0)}, \dots, x_j^{(t-1)}), \text{Ber}(c)^{\otimes t}) \leq \epsilon \triangleq \gamma/N$, and so that

$$\sum_{i=0}^{t-1} x_j^{(i)} = x_j \pmod{p}.$$

By Lemma 4.5 (deferred but analogous to Lemma 3.8), we may choose $t = O(c^{-1}(1-c)^{-1} \log(N/\gamma))$, and we may sample in time $O(Nt \log(N/\gamma))$. Again, we couple the $x_j^{(i)}$ variables with variables $Z_j^{(i)} \sim \text{Ber}(c)$ such that the event E that $x_j^{(i)} = Z_j^{(i)}$ for all i, j has probability $\mathbb{P}[E] \geq 1 - \gamma$ and such that $Z_j^{(i)}$ is independent of $Z_l^{(k)}$ whenever $j \neq l$. By a similar, and simpler, calculation to the one for the case $p > 2$, we have that $\tilde{P}_{n, k, s, \mathbb{F}_2}(Z) = P_{n, k, s, \mathbb{F}_2}(x)$ conditioned on E , where

$$\tilde{P}_{n, k, s, \mathbb{F}_2}(Z) \triangleq \sum_{a: \binom{[k]}{s} \rightarrow \{0, \dots, t-1\}} P_{n, k, s, \mathbb{F}_2}(Z^{(a \circ L)}).$$

This can be calculated using the algorithm A similarly to the $p > 2$ case, because each $Z^{(a \circ L)}$ is distributed as $G(nk, c, s, k)$. \square

3.5 Reduction to Counting k -Cliques in $G(n, c, s)$

So far, we have reduced PARITY- (k, s) -CLIQUE and $\#(k, s)$ -CLIQUE for worst-case input hypergraphs to average-case inputs drawn from the k -partite Erdős-Rényi distribution $G(nk, c, s, k)$. We now carry out the final step of the reduction, showing that PARITY- (k, s) -CLIQUE and $\#(k, s)$ -CLIQUE on inputs drawn from $G(nk, c, s, k)$ reduce to inputs drawn from the non- k -partite Erdős-Rényi distribution $G(n, c, s)$. Recall that a hypergraph G drawn from $G(nk, c, s, k)$ has vertex set $V(G) = [n] \times [k]$ and vertex partition given by the labels $L : (i, j) \in [n] \times [k] \mapsto j \in [k]$.

Lemma 3.10. *Let $\delta = \delta(n) \in (0, 1)$ be a non-increasing function of n and let $c = c(n) \in (0, 1)$. Suppose that A is a randomized algorithm for $\#(k, s)$ -CLIQUE such that for any n , A has error probability less than $\delta(n)$ on hypergraphs drawn from $G(n, c, s)$ in $T(A, n)$ time. Then there exists an algorithm B solving $\#(k, s)$ -CLIQUE that has error probability less than $2^k \cdot \delta(n)$ on hypergraphs drawn from $G(nk, c, s, k)$ and that runs in $T(B, n) = O(2^k \cdot T(A, nk) + k^s n^s + s^2 k^3 2^k \log^2(nk))$ time.*

Proof. It suffices to count the number of k -cliques in $G \sim G(nk, c, s, k)$ given blackbox access to A . Construct the hypergraph H over the same vertex set $V(H) = [n] \times [k]$ by starting with G and adding every edge $e = \{v_1, v_2, \dots, v_s\} \in \binom{[n] \times [k]}{s}$ satisfying the condition $|\{L(v_1), \dots, L(v_s)\}| < s$ independently with probability c . In other words, independently add each edge to G containing two vertices from the same part of G . It follows that H is distributed according to $G(nk, c, s)$. More generally, for every $S \subset [k]$, H_S is distributed according to $G(|L^{-1}(S)|, c, s)$ where H_S is the restriction of H to the vertices $L^{-1}(S) \subset V(H)$ with labels in S . Note that H can be constructed in $O(k^s n^s)$ time.

Now observe that for each $S \neq \emptyset$, it holds that $n \leq |L^{-1}(S)| \leq nk$ and the algorithm A succeeds on each H_S with probability at least $1 - \delta(n)$. By a union bound, we may compute the number of k -cliques $|\text{cl}_k(H_S)|$ in H_S for all $S \subset [k]$ with error probability less than $2^k \cdot \delta(n)$. Note that this can be done in $O(2^k \cdot T(A, nk))$ time. From these counts $|\text{cl}_k(H_S)|$, we now to inductively compute

$$t_d \triangleq |\{S \in \text{cl}_k(H) : |L(S)| = d\}|$$

for each $d \in [k]$. Note that $t_0 = 0$ in the base case $d = 0$. Given t_0, t_1, \dots, t_d , the next count t_{d+1} can be expressed by inclusion-exclusion as

$$\begin{aligned} t_{d+1} &= \sum_{T \subset [k], |T|=d+1} |\{S \in \text{cl}_k(H) : L(S) = T\}| \\ &= \sum_{T \subset [k], |T|=d+1} \left(|\text{cl}_k(H_T)| - \sum_{i=0}^d \sum_{U \subset T, |U|=i} |\{S \in \text{cl}_k(H) : L(S) = U\}| \right) \\ &= \left(\sum_{T \subset [k], |T|=d+1} |\text{cl}_k(H_T)| \right) - \sum_{i=0}^d \binom{k-i}{d+1-i} |\{S \in \text{cl}_k(H) : |L(S)| = i\}| \\ &= \sum_{T \subset [k], |T|=d+1} |\text{cl}_k(H_T)| - \sum_{i=0}^d \binom{k-i}{d+1-i} t_i \end{aligned}$$

After $O(k2^k)$ operations, this recursion yields the number of k -cliques $t_k = |\{S \in \text{cl}_k(H) : |L(S)| = k\}| = |\text{cl}_k(G)|$ in the original k -partite hypergraph G . The sizes of the integers manipulated are always at most $\max(2^k \binom{nk}{s})$, so each arithmetic operation takes $O((ks \log(nk))^2)$ time. \square

Repeating the same proof over \mathbb{F}_2 yields an analogue of Lemma 3.10 for PARITY- (k, s) -CLIQUE, as stated below.

Lemma 3.11. *Lemma 3.10 holds when $\#(k, s)$ -CLIQUE is replaced by PARITY- (k, s) -CLIQUE.*

3.6 Proofs of Theorems 2.8 and 2.9

We now combine Steps 1-5 formally in order to prove Theorems 2.8 and 2.9.

Proof of Theorem 2.8. Our goal is to construct an algorithm B solving $\#(k, s)$ -CLIQUE with error probability $< 1/3$ on any s -uniform hypergraph x . We are given an algorithm A that solves $\#(k, s)$ -CLIQUE with probability of error $< 1/\Upsilon_{\#}$ on hypergraphs drawn from $G(n, c, s)$. We will construct the following intermediate algorithms in our reduction:

- Algorithm A_0 that solves $\#(k, s)$ -CLIQUE with error probability $< 1/3$ for any worst-case k -partite hypergraph.
- Algorithm $A_1(x, p)$ that computes $P_{n,k,s,\mathbb{F}_p}(x)$ for any $x \in \mathbb{F}_p^N$ and for any prime p such that $12\binom{k}{s} < p < 10 \log n^k$, with worst-case error probability $< 1/3$.
- Algorithm $A_2(y, p)$ for primes $12\binom{k}{s} < p < 10 \log n^k$ computing $P_{n,k,s,\mathbb{F}_p}(y)$ on inputs $y \sim \text{Unif}[\mathbb{F}_p^N]$ with error probability $< 1/3$.
- Algorithm $A_3(z)$ that computes $P_{n,k,s}(z)$ on inputs $z \sim G(nk, c, s, k)$ with error probability $< \delta$. (The required value of δ will be determined later on.)

We construct algorithm B from A_0 , A_0 from A_1 , A_2 from A_3 , and A_3 from A .

1. *Reduce to computing $\#(k, s)$ -CLIQUE for k -partite hypergraphs.* We use Lemma 3.3 to construct B from A_0 , such that B runs in time

$$T(B, n) = T(A_0, n) + O((nk)^s).$$

2. *Reduce to computing P_{n,k,s,\mathbb{F}_p} on worst-case inputs.* We use Proposition 3.4 to construct A_0 from A_1 such that A_0 runs in time

$$T(A_0, n) \leq O(T(A_1, n) \cdot \log n^k + (\log n^k)^2).$$

The algorithm A_0 starts by using a sieve to find the first T primes $12\binom{k}{s} < p_1 < \dots < p_T$ such that $\prod_{i=1}^T p_i > n^k$. Notice that $p_T \leq 10 \log n^k$, so this step takes time $O((\log n^k)^2)$. Then, given a k -partite hypergraph $x \in \{0, 1\}^N$, the algorithm A_0 computes $P_{n,k,s}(x)$ by computing $P_{n,k,s,\mathbb{F}_{p_i}}(x)$ for all p_i , boosting the error of A_1 by repetition and majority vote. Since $T = O((\log n^k)/(\log \log n^k))$, we only need to repeat $O(\log \log n^k)$ times per prime; this yields a total slowdown factor of $O(\log n^k)$. Once we have computed $P_{n,k,s}(x)$, we recall that it is equal to the number of k -cliques in x .

3. *Reduce to computing P_{n,k,s,\mathbb{F}_p} on random inputs in \mathbb{F}_p^N .* We use Lemma 3.6 to construct A_1 from A_2 such that A_1 runs in time

$$\begin{aligned} T(A_1, n) &= O((N + D)D^2 \log^2 p + D \cdot T(A_2, n)) \\ &= O(n^s \binom{k}{s}^3 \log^2 \log n^k + \binom{k}{s} \cdot T(A_2, n)). \end{aligned}$$

4. *Reduce to computing $P_{n,k,s}$ on random inputs in $\{0, 1\}^N$.* We use Lemma 3.9 to construct A_2 from A_3 such that A_2 runs in time

$$T(A_2, n) = O(Npt \log(Np) + t^{\binom{k}{s}} \cdot T(A_3, n)),$$

for some $t = O(c^{-1}(1-c)^{-1}s(\log n)(\log p))$. For this step, we require the error probability δ of algorithm $A_3(z)$ on inputs $z \sim G(nk, c, s, k)$ to be at most $1/(4t^D) = 1/(4t^{\binom{k}{s}})$.

5. *Reduce to computing $\#(k, s)$ -CLIQUE for $G(n, c, s)$ hypergraphs.* We use Lemma 3.10 to construct A_3 from A such that A_3 runs in time

$$T(A_3, n) = O((nk)^s + s^2 k^3 2^k \log(nk) + 2^k \cdot T(A, nk)),$$

and such that A_3 has error probability at most $\delta < 2^k / \Upsilon_{\#}$.

As in the theorem statement, let $\Upsilon_{\#}(n, c, s, k) \triangleq (C(c^{-1}(1-c)^{-1})s(\log n)(\log k + \log \log n))^{\binom{k}{s}}$, where $C > 0$ is a large constant to be determined. If we take C large enough, then $4t^{\binom{k}{s}} \cdot 2^k \leq \Upsilon_{\#}$. In this case, the error δ of A_3 will be at most $1/(4t^{\binom{k}{s}})$, which is what we needed for the fourth step. Putting the runtime bounds together,

$$\begin{aligned} T(B, n) &= O\left((nk)^s + (\log n^k)^2 \right. \\ &\quad \left. + (\log n^k) \cdot \left(n^s t k \binom{k}{s}^3 (\log n)^2 + \binom{k}{s} \cdot (4t)^{\binom{k}{s}} \cdot (T(A, nk) + (nk)^s)\right)\right) \\ &= O\left(n^s k^3 \binom{k}{s}^3 (c^{-1}(1-c)^{-1})(\log k + \log \log n) \log^4 n \right. \\ &\quad \left. + (\log n) \cdot \Upsilon_{\#} \cdot (T(A, nk) + (nk)^s)\right), \end{aligned}$$

if we choose $C > 0$ large enough. Hence, the second term dominates and

$$T(B, n) = O((\log n) \cdot \Upsilon_{\#} \cdot (T(A, nk) + (nk)^s)),$$

as $\binom{k}{s} \geq 3$ without loss of generality. \square

Proof of Theorem 2.9. The proof of item 1 of Theorem 2.9 is analogous to the proof of Theorem 2.8, except that it does not use the Chinese remainder theorem. Moreover, special care is needed in order to ensure that the field \mathbb{F} over which we compute the polynomial $P_{n,k,s,\mathbb{F}}$ in the intermediate steps is large enough that we may use the random self-reducibility of polynomials.

Our goal is to construct an algorithm B that solves PARITY- (k, s) -CLIQUE with error probability $< 1/3$ on any s -uniform hypergraph x . We are given an algorithm A that solves PARITY- (k, s) -CLIQUE with probability of error $< 1/\Upsilon_{P,1}$ on hypergraphs drawn from $G(n, c, s)$. We will construct the following intermediate algorithms in our reduction:

- Algorithm A_0 that solves PARITY- (k, s) -CLIQUE with error probability $< 1/3$ for any worst-case k -partite hypergraph.
- Algorithm $A_1(w)$ that computes $P_{n,k,s,\mathbb{F}_{2^\kappa}}(w)$ on inputs $w \sim \text{Unif}[\mathbb{F}_{2^\kappa}^N]$ for $\kappa = \lceil \log_2(12 \binom{k}{s}) \rceil$, with error probability $< 1/3$.
- Algorithm $A_2(y)$ that computes $P_{n,k,s,\mathbb{F}_2}(y)$ on inputs $y \sim \text{Unif}[\mathbb{F}_2^N]$ with error probability $< \delta_2$. (The required value of δ_2 will be determined later on.)

- Algorithm $A_3(z)$ that computes $P_{n,k,s,\mathbb{F}_2}(z)$ on inputs $z \sim G(nk, c, s, k)$ with error probability $< \delta_3$. (The required value of δ_3 will be determined later on.)

We construct algorithm B from A_0 , A_0 from A_1 , A_2 from A_3 , and A_3 from A .

1. *Reduce to computing PARITY- (k, s) -CLIQUE for k -partite hypergraphs.* We use Lemma 3.3 to construct B from A_0 , such that B runs in time

$$T(B, n) = T(A_0, n) + O((nk)^s).$$

2. *Reduce to computing $P_{n,k,s,\mathbb{F}_{2^\kappa}}$ on random inputs in $\mathbb{F}_{2^\kappa}^N$.* Note that by Proposition 3.5 if we can compute $P_{n,k,s,\mathbb{F}_{2^\kappa}}$ for worst-case inputs, then we can solve PARITY- (k, s) -CLIQUE. We use Lemma 3.6 to construct A_0 from A_1 such that A_0 runs in time

$$T(A_0, n) = O(\kappa^2(N + D)D^2 + D \cdot T(A_1, n)) = O(n^s \binom{k}{s}^2 \kappa^2 + \binom{k}{s} \cdot T(A_1, n))$$

3. *Reduce to computing P_{n,k,s,\mathbb{F}_2} on random inputs in \mathbb{F}_2^N .* We use Lemma 3.7 to construct A_1 from A_2 such that A_1 runs in time

$$T(A_1, n) \leq O(N\kappa^4 + \kappa \binom{k}{s} \cdot T(A_2, n)),$$

and has error probability at most $\delta_2 \cdot \kappa \binom{k}{s}$ on random inputs $w \sim \text{Unif}[\mathbb{F}_{2^\kappa}^N]$. Thus, A_2 must have error probability at most $\delta_2 < 1/(3\kappa \binom{k}{s})$ on random inputs in $y \sim \text{Unif}[\mathbb{F}_2^N]$ for this step of the reduction to work.

4. *Reduce to computing P_{n,k,s,\mathbb{F}_2} on random inputs in $\{0, 1\}^N$.* We use Lemma 3.9 to construct A_2 from A_3 such that A_2 runs in time

$$T(A_2, n) = O(Nt \log(N/\gamma) + t \binom{k}{s} \cdot T(A_3, n)),$$

for some $t = O(c^{-1}(1 - c)^{-1}(s \log(n) + \log(1/\gamma)))$. The error probability of A_2 on random inputs $z \sim G(nk, c, s, k)$ will be at most $\delta_2 < \delta_3 \cdot t \binom{k}{s} + \gamma$. Since we require error probability at most $\delta_2 \leq 1/(3\kappa \binom{k}{s})$ of algorithm $A_2(z)$ on inputs $z \sim G(nk, c, s, k)$, we set $\gamma = 1/(10\kappa \binom{k}{s})$ and require $\delta_3 \leq 1/(10(t\kappa) \binom{k}{s})$, which is sufficient. For this choice of γ , we have $t = O(c^{-1}(1 - c)^{-1}(s \log(n) + \binom{k}{s} \log \log \binom{k}{s}))$.

5. *Reduce to computing PARITY- (k, s) -CLIQUE for $G(n, c, s)$ hypergraphs.* We use Lemma 3.11 to construct A_3 from A such that A_3 runs in time

$$T(A_3, n) = O((nk)^s + s^2 k^3 2^k \log(nk) + 2^k \cdot T(A, nk)),$$

and such that A_3 has error probability at most $\delta_3 < 2^k/\Upsilon_{P,1}$.

As in the theorem statement, let

$$\Upsilon_{P,1}(n, c, s, k) \triangleq \left(C(c^{-1}(1 - c)^{-1})s(\log k) \left(s \log n + \binom{k}{s} \left(\log \log \binom{k}{s} \right) \right) \right) \binom{k}{s}$$

for some large enough constant C .

If we take C large enough, then $(\kappa t)^{\binom{k}{s}} \leq \frac{1}{10} \cdot 2^{-k} \cdot \Upsilon_{P,1}$, as desired. In this case, the error of A_0 on uniformly random inputs will be at most $1/3$, which is what we needed. Putting the runtime bounds together,

$$\begin{aligned} T(B, n) &= O\left(n^s \binom{k}{s}^2 \log^2 \kappa + n^s \binom{k}{s}^2 \kappa^{\binom{k}{s}} t \log\left(n^s \kappa^{\binom{k}{s}}\right)\right) \\ &\quad + n^s \binom{k}{s}^2 \kappa^4 + \binom{k}{s} \cdot (4\kappa t)^{\binom{k}{s}} \cdot (T(A, nk) + (nk)^s) \\ &= O\left(n^s \binom{k}{s}^2 (t\kappa^{\binom{k}{s}} \log^2 \kappa + \kappa^4) + \Upsilon_{P,1} \cdot (T(A, nk) + (nk)^s)\right), \end{aligned}$$

if we choose $C > 0$ large enough. Since $\binom{k}{s} \geq 3$ without loss of generality, the second term dominates and

$$T(B, n) = O(\Upsilon_{P,1} \cdot (T(A, nk) + (nk)^s)).$$

For item 2 of the theorem, we restrict the inputs to come from $G(n, 1/2, s)$, and we achieve a better error tolerance because algorithm A_3 is the same as A_2 . This means that we may skip step 4 of the proof of item 1. In particular, we only need $\delta_3 = \delta_2 \leq 1/(3\kappa^{\binom{k}{s}})$. So algorithm A only needs to have error $< 1/\Upsilon_{P,2}$, for $\Upsilon_{P,2}(k, s) \triangleq (Cs \log k)^{\binom{k}{s}}$. It is not hard to see that, skipping step 4, the algorithm B that we construct takes time $T(B, n) = O(\Upsilon_{P,2} \cdot (T(A, nk) + (nk)^s))$. \square

4 Random Binary Expansions Modulo p

In this section, we consider the distributions of random binary expansions of the form

$$Z_t \cdot 2^t + Z_{t-1} \cdot 2^{t-1} + \dots + Z_0 \pmod{p}$$

for some prime p and independent, possibly biased, Bernoulli random variables $Z_i \in \{0, 1\}$. We show that for t polylogarithmic in p , these distributions become close to uniformly distributed over \mathbb{F}_p , more or less regardless of the biases of the Z_i . This is then used to go in the other direction, producing approximately independent Bernoulli variables that are the binary expansion of a number with a given residue. The special case of this argument in which the Bernoulli variables are unbiased has already appeared in an earlier work by Goldreich and Rothblum [GR17]. In that case, the proof of correctness is much simpler, because the Fourier-analytic tools used below can be avoided.

For $p > 2$, the main result of the section is the following slightly more general restatement of Lemma 3.8. It implies that we can efficiently sample biased binary expansions, conditioned on the expansion being equivalent to some x modulo p .

Lemma 4.1 (Restatement of Lemma 3.8). *There is an absolute constant K such that the following holds. Let $p > 2$ be prime, $t \geq K \cdot c^{-1}(1-c)^{-1} \log(p/\epsilon^2) \log p$, $c \leq q_0, q_1, \dots, q_t \leq 1-c$ for some $c \in (0, 1/2]$ and $\epsilon > 0$. Let $Z_i \sim \text{Ber}(q_i)$ be independent. Then there is an $O(pt \log(p/\epsilon))$ -time algorithm that given $x \in [p]$ samples $X_0, \dots, X_t \in \{0, 1\}$ satisfying $\sum_{i=0}^t 2^i \cdot X_i \equiv x \pmod{p}$ almost surely. Moreover, if x is chosen uniformly in $[p]$, then $d_{TV}(\mathcal{L}(X_0, \dots, X_t), \mathcal{L}(Z_0, \dots, Z_t)) \leq \epsilon$.*

Our argument uses finite Fourier analysis on \mathbb{F}_p . Given a function $f : \mathbb{F}_p \rightarrow \mathbb{R}$, define its Fourier transform to be $\hat{f} : \mathbb{F}_p \rightarrow \mathbb{C}$, where $\hat{f}(t) = \sum_{x=0}^{p-1} f(x) \omega^{tx}$ and $\omega = e^{2\pi i/p}$. In this section, we endow \mathbb{F}_p with the total ordering of $\{0, 1, \dots, p-1\}$ as elements of \mathbb{Z} . Given a set S , let $2S = \{2s : s \in S\}$. We begin with a simple claim showing that sufficiently long geometric progressions with ratio 2 in \mathbb{F}_p contain a middle residue modulo p .

Claim 4.2. *Suppose that $a_1, \dots, a_k \in \mathbb{F}_p$ is a sequence with $a_1 \neq 0$ and $a_{i+1} = 2a_i$ for each $1 \leq i \leq k-1$. Then if $k \geq 1 + \log_2(p/3)$, there is some j with $\frac{p}{3} \leq a_j \leq \frac{2p}{3}$.*

Proof. Let $S = \{x \in \mathbb{F}_p : x < p/3\}$ and $T = \{x \in \mathbb{F}_p : x > 2p/3\}$. Observe that $2S \cap T = \emptyset$ and $S \cap 2T = \emptyset$, which implies that there is no i such that a_i and a_{i+1} are both in S and T . Therefore if (a_1, a_2, \dots, a_k) contains elements of both S and T , there must be some j with $a_j \in (S \cup T)^C$ and the claim follows. It thus suffices to show that (a_1, a_2, \dots, a_k) cannot be entirely contained in one of S or T . First consider the case that it is contained in S . Define the sequence $(a'_1, a'_2, \dots, a'_k)$ of integers by $a'_{i+1} = 2a'_i$ for each $1 \leq i \leq k-1$ and $a'_1 \in [1, p/3)$ is such that $a'_1 \equiv a_1 \pmod{p}$. It follows that $a'_i \equiv a_i \pmod{p}$ for each i and $a'_k \geq 2^{k-1} \geq p/3$. Now consider the smallest j with $a'_j > p/3$. Then $p/3 \geq a'_{j-1} = a'_j/2$ by the minimality of j , and $p/3 \leq a_j \leq 2p/3$ which is a contradiction. If the sequence is contained in T , then $(-a_1, -a_2, \dots, -a_k)$ is contained in S and applying the same argument to this sequence proves the claim. \square

We now bound the total variation between the distribution of random binary expansions modulo p and the uniform distribution. In Appendix C, we show Lemma 4.3 is tight assuming there are infinitely-many Mersenne primes.

Lemma 4.3. *There is an absolute constant $K > 0$ such that the following holds. Let $p > 2$ be a prime, let $c, \epsilon \in (0, 1/2]$ be arbitrary and suppose that $c \leq q_0, q_1, \dots, q_t \leq 1 - c$. Then if $t \geq K \cdot c^{-1}(1-c)^{-1} \log(p/\epsilon^2) \log p$ and $Z_i \sim \text{Ber}(q_i)$ are independent, the distribution of $S = \sum_{i=0}^t Z_i \cdot 2^i \pmod{p}$ is within ϵ total variation distance of the uniform distribution on \mathbb{F}_p .*

Proof. Let $f : \mathbb{F}_p \rightarrow \mathbb{R}$ be the probability mass function of $\sum_{i=0}^t 2^i Z_i \pmod{p}$. By definition, we have that

$$f(x) = \sum_{z \in \{0,1\}^{t+1}} \left(\prod_{i=0}^t q_i^{z_i} (1-q_i)^{1-z_i} \right) \mathbf{1} \left\{ \sum_{i=0}^t z_i \cdot 2^i \equiv x \pmod{p} \right\}$$

Now observe that $\hat{f}(s)$ is given by

$$\hat{f}(s) = \sum_{x=0}^{p-1} f(x) \omega^{sx} = \prod_{i=0}^t \left(1 - q_i + q_i \cdot \omega^{2^i \cdot s} \right)$$

The last equality follows directly from expanding the product and the definition of $f(x)$. Note that the constant function $\mathbf{1}$ has Fourier transform $p \cdot \mathbf{1}_{\{s=0\}}$. By Cauchy-Schwarz and Parseval's theorem, we have that

$$\begin{aligned} 4 \cdot d_{\text{TV}}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p])^2 &= \|f - p^{-1} \cdot \mathbf{1}\|_1^2 \leq p \cdot \|f - p^{-1} \cdot \mathbf{1}\|_2^2 = \|\hat{f} - \mathbf{1}_{\{s=0\}}\|_2^2 \\ &= \sum_{s \neq 0} \prod_{i=0}^t \left| 1 - q_i + q_i \cdot \omega^{2^i \cdot s} \right|^2 \end{aligned}$$

where $\mathcal{L}(S)$ denotes the law of S . Note that $|1 - q + q \cdot \omega^a| \leq 1$ by the triangle inequality for all $a \in \mathbb{F}_p$ and $q \in (0, 1)$. Furthermore, if $a \in \mathbb{F}_p$ is such that $p/3 \leq a \leq 2p/3$ and $q \in [c, 1 - c]$, then we have that

$$\begin{aligned} |1 - q + q \cdot \omega^a|^2 &= (1 - q)^2 + q^2 + 2q(1 - q) \cos(2\pi a/p) \\ &= 1 - 2q(1 - q) (1 - \cos(2\pi a/p)) \\ &\leq 1 - 2c(1 - c) (1 - \cos(4\pi/3)) \\ &= 1 - 3c(1 - c) \end{aligned}$$

since $\cos(x)$ is maximized at the endpoints on the interval $x \in [2\pi/3, 4\pi/3]$ and $q(1-q)$ is minimized at the endpoints on the interval $[c, 1-c]$. Now suppose that t is such that

$$t \geq \left\lceil \frac{\log(4\epsilon^2/p)}{\log(1-3c(1-c))} \right\rceil \cdot \lceil 1 + \log_2(p/3) \rceil = \Theta(c^{-1}(1-c)^{-1} \log(p/\epsilon^2) \log p)$$

Fix some $s \in \mathbb{F}_p$ with $s \neq 0$. By Claim 4.2, any $\lceil 1 + \log_2(p/3) \rceil$ consecutive terms of the sequence $s, 2s, \dots, 2^t s \in \mathbb{F}_p$ contain an element between $p/3$ and $2p/3$. Therefore this sequence contains at least $m = \left\lceil \frac{\log(4\epsilon^2/p)}{\log(1-3c(1-c))} \right\rceil$ such terms, which implies that

$$\prod_{i=0}^t \left| 1 - q_i + q_i \cdot \omega^{2^i \cdot s} \right|^2 \leq (1 - 3c(1-c))^m \leq \frac{4\epsilon^2}{p}$$

by the inequality above and the fact that each term in this product is at most 1. Since this holds for each $s \neq 0$, it now follows that

$$4 \cdot d_{\text{TV}}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p])^2 \leq \sum_{s \neq 0} \prod_{i=0}^t \left| 1 - q_i + q_i \cdot \omega^{2^i \cdot s} \right|^2 < 4\epsilon^2$$

and thus $d_{\text{TV}}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p]) < \epsilon$, proving the lemma. \square

Using the lemma, we can sample (Z_0, \dots, Z_t) conditioned on $S \equiv x \pmod{p}$:

Lemma 4.4. *Let $p > 2$ be prime. Suppose that $c \leq q_1, q_2, \dots, q_t \leq 1 - c$ for some $c \in (0, 1/2]$ and that $Z_i \sim \text{Ber}(q_i)$ are independent. Let $Y = \sum_{i=0}^t Z_i \cdot 2^i$ and for each $x \in \mathbb{F}_p$, let $Y_x \sim \mathcal{L}(Y|Y \equiv x \pmod{p})$. Consider Y_R , where R is chosen uniformly at random with $R \sim \text{Unif}[\mathbb{F}_p]$. If $S = Y \pmod{p}$ and $\Delta = d_{\text{TV}}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p]) < p^{-1}$, then it holds that*

1. *Given $x \in \mathbb{F}_p$, we may sample $\mathcal{L}(Y_x)$ within δ total variation distance in $O\left(\frac{t \log(1/\delta)}{p^{-1} - \Delta}\right)$ time.*
2. $d_{\text{TV}}(\mathcal{L}(Y), \mathcal{L}(Y_R)) \leq \Delta$.

Proof. Note that the $x \rightarrow Y_x$ defines a Markov transition sending $S \rightarrow Y$ and $R \rightarrow Y_R$. The data-processing inequality yields $d_{\text{TV}}(\mathcal{L}(Y), \mathcal{L}(Y_R)) \leq d_{\text{TV}}(\mathcal{L}(S), \mathcal{L}(R)) = \Delta$, implying the second item.

The second item can be achieved by rejection sampling from the distribution $\mathcal{L}(Y)$ until receiving an element congruent to x modulo p or reaching the cutoff of

$$m = \left\lceil \frac{\log \delta}{\log(1 - p^{-1} + \Delta)} \right\rceil = O\left(\frac{\log(1/\delta)}{p^{-1} - \Delta}\right)$$

rounds. Each sample from $\mathcal{L}(Y)$ can be obtained in $O(t)$ by sampling Z_0, Z_1, \dots, Z_t and forming the number Y with binary digits Z_t, Z_{t-1}, \dots, Z_0 . If we receive a sample by the m th round, then it is exactly sampled from the conditional distribution $\mathcal{L}(Y_x) = \mathcal{L}(Y|Y \equiv x \pmod{p})$. Therefore the total variation between the output of this algorithm and $\mathcal{L}(Y_x)$ is upper bounded by the probability that the rejection sampling scheme fails to output a sample. Now note that the probability that a sample is output in a single round is

$$\mathbb{P}[S = x] \geq p^{-1} - d_{\text{TV}}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p]) = p^{-1} - \Delta$$

by the definition of total variation. By the independence of sampling in different rounds, the probability that no sample is output is at most

$$(1 - \mathbb{P}[S = x])^m \leq (1 - p^{-1} + \Delta)^m \leq \delta$$

which completes the proof of the first item. \square

Lemma 4.1 now follows by combining Lemmas 4.3 and 4.4:

Proof of Lemma 4.1. Let $\epsilon' = \epsilon/(2p)$. Given x , we approximately sample Y_x up to ϵ' total variation error as in Lemma 4.4. Then we return the (X_0, \dots, X_t) which are uniquely determined by $Y_x = \sum_{i=0}^t 2^i X_i$.

We prove correctness of the algorithm. $Y_x \equiv x \pmod{p}$ by construction. And by Lemma 4.3, when $t \geq K \cdot c^{-1}(1-c)^{-1} \log(p/\epsilon^2) \log p$ for a large enough constant K , we have $d_{TV}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p]) \leq \epsilon'$. So by the first item of Lemma 4.4, we may sample from $\mathcal{L}(Y_x)$ within ϵ' total variation distance in $O(pt \log(p/\epsilon))$ time. Finally, by the second item of Lemma 4.4 we also have $d_{TV}(\mathcal{L}(Y), \mathcal{L}(Y_R)) \leq \epsilon'$ if $x = R \sim \text{Unif}[\mathbb{F}_p]$ is chosen uniformly at random in $[p]$. So overall, we have

$$d_{TV}(\mathcal{L}(X_0, \dots, X_t), \mathcal{L}(Z_0, \dots, Z_t)) \leq d_{TV}(\mathcal{L}(Y_R), \mathcal{L}(Y)) + \epsilon' \leq 2\epsilon' \leq \epsilon$$

which concludes the proof of the lemma. \square

We conclude with a sampling result analogous to Lemma 4.1, but for $p = 2$.

Lemma 4.5. *There is a constant $K > 0$ such that the following holds. Let $\epsilon > 0$, $t \geq Kc^{-1}(1-c)^{-1} \log(1/\epsilon)$, and $c \leq q_0, q_1, \dots, q_t \leq 1 - c$ for some $c \in (0, 1/2]$. And let $Z_i \sim \text{Ber}(q_i)$ be independent. Given a random variable $R \sim \text{Unif}[\mathbb{F}_2]$, in $O(t \log(1/\epsilon))$ time one may sample X_0, \dots, X_t supported on $\{0, 1\}$, such that $R = \sum_{i=0}^t X_i \pmod{2}$ and $d_{TV}(\mathcal{L}(X_0, \dots, X_t), \mathcal{L}(Z_0, \dots, Z_t)) < \epsilon$.*

Proof. By induction on t , one may show that

$$\mathbb{P} \left[\sum_{i=0}^t Z_i \equiv 0 \pmod{2} \right] = \frac{1}{2} - \frac{\prod_{i=0}^t (1 - 2q_i)}{2}$$

If t satisfies the lower bound $t \geq \lceil \log(\epsilon/2) / \log(|1 - 2c|) \rceil + 1 = \Omega(c^{-1}(1-c)^{-1} \log(1/\epsilon))$, it holds that $d_{TV}(\mathcal{L}(\sum_{i=0}^t Z_i), \mathcal{L}(R)) \leq \epsilon/2$. Sample the distribution

$$X \sim \mathcal{L} \left(Z \mid \sum_{i=0}^t Z_i \equiv R \pmod{2} \right)$$

within $\epsilon/2$ total variation distance by rejection sampling. This takes $O(t \log(1/\epsilon))$ time, because it consists of at most $O(\log(1/\epsilon))$ rounds of sampling fresh copies of $Z_i \sim \text{Ber}(q_i)$ for all $i \in \{0, \dots, t\}$ and checking if $\sum_{i=0}^t Z_i = R$. By triangle inequality, it suffices to show that $d_{TV}(\mathcal{L}(X), \mathcal{L}(Z)) \leq \epsilon/2$. This is true by data processing inequality, since R is uniform and hence is within $\epsilon/2$ total variation distance of $\sum_{i=0}^t Z_i \pmod{2}$. \square

5 Algorithms for Counting k -Cliques in $G(n, c, s)$

In this section, we consider several natural algorithms for counting k -cliques in $G(n, c, s)$ with $c = \Theta(n^{-\alpha})$ for some $\alpha \in (0, 1)$. The main objective of this section is to show that, when k and s are constant, these algorithms all run faster than all known algorithms for $\#(k, s)$ -CLIQUE on worst-case hypergraphs and nearly match the lower bounds from our reduction for certain k, c and s . This demonstrates that the average-case complexity of $\#(k, s)$ -CLIQUE on Erdős-Rényi hypergraphs is intrinsically different from its worst-case complexity. As discussed in Section 2.2, this also shows the necessity of a slowdown term comparable to $\Upsilon_{\#}$ in our worst-case to average-case reduction for $\#(k, s)$ -CLIQUE. We begin with a randomized sampling-based algorithm for counting k -cliques in $G(n, c, s)$, extending well-known greedy heuristics for finding k -cliques in random graphs. We then present an improvement to this algorithm in the graph case and a deterministic alternative.

5.1 Greedy Random Sampling

In this section, we consider a natural greedy algorithm GREEDY-RANDOM-SAMPLING for counting k -cliques in a s -uniform hypergraph $G \sim G(n, c, s)$ with $c = \Theta(n^{-\alpha})$. Given a subset of vertices $A \subseteq [n]$ of G , define $\text{CN}_G(A)$ to be

$$\text{CN}_G(A) = \{v \in V(G) \setminus A : B \cup \{v\} \in E(G) \text{ for all } (s-1)\text{-subsets } B \subseteq A\}$$

or, in other words, the set of common neighbors of A . The algorithm GREEDY-RANDOM-SAMPLING maintains a set S of k -subsets of $[n]$ and for T iterations does the following:

1. Sample distinct starting vertices v_1, v_2, \dots, v_{s-1} uniformly at random and proceed to sample the remaining vertices v_s, v_{s+1}, \dots, v_k iteratively so that v_{i+1} is chosen uniformly at random from $\text{CN}_G(v_1, v_2, \dots, v_i)$ if it is nonempty.
2. If k vertices $\{v_1, v_2, \dots, v_k\}$ are chosen then add $\{v_1, v_2, \dots, v_k\}$ to S if it is not already in S .

This algorithm is an extension of the classical greedy algorithm for finding $\log_2 n$ sized cliques in $G(n, 1/2)$ in [Kar76; GM75], the Metropolis process examined in [Jer92] and the greedy procedure solving k -CLIQUE on $G(n, c)$ with $c = \Theta(n^{-2/(k-1)})$ discussed by Rossman in [Ros16]. These and other natural polynomial time search algorithms fail to find cliques of size $(1+\epsilon)\log_2 n$ in $G(n, 1/2)$, even though its clique number is approximately $2\log_2 n$ with high probability [McD84; Pit82]. Our algorithm GREEDY-RANDOM-SAMPLING extends this greedy algorithm to count k -cliques in $G(n, c, s)$. In our analysis, we will see a phase transition in the behavior of this algorithm at $k = \tau$ for some τ smaller than the clique number of $G(n, c, s)$. This is analogous to the breakdown of the natural greedy algorithm at cliques of size $\log_2 n$ on $G(n, 1/2)$.

Before analyzing GREEDY-RANDOM-SAMPLING, we state a simple classical lemma counting the number of k -cliques in $G(n, c, s)$. This lemma follows from linearity of expectation and Markov's inequality. Its proof is included in Appendix D for completeness.

Lemma 5.1. *For fixed $\alpha \in (0, 1)$ and s , let $\kappa \geq s$ be the largest positive integer satisfying $\alpha \binom{\kappa}{s-1} < s$. If $G \sim G(n, c, s)$ where $c = O(n^{-\alpha})$, then $\mathbb{E}[|\text{cl}_k(G)|] = \binom{n}{k} c^{\binom{k}{s}}$ and $\omega(G) \leq \kappa + 1 + t$ with probability at least $1 - O\left(n^{-\alpha t(1-s^{-1})\binom{\kappa+2}{s-1}}\right)$ for any fixed nonnegative integer t .*

In particular, this implies that the clique number of $G(n, c, s)$ is typically at most $(s!\alpha^{-1})^{\frac{1}{s-1}} + s - 1$. In the graph case of $s = 2$, this simplifies to $1 + 2\alpha^{-1}$. In the next subsection, we give upper bounds on the number of iterations T causing all k -cliques in G to end up in S and analyze the runtime of

the algorithm. The subsequent subsection improves the runtime of GREEDY-RANDOM-SAMPLING for graphs when $s = 2$ through a matrix multiplication post-processing step. The last subsection gives an alternative deterministic algorithm with a similar performance to GREEDY-RANDOM-SAMPLING.

5.2 Sample Complexity and Runtime of Greedy Random Sampling

In this section, we analyze the runtime of GREEDY-RANDOM-SAMPLING and give upper bounds on the number of iterations T needed for the algorithm to terminate with $S = \text{cl}_k(G)$. The dynamic set S needs to support search and insertion of k -cliques. Consider labelling the vertices of G with elements of $[n]$ and storing the elements of S in a balanced binary search tree sorted according to the lexicographic order on $[n]^k$. Search and insertion can each be carried out in $O(\log |\text{cl}_k(G)|) = O(k \log n)$ time. It follows that each iteration of GREEDY-RANDOM-SAMPLING therefore takes $O(n + k \log n) = O(n)$ time as long as $k = O(1)$. Outputting $|S|$ in GREEDY-RANDOM-SAMPLING therefore yields a $O(nT)$ time algorithm for $\#(k, s)$ -CLIQUE on $G(n, c, s)$ that succeeds with high probability.

The following theorem provides upper bounds on the minimum number of iterations T needed for this algorithm to terminate with $S = \text{cl}_k(G)$ and therefore solve $\#(k, s)$ -CLIQUE. Its proof is deferred to Appendix E.

Theorem 5.2. *Let k and s be constants and $c = \Theta(n^{-\alpha})$ for some $\alpha \in (0, 1)$. Let τ be the largest integer satisfying $\alpha \binom{\tau}{s-1} < 1$ and suppose that*

$$T \geq \begin{cases} 2n^{\tau+1} c^{\binom{\tau+1}{s}} (\log n)^{3(k-\tau)(1+\epsilon)} & \text{if } k \geq \tau + 1 \\ 2n^k c^{\binom{k}{s}} (\log n)^{1+\epsilon} & \text{if } k < \tau + 1 \end{cases}$$

for some $\epsilon > 0$. Then GREEDY-RANDOM-SAMPLING run with T iterations terminates with $S = \text{cl}_k(G)$ with probability $1 - n^{-\omega(1)}$ over the random bits of the algorithm GREEDY-RANDOM-SAMPLING and with probability $1 - n^{-\omega(1)}$ over the choice of random hypergraph $G \sim G(n, c, s)$.

Implementing S as a balanced binary search tree and outputting $|S|$ in GREEDY-RANDOM-SAMPLING yields the following algorithmic upper bounds for $\#(k, s)$ -CLIQUE with inputs sampled from $G(n, c, s)$.

Corollary 5.3. *Suppose that k and s are constants and $c = \Theta(n^{-\alpha})$ for some $\alpha \in (0, 1)$. Let τ be the largest integer satisfying $\alpha \binom{\tau}{s-1} < 1$. Then it follows that*

1. *If $k \geq \tau + 1$, there is an $\tilde{O}\left(n^{\tau+2-\alpha \binom{\tau+1}{s}}\right)$ time randomized algorithm solving $\#(k, s)$ -CLIQUE on inputs sampled from $G(n, c, s)$ with probability at least $1 - n^{-\omega(1)}$.*
2. *If $k < \tau + 1$, there is an $\tilde{O}\left(n^{k+1-\alpha \binom{k}{s}}\right)$ time randomized algorithm solving $\#(k, s)$ -CLIQUE on inputs sampled from $G(n, c, s)$ with probability at least $1 - n^{-\omega(1)}$.*

By Lemma 5.1, the hypergraph $G \sim G(n, c, s)$ has clique number $\omega(G) \leq \kappa + 2$ with probability $1 - 1/\text{poly}(n)$ if where $\kappa \geq s$ is the largest positive integer satisfying $\alpha \binom{\kappa}{s-1} < s$. In particular, when $k > \kappa + 2$ in the theorem above, the algorithm outputting zero succeeds with probability $1 - 1/\text{poly}(n)$ and $\#(k, s)$ -CLIQUE is trivial. For there to typically be a nonzero number of k -cliques in $G(n, c, s)$, it should hold that $0 < \alpha \leq s \binom{k-1}{s-1}^{-1}$. In the graph case of $s = 2$, this simplifies to the familiar condition that $0 < \alpha \leq \frac{2}{k-1}$. We also remark that when $k < \tau + 1$, the runtime of this algorithm is an $\tilde{O}(n)$ factor off from the expectation of the quantity being counted, the number of k -cliques in $G \sim G(n, c, s)$.

5.3 Post-Processing with Matrix Multiplication

In this section, we improve the runtime of GREEDY-RANDOM-SAMPLING as an algorithm for $\#(k, s)$ -CLIQUE in the graph case of $s = 2$. The improvement comes from the matrix multiplication step of Nešetřil and Poljak from their $O(n^{\omega \lfloor k/3 \rfloor + (k \pmod{3})})$ time worst-case algorithm for $\#(k, 2)$ -CLIQUE [NP85]. Our improved runtime for the algorithm GREEDY-RANDOM-SAMPLING is stated in the following theorem.

Theorem 5.4. *Suppose that $k > 2$ is a fixed positive integer and $c = \Theta(n^{-\alpha})$ where $0 < \alpha \leq \frac{2}{k-1}$ is also fixed. Then there is a randomized algorithm solving $\#(k, 2)$ -CLIQUE on inputs sampled from $G(n, c)$ with probability $1 - n^{-\omega(1)}$ that runs in $\tilde{O}\left(n^{\omega \lfloor k/3 \rfloor + \omega - \omega \alpha \binom{k/3}{2}}\right)$ time.*

Proof. Label the vertices of an input graph $G \sim G(n, c)$ with the elements of $[n]$. Consider the following application of GREEDY-RANDOM-SAMPLING with post-processing:

1. Run GREEDY-RANDOM-SAMPLING to compute the two sets of cliques $S_1 = \text{cl}_{\lfloor k/3 \rfloor}(G)$ and $S_2 = \text{cl}_{\lfloor k/3 \rfloor}(G)$ with the number of iterations T as given in Theorem 5.2.
2. Construct the matrix $M_1 \in \{0, 1\}^{|S_1| \times |S_1|}$ with rows and columns indexed by the elements of S_1 such that $(M_1)_{A,B} = 1$ for $A, B \in S_1$ if $A \cup B$ forms a clique of G and all labels in A are strictly less than all labels in B .
3. Construct the matrix $M_2 \in \{0, 1\}^{|S_1| \times |S_2|}$ with rows indexed by the elements of S_1 and columns indexed by the elements of S_2 such that $(M_2)_{A,B} = 1$ for $A \in S_1$ and $B \in S_2$ under the same rule that $A \cup B$ forms a clique of G and all labels in A are strictly less than all labels in B . Construct the matrix M_3 with rows and columns indexed by S_2 analogously.
4. Compute the matrix product

$$M_P = \begin{cases} M_1^2 & \text{if } k \equiv 0 \pmod{3} \\ M_1 M_2 & \text{if } k \equiv 1 \pmod{3} \\ M_2 M_3 & \text{if } k \equiv 2 \pmod{3} \end{cases}$$

5. Output the sum of entries

$$\sum_{(A,B) \in \mathcal{S}} (M_P)_{A,B}$$

where \mathcal{S} is the support of M_1 if $k \equiv 0 \pmod{3}$ and \mathcal{S} is the support of M_2 if $k \not\equiv 0 \pmod{3}$.

We will show that this algorithm solves $\#(k, 2)$ -CLIQUE with probability $1 - n^{-\omega(1)}$ when $k \equiv 1 \pmod{3}$. The cases when $k \equiv 0, 2 \pmod{3}$ follow from a nearly identical argument. By Theorem 5.2, the first step applying GREEDY-RANDOM-SAMPLING succeeds with probability $1 - n^{-\omega(1)}$. Note that $(M_P)_{A,B}$ counts the number of $\lfloor k/3 \rfloor$ -cliques C in G such that the labels of C are strictly greater than those of A and less than those of B and such that $A \cup C$ and $C \cup B$ are both cliques. If it further holds that $(M_2)_{A,B} = 1$, then $A \cup B$ is a clique and $A \cup B \cup C$ is also clique. Therefore the sum output by the algorithm exactly counts the number of triples (A, B, C) such that $A \cup B \cup C$ is a clique, $|A| = |C| = \lfloor k/3 \rfloor$, $|B| = \lceil k/3 \rceil$ and the labels of C are greater than those of A and less than those of B . Observe that any clique $\mathcal{C} \in \text{cl}_k(G)$ is counted in this sum exactly once by the triple (A, B, C) where A consists of the lowest $\lfloor k/3 \rfloor$ labels in \mathcal{C} , B consists of the highest $\lceil k/3 \rceil$ labels in \mathcal{C} and C contains the remaining vertices of \mathcal{C} . Therefore this algorithm solves $\#(k, 2)$ -CLIQUE as long as Step 1 succeeds.

It suffices to analyze the additional runtime incurred by this post-processing. Observe that the number of cliques output by a call to greedy-random-sampling with T iterations is at most T . Also note that if $\alpha \leq \frac{2}{k-1}$, then $\tau \geq \lfloor \frac{k}{2} \rfloor - 1$. If $k \geq 3$, then it follows that $\tau + 1 \geq \lfloor \frac{k}{2} \rfloor \geq \lceil \frac{k}{3} \rceil$. It follows by Theorem 5.2 that $\max\{|S_1|, |S_2|\} = \tilde{O}\left(n^{\lceil k/3 \rceil + 1 - \alpha \binom{\lceil k/3 \rceil}{s}}\right)$. Note that computing the matrix M_P takes $\tilde{O}(\max\{|S_1|, |S_2|\}^\omega) = \tilde{O}\left(n^{\omega \lceil k/3 \rceil + \omega - \omega \alpha \binom{\lceil k/3 \rceil}{s}}\right)$ time. Now observe that all other steps of the algorithm run in $\tilde{O}\left(n^{2\lceil k/3 \rceil - 2\alpha \binom{\lceil k/3 \rceil}{s}}\right)$ time, which completes the proof of the theorem since the matrix multiplication constant satisfies $\omega \geq 2$. \square

We remark that for simplicity, we have ignored minor improvements in the runtime that can be achieved by more carefully analyzing Step 4 in terms of rectangular matrix multiplication constants if $k \not\equiv 0 \pmod{3}$. Note that the proof above implicitly used a weak large deviations bound on $|\text{cl}_k(G)|$. More precisely, it used the fact that if GREEDY-RANDOM-SAMPLING with T iterations succeeds, then $|\text{cl}_k(G)| \leq T$. Theorem 5.2 thus implies that $|\text{cl}_k(G)|$ is upper bounded by the minimal settings of T in the theorem statement with probability $1 - n^{-\omega(1)}$ over $G \sim G(n, c, s)$.

When $k \leq \tau + 1$, these upper bounds are a polylog(n) factor from the expectation of $|\text{cl}_k(G)|$. While this was sufficient in the proof of Theorem 5.4, stronger upper bounds will be needed in the next subsection to analyze our deterministic iterative algorithm. The upper tails of $|\text{cl}_k(G)|$ and more generally of the counts of small subhypergraphs in $G(n, c, s)$ have been studied extensively in the literature. We refer to [Vu01; JR02; JOR04; DK12] for a survey of the area and recent results. Given a hypergraph H , let $N(n, m, H)$ denote the largest number of copies of H that can be constructed in an s -uniform hypergraph with at most n vertices and m hyperedges. Define the quantity

$$M_H(n, c) = \max \left\{ m \leq \binom{n}{s} : N(n, m, H') \leq n^{|V(H')|} c^{|E(H')|} \text{ for all } H' \subseteq H \right\}$$

The following large deviations result from [DPR10] generalizes a graph large deviations bound from [JOR04] to hypergraphs to obtain the following result.

Theorem 5.5 (Theorem 4.1 from [DPR10]). *For every s -uniform hypergraph H and every fixed $\epsilon > 0$, there is a constant $C(\epsilon, H)$ such that for all $n \geq |V(H)|$ and $c \in (0, 1)$, it holds that*

$$\mathbb{P}[X_H \geq (1 + \epsilon)\mathbb{E}[X_H]] \leq \exp(-C(\epsilon, H) \cdot M_H(n, c))$$

where X_H is the number of copies of H in $G \sim G(n, c, s)$.

Proposition 4.3 in [DPR10] shows that if H is a d -regular s -uniform hypergraph and $c \geq n^{-s/d}$ then $M_H(n, c) = \Theta(n^s c^d)$. This implies that

$$\mathbb{P}\left[|\text{cl}_k(G)| \geq (1 + \epsilon) \binom{n}{k} c^{\binom{k}{s}}\right] \leq \exp\left(-C'(\epsilon) \cdot n^s c^{\binom{k-1}{s-1}}\right)$$

as long as $c \geq n^{-s!(k-s)!/(k-1)!}$. This provides strong bounds on the upper tails of $|\text{cl}_k(G)|$ that will be useful in the next subsection.

5.4 Deterministic Iterative Algorithm for Counting in $G(n, c, s)$

In this section, we present an alternative deterministic algorithm IT-GEN-CLIQUEs achieving a similar runtime to GREEDY-RANDOM-SAMPLING. Although they have very different analyses, the algorithm IT-GEN-CLIQUEs can be viewed as a deterministic analogue of GREEDY-RANDOM-SAMPLING.

Both are constructing cliques one vertex at a time. The algorithm IT-GEN-CLIQUEs takes in cutoffs C_{s-1}, C_s, \dots, C_k and generates sets S_{s-1}, S_s, \dots, S_k as follows:

1. Initialize S_{s-1} to be the set of all $(s-1)$ -subsets of $[n]$.
2. Given the set S_i , for each vertex $v \in [n]$, iterate through all subsets $A \in S_i$ and add $A \cup \{v\}$ to S_{i+1} if $A \cup \{v\}$ is a clique and v is larger than the labels of all of the vertices in A . Stop if ever $|S_{i+1}| \geq C_{i+1}$.
3. Stop once S_k has been generated and output S_k .

Suppose that C_t are chosen to be any high probability upper bounds on the number of t -cliques in $G \sim G(n, c, s)$ such as the bounds in Theorem 5.5. Then we have the following guarantees for the algorithm IT-GEN-CLIQUEs.

Theorem 5.6. *Suppose that k and s are constants and $c = \Theta(n^{-\alpha})$ for some $\alpha \in (0, 1)$. Let τ be the largest integer satisfying $\alpha \binom{\tau}{s-1} < 1$ and $C_t = 2n^t c^{\binom{t}{s}}$ for each $s \leq t \leq k$. Then IT-GEN-CLIQUEs with the cutoffs C_t outputs $S_k = \text{cl}_k(G)$ with probability $1 - n^{-\omega(1)}$ where*

1. *The runtime of IT-GEN-CLIQUEs is $O\left(n^{\tau+2-\alpha \binom{\tau+1}{s}}\right)$ if $k \geq \tau + 2$.*
2. *The runtime of IT-GEN-CLIQUEs is $O\left(n^{k-\alpha \binom{k-1}{s}}\right)$ if $k < \tau + 2$.*

Proof. We first show that $S_k = \text{cl}_k(G)$ with probability $1 - n^{-\omega(1)}$ in the algorithm IT-GEN-CLIQUEs. By a union bound and Theorem 5.5, it follows that $|\text{cl}_t(G)| < C_t$ for each $s \leq t \leq k$ with probability at least $1 - (k-s+1)n^{-\omega(1)}$. The following simple induction argument shows that $S_t = \text{cl}_t(G)$ for each $s-1 \leq t \leq k$ conditioned on this event. Note that $\text{cl}_{s-1}(G)$ is by definition the set of all $(s-1)$ -subsets of $[n]$ and thus $S_{s-1} = \text{cl}_{s-1}(G)$. If $S_t = \text{cl}_t(G)$, then each $(t+1)$ -clique \mathcal{C} of G is added exactly once to S_{t+1} as $A \cup \{v\}$ where v is the vertex of \mathcal{C} with the largest label and $A = \mathcal{C} \setminus \{v\} \in \text{cl}_t(G)$ are the remaining vertices. Now note that the runtime of IT-GEN-CLIQUEs is

$$O\left(\sum_{t=s-1}^{k-1} nC_t\right) = O\left(\max_{s-1 \leq t \leq k-1} (nC_t)\right) = \begin{cases} O\left(n^{\tau+2-\alpha \binom{\tau+1}{s}}\right) & \text{if } k \geq \tau + 2 \\ O\left(n^{k-\alpha \binom{k-1}{s}}\right) & \text{if } k < \tau + 2 \end{cases}$$

since $k = O(1)$. To see the second inequality, note that $\log_n(C_{t+1}/C_t) = 1 - \alpha \binom{t}{s-1}$. This implies that $C_{t+1} > C_t$ if $t \leq \tau$ and C_t is maximized on $s \leq t \leq k$ when $t = \tau + 1$. This completes the proof of the theorem. \square

We remark that in the case of $k < \tau + 1$, IT-GEN-CLIQUEs attains a small runtime improvement over GREEDY-RANDOM-SAMPLING. However, the algorithm GREEDY-RANDOM-SAMPLING can be modified to match this runtime up to a $\text{polylog}(n)$ factor by instead generating the $(k-1)$ -cliques of G and applying the last step of IT-GEN-CLIQUEs to generate the k -cliques of G . We also remark that IT-GEN-CLIQUEs can also be used instead of GREEDY-RANDOM-SAMPLING in Step 1 of the algorithm in Theorem 5.4, yielding a nearly identical runtime of $\tilde{O}\left(n^{\omega \lceil k/3 \rceil - \omega \alpha \binom{\lceil k/3 \rceil - 1}{2}}\right)$ for $\#(k, 2)$ -CLIQUE on inputs sampled from $G(n, c)$.

6 Extensions and Open Problems

In this section, we outline several extensions of our methods and problems left open after our work.

Improved Average-Case Lower Bounds. A natural question is if tight average-case lower bounds for $\#(k, s)$ -CLIQUE can be shown above the k -clique percolation threshold when $s \geq 3$ and if the constant C in the exponent of our lower bounds for the graph case of $s = 2$ can be improved from 1 to $\omega/9$.

Raising Error Tolerance for Average-Case Hardness. A natural question is if the error tolerance of the worst-case to average-case reductions in Theorems 2.8 and 2.9 can be increased. We remarked in the introduction that for certain choices of k , the error tolerance cannot be significantly increased – for example, when $k = 3 \log_2 n$, the trivial algorithm that outputs 0 on any graph has subpolynomial error on graphs drawn from $G(n, 1/2)$, but is useless for reductions from worst-case graphs. Nevertheless, for other regimes of k , such as when $k = O(1)$ is constant, counting k -cliques with error probability less than $1/4$ on graphs drawn from $G(n, 1/2)$ appears to be nontrivial. It is an open problem to prove hardness for such a regime. In general, one could hope to understand the tight tradeoffs between computation time, error tolerance, k , c , and s for k -clique-counting on $G(n, c, s)$.

Hardness of Approximating Clique Counts. Another interesting question is if it is hard to approximate the k -clique counts, within some additive error ϵ , of hypergraphs drawn from $G(n, c, s)$. Since the number of k -cliques in $G(n, c, s)$ concentrates around the mean $\mu \approx c^{\binom{k}{s}} n^k$ with standard deviation σ , one would have to choose $\epsilon \ll \sigma$ for approximation to be hard.

Inhomogeneous Erdős-Rényi Hypergraphs. Consider an inhomogeneous Erdős-Rényi hypergraph model, where each hyperedge e is independently chosen to be in the hypergraph with probability $c(e)$. Also suppose that we may bound $c(e)$ uniformly away from 0 and 1 (that is, $c(e) \in [c, 1 - c]$ for all possible hyperedges e and for some constant c). We would like to prove that $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE are hard on average for inhomogeneous Erdős-Rényi hypergraphs. Unfortunately, this does not follow directly from our proof techniques because step 5 in the proof of Theorems 2.8 and 2.9 breaks down due to the inhomogeneity of the model. Nevertheless, steps 1-4 still hold, and therefore we can show that $\#(k, s)$ -CLIQUE and PARITY- (k, s) -CLIQUE are average-case hard for k -partite inhomogeneous Erdős-Rényi hypergraphs – when only the edges e that respect the k -partition are chosen to be in the hypergraph with inhomogeneous edge-dependent probability $c(e) \in [c, 1 - c]$.

General Subgraph Counts. Given a hypergraph H on k vertices, let H -COUNTING be the problem of counting the number of occurrences (as an induced subgraph) of H in a hypergraph G . Can one show that H -COUNTING in the worst case reduces to H -COUNTING in the average case on Erdős-Rényi hypergraphs?

Our reduction (Theorem 2.8) applies to the special case when H is a clique. Unfortunately, the proof of Theorem 2.8 breaks down when counting general hypergraphs. First, the reductions to and from k -partite hypergraphs (steps 1 and 5) no longer work, because H contains non-edges, and therefore there may be a copy of H that contains more than one vertex in a given k -partition. In order to remedy this, we could consider the modification H -COUNTING' of the H -COUNTING problem that respects k -partite structure, by only counting the copies of H in a k -partite hypergraph G , such that the k vertices of the copy of H lie in the k different parts of the vertex partition of G . For this modified problem, the strategy of our reduction still fails – this time at Step 4, because the polynomial that counts copies of H in G is not homogeneous. Indeed, for clique-counting, Step 4

of the reduction uses the fact that the variables of the clique-counting polynomial can be split up into $\binom{k}{s}$ groups, such that each monomial contained exactly one variable from each group.

References

- [AB87] Noga Alon and Ravi B Boppana. “The monotone circuit complexity of Boolean functions”. In: *Combinatorica* 7.1 (1987), pp. 1–22.
- [Ajt96] Miklós Ajtai. “Generating hard instances of lattice problems”. In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*. ACM. 1996, pp. 99–108.
- [AKS98] Noga Alon, Michael Krivelevich, and Benny Sudakov. “Finding a large hidden clique in a random graph”. In: *Random Structures and Algorithms* 13.3-4 (1998), pp. 457–466.
- [AM05] Kazuyuki Amano and Akira Maruoka. “A superpolynomial lower bound for a circuit computing the clique function with at most $(1/6) \log \log n$ negation gates”. In: *SIAM Journal on Computing* 35.1 (2005), pp. 201–216.
- [Ats+18] Albert Atserias, Ilario Bonacina, Susanna F de Rezende, Massimo Lauria, Jakob Nordström, and Alexander Razborov. “Clique is hard on average for regular resolution”. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. ACM. 2018, pp. 866–877.
- [AV11] Brendan PW Ames and Stephen A Vavasis. “Nuclear norm minimization for the planted clique and biclique problems”. In: *Mathematical programming* 129.1 (2011), pp. 69–89.
- [Bal+17] Marshall Ball, Alon Rosen, Manuel Sabin, and Prashant Nalini Vasudevan. “Average-case fine-grained hardness”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM. 2017, pp. 483–496.
- [Bar+16] Boaz Barak, Samuel B Hopkins, Jonathan Kelner, Pravesh Kothari, Ankur Moitra, and Aaron Potechin. “A nearly tight sum-of-squares lower bound for the planted clique problem”. In: *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*. IEEE. 2016, pp. 428–437.
- [BB19] Matthew Brennan and Guy Bresler. “Optimal Average-Case Reductions to Sparse PCA: From Weak Assumptions to Strong Hardness”. In: *arXiv preprint arXiv:1902.07380* (2019).
- [BBH18] Matthew Brennan, Guy Bresler, and Wasim Huleihel. “Reducibility and Computational Lower Bounds for Problems with Planted Sparse Structure”. In: *Conference On Learning Theory*. 2018, pp. 48–166.
- [BBH19] Matthew Brennan, Guy Bresler, and Wasim Huleihel. “Universality of Computational Lower Bounds for Submatrix Detection”. In: *arXiv preprint arXiv:1902.06916* (2019).
- [BDS93] Eric Bach, James Driscoll, and Jeffrey Shallit. “Factor refinement”. In: *Journal of Algorithms* 15.2 (1993), pp. 199–222.
- [BR09] Béla Bollobás and Oliver Riordan. “Clique percolation”. In: *Random Structures & Algorithms* 35.3 (2009), pp. 294–322.
- [BR13] Quentin Berthet and Philippe Rigollet. “Complexity Theoretic Lower Bounds for Sparse Principal Component Detection.” In: *COLT*. 2013, pp. 1046–1066.

- [BT06a] Andrej Bogdanov and Luca Trevisan. “Average-case complexity”. In: *Foundations and Trends® in Theoretical Computer Science* 2.1 (2006), pp. 1–106.
- [BT06b] Andrej Bogdanov and Luca Trevisan. “On worst-case to average-case reductions for NP problems”. In: *SIAM Journal on Computing* 36.4 (2006), pp. 1119–1159.
- [Cal+08] Chris Calabro, Russell Impagliazzo, Valentine Kabanets, and Ramamohan Paturi. “The complexity of unique k-SAT: An isolation lemma for k-CNFs”. In: *Journal of Computer and System Sciences* 74.3 (2008), pp. 386–393.
- [CE15] Amin Coja-Oghlan and Charilaos Efthymiou. “On independent sets in random graphs”. In: *Random Structures & Algorithms* 47.3 (2015), pp. 436–486.
- [Che+06] Jianer Chen, Xiuzhen Huang, Iyad A Kanj, and Ge Xia. “Strong computational lower bounds via parameterized complexity”. In: *Journal of Computer and System Sciences* 72.8 (2006), pp. 1346–1367.
- [Che15] Yudong Chen. “Incoherence-optimal matrix completion”. In: *IEEE Transactions on Information Theory* 61.5 (2015), pp. 2909–2923.
- [CPS99] Jin-Yi Cai, Aduri Pavan, and D Sivakumar. “On the hardness of permanent”. In: *Annual Symposium on Theoretical Aspects of Computer Science*. Springer, 1999, pp. 90–99.
- [CX16] Yudong Chen and Jiaming Xu. “Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices”. In: *Journal of Machine Learning Research* 17.27 (2016), pp. 1–57.
- [DF95] Rod G Downey and Michael R Fellows. “Fixed-parameter tractability and completeness II: On completeness for W [1]”. In: *Theoretical Computer Science* 141.1-2 (1995), pp. 109–131.
- [DGM08] Sergey N Dorogovtsev, Alexander V Goltsev, and José FF Mendes. “Critical phenomena in complex networks”. In: *Reviews of Modern Physics* 80.4 (2008), p. 1275.
- [DGP14] Yael Dekel, Ori Gurel-Gurevich, and Yuval Peres. “Finding hidden cliques in linear time with high probability”. In: *Combinatorics, Probability and Computing* 23.1 (2014), pp. 29–49.
- [DK12] Robert DeMarco and Jeff Kahn. “Tight upper tail bounds for cliques”. In: *Random Structures & Algorithms* 41.4 (2012), pp. 469–487.
- [DM15] Yash Deshpande and Andrea Montanari. “Finding hidden cliques of size $\sqrt{N/e}$ in nearly linear time”. In: *Foundations of Computational Mathematics* 15.4 (2015), pp. 1069–1128.
- [DPR10] Andrzej Dudek, Joanna Polcyn, and Andrzej Ruciński. “Subhypergraph counts in extremal and random hypergraphs and the fractional q-independence”. In: *Journal of combinatorial optimization* 19.2 (2010), pp. 184–199.
- [DPV05] Imre Derényi, Gergely Palla, and Tamás Vicsek. “Clique percolation in random networks”. In: *Physical review letters* 94.16 (2005), p. 160202.
- [Fei+18] Uriel Feige, David Gamarnik, Joe Neeman, Miklós Z Rácz, and Prasad Tetali. “Finding cliques using few probes”. In: *arXiv preprint arXiv:1809.06950* (2018).
- [Fel+13] Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh Vempala, and Ying Xiao. “Statistical algorithms and a lower bound for detecting planted cliques”. In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, 2013, pp. 655–664.

- [FF93] Joan Feigenbaum and Lance Fortnow. “Random-self-reducibility of complete sets”. In: *SIAM Journal on Computing* 22.5 (1993), pp. 994–1005.
- [FK00] Uriel Feige and Robert Krauthgamer. “Finding and certifying a large hidden clique in a semirandom graph”. In: *Random Structures and Algorithms* 16.2 (2000), pp. 195–208.
- [FL92] Uriel Feige and Carsten Lund. “On the hardness of computing the permanent of random matrices”. In: *Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*. ACM. 1992, pp. 643–654.
- [FR10] Uriel Feige and Dorit Ron. “Finding hidden cliques in linear time”. In: *21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (AofA’10)*. Discrete Mathematics and Theoretical Computer Science. 2010, pp. 189–204.
- [Gam18] David Gamarnik. “Computing the partition function of the Sherrington-Kirkpatrick model is hard on average”. In: *arXiv preprint arXiv:1810.05907* (2018).
- [Gao93] Shuhong Gao. “Normal Bases over Finite Fields”. In: *Doctoral thesis, Waterloo* (1993).
- [Gem+91] Peter Gemmell, Richard Lipton, Ronitt Rubinfeld, Madhu Sudan, and Avi Wigderson. “Self-testing/correcting for polynomials and for approximate functions”. In: *Proceedings of the twenty-third annual ACM symposium on Theory of computing*. ACM. 1991, pp. 33–42.
- [GM75] Geoffrey R Grimmett and Colin JH McDiarmid. “On colouring random graphs”. In: *Mathematical Proceedings of the Cambridge Philosophical Society*. Vol. 77. 2. Cambridge University Press. 1975, pp. 313–324.
- [GR17] Oded Goldreich and Guy N Rothblum. “Worst-case to Average-case reductions for subclasses of P”. In: 2017.
- [GR18] Oded Goldreich and Guy Rothblum. “Counting t-cliques: Worst-case to average-case reductions and Direct interactive proof systems”. In: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2018, pp. 77–88.
- [GS14] David Gamarnik and Madhu Sudan. “Limits of local algorithms over sparse random graphs”. In: *Proceedings of the 5th conference on Innovations in theoretical computer science*. ACM. 2014, pp. 369–376.
- [GS92] Peter Gemmell and Madhu Sudan. “Highly resilient correctors for polynomials”. In: *Inf. Process. Lett.* 43.4 (1992), pp. 169–174.
- [HWX15] Bruce E Hajek, Yihong Wu, and Jiaming Xu. “Computational Lower Bounds for Community Detection on Random Graphs.” In: *COLT*. 2015, pp. 899–928.
- [IR78] Alon Itai and Michael Rodeh. “Finding a minimum circuit in a graph”. In: *SIAM Journal on Computing* 7.4 (1978), pp. 413–423.
- [Jer92] Mark Jerrum. “Large cliques elude the Metropolis process”. In: *Random Structures & Algorithms* 3.4 (1992), pp. 347–359.
- [JOR04] Svante Janson, Krzysztof Oleszkiewicz, and Andrzej Ruciński. “Upper tails for subgraph counts in random graphs”. In: *Israel Journal of Mathematics* 142.1 (2004), pp. 61–92.
- [JP00] Ari Juels and Marcus Peinado. “Hiding cliques for cryptographic security”. In: *Designs, Codes and Cryptography* 20.3 (2000), pp. 269–280.
- [JR02] Svante Janson and Andrzej Ruciński. “The infamous upper tail”. In: *Random Structures & Algorithms* 20.3 (2002), pp. 317–342.

- [Kar72] Richard M Karp. “Reducibility among combinatorial problems”. In: *Complexity of computer computations*. Springer, 1972, pp. 85–103.
- [Kar76] Richard M Karp. “Probabilistic analysis of some combinatorial search problems”. In: *Algorithms and Complexity: New Directions and Recent Results*. Academic Press. 1976.
- [Kuc95] Luděk Kučera. “Expected complexity of graph partitioning problems”. In: *Discrete Applied Mathematics* 57.2-3 (1995), pp. 193–212.
- [KZ14] Pascal Koiran and Anastasios Zouzias. “Hidden cliques and the certification of the restricted isometry property”. In: *IEEE Transactions on Information Theory* 60.8 (2014), pp. 4999–5006.
- [Lip89] Richard J Lipton. “New Directions In Testing.” In: *Distributed Computing and Cryptography* 2 (1989), pp. 191–202.
- [McD84] Colin McDiarmid. “Colouring random graphs”. In: *Annals of Operations Research* 1.3 (1984), pp. 183–200.
- [McS01] Frank McSherry. “Spectral partitioning of random graphs”. In: *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*. IEEE. 2001, pp. 529–537.
- [Mul54] David E Muller. “Application of Boolean algebra to switching circuit design and to error detection”. In: *Transactions of the IRE professional group on electronic computers* 3 (1954), pp. 6–12.
- [MW15] Zongming Ma and Yihong Wu. “Computational barriers in minimax submatrix detection”. In: *The Annals of Statistics* 43.3 (2015), pp. 1089–1116.
- [NP85] Jaroslav Nešetřil and Svatopluk Poljak. “On the complexity of the subgraph problem”. In: *Commentationes Mathematicae Universitatis Carolinae* 26.2 (1985), pp. 415–419.
- [PDV07] Gergely Palla, Imre Derényi, and Tamás Vicsek. “The critical point of k-Clique percolation in the Erdős–Rényi graph”. In: *Journal of Statistical Physics* 128.1-2 (2007), pp. 219–227.
- [Pit82] B Pittel. “On the probable behaviour of some algorithms for finding the stability number of a graph”. In: *Mathematical Proceedings of the Cambridge Philosophical Society*. Vol. 92. 3. Cambridge University Press. 1982, pp. 511–526.
- [Raz85] Alexander A Razborov. “Lower bounds for the monotone complexity of some Boolean functions”. In: *Soviet Math. Dokl.* Vol. 31. 1985, pp. 354–357.
- [Reg] Oded Regev. “The learning with errors problem”. In: ().
- [Reg09] Oded Regev. “On lattices, learning with errors, random linear codes, and cryptography”. In: *Journal of the ACM (JACM)* 56.6 (2009), p. 34.
- [Ros08] Benjamin Rossman. “On the constant-depth complexity of k-clique”. In: *Proceedings of the fortieth annual ACM symposium on Theory of computing*. ACM. 2008, pp. 721–730.
- [Ros10] Benjamin Rossman. “The monotone complexity of k-clique on random graphs”. In: *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE. 2010, pp. 193–201.
- [Ros16] Benjamin Rossman. *Lower bounds for subgraph isomorphism*. 2016.
- [RV17] Mustazee Rahman and Balint Virag. “Local algorithms for independent sets are half-optimal”. In: *The Annals of Probability* 45.3 (2017), pp. 1543–1577.

- [Sud97] Madhu Sudan. “Decoding of Reed Solomon codes beyond the error-correction bound”. In: *Journal of complexity* 13.1 (1997), pp. 180–193.
- [Val79] Leslie G Valiant. “The Complexity of Enumeration and Reliability Problems”. In: *SIAM Journal on Computing* 8.3 (1979), p. 410.
- [Vu01] Van H Vu. “A large deviation result on the number of small subgraphs of a random graph”. In: *Combinatorics, Probability and Computing* 10.1 (2001), pp. 79–94.
- [Yus06] Raphael Yuster. “Finding and counting cliques and independent sets in r-uniform hypergraphs”. In: *Information Processing Letters* 99.4 (2006), pp. 130–134.

A Reduction from DECIDE- (k, s) -CLIQUE to PARITY- (k, s) -CLIQUE

The following is a precise statement and proof of the reduction from DECIDE- (k, s) -CLIQUE to PARITY- (k, s) -CLIQUE claimed in Section 2.1.

Lemma A.1. *Given an algorithm A for PARITY- (k, s) -CLIQUE with error probability $< 1/3$ on any s -uniform hypergraph G , there is an algorithm B that runs in time $O(k2^k|A|)$ and solves DECIDE- (k, s) -CLIQUE with error $< 1/3$ on any s -uniform hypergraph G .*

Proof. Let $\text{cl}_k(G)$ denote the set of k -cliques in hypergraph $G = (V, E)$. Consider the polynomial

$$P_G(x_V) = \sum_{S \in \text{cl}_k(G)} \prod_{v \in S} x_v \pmod{2},$$

over the finite field \mathbb{F}_2 . If G has a k -clique at vertices $S \subset V$, then P_G is nonzero, because $P_G(1_S) = 1$. If G has no k -clique, then P_G is zero. Therefore, deciding whether G has a k -clique reduces to testing whether or not P_G is identically zero. P_G is of degree at most k , so if P_G is nonzero on at least one input, then it is nonzero on at least a 2^{-k} fraction of inputs. One way to see this is that if we evaluate P_G at all points $a \in \{0, 1\}^m$, the result is a non-zero Reed-Muller codeword in $RM(k, m)$. Since the distance of the $RM(k, m)$ code is 2^{m-k} , and the block-length is 2^m , the claim follows [Mul54]. We therefore evaluate P_G at $c \cdot 2^k$ independent random inputs for some large enough $c > 0$, accept if any of the evaluations returns 1, and reject if all of the evaluations return 0. Each evaluation corresponds to calculating PARITY- (k, s) -CLIQUE on a hypergraph G' formed from G by removing each vertex independently with probability $1/2$. As usual, we boost the error of A by running the algorithm $O(k)$ times for each evaluation, and using the majority vote. \square

B Proof of Lemma 3.6

We restate and prove Lemma 3.6.

Lemma B.1 (Theorem 4 of [GS92]). *Let \mathbb{F} be a finite field with $|\mathbb{F}| = q$ elements. Let $N, D > 0$. Suppose $9 < D < q/12$. Let $f : \mathbb{F}^N \rightarrow \mathbb{F}$ be a polynomial of degree at most D . If there is an algorithm A running in time $T(A, N)$ such that*

$$\mathbb{P}_{x \sim \text{Unif}[\mathbb{F}^N]}[A(x) = f(x)] > 2/3,$$

then there is an algorithm B running in time $O((N + D)D^2 \log^2 q + T(A, N) \cdot D)$ such that for any $x \in \mathbb{F}^N$,

$$\mathbb{P}[B(x) = f(x)] > 2/3.$$

Proof. Our proof of the lemma is based off of the proof that appears in [Bal+17]. The only difference is that in [Bal+17], the lemma is stated only for finite fields whose size is a prime. Suppose we wish to calculate $f(x)$ for $x \in \mathbb{F}^N$. In order to do this, choose $y_1, y_2 \stackrel{i.i.d}{\sim} \text{Unif}[\mathbb{F}^N]$, and define the polynomial $g(t) = x + ty_1 + t^2y_2$ where $t \in \mathbb{F}$. We evaluate $A(g(t))$ at m different values $t_1, \dots, t_m \in \mathbb{F}$. This takes $O(mND \log^2 q + m \cdot T(A, N))$ time. Suppose that we have the guarantee that at most $(m - 2D)/2$ of these evaluations are incorrect. Then, since $f(g(t))$ is a univariate polynomial of degree at most $2D$, we may use Berlekamp-Welch to recover $f(g(0)) = A(x)$ in $O(m^3)$ arithmetic operations over \mathbb{F} , each of which takes $O(\log^2 q)$ time. Since $g(t_i)$ and $g(t_j)$ are pairwise independent and uniform in \mathbb{F}^N for any distinct $t_i, t_j \neq 0$, by the second-moment method, with probability $> 2/3$, at most $(m - 2D)/2$ evaluations of $A(g(t))$ will be incorrect if we take $m = 12D$. \square

C Tightness of Bounds in Section 4

In this appendix, we briefly discuss the tightness of the bounds on t in Lemma 4.3 and how the case of $c = 1/2$ differs from $c \neq 1/2$. Note that if $q_i = 1/2$ for each i , then $\sum_{i=0}^t Z_i \cdot 2^i$ is uniformly distributed on $\{0, 1, \dots, 2^{t+1} - 1\}$. It follows that

$$d_{\text{TV}}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p]) = \sum_{x \in \mathbb{F}_p} |p^{-1} - \mathbb{P}[S = x]|_+ = \frac{a(p-a)}{2^{t+1}p} \leq \frac{p}{2^{t+1}}$$

if $0 \leq a \leq p-1$ is such that $2^{t+1} \equiv a \pmod{p}$. Here, $|\cdot|_+$ denotes $|x|_+ = \max(x, 0)$. Therefore S is within total variation of $1/\text{poly}(p)$ of $\text{Unif}[\mathbb{F}_p]$ if $t = \Omega(\log p)$. However, note that for c constant and $\epsilon = 1/\text{poly}(p)$, our lemma requires that $t = \Omega(\log^2 p)$. This raises the question: is the additional factor of $\log p$ necessary or an artifact of our analysis? We answer this question with an example suggesting that the extra $\log p$ factor is in fact necessary and that the case $c = 1/2$ is special.

Suppose that p is a Mersenne prime with $p = 2^r - 1$ for some prime r and for simplicity, take $q_i = 1/3$ for each i . Observe by the triangle inequality that

$$\left| \hat{f}(1) \right| = \left| \sum_{x \in \mathbb{F}_p} (f(x) - p^{-1}) \cdot \omega^x \right| \leq \|f - p^{-1} \cdot \mathbf{1}\|_1 = 2 \cdot d_{\text{TV}}(\mathcal{L}(S), \text{Unif}[\mathbb{F}_p])$$

Now suppose that $t = ar - 1$ for some positive integer a . As shown in the lemma, we have

$$\left| \hat{f}(1) \right|^2 = \prod_{i=0}^t \left| \frac{2}{3} + \frac{1}{3} \cdot \omega^{2^i} \right|^2 = \left[\prod_{i=0}^{r-1} \left(\frac{5}{9} + \frac{4}{9} \cdot \cos \left(\frac{2\pi}{p} \cdot 2^i \right) \right) \right]^a$$

where the second equality is due to the fact that the sequence 2^i has period r modulo p . Now observe that since $\frac{5}{9} + \frac{4}{9} \cdot \cos(x) \geq e^{-x^2}$, we have that

$$\prod_{i=0}^{r-1} \left(\frac{5}{9} + \frac{4}{9} \cdot \cos \left(\frac{2\pi}{p} \cdot 2^i \right) \right) \geq \exp \left(-\frac{4\pi^2}{p^2} \sum_{i=0}^{r-1} 2^{2i} \right) = \exp \left(-\frac{4\pi^2}{p^2} \cdot \frac{2^{2r} - 1}{3} \right) = \Omega(1)$$

which implies that a should be $\Omega(r)$ for $\hat{f}(1)$ to be polynomially small in p . Thus the extra $\log p$ factor is necessary in this case and our analysis is tight. Note that in the special case of $c = 1/2$, the factors in the expressions for $\hat{f}(s)$ are of the form $\frac{1}{2} + \frac{1}{2} \cdot \omega^{2^i \cdot s}$ which can be arbitrarily close to zero. We remark that the construction, as stated, relies on there being infinitely many Mersenne primes. However, it seems to suggest that the extra $\log p$ factor is necessary. Furthermore, similar examples can be produced with p that are not Mersenne, as long as the order of 2 modulo p is relatively small.

D Clique Counts in Sparse Erdős-Rényi Hypergraphs

We prove the following classical lemma from Section 5.1.

Lemma D.1. *For fixed $\alpha \in (0, 1)$ and s , let $\kappa \geq s$ be the largest positive integer satisfying $\alpha \binom{\kappa}{s-1} < s$. If $G \sim G(n, c, s)$ where $c = O(n^{-\alpha})$, then $\mathbb{E}[\text{cl}_k(G)] = \binom{n}{k} c^{\binom{k}{s}}$ and $\omega(G) \leq \kappa + 1 + t$ with probability at least $1 - O\left(n^{-\alpha t(1-s^{-1})\binom{\kappa+2}{s-1}}\right)$ for any fixed nonnegative integer t .*

Proof. Let $C > 0$ be such that $c \leq Cn^{-\alpha}$ for sufficiently large n . For any given set $\{v_1, v_2, \dots, v_k\}$ of k vertices in $[n]$, the probability that all hyperedges are present among $\{v_1, v_2, \dots, v_k\}$ and thus these vertices form a k -clique in G is $c^{\binom{k}{s}}$. Linearity of expectation implies that the expected number of k -cliques is $\mathbb{E}[\text{cl}_k(G)] = \binom{n}{k} c^{\binom{k}{s}}$. Now consider taking $k = \kappa + 2 + t$ and note that

$$\begin{aligned} \mathbb{E}[\text{cl}_k(G)] &= \binom{n}{k} c^{\binom{k}{s}} \\ &\leq n^k c^{\binom{k}{s}} \leq C^{\binom{k}{s}} \cdot \exp\left(\left(1 - \frac{\alpha}{s} \binom{k-1}{s-1}\right) k \log n\right) \\ &\leq C^{\binom{k}{s}} \cdot \exp\left(\left(1 - \frac{\alpha}{s} \binom{\kappa+1}{s-1}\right) k \log n - \frac{\alpha}{s} \cdot t \binom{\kappa+1}{s-2} k \log n\right) \\ &\leq C^{\binom{k}{s}} n^{-\alpha t(1-s^{-1})\binom{\kappa+2}{s-1}} \end{aligned}$$

since $k \geq \kappa + 2$ and $\binom{\kappa+1+t}{s-1} \geq \binom{\kappa+1}{s-1} + t \binom{\kappa+1}{s-2}$ by iteratively applying Pascal's identity. Observe that $\kappa = O(1)$ and thus $C^{\binom{k}{s}} = O(1)$. Now by Markov's inequality, it follows that $\mathbb{P}[\omega(G) \geq k] = \mathbb{P}[\text{cl}_k(G) \geq 1] \leq \mathbb{E}[\text{cl}_k(G)]$, completing the proof of the lemma. \square

E Analysis of Greedy Random Sampling

This section is devoted to proving Theorem 5.2, which is restated below for convenience.

Theorem E.1. *Let k and s be constants and $c = \Theta(n^{-\alpha})$ for some $\alpha \in (0, 1)$. Let τ be the largest integer satisfying $\alpha \binom{\tau}{s-1} < 1$ and suppose that*

$$T \geq \begin{cases} 2n^{\tau+1} c^{\binom{\tau+1}{s}} (\log n)^{3(k-\tau)(1+\epsilon)} & \text{if } k \geq \tau + 1 \\ 2n^k c^{\binom{k}{s}} (\log n)^{1+\epsilon} & \text{if } k < \tau + 1 \end{cases}$$

for some $\epsilon > 0$. Then GREEDY-RANDOM-SAMPLING run with T iterations terminates with $S = \text{cl}_k(G)$ with probability $1 - n^{-\omega(1)}$ over the random bits of the algorithm GREEDY-RANDOM-SAMPLING and with probability $1 - n^{-\omega(1)}$ over the choice of random hypergraph $G \sim G(n, c, s)$.

Proof. We first consider the case where $k \geq \tau + 1$. Fix some $\epsilon > 0$ and let $v = (v_1, v_2, \dots, v_k)$ be an ordered tuple of distinct vertices in $[n]$. Define the random variable

$$Z_v = n(n-1) \cdots (n-s+2) \prod_{i=s-1}^{k-1} |\text{CN}_G(v_1, v_2, \dots, v_i)|$$

The key property of Z_v is that, in each iteration of GREEDY-RANDOM-SAMPLING, the probability that the k vertices v_1, v_2, \dots, v_k are chosen in that order is exactly $1/Z_v$. The proof of this theorem

will proceed by establishing upper bounds on Z_v that hold for all k -cliques v with high probability over the randomness of G , which will yield a bound on the number of iterations T needed to exhaust all such k -cliques in G .

Consider the following event over the sampling $G \sim G(n, c, s)$

$$A_v = \left\{ Z_v \geq 2n^{\tau+1} c^{\binom{\tau+1}{s}} (\log n)^{3(k-1-\tau)(1+\epsilon)} \quad \text{and} \quad \{v_1, v_2, \dots, v_k\} \in \text{cl}_k(G) \right\}$$

We now proceed to bound the probability of A_v through simple Chernoff and union bounds over G . In the next part of the argument, we condition on the event that $\{v_1, v_2, \dots, v_k\}$ forms a clique in G . For each $i \in \{s-1, s, \dots, k-1\}$, let $Y_{v,i}$ be the number of common neighbors of v_1, v_2, \dots, v_i in $V(G) \setminus \{v_1, v_2, \dots, v_k\}$. Note that $Y_{v,i} \sim \text{Bin}\left(n-k, c^{\binom{i}{s-1}}\right)$ and that $|\text{CN}_G(v_1, v_2, \dots, v_i)| = k-i + Y_{v,i}$. The standard Chernoff bound for the binomial distribution implies that for all $\delta_i > 0$,

$$\begin{aligned} & \mathbb{P} \left[|\text{CN}_G(v_1, v_2, \dots, v_i)| \geq k-i + (1+\delta_i)(n-k)c^{\binom{i}{s-1}} \right] \\ & \leq \exp \left(-\frac{\delta_i^2}{2+\delta_i} \cdot (n-k)c^{\binom{i}{s-1}} \right) \end{aligned}$$

Now define κ_i to be

$$\kappa_i = (n-k)^{-1} c^{-\binom{i}{s-1}} \cdot (\log n)^{1+\epsilon}$$

for each $i \in \{s-1, s, \dots, k-1\}$. Let $\delta_i = \sqrt{\kappa_i}$ if $i \leq \tau$ and $\delta_i = \kappa_i$ if $i > \tau$. Note that for sufficiently large n , $\delta_i < 1$ if $i \leq \tau$ and $\delta_i \geq 1$ if $i > \tau$. These choices of δ_i ensure that the Chernoff upper bounds above are each at most $\exp(-\frac{1}{3}(\log n)^{1+\epsilon})$ for each i . A union bound implies that with probability at least $1 - k \exp(-\frac{1}{3}(\log n)^{1+\epsilon})$, it holds that

$$|\text{CN}_G(v_1, v_2, \dots, v_i)| < k-i + (1+\delta_i)(n-k)c^{\binom{i}{s-1}} < (1+2\delta_i)(n-k)c^{\binom{i}{s-1}}$$

for all i and sufficiently large n . Here, we used the fact that $\delta_i(n-k)c^{\binom{i}{s-1}} = \omega(1)$ for all i by construction and $k = O(1)$. Observe that $(1+2\delta_i)(n-k)c^{\binom{i}{s-1}} \leq 3(\log n)^{1+\epsilon}$ for all $i \geq \tau+1$. These inequalities imply that

$$\begin{aligned} \log Z_v & < \log n^{s-1} + \sum_{i=s-1}^{\tau} \log \left((1+2\delta_i)(n-k)c^{\binom{i}{s-1}} \right) + 3(k-1-\tau)(1+\epsilon) \log \log n \\ & < \log n^{\tau+1} + (\log c) \sum_{i=s-1}^{\tau} \binom{i}{s-1} \\ & \quad + \sum_{i=s-1}^{\tau} \log(1+2\delta_i) + 3(k-1-\tau)(1+\epsilon) \log \log n \\ & \leq \log \left(n^{\tau+1} c^{\binom{\tau+1}{s}} \right) + 3(k-1-\tau)(1+\epsilon) \log \log n + 2 \sum_{i=s-1}^{\tau} \delta_i \\ & \leq \log \left(n^{\tau+1} c^{\binom{\tau+1}{s}} \right) + 3(k-1-\tau)(1+\epsilon) \log \log n + o(1) \end{aligned}$$

The last inequality holds since $\tau = O(1)$ and since $\delta_i \lesssim (\log n)^{\frac{1}{2}+\frac{\epsilon}{2}} n^{-\frac{1}{2}+\frac{1}{2}\alpha\binom{\tau}{s-1}} = o(1)$ for all $i \leq \tau$ because of the definition that $\alpha\binom{\tau}{s-1} < 1$. In summary, we have shown that for sufficiently large n

$$\begin{aligned} & \mathbb{P} \left[Z_v \geq 2n^{\tau+1} c^{\binom{\tau+1}{s}} (\log n)^{3(k-1-\tau)(1+\epsilon)} \mid \{v_1, v_2, \dots, v_k\} \in \text{cl}_k(G) \right] \\ & \leq k \exp \left(-\frac{1}{3}(\log n)^{1+\epsilon} \right) = n^{-\omega(1)} \end{aligned}$$

for any k -tuple of vertices $v = (v_1, v_2, \dots, v_k)$. Since $\mathbb{P}[\{v_1, v_2, \dots, v_k\} \in \text{cl}_k(G)] = c^{\binom{k}{s}}$, we have that $\mathbb{P}[A_v] \leq c^{\binom{k}{s}} n^{-\omega(1)} = n^{-\omega(1)}$ for each k -tuple v . Now consider the event

$$B = \left\{ Z_v < 2n^{\tau+1} c^{\binom{\tau+1}{s}} (\log n)^{3(k-1-\tau)(1+\epsilon)} \text{ for all } v \right. \\ \left. \text{such that } \{v_1, v_2, \dots, v_k\} \in \text{cl}_k(G) \right\}$$

Note that $\overline{B} = \bigcup_{k\text{-tuples } v} A_v$ and a union bound implies that $\mathbb{P}[B] \geq 1 - \sum_v \mathbb{P}[A_v] \geq 1 - n^k \cdot n^{-\omega(1)} = 1 - n^{-\omega(1)}$ since there are fewer than n^k k -tuples v .

We now show that as long as B holds over the random choice of G , then the algorithm GREEDY-RANDOM-SAMPLING terminates with $S = \text{cl}_k(G)$ with probability $1 - n^{-\omega(1)}$ over the random bits of GREEDY-RANDOM-SAMPLING, which completes the proof of the lemma in the case $k > \tau + 1$. In the next part of the argument, we consider G conditioned on the event B . Fix some ordering $v = (v_1, v_2, \dots, v_k)$ of some k -clique $\mathcal{C} = \{v_1, v_2, \dots, v_k\}$ in G . Recall that in any one of the T iterations of GREEDY-RANDOM-SAMPLING, the probability that the k vertices v_1, v_2, \dots, v_k are chosen in that order is exactly $1/Z_v$. Since the T iterations of GREEDY-RANDOM-SAMPLING are independent, we have that

$$\mathbb{P}[v \text{ is never chosen in a round}] = \left(1 - \frac{1}{Z_v}\right)^T \leq \exp\left(-\frac{T}{Z_v}\right) = n^{-\omega(1)}$$

since T is chosen so that $T \geq Z_v (\log n)^{3(1+\epsilon)}$ for all k -tuples v , given the event B . Since there are at most n^k possible v , a union bound implies that every such v is chosen in a round of GREEDY-RANDOM-SAMPLING with probability at least $1 - n^k \cdot n^{-\omega(1)} = 1 - n^{-\omega(1)}$ over the random bits of the algorithm. In this case, $S = \text{cl}_k(G)$ after the T rounds of GREEDY-RANDOM-SAMPLING. This completes the proof of the theorem in the case $k \geq \tau + 1$.

We now handle the case $k < \tau + 1$ through a nearly identical argument. Define κ_i as in the previous case and set $\delta_i = \sqrt{\kappa_i}$ for all $i \in \{s-1, s, \dots, k-1\}$. By the same argument, for each k -tuple v we have with probability $1 - n^{-\omega(1)}$ over the choice of G that

$$\begin{aligned} \log Z_v &< \log n^{s-1} + \sum_{i=s-1}^{k-1} \log \left((1 + 2\delta_i)(n-k)c^{\binom{i}{s-1}} \right) \\ &< \log n^k + (\log c) \sum_{i=s-1}^{k-1} \binom{i}{s-1} + 2 \sum_{i=s-1}^{k-1} \delta_i \\ &= \log \left(n^k c^{\binom{k}{s}} \right) + o(1) \end{aligned}$$

where again $\delta_i \lesssim (\log n)^{\frac{1}{2} + \frac{\epsilon}{2}} n^{-\frac{1}{2} + \frac{1}{2}\alpha \binom{\tau}{s-1}} = o(1)$ for all $i \leq k-1 < \tau$. Define the event

$$B' = \left\{ Z_v < 2n^k c^{\binom{k}{s}} (\log n)^{1+\epsilon} \text{ for all } v \text{ such that } \{v_1, v_2, \dots, v_k\} \in \text{cl}_k(G) \right\}$$

Note that T is such that $T \geq Z_v (\log n)^{1+\epsilon}$ for all v if B' holds. Now repeating the rest of the argument from the $k \geq \tau + 1$ case shows that $\mathbb{P}[B'] \geq 1 - n^{-\omega(1)}$ and that GREEDY-RANDOM-SAMPLING terminates with $S = \text{cl}_k(G)$ with probability $1 - n^{-\omega(1)}$ over its random bits if G is such that B' holds. This completes the proof of the theorem. \square