Electronic Colloquium on Computational Complexity, Report No. 22 (2025)

Boosting uniformity in quasirandom groups: fast and simple

Harm Derksen*
Northeastern University
ha.derksen@northeastern.edu

Chin Ho Lee
North Carolina State University
chinho.lee@ncsu.edu

Emanuele Viola†
Northeastern University
viola@ccs.neu.edu

March 3, 2025

Abstract

We study the communication complexity of multiplying $k \times t$ elements from the group $H = \mathrm{SL}(2, q)$ in the number-on-forehead model with k parties. We prove a lower bound of $(t \log H)/c^k$. This is an exponential improvement over previous work, and matches the state-of-the-art in the area.

Relatedly, we show that the convolution of k^c independent copies of a 3-uniform distribution over H^m is close to a k-uniform distribution. This is again an exponential improvement over previous work which needed c^k copies.

The proofs are remarkably simple; the results extend to other quasirandom groups.

We also show that for any group H, any distribution over H^m whose weight-k Fourier coefficients are small is close to a k-uniform distribution. This generalizes previous work in the abelian setting, and the proof is simpler.

1 Introduction and our results

Iterated multiplication of elements in a group is a fundamental problem that has a long history and wide-ranging applications, and is linked to long-standing open problems. Already in [LZ77] it has been pivotal to provide space-efficient algorithms for Dyck languages. Depending on the underlying group, iterated multiplication is complete for various complexity classes [KMR66, MC87, Mix89, BC92, IL95, Mil14]. For example, Barrington's famous result [Mix89]

*Partially supported by NSF grant DMS 2147769.
†Supported by NSF grant CCF-2114116.

ISSN 1433-8092

shows that it is complete for NC^1 if and only if the underlying group is non-solvable. This in particular disproved a conjecture about the complexity of majority [BDFP83]. This type of results has then been taken further in the study of *catalytic computation* [BCK+14], leading to other surprising discoveries [BCK+14, CM20].

The focus of this paper is on number-on-forehead communication complexity [CFL83]. For a survey on the communication complexity of group products, see [Vio19], and see [KN97, RY19] for general background on communication complexity. Concretely, the input is a matrix of $k \times t$ elements a_{ij} from a group H, and the goal is computing $\prod_{j=1}^{t} a_{1j} \cdots a_{kj}$. There are k collaborating parties, with Party i knowing all the input except row i.

This problem is also linked to central open problems in communication complexity. Specifically, [GV19] conjectured that over certain groups this problem remains hard even for k larger than $\log n$. Establishing such bounds is arguably the most significant open problem in the area. A number of candidates have been put forward over the years, but many have been ruled out via ingenious protocols, e.g. in [PRS97, BGKL03, BC08, ACFN15]. Interestingly, for the iterated-product candidate proposed in [GV19], no non-trivial protocol is known.

Iterated group products are also candidate for providing strong separations between randomized and deterministic number-on-forehead communication. The current bounds (see [Vio19]) give a separation matching one in [BDPW07]. Stronger bounds could simplify and strengthen the recent exciting separation [KLM23].

Returning to the problem, we note that its complexity heavily depends on the underlying group. If it is abelian, then the problem can be solved with constant communication, using the public-coin protocol for equality. Over certain other groups a communication lower bound of t/c^k follows via [Mix89] from the landmark lower bound in [BNS92] for generalized inner product; cf. [MV13]. However, this bound does not improve with the size of the group. In particular it is far from the (trivial) upper bound of $t \log H$, and it gives nothing when t is constant. Motivated by a cryptographic application, [MV13] asked whether a lower bound that grows with the size of the group, ideally $ct \log H$, can be established over some group H.

Gowers and Viola [GV15, GV19] proved a bound of $(t \log H)/c^{c^k}$ for the group $\mathrm{SL}(2, q)$ of 2×2 invertible matrices over $\mathbb{F}_q$, which enables the motivating application from [MV13]. Subsequent work [DV23] simplified the proof and generalized it to any quasi-random group [Gow08], see also [GV19, Sha16]. While such bounds do grow with the size of the group, thus answering the question in [MV13] and enabling the motivating application in cryptography, the dependency on the number k of parties is weak: One can only afford k doubly-logarithmic in the input length.

In this work we give an exponential improvement and obtain bounds of the form $(t \log H)/c^k$, thus matching the state-of-the-art in number-on-forehead communication [BNS92]. As in [GV19], we prove stronger results that even bound the advantage such protocols have when the input is promised to multiply to one of two fixed elements.

Theorem 1. *Let $H = \mathrm{SL}(2, q)$. Let $P \colon H^{k \times t} \to [2]$ be a number-on-forehead communication protocol with k parties and communication b bits. For $g \in H$ denote by p_g the probability that P outputs 1 over a uniform input $(a_{i,j})_{i \le k, j \le t}$ such that $\prod_{j=1}^{t} a_{1j} \cdots a_{kj} = g$. For any k and any two $g, h \in H$, if $t \ge c^k$ then $|p_g - p_h| \le 2^b \cdot H^{-t/c^k}$.*

The high-level proof technique is the same as in [GV19]. They reduced the problem to *boosting uniformity* over m copies of H.

Definition 2. A distribution p over a set S is ϵ-uniform if $|p(x) - 1/S| \leq \epsilon/S$. If $S = H^t$ and $k \leq t$ we say p is (ϵ, k)-uniform if for any k coordinates, the induced distribution over those coordinates is ϵ-uniform. We say p is k-uniform if it is $(0, k)$-uniform.

[GV19] showed that if s is a 2-uniform distribution over H^m then the convolution (a.k.a. component-wise product) of some ℓ independent copies of s is H^{-m}-uniform over the whole space H^m. Note that such a result is false for abelian groups – the convolution can remain only 2-uniform. Quantitatively, they show that $\ell = c^m$ copies suffice. In the application to Theorem 1 one has $m = 2^k$, which gives the doubly logarithmic dependence on k.

In this work we give a corresponding exponential improvement on the number of copies required to boost uniformity: we show that in fact $\ell = m^c$ copies suffice. Our proof is remarkably simple, especially if we start with 3-uniform distributions, which we note suffices for Theorem 1. (We discuss below extensions to 2-uniform and other groups.) We state this result next.

Theorem 3. *Let $H = \mathrm{SL}(2, q)$. Let p be a 3-uniform distribution over H^m. The convolution of m^c independent copies of p is H^{-m}-uniform.*

Our approach allows us to double the uniformity, i.e., go from k-uniform to $2k$-uniform using only a *constant* number of convolutions, independently of k, whereas [GV19] would use $\geq k$ convolutions. This points to a key difference in the techniques. In [GV19], boosting uniformity is achieved by reduction to *interleaved products,* and appears tailored to going from 2-uniform to 3-uniform. Our approach is different, and simpler, even taking into account the simple proof of interleaved mixing from [DV23]. It can be seen as a k-uniform version of the flattening lemmas discovered in [Gow08, BNP08, GV19]. In a nutshell, the k-uniformity assumption allows us to remove all "low-degree" Fourier coefficients, leaving only those of degree $> k$. Then the quasi-randomness assumption, combined with the tensor-product structure of the Fourier coefficients allows us to "flatten" distributions at a rate proportional to H^{-ck}, instead of H^{-c} as in previous work. We note that while using k-uniformity to remove low-degree coefficients is a common proof technique, we are not aware of previous work where this is done in the non-abelian setting. This might indicate that our techniques might find other applications, and in general we advocate a systematic study of non-abelian analogues of the Fourier toolkit. Another step in this direction is discussed next.

(ϵ, k)-uniformity vs. k-uniformity. Extending the classic work [AGM03], Rubinfeld and Xie [RX13] showed that every almost k-uniform distribution over any Abelian product group is statistically close to some k-uniform distribution. We generalize their result to any product group. Our approach is significantly simpler. [RX13] decomposes the given k-uniform distribution in a real orthogonal basis instead of the Fourier basis; we show that in fact the same argument can be carried out directly over the Fourier basis. A critical observation is that removing Fourier coefficients of a fixed weight from a real function keeps the function real.

This generalization, in combination with Theorem 3, gives a refinement of Theorem 3 where the number of copies is k^c and the final distribution is statistically close to a k-uniform distribution (whereas a direct application of Theorem 3 would just give an (H^{-k^c}, k)-uniform distribution).

Extensions. Theorem 1 and Theorem 3 above can be generalized to any quasi-random group and to distributions which are 2-uniform. This can be done by first using the results in [GV19, DV23] to boost 2-uniformity to v-uniformity for a sufficiently large constant v depending on the quasirandomness of the group (for $\mathrm{SL}(2,q)$, $v=3$ suffices). This requires a number of convolutions that is exponential in v, but since v is constant it can be afforded. After that, our results kick in and allow to boost faster.

2 Preliminaries

In this section we fix some notation, especially about representation theory.

For a set X, we also write X for its size $|X|$. We write $[i]$ for the set $\{0,1,\dots,i-1\}$. Every occurrence of "c" denotes a possibly different universal constant. Replacing "c" with $O(1)$ everywhere is consistent with a common interpretation of the latter. For a function $f\colon G\to\mathbb{C}$ we denote by $|f|_2^2$ the un-normalized quantity $\sum_{x\in G}|f(x)|^2$.

Next we present the standard framework of representation theory. The books by Serre [Ser77], Diaconis [Dia88], and Terras [Ter99] are good references for representation theory and non-abelian Fourier analysis. The Barbados notes [Wig10] or Section 13 of [Gow17] or [GV22] provide briefer introductions. The exposition in these sources is not always consistent, and often has different aims from ours. So let us give a quick account of the theory that is most relevant for this work.

Matrices. Let M be a square complex matrix. We denote by $\mathrm{tr}(M)$ the trace of M, by $\overline{M}$ the conjugate of M, by M^T the transpose of M, and by M^* the conjugate transpose $\overline{M^T}$ (aka adjoint, Hermitian conjugate, etc.). The matrix M is *unitary* if the rows and the columns are orthonormal; equivalently $M^{-1}=M^*$.

We denote by

$$|M|_2^2 := \sum_{i,j}|M_{i,j}|^2 = \mathrm{tr}(MM^*).$$

This is known as the Frobenius norm, or Schatten 2-norm, or Hilbert-Schmidt operator, etc.

If $M=AB$ we have

$$|M|_2^2 = \sum_{i,j}\Big|\sum_k A_{i,k}B_{k,j}\Big|^2 \le \sum_{i,j}\Big(\sum_k |A_{i,k}|^2\Big)\Big(\sum_k |B_{k,j}|^2\Big) = |A|_2^2|B|_2^2, \tag{1}$$

where the inequality is Cauchy–Schwarz.

Representation theory. Let G be a group. A *representation* ρ of G with dimension d maps elements of G to $d \times d$ unitary, complex matrices so that $\rho(xy) = \rho(x)\rho(y)$. Thus, ρ is a homomorphism from G to the group of linear transformations of the vector space $\mathbb{C}^d$. We denote by d_ρ the dimension of ρ.

If there is a non-trivial subspace W of $\mathbb{C}^d$ that is invariant under ρ, that is, $\rho(x)W \subseteq W$ for every $x \in G$, then ρ is *reducible*; otherwise it is *irreducible.* Irreducible representations are abbreviated *irreps* and play a critical role in Fourier analysis. We denote by $\widehat{G}$ a complete set of inequivalent irreducible representations of G.

In every group we have

$$\sum_{\rho \in \widehat{G}} d_\rho^2 = G. \tag{2}$$

We have the following fundamental orthogonality principle.

Lemma 4 (Schur's lemma, see [Dia88], Page 11 or Lemma 2.3.3 in [Wig10])**.** *Let ρ, ψ be irreps. Then $\mathbb{E}_x \rho(x)_{k,h}\overline{\psi(x)}_{i,j}$ is 0 unless $\rho = \psi$ and $k = i$ and $h = j$, in which case it is $1/d_\rho$. In particular, $\mathbb{E}_x|\rho(x)_{i,j}|^2 = 1/d_\rho$.*

Let $f\colon G \to \mathbb{C}$. The *ρ-th Fourier coefficient* of f is

$$\widehat{f}(\rho) := \mathbb{E}_x f(x)\overline{\rho(x)}.$$

The Fourier inversion formula is then

$$f(x) = \sum_{\rho \in \widehat{G}} d_\rho \operatorname{tr}\big(\widehat{f}(\rho)\rho(x)^T\big).$$

We define the *convolution* as follows:

$$p * q(x) := \sum_y p(y)q(y^{-1}x).$$

Note that if p and q are distributions then $p * q$ is the distribution obtained by sampling x from p, y from q, and then outputting xy.

We note that under this normalization we have

$$\widehat{p * q}(\alpha) = G \cdot \widehat{p}(\alpha)\widehat{q}(\alpha).$$

Combining this with 1 we obtain

$$|\widehat{p * q}(\alpha)|_2^2 \le G^2 \cdot |\widehat{p}(\alpha)|_2^2 |\widehat{q}(\alpha)|_2^2. \tag{3}$$

Parseval's identity is

$$\mathbb{E} f(x)\overline{g(x)} = \sum_\rho d_\rho \operatorname{tr}\big(\widehat{f}(\rho)\widehat{g}(\rho)^*\big).$$

In case $f = g$ this becomes

$$\mathbb{E}|f(x)|^2 = \sum_\rho d_\rho \operatorname{tr}\big(\widehat{f}(\rho)\widehat{f}(\rho)^*\big) = \sum_\rho d_\rho |\widehat{f}(\rho)|_2^2.$$

Fact 5 (Theorem 10 in Section 3.2 in [Ser77], or Theorem 9 in [Dia88])**.** *Any irrep ρ of H^n is the tensor product $\otimes_{i=1}^n \rho_i$ of n irreps ρ_i of H. In particular, the dimension of ρ is the product of the dimensions of the ρ_i.*

For $\rho = \otimes_{i=1}^n \rho_i$ we denote by $|\rho|$ the number of i s.t. ρ_i is not the trivial representation 1.

Definition 6 ([Gow08])**.** A group H is d-quasirandom if every non-trivial irrep of H has dimension $\geq d$.

3 Boosting uniformity

In this section we prove Theorem 3. The proof follows by repeated applications of the following theorem.

Theorem 7. *Let $H = \mathrm{SL}(2, q)$. Let p be a distribution over H^t that is (H^{-k}, k)-uniform for $k \geq 3$ and $m = \lceil (1+c)k \rceil$. Then the convolution of c independent copies of p is H^{-m}-uniform.*

Note for small k we may have $m = k + 1$. But if $k \geq c$ then m is a constant factor larger than k.

The choice of the error parameter is not too important because it can be boosted with convolutions:

Lemma 8 (Lemma 3.3 in [GV19])**.** *Let p and q be (ϵ, k)-uniform distributions over H^m. Then $p * q$ is (ϵ^2, k)-uniform.*

Proof. It is enough to consider the case $m = k$. We have

$$\left|p * q(x) - 1/H^t\right| = \left|\sum_y \left(p(y^{-1}) - 1/H^t\right)\left(q(yx) - 1/H^t\right)\right| \leq \sum_y (\epsilon/G)^2 = \epsilon^2/H^t. \qquad \square$$

In the rest of this section, we prove Theorem 7. The proof involves an excursion to 2-norms. The main step is the following new flattening lemma which can be seen as a k-wise variant of the flattening lemmas discovered in [Gow08, BNP08, GV19].

Lemma 9. *Let p be a distribution over H^m where H is d-quasirandom. Suppose p is (H^{-k}, k)-uniform. Then $|p * p - u|_2^2 \leq |p - u|_2^2 \cdot 2 \cdot H^{m-k} d^{-(k+1)}$.*

We need the following couple of claims to go back-and-forth between ϵ-uniform and 2-norms.

Claim 10. $|p * p - u|_\infty \leq |p - u|_2^2$.

Proof. $(p * p - 1/G)(x) = \sum_y (p(y) - 1/G)(p(y^{-1}x) - 1/G) \leq \sum_x (p(x) - 1/G)^2$. The last inequality is Cauchy–Schwarz. $\square$

Claim 11. Let p be an ϵ-uniform distribution over a group G, and let ρ be a non-trivial representation of ρ with dimension d_ρ. Then $|\widehat{p}(\rho)|_2^2 \leq d_\rho \epsilon^2 G^{-2}$

Proof. The LHS is

$$\begin{aligned}
\sum_{i,j} |\widehat{p}(\rho)_{i,j}|^2 &= \sum_{i,j} \left|\mathbb{E}_x\left[p(x)\overline{\rho(x)}\right]_{i,j}\right|^2 \\
&= \sum_{i,j} \left|\mathbb{E}_x\left[(G^{-1}+\epsilon_x G^{-1})\overline{\rho(x)}\right]_{i,j}\right|^2 && \text{(for some } \epsilon_x \text{ with } |\epsilon_x| \le \epsilon\text{)} \\
&= \sum_{i,j} \left|\mathbb{E}_x\left[\epsilon_x G^{-1}\overline{\rho(x)}_{i,j}\right]\right|^2 && \text{(by Lemma 4 with } \psi := 1\text{, using that } \rho \text{ is non-trivial)} \\
&\le G^{-2} \sum_{i,j} \mathbb{E}_x\left[\epsilon_x^2 \cdot |\overline{\rho(x)}_{i,j}|\right]^2 \\
&\le G^{-2} \sum_{i,j} \epsilon^2 \mathbb{E}_x |\overline{\rho(x)}_{i,j}|^2 \\
&= G^{-2} \sum_{i,j} \epsilon^2/d_\rho && \text{(by Lemma 4, see "in particular" part)} \\
&= G^{-2} d_\rho \epsilon^2. && \square
\end{aligned}$$

Proof of Lemma 9. Write G for the group H^m. For any distribution q we have

$$|q-u|_2^2 = |q|_2^2 - 1/G = G\sum_{\rho} d_\rho |\widehat{q}(\rho)|_2^2 - 1/G = G\sum_{\rho\neq 1} d_\rho |\widehat{q}(\rho)|_2^2.$$

In our case $q = p * p$, and using 3 and the above equality we bound the RHS by

$$\le G^3 \sum_{\rho\neq 1} d_\rho |\widehat{p}(\rho)|_2^4 \le G^2 \cdot |p-u|_2^2 \cdot \max_{\rho\neq 1} |\widehat{p}(\rho)|_2^2.$$

It remains to bound $G^2 \max_{\rho\neq 1} |\widehat{p}(\rho)|_2^2$. We consider two cases:

If $|\rho| > k$, then $d_\rho \ge d^{k+1}$ by Fact 5, so we simply use Parseval to bound

$$G^2 |\widehat{p}(\rho)|_2^2 \le G|p|_2^2/d_\rho \le G|p|_2^2/d^{k+1}.$$

We also have $|p|_2^2 \le (\max_x p(x)) \cdot \sum_x p(x) = \max_x p(x) \le 2/H^k$, because p is in particular $(1,k)$-uniform. Hence, we get a bound of $G \cdot 2 \cdot H^{-k} d^{-(k+1)}$, as desired.

If $|\rho| \le k$, then restrict to the non-trivial coordinates of ρ. On those coordinates, p induces a distribution that is H^{-k}-uniform. By Claim 11, we have

$$G^2 |\widehat{p}(\rho)|_2^2 \le d_\rho H^{-2k}.$$

Note $d_\rho \le H^{k/2}$ by 2. Thus, we obtain a bound of $H^{-1.5k} \le d^{-(k+1)}$. $\square$

We can now present the proof of Theorem 7.

Proof of Theorem 7. It is known that H is $\ge cH^{1/3}$-quasirandom, a proof can be found in [DSV03]. Hence, the parameter $d^{-(k+1)}$ in Lemma 9 is $\le cH^{-(k+1)/3} \le H^{-ck}$ for any $k \ge 3$. Also, we have $|p-u|_2^2 = |p|_2^2 - 1/G$. If p is (H^{-k}, k)-uniform then $|p|_2^2 \le \max_x p(x) \le 2/H^k$. Moreover, the uniformity parameter is maintained when taking convolutions. So one can apply the lemma a constant number of times to drive the L_2 norm to H^{-m}, and then convolve one more time to obtain a distribution that is H^{-m}-uniform via Claim 10. $\square$

4 Proof of Theorem 1

Let $m := 2^k$. As noted in [GV19], an application of the box norm (Corollary 3.11 in [VW08]) shows that the LHS in the conclusion is $\leq cH2^d$ times the statistical distance between the uniform distribution over H^m and the convolution of t independent copies of the following distribution s over H^m.

Definition 12. Pick u_i^0, u_i^1 for $i \in [k]$ uniformly from H. For $x \in [2]^k$ the x coordinate $s(x)$ of s is defined to be $\prod_{i\in[k]} u_i^{x_i}$.

Claim 13. s is 3-uniform.

Proof. Pick a coordinate i s.t. $x_i \neq y_i$. W.l.o.g. let $i = 0$, $x_0 = 0$, and $y_0 = 1$. Now z_0 is equal to either x_0 or y_0. Assume w.l.o.g. that $z_0 = y_0$. Consider a coordinate j where $z_j \neq y_j$. Assume again w.l.o.g. that $j = 1$. We can fix all other u_i with $i \geq 2$ and prove 3-uniformity just considering those two coordinates. For concreteness, details follow.

Up to swapping y and z there are only two cases to consider. The first is

$$x = 00$$
$$y = 10$$
$$z = 11.$$

In this case we can fix arbitrarily the u corresponding to y, and then x is uniform because of u_0^0 and z because of u_1^1.

Alternatively,

$$x = 01$$
$$y = 10$$
$$z = 11.$$

In this case we can similarly fix arbitrarily the u corresponding to z. □

We note that s is not 4-uniform, again just considering two coordinates.

To conclude the proof of Theorem 1, note that the convolution of m^c copies of s is (H^{-m}, m)-uniform by Theorem 3. By Lemma 8 the convolution of t copies is then $(H^{-m\cdot t/m^c}, m)$-uniform, and the result follows.

5 From (ϵ, k)-uniform to k-uniform

In this section we prove the following generalization of [RX13].

Theorem 14. *Let p be a distribution on $G = H^m$ s.t. $|\widehat{p}(\rho)|_2 \leq \epsilon/G$ for every $\rho : |\rho| \in [1, k]$. Then p has distance at most $3(mH)^{2k}\epsilon$ from a k-uniform distribution.*

First we note the following converse to Claim 11.

Claim 15. Let p be a distribution over H^m. Suppose $\widehat{p}(\rho) = 0$ whenever $|\rho| \in \{1, \dots, k\}$. Then p is k-uniform.

Proof. Consider any k coordinates; assume they are the first k w.l.o.g. The probability of a string $x \in H^k$ is

$$\sum_{y \in H^{m-k}} p(xy).$$

By the inversion formula and the assumption this is

$$\sum_y \sum_\rho d_\rho \operatorname{tr}\big(\widehat{p}(\rho)\rho(xy)^T\big) = \sum_y \widehat{p}(1) + \sum_{\rho:|\rho|>k} d_\rho \operatorname{tr}\Big(\widehat{p}(\rho) \sum_y \rho(xy)\Big).$$

We have $\widehat{p}(1) = \mathbb{E}_x p(x) = 1/H^m$ and so the first summand is $1/H^k$. We show that the second summand is 0 by showing that $\sum_y \rho(xy)$ is the zero matrix. To verify this, write ρ as a tensor product of ρ_i using Fact 5. Then one entry of $\rho(xy)$ is the product of the entries of the ρ_i. There is a non-trivial ρ_i corresponding to a y coordinate. The sum over that coordinate of y yields 0 by Lemma 4. □

Proof of Theorem 14. Let

$$\ell(x) := \sum_{\rho:|\rho|\in[1,k]} d_\rho \operatorname{tr}\big(\widehat{p}(\rho)\rho(x)^T\big)$$

be the "low-degree" part of p, and let

$$p'(x) := p(x) - \ell(x) = \sum_{\rho:|\rho|\notin[1,k]} d_\rho \operatorname{tr}\big(\widehat{p}(\rho)\rho(x)^T\big).$$

We first observe that ℓ and hence p' is real. This is because the conjugate $\overline{\rho}$ of an irrep is also an irrep, and over H^t the number of non-trivial coordinates is the same. Hence the sum over ρ is the same as the sum over $\overline{\rho}$. Moreover, $\widehat{p}(\overline{\rho}) \cdot \overline{\rho} = \overline{\widehat{p}(\rho) \cdot \rho}$. Therefore, we can write

$$2\ell(x) = \sum_{\rho:|\rho|\in[1,k]} \Big(d_\rho \operatorname{tr}\big(\widehat{p}(\rho)\rho(x)^T\big) + d_\rho \operatorname{tr}\big(\overline{\widehat{p}(\rho)\rho(x)^T}\big)\Big).$$

The expression inside the brackets is real and so $p'(x)$ is real as well.

Also, $\sum_x p'(x) = G \cdot \widehat{p'}(1) = G \cdot \widehat{p}(1) = 1$ by Lemma 4.

However, p' may be < 0 on some x. To remedy that, following previous work, we will "mix" p' with the uniform distribution so that the mixture becomes a distribution. Note that the mixture is k-uniform by Claim 15 as the low-degree non-trivial Fourier coefficients of both p' and uniform are zero.

Concretely, let

$$q := (1 - \beta)p' + \beta \frac{1}{H^m}.$$

for β to be determined. Note that q sums to 1 as p' and $1/H^m$ both do.

To determine β, first note that by definition

$$p'(x) \geq -\ell(x).$$

Crudely, $\ell(x) \leq (mH)^k \max_\rho d_\rho |\operatorname{tr}(\widehat{p}(\rho)\rho(x)^T)|$. Now, each absolute value is $\leq |\widehat{p}(\rho)|_2 |\rho(x)|_2$. The first term is ϵ/G by assumption. For the second, we use the fact that $\rho(x)$ is unitary, and so $|\rho(x)|_2 = |I|_2 = \sqrt{d_\rho}$. So $d_\rho |\operatorname{tr}(\widehat{p}(\rho)\rho(x)^T)| \leq d_\rho^{3/2}\epsilon/G \leq H^{3k/4}\epsilon/G$, using, in the last inequality, 2 over the underlying group H^k.

Hence, $|\ell(x)| \leq (mH)^{2k}\epsilon/G$ and we can set $\beta := (tH)^{2k}\epsilon/G$ and obtain that $q \geq 0$. As remarked earlier, q sums to 1, and so q is a distribution. It remains to bound the distance between p and q. We have

$$|p - q|_1 = \left|p - (1-\beta)p' + \beta/G\right|_1 \leq |\ell(x)|_1 + \beta|p'|_1 + \beta \cdot |1/G|_1.$$

The last two summands are β each. The first one is $\leq (mH)^{2k}\epsilon = \beta$ by the bound on $|\ell(x)|$ above. Hence the distance is $\leq 3\beta = 3(mH)^{2k}\epsilon$. □

References

[ACFN15] Anil Ada, Arkadev Chattopadhyay, Omar Fawzi, and Phuong Nguyen. The NOF multiparty communication complexity of composed functions. *Comput. Complex.*, 24(3):645–694, 2015. 2

[AGM03] Noga Alon, Oded Goldreich, and Yishay Mansour. Almost k-wise independence versus k-wise independence. *Inf. Process. Lett.*, 88(3):107–110, 2003. 3

[BC92] Michael Ben-Or and Richard Cleve. Computing algebraic formulas using a constant number of registers. *SIAM J. on Computing*, 21(1):54–58, 1992. 1

[BC08] Joshua Brody and Amit Chakrabarti. Sublinear communication protocols for multi-party pointer jumping and a related lower bound. In *25th Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 145–156, 2008. 2

[BCK+14] Harry Buhrman, Richard Cleve, Michal Kouckỳ, Bruno Loff, and Florian Speelman. Computing with a full memory: catalytic space. In *ACM Symp. on the Theory of Computing (STOC)*, pages 857–866, 2014. 2

[BDFP83] Allan Borodin, Danny Dolev, Faith E. Fich, and Wolfgang J. Paul. Bounds for width two branching programs. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 87–93, 1983. 2

[BDPW07] Paul Beame, Matei David, Toniann Pitassi, and Philipp Woelfel. Separating deterministic from nondeterministic nof multiparty communication complexity. In *34th Coll. on Automata, Languages and Programming (ICALP)*, pages 134–145. Springer, 2007. 2

[BGKL03] László Babai, Anna Gál, Peter G. Kimmel, and Satyanarayana V. Lokam. Communication complexity of simultaneous messages. *SIAM J. on Computing*, 33(1):137–166, 2003. 2

[BNP08] László Babai, Nikolay Nikolov, and László Pyber. Product growth and mixing in finite groups. In *ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 248–257, 2008. 3, 6

[BNS92] László Babai, Noam Nisan, and Márió Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *J. of Computer and System Sciences*, 45(2):204–232, 1992. 2

[CFL83] Ashok K. Chandra, Merrick L. Furst, and Richard J. Lipton. Multi-party protocols. In *15th ACM Symp. on the Theory of Computing (STOC)*, pages 94–99, 1983. 2

[CM20] James Cook and Ian Mertz. Catalytic approaches to the tree evaluation problem. In *STOC*, pages 752–760. ACM, 2020. 2

[Dia88] Persi Diaconis. *Group representations in probability and statistics*, volume 11 of *Institute of Mathematical Statistics Lecture Notes—Monograph Series*. Institute of Mathematical Statistics, Hayward, CA, 1988. 4, 5, 6

[DSV03] G. Davidoff, P. Sarnak, and A. Valette. *Elementary Number Theory, Group Theory and Ramanujan Graphs*. Elementary Number Theory, Group Theory, and Ramanujan Graphs. Cambridge University Press, 2003. 7

[DV23] Harm Derksen and Emanuele Viola. Quasirandom groups enjoy interleaved mixing. *Discrete Analysis*, 2023. 2, 3, 4

[Gow08] W. T. Gowers. Quasirandom groups. *Combinatorics, Probability & Computing*, 17(3):363–387, 2008. 2, 3, 6

[Gow17] W. T. Gowers. Generalizations of Fourier analysis, and how to apply them. *Bull. Amer. Math. Soc. (N.S.)*, 54(1):1–44, 2017. 4

[GV15] W. T. Gowers and Emanuele Viola. The communication complexity of interleaved group products. In *ACM Symp. on the Theory of Computing (STOC)*, 2015. 2

[GV19] W. T. Gowers and Emanuele Viola. Interleaved group products. *SIAM J. on Computing*, 48(3):554–580, 2019. Special issue of FOCS 2016. 2, 3, 4, 6, 8

[GV22] W. T. Gowers and Emanuele Viola. Mixing in non-quasirandom groups. In *ACM Innovations in Theoretical Computer Science conf. (ITCS)*, 2022. 4

[IL95] Neil Immerman and Susan Landau. The complexity of iterated multiplication. *Inf. Comput.*, 116(1):103–116, 1995. 1

[KLM23] Zander Kelley, Shachar Lovett, and Raghu Meka. Explicit separations between randomized and deterministic number-on-forehead communication. *CoRR*, abs/2308.12451, 2023. 2

[KMR66] Kenneth Krohn, W. D. Maurer, and John Rhodes. Realizing complex Boolean functions with simple groups. *Information and Control*, 9:190–195, 1966. 1

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997. 2

[LZ77] Richard J. Lipton and Yechezkel Zalcstein. Word problems solvable in logspace. *J. ACM*, 24(3):522–526, 1977. 1

[MC87] Pierre McKenzie and Stephen A. Cook. The parallel complexity of abelian permutation group problems. *SIAM J. Comput.*, 16(5):880–909, 1987. 1

[Mil14] Eric Miles. Iterated group products and leakage resilience against NC^1. In *ACM Innovations in Theoretical Computer Science conf. (ITCS)*, 2014. 1

[Mix89] David A. Mix Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1. *J. of Computer and System Sciences*, 38(1):150–164, 1989. 1, 2

[MV13] Eric Miles and Emanuele Viola. Shielding circuits with groups. In *ACM Symp. on the Theory of Computing (STOC)*, 2013. 2

[PRS97] Pavel Pudlák, Vojtěch Rödl, and Jiří Sgall. Boolean circuits, tensor ranks, and communication complexity. *SIAM J. on Computing*, 26(3):605–633, 1997. 2

[RX13] Ronitt Rubinfeld and Ning Xie. Robust characterizations of k-wise independence over product spaces and related testing results. *Random Struct. Algorithms*, 43(3):265–312, 2013. 3, 8

[RY19] Anup Rao and Amir Yehudayoff. *Communication complexity*. 2019. https://homes.cs.washington.edu/ anuprao/pubs/book.pdf. 2

[Ser77] Jean Pierre Serre. *Linear Representations of Finite Groups*. Springer, 1977. 4, 6

[Sha16] Aner Shalev. Mixing, communication complexity and conjectures of Gowers and Viola. *Combinatorics, Probability and Computing*, pages 1–13, 6 2016. arXiv:1601.00795. 2

[Ter99] Audrey Terras. *Fourier analysis on finite groups and applications*, volume 43 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1999. 4

[Vio19] Emanuele Viola. Non-abelian combinatorics and communication complexity. *SIGACT News, Complexity Theory Column*, 50(3), 2019. 2

[VW08] Emanuele Viola and Avi Wigderson. Norms, XOR lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4:137–168, 2008. 8

[Wig10] Avi Wigderson. Representation theory of finite groups, and applications. Available at http://www.math.ias.edu/∼avi/TALKS/Green_Wigderson_lecture.pdf, 2010. 4, 5

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il