Electronic Colloquium on Computational Complexity, Report No. 54 (2026)

Optimal Single-Pass Streaming Lower Bounds for Approximating CSPs

Noah G. Singer* Madhur Tulsiani† Santhoshini Velusamy‡

April 9, 2026

Abstract

For an arbitrary family of predicates $\mathcal{F} \subseteq \{0,1\}^{[q]^k}$ and any $\varepsilon > 0$, we prove a single-pass, linear-space streaming lower bound against the gap promise problem of distinguishing instances of MAX-CSP($\mathcal{F}$) with at most $\beta + \varepsilon$ fraction of satisfiable constraints from instances of with at least $\gamma - \varepsilon$ fraction of satisfiable constraints, whenever MAX-CSP($\mathcal{F}$) admits a (γ, β)-integrality gap instance for the basic LP. This subsumes the linear-space lower bound of Chou, Golovnev, Sudan, Velingker, and Velusamy (STOC 2022), which applies only to a special subclass of CSPs with linear-algebraic structure. (Their result itself generalizes work of Kapralov and Krachun (STOC 2019) for MAX-CUT.) Our approach identifies the right "analytic" analogues of previously-used linear-algebraic conditions; this yields substantial simplifications while capturing a much larger class of problems.

Our lower bound is essentially optimal for single-pass streaming, since: (1) All CSPs admit $(1-\varepsilon)$-approximations in quasilinear space, and (2) sublinear-space streaming algorithms can simulate the LP (on bounded-degree instances), giving approximation algorithms when integrality gap instances do *not* exist.

The starting point for our lower bound is a reduction from a "distributional implicit hidden partition" problem defined by Fei, Minzer, and Wang (STOC 2026) in the context of multi-pass streaming. Our result is an analogue of theirs in the single-pass setting, where we obtain a much stronger (and tight) space lower bound.

Independent work. Very recently, similar linear-space lower bounds for CSPs were also obtained by an independent work of Fei, Minzer, and Wang (arXiv, April 2026). Their recent result also implies almost-linear space lower bounds for multi-pass streaming algorithms with $o(\log n)$ passes.

*Carnegie Mellon University, Pittsburgh, PA, USA. Email: ngsinger@cs.cmu.edu.

†Toyota Technological Institute, Chicago, IL, USA. Email: madhurt@ttic.edu. Work supported in part by NSF award CCF-2326685.

‡University of Waterloo, ON, Canada. Email: santhoshini.velusamy@uwaterloo.ca. Work supported in part by NSF award CCF 2348475 when the author was affiliated with Toyota Technological Institute at Chicago.

ISSN 1433-8092

Contents

1 Introduction

Constraint Satisfaction Problems (CSPs) are some of the most fundamental and expressive problems in discrete optimization, which have been a tested for algorithmic techniques and have played a foundational role in the theory of approximation algorithms and inapproximability. Investigation of the complexity of CSPs in various computational models have not only led to a new understanding of the limits of these computational models, but the finite nature of the templates for CSPs have led to surprising classifications and canonical algorithms. In this work, we study the complexity of approximating CSPs in the model of single-pass streaming algorithms.

Constraint satisfaction problems. Fix $q, k \geq 2 \in \mathbb{N}$ and let $\mathcal{V}$ be a finite set, elements of which we call *variables* (typically, $\mathcal{V} = [n] = \{1, \ldots, n\}$). A *constraint* on $\mathcal{V}$ is a pair $C = (f, e)$, where $f : [q]^k \to \{0,1\}$ is a function called a *predicate* and $e = (i_1, \ldots, i_k) \in \mathcal{V}^k$ is a k-tuple of *distinct* variables. (q is called the *alphabet size* and k the *arity*.) An *assignment* to $\mathcal{V}$ is a vector $x \in [q]^{\mathcal{V}}$, and an assignment x *satisfies* a constraint $C = (f, e = (i_1, \ldots, i_k))$ if $f(x_{i_1}, \ldots, x_{i_k}) = 1$.

A *constraint satisfaction problem (CSP)* is specified by a *predicate family* $\mathcal{F} \subseteq \{0,1\}^{[q]^k}$, i.e., a set of functions from $[q]^k \to \{0,1\}$. An *instance* Φ of MAX-CSP($\mathcal{F}$) on a set of variables $\mathcal{V}$ is list of constraints (f, e) on $\mathcal{V}$ with $f \in \mathcal{F}$. The *value* of an assignment $x \in [q]^{\mathcal{V}}$ is $\mathsf{val}_\Phi(x) := \mathbb{E}_{C \sim \Phi}[x \text{ satisfies } C]$, where the distribution over constraints in Φ is taken to be uniform. The *optimum value* of an instance Φ is $\mathsf{opt}^{\mathrm{CSP}}(\Phi) := \max_{x \in [q]^{\mathcal{V}}} \mathsf{val}_\Phi(x)$.

Example 1.1. For $k = 2$ and $q = 2$, and the set $[q]$ identified with $\{0,1\}$, consider the CSPs:

- MAX-CUT: defined using the single predicate $f_{\mathrm{CUT}}(x_1, x_2) = x_1 \oplus x_2$
- MAX-DICUT: defined using the predicate $f_{\mathrm{DICUT}}(x_1, x_2) = x_1 \wedge \neg x_2$.
- MAX-2AND: defined using the *predicate family* $\mathcal{F} = \{f^{b_1,b_2}_{\mathrm{2AND}}(x_1, x_2) = (b_1 \oplus x_1) \wedge (b_2 \oplus x_2) \mid b_1, b_2 \in \{0,1\}\}$.

We note that MAX-2AND is an example of a "CSP with literals" i.e., a predicate family closed under negations of variables. ◊

There are two natural approximation variants of the problem MAX-CSP($\mathcal{F}$). The first, denoted MAX-CSP$_\alpha(\mathcal{F})$, is to distinguish between the cases (for a given instance Φ and threshold v) between the cases $\mathsf{opt}^{\mathrm{CSP}}(\Phi) \geq v$ and $\mathsf{opt}^{\mathrm{CSP}}(\Phi) < \alpha \cdot v$. The second problem, denoted MAX-CSP$_{\beta,\gamma}(\mathcal{F})$, a more fine-grained version of the above and requires distinguishing for a given instance Φ between the cases $\mathsf{opt}^{\mathrm{CSP}}(\Phi) \geq \gamma$ and $\mathsf{opt}^{\mathrm{CSP}}(\Phi) \leq \beta$. We will consider randomized algorithms in this work, for which it suffices to solve the above tasks with probability (say) $2/3$.

Streaming algorithms. We consider *single-pass streaming algorithms* for solving the problems MAX-CSP$_\alpha(\mathcal{F})$ and MAX-CSP$_{\beta,\gamma}(\mathcal{F})$. A *single-pass streaming algorithm* for MAX-CSP$_\alpha(\mathcal{F})$ or MAX-CSP$_{\beta,\gamma}(\mathcal{F})$ is given as input an ordered list (stream) of (not necessarily distinct) constraints (f, e); this list defines an associated instance of MAX-CSP($\mathcal{F}$), with the associated constraint distribution being uniform on the list. The stream is presented to the algorithm one constraint at a time, and the algorithm can only use a bounded amount of memory space to store its current state between each successive pair of constraints.

Formally, *space-s single-pass streaming algorithm* on instances of MAX-CSP($\mathcal{F}$) with n variables and m constraints may be specified by a function $\Gamma : \mathcal{C}_n \times \{0,1\}^s \to \{0,1\}^s$, where $\mathcal{C}_n$ denotes the set of all possible constraints on n variables. The initial state is taken to be $S_0 \leftarrow 0^s \in \{0,1\}^s$, and

j-th constraint leads to the update $S_j \leftarrow \Gamma(C_j, S_{j-1})$. We take (say, the first bit of) the final state S_m as the output for the relevant distinguishing problem. Note that the algorithm is allowed to depend on n and (for the purpose of lower bounds) we also allow it to depend on m. There is no constraint on the time-complexity or randomness used by the algorithm.

1.1 Prior work on streaming approximability of CSPs

Streaming approximability of CSPs has been extensively studied over the past decade; we recall some of the most relevant results below. We refer the reader to the excellent surveys by Sudan [Sud22] and Assadi [Ass23], and the column [Sin25] for more detailed accounts.

Approximability in $O(\sqrt{n})$ space. The first lower bounds against streaming algorithms for CSPs were studied by Kogan and Krauthgamer [KK15], and Kapralov, Khanna, and Sudan [KKS15]. In particular, [KKS15] showed that the MAX-CUT$_{1,1/2+\varepsilon}$ problem requires $\Omega_\varepsilon(\sqrt{n})$ space for every $\varepsilon > 0$ (while a random assignment trivially achieves an approximation ratio $1/2$). Follow-up works [CGV20; GT19] proved optimal $\Omega(\sqrt{n})$-space lower bounds for MAX-UNIQUE-GAMES and MAX-DICUT. We highlight that, unlike MAX-CUT and MAX-UNIQUE-GAMES, the MAX-DICUT problem turns out to admit nontrivial streaming approximations in $\log n \ll \sqrt{n}$ space [GVV17]. This line of work culminated in the work of Chou, Golovnev, Sudan, and Velusamy [CGSV24], who proved a *dichotomy theorem* stating that for every predicate family $\mathcal{F}$ and $\gamma > \beta$, either MAX-CSP$_{\gamma+\varepsilon,\beta-\varepsilon}(\mathcal{F})$ admits a polylog$_\varepsilon(n)$-space algorithm for every $\varepsilon > 0$ (using the "bias" of the variables in the CSP instance) or MAX-CSP$_{\gamma-\varepsilon,\beta+\varepsilon}(\mathcal{F})$ requires $\Omega_\varepsilon(\sqrt{n})$ space for every $\varepsilon > 0$.[1]

Approximability in sublinear space. Beyond $\sqrt{n}$ space, the picture becomes more complex. On the algorithmic side, a series of works [SSSV23b; SSSV23a; SSSV25; ABFS26] established that *better* approximations are attainable for the MAX-DICUT problem in $o(n)$ space (in comparison to the $o(\sqrt{n})$-space regime). This opened up the exciting possibility that the right approximation thresholds for many other CSPs may be *different* in the (arguably, more natural) setting of sublinear space, compared to the ones obtained for $O(\sqrt{n})$ space above. On the hardness side, Kapralov and Krachun [KK19] showed (following up on [KKSV17]) that the MAX-CUT$_{1,1/2+\varepsilon}$ problem requires $\Omega_\varepsilon(n)$ space for every $\varepsilon > 0$, improving the [KKS15] result (which had an $\Omega_\varepsilon(\sqrt{n})$-space lower bound). This result was generalized by Chou, Golovnev, Sudan, Velingker, and Velusamy [CGS+22], who gave $\Omega(n)$-space lower bounds against CSPs satisfying a certain (linear) structural condition called "width". While the work [CGS+22] greatly clarified the technical aspects of proving $\Omega(n)$-space lower bounds (especially in the case $k > 2$), their techniques relied heavily on the presence of such linear structure, leaving open the approximability of many CSPs of interest. Moreover the"width" condition itself is quite brittle, and some CSPs which are not "wide" were still known to admit lower bounds via reduction from wide CSPs (but were not captured by the above result).

Approximability for multi-pass streaming. While the above results concern single-pass streaming algorithms, another line of work has also studied the *multi-pass* approximability of CSPs [BDV18; AKSY20; AN21; CKP+23; KPSY23; SSSV23b; KPSY23; SSSV25; FMW25; Vel25]. Recently, a beautiful result of Fei, Minzer and Wang [FMW26a] obtained a dichotomy theorem for multi-pass streaming using a linear programming relaxation called the "basic LP". They proved that for every predicate family $\mathcal{F}$, if there exists a (γ, β)-LP integrality gap instance (an instance certifying inapproximability via the basic LP) for MAX-CSP$(\mathcal{F})$, then for every $\varepsilon > 0$, the MAX-CSP$_{\gamma-\varepsilon,\beta+\varepsilon}(\mathcal{F})$ problem

[1] The lower bound in [CGSV24] is currently known only to hold against "sketching" algorithms, which are a special subclass of streaming algorithms. Whether this distinction is only a technical artifact, or whether there are streaming algorithms outperforming sketching algorithms, is an interesting open problem.

requires $p \cdot s \geq \Omega_\varepsilon(n^{1/3})$, where p is the number of passes and s is the amount of space. Conversely, if no such instance exists, there is an algorithm for $\text{MAX-CSP}_{\gamma+\varepsilon,\beta-\varepsilon}(\mathcal{F})$ with $p \cdot s \leq \text{polylog}_\varepsilon(n)$. While this result completely characterizes polynomial space approximability in the the multi-pass setting, it does not imply bounds more than $\Omega(n^{1/3})$ for the single-pass setting.

1.2 Results

We show that the basic LP also characterizes the sublinear-space approximability of (bounded-degree instances) of all CSPs in the single-pass setting. In particular, we show that basic LP lower bounds for all CSPs can be converted to similar lower bound on approximation achieved by any algorithm in $o(n)$ space. Before describing the formal statements of the results, we briefly describe the basic LP relaxation for $\text{MAX-CSP}(\mathcal{F})$.

The basic linear programming relaxation for a CSP. Let Φ be an instance of $\text{MAX-CSP}(\mathcal{F})$. The *basic linear programming (LP) relaxation* of Φ is formulated in its "distributional" version in Figure 1 below. The "variables" of this LP are distributions $\mathcal{Y}_C$ over $[q]^k$ for every constraint $C \in \text{supp}(\Phi)$. (Each such distribution can be encoded as $q^k - 1$ standard LP variables.) An important fact about this LP is that it is a *relaxation* of the underlying CSP instance Φ: we always have $1 \geq \text{opt}^{\text{LP}}(\Phi) \geq \text{opt}^{\text{CSP}}(\Phi)$, where $\text{opt}^{\text{LP}}(\Phi)$ is the maximum feasible value for the LP.

$$
\begin{array}{lll}
\text{maximize} & \mathop{\mathbb{E}}_{C=(f,e)\sim\Phi}\left[\mathop{\mathbb{E}}_{\boldsymbol{a}\sim\mathcal{Y}_C} f(\boldsymbol{a})\right] & \\
\text{such that} & \Pr_{\boldsymbol{a}\sim\mathcal{Y}_C}[\boldsymbol{a}_\ell = b] = \Pr_{\boldsymbol{a}\sim\mathcal{Y}_{C'}}[\boldsymbol{a}_{\ell'} = b], & \forall C = (f,e), C' = (f',e') \in \Phi, i_\ell = i'_{\ell'}, b \in [q]. \\
& \mathcal{Y}_C \in \Delta([q]^k) & \forall C \in \Phi.
\end{array}
$$

Figure 1: Distributional view of basic LP.

A (γ,β)-LP Integrality Gap instance for $\text{MAX-CSP}(\mathcal{F})$ is an instance Φ such that $\text{opt}^{\text{CSP}}(\Phi) \leq \beta$ and $\text{opt}^{\text{LP}}(\Phi) \geq \gamma$. Our main result is the following:

Theorem 1.2 (Linear-space lower bounds against single-pass streaming algorithms for CSPs)**.** *Let $\mathcal{F} \subseteq \{0,1\}^{[q]^k}$. For every $1 \geq \gamma > \beta \geq 0$, if there exists a (γ,β)-LP integrality gap instance of $\text{MAX-CSP}(\mathcal{F})$, then for every $\varepsilon > 0$, any single-pass streaming algorithm for $\text{MAX-CSP}_{\gamma-\varepsilon,\beta+\varepsilon}(\mathcal{F})$ on n variables requires $\Omega_\varepsilon(n)$ space. Moreover, $\text{MAX-CSP}_{1,\beta+\varepsilon}(\mathcal{F})$ also requires space $\Omega_\varepsilon(n)$ for the case $\gamma = 1$.*

We emphasize that Theorem 1.2 subsumes all previous known $\Omega_\varepsilon(n)$-space streaming lower bounds [KK19; CGS+22], and is an analogue of (though formally incomparable with) the multi-pass lower bound in [FMW26a]. We also remark that Theorem 1.2 is essentially an *optimal lower bound* against single-pass streaming algorithms for all CSPs, in the following sense:

- With respect to space complexity, as observed in e.g. [CGS+22, §1.1], algorithms with $\widetilde{O}_\varepsilon(n)$ space can solve the $\text{MAX-CSP}_{1-\varepsilon}(\mathcal{F})$ problem for every predicate family $\mathcal{F}$ and $\varepsilon > 0$. Hence, one cannot hope to improve the space lower bound in the theorem (for any $\mathcal{F}$).
- With respect to approximation factors, when a (γ,β)-integrality gap instance does *not* exist, [STV25] observe that for every $\varepsilon > 0$ and $D \in \mathbb{N}$, the $\text{MAX-CSP}_{\gamma+\varepsilon,\beta-\varepsilon}(\mathcal{F})$ problem admits an $o_{\varepsilon,D}(n)$-space algorithm on instances of maximum degree D. Removing the bounded-degree assumption remains an interesting open question.[2]

[2]We feel that the maximum degree assumption can likely be removed, as was true for Max-Dicut [ABFS26].

- The hard instances we use in the proof of Theorem 1.2 have *bounded maximum degree*, and so Theorem 1.2 already proves a *linear-space dichotomy theorem for bounded-degree streaming CSPs*.

Thus, Theorem 1.2 yields optimal results for problems considered in several previous works on single-pass streaming lower bounds [KK15; KKS15; KKSV17; GT19; KK19; CGV20; CGS+22; CGSV24].

Translating Theorem 1.2 to MAX-CSP$_\alpha(\mathcal{F})$-type problems yields the following. For a predicate family $\mathcal{F} \subseteq \{0,1\}^{[q]^k}$, let $\alpha_{\mathcal{F}}^{\mathrm{LP}}$ be the best possible approximation ratio over all instances Φ i.e.,

$$\alpha_{\mathcal{F}}^{\mathrm{LP}} := \inf_{\Phi} \frac{\mathsf{opt}^{\mathrm{CSP}}(\Phi)}{\mathsf{opt}^{\mathrm{LP}}(\Phi)}.$$

Then Theorem 1.2 immediately yields the following:

Corollary 1.3. *For every $\varepsilon > 0$ and every $\mathcal{F} \subseteq \{0,1\}^{[q]^k}$, any single-pass streaming algorithm for* MAX-CSP$_{\alpha_{\mathcal{F}}^{\mathrm{LP}}+\varepsilon}(\mathcal{F})$ *on n variables requires $\Omega_\varepsilon(n)$ space.*

As a special case of our results, we also get new lower bounds against many specific CSPs. In addition to previous examples for $k = 2$, we consider some examples with higer arity below.

Example 1.4. Let $k > 2$ but fix $q = 2$ (with $[q]$ identified with $\{0,1\}$).

- MAX-k-XOR: defined by predicates $f^b_{k\text{-XOR}}(x_1,\dots,x_k) = x_1 \oplus \cdots \oplus x_k \oplus b$ for $b \in \{0,1\}$.
- MAX-LTF^w: For $w \in \mathbb{R}^k$, the linear threshold function (LTF) predicates are the 2^k functions

$$f^{w,b}_{\mathrm{LTF}}(x_1,\dots,x_k) = \begin{cases} 1 & \text{if } (-1)^{x_1+b_1}\cdot w_1 + \cdots + (-1)^{x_k+b_k}\cdot w_k > 0, \\ 0 & \text{if } (-1)^{x_1+b_1}\cdot w_1 + \cdots + (-1)^{x_k+b_k}\cdot w_k < 0. \end{cases}$$

  The LTF is said to be balanced if $\mathbb{E}_x[f^{w,b}_{\mathrm{LTF}}] = 1/2$ (for all $b \in \{0,1\}^k$).
- MAX-EXACTLY-ℓ-OF-k: defined by the predicate $f_{\text{EXACTLY-}\ell\text{-OF-}k}(x) = 1[\mathrm{wt}(x) = \ell]$ ($\mathrm{wt}(\cdot)$ denotes the Hamming weight).

$\Diamond$

Corollary 1.5. *In all of the following, n denotes the number of variables.*

- *For every $k \in \mathbb{N}$ odd and $\varepsilon > 0$, any single-pass streaming algorithm for* MAX-kXOR$_{1,1/2+\varepsilon}$ *requires $\Omega_\varepsilon(n)$ space.*
- *If w denotes the weights of the balanced (approximation resistant) LTF constructed by Potechin [Pot19], then any single-pass streaming algorithm for* MAX-LTF$^w_{1,1/2+\varepsilon}$ *requires $\Omega_\varepsilon(n)$ space.*
- *For all $k \in \mathbb{N}$, $\ell \neq k/2$, and $\varepsilon > 0$, any single-pass streaming algorithm for* MAX-EXACTLY-ℓ-OF-$k_{1,\rho(\ell,k)+\varepsilon}$ *requires $\Omega_\varepsilon(n)$ space, where $\rho(\ell,k) := \max_{p\in[0,1]} p^\ell(1-p)^{k-\ell}$.*
- *For all $\gamma \in (\frac{1}{2}, 1]$ and $\beta = (3\gamma-1)/2$, any single-pass streaming algorithm for* MAX-DICUT$_{\gamma-\varepsilon,\beta+\varepsilon}$ *requires $\Omega_\varepsilon(n)$ space.*
- *For all $\gamma \in (\frac{1}{2}, 1]$ and $\beta = (2\gamma+1)/4$, any single-pass streaming algorithm for* MAX-2SAT$_{\gamma-\varepsilon,\beta+\varepsilon}$ *requires $\Omega_\varepsilon(n)$ space.*

We also present a more detailed comparison to previous work below.

1.3 Comparison with the prior work

The most general linear-space lower bounds for CSPs in prior work are by Chou, Golovnev, Sudan, Velingker, and Velusamy [CGS+22], based on a structural property that they called *width*, which we now describe. Put simply, a predicate $f : \mathbb{Z}_q^k \to \{0,1\}$ is *wide* if its support contains a translate of the "diagonal" line $\{b \cdot \mathbb{1} : b \in \mathbb{Z}_q\}$, where $\mathbb{1} = (1,\ldots,1) \in \mathbb{Z}_q^k$ is the all-1's vector. More generally, the *width* of f is

$$\omega(f) := \max_{y \in \mathbb{Z}_q^k} \left\{ \Pr_{b \in \mathbb{Z}_q} [f(y + b \cdot \mathbb{1}) = 1] \right\},$$

and the width of $\mathcal{F} \subseteq \{0,1\}^{[q]^k}$ is $\omega(\mathcal{F}) := \min_{f \in \mathcal{F}} \omega(f)$. The result of [CGS+22] is the following:

Theorem 1.6 ([CGS+22]). *For every $\mathcal{F} \subseteq \{0,1\}^{[q]^k}$ and $\varepsilon > 0$, every single-pass streaming algorithm for $\text{MAX-CSP}_{\omega(\mathcal{F})-\varepsilon,\rho(\mathcal{F})+\varepsilon}(\mathcal{F})$ requires $\Omega_\varepsilon(n)$ space, where*

$$\rho(\mathcal{F}) := \lim_{n \to \infty} \inf_{\Phi \text{ instance of } \text{MAX-CSP}(\mathcal{F}) \text{ on } n \text{ variables}} \left\{ \mathsf{opt}^{\text{CSP}}(\Phi) \right\}$$

is the best possible uniform lower bound on the optimum value of instances of $\text{MAX-CSP}(\mathcal{F})$.

The authors of [CGS+22] use this lower bound to prove linear-space streaming inapproximability for various CSPs, including MAX-qCOLORING, MAX-qUNIQUEGAMES, and MAX-qLESSTHAN (see [CGS+22, §4.1]). For instance, MAX-qCOLORING is wide (i.e., $\omega = 1$ for every member of the predicate family).

Note that the basic LP framework for approximability is stronger than the width criterion, in the following sense: every instance Φ of $\text{MAX-CSP}(\mathcal{F})$ has $\mathsf{opt}^{\text{LP}}(\Phi) \geq \omega(\mathcal{F})$. (The argument is just to set the local distribution $\mathcal{Y}_C$ for each constraint C to be uniform over the corresponding translated line; the marginals are consistent because they are in fact uniform.) Thus, all linear-space lower bounds from the work of [CGS+22] (i.e., via Theorem 1.6) can alternatively be derived from our Theorem 1.2.

Remark 1.7. For CSPs including MAX-qUNIQUEGAMES [CGS+22] do not apply Theorem 1.6 directly; rather, they identify a subfamily of predicates $\mathcal{F}' \subseteq \mathcal{F}$ and observe that the resulting lower bound for $\text{MAX-CSP}_{\omega(\mathcal{F}')-\varepsilon,\rho(\mathcal{F}')+\varepsilon}(\mathcal{F}')$ is a lower bound for $\text{MAX-CSP}_{\omega(\mathcal{F}')-\varepsilon,\rho(\mathcal{F}')+\varepsilon}(\mathcal{F})$ too. This trick is not needed to apply our Theorem 1.2, which (unlike Theorem 1.6) is "monotone" with respect to the predicate family. ◊

We also consider below concrete examples of CSPs where our bounds (i.e., Theorem 1.2) are substantially stronger than the bounds of [CGS+22] (i.e., Theorem 1.6). A predicate family $\mathcal{F}$ is said to be approximation resistant for a family of algorithms, if it is hard for the algorithms to distinguish the two cases in $\text{MAX-CSP}_{1,\rho(\mathcal{F})+\varepsilon}(\mathcal{F})$ (where $\rho(\mathcal{F})$ is a trivial lower bound). Note that in the case $q = 2$, the width of a predicate $f : \{0,1\}^k \to \{0,1\}$ can only be either $1/2$ or 1. Correspondingly, the [CGS+22] criterion (Theorem 1.2), applied to any family $\mathcal{F}$, either shows that for every $\varepsilon > 0$, $\text{MAX-CSP}_{1,\rho(\mathcal{F})+\varepsilon}(\mathcal{F})$ requires $\Omega_\varepsilon(n)$ space (approximation resistance), or it only shows that for every $\varepsilon > 0$, $\text{MAX-CSP}_{1/2-\varepsilon,\rho(\mathcal{F})+\varepsilon}(\mathcal{F})$ requires $\Omega_\varepsilon(n)$ space. We argue that this bound is sometimes too crude for many interesting approximation resistant CSPs.

Approximation-resistant CSPs which evade [CGS+22]. We observe that a predicate $f : \{0,1\}^k \to \{0,1\}$ is wide (i.e., has $\omega(f) = 1$) iff there is a point $x \in \{0,1\}^k$ such that $f(x) = f(\neg x) = 1$; otherwise, $\omega(f) = \frac{1}{2}$. We claim that all the CSPs in Example 1.4 have $\omega(\mathcal{F}) = \frac{1}{2}$, and therefore the [CGS+22] condition proves hardness only for the $\text{MAX-CSP}_{1/2-\varepsilon,\rho(\mathcal{F})+\varepsilon}(\mathcal{F})$ problem. Indeed:

- If k is odd and x is a string where the XOR of all bits is 1, then the XOR of all bits in $\neg x$ will be 0. Therefore, the MAX-kXOR problem has width $1/2$ (for odd k).
- If $\ell \neq k/2$ and x is a k-bit string with Hamming weight ℓ, then the Hamming weight of $\neg x$ cannot be ℓ (indeed, it must be $k - \ell$). Hence, the MAX-EXACTLY-ℓ-OF-k problem has width $1/2$ (for $\ell \neq k/2$).
- If w is a weight vector for any balanced LTF, then if $\sum_{\ell=1}^{k} (-1)^{x_\ell + b_\ell} w_\ell > 0$, we have

$$\sum_{\ell=1}^{k} (-1)^{-x_\ell + b_\ell} w_\ell < 0 \,.$$

  Therefore, the MAX-LTF^w problem has width $1/2$.

Therefore, Theorem 1.6 only asserts a linear-space lower bound against MAX-CSP$_{\frac{1}{2}-\varepsilon,\rho(\mathcal{F})+\varepsilon}(\cdot)$. In the cases of MAX-$k$XOR and MAX-LTF^w ($w$ balanced), this is trivial [3] , since already $\rho(\mathcal{F}) = \frac{1}{2}$. In the case of MAX-EXACTLY-ℓ-OF-k, one can show $\rho(\mathcal{F}) = \binom{k}{\ell} \max_{p\in[0,1]} p^\ell (1-p)^{k-\ell}$. For instance, $\rho(f_{\text{EXACTLY-2-OF-3}}) = 4/9$, and so [CGS+22] does give a linear-space lower bound for MAX-EXACTLY-2-OF-3$_{1/2-\varepsilon,4/9+\varepsilon}$. However, all of these problems are known to admit $(1, \rho(\mathcal{F}) + \varepsilon)$-LP integrality gaps (they are approximation resistant for the basic LP) [Pot19; STV25], and therefore we prove hardness for the MAX-CSP$_{1,\rho(\mathcal{F})+\varepsilon}(\mathcal{F})$ problem in each case (establishing the first three items in Corollary 1.5).

Gap problems and approximation curves. For other interesting problems, the existing lower bounds of [KK19; CGS+22] do yield the best possible linear space lower bounds in terms of the overall approximation factor, but, as discussed in [CGS+22, §1.2.1], LP-based lower bounds (as in our Theorem 1.2, or the lower bounds of [FMW26a]) give "finer-grained" lower bounds. For instance, the MAX-DICUT predicate has width $\omega(f_{\text{DICUT}}) = 1/2$ (since it only has one satisfying assignment), and since its trivial approximability is $\rho(f_{\text{DICUT}}) = 1/4$, the [CGS+22] result gives a linear-space lower bound for MAX-DICUT$_{1/4+\varepsilon,1/2-\varepsilon}$.[4] Based on the LP analysis in [FMW26a, §1.2.1], our new lower bound (Theorem 1.2) gives a linear-space lower bound against MAX-DICUT$_{\gamma-\varepsilon,(3/2)\gamma-1/2+\varepsilon}$ for every $\gamma \in (\frac{1}{2}, 1]$; for instance, at $\gamma = \frac{3}{4}$, we get a linear-space lower bound against MAX-DICUT$_{3/4-\varepsilon,5/8+\varepsilon}$.

1.4 Concurrent and independent work

Very recently, two beautiful results, independent of the current work, have also made significant progress towards resolving the approximability of CSPs in the streaming model.

The first is a work of Fei, Minzer, and Wang [FMW26b] which proves a linear-space lower bound which also applies to the multi-pass model. For the single-pass model, their results are analogous to the ones shown here, yielding a linear-space lower bound from integrality gaps for the basic LP. Moreover, both our results and [FMW26b] rely on the Fourier-pseudorandomness framework of Kapralov and Krachun [KK19], using the analytic conditions obtained from the integrality gap

[3] Actually, there is an *ad hoc* reduction from MAX-$(2k-1)$XOR to MAX-$2k$XOR which does yield approximation resistance of MAX-$(2k-1)$XOR, since Theorem 1.6 does give approximation resistance for even k: Create a single new dummy variable v' and add it into every constraint. to get a new instance Φ' of MAX-$2k$XOR (Φ' has exactly the same optimal value as Φ). The other CSPs we have listed here do not appear to be handled by such *ad hoc* arguments.

[4] In fact, [CGS+22] is not needed; the same lower bound can be achieved simply by taking the lower bound instances for MAX-CUT$_{1/2+\varepsilon,1}$ from [KK19] and randomly directing the edges.

instances to obtain Fourier ℓ_1 bounds on the distributions over assignments maintained by the streaming algorithm (see proof overview below). However, the results of [FMW26b] also imply a $n \cdot 2^{-O(p)}$ space lower bound for algorithms with $p = o(\log n)$ passes. In addition to Fourier bounds above, their work also uses a decomposition of protocols into structured vs. pseudorandom components (following [FMW25; FMW26a]).

The second result, by Azarmehr, Behnezhad, and Ferrante [ABF26], shows that if the approximation ratio achieved by the basic LP is $\alpha < 1$, then an $\alpha - \varepsilon$ approximation can be obtained $O(n^{1-\Omega_\varepsilon(1)})$ space in a single pass, *even for instances with unbounded degree*. This can be seen as complimentary to our results (linear-space lower bounds for any better approximations) and proves that the characterization obtained in Theorem 1.2 for single-pass algorithms is indeed optimal.

1.5 Overview of proofs and techniques

We present a brief overview of the techniques used in our proof, and also discuss their connection to related works on lower bounds for streaming problems. Particularly relevant to our work are the recent dichotomy results for *multi-pass* streaming algorithms by Fei, Minzer and Wang [FMW26a] which prove an $\Omega(n^{1/3}/p)$ space lower bound for p-pass streaming algorithms using LP integrality gaps, and the linear space (single-pass) lower bounds by Chou et al. [CGS+22] for a *subclass* of k-CSPs with certain algebraic structure. While the setup for the problem we consider is similar to [FMW26a], the structure of our proof significantly extends the framework of Chou et al. [CGS+22], which itself builds on an earlier result of Kapralov and Krachun [KK19] for Max-CUT.

Distribution labeled hypergraphs from LP instances. Recall that for a predicate family $\mathcal{F} \subseteq \{f : [q_0]^k \to \{0,1\}\}$, an instance Φ of MAX-CSP($\mathcal{F}$) on (say) n_0 variables is specified by a collection of constraints $C = (f, e)$, where $e = (i_1, \ldots, i_k) \in [n_0]^k$ denotes an ordered k-uniform hyperedge on which the predicate f is applied. The basic LP corresponding to Φ (see Fig. 1) can be seen to search for *local distributions* $\mathcal{Y}_C$ on $[q_0]^k$, for each constraint $C \in \Phi$, such that the any two local distribution sharing a variable agree on the marginal distribution of the variable.

For $0 < \beta < \gamma \leq 1$, a (γ, β)-integrality gap instance Φ is one where the LP optimum is γ but no true assignment to the variables achieves a value higher than β. As is often the case for translating LP/SDP lower bounds to other computational models [Rag08; CLRS16; Lee15; KMR22; GT17; FMW26a], the integrality gap instance can be used as a (constant-sized) "gadget" or "template" to produce an infinite family of hard instances for the streaming problem. We view the integrality gap instances as ordered k-uniform hypergraphs on the vertex set $[n_0]$ with each edge corresponding to a constraint C labeled by the corresponding distribution $\mathcal{Y}_C$.

Reducing to one-wise uniform distributions. A very useful observation of Fei, Minzer and Wang [FMW26a] is that by suitably increasing the alphabet size q_0 to (say) q, the marginal distribution of each variable can be thought of as *uniform* on $[q]$ (a similar construction is also used by Yoshida [Yos11]). This is because for rational LP solutions with all probabilities multiples of (say) $1/q$, the uniform distribution on $[q]$ can be used to generate any distribution with probabilities $p_1, \ldots, p_{q_0}$: we partition $[q]$ into sets $S_1, \ldots S_{q_0}$ of size $|S_j| = p_j \cdot q$ and use the induced distribution on the sets. Each $\mathcal{Y}_C$ then gives a distribution $\mathcal{Y}'_C$ on $[q]^k$, where $a \sim \mathcal{Y}_C$ corresponds to a random sample from $S_{a_1} \times \cdots \times S_{a_k}$, with the marginals of all k variables in $\mathcal{Y}'_C$ being uniform on $[q]$. We refer to such distributions as *one-wise uniform*, and will only need to work with these for our proof.

Note that this reduction formally *changes* the predicate family to $\mathcal{F}' = \{f \circ \phi \mid f \in \mathcal{F}\}$, where $\phi : [q] \to [q_0]$ denotes the partition $\phi(z_\ell) = j \Leftrightarrow z_\ell \in S_j$. However, instances of MAX-CSP($\mathcal{F}'$)

generated via this reduction encode instances of MAX-CSP($\mathcal{F}$) and a good (streaming) approximation for instances of the former implies an approximation for MAX-CSP($\mathcal{F}$). We will suppress this distinction below, and will simply refer to the (new) predicate family as $\mathcal{F}$ and one-wise uniform distributions as $\mathcal{Y}_C$ for ease of notation, and will only consider these in the discussion below. We also identify the new alphabet $[q]$ with the set $\mathbb{Z}_q$.

Signal vs. noise problems. As in all previous works [KK19; CGSV24; CGS+22; FMW26a], we consider an average-case version of the problem MAX-CSP$_{\gamma,\beta}(\mathcal{F})$ where the goal is to distinguish if the random instances are from a **YES** distribution corresponding to a hidden signal (a planted assignment) or from a **NO** distribution with random noise. Following [FMW26a], who also considered the same average-case problem, we refer to it as the "Distributional Implicit Hidden Partition" (DIHP) problem.

To define the relevant distributions, we consider "literal shifts" $f_z(x) = f(x - z)$ for predicates $f \in \mathcal{F}$ (note that $\mathcal{F}$ *may not* be closed under literal shifts). For a fixed (hidden) assignment $X \in \mathbb{Z}_q^n$ and $y \in [q]^k$, note that if we consider a constraint (f_z, e) for $z = X_e - y$ (where X_e denotes the restriction of X to e) then $f_z(X_e) = f(y)$. Consider a random collection Λ of such constraints, with each constraint being generated independently as follows: (i) choose a random $\boldsymbol{C} = (\boldsymbol{f}, \boldsymbol{e}_0) \in \Phi$ from the gadget LP integrality gap instance and a random $\boldsymbol{y} \sim \mathcal{Y}_C$, (ii) for a randomly drawn $\boldsymbol{e} \in [n]^k$ include the constraint (e, f_z) with $z = X_e - y$. Then expected fraction of constraints in Λ satisfied by X equals

$$\mathop{\mathbb{E}}_{(\boldsymbol{e},\boldsymbol{f}_{X_e-y})\sim\Lambda}[f_{X_e-y}(X_e)] \;=\; \mathop{\mathbb{E}}_{C\sim\Phi}\mathop{\mathbb{E}}_{y\sim\mathcal{Y}_C}[f_{y-X_e}(X_e)] \;=\; \mathop{\mathbb{E}}_{C\sim\Phi}\mathop{\mathbb{E}}_{y\sim\mathcal{Y}_C}[f(y)]\,,$$

which equals the LP value γ on the instance Φ. Moreover, this remains true in expectation over the choice of $\boldsymbol{X}$, *even if restrict to the subcollection* Λ_0 *of constraints where* $y = X_e$ (so that the literal shift is 0 and we remain within the family $\mathcal{F}$). This is because we will have $|\Lambda_0| \approx |\Lambda|/q^k$ with high probability, and the expected fraction of constraints from Λ included in Λ_0 and satisfied by $\boldsymbol{X}$ is

$$\mathop{\mathbb{E}}_{(\boldsymbol{e},\boldsymbol{f}_{X_e-y})\sim\Lambda}\left[f_{X_e-y}(X_e)\cdot\mathbb{1}_{(\boldsymbol{e},\boldsymbol{f}_{\boldsymbol{y}-X_e})\in\boldsymbol{\Lambda}_0}\right] \;=\; \mathop{\mathbb{E}}_{C\sim\Phi}\mathop{\mathbb{E}}_{y\sim\mathcal{Y}_C}\left[f(y)\cdot\mathop{\mathbb{E}}_{X\sim\mathbb{Z}_q^n}\mathbb{1}_{y=X_e}\right] \;=\; \frac{1}{q^k}\cdot\mathop{\mathbb{E}}_{C\sim\Phi}\mathop{\mathbb{E}}_{y\sim\mathcal{Y}_C}[f(y)]$$

We consider an average-case version of the problem MAX-CSP$_{\gamma,\beta}(\mathcal{F})$, where the goal is to distinguish between the **YES** case of instances generated as above, with a "planted" random $X \sim \mathsf{Unif}(\mathbb{Z}_q^n)$ satisfying γ fraction of the constraints), and the **NO** case where each literal shift z is chosen uniformly from $\mathbb{Z}_q^k$. Also as in previous works [CGSV24; FMW26a] and with an eye towards the relevant communication problem described next, we will think of instances as "padded" where we specify the entire collection of constraints in Λ, but the actual CSP instance only consists of the constraints in Λ_0 (which are easily identifiable). The value of the instances in the **NO** case can be seen to be equal to that of an assignment for the LP instance Φ, with with each variable rounded independently according to its marginal distribution (which is at most β). We will restrict the discussion below to the special case with a *single* one-wise uniform distribution $\mathcal{Y}$, which contains all our conceptual ideas.

Such "planted vs. random" problems are also considered in earlier works on lower bounds for streaming problems (including [KKS15; KK19; CGSV24; CGS+22]) and many other computational models for approximation and refutation of random CSPs (see [KMOW17] for an excellent summary). Of particular relevance to us are the parallels with works on lower bounds for Sum-of-Squares (SoS) relaxations of CSPs [BCK15; KMOW17], which similarly extend previous results based on algebraic structure of XOR predicates [Gri01; Sch08; Tul09] using the analytic structure of *pairwise* uniformity.

Communication games. As in the case of most streaming lower bounds [KKS15; KK19; CGSV24; CGS+22; FMW26a], we reduce to proving lower bounds for a related communication problem. We consider random instances on n variables above as consisting of T partial k-uniform hypermatchings (with each edge in $[n]^k$) of size $m = \alpha \cdot n$, for suitable constants $T \gg 1$ and $\alpha \ll 1$. This is used to define a communication game with T players, where for each $t \in [T]$, $\texttt{Player}_t$ receives a pair of random variables $(\boldsymbol{M}_t, \boldsymbol{Z}_t)$ where the variables $\boldsymbol{M}_t$ denote the random hypermatchings, and the variables $\boldsymbol{Z}_t \in \mathbb{Z}_q^{m \times k}$ are distributed differently in the **YES** and **NO** distributions:

YES: For a common $\boldsymbol{X} \sim \mathsf{Unif}(\mathbb{Z}_q^n)$ and independent $\boldsymbol{Y}_t \sim \mathcal{Y}^{\otimes m}$, we have $\boldsymbol{Z}_t = \Pi_{\boldsymbol{M}_t}\boldsymbol{X} - \boldsymbol{Y}_t$. Here, Π_M denotes the projection of $\boldsymbol{X}$ to the vertices in M and $\mathcal{Y}^{\otimes m}$ denotes the m-wise product of the distribution $\mathcal{Y}$ on $\mathbb{Z}_q^k$.

NO: For each $t \in [T]$, we independently sample $\boldsymbol{Z}_t \sim \mathsf{Unif}(\mathbb{Z}_q^{m \times k})$.

Note that the above is specialized to the case of a single one-wise uniform distribution $\mathcal{Y}$. For the general case, each matching $\boldsymbol{M}_t$ is labeled with a randomly chosen edge e_0 from the gadget hypergraph from the LP instance Φ with a corresponding distribution $\mathcal{Y}_{e_0}$, and we take $\boldsymbol{Z}_t = \Pi_{\boldsymbol{M}_t}\boldsymbol{X} - \boldsymbol{Y}_t$ for $\boldsymbol{Y}_t \sim \mathcal{Y}_{e_0}^{\otimes m}$. We also take the set of n variables to be n_0-partite, where n_0 denotes the number of vertices in the gadget hypergraph.

The communication model we consider is that of one-way broadcast communication, where the players speak in sequence $1, \ldots, T$, and their messages are then visible to all other players. We show a lower bound of $\Omega(n)$ on the total communication for any protocol for distinguishing the above distributions. This is the technical contribution of our work, and also where argument diverges from [FMW26a], who prove an $\Omega(n^{1/3})$ communication lower bound in the stronger number-in-hand model relevant for multi-pass streaming (and in fact there is an $O(\sqrt{n})$ *upper bound* in this model). Our goal is thus to show a much stronger lower bound in a weaker communication model.

Hybrid argument and posterior distributions. Let $\boldsymbol{Z}_t^{\mathcal{Y}}$ and $\boldsymbol{Z}_t^{N}$ denote the samples from $\mathbb{Z}_q^{m \times k}$ for the t-th player, drawn respectively from the **YES** and **NO** distributions. Also, let $\boldsymbol{S}_1^{\mathcal{Y}}, \ldots, \boldsymbol{S}_t^{\mathcal{Y}}$ denote the sequence of messages sent by the first t players when given inputs $(\boldsymbol{M}_1, \boldsymbol{Z}_1^{\mathcal{Y}}), \ldots, (\boldsymbol{M}_1, \boldsymbol{Z}_1^{\mathcal{Y}})$ from the **YES** distribution. We denote this sequence as $\boldsymbol{S}_{1:t}^{\mathcal{Y}}$ and define $\boldsymbol{S}_{1:t}^{N}$ analogously. We also assume that the in addition to their messages, the players can publish their matchings $\boldsymbol{M}_t$ (but not the samples $\boldsymbol{Z}_t$) "for free", which only makes our lower bound stronger.

Our goal is to show that for any protocol (which can be assumed to be deterministic using Yao's principle) with $o(n)$ communication, we must have that $\|\boldsymbol{S}_T^{\mathcal{Y}} - \boldsymbol{S}_T^{N}\|_{\mathsf{tvd}} \leq \delta$ for a suitably small δ (since the T-th player can be assumed to simply output **YES** or **NO**). We will instead show the stronger bound that $\|(\boldsymbol{M}_{1:T}, \boldsymbol{S}_{1:T}^{\mathcal{Y}}) - (\boldsymbol{M}_{1:T}, \boldsymbol{S}_{1:T}^{N})\|_{\mathsf{tvd}} \leq \delta$ where $\boldsymbol{M}_{1:T}$ denote the aggregated matchings. Analogous to [KK19; CGS+22], we proceed via a hybrid argument, showing that for each $t \in [T]$, when the first $t-1$ messages are drawn from the **YES** distribution, we must have

$$\|\texttt{Player}_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}_{1:t-1}^{\mathcal{Y}}, \boldsymbol{Z}_t^{\mathcal{Y}}) - \texttt{Player}_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}_{1:t-1}^{\mathcal{Y}}, \boldsymbol{Z}_t^{N})\|_{\mathsf{tvd}} \;\leq\; \delta/T\,.$$

Consider the case $t = 1$. In this case, the distributions of the random variables $\boldsymbol{Z}_1^{\mathcal{Y}} = \Pi_{\boldsymbol{M}_1}\boldsymbol{X} - \boldsymbol{Y}_1$ and $\boldsymbol{Z}_1^{N} \sim \mathsf{Unif}(\mathbb{Z}_q^{m \times k})$ are clearly identical, since we have $\boldsymbol{X} \sim \mathsf{Unif}(\mathbb{Z}_q^n)$. However, this is no longer true at $t = 2$ since the distribution of $\boldsymbol{X}$ is no longer uniform *conditioned on the output of* $\texttt{Player}_1$. The proofs in [KK19] and [CGS+22] track the *posterior* distribution of $\boldsymbol{X}$ conditioned on the output of players at various steps, and show that as long as the Fourier tails for the (density function of) the posterior distribution are bounded in ℓ_1 norm, the distributions of $\boldsymbol{Z}_t^{\mathcal{Y}}$ and $\boldsymbol{Z}_t^{N}$ are close, even conditioned on outputs of the previous players. However, they crucially rely on the

(linear) structure of the underlying predicates generating the distribution of $\boldsymbol{Y}_t$ for proving such Fourier bounds. While the results of [KK19] are for Max-CUT, the ones in [CGS+22] only hold for predicates $f : \mathbb{Z}_q^k \to \{0,1\}$ where $f^{-1}(1)$ contains a "line" $L_y = \{y + b \cdot \mathbb{1} \mid b \in \mathbb{Z}_q\}$, where $\mathbb{1} = (1,\ldots,1)$.

No such linear structure is available in our case since the goal is to prove lower bounds for *all* CSPs. The technical core of our argument is obtaining such Fourier bounds only relying on the (weak) *analytic* property that the distributions $\mathcal{Y}$ generating $\boldsymbol{Y}_t$ can be taken to be one-wise uniform. As mentioned earlier, this phenomenon also has a parallel in the literature on SoS lower bounds [KMOW17], where the technical challenge was working with the weaker analytic Fourier structure provided by pairwise uniformity (instead of the linear structure of XOR predicates).

Fourier bounds via one-wise uniformity. We now consider the key technical challenge, which arises in understanding the Fourier structure for the posterior distribution. For the case $t = 2$, let the output of $\texttt{Player}_1(\boldsymbol{M}_1, \boldsymbol{Z}_1)$ be in $\mathbb{Z}_q^r$ (it will be convenient to think of output also as q-ary). For a fixed matching M (value of $\boldsymbol{M}_1$) the outputs of $\texttt{Player}_1$ partition $\mathbb{Z}_q^{m\times k}$ (the domain of $\boldsymbol{Z}_1$). Fixing an output then defines a set $\mathcal{B} \subseteq \mathbb{Z}_q^{m\times k}$, which is typically of size q^{mk-r}. This induces a posterior distribution $\mathcal{D}_{M,\mathcal{B}}$ on $\mathbb{Z}_q^n$ with probabilities

$$\mathcal{D}_{M,\mathcal{B}}(X_0) \;=\; \Pr_{\boldsymbol{X}\sim\mathbb{Z}_q^n, \boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}} [\boldsymbol{X} = X_0 \;\mid\; \Pi_M \boldsymbol{X} - \boldsymbol{Y} \in \mathcal{B}] \;.$$

For a distribution $\mathcal{D}$ on $\mathbb{Z}_q^N$, let $\mu_{\mathcal{D}}$ denote its density $q^N \cdot \mathcal{D}$ (so that $\mathbb{E}[\mu_{\mathcal{D}}] = 1$) and let $\mu_{\mathcal{B}}$ be the density of $\mathsf{Unif}(\mathcal{B})$. The density $\mu_{\mathcal{D}_{M,\mathcal{B}}}$ can be expressed as $\mu_{\mathcal{D}_{M,\mathcal{B}}} = \mu_{\mathcal{B}} * \mu_{\mathcal{Y}^{\otimes m}}$. The Fourier coefficient for any "frequency" $U \in \mathbb{Z}_q^n$ is then given by

$$\widehat{\mu_{\mathcal{D}_{M,\mathcal{B}}}}(U) \;=\; \begin{cases} \widehat{\mu_{\mathcal{B}}}(\Pi_M U) \cdot \widehat{\mu_{\mathcal{Y}^{\otimes m}}}(-\Pi_M U) & \text{if } \mathsf{supp}(U) \subseteq M \\ 0 & \text{otherwise} \end{cases} .$$

The Fourier conditions needed on the posterior distribution require understanding, for each h, the sum $\sum_{\mathrm{wt}(U)=h} |\widehat{\mu_{\mathcal{D}_{M,\mathcal{B}}}}(U)|$. Writing $V = \Pi_M U \in \mathbb{Z}_q^{m\times k}$ as a matrix with rows $V_1,\ldots,V_m$, the argument of [CGS+22] can be seen as decomposing this sum based on the number of nonzero rows (say) ℓ in V. They then study the ℓ_2 norm of Fourier coefficients for V with exactly h nonzero entries and ℓ nonzero rows (together with Cauchy-Schwarz and counting to obtain ℓ_1 bounds). Note that this "(h,ℓ)-inequality" is a refinement of the usual level-h inequality, which bounds the Fourier mass of coefficients with h nonzero entries. Such inequalities are obtained in [CGS+22] by using the algebraic structure of predicates defining $\mathcal{Y}$ to eliminate one nonzero entry for each nonzero row.

While we do not have such algebraic structure, the one-wise uniformity of $\mathcal{Y}$ implies that

$$\widehat{\mu_{\mathcal{Y}^{\otimes m}}}(-V) = \prod_{i\in[m]} \widehat{\mu_{\mathcal{Y}}}(-V_i) \neq 0 \quad\Longrightarrow\quad \forall i \in [m],\ |\mathsf{supp}(V_i)| \neq 1.$$

Due to the product structure of Fourier coefficients in $\mu_{\mathcal{D}_{M,\mathcal{B}}}$, it then suffices to bound the Fourier mass of the function $\mu_{\mathcal{B}} : \mathbb{Z}_q^{m\times k} \to \mathbb{R}_{\geq 0}$ (with $\|\mu_{\mathcal{B}}\|_\infty \leq q^r$) on Fourier coefficients corresponding to matrices $V = \Pi_M U$ with "singleton-free" rows (since $\widehat{\mu_{\mathcal{Y}^{\otimes m}}}(-V) = 0$ otherwise).

To prove such a refinement of the level-h inequality, we first recall that usual inequality is proved by applying the Bonami–Beckner operator T_ρ, which has the same eigenvalue ρ^h for all Fourier coefficients of weight h, and then using hypercontractivity to bound $\|T_\rho f\|_2 \leq \|f\|_p$ for an appropriate p. To prove our "singleton-free" inequality, we now design a *custom noise operator*

with eigenvalues taking advantage of this singleton-free structure. In particular, consider the linear operator S_ρ defined on $\mathbb{Z}_q^k$ with characters $\chi_u(x) = \omega_q^{\langle u,x\rangle}$ as eigenvectors, and eigenvalues

$$S_\rho(\chi_u) \;=\; \lambda_\rho(u)\cdot\chi_u \qquad \text{where} \qquad \lambda_\rho(u) \;=\; \begin{cases} 1 & \text{if } u = 0 \\ 0 & \text{if } |\mathsf{supp}(u)| = 1 \\ \rho^{|\mathsf{supp}(u)|-1} & \text{otherwise} \end{cases}.$$

One can check that for characters χ_V corresponding to $V \in \mathbb{Z}_q^{m\times k}$, we have $S_\rho^{\otimes m}(\chi_V) = \rho^{h-\ell}\cdot\chi_V$, where $h = \mathsf{supp}(V)$ and ℓ is the number of (singleton-free) nonzero rows. We can also show that if the Fourier spectrum of $g : \mathbb{Z}_q^{m\times k} \to \mathbb{R}$ is singleton-free, then $\|S_\rho^{\otimes m} g\|_2 \le \|T_\rho^* g\|_2$, where T_ρ^* is a "row-wise" Bonami-Beckner operator, resampling inputs corresponding to entire rows in $\mathbb{Z}_q^k$, with probability $1-\rho$. This comparison also yields a hypercontractive inequality for our new operator, and we use this obtain the required (h,ℓ)-inequality.

This argument (described in §6) is a novel technical contribution of our work, and may also be useful for other Fourier-analytic applications. Beyond generalizing the bound in [CGS+22], it also significantly simplifies their proof, avoiding the careful elimination of entries in nonzero rows.

Covering with structured centers. While the above argument suffices for analyzing the posterior distribution at $t = 2$ (conditioning on the output of $\texttt{Player}_1$), understanding the Fourier structure at later times presents a slight twist on the above problem. If $\mathcal{D}_t$ is the posterior distribution after $\texttt{Player}_t$ speaks, M is the revealed matching for $\texttt{Player}_{t+1}$ and $\mathcal{B} \subseteq \mathbb{Z}^{m\times k}$ is a set corresponding to the output of $\texttt{Player}_{t+1}$, the posterior distribution $\mathcal{D}_{t+1,M,\mathcal{B}}$ is given by

$$\mathcal{D}_{t+1,M,\mathcal{B}}(X_0) \;=\; \Pr_{\boldsymbol{X}\sim\mathcal{D}_t,\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}\left[\boldsymbol{X} = X_0 \;\middle|\; \Pi_M\boldsymbol{X} - \boldsymbol{Y} \in \mathcal{B}\right],$$

where we now condition the distribution $\mathcal{D}_t$ instead of the uniform distribution on $\mathbb{Z}_q^n$. The new density is now given by $\mu_{\mathcal{D}_{t+1,M,\mathcal{B}}} = \nu_{M,\mathcal{B}}\cdot\mu_{\mathcal{D}_t}\cdot(\mu_{\mathcal{B}} * \mu_{\mathcal{Y}^{\otimes m}})$, where $\nu_{M,\mathcal{B}}$ is a normalizing factor, and the Fourier coefficients are given by

$$\widehat{\mu_{\mathcal{D}_{t+1,M,\mathcal{B}}}}(W) \;=\; C_{M,\mathcal{B}}\cdot \sum_{\substack{U\in\mathbb{Z}_q^n \\ \mathsf{supp}(W-U)\subseteq M}} \widehat{\mu_{\mathcal{D}_t}}(U)\cdot\widehat{\mu_{\mathcal{B}}}(\Pi_M(W-U))\cdot\widehat{\mu_{\mathcal{Y}^{\otimes m}}}(-\Pi_M(W-U))$$

Assuming the Fourier ℓ_1 norm of $\mu_{\mathcal{D}_t}$ can be controlled using an induction hypothesis, understanding the norm for $\mu_{t+1,M,\mathcal{B}}$ now requires bounding, for each U, expressions of the form

$$\sum_{\substack{\mathrm{wt}(W)=h \\ \mathsf{supp}(W-U)\subseteq M}} \left|\widehat{\mu_{\mathcal{B}}}(\Pi_M(W-U))\cdot\widehat{\mu_{\mathcal{Y}^{\otimes m}}}(-\Pi_M(W-U))\right|$$

While we can again study this expression by using a fine-grained weight inequality in terms of the number of non-zero rows in $\Pi_M(W-U)$, the weight constraint is now on W, while the relevant row structure is only available for $\Pi_M(W-U)$. To handle this mismatch, we need to "cover" such a sum for an arbitrary U using a similar sum for "structured centers" U', such that for W' defined by $W'-U' = W-U$, we have that W' itself is singleton-free and we can also bound $\mathrm{wt}(W')$ (with smaller values of $h' < h$). [5] We rely on the properties of the random matching M to bound the number of structured centers used in our covering, leading to the desired inductive Fourier bound.

[5] A similar covering argument was actually also used in [CGS+22], though it's purpose there was only to reduce the weight h. We need to modify this argument to provide the required structure for U' and W'.

2 Preliminaries

2.1 General notation

We use **boldface** to represent random variables. For an event E, we also use $[\![E]\!]$ to denote the indicator for the event. For a function $g : \mathbb{Z}_q^N \to \mathbb{C}$ and $p \in [1, \infty]$, we use the convention that

$$\|g\|_p := \left(\mathop{\mathbb{E}}_{\boldsymbol{X} \in \mathbb{Z}_q^N} [|g(\boldsymbol{X})|^p] \right)^{1/p}.$$

In particular, $\|g\|_1 = \mathbb{E}_{\boldsymbol{X} \in \mathbb{Z}_q^N}[|g(X)|]$ and $\|g\|_\infty = \max_{X \in \mathbb{Z}_q^N} |g(X)|$.

2.2 Distributions

For a finite set S, we use $\Delta(S)$ to denote the set of all distributions over S. If S is a finite set and S^k its k-fold product set, a distribution $\mathcal{D} \in \Delta(S^k)$ is *one-wise uniform* if for $(\boldsymbol{s}_j)_{j \in [k]} \sim \mathcal{D}$, the marginal distribution of $\boldsymbol{s}_j$ is uniform on S for every $j \in [k]$. We use $\Delta_1(S^k)$ to denote the set of all one-wise uniform distributions.

For a distribution $\mathcal{D}$, $\mathcal{D}^{\otimes \ell}$ denotes the ℓ-wise product distribution $\mathcal{D} \times \cdots \times \mathcal{D}$. In particular, if $\mathcal{D} \in \Delta(\mathbb{Z}_q^k)$, we can view $\mathcal{D}^{\otimes m} \in \Delta(\mathbb{Z}_q^{m \times k})$ as a distribution on matrices.

Abusing notation, we can view a distribution $\mathcal{D} \in \Delta(S)$ as a function $\mathcal{D} : S \to \mathbb{R}$, where $\mathcal{D}(x)$ represents the probability that a sample $\boldsymbol{x} \sim \mathcal{D}$ satisfies $\boldsymbol{x} = x$. We can then define the *density* function $\mu_{\mathcal{D}} : S \to \mathbb{R}$ via

$$\mu_{\mathcal{D}}(x_0) := \mathcal{D}(x_0) \cdot |S|.$$

One useful property of density functions is:

Proposition 2.1. *Let $\mathcal{D} \in \Delta(S)$ be a distribution and $\phi : S \to \mathbb{C}$ a function. Then*

$$\mathop{\mathbb{E}}_{\boldsymbol{x} \in S}[\mu_{\mathcal{D}}(\boldsymbol{x}) \cdot \phi(\boldsymbol{x})] = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{D}}[\phi(\boldsymbol{x})].$$

Proof. We can expand

$$\mathop{\mathbb{E}}_{\boldsymbol{x} \in S}[\mu_{\mathcal{D}}(\boldsymbol{x}) \cdot \phi(\boldsymbol{x})] = |S| \mathop{\mathbb{E}}_{\boldsymbol{x} \in S}[\mathcal{D}(\boldsymbol{x}) \cdot \phi(\boldsymbol{x})] = \sum_{x \in S} \mathcal{D}(x) \cdot \phi(x) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{D}}[\phi(\boldsymbol{x})]. \qquad \square$$

2.3 Total variation distance

Proposition 2.2 (Data processing inequality)**.** *For random variables $\boldsymbol{X}, \boldsymbol{Y}$ and $\boldsymbol{W}$, if $\boldsymbol{W}$ is independent of both $\boldsymbol{X}$ and $\boldsymbol{Y}$, and f is a function, then $\|f(\boldsymbol{X}, \boldsymbol{W}) - f(\boldsymbol{Y}, \boldsymbol{W})\|_{\mathsf{tvd}} \le \|\boldsymbol{X} - \boldsymbol{Y}\|_{\mathsf{tvd}}$.*

Proposition 2.3 ([CGS+22, Lemma 2.3])**.** *Let $\boldsymbol{X}$, $\boldsymbol{Y}$, $\boldsymbol{W}$ be random variables and let f be a function. If there exists $\delta > 0$ such that for every fixed X_0 in the support of $\boldsymbol{X}$, we have*

$$\|f(X_0, \boldsymbol{Y}) - f(X_0, \boldsymbol{W})\|_{\mathsf{tvd}} \le \delta\,,$$

then the following holds:

$$\|(\boldsymbol{X}, f(\boldsymbol{X}, \boldsymbol{Y})) - (\boldsymbol{X}, f(\boldsymbol{X}, \boldsymbol{W}))\|_{tvd} \le \delta\,.$$

Proposition 2.4 ([KK19, Lemma B.2])**.** *Let $\boldsymbol{X}^1, \boldsymbol{X}^2$ be random variables taking values on the same sample space, let $\boldsymbol{Z}^1, \boldsymbol{Z}^2$ be random variables taking values on the same sample space, and let f be a function. If $\boldsymbol{Z}^2$ is independent of $\boldsymbol{X}^1, \boldsymbol{X}^2$, then*

$$\|(\boldsymbol{X}^1, f(\boldsymbol{X}^1, \boldsymbol{Z}^1)) - (\boldsymbol{X}^2, f(\boldsymbol{X}^2, \boldsymbol{Z}^2))\|_{\mathsf{tvd}} \le \|(\boldsymbol{X}^1, f(\boldsymbol{X}^1, \boldsymbol{Z}^1)) - (\boldsymbol{X}^1, f(\boldsymbol{X}^1, \boldsymbol{Z}^2))\|_{\mathsf{tvd}} + \|\boldsymbol{X}^1 - \boldsymbol{X}^2\|_{\mathsf{tvd}}\,.$$

2.4 Convolutions

If $f, g : \mathbb{Z}_q^N \to \mathbb{C}$ are functions, their *convolution* $f * g : \mathbb{Z}_q^N \to \mathbb{C}$ is defined as

$$(f * g)(z) := \mathop{\mathbb{E}}_{\boldsymbol{x} \in \mathbb{Z}_q^N}[f(\boldsymbol{x}) \cdot g(z - \boldsymbol{x})].$$

We now give two propositions which help interpret convolutions. The first shows that convolving a function f with a density function can be viewed as "noising" f; the second shows that convolving two density functions corresponds to the density function of a sum distribution.

Proposition 2.5. *Let $g : \mathbb{Z}_q^N \to \mathbb{C}$ and $\mathcal{D} \in \Delta(\mathbb{Z}_q^N)$. Then*

$$(g * \mu_{\mathcal{D}})(z) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{D}}[g(z - \boldsymbol{x})].$$

Proof. Apply Proposition 2.1 with the function $\phi(x) := g(z - x)$. □

Proposition 2.6. *Let $\mathcal{D}_1, \mathcal{D}_2 \in \Delta(\mathbb{Z}_q^N)$. Let $\mathcal{D} \in \Delta(\mathbb{Z}_q^N)$ denote the marginal distribution of $\boldsymbol{x} + \boldsymbol{y}$ when $\boldsymbol{x} \sim \mathcal{D}_1, \boldsymbol{y} \sim \mathcal{D}_2$ independently, so that for $z \in \mathbb{Z}_q^N$, $\mathcal{D}(z) = \Pr_{\boldsymbol{x} \sim \mathcal{D}_1, \boldsymbol{y} \sim \mathcal{D}_2}[\boldsymbol{x} + \boldsymbol{y} = z]$. Then the density function of $\mathcal{D}$ is the convolution of the density functions of $\mathcal{D}_1$ and $\mathcal{D}_2$: For every $z \in \mathbb{Z}_q^N$,*

$$\mu_{\mathcal{D}}(z) = (\mu_{\mathcal{D}_1} * \mu_{\mathcal{D}_2})(z).$$

Proof. We have $\mathcal{D}(z) = \sum_{x,y \in \mathbb{Z}_q^N} \mathcal{D}_1(x) \cdot \mathcal{D}_2(y) \cdot [\![x + y = z]\!] = \sum_{x \in \mathbb{Z}_q^N} \mathcal{D}_1(z) \cdot \mathcal{D}_2(z - x)$. Correspondingly,

$$\mu_{\mathcal{D}}(z) = q^N \cdot \sum_{x \in \mathbb{Z}_q^N} \mu_{\mathcal{D}_1}(x) \cdot \mu_{\mathcal{D}_2}(z - x) = \mathop{\mathbb{E}}_{\boldsymbol{x} \in \mathbb{Z}_q^N}[\mu_{\mathcal{D}_1}(\boldsymbol{x}) \cdot \mu_{\mathcal{D}_2}(z - \boldsymbol{x})] = (\mu_{\mathcal{D}_1} * \mu_{\mathcal{D}_2})(z). \qquad \square$$

2.5 Fourier analysis

In this subsection, we use N for the dimension of a $\mathbb{Z}_q$-vector space (or module, if q is not a prime power) over which we perform Fourier analysis. We will use various settings of N within the paper, e.g., n, k, nk, km, nk', etc.

Definition 2.7 (Fourier characters)**.** For $u, x \in \mathbb{Z}_q^N$, we define the *character function*

$$\chi_u(x) := \omega_q^{\langle u, x \rangle},$$

where we fix the q-th root of unity $\omega_q := e^{2\pi i/q}$ and $\langle \cdot, \cdot \rangle$ represents the standard inner product modulo q. We will write $\omega = \omega_q$ when q is clear from context. ♢

Definition 2.8 (Fourier coefficients)**.** For any function $g : \mathbb{Z}_q^N \to \mathbb{C}$ and $u \in \mathbb{Z}_q^N$, we define the *Fourier coefficient*

$$\widehat{g}(u) := \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathbb{Z}_q^N}[g(\boldsymbol{x}) \cdot \overline{\chi_u(\boldsymbol{x})}],$$

where $\bar{\cdot}$ denotes complex conjugation, so that $\overline{\chi_u(\boldsymbol{x})} = \omega^{-\langle u, x \rangle}$. We sometimes refer to the vector of u as a *frequency*. ♢

Proposition 2.9 (Fourier coefficients of convolutions, e.g., [O'D14, Theorem 8.60])**.** *For every $f, g : \mathbb{Z}_q^N \to \mathbb{C}$ and $u \in \mathbb{Z}_q^N$,*

$$\widehat{f * g}(u) = \widehat{f}(u) \cdot \widehat{g}(u).$$

Conversely, $\widehat{f \cdot g}(u) = q^N (\widehat{f} * \widehat{g})(u) = \sum_{v \in \mathbb{Z}_q^N} f(v) g(u - v)$.

Proposition 2.10 (Expansion in the Fourier basis). *For every $f : \mathbb{Z}_q^N \to \mathbb{C}$ and $x \in \mathbb{Z}_q^N$, we have*

$$f(x) = \sum_{u \in \mathbb{Z}_q^N} \widehat{f}(u) \cdot \chi_u(x).$$

Proposition 2.11 (Parseval's identity). *For every $g : \mathbb{Z}_q^N \to \mathbb{C}$,*

$$\|g\|_2^2 = \sum_{U \in \mathbb{Z}_q^N} |\widehat{g}(U)|^2.$$

Proposition 2.12 (Fourier coefficients of (density functions of) distributions). *If $\mathcal{D} \in \Delta(\mathbb{Z}_q^N)$, then*

$$\widehat{\mu_{\mathcal{D}}}(u) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{D}} \left[\overline{\chi_u(\boldsymbol{x})} \right].$$

Proof. Combine Proposition 2.1 and Definition 2.8. □

Proposition 2.13. *For every $q \geq 2 \in \mathbb{N}$,*

$$\mathop{\mathbb{E}}_{\boldsymbol{b} \in \mathbb{Z}_q} [\omega_q^{\boldsymbol{b}}] = 0.$$

Proof. The geometric sum formula gives $\sum_{b=0}^{q-1} \omega^b = \frac{1-\omega^q}{1-\omega} = 0$ since $\omega^q = 1$. □

Proposition 2.14 (Fourier coefficients of one-wise uniform distributions). *Let $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^N)$ be a one-wise uniform distribution. If $u \in \mathbb{Z}_q^N$ has $|\mathsf{supp}(u)| = 1$, then*

$$\widehat{\mu_{\mathcal{Y}}}(u) = 0.$$

Proof. Suppose $j \in [N]$ is the single nonzero coordinate in u. Then,

$$\widehat{\mu_{\mathcal{Y}}}(u) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{Y}} \left[\overline{\chi_u(\boldsymbol{x})} \right] = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{Y}} \left[\omega^{-\langle u, \boldsymbol{x} \rangle} \right] = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{Y}} \left[\omega^{-u_j \cdot \boldsymbol{x}_j} \right] = \mathop{\mathbb{E}}_{\boldsymbol{b} \in \mathbb{Z}_q} \left[\omega^{-u_j \cdot \boldsymbol{b}} \right] = \mathop{\mathbb{E}}_{\boldsymbol{b} \in \mathbb{Z}_q} [\omega^{\boldsymbol{b}}] = 0,$$

where we used, respectively, Proposition 2.1, the definition of χ, that u is supported only on j, one-wise uniformity, the nonzeroness of u_j, and Proposition 2.13. □

Proposition 2.15 (Fourier coefficients of product distributions). *Let $\mathcal{Y} \in \Delta(\mathbb{Z}_q^k)$. Then for every $U \in \mathbb{Z}_q^{k \times m}$, we have:*

$$\widehat{\mu_{\mathcal{Y}^{\otimes m}}}(U) = \prod_{j=1}^{m} \widehat{\mu_{\mathcal{Y}}}(U_j).$$

Proposition 2.16 (Fourier coefficients of marginals). *Let $\widetilde{N}, N \in \mathbb{N}$, $\Pi : \mathbb{Z}_q^{\widetilde{N}} \to \mathbb{Z}_q^N$ be ($\mathbb{Z}_q$-)linear and surjective, $\widetilde{g} : \mathbb{Z}_q^{\widetilde{N}} \to \mathbb{C}$ be any function, and $g : \mathbb{Z}_q^N \to \mathbb{C}$ be defined by $g(x) := \mathbb{E}_{\widetilde{\boldsymbol{x}} \in \mathbb{Z}_q^{\widetilde{N}}} [g(\widetilde{\boldsymbol{x}}) \mid \Pi(\widetilde{\boldsymbol{x}}) = x]$. Then for every $u \in \mathbb{Z}_q^N$,*

$$\widehat{g}(u) = \widehat{\widetilde{g}}(\Pi^{\intercal}(u)),$$

where $\Pi^{\intercal}$ is the (standard) transpose of Π (which is adjoint with respect to the standard inner product).

Proof. Let $u \in \mathbb{Z}_q^N$. Note that for every $\widetilde{x} \in \mathbb{Z}_q^{\widetilde{N}}$, we have $\langle u, \Pi(\widetilde{x})\rangle = \langle \Pi^\intercal(u), \widetilde{x}\rangle$ by adjointness, and hence $\chi_u(\Pi(\widetilde{x})) = \chi_{\Pi^\intercal(u)}(\widetilde{x})$. Hence,

$$\widehat{g}(u) = \mathop{\mathbb{E}}_{x\in\mathbb{Z}_q^N}\left[g(x)\cdot\overline{\chi_u(x)}\right] = \mathop{\mathbb{E}}_{x\in\mathbb{Z}_q^N}\left[\mathop{\mathbb{E}}_{\widetilde{x}\in\mathbb{Z}_q^{\widetilde{N}}}[\widetilde{g}(\widetilde{x}) \mid \Pi(\widetilde{x}) = x]\cdot\overline{\chi_u(x)}\right] = \mathop{\mathbb{E}}_{\widetilde{x}\in\mathbb{Z}_q^{\widetilde{N}}}\left[\widetilde{g}(\widetilde{x})\cdot\overline{\chi_u(\Pi(\widetilde{x}))}\right]$$
$$= \mathop{\mathbb{E}}_{\widetilde{x}\in\mathbb{Z}_q^N}\left[\widetilde{g}(\widetilde{x})\cdot\overline{\chi_{\Pi^\intercal(u)}(\widetilde{x})}\right] = \widehat{\widetilde{g}}(\Pi^\intercal(u)),$$

where the third equality uses that for $\Pi : \mathbb{Z}_q^{\widetilde{N}} \to \mathbb{Z}_q^N$ surjective, the marginal distribution of $\Pi(\widetilde{x})$ for uniform $\widetilde{x} \in \mathbb{Z}_q^{\widetilde{N}}$ is uniform on $\mathbb{Z}_q^N$. □

Proposition 2.17 (Fourier coefficients of coordinate-projected functions)**.** *Let $N \in \mathbb{N}$, $\mathcal{S} \subseteq [N]$ a subset of the coordinates, $g : \mathbb{Z}_q^{\mathcal{S}} \to \mathbb{C}$ be any function, and $\Pi_{\mathcal{S}} : \mathbb{Z}_q^N \to \mathbb{Z}_q^{\mathcal{S}}$ the standard projection map onto the coordinates in $\mathcal{S}$. Then for every $u \in \mathbb{Z}_q^N$,*

$$\widehat{g\circ\Pi_{\mathcal{S}}}(u) = \begin{cases}\widehat{g}(\Pi_{\mathcal{S}}(u)) & \text{if } \mathsf{supp}(u)\subseteq\mathcal{S},\\ 0 & \text{otherwise.}\end{cases}$$

Proof. Let $u \in \mathbb{Z}_q^N$. We write e.g. $u_{\mathcal{S}} := \Pi_{\mathcal{S}}(u)$ for short. We have

$$\widehat{g\circ\Pi_{\mathcal{S}}}(u) = \mathop{\mathbb{E}}_{x\in\mathbb{Z}_q^N}[g(x_{\mathcal{S}})\cdot\omega^{-\langle u,x\rangle}] = \mathop{\mathbb{E}}_{\overline{x}\in\mathbb{Z}_q^{\mathcal{S}}}\left[g(\overline{x})\cdot\mathop{\mathbb{E}}_{x\in\mathbb{Z}_q^N}[\omega^{-\langle u,x\rangle} \mid x_{\mathcal{S}} = \overline{x}]\right]$$

where we used the fact that the marginal distribution of $x_{\mathcal{S}}$ for $x \in \mathbb{Z}_q^N$ is uniform on $\mathbb{Z}_q^{\mathcal{S}}$. Now fix $\overline{x} \in \mathbb{Z}_q^{\mathcal{S}}$; we claim that

$$\mathop{\mathbb{E}}_{x\in\mathbb{Z}_q^N}[\omega^{-\langle u,x\rangle} \mid x_{\mathcal{S}} = \overline{x}] = \begin{cases}\omega^{-\langle u_{\mathcal{S}},\overline{x}\rangle} & \text{if } \mathsf{supp}(u)\subseteq\mathcal{S}\\ 0 & \text{otherwise,}\end{cases}$$

which obviously proves the proposition. Indeed, if $\mathsf{supp}(u) \subseteq \mathcal{S}$, then $\langle u, x\rangle = \langle u_{\mathcal{S}}, x_{\mathcal{S}}\rangle$ and hence $\omega^{-\langle u,x\rangle} = \omega^{-\langle u_{\mathcal{S}},x_{\mathcal{S}}\rangle}$ (for every $x \in \mathbb{Z}_q^N$). On the other hand, if $\mathsf{supp}(u) \not\subseteq \mathcal{S}$, then picking some $i \in \mathsf{supp}(u) \setminus \mathcal{S}$, we can write $\langle u, x\rangle = u_i \cdot x_i + \langle u_{[N]\setminus i}, x_{[N]\setminus i}\rangle$, and therefore

$$\mathop{\mathbb{E}}_{x\in\mathbb{Z}_q^N}[\omega^{-\langle u,x\rangle} \mid x_{\mathcal{S}} = \overline{x}] = \mathop{\mathbb{E}}_{b\in\mathbb{Z}_q}\left[\omega^{-u_i\cdot b}\mathop{\mathbb{E}}_{x\in\mathbb{Z}_q^N}[\omega^{-\langle u_{[N]\setminus\{i\}},x_{[N]\setminus\{i\}}\rangle} \mid x_{\mathcal{S}} = \overline{x} \text{ and } x_j = b]\right]$$

which again vanishes by Proposition 2.13 (and the fact that $u_i \neq 0$). □

Remark 2.18. Proposition 2.17 can be easily generalized to the case where $\Pi : \mathbb{Z}_q^N \to \mathbb{Z}_q^{N'}$ is any surjective linear map, but we will not need it in this generality, so choose to keep its current form for simplicity in notation. ◊

3 The communication problem and streaming reduction

We now turn to defining the communication problem used in our lower bound. This is based on the communication problem defined in [FMW26a], as modified by [STV25] for the single-pass setting.

Definition 3.1 (Partite hypermatchings)**.** For $n, m, k \in \mathbb{N}$, we define the set of k-uniform partite hypermatchings with m edges on a fixed k-partite vertex-set $[n] \times [k]$ as:

$$\mathcal{M}_n^{m,k} := \{M \in [n]^{m \times k} : \text{all entries in each column of } M \text{ are distinct}\}.$$

A k-partite hypermatching $M \in \mathcal{M}_n^{m,k}$ gives rise to an injective mapping $\iota_M : [m] \times [k] \to [n] \times [k]$ defined via $\iota_M(j, \ell) := (M_{j\ell}, \ell)$. (Indeed, the mapping $M \mapsto \iota_M$ yields a one-to-one correspondence between the set of k-partite hypermatchings $\mathcal{M}_n^{m,k}$ and the set of injective mappings $\iota : [m] \times [k] \to [n] \times [k]$ with the property that $\iota([m] \times \{\ell\}) \subseteq \iota([n] \times \{\ell\})$ for every $\ell \in [k]$.) If $e = (e_1, \dots, e_k) \in [n]^k$ is a row of such a matrix, we write $\mathsf{supp}(e) := \{(e_\ell, \ell) : \ell \in [k]\} \subset [n] \times [k]$ as the set of vertices participating in the edge e. We similarly write $\mathsf{supp}(M) := \bigcup_{j=1}^m \mathsf{supp}(M_j)$, where M_j is the j-th row of M. ♢

Definition 3.2 (Gadget hypergraphs)**.** For $T, k, k' \in \mathbb{N}$, we define the set of k-uniform hypergraphs with T edges on a fixed vertex-set $[k']$ as:

$$\mathcal{H}_{k'}^{T,k} := \{G \in [k']^{T \times k} : \text{all entries in each row of } G \text{ are distinct}\}.$$

A vector $e \in [k']^k$ where every entry is distinct may equivalently be viewed as an injection $\phi : [k] \to [k']$. By (slight) abuse of notation, we will also defined $\phi(j, \ell)$ as $(j, \phi(\ell))$ for $(j, \ell) \in [m] \times [k]$. Thus, we can use such an injection ϕ to embed k-partite matchings into k'-partite graphs. ♢

Definition 3.3. For an injection $\phi : [k] \to [k']$ and a k-partite hypermatching $M \in \mathcal{M}_n^{m,k}$, we define a corresponding projection operator $\Pi_M^\phi : \mathbb{Z}_q^{n \times k'} \to \mathbb{Z}_q^{m \times k}$ via $(\Pi_M^\phi X)_{j,\ell} := X_{\phi(\iota_M(j,\ell))}$. For a matrix $X \in \mathbb{Z}_q^{n \times k'}$, viewed as a labeling on a set of vertices $[n] \times [k']$, $\Pi_M^\phi X$ is precisely the *induced* labeling on $[m] \times [k]$. ♢

We now define the following communication problem, which is a single-pass variant of the problem in [FMW26a].

Definition 3.4 ("Distributional implicit hidden partition" problem)**.** Suppose we have a k-uniform hypergraph $G \in \mathcal{H}_k^{T,k'}$ and a sequence of one-wise uniform distributions $\mathcal{Y}_1, \dots, \mathcal{Y}_T \in \Delta_1(\mathbb{Z}_q^k)$. Let $m \le n \in \mathbb{N}$. Consider the following communication problem DIHP.

- *Communication structure*: There are T players named $\texttt{Player}_t$ for $t \in [T]$. This is a one-way communication game; the players speak in order from 1 to T.
- *Input type:* Each player $\texttt{Player}_t$ gets a k-partite hypermatching $M_t \in \mathcal{M}_n^{m,k}$ and a noisy signal $Z_t \in \mathbb{Z}_q^{m \times k}$. (I.e., each player can be viewed as a function $\texttt{Player}_t : \mathcal{M}_n^{m,k} \times \mathbb{Z}_q^{m \times k} \times \{0,1\}^s \to \{0,1\}^s$.)
- *Input distributions:* The goal is to distinguish between the **YES** and **NO** input distributions. In both cases, the inputs $\boldsymbol{M}_t$ are uniformly random and independent. The two cases are defined as follows:
  - In the **NO** case, $\boldsymbol{Z}_t$ is also picked uniformly and independently for every $t \in [T]$.
  - In the **YES** case, we sample an underlying hidden assignment $\boldsymbol{X}^* \in \mathbb{Z}_q^{n \times k'}$ (uniformly at random). Then, for every $t \in [T]$, we sample $\boldsymbol{Y}_t \sim \mathcal{Y}_t^{\otimes m}$ and set $\boldsymbol{Z}_t := \Pi_{\boldsymbol{M}_t}^{\phi_t} \boldsymbol{X}^* - \boldsymbol{Y}_t$. ♢

We interpret these distributions in the following way: In the **YES** case, X is $\mathbb{Z}_q$-labeling on $[n] \times [k']$. À la Definition 3.3, the k-partite hypermatching M_t and the injection $\phi_t : [k] \to [k']$ give a correspondence by which X also "induces" a $\mathbb{Z}_q$-labeling on $[m] \times [k]$, namely, the labeling $\Pi_{M_t}^{\phi_t} X$.

The players want to determine whether their inputs Z_t are consistent with a global labeling X (**YES** case) or are just independent and random (**NO** case). However, in the **YES** case, the players do not receive the induced labelings $\Pi_{M_t}^{\phi_t} X$ directly; instead, they are "masked" by randomized noise drawn from the corresponding distribution $\mathcal{Y}_t$.

The goal of the players is to determine whether their input is drawn from the **YES** or the **NO** distribution. As in previous works [CGS+22; KK19], we consider the blackboard communication model, where players communicate via a shared blackboard in a fixed sequential order beginning with $\texttt{Player}_1$. Without loss of generality, we assume that the last player $\texttt{Player}_t$ outputs a single bit indicating which distribution their input is drawn from. In the i-th round, $\texttt{Player}_i$ observes the current contents of the blackboard and writes their message $\boldsymbol{S}_i$. We further assume that each $\texttt{Player}_i$ reveals their hypermatching $\boldsymbol{M}_i$ alongside their message $\boldsymbol{S}_i$, at no cost.[6] The communication cost of a protocol is then defined as the worst-case total length (in bits) of all *messages* written on the blackboard, i.e., $\sum_{i\in[T]} |\boldsymbol{S}_i|$.

We prove the following linear lower bound on the communication complexity of DIHP in §4.

Theorem 3.5 (Linear lower bound for DIHP)**.** *For every $q, k \in \mathbb{N}$ and $\delta \in (0, 1/2)$, there exists $\alpha_0 > 0$ such that for every $\alpha \in (0, \alpha_0]$ and $T \in \mathbb{N}$, there exists $n_0 \in \mathbb{N}$ and $\tau \in (0, 1)$ such that the following holds. When $n \geq n_0$, the communication complexity of any protocol for DIHP that succeeds with advantage δ is at least τn.*

3.1 The streaming reduction

We now present a reduction from DIHP to Max-CSP due to [FMW26a], which (as we will see later in this section) can be implemented in the streaming setting.

Definition 3.6 (Max-CSP($\mathcal{F}$) instances from labeled matchings)**.** Let $k, k', q_0, q \in \mathbb{N}$, $\phi : [k] \to [k']$ be an injection, and $f : [q_0]^k \to \{0, 1\}$ be a predicate.

Let $M \in \mathcal{M}_n^{m,k}$ be a k-partite hypermatching and $Z \in \mathbb{Z}_q^{m\times k}$. We define the instance $\Phi_{\phi,f,M,Z}$ of Max-CSP($\mathcal{F}$) on the variable-set $[n] \times [k']$, whose constraints are the following: For each row $j \in [m]$ such that $Z_j = 0$, create a constraint with predicate f and variable sequence $e_j := ((M_{j,1}, \phi(1)), \ldots, (M_{j,k}, \phi(k))) = (\phi(\iota_M(j, \ell)))_{\ell\in[k]}$. ◊

Note that only variables in $\phi(\mathsf{supp}(M))$ are used in the instance $\Phi_{\phi,f,M,Z}$ and all constraints in this instance use the same predicate (f). Also, q_0, which appears in the domain of f, may differ from q, which appears in the domain of Z.

Proposition 3.7 ([FMW26a, Proof of Lemma 3.1])**.** *Let $q_0, k \in \mathbb{N}$ and $\mathcal{F} \subseteq \{0, 1\}^{[q_0]^k}$. Suppose that* Max-CSP($\mathcal{F}$) *admits a (γ, β)-integrality gap instance. Then there exists $q, k' \in \mathbb{N}$ such that for every $\varepsilon > 0$, there exists $\alpha_0 \in (0, 1)$ such that the following holds.*

For every $\alpha \in (0, \alpha_0)$, there exists $T \in \mathbb{N}$, injections $\phi_1, \ldots, \phi_T : [k] \to [k']$ corresponding to a gadget hypergraph $G \in \mathcal{H}_k^{T,k'}$, one-wise uniform distributions $\mathcal{Y}_1, \ldots, \mathcal{Y}_T \in \Delta_1(\mathbb{Z}_q^k)$, and predicates $f_1, \ldots, f_T : [q_0]^k \to \{0, 1\}$ such that the following holds: Suppose we sample inputs $(\boldsymbol{M}_1, \boldsymbol{Z}_1), \ldots, (\boldsymbol{M}_T, \boldsymbol{Z}_T)$ from DIHP with parameters G, $\mathcal{Y}_1, \ldots, \mathcal{Y}_T$, n, and $m := \alpha n$. Then, defining the random CSP instance $\boldsymbol{\Phi} := \bigsqcup_{t=1}^T \Phi_{\phi_t,f_t,\boldsymbol{M}_t,\boldsymbol{Z}_t}$, we have:

- *In the* **YES** *case,* $\Pr[\mathsf{opt}^{\mathrm{CSP}}(\boldsymbol{\Phi}) \geq \gamma - \varepsilon] \geq 1 - o_n(1)$. *Further, if $\gamma = 1$, then* $\mathsf{opt}^{\mathrm{CSP}}(\Phi) = 1$ *deterministically.*
- *In the* **NO** *case,* $\Pr[\mathsf{opt}^{\mathrm{CSP}}(\boldsymbol{\Phi}) \leq \beta + \varepsilon] \geq 1 - o_n(1)$.

[6] Note that this assumption only yields a stronger lower bound.

Proof sketch. Since the coefficients of the basic LP are rational (see Figure 1) and its feasible region is nonempty, there is a (γ, β)-integrality gap instance Φ_0 for MAX-CSP($\mathcal{F}$) in which all the local distributions $\mathcal{Y}_C^0 \in \Delta([q_0]^k)$ are rational.

Let $\mathcal{V}$ denote the variable set of this instance. For each $v \in \mathcal{V}$, let $\mathcal{X}_v \in \Delta([q_0])$ denote the marginal distribution on the variable v which is consistent with all the local distributions. Take $q \in \mathbb{N}$ so that for every $v \in \mathcal{V}$ and $a \in [q_0]$, $q \cdot \mathcal{X}_v(a) \in \mathbb{N}$ (all marginal probabilities are multiples of $1/q$).

For every variable $v \in \mathcal{V}$, consider an arbitrary partition $\mathcal{P}_v$ of $[q]$ into sets $S_1^{(v)}, \dots, S_{q_0}^{(v)}$ with $|S_a^{(v)}| = q \cdot \mathcal{X}_v(a)$ for all $a \in [q_0]$. Let $\kappa_v : [q] \to [q_0]$ be the function mapping $S_a^{(v)}$ to a for all $a \in [q_0]$, so that sampling $\boldsymbol{b} \sim [q]$ uniformly, the random variable $\kappa_v(\boldsymbol{b})$ has distribution $\mathcal{X}_v$. Observe that if we sample $\boldsymbol{a} \sim \mathcal{X}_v$ and then sample $\boldsymbol{b} \sim \kappa_v^{-1}(\boldsymbol{a})$, the marginal distribution of $\boldsymbol{b}$ is uniform on $[q]$.

Let $C = (f, e)$ be a constraint in the basic LP integrality gap instance Φ_0 (with $e = (i_1, \dots, i_k)$). Let $\mathcal{Y}_C^0 \in \Delta([q_0]^k)$ denote the local distribution corresponding to C. We define a corresponding distribution $\mathcal{Y}_C \in \Delta([q]^k)$ by sampling $\boldsymbol{a} = (\boldsymbol{a}_1, \dots, \boldsymbol{a}_k) \sim \mathcal{Y}_C^0$ and then outputting a random sample from $\kappa_{i_1}^{-1}(\boldsymbol{a}_1) \times \cdots \times \kappa_{i_k}^{-1}(\boldsymbol{a}_k)$. $\mathcal{Y}_C$ is one-wise uniform by the observation in the preceding paragraph (and using the fact that for each $\ell \in [k]$, $\boldsymbol{a}_\ell$ is distributed as $\mathcal{X}_{i_\ell}$). Conversely, sampling $\boldsymbol{b} = (\boldsymbol{b}_1, \dots, \boldsymbol{b}_k) \sim \mathcal{Y}_C$, the random variable $(\kappa_{i_1}(\boldsymbol{b}_1), \dots, \kappa_{i_k}(\boldsymbol{b}_k))$ has distribution $\mathcal{Y}_C^0$.

We will set $k' := |\mathcal{V}|$ and eventually $T := T_0 \cdot K$, where T_0 is the number of constraints in the original instance Φ_0 and K is a large enough constant (enough to ensure concentration below). The gadget hypergraph G on k' vertices is constructed by adding K copies of e for each constraint $C = (f, e)$ in Φ_0, with all copies labeled with the same (one-wise uniform) distribution $\mathcal{Y}_C$ and corresponding to the same predicate f.

The YES case analysis. Given a signal $X \in \mathbb{Z}_q^{n \times k'}$, define $\kappa(X) \in [q_0]^{n \times k'}$ via applying the functions κ_v appropriately, i.e., $(\kappa(X))_{i,v} := \kappa_v(X_{i,v})$, which we view as an assignment to the CSP instance Φ. We argue that in the **YES** case, even fixing the matchings $M_1, \dots, M_T$, we have

$$\frac{\mathbb{E}[\#\text{ constraints in } \boldsymbol{\Phi} \text{ satisfied by } \kappa(\boldsymbol{X})]}{\mathbb{E}[\#\text{ constraints in } \boldsymbol{\Phi}]} = \gamma,$$

where the above expectation is over the choice of a random X. (To carry out the full argument of [FMW26a], we would then use concentration bounds over the choice of $\boldsymbol{Y}_1, \dots, \boldsymbol{Y}_T$ and $\boldsymbol{X}$ to get that the actual ratio is roughly γ with high probability.)

Recall that, for $t \in [T]$ and $j \in [m]$, an injection $\phi_t : [k] \to [k']$, a matching $M_t \in \mathcal{M}_n^{m,k}$, and labeling $Z_t \in \mathbb{Z}_q^{m \times k}$, we inject the constraint $(f_t, (M_t)_j)$ into Φ iff $(Z_t)_j = 0$. In the **YES** case, $\boldsymbol{Z}_t = \Pi_{M_t}^{\phi_t} \boldsymbol{X} - \boldsymbol{Y}_t$, where $\boldsymbol{X} \sim \mathbb{Z}_q^{n \times k}$ and $\boldsymbol{Y}_t \sim \mathcal{Y}_t^{\otimes m}$ independently. Hence, in particular, $(\boldsymbol{Z}_t)_j = \Pi_{M_j}^{\phi_t} \boldsymbol{X} + (\boldsymbol{Y}_t)_j$, where $\boldsymbol{X} \sim \mathbb{Z}_q^{n \times k}$ and $(\boldsymbol{Y}_t)_j \sim \mathcal{Y}_t$ independently. Note that $\Pr[(\boldsymbol{Z}_t)_j = 0] = q^{-k}$; this only uses the uniformity of $\boldsymbol{X}$, and so is true even conditioned on any fixing of $(Y_t)_j$. Hence, by linearity, the expected number of constraints in $\boldsymbol{\Phi}$ is $q^{-k} \cdot mT$. On the other hand,

$$\begin{aligned} &\Pr[f_t(\kappa(\boldsymbol{X})_{(M_t)_{j,1}, \phi_t(1)}, \dots, \kappa(\boldsymbol{X})_{(M_t)_{j,k}, \phi_t(k)}) = 1 \wedge (\boldsymbol{Z}_t)_j = 0] \\ = \; &\Pr[f_t(\kappa_1(\boldsymbol{X}_{(M_t)_{j,1}, \phi_t(1)}), \dots, \kappa_k(\boldsymbol{X}_{(M_t)_{j,k}, \phi_t(k)})) = 1 \mid (\boldsymbol{Z}_t)_j = 0] \cdot \Pr[(\boldsymbol{Z}_t)_j = 0]. \end{aligned}$$

The second factor is again q^{-k}. For the first factor, we use that $(X_{(M_t)_{j,1}, \phi_t(1)}, \dots, X_{(M_t)_{j,k}, \phi_t(k)}) = \Pi_{(M_t)_j}^{\phi_t} X$ and conditioned on $(\boldsymbol{Z}_t)_j = 0$, the marginal distribution of $\Pi_{(M_t)_j}^{\phi_t} X$ is $\mathcal{Y}_C$. Hence, the marginal distribution of $\kappa_1(\boldsymbol{X}_{(M_t)_{j,1}, 1}), \dots, \kappa_k(\boldsymbol{X}_{(M_t)_{j,k}, k})$ is uniform on $\mathcal{Y}_C^0$, so that the first factor

is just $\mathbb{E}_{\boldsymbol{a}\sim\mathcal{Y}_C^0}[f_t(\boldsymbol{a})]$. Now summing over (all K copies of) all T_0 constraints C, we recover the expression $q^{-k}\cdot\sum_{C=(f,e)\in\Phi}\mathbb{E}_{\boldsymbol{a}\sim\mathcal{Y}_C^0}[f(\boldsymbol{a})]$, which is exactly the (unnormalized) LP objective.

The NO case analysis. For the **NO** case, we can use that $\boldsymbol{\Phi}$ is essentially a random blowup of the base instance Φ_0. We can therefore show that the expected value of every fixed assignment to $\boldsymbol{\Phi}$ is at most β (since any such assignment projects down to an assignment to Φ_0), then use union+concentration bounds (assuming α_0 is sufficiently small and K sufficiently large). □

We are now ready to prove our main theorem, as restated below.

Theorem 1.2 (Linear-space lower bounds against single-pass streaming algorithms for CSPs)**.** *Let $\mathcal{F}\subseteq\{0,1\}^{[q]^k}$. For every $1\geq\gamma>\beta\geq 0$, if there exists a (γ,β)-LP integrality gap instance of* Max-CSP$(\mathcal{F})$*, then for every $\varepsilon>0$, any single-pass streaming algorithm for* Max-CSP$_{\gamma-\varepsilon,\beta+\varepsilon}(\mathcal{F})$ *on n variables requires $\Omega_\varepsilon(n)$ space. Moreover,* Max-CSP$_{1,\beta+\varepsilon}(\mathcal{F})$ *also requires space $\Omega_\varepsilon(n)$ for the case $\gamma=1$.*

Proof. Consider any (γ,β)-LP integrality gap instance of Max-CSP$(\mathcal{F})$. Suppose for contradiction that there exists an $\varepsilon_0>0$ and a streaming algorithm $\mathcal{A}$ for Max-CSP$_{\gamma-\varepsilon_0,\beta+\varepsilon_0}(\mathcal{F})$ that requires only $o(n)$ space. Using Proposition 3.7, we will derive a $o(n)$-bit communication protocol for DIHP, contradicting Theorem 3.5. To apply Proposition 3.7, we set $(q_0,k,\varepsilon)_{\text{Proposition 3.7}}=(q,k,\varepsilon_0)$. For these parameters, consider the DIHP instance produced in Proposition 3.7 and the corresponding random CSP instance $\boldsymbol{\Phi}:=\bigsqcup_{t=1}^T\Phi_{\phi_t,f_t,\boldsymbol{M}_t,\boldsymbol{Z}_t}$. Since the reduction described in Definition 3.6 is local, each player $\texttt{Player}_t$ can independently sample $\Phi_{\phi_t,f_t,\boldsymbol{M}_t,\boldsymbol{Z}_t}$. Starting from $\texttt{Player}_1$, the players can sequentially run $\mathcal{A}$ on their local instance and pass the memory state to the next player; specifically, starting from the memory state passed by $\texttt{Player}_{t-1}$, player $\texttt{Player}_t$ runs $\mathcal{A}$ on $\Phi_{\phi_t,f_t,\boldsymbol{M}_t,\boldsymbol{Z}_t}$. Finally, $\texttt{Player}_T$ outputs **YES** if $\mathcal{A}$ categorizes the instance as having value at least $\gamma+\varepsilon_0$, and **NO** otherwise. Correctness follows immediately from Proposition 3.7. Since $\mathcal{A}$ uses only $o(n)$ memory, this yields a $o(n)$-bit communication protocol for DIHP. The case $\gamma=1$ is analogous. □

4 Lower bound on the communication complexity

In this section, we prove the main communication lower bound for DIHP, as restated below.

Theorem 3.5 (Linear lower bound for DIHP)**.** *For every $q,k\in\mathbb{N}$ and $\delta\in(0,1/2)$, there exists $\alpha_0>0$ such that for every $\alpha\in(0,\alpha_0]$ and $T\in\mathbb{N}$, there exists $n_0\in\mathbb{N}$ and $\tau\in(0,1)$ such that the following holds. When $n\geq n_0$, the communication complexity of any protocol for DIHP that succeeds with advantage δ is at least τn.*

By Yao's minimax principle [Yao77], it suffices to prove such a lower bound against all *deterministic* protocols for DIHP. Namely, a protocol for DIHP can be specified by deterministic message functions $\texttt{Player}_1,\ldots,\texttt{Player}_T$ so that $S_t=\texttt{Player}_t(M_{1:t},S_{1:t-1},Z_t)$ denotes the message sent by the t-th player on input $(M_{1:t},S_{1:t-1},Z_t)$. Without loss of generality, the final function $\texttt{Player}_T$ outputs a simple bit "**YES**/**NO**" indicating the output of the protocol.

4.1 Informal overview of proof

To prove Theorem 3.5, we consider the following "hybrid experiment": Independently sample a common sequence of random k-partite hypermatchings $\boldsymbol{M}_{1:T}$, a hidden signal $\boldsymbol{X}\in\mathbb{Z}_q^{n\times k'}$, and then sample **YES** inputs $\boldsymbol{Z}_{1:T}^{Y}$ (i.e., $\boldsymbol{Z}_t^Y=\Pi_{\boldsymbol{M}_t}^{\phi_t}\boldsymbol{X}^*-\boldsymbol{Y}_t$ for hidden signal for $\boldsymbol{Y}_t\sim\mathcal{Y}_t^{\otimes k}$) and **NO** inputs

$\boldsymbol{Z}^N_{1:T}$ (uniformly random). Note that $(\boldsymbol{M}_{1:T}, \boldsymbol{Z}^Y_{1:T})$ have the marginal distribution of a **YES** input, and $(\boldsymbol{M}_{1:T}, \boldsymbol{Z}^N_{1:T})$ have the marginal distribution of a **NO** input. Thus, this experiment is a *coupling* between the **YES** and **NO** input distributions.

In the above experiment, let $\boldsymbol{S}^Y_{1:T}$ denote the messages sent in the **YES** case (so that $\boldsymbol{S}^Y_t = r_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}^Y_{1:t-1}, \boldsymbol{Z}^Y_t)$) and $\boldsymbol{S}^N_{1:T}$ denote the messages in the **NO** case (so that $\boldsymbol{S}^N_t = r_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}^N_{1:t-1}, \boldsymbol{Z}^N_t)$). To prove Theorem 3.5 we need to show that the distributions of the random variables $\boldsymbol{S}^Y_T$ and $\boldsymbol{S}^N_T$ are close in total variation distance. For the induction, we prove the much stronger statement that the distributions of $(\boldsymbol{M}_{1:T}, \boldsymbol{S}^Y_{1:T})$ and $(\boldsymbol{M}_{1:T}, \boldsymbol{S}^N_{1:T})$ are close in total variation distance, i.e.,

$$\|(\boldsymbol{M}_{1:T}, \boldsymbol{S}^Y_{1:T}) - (\boldsymbol{M}_{1:T}, \boldsymbol{S}^N_{1:T})\|_{tvd} \le \delta\,.$$

We prove the above bound via a hybrid lemma analogous to those in [CGS+22; KK19]. Roughly speaking, this lemma states that if the first $t-1$ players' inputs are drawn from the **YES** distribution, then the t-th player's output on a **YES** input is distributed very similarly to their output on a **NO** input, even conditioned on all previous hypermatchings and messages. Formally, the lemma identifies a *filtration* (nested sequence) of events $\mathcal{E}_1 \supset \mathcal{E}_2 \supset \cdots \supset \mathcal{E}_T$ in the **YES** experiment such that: (i) $\mathcal{E}_t$ enforces a typicality condition on the k-partite hypermatchings received by the first t players and the messages of the first $t-1$ players; and (ii) whenever these random variables take typical values, player $\texttt{Player}_t$ cannot distinguish whether their input is sampled from the **YES** distribution or the **NO** distribution, assuming all previous players' inputs were drawn from the **YES** distribution.

4.2 Formal statement of hybrid lemma

The probability space underlying the hybrid lemma (Lemma 4.1) will be the following distribution. Let $\Omega^Y := \Omega_{k,q,\alpha,n,T,G,\mathcal{Y}_{1:T}}$ denote the distribution over tuples $(\boldsymbol{X}^*, \boldsymbol{M}_{1:T}, \boldsymbol{Z}_{1:T})$ where $\boldsymbol{X}^* \sim \mathsf{Unif}(\mathbb{Z}_q^n)$, $\boldsymbol{M}_t \in \mathcal{M}_n^{m,k}$ is a k-partite hypermatching with m edges, and $\boldsymbol{Z}_t \in \mathbb{Z}_q^{m\times k}$ is a noisy signal, as defined in the **YES** case of Definition 3.4 for every $t \in [T]$, i.e., for each $t \in [T]$, we sample $\boldsymbol{Y}_t \sim \mathcal{Y}_t^{\otimes m}$ and set $\boldsymbol{Z}_t := \Pi^{\phi_t}_{\boldsymbol{M}_t}\boldsymbol{X}^* - \boldsymbol{Y}_t$.

These variables along with a deterministic protocol given by $\texttt{Player}_1, \ldots, \texttt{Player}_T$ specify additional random variables that are determined by $(\boldsymbol{X}^*, \boldsymbol{M}_{1:T}, \boldsymbol{Z}_{1:T})$ including $\boldsymbol{\mathcal{D}}_{1:t}$, $\boldsymbol{S}^Y_t$.

Lemma 4.1 (Hybrid lemma)**.** *For every $q, k \in \mathbb{N}$, there exists $\alpha_0 > 0$ such that for every $T \in \mathbb{N}$, and $\delta \in (0,1)$, there exists $\theta \in (0,1)$ and $n_0 < \infty$ such that the following holds for every $n \ge n_0$:*

Let $\Pi = (\texttt{Player}_t)_{t\in[T]}$ be a deterministic protocol for DIHP where $m \le \alpha_0 n$ and each message function r_t outputs a message of at most θn bits. Let $(\boldsymbol{X}^*, \boldsymbol{M}_{1:T}, \boldsymbol{Z}_{1:T}) \sim \Omega$. Then there exists a sequence of events $\{\mathcal{E}_t\}_{t\in[T]}$ such that:

1. *For $t \in [T]$, $\mathcal{E}_t$ only depends on $\boldsymbol{M}_{1:t}$ and $\boldsymbol{S}^Y_{1:t-1}$ (with $\boldsymbol{S}^Y_{1:0}$ denoting an empty set of variables).*
2. *For every $t \in [T]$, $\mathcal{E}_t \Rightarrow \mathcal{E}_{t-1}$ and $\Pr[\overline{\mathcal{E}_t} \mid \mathcal{E}_{t-1}] \le (\delta/(2T))$ (where $\mathcal{E}_0$ is the trivial event occurring with probability 1).*
3. *For every fixed $M_{1:t}$ and $S^Y_{1:t-1}$ satisfying $\mathcal{E}_t$, one has*

$$\|\boldsymbol{S}^Y_t - \texttt{Player}_t(M_{1:t}, S^Y_{1:t-1}, \boldsymbol{U})\|_{\mathsf{tvd}} \le \delta/(2T), \tag{4.2}$$

where $\boldsymbol{U} \sim \mathsf{Unif}(\mathbb{Z}_q^{m\times k})$.

Using Lemma 4.1, we can prove Theorem 3.5 using a simple inductive argument. The proof is analogous to the proof of Lemma 6.3 in [KK19] and Theorem 3.5 in [CGS+22]. So we skip this proof here and include it in §A for completeness sake. Lemma 4.1 is proved in §4.5 below.

4.3 Informal proof of Lemma 4.1

We prove Lemma 4.1 formally in §4.5. Here, we give a brief overview of the proof along with the definitions and lemmas that we will need. At a high level, the main idea is to show that for every $t \in [T]$, after the first t rounds of communication, sufficient randomness remains in $\boldsymbol{X}^*$ so that $\texttt{Player}_{t+1}$ cannot determine whether their input is drawn from the **YES** or the **NO** distribution. As in [CGS+22; KK19], we formally capture the randomness in $\boldsymbol{X}^*$ through the so-called *boundedness* of its posterior distribution, conditioned on the hypermatchings and the messages of the first t players. The boundedness condition (stated formally in Definition 4.5) bounds the ℓ_1-norm of the Fourier coefficients at each Hamming weight. The subsequent *boundedness implies uniformity* lemma (Lemma 4.7, proved in §5) then implies that if the posterior distribution of $\boldsymbol{X}^*$ after t rounds of communication is bounded, the total variation distance between $\texttt{Player}_{t+1}$'s inputs under the **YES** and **NO** distributions is very small.

The next step is to inductively show that if the posterior distribution of $\boldsymbol{X}^*$ is bounded after the first t rounds of communication (roughly captured by event $\mathcal{E}_t$), then it remains bounded after the $(t+1)$-th round. This is the step where we depart significantly from prior work [CGS+22; KK19]. In previous works, the players' inputs in the **YES** case were noiseless, i.e., they received the induced labelings $\Pi^{\phi_t}_{M_t} X$ directly.[7] Consequently, the posterior distributions in those works were always uniform over a set, and showing that this set is large and well-structured sufficed to establish boundedness of the posterior distribution. The analysis in our setting is more intricate because the posterior distribution of $\boldsymbol{X}^*$ is no longer uniform over a set. While the strong structural guarantees enjoyed by the posterior sets in previous works do not hold here, we show that weaker properties suffice to carry out the induction. The inductive step lemma is stated formally as Lemma 4.9 and proved in §7. Assuming Lemmas 4.7 and 4.9, we formally prove Lemma 4.1 in §4.5.

4.4 Definitions and statements of key lemmas

Definition 4.3 ((One-step) posterior distribution)**.** Let $n, k, k', m \in \mathbb{N}$, $\phi : [k] \to [k']$ be an injection, and $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$ a one-wise uniform distribution. For a "prior" distribution $\mathcal{D} \in \Delta(\mathbb{Z}_q^{n \times k'})$, a k-partite hypermatching $M \in \mathcal{M}_n^{m,k}$, and a set $\mathcal{B} \subseteq \mathbb{Z}_q^{m \times k}$, we define the *posterior distribution* $\mathrm{Post}_{\phi,\mathcal{Y}}(\mathcal{D}; M, \mathcal{B})$ as the conditional distribution on $\boldsymbol{X} \in \mathbb{Z}_q^{n \times k'}$ when we sample $\boldsymbol{X} \sim \mathcal{D}$ and $\boldsymbol{Y} \sim \mathcal{Y}^{\otimes m}$ independently and condition on the event that $\Pi^{\phi}_{M} \boldsymbol{X} - \boldsymbol{Y} \in \mathcal{B}$. The density function of this distribution is

$$\mu_{\mathrm{Post}_{\phi,\mathcal{Y}}(\mathcal{D};M,\mathcal{B})}(X_0) = q^{n \cdot k'} \cdot \Pr_{\substack{\boldsymbol{X} \sim \mathcal{D}, \\ \boldsymbol{Y} \sim \mathcal{Y}^{\otimes m}}} \left[\boldsymbol{X} = X_0 \mid \Pi^{\phi}_{M} \boldsymbol{X} - \boldsymbol{Y} \in \mathcal{B} \right].$$

We omit ϕ and $\mathcal{Y}$ and write $\mathrm{Post}(\mathcal{D}; M, \mathcal{B})$ when clear from context. ◊

Definition 4.4 (Input distribution)**.** Let $n, k, k', m \in \mathbb{N}$, $\phi : [k] \to [k']$ an injection, and $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$ a one-wise uniform distribution. For a distribution $\mathcal{D} \in \Delta(\mathbb{Z}_q^{n \times k'})$ and a k-partite hypermatching $M \in \mathcal{M}_n^{m,k}$, we define the *input distribution* $\mathrm{Input}_{\phi,\mathcal{Y}}(\mathcal{D}, M)$ as the distribution of $\Pi^{\phi}_{M} \boldsymbol{X} - \boldsymbol{Y}$ when $\boldsymbol{X} \sim \mathcal{D}$ and $\boldsymbol{Y} \sim \mathcal{Y}^{\otimes m}$ independently. We omit ϕ and $\mathcal{Y}$ and write $\mathrm{Input}(\mathcal{D}, M)$ when clear from context. ◊

[7] In [CGS+22], the authors specifically consider a "shifted noise" distribution that applies a uniform shift to all coordinates; via a recentering argument, they reduce this to the noiseless case.

Definition 4.5 (Fourier bound function). A distribution $\mathcal{D} \in \Delta(\mathbb{Z}_q^N)$ is (C,s)-*bounded* if, for every $h \in [N]$,

$$\sum_{\substack{u \in \mathbb{Z}_q^N, \\ \mathrm{wt}(u)=h}} |\widehat{\mu_{\mathcal{D}}}(u)| \le \mathrm{U}_{C,s,N}(h),$$

where for $C < \infty$ and $s, h, N \in \mathbb{N}$, we define:

$$\mathrm{U}_{C,s,N}(h) := \begin{cases} 1, & \text{if } h = 0 \\ \left(\frac{C\sqrt{sN}}{h}\right)^{h/2}, & \text{if } 1 \le h \le s \\ \min\left\{\left(\frac{C\sqrt{N}}{\sqrt{h}}\right)^{h/2}, \left(\frac{eq^2N}{h}\right)^{h/2}\right\}, & \text{if } h > s. \end{cases} \qquad \lozenge$$

Observation 4.6 (Monotonicity of boundedness with respect to C)**.** *The function $\mathrm{U}_{C,s,N}(h)$ increases monotonically with respect to the parameter s. Hence, if a distribution $\mathcal{D} \in \Delta(\mathbb{Z}_q^N)$ is (C,s)-bounded, then for every $C' \ge C$, $\mathcal{D}$ is also (C',s)-bounded.*

Lemma 4.7 (Boundedness implies uniformity)**.** *For every $k, q, k' \ge 2 \in \mathbb{N}$ ($k' \ge k$), there exists $\alpha_0 > 0$ such that for every $\delta \in (0, 1/2)$ and $C < \infty$, there exists $\tau > 0$ such that the following holds.*

For every injection $\phi : [k] \to [k']$, one-wise uniform distribution $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$, $n \in \mathbb{N}$ sufficiently large, $\mathcal{D} \in \Delta(\mathbb{Z}_q^{n \times k'})$ a (C,s)-bounded distribution satisfying $\|\mu_{\mathcal{D}}\|_\infty \le q^b$ and $4\log_q(3/\delta) \le b \le s \le \tau n$, $m \le \alpha_0 n$,

$$\Pr_{\boldsymbol{M} \in \mathcal{M}_n^{m,k}}\left[\|\mu_{\mathrm{Input}(\mathcal{D},\boldsymbol{M})} - 1\|_\infty \le \delta\right] \ \ge \ 1 - \delta,$$

where $\boldsymbol{M} \in \mathcal{M}_n^{m,k}$ is a randomly sampled matching.

Observe that the parameter b in Lemma 4.7 can be eliminated, if the inequality is replaced with $4\log_q \max\{3/\delta, \|\mu_{\mathcal{D}}\|_\infty\} \le s \le \tau n$.

Observation 4.8. *For $\mathcal{Z} \in \Delta(\mathbb{Z}_q^{m \times k})$, the condition $\|\mu_{\mathcal{Z}} - 1\|_\infty \le \delta$ is equivalent to that for every $Z_0 \in \mathbb{Z}_q^{m \times k}$, it holds that $1 - \delta \le \mu_{\mathcal{Z}}(Z_0) \le 1 + \delta$. In this case, we also have $\|\mathcal{Z} - \mathsf{Unif}(\mathbb{Z}_q^{m \times k})\|_{\mathrm{tvd}} \le \delta$, and for every function $f : \mathbb{Z}_q^{m \times k} \to \mathbb{R}$, $\left|\mathbb{E}_{\boldsymbol{Z} \sim \mathcal{Z}}[f(\boldsymbol{Z})] - \mathbb{E}_{\boldsymbol{Z} \in \mathbb{Z}_q^{m \times k}}[f(\boldsymbol{Z})]\right| \le \delta \|f\|_1$.*

Lemma 4.9 (Inductive step)**.** *For every $k, q, k' \in \mathbb{N}$ there exist $\alpha_0 > 0$ and $C_0 < \infty$ such that for every $C \ge C_0$ and $\delta \in (0, 1/2)$ there exist $\sigma \in (0,1)$ and $C'' > 0$ such that the following holds.*

For every injection $\phi : [k] \to [k']$ and one-wise uniform distribution $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$, for every $n, b, b', s, s', m \in \mathbb{N}$ satisfying $m \le \alpha_0 n$, $0 < b, b', s < \sigma n$ and $b + b' + \log_q(1 - 1/\delta) \le s$ and every (C,s)-bounded distribution $\mathcal{D} \in \Delta(\mathbb{Z}_q^{n \times k'})$ satisfying $\|\mu_{\mathcal{D}}\|_\infty \le q^b$, we have that

$$\Pr_{\boldsymbol{M} \in \mathcal{M}_n^{m,k}}\left[\begin{array}{l} \forall \mathcal{B} \subset \mathbb{Z}_q^{m \times k} \text{ such that } |\mathcal{B}| \ge q^{m \cdot k - b'}, \\ \qquad \mathrm{Post}(\mathcal{D}; \boldsymbol{M}, \mathcal{B}) \text{ is } (C'', s)\text{-bounded and } \|\mu_{\mathrm{Post}(\mathcal{D};\boldsymbol{M},\mathcal{B})}\|_\infty \le \dfrac{1}{1-\delta} \cdot q^{b+b'} \end{array}\right] \ge 1 - 5\delta$$

where $\boldsymbol{M} \in \mathcal{M}_n^{m,k}$ is a randomly sampled k-partite hypermatching.

4.5 Hybrid lemma: Proof of Lemma 4.1

In this section, we formally prove Lemma 4.1.

Proof of Lemma 4.1. Let $\texttt{Player}_1, \ldots, \texttt{Player}_T$ be the message functions corresponding to some deterministic protocol for DIHP. Recall that the input of $\texttt{Player}_t$ (in addition to the prior message) is a k-partite hypermatching $\boldsymbol{M}_t$ and a noisy signal $\boldsymbol{Z}_t$. (We also have fixed $\phi_t, \mathcal{Y}_t$, which are, respectively the injective map and the one-wise uniform distribution corresponding to $\texttt{Player}_t$.)

Posterior set and distribution. We define some auxiliary random variables determined by the randomness of $(\boldsymbol{M}_{1:t}, \boldsymbol{S}^{\mathcal{Y}}_{1:t})$. For $t \in [T]$, we define:

$$\boldsymbol{\mathcal{B}}_t \coloneqq \{Z \in \mathbb{Z}_q^{m \times k} : \boldsymbol{S}^{\mathcal{Y}}_t = \texttt{Player}_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}^{\mathcal{Y}}_{1:t-1}, Z)\}$$

and, recursively:

$$\boldsymbol{\mathcal{D}}_t \coloneqq \mathrm{Post}_{\phi_t, \mathcal{Y}_t}(\boldsymbol{\mathcal{D}}_{t-1}; \boldsymbol{M}_t, \boldsymbol{\mathcal{B}}_t).$$

where $\boldsymbol{\mathcal{D}}_0 = \mathcal{D}_0 = \mathsf{Unif}\{\mathbb{Z}_q^{n \times k'}\}$. $\boldsymbol{\mathcal{B}}_t$ denotes the set of all inputs Z_t that would result in the message $\boldsymbol{S}^{\mathcal{Y}}_t$ (conditioned on the previous messages $\boldsymbol{S}^{\mathcal{Y}}_1, \ldots, \boldsymbol{S}^{\mathcal{Y}}_{t-1}$ and the hypermatchings $\boldsymbol{M}_1, \ldots, \boldsymbol{M}_t$), while $\boldsymbol{\mathcal{D}}_t$ denotes the posterior distribution on $\boldsymbol{X}^*$ conditioned on $\boldsymbol{M}_{1:t}$, $\boldsymbol{S}^{\mathcal{Y}}_{1:t}$ in the following formal sense:

Claim 4.10. *For every $0 \le t \le T$, $X_0 \in \mathbb{Z}_q^{n \times k'}$, and fixed $M_{1:t}$, $S^{\mathcal{Y}}_{1:t}$, we have*

$$\Pr_{\Omega}[\boldsymbol{X}^* = X_0 \mid M_{1:t}, S^{\mathcal{Y}}_{1:t}] = \mathcal{D}_t(X_0)$$

(here $M_{1:t}$ and $S^{\mathcal{Y}}_{1:t}$ are both tuples of length t, so that $M_{1:0}, S^{\mathcal{Y}}_{1:0}$ denote empty tuples). In particular, for every fixed $M_{1:t}$ and $S^{\mathcal{Y}}_{1:t-1}$, the marginal distribution of $\boldsymbol{S}^{\mathcal{Y}}_t$ is the same as the marginal distribution of $\texttt{Player}_t(M_{1:t}, S^{\mathcal{Y}}_{1:t-1}, \boldsymbol{Z})$, where $\boldsymbol{Z} = \Pi^{\phi_t}_{M_t} \boldsymbol{X} - \boldsymbol{Y}_t$ with $\boldsymbol{X} \sim \mathcal{D}_{t-1}$ and $\boldsymbol{Y}_t \sim \mathcal{Y}_t^{\otimes m}$.

Proof. We will prove by induction. The base case when $t = 0$ is trivially satisfied since $\mathcal{D}_0$ is the uniform distribution. Inductively, we have:

$$\begin{aligned}
&\Pr_{\Omega}[\boldsymbol{X}^* = X_0 \mid M_{1:t}, S^{\mathcal{Y}}_{1:t}] \\
&= \Pr_{\substack{\boldsymbol{X} \in \mathbb{Z}_q^{n \times k'}, \\ \forall t' \in [t],\, \boldsymbol{Y}_{t'} \sim \mathcal{Y}_{t'}^{\otimes m}}} [\boldsymbol{X} = X_0 \mid \forall t' \in [t],\ S^{\mathcal{Y}}_{t'} = \texttt{Player}_{t'}(M_{1:t'}, S^{\mathcal{Y}}_{1:t'-1}, \Pi^{\phi_{t'}}_{M_{t'}} \boldsymbol{X} - \boldsymbol{Y}_{t'})] \\
&= \Pr_{\substack{\boldsymbol{X} \sim \mathcal{D}_{t-1}, \\ \boldsymbol{Y}_t \sim \mathcal{Y}_t^{\otimes m}}} [\boldsymbol{X} = X_0 \mid S^{\mathcal{Y}}_t = \texttt{Player}_t(M_{1:t}, S^{\mathcal{Y}}_{1:t-1}, \Pi^{\phi_t}_{M_t} \boldsymbol{X} - \boldsymbol{Y}_t)] \qquad \text{(induction hypothesis)} \\
&= \Pr_{\substack{\boldsymbol{X} \sim \mathcal{D}_{t-1}, \\ \boldsymbol{Y}_t \sim \mathcal{Y}_t^{\otimes m}}} [\boldsymbol{X} = X_0 \mid \Pi^{\phi_t}_{M_t} \boldsymbol{X} - \boldsymbol{Y}_t \in \mathcal{B}_t] = \mathcal{D}_t(X_0)\,. \qquad \square
\end{aligned}$$

Parameters. The required parameters that we need to specify for Lemma 4.1 are α_0, θ, n_0, and δ'. In addition to setting these parameters below, we also set the following parameters required in the proof: (i) $\{C_t\}_{0 \le t \le T}$ and s to quantify the boundedness of the posterior distributions, (ii) s^* to upper bound the length of each message, and (iii) b to upper bound the ℓ_∞ norms of the density functions of the posterior distributions, which is required to apply Lemmas 4.7 and 4.9.

Given k, q, and k', let $\alpha_{0,1}$ be the $\alpha_0(k,q,k')$ parameter from Lemma 4.7 and $\alpha_{0,2}$ be the $\alpha_0(k,q,k')$ parameter from Lemma 4.9. We set α_0 to be $\min\{\alpha_{0,1}, \alpha_{0,2}\}$. Let C_0 be the $C_0(k,q,k')$ parameter from Lemma 4.9. Now, given T and δ, we need to specify $\tau > 0$, $n_0 < \infty$, and δ'. Along the way, we also specify the intermediary parameters that we mentioned above. Observe that Lemma 4.9 takes as input parameters $k, q, k'\delta$, and $C \geq C_0$, and specifies parameters $C'' = C''(k,q,k',\delta,C)$ and $\sigma = \sigma(k,q,k',\delta,C)$ for which the lemma holds. Observe also that Lemma 4.7 takes as input parameters $k, q, k'\delta$, and $C \geq C_0$, and specifies $\tau = \tau(k,q,k',\delta,C)$ for which the lemma holds. We set $\delta' := \delta/(30T)$. We define $\{C_t\}_{0\leq t\leq T}$ recursively as $C_t := \max\{C_0, C''(k,q,k',\delta',C_{t-1})\}$. For $t \in [T]$, let $\sigma_t := \sigma(k,q,k',\delta',C_t)$, $\tau_t := \tau(k,q,k',\delta',C_t)$, and $\gamma := \min\{\{\sigma_t\}_{t\in[T]}, \{\tau_t\}_{t\in[T]}\}$. We set $s := \gamma n$, $b := s/2T$, $b' := b$, $\theta := \gamma/8T$, and $n_0 := (1/\gamma)\cdot\max\{4\log_q(3/\delta'), 8T\log_q(1/(1-\delta'))\}$. Let $m, n \in \mathbb{N}$ be such that $n \geq n_0, m \leq \alpha_0 n$. We set $s^* := \theta n$.

Events. We define the events $\mathcal{E}_1, \ldots, \mathcal{E}_T$ (in the probability space Ω^Y) as follows. For $t \in [T]$, we define $\mathcal{E}_t^{\text{post}}$ to be the event that the distribution $\boldsymbol{\mathcal{D}}_{t-1}$ is (C_t, s)-bounded and $\|\mu_{\boldsymbol{\mathcal{D}}_{t-1}}\|_\infty \leq q^{2(t-1)b}$; $\mathcal{E}_t^{\text{in}}$ to be the event that $\|\text{Input}_{\phi_t,\mathcal{Y}_t}(\boldsymbol{\mathcal{D}}_{t-1}, \boldsymbol{M}_t) - \text{Unif}(\mathbb{Z}_q^{m\times k})\|_{\text{tvd}} \leq \delta'$; and $\mathcal{E}_t := \mathcal{E}_{t-1} \cap \mathcal{E}_t^{\text{post}} \cap \mathcal{E}_t^{\text{in}}$, where $\mathcal{E}_0$ is the trivial event occuring with probability 1.

Condition 1. It immediately follows from the above definitions that for all $t \in [T]$, $\mathcal{E}_t$ is fully determined by $\boldsymbol{M}_{1:t}$ and $\boldsymbol{S}^Y_{1:t-1}$. (Further, $\mathcal{E}_t^{\text{post}}$ itself is determined only by $\boldsymbol{M}_{1:t-1}$ and $\boldsymbol{S}^Y_{1:t-1}$.)

Condition 2. It follows also from the definitions that $\mathcal{E}_t \implies \mathcal{E}_{t-1}$, for all $t \in [T]$. We upper bound the failure probabilities of the events using induction. For the base case, we first prove that $\Pr[\mathcal{E}_1] = 1$. Since $\mathcal{D}_0$ is the uniform distribution, $\widehat{\mu_{\mathcal{D}_0}(u)} = 1$ if $u = 0^{n\times k'}$, and $\widehat{\mu_{\mathcal{D}_0}(u)} = 0$ otherwise. Therefore, $\mathcal{D}_0$ is trivially (C_0, s)-bounded. We also have $\|\mu_{\mathcal{D}_0}\|_\infty \leq 1$. For every $M \in \mathcal{M}_n^{m,k}$ and injection $\phi : [k] \to [k']$, $\Pi_M^\phi \boldsymbol{X}$ is uniformly distributed in $\mathbb{Z}_q^{m\times k}$, since $\boldsymbol{X} \sim \mathcal{D}_0$. Therefore, $\Pi_M^\phi \boldsymbol{X} - \boldsymbol{Y}$ is also uniformly distributed for $\boldsymbol{Y} \sim \mathcal{D}$ for any distribution $\mathcal{D} \in \Delta(\mathbb{Z}_q^{m\times k})$. Thus, $\mathcal{E}_1$ holds with probability 1.

We now prove that for any $t \in [T-1]$, $\Pr[\overline{\mathcal{E}_{t+1}} \mid \mathcal{E}_t] \leq \delta/(2T)$. Specifically, we will show

$$\Pr\left[\overline{\mathcal{E}_{t+1}^{\text{in}}} \mid \mathcal{E}_{t+1}^{\text{post}} \wedge \mathcal{E}_t\right] \leq \delta', \tag{4.11}$$

$$\Pr\left[\overline{\mathcal{E}_{t+1}^{\text{post}}} \mid \mathcal{E}_t^{\text{post}} \wedge \mathcal{E}_{t-1}\right] \leq 7\delta'. \tag{4.12}$$

We proceed modulo these. Noting that

$$\Pr[\mathcal{E}_{t+1} \mid \mathcal{E}_t] = \Pr\left[\mathcal{E}_{t+1}^{\text{post}} \wedge \mathcal{E}_{t+1}^{\text{in}} \mid \mathcal{E}_t\right] = \Pr\left[\mathcal{E}_{t+1}^{\text{in}} \mid \mathcal{E}_{t+1}^{\text{post}} \wedge \mathcal{E}_t\right] \cdot \Pr\left[\mathcal{E}_{t+1}^{\text{post}} \mid \mathcal{E}_t\right],$$

the elementary inequality $1 - ab \leq (1-a) + (1-b)$ (for $a, b \leq 1$) implies

$$\Pr[\overline{\mathcal{E}_{t+1}} \mid \mathcal{E}_t] \leq \Pr\left[\overline{\mathcal{E}_{t+1}^{\text{in}}} \mid \mathcal{E}_{t+1}^{\text{post}} \wedge \mathcal{E}_t\right] + \Pr\left[\overline{\mathcal{E}_{t+1}^{\text{post}}} \mid \mathcal{E}_t\right]. \tag{4.13}$$

We can upper-bound the second term on the RHS of Equation (4.13) by (using that $\mathcal{E}_t \implies \mathcal{E}_{t-1}$):

$$\begin{aligned}\Pr\left[\overline{\mathcal{E}_{t+1}^{\text{post}}} \mid \mathcal{E}_t \wedge \mathcal{E}_{t-1}\right] &= \frac{\Pr\left[\overline{\mathcal{E}_{t+1}^{\text{post}}} \wedge \mathcal{E}_t \mid \mathcal{E}_{t-1}\right]}{\Pr[\mathcal{E}_t \mid \mathcal{E}_{t-1}]} \leq \frac{\Pr\left[\overline{\mathcal{E}_{t+1}^{\text{post}}} \wedge \mathcal{E}_t^{\text{post}} \mid \mathcal{E}_{t-1}\right]}{\Pr[\mathcal{E}_t \mid \mathcal{E}_{t-1}]} \\ &= \frac{\Pr\left[\overline{\mathcal{E}_{t+1}^{\text{post}}} \mid \mathcal{E}_t^{\text{post}} \wedge \mathcal{E}_{t-1}\right] \cdot \Pr\left[\mathcal{E}_t^{\text{post}} \mid \mathcal{E}_{t-1}\right]}{\Pr[\mathcal{E}_t \mid \mathcal{E}_{t-1}]} \leq \frac{\Pr\left[\overline{\mathcal{E}_{t+1}^{\text{post}}} \mid \mathcal{E}_t^{\text{post}} \wedge \mathcal{E}_{t-1}\right]}{\Pr[\mathcal{E}_t \mid \mathcal{E}_{t-1}]}\end{aligned}$$

where the first inequality used $\mathcal{E}_t \implies \mathcal{E}_t^{\text{post}}$. Now, we conclude using Equations (4.11) to (4.13):

$$\Pr\left[\overline{\mathcal{E}_{t+1}} \mid \mathcal{E}_t\right] \le \delta' + (7\delta')/(1-\delta) \le 15\delta' \le \delta/(2T),$$

proving Condition 2.

Proof of Equation (4.11). We prove that for every fixing of $M_{1:t}, S^{\Upsilon}_{1:t}$ such that $\mathcal{E}^{\text{post}}_{t+1}$ and $\mathcal{E}_t$ hold, the event $\mathcal{E}^{\text{in}}_{t+1}$ is likely. The probability is over the choice of $\boldsymbol{M}_{t+1}$, the one additional variable needed to determine $\mathcal{E}^{\text{in}}_{t+1}$; the conditional distribution of $\boldsymbol{M}_{t+1}$ is still uniformly random.

By definition of $\mathcal{E}^{\text{post}}_{t+1}$, we have that $\mathcal{D}_t$ is (C_t, s)-bounded and $\|\mu_{\mathcal{D}_t}\|_\infty \le q^{2tb}$. Therefore, applying the boundedness implies uniformity lemma (Lemma 4.7) with parameters $(\delta, C, \tau, \mathcal{D}, b)_{\text{Lemma 4.7}} = (\delta', C_t, \gamma, \mathcal{D}_t, 2tb)$ and Observation 4.8, we get

$$\Pr_{\boldsymbol{M}_{t+1}}\left[\|\text{Input}_{\phi_t,\mathcal{Y}_t}(\mathcal{D}_t, \boldsymbol{M}_{t+1}) - \text{Unif}(\mathbb{Z}_q^{m\times k})\|_{\text{tvd}} \le \delta'\right] \ge 1-\delta'.$$

Proof of Equation (4.12). We prove that for every fixing of $M_{1:t-1}, S^{\Upsilon}_{1:t-1}$ such that $\mathcal{E}^{\text{post}}_{t}$ and $\mathcal{E}_{t-1}$ holds, the event $\mathcal{E}^{\text{post}}_{t+1}$ is likely. The probability is over the choice of $\boldsymbol{M}_t, \boldsymbol{X}^*, \boldsymbol{Z}_t$, and the conditional marginal distribution of $\boldsymbol{M}_t$ is still uniformly random.

By definition of $\mathcal{E}^{\text{post}}_{t}$, $\mathcal{D}_{t-1}$ is (C_{t-1}, s) bounded and $\|\mu_{\mathcal{D}_{t-1}}\|_\infty \le q^{2(t-1)b}$. Applying the inductive step (Lemma 4.9) with parameters $(C, \delta, \sigma, C'', \mathcal{D}, b)_{\text{Lemma 4.9}} = (C_{t-1}, \delta', \gamma, C_t, \mathcal{D}_{t-1}, 2(t-1)b)$, we have

$$\Pr_{\boldsymbol{M}_t}\left[\begin{array}{l}\forall \mathcal{B} \subset \mathbb{Z}_q^{m\times k} \text{ such that } |\mathcal{B}| \ge q^{m\cdot k - b'},\\ \qquad \text{Post}(\mathcal{D}_{t-1}; \boldsymbol{M}_t, \mathcal{B}) \text{ is } (C_t, s)\text{-bounded and } \|\mu_{\text{Post}(\mathcal{D}_{t-1};\boldsymbol{M}_t,\mathcal{B})}\|_\infty \le q^{2tb}\end{array}\right] \ge 1 - 5\delta',$$

where we used the fact for our choice of parameters b, b', and δ', $q^{2(t-1)b + b' + \log_q(1/(1-\delta'))} \le q^{2tb}$.

Applying the boundedness implies uniformity lemma (Lemma 4.7) with the choice of parameters $(\delta, C, \tau, \mathcal{D}, b)_{\text{Lemma 4.7}} = (\delta', C_{t-1}, \gamma, \mathcal{D}_{t-1}, 2(t-1)b)$, we have

$$\Pr_{\boldsymbol{M}_t}\left[\|\mu_{\text{Input}(\mathcal{D}_{t-1}, \boldsymbol{M}_t)} - 1\|_\infty \le \delta'\right] \ge 1-\delta'.$$

Therefore, with probability at least $1-6\delta'$ (over the randomness of $\boldsymbol{M}_t$), we have

$$\Pr_{\boldsymbol{M}_t}\left[\begin{array}{l}\forall \mathcal{B} \subset \mathbb{Z}_q^{m\times k} \text{ such that } |\mathcal{B}| \ge q^{m\cdot k - b'},\\ \qquad \text{Post}(\mathcal{D}_{t-1}; \boldsymbol{M}_t, \mathcal{B}) \text{ is } (C_t, s)\text{-bounded and } \|\mu_{\text{Post}(\mathcal{D}_{t-1};\boldsymbol{M}_t,\mathcal{B})}\|_\infty \le q^{2tb}\\ \text{and } \|\mu_{\text{Input}(\mathcal{D}_{t-1},\boldsymbol{M}_t)} - 1\|_\infty \le \delta'\end{array}\right] \ge 1 - 6\delta'. \quad (4.14)$$

Now *fix* such an M_t. By Claim 4.15 below, with probability at least $1-\delta'$, we have $|\boldsymbol{\mathcal{B}}_t| \ge q^{m\cdot k - b'}$, which when applied to Equation (4.14) gives that $\boldsymbol{\mathcal{D}}_t$ is (C_t, s)-bounded and $\|\mu_{\boldsymbol{\mathcal{D}}_t}\|_\infty \le q^{2tb}$. Thus, $\mathcal{E}^{\text{post}}_t$ holds, as desired.

Claim 4.15. *For any $t \in [T]$, consider any fixed $M_{1:t}$ and $S^{\Upsilon}_{1:t-1}$ such that $\mathcal{D}_{t-1}$ is (C_{t-1}, s)-bounded and $\|\mu_{\text{Input}(\mathcal{D}_{t-1}, M_t)} - 1\|_\infty \le \delta'$. Then:*

$$\Pr[|\boldsymbol{\mathcal{B}}_t| \ge q^{m\cdot k - b'} \mid M_{1:t}, S^{\Upsilon}_{1:t-1}] \ge 1-\delta'.$$

Proof. For fixed $M_{1:t}, S^{\mathcal{Y}}_{1:t-1}$, the message $S^{\mathcal{Y}}_t = \mathtt{Player}_t(M_{1:t}, S^{\mathcal{Y}}_{1:t-1}, \boldsymbol{Z}_t)$ partitions the space $\mathbb{Z}_q^{m\times k}$ into at most 2^{s^*} sets. Let us denote this partition by $\mathcal{P}$. We now argue that in expectation, $\frac{1}{|\boldsymbol{\mathcal{B}}_t|}$ is small and then apply Markov's inequality to get the desired result. We have

$$\begin{aligned}\mathop{\mathbb{E}}_{\substack{\boldsymbol{X}^*\sim\mathcal{D}_{t-1},\\ \boldsymbol{Y}_t\sim\mathcal{Y}_t^{\otimes m}}}\left[\frac{1}{|\boldsymbol{\mathcal{B}}_t|}\right] &= \sum_{P\in\mathcal{P}} \frac{1}{|P|}\cdot \Pr_{\substack{\boldsymbol{X}^*\sim\mathcal{D}_{t-1},\\ \boldsymbol{Y}_t\sim\mathcal{Y}_t^{\otimes m}}}\left[\Pi_{M_t}^{\phi_t}\boldsymbol{X}^* - \boldsymbol{Y}_t \in P\right]\\ &\le \sum_{P\in\mathcal{P}} \frac{1}{|P|}\cdot |P|\,(1+\delta')\,q^{-m\cdot k}\\ &\le (1+\delta')\,q^{-m\cdot k}\,2^{s^*}\,.\end{aligned}$$

Therefore, applying Markov's inequality:

$$\Pr_{\substack{\boldsymbol{X}^*\sim\mathcal{D}_{t-1},\\ \boldsymbol{Y}_t\sim\mathcal{Y}_t^{\otimes m}}}[|\boldsymbol{\mathcal{B}}_t| \ge (1-\delta')q^{m\cdot k}2^{-s^*}(1+\delta')^{-1}] \ge 1-\delta'\,.$$

Finally, we use that $(1-\delta')q^{m\cdot k}2^{-s^*}(1+\delta')^{-1} \ge q^{m\cdot k-b'}$ from the choice of s^*. □

Condition 3. For every fixed $M_{1:t}$ and $S^{\mathcal{Y}}_{1:t-1}$ satisfying $\mathcal{E}_t$, we have that (by definition of $\mathcal{E}_t^{\text{in}}$ and Observation 4.8):

$$\|\mathrm{Input}_{\phi_t,\mathcal{Y}_t}(\mathcal{D}_{t-1}, M_t) - \mathsf{Unif}(\mathbb{Z}_q^{m\times k})\|_{\mathsf{tvd}} \le \delta'\,.$$

Applying the data-processing inequality (Proposition 2.2) and Claim 4.10, we conclude that

$$\|S^{\mathcal{Y}}_t - \mathtt{Player}_t(M_{1:t}, S^{\mathcal{Y}}_{1:t-1}, \boldsymbol{U})\|_{\mathsf{tvd}} \le \delta'\,,$$

as desired. □

5 Boundedness implies uniformity: Proof of Lemma 4.7

In this section, we prove Lemma 4.7, restated as follows:

Lemma 4.7 (Boundedness implies uniformity). *For every $k, q, k' \ge 2 \in \mathbb{N}$ ($k' \ge k$), there exists $\alpha_0 > 0$ such that for every $\delta \in (0, 1/2)$ and $C < \infty$, there exists $\tau > 0$ such that the following holds.*

For every injection $\phi : [k] \to [k']$, one-wise uniform distribution $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$, $n \in \mathbb{N}$ sufficiently large, $\mathcal{D} \in \Delta(\mathbb{Z}_q^{n\times k'})$ a (C,s)-bounded distribution satisfying $\|\mu_{\mathcal{D}}\|_\infty \le q^b$ and $4\log_q(3/\delta) \le b \le s \le \tau n$, $m \le \alpha_0 n$,

$$\Pr_{\boldsymbol{M}\in\mathcal{M}_n^{m,k}}\left[\|\mu_{\mathrm{Input}(\mathcal{D},\boldsymbol{M})} - 1\|_\infty \le \delta\right] \;\ge\; 1-\delta,$$

where $\boldsymbol{M} \in \mathcal{M}_n^{m,k}$ is a randomly sampled matching.

In order to prove the lemma, we begin with some definitions.

Definition 5.1 (Singleton-freeness). We say a set $S \subseteq [m]\times[k]$ is *singleton-free* if for every $j \in [m]$, $|S \cap (\{j\}\times[k])| \ne 1$. (If we view S equivalently as a matrix in $\{0,1\}^{m\times k}$, then S is singleton-free if no row has weight exactly 1.) We let $\mathrm{SF}^{m,k} := \{S \subseteq [m]\times[k] : S \text{ singleton-free}\}$; when m and k are clear from context, we write $\mathrm{SF} = \mathrm{SF}^{m,k}$. ◊

Proof. Follows immediately from Propositions 2.14 and 2.15. □

The following simple proposition then forms the heart of our analysis:

Proposition 5.2. *Let $k, m \in \mathbb{N}$, $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$ be a one-wise uniform distribution, and $U \in \mathbb{Z}_q^{m\times k}$. If $\mathsf{supp}(U) \notin \mathrm{SF}^{m,k}$ (i.e., there exists $j \in [m]$ with $|\mathsf{supp}(U_j)| = 1$, where U_j is the j-th row of U), then $\widehat{\mu_{\mathcal{Y}^{\otimes m}}}(U) = 0$.*

Recall the definition of the embedding $\iota_M : [m] \times [k] \to [n] \times [k]$ corresponding to a matching M (Definition 3.1). In particular,

$$\iota_M(\mathrm{SF}) = \{U \subseteq [n] \times [k] : U \subseteq \mathsf{supp}(M) \text{ and } \forall j \in [m],\ |U \cap \mathsf{supp}(M_j))| \neq 1\}. \tag{5.3}$$

For an injection $\phi : [k] \to [k']$, we also consider the set $\phi(\iota_M(\mathrm{SF})) = \{(j, \phi(\ell)) : (j, \ell) \in \mathcal{S}\}$ (as in Definition 3.2).

5.1 Bounding ∞-distance from uniform via 1-norm of "singleton-free" Fourier mass

We first prove the following upper-bound on the ∞-distance from uniform of $\mathrm{Input}(\mathcal{D}, M)$'s density function in terms of the (1-norm) Fourier mass of $\mu_{\mathcal{D}}$ on "singleton-free" coefficients. This is the key reason we need to consider singleton-free coefficients in this paper.

Lemma 5.4. *For every $q, m, k, n, k' \in \mathbb{N}$, injection $\phi : [k] \to [k']$, k-partite hypermatching $M \in \mathcal{M}_n^{m,k}$, and one-wise uniform distribution $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$, we have:*

$$\|\mu_{\mathrm{Input}(\mathcal{D},M)} - 1\|_\infty \leq \sum_{\substack{U \neq 0 \in \mathbb{Z}_q^{n\times n'},\\ \mathsf{supp}(U) \in \phi(\iota_M(\mathrm{SF}))}} |\widehat{\mu_{\mathcal{D}}}(U)|$$

where, as in Lemma 4.7, $\mathrm{Input}(\mathcal{D}, M)$ is the distribution of $\Pi_M^\phi \boldsymbol{X} - \boldsymbol{Y}$ for $\boldsymbol{X} \sim \mathcal{D}$ and $\boldsymbol{Y} \sim \mathcal{Y}^{\otimes m}$, and $\iota_M(\mathrm{SF})$ is as in Equation (5.3).

Proof. Let $\mathcal{P}_M \in \Delta(\mathbb{Z}_q^{m\times k})$ denote the distribution of the projection $\Pi_M^\phi \boldsymbol{X}$ for $\boldsymbol{X} \sim \mathcal{D}$. Hence $\mathrm{Input}(\mathcal{D}, M)$ is the distribution of sum of independent samples from $\mathcal{P}_M$ and $-\mathcal{Y}^{\otimes m}$. Consequently, by Proposition 2.6, we have $\mu_{\mathrm{Input}(\mathcal{D},M)} = \mu_{\mathcal{P}_M} * \mu_{-\mathcal{Y}^{\otimes m}}$ and therefore by Propositions 2.9 and 2.10, for every $Z_0 \in \mathbb{Z}_q^{m\times k}$, we have:

$$\mu_{\mathrm{Input}(\mathcal{D},M)}(Z_0) = \sum_{U \in \mathbb{Z}_q^{m\times k}} \widehat{\mu_{\mathcal{P}_M} * \mu_{-\mathcal{Y}^{\otimes m}}}(U) \cdot \chi_U(Z_0) = \sum_{U \in \mathbb{Z}_q^{m\times k}} \widehat{\mu_{\mathcal{P}_M}}(U) \cdot \widehat{-\mu_{\mathcal{Y}^{\otimes m}}}(U) \cdot \chi_U(Z_0).$$

Hence, using the triangle inequality, that the zero Fourier coefficient is 1 for a density function, and $|\chi_U(Z_0)| = 1$:

$$|\mu_{\mathrm{Input}(\mathcal{D},M)}(Z_0) - 1| \leq \sum_{U \neq 0 \in \mathbb{Z}_q^{m\times k}} |\widehat{\mu_{\mathcal{P}_M}}(U)| \cdot |\widehat{\mu_{-\mathcal{Y}^{\otimes m}}}(U)|.$$

The quantity on the RHS is independent of Z_0 and thus gives an upper-bound on $\|\mu_{\mathrm{Input}(\mathcal{D},M)} - 1\|_\infty$.

By Proposition 2.16 and the definition of $\mathcal{P}_M$, we know that for every $U \in \mathbb{Z}_q^{m\times k}$,

$$\widehat{\mu_{\mathcal{P}_M}}(U) = \widehat{\mu_{\mathcal{D}}}((\Pi_M^\phi)^\intercal(U)).$$

Consequently, using the trivial estimate $|\widehat{\mu_{-\mathcal{Y}^{\otimes m}}}(U)| \leq 1$, we get by Proposition 5.2:

$$\|\mu_{\mathrm{Input}(\mathcal{D},M)} - 1\|_\infty \leq \sum_{\substack{U \neq 0 \in \mathbb{Z}_q^{m\times k},\\ \mathsf{supp}(U) \in \mathrm{SF}}} |\widehat{\mu_{\mathcal{D}}}((\Pi_M^\phi)^\intercal(U))|.$$

Finally, we observe that $\{(\Pi_M^\phi)^\intercal U : U \neq 0 \in \mathbb{Z}_q^{m\times k} \text{ and } \mathsf{supp}(U) \in \mathrm{SF}\} = \{U \neq 0 \in \mathbb{Z}_q^{n\times k'} : \mathsf{supp}(U) \in \phi(\iota_M(\mathrm{SF}))\}$ (just by unpacking definitions). □

5.2 Combinatorial bound

We next prove an upper-bound on the probability that a fixed set U is in $\iota_{\boldsymbol{M}}(\mathrm{SF})$ when $\boldsymbol{M}$ is sampled randomly in terms of the cardinality $|U|$ of U.

Lemma 5.5. *For every $k \in \mathbb{N}$, there exists $C < \infty$ and $\alpha_0 > 0$ such that the following holds. Let $n \in \mathbb{N}$, $m \leq \alpha_0 n \in \mathbb{N}$, $U \subset [n] \times [k]$, and $h := |U|$. Then for every $2 \leq h \leq km \in \mathbb{N}$, sampling a random k-partite hypermatching $\boldsymbol{M} \in \mathcal{M}_n^{m,k}$, we have*

$$\Pr[U \in \iota_{\boldsymbol{M}}(\mathrm{SF})] \leq \left(\frac{Ch}{n}\right)^{h/2}.$$

(The probability vanishes at $h = 1$ and $h > km$.)

Proof. Let $\mathcal{V}_\ell := [n] \times \{\ell\}$ so that $[n] \times [k] = \bigsqcup_{\ell=1}^{k} \mathcal{V}_\ell$. Let $h_\ell := |U \cap \mathcal{V}_\ell|$.

Step 1: Symmetry reduction. Define $\mathcal{U} := \{U' \subset [n] \times [k] : |U' \cap \mathcal{V}_\ell| = h_\ell\}$. To calculate $\Pr[U \in \iota_{\boldsymbol{M}}(\mathrm{SF})]$, by column-wise permutation symmetry, it is equivalent to instead fix one matching, such as

$$M := \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 2 & 2 & \cdots & 2 \\ \vdots & \vdots & \ddots & \vdots \\ m & m & \cdots & m \end{bmatrix},$$

sample $\boldsymbol{U}$ uniformly at random from the set $\mathcal{U} := \{U \subset [n] \times [k] : \forall \ell \in [k],\ |U \cap \mathcal{V}_\ell| = h_\ell\}$, and calculate $\Pr[\boldsymbol{U} \in \iota_M(\mathrm{SF})]$. That is, defining $\mathcal{N} := \mathcal{U} \cap \iota_M(\mathrm{SF})$, we want to calculate $|\mathcal{N}|/|\mathcal{U}|$.

Step 2: Counting and partitioning. For $U \in \mathcal{N}$, we define $D(U) := \{j \in [m] : U \cap \mathsf{supp}(M_j) \neq \emptyset\}$. Note that if $U \in \mathcal{N}$ and $j \in D(U)$ then $|U \cap \mathsf{supp}(M_j)| \geq 2$ by definition of $\mathcal{N}$. Hence, for $d(u) := |D(U)|$, we have $d(U) \leq |U|/2 = h/2$.

For a fixed size d, there are $\binom{m}{d}$ sets $D \subseteq [m]$ of size d, and for a fixed such set D, there are at most $\binom{dk}{h}$ possible sets $U \in \mathcal{N}$ with $D(U) = U$ (since U must be fully supported on the rows in D). Hence we have:

$$|\mathcal{N}| \leq \sum_{d=1}^{h/2} \binom{\alpha n}{d} \binom{dk}{h}.$$

Conversely, we have

$$|\mathcal{U}| = \prod_{\ell=1}^{k} \binom{n}{h_\ell} \geq \binom{n}{h}.$$

Step 3: Estimation. Note that for every $h \in [1, km]$ and $d \leq h/2$, we have $\binom{m}{d} \leq \left(\frac{4km}{h}\right)^{h/2}$, since by cases:

$$h \in [1, m] \implies \binom{m}{d} \leq \left(\frac{em}{d}\right)^{d} \leq \left(\frac{em}{h/2}\right)^{h/2} \leq \left(\frac{4km}{h}\right)^{h/2},$$
$$h \in [m, km] \implies \binom{m}{d} \leq 2^m \leq 4^{h/2} \leq \left(\frac{4km}{h}\right)^{h/2}.$$

For the second inequality on the first line, we used that $(C/x)^x$ is an increasing function of x on the interval $x \in (0, C/e)$. We also have, by binomial estimates, $\binom{dk}{h} \le \left(\frac{edk}{h}\right)^h \le (2k)^h$ (using that $ed \le eh/2 \le 2h$) and $\binom{n}{h} \ge \left(\frac{n}{h}\right)^h$. Using $m \le \alpha_0 n$ and the trivial bound $h/2 \le 2^{h/2}$, we altogether get a bound of

$$\frac{|\mathcal{N}|}{|\mathcal{U}|} \le \frac{h}{2} \cdot \frac{\left(\frac{4k\alpha_0 n}{h}\right)^{h/2} \cdot (2k)^h}{\left(\frac{n}{h}\right)^h} = 2^{h/2} \cdot \frac{\left(\frac{4k\alpha_0 n}{h}\right)^{h/2} \cdot (4k^2)^{h/2}}{\left(\frac{n^2}{h^2}\right)^{h/2}} = \left(\frac{32k^3\alpha h}{n}\right)^{h/2},$$

as desired (setting $C := 32k^3\alpha_0$). □

5.3 Finishing the proof

We now use the following fact from [CGS+22, Lemma 5.18]:

Lemma 5.6. *For every $k, q \ge 2 \in \mathbb{N}$, $C_1, C_2 < \infty$ and $\delta \in (0,1)$, there exists $\tau > 0$ such that the following holds.*

For all $n, s \in \mathbb{N}$ satisfying $4\log_q(3/\delta) \le s \le \tau n$, we have:

$$\sum_{h=2}^{km} \mathrm{U}_{C_1,s,n}(h) \cdot \left(\frac{C_2 h}{n}\right)^{h/2} \le \delta^2.$$

Proof of Lemma 4.7. Using Markov's inequality, it suffices to prove that

$$\mathop{\mathbb{E}}_{\boldsymbol{M}}\left[\|\mu_{\mathrm{Input}(\mathcal{D},\boldsymbol{M})} - 1\|_\infty\right] \le \delta^2. \tag{5.7}$$

By Lemma 5.4 and linearity of expectation, we have:

$$\begin{aligned}\mathop{\mathbb{E}}_{\boldsymbol{M}}\left[\|\mu_{\mathrm{Input}(\mathcal{D},\boldsymbol{M})} - 1\|_\infty\right] &\le \mathop{\mathbb{E}}_{\boldsymbol{M}}\left[\sum_{\substack{U \neq 0 \in \mathbb{Z}_q^{n\times k'},\\ \mathsf{supp}(U) \in \phi(\iota_{\boldsymbol{M}}(\mathrm{SF}))}} |\widehat{\mu_{\mathcal{D}}}(U)|\right] \\ &= \sum_{U \neq 0 \in \mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(U)| \cdot \Pr_{\boldsymbol{M}}[\mathsf{supp}(U) \in \phi(\iota_{\boldsymbol{M}}(\mathrm{SF}))].\end{aligned}$$

Now, consider a fixed $U \neq 0 \in \mathbb{Z}_q^{n\times k'}$ with size $h := |U|$. We know that $\Pr_{\boldsymbol{M}}[\mathsf{supp}(U) \in \phi(\iota_{\boldsymbol{M}}(\mathrm{SF}))] = 0$ if $h < 2$ or $h > km$. On the other hand, if $2 \le h \le km$, then we claim that $\Pr_{\boldsymbol{M}}[\mathsf{supp}(U) \in \phi(\iota_{\boldsymbol{M}}(\mathrm{SF}))] \le \left(\frac{C_2 h}{n}\right)^{h/2}$ (where C_2 is the constant coming from Lemma 5.5). Indeed, we observe that if $\mathsf{supp}(U) \not\subseteq \phi([n] \times [k])$ then if $\Pr_{\boldsymbol{M}}[\mathsf{supp}(U) \in \phi(\iota_{\boldsymbol{M}}(\mathrm{SF}))] = 0$, while if $\mathsf{supp}(U) \subseteq \phi([n] \times [k])$ then $\Pr_{\boldsymbol{M}}[\mathsf{supp}(U) \in \phi(\iota_{\boldsymbol{M}}(\mathrm{SF}))] = \Pr_{\boldsymbol{M}}[\phi^{-1}(\mathsf{supp}(U)) \in \iota_{\boldsymbol{M}}(\mathrm{SF})] \le \left(\frac{C_2 h}{n}\right)^{h/2}$ by Lemma 5.5. Altogether, these calculations give

$$\mathop{\mathbb{E}}_{\boldsymbol{M}}[\|\mu_{\mathrm{Input}(\mathcal{D},\boldsymbol{M})} - 1\|_\infty] \le \sum_{h=2}^{km} \left(\sum_{\substack{U \neq 0 \in \mathbb{Z}_q^{n\times k'},\\ \mathrm{wt}(U)=h}} |\widehat{\mu_{\mathcal{D}}}(U)|\right) \cdot \left(\frac{C_2 h}{n}\right)^{h/2} \le \sum_{h=2}^{km} \mathrm{U}_{C_1,s,k'n}(h) \cdot \left(\frac{(C_2 k')h}{k'n}\right)^{h/2},$$

which is then bounded by the foregoing lemma (Lemma 5.6). □

6 A "singleton-free" level inequality

In this section, we derive upper bounds on the ℓ_1 norm of a certain subset of Fourier coefficients over Hamming balls — centered at the origin at first, and then alter at arbitrary centers in §6.3 below. These bounds are later used in the proof of Lemma 4.9 in §7. We begin with some definitions.

The *weight* of a frequency matrix $U \in \mathbb{Z}_q^{m\times k}$ is $\mathrm{wt}(U) := |\mathsf{supp}(U)|$. As a convenient abbreviation, we define the notation below for sets of "singleton-free" frequency matrices (where the set $\mathrm{SF}^{m,k}$ is as in Definition 5.1.)

$$\begin{aligned}\mathcal{U}_q^{m,k} &:= \{U \in \mathbb{Z}_q^{m,k} : \mathsf{supp}(U) \in \mathrm{SF}^{m,k}\}\\ \mathcal{U}_q^{m,k}(h) &:= \{U \in \mathcal{U}_q^{m,k} : \mathrm{wt}(U) = h\}\end{aligned}$$

In particular, $\mathcal{U}_q^{m,k} = \{U \in \mathbb{Z}_q^{m\times k} : \forall j \in [m],\ |\mathsf{supp}(U_j)| \neq 1\}$, where U_j is the j-th row of U. The goal of this section is to prove the following "refined" level inequality bounding the ℓ_1-norm of singleton-free Fourier coefficients:

Theorem 6.1. *For every $q, k \in \mathbb{N}$ and $\theta < \infty$, there exists $\zeta < \infty$ and $\varepsilon_0 > 0$ such that the following holds. For every $m, b \in \mathbb{N}$, $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ with $\|g\|_1 = 1$, and $h \in \mathbb{N}$ with $h < \theta b$ and $\log_q(\|g\|_\infty) \le b \le \varepsilon_0 km$, we have:*

$$\sum_{U\in\mathcal{U}_q^{m,k}(h)} |\widehat{g}(U)| \le \left(\frac{\zeta\sqrt{bm}}{h}\right)^{h/2}.$$

Writing $U_j \in \mathbb{Z}_q^k$ for row j of a matrix $U \in \mathbb{Z}_q^{m\times k}$, we define the *row weight* of $U \in \mathbb{Z}_q^{m\times k}$ as the number of non-zero rows $\mathrm{rwt}(U) := |\{j : U_j \neq 0\}|$ (i.e., the number of nonzero rows). We also denote by $\mathcal{U}_q^{m,k}(h,\ell)$, the set of set of singleton-free frequencies with weight h and row-weight ℓ

$$\mathcal{U}_q^{m,k}(h,\ell) := \{U \in \mathcal{U}_q^{m,k}(h) : \mathrm{rwt}(U) = \ell\}.$$

We will prove the above theorem by bounding the ℓ_2 Fourier mass for the sets $\mathcal{U}_q^{m,k}(h,\ell)$, together with a counting argument and Cauchy-Schwarz to obtain ℓ_1 bounds.

6.1 Bounding the squared Fourier mass on $\mathcal{U}_q^{m,k}(h,\ell)$

We now prove the following bound on the ℓ_2 Fourier mass of the set $\mathcal{U}_q^{m,k}(h,\ell)$.

Lemma 6.2. *For every $q, k \in \mathbb{N}$ and $\theta > 0$, there exists $\zeta > 0$ such that the following holds. For every $m, b \in \mathbb{N}$, $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ with $\|g\|_1 = 1$, and $h > \ell \in \mathbb{N}$ with $h - \ell < \theta b$ and $\log_q(\|g\|_\infty) \le b \le \varepsilon_0 km$, we have:*

$$\sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} |\widehat{g}(U)|^2 \le \left(\frac{\zeta b}{h-\ell}\right)^{h-\ell}.$$

We will prove the lemma using hypercontractivity for a noisy version of g. We consider two different noise operators below.

6.1.1 The row noise operator and row hypercontractivity

Definition 6.3. A *multiplier operator* is an operator on $L^2(\mathbb{Z}_q^{m\times k})$ (the vector space of functions $\mathbb{Z}_q^{m\times k} \to \mathbb{C}$) which is diagonalizable in the Fourier basis. That is, letting $\Lambda : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ be a *multiplier function*, the corresponding *multiplier operator*, denoted $T_\Lambda : L^2(\mathbb{Z}_q^{m\times k}) \to L^2(\mathbb{Z}_q^{m\times k})$ is defined on the Fourier basis via $T_\Lambda \chi_U := \Lambda(U)\chi_U$ for every frequency $U \in \mathbb{Z}_q^{m\times k}$. Hence, for arbitrary $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$, we have $\widehat{T_\Lambda g}(U) = \Lambda(U) \cdot \widehat{g}(U)$. Equivalently, $T_\Lambda g = g * \Lambda'$, where $*$ denotes convolution and Λ' is the inverse Fourier transform of Λ. ◊

Definition 6.4 (Row noise)**.** For $\rho \in [0,1]$, we define:

- $\mathcal{D}_\rho^* \in \Delta(\mathbb{Z}_q^k)$, is the distribution which outputs 0 with probability ρ and a uniformly random element of $\mathbb{Z}_q^k$ with probability $1-\rho$.
- The Fourier transform $\lambda_\rho^* := \widehat{\mu_{\mathcal{D}_\rho^*}}$, which we view as a multiplier function. We can calculate $\lambda_\rho^*(u) = \begin{cases} 1 & \text{if } u = 0, \\ \rho & \text{if } u \neq 0 \end{cases}$.
- For $U \in \mathbb{Z}_q^{m\times k}$, $\Lambda^* := \widehat{\mu_{(\mathcal{D}_\rho^*)^{\otimes m}}}$. (Equivalently, $\Lambda^*(U) = \prod_{j=1}^m \lambda_\rho^*(U_j) = \rho^{\mathrm{rwt}(U)}$ where $\mathrm{rwt}(U) = |\{i : U_i \neq 0\}|$ is the row weight.)
- The corresponding *row noise operator* $T_\rho^* := T_{\lambda_\rho^*}$, which acts by convolution with $\mu_{\mathcal{D}_\rho^*}$. That is,

$$(T_\rho^* g)(X) = \mathop{\mathbb{E}}_{\boldsymbol{Y}\sim(\mathcal{D}_\rho^*)^{\otimes m}}[g(X - \boldsymbol{Y})].$$ ◊

Viewing each row as a single coordinate supported on q^k different values, we get the following hypercontractive inequality:

Lemma 6.5 (Row hypercontractive inequality, e.g., [O'D14])**.** *Let $Q := q^k$. For any $1 < p < 2$ and $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$, it holds that*

$$\|T_\rho^* g\|_2 \le \|g\|_p \qquad \textit{where} \qquad \rho := \sqrt{p-1} \cdot Q^{1/2-1/p}.$$

6.1.2 A modified noise operator

Definition 6.6. For $\rho \in [0,1]$, we define:

- The *modified noise multiplier* (for a single row)

$$\lambda_\rho(u) := \begin{cases} 1 & \text{if } u = 0, \\ 0 & \text{if } |\mathsf{supp}(u)| = 1, \\ \rho^{|\mathsf{supp}(u)|-1} & \text{otherwise.} \end{cases}$$

- The corresponding modified noise multiplier for a matrix: $\Lambda_\rho(U) := \prod_{j=1}^m \lambda_\rho(U_j)$.
- The corresponding *modified noise operator* $T_\rho := T_{\Lambda_\rho}$. ◊

Claim 6.7. *For every $U \in \mathcal{U}_q^{m,k}(h,\ell)$ and $\rho \in [0,1]$, $\Lambda_\rho(U) = \rho^{h-\ell}$.*

Proof. Taking the logarithm base ρ of both sides, on the LHS, rows with support 0 contribute 0, and a row j with support at least 2 contributes $|\mathsf{supp}(U_j)| - 1$. We then calculate $\sum_{j:|\mathsf{supp}(U_j)|\geq 2}(|\mathsf{supp}(U_j)| - 1) = h - \ell$. □

Claim 6.8. *For every $u \in \mathbb{Z}_q^k$ and $\rho \in [0,1]$, it holds that $\lambda_\rho(u) \leq \lambda_\rho^*(u)$.*

Proof. If $u = 0$, we have $\lambda_\rho(u) = \lambda_\rho^*(u) = 1$. Otherwise, $\lambda_\rho(u) \leq \rho$ by inspection, whereas $\lambda_\rho(u) = \rho$. □

Claim 6.9. *For every $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ and $\rho \in [0,1]$, it holds that $\|T_\rho g\|_2 \leq \|T_\rho^* g\|_2$.*

Proof. We calculate

$$\|T_\rho g\|_2^2 = \sum_{U\in\mathbb{Z}_q^{m\times k}} (\Lambda_\rho(U))^2 |\widehat{g}(U)|^2 \leq \sum_{U\in\mathbb{Z}_q^{m\times k}} (\Lambda_\rho^*(U))^2 |\widehat{g}(U)|^2 = \|T_\rho^* g\|_2^2. \qquad \square$$

Claim 6.10. *For every $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ and $r \in [0,1]$, it holds that $\sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} |\widehat{g}(U)|^2 \leq \rho^{-2(h-\ell)} \|T_\rho^* g\|_2^2$.*

Proof. By Claims 6.7 and 6.9, we have:

$$\begin{aligned}\rho^{2(h-\ell)} \sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} |\widehat{g}(U)|^2 &= \sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} \rho^{2(h-\ell)} |\widehat{g}(U)|^2 = \sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} (\Lambda_\rho(U))^2 |\widehat{g}(U)|^2 \\ &\leq \sum_{U\in\mathbb{Z}_q^{m\times k}} (\Lambda_\rho(U))^2 |\widehat{g}(U)|^2 = \|T_\rho g\|_2^2 \leq \|T_\rho^* g\|_2^2. \qquad \square\end{aligned}$$

6.1.3 Proof of Lemma 6.2

We are now equipped to prove Lemma 6.2.

Proof of Lemma 6.2. We proceed in two steps.

Step 1: Row hypercontractivity. Let $p \in (1,2)$. By Lemma 6.5,

$$\|T_\rho^* g\|_2^2 \leq \|g\|_p^2.$$

Since $\|g\|_1 = 1$ and $\|g\|_\infty \leq q^b$:

$$\|g\|_p^p = \mathbb{E}[|g|^p] = \mathbb{E}[|g| \cdot |g|^{p-1}] \leq \|g\|_\infty^{p-1} \mathbb{E}[|g|] = q^{b(p-1)}.$$

Hence $\|g\|_p^2 \leq q^{2b(1-1/p)}$.

Step 2: Setting parameters. Set $Q := q^k$ and $p := 1 + \frac{h-\ell}{\theta b}$. This is less than 2 since $h - \ell < \theta b$. Combining Claim 6.10 and Step 1:

$$\sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} |\widehat{g}(U)|^2 \leq \underbrace{\rho^{-2(h-\ell)}}_{=:W_1} \cdot \underbrace{q^{2b(1-1/p)}}_{=:W_2}.$$

We bound the terms separately as:

$$W_1 = \left(\sqrt{p-1}\cdot Q^{1/2-1/p}\right)^{-2(h-\ell)} = \left(\frac{\theta b}{h-\ell}\cdot Q^{2/p-1}\right)^{h-\ell} \leq \left(\frac{(\theta q^k) b}{h-\ell}\right)^{h-\ell},$$

where we used $p \geq 1 \implies Q^{2/p-1} \leq Q$ and $Q = q^k$, and

$$W_2 = q^{2b(1-1/p)} \leq q^{2b(p-1)} = (q^{2/\theta})^{h-\ell},$$

where the inequality used that $1 - 1/p \leq p - 1$ for $p > 0$. This gives the desideratum with $\zeta := \theta q^{k+2/\theta}$. □

6.2 Bounding the Fourier mass on $\mathcal{U}_q^{m,k}(h)$

We now return to proving Theorem 6.1. This first requires the following lemma:

Lemma 6.11. *For every $q, k \in \mathbb{N}$, there exists $\zeta > 0$ such that the following holds. For every $m \in \mathbb{N}$ and $h > \ell \in \mathbb{N}$, we have:*

$$|\mathcal{U}_q^{m,k}(h,\ell)| \leq \left(\frac{\zeta m}{\ell}\right)^{\ell}.$$

Proof. We severely overcount and give an upper bound on the set $\{U \in \mathbb{Z}_q^{m\times k} : \mathrm{rwt}(U) = \ell\}$. Each $U \in \mathcal{U}_q^{m,k}(h,\ell)$ is uniquely determined by specifying, for every nonzero row j (i.e., $U_j \neq 0$), the row vector U_j. It therefore suffices to count the number of ways to choose a subset of nonzero rows and then assign these rows to (nonzero) row vectors. There are $\binom{m}{\ell}$ ways to choose a set of ℓ rows to be nonzero. Each row then can be assigned to at most $q^k - 1$ vectors. Hence, altogether we get a bound of

$$\binom{m}{\ell} \cdot (q^k-1)^{\ell} \leq \left(\frac{em}{\ell}\right)^{\ell} \cdot (q^k-1)^{\ell} = \left(\frac{\zeta m}{\ell}\right)^{\ell}$$

for $\zeta := e(q^k - 1)$, as desired. □

We can now prove Theorem 6.1.

Proof of Theorem 6.1. Let $\zeta_1, \zeta_2 > 0$ be the two constants from Lemmas 6.2 and 6.11, respectively. By Cauchy-Schwarz, and since $h/k \leq \mathrm{rwt}(U) \leq h/2$ for every $U \in \mathcal{U}_q^{m,k}(h)$, we have:

$$\sum_{U\in\mathcal{U}_q^{m,k}(h)} |\widehat{g}(U)| = \sum_{\ell=h/k}^{h/2} \sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} |\widehat{g}(U)| \leq \sum_{\ell=h/k}^{h/2} \left(\underbrace{|\mathcal{U}_q^{m,k}(h,\ell)| \cdot \sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} |\widehat{g}(U)|^2}_{=:S_\ell} \right)^{1/2}.$$

We can bound each individual S_ℓ via Lemmas 6.2 and 6.11:

$$\begin{aligned} S_\ell \;=\; |\mathcal{U}_q^{m,k}(h,\ell)| \cdot \sum_{U\in\mathcal{U}_q^{m,k}(h,\ell)} |\widehat{g}(U)|^2 \;&\leq\; \left(\frac{\zeta_1 b}{h-\ell}\right)^{h-\ell} \cdot \left(\frac{\zeta_2 m}{\ell}\right)^{\ell} \\ &\leq \left(\frac{2\zeta_1 b}{h}\right)^{h-\ell} \cdot \left(\frac{\zeta_2 m}{\ell}\right)^{\ell} = \left(\frac{2\zeta_1 b}{h}\right)^{h} \cdot \left(\frac{\zeta_2 hm}{2\zeta_1 b\ell}\right)^{\ell} \\ &\leq \left(\frac{2\zeta_1 b}{h}\right)^{h} \cdot \left(\frac{\zeta_2 km}{2\zeta_1 b}\right)^{\ell}, \end{aligned}$$

where the last two inequalities use $\ell \leq h/2$ and $\ell \geq h/k$, respectively. Hence

$$\begin{aligned} \sum_{U\in\mathcal{U}_q^{m,k}(h)} |\widehat{g}(U)| \;&\leq\; \sum_{\ell=h/k}^{h/2} \sqrt{S_\ell} \;\leq\; \left(\frac{2\zeta_1 b}{h}\right)^{h/2} \cdot \sum_{\ell=h/k}^{h/2} \left(\sqrt{\frac{\zeta_2 km}{2\zeta_1 b}}\right)^{\ell} \\ &\leq \left(\frac{2\zeta_1 b}{h}\right)^{h/2} \cdot 2\left(\sqrt{\frac{\zeta_2 km}{2\zeta_1 b}}\right)^{h/2} = 2\left(\frac{\sqrt{2\zeta_1\zeta_2 kbm}}{h}\right)^{h/2}, \end{aligned}$$

where the final inequality sets $\varepsilon_0 := \frac{\zeta_2}{8\zeta_1}$, forcing $\sqrt{\frac{\zeta_2 km}{2\zeta_1 b}} \geq 2$ (using that $b \leq \varepsilon_0 km$).[8] □

[8] Note that if $X \geq 2$ then $\sum_{j=0}^{H} X^j = \frac{X^{j+1}-1}{X-1} \leq \frac{X^{j+1}}{X/2} = 2X$, where we used $X - 1 \geq X/2$.

6.3 "Singleton-free" level inequality around arbitrary centers

In this section, we "boost" the theorem which we stated at the beginning of this section (Theorem 6.1) to give bounds for singleton-free matrices in balls around arbitrary points.

We begin with the following simple corollary of Theorem 6.1:

Corollary 6.12. *For every $q, k \in \mathbb{N}$ and $\theta < \infty$, there exists $\zeta < \infty$ and $\varepsilon_0 > 0$ such that the following holds. For every $m, b \in \mathbb{N}$, $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ with $\|g\|_1 = 1$, $V \in \mathbb{Z}_q^{m\times k}$, $h \in \mathbb{N}$ with $h < \theta b$ and $\log_q(\|g\|_\infty) \le b \le \varepsilon_0 km$, we have:*

$$\sum_{U-V\in\mathcal{U}_q^{m,k}(h)} |\widehat{g}(U)| \le \mathrm{U}_{C,s,km}(h).$$

Proof. Define $g' : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ via $g'(U) := g(U)\cdot\overline{\chi_V(U)}$, so that $\widehat{g'}(U) = \mathbb{E}_{\boldsymbol{X}\in\mathbb{Z}_q^{m\times k}}[g'(\boldsymbol{X})\cdot\overline{\chi_V(\boldsymbol{X})}] = \mathbb{E}_{\boldsymbol{X}\in\mathbb{Z}_q^{m\times k}}[g(\boldsymbol{X})\cdot\overline{\chi_{U+V}(\boldsymbol{X})}] = \widehat{g}(U+V)$. Hence,

$$\sum_{U-V\in\mathcal{U}_q^{m,k}(h)} |\widehat{g}(U)| = \sum_{U\in\mathcal{U}_q^{m,k}(h)} |\widehat{g}(U+V)| = \sum_{U\in\mathcal{U}_q^{m,k}(h)} |\widehat{g'}(U)|,$$

at which point we can apply Theorem 6.1 (since g' is just a phase shift of g, we have $\|g\|_1 = \|g'\|_1$ and $\|g\|_\infty = \|g'\|_\infty$). □

Lemma 6.13. *For every $q, k \in \mathbb{N}$ and $\theta < \infty$, there exists $\zeta < \infty$ and $\varepsilon_0 > 0$ such that the following holds. For every $m \in \mathbb{N}$, $g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ with $\|g\|_1 = 1$, $V \in \mathbb{Z}_q^{m\times k}$, and $h \in \mathbb{N}$ with $h < \theta b$, we have:*

$$\sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U-V)=h}} |\widehat{g}(U)| \le q^{k\cdot\mathrm{rwt}(V)} \cdot \left(\frac{\zeta\sqrt{bm}}{h-\kappa}\right)^{(h-\kappa)/2},$$

where $\log_q(\|g\|_\infty) \le b \le \varepsilon_0 km$ and $\kappa := |\{j \in [m] : |\mathsf{supp}(V_j)| = 1\}|$.

Proof. Fix $V \in \mathbb{Z}_q^{m\times k}$. For $U \in \mathbb{Z}_q^{m\times k}$, define a matrix $Z(U) \in Z_q^{m\times k}$ row-wise, via

$$(Z(U))_j := -(U+V)_j \cdot [\![|\mathsf{supp}(V_j)| = 1 \vee |\mathsf{supp}((U+V)_j)| = 1]\!].$$

We have a few basic facts:

- For every $U \in \mathbb{Z}_q^{m\times k}$, $U + V + Z(U) \in \mathrm{SF}$ by definition.
- If $U \in \mathrm{SF}$, then for every $j \in [m]$, $(Z(U))_j \neq 0 \implies V_j \neq 0$. Indeed, by contrapositive, if $V_j = 0$, then $(U+V)_j = U_j$, and hence $|\mathsf{supp}((U+V)_j)| = |\mathsf{supp}(U_j)| \neq 1$ since $U \in \mathrm{SF}$.
- $\mathrm{wt}(Z(U)) \ge \kappa$.
- If $\mathrm{wt}(U+V) = h$, then $\mathrm{wt}(U+V+Z(U)) = h - \mathrm{wt}(Z(U))$.

Let $\mathcal{Z} := \{Z \in \mathbb{Z}_q^{m\times k} : \forall j \in [m], Z_j \neq 0 \implies V_j \neq 0\}$ denote the set of frequencies supported only on rows in V's support. By the second item above, we have $Z(U) \in \mathcal{Z}$ for every $U \in \mathrm{SF}$. Also, note that $|\mathcal{Z}| \le q^{k\cdot\mathrm{rwt}(V)}$.

We can now bound, using the preceding Corollary 6.12:

$$\begin{aligned}
\sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U+V)=h}} |\widehat{g}(U)| &= \sum_{Z\in\mathcal{Z}} \sum_{\substack{U\in\mathrm{SF},Z(U)=Z,\\ \mathrm{wt}(U+V+Z)=h-\mathrm{wt}(Z)}} |\widehat{g}(U)| \\
&\le \sum_{Z\in\mathcal{Z}} \sum_{\substack{U+V+Z\in\mathrm{SF},Z(U)=Z,\\ \mathrm{wt}(U+V+Z)=h-\mathrm{wt}(Z)}} |\widehat{g}(U)| \\
&\le \sum_{Z\in\mathcal{Z}} \sum_{\substack{U+V+Z\in\mathrm{SF},\\ \mathrm{wt}(U+V+Z)=h-\mathrm{wt}(Z)}} |\widehat{g}(U)| \\
&\le \sum_{Z\in\mathcal{Z}} \left(\frac{\zeta\sqrt{bm}}{h-\mathrm{wt}(Z)}\right)^{(h-\mathrm{wt}(Z))/2} \\
&\le \sum_{Z\in\mathcal{Z}} \left(\frac{\zeta\sqrt{bm}}{h-\kappa}\right)^{(h-\kappa)/2} \le q^{k\cdot\mathrm{rwt}(V)}\cdot\left(\frac{\zeta\sqrt{bm}}{h-\kappa}\right)^{(h-\kappa)/2}.
\end{aligned}$$

□

7 The inductive argument: Proof of Lemma 4.9

Finally, in this section, we prove Lemma 4.9, restated as follows:

Lemma 4.9 (Inductive step). *For every $k, q, k' \in \mathbb{N}$ there exist $\alpha_0 > 0$ and $C_0 < \infty$ such that for every $C \ge C_0$ and $\delta \in (0, 1/2)$ there exist $\sigma \in (0,1)$ and $C'' > 0$ such that the following holds.*

For every injection $\phi : [k] \to [k']$ and one-wise uniform distribution $\mathcal{Y} \in \Delta_1(\mathbb{Z}_q^k)$, for every $n, b, b', s, s', m \in \mathbb{N}$ satisfying $m \le \alpha_0 n$, $0 < b, b', s < \sigma n$ and $b + b' + \log_q(1-1/\delta) \le s$ and every (C,s)-bounded distribution $\mathcal{D} \in \Delta(\mathbb{Z}_q^{n\times k'})$ satisfying $\|\mu_{\mathcal{D}}\|_\infty \le q^b$, we have that

$$\Pr_{\boldsymbol{M}\in\mathcal{M}_n^{m,k}}\left[\begin{array}{l}\forall \mathcal{B}\subset\mathbb{Z}_q^{m\times k} \text{ such that } |\mathcal{B}| \ge q^{m\cdot k-b'},\\ \qquad \mathrm{Post}(\mathcal{D};\boldsymbol{M},\mathcal{B}) \text{ is } (C'',s)\text{-bounded and } \|\mu_{\mathrm{Post}(\mathcal{D};\boldsymbol{M},\mathcal{B})}\|_\infty \le \frac{1}{1-\delta}\cdot q^{b+b'}\end{array}\right] \ge 1-5\delta$$

where $\boldsymbol{M} \in \mathcal{M}_n^{m,k}$ is a randomly sampled k-partite hypermatching.

7.1 An expression for the density function

We first give an analytic expression for the density function of the distribution $\mathrm{Post}(\mathcal{D}; M, \mathcal{B})$ in terms of the density functions of $\mathcal{D}$, $\mathcal{Y}^{\otimes m}$, and $\mathcal{B}$ (note that $\Pr_{\boldsymbol{Z}\sim\mathbb{Z}_q^{m\times k}}[\boldsymbol{Z}\in\mathcal{B}] = |\mathcal{B}|/q^{m\cdot k}$ below).

Lemma 7.1. *Let $\mathcal{D} \in \Delta(\mathbb{Z}_q^{n\times k'})$ be a distribution, $\phi : [k] \to [k']$ be an injection, $\mathcal{Y} \in \Delta(\mathbb{Z}_q^k)$ a distribution, $M \in \mathcal{M}_n^{m,k}$ a k-partite hypermatching, and $\mathcal{B} \subseteq \mathbb{Z}_q^{m\times k}$ a set. Then for every $X_0 \in \mathbb{Z}_q^{n\times k'}$:*

$$\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}(X_0) = \mu_{\mathcal{D}}(X_0)\cdot(\mu_{\mathcal{Y}^{\otimes m}} * \mu_{\mathcal{B}})(\Pi_M^\phi(X_0))\cdot\nu_{M,\mathcal{B}},$$

where

$$\nu_{M,\mathcal{B}} := \frac{\Pr_{\boldsymbol{Z}\sim\mathbb{Z}_q^{m\times k}}[\boldsymbol{Z}\in\mathcal{B}]}{\Pr_{\boldsymbol{X}\sim\mathcal{D},\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\Pi_M^\phi\boldsymbol{X}-\boldsymbol{Y}\in\mathcal{B}]}$$

is a normalization factor which does not depend on X_0.

Remark 7.2. Though we will not need this fact, it is easy to check that $(\mu_{\mathcal{Y}^{\otimes m}} * \mu_{\mathcal{B}})(\Pi_M^\phi(X_0)) = \mu_{\mathrm{Post}(\mathcal{U};M,\mathcal{B})}(X_0)$, where $\mathcal{U} := \mathsf{Unif}\mathbb{Z}_q^{n\times k'}$ is the *uniform* distribution on $\mathbb{Z}_q^{n\times k'}$. That is, $\mathrm{Post}(\mathcal{U}; M, \mathcal{B})$ is the posterior distribution of $\boldsymbol{X}$ when sampling $\boldsymbol{X}$ from the *uniform* prior ($\mathcal{U}$) conditioned on $\Pi_M^\phi \boldsymbol{X} - \boldsymbol{Y} \in \mathcal{B}$ for $\boldsymbol{Y} \sim \mathcal{Y}^{\otimes m}$.) Thus, Lemma 7.1 states that (up to normalization) the posterior on $\boldsymbol{X}$ with respect to the prior $\mathcal{D}$ is the product of the prior $\mathcal{D}$ and the posterior with respect to the uniform prior. ◊

Proof. Let $X_0 \in \mathbb{Z}_q^{n\times k'}$. By Bayes' rule,

$$\Pr_{\boldsymbol{X}\sim\mathcal{D},\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\boldsymbol{X} = X_0 \mid \Pi_M^\phi \boldsymbol{X} - \boldsymbol{Y} \in \mathcal{B}]$$
$$= \Pr_{\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\Pi_M^\phi X_0 - \boldsymbol{Y} \in \mathcal{B}] \cdot \Pr_{\boldsymbol{X}\sim\mathcal{D}}[\boldsymbol{X} = X_0] \cdot \left(\Pr_{\boldsymbol{X}\sim\mathcal{D},\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\Pi_M^\phi \boldsymbol{X} - \boldsymbol{Y} \in \mathcal{B}]\right)^{-1}.$$

Consequently,

$$\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}(X_0) = \Pr_{\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\Pi_M^\phi X_0 - \boldsymbol{Y} \in \mathcal{B}] \cdot \mu_{\mathcal{D}}(X_0) \cdot \nu_{M,\mathcal{B}} \cdot \frac{q^{m\cdot k}}{|\mathcal{B}|}.$$

Finally,

$$\frac{\Pr_{\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\Pi_M^\phi X_0 - \boldsymbol{Y} \in \mathcal{B}]}{|\mathcal{B}|} = \Pr_{\boldsymbol{B}\sim\mathcal{B},\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\boldsymbol{B} + \boldsymbol{Y} = \Pi_M^\phi X_0]$$

and therefore

$$\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}(X_0) = \Pr_{\boldsymbol{B}\sim\mathcal{B},\boldsymbol{Y}\sim\mathcal{Y}^{\otimes m}}[\boldsymbol{B} + \boldsymbol{Y} = \Pi_M^\phi X_0] \cdot \mu_{\mathcal{D}}(X_0) \cdot \nu_{M,\mathcal{B}} \cdot q^{m\cdot k}$$
$$= (\mu_{\mathcal{B}} * \mu_{\mathcal{Y}^{\otimes m}})(\Pi_M^\phi X_0) \cdot \mu_{\mathcal{D}}(X_0) \cdot \nu_{M,\mathcal{B}}.$$ □

Lemma 7.1 lets us derive a useful upper bound on the 1-norm of Fourier coefficients in a ball around the origin in a posterior distribution:

Corollary 7.3. *In the setup of Lemma 7.1, we have*

$$\sum_{\substack{W\in\mathbb{Z}_q^{n\times k'},\\ \mathrm{wt}(W)=h}} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le \nu_{M,\mathcal{B}} \cdot \sum_{V\in\mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(U)| \sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U+\Pi_M^\phi U)=h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))}} |\widehat{\mu_{\mathcal{B}}}(W)|.$$

Note that $\mathrm{wt}(V) - \mathrm{wt}(\Pi_M^\phi V)$ equals the size of the support of V *outside* of $\phi(\mathsf{supp}(M))$.

Proof. We apply Lemma 7.1 and Proposition 2.9 to get:

$$\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W) = \nu_{M,\mathcal{B}} \cdot \sum_{V\in\mathbb{Z}_q^{n\times k'}} \widehat{\mu_{\mathcal{D}}}(V) \cdot ((\widehat{\mu_{\mathcal{Y}^{\otimes m}} * \mu_{\mathcal{B}}}) \circ \Pi_M^\phi)(W - V)$$

and then Propositions 2.9 and 2.17 to get:

$$= \nu_{M,\mathcal{B}} \cdot \sum_{\substack{V\in\mathbb{Z}_q^{n\times k'},\\ \mathsf{supp}(W-V)\subseteq\phi(\mathsf{supp}(M))}} \widehat{\mu_{\mathcal{D}}}(V) \cdot \widehat{\mu_{\mathcal{Y}^{\otimes m}}}(\Pi_M^\phi(W - V)) \cdot \widehat{\mu_{\mathcal{B}}}(\Pi_M^\phi(W - V)).$$

Hence by Proposition 2.14, the triangle inequality, and upper bounding $|\widehat{\mu_{\mathcal{Y}^{\otimes m}}}(\Pi_M^\phi(W-V))|$ by 1,

$$|\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le \nu_{M,\mathcal{B}} \cdot \sum_{\substack{V\in\mathbb{Z}_q^{n\times k'},\\ \mathsf{supp}(W-V)\subseteq\phi(\mathsf{supp}(M)),\\ \Pi_M^\phi(W-V)\in\mathrm{SF}}} |\widehat{\mu_{\mathcal{D}}}(V)| \cdot |\widehat{\mu_{\mathcal{B}}}(\Pi_M^\phi(W-V))|.$$

Hence

$$\sum_{\substack{W\in\mathbb{Z}_q^{n\times k'},\\ \mathrm{wt}(W)=h}} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le \nu_{M,\mathcal{B}} \cdot \sum_{V\in\mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(V)| \sum_{\substack{W\in\mathbb{Z}_q^{n\times k'},\\ \mathrm{wt}(W)=h,\\ \mathsf{supp}(W-V)\subseteq\phi(\mathsf{supp}(M)),\\ \Pi_M^\phi(W-V)\in\mathrm{SF}}} |\widehat{\mu_{\mathcal{B}}}(\Pi_M^\phi(W-V))|$$

and substituting $U := \Pi_M^\phi(W-V)$,

$$\le \nu_{M,\mathcal{B}} \cdot \sum_{V\in\mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(V)| \sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U+\Pi_M^\phi V)=h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))}} |\widehat{\mu_{\mathcal{B}}}(U)|,$$

where we used that for $W\in\mathbb{Z}_q^{n\times k'}$ with $\mathsf{supp}(W-V)\subseteq\phi(\mathsf{supp}(M))$, we have $\mathrm{wt}(W)=h \iff \mathrm{wt}(\Pi_M^\phi W)=h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))$ (because $\mathsf{supp}(W-V)\subseteq\phi(\mathsf{supp}(M))$, implies W and V agree on all coordinates outside $\phi(\mathsf{supp}(M))$, and thus $\mathrm{wt}(W)-\mathrm{wt}(\Pi_M^\phi W)=\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V)$.) □

7.2 The inductive step "in expectation"

Corollary 7.3 gives a natural approach to proving Lemma 4.9: We formulate and prove an "in-expectation" version of the lemma, which then will imply the lemma by Markov's inequality.

Lemma 7.4 (Inductive step in expectation)**.** *For every $q,k,k'\in\mathbb{N}$, there exist $\alpha_0>0$ and $C_0<\infty$ such that for every $C>C_0$, there exist $\sigma\in(0,1)$ and $C'<\infty$ such that the following holds.*

For every injection $\phi:[k]\to[k']$, one-wise uniform distribution $\mathcal{Y}\in\Delta_1(k)$, and every $n,m,s,h\in\mathbb{N}$ satisfying $m\le\alpha_0 n$ and $1\le h\le s\le\sigma n$, and every $\mathcal{D}\in\Delta(\mathbb{Z}_q^{n\times k'})$ which is (C,s)-bounded:

$$\sum_{V\in\mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(V)| \cdot \mathop{\mathbb{E}}_{\boldsymbol{M}\in\mathcal{M}_n^{m,k}} \left[\max_{\substack{g:\mathbb{Z}_q^{m\times k}\to\mathbb{C},\\ \|g\|_1=1,\\ \|g\|_\infty\le q^s}} \sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U+\Pi_{\boldsymbol{M}}^\phi V)=h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_{\boldsymbol{M}}^\phi V))}} |\widehat{g}(U)| \right] \le \left(\frac{C'\sqrt{snk'}}{h}\right)^{h/2}.$$

Before proving this lemma, we show how it will imply Lemma 4.9. To do so, we will require the following simple analogue of [CGS+22, Lemma 6.3], which is a trivial upper-bound on the 1-norm mass in high-Hamming weight balls:

Lemma 7.5. *For every $q,N\in\mathbb{N}$ and $g:\mathbb{Z}_q^N\to\mathbb{C}$ with $\|g\|_1=1$ and $\log_q(\|g\|_\infty)\le h$, we have*

$$\sum_{\substack{U\in\mathbb{Z}_q^N,\\ \mathrm{wt}(U)=h}} |\widehat{g}(U)| \le \left(\frac{q^2eN}{h}\right)^{h/2}.$$

Proof. By Parseval's identity (Proposition 2.11) and the assumptions, we have:

$$\sum_{\substack{U\in\mathbb{Z}_q^N,\\ \mathrm{wt}(U)=h}} |\widehat{g}(U)|^2 \le \sum_{U\in\mathbb{Z}_q^N} |\widehat{g}(U)|^2 = \|g\|_2^2 \le \|g\|_1\cdot\|g\|_\infty \le q^h. \tag{7.6}$$

Also, there are (exactly) $(q-1)^h\cdot\binom{N}{h}$ frequencies $U\in\mathbb{Z}_q^N$ with $\mathrm{wt}(U)=h$. By Cauchy-Schwarz,

$$\sum_{\substack{U\in\mathbb{Z}_q^N,\\ \mathrm{wt}(U)=h}} |\widehat{g}(U)| \le \sqrt{(q-1)^h\binom{N}{h}\sum_{\substack{U\in\mathbb{Z}_q^N,\\ \mathrm{wt}(U)=h}} |\widehat{g}(U)|^2} \le \sqrt{(q-1)^h\left(\frac{eN}{h}\right)^h\cdot q^h} \le \left(\frac{eq^2N}{h}\right)^{h/2},$$

where the second inequality uses the standard binomial estimate and Equation (7.6). □

We can now give a proof of Lemma 4.9.

Proof of Lemma 4.9 (modulo Lemma 7.4). Let α_0 and C_0 be as defined in Lemma 7.4. Given parameters C and δ, we then invoke Lemma 7.4 and let σ and C' denote the resulting values.

Let $\mathcal{E}(M)$ denote the event that the matching $M\in\mathcal{M}_n^{m,k}$ is such that $\|\mu_{\mathrm{Input}(\mathcal{D},M)}-1\|_\infty>\delta$. For $h\in[\sigma n]$, let $\mathcal{F}_h(M)$ denote the event that

$$\frac{1}{1-\delta}\sum_{V\in\mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(V)|\cdot\left(\max_{\substack{g:\mathbb{Z}_q^{m\times k}\to\mathbb{C},\\ \|g\|_1=1,\\ \|g\|_\infty\le q^s}}\ \sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U+\Pi_M^\phi V)=h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))}} |\widehat{g}(U)|\right) > \frac{1}{\delta^h}\left(\frac{C'\sqrt{snk'}}{\max\{h,s\}}\right)^{h/2}.$$

Sufficiency of the events. We first claim that whenever $\overline{\mathcal{E}(M)\cup\bigcup_{h=1}^{\sigma n}\mathcal{F}_h(M)}$ holds, then the desideratum holds, i.e., for every $\mathcal{B}\subseteq\mathbb{Z}_q^{m\times k}$ such that $|\mathcal{B}|\ge q^{n-b'}$, we have that $\mathrm{Post}(\mathcal{D};M,\mathcal{B})$ is (C',s)-bounded and $\|\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}-1\|_\infty<q^s$.

We first observe that $\nu_{M,\mathcal{B}}^{-1}=\frac{\Pr_{\mathbf{Z}\sim\mathrm{Input}(\mathcal{D},M)}[\mathbf{Z}\in\mathcal{B}]}{\Pr_{\mathbf{Z}\in\mathbb{Z}_q^{m\times k}}[\mathbf{Z}\in\mathcal{B}]}$, which is at least $1-\delta$ by Observation 4.8 and the definition of $\overline{\mathcal{E}(M)}$. Hence, $\nu_{M,\mathcal{B}}\le\frac{1}{1-\delta}$.

Now, Lemma 7.1 and Proposition 2.5 give that $\|\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}\|_\infty\le\|\mu_{\mathcal{D}}\|_\infty\cdot\|\mu_{\mathcal{B}}\|_\infty\cdot\nu_{M,\mathcal{B}}$. We have $\|\mu_{\mathcal{D}}\|_\infty\le q^b$ and $\|\mu_{\mathcal{B}}\|_\infty\le q^{b'}$ (by assumption), and $\nu_{M,\mathcal{B}}\le\frac{1}{1-\delta}$, and therefore the product is at most $q^{b+b'+\log_q(1/1-\delta')}\le q^s$, by our assumption.

We now choose a constant $C''<\infty$ such that for every $h\in[nk']$, we can guarantee the inequality $\sum_{W\in\mathbb{Z}_q^{n\times k'},\mathrm{wt}(W)=h}|\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)|\le \mathrm{U}_{C'',s,nk'}(h)$. We do so using a few cases on h:

- If $h\le s$, by Corollary 7.3 and the inequality on $\nu_{M,B}$, since $\overline{\mathcal{F}_h(M)}$ holds, it is the case that

$$\sum_{W\in\mathbb{Z}_q^{n\times k'},\mathrm{wt}(W)=h}|\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)|\le\frac{1}{\delta^h}\left(\frac{C'\sqrt{snk'}}{h}\right)^{h/2}=\left(\frac{(C'/\sqrt{\delta})\sqrt{snk'}}{h}\right)^{h/2}.$$

  Since $h\le s$, the RHS is at most $\mathrm{U}_{C'',s,nk'}(h)$ as long as $C''\ge C'/\sqrt{\delta}$.

- If $s \le h \le \sigma n$, the same argument gives $\sum_{W \in \mathbb{Z}_q^{n\times k'}, \mathrm{wt}(W)=h} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le \left(\frac{(C'/\sqrt{\delta})\sqrt{nk'}}{\sqrt{h}}\right)^{h/2}$. By Lemma 7.5 , we also have that $\sum_{W \in \mathbb{Z}_q^{n\times k'}, \mathrm{wt}(W)=h} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le (eq^2nk'/h)^{h/2}$. (This uses $\|\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}\|_\infty \le q^s$ as already proved above.) Hence, for $C'' \ge C'/\sqrt{\delta}$ we have,

$$\sum_{W \in \mathbb{Z}_q^{n\times k'}, \mathrm{wt}(W)=h} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le \mathrm{U}_{C'',s,nk'}(h).$$

- Finally, if $\sigma n \le h \le nk'$, we still have $\sum_{W \in \mathbb{Z}_q^{n\times k'}, \mathrm{wt}(W)=h} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le (eq^2nk'/h)^{h/2}$. To guarantee that $\sum_{W \in \mathbb{Z}_q^{n\times k'}, \mathrm{wt}(W)=h} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le \left(\frac{C''\sqrt{nk'}}{\sqrt{h}}\right)^{h/2}$, it therefore suffices to show that $\frac{C''\sqrt{nk'}}{\sqrt{h}}^{h/2} \ge (eq^2nk'/h)^{h/2}$. This is achieved by setting $C'' \ge q^2e/\sqrt{\sigma}$, so that we have $h \ge \sigma nk' \implies C''\sqrt{h} \ge q^2e\sqrt{nk'} \implies C''\sqrt{nk'}/\sqrt{h} \ge q^2enk'/h$.

Hence, if we set $C'' := \max\{C'/\sqrt{\delta}, q^2e/\sqrt{\sigma}\}$, we get that in all cases we have

$$\sum_{W \in \mathbb{Z}_q^{n\times k'}, \mathrm{wt}(W)=h} |\widehat{\mu_{\mathrm{Post}(\mathcal{D};M,\mathcal{B})}}(W)| \le \mathrm{U}_{C'',s,nk'}(h).$$

Bounding the probability over events. Since $\mathcal{D}$ is (C,s)-bounded, the boundedness implies uniformity lemma (Lemma 4.7) posits directly that $\Pr[\mathcal{E}(\boldsymbol{M})] \le \delta$. Further, applying Markov's inequality to Lemma 7.4, we conclude that $\Pr[\mathcal{F}_h(\boldsymbol{M})] \le \frac{\delta^h}{1-\delta}$. Thus a union bound gives $\Pr[\mathcal{F}_h(\boldsymbol{M})] \le \frac{1}{1-\delta}\sum_{h=1}^{\infty}\delta^h \le 4\delta$. □

7.3 Combinatorial bound

Consider the set $\mathcal{V} := [n] \times [k]$ as a set of $k \cdot n$ vertices partitioned into k parts $\mathcal{V}_\ell := [n] \times \{j\}$, each of size n. If $U \subseteq \mathcal{V}$ is a set of vertices, and $M \in \mathcal{M}_n^{m,k}$, we define the following:

- Sets of edges: $K(M,U) := \{j \in [m] : |e_j \cap U| = 1\}$ and $D := \{j \in [m] : |e_j \cap U| \ge 2\}$, where $e_j \subseteq \mathcal{V}$ is the j-th edge of M.
- Size parameters: $\kappa(M,U) := |K|$, $d(M,U) := |D(M,U)|$, and $\eta(M,U) := |U \cap \bigcup_{j\in D} e_j|$.

We suppress the dependence on M and U when clear from context. Note that $\eta/k \le d \le \eta/2$ and $d \le m = \alpha n \le n/k$, so $\eta \le kd \le n$. In particular $d \le n/k < n$.

Proposition 7.7. *For every $x, y > 0$, it holds that $(x/y)^y \le e^x$.*

Proposition 7.8. *For every $x \in \mathbb{R}$, it holds that $x \le (e^{1/e})^x$.*

We now define the following function:

$$p_{C,\alpha}(n,u,\kappa,\eta) := \alpha^\kappa \cdot C^u \cdot \left(\frac{u}{n}\right)^{\eta/2}. \tag{7.9}$$

We note that this bound looks drastically simpler than the comparable bound in [CGS+22, Lemma 6.24].

Lemma 7.10. *For every $k \in \mathbb{N}$ there exists a constant $C < \infty$ such that the following holds. Let $U \subseteq \mathcal{V} = [n] \times [k]$ be a fixed set of $|U| = u$ marked vertices. For a random matching $\boldsymbol{M} \in \mathcal{M}_n^{m,k}$ with $m = \alpha n$ hyperedges, we have:*

$$\Pr[\kappa(U, \boldsymbol{M}) = \kappa,\ \eta(U, \boldsymbol{M}) = \eta] \leq p_{C,\alpha}(n, u, \kappa, \eta)$$

assuming $k + \eta \leq u$, and 0 otherwise.

Proof. We follow and adapt (simplify) the proof of [CGS+22, Lemma 6.24].

Step 1: Symmetry reduction. For $j \in [k]$, let $u_\ell := |U \cap \mathcal{V}_\ell|$ denote the support of U on the j-th part. To calculate $\Pr[\kappa(U, \boldsymbol{M}) = \kappa,\ \eta(U, \boldsymbol{M}) = \eta]$, by column-wise permutation symmetry, it is equivalent to instead fix one matching, such as

$$M := \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 2 & 2 & \cdots & 2 \\ \vdots & \vdots & \ddots & \vdots \\ m & m & \cdots & m \end{bmatrix},$$

sample $\boldsymbol{U}$ uniformly at random from the set $\mathcal{U} := \{U \subseteq \mathcal{V} : \forall j \in [k],\ |U \cap \mathcal{V}_j| = u_\ell\}$, and calculate $\Pr[\kappa(\boldsymbol{U}, M) = \kappa,\ \eta(\boldsymbol{U}, M) = \eta]$.

Step 2: Partitioning by d. For fixed $d \in [\eta/k, \eta/2]$, let $\mathcal{N}(d) := \{U \in \mathcal{U} : \kappa(U, M) = \kappa,\ d(U, M) = d,\ \eta(U, M) = \eta\}$. Then for $\boldsymbol{U} \sim \mathcal{U}$,

$$\Pr[\kappa(\boldsymbol{U}, M) = \kappa,\ \eta(\boldsymbol{U}, M) = \eta] \leq \sum_{d=\eta/k}^{\eta/2} \frac{|\mathcal{N}(d)|}{\prod_{j=1}^{k} \binom{n}{u_\ell}}.$$

Step 3: Counting $|\mathcal{N}(d)|$. We now upper-bound $|\mathcal{N}(d)|$.

To specify an element in $U \in \mathcal{N}(d)$, it suffices to specify the edge-sets $K = K(U, M) \subseteq [m]$ and $D = D(U, M) \subseteq [m]$, and then the vertex-sets $U \cap \bigcup_{j \in K} e_j$, $U \cap \bigcup_{j \in D} e_j$, and $U \cap \bigcup_{j \in [m] \setminus (K \cup D)} e_j$.

Pessimistically, there are $\binom{m}{\kappa}$ choices for K and $\binom{m}{d}$ choices for D. Once these have been specified, there are now k^κ choices for $U \cap \bigcup_{j \in K} e_j$ (once K has been fixed, $U \cap \bigcup_{j \in K} e_j$ contains exactly one vertex per edge in K and $|K| = \kappa$) and (pessimistically) $\binom{kd}{\eta}$ choices for $U \cap \bigcup_{j \in D} e_j$ (since $|U \cap \bigcup_{j \in D} e_j| = \eta$ and $|\bigcup_{j \in D} e_j| = kd$). Once these have all been fixed, it remains to pick $U \cap \bigcup_{j \in [m] \setminus (K \cup D)} e_j$, and we do so by parts; i.e., we specify the intersections $U \cap \mathcal{V}_\ell \cap \bigcup_{j \in [m] \setminus (K \cup D)} e_j$ with each part. Since we must have $|U \cap V_\ell| = u_\ell$, letting $\mu_\ell := |U \cap \mathcal{V}_\ell \cap \bigcup_{j \in (K \cup D)} e_j|$ denote the number of already-selected vertices in column ℓ, we conclude that $|U \cap \mathcal{V}_\ell \cap \bigcup_{j \in [m] \setminus (K \cup D)} e_j| = u_\ell - \mu_\ell$. Since $U \cap \mathcal{V}_\ell \cap \bigcup_{j \in [m] \setminus (K \cup D)} e_j \subseteq \mathcal{V}_\ell \cap \bigcup_{j \in [m] \setminus (K \cup D)} e_j$ and $|\mathcal{V}_\ell \cap \bigcup_{j \in [m] \setminus (K \cup D)} e_j| = n - (\kappa + d)$, we deduce:

$$|\mathcal{N}(d)| \leq \binom{m}{\kappa} \cdot \binom{m}{d} \cdot k^\kappa \cdot \binom{kd}{\eta} \cdot \prod_{\ell=1}^{k} \binom{n - (\kappa + d)}{u_\ell - \mu_\ell}.$$

Step 4: Applying estimates and collecting like terms. Note that for every ℓ, we have $\mu_\ell \leq \kappa + d$, hence $\binom{n-(\kappa+d)}{u_\ell - \mu_\ell} \leq \binom{n-\mu_\ell}{u_\ell - \mu_\ell}$. We can therefore apply $\frac{\binom{a-c}{b-c}}{\binom{a}{b}} \leq \left(\frac{b}{a}\right)^c$, as well as $\binom{a}{b} \leq (ea/b)^b$ to the remaining binomials, to get:

$$\frac{|\mathcal{N}(d)|}{\prod_{\ell=1}^{k} \binom{n}{u_\ell}} \leq \left(\frac{e\alpha n}{\kappa}\right)^\kappa \cdot k^\kappa \cdot \left(\frac{e\alpha n}{d}\right)^d \cdot \left(\frac{ekd}{\eta}\right)^\eta \cdot \left(\frac{u}{n}\right)^{\kappa+\eta}.$$

Regrouping terms and using that $\alpha^d \le 1$ since $\alpha \le 1$ and $d \ge 0$ gives:

$$\frac{|\mathcal{N}(d)|}{\prod_{\ell=1}^{k}\binom{n}{u_\ell}} \le (ek)^{\kappa+\eta}\cdot\alpha^\kappa\cdot\left(\frac{u}{\kappa}\right)^\kappa\cdot\left(\frac{u}{n\eta}\right)^\eta\cdot\left(\frac{en}{d}\right)^d\cdot d^\eta.$$

Writing $(u/(n\eta))^\eta = (u/(\eta/2))^{\eta/2}\cdot(u/(2\eta))^{\eta/2}\cdot(1/n)^\eta$, we can now apply Proposition 7.7 and use $\kappa+\eta\le u$, giving:

$$\frac{|\mathcal{N}(d)|}{\prod_{\ell=1}^{k}\binom{n}{u_\ell}} \le (e^{3u}\cdot k^u)\cdot\alpha^\kappa\cdot\left(\frac{u}{2\eta}\right)^{\eta/2}\cdot\left(\frac{1}{n}\right)^\eta\cdot\left(\frac{en}{d}\right)^d\cdot d^\eta.$$

Step 5: Summing over d. Consider the function $g(t) := \ln((en)^t\cdot t^{\eta-t}) = t\ln(en)+(\eta-t)\ln t$. Analytically, we have $\frac{d}{dt}g(t) = \ln(en)+\eta/t-(\ln t+1) = \ln(n)-\ln(t)+\eta/t \ge \ln(n)-\ln(t)$. In particular, g is increasing on the interval $(0,n]$. This gives, using that $\eta\le u$:

$$\sum_{d=\eta/k}^{\eta/2}\left(\frac{en}{d}\right)^d\cdot d^\eta \le \eta/2\cdot\left(\frac{en}{\eta/2}\right)^{\eta/2}\cdot(\eta/2)^\eta \le (u\cdot e^{\eta/2}\cdot 2^{-\eta/2})\cdot(n\cdot\eta)^{\eta/2}.$$

Combining and regrouping gives:

$$\sum_{d=\eta/k}^{\eta/2}\frac{|\mathcal{N}(d)|}{\prod_{\ell=1}^{k}\binom{n}{u_\ell}} \le (e^{3u}\cdot k^u\cdot u\cdot e^{\eta/2}\cdot 2^{-\eta/2})\cdot\alpha^\kappa\cdot\left(\frac{u}{n}\right)^{\eta/2}.$$

We can now finish by setting $C := e^{3+1/e+1/2}\cdot k/\sqrt{2}$, where we used that $u\le(e^{1/e})^u$ and $\eta\le u$. □

7.4 The inductive step in expectation: Proving Lemma 7.4

We finally turn to proving Lemma 7.4. To do so, we require the following analytic lemma, which follows from the proof of [CGS+22, Lemma 6.24]; we give a dramatically simplified proof in §B below.

Lemma 7.11. *For every $q,k\in\mathbb{N}$, there exists $\alpha_0\in(0,1/k)$ such that the following holds. For every $C_1,C_2,C_3,C_4<\infty$, there exists $\varepsilon_0>0$ and $C_5<\infty$ such that for every $s,n,h,u,\kappa,\eta\in\mathbb{N}$ with $\kappa\le h\le s\le\varepsilon_0 n$, and $\kappa+\eta\le u\le\min\{n,h+\eta\}$, we have*

$$\mathrm{U}_{C_1,s,n}(u)\cdot C_2^{\kappa+\eta}\cdot p_{C_3,\alpha_0}(n,u,\kappa,\eta)\cdot\left(\frac{C_4\sqrt{sn}}{h+\eta-u}\right)^{(h-\eta+u)/2}\le\left(\frac{C_5\sqrt{sn}}{h}\right)^{h/2}.$$

Finally, we give:

Proof of Lemma 7.4. By Lemma 6.13 with $\theta := 1.1$ (so that $h<\theta s$), there exists $C_4<\infty$ such that for every injection $\phi:[k]\to[k']$, fixed $M\in\mathcal{M}_n^{m,k}$, and $V\in\mathbb{Z}_q^{n\times k'}$ (equivalently, $\Pi_M^\phi V\in\mathbb{Z}_q^{m\times k}$), and

$g : \mathbb{Z}_q^{m\times k} \to \mathbb{C}$ satisfying $\|g\|_1 = 1$ and $\|g\|_\infty \le q^s$:

$$\begin{aligned}
&\sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U+\Pi_M^\phi V)=h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))}} |\widehat{g}(U)| \\
&\le q^{k\cdot\mathrm{wt}(V)} \cdot \left(\frac{C_4\sqrt{sm}}{h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))-\kappa(M,\Pi_M^\phi V))}\right)^{(h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))-\kappa(M,\Pi_M^\phi V))/2} \\
&= q^{k\cdot\mathrm{wt}(V)} \cdot \left(\frac{C_4\sqrt{sm}}{h-(\mathrm{wt}(V)-\eta(M,\Pi_M^\phi V))}\right)^{(h-(\mathrm{wt}(V)-\eta(M,\Pi_M^\phi V)))/2},
\end{aligned}$$

where we used the pessimistic upper bounds $\mathrm{rwt}(\Pi_M^\phi V) \le \mathrm{wt}(V)$ and the equality $\mathrm{wt}(\Pi_M^\phi V) = \kappa(M,\Pi_M^\phi V) + \eta(M,\Pi_M^\phi V)$. We now take a maximum over g, sample $\boldsymbol{M}$ randomly, and sum over V. This gives:

$$\begin{aligned}
&\sum_{V\in\mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(V)| \cdot \mathop{\mathbb{E}}_{\boldsymbol{M}}\left[\max_{\substack{g:\mathbb{Z}_q^{m\times k}\to\mathbb{C},\\ \|g\|_1=1,\\ \|g\|_\infty\le q^s}} \sum_{\substack{U\in\mathrm{SF},\\ \mathrm{wt}(U+\Pi_M^\phi V)=h-(\mathrm{wt}(V)-\mathrm{wt}(\Pi_M^\phi V))}} |\widehat{g}(U)|\right] \\
&\le \sum_{V\in\mathbb{Z}_q^{n\times k'}} |\widehat{\mu_{\mathcal{D}}}(V)| \cdot q^{k\cdot\mathrm{wt}(V)} \cdot \mathop{\mathbb{E}}_{\boldsymbol{M}}\left[\left(\frac{C_4\sqrt{sm}}{h-(\mathrm{wt}(V)-\eta(\boldsymbol{M},\Pi_{\boldsymbol{M}}^\phi V))}\right)^{(h-(\mathrm{wt}(V)-\eta(M,\Pi_{\boldsymbol{M}}^\phi V)))/2}\right]
\end{aligned}$$

Partitioning the sum over V by $\mathrm{wt}(V)$, and the expectation over $\boldsymbol{M}$ by $\kappa(M,\Pi_M^\phi V)$ and $\eta(M,\Pi_M^\phi V)$:

$$\le \sum_{u=0}^{nk'} \sum_{\substack{V\in\mathbb{Z}_q^{n\times k'},\\ \mathrm{wt}(V)=u}} |\widehat{\mu_{\mathcal{D}}}(V)| \cdot q^{ku} \cdot \sum_{\substack{\kappa,\eta\in\mathbb{N},\\ \kappa+\eta\le u\le h+\eta}} \left(\frac{C_4\sqrt{sm}}{h-(u-\eta)}\right)^{(h-(u-\eta))/2} \cdot \Pr_{\boldsymbol{M}}[\kappa(\boldsymbol{M},\Pi_{\boldsymbol{M}}^\phi V)=\kappa,\ \eta(\boldsymbol{M},\Pi_{\boldsymbol{M}}^\phi V)=\eta]$$

By the (C_1,s)-boundedness of $\mathcal{D}$ and Lemma 7.10 (letting C_3 denote the constant from the latter), and using $m \le nk$:

$$\le \sum_{u=0}^{nk'} \sum_{\substack{\kappa,\eta\in\mathbb{N},\\ \kappa+\eta\le u\le h+\eta}} \mathrm{U}_{C_1,s,n}(u) \cdot q^{ku} \cdot \left(\frac{C_4\sqrt{k}\sqrt{sn}}{h-(u-\eta)}\right)^{(h-(u-\eta))/2} \cdot p_{C_3,\alpha}(n,u,\kappa,\eta).$$

Defining $C_2 := 2q^k$ and applying the foregoing Lemma 7.11:

$$\le \sum_{u=0}^{nk'} \sum_{\substack{\kappa,\eta\in\mathbb{N},\\ \kappa+\eta\le u\le h+\eta}} 2^{-u}\left(\frac{C_5\sqrt{sn}}{h}\right)^{h/2} \le \sum_{u=0}^{nk'} u^2\cdot 2^{-u}\left(\frac{C_5\sqrt{sn}}{h}\right)^{h/2} \le \left(\frac{169C_5\sqrt{sn}}{h}\right)^{h/2}$$

where for the final inequality we used that $\sum_{u=0}^{nk'}(u+1)^2\cdot 2^{-u} \le \sum_{u=0}^{\infty}(u+1)^2\cdot 2^{-u} \le 1 + \int_0^\infty (u+1)^2\cdot 2^{-u}\,du < 13$ and $h \ge 1$. □

AI Disclosure

We used Claude to assist with the proofs in §6, though all content in the paper is written by the authors. The authors verified the correctness and originality of all content including references.

References

[ABF26] Amir Azarmehr, Soheil Behnezhad, and Shane Ferrante. *Single-Pass Streaming CSPs via Two-Tier Sampling*. Apr. 2, 2026. arXiv: 2604.01575 [cs]. Pre-published.

[ABFS26] Amir Azarmehr, Soheil Behnezhad, Shane Ferrante, and Mohammad Saneian. "Half-Approximating Maximum Dicut in the Streaming Setting". In: *Proceedings of the 58th ACM Symposium on Theory of Computing*. STOC 2026 (June 22–26, 2026). Association for Computing Machinery, 2026.

[AKSY20] Sepehr Assadi, Gillat Kol, Raghuvansh R. Saxena, and Huacheng Yu. "Multi-Pass Graph Streaming Lower Bounds for Cycle Counting, MAX-CUT, Matching Size, and Other Problems". In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science*. FOCS 2020 (Nov. 16–19, 2020). IEEE Computer Society, Nov. 2020, pp. 354–364. DOI: 10.1109/FOCS46700.2020.00041.

[AN21] Sepehr Assadi and Vishvajeet N. "Graph Streaming Lower Bounds for Parameter Estimation and Property Testing via a Streaming XOR Lemma". In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2021 (June 21–25, 2021). Association for Computing Machinery, June 15, 2021, pp. 612–625. DOI: 10.1145/3406325.3451110.

[Ass23] Sepehr Assadi. "Recent Advances in Multi-Pass Graph Streaming Lower Bounds". In: *ACM SIGACT News* 54.3 (Sept. 7, 2023), pp. 48–75. ISSN: 0163-5700. DOI: 10.1145/3623800.3623808.

[BCK15] Boaz Barak, Siu On Chan, and Pravesh K. Kothari. "Sum of Squares Lower Bounds from Pairwise Independence". In: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*. STOC 2015 (June 15–17, 2015). Portland, OR, USA: ACM, June 14, 2015, pp. 97–106. DOI: 10.1145/2746539.2746625.

[BDV18] Aditya Bhaskara, Samira Daruki, and Suresh Venkatasubramanian. "Sublinear Algorithms for MAXCUT and Correlation Clustering". In: *45th International Colloquium on Automata, Languages, and Programming*. ICALP 2018 (July 9–13, 2018). Vol. 107. LIPIcs. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2018, 16:1–16:14. DOI: 10.4230/LIPICS.ICALP.2018.16.

[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 (June 20–24, 2022). Association for Computing Machinery, 2022. DOI: 10.1145/3519935.3519983.

[CGSV24] Chi-Ning Chou, Alexander Golovnev, Madhu Sudan, and Santhoshini Velusamy. "Sketching Approximability of All Finite CSPs". In: *Journal of the ACM* 71.2 (Apr. 12, 2024). Conference version in FOCS 2021, 15:1–15:74. DOI: 10.1145/3649435.

[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 (Nov. 16–19, 2020). IEEE Computer Society, Nov. 2020, pp. 330–341. DOI: 10.1109/FOCS46700.2020.00039.

[CKP+23] Lijie Chen, Gillat Kol, Dmitry Paramonov, Raghuvansh Saxena, Zhao Song, and Huacheng Yu. "Towards Multi-Pass Streaming Lower Bounds for Optimal Approximation of Max-Cut". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA 2023 (Jan. 22–25, 2023). 2023. DOI: 10.1137/1.9781611977554.ch35.

[CLRS16] Siu On Chan, James R. Lee, Prasad Raghavendra, and David Steurer. "Approximate Constraint Satisfaction Requires Large LP Relaxations". In: *Journal of the ACM* 63.4 (Nov. 8, 2016). Conference version in FOCS 2013, pp. 1–22. DOI: 10.1145/2811255.

[FMW25] Yumou Fei, Dor Minzer, and Shuo Wang. "Multi-Pass Streaming Lower Bounds for Approximating Max-Cut". In: *Proceedings of the 66th IEEE Symposium on Foundations of Computer Science*. FOCS 2025 (Dec. 14–17, 2025). IEEE Computer Society, Dec. 2025.

[FMW26a] Yumou Fei, Dor Minzer, and Shuo Wang. "A Dichotomy Theorem for Multi-Pass Streaming CSPs". In: *Proceedings of the 58th ACM Symposium on Theory of Computing*. STOC 2026 (June 22–26, 2026). To appear. Association for Computing Machinery, 2026.

[FMW26b] Yumou Fei, Dor Minzer, and Shuo Wang. *Near-Optimal Space Lower Bounds for Streaming CSPs*. Apr. 1, 2026. arXiv: 2604.01400 [cs]. Pre-published.

[Gri01] D. Grigoriev. "Complexity of Positivstellensatz Proofs for the Knapsack". In: *computational complexity* 10.2 (Dec. 1, 2001), pp. 139–154. DOI: 10.1007/s00037-001-8192-0.

[GT17] Mrinalkanti Ghosh and Madhur Tulsiani. "From Weak to Strong LP Gaps for All CSPs". In: CCC 2017 (July 6–9, 2017). Ed. by Ryan O'Donnell. Vol. 79. LIPIcs. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, 11:1–11:27. DOI: 10.4230/LIPICS.CCC.2017.11.

[GT19] Venkatesan Guruswami and Runzhou Tao. "Streaming Hardness of Unique Games". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. APPROX 2019 (Sept. 20–22, 2019). Ed. by Dimitris Achlioptas and László A. Végh. Vol. 145. LIPIcs. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, Sept. 2019, 5:1–5:12. DOI: 10.4230/LIPIcs.APPROX-RANDOM.2019.5.

[GVV17] Venkatesan Guruswami, Ameya Velingker, and Santhoshini Velusamy. "Streaming Complexity of Approximating Max 2CSP and Max Acyclic Subgraph". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. APPROX 2017 (Aug. 16–18, 2017). Ed. by Klaus Jansen, José D. P. Rolim, David Williamson, and Santosh S. Vempala. Vol. 81. LIPIcs. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, Aug. 2017, 8:1–8:19. DOI: 10.4230/LIPIcs.APPROX-RANDOM.2017.8.

[KK15] Dmitry Kogan and Robert Krauthgamer. "Sketching Cuts in Graphs and Hypergraphs". In: *Proceedings of the 6th Annual Conference on Innovations in Theoretical Computer Science*. ITCS 2015 (Jan. 11–13, 2015). Association for Computing Machinery, 2015, pp. 367–376. DOI: 10.1145/2688073.2688093.

[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 (June 23–26, 2019). Association for Computing Machinery, June 2019, pp. 277–288. DOI: 10.1145/3313276.3316364.

[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 (Jan. 4–6, 2015). Society for Industrial and Applied Mathematics, Jan. 2015, pp. 1263–1282. DOI: 10.1137/1.9781611973730.84.

[KKSV17] Michael Kapralov, Sanjeev Khanna, Madhu Sudan, and Ameya Velingker. "$(1+\omega(1))$-Approximation to MAX-CUT Requires Linear Space". In: *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA 2017 (Jan. 16–19, 2017). Society for Industrial and Applied Mathematics, Jan. 2017, pp. 1703–1722. DOI: 10.5555/3039686.3039798.

[KMOW17] Pravesh K. Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer. "Sum of Squares Lower Bounds for Refuting Any CSP". In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2017. To appear. Association for Computing Machinery, June 19, 2017, pp. 132–145. DOI: 10.1145/3055399.3055485.

[KMR22] Pravesh K. Kothari, Raghu Meka, and Prasad Raghavendra. "Approximating Rectangles by Juntas and Weakly Exponential Lower Bounds for LP Relaxations of CSPs". In: *SIAM Journal on Computing* 51.2 (Apr. 2022). Conference version in STOC 2017, STOC17-305-STOC17–332. DOI: 10.1137/17M1152966.

[KPSY23] Gillat Kol, Dmitry Paramonov, Raghuvansh R. Saxena, and Huacheng Yu. "Characterizing the Multi-Pass Streaming Complexity for Solving Boolean CSPs Exactly". In: *14th Innovations in Theoretical Computer Science Conference*. ITCS 2023 (Jan. 10–13, 2023). Ed. by Yael Tauman Kalai. Vol. 251. LIPIcs. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 80:1–80:15. DOI: 10.4230/LIPIcs.ITCS.2023.80.

[Lee15] Euiwoong Lee. "Hardness of Graph Pricing Through Generalized Max-Dicut". In: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*. STOC 2015 (June 15–17, 2015). Portland, OR, USA: ACM, June 14, 2015, pp. 391–399. DOI: 10.1145/2746539.2746549.

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions*. 1st edition. New York, NY: Cambridge University Press, June 1, 2014. 444 pp. ISBN: 978-1-107-03832-5.

[Pot19] Aaron Potechin. "On the Approximation Resistance of Balanced Linear Threshold Functions". In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019 (June 23–26, 2019). Association for Computing Machinery, June 23, 2019, pp. 430–441. DOI: 10.1145/3313276.3316374.

[Rag08] Prasad Raghavendra. "Optimal Algorithms and Inapproximability Results for Every CSP?" In: *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*. STOC 2008 (May 17–20, 2008). 2008, pp. 245–254. DOI: 10.1145/1374376.1374414.

[Sch08] Grant Schoenebeck. "Linear Level Lasserre Lower Bounds for Certain K-CSPs". In: *49th Annual IEEE Symposium on Foundations of Computer Science*. FOCS 2008 (May 25–28, 2008). Philadelphia, PA, USA: IEEE Computer Society, Oct. 2008, pp. 593–602. DOI: 10.1109/FOCS.2008.74.

[Sin25] Noah G. Singer. "Nine Lower Bound Conjectures on Streaming Approximation Algorithms for CSPs". In: *SIGACT News* 56.4 (Dec. 2025), pp. 25–36. DOI: 10.1145/3785512.3785516.

[SSSV23a] Raghuvansh R. Saxena, Noah Singer, Madhu Sudan, and Santhoshini Velusamy. "Improved Streaming Algorithms for Maximum Directed Cut via Smoothed Snapshots". In: *63rd Annual Symposium on Foundations of Computer Science*. FOCS 2023 (Nov. 6–9, 2023). IEEE Computer Society, 2023, pp. 855–870. DOI: 10.1109/FOCS57990.2023.00055.

[SSSV23b] Raghuvansh R. Saxena, Noah G. Singer, Madhu Sudan, and Santhoshini Velusamy. "Streaming Complexity of CSPs with Randomly Ordered Constraints". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA 2023 (Jan. 22–25, 2023). Society for Industrial and Applied Mathematics, Jan. 2023, pp. 4083–4103. DOI: 10.1137/1.9781611977554.ch156.

[SSSV25] Raghuvansh Saxena, Noah G. Singer, Madhu Sudan, and Santhoshini Velusamy. "Streaming Algorithms via Local Algorithms for Maximum Directed Cut". In: *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA 2025 (Jan. 12–15, 2025). Society for Industrial and Applied Mathematics, 2025, pp. 3392–3408. DOI: 10.1137/1.9781611978322.111.

[STV25] Noah G. Singer, Madhur Tulsiani, and Santhoshini Velusamy. *Sketching Approximations and LP Approximations for Finite CSPs Are Related*. Sept. 23, 2025. arXiv: 2509.17926 [cs]. Pre-published.

[Sud22] Madhu Sudan. "Streaming and Sketching Complexity of CSPs: A Survey (Invited Talk)". In: *49th International Colloquium on Automata, Languages, and Programming*. ICALP 2022 (July 4–8, 2022). Ed. by Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff. Vol. 229. LIPIcs. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2022, 5:1–5:20. DOI: 10.4230/LIPIcs.ICALP.2022.5.

[Tul09] Madhur Tulsiani. "CSP Gaps and Reductions in the Lasserre Hierarchy". In: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*. STOC 2009 (May 31–June 2, 2009). Bethesda, MD, USA: ACM, May 31, 2009, pp. 303–312. ISBN: 978-1-60558-506-2. DOI: 10.1145/1536414.1536457.

[Vel25] Santhoshini Velusamy. *Near-Optimal Streaming Approximation for Max-DICUT in Sublinear Space Using Two Passes*. Comment: 27 pages. Dec. 22, 2025. arXiv: 2512.19521 [cs]. Pre-published.

[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 (Oct. 31–Nov. 2, 1977). IEEE Computer Society, Sept. 30, 1977, pp. 222–227. DOI: 10.1109/SFCS.1977.24.

[Yos11] Yuichi Yoshida. "Optimal Constant-Time Approximation Algorithms and (Unconditional) Inapproximability Results for Every Bounded-Degree CSP". In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*. STOC 2011 (June 6–8, 2011). Association for Computing Machinery, June 6, 2011, pp. 665–674. DOI: 10.1145/1993636.1993725.

A Proof of Theorem 3.5

We inductively show that for every $t \in [T]$,

$$\|(\boldsymbol{M}_{1:t}, \boldsymbol{S}^Y_{1:t}) - (\boldsymbol{M}_{1:t}, \boldsymbol{S}^N_{1:t})\|_{\mathsf{tvd}} \le (t/T)\delta \qquad \text{(Induction hypothesis)}$$

where $\mathcal{E}_0$ is the trivial event that is always true.

First, we prove the base case $t = 1$. Recalling that $\boldsymbol{S}^Y_0 = \boldsymbol{S}^N_0 = \perp$, we have

$$\begin{aligned}\|(\boldsymbol{M}_1, \boldsymbol{S}^Y_1) - (\boldsymbol{M}_1, \boldsymbol{S}^N_1)\|_{\mathsf{tvd},\mathcal{E}_1} &= \|(\boldsymbol{M}_1, \boldsymbol{S}^Y_1) - (\boldsymbol{M}_1, \mathtt{Player}_1(\boldsymbol{M}_1, \boldsymbol{S}^N_0, \boldsymbol{U}))\|_{\mathsf{tvd},\mathcal{E}_1} \\ &= \|(\boldsymbol{M}_1, \boldsymbol{S}^Y_1) - (\boldsymbol{M}_1, \mathtt{Player}_1(\boldsymbol{M}_1, \boldsymbol{S}^Y_0, \boldsymbol{U}))\|_{\mathsf{tvd},\mathcal{E}_1},\end{aligned}$$

where $\boldsymbol{U} \sim \mathsf{Unif}(\mathbb{Z}_q^{\alpha n \times k})$ and $\|\cdot\|_{\mathsf{tvd},\mathcal{E}}$ denotes the total variation distance, conditioned on event $\mathcal{E}$. Observe that for every fixed M_1 and S^Y_0 satisfying $\mathcal{E}_1$, we have $\|\boldsymbol{S}^Y_1 - \mathtt{Player}_1(M_1, S^Y_0, \boldsymbol{U})\|_{\mathsf{tvd}} \le \delta/(2T)$. It follows from data-processing inequality (Proposition 2.2) that

$$\|(\boldsymbol{M}_1, \boldsymbol{S}^Y_1) - (\boldsymbol{M}_1, \mathtt{Player}_1(\boldsymbol{M}_1, \boldsymbol{S}^Y_0, \boldsymbol{U}))\|_{\mathsf{tvd},\mathcal{E}_1} \le \delta/2T.$$

Therefore,

$$\begin{aligned}\|(\boldsymbol{M}_1, \boldsymbol{S}^Y_1) - (\boldsymbol{M}_1, \mathtt{Player}_1(\boldsymbol{M}_1, \boldsymbol{S}^Y_0, \boldsymbol{U}))\|_{\mathsf{tvd}} &\le \|(\boldsymbol{M}_1, \boldsymbol{S}^Y_1) - (\boldsymbol{M}_1, \mathtt{Player}_1(\boldsymbol{M}_1, \boldsymbol{S}^Y_0, \boldsymbol{U}))\|_{\mathsf{tvd},\mathcal{E}_1} + \Pr[\overline{\mathcal{E}_1}] \\ &\le \delta/(2T) + \Pr[\overline{\mathcal{E}_1}] \le \delta/T,\end{aligned}$$

which completes the base case.

For every $t = 2, \ldots, T$, we have

$$\begin{aligned}&\|(\boldsymbol{M}_{1:t}, \boldsymbol{S}^Y_{1:t}) - (\boldsymbol{M}_{1:t}, \boldsymbol{S}^N_{1:t})\|_{\mathsf{tvd}} \\ &\quad = \|(\boldsymbol{M}_{1:t}, \boldsymbol{S}^Y_{1:t-1}, \mathtt{Player}_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}^Y_{t-1}, \boldsymbol{Z}_t)) - (\boldsymbol{M}_{1:t}, \boldsymbol{S}^N_{1:t-1}, \mathtt{Player}_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}^N_{t-1}, \boldsymbol{U}))\|_{\mathsf{tvd}}.\end{aligned}$$

Let us define $\boldsymbol{Q}^Y_{t-1} = (\boldsymbol{M}_{1:t-1}, \boldsymbol{S}^Y_{1:t-1})$ and $\boldsymbol{Q}^N_{t-1} = (\boldsymbol{M}_{1:t-1}, \boldsymbol{S}^N_{1:t-1})$. Then, we can rewrite the above expression for total variation distance in terms of the new notation as follows:

$$\begin{aligned}&\|(\boldsymbol{M}_{1:t}, \boldsymbol{S}^Y_{1:t-1}, \mathtt{Player}_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}^Y_{t-1}, \boldsymbol{Z}_t)) - (\boldsymbol{M}_{1:t}, S^N_{1:t-1}, \mathtt{Player}_t(\boldsymbol{M}_{1:t}, \boldsymbol{S}^N_{t-1}, \boldsymbol{U}))\|_{\mathsf{tvd}} \\ &\quad = \|(\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \boldsymbol{Z}_t)) - (\boldsymbol{Q}^N_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^N_{t-1}, \boldsymbol{M}_t, \boldsymbol{U}))\|_{\mathsf{tvd}}. \qquad \text{(A.1)}\end{aligned}$$

We now apply Proposition 2.4 to Equation (A.1). Applying this lemma with $\boldsymbol{X}^1 = \boldsymbol{Q}^Y_{t-1}$, $\boldsymbol{X}^2 = \boldsymbol{Q}^N_{t-1}$, $\boldsymbol{Z}^1 = (\boldsymbol{M}_t, \boldsymbol{Z}_t)$, $\boldsymbol{Z}^2 = (\boldsymbol{M}_t, \boldsymbol{U})$, and f as the function that maps the tuple $(X, (B, C))$ to $(B, \mathtt{Player}_t(X, B, C))$, we get

$$\begin{aligned}&\|(\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \boldsymbol{Z}_t)) - (\boldsymbol{Q}^N_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^N_{t-1}, \boldsymbol{M}_t, \boldsymbol{U}))\|_{\mathsf{tvd}} \\ &\le \|\boldsymbol{Q}^Y_{t-1} - \boldsymbol{Q}^N_{t-1}\|_{\mathsf{tvd}} + \|(\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \boldsymbol{Z}_t)) - (\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^Y_{t-1}, \boldsymbol{M}_t, \boldsymbol{U}))\|_{\mathsf{tvd}}. \qquad \text{(A.2)}\end{aligned}$$

Now, by applying the induction hypothesis, we have that

$$\|\boldsymbol{Q}^{Y}_{t-1} - \boldsymbol{Q}^{N}_{t-1}\|_{\mathsf{tvd}} \leq (t-1)\delta/T. \tag{A.3}$$

Next, we bound the second term on the right hand side of (A.2), i.e.,

$$\|(\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \boldsymbol{Z}_t)) - (\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \boldsymbol{U}))\|_{\mathsf{tvd}},$$

by applying condition (iii) from Lemma 4.1. According to this condition, for every *fixed* $(M_{1:t})$ and $S^{Y}_{1:t-1}$ satisfying $\mathcal{E}_t$, we have

$$\|\mathtt{Player}_t(M_{1:t}, S^{Y}_{1:t-1}, \boldsymbol{Z}_t) - \mathtt{Player}_t(M_{1:t}, S^{Y}_{1:t-1}, \boldsymbol{U})\|_{\mathsf{tvd}} \leq \delta/(2T)$$

where $U \sim \mathsf{Unif}(\mathbb{Z}_q^{(k-1)\alpha n})$. By Proposition 2.3, it follows that

$$\|(\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \boldsymbol{Z}_t)) - (\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \mathtt{Player}_t(\boldsymbol{Q}^{Y}_{t-1}, \boldsymbol{M}_t, \boldsymbol{U}))\|_{\mathsf{tvd},\mathcal{E}_t} \leq \delta/(2T)\,. \tag{A.4}$$

Combining Eqs. (A.1) to (A.4), we have

$$\|(\boldsymbol{M}_{1:t}, \boldsymbol{S}^{Y}_{1:t}) - (\boldsymbol{M}_{1:t}, \boldsymbol{S}^{N}_{1:t})\|_{\mathsf{tvd}} \leq (t-1)\delta/T + \delta/(2T) + \Pr[\overline{\mathcal{E}_t} \mid \mathcal{E}_{t-1}] \leq t\delta/T,$$

which completes the induction.

Thus,

$$\|(\boldsymbol{M}_{1:T}, \boldsymbol{S}^{Y}_{1:T}) - (\boldsymbol{M}_{1:T}, \boldsymbol{S}^{N}_{1:T})\|_{\mathsf{tvd}} \leq \delta\,.$$

This implies that Π cannot have advantage more than δ, which contradicts the assumptions of the theorem statement. Therefore, we conclude that any protocol for DIHP with advantage δ requires τn bits of communication, as desired.

B Proof of Lemma 7.11

In this appendix, we prove Lemma 7.11, restated as follows:

Lemma 7.11. *For every $q, k \in \mathbb{N}$, there exists $\alpha_0 \in (0, 1/k)$ such that the following holds. For every $C_1, C_2, C_3, C_4 < \infty$, there exists $\varepsilon_0 > 0$ and $C_5 < \infty$ such that for every $s, n, h, u, \kappa, \eta \in \mathbb{N}$ with $\kappa \leq h \leq s \leq \varepsilon_0 n$, and $\kappa + \eta \leq u \leq \min\{n, h + \eta\}$, we have*

$$\mathrm{U}_{C_1,s,n}(u) \cdot C_2^{\kappa+\eta} \cdot p_{C_3,\alpha_0}(n,u,\kappa,\eta) \cdot \left(\frac{C_4\sqrt{sn}}{h+\eta-u}\right)^{(h-\eta+u)/2} \leq \left(\frac{C_5\sqrt{sn}}{h}\right)^{h/2}.$$

We recall the definitions:

$$\mathrm{U}_{C_1,s,n}(u) = \begin{cases} \left(\frac{C_1\sqrt{sn}}{u}\right)^{u/2} & \text{if } 0 \leq u \leq s, \\ \left(\min\left\{\frac{C_1\sqrt{n}}{\sqrt{u}}, \frac{eq^2 n}{u}\right\}\right)^{u/2} & \text{if } u > s, \end{cases} \tag{B.1}$$

$$\mathrm{U}_{C_2,s,n}(h) = \left(\frac{C_2\sqrt{sn}}{h}\right)^{h/2}, \tag{B.2}$$

$$p_{C,\alpha}(n,u,\kappa,\eta) = C^u \cdot \alpha^\kappa \cdot \left(\frac{u}{n}\right)^{\eta/2}, \tag{B.3}$$

where for Equation (B.2) we used the assumption that $h \leq s$.

Our first step is showing that the following lemma implies Lemma 7.11:

Lemma B.4. *There exists $\alpha_0 \in (0, 1/k)$ such that for every $C_1, C_{\mathrm{LHS}} < \infty$, there exist $\varepsilon_0 > 0$ and $C_{\mathrm{RHS}} < \infty$ such that the following holds.*

For every $\alpha \in (0, \alpha_0)$, and $s, n, u, h, \kappa, \eta \in \mathbb{N}$ with $h \le s \le \varepsilon_0 n$ and $\kappa + \eta \le u \le \min\{n, h + \eta\}$,

$$\mathrm{U}_{C_1,s,n}(u) \cdot C_{\mathrm{LHS}}^u \cdot \alpha^\eta \cdot \frac{u^{\eta/2} \cdot h^{u/2-\eta/2}}{n^{u/4+\eta} \cdot s^{u/4-\eta}} \le C_{\mathrm{RHS}}^h,$$

where $\mathrm{U}_{C_1,s,n}(u)$ is defined as in Equation (B.3) above.

Proof of Lemma 7.11 using Lemma B.4. Let $h' := h + \eta - u$. We invoke Lemma B.4 with $C_{\mathrm{LHS}} := C_2 C_3$.

To prove Lemma 7.11, it suffices to pick $C_5 < \infty$ such that:

$$\mathrm{U}_{C_1,s,n}(u) \cdot C_2^u \cdot \left(\frac{C_3^u \cdot \alpha^\kappa \cdot u^{\eta/2}}{n^{\eta/2}}\right) \cdot \left(\frac{C_4^{h'/2} \cdot s^{h'/4} \cdot n^{h'/4}}{(h')^{h'/2}}\right) \le \left(\frac{C_5^{h/2} \cdot s^{h/4} \cdot n^{h/4}}{h^{h/2}}\right).$$

Simplifying and moving factors between sides gives the equivalent desideratum:

$$\mathrm{U}_{C_1,s,n}(u) \cdot C_{\mathrm{LHS}}^u \cdot \alpha^\kappa \cdot \frac{u^{\eta/2}}{n^{\eta/2+(h-h')/4} \cdot s^{(h-h')/4}} \cdot \frac{h^{h/2}}{(h')^{h'/2}} \le \frac{C_5^{h/2}}{C_4^{h'/2}}.$$

Now, we apply several simple inequalities. We observe $h^{h/2}/(h')^{h'/2} = (h/(h'/2))^{h'/2} \cdot 2^{-h'/2} \cdot h^{(h-h')/2} \le (e/\sqrt{2})^{h'} \cdot h^{(h-h')/2}$ by Proposition 7.7. This, together with the fact that $h' \le h$ and (WLOG) $C_4 > 1$, means that the following inequality implies our desideratum:

$$\mathrm{U}_{C_1,s,n}(u) \cdot C_{\mathrm{LHS}}^u \cdot \alpha^\kappa \cdot \frac{u^{\eta/2}}{n^{\eta/2+(h-h')/4} \cdot s^{(h-h')/4}} \cdot h^{(h-h')/2} \le \frac{C_5^{h/2}}{C_4^{h/2}} \cdot (2/\sqrt{e})^{h/2}.$$

Substituting in $h - h' = u - \eta$, this is exactly the hypothesis of Lemma B.4. Letting C_{RHS} denote the resulting constant from Lemma B.4, we can pick $C_5 := C_{\mathrm{RHS}}^2 \cdot C_4 \cdot 2/\sqrt{e}$. □

Proof of Lemma B.4. Let $\varepsilon := s/n$ (so that $s = \varepsilon n$).

We split into five cases based on u: (1a) $1 \le u \le h$, (1b) $h < u \le s$, (2a) $s < u \le 16s$, (2b) $16s < u \le \sqrt{\varepsilon} n$, (3) $\sqrt{\varepsilon} n < u \le n$. Let

$$Q := \mathrm{U}_{C_1,s,n}(u) \cdot C_{\mathrm{LHS}}^u \cdot \alpha^{\eta/2} \cdot \frac{u^{\eta/2} \cdot h^{u/2-\eta/2}}{n^{u/4+\eta/4} \cdot s^{u/4-\eta/4}}$$

denote the quantity we want to bound. We determine C_{RHS} later.

Case 1: $1 \le u \le s$. In this regime,

$$\mathrm{U}_{C_1,s,u}(n) = C_1^{u/2} \cdot \frac{s^{u/4} \cdot n^{u/4}}{u^{u/2}}.$$

Using the trivial bound $\alpha \le 1$, we get:

$$Q \le C_1^{u/2} \cdot \frac{s^{u/4} \cdot n^{u/4}}{u^{u/2}} \cdot C_{\mathrm{LHS}}^u \cdot \frac{u^{\eta/2} \cdot h^{u/2-\eta/2}}{n^{u/4+\eta/4} \cdot s^{u/4-\eta/4}} = \underbrace{(C_1 C_{\mathrm{LHS}})^u \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} \cdot \frac{s^{\eta/4}}{n^{\eta/4}}}_{=:Q_1}. \tag{B.5}$$

We continue bounding Q_1 in subcases.

Subcase 1a: $1 \le u \le h$. We have:

$$Q_1 \le (C_1 C_{\text{LHS}})^u \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} \le (C_1 C_{\text{LHS}})^h \cdot \frac{h^u}{u^u} \le (eC_1 C_{\text{LHS}})^h.$$

where the first equality uses the trivial bound $s \le n$, the second uses the assumption $h \ge u$ (and $\eta/2 \ge 0$), and the third is Proposition 7.7. This suffices as long as $C_{\text{RHS}} \ge eC_1 C_{\text{LHS}}$.

Case 1b: $h < u \le s$. We now have:

$$\begin{aligned} Q_1 = (C_1 C_{\text{LHS}})^u \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} \cdot \varepsilon^{\eta/4} &\le (C_1 C_{\text{LHS}})^u \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} \cdot \varepsilon_0^{\eta/4} \le (C_1 C_{\text{LHS}})^u \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} \cdot \varepsilon_0^{(u-h)/4} \\ &= (C_1 C_{\text{LHS}} \varepsilon_0^{1/4})^u \cdot (1/\varepsilon_0^{1/4})^h \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} \le (C_1 C_{\text{LHS}} \varepsilon_0^{1/4})^u \cdot (1/\varepsilon_0^{1/4})^h, \end{aligned}$$

using, respectively, the definition of s, the assumed bound on ε_0, the assumption that $u - h \le 2\eta/2$, rearranging, and finally, that $h \le u$ and $u \ge 2\eta/2$. This suffices as long as $\varepsilon_0 \le 1/(C_1 C_{\text{LHS}})^4$ and $C_{\text{RHS}} \ge 1/\varepsilon_0^{1/4}$. (We note that the inequality on Q_1 we showed here did not need $u \le s$, only that $h < u$; we will use this fact in a subsequent subcase.)

Case 2: $s < u \le \sqrt{\varepsilon} n$. In this regime, we use:

$$\mathrm{U}_{C_1,s,u}(n) \le \frac{C_1^{u/2} \cdot u^{u/4} \cdot n^{u/4}}{u^{u/2}} = \frac{C_1^{u/2} \cdot n^{u/4}}{u^{u/4}}$$

This gives

$$Q \le C_1^{u/2} \cdot \frac{n^{u/4}}{u^{u/4}} \cdot C_{\text{LHS}}^u \cdot \frac{u^{\eta/2} \cdot h^{u/2-\eta/2}}{n^{u/4+\eta/4} \cdot s^{u/4-\eta/4}} = \underbrace{(\sqrt{C_1} C_{\text{LHS}})^u \cdot \frac{u^{u/4} \cdot s^{\eta/4}}{n^{\eta/4} \cdot s^{u/4}} \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}}}_{=:Q_2},$$

and we proceed to bound Q_2. Note that we artificially split $\eta/2 - u/4 = u/4 - (u/2 - \eta/2)$ (this will be useful for subsequent calculations).

Case 2a: $s < u \le 16s$. We have

$$Q_2 \le (\sqrt{C_1} C_{\text{LHS}})^u \cdot \frac{(16s)^{u/4} \cdot s^{\eta/4}}{n^{\eta/4} \cdot s^{u/4}} \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} = (2\sqrt{C_1} C_{\text{LHS}})^u \cdot \frac{s^{\eta/4}}{n^{\eta/4}} \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} = 2^u Q_1,$$

where the equality uses the assumption $u \le 16s$. In the final RHS, Q_1 is the quantity considered in case (1b), wherein we showed — without assuming $u < s$ — that $Q_1 \le (C_1 C_{\text{LHS}} \varepsilon_0^{1/4})^u \cdot (1/\varepsilon_0^{1/4})^h$. Hence $2^u Q_1 \le (2C_1 C_{\text{LHS}} \varepsilon_0^{1/4})^u \cdot (1/\varepsilon_0^{1/4})^h$, which is sufficiently bounded if $\varepsilon_0 \le 1/(2C_1 C_{\text{LHS}})^4$.

Case 2b: $16s \le u < \sqrt{\varepsilon} n$. In this case,

$$Q_2 \le (\sqrt{C_1} C_{\text{LHS}})^u \cdot \frac{(\sqrt{\varepsilon} n)^{u/4} \cdot (\varepsilon n)^{\eta/4}}{n^{\eta/4} \cdot (\varepsilon n)^{u/4}} \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} \le (\sqrt{C_1} C_{\text{LHS}})^u \cdot \varepsilon^{\eta/4-u/8}.$$

Here, the first inequality uses the assumption on u (and rewriting the definition of s), and the second uses that $h \le u$ (and $u \ge 2\eta/2$, and consolidates powers of ε). Now, $\eta/4 - u/8 \ge (u-h)/4 - u/8 = u/8 - h/4 \ge 2s - h/4 \ge 2h - h/4 \ge 0$, and therefore $\varepsilon^{\eta/4-u/8} \le \varepsilon^{u/8-h/4} \le \varepsilon_0^{u/8-h/4}$, and so

$$Q_2 \le (\sqrt{C_1} C_{\text{LHS}} \varepsilon_0^{1/8})^u \cdot (1/\varepsilon_0^{1/4})^h.$$

Thus, we get the required bound as long as $\varepsilon_0 \le 1/(\sqrt{C_1} C_{\text{LHS}})^8$ and $C_{\text{RHS}} \ge 1/\varepsilon_0^{1/4}$.

Case 3: $\sqrt{\varepsilon}n \le u \le n$**.** Here, we use the final bound

$$\mathrm{U}_{C_1,s,u}(n) \le \left(\frac{8en}{u}\right)^{u/2} = C_{\mathrm{W}}^{u/2} \cdot \frac{n^{u/2}}{u^{u/2}}.$$

Thus, we have:

$$Q \le C_{\mathrm{W}}^{u/2} \cdot \frac{n^{u/2}}{u^{u/2}} \cdot C_{\mathrm{LHS}}^{u} \cdot \alpha^{b+\eta/2} \cdot \frac{u^{\eta/2} \cdot h^{u/2-\eta/2}}{n^{u/4+\eta/4} \cdot s^{u/4-\eta/4}} \le (\sqrt{C_{\mathrm{W}}} C_{\mathrm{LHS}})^{u} \cdot \alpha_0^{(u-h)/2} \cdot \underbrace{\frac{n^{u/4-\eta/4}}{s^{u/4-\eta/4}} \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}}}_{=:R_3},$$

where the equality uses $\alpha \le \alpha_0 \le 1$ and $b \ge 0$, $\eta/2 \ge (u-h)/2 \ge 0$. We will verify that $R_3 \le 1$. Assuming this, we get

$$Q \le (\sqrt{C_{\mathrm{W}}} C_{\mathrm{LHS}})^{u} \cdot \alpha_0^{(u-h)/2} = (\sqrt{C_{\mathrm{W}}} C_{\mathrm{LHS}} \alpha_0^{1/2})^{u} \cdot (1/\alpha_0^{1/2})^{h},$$

which is sufficiently small as long as $\alpha_0 \le 1/(\sqrt{C_{\mathrm{W}}} C_{\mathrm{LHS}})^2$ and $C_{\mathrm{RHS}} \ge 1/\alpha_0^{1/2}$.

Finally, we bound

$$R_3 = \frac{n^{u/4-\eta/4}}{(\varepsilon n)^{u/4-\eta/4}} \cdot \frac{h^{u/2-\eta/2}}{u^{u/2-\eta/2}} = \left(\frac{h}{\sqrt{\varepsilon} u}\right)^{u/2-\eta/2},$$

and we use $h \le s = \varepsilon n$ and $u \ge \sqrt{\varepsilon} n$ to get $(h/(\sqrt{\varepsilon} u)) \le 1$. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il