Electronic Colloquium on Computational Complexity, Report No. 100 (2022)

Streaming complexity of CSPs with randomly ordered constraints

Raghuvansh R. Saxena* Noah Singer† Madhu Sudan‡ Santhoshini Velusamy§

Abstract

We initiate a study of the streaming complexity of constraint satisfaction problems (CSPs) when the constraints arrive in a random order. We show that there exists a CSP, namely Max-DICUT, for which random ordering makes a provable difference. Whereas a $4/9 \approx 0.445$ approximation of DICUT requires $\Omega(\sqrt{n})$ space with adversarial ordering, we show that with random ordering of constraints there exists a 0.48-approximation algorithm that only needs $O(\log n)$ space. We also give new algorithms for Max-DICUT in variants of the adversarial ordering setting. Specifically, we give a two-pass $O(\log n)$ space 0.48-approximation algorithm for general graphs and a single-pass $\widetilde{O}(\sqrt{n})$ space 0.48-approximation algorithm for bounded degree graphs.

On the negative side, we prove that CSPs where the satisfying assignments of the constraints support a one-wise independent distribution require $\Omega(\sqrt{n})$-space for any non-trivial approximation, even when the constraints are randomly ordered. This was previously known only for adversarially ordered constraints. Extending the results to randomly ordered constraints requires switching the hard instances from a union of random matchings to simple Erdös-Renyi random (hyper)graphs and extending tools that can perform Fourier analysis on such instances.

The only CSP to have been considered previously with random ordering is Max-CUT where the ordering is not known to change the approximability. Specifically it is known to be as hard to approximate with random ordering as with adversarial ordering, for $o(\sqrt{n})$ space algorithms. Our results show a richer variety of possibilities and motivate further study of CSPs with randomly ordered constraints.

*Microsoft Research. Email: raghuvansh.saxena@gmail.com

†Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA and Harvard College, Harvard University, Cambridge, MA, USA. Supported by an NSF Graduate Research Fellowship (Award DGE2140739). Email: ngsinger@andrew.cmu.edu.

‡School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Supported in part by a Simons Investigator Award and NSF Awards CCF 1715187 and CCF 2152413. Email: madhu@cs.harvard.edu.

§School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Supported in part by a Google Ph.D. Fellowship, a Simons Investigator Award to Madhu Sudan, and NSF Awards CCF 1715187 and CCF 2152413. Email: svelusamy@g.harvard.edu.

ISSN 1433-8092

Contents

1 Introduction

In this paper we consider the streaming complexity of solving constraint satisfaction problems (CSPs) approximately with randomly ordered constraints. We introduce these terms below before turning to the context and our work. Readers familiar with these topics may safely skip to Section 1.1.

Constraint satisfaction problems: A *constraint satisfaction problem (CSP)* is described by a family of predicates $\mathcal{F} \subseteq \{f : \mathbb{Z}_q^k \to \{0,1\}\}$ where $k, q \in \mathbb{N}$ and $\mathbb{Z}_q = \{0, \ldots, q-1\}$. Given such a family $\mathcal{F}$, an instance Ψ of the problem $\mathsf{Max\text{-}CSP}(\mathcal{F})$ on n variables is described by m constraints $C_1, \ldots, C_m$ where for $i \in [m]$, $C_i = (f_i, \mathbf{j}(i) = (j_1(i), \ldots, j_k(i)))$ with $f_i \in \mathcal{F}$ and $\mathbf{j}(i)$ is a sequence of k distinct elements of $[n]$. An assignment to the n variables is given by $\mathbf{a} \in \mathbb{Z}_q^n$. The assignment satisfies C_i if $C_i(\mathbf{a}) := f_i(a_{j_1(i)}, \ldots, a_{j_k(i)}) = 1$ and the value of the assignment on the instance Ψ is given by $\mathsf{val}_\Psi(\mathbf{a}) = \frac{1}{m}\sum_{i\in[m]} C_i(\mathbf{a})$. The goal of $\mathsf{Max\text{-}CSP}(\mathcal{F})$ is to compute $\mathsf{val}_\Psi := \max_{\mathbf{a}\in\mathbb{Z}_q^n}\{\mathsf{val}_\Psi(\mathbf{a})\}$. We will also be interested in approximation algorithms **ALG**: Given $\alpha \in [0,1]$, an α-approximation algorithm to $\mathsf{Max\text{-}CSP}(\mathcal{F})$ is one whose output satisfies $\alpha \cdot \mathsf{val}_\Psi \leq \mathbf{ALG}(\Psi) \leq \mathsf{val}_\Psi$ for every instance Ψ.

Many natural problems can be expressed as CSPs. One example of particular interest to this paper is the $\mathsf{Max\text{-}DICUT}$ problem which is $\mathsf{Max\text{-}CSP}(\{\mathsf{DICUT}\})$ where $\mathsf{DICUT} : \mathbb{Z}_2^2 \to \{0,1\}$ is the predicate $\mathsf{DICUT}(x,y) = (1-x)y$ (with the arithmetic being over $\mathbb{Z}_2$). $\mathsf{Max\text{-}DICUT}$ can equivalently be viewed as a *graph* problem in which variables correspond to vertices and constraints correspond to edges. The goal is then to estimate the size of the highest-value "directed partition" (i.e., $\{0,1\}$-assignment) of the vertices, where the value of a partition is the number of edges from 0-vertices to 1-vertices.

Streaming Algorithms: The class of algorithms we consider (and rule out) are randomized streaming algorithms. Inputs to these algorithms arrive as a stream of elements, in our case a stream of constraints. We consider algorithms that use some bounded amount of space, denoted $s(n)$, to process the stream and produce their output. They may toss their own coins to process the stream. In this work we focus mainly on algorithms whose inputs are *randomly ordered*, i.e., given an instance m on variables with constraint $C_1, \ldots, C_m$ a permutation $\pi : [m] \to [m]$ is chosen uniformly at random and the constraints arrive in the order $C_{\pi(1)}, \ldots, C_{\pi(m)}$. We say that an algorithm is correct if it outputs a correct answer[1] with probability 2/3, where the probability is both over internal coin tosses and over the random arrival order of the input.

1.1 Previous work

The recent years have seen a significant amount of research on the streaming complexity of approximating CSPs with adversarial order of arrival. We refer the reader to Chou, Golovnev, Sudan and Velusamy [CGSV21b] for some of the history. (See also [Sin22] and [Sud22] for some broader surveys.) The summary of this line of research is a dichotomy result for "sketching algorithms" to approximate all CSPs, while getting dichotomies in the more general streaming context for many subclasses. A sketching algorithm is a streaming algorithm that works by compressing substreams into small summaries called sketches with the feature that the sketch of a concatenation of two streams can be obtained from sketches of the two component streams. All known algorithms for CSPs (with proven guarantees on approximation) are sketching algorithms motivating the current

[1] Recall that approximation algorithms are not required to output any one fixed answer. An answer is correct on input Ψ if it lies in the interval $[\alpha \cdot \mathsf{val}_\Psi, \mathsf{val}_\Psi]$.

work. In this work we consider a weakening of the input space, to random ordering of constraints, to explore the possibility of other algorithms, or to rule them out.

Turning to random order in graph streaming problems, [KKS14] gave a polylog(n)-space random-order streaming algorithm for polylog(n)-approximating the maximum matching problem; [KMNT20] improved the exponent in the approximation factor. Another line of works [MMPS17, PS18] explores "generic" ways in which sublinear-time algorithms for graph problems can be transformed into random-ordering streaming algorithms; the latter work establishes provable separations for random-ordering streaming from adversarial-order streaming for problems including estimating the number of connected components and the minimum spanning tree weight. Most relevantly, Kapralov, Khanna, and Sudan [KKS15] showed that the CSP $\mathsf{Max\text{-}CUT} = \mathsf{Max\text{-}CSP}(\{\mathsf{CUT}\})$ where $\mathsf{CUT} : \mathbb{Z}_2^2 \to \{0,1\}$ is defined by $\mathsf{CUT}(a,b) = a+b$ cannot be nontrivially approximated by $o(\sqrt{n})$-space streaming algorithms even in the random-order setting. Thus, other than [KKS15], the previous works on random-order streaming have not studied CSPs; and in particular, none of the previous works suggest that random order of arrival could lead to any algorithmic improvement.

1.2 Main results

In this paper, we present both positive (algorithmic) and negative (hardness) on the usefulness of randomly-ordered streams for approximating CSPs, in comparison to adversarially-ordered streams.

1.2.1 Positive results

Our main positive result asserts that there exists a constraint satisfaction problem where random arrival of constraints provably leads to better approximation with $o(\sqrt{n})$ space.

Theorem 1.1. *There exists a $O(\log n)$-space streaming algorithm that outputs a .483-approximation to the $\mathsf{Max\text{-}DICUT}$ value of directed graphs on n vertices whose edges arrive in a random order.*

This theorem is restated as Theorem 3.1 and proved in Section 3.1.

The result above should be contrasted with the result of Chou, Golovnev and Velusamy [CGV20] who show that for every $\epsilon > 0$, a streaming algorithm that achieves a $(4/9+\epsilon)$-approximation of $\mathsf{Max\text{-}DICUT}$ requires $\Omega(\sqrt{n})$ space when the constraints are ordered *adversarially*. (Note $4/9 = 0.444\ldots$.) Their lower bound holds in the general setting of streaming algorithms, with a matching upper bound using a sketching algorithm. Our algorithm is not a sketching algorithm. This is the only result to our knowledge for a streaming CSP (even with assumptions on arrival order) where a non-sketching algorithm outperforms known sketching algorithms. Indeed the ideas from this algorithm help in contexts other than just the random arrival order and we describe some of these consequences next.

1.2.2 Positive results in other streaming models

The algorithm used to prove Theorem 1.1 can also be modified to the setting of 2-pass algorithms with adversarial order as asserted below.

Theorem 1.2. *There exists a $O(\log n)$-space 2-pass streaming algorithm that outputs a .483-approximation to the $\mathsf{Max\text{-}DICUT}$ value of directed graphs on n vertices under adversarial ordering of edges.*

This theorem is restated as Theorem 3.2 and proved in Section 3.2. The 2-pass algorithm answers an open question in [CGSV21b], perhaps with an unexpected answer.

Finally, we also show how the algorithm can be further modified to get the same approximation to $\mathsf{Max\text{-}DICUT}$ using $\widetilde{O}(\sqrt{n})$ space with a single-pass streaming algorithm in *bounded degree* graphs with adversarial ordering of edges.

Theorem 1.3. *There exists a $\widetilde{O}(\sqrt{n})$-space streaming algorithm that outputs a .483-approximation to the $\mathsf{Max\text{-}DICUT}$ value of bounded-degree directed graphs on n vertices under adversarial ordering of edges.*

Theorem 3.3 states a more detailed relationship between the space needed and the maximum degree of the graph. It implies the theorem above and is proved in Section 3.3. We remark that [CGV20] show that $o(\sqrt{n})$ space algorithms cannot get better than a 4/9-approximation and their proof actually holds even when the input graphs are of bounded degree. Thus Theorem 1.3 establishes the significance of the $\sqrt{n}$-space threshold — again a result that may be somewhat surprising.

1.2.3 Negative results

Returning to our main quest of understanding streaming CSPs in the random-ordering setting and motivated by the algorithmic potential demonstrated by Theorem 1.1 above, we re-explore negative results on streaming to see when they apply also to random arrival ordering. We show that for a broad class of constraint satisfaction problems, the known hardness results on streaming algorithms with adversarial ordering, also extend (with non-trivial analysis) to the case of randomly ordered constraints. We define the class of problems considered and the approximation lower bound achieved below, starting with the latter.

We say that an algorithm is *trivial* if its output is a constant (independent of the input). For a class of constraints $\mathcal{F}$, define $\rho_{\min}(\mathcal{F})$ to be the minimum (strictly, infimum) value val_Ψ over all instance Ψ of $\mathsf{Max\text{-}CSP}(\mathcal{F})$. (A priori, $\rho_{\min}(\mathcal{F})$ might not be computable given $\mathcal{F}$, but [CGSV21b] show it is computable.) Clearly an algorithm that outputs $\rho = \rho_{\min}(\mathcal{F})$ on every instance is a valid, but trivial, ρ-approximation algorithm for $\mathsf{Max\text{-}CSP}(\mathcal{F})$. Motivated by this [CGSV21b] define a problem to be *approximation-resistant* to a class of algorithms if for every $\epsilon > 0$ it does not have a $(\rho + \epsilon)$-approximation within the class. Our next theorem proves a broad class of CSPs to be approximation-resistant to $o(\sqrt{n})$-space single pass streaming algorithms, even with a random ordering of constraints.

We now turn to the class of problems covered by our theorem. We say a predicate $f : \mathbb{Z}_q^k \to \{0,1\}$ *supports one-wise independence* if there exists a distribution $\mathcal{D}$ supported on $f^{-1}(1)$ whose marginals are uniform (i.e., if $\mathbf{a} = (a_1, \ldots, a_k) \sim \mathcal{D}$ then for every i, a_i is distributed uniformly over $\mathbb{Z}_q$). We say a family $\mathcal{F}$ *supports one-wise independence* if every $f \in \mathcal{F}$ supports one-wise independence. We say a family $\mathcal{F}$ *weakly supports one-wise independence* if there exists $\mathcal{F}' \subseteq \mathcal{F}$ supporting one-wise independence with $\rho_{\min}(\mathcal{F}') = \rho_{\min}(\mathcal{F})$. Our theorem below asserts the approximation resistance of $\mathsf{Max\text{-}CSP}(\mathcal{F})$ on randomly ordered instances when $\mathcal{F}$ weakly supports one-wise independence.

Theorem 1.4. *For every $k, q \in \mathbb{N}$ and $\mathcal{F}$ s.t. $\mathcal{F} \subseteq \{f : \mathbb{Z}_q^k \to \{0,1\}\}$ that weakly supports one-wise independence, $\mathsf{Max\text{-}CSP}(\mathcal{F})$ is approximation resistant to $o(\sqrt{n})$-space streaming algorithms in the random order model. That is, for every $\epsilon > 0$, there exists $\tau > 0$ such that every streaming algorithm which $(\rho_{\min}(\mathcal{F}) + \epsilon)$-approximates $\mathsf{Max\text{-}CSP}(\mathcal{F})$ in the random-order model uses at least $\tau\sqrt{n}$ space on instances with n variables.*

We assert that all known families that are known to be approximation-resistant to $o(\sqrt{n})$-space single pass streaming algorithms, even under adversarial ordering, weakly support one-wise independence [CGSV21b]. Such problems include $\mathsf{Max\text{-}CUT}$ (and thus our result subsumes that of [KKS15]), $\mathsf{Max\text{-}}q\mathsf{UniqueGames}$, $\mathsf{Max\text{-}}q\mathsf{Coloring}$, and $\mathsf{Max\text{-}}k\mathsf{OR}$. The question of proving random-ordering approximation-resistance for $\mathsf{Max\text{-}}q\mathsf{UniqueGames}$ was posed by Guruswami and Tao [GT19, §5]. Our result thus strengthens our understanding of approximation resistance for the broadest class of problems where it was previously understood.

1.3 Technical contributions

1.3.1 Positive results

All streaming algorithms for CSPs in previous works [GVV17, CGV20, CGSV21a, CGSV21b, BHP^{+}22] have been based on measuring generalizations of the "total bias" of CSP instances defined originally in [GVV17]; this quantity, even in its richest form from [CGSV21b], is a sum over the variables of some form of "bias", and can be computed using norm-sketching algorithms [Ind06, KNW10, AKO11]. Bias, in turn, roughly measures whether, considering each constraint in which a variable appears independently, the variable prefers to take one value more often than others. In the specific case of $\mathsf{Max\text{-}DICUT}$, the bias $\mathsf{bias}(i)$ of vertex i is simply $\frac{\mathsf{out\text{-}deg}(i)-\mathsf{in\text{-}deg}(i)}{\mathsf{out\text{-}deg}(i)+\mathsf{in\text{-}deg}(i)}$, where $\mathsf{out\text{-}deg}(i)$ and $\mathsf{in\text{-}deg}(i)$ denote the out- and in-degrees of i, respectively. Thus, if $\mathsf{bias}(i) \approx 1$, i has mostly out-edges, so we should assign it to 0, while if $\mathsf{bias}(i) \approx -1$, it has mostly in-edges, so we should assign i to 1.

Thus, for the random-ordering algorithmic result, a key contribution of our work is the first new *algorithmic* paradigm for streaming CSPs since [GVV17]. This should be contrasted with the fact that many works [GT19, KK19, CGV20, CGSV21a, CGSV21b, SSV21, CGS^{+}22] have made significant progress on the hardness front. Instead of estimating the *total* bias of the input graph, we build a *snapshot* of the graph: Specifically we merge vertices with (roughly) the same bias and estimate the fraction of edges that go from vertices of different bias. To get this snapshot information, we look at a representative sample of edges and consider the biases of their endpoints. Here is where we use the random arrival order of edges: We can sample typical edges at the beginning of the stream, and then we measure the bias of their endpoints over the rest of the stream. (So really our algorithm just needs the first few edges to be random, and the rest of the stream could even be ordered adversarially!)

Using this bias information to produce a cut is not trivial, but fortunately for us a previous work of Feige and Jozeph [FJ15] analyzed exactly this question. They studied "oblivious algorithms" for $\mathsf{Max\text{-}DICUT}$, which are algorithms which randomly assign each vertex independently based solely on its bias, and showed the existence of an α_{FJ}-approximation algorithm for some $\alpha_{\mathrm{FJ}} \in (0.483, 0.4899)$. Our theorem follows by appealing to their result. We remark that based on the trivial reduction from $\mathsf{Max\text{-}CUT}$, $\mathsf{Max\text{-}DICUT}$'s approximability for $o(\sqrt{n})$-space algorithms with randomly ordered constraints is at most $1/2$. And while [FJ15] showed that oblivious algorithms cannot do better than 0.4899-approximations, it is quite possible that other quantities that can be easily estimated with random arrival orders (such as the number of copies of $O(1)$-vertex subgraphs, such as paths) could lead to $1/2$-approximation algorithms.

The idea of computing a snapshot of the graph and then using that (via the Feige-Jozeph analysis) to approximate the Dicut value of a graph turns out to work in other streaming settings as well. For instance in the two-pass setting with adversarial ordering of the edges, we can pick a random sample of edges in the first pass and then use the second pass to compute the bias of the endpoints of the edges. This leads to a polylog space streaming 2-pass algorithm achieving the same approximation for $\mathsf{Max\text{-}DICUT}$ even in the adversarial arrival setting. In the case of bounded degree graphs also we are able to compute snapshots with $\widetilde{O}(\sqrt{n})$-space when the edge arrival order is adversarial. While this requires some additional care, to deal with very sparse graphs (with most vertices being isolated), the general plan can be implemented leading to a single-pass $\widetilde{O}(\sqrt{n})$-space algorithm achieving the same approximation for Dicut.

1.3.2 Negative results

Turning to the negative results that form the technical meat of this paper, we comment briefly on where previous works used the adversarial ordering and what we need to do to overcome it. Starting with [KKS15], all hardness results for streaming $\mathsf{Max\text{-}CSP}(\mathcal{F})$ problems have been based on constructing so-called "**YES**" and "**NO**" distributions over instances which have high and low values, respectively (with high probability), and showing that these are indistinguishable by reducing from a one-way communication problem. Designing these distributions is typically a tradeoff between desired properties for the streaming lower bound (e.g., optimizing the value gap between **YES** and **NO** instances) and technical considerations in terms of how to prove the appropriate communication lower bounds (and whether they even hold at all!). The distributions themselves result from a two-fold process: First, sample a random hypergraph, and then treat each hyperedge as a CSP constraint by labeling it with an appropriate predicate $f \in \mathcal{F}$. Indeed, this "labeling" is the only difference between the **YES** and **NO** distributions; typically, in the **NO** distribution the labels are completely random, while in the **YES** distribution they are selected to be consistent with some global assignment.

Now, consider the communication problem in which we split up hypergraph edges and labels among $T = O(1)$ of "players", and the players must distinguish between the **YES** and **NO** cases. At a high level, the technical complexity of such problems is closely connected to the structure of the hypergraphs that the players receive. In particular, it becomes necessary to analyze a counting problem involving $\mathbb{Z}_q$-labelings of edge-vertex incidences with sum constraints at vertices and density constraints on edges (see Eq. (5.1) below for a technical statement). In previous works aside from [KKS15], each player's input hypergraph was a random (partial) *hypermatching*. Crucially, hypermatchings (of any particular size) are unique up to renaming of vertices. While this significantly simplifies the combinatorial analysis, it is not appropriate for proving random-ordering streaming lower bounds. This is because, in the communication-to-streaming reduction, the resultant stream of constraints is the concatenation of constraints contributed by each player; these streams will have the property that in each successive "chunk" of $\approx 1/T$ constraints, no variables are repeated, which is unlikely in a randomly-ordered stream. Thus, it is necessary to draw the players' input hypergraphs from a different distribution. In the case of $\mathsf{Max\text{-}CUT}$, with alphabet size $q = 2$ and arity $k = 2$, Kapralov *et al.* [KKS15] instead worked with general random graphs. Such graphs are no longer unique up to renaming of vertices; there are many different equivalence classes, and each behaves differently in the proof of the lower bound. However, [KKS15] manages this difficulty by showing that (1) cycles are unlikely, and (2) conditioned on cycle-freeness, each equivalence class corresponds to a union of paths with a certain length profile. It turns out that both the $k = 2$ and $q = 2$ assumptions are significantly helpful the analysis of [KKS15]. If $k > 2$, we lose the decomposition into unions of paths, while if $q > 2$, we need to worry about different $\mathbb{Z}_q$-labelings even of the same path, and thus the length of paths comes into play.

Nevertheless, in our work, we manage to generalize to arbitrary $k, q \in \mathbb{N}$ by conducting a careful combinatorial analysis of connected component sizes in random hypergraphs (see Section 6). This allows us to develop streaming hardness results for all CSPs weakly supporting one-wise independence (Theorem 1.4). Indeed, we show that perfectly satisfiable instances (i.e., those with value 1) are indistinguishable from random instances with independent, uniformly random constraints!

2 Preliminaries

For $n > 0$, we use 0^n to denote the all zeros vector of length n and $\mathcal{S}(n)$ to denote the set of all permutations mapping the set $[n]$ to itself. Let Σ be a set, $n \in \mathbb{N}$, and $\pi \in \mathcal{S}(n)$ be a permutation. For $\sigma \in \Sigma^n$ and $i \in [n]$, we use σ_i to denote coordinate i of σ and $\pi(\sigma)$ to denote the vector $\sigma_{\pi(1)}, \sigma_{\pi(2)}, \ldots, \sigma_{\pi(n)}$. For $\sigma \in \Sigma^*$, we use $|\sigma|$ to denote the number of coordinates in σ.

For a set S, we use $\Delta(S)$ to denote the set of all distributions whose support is S. For $k > 0$ and sets $S_1, S_2, \ldots, S_k$, we use $\Delta_{\mathsf{unif}}(S_1, S_2, \ldots, S_k)$ to denote the set of all distributions on the product set $S = S_1 \times S_2 \times \ldots \times S_k$ for which the marginal distribution on the set S_i, for all $i \in [k]$ is uniform. We simply write $\Delta_{\mathsf{unif}}(S)$ if the decomposition into the sets S_i is clear from context.

2.1 Definitions

2.1.1 The Random-Order Streaming Model

Let Σ be an alphabet set. A deterministic streaming algorithm $\mathbf{ALG}$ for Σ-streams is defined by the tuple:

$$\mathbf{ALG} = (S, \mathsf{mdfy}, \mathsf{out}),$$

where: (1) $S = \|\mathbf{ALG}\|$ is the space/memory required by the algorithm $\mathbf{ALG}$. (2) $\mathsf{mdfy} = \Sigma \times \{0,1\}^S \to \{0,1\}^S$ is the function the algorithm uses to update its state upon reading a symbol from the stream. (3) $\mathsf{out} = \{0,1\}^S \to \{0,1\}^S$ is the function the algorithm uses to compute its output from its state at the end of the stream. We shall suppress arguments on the right hand side when they are clear from context. We define a randomized streaming algorithm on Σ-streams to be a distribution over deterministic streaming algorithms. Additionally, the space required by a randomized streaming algorithm is the maximum space required by a deterministic algorithm in its support.

Execution of a streaming algorithm. Let Σ be an alphabet set and $\mathbf{ALG}$ be a (deterministic) algorithm for Σ-streams. For an element $\sigma \in \Sigma^*$ with $m = |\sigma|$, the algorithm $\mathbf{ALG}$ acts on σ in m steps as follows. At the beginning (before step 1), the algorithm is the state $s_0 = 0^S$. Then, for $i \in [m]$, the algorithm reads the symbol σ_i and uses it to update its state by defining $s_i = \mathsf{mdfy}(\sigma_i, s_{i-1})$. Finally, after m steps, the algorithm outputs the value $\mathsf{out}(s_m)$.

Note that all the states of the algorithm and its final output are determined by its input σ. For $i \in [m]$, we write $\mathbf{ALG}(\sigma, i) \in \{0,1\}^S$ to denote the state after step i of the algorithm on input σ. We define $\mathbf{ALG}(\sigma, 0) = 0^S$ for convenience. Finally, we write $\mathbf{ALG}(\sigma) \in \{0,1\}$ to denote the output of the algorithm on input σ.

Computation using streaming algorithms. Let Σ be an alphabet set and $f : \Sigma^* \to \{0,1\}$ be a (possibly partial) function. For $p > 0$, we say that a randomized streaming algorithm $\mathcal{A}$ computes the function f in the random-order streaming model with probability p if for all $\sigma \in \Sigma^*$, we have:

$$\Pr_{\mathbf{ALG} \sim \mathcal{A}, \pi \sim \mathcal{S}(|\sigma|)} (\mathbf{ALG}(\pi(\sigma)) = f(\sigma)) \geq p.$$

Distinguishing using streaming algorithms. Let Σ be an alphabet set and $(\mathcal{Y}, \mathcal{N})$ be a pair of distributions over Σ^*. For $\delta \geq 0$, we say that a deterministic streaming algorithm $\mathbf{ALG}$ distinguishes between $\mathcal{Y}$ and $\mathcal{N}$ with advantage δ in the random-order streaming model if:

$$\left| \Pr_{\sigma \sim \mathcal{Y}, \pi \sim \mathcal{S}(|\sigma|)} (\mathbf{ALG}(\pi(\sigma)) = 1) - \Pr_{\sigma \sim \mathcal{N}, \pi \sim \mathcal{S}(|\sigma|)} (\mathbf{ALG}(\pi(\sigma)) = 1) \right| \geq \delta.$$

We say that **ALG** distinguishes between $\mathcal{Y}$ and $\mathcal{N}$ with advantage δ in the worst case streaming model if:

$$\left| \Pr_{\sigma \sim \mathcal{Y}}(\mathbf{ALG}(\sigma) = 1) - \Pr_{\sigma \sim \mathcal{N}}(\mathbf{ALG}(\sigma) = 1) \right| \geq \delta.$$

We may sometimes refer to a pair $(\mathcal{Y}, \mathcal{N})$ of distributions as a streaming problem and say that "**ALG** solves the $(\mathcal{Y}, \mathcal{N})$-problem" instead of saying that "**ALG** distinguishes between $\mathcal{Y}$ and $\mathcal{N}$". We also note that the two notions of distinguishability are equivalent if the distributions $\mathcal{Y}$ and $\mathcal{N}$ are sufficiently symmetric.

Lemma 2.1. *Let Σ be an alphabet set and $\mathcal{D}$ be a distribution over Σ^* such that for all $\sigma \in \Sigma^*$ and $\pi \sim \mathcal{S}(|\sigma|)$, we have $\mathcal{D}(\sigma) = \mathcal{D}(\pi(\sigma))$. Then, for all $\tau \in \Sigma^*$, we have:*

$$\Pr_{\sigma \sim \mathcal{D}}(\sigma = \tau) = \Pr_{\sigma \sim \mathcal{D}, \pi \sim \mathcal{S}(|\sigma|)}(\pi(\sigma) = \tau).$$

Proof. Let $\mathcal{D}'$ be the distribution on $\mathbb{N}$ obtained by sampling σ from $\mathcal{D}$ and outputting $|\sigma|$. We can view the process of sampling σ from $\mathcal{D}$ and then sampling π from $\mathcal{S}(|\sigma|)$ as the process of first sampling an integer $m \geq 0$ from $\mathcal{D}'$, then sampling a permutation π from $\mathcal{S}(m)$ and finally, a string σ from $\mathcal{D}$ conditioned on the fact that $|\sigma| = m$. Moreover, as $\pi(\sigma) = \tau$ can happen only if $m = |\tau|$, we get (using $m = |\tau|$):

$$\begin{aligned}
\Pr_{\sigma \sim \mathcal{D}, \pi \sim \mathcal{S}(|\sigma|)}(\pi(\sigma) = \tau) &= \mathcal{D}'(m) \cdot \Pr_{\pi \sim \mathcal{S}(m), \sigma \sim \mathcal{D}|_{|\sigma| = m}}(\pi(\sigma) = \tau) \\
&= \mathcal{D}'(m) \cdot \frac{1}{m!} \cdot \sum_{\pi \in \mathcal{S}(m)} \Pr_{\sigma \sim \mathcal{D}|_{|\sigma| = m}}(\pi(\sigma) = \tau) \\
&= \mathcal{D}'(m) \cdot \frac{1}{m!} \cdot \sum_{\pi \in \mathcal{S}(m)} \Pr_{\sigma \sim \mathcal{D}|_{|\sigma| = m}}\left(\sigma = \pi^{-1}(\tau)\right) \\
&= \mathcal{D}'(m) \cdot \frac{1}{m!} \cdot \sum_{\pi \in \mathcal{S}(m)} \Pr_{\sigma \sim \mathcal{D}|_{|\sigma| = m}}(\sigma = \tau) \\
&= \mathcal{D}'(m) \cdot \Pr_{\sigma \sim \mathcal{D}|_{|\sigma| = m}}(\sigma = \tau) \\
&= \Pr_{\sigma \sim \mathcal{D}}(\sigma = \tau).
\end{aligned}$$

□

Corollary 2.2 (Random order to worst-case)**.** *Let Σ be an alphabet set and $(\mathcal{Y}, \mathcal{N})$ be a pair of distributions over Σ^* such that for all $\sigma \in \Sigma^*$ and $\pi \sim \mathcal{S}(|\sigma|)$, we have $\mathcal{Y}(\sigma) = \mathcal{Y}(\pi(\sigma))$ and $\mathcal{N}(\sigma) = \mathcal{N}(\pi(\sigma))$. Then, for all $\delta \geq 0$ and any deterministic streaming algorithm* **ALG** *from Σ-streams, we have that* **ALG** *distinguishes between $\mathcal{Y}$ and $\mathcal{N}$ with advantage δ in the random-order streaming model if and only if* **ALG** *distinguishes between $\mathcal{Y}$ and $\mathcal{N}$ with advantage δ in the worst case streaming model.*

We shall also need the following connection between computation and distinguishing using streaming algorithms.

Fact 2.3. *Let Σ be an alphabet set, $f : \Sigma^* \to \{0, 1\}$ be a partial function, and $p > 0$. If there exists a randomized streaming algorithm $\mathcal{A}$ that computes the function f in the random-order streaming model with probability p, then for all distributions $\mathcal{Y}$ and $\mathcal{N}$ supported on $f^{-1}(1)$ and $f^{-1}(0)$ respectively, we have a deterministic streaming algorithm* **ALG**, $\|\mathbf{ALG}\| \leq \|\mathcal{A}\|$ *such that* **ALG** *distinguishes between $\mathcal{Y}$ and $\mathcal{N}$ with advantage $2 \cdot \left(p - \frac{1}{2}\right)$ in the random-order streaming model.*

2.2 The $\mathsf{Max\text{-}CSP}(\cdot)$ Problem

Throughout this subsection, we let $q, k \in \mathbb{N}$ and $\mathcal{F}$ be a non-empty set of functions mapping $\mathbb{Z}_q^k \to \{0,1\}$. Let $n \geq k \in \mathbb{N}$. An instance Ψ of $\mathsf{Max\text{-}CSP}_n(\mathcal{F})$ is given by a sequence:

$$\Psi = (f_i, M_i)_{i>0} \in \left(\mathcal{F} \times \{0,1\}^{k\times n}\right)^*,$$

where, for all $i \in [|\Psi|]$, the matrix M_i is partial permutation matrix, *i.e.*, a matrix with $0,1$ entries and exactly one 1 in each row and at most one 1 in every column. Let $m = |\Psi|$. Intuitively, Ψ can be seen as a sequence of m constraints, with constraint $i \in [m]$ requiring that the function f_i when applied to the k variables indicated by M_i evaluates to 1. Here, for $j \in [k]$ the j^{th} variable indicated by M_i is the unique column that has the 1 in row j of M_i.

Value of Ψ. For an assignment $\mathbf{x} \in \mathbb{Z}_q^n$ of the n variables, the fraction of satisfied constraints is given by:

$$\mathsf{val}_\Psi(\mathbf{x}) = \frac{1}{L} \cdot \sum_{i\in[L]} f_i(M_i\mathbf{x}). \tag{2.4}$$

We define the value of Ψ to be the largest fraction of the constraints that can be satisfied by an assignment. Thus,

$$\mathsf{val}_\Psi = \max_{\mathbf{x}\in\mathbb{Z}_q^n} \mathsf{val}_\Psi(\mathbf{x}). \tag{2.5}$$

The function $\rho_{\min}(\cdot)$. The minimum value of an instance of $\mathsf{Max\text{-}CSP}(\mathcal{F})$ is given by:

$$\rho_{\min}(\mathcal{F}) = \inf_{\substack{n\in\mathbb{N} \\ \Psi \text{ instance of } \mathsf{Max\text{-}CSP}_n(\mathcal{F})}} \mathsf{val}_\Psi. \tag{2.6}$$

The following lemma, taken from [CGSV21b], gives an equivalent formulation of the function $\rho(\cdot)$ above that is slightly more amenable to analysis.

Lemma 2.7 ([CGSV21b], Proposition 2.12)**.** *Let $q, k \in \mathbb{N}$ be given and $\mathcal{F}$ be a non-empty set of functions mapping $\mathbb{Z}_q^k \to \{0,1\}$. It holds that:*

$$\rho_{\min}(\mathcal{F}) = \min_{D\in\Delta(\mathcal{F})} \max_{D'\in\Delta(\mathbb{Z}_q)} \mathop{\mathbb{E}}_{\substack{f\sim D \\ \mathbf{a}\sim D'^k}}[f(\mathbf{a})].$$

Approximation resistance. Let $n \geq k \in \mathbb{N}$ and $\epsilon > 0$. Define the partial function $\mathsf{aprx}_{\mathcal{F},n,\epsilon}$ on instances Ψ of $\mathsf{Max\text{-}CSP}_n(\mathcal{F})$ to be 1 if $\mathsf{val}_\Psi = 1$ and 0 if $\mathsf{val}_\Psi \leq \rho_{\min}(\mathcal{F}) + \epsilon$. We are now ready to define the notion of approximation resistance.

Definition 2.8 (Approximation resistance)**.** *Let $q, k \in \mathbb{N}$ be given and $\mathcal{F}$ be a non-empty set of functions mapping $\mathbb{Z}_q^k \to \{0,1\}$. Let $s : \mathbb{N} \to \mathbb{R}$ be a monotone function. We say that $\mathsf{Max\text{-}CSP}(\mathcal{F})$ is approximation resistant to $o(s)$ space in the random order streaming model if for all $\epsilon > 0$ and $p > \frac{1}{2}$, there exists $\tau > 0$ such that for all $n \in \mathbb{N}$ and all randomized streaming algorithms $\mathcal{A}$ that compute $\mathsf{aprx}_{\mathcal{F},n,\epsilon}$ in the random-order streaming model with probability p, we have $\|\mathcal{A}\| \geq \tau \cdot s(n)$.*

One-wise independence. We say that a function $f : \mathbb{Z}_q^k \to \{0,1\}$ *supports one-wise independence* if there exists a distribution $D \in \Delta_{\mathsf{unif}}(\mathbb{Z}_q^k)$ that is supported on $f^{-1}(1)$. Similarly, we say that a family $\mathcal{F}$ of functions *(strongly) supports one-wise independence* if all functions in the family support one-wise independence. Finally, we say that a family $\mathcal{F}$ *weakly supports one-wise independence* if there exists a non-empty sub-family $\mathcal{F}' \subseteq \mathcal{F}$ that strongly supports one-wise independence and satisfies $\rho_{\min}(\mathcal{F}) = \rho_{\min}(\mathcal{F}')$.

2.3 One Way Communication Protocols

Let $\mathcal{X}^A$ and $\mathcal{X}^B$ be two sets. We will treat these sets as the inputs sets for Alice and Bob respectively. We now define one-way communication protocols between Alice and Bob, where the inputs of the parties come from the sets $\mathcal{X}^A$ and $\mathcal{X}^B$ respectively, and Alice sends a single message to Bob. We start by defining deterministic protocols. Such a protocol is defined by a tuple:

$$\Pi = (L, \mathsf{msg}, \mathsf{out}),$$

where: (1) $L = \|\Pi\|$ is the length of the protocol Π. (2) $\mathsf{msg} : \mathcal{X}^A \to \{0,1\}^L$ is the function Alice uses to compute her message. (3) $\mathsf{out} : \mathcal{X}^B \times \{0,1\}^L \to \{0,1\}$ is the function Bob uses to compute his output. We shall suppress the arguments on the right hand side when they are clear from context. We define a randomized protocol to be a distribution over deterministic protocols with the same input sets. The length of a randomized protocol is defined to be the maximum length of the deterministic protocols in its support.

Execution of a protocol. Let $\mathcal{X}^A$ and $\mathcal{X}^B$ be sets and Π be a deterministic protocol with inputs sets $\mathcal{X}^A$ and $\mathcal{X}^B$. For $x^A \in \mathcal{X}^A$ and $x^B \in \mathcal{X}^B$, we define the output $\Pi(x^A, x^B) \in \{0,1\}$ of the protocol Π on inputs x^A and x^B as:

$$\Pi(x^A, x^B) = \mathsf{out}\big(x^B, \mathsf{msg}\big(x^A\big)\big).$$

This is because, when the inputs are x^A and x^B, the string $\mathsf{msg}\big(x^A\big)$ is the message sent by Alice to Bob, and therefore, $\mathsf{out}\big(x^B, \mathsf{msg}\big(x^A\big)\big)$ is the output computed by Bob upon receiving this message.

One-way communication problems. We define a communication problem to be a pair of distributions[2] $(\mathcal{Y}, \mathcal{N})$ on the same product set $\mathcal{X}^A \times \mathcal{X}^B$. A protocol for the $(\mathcal{Y}, \mathcal{N})$-problem is a one way communication protocol where Alice's input comes from the set $\mathcal{X}^A$ and Bob's input comes from the set $\mathcal{X}^B$. Let $(\mathcal{Y}, \mathcal{N})$ be a communication problem and $\mathsf{\Pi}$ be a randomized communication protocol for the $(\mathcal{Y}, \mathcal{N})$-problem. For $\delta \geq 0$, we say that $\mathsf{\Pi}$ solves the $(\mathcal{Y}, \mathcal{N})$-problem with advantage δ if we have:

$$\left| \Pr_{\substack{(x^A, x^B) \sim \mathcal{Y} \\ \Pi \sim \mathsf{\Pi}}} \big(\Pi(x^A, x^B) = 1\big) - \Pr_{\substack{(x^A, x^B) \sim \mathcal{N} \\ \Pi \sim \mathsf{\Pi}}} \big(\Pi(x^A, x^B) = 1\big) \right| \geq \delta.$$

2.4 Analytical tools

2.4.1 Random variables

Lemma 2.9 (Triangle inequality)**.** *Let $\mathcal{Y}, \mathcal{N}, \mathcal{Z} \in \Delta(\Omega)$. Then*

$$\|\mathcal{Y} - \mathcal{N}\|_{\mathrm{tv}} \geq \|\mathcal{Y} - \mathcal{Z}\|_{\mathrm{tv}} - \|\mathcal{Z} - \mathcal{N}\|_{\mathrm{tv}}.$$

Lemma 2.10 (Data processing inequality)**.** *Let Y, N be random variables with sample space Ω, and let Z be a random variable with sample space Ω' which is independent of Y and N. If $g : \Omega \times \Omega' \to \Omega''$ is any function, then*

$$\|Y - N\|_{\mathrm{tv}} \geq \|g(Y, Z) - g(N, Z)\|_{\mathrm{tv}}.$$

We will use the following concentration inequality from [KK19].

[2] Note that this matches our notation for distributional streaming problems. Nonetheless, the difference will be clear from context.

Lemma 2.11 ([KK19, Lemma 2.5])**.** *Let $X = \sum_{i=1}^n X_i$, where X_i are Bernoulli $\{0,1\}$-valued random variables satisfying, for every $k \in [n]$, $\mathbb{E}[X_k \mid X_1, \ldots, X_{k-1}] \le p$ for some $p \in (0,1)$. Let $\mu = np$. Then for all $\Delta > 0$,*

$$\Pr[X \ge \mu + \Delta] \le \exp\left(-\frac{\Delta^2}{2(\mu+\Delta)}\right) .$$

We also need the following concentration inequality that we prove using Lemma 2.11.

Lemma 2.12. *Let $X = \sum_{i=1}^n X_i$, where X_i are Bernoulli $\{0,1\}$-valued random variables satisfying, for every $k \in [n]$, $\mathbb{E}[X_k \mid X_1, \ldots, X_{k-1}] \ge p$ for some $p \in (0,1)$. Let $\mu = np$. Then for all $\Delta > 0$,*

$$\Pr[X \le \mu - \Delta] \le \exp\left(-\frac{\Delta^2}{2(n-(\mu-\Delta))}\right) .$$

Proof. Follows immediately from Lemma 2.11 on the random variables $Y_i = 1 - X_i$, $q = 1 - p$, and $\nu = nq$ (since $X \le \mu - \Delta$ is equivalent to $Y \ge \nu + \Delta$). □

2.4.2 Fourier analysis over $\mathbb{Z}_q$

Let $q \ge 2 \in \mathbb{N}$, and let $\omega \overset{\text{def}}{=} e^{2\pi i/q}$ denote a (fixed primitive) q-th root of unity. Here, we summarize relevant aspects of Fourier analysis over $\mathbb{Z}_q^n$; see e.g. [O'D14, §8] for details.[3] Given a function $f : \mathbb{Z}_q^n \to \mathbb{C}$ and $\mathbf{s} \in \mathbb{Z}_q^n$, we define the *Fourier coefficient*

$$\widehat{f}(\mathbf{s}) \overset{\text{def}}{=} \sum_{\mathbf{x} \in \mathbb{Z}_q^n} \omega^{-\mathbf{s}\cdot\mathbf{x}} f(\mathbf{x})$$

where $\cdot$ denotes the inner product over $\mathbb{Z}_q$. For $p \in (0, \infty)$, we define f's *p-norm*

$$\|f\|_p \overset{\text{def}}{=} \left(\sum_{\mathbf{x} \in \mathbb{Z}_q^n} |f(\mathbf{x})|^p\right)^{1/p} .$$

We also define f's 0-norm

$$\|f\|_0 \overset{\text{def}}{=} \sum_{\mathbf{x} \in \mathbb{Z}_q^n} \mathbb{1}_{f(\mathbf{x}) \neq 0}$$

(a.k.a. the size of its support and the Hamming weight of its "truth table"). Also, for $\ell \in \{0\} \cup [n]$, we define the *level-ℓ Fourier (2-)weight* as

$$\mathsf{W}^{\ell}[f] \overset{\text{def}}{=} \sum_{\mathbf{s} \in \mathbb{Z}_q^n : \|\mathbf{s}\|_0 = \ell} |\widehat{f}(\mathbf{s})|^2.$$

These weights are closely connected to f's 2-norm:

Proposition 2.13 (Parseval's identity)**.** *For every $q, n \in \mathbb{N}$ and $f : \mathbb{Z}_q^n \to \mathbb{C}$, we have*

$$\|f\|_2^2 = q^n \sum_{\ell=0}^{n} \mathsf{W}^{\ell}[f].$$

[3] [O'D14] uses a different normalization for norms and inner products, essentially because it considers expectations instead of sums over inputs.

Moreover, let $\mathbb{D} \overset{\text{def}}{=} \{w \in \mathbb{C} : |w| \leq 1\}$ denote the (closed) unit disk in the complex plane. The following lemma bounding the low-level Fourier weights for functions mapping into $\mathbb{D}$ is derived from hypercontractivity theorems in [CGS+22]:

Lemma 2.14 ([CGS+22, Lemma 2.11])**.** *There exists $\zeta > 0$ such that the following holds. Let $q \geq 2, n \in \mathbb{N}$ and consider any function $f : \mathbb{Z}_q^n \to \mathbb{D}$. If for $c \in \mathbb{N}$, $\|f\|_0 \geq q^{n-c}$, then for every $\ell \in \{1, \ldots, 4c\}$, we have*

$$\frac{q^{2n}}{\|f\|_0^2} \mathsf{W}^\ell[f] \leq \left(\frac{\zeta c}{\ell}\right)^\ell.$$

Lemma 2.15. *Let $\mathcal{U} = \mathcal{U}(\mathbb{Z}_q^m)$. Then for all $\mathcal{Z} \in \Delta(\mathbb{Z}_q^m)$,*

$$\|\mathcal{Z} - \mathcal{U}\|_{\text{tv}}^2 \leq q^{2m} \sum_{\ell=1}^{m} \mathsf{W}^\ell[\mathcal{Z}].$$

Proof. We have

$$\|\mathcal{Z} - \mathcal{U}\|_{\text{tv}} = \frac{q^m}{2} \|\mathcal{Z} - \mathcal{U}\|_1.$$

Thus by Cauchy-Schwartz,

$$\|\mathcal{Z} - \mathcal{U}\|_{\text{tv}}^2 \leq q^{2m} \|\mathcal{Z} - \mathcal{U}\|_2^2.$$

Finally, we apply Parseval and observe that $\widehat{\mathcal{Z}}(\mathbf{0}) = \widehat{\mathcal{U}}(\mathbf{0}) = 1$ while for all $\mathbf{s} \neq \mathbf{0}$, $\widehat{\mathcal{U}}(\mathbf{s}) = 0$ by symmetry. □

2.4.3 Hypergraphs

Let $2 \leq k, n \in \mathbb{N}$. A *$k$-hyperedge* on $[n]$ is a k-tuple $\mathbf{e} = (e_1, \ldots, e_k) \in [n]^k$ of distinct indices, and a *k-hypergraph* (a.k.a. "k-uniform hypergraph") G on $[n]$ is a sequence $(\mathbf{e}(1), \ldots, \mathbf{e}(m))$ of (not necessarily distinct) k-hyperedges. For $\alpha \in (0,1), n \in \mathbb{N}$, let $\mathcal{G}_{k,\alpha}(n)$ denote the uniform distribution over k-hypergraphs on $[n]$ with αn hyperedges.

Given a graph G with m edges $\mathbf{e}(1), \ldots, \mathbf{e}(m)$, we associate each hyperedge $\mathbf{e}(i)$ with a partial permutation matrix $M_i \in \{0,1\}^{k \times n}$, such that for each $j \in [k]$, row j has a 1 only in position $e(j)_i$. We associate to G an *adjacency matrix* $M \in \{0,1\}^{km \times n}$ by stacking together $M_1, \ldots, M_m$. Since they encode the same information, we will often treat adjacency matrices and k-hypergraphs as interchangeable (and speak of drawing a matrix M from $\mathcal{G}_{k,\alpha}(n)$.

For a k-hypergraph G on vertex-set $[n]$ with hyperedges $(\mathbf{e}(1), \ldots, \mathbf{e}(m))$, we define the *vertex-hyperedge incidence graph B_G*, which is a bipartite graph (i.e., 2-hypergraph) defined as follows: The left vertex-set is $[n]$, the right vertex-set is $[m]$, and there is an edge between $i \in [n]$ and $j \in [m]$ iff $i \in \mathbf{e}(j)$.

2.5 Reservoir sampling in the streaming setting

Reservoir sampling is a term used to refer to a family of randomized streaming algorithms that are used to sample uniform k random elements from the stream without prior knowledge on the length of the stream. The simplest algorithm, known as *Algorithm R*, was created by Alan Waterman in 1975. The algorithm runs in $O(k)$ space and works as follows: it maintains a "reservoir" of size k. Initially, the first k elements in the stream are stored in the reservoir. For $i > k$, when the i-th element of the stream, denoted by a_i, arrives, the algorithm generates a random number j between 1 and i, and if $j \leq k$, it replaces the j-th element in the reservoir with a_i. It is not hard to show that if m elements have arrived in the stream so far, then the probability of any one of them being in the reservoir is exactly k/m (see [Vit85] for more details).

2.6 k-wise independent hash family

A k-wise independent hash family is a family of hash functions $\mathsf{H}(n,m) = \{h : [n] \to [m]\}$ that satisfies the following properties: For a hash function h drawn uniformly at random from H,

- for every $x \in [n]$ and $a \in [m]$, $\Pr[h(x) = a] = \frac{1}{m}$, and
- for every distinct $x_1, \dots, x_k \in [n]$, $h(x_1), \dots, h(x_k)$ are independent random variables.

We give a construction of $\mathsf{H}(n,m)$ for $m = 2^\ell$ for some $\ell \in \mathbb{N}$. Let $r \in \mathbb{N}$ be the smallest integer such that $2^r \geq \max\{n, m\}$. Let $\mathbb{F}$ be a field of size 2^r. Consider the hash family $\mathsf{H} = \{h_{a_1,\dots,a_k} : a_i \in \mathbb{F}\}$, where $h_{a_1,\dots,a_k}$ is the hash function defined as follows. Let $h'_{a_1,\dots,a_k} : \mathbb{F} \to \mathbb{F}$ be the function defined as $h'(x) = \sum_{i=1}^k a_i x^{i-1}$. Let $f : [n] \to \mathbb{F}$ be any injective function and $g : \mathbb{F} \to [m]$ be a function such that for every $a \in [m]$, $|g^{-1}(a)| = 2^{r-\ell}$. We define $h_{a_1,\dots,a_k} = g \circ h'_{a_1,\dots,a_k} \circ f$.

To show that H is a k-wise independent family, observe that it suffices to show that $\mathsf{H}' = \{h'_{a_1,\dots,a_k} : a_i \in \mathbb{F}\}$ is a k-wise independent hash family. Indeed, for $x \in [n]$ and $a \in [m]$,

$$\Pr_{a_1,\dots,a_k \in \mathbb{F}}[h_{a_1,\dots,a_k}(x) = a] = \Pr_{a_1,\dots,a_k \in \mathbb{F}}[h'_{a_1,\dots,a_k}(x) \in g^{-1}(a)] = 2^{r-\ell} \cdot 2^{-r} = 2^{-\ell}\,.$$

The independence of $h(x_1), \dots, h(x_k)$ follows from the independence of $h'(x_1), \dots, h'(x_k)$. It is a standard exercise to show that H' is a k-wise independent family (see [Vad12] for instance).

3 Algorithms for Max-DICUT

We review the definition of Max-DICUT as an optimization problem on unweighted directed graphs. Let $\mathcal{G} = (V, E)$ be an unweighted directed (multi)graph. $\mathcal{G}$'s Max-DICUT value, denoted $\mathsf{val}_\mathcal{G}$, is defined as the size of the largest directed cut in the graph. Formally,

$$\mathsf{val}_\mathcal{G} \stackrel{\text{def}}{=} \max_{L,R : V = L \sqcup R} |E_{L \to R}|\,,$$

where $E_{L\to R} = \{(i,j) \in E : i \in L \text{ and } j \in R\}$. In this section, we prove the following three theorems for a constant $\alpha_{\mathrm{FJ}} \geq 0.483$:

Theorem 3.1 (Random-ordering algorithm)**.** *Let $\epsilon > 0$ and $c > 0$ be constants. There exists an $O(\log n)$-space single-pass streaming algorithm* **ALG** *such that for every directed graph $\mathcal{G} = (V, E)$ with $|V| = n$ and $|E| \leq n^c$, the following holds: On input the edges of $\mathcal{G}$ in a uniformly random order,* **ALG** *outputs an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation to $\mathsf{val}_\mathcal{G}$ with probability at least $2/3$.*

Theorem 3.2 (Two-pass algorithm)**.** *Let $\epsilon > 0$ and $c > 0$ be constants. There exists an $O(\log n)$-space* two-pass *streaming algorithm* **ALG** *such that for every directed graph $\mathcal{G} = (V, E)$ with $|V| = n$ and $|E| \leq n^c$, the following holds: On input the edges of $\mathcal{G}$ in adversarial order,* **ALG** *outputs an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation to $\mathsf{val}_\mathcal{G}$ with probability at least $2/3$.*

Theorem 3.3 (Bounded-degree algorithm)**.** *Let $\epsilon > 0$, $c > 0$ be constants. There exists an $O(D^{3/2}\sqrt{n}\log^2 n)$-space single-pass streaming algorithm* **ALG** *such that for every directed graph $\mathcal{G} = (V, E)$ with $|V| = n$, $|E| \leq n^c$, and max-degree at most D, the following holds: On input the edges of $\mathcal{G}$ in adversarial order,* **ALG** *outputs an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation to $\mathsf{val}_\mathcal{G}$ with probability at least $2/3$.*

But first, we build some notation. The *bias* of a vertex $i \in V$ with respect to a directed graph $\mathcal{G} = (V, E)$, denoted $\mathsf{bias}_{\mathcal{G}}(i)$, is defined as $\mathsf{bias}_{\mathcal{G}}(i) = \frac{\mathsf{out\text{-}deg}_{\mathcal{G}}(i) - \mathsf{in\text{-}deg}_{\mathcal{G}}(i)}{\mathsf{out\text{-}deg}_{\mathcal{G}}(i) + \mathsf{in\text{-}deg}_{\mathcal{G}}(i)}$, where $\mathsf{out\text{-}deg}_{\mathcal{G}}(i), \mathsf{in\text{-}deg}_{\mathcal{G}}(i)$ respectively denote the out-degree and in-degree of i in $\mathcal{G}$. We now define a quantity called the "density matrix" of a graph with respect to a partition of its vertices into bias intervals. Given any vector $\mathbf{t} = (t_1, \ldots, t_\ell) \in [-1, 1]^\ell$ satisfying $-1 = t_1 < \cdots < t_\ell = 1$, we let $\mathcal{P}_{\mathcal{G},\mathbf{t}}$ denote the "canonical" partition partition of V into blocks of vertices $V = V_1 \sqcup \cdots \sqcup V_\ell$ where for every $r \in [\ell - 1]$, $V_r = \{i : \mathsf{bias}_{\mathcal{G}}(i) \in [t_r, t_{r+1})\}$, and $V_\ell = \{i : \mathsf{bias}_{\mathcal{G}}(i) = 1\}$. Now the density matrix of $\mathcal{G}$ with respect to $\mathbf{t}$, denoted by $M_{\mathcal{G},\mathbf{t}}$, is an $\ell \times \ell$ matrix of natural numbers defined as $M_{\mathcal{G},\mathbf{t}}(i, j) = |E_{V_i \to V_j}|$, for every $i, j \in [\ell]$, i.e., the (i, j)-th entry of $M_{\mathcal{G},\mathbf{t}}$ counts the number of edges in $\mathcal{G}$ between vertices with biases in the intervals $[t_i, t_{i+1})$ (or $\{1\}$ if $i = \ell$) and $[t_j, t_{j+1})$ (or $\{1\}$ if $j = \ell$).

The following lemma was proved in [FJ15] and it shows that there exists a vector $\mathbf{t}$ such that for every directed graph $\mathcal{G}$, the density matrix of $\mathcal{G}$ with the respect to the canonical partition $\mathcal{P}_{\mathcal{G},\mathbf{t}}$ can be used to get a good approximation to the $\mathsf{Max\text{-}DICUT}$ value of $\mathcal{G}$.

Lemma 3.4 ([FJ15])**.** *There exists a constant $\alpha_{\mathrm{FJ}} \in (0.483, 0.4899)$, $\ell_{\mathrm{FJ}} \in \mathbb{N}$, a vector of bias thresholds $\mathbf{t}_{\mathrm{FJ}} = (t_1, \ldots, t_\ell) \in [-1, 1]^{\ell_{\mathrm{FJ}}}$, and a vector of probabilities $\mathbf{p}_{\mathrm{FJ}} = (p_1, \ldots, p_\ell) \in [0, 1]^\ell$ such that for every directed graph $\mathcal{G}$,*

$$\alpha_{\mathrm{FJ}} \cdot \mathsf{val}_{\mathcal{G}} \leq \sum_{i,j=1}^{\ell_{\mathrm{FJ}}} p_i(1 - p_j) M_{\mathcal{G},\mathbf{t}}(i, j) \leq \mathsf{val}_{\mathcal{G}}.$$

We observe that algorithmically, the estimate for $\mathsf{val}_{\mathcal{G}}$ in this lemma corresponds to assigning each vertex in block V_i to L w.p. p_i and R w.p. $1 - p_i$, independently of all other vertices.

As a corollary of Lemma 3.4, we show that in order to get an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation for the $\mathsf{Max\text{-}DICUT}$ value of $\mathcal{G}$, it suffices to obtain an additive $\pm\epsilon' m$ approximation for every element of $M_{\mathcal{G},\mathbf{t}}$, for $\epsilon' = O(\epsilon)$.

Corollary 3.5. *Let $\alpha_{\mathrm{FJ}}, \ell_{\mathrm{FJ}}, \mathbf{t}_{\mathrm{FJ}}, \mathbf{p}_{\mathrm{FJ}}$ be as in Lemma 3.4. Let $\mathcal{G}$ be a directed graph and let m denote the number of edges in $\mathcal{G}$. Let $\epsilon \in (0, \alpha_{\mathrm{FJ}})$ and $\epsilon' = \frac{\epsilon}{8(\ell_{\mathrm{FJ}})^2}$. If there exists $N \in \mathbb{R}^{\ell_{\mathrm{FJ}} \times \ell_{\mathrm{FJ}}}$ such that for every $i, j \in [\ell_{\mathrm{FJ}}]$,*

$$M_{\mathcal{G},\mathbf{t}}(i, j) - \epsilon' m \leq N(i, j) \leq M_{\mathcal{G},\mathbf{t}}(i, j) + \epsilon' m\,,$$

then

$$(\alpha_{\mathrm{FJ}} - \epsilon)\mathsf{val}_{\mathcal{G}} \leq \sum_{i,j \in [\ell_{\mathrm{FJ}}]} p_i(1 - p_j) N(i, j) - \frac{\epsilon}{8} m \leq \mathsf{val}_{\mathcal{G}}.$$

Proof. For the upper bound, we have

$$\begin{aligned} \sum_{i,j \in [\ell_{\mathrm{FJ}}]} p_i(1 - p_j) N(i, j) - \frac{\epsilon}{8} m &\leq \sum_{i,j \in [\ell_{\mathrm{FJ}}]} p_i(1 - p_j)(M_{\mathcal{G},\mathbf{t}}(i, j) + \epsilon' m) - \frac{\epsilon}{8} m \\ &\qquad\qquad \text{(assumption on } N(i,j)) \\ &\leq \mathsf{val}_{\mathcal{G}} + (\ell_{\mathrm{FJ}})^2 \epsilon' m - \frac{\epsilon}{8} m \qquad \text{(Lemma 3.4)} \\ &\leq \mathsf{val}_{\mathcal{G}}. \qquad \text{(choice of } \epsilon') \end{aligned}$$

For the lower bound, we have

$$\begin{aligned} \sum_{i,j \in [\ell_{\mathrm{FJ}}]} p_i(1 - p_j) N(i, j) - \frac{\epsilon}{8} m &\geq \sum_{i,j \in [\ell_{\mathrm{FJ}}]} p_i(1 - p_j)(M_{\mathcal{G},\mathbf{t}}(i, j) - \epsilon' m) - \frac{\epsilon}{8} m \\ &\qquad\qquad \text{(assumption on } N(i,j)) \end{aligned}$$

$$
\begin{aligned}
&\geq \alpha_{\mathrm{FJ}}\mathsf{val}_{\mathcal{G}} - (\ell_{\mathrm{FJ}})^2\epsilon' m - \frac{\epsilon}{8}m && \text{(Lemma 3.4)}\\
&\geq \alpha_{\mathrm{FJ}}\mathsf{val}_{\mathcal{G}} - \frac{\epsilon}{4}m && \text{(choice of } \epsilon')\\
&\geq (\alpha_{\mathrm{FJ}} - \epsilon)\mathsf{val}_{\mathcal{G}}. && (\mathsf{val}_{\mathcal{G}} \geq \tfrac{m}{4})
\end{aligned}
$$

□

In the following subsections, we describe how to estimate $M_{\mathcal{G},\mathbf{t}}$ in a number of different settings: $O(\log n)$-space single-pass streaming algorithm under random ordering of edges (Section 3.1), $O(\log n)$-space two-pass streaming algorithm under adversarial ordering (Section 3.2), and $O(D^{3/2}\sqrt{n}\log^2 n)$-space single-pass streaming algorithm for degree-D bounded graphs under adversarial ordering (Section 3.3). These algorithms share the same central principle: First, let $\mathcal{H} = (V, E')$ be a subgraph of $\mathcal{G} = (V, E)$ (i.e., $E' \subseteq E$). Given bias thresholds $-1 = t_1 < \cdots < t_\ell = 1$, let $M_{\mathcal{H}\subseteq\mathcal{G},\mathbf{t}} \in \mathbb{N}^{\ell\times\ell}$ denote the matrix with entries $M_{\mathcal{H}\subseteq\mathcal{G},\mathbf{t}}(i,j) = |E'_{V_i\to V_j}|$ where $\mathcal{P}_{\mathcal{G},\mathbf{t}} = V_1 \sqcup \cdots \sqcup V_\ell$ is the canonical partition of V with respect to bias in $\mathcal{G}$. (Note that this is distinct from the matrices $M_{\mathcal{G},\mathbf{t}}$ and $M_{\mathcal{H},\mathbf{t}}$ because it counts *edges* in $\mathcal{H}$ but measures *bias* with respect to $\mathcal{G}$.) Now the strategy of all three algorithms is to somehow sample a "representative" subgraph $\mathcal{H}$ of $\mathcal{G}$, and then estimate $M_{\mathcal{G},\mathbf{t}}$ from $M_{\mathcal{H}\subseteq\mathcal{G},\mathbf{t}}$ simply by multiplying every entry by a scale factor $\frac{m(\mathcal{G})}{m(\mathcal{H})}$ (where $m(\mathcal{G}) = |E|$ and $m(\mathcal{H}) = |E'|$). There are two questions associated with this approach, which we answer differently in each setting:

1. *How do we sample a "representative" subgraph $\mathcal{H}$, which doesn't oversample edges from $E_{V_i\to V_j}$ for any $i, j \in [\ell]$?* In Sections 3.1 and 3.2, $\mathcal{H}$ consists of random edges from $\mathcal{G}$, while in Section 3.3, $\mathcal{H}$ is the subgraph induced on random vertices from $\mathcal{G}$. In both cases, we show that (for a sufficiently large sample size), $\mathcal{H}$ is "sufficiently representative" with high probability using concentration bounds.

2. *How do we remember the "global bias" (i.e., the bias in $\mathcal{G}$) of vertices we sample in $\mathcal{H}$?* In the single-pass setting, we measure biases "online": Each time we see a new vertex appear as an endpoint in an edge, we decide whether to track its bias over the rest of the stream or not, and if we decide not to, it cannot have positive degree in $\mathcal{H}$. The two-pass setting obviates this limitation, since we can decide which vertices to track in the first pass and then actually track them in the second pass.

3.1 $O(\log n)$-space random-ordering (single-pass) algorithm

In this subsection, we prove Theorem 3.1 by showing that Algorithm 1 is an $(\alpha_{\mathrm{FJ}}-\epsilon)$-approximation streaming algorithm for computing Max-DICUT value when the edges of the input graph $\mathcal{G}$ are randomly ordered and uses space at most $O(\log n)$. Algorithm 1 uses Algorithm 2 as a subroutine to estimate $M_{\mathcal{G},\mathbf{t}}$ within a small additive error and then uses this estimate to compute an $(\alpha_{\mathrm{FJ}}-\epsilon)$-approximation to the Max-DICUT value of $\mathcal{G}$. We now describe and analyse Algorithm 1 and Algorithm 2.

Algorithm 1 Random-Order-Dicut$_\epsilon(n, \boldsymbol{\sigma})$:

```
Input: $n \in \mathbb{N}$ and a stream $\boldsymbol{\sigma} = (\mathbf{e}(1), \dots, \mathbf{e}(m))$ representing randomly ordered edges of $\mathcal{G}$ on $n$ vertices.
Let $\ell_{\mathrm{FJ}}$, $\mathbf{t}_{\mathrm{FJ}}$, $\mathbf{p}_{\mathrm{FJ}}$ be from Lemma 3.4. Let $k$ and $m_0$ be fixed according to Lemma 3.6 corresponding to $\ell_{\mathrm{FJ}}, \mathbf{t}_{\mathrm{FJ}}$, and $\epsilon' = \frac{\epsilon}{8(\ell_{\mathrm{FJ}}^2)}$.
Store the first $m_0$ edges that arrive in the stream.
Let $N \leftarrow$ Random-Order-Estimate-$M_{\mathcal{G},\mathbf{t}}(n, \boldsymbol{\sigma}, \mathbf{t}_{\mathrm{FJ}}, k)$.
if $m < m_0$ then
    Compute $M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$ directly from the stored edges and $N \leftarrow M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$.
Output $\sum_{i,j=1}^{\ell_{\mathrm{FJ}}} p_i(1-p_j)N(i,j) - \frac{\epsilon}{8}m$.
```

We are now ready to describe our first algorithm for estimating $M_{\mathcal{G},\mathbf{t}}$.

Algorithm 2 Random-Order-Estimate-$M_{\mathcal{G},\mathbf{t}}(n, \boldsymbol{\sigma}, \mathbf{t}, k)$

```
Input: the number $n$ of vertices of a directed graph $\mathcal{G}$, a stream $\boldsymbol{\sigma} = (\mathbf{e}(1), \dots, \mathbf{e}(m))$ representing randomly ordered edges of $\mathcal{G}$, bias thresholds $-1 = t_1 < \cdots < t_\ell = 1$, and a parameter $k \in \mathbb{N}$.
Store the first $k$ edges $(\mathbf{e}(1), \dots, \mathbf{e}(k))$ of the stream. Let $\mathcal{H}$ denote the corresponding subgraph.
Over the remainder of the stream, track the following:
    • for every vertex $i$ with positive degree in $\mathcal{H}$, the degrees out-deg$_{\mathcal{G}}(i)$ and in-deg$_{\mathcal{G}}(i)$,
    • and the total number $m$ of edges in the stream.
After the stream ends, compute the following:
    • for every $i$ with positive degree in $\mathcal{H}$, bias$_{\mathcal{G}}(i)$,
    • and the matrix $M_{\mathcal{H}\subseteq\mathcal{G},\mathbf{t}}$.
Output: $N \in \mathbb{R}^{\ell\times\ell}$, where for every $i,j \in [\ell]$, $N(i,j) = \frac{m}{k} M_{\mathcal{H}\subseteq\mathcal{G},\mathbf{t}}(i,j)$.
```

Now the following lemma asserts the correctness of the estimate in Algorithm 2 for a sufficiently large choice of k:

Lemma 3.6. *For every $\ell \in \mathbb{N}$ and threshold vector $\mathbf{t} \in [-1,1]^\ell$ and $\epsilon' > 0$, there exists $k, m_0 \in \mathbb{N}$ such that for every directed graph $\mathcal{G} = (V, E)$ with $m = |E| \geq m_0$ edges, with probability $\frac{2}{3}$, the matrix N output by Algorithm 2 on input $\mathcal{G}$ satisfies, for every $i, j \in [\ell]$, the inequalities*

$$M_{\mathcal{G},\mathbf{t}}(i,j) - \epsilon' m \leq N(i,j) \leq M_{\mathcal{G},\mathbf{t}}(i,j) + \epsilon' m.$$

Proof. Consider the canonical partition $\mathcal{P}_{\mathcal{G},\mathbf{t}} : V_1 \sqcup \cdots \sqcup V_\ell = V$ of $\mathcal{G}$ with respect to $\mathbf{t}$. Fix some $i, j \in [\ell]$ (over which we'll take a union bound) and let $T = M_{\mathcal{G},\mathbf{t}}(i,j)$ denote the total number of edges in $E_{V_i \to V_j}$.

Now consider random variables $X_1, \dots, X_k$, where X_s is the indicator for the event that $\mathbf{e}(s)$ belongs to $E_{V_i\to V_j}$. Let $X = X_1 + \cdots + X_k$ denote the number of observed edges (i.e., edges in $\{\mathbf{e}(1), \dots, \mathbf{e}(k)\}$) that belong to $E_{V_i\to V_j}$; thus, $X = M_{\mathcal{H}\subseteq\mathcal{G},\mathbf{t}}(i,j)$. Note that $\mathbb{E}[X_s] = T/m$ and so $\mathbb{E}[X] = Tk/m$ and $\mathbb{E}[N(i,j)] = T$. Our goal is to prove that w.h.p., $|N(i,j) - T| \leq \epsilon' m$; rescaling by k/m, we seek to prove that $|X - Tk/m| \leq \epsilon' k$ w.h.p.

For this, we apply the concentration inequalities in Lemmas 2.11 and 2.12 to show that the inequalities $X - Tk/m \leq \epsilon' k$, $X - Tk/m \geq -\epsilon' k$ are violated with probability at most $\exp(-O_{\epsilon'}(k))$. This is sufficient to take a union bound over $i, j \in [\ell]$ if we pick k sufficiently large in terms of ϵ', ℓ and then m_0 sufficiently large in terms of k.

Upper bound. Since $\mathbf{e}(1), \dots, \mathbf{e}(s)$ are sampled from $E(\mathcal{G})$ without replacement, for each $s \in [k]$, we have

$$\mathbb{E}[X_s \mid X_1, \dots, X_{s-1}] = \frac{T - (X_1 + \cdots + X_{s-1})}{m - (s-1)} \leq \frac{T}{m-k}.$$

Setting $p = T/(m-k)$, $\mu = kp$, and $\Delta = \epsilon' k/2$, for sufficiently large m, we claim that $\mu - \frac{T}{m}k \leq \Delta$, and thus that $\mu + \Delta \leq \frac{T}{m}k + \epsilon' k$. The claim follows because, canceling k's and cross-multiplying by m and $m-k$, we get the inequality $kT \leq \epsilon' m(m-k)/2$, which since $T \leq m$ holds whenever $k \leq \epsilon'/(2+\epsilon')m$ (which holds for $m_0 \geq (2+\epsilon')k/\epsilon'$).

Now Lemma 2.11 implies that $X \geq \mu + \Delta$ with probability at most

$$\exp\left(-\frac{\Delta^2}{2(\mu+\Delta)}\right) \leq \exp\left(-\frac{(\epsilon')^2 k^2}{8k(m/(m-k) + \epsilon'/2)}\right) \leq \exp\left(-\frac{(\epsilon')^2}{8(1+\epsilon'/2)}k\right)$$

(using $T \leq m$ and setting $m_0 \geq 2k$).

Lower bound. As in the upper bound, we get $\mathbb{E}[X_s \mid X_1, \dots, X_{s-1}] \geq \frac{T-k}{m-k}$; setting this time $p = (T-k)/(m-k)$, and again $\mu = pk$ and $\Delta = \epsilon' k/2$, we now claim that $\mu - \Delta \geq \frac{T}{m}k - \epsilon' k$; this holds because it's implied by the inequality $k(m-T) \leq \epsilon' m(m-k)/2$, which again holds whenever $k \leq \epsilon'/(2+\epsilon')m$ (now since $T \geq 0$). Now Lemma 2.12 implies that $X \leq \mu - \Delta$, again with probability at most

$$\exp\left(-\frac{\Delta^2}{2(k-(\mu-\Delta))}\right) \leq \exp\left(-\frac{(\epsilon')^2 k}{8(1-(T-k)/(m-k)+\epsilon'/2)}\right) \leq \exp\left(-\frac{(\epsilon')^2}{8(1+\epsilon'/2)}k\right)$$

(using $T \geq 0$ and again $m_0 \geq 2k$). □

Finally, we prove Theorem 3.1.

Proof of Theorem 3.1. Consider Algorithm 1. We fix $\ell_{\mathrm{FJ}}, \mathbf{t}_{\mathrm{FJ}}, \mathbf{p}_{\mathrm{FJ}}, \alpha_{\mathrm{FJ}}$ according to Lemma 3.4. For the choice of $k \in \mathbb{N}$ in Lemma 3.6 that corresponds to $\ell_{\mathrm{FJ}}, \mathbf{t}_{\mathrm{FJ}}$, and $\epsilon' = \frac{\epsilon}{8(\ell_{\mathrm{FJ}})^2}$, we run Algorithm 2 with the parameters $\mathbf{t}_{\mathrm{FJ}}, k$ on the input graph $\mathcal{G}$. For $m \geq m_0$, Lemma 3.6 implies that with probability $\frac{2}{3}$, the output N of Algorithm 2 entrywise approximates $M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$ up to an additive $\pm\epsilon' m$. For $m < m_0$, Algorithm 1 computes $M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$ exactly. Now Corollary 3.5 implies that the output of Algorithm 1 is an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation to the Max-DICUT value of $\mathcal{G}$ as desired.

Finally, we show that Algorithm 1 can be implemented in $O(\log n)$ space. Since m_0 is a constant, it takes only $O(\log n)$ space to store the first m_0 edges. Algorithm 2 can be implemented in $O(\log n)$ space since it takes $O(\log n)$ space to store k edges and we use a simple counter in step 2 that uses $O(\log n)$ space for m that is bounded by $\mathrm{poly}(n)$. □

3.2 Two-pass $O(\log n)$-space adversarial-ordering algorithm

In this subsection, we show how the random-ordering algorithm presented in Section 3.1 can be modified to work with adversarial input ordering given *two* passes over the input stream to prove Theorem 3.2.

Proof of Theorem 3.2. Let $\mathbf{ALG}$ denote the $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation algorithm for Max-DICUT in the random ordering setting (Algorithm 1). Consider the following algorithm $\mathbf{ALG}'$: In the first pass $\mathbf{ALG}'$ uses reservoir sampling (see Section 2.5) to randomly sample k edges from the stream;

this requires $O(k)$ space.[4] In the second pass, it runs the remainder of Algorithm 2 with parameters $\mathbf{t}_{\mathrm{FJ}}, k$ to obtain N and outputs $\sum_{i,j=1}^{\ell_{\mathrm{FJ}}} p_i(1-p_j)N(i,j) - \frac{\epsilon}{8}m$. The same proof of correctness, as well as the space analysis for Algorithm 1 works here as well. We conclude that with probability at least $2/3$, $\mathbf{ALG}'$ outputs an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation to the Max-DICUT value of $\mathcal{G}$. □

3.3 $O(D^{3/2}\sqrt{n}\log^2 n)$-space adversarial-ordering algorithm for degree-D bounded graphs

In this subsection, we prove Theorem 3.3 by showing that Algorithm 3 is an $(\alpha_{\mathrm{FJ}}-\epsilon)$-approximation streaming algorithm for computing Max-DICUT value of degree-D bounded graphs and uses space at most $O(D^{3/2}\sqrt{n}\log^2 n)$. The basic idea is to sample a subset of the vertices of the input graph $\mathcal{G}$ and estimate $M_{\mathcal{G},\mathbf{t}}$ using the density matrix for the induced subgraph $M_{\mathcal{H}\subseteq\mathcal{G},\mathbf{t}}$. However, there are a few issues that ensue. Firstly, we need to deal with the case where most of $\mathcal{G}$'s vertices are isolated (i.e., they have degree zero); we manage this by only sampling vertices which have positive degree, by using a hash function on these vertices. This, in turn, requires estimating the number m of edges in the stream, which is not known *a priori*. For an estimate $\widehat{m}$ that satisfies $\widehat{m} \le m < 2\widehat{m}$, with high probability, Algorithm 4 estimates $M_{\mathcal{G},\mathbf{t}}$ correctly within a small additive error. Algorithm 3 runs Algorithm 4 for various estimates of m and using the correct output from Algorithm 4, it computes an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation to the Max-DICUT value of $\mathcal{G}$. We now describe and analyse Algorithm 3 and Algorithm 4.

Algorithm 3 Bounded-Degree-Dicut$_D(n, \boldsymbol{\sigma})$:

```
Input: $n \in \mathbb{N}$ and a stream $\boldsymbol{\sigma} = (\mathbf{e}(1), \ldots, \mathbf{e}(m))$ representing randomly ordered edges of $\mathcal{G}$ on $n$ vertices.
Let $\ell_{\mathrm{FJ}}$, $\mathbf{t}_{\mathrm{FJ}}$, $\mathbf{p}_{\mathrm{FJ}}$ be from Lemma 3.4. Let $C_1$ and $k$ be fixed according to Lemma 3.7 corresponding to $\ell_{\mathrm{FJ}}, \mathbf{t}_{\mathrm{FJ}}$, and $\epsilon' = \frac{\epsilon}{8(\ell_{\mathrm{FJ}}^2)}$.
Store the first $2C_1^2 D$ edges that arrive in the stream.
for every integer $b$ from 0 to $\lfloor \log(nD/2) \rfloor$ do
    $\widehat{N}_b \leftarrow$ Bounded-Degree-Estimate-$M_{\mathcal{G},\mathbf{t}}(n, \boldsymbol{\sigma}, \mathbf{t}_{\mathrm{FJ}}, k, 2^b)$
    if $\widehat{N}_b$ is not Fail then
        $N \leftarrow \widehat{N}_b$.
if $m < 2C_1^2 D$ then
    Compute $M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$ directly from the stored edges and $N \leftarrow M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$.
Output $\sum_{i,j=1}^{\ell_{\mathrm{FJ}}} p_i(1-p_j)N(i,j) - \frac{\epsilon}{8}m$.
```

[4] Note that if the length of the stream is known *a priori*, there is a simpler sampling procedure. In the first pass, $\mathbf{ALG}'$ can sample every edge in the stream with probability $\frac{2k}{m}$. Let S denote the number of edges that were sampled. With high probability, $|S| \ge k$. Now, $\mathbf{ALG}'$ can choose a random subset of k edges from S.

Algorithm 4 Bounded-Degree-Estimate-$M_{\mathcal{G},\mathbf{t}}(n, \boldsymbol{\sigma}, \mathbf{t}, k, \widehat{m})$

```
Input: the number $n$ of vertices of a directed graph $\mathcal{G}$, a stream $\boldsymbol{\sigma} = (\mathbf{e}(1), \ldots, \mathbf{e}(m))$ representing
    adversarially ordered edges, a vector $\mathbf{t} = (t_1, \ldots, t_\ell) \in [-1, 1]^\ell$, and parameters $k, \widehat{m} \in \mathbb{N}$, where
    $\widehat{m}$ is a power of 2.
Sample a random hash function $\pi : [n] \to [\widehat{m}]$ from a 4-wise independent hash family $\mathsf{H}(n, \widehat{m})$
    (see Section 2.6).
For the remainder of the stream, track the number of edges $m$ that arrive.
Define $s \leftarrow k\sqrt{\widehat{m}}$.
Initialize $\widehat{n} \leftarrow 0$.
Initialize $\mathcal{H} \leftarrow (V, \emptyset)$, where $V$ is the vertex set of $\mathcal{G}$.
for each edge $\mathbf{e}(t) = (u, v)$ in the stream do
    if $\pi(u) \le s$ then
        Track the bias of $u$. Increase $\widehat{n}$ by 1 if this is the first edge incident on $u$.
    if $\pi(v) \le s$ then
        Track the bias of $v$. Increase $\widehat{n}$ by 1 if this is the first edge incident on $v$.
    if $\pi(u) \le s$ and $\pi(v) \le s$ then
        Add $\mathbf{e}$ to $\mathcal{H}$.
    if $\widehat{n} > (5s \cdot \min\{n, 4\widehat{m}\})/\widehat{m}$ then
        Halt and output Fail.
if $m < \widehat{m}$ or $m \ge 2\widehat{m}$ then
    Halt and output Fail.
Output: $N \in \mathbb{R}^{\ell \times \ell}$, where for every $i, j \in [\ell]$, $N(i, j) = \frac{m}{\mu} M_{\mathcal{H} \subseteq \mathcal{G},\mathbf{t}}(i, j)$ where $\mu = ms^2/\widehat{m}^2$.
```

The correctness of Algorithm 4 conditioned on the estimate $\widehat{m}$ being approximately accurate is asserted in the following lemma:

Lemma 3.7. *For every ℓ, threshold vector $\mathbf{t} \in [-1, 1]^\ell$, and $\epsilon' > 0$, there exists $C_1 = C_1(\epsilon') > 0$ such that the following holds. Let $\mathcal{G}$ be a graph with n vertices, m edges, and max-degree $\le D$ such that $m \ge 2C_1^2 D$, and let $\widehat{m} \in \mathbb{N}$ be such that $\widehat{m} \le m < 2\widehat{m}$. Then with probability $\frac{2}{3}$ (over the choice of the permutation π), the matrix N output by Algorithm 4 on input $\mathcal{G}$ (with parameters $k = C_1\sqrt{D}, \widehat{m}$) satisfies, for every $i, j \in [\ell]$, the inequalities*

$$M_{\mathcal{G},\mathbf{t}}(i, j) - \epsilon' m \le N(i, j) \le M_{\mathcal{G},\mathbf{t}}(i, j) + \epsilon' m.$$

Proof. Let $p = s/\widehat{m} = k/\sqrt{m}$[5] and $\mu = p^2 m$. Conditioned on $\widehat{m} \le m < 2\widehat{m}$, we first bound the probability that Algorithm 4 halts and outputs Fail. Observe that $\widehat{n}$ is the number of non-isolated vertices with hash value at most s. Let S denote the set of non-isolated vertices in $\mathcal{G}$. We have $|S| \le \min\{n, 2m\} \le \min\{n, 4\widehat{m}\}$. For vertex $i \in S$, let Y_i be the event that $\pi(i) \le s$. Let $Y = \sum_{i \in S} Y_i = \widehat{n}$. Let $p = s/\widehat{m}$.We have $\mathbb{E}[Y_i] = p$ for every $i \in [n]$ and hence $\mathbb{E}[Y] = p|S|$.[6] Since Y_i, Y_j are independent for $i \ne j$, the variance of Y is given by

$$\mathsf{Var}[Y] = p|S| + p^2(|S|^2 - |S|) - p^2|S|^2 \le p|S|\,.$$

So by Chebyshev's inequality,

$$\Pr\left[\left|Y - p|S|\right| \ge a\sqrt{p|S|}\right] \le \frac{1}{a^2}\,.$$

[5] Note that $p \le 1$ since $\frac{s}{\widehat{m}} = C_1\sqrt{\frac{D}{\widehat{m}}} \le C_1\sqrt{\frac{2D}{m}} \le 1$, by assumption.

[6] Note that $p|S| \ge 1$ since $|S| \ge m/D$ and $p|S| \ge C_1\sqrt{\frac{m}{D}} \ge 1$.

By setting a to be $\sqrt{10}$, we conclude that $\widehat{n} = Y \le 5p|S| \le (5s \cdot \min\{n, 4\widehat{m}\})/\widehat{m}$ with probability at least $\frac{9}{10}$.

Therefore with probability at least $9/10$, conditioned on $\widehat{m} \le m < 2\widehat{m}$, Algorithm 4 does not halt and output `Fail`. Now conditioned on this event, we show that with high probability, the matrix N output by Algorithm 4 on input $\mathcal{G}$ (with parameters $k = C_1\sqrt{D}, \widehat{m}$) satisfies, for every $i, j \in [\ell]$, the inequalities

$$M_{\mathcal{G},\mathbf{t}}(i,j) - \epsilon' m \le N(i,j) \le M_{\mathcal{G},\mathbf{t}}(i,j) + \epsilon' m.$$

Fix $i, j \in [\ell]$, and let $T = |E_{V_i \to V_j}| = M_{\mathcal{G},\mathbf{t}}(i,j)$. Enumerate the edges of $E_{V_i \to V_j}$ as $\mathbf{e}(e_1), \dots, \mathbf{e}(e_T)$ with $\mathbf{e}(e_t) = (u_t, v_t)$. For $t \in [T]$, let X_t be the indicator variable for the event that $\pi(u_t) \le s$ and $\pi(v_t) \le s$. The events $\pi(u_t) \le s$ and $\pi(v_t) \le s$ each occur with probability $s/\widehat{m} = p$, and they are independent (since H is 4- and thus 2-wise independent). Hence $\mathbb{E}[X_t] = p^2$ and, defining $X = X_1 + \cdots + X_T = M_{\mathcal{H} \subseteq \mathcal{G},\mathbf{t}}(i,j)$, we have $\mathbb{E}[X] = p^2 T$ and so $\mathbb{E}[N(i,j)] = m\,\mathbb{E}[X]/\mu = m(p^2T)/(p^2m) = T$. Now observe that the desired inequality can be restated as $|N(i,j) - T| \le \epsilon' m$ which, rescaling by $\mu/m = p^2$, is equivalent to the inequality $|X - Tp^2| \le \epsilon'\mu$. We prove that this holds with high probability using Chebyshev's inequality.

First, we calculate that

$$\mathsf{Var}[X] = \sum_{t,t'=1}^{T} \mathbb{E}[X_t X_{t'}] - (Tp^2)^2.$$

Also, when $\mathbf{e}(e_t)$ and $\mathbf{e}(e_{t'})$ do not share a vertex, the events $\pi(u_t) \le s, \pi(v_t) \le s, \pi(u_{t'}) \le s$, and $\pi(v_{t'}) \le s$ are all independent by 4-wise independence of π, and so $\mathbb{E}[X_t X_{t'}] = \mathbb{E}[X_t]\,\mathbb{E}[X_{t'}] = p^4$. On the other hand, when they are dependent, we can upper-bound $\mathbb{E}[X_t X_{t'}] \le \mathbb{E}[X_t] = p^2$. Since $p \le 1$ and each X_t is dependent on at most $D' = 2D - 1$ variables $X_{t'}$ (by the max-degree assumption), we have

$$\mathsf{Var}[X] \le (T^2 - D'T)p^4 + D'Tp^2 - T^2p^4 \le D'Tp^2.$$

So by Chebyshev's inequality,

$$\Pr[|X - Tp^2| \ge ap\sqrt{D'T}] \le \frac{1}{a^2}.$$

Setting $ap\sqrt{D'T} = \epsilon' p^2 m$, squaring, and simplifying, we get $a^2 D'T = (\epsilon')^2 p^2 m^2$, so

$$\frac{1}{a^2} = \frac{D'T}{(\epsilon')^2 p^2 m^2} = \frac{D'T\widehat{m}}{(\epsilon')^2 k^2 m^2}$$

by the definition of p. Now $D' < 2D$, $T \le m$, and $\widehat{m} \le m$ by assumption, and recalling $k = C_1\sqrt{D}$, we can upper-bound the probability by $\frac{2}{(\epsilon')^2 C_1^2}$, which can be made arbitrarily small (in particular, less than, say, $1/(100\ell^2)$) for a sufficiently large choice of C_1. □

Finally, we prove Theorem 3.3.

Proof of Theorem 3.3. Consider Algorithm 3. We fix $\ell_{\mathrm{FJ}}, \mathbf{t}_{\mathrm{FJ}}, \mathbf{p}_{\mathrm{FJ}}, \alpha_{\mathrm{FJ}}$ according to Lemma 3.4 and k according to Lemma 3.7 corresponding to $\ell_{\mathrm{FJ}}, \mathbf{t}_{\mathrm{FJ}}$, and $\epsilon' = \frac{\epsilon}{8(\ell_{\mathrm{FJ}})^2}$. Since the max-degree of $\mathcal{G}$ is at most D, the number of edges m is at most $nD/2$. Observe that for every m, there is a unique $b \in [0, \lfloor \log(nD/2) \rfloor]$ such that $2^b \le m < 2^{b+1}$. Namely, for $\widehat{b} = \lfloor \log m \rfloor$, we have $2^{\widehat{b}} \le m < 2^{\widehat{b}+1}$. For $b = \widehat{b}$, the algorithm executes Algorithm 4 with $\widehat{m} = 2^{\widehat{b}}$. For $m \ge 2C_1^2 D$, Lemma 3.7 implies that with probability $\frac{2}{3}$, the output N of Algorithm 4 entrywise approximates $M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$ up to an additive $\pm\epsilon' m$. For $m < 2C_1^2 D$, Algorithm 3 computes $M_{\mathcal{G},\mathbf{t}_{\mathrm{FJ}}}$ exactly. Now Corollary 3.5 implies that output of Algorithm 3 is an $(\alpha_{\mathrm{FJ}} - \epsilon)$-approximation to the Max-DICUT value of $\mathcal{G}$ as desired.

Finally, we show that Algorithm 3 can be implemented in $O(D^{3/2}\sqrt{n}\log^2 n)$ space. The first $2C_1^2 D$ edges in the stream can be stored in $O(D\log n)$ space. Since Algorithm 3 executes Algorithm 4 $O(\log n)$ times, it suffices to prove that Algorithm 4 can be implemented in $O(D^{3/2}\sqrt{n}\log n)$ space. Firstly, it takes $O(\log n)$ space to store π (see Section 2.6 for an example construction). Moreover, we can maintain the counter for the number of edges using $O(\log m)$ space. We have $\widehat{n} \leq (5s \cdot \min\{n, 4\widehat{m}\})/\widehat{m}$. Every tracked vertex contributes only $O(D\log n)$ space to store its degree and neighborhood. Therefore, Algorithm 4 requires at most $O\left(D^{3/2}\log n \cdot \min\{n, \widehat{m}\}/\sqrt{\widehat{m}}\right) \leq O(D^{3/2}\log n \cdot \sqrt{n})$ space. Hence, Algorithm 3 can be implemented in $O(D^{3/2}\log^2 n\sqrt{n})$ space. □

4 Lower bounds for Max-CSP in the random-ordering setting

4.1 The Generalized Uniform Randomized Mask Detection (RMD) Problem

We now define the Generalized-Uniform-RMD problem, the main focus of our lower bound. We shall define both a communication version and a streaming version. In either case, we need to define a pair of distributions. As the two pairs are rather closely related, we define them together.

Definition 4.1 (Generalized-Uniform-RMD)**.** *Let $q, k \in \mathbb{N}$ be given and $\mathcal{F}$ be a non-empty set of functions mapping $\mathbb{Z}_q^k \to \{0,1\}$. Let $\alpha > 0$ and $n \in \mathbb{N}$ be parameters and $\mathcal{D}_Y \in \Delta\big(\mathcal{F} \times \Delta_{\mathsf{unif}}(\mathbb{Z}_q^k)\big)$ be a distribution with finite support*[7]*. For all integers $0 \leq t \leq \alpha n$ and both versions, we now define a distribution $\mathcal{H}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n,t)$ as follows:*

1. *For both versions:*
   - (a) *Sample a vector $\mathbf{x}^*$ uniformly at random from $\mathbb{Z}_q^n$.*
   - (b) *For all $i \in [\alpha n]$, sample a matrix $M_i \in \{0,1\}^{k\times n}$ uniformly and independently from the set of all partial permutation matrices*[8]*.*
   - (c) *For all $i \in [\alpha n]$, sample a pair (f_i, D_i) independently from $\mathcal{D}_Y$.*
   - (d) *For all $i \in [\alpha n]$, sample a vector $\mathbf{b}(i) \in \mathbb{Z}_q^k$ independently from D_i if $i \leq t$ and uniformly and independently from the set $\mathbb{Z}_q^k$ if $i > t$.*
   - (e) *For all $i \in [\alpha n]$, set $\mathbf{z}(i) = M_i\mathbf{x}^* - \mathbf{b}(i)$.*
2. *Output as follows:*
   - (a) *For the communication version, define M (respectively, $\mathbf{z}$) to be the matrix (resp., vector) obtained by stacking all the M_i (resp., $\mathbf{z}(i)$) on top of each other. Also, define the vector $\mathbf{D}$ to be the vector consisting of the pairs $(f_i, D_i)_{i\in[\alpha n]}$. Output the pair $(\mathbf{x}^*, (M, \mathbf{z}, \mathbf{D}))$. (The first element of the pair $\mathbf{x}^*$ forms the input for Alice and the second element $(M, \mathbf{z}, \mathbf{D})$ forms the input for Bob.)*
   - (b) *For the streaming version, output the stream $(f_i, M_i, \mathbf{z}(i))_{i\in[\alpha n]}$. (Note that the length of the stream is αn and each symbol is a triple $(f_i, M_i, \mathbf{z}(i))$.)*

[7] Observe that $\mathcal{D}_Y$ is a finite support distribution over pairs, the second element of which is itself a distribution.

[8] Recall that a partial permutation matrix is a matrix with $0, 1$ entries and exactly one 1 in each row and at most one 1 in every column.

For both versions, the problem $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$ *is defined to be the pair of distributions* $(\mathcal{H}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n,\alpha n),\mathcal{H}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n,0))$. *We shall often refer to* $\mathcal{H}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n,\alpha n)$ *as the "yes" distribution and denote it by* $\mathcal{Y}$ *and* $\mathcal{H}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n,0)$ *as the "no" distribution and denote it by* $\mathcal{N}$. *The remaining distributions will only be needed for the streaming version and will be used as "hybrids".*

We note that in the communication version of Definition 4.1, the matrix M given to Bob is the adjacency matrix of a graph sampled from the distribution $\mathcal{G}_{k,\alpha}(n)$ (see Section 2.4.3).

We now define what it means to solve the Generalized-Uniform-RMD communication problem arising from the pair $(\mathcal{F},\mathcal{D}_Y)$ with *non-trivial advantage*. The main emphasis of the definition is the advantage one can get as $\alpha\to 0$. It is natural to expect the advantage to shrink with α, and the definition below requires that the advantage only shrinks linearly with α.

Definition 4.2 (Solving Generalized-Uniform-RMD with non-trivial advantage)**.** *Let* $q,k\in\mathbb{N}$ *be given and* $\mathcal{F}$ *be a non-empty set of functions mapping* $\mathbb{Z}_q^k\to\{0,1\}$. *Let* $\mathcal{D}_Y\in\Delta\big(\mathcal{F}\times\Delta_{\mathsf{unif}}\big(\mathbb{Z}_q^k\big)\big)$ *be a distribution with finite support and* $s:\mathbb{N}\to\mathbb{R}$ *be a function. We say that the pair* $(\mathcal{F},\mathcal{D}_Y)$ *can be solved with non-trivial advantage using* $o(s)$ *communication if there exists* $\delta>0$ *such that for all* $\alpha,\tau>0$, *there exist infinitely many* $n\in\mathbb{N}$ *for which there exists a (randomized) protocol* Π *that solves the* $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$*-problem with advantage* $\delta\cdot\alpha$ *and satisfies* $\|\Pi\|\le\tau\cdot s(n)$.

4.2 Proof of Theorem 1.4

In this section, we state two theorems that together imply Theorem 1.4. These theorems are then proved in the following sections. First, we have the following communication lower bound on the Generalized-Uniform-RMD problem.

Theorem 4.3. *Let* $q,k\in\mathbb{N}$ *be given and* $\mathcal{F}$ *be a non-empty set of functions mapping* $\mathbb{Z}_q^k\to\{0,1\}$. *Let* $\mathcal{D}_Y\in\Delta\big(\mathcal{F}\times\Delta_{\mathsf{unif}}\big(\mathbb{Z}_q^k\big)\big)$ *be a distribution with finite support. Then,* $(\mathcal{F},\mathcal{D}_Y)$ *cannot be solved with non-trivial advantage using* $o(\sqrt{n})$ *communication.*

We also show why the above communication lower bound implies that certain CSPs are approximation resistant.

Theorem 4.4. *Let* $q,k\in\mathbb{N}$ *be given and* $\mathcal{F}$ *be a non-empty set of functions mapping* $\mathbb{Z}_q^k\to\{0,1\}$ *and weakly supporting one-wise independence. There exists a distribution* $\mathcal{D}_Y\in\Delta\big(\mathcal{F}\times\Delta_{\mathsf{unif}}\big(\mathbb{Z}_q^k\big)\big)$ *with a finite support such that if* $(\mathcal{F},\mathcal{D}_Y)$ *cannot be solved with non-trivial advantage using* $o(\sqrt{n})$ *communication, then Max-CSP*$(\mathcal{F})$ *is approximation resistant to* $o(\sqrt{n})$ *space in the random order streaming model.*

4.3 Proof of Theorem 4.4

We now prove Theorem 4.4. The proof of Theorem 4.3 is in the following section. This proof closely follows arguments in [KKS15, CGSV21b].

Proof of Theorem 4.4. As $\mathcal{F}$ weakly supports one wise independence, there exists a non-empty subfamily $\mathcal{F}'\subseteq\mathcal{F}$ that such that $\rho_{\min}(\mathcal{F})=\rho_{\min}(\mathcal{F}')$ and for all $f\in\mathcal{F}'$, there exists a distribution $D_f\in\Delta_{\mathsf{unif}}\big(\mathbb{Z}_q^k\big)$ that is supported on $f^{-1}(1)$. Fix such a family $\mathcal{F}'$ and note by Lemma 2.7 that there exists a distribution $D\in\Delta(\mathcal{F}')$ such that

$$\rho_{\min}(\mathcal{F})=\rho_{\min}(\mathcal{F}')=\max_{D'\in\Delta(\mathbb{Z}_q)}\underset{\substack{f\sim D\\ \mathbf{a}\sim D'^k}}{\mathbb{E}}[f(\mathbf{a})]. \tag{4.5}$$

Define the distribution $\mathcal{D}_Y$ to be the one that first samples $f \sim D$ and then outputs the pair (f, D_f). Clearly, the support of $\mathcal{D}_Y$ is finite and all that remains to be shown is that if $(\mathcal{F}, \mathcal{D}_Y)$ cannot be solved with non-trivial advantage using $o(\sqrt{n})$ communication, then $\mathsf{Max\text{-}CSP}(\mathcal{F})$ is approximation resistant to $o(\sqrt{n})$ space in the random order streaming model. We shall show this in the contrapositive.

Suppose that $\mathsf{Max\text{-}CSP}(\mathcal{F})$ is not approximation resistant to $o(\sqrt{n})$ space in the random order streaming model, and let $\epsilon > 0, p > \frac{1}{2}$ be the parameters promised by Definition 2.8 in this case. Thus, we have for all $\tau > 0$ that there exists $n \in \mathbb{N}$ for which:

$$\begin{aligned}&\text{There exists a randomized streaming algorithms } \mathcal{A},\ \|\mathcal{A}\| < \tau \cdot \sqrt{n} \text{ that}\\ &\text{computes } \mathsf{aprx}_{\mathcal{F},n,\epsilon} \text{ in the random-order streaming model with probability } p.\end{aligned} \qquad (\star)$$

In fact, for any $\tau > 0$, we must have infinitely many values of n such that $(\star)$ holds. Indeed, if there is a τ for which there only finitely many such n, as any non-trivial algorithm must have $\|\mathbf{ALG}\| \geq 1$, we can construct a smaller τ for which there is no value of n satisfying $(\star)$, a contradiction.

To show that $(\mathcal{F}, \mathcal{D}_Y)$ can be solved with non-trivial advantage using $o(\sqrt{n})$ communication, we will show Definition 4.2 with the parameter $\delta = \theta^{20}$, where we define $\theta = \frac{\epsilon}{100} \cdot \frac{(p-1/2)\cdot\rho_{\min}(\mathcal{F})}{q^k}$. Let $\alpha, \tau > 0$ be arbitrary. Applying the reasoning in the foregoing paragraph with this value of τ, we get that there are infinitely many $n \in \mathbb{N}$ for which $(\star)$ holds. Fix any such n that is also larger than $\left(\frac{k}{\theta}\right)^5$ (this only excludes finitely many values). We will show that there exists a protocol Π that solves the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$-problem with advantage $\delta \cdot \alpha$ and satisfies $\|\Pi\| \leq \tau \cdot \sqrt{n})$. We do this in two steps.

Streaming algorithm for $\mathsf{Generalized\text{-}Uniform\text{-}RMD}$. As a first step we define $T = \frac{1}{\alpha\cdot\theta^{10}}$ and show that there exists a deterministic streaming algorithm $\mathbf{ALG}$ that solves the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha T}(n)$ problem with advantage θ in the worst case streaming model. To this end, for $0 \leq t \leq T$, we let $\mathsf{Hyb}^{\mathsf{Str}}(t)$ be the $(\alpha n t)^{\text{th}}$ hybrid distribution of $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha T}(n)$, as defined in Definition 4.1. We also define the distributions $\mathcal{Y}^{\mathsf{Str}} = \mathsf{Hyb}^{\mathsf{Str}}(T)$ and $\mathcal{N}^{\mathsf{Str}} = \mathsf{Hyb}^{\mathsf{Str}}(0)$.

For an instance $\Psi = (f_i, M_i, \mathbf{z}(i))_{i\in[\alpha T n]}$, we define an instance $\mathsf{Clean}(\Psi)$ of $\mathsf{Max\text{-}CSP}_n(\mathcal{F})$ so that for each $i \in [\alpha T n]$ for which $\mathbf{z}(i) = 0^k$, the instance $\mathsf{Clean}(\Psi)$ has (in order) the tuple (f_i, M_i). Also define the distribution $\mathcal{Y}^{\mathsf{CSP}}$ (respectively, $\mathcal{N}^{\mathsf{CSP}}$) to be the distribution that samples an instance Ψ from $\mathcal{Y}^{\mathsf{Str}}$ (resp. $\mathcal{N}^{\mathsf{Str}}$) and outputs $\mathsf{Clean}(\Psi)$. We show that

Claim 4.6. *We have $\mathsf{val}_{\Psi'} = 1$ for all Ψ' in the support of $\mathcal{Y}^{\mathsf{CSP}}$.*

Proof. It suffices to show that $\mathsf{val}_{\Psi'} \geq 1$. If Ψ' is in the support of $\mathcal{Y}^{\mathsf{CSP}}$, there exists Ψ in the support of $\mathcal{Y}^{\mathsf{Str}}$ such that $\mathsf{Clean}(\Psi) = \Psi'$. Let L' be the length of Ψ' and $\left(f'_{i'}, M'_{i'}\right)_{i'\in[L']}$ be the constraints in Ψ'. By definition, we get that for all $i' \in [L']$, there exists an $i = i(i') \in [\alpha T n]$ such that $(f_i, M_i, \mathbf{z}(i)) = \left(f'_{i'}, M'_{i'}, 0^k\right)$. Let $\mathbf{x}^*$ as in definition Definition 4.1 be the one that gave rise to Ψ. We have:

$$\begin{aligned}\mathsf{val}_{\Psi'} &\geq \mathsf{val}_{\Psi'}(\mathbf{x}^*)\\ &= \frac{1}{L'} \cdot \sum_{i'\in[L']} f'_{i'}\left(M'_{i'}\mathbf{x}^*\right)\\ &= \frac{1}{L'} \cdot \sum_{i'\in[L']} f_{i(i')}\left(M_{i(i')}\mathbf{x}^*\right)\end{aligned}$$

$$= \frac{1}{L'} \cdot \sum_{i' \in [L']} f_{i(i')}\big(\mathbf{b}(i(i'))\big) \qquad \text{(As } \mathbf{z}(i(i')) = 0^k\text{)}$$
$$= 1,$$

where the final step uses the fact that $\mathcal{Y}^{\mathsf{Str}} = \mathsf{Hyb}^{\mathsf{Str}}(T)$ is the yes distribution in $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha T}(n)$, which implies that $\mathbf{b}(i(i')) \in f_{i(i')}^{-1}(1)$ by our choice of $\mathcal{D}_Y$. □

Claim 4.7. *For all $i \in [\alpha T n]$ and all $\mathbf{x} \in \mathbb{Z}_q^n$, we have*

$$\Pr_{\Psi' \sim \mathcal{N}^{\mathsf{CSP}}}(f_i(M_i\mathbf{x}) = 1) \leq \rho_{\min}(\mathcal{F}) \cdot \big(1 + \theta^2\big).$$

Proof. Let D_{perm} be the distribution that outputs a uniformly random partial permutation matrix $M \in \{0,1\}^{k \times n}$ and D_{row} be the distribution that outputs a uniformly random matrix $M \in \{0,1\}^{k \times n}$ with exactly one 1 in every row (but a column may have more than one 1). Clearly, D_{perm} is D_{row} conditioned on the event that each row has its 1 in a different column. This means that

$$\|D_{\mathsf{perm}} - D_{\mathsf{row}}\|_{\mathrm{tv}} \leq \Pr_{M \sim D_{\mathsf{row}}}(\text{ Exists two rows with 1 in the same column }) \leq \frac{k^2}{n} \leq \theta^2 \cdot \rho_{\min}(\mathcal{F}),$$

by our choice of n, θ We get:

$$\Pr_{\Psi' \sim \mathcal{N}^{\mathsf{CSP}}}(f_i(M_i\mathbf{x}) = 1) = \Pr_{f \sim D, M \sim D_{\mathsf{perm}}}(f(M\mathbf{x}) = 1) \qquad \text{(Items 1b and 1c)}$$
$$\leq \Pr_{f \sim D, M \sim D_{\mathsf{row}}}(f(M\mathbf{x}) = 1) + \theta^2 \cdot \rho_{\min}(\mathcal{F}).$$
$$\text{(As } \|D_{\mathsf{perm}} - D_{\mathsf{row}}\|_{\mathrm{tv}} \leq \theta^2 \cdot \rho_{\min}(\mathcal{F})\text{)}$$

Now, let D' be the distribution over $\mathbb{Z}_q$ that samples a uniformly random $i \in [n]$ and outputs x_i. Observe that distribution of $M\mathbf{x}$ when $M \sim D_{\mathsf{row}}$ is the same as D'^k. We get:

$$\Pr_{\Psi' \sim \mathcal{N}^{\mathsf{CSP}}}(f_i(M_i\mathbf{x}) = 1) \leq \Pr_{\substack{f \sim D \\ \mathbf{a} \sim D'^k}}(f(\mathbf{a}) = 1) + \theta^2 \cdot \rho_{\min}(\mathcal{F})$$
$$\leq \rho_{\min}(\mathcal{F}) \cdot \big(1 + \theta^2\big). \qquad \text{(Eq. (4.5))}$$

□

Claim 4.8. *We have:*

$$\Pr_{\Psi' \sim \mathcal{N}^{\mathsf{CSP}}}(\mathsf{val}_{\Psi'} > \rho_{\min}(\mathcal{F}) + \epsilon) \leq \theta^2.$$

Proof. Note that:

$$\Pr_{\Psi' \sim \mathcal{N}^{\mathsf{CSP}}}(\mathsf{val}_{\Psi'} > \rho_{\min}(\mathcal{F}) + \epsilon) \leq \Pr_{\Psi' \sim \mathcal{N}^{\mathsf{CSP}}}\big(\exists \mathbf{x} \in \mathbb{Z}_q^n : \mathsf{val}_{\Psi'}(\mathbf{x}) > \rho_{\min}(\mathcal{F}) + \epsilon\big) \qquad \text{(Eq. (2.5))}$$
$$\leq q^n \cdot \max_{\mathbf{x} \in \mathbb{Z}_q^n} \Pr_{\Psi' \sim \mathcal{N}^{\mathsf{CSP}}}(\mathsf{val}_{\Psi'}(\mathbf{x}) > \rho_{\min}(\mathcal{F}) + \epsilon) \qquad \text{(Union bound)}$$
$$\leq q^n \cdot \max_{\mathbf{x} \in \mathbb{Z}_q^n} \Pr_{\Psi \sim \mathcal{N}^{\mathsf{Str}}}\big(\mathsf{val}_{\mathsf{Clean}(\Psi)}(\mathbf{x}) > \rho_{\min}(\mathcal{F}) + \epsilon\big).$$

To finish the proof, we now fix an arbitrary $\mathbf{x} \in \mathbb{Z}_q^n$ and upper bound the probability term above. We shall omit writing $\Psi \sim \mathcal{N}^{\mathsf{Str}}$ for brevity of notation. Note that $\mathsf{val}_{\mathsf{Clean}(\Psi)}(\mathbf{x}) > \rho_{\min}(\mathcal{F}) + \epsilon$ implies by our choice of θ that either $\mathsf{Clean}(\Psi)$ has at most $\big(1 - \theta^2\big) \cdot q^{-k} \cdot \alpha T n$ constraints or it has at least $\big(1 + \theta^2\big) \cdot (\rho_{\min}(\mathcal{F}) + \epsilon/2) \cdot q^{-k} \cdot \alpha T n$ that are satisfied by $\mathbf{x}$. For all $i \in [\alpha T n]$, define indicator random variables X_i and Y_i such that X_i is 1 if and only if $\mathbf{z}(i) = 0^k$ and Y_i is 1 if and only if $\mathsf{X}_i = 1$ and $f_i(M_i\mathbf{x}) = 1$. We get using a union bound:

$$\Pr\left(\mathsf{val}_{\mathsf{Clean}(\Psi)}(\mathbf{x}) > \rho_{\min}(\mathcal{F}) + \epsilon\right) \leq \Pr\left(\sum_{i\in[\alpha Tn]} \mathsf{X}_i \leq \left(1-\theta^2\right)\cdot q^{-k}\cdot \alpha Tn\right) + \Pr\left(\sum_{i\in[\alpha Tn]} \mathsf{Y}_i \geq \left(1+\theta^2\right)\cdot (\rho_{\min}(\mathcal{F}) + \epsilon/2)\cdot q^{-k}\cdot \alpha Tn\right).$$

It is therefore sufficient to bound the probability terms on the right. We will do this using Chernoff bounds. We first claim that the random variables X_i are mutually independent and so are the random variables Y_i. For this, note that both these random variables are determined by the triple $(f_i, M_i, \mathbf{z}(i))$ and (1) For each $i \in [\alpha Tn]$, the triple $(f_i, M_i, \mathbf{z}(i))$ is independent of $\mathbf{x}^*$. This is because, in the distribution $\mathcal{N}^{\mathsf{Str}}$, the vector sampled in Item 1d is uniform over $\mathbb{Z}_q^k$. (2) Conditioned on $\mathbf{x}^*$, the triples $(f_i, M_i, \mathbf{z}(i))$ are mutually independent. This can be observed from Definition 4.1.

Next, we analyze $\Pr(\mathsf{X}_i = 1)$ and $\Pr(\mathsf{Y}_i = 1)$ for $i \in [\alpha Tn]$. For the former, we simply observe from Item 1d that $\Pr(\mathsf{X}_i = 1) = q^{-k}$. For the latter, we have from the definition of Y_i that:

$$\begin{aligned}
\Pr(\mathsf{Y}_i = 1) &= \Pr\left(f_i(M_i\mathbf{x}) = 1 \wedge \mathbf{z}(i) = 0^k\right) \\
&= \Pr(f_i(M_i\mathbf{x}) = 1 \wedge \mathbf{b}(i) = M_i\mathbf{x}^*) && \text{(Item 1e)} \\
&= q^{-k}\cdot \Pr(f_i(M_i\mathbf{x}) = 1) && \text{(Item 1d)} \\
&\leq q^{-k}\cdot \rho_{\min}(\mathcal{F})\cdot \left(1+\theta^2\right) && \text{(Claim 4.7)} \\
&\leq q^{-k}\cdot (\rho_{\min}(\mathcal{F}) + \epsilon/2). && \text{(Claim 4.7)}
\end{aligned}$$

We can now use Chernoff bounds to get:

$$\Pr\left(\mathsf{val}_{\mathsf{Clean}(\Psi)}(\mathbf{x}) > \rho_{\min}(\mathcal{F}) + \epsilon\right) \leq 2^{-\theta^5\cdot q^{-k}\cdot\alpha Tn} + 2^{-\theta^5\cdot(\rho_{\min}(\mathcal{F})+\epsilon/2)\cdot q^{-k}\cdot\alpha Tn} \leq \theta^2\cdot q^{-n},$$

by our choice of T and θ.

□

Define $\mathcal{N}^{\mathsf{CSP}}_{\mathsf{good}}$ to be the same as the distribution $\mathcal{N}^{\mathsf{CSP}}$ conditioned on the event in Claim 4.8 not happening. It follows that $\mathsf{val}_{\Psi'} \leq \rho_{\min}(\mathcal{F}) + \epsilon$ for all Ψ' in the support of $\mathcal{N}^{\mathsf{CSP}}_{\mathsf{good}}$ and that $\|\mathcal{N}^{\mathsf{CSP}}_{\mathsf{good}} - \mathcal{N}^{\mathsf{CSP}}\|_{\mathrm{tv}} \leq \theta^2$. Using the former, Claim 4.6, ($\star$) and Fact 2.3, we get that there is a deterministic streaming algorithm $\mathbf{ALG}'$ with $\|\mathbf{ALG}'\| \leq \tau\cdot\sqrt{n}$ that distinguishes between $\mathcal{Y}^{\mathsf{CSP}}$ and $\mathcal{N}^{\mathsf{CSP}}_{\mathsf{good}}$ with advantage $2\cdot\left(p-\frac{1}{2}\right)$ in the random-order streaming model. This means that

$$\left|\Pr_{\Psi'\sim\mathcal{Y}^{\mathsf{CSP}},\pi\sim\mathcal{S}(|\Psi'|)}\left(\mathbf{ALG}'\left(\pi\left(\Psi'\right)\right) = 1\right) - \Pr_{\Psi'\sim\mathcal{N}^{\mathsf{CSP}}_{\mathsf{good}},\pi\sim\mathcal{S}(|\Psi'|)}\left(\mathbf{ALG}'\left(\pi\left(\Psi'\right)\right) = 1\right)\right| \geq 2\theta.$$

Using $\|\mathcal{N}^{\mathsf{CSP}}_{\mathsf{good}} - \mathcal{N}^{\mathsf{CSP}}\|_{\mathrm{tv}} \leq \theta^2$, we get:

$$\left|\Pr_{\Psi'\sim\mathcal{Y}^{\mathsf{CSP}},\pi\sim\mathcal{S}(|\Psi'|)}\left(\mathbf{ALG}'\left(\pi\left(\Psi'\right)\right) = 1\right) - \Pr_{\Psi'\sim\mathcal{N}^{\mathsf{CSP}},\pi\sim\mathcal{S}(|\Psi'|)}\left(\mathbf{ALG}'\left(\pi\left(\Psi'\right)\right) = 1\right)\right| \geq \theta.$$

Next, use Corollary 2.2 to get:

$$\left|\Pr_{\Psi'\sim\mathcal{Y}^{\mathsf{CSP}}}\left(\mathbf{ALG}'\left(\Psi'\right) = 1\right) - \Pr_{\Psi'\sim\mathcal{N}^{\mathsf{CSP}}}\left(\mathbf{ALG}'\left(\Psi'\right) = 1\right)\right| \geq \theta.$$

By definition of $\mathcal{Y}^{\mathsf{CSP}}, \mathcal{N}^{\mathsf{CSP}}$, we have:

$$\left| \Pr_{\Psi \sim \mathcal{Y}^{\mathsf{Str}}} \left(\mathbf{ALG}'(\mathsf{Clean}(\Psi)) = 1 \right) - \Pr_{\Psi \sim \mathcal{N}^{\mathsf{Str}}} \left(\mathbf{ALG}'(\mathsf{Clean}(\Psi)) = 1 \right) \right| \geq \theta.$$

Now consider a streaming algorithm $\mathbf{ALG}$ for the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha T}(n)$ problem that goes over all triples $(f_i, M_i, \mathbf{z}(i))$ for $i \in [\alpha T n]$, and applies $\mathbf{ALG}'$ on the triples for which $\mathbf{z}(i) = 0^k$. By definition of $\mathbf{ALG}$, we have

$$\left| \Pr_{\Psi \sim \mathcal{Y}^{\mathsf{Str}}} (\mathbf{ALG}(\Psi) = 1) - \Pr_{\Psi \sim \mathcal{N}^{\mathsf{Str}}} (\mathbf{ALG}(\Psi) = 1) \right| \geq \theta. \tag{4.9}$$

Protocol for $\mathsf{Generalized\text{-}Uniform\text{-}RMD}$. We now use our algorithm $\mathbf{ALG}$ to define a (randomized) protocol Π that solves the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$-problem with advantage $\delta \cdot \alpha$ and satisfies $\|\Pi\| \leq \tau \cdot \sqrt{n}$). To start, note that Eq. (4.9) together with the fact that $\mathsf{Hyb}^{\mathsf{Str}}(0) = \mathcal{Y}^{\mathsf{Str}}$ and $\mathsf{Hyb}^{\mathsf{Str}}(T) = \mathcal{N}^{\mathsf{Str}}$ and the triangle inequality, implies there exists a $t \in [T]$ such that

$$\left| \Pr_{\Psi \sim \mathsf{Hyb}^{\mathsf{Str}}(t)} (\mathbf{ALG}(\Psi) = 1) - \Pr_{\Psi \sim \mathsf{Hyb}^{\mathsf{Str}}(t-1)} (\mathbf{ALG}(\Psi) = 1) \right| \geq \frac{\theta}{T} \geq \delta \cdot \alpha. \tag{4.10}$$

Fix such a t and using it to define a Π for the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$-problem as in Algorithm 5. Recall from Section 2.1.1 that notation $\mathbf{ALG}(\sigma, t)$ to denote the state of the streaming algorithm $\mathbf{ALG}$ on input σ after it has processed t symbols from the stream.

Algorithm 5 The protocol Π for the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$-problem.

```
Input: Alice's input is a vector $\mathbf{x}^* \in \mathbb{Z}_q^n$. Bob's input is a triple $(M, \mathbf{z}, \mathbf{D})$ as in Definition 4.1.
  Sampling phase:
  1: Alice samples an instance $\Psi^A$ from the yes distribution of the streaming version of
     $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha(t-1)}(n)$ conditioned on the value $\mathbf{x}^*$.
  2: Bob uses his input to construct $\Psi^{B,1} = (f_i, M_i, \mathbf{z}(i))_{i \in [\alpha n]}$. Next, he samples an instance $\Psi^{B,2}$
     from the no distribution of the streaming version of $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha(T-t)}(n)$.
     He appends this to $\Psi^{B,1}$ to get an instance $\Psi^B = \left(\Psi^{B,1}, \Psi^{B,2}\right)$.
  Communication phase:
  3: Alice and Bob together run $\mathbf{ALG}$ on the instance $\left(\Psi^A, \Psi^B\right)$ as follows:
     (a) Alice runs $\mathbf{ALG}$ on $\Psi^A$, and sends the final state $\mathbf{ALG}\left(\Psi^A, \alpha(t-1)n\right)$ to Bob.
     (b) Bob receives a message $M$ from Alice, and runs $\mathbf{ALG}$ on $\Psi^B$ starting from the state $M$
         and outputting what $\mathbf{ALG}$ outputs.
```

We now analyze the protocol Π and show that it solves the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$-problem with advantage $\delta \cdot \alpha$. For an input $\Phi = (\mathbf{x}^*, (M, \mathbf{z}, \mathbf{D}))$ to the parties in the protocol Π, we define $\Psi^A(\Phi)$ to be the random variable (over Alice's randomness in Π) that equals the instance sampled by Alice in Line 1. Similarly, we define $\Psi^B(\Phi)$ to be the random variable (over Bob's randomness in Π) that equals the instance sampled by Bob in Line 2. Let $\mathcal{Y}^{\mathsf{CC}}$ and $\mathcal{N}^{\mathsf{CC}}$ be the yes and no distributions in the communication version of $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$. We show that:

Lemma 4.11. *It holds for all instances Ψ' that:*

$$\Pr_{\Psi \sim \mathsf{Hyb}^{\mathsf{Str}}(t)}\big(\Psi = \Psi'\big) = \Pr_{\substack{\Phi \sim \mathcal{Y}^{\mathsf{CC}} \\ \Pi \sim \mathsf{\Pi}}}\big(\big(\Psi^A(\Phi), \Psi^B(\Phi)\big) = \Psi'\big).$$

$$\Pr_{\Psi \sim \mathsf{Hyb}^{\mathsf{Str}}(t-1)}\big(\Psi = \Psi'\big) = \Pr_{\substack{\Phi \sim \mathcal{N}^{\mathsf{CC}} \\ \Pi \sim \mathsf{\Pi}}}\big(\big(\Psi^A(\Phi), \Psi^B(\Phi)\big) = \Psi'\big).$$

Before proving Lemma 4.11, we use it to finish the proof of Theorem 4.4 by showing that $\mathsf{\Pi}$ solves the $\mathsf{Generalized\text{-}Uniform\text{-}RMD}_{\mathcal{F},\mathcal{D}_Y,\alpha}(n)$-problem with advantage $\delta \cdot \alpha$. As Line 3 simply runs **ALG** on the sampled instance $\big(\Psi^A(\Phi), \Psi^B(\Phi)\big)$, we have:

$$\delta \cdot \alpha \le \left| \Pr_{\Psi \sim \mathsf{Hyb}^{\mathsf{Str}}(t)}(\mathbf{ALG}(\Psi) = 1) - \Pr_{\Psi \sim \mathsf{Hyb}^{\mathsf{Str}}(t-1)}(\mathbf{ALG}(\Psi) = 1) \right| \qquad \text{(Eq. (4.10))}$$

$$= \left| \Pr_{\substack{\Phi \sim \mathcal{Y}^{\mathsf{CC}} \\ \Pi \sim \mathsf{\Pi}}}\big(\mathbf{ALG}\big(\big(\Psi^A(\Phi), \Psi^B(\Phi)\big)\big) = 1\big) - \Pr_{\substack{\Phi \sim \mathcal{N}^{\mathsf{CC}} \\ \Pi \sim \mathsf{\Pi}}}\big(\mathbf{ALG}\big(\big(\Psi^A(\Phi), \Psi^B(\Phi)\big)\big) = 1\big) \right| \qquad \text{(Lemma 4.11)}$$

$$= \left| \Pr_{\substack{\Phi \sim \mathcal{Y}^{\mathsf{CC}} \\ \Pi \sim \mathsf{\Pi}}}(\Pi(\Phi) = 1) - \Pr_{\substack{\Phi \sim \mathcal{N}^{\mathsf{CC}} \\ \Pi \sim \mathsf{\Pi}}}(\Pi(\Phi) = 1) \right|,$$

as desired. We now show Lemma 4.11.

Proof of Lemma 4.11. We only show the first statement as the proof for the second one is analogous. Let $\mathsf{Hyb}^{\mathsf{CC}}$ be the distribution obtained by first sampling a $\Phi \sim \mathcal{Y}^{\mathsf{CC}}$ and then outputting $\big(\mathbf{x}^*, \Psi^A(\Phi), \Psi^B(\Phi)\big)$ as in the protocol $\mathsf{\Pi}$. Viewing $\mathsf{Hyb}^{\mathsf{Str}}(t)$ as a distribution over $\Big(\mathbf{x}^*, (f_i, M_i, \mathbf{z}(i))_{i \in [\alpha T n]}\Big)$ as in Definition 4.1, we shall show the stronger statement that the distributions $\mathsf{Hyb}^{\mathsf{CC}}$ and $\mathsf{Hyb}^{\mathsf{Str}}(t)$ are the same. We do this in steps.

The marginal distribution of $\mathbf{x}^*$ is the same. We first show that the marginal distribution of $\mathbf{x}^*$ is the same in both distributions. This is because by Definition 4.1, $\mathbf{x}^* \in \mathbb{Z}_q^n$ is uniformly random in both cases.

Conditioned on $\mathbf{x}^*$, the marginals $\{(f_i, M_i, \mathbf{z}(i))\}_{i \in [\alpha T n]}$ are mutually independent. For the case of $\mathsf{Hyb}^{\mathsf{Str}}(t)$, this follows immediately from Definition 4.1. Thus, we only analyze the case of $\mathsf{Hyb}^{\mathsf{CC}}$. In this case, note first from Lines 1 and 2 that conditioned on $\mathbf{x}^*$ the three marginals corresponding to:

$$(f_i, M_i, \mathbf{z}(i))_{0 < i \le \alpha(t-1)n} \qquad (f_i, M_i, \mathbf{z}(i))_{\alpha(t-1)n < i \le \alpha t n} \qquad (f_i, M_i, \mathbf{z}(i))_{\alpha t n < i \le \alpha T n},$$

are mutually independent. This is because conditioned on $\mathbf{x}^*$, the second vector above is Bob's input in $\mathsf{\Pi}$ that Alice does not need to see to sample the first vector in Line 1, and also because the third vector is what Bob samples in Line 2, for which he does not need to see anything (including his input). Thus, it is enough to show that the marginal distribution of all the coordinates in each of the three vectors above are mutually independent conditioned on $\mathbf{x}^*$.

For the first two vectors, this is because of Definition 4.1. For the third vector, this is also because of Definition 4.1 and the fact that in the no distribution of $\mathsf{Generalized\text{-}Uniform\text{-}RMD}$, the triples $(f_i, M_i, \mathbf{z}(i))$ are independent and identically distributed.

For all $i \in [\alpha T n]$, the marginal distribution of $(f_i, M_i, \mathbf{z}(i))$ conditioned on $\mathbf{x}^*$ is the same. For $0 < i \leq \alpha(t-1)n$, this is because of the way Alice samples her Ψ^A in Line 1. For $\alpha(t-1)n < i \leq \alpha tn$, this is by definition of $\mathcal{Y}^{\mathsf{CC}}$. For $\alpha tn < i \leq \alpha Tn$, this is because of the way Bob samples his $\Psi^{B,2}$ in Line 2. Note that in this case as $\mathbf{b}(i)$ is chosen uniformly from $\mathbb{Z}_q^k$, the marginal distribution is actually independent of $\mathbf{x}^*$.

□

□

5 Proof of Theorem 4.3

In this section we prove that the Generalized-Uniform-RMD communication problem arising from $(\mathcal{F}, \mathcal{D}_Y)$ cannot be solved with non-trivial advantage using $o(\sqrt{n})$ communication. The central element in the proof is to look at the distribution of Bob's input $\mathbf{z}$ conditioned on Alice's message and the matrix M, and to argue that the distributions are close in the **YES** and **NO** cases. By definition, the distribution in the **NO** case is uniform over $\mathbb{Z}_q^{km}$ and so what needs to be really shown is that in the **YES** case also this distribution is close to uniform.

Note that Alice's message specifies a set $A \subseteq \mathbb{Z}_q^n$ such that $\mathbf{x}^* \sim \mathsf{Unif}(A)$. Lemma 5.2 roughly relates the distance of the conditional distribution of $\mathbf{z}$ (in the **YES** case) to the Fourier spectrum of the indicator of the set A and to a somewhat complex combinatorial parameter associated with the random hypergraph described by M (see Eq. (5.1)). More precisely Lemma 5.2 bounds this distance provided M is "cycle-free" according to a natural notion of cycle-freeness for hypergraphs that we introduce below. We then state two lemmas upper-bounding the expectation of the combinatorial parameter (Lemma 5.3) and the probability of a cycle (Lemma 5.4), whose proofs are deferred to Section 6. We use these bounds to complete the proof of Theorem 4.3.

The proof outline described above follows the same structure as that of [KKS15] with two significant differences. First, the definition of cycle-freeness is different in our work and this difference has a quantitative effect in that the probability of being cycle-free increases to $\Theta(\alpha^2)$ in our setting compared to $\Theta(\alpha^3)$ in their work. This difference is significant in the context of "non-trivial advantage". Directly following the proof in [KKS15] would have led to a $\Theta(\alpha)$ advantage and we make some changes in the proof of Theorem 4.3 to show that despite the higher probability of cycle-freeness, protocols with non-trivial advantage require $\Omega(\sqrt{n})$ communication. The second difference is in the combinatorial quantity of interest which sees differences due to the higher values of k and q, and the richness of the distributions $\mathcal{D}_Y$ that we need to handle. The analysis of the combinatorial quantity is also more complex and we describe the differences in the next section.

5.1 Indististinguishability via Fourier Analysis

Conditioned on a set $A \subset \mathbb{Z}_q^n$ of $\mathbf{x}^*$'s corresponding to an Alice message, a k-hypergraph $M \in \{0,1\}^{k\alpha n \times n}$, and a vector $\mathbf{D} = ((f_1, D_1), \ldots, (f_m, D_m)) \in (\mathcal{F} \times \Delta_{\mathsf{unif}}(\mathbb{Z}_q^k))^m$, let $\mathcal{Z}_{A,M,\mathbf{D}} \in \Delta(\mathbb{Z}_q^{k\alpha n})$ denote the conditional distribution of Bob's input $\mathbf{z}$ in the **YES** case, i.e.,

$$\mathcal{Z}_{A,M,\mathbf{D}}(\mathbf{z}) = \Pr_{\mathbf{x}^* \sim \mathcal{U}(A), \mathbf{b} \sim D_1 \times \cdots \times D_m}[\mathbf{z} = M\mathbf{x}^* - \mathbf{b}].$$

For a k-hypergraph G, let $\mathsf{cf}(G)$ denote the event that G is *cycle-free* in the sense that its point-hyperplane incidence graph B_G contains no cycles. Let $S_{\neq 1} \overset{\text{def}}{=} \{\mathbf{s} \in (\mathbb{Z}_q^k)^{\alpha n} : \forall i \in [\alpha n], \|\mathbf{s}(i)\|_0 \neq 1\}$.

Then for $\ell \in [n]$, we define the quantity

$$h_{k,\alpha}(\ell, n) \overset{\text{def}}{=} \max_{\mathbf{v} \in \mathbb{Z}_q^n, \|\mathbf{v}\|_0 = \ell} \left(\mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)} \left[\mathbb{1}_{\mathsf{cf}(M)} \cdot \left| \left\{ \mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v} \right\} \right| \right] \right). \tag{5.1}$$

Lemma 5.2 (Fourier-analytic reduction)**.** *Fix $n \in \mathbb{N}$, $\alpha \in (0, 1/100k)$, and a vector*

$$\mathbf{D} = ((f_1, D_1), \ldots, (f_m, D_M)) \in (\mathcal{F} \times \Delta_{\mathsf{unif}}(\mathbb{Z}_q^k))^m.$$

Then

$$\mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)} [\mathbb{1}_{\mathsf{cf}(M)} \cdot \|\mathcal{Z}_{A,M,\mathbf{D}} - \mathcal{U}(\mathbb{Z}_q^{k\alpha n})\|_{\mathrm{tv}}^2] \leq \frac{q^{2n}}{|A|^2} \sum_{\ell=1}^{k\alpha n} h_{k,\alpha}(\ell, n) \mathsf{W}^\ell[\mathbb{1}_A]$$

where $h_{k,\alpha}(\ell, n)$ is defined as in Eq. (5.1).

Proof. Fix $\mathbf{s} \neq \mathbf{0} \in \mathbb{Z}_q^{\alpha kn}$ and let $D = D_1 \times \cdots \times D_m$. We have

$$\begin{aligned}
\widehat{\mathcal{Z}_{A,M,\mathbf{D}}}(\mathbf{s}) &= \frac{1}{q^{\alpha kn}} \sum_{\mathbf{z} \in \mathbb{Z}_q^{k\alpha n}} \mathcal{Z}_{A,M,\mathbf{D}}(\mathbf{z})\, \omega^{-\mathbf{s} \cdot \mathbf{z}} && (\text{definition of } \widehat{\mathcal{Z}_{A,M,\mathbf{D}}}) \\
&= \frac{1}{q^{\alpha kn}} \sum_{\mathbf{z} \in \mathbb{Z}_q^{k\alpha n}} \left(\mathop{\mathbb{E}}_{\mathbf{x}^* \sim A, \mathbf{b} \sim D} [\mathbb{1}_{\mathbf{z} = M\mathbf{x}^* - \mathbf{b}}] \right) \omega^{-\mathbf{s} \cdot \mathbf{z}} && (\text{definition of } \mathcal{Z}_{A,M,\mathbf{D}}) \\
&= \frac{1}{q^{\alpha kn}} \mathop{\mathbb{E}}_{\mathbf{x}^* \sim A, \mathbf{b} \sim D} [\omega^{-\mathbf{s} \cdot (M\mathbf{x}^* - \mathbf{b})}] && (\text{linearity of expectation}) \\
&= \frac{1}{q^{\alpha kn}} \left(\mathop{\mathbb{E}}_{\mathbf{x}^* \sim A} [\omega^{-\mathbf{s} \cdot (M\mathbf{x}^*)}] \right) \left(\prod_{i=1}^{\alpha n} \left(\mathop{\mathbb{E}}_{\mathbf{b}(i) \sim D_i} [\omega^{\mathbf{s}(i) \cdot \mathbf{b}(i)}] \right) \right). && (\text{independence and linearity})
\end{aligned}$$

Now if $\mathbf{s} \notin S_{\neq 1}$, there exists i such that $\|\mathbf{s}(i)\|_0 = 1$, so for some $j \in [k]$, $s(i)_j \neq 0$ while $s(i)_{j'} = 0$ for all $j' \neq j$. Thus, we have $\mathbb{E}_{\mathbf{b}(i) \sim D_i}[\omega^{\mathbf{s}(i) \cdot \mathbf{b}(i)}] = \omega^{s(i)_j} \mathbb{E}_{\mathbf{b}(i) \sim D_i}[\omega^{b(i)_j}] = 0$ because $b(i)_j$ is uniformly distributed on $\mathbb{Z}_q$ by one-wise independence of D_i, and so $\widehat{\mathcal{Z}_{A,M,\mathbf{D}}}(\mathbf{s}) = 0$. Otherwise, using the trivial upper bound $\left|\mathbb{E}_{\mathbf{b}(i) \sim D_i}[\omega^{\mathbf{s}(i) \cdot \mathbf{b}(i)}]\right| \leq 1$, we have

$$\begin{aligned}
|\widehat{\mathcal{Z}_{A,M,\mathbf{D}}}(\mathbf{s})| &\leq \frac{1}{q^{k\alpha n}} \left| \mathop{\mathbb{E}}_{\mathbf{x}^* \sim A} [\omega^{-\mathbf{s} \cdot (M\mathbf{x}^*)}] \right| \\
&= \frac{1}{q^{k\alpha n}} \left| \mathop{\mathbb{E}}_{\mathbf{x}^* \sim A} [\omega^{-(M^\top \mathbf{s}) \cdot \mathbf{x}^*}] \right| && (\text{adjointness}) \\
&= \frac{q^n}{q^{\alpha kn}|A|} |\widehat{\mathbb{1}_A}(M^\top \mathbf{s})|. && (\text{definition of } \widehat{\mathbb{1}_A})
\end{aligned}$$

Thus, by Lemma 2.15 and taking expectation over M, we have

$$\mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)} [\mathbb{1}_{\mathsf{cf}(M)} \cdot \|\mathcal{Z}_{A,M,\mathbf{D}} - \mathcal{U}(\mathbb{Z}_q^m)\|_{\mathrm{tv}}^2] \leq \frac{q^{2n}}{|A|^2} \sum_{\mathbf{s} \neq \mathbf{0} \in S_{\neq 1}} \mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)} [|\widehat{\mathbb{1}_A}(M^\top \mathbf{s})|^2].$$

Rewriting as a sum over $\mathbf{v} = M^\top \mathbf{s}$ gives exactly the desired inequality. □

5.2 Properties of random hypergraphs

Now we state two lemmas about the distribution $\mathcal{G}_{k,\alpha}(n)$ which we will prove in Section 6 below:

Lemma 5.3. *For all $2 \le q, k \in \mathbb{N}$, there exists $c_h < \infty$ and $\alpha_0 > 0$ such that for all $\alpha \in (0, \alpha_0)$,*

$$h_{k,\alpha}(\ell, n) \le \left(\frac{c_h \ell}{n}\right)^{\ell/2}.$$

Lemma 5.4. *For every $k \ge 2$, there exists $c_{\mathsf{cf}} < \infty$ and $\alpha_0 \in (0,1)$ such that for all $n \ge k$ and $\alpha \in (0, \alpha_0)$,*

$$\Pr_{G \sim \mathcal{G}_{k,\alpha}(n)}[\neg \mathsf{cf}(G)] \le c\alpha^2.$$

5.3 Putting the ingredients together

Modulo these lemmas, we can now prove Theorem 4.3:

Proof of Theorem 4.3. Suppose Alice and Bob use a one-way communication protocol Π for Generalized-Uniform-RMD$_{q,k,\mathcal{F},\mathcal{D}_Y,\alpha}$ which uses at most $s = \tau\sqrt{n}$ communication and achieves advantage greater than $\alpha\delta$, where τ is a constant to be determined later. By Yao's principle [Yao77], we may assume WLOG that Π is deterministic and that, from Bob's perspective, Alice's message partitions the set of possible $\mathbf{x}^*$'s into sets $\{A_i \subseteq \mathbb{Z}_q^n\}_{i \in [2^s]}$.

Conditioned on a fixed set $A \subseteq \mathbb{Z}_q^n$, we can view Bob's input $(M, \mathbf{z}, \mathbf{D})$ in both the **YES** and **NO** cases as being sampled by the following process: We sample $M \sim \mathcal{G}_{k,\alpha}(n)$ and $\mathbf{D} \sim \mathcal{D}_Y^{\alpha n}$, and then sample $\mathbf{z}$ either uniformly from $\mathcal{U}(\mathbb{Z}_q^{k\alpha n})$ in the **NO** case or from the conditional distribution $\mathcal{Z}_{A,M,\mathbf{D}}$ in the **YES** case. Thus, Π achieves advantage at most

$$\delta_A \overset{\text{def}}{=} \mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n), \mathbf{D} \sim \mathcal{D}_Y^{\alpha n}}[\|\mathcal{Z}_{A,M,\mathbf{D}} - \mathcal{U}(\mathbb{Z}_q^{k\alpha n})\|_{\mathrm{tv}}].$$

Letting $\mathcal{A}$ denote the distribution which samples each A_i w.p. $|A_i|/q^n$, we have

$$\alpha\delta \le \mathop{\mathbb{E}}_{A \sim \mathcal{A}}[\delta_A]. \tag{5.5}$$

Our goal is to contradict Eq. (5.5) for a sufficiently small choice of τ. We set $\tau = 2\tau'$, where $\tau' > 0$ is to be determined later, and let $s' = \tau'\sqrt{n}$. Also, let $\delta' = \frac{\alpha\delta}{2}$, and let α_0 be the minimum of $\frac{\delta}{2c_{\mathsf{cf}}}$ and the α_0's from Lemmas 5.3 and 5.4. Since $\alpha \le \alpha_0$, we have $c_{\mathsf{cf}}\alpha^2 + \delta' \le \alpha\delta$, so Eq. (5.5) implies

$$c_{\mathsf{cf}}\alpha^2 + \delta' \le \mathop{\mathbb{E}}_{A \sim \mathcal{A}}[\delta_A]. \tag{5.6}$$

A "typical" $A \sim \mathcal{A}$ is large, so to contradict Eq. (5.6), we want to show that δ_A is small for large A. Indeed, since $s' < s - \log_q(2/\delta')$ (for sufficiently large n), we have $\Pr_{A \sim \mathcal{A}}[|A| \le q^{n-s'}] \le \frac{\delta'}{2}$, and it therefore suffices to prove that if $|A| \ge q^{n-s'}$, then $\delta_A \le c\alpha^2 + \frac{\delta'}{2}$.

Let $A \subseteq \mathbb{Z}_q^n$ with $|A| \ge q^{n-s'}$. Conditioning on $\mathsf{cf}(M)$ and using Jensen's inequality and Lemma 5.4, we have

$$\begin{aligned} \delta_A &\le \Pr[\neg\mathsf{cf}(M)] + \mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)}[\mathbb{1}_{\mathsf{cf}(M)} \cdot \|\mathcal{Z}_{A,M,\mathbf{D}} - \mathcal{U}(\mathbb{Z}_q^{k\alpha n})\|_{\mathrm{tv}}] \\ &\le c_{\mathsf{cf}}\alpha^2 + \sqrt{\mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)}[\mathbb{1}_{\mathsf{cf}(M)} \cdot \|\mathcal{Z}_{A,M,\mathbf{D}} - \mathcal{U}(\mathbb{Z}_q^{k\alpha n})\|_{\mathrm{tv}}^2]}. \end{aligned} \tag{5.7}$$

Now we apply Lemma 5.2:

$$\mathop{\mathbb{E}}_{M\sim\mathcal{G}_{k,\alpha}(n)}[\mathbb{1}_{\mathsf{cf}(M)}\cdot\|\mathcal{Z}_{A,M,\mathbf{D}}-\mathcal{U}(\mathbb{Z}_q^{k\alpha n})\|_{\mathrm{tv}}^2]\le\frac{q^{2n}}{|A|^2}\sum_{\ell=1}^{k\alpha n}h_{k,\alpha}(\ell,n)\mathsf{W}^{\ell}[\mathbb{1}_A]$$

We split the sum at $\ell=4s'$, using Lemma 2.14 for the first term and Parseval's identity (Proposition 2.13) for the second:

$$\begin{aligned}&=\frac{q^{2n}}{|A|^2}\sum_{\ell=1}^{4s'}h_{k,\alpha}(\ell,n)\mathsf{W}^{\ell}[\mathbb{1}_A]+\frac{q^{2n}}{|A|^2}\sum_{\ell=4s'}^{k\alpha n}h_{k,\alpha}(\ell,n)\mathsf{W}^{\ell}[\mathbb{1}_A]\\&\le\sum_{\ell=1}^{4s'}h_{k,\alpha}(\ell,n)\left(\frac{\zeta s'}{\ell}\right)^{\ell}+\frac{q^{2n}}{|A|^2}\max_{4s'\le\ell\le k\alpha n}h_{k,\alpha}(\ell,n)\end{aligned}$$

Since $|A|\ge q^{n-s'}$ and $s'=\tau'\sqrt{n}$:

$$\le\sum_{\ell=1}^{4s'}h_{k,\alpha}(\ell,n)\left(\frac{\zeta\tau'\sqrt{n}}{\ell}\right)^{\ell}+q^{2s'}\max_{4s'\le\ell\le k\alpha n}h_{k,\alpha}(\ell,n)$$

Applying Lemma 5.3 and $s'=\tau'\sqrt{n}$:

$$\le\sum_{\ell=1}^{4s'}\left(\zeta\tau'\sqrt{c_h}\right)^{\ell}+\left(16c_hq(\tau')^2\right)^{2s'}$$

where c_h is the constant from Lemma 5.3. Upper-bounding with a geometric series and using the fact that $s'\ge1$ for sufficiently large n:

$$\begin{aligned}&\le\sum_{\ell=1}^{\infty}\left(\zeta\tau'\sqrt{c_h}\right)^{\ell}+16c_hq(\tau')^2\\&=\frac{\zeta\tau'\sqrt{c_h}}{1-\zeta\tau'\sqrt{c_h}}+16cq(\tau')^2\end{aligned}$$

Finally, we set $\tau'>0$ sufficiently small such that both of these terms are at most $\frac{(\delta')^2}{4}$. So plugging in to Eq. (5.7) we get:

$$\delta_A\le c_{\mathsf{cf}}\alpha^2+\frac{\delta'}{2},$$

as desired. □

Remark. *Even a weaker bound in Lemma 5.3 of $(c_h\ell^2/n)^{\ell/2}$ would have sufficed for us to prove Theorem 4.3. On the other hand, we also note that the lemma can be strengthened even further and our proof could actually yield any $c_h>0$ by choosing α_0 small enough. We omit this optimization in Section 6.*

6 Hypergraph analyses

In this section we analyze the quantities of interest in random hypergraphs. In Section 6.1 we analyze the probability that a random hypergraph has a cycle — this analysis is straightforward (and included mainly for completeness). In Section 6.2 we analyze the quantity $h_{k,\alpha}(\ell,n)$ which takes more work. An overview is included in the beginning of Section 6.2.

6.1 Proving Lemma 5.4: Upper-bounding the probability of cycles

Proposition 6.1. *Let $2 \leq k \leq n \in$ and $\alpha \in (0,1)$. For every $u, v \in [n]$ and $j \in [\alpha n]$,*

$$\Pr_{G \sim \mathcal{G}_{k,\alpha}(n)}[u, v \in \mathbf{e}(j)] = \frac{\binom{k}{2}}{\binom{n}{2}},$$

where G has hyperedges $\mathbf{e}(1), \ldots, \mathbf{e}(\alpha n)$.

Proof. By definition, $\mathbf{e}(j)$ is a uniformly random k-tuple of distinct vertices in $[n]$. Consider the following equivalent process for sampling $\mathbf{e}(j)$: Let $\mathbf{e}' = (e'_1, \ldots, e'_n)$ be a uniformly random permutation of $[n]$, and then set $\mathbf{e}(j_i) = (e'_1, \ldots, e'_k)$. We wish to bound the probability that u and v both occur in the first k positions in $\mathbf{e}'$; there are $\binom{n}{2}$ equiprobable pairs of indices at which they can occur, $\binom{k}{2}$ of which satisfy the desired property. □

Proof of Lemma 5.4. First, fix $\ell \geq 2$. Let G have hyperedges $(\mathbf{e}(1), \ldots, \mathbf{e}(\alpha n))$. Fix a sequence $(v_1, \ldots, v_\ell) \in [n]^k$ of distinct vertices and $(j_1, \ldots, j_\ell) \in [\alpha n]^k$ of distinct edge-indices. Consider the event C that $(v_1, \ldots, v_\ell)$ and $(\mathbf{e}(j_1), \ldots, \mathbf{e}(j_\ell))$ form a cycle in $G \sim \mathcal{G}_{k,\alpha}(n)$. Let E_i denote the event that $v_i, v_{i+1} \in \mathbf{e}(j_i)$ (for $i \in [\ell - 1]$) or $v_n, v_1 \in \mathbf{e}(j_\ell)$ (for $i = \ell$). We have $C = E_1 \wedge \cdots \wedge E_\ell$, and since each edge $\mathbf{e}(j_i)$ is selected independently, $E_1, \ldots, E_\ell$ are independent. Thus, we can apply Proposition 6.1 to each E_i to conclude that

$$\Pr[C] = \left(\frac{\binom{k}{2}}{\binom{n}{2}} \right)^\ell \leq \left(\frac{k}{n} \right)^{2\ell}.$$

Now there are $\binom{n}{\ell}\ell! \leq n^\ell$ sequences $(v_1, \ldots, v_\ell)$ and $\binom{\alpha n}{\ell}\ell! \leq (\alpha n)^\ell$ sequences $(j_1, \ldots, j_\ell)$; union bounding over all, we have

$$\Pr_{G \sim \mathcal{G}_{k,\alpha}(n)}[G \text{ contains a cycle of length } \ell] \leq n^\ell (\alpha n)^\ell \left(\frac{k}{n} \right)^{2\ell} = (k^2 \alpha)^\ell.$$

Now we set $\alpha_0 = \frac{1}{2k^2}$, take a union bound over ℓ, and use the geometric series formula:

$$\Pr_{G \sim \mathcal{G}_{k,\alpha}(n)}[\neg \mathsf{cf}(G)] \leq \sum_{\ell=2}^{n} (k^2 \alpha)^\ell \leq \sum_{\ell=2}^{\infty} (k^2 \alpha)^\ell = \frac{(k^2 \alpha)^2}{1 - k^2 \alpha} \leq 2k^4 \alpha^2.$$

Taking $c_{\mathsf{cf}} = 2k^4$ is thus sufficient. □

6.2 Proving Lemma 5.3: Upper-bounding $h_{k,\alpha}(\ell, n)$

In what follows we fix a vector $\mathbf{v} \in \mathbb{Z}_q^n$ with support $U \subseteq [n]$ and upper bound the quantity $\mathbb{E}_{M \sim \mathcal{G}_{k,\alpha}(n)} \left[\mathbb{1}_{\mathsf{cf}(M)} \cdot \left| \left\{ \mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v} \right\} \right| \right]$. For $M \in \mathsf{supp}(\mathcal{G}_{k,\alpha}(n))$ let $X(M) = \mathbb{1}_{\mathsf{cf}(M)} \cdot \left| \left\{ \mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v} \right\} \right|$ so that the quantity of interest is $\mathbb{E}_{M \sim \mathcal{G}_{k,\alpha}(n)} [X(M)]$. To analyze this expectation, first in Proposition 6.2 we give combinatorial conditions on M under which $X(M) = 0$. Further we give a simpler upper bound on $X(M)$ in terms of the connected component structure of M when $X(M)$ is potentially non-zero. Roughly, this proposition bounds $X(M)$ by some function of the size of the connected components of M that are incident to the set U. Lemmas 6.3 to 6.6 then analyze the probability that the components have large size. The resulting bounds are put together to prove Lemma 5.3 at the end of this section.

We now turn to proving Lemma 5.3. Throughout this section, the vertex-hyperedge incidence graph $B = B_M$ corresponding to a k-hypergraph M (from Section 6) will be the central object of interest. While we refer to vertices of M as "vertices", the vertices of B are referred to as either "left vertices" (corresponding to vertices of M) or "right vertices" (corresponding to hyperedges of M). Similarly we use "hyperedges" to refer to edges of M and "edges" to refer to edges of B. In this interpretation, the i-th hyperedge $\mathbf{e}(i)$ of M is the neighborhood of the i-th right vertex of B. Thus, sampling a random hypergraph $M \sim \mathcal{G}_{k,\alpha}(n)$ is equivalent to sampling B by setting each right vertex's neighborhood to be a uniform and independent subset of k left vertices. The vector $\mathbf{v}$ can be viewed as a $\mathbb{Z}_q$-labelling of the left vertices of B, while the vector $\mathbf{s}$ is a $\mathbb{Z}_q$-labelling of B's edges. The condition $\mathbf{s} \in S_{\neq 1}$ means that no right-vertex of B has degree exactly one, and the condition $M^\top \mathbf{s} = \mathbf{v}$ implies that the left vertices of B are each labelled by the sum (modulo q) of the labels of incident edges of B. The condition that U is the support of $\mathbf{v}$ implies that U is exactly the set of left vertices with non-zero labels.

Now consider the connected component decomposition of B, which induces a partition $V_1, \dots, V_{t'}$ of B's left vertices $[n]$. Since $U \subseteq [n]$ is a subset of B's left vertices, B's partition of $[n]$ further induces a partition of U into subsets $U_1, \dots, U_t$ for $t \leq t'$. (This partition is given by intersecting each V_i with U and throwing it away if the intersection is empty. Thus, each component U_i of U is contained in a single connected component of B.)

Note that this partition (given U and B) is essentially unique up to renaming of the parts. We formalize this as follows. We say that $U_1, \dots, U_t$ is a *canonical partition* of U if each U_i contains the least numbered vertex of U that is not contained in $\cup_{j<i} U_j$. (Note that every partition $U_1, \dots, U_t$ can be converted into a canonical one by renumbering the parts. Furthermore given U and B this partition is unique.) We let $\mathsf{cc\text{-}part}(B, U)$, for "connected component partition", denote this canonical partition of U induced by B. We say that B *partitions* U *into* t *connected components* if $\mathsf{cc\text{-}part}(B, U)$ has t parts.

Given a subset $U' \subseteq U$ contained in a unique connected component of B, we say it has L-*type* ℓ if $\ell = |U'|$, and R-*type* r if the connected component of B containing U' has exactly r right vertices. These numbers satisfy the inequality $\ell \leq kr$ since every left vertex must touch at least one right vertex. More generally, if B partitions U into connected components $\mathsf{cc\text{-}part}(B, U) = (U_1, \dots, U_t)$, we say $\mathsf{cc\text{-}part}(B, U)$ is of L-*type* $(\ell_1, \dots, \ell_t)$ if $\ell_i = |U_i|$ for every $i \in [t]$. We say $\mathsf{cc\text{-}part}(B, U)$ is *valid* if $\ell_i \geq 2$ for every i. We say $\mathsf{cc\text{-}part}(B, U)$ is of R-*type* $(r_1, \dots, r_t)$ if in B, the connected component containing U_i has exactly r_i right vertices for every $i \in [t]$, and $\mathsf{cc\text{-}part}(B, U)$ is of R-*total-type* r if $\sum_{i \in [t]} r_i = r$.

The following proposition fixes a graph M and give conditions on when the quantity $\mathbb{1}_{\mathsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right|$ is non-zero; moreover, when it is non-zero, we give an upper bound on it.

Proposition 6.2. *For a fixed* $\mathbf{v} \in \mathbb{Z}_q^n$ *with support* $U \subseteq [n]$ *and a fixed* k*-hypergraph* M*, the quantity* $\mathbb{1}_{\mathsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right|$ *is non-zero only if* M *is cycle-free, and* $\mathsf{cc\text{-}part}(B, U)$ *is a valid partition. Furthermore, for every* $r \in \mathbb{N}$*, if* M *is cycle-free and* $\mathsf{cc\text{-}part}(B, U)$ *is a valid partition of R-total-type* r*, we have* $\mathbb{1}_{\mathsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right| \leq q^{kr}$.

Proof. For the quantity $\mathbb{1}_{\mathsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right|$ to be non-zero, clearly it is necessary that M is cycle-free, which is equivalent to requiring that B is acyclic.

Fix $\mathbf{s} \in S_{\neq 1}$ with $M^\top \mathbf{s} = \mathbf{v}$. Let $B_{\neq}$ be the subgraph of B consisting of the edges with non-zero labels. Recall that we view $\mathbf{v}$ and $\mathbf{s}$ as $\mathbb{Z}_q$-labelings of B's left vertices and edges, respectively, such that the sum of edge labels at every left vertex equals the vertex's label (in $\mathbb{Z}_q$). Thus, every left vertex which has degree zero in $B_{\neq}$ must be labelled 0. Thus, every vertex of $U = \mathsf{supp}(\mathbf{v})$ must have degree at least 1 in $B_{\neq}$, and conversely, defining a *leaf* of $B_{\neq}$ as a vertex with degree exactly 1 in $B_{\neq}$, we see that every left vertex which is a leaf of $B_{\neq}$ must be in U.

Now the condition $\mathbf{s} \in S_{\neq 1}$ implies that no right vertex is a leaf in $B_{\neq}$. Fix a vertex $j \in U$. By the previous paragraph, j has degree at least 1 in $B_{\neq}$. Now consider the connected component of j in $B_{\neq}$. This component is a tree (since B is acyclic and $B_{\neq}$ is a subgraph of B), and so it must have at least two leaves. Since right vertices cannot be leaves in $S_{\neq}$, these leaves must be left vertices. At most one of these leaves can be j, so it follows that the component containing j in $B_{\neq}$ must contain at least one more vertex of U. Thus, the component containing j in B, which is a superset of j's component in $B_{\neq}$, must also contain at least one more vertex of U. Since this holds for every $j \in U$, it follows that U is partitioned into connected components by B with each component containing at least two vertices. In other words, $\mathsf{cc\text{-}part}(B, U)$ is a valid partition of U.

We now turn to bounding the number of vectors $\mathbf{s}$ satisfying $\mathbf{s} \in S_{\neq 1}$ and $M^\top \mathbf{s} = \mathbf{v}$ assuming M is cycle-free and $\mathsf{cc\text{-}part}(B, U)$ is a valid partition of U. Consider a right vertex of B whose connected component does not contain any vertex of U. We claim that all edges of B in this connected component must have a label of zero: this is so since if there is an edge with a non-zero label, the component of $B_{\neq}$ containing this edge must have a leaf, but all of $B_{\neq}$'s leaves are in U. We thus conclude that only edges of B from components containing vertices of U can have non-zero labels. By definition of R-total-type we have that the number of right vertices of B in components containing vertices of U is r, and so the number of edges of B in components containing vertices of U is at most kr. It follows that the number of vectors $\mathbf{s}$ satisfying $\mathbf{s} \in S_{\neq 1}$ and $M^\top \mathbf{s} = \mathbf{v}$ (assuming B is cycle-free and $\mathsf{cc\text{-}part}(B, U)$ is a valid partition of U) is at most q^{kr}. □

Now, we prove several lemmas regarding the probability of a random graph M partitioning sets in various ways, building towards Lemma 6.6 below which bounds the probability that $\mathsf{cc\text{-}part}(B_M, U)$ is a valid partition of R-total-type r.

Lemma 6.3. *Let $n/2 + 1 \leq n' \leq n$ and $\alpha \in (0, 1)$. Let $M \sim \mathcal{G}_{k,\alpha'}(n')$ for $\alpha' = \alpha n/n'$ and $B = B_M$. Then for every $u \in [n']$,*

$$\Pr_M[B \textit{ places } u \textit{ in a component of R-type at least } r_1] \leq (2ek^2\alpha)^{r_1}.$$

Proof. Let B's right vertices have neighborhoods $\mathbf{e}(1), \ldots, \mathbf{e}(\alpha n)$ (corresponding to M's hyperedges). For fixed $j \in [\alpha n]$, the probability that $u \in \mathbf{e}(j)$ is exactly k/n'. Thus, the probability that there exists $j \in [\alpha n]$ such that $u \in \mathbf{e}(j)$ is at most $\alpha n k/n' \leq 2k\alpha \leq 2k^2\alpha$.

Now, condition on the event that there exists $j_1 \in [\alpha n]$ such that $u \in \mathbf{e}(j_1)$. We bound the probability that there exist $r_1 - 1$ additional right vertices in B forming a connected component with j_1. For this to happen there must exist a set of distinct right vertices $\{j_2, \ldots, j_{r_1}\} \subseteq [\alpha n]$ and a spanning tree T on $\{j_1, \ldots, j_{r_1}\}$ such that if $(j_i, j_{i'}) \in T$ then their neighborhoods intersect, i.e., $\mathbf{e}(j_i) \cap \mathbf{e}(j_{i'}) \neq \emptyset$ (or, in M, the hyperedges $\mathbf{e}(j_i)$ and $\mathbf{e}(j_{i'})$ share a common vertex). For a fixed set $\{j_2, \ldots, j_{r_1}\}$ and spanning tree T, this occurs with probability at most $(2k^2/n)^{r_1 - 1}$, since we can do a "depth-first search" on T: Each new right vertex's neighborhood is selected independently of all previous neighborhoods, and intersects its parent's neighborhood with probability $k^2/n' \leq 2k^2/n$.

Now, we do a union bound over all possible subsets $\{j_2, \ldots, j_{r_1}\}$ and spanning trees T. There are $\binom{\alpha n}{r_1 - 1}$ possible subsets and $r_1^{r_1 - 1}$ spanning trees. Thus the probability that there exists a connected component of R-type r_1 including j_1 is at most

$$\binom{\alpha n}{r_1 - 1} \cdot r_1^{r_1 - 1} \cdot (2k^2/n)^{r_1 - 1} \leq \left(\frac{2ek^2\alpha r_1}{(r_1 - 1)}\right)^{r_1 - 1} \leq e^{r_1}(2ek^2\alpha)^{r_1 - 1}.$$

Factoring in the probability that there exists a right vertex j_1 connecting to u gives the desired conclusion. □

Lemma 6.4. *Let $\alpha \leq 1/(2e^3k^2)$, $n/2+1 \leq n' \leq n$ and $r_1 \in \mathbb{N}$. Fix a set $U_1 \subseteq [n']$ with $|U_1| = \ell_1$. Let $M \sim \mathcal{G}_{k,\alpha'}(n')$ for $\alpha' = \alpha n/n'$ and let $B = B_M$. Then*

$$\Pr_M\left[B \textit{ partitions } U_1 \textit{ into a single connected component of R-type } r_1\right] \leq (2ek^2\alpha)^{r_1/2}(2k(\ell_1 - 1)/n)^{\ell_1 - 1}.$$

Proof. We first upper bound the LHS above by $(2ek^2\alpha)^{r_1}(k^2r_1/n)^{\ell_1-1}$, and then show that this is upper bounded by the RHS for $\alpha \leq 1/(2e^3k^2)$.

Let B's right vertices have neighborhoods $\mathbf{e}(1), \ldots, \mathbf{e}(\alpha n)$. Fix a left vertex $u \in U_1$. We condition on the event that, as in the previous lemma (Lemma 6.3), when B partitions $[n]'$, the connected component containing u has R-type r_1. We now bound the probability that the rest of U_1 is contained in this same component. Let $S \subseteq [n']$ be the set of left vertices in the connected component containing u. Since this component has R-type r_1, we have $|S| \leq kr_1$. Our goal is to analyze the probability that $U_1 \setminus \{u\} \subseteq S$. Since the conditioning is symmetric with respect to renaming the vertices of $U_1 \setminus \{u\}$, we can instead consider the probability that $\ell_1 - 1$ random independent left vertices are in S. There are $\binom{|S|}{\ell_1-1}$ ways of choosing $\ell_1 - 1$ vertices in S, out of the possible universe of $\binom{n'-1}{\ell_1-1} \geq \binom{n/2}{\ell_1-1}$ ways of choosing $\ell_1 - 1$ vertices. We thus get that the probability that $U_1 \setminus \{u\} \subseteq S$ is at most

$$\frac{\binom{|S|}{\ell_1-1}}{\binom{n/2}{\ell_1-1}} \leq \left(\frac{2|S|}{n}\right)^{\ell_1-1} \leq \left(\frac{2kr_1}{n}\right)^{\ell_1-1}.$$

Combining this bound with the result of Lemma 6.3, we get that the probability that U_1 is in a connected component of R-type r_1 is at most

$$(2ek^2\alpha)^{r_1}\left(\frac{2kr_1}{n}\right)^{\ell_1-1}.$$

To conclude we need to show that the expression above is upper bounded by the RHS in the statement of the claim.

We consider two cases. If $r_1 \leq \ell_1$ then the bound is immediate assuming $2ek^2\alpha \leq 1$ since we have

$$(2ek^2\alpha)^{r_1}\left(\frac{2kr_1}{n}\right)^{\ell_1-1} \leq (2ek^2\alpha)^{r_1}\left(\frac{2k\ell_1}{n}\right)^{\ell_1-1} \leq (2ek^2\alpha)^{r_1/2}\left(\frac{2k\ell_1}{n}\right)^{\ell_1-1}.$$

When $r_1 > \ell_1$ we note that the expression $a^x x^b$ is non-increasing in x for integer $x \geq b$ and $a \leq 1/e$ and hence is upper bounded by $(ab)^b \leq b^b$. (Incrementing x by 1 multiplies the first term by $a \leq 1/e$ while multiplying the second term by $(1+1/x)^b \leq (1+1/b)^b \leq e$.) We thus get

$$\begin{aligned}(2ek^2\alpha)^{r_1}\left(\frac{2kr_1}{n}\right)^{\ell_1-1} &= (2ek^2\alpha)^{r_1/2}(2ek^2\alpha)^{r_1/2}\left(\frac{2kr_1}{n}\right)^{\ell_1-1} \\ &\leq (2ek^2\alpha)^{r_1/2}\left(\frac{2k(\ell_1-1)}{n}\right)^{\ell_1-1}.\end{aligned}$$

(The first inequality above applies $a^x b^x \leq b^b$ when $a \leq 1$ and $x \leq b$, $x = r_1$, $a = (2k^2\alpha)^{1/2}$, and $b = \ell_1 - 1$.) This concludes the proof of the lemma. □

Lemma 6.5. *Let $\alpha \leq 1(2e^3k^2)$ and $n \geq 4$. Fix $r \in \mathbb{N}$, a set $U \subseteq [n]$ and a canonical partition $U_1, \ldots, U_t$ of U. Let $\ell = |U|$ and $\ell_i = |U_i|$. Let $M \sim \mathcal{G}_{k,\alpha}(n)$. We have*

$$\Pr[\textsf{cc-part}(B,U) = (U_1, \ldots, U_t) \textit{ with R-total-type } r] \leq (32ek^2\alpha)^{r/2}(2k/n)^{\ell-t}\prod_{i=1}^{t}(\ell_i - 1)^{\ell_i - 1}.$$

Proof. Fix $r_1, \dots, r_t$ such that $\sum_i r_i = r$. For every $i \in [t]$ we claim that conditioned on $U_1, \dots, U_{i-1}$ being the first $i-1$ components in the canonical partition $\textsf{cc-part}(B, U)$ of U induced by B, the probability that U_i is the i-th component and has R-type r_i is at most $(2ek^2\alpha)^{r_i/2}(2k(\ell_i - 1)/n)^{\ell_i - 1}$. This follows essentially immediately from Lemma 6.4.

Indeed, observe that conditioned on $U_1, \dots, U_{i-1}$ being the first $i-1$ components of the canonical partition induced by $\textsf{cc-part}(B, U)$, B is "random on the remaining vertices", i.e., the neighborhood of every remaining right vertex is a uniform and independent subset of k remaining left vertices, where "remaining" means not in any of the connected components containing $U_1, \dots, U_{i-1}$. Let n' denote the number of remaining left vertices. We have $n' \geq n/2 + 1$ since the total number of right vertices of B is αn, each touches k left vertices, and $k\alpha n \leq n/2 - 1$ for every $n \geq 4$ and $\alpha \leq 1/(4k)$. Thus we can apply Lemma 6.4 to the remaining hypergraph which has at most αn edges and n' vertices. We conclude that the probability that U_i is the i-th component in $\textsf{cc-part}(B, U)$ and has R-type r_i is at most $(2ek^2\alpha)^{r_i/2}(2k(\ell_i - 1)/n)^{\ell_i - 1}$.

Taking the product of these conditional probabilities, it follows that the probability that $(U_1, \dots, U_t)$ is the partition of U induced by B and has R-type $(r_1, \dots, r_t)$ is at most

$$\prod_{i=1}^{t} (2ek^2\alpha)^{r_i/2}(2k(\ell_i - 1)/n)^{\ell_i - 1} = (2ek^2\alpha)^{r/2}(2k/n)^{\ell - t}\prod_{i=1}^{t}(\ell_i - 1)^{\ell_i - 1}.$$

Finally to conclude the lemma we take a union bound over all possible ways of obtaining r_i's that sum to r. There are at most $\binom{r+t}{t} \leq 4^r$ such ways and thus we get that:

$$\begin{aligned}\Pr[\textsf{cc-part}(B, U) = (U_1, \dots, U_t) \text{ and has R-total-type } r] &\leq 4^r \cdot (2ek^2\alpha)^{r/2}(2k/n)^{\ell - t}\prod_{i=1}^{t}(\ell_i - 1)^{\ell_i - 1} \\ &= (32ek^2\alpha)^{r/2}(2k/n)^{\ell - t}\prod_{i=1}^{t}(\ell_i - 1)^{\ell_i - 1}.\end{aligned}$$

□

Lemma 6.6. *Let $\alpha \leq 1(2e^3k^2)$, $n \geq 4$ and $\ell \leq n/(4ek)$. Fix $r \in \mathbb{N}$, a set $U \subseteq [n]$ with $|U| = \ell$. Let $M \sim \mathcal{G}_{k,\alpha}(n)$ and $B = B_M$. Then*

$$\Pr_M[\textsf{cc-part}(B, U) \text{ is valid and has R-total-type } r] \leq 2(32ek^2\alpha)^{r/2}(32ek\ell/n)^{\ell/2}.$$

Proof. The lemma follows by using Lemma 6.5 and a union bound of all valid canonical partitions of U. Fix t and $\ell_1, \dots, \ell_t$ such that $\sum_i \ell_i = \ell$ and $\ell_i \geq 2$ for all i. Let $N(\ell_1, \dots, \ell_t)$ denote the number of canonical partitions of U of L-type $(\ell_1, \dots, \ell_t)$. We have:

$$N(\ell_1, \dots, \ell_t) = \binom{\ell - 1}{\ell_1 - 1} \cdot \binom{\ell - \ell_1 - 1}{\ell_2 - 1} \cdots \binom{\ell - (\sum_{i<t} \ell_i) - 1}{\ell_t - 1} \leq \frac{\ell^{\ell - t}}{\prod_{i=1}^{t}(\ell_i - 1)!} \leq \frac{(e\ell)^{\ell - t}}{\prod_{i=1}^{t}(\ell_i - 1)^{\ell_i - 1}}$$

For every such partition $U_1, \dots, U_t$ of L-type $(\ell_1, \dots, \ell_t)$, Lemma 6.5 gives an upper bound on the probability that the canonical partition of U under B is $U_1, \dots, U_t$ and has R-total-type r. Taking the union over all such $U_1, \dots, U_t$ we get:

$$\begin{aligned}&\Pr_M[\textsf{cc-part}(B, U) \text{ is of R-total-type } r \text{ and of L-type}(\ell_1, \dots, \ell_t)] \\ &\leq N(\ell_1, \dots, \ell_t) \cdot (32ek^2\alpha)^{r/2}(2k/n)^{\ell - t}\prod_{i=1}^{t}(\ell_i - 1)^{\ell_i - 1}\end{aligned}$$

$$
\begin{aligned}
&\le (e\ell)^{\ell-t} \cdot (32ek^2\alpha)^{r/2}(2k/n)^{\ell-t} \\
&\le (32ek^2\alpha)^{r/2}(2ek\ell/n)^{\ell-t}
\end{aligned}
$$

To conclude the lemma we need to take a union bound over all $(\ell_1, \ldots, \ell_t)$ that are valid. The number of these is at most 4^ℓ for any give t. Furthermore we have $t \le \ell/2$ since $\ell_i \ge 2$ for every i. We conclude

$$
\begin{aligned}
\Pr_M [\textsf{cc-part}(B,U) \text{ is valid of R-total-type } r] &\le \sum_{t=1}^{\ell/2} 4^\ell (32ek^2\alpha)^{r/2}(2ek\ell/n)^{\ell-t} \\
&\le 2 \cdot 4^\ell (32ek^2\alpha)^{r/2}(2ek\ell/n)^{\ell/2} \\
&= 2(32ek^2\alpha)^{r/2}(32ek\ell/n)^{\ell/2}.
\end{aligned}
$$

□

We are now ready to prove Lemma 5.3.

Proof of Lemma 5.3. We prove the lemma for $\alpha_0 = 1/(128e^3k^2q^{2k})$ and $c_h = 128ek$.

Fix $\mathbf{v} \in \mathbb{Z}_q^n$ with support U of cardinality ℓ. Let $B = B_M$. By Proposition 6.2 we have that $\mathbb{1}_{\textsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right|$ is zero unless M is cycle-free and $\textsf{cc-part}(B,U)$ is a valid partition. If $\textsf{cc-part}(B,U)$ is a valid partition it must have R-total-type r for some $r \le \alpha n$. But since every vertex in U has nonzero degree in B, we must also have $r \ge \ell/k$. For any given r in this range, by Lemma 6.6 we have that $\textsf{cc-part}(B,U)$ is a valid partition of R-total-type r with probability at most $2(32ek^2\alpha)^{r/2}(32ek\ell/n)^{\ell/2}$. Conditioned on this event we have (again from Proposition 6.2) that $\mathbb{1}_{\textsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right| \le q^{kr}$. Combining these expressions we have that

$$
\begin{aligned}
&\mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)} \left[\mathbb{1}_{\textsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right|\right] \\
&\le \sum_{r=\ell/k}^{\alpha n} 2q^{kr}(32ek^2\alpha)^{r/2}(32ek\ell/n)^{\ell/2} \\
&\le \sum_{r=\ell/k}^{\infty} 2(32ek^2q^{2k}\alpha)^{r/2}(32ek\ell/n)^{\ell/2} \\
&\le 4(32ek^2q^{2k}\alpha)^{\ell/2k}(32ek\ell/n)^{\ell/2},
\end{aligned}
$$

where the final inequality uses the fact that for $\alpha \le \alpha_0$ we have $32ek^2q^{2k}\alpha \le 1/4$ and so the sum telescopes to at most twice the first term in the series. We simply the final expression further using $4 \le 4^{\ell/2}$ (which holds for every $\ell \ge 2$) and $32ek^2q^{2k}\alpha \le 1$ to get

$$
h_{k,\alpha}(\ell,n) \overset{\text{def}}{=} \max_{\mathbf{v} \in \mathbb{Z}_q^n, \|\mathbf{v}\|_0 = \ell} \left(\mathop{\mathbb{E}}_{M \sim \mathcal{G}_{k,\alpha}(n)} \left[\mathbb{1}_{\textsf{cf}(M)} \cdot \left|\left\{\mathbf{s} \in S_{\neq 1} : M^\top \mathbf{s} = \mathbf{v}\right\}\right|\right] \right) \le (c_h\ell/n)^{\ell/2},
$$

for $c_h = 128ek$. (We note that we could have got any $c_h > 0$ by choosing α small enough, but we don't seem to need this in the application of this lemma, so omit this easy step.) □

References

[AKO11] Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Streaming Algorithms via Precision Sampling. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS 2011, Palm Springs, CA, USA, October 23-25, 2011)*, pages 363–372, October 2011.

[BHP$^+$22] Joanna Boyland, Michael Hwang, Tarun Prasad, Noah Singer, and Santhoshini Velusamy. Sketching approximations for (some) symmetric Boolean CSPs: Closed-form ratios and simple algorithms. February 2022.

[CGS$^+$22] Chi-Ning Chou, Alexander Golovnev, Madhu Sudan, Ameya Velingker, and Santhoshini Velusamy. Linear Space Streaming Lower Bounds for Approximating CSPs. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC 2022, Rome, Italy, June 20-24, 2022)*, 2022. To appear.

[CGSV21a] Chi-Ning Chou, Alexander Golovnev, Madhu Sudan, and Santhoshini Velusamy. Approximability of all Boolean CSPs with linear sketches. February 2021.

[CGSV21b] Chi-Ning Chou, Alexander Golovnev, Madhu Sudan, and Santhoshini Velusamy. Approximability of all finite CSPs with linear sketches. In *Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2021, Denver, CO, USA, February 7-10, 2022)*. IEEE Computer Society, 2021.

[CGV20] Chi-Ning Chou, Alexander Golovnev, and Santhoshini Velusamy. Optimal Streaming Approximations for all Boolean Max-2CSPs and Max-kSAT. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS 2020, Virtual, November 16-19, 2020)*, pages 330–341. IEEE Computer Society, November 2020.

[FJ15] Uriel Feige and Shlomo Jozeph. Oblivious Algorithms for the Maximum Directed Cut Problem. *Algorithmica*, 71(2):409–428, February 2015.

[GT19] Venkatesan Guruswami and Runzhou Tao. Streaming Hardness of Unique Games. In Dimitris Achlioptas and László A. Végh, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX 2019, Cambridge, MA, USA, September 20-22, 2019)*, volume 145 of *LIPIcs*, pages 5:1–5:12. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, September 2019.

[GVV17] Venkatesan Guruswami, Ameya Velingker, and Santhoshini Velusamy. Streaming Complexity of Approximating Max 2CSP and Max Acyclic Subgraph. In Klaus Jansen, José D. P. Rolim, David Williamson, and Santosh S. Vempala, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX 2017, Berkeley, CA, USA, August 16-18, 2017)*, volume 81 of *LIPIcs*, pages 8:1–8:19. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, August 2017.

[Ind06] Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *Journal of the ACM*, 53(3):307–323, May 2006. Conference version in FOCS 2000.

[KK19] Michael Kapralov and Dmitry Krachun. An optimal space lower bound for approximating MAX-CUT. In *Proceedings of the 51st Annual ACM SIGACT Symposium* *on Theory of Computing (STOC 2019, Phoenix, AZ, USA, June 23-26, 2019)*, pages 277–288. Association for Computing Machinery, June 2019.

[KKS14] Michael Kapralov, Sanjeev Khanna, and Madhu Sudan. Approximating matching size from random streams. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014, Portland, OR, USA, January 5-7, 2014)*, pages 734–751, USA, January 2014. Society for Industrial and Applied Mathematics.

[KKS15] Michael Kapralov, Sanjeev Khanna, and Madhu Sudan. Streaming lower bounds for approximating MAX-CUT. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015, San Diego, California, USA, January 4-6, 2015)*, pages 1263–1282. Society for Industrial and Applied Mathematics, January 2015.

[KMNT20] Michael Kapralov, Slobodan Mitrović, Ashkan Norouzi-Fard, and Jakab Tardos. Space efficient approximation to maximum matching size from uniform edge samples. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1753–1772. Society for Industrial and Applied Mathematics, January 2020.

[KNW10] Daniel M. Kane, Jelani Nelson, and David P. Woodruff. On the Exact Space Complexity of Sketching and Streaming Small Norms. In *Proceedings of the 2010 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010, Austin, TX, USA, January 17-19, 2010)*, pages 1161–1178. Society for Industrial and Applied Mathematics, 2010.

[MMPS17] Morteza Monemizadeh, S. Muthukrishnan, Pan Peng, and Christian Sohler. Testable Bounded Degree Graph Properties Are Random Order Streamable. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017, Warsaw, Poland, July 10-14, 2017)*, volume 80 of *LIPIcs*, pages 131:1–131:14. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2017.

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions.* Cambridge University Press, New York, NY, 1st edition edition, June 2014.

[PS18] Pan Peng and Christian Sohler. Estimating Graph Parameters from Random Order Streams. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2018, New Orleans, LA, USA, January 7-10, 2018).* Society for Industrial and Applied Mathematics, January 2018.

[Sin22] Noah Singer. *On Streaming Approximation Algorithms for Constraint Satisfaction Problems.* Bachelor's thesis, Harvard University, Cambridge, MA, March 2022.

[SSV21] Noah Singer, Madhu Sudan, and Santhoshini Velusamy. Streaming approximation resistance of every ordering CSP. In Mary Wootters and Laura Sanità, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX 2021, August 16-18, 2021)*, volume 207 of *LIPIcs*, pages 17:1–17:19. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, September 2021.

[Sud22] Madhu Sudan. Streaming and Sketching Complexity of CSPs: A survey. To appear as invited talk at ICALP 2022, 2022.

[Vad12] Salil Vadhan. Pesudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1–3):1–336, 2012.

[Vit85] Jeffrey S. Vitter. Random sampling with a reservoir. *ACM Trans. Math. Softw.*, 11(1):37–57, mar 1985.

[Yao77] Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (SFCS 1977, Providence, RI, USA, October 31-November 2, 1977)*, pages 222–227. IEEE Computer Society, September 1977.