



Linear Systems Over Finite Abelian Groups

Arkadev Chattopadhyay *
 University of Toronto, Toronto
 arkadev@cs.toronto.edu

Shachar Lovett †
 IAS, Princeton
 slovett@math.ias.edu

Abstract

We consider a system of linear constraints over any finite Abelian group G of the following form: $\ell_i(x_1, \dots, x_n) \equiv \ell_{i,1}x_1 + \dots + \ell_{i,n}x_n \in A_i$ for $i = 1, \dots, t$ and each $A_i \subset G$, $\ell_{i,j}$ is an element of G and x_i 's are Boolean variables. Our main result shows that the subset of the Boolean cube that satisfies these constraints has exponentially small correlation with the MOD_q boolean function, when the order of G and q are co-prime numbers.

Our work extends the recent result of Chattopadhyay and Wigderson (FOCS'09) who obtain such a correlation bound for linear systems over cyclic groups whose order is a product of two distinct primes or has at most one prime factor. Our result also immediately yields the first exponential bounds on the size of boolean depth-four circuits of the form $\text{MAJ} \circ \text{AND} \circ \text{ANY}_{O(1)} \circ \text{MOD}_m$ for computing the MOD_q function, when m, q are co-prime. No superpolynomial lower bounds were known for such circuits for computing any explicit function.

This completely solves an open problem posed by Beigel and Maciél (Complexity'97).

1 Introduction

A fundamental open problem in theoretical computer science is to understand the computational power of counting modulo composite numbers. For example, we do not know if hard problems like SATISFIABILITY have efficient depth-three circuits comprising only MOD_6 gates. This is in contrast to the classical theorem of Razborov [Raz87] and Smolensky [Smo87] that says constant-depth circuits having only AND, OR and MOD_m gates cannot compute even the MOD_q function in sub-exponential size, i.e. in size $2^{n^{o(1)}}$, when m, q are co-prime and when m has only one prime factor. Smolensky [Smo87] conjectured that this theorem extends to all m , but despite a series of attempts over two decades, this conjecture remains wide open. While Smolensky's conjecture easily implies that not all functions computable in deterministic linear time have efficient ACC^0 circuits, the best one can prove is the recent breakthrough result of Williams [Wil11] who showed that non-deterministic exponential time does not have efficient ACC^0 circuits.

A recent result of Hansen and Koucky [HK09] shows that every function in $\text{ACC}^0[m]$ is in fact computed by a quasipolynomial size circuit of the canonical form $\text{OR} \circ \text{AND} \circ \text{CC}^0[m]$, where $\text{CC}^0[m]$ refers to the class of constant-depth circuits just comprising MOD_m gates. Given this characterization, a natural first step towards proving Smolensky's conjecture is to verify it for depth-three circuits $\text{OR} \circ \text{AND} \circ \text{MOD}_m$ with just one layer of MOD_m gates at the base. However,

*supported partially by a NSERC postdoctoral fellowship, an Ontario Ministry of Innovation fellowship and NSERC discovery grants 482686 and 480508.

†Research supported by NSF grant DMS-0835373.

this step was long identified by Beigel and Maciel [BM97] as a barrier. They observed that there are no known techniques to prove strong lower bounds on the size of such depth-three circuits when the MOD_m gates at the base are generalized in the following sense: each such gate has an associated *accepting set* A and the gates output 1 iff the sum of the input bits evaluates to an element in A modulo m . Interestingly, if each gate at the base had a singleton accepting set then they could prove very strong lower bounds, but their methods failed for general accepting sets. The problem of handling such general accepting sets is not specific to their work but is well known to researchers. For example, it is not hard to show that depth-two circuits having MOD_m gates cannot compute even the AND function in sub-exponential size when the output gate has a singleton accepting set, while it is consistent with our current knowledge that such circuits in linear size compute SATISFIABILITY for an appropriate choice of accepting set for the output gate (see Caussinus [Cau96]).

It is known that the choice of accepting sets makes a non-trivial difference in the closely related world of polynomial representation of boolean functions. For example, polynomials, over the ring \mathbb{Z}_m , need $\Omega(n)$ degree to compute simple functions like AND, OR and MOD_q when the accepting set is a singleton. However, Barrington, Beigel and Rudich [BBR94] gave an elegant and surprising construction for computing AND and OR with polynomials of degree $O(n^{1/t})$ having proper accepting sets, where m has t distinct prime factors. Hansen [Han06] showed that judicious choice of accepting sets affords similar advantage for computing MOD_q . No superlogarithmic lower bound on the degree of polynomials with general accepting sets is known for computing any function in NP. On the positive side, the construction of Barrington *et al* has led to other interesting constructions outside of circuit complexity. For example, all known bounds on explicit constructions of Ramsey graphs can be achieved using it [Gop06]. Further, a series of recent breakthroughs in constructing more efficient locally decodable codes [Efr09, DGY10] have crucially relied on the construction of Barrington *et al*.

Recently, Chattopadhyay and Wigderson [CW09] attacked this depth-three question by naturally considering a system of linear constraints of the form $\ell_i \in A_i$ for $i = 1, \dots, N$, where each $A_i \subset \mathbb{Z}_m$ and ℓ_i 's are linear forms. Their main result gives an exponentially small upper bound on the correlation of the Boolean solution set of any such system with the MOD_q function, when m is either a prime power or is a product of two distinct primes like 6. This implied the first exponential lower bounds on the size of depth-three circuits of the form $\text{MAJ} \circ \text{AND} \circ \text{MOD}_m^A$ for computing MOD_q , for such m .

Our Work: We extend the result of Chattopadhyay and Wigderson to arbitrary, fixed m . More generally, our main result is the following: for any system of linear constraints \mathcal{L} , let $B_{\mathcal{L}}$ be the set of points in the Boolean cube that satisfies \mathcal{L} . The *correlation* of a set S of boolean points with the MOD_q function, denoted by $\text{Corr}(S, \text{MOD}_q)$, is defined as $\max_{a,b} |\Pr_x[x \in S \wedge \text{MOD}_q^{\{a\}}(x) = 1] - \Pr_x[x \in S \wedge \text{MOD}_q^{\{b\}}(x) = 1]|$.

Theorem 1 (Main Theorem). *Let m be an arbitrary fixed positive integer and \mathcal{L} be a system of linear constraints over n variables of the following form: $\ell_i(x_1, \dots, x_n) \in A_i$ for $i = 1, \dots, t$, where each $A_i \subset \mathbb{Z}_m$ and ℓ_i is a linear form over \mathbb{Z}_m . Then, $\text{Corr}(B_{\mathcal{L}}, \text{MOD}_q) \leq 2^{-\Omega(n)}$ when m and q are co-prime.*

A direct consequence of Theorem 1, obtained by an easy application of the so call ϵ -discriminator Lemma of Hajnal *et al* [HMP⁺93] (restated equivalently in Section 2 of this article), is the following

exponential lower bound on the size of boolean circuits:

Corollary 2. *Let GMOD_m denote mod- m gates with general accepting sets. Then, depth-four circuits of the form $\text{MAJ} \circ \text{AND} \circ \text{GMOD}_m$ require exponential fan-in at the output Majority gate to compute the MOD_q function, if m, q are co-prime.*

We can generalize Theorem 1 to arbitrary fixed Abelian groups.

Theorem 3. *Let G be any finite and fixed Abelian group and \mathcal{L} be a system of linear constraints with n boolean variables where the coefficients of each constraint are elements of G . Then, $\text{Corr}(B_{\mathcal{L}}, \text{MOD}_q) \leq 2^{-\Omega(n)}$ when the order of G and q are co-prime.*

Applying Theorem 3 to the group $G = \mathbb{Z}_m^{O(1)}$, allows us to generalize Corollary 2 by replacing GMOD_m gates by an arbitrary composition of a fixed number of MOD_m gates (or equivalently, a fixed number of GMOD_m gates).

Corollary 4. *Depth-four circuits of the form $\text{MAJ} \circ \text{AND} \circ \text{ANY}_{O(1)} \circ \text{MOD}_m$ require exponential fan-in at the output Majority gate to compute the MOD_q function, if m, q are co-prime.*

Proof. Note that for any $s = O(1)$, $\text{ANY}_s \circ \text{MOD}_m \subset \text{GMOD}_{(\mathbb{Z}_m)^s}$. Let $G = (\mathbb{Z}_m)^s$. Theorem 3 gives that the MOD_q has exponentially small correlation with any function in $\text{AND} \circ \text{GMOD}_G$. Hence $\text{MAJ} \circ \text{AND} \circ \text{GMOD}_G$ circuits that compute the MOD_q function require exponential size. \square

Beigel and Maciel [BM97] identified the problem of proving lower bounds for depth-three circuits of the form $\text{MAJ} \circ \text{AND} \circ \text{GMOD}_m$ as an important next step towards understanding circuits having modular gates. Corollary 2 completely solves this problem by obtaining the first strong lower bounds for such circuits.

In the language of Barrington and Thérien [BT87], the result of Hansen and Koucky implies that one way of proving Smolensky’s conjecture is to show that functions computed by systems of programs over finite solvable groups do not correlate well with the MOD_q function if q is co-prime with the order of the group. Our result takes the first step in this direction by verifying this for Abelian groups.

Our Technique In the world of arithmetic circuits, Grigoriev and Razborov [GR98] introduced the ingenious notion of communication rank for linear systems over a finite fields. Chattopadhyay and Wigderson [CW09] generalized this notion to systems over \mathbb{Z}_m for an arbitrary composite m . Using this notion, [CW09] showed that if a system has high rank then it is highly unsatisfiable over the boolean cube and if it has low rank, exploiting estimates of exponential sums by Bourgain, they showed that the correlation of the solution set to the MOD_q function is small. For technical reasons, their analysis of the low rank case only worked if m was a product of two distinct primes or had just one prime factor.

In this work, we realize that in order to work with arbitrary composites, it is convenient to consider more general systems of linear constraints. We consider constraints in which the accepting set is itself a function of a constant number of variables as opposed to being a constant as in the work of Chattopadhyay and Wigderson. This leads us to further generalize the notion of communication rank to facilitate analysis of such linear systems. In particular, we invent an iterative simplification of linear systems where this new notion of rank plays a crucial role. This simplification, driven by our Lemma 9 in Section 3, is the key new ingredient of our work that allows us to work with

arbitrary modulus m . A rough description of the main idea is as follows: either our system has large communication rank in which case it is highly unsatisfiable or it has low rank in which case we simplify in the following sense. Each constraint in the simplified system has either a singleton accepting set or the number of variables on which the accepting set depends is one less than before or the system is over a modulus m' that is less than m . A repeated application of this procedure yields a nice structural result: every generalized linear system L over \mathbb{Z}_m over n variables can be decomposed into at most $t = 2^{\epsilon n}$ linear systems L_1, \dots, L_t where each L_i is either satisfied by an exponentially small fraction of the points in the boolean cube or L_i is the intersection of L_i^0 and L_i^1 where every constraint in L_i^0 has a singleton accepting set and each constraint of L_i^1 corresponds to a $k = k(m)$ -junta.

It is already known from the work of Chattopadhyay and Wigderson, restated in Lemma 8 of this work, that subsystems that are intersections of singleton systems over \mathbb{Z}_m and junta systems have exponentially small correlation with MOD_q . The subsystems of our decomposition that have poor satisfiability cannot, by definition, correlate with a much more balanced function like MOD_q . Since there are only few subsystems in the decomposition, an easy application of the union bound finishes the argument.

Paper organization We give basic definitions and recall the necessary background in Section 2. We prove our result for arbitrary cyclic groups \mathbb{Z}_m in Section 3. We highlight the required generalization for arbitrary Abelian groups in Appendix A. We prove some technical generalizations of known lemmas for Abelian groups in Appendix B.

2 Preliminaries

Let $\mathbb{Z}_q := \{0, \dots, q-1\}$ and $\mathbb{N} = \{0, 1, 2, \dots\}$ denote the natural numbers. We study the correlation of subsets $S \subset \{0, 1\}^N$ with the sum modulo q . It is natural and convenient to estimate this quantity using the q -th roots of unity. Let $e_q(y) := \exp(2\pi iy/q)$, where i denotes the complex square-root of unity. We will use in the paper the following definition for correlation:

$$\text{Corr}(S, \text{MOD}_q) := \max_{b \in \mathbb{Z}_q \setminus \{0\}} \left| \mathbb{E}_{x \in \{0,1\}^n} [1_S(x) \cdot e_q(b(x_1 + \dots + x_n))] \right|,$$

where 1_S is the indicator function of S . It is straightforward to verify that

$$\max_{a, b \in \mathbb{Z}_q} \left| \Pr_x [x \in S \wedge \text{MOD}_q^{\{a\}}(x) = 1] - \Pr_x [x \in S \wedge \text{MOD}_q^{\{b\}}(x) = 1] \right| \leq 2 \cdot \text{Corr}(S, \text{MOD}_q),$$

so our definition indeed captures the more intuitive definition of having elements of S being approximately equidistributed modulo q . For a family of subsets $\mathcal{S}_N = \{S \subset \{0, 1\}^N\}$ we define their correlation with sums modulo q as the maximal correlation for $S \in \mathcal{S}_N$.

The simple tool that we use for lower bounding the size of our circuits for computing MOD_q is the so-called ϵ -Discriminator Lemma, introduced by Hajnal *et al* [HMP⁺93]. We state here a specialized version of it that is particularly convenient for our work, and has been also used in earlier works (see for example [Bou05, CGPT06]).

Lemma 5 (Discriminator Lemma). *Let C be a circuit that has a MAJORITY gate at its output that is being fed by t subcircuits C_1, \dots, C_t . If C computes MOD_q , then there exists a subcircuit C_i , such that $\text{Corr}(C_i, \text{MOD}_q) = \Omega(1/t)$.*

In order to estimate the correlation of the solution sets of linear systems with the MOD_q function, we will need estimates of exponential sums that were first obtained in the work of Bourgain [Bou05] and refined progressively in further works [GRS05, VW08, Cha07]. We state the most refined estimate below:

Theorem 6 ([Cha07]). *Let m, q be two fixed positive co-prime integers and let P be any n -variate multilinear polynomial of degree d with coefficients in \mathbb{Z}_m and b be any number non-zero modulo q . Then, there exists a constant $\beta = \beta(m, q)$ such that the following holds:*

$$\left| \mathbb{E}_{x \in \{0,1\}^n} \left[e_m(P(x)) e_q \left(b \sum_i x_i \right) \right] \right| \leq \exp(-\beta^d n). \quad (1)$$

We point out that the above estimate fails to give anything non-trivial when the degree d of the polynomial P is more than $\log n$. Finding exponentially small upper bounds for the exponential sum in (1) for $d > \log n$, even when m is prime, remains a very interesting open problem.

3 Linear systems over cyclic groups

We study in this section systems of linear equations with arbitrary accepting sets over arbitrary (constant) moduli.

Definition 1 (Linear system with accepting sets). *A linear equation in n binary variables over \mathbb{Z}_m with an accepting set is the set of solutions (over $\{0,1\}^n$) to an equation of the form*

$$E = \{x \in \{0,1\}^n : \sum a_i x_i \pmod{m} \in A\},$$

where $a_1, \dots, a_n \in \mathbb{Z}_m$ and $A \subsetneq \mathbb{Z}_m$.

A linear system in n binary variables over \mathbb{Z}_m with accepting sets is the set of common solutions to several such equations, i.e. it is a subset of $\{0,1\}^n$ of the form

$$\begin{aligned} L &= E_1 \cap \dots \cap E_N \\ &= \{x \in \{0,1\}^n : \sum a_{i,j} x_j \pmod{m} \in A_i \quad \forall i = 1, \dots, N\}, \end{aligned}$$

where $a_{i,j} \in \mathbb{Z}_m$ and $A_i \subsetneq \mathbb{Z}_m$. We denote by $\mathcal{L}_n(m)$ the family of all such linear systems (where we do not distinguish the number of equations).

We aim to bound the correlation of the solution set of linear systems with the MOD_q function. Our main approach is to iteratively simplify the system. In order for us to define these simplifications, we need some further definitions of more general systems of linear forms. We first define the special case of a linear system all of whose accepting sets are singletons, i.e. consist of a single value.

Definition 2 (Singleton linear systems). *A linear equation in n binary variables over \mathbb{Z}_m with a single accepting value is the set of solutions (over $\{0,1\}^n$) to an equation of the form*

$$E = \{x \in \{0,1\}^n : \sum a_i x_i \equiv b \pmod{m}\},$$

where $a_1, \dots, a_n, b \in \mathbb{Z}_m$.

A singleton linear system in n binary variables over \mathbb{Z}_m is the set of common solutions to several such equations, i.e. it is a subset of $\{0, 1\}^n$ of the form

$$\begin{aligned} L &= E_1 \cap \dots \cap E_N \\ &= \{x \in \{0, 1\}^n : \sum a_{i,j} x_j \equiv b_i \pmod{m} \quad \forall i = 1, \dots, N\}, \end{aligned}$$

where $a_{i,j}, b_j \in \mathbb{Z}_m$. We denote by $\mathcal{L}_n^{Sing}(m)$ the family of all such linear systems (where we do not distinguish the number of equations).

We will also need the following generalized notions of linear systems with accepting sets which depend on a few variables in an arbitrary manner.

Definition 3 (Linear systems with accepting sets of sparsity k). A linear equation in n binary variables over \mathbb{Z}_m with an accepting set of sparsity k is the set of solutions (over $\{0, 1\}^n$) to a linear equation with an accepting set which depends on k of the variables, i.e. to an equation of the form

$$E = \{x \in \{0, 1\}^n : \sum a_i x_i \pmod{m} \in A(x_{i_1}, \dots, x_{i_k})\},$$

where $a_1, \dots, a_n \in \mathbb{Z}_m$, $i_1, \dots, i_k \in [n]$ and each set $A(x_{i_1}, \dots, x_{i_k})$ is a subset of \mathbb{Z}_m for every setting of $x_{i_1}, \dots, x_{i_k} \in \{0, 1\}^k$. Moreover, we require that the accepting-set function A is not trivial, i.e. $A(x_{i_1}, \dots, x_{i_k}) \subsetneq \mathbb{Z}_m$ for at least one setting of x_{i_1}, \dots, x_{i_k} .

A linear system in n binary variables over \mathbb{Z}_m with accepting sets of sparsity k is the set of common solutions to several such equations, i.e. it is a subset of $\{0, 1\}^n$ of the form

$$\begin{aligned} L &= E_1 \cap \dots \cap E_N \\ &= \{x \in \{0, 1\}^n : \sum a_{i,j} x_j \pmod{m} \in A_i(x_{i_1}, \dots, x_{i_k}) \quad \forall i = 1, \dots, N\}, \end{aligned}$$

where $a_{i,j}, b_j \in \mathbb{Z}_m$, and each accepting set-function A_i is not trivial. We denote by $\mathcal{L}_n(m, k)$ the family of all such linear systems (where we do not distinguish the number of equations in the system).

Note that $\mathcal{L}_n(m, 0) = \mathcal{L}_n(m)$. For $k > 1$, we also allow modulus $m = 1$, in which case we interpret each equation in $\mathcal{L}_n(1, k)$ as $E = \{x : 0 \in A(x_{i_1}, \dots, x_{i_k})\}$, where set function A is not trivial. That is, the equation is a k -junta. The linear system is the set of common solutions to several such equations.

We now define the most general linear system, which will be simplified iteratively in the proof of our main theorem. It will be the intersection of systems over several moduli ℓ which divide m . Let $\ell \div m$ denote " ℓ divides m ". We define linear systems which are intersections of linear systems in $\mathcal{L}_n^{Sing}(m)$ and $\mathcal{L}_n(\ell, k)$ for several $\ell \div m$. For a modulus m define $\text{div}(m) = \{1 \leq \ell \leq m : \ell \div m\}$ to be the set of (not necessarily prime) factors of m . We will maintain a sparsity function $\kappa : \text{div}(m) \rightarrow \mathbb{N} \cup \{-\infty\}$ which will specify the allowed sparsity for each $\ell \div m$. That is, we will have $\mathcal{L}_n(\ell, \kappa(\ell))$ systems for all $\ell \in \text{div}(m)$ such that $\kappa(\ell) \geq 0$, where $\kappa(\ell) = -\infty$ means we have no $\mathcal{L}_n(\ell, \cdot)$ system. Note that if $\ell \div m$ and $\kappa : \text{div}(\ell) \rightarrow \mathbb{N} \cup \{-\infty\}$ then $\mathcal{L}_n(\ell, \kappa(\cdot)) \subset \mathcal{L}_n(m, \kappa(\cdot))$ since equations modulo ℓ can always be lifted to equations modulo m by multiplying them by m/ℓ .

Definition 4 (Linear systems with general accepting sets over several moduli). Let m be a modulus and let $\kappa : \text{div}(m) \rightarrow \mathbb{N} \cup \{-\infty\}$. We define $\mathcal{L}_n(m, \kappa(\cdot))$ as follows: $L \in \mathcal{L}_n(m, \kappa(\cdot))$ if there exists

$L^{Sing} \in \mathcal{L}_n^{Sing}(m)$ and $L^\ell \in \mathcal{L}_n(\ell, \kappa(\ell))$ for all $\ell \in \text{div}(m)$ such that $\kappa(\ell) \geq 0$ and

$$L = L^{Sing} \cap \bigcap_{\ell \in \text{div}(m): \kappa(\ell) \geq 0} L^\ell.$$

Theorem 1 follows from the following theorem for κ defined as $\kappa(m) = 0$ and $\kappa(\ell) = -\infty$ for all $\ell \in \text{div}(m) \setminus \{m\}$.

Theorem 7 (Correlation bound for $\mathcal{L}_n(m, \kappa(\cdot))$ and MOD_q). *Let m, q be co-prime and $\kappa : \text{div}(m) \rightarrow \mathbb{N} \cup \{-\infty\}$. Let $L \in \mathcal{L}_n(m, \kappa(\cdot))$ be any linear system. Then*

$$\text{Corr}(L, \text{MOD}_q) \leq \exp(-n/c),$$

where $c = c_7(m, q, \kappa(\cdot))$. Crucially, c does not depend on n .

The proof of Theorem 7 follows from induction over $\kappa(\cdot)$. The following two Lemmas specify the base case and the inductive step.

Lemma 8 (Base case). *Let m, q be co-prime and let $k \geq 0$ be a sparsity. Let $L = L' \cap L''$ where $L' \in \mathcal{L}_n^{Sing}(m)$ and $L'' \in \mathcal{L}_n(1, k)$. Then*

$$\text{Corr}(L, \text{MOD}_q) \leq \exp(-n/\beta^k),$$

where $\beta = \beta(m, k)$ is as given in Theorem 6.

The Lemma above is implicit in the work of Chattopadhyay and Wigderson and points out why we call such linear systems simple. It is obvious that if we could decompose a given linear system into unions of a few such simple systems, we would obtain our desired correlation bounds by the union bound. The next lemma, the main inductive step, roughly shows that the only obstacle from having such a nice decomposition is the existence of subsystems that are satisfied by an exponentially small fraction of the points of the cube.

Lemma 9 (Simplification process for $\mathcal{L}_n(m, k)$). *For any $m, k \geq 0$, there exists $c = c_9(m, k)$ such that for any $L \in \mathcal{L}_n(m, k)$ and any $1 \leq r \leq n$, one of the following must hold:*

1. $\Pr_{x \in \{0,1\}^n} [x \in L] \leq \exp(-r/c)$.
2. There exist $L_1, \dots, L_R \in \mathcal{L}_n(m, \kappa(\cdot))$ such that $L = L_1 \cup \dots \cup L_R$ for $R \leq \exp(cr)$, and $\kappa : \text{div}(m) \rightarrow \mathbb{N} \cup \{-\infty\}$ is given as
 - (a) If $k > 0$ then $\kappa(m) = k - 1$ and $\kappa(\ell) = k + m \log m$ for all $\ell \in \text{div}(m) \setminus \{m\}$.
 - (b) If $k = 0$ then $\kappa(m) = -\infty$ and $\kappa(\ell) = k + m \log m$ for all $\ell \in \text{div}(m) \setminus \{m\}$.

We first prove Theorem 7 given Lemmas 8 and 9. We then proceed to prove Lemmas 8 and 9.

Proof of Theorem 7 given Lemmas 8 and 9. Define a lexicographic order on $\kappa : \text{div}(m) \rightarrow \mathbb{N} \cup \{-\infty\}$: $\kappa > \kappa'$ if there exists $\ell_0 \in \text{div}(m)$ such that $\kappa(\ell) = \kappa'(\ell)$ for all $\ell > \ell_0$ and $\kappa(\ell_0) > \kappa'(\ell_0)$ ¹.

¹As \mathbb{N}^d is a well founded set for all $d \geq 1$ this defines a proper Noetherian induction. An explicit bound can be derived using the explicit bounds on the growth of $\kappa(\cdot)$ given by Lemma 9.

The base case of $\kappa(\ell) = -\infty$ for all $\ell > 1$ is given by Lemma 8. For the inductive step, let $\ell_{\max} > 1$ be maximal such that $\kappa(\ell_{\max}) \geq 0$. Let $L \in \mathcal{L}_n(m, \kappa(\cdot))$. Then

$$L = L^{Sing} \cap \bigcap_{\ell \in \text{div}(m): \kappa(\ell) \geq 0} L^\ell,$$

where $L^{Sing} \in \mathcal{L}_n^{Sing}(m)$ and $L^\ell \in \mathcal{L}_n(\ell, \kappa(\ell))$. Apply Lemma 9 for $L^{\ell_{\max}}$. Let $c = c_9(\ell_{\max}, \kappa(\ell_{\max}))$ and let $r = n/c^*$ be a parameter to be determined later. One of the following must hold:

1. $\Pr_{x \in \{0,1\}^n} [x \in L^{\ell_{\max}}] \leq \exp(-r/c) = \exp(-n/(cc^*))$. Hence, $\Pr_{x \in \{0,1\}^n} [x \in L] \leq \exp(-n/(cc^*))$ and $\text{Corr}(L, \text{MOD}_q) \leq \exp(-n/(cc^*))$.
2. There exists $L_1^{\ell_{\max}}, \dots, L_R^{\ell_{\max}} \in \mathcal{L}_n(\ell_{\max}, \kappa_1(\cdot)) \subset \mathcal{L}_n(m, \kappa_1(\cdot))$ such that $L^{\ell_{\max}} = L_1^{\ell_{\max}} \cup \dots \cup L_R^{\ell_{\max}}$ for $R \leq \exp(cr) = \exp(n(c/c^*))$, where $\kappa_1 < \kappa$ is given by $\kappa_1(\ell_{\max}) < \kappa(\ell_{\max})$ and $\kappa_1(\ell) = \kappa(\ell_{\max}) + \ell_{\max} \log \ell_{\max}$ for $\ell < \ell_{\max}$. Define

$$L_i := L^{Sing} \cap \bigcap_{\ell \in \text{div}(m): \ell < \ell_{\max}, \kappa(\ell) \geq 0} L^\ell \cap L_i^{\ell_{\max}}.$$

We have $L = L_1 \cup \dots \cup L_R$ where $L_i \in \mathcal{L}_n(m, \kappa'(\cdot))$, with $\kappa'(\cdot)$ defined as $\kappa'(\ell_{\max}) = \kappa_1(\ell_{\max})$ and $\kappa'(\ell) = \max(\kappa(\ell), \kappa_1(\ell))$ for $\ell < \ell_{\max}$. Note that $\kappa' < \kappa$, so we can apply the induction hypothesis for L_1, \dots, L_R :

$$\begin{aligned} \text{Corr}(L, \text{MOD}_q) &= \max_{b \in \mathbb{Z}_q \setminus \{0\}} \left| \mathbb{E}_{x \in \{0,1\}^n} [1_L(x) \cdot \omega_q(b(x_1 + \dots + x_n))] \right| \\ &= \max_{b \in \mathbb{Z}_q \setminus \{0\}} \left| \sum_{i=1}^R \mathbb{E}_{x \in \{0,1\}^n} [1_{L_i}(x) \cdot \omega_q(b(x_1 + \dots + x_n))] \right| \\ &\leq \sum_{i=1}^R \max_{b \in \mathbb{Z}_q \setminus \{0\}} \left| \mathbb{E}_{x \in \{0,1\}^n} [1_{L_i}(x) \cdot \omega_q(b(x_1 + \dots + x_n))] \right| \\ &\leq |R| \cdot \text{Corr}(\mathcal{L}_n(m, \kappa'(\cdot)), \text{MOD}_q) \\ &\leq \exp((c/c^* - 1/c_7(m, \kappa'(\cdot)))n), \end{aligned}$$

where crucially we used the fact that L_1, \dots, L_R are disjoint.

Setting c^* to be a large enough constant (say $c^* = 2c \cdot c_7(m, \kappa'(\cdot))$) concludes the proof. \square

3.1 Proof of base case: Lemma 8

Our argument here essentially is taken from [CW09]. Let $L = L' \cap L''$ with $L' \in \mathcal{L}_n^{Sing}(m)$ and $L'' \in \mathcal{L}_n(1, k)$. That is,

$$\begin{aligned} L' &= \{x \in \{0,1\}^n : \sum a_{i,j} x_j \equiv b_i \pmod{m} \quad \forall i = 1, \dots, N'\} \\ L'' &= \{x \in \{0,1\}^n : 0 \in A_i(x_{v(i,1)}, \dots, x_{v(i,k)}) \quad \forall i = 1, \dots, N''\}, \end{aligned}$$

where $a_{i,j}, b_i \in \mathbb{Z}_m$, $v(i, j) \in [n]$ and $A_i(z_1, \dots, z_k) \subset \mathbb{Z}_m$.

Define $P_i(x) := \sum a_{i,j}x_j - b_i$ to be linear functions over \mathbb{Z}_m for $i \in [N']$ so that

$$1_{L'}(x) = \prod_{i=1}^{N'} 1_{P_i(x)=0}.$$

Define $Q_i(x)$ to be polynomials over \mathbb{Z}_m of degree at most k such that $Q_i(x) = 0$ iff $0 \in A_i(x_{v(i,1)}, \dots, x_{v(i,k)})$, so that

$$1_{L''}(x) = \prod_{i=1}^{N''} 1_{Q_i(x)=0}.$$

Using the fact that for $z \in \mathbb{Z}_m$ we have $1_{z=0} = \frac{1}{m} \sum_{a=0}^{m-1} e_m(a \cdot z)$ we get

$$\begin{aligned} 1_L(x) &= \left(\prod_{i=1}^{N'} \frac{1}{m} \sum_{a=0}^{m-1} e_m(a \cdot P_i(x)) \right) \times \left(\prod_{j=1}^{N''} \frac{1}{m} \sum_{b=0}^{m-1} e_m(b \cdot Q_j(x)) \right) \\ &= \frac{1}{m^{N'+N''}} \sum e_m \left(\sum_{i=1}^{N'} a_i \cdot P_i(x) + \sum_{j=1}^{N''} b_j \cdot Q_j(x) \right). \end{aligned}$$

where the last summation is over $a_1, \dots, a_{N'}, b_1, \dots, b_{N''} \in \mathbb{Z}_m$. The bound for the correlation between L and MOD_q now follows from Theorem 6 since all $P_i(x), Q_j(x)$ are polynomials of degree at most k , and so are all linear combinations of them.

3.2 Proof of inductive step: Lemma 9

We define a notion of rank of $L \in \mathcal{L}_n(m, k)$ which is appropriate for our purposes and whose origins lie in the elegant work of Grigoriev and Razborov [GR98]. Chattopadhyay and Wigderson [CW09] generalized the Grigoriev-Razborov notion to deal with linear systems of type $\mathcal{L}_n(m)$. We further generalize it to deal with systems in $\mathcal{L}_n(m, k)$, where k is a constant non-negative integer. For the sake of consistency with earlier work, we call this notion the *communication rank* of the linear system. Fix some equations E_1, \dots, E_N of sparsity k such that $L = E_1 \cap \dots \cap E_N$, where each E_i is given by

$$E_i = \{x \in \{0, 1\}^n : \sum a_{i,j}x_j \pmod{m} \in A_i(x_{v(i,1)}, \dots, x_{v(i,k)})\}.$$

The definition of communication rank will in fact depend on the specific E_1, \dots, E_N chosen.

Definition 5 (Communication Rank). *Let $L \in \mathcal{L}_n(m, k)$ given by $L = E_1 \cap \dots \cap E_N$. We say that a subset of equations $I = \{i_1, \dots, i_r\} \subset [N]$ is s -wise independent if the following conditions hold. Let $V_i = \{v(i, 1), \dots, v(i, k)\}$ for $i \in I$ be the set of variables on which A_i depends. We first require that all sets V_{i_1}, \dots, V_{i_r} be disjoint. We also require that there exist subsets of variables $J_{p,t} \subset [n]$ of size $|J_{p,t}| = |I| = r$, where p ranges over the distinct prime factors of m and $t = 1, \dots, s$, such that:*

1. All sets $J_{p,t}$ and V_{i_1}, \dots, V_{i_r} are pair-wise disjoint.
2. Let $M_{p,t}$ be the following $r \times r$ matrix over \mathbb{F}_p : if $J_{p,t} = \{j_1, \dots, j_r\}$ then the (x, y) -entry of $M_{p,t}$ is given by a_{i_x, j_y} (modulo p), i.e. $M_{p,t}$ is the $r \times r$ minor given by the rows of I and the columns of $J_{p,t}$. We require that for any prime factor p of m , and any $t = 1, \dots, s$, the matrix $M_{p,t}$ has full rank modulo p .

The s -wise communication rank of L modulo m , denoted by $\text{ccrank}_s^m(L)$, is the maximal r for which this holds for some $I \subset [N]$ of size $|I| = r$.

If $k = 0$, then the above definition exactly corresponds to the notion of communication rank used by Chattopadhyay and Wigderson. We use the following result that follows from their work:

Lemma 10 (Implicit in Chattopadhyay-Wigderson [CW09]). *Let $L \in \mathcal{L}_n(m)$ have ccrank_m^m communication rank at least r . Then*

$$\Pr_{x \in \{0,1\}^n} \left[\bigwedge_{i=1}^N \ell_i(x) \in A_i \right] \leq \exp(-r/c_{10}(m)),$$

where each $A_i \subsetneq \mathbb{Z}_m$ is an arbitrary set.

Remark 1. *This lemma appears in [CW09] with the restriction that m has no repeated prime factors. We show in the appendix that this restriction can be lifted by a slight modification of the argument in [CW09]. In fact, we generalize it to all Abelian groups.*

We next show an easy corollary of the above lemma for systems in $\mathcal{L}_n(m, k)$, when $k > 0$.

Lemma 11. *Let $L \in \mathcal{L}_n(m, k)$ have m -wise communication rank at least r . Then*

$$\Pr_{x \in \{0,1\}^n} [x \in L] \leq \exp(-r/c_{11}(m, k)).$$

Proof. Let $I = \{i_1, \dots, i_r\}$ be a set of indices corresponding to independent equations. We focus entirely on the sub-system indexed by this set. Since the sets V_{i_1}, \dots, V_{i_r} are disjoint, we can sample all $x_j \in \cup_{i \in I} V_i$ and guarantee, by the Chernoff bound, that with probability at least $1 - \exp(-r/2^k)$ we will get $\Omega(r/2^k)$ non-trivial accepting sets left. Thus, after sampling, we are left with an ordinary sub-system in $\mathcal{L}_{n'}(m)$ whose rank is $\Omega(r/2^k)$, where $n' \geq n - rk$. Applying Lemma 10 to this sub-system, the argument follows by setting $c_9(m, k) = O(2^k c_8(m))$. \square

Lemma 11 shows that if the linear system has high communication rank, then its correlation with MOD_q is small as the size of the solution set is a very small fraction of the boolean cube. We next deal with the complementary case, where the communication rank is small.

We start off by a convenient structural result about such systems, stated in Chattopadhyay and Wigderson that generalizes a lemma of Grigoriev and Razborov.

Lemma 12 (Restatement of Lemma 13 of [CW09]). *Let m have w distinct prime factors p_1, \dots, p_w . Consider an ordinary linear system $L \in \mathcal{L}_n(m)$ such that $\text{ccrank}_s^m(L) = r$. Then, there exists a set I of r linear forms satisfying the following condition: for every linear form ℓ in L , there exists a prime p_j such that $\ell \equiv \ell_I + \ell_0 \pmod{p_j}$, where ℓ_I is in the \mathbb{Z}_{p_j} -linear span of I and ℓ_0 is ws -sparse.*

In our formulations we take the set I for all primes. Thus $|I| = r$ and not $(ws + 1)r$

We will also need the following simple claim.

Claim 13. *Let $\ell(x_1, \dots, x_n) \in A \pmod{m}$ be any linear constraint. Let p be a prime factor of m , and assume that $\ell \equiv \ell_I + \ell_0 \pmod{p}$ where ℓ_0 is supported on variables x_{i_1}, \dots, x_{i_k} . For $a \in \mathbb{Z}_m$ let $B_a = \{x \in \{0, 1\}^n : \ell_I(x) \equiv a \pmod{m}\}$. Then for every $a \in \mathbb{Z}_m$ and values for $x_{i_1}, \dots, x_{i_k} \in \{0, 1\}$, there exists a linear form $\ell'(x)$ and a set $A_a(x_{i_1}, \dots, x_{i_k}) \subset \mathbb{Z}_m/p$ such that*

$$\begin{aligned} B_a \cap \{x \in \{0, 1\}^n : \ell(x) \pmod{m} \in A\} \\ = B_a \cap \{x \in \{0, 1\}^n : \ell'(x) \pmod{m/p} \in A_a(x_{i_1}, \dots, x_{i_k})\}, \end{aligned}$$

Proof. Let $S = \{x_1, \dots, x_k\}$ be the variables in the support of ℓ_0 , and $T = [n] \setminus S$ be the remaining variables. For $x \in \{0, 1\}^n$ let $x^S \in \{0, 1\}^S$ and $x^T \in \{0, 1\}^T$ denote its restriction to the corresponding variables sets. Note that $\ell_0(x) = \ell_0(x^S)$. Partition $\ell(x) = \ell^S(x^S) + \ell^T(x^T)$ to a linear form over x^S and a linear form over x^T , and similarly $\ell_I(x) = \ell_I^S(x^S) + \ell_I^T(x^T)$. Note that by assumption $\ell^T \equiv \ell_I^T \pmod{p}$, hence all the coefficients of $\ell^T - \ell_I^T$ are divisible by p . We now define $\ell'(x) = \ell'(x^T) = (\ell^T(x^T) - \ell_I^T(x^T))/p \pmod{m/p}$. Note that $\ell(x) \equiv p\ell'(x^T) + \ell_I(x) + \ell^S(x^S) - \ell_I^S(x^S) \pmod{m}$. Consider any assignment for x_{i_1}, \dots, x_{i_k} , and let $b := a + \ell^S(x^S) - \ell_I^S(x^S) \pmod{m}$. We set $A_a(x_{i_1}, \dots, x_{i_k}) = \{(z - b)/p \pmod{m/p} : z \in A, z \equiv b \pmod{p}\}$. \square

We now state the simplification lemma for $\mathcal{L}_n(m, k)$ systems of low communication rank.

Lemma 14. *Let $L \in \mathcal{L}_n(m, k)$ have m -wise communication rank at most r . Then there exist $L_1, \dots, L_R \in \mathcal{L}_n(m, \kappa(\cdot))$ such that $L = L_1 \cup \dots \cup L_R$ where $R \leq \exp((k + \log m + m \log m)r)$ and $\kappa : \text{div}(m) \rightarrow \mathbb{N} \cup \{-\infty\}$ is defined as*

1. *If $k > 0$ then $\kappa(m) = k - 1$ and $\kappa(\ell) = k + m \log m$ for all $\ell \in \text{div}(m) \setminus \{m\}$.*
2. *If $k = 0$ then $\kappa(m) = -\infty$ and $\kappa(\ell) = k + m \log m$ for all $\ell \in \text{div}(m) \setminus \{m\}$.*

Proof. Let $L = \bigwedge_{i=1}^N \{x \in \{0, 1\}^n : \ell_i(x) \pmod{m} \in A_i(x_{v(i,1)}, \dots, x_{v(i,k)})\}$. Denote the number of distinct prime factors of m by $w \leq \log m$. Assume that L has m -wise communication rank of at most r . Let I be the set of r equations given by Lemma 12. Consider the set of variables $W \subset [n]$ given by

$$W = \bigcup_{i \in I} V_i \cup \bigcup_{p,t} J_{p,t}.$$

We have $|W| \leq r(k + mw)$. We will consider all possible assignments to variables in W and all possible values for equations in I . For $\alpha \in \{0, 1\}^W$ and $\beta \in \mathbb{Z}_m^I$ define

$$B_{\alpha, \beta} := \{x \in \{0, 1\}^n : \forall w \in W, x_w = \alpha_w \text{ and } \forall i \in I, \ell_i(x) \equiv \beta_i \pmod{m}\}.$$

Note that $\{0, 1\}^n = \cup_{\alpha \in \{0, 1\}^W, \beta \in \mathbb{Z}_m^I} B_{\alpha, \beta}$. We will show that for any setting of α, β we have

$$B_{\alpha, \beta} \cap L \in \mathcal{L}_n(m, \kappa(\cdot)).$$

Hence, $L = L_1 \cup \dots \cup L_R$ for $R = 2^{|W|} m^{|I|} \leq \exp((k + \log m + m \log m)r)$.

We partition the set of rows outside of I into two parts. I' is the set of all rows i' such that $V_{i'}$ intersects W . Let I'' be the set of all other rows not in I' . Note that if $k = 0$ then I' is empty.

Consider first rows $i' \in I'$. Note that after fixing values for elements in W the sparsity of $A_{i'}$ reduces by at least one, hence they are equivalent to linear forms in $\mathcal{L}_n(m, k - 1)$. Consider next a row $i'' \in I''$. By Lemma 12 there exists a prime factor p of m with $\ell_{i''} \equiv \ell_I + \ell_0 \pmod{p}$, where ℓ_I is in the \mathbb{Z}_p -span of rows in I and ℓ_0 is wm -sparse and supported on variables in $x_{u(i'', 1)}, \dots, x_{u(i'', wm)}$. We will add these wm variables to the accepting set of row i'' . Thus, under every fixing of linear forms in I over \mathbb{Z}_m , by Claim 13 every linear constraint in I'' simplifies to one over a modulus m/p for some factor p of m with an accepting set of sparsity at most $k + mw$. Combining these arguments, for any row $i \in [N]$ and any assignment $\alpha \in \{0, 1\}^W, \beta \in \mathbb{Z}_m^I$ we have

$$L \cap B_{\alpha, \beta} = L_{\text{Sing}}^{(\alpha, \beta)} \cap L_m^{(\alpha, \beta)} \cap \bigcap_{p \text{ prime factor of } m} L_{m/p}^{(\alpha, \beta)},$$

where $L_{\text{Sing}}^{(\alpha, \beta)} \in \mathcal{L}_m^{\text{Sing}}$ is the constraints $x_w = \alpha_w$ for $w \in W$ (which can equivalently be stated modulo m as $x_w \in \{0, 1\}$) and $l_i(x) \equiv \beta_i \pmod{m}$ for $i \in I$; $L_m^{(\alpha, \beta)} \in \mathcal{L}_n(m, k-1)$ is given by simplified equations for $i' \in I'$; and $L_{m/p}(\alpha, \beta) \in \mathcal{L}_n(m/p, k+mw)$ is given by equations for $i'' \in I''$. Note that if $k=0$ then L_m does not appear. Thus $L \cap B_{\alpha, \beta} \in \mathcal{L}_n(m, \kappa(\cdot))$ and the lemma follows. \square

Lemma 9 follows immediately by combining Lemma 11 and Lemma 14, concluding this section.

References

- [BBR94] David Barrington, Richard Beigel, and Steven Rudich. Representing boolean functions as polynomials modulo composite numbers. *Computational Complexity*, 4:367–382, 1994. 10.1007/BF01263424.
- [BM97] Richard Beigel and Alexis Maciel. Upper and lower bounds for some depth-3 circuit classes. In *Proceedings of the 12th Annual IEEE Conference on Computational Complexity*, pages 149–157, Washington, DC, USA, 1997. IEEE Computer Society.
- [Bou05] Jean Bourgain. Estimation of certain exponential sums arising in complexity theory. *Comptes Rendus Mathématique*, 340(9):627 – 631, 2005.
- [BT87] D. Barrington and D. Thérien. Finite monoids and the fine structure of nc1 . In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, STOC '87, pages 101–109, New York, NY, USA, 1987. ACM.
- [Cau96] Hervé Caussinus. A note on a theorem of barrington, straubing and thrien. *Information Processing Letters*, 58(1):31 – 33, 1996.
- [CGPT06] A. Chattopadhyay, N. Goyal, P. Pudlák, and D. Thérien. Lower bounds for circuits with MOD_m gates. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 709–718, 2006.
- [Cha07] Arkadev Chattopadhyay. Discrepancy and the power of bottom fan-in in depth-three circuits. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 449–458, Washington, DC, USA, 2007. IEEE Computer Society.
- [CW09] A. Chattopadhyay and A. Wigderson. Linear systems over composite moduli. In *Foundations of Computer Science, 2009. FOCS '09. 50th Annual IEEE Symposium on*, pages 43 –52, 2009.
- [DGY10] Z. Dvir, P. Gopalan, and S. Yekhanin. Matching vector codes. In *IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 705–714, 2010.
- [Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pages 39–44, New York, NY, USA, 2009. ACM.

- [Gop06] P. Gopalan. Constructing ramsey graphs from boolean function representations. In *Computational Complexity, 2006. CCC 2006. Twenty-First Annual IEEE Conference on*, 0 2006.
- [GR98] D. Grigoriev and A.A. Razborov. Exponential complexity lower bounds for depth 3 arithmetic circuits in algebras of functions over finite fields. In *Foundations of Computer Science, 1998. Proceedings.39th Annual Symposium on*, pages 269–278, November 1998.
- [GRS05] Frederic Green, Amitabha Roy, and Howard Straubing. Bounds on an exponential sum arising in boolean circuit complexity. *Comptes Rendus Mathematique*, 341(5):279–282, 2005.
- [Han06] K.A. Hansen. On modular counting with polynomials. In *Computational Complexity, 2006. CCC 2006. Twenty-First Annual IEEE Conference on*, 0 2006.
- [HK09] Kristoffer Arnsfelt Hansen and Michal Koucký. A new characterization of ACC^0 and probabilistic CC^0 . *Computational Complexity, Annual IEEE Conference on*, 0:27–34, 2009.
- [HMP⁺93] A. Hajnal, W. Maass, P. Pudlák, M. Szegedy, and G. Turán. Threshold circuits of bounded depth. *J.Computer.System.Sciences*, 46(2):129–154, 1993.
- [Raz87] A. A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes*, 41:333–338, 1987. 10.1007/BF01137685.
- [Smo87] R. Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, STOC '87, pages 77–82, New York, NY, USA, 1987. ACM.
- [VW08] Emanuele Viola and Avi Wigderson. Norms, XOR lemmas, and lower bounds for $GF(2)$ polynomials and multiparty protocols. *Theory of Computing*, pages 137–168, 2008.
- [Wil11] Ryan Williams. Non-uniform ACC circuit lower bounds. to appear in 26th IEEE Conference in Computational Complexity, 2011.

A Linear systems over Abelian groups

In this section we generalize the argument to arbitrary Abelian groups of constant size. Most of the argument generalizes in a straightforward manner, and we highlight the required changes.

Let $G = \mathbb{Z}_{m_1} \oplus \dots \oplus \mathbb{Z}_{m_t}$ be a general Abelian group. We start by generalizing the definitions of systems of linear forms. A linear equation in n binary variables over G with an accepting set A is the set of solutions of

$$E = \{x \in \{0, 1\}^n : \sum g_i x_i \in A\}$$

where $g_1, \dots, g_n \in G$ and $A \subset G$. Let $\mathcal{L}_n(G)$ denote the set of common solutions to several such equations. Similarly, let $\mathcal{L}_n^{Sing}(G)$ denote the common solutions to several equations with a singleton accepting set and $\mathcal{L}_n(G, k)$ denote the set of common solutions to equations with accepting sets of sparsity k . We next generalize the definition of several linear systems with different moduli to the more general setting of Abelian groups. Let $\text{div}(G) = \{H < G\}$ be the set of all subgroups of G .

Note that for $G = \mathbb{Z}_m$ this is equivalent to our former definition, as $\text{div}(\mathbb{Z}_m) = \{\mathbb{Z}_\ell : \ell \div m\}$. For $\kappa : \text{div}(G) \rightarrow \mathbb{N} \cup \{-\infty\}$ we define $\mathcal{L}_n(G, \kappa(\cdot))$ to be the intersection of a system in $\mathcal{L}_n^{\text{Sing}}(G)$ with systems $\mathcal{L}_n(H, \kappa(H))$ for $H \in \text{div}(G)$ for which $\kappa(H) \geq 0$.

The following is the natural generalization of Theorem 7, which we will shortly prove.

Theorem 15 (Correlation bound for $\mathcal{L}_n(G, \kappa(\cdot))$ and MOD_q). *Let G be an Abelian group, and let q be co-prime to $|G|$. Let $\kappa : \text{div}(G) \rightarrow \mathbb{N} \cup \{-\infty\}$. Then*

$$\text{Corr}(\mathcal{L}_n(G, \kappa(\cdot)), \text{MOD}_q) \leq \exp(-n/c),$$

where $c = c_{15}(G, q, \kappa(\cdot))$.

Similar to Theorem 7, the proof follows from a base case and an inductive argument. Note that $\mathbb{Z}_1 = \{0\}$ corresponds to $\ell = 1$ in the cyclic case.

Lemma 16 (Base case for Abelian groups). *Let G be an Abelian group, and let q be co-prime to $|G|$. Let $k \geq 0$ be a sparsity. Let $L = L' \cap L''$ where $L' \in \mathcal{L}_n^{\text{Sing}}(G)$ and $L'' \in \mathcal{L}_n(\mathbb{Z}_1, k)$. Then*

$$\text{Corr}(L, \text{MOD}_q) \leq \exp(-n/\beta^k),$$

where $\beta = \beta(|G|, k)$ is as given in Theorem 6.

Proof. Let $m = |G| = m_1 \dots m_t$. The proof will follow from Lemma 8 for \mathbb{Z}_m . Let $E = \{x \in \{0, 1\}^n : \sum g_i x_i = b\}$ be an equation with a singleton accepting set. Let $g_i = (a_{i,1}, \dots, a_{i,t}) \in \mathbb{Z}_{m_1} \oplus \dots \oplus \mathbb{Z}_{m_t}$ and similarly let $b = (b_1, \dots, b_t)$. Then E can be equivalently described as

$$E = \bigcap_{j=1}^t \{x \in \{0, 1\}^n : \sum a_{i,j} x_i \pmod{m_j} \equiv b_j\}.$$

As $\mathbb{Z}_{m_1}, \dots, \mathbb{Z}_{m_t}$ are all subgroups of \mathbb{Z}_m we have that $E \in \mathcal{L}_n^{\text{Sing}}(\mathbb{Z}_m)$ and hence $\mathcal{L}_n^{\text{Sing}}(G) \subset \mathcal{L}_n^{\text{Sing}}(\mathbb{Z}_m)$. An equation $E \in \mathcal{L}_n(1, k)$ can be described by a degree k polynomial over *any* abelian group and in particular over \mathbb{Z}_m . \square

Lemma 17 (Simplification process for $\mathcal{L}_n(G, k)$). *Let G be an Abelian group, q be co-prime to $|G|$ and $k \geq 0$ be a sparsity. Then there exists $c = c_{17}(G, k)$ such that for any $L \in \mathcal{L}_n(G, k)$ and any $1 \leq r \leq n$, one of the following must hold:*

1. $\Pr_{x \in \{0,1\}^n} [x \in L] \leq \exp(-r/c)$.
2. There exist $L_1, \dots, L_R \in \mathcal{L}_n(G, \kappa(\cdot))$ such that $L = L_1 \cup \dots \cup L_R$ for $R \leq \exp(cr)$, and $\kappa : \text{div}(G) \rightarrow \mathbb{N} \cup \{-\infty\}$ is given as

- (a) If $k > 0$ then $\kappa(G) = k - 1$ and $\kappa(H) = k + |G|^2$ for all $H \in \text{div}(G) \setminus \{G\}$.
- (b) If $k = 0$ then $\kappa(G) = -\infty$ and $\kappa(H) = k + |G|^2$ for all $H \in \text{div}(G) \setminus \{G\}$.

The proof of Lemma 17 requires a generalization of communication rank to Abelian groups. Let $\text{Hom}(G)$ denote the set of all nonzero homomorphism to prime fields. That is,

$$\text{Hom}(G) = \{\varphi : G \rightarrow \mathbb{F}_p\}.$$

Note that these exist nonzero homomorphisms $\varphi : G \rightarrow \mathbb{F}_p$ iff $p \div |G|$. It is also not hard to show that $|\text{Hom}(G)| < |G|$, as we show in the appendix.

Definition 6 (Communication Rank for Abelian groups). Let $L \in \mathcal{L}_n(G, k)$ given by $L = E_1 \cap \dots \cap E_N$. We say that a subset of equations $I = \{i_1, \dots, i_r\} \subset [N]$ is s -wise independent if the following conditions hold. Let $V_i = \{v(i, 1), \dots, v(i, k)\}$ for $i \in I$ be the set of variables on which A_i depends. We first require that all the sets V_{i_1}, \dots, V_{i_r} are disjoint. We also require that there exist subsets of variables $J_{\varphi, t} \subset [n]$ of size $|J_{\varphi, t}| = |I| = r$, where $\varphi \in \text{Hom}(G)$ and $t = 1, \dots, s$, such that:

1. All sets $J_{\varphi, t}$ and V_{i_1}, \dots, V_{i_r} are pair-wise disjoint.
2. Assume $\varphi : G \rightarrow \mathbb{F}_p$. Let $M_{\varphi, t}$ be the following $r \times r$ matrix over \mathbb{F}_p : if $J_{\varphi, t} = \{j_1, \dots, j_r\}$ then the (x, y) -entry of $M_{\varphi, t}$ is given by $\varphi(g_{i_x, j_y})$, i.e. $M_{\varphi, t}$ is the image under φ of the $r \times r$ minor given by the rows of I and the columns of $J_{\varphi, t}$. We require that for any $\varphi \in \text{Hom}(G)$, and any $t = 1, \dots, s$, the matrix $J_{\varphi, t}$ has full rank modulo p .

The s -wise communication rank of L modulo G , denoted by $\text{ccrank}_s^G(L)$, is the maximal r for which this holds for some $I \subset [N]$ of size $|I| = r$.

We will need the following generalization of Lemma 10, which we prove in the appendix.

Lemma 18. Let $L \in \mathcal{L}_n(G)$ have $\text{ccrank}_{|G|}^G$ communication rank at least r . Then

$$\Pr_{x \in \{0,1\}^n} \left[\bigwedge_{i=1}^N \ell_i(x) \in A_i \right] \leq \exp(-r/c_{18}(G)),$$

where each $A_i \subsetneq \mathbb{Z}_m$ is an arbitrary set.

Given Lemma 18, it is straightforward to obtain the following generalization of Lemma 11.

Lemma 19. Let $L \in \mathcal{L}_n(G, k)$ have $|G|$ -wise communication rank at least r . Then

$$\Pr_{x \in \{0,1\}^n} [x \in L] \leq \exp(-r/c_{19}(|G|, k)).$$

The following is a generalization of Lemma 12. The proof is identical to the original proof in [CW09] given the natural changes required.

Lemma 20. Let G be an Abelian group. Consider an ordinary linear system $L \in \mathcal{L}_n(G)$ such that $\text{ccrank}_s^G(L) = r$. Then, there exists a set I of at r linear forms satisfying the following condition: for every linear form ℓ in L , there exists a nonzero homomorphism $\varphi : G \rightarrow \mathbb{F}_p$ such that $\varphi(\ell - (\ell_I + \ell_0)) \equiv 0 \pmod{p}$, where ℓ_I is in the \mathbb{Z} -span of I and ℓ_0 is $(|\text{Hom}(G)|s)$ -sparse.

The following is the generalization of Claim 13. Its proof is identical.

Claim 21. Let $\ell(x_1, \dots, x_n) \in A$ be any linear constraint in G . Let $\varphi : G \rightarrow \mathbb{F}_p$ be a nonzero homomorphism and let $H = \text{Ker}(\varphi)$ be the kernel of φ ($H < G$). Assume that $\varphi(\ell - (\ell_I + \ell_0)) \equiv 0 \pmod{p}$ where ℓ_0 is supported on variables x_{i_1}, \dots, x_{i_k} . For $a \in G$ let $B_a = \{x \in \{0, 1\}^n : \ell_I(x) \equiv a \pmod{G}\}$. Then for every $a \in G$ and values for $x_{i_1}, \dots, x_{i_k} \in \{0, 1\}$, there exists a linear form $\ell'(x)$ in H and a set $A_a(x_{i_1}, \dots, x_{i_k}) \subset H$ such that

$$\begin{aligned} B_a \cap \{x \in \{0, 1\}^n : \ell(x) \in A\} \\ = B_a \cap \{x \in \{0, 1\}^n : \ell'(x) \in A_a(x_{i_1}, \dots, x_{i_k})\}, \end{aligned}$$

The following is the natural generalization of Lemma 14. Its proof is identical.

Lemma 22. *Let $L \in \mathcal{L}_n(G, k)$ have $|G|$ -wise communication rank at most r . Then there exist $L_1, \dots, L_R \in \mathcal{L}_n(G, \kappa(\cdot))$ such that $L = L_1 \cup \dots \cup L_R$ where $R \leq \exp((k+2|G|^2)r)$ and $\kappa : \text{div}(G) \rightarrow \mathbb{N} \cup \{-\infty\}$ is defined as*

1. *If $k > 0$ then $\kappa(G) = k - 1$ and $\kappa(H) = k + |G|^2$ for all $H \in \text{div}(G) \setminus \{G\}$.*
2. *If $k = 0$ then $\kappa(G) = -\infty$ and $\kappa(H) = k + |G|^2$ for all $H \in \text{div}(G) \setminus \{G\}$.*

B Generalizations of technical lemmas to Abelian groups

We outline how the proof of Lemmas 10 and 18 follow along the lines of the original proof of [CW09, Lemma 13]. The only ingredient which doesn't generalize immediately is [CW09, Lemma 14], which we give below for general Abelian groups. Say a subset $S \subset G$ is *nontrivial* if for any homomorphism $\varphi \in \text{Hom}(G)$ there exist $g \in S$ such that $\varphi(g) \neq 0$. For two sets $S', S'' \subset G$ define their sum as $S' + S'' = \{g' + g'' : g' \in S', g'' \in S''\}$.

Claim 23. *Let G be an Abelian group with $|G| = m$. Let S_1, \dots, S_m be nontrivial subsets of G . Let $T_i = S_i \cup \{0\}$. Then $T_1 + \dots + T_m = G$. In particular any element of G can be written as a sum of subsets of S_1, \dots, S_m .*

Proof. It is enough to prove that if $A \subsetneq G$ and S is a nontrivial subset of G then $|A \cup (A+S)| > |A|$. Assume otherwise. For every $g \in S$ we have $A + g = A$, hence A is closed to \mathbb{Z} -sums of elements of S . That is, if $S = \{g_1, \dots, g_s\}$ then for any element $a \in A$,

$$\left\{ a + \sum_{i=1}^s c_i g_i : c_1, \dots, c_s \in \mathbb{Z} \right\} \subset A.$$

Let $H = \left\{ \sum_{i=1}^s c_i g_i : c_1, \dots, c_s \in \mathbb{Z} \right\}$, and note that $H < G$. We will next show that $H = G$ which contradicts the assumption that $A \subsetneq G$. If H is a proper subgroup of G then there is a homomorphism $\varphi' : G \rightarrow G/H$ such that $\varphi'(H) = 0$. Let p be a prime factor of $|G/H|$. There exists a nonzero homomorphism $\varphi'' : G/H \rightarrow \mathbb{F}_p$. Let $\varphi = \varphi'' \circ \varphi'$ be their composition. Then $\varphi : G \rightarrow \mathbb{F}_p$ is a nonzero homomorphism for which $\varphi(H) = \{0\}$. In particular $\varphi(S) = \{0\}$, which contradicts the assumption that S is nontrivial. \square

Claim 24. *Let G be an Abelian group. Then $|\text{Hom}(G)| < |G|$.*

Proof. Decompose G into a sum of cyclic groups of prime order, $G = \bigoplus_{p \text{ prime}, t \geq 1} (\mathbb{Z}_{p^t})^{e_{p,t}}$. We have $|G| = \prod |p^t|^{e_{p,t}}$. Fix a prime p , and consider all homomorphisms $\varphi : G \rightarrow \mathbb{F}_p$. They are identically zero on any \mathbb{Z}_q for q co-prime to p , and are determined by the image of the identity on any \mathbb{Z}_{p^t} . Thus the number of all homomorphism $\varphi : G \rightarrow \mathbb{F}_p$ is given by $\prod_{t \geq 1} p^{e_{p,t}}$. This includes also the identically zero homomorphism, which we need to subtract. We thus get that

$$|\text{Hom}(G)| = \sum_{p \text{ prime}} \left(\prod_{t \geq 1} p^{e_{p,t}} - 1 \right) \leq |G| - 1.$$

This is tight in general, as for $G = (\mathbb{Z}_p)^t$ we have $|G| = p^t$ and $|\text{Hom}(G)| = p^t - 1$. \square

The above two claims allow us to extend the argument of [CW09] to the setting of abelian groups to prove Lemma 18. We now describe the outline of this argument, referring the reader to [CW09] for the missing details. Let $L \in \mathcal{L}_n(G)$ with $\text{ccrank}_m^G = r$, where $|G| = m$. By definition of communication rank, there exists a set of r rows, I , and for each $\phi \in \text{Hom}(G)$ sets of pairwise disjoint columns $J_{\phi,1}, \dots, J_{\phi,m}$, $|J_{\phi,i}| = r$ for each i , with the following property: the rank of the sub-matrix corresponding to rows in I and columns in $J_{\phi,i}$, denoted by $M_{\phi,i}$ has full rank r over the field \mathbb{F}_p that is the range of ϕ . In order to prove Lemma 18, just as in [CW09] we will prove the following sharper result: let $d = |\text{Hom}(G)|m$ and let us enumerate the various sets of r columns simply as J_1, \dots, J_d . Further, let $\mathcal{J} = \cup_{i=1}^d J_i$. Let x^i be the variable representing the boolean assignment to variables corresponding to columns in J_i . For any $\ell \in L$, let $\ell(x^1, \dots, x^d)$ represent its natural projection to variables indexed by \mathcal{J} . We establish

$$\Pr_{x^1, \dots, x^d} \left[\bigwedge_{i \in I} \ell_i(x^1, \dots, x^d) \in A_i \right] \leq \exp(-\Omega(r)),$$

where each $A_i \subsetneq G$ is an arbitrary subset. This immediately implies Lemma 18. Following [CW09, GR98],

$$\left(\Pr_{x^1, \dots, x^d} \left[\bigwedge_{i \in I} \ell_i(x^1, \dots, x^d) \in A_i \right] \right)^{2^d} \leq \Pr_{x_0^1, x_1^1, \dots, x_0^d, x_1^d} [\forall i \in I; \forall u \in \{0, 1\}^d; \ell_i(x_{u_1}^1, \dots, x_{u_d}^d) \in A_i] \quad (2)$$

Remarking that Claim 24 implies d to be a constant, we bound the RHS of (2) for proving our Lemma. We note an obvious thing: let $\phi_d : G \rightarrow \mathbb{F}_{p_d}$ be the relevant homomorphism for the d th matrix. Then, for any i , $\phi_d(\ell_i(x^d)) = M_d x^d[i]$. Recalling that each M_d has full rank over \mathbb{F}_{p_d} , we directly use the following claim from the argument in [CW09].

Claim 25. *There exists constants γ and ν such that with probability at least $1 - \exp(-\nu r)$, there exists a set of rows I' of size r/γ^d , such that for each $i \in I'$ and $j \leq d$, we have $\phi_j(\ell_i(x_0^j)) \neq \phi_j(\ell_j(x_1^j))$.*

We refer the reader to [CW09] for a proof of this claim. Assume r is large so that $r/\gamma^d \geq 1$. Then, using above claim, pick a row $k \in I'$ such that $\phi_j(\ell_k(x_0^j)) \neq \phi_j(\ell_k(x_1^j))$ for all j . Consider sets T_1, \dots, T_m , defined as follows: for any $\phi \in \text{Hom}(G)$ and $i \leq m$, there is a $u_{\phi,i} \leq d$ such that $J_{u_{\phi,d}} = J_{\phi,i}$. Set,

$$T_i \equiv \{ \ell_k(x_0^u), \ell_k(x_1^u) \mid u = u_{\phi,i}, \phi \in \text{Hom}(G) \}$$

The sets T_1, \dots, T_m clearly satisfy the conditions required to invoke Claim 23. Hence, $T_1 + \dots + T_m = G$. So, there is a setting of $u \in \{0, 1\}^d$ such that $\ell_k(x_{u_1}^1, \dots, x_{u_d}^d)$ cannot lie in A_k . Such a k exists by Claim 25, with probability at least $(1 - \exp(-\nu r))$, providing the required upper bound for the RHS of (2). This completes the proof of Lemma 18.

We end by remarking that the main simplification tool for abelian groups, Lemma 17, follows by combining Lemma 19 and Lemma 22. This combination essentially mimicks the combination of Lemma 11 and Lemma 14 to yield the simplification tool for cyclic groups, Lemma 9, in the main body of the paper.