



CSPs with Global Modular Constraints: Algorithms and Hardness via Polynomial Representations

Joshua Brakensiek* Sivakanth Gopi† Venkatesan Guruswami‡

Abstract

We study the complexity of Boolean constraint satisfaction problems (CSPs) when the assignment must have Hamming weight in some congruence class modulo M , for various choices of the modulus M . Due to the known classification of tractable Boolean CSPs, this mainly reduces to the study of three cases: 2-SAT, HORN-SAT, and LIN-2 (linear equations mod 2). We classify the moduli M for which these respective problems are polynomial time solvable, and when they are not (assuming the ETH). Our study reveals that this modular constraint lends a surprising richness to these classic, well-studied problems, with interesting broader connections to complexity theory and coding theory. The HORN-SAT case is connected to the covering complexity of polynomials representing the NAND function mod M . The LIN-2 case is tied to the sparsity of polynomials representing the OR function mod M , which in turn has connections to modular weight distribution properties of linear codes and locally decodable codes. In both cases, the analysis of our algorithm as well as the hardness reduction rely on these polynomial representations, highlighting an interesting algebraic common ground between hard cases for our algorithms and the gadgets which show hardness. These new complexity measures of polynomial representations merit further study.

The inspiration for our study comes from a recent work by Nägele, Sudakov, and Zenklusen on submodular minimization with a global congruence constraint. Our algorithm for HORN-SAT has strong similarities to their algorithm, and in particular identical kind of set systems arise in both cases. Our connection to polynomial representations leads to a simpler analysis of such set systems, and also sheds light on (but does not resolve) the complexity of submodular minimization with a congruency requirement modulo a composite M .

Contents

1	Introduction	1
1.1	CSPs and modular CSPs	1
1.2	Our Results	3
1.3	Our Techniques and Connections	4
1.3.1	2-SAT-MOD $_M$ and recursive methods	4
1.3.2	HORN-SAT-MOD $_M$ and covering number of NAND mod M over $\{0, 1\}$ basis	5
1.3.3	LIN-2-MOD $_M$ and sparsity of the OR mod M over $\{-1, 1\}$ basis	6

*Department of Computer Science, Stanford University, Stanford, CA. Email: jbrakens@stanford.edu. Portions of this work were done while at Carnegie Mellon University and during a visit to Microsoft Research, Redmond. Research supported in part by NSF CCF-1526092 and an NSF Graduate Research Fellowship.

†Microsoft Research, Redmond, WA. Email: sigopi@microsoft.com.

‡Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213. Email: venkatg@cs.cmu.edu. Research supported in part by NSF grants CCF-1422045 and CCF-1526092.

1.3.4	Connections to submodular minimization and integer programming.	8
1.4	Future directions	8
1.5	Organization	9
2	Preliminaries	10
2.1	CSP basics	10
2.2	Mod-CSPs	11
2.3	Lucas theorem and other polynomial preliminaries	12
3	HORN-SAT-MOD_M	13
3.1	Algorithm for HORN-SAT-MOD _M	13
3.2	HORN-SAT-MOD _M with multiple modular constraints	15
3.3	Covering number of NAND _d mod M over $\{0, 1\}^d$	17
3.3.1	When M is a prime power	18
3.3.2	When M has multiple prime factors	19
3.4	Hardness of HORN-SAT-MOD _M	21
3.5	Submodular minimization with modular constraints	22
4	LIN-2-MOD_M	24
4.1	Algorithm for LIN-2-MOD _M	24
4.2	LIN-2-MOD _M with multiple modular constraints	26
4.3	Sparsity of polynomials representing OR _d mod M over $\{-1, 1\}^d$	28
4.3.1	M is a power of 2	28
4.3.2	When $M = 2^\ell M'$ for some odd M'	28
4.3.3	Reduction of $\mathcal{D}(n, M)$ to sparsity of OR mod M over $\{-1, 1\}$ basis for odd M	29
4.3.4	When M is an odd prime power	31
4.3.5	When M has multiple odd prime factors	32
4.4	Hardness of LIN-2-MOD _M	33
5	2-SAT-MOD_M	35
5.1	Preprocessing	36
5.2	Acyclic case	37
5.2.1	Correctness of Algorithm 5.2	38
5.2.2	Run-time of Algorithm 5.2	39
A	Establishing the Boolean Mod-CSP dichotomy	39
A.1	PP-reductions, Polymorphisms, and Galois correspondence	39
A.2	Extension to Mod-CSPs	41
A.3	Classification for prime powers	44
A.4	Partial classification for non-prime powers	44

1 Introduction

We study how the complexity of tractable cases of Boolean constraint satisfaction problems, namely 2-SAT, HORN-SAT and LIN-2 (linear equations over \mathbb{F}_2), is affected when we seek a solution that obeys a global modular constraint, such as Hamming weight being divisible by some modulus M . As our work reveals, this seemingly simple twist lends a remarkable amount of richness to these classic problems, raising new questions concerning polynomial representations of simple Boolean functions that form the common meeting ground of both algorithmic and hardness results.

The inspiration for our study comes from a beautiful recent work on minimizing a submodular function in the presence of a global modular constraint [NSZ18]. This framework captures questions such as: Given a graph G , find the minimum cut one of whose sides has size divisible by 6. The complexity of this basic question remains open. Surprisingly, the same combinatorial set system that governed the complexity of the algorithms in [NSZ18] arises in our study of HORN-SAT with a modular constraint. We connect such set systems to polynomial representations of the NAND function, thereby shedding light on submodular minimization with a global constraint involving a composite modulus, a case not handled by [NSZ18]. We describe this connection, as well as the relation to integer programs of bounded modularity that partly motivated [NSZ18] in Section 1.3.4.

1.1 CSPs and modular CSPs

We now describe some of our other, intrinsic motivations to study the modular variant of constraint satisfaction problems (CSPs). CSPs have a storied place in computational complexity, spurring several of its most influential developments such as NP-completeness, Schaefer’s dichotomy theorem [Sch78], the PCP theorem and the Unique Games conjecture (which together have led to the very rich field of inapproximability), and the algebraic program for studying CSPs inspired by the Feder-Vardi [FV98] and crystallized by Bulatov, Jeavons and Krokhin [BJK05], which recently culminated in the resolution of the CSP dichotomy conjecture [Bul17, Zhu17].

One reason that CSPs receive so much attention is that the local nature of their constraints offers just the right amount of structure to aid the development of novel algorithmic and hardness techniques, which then often extend to more general settings. For instance, semidefinite programming which was first used in approximating the Max-Cut CSP, has been one of the most influential algorithmic tools in approximating a whole variety of problems. On the hardness side, the PCP theorem, which is a statement about hardness of approximating Max-CSP, in combination with clever gadgets has led to inapproximability results for covering, packing, cut, routing, and other classes of problems. Further, for the problem of satisfiability of CSPs, we have a precise understanding of the interplay between mathematical structure and tractability: efficient algorithms exist iff the problem admits non-trivial “polymorphisms” which are operations under which the solution space is closed.

One enhancement to the CSP framework would be to impose some *global* constraint on the solution. For example, one could impose a global cardinality constraint, such as an equal number of 0’s and 1’s in the solution, or more generally a specified frequency for each value in the domain. This global condition is quite strict, often making many tractable CSPs NP-complete once these constraints are added. In fact, in the Boolean case the “hardest” problem that can be solved in polynomial time is (weighted)graph 2-coloring (by doing a simple dynamic program on the connected components) [CSS10].

In the case of approximating a 2-SAT instance with a balanced cardinality constraint (an equal number of 0s and 1s), it is NP-hard to solve [GL16]. In fact, the authors show that it is NP-hard

to find an assignment satisfying a $(1 - \varepsilon_0)$ fraction of clauses for some absolute constant $\varepsilon_0 > 0$. Further, this inapproximability holds for 2-SAT instances with Horn clauses, and thus also implies hardness of HORN-SAT with a balanced cardinality constraint.

This strictness of the constraints allowed for a full dichotomy (for any sized domain) to be proved long before the Feder-Vardi dichotomy was resolved, as most problems become NP-hard [BM10]. Bulatov and Marx showed that the only tractable problems are those that are “non-crossing decomposable.” Although the formal definition is a bit technical, informally such problems need to both be “convex” in that they are tractable in the second round of the Sherali-Adams hierarchy as well as “linear” in that they are solvable using a variant of Gaussian elimination (c.f., [BKW17]). As those two types of algorithms typically solve quite different problems (think 2-SAT versus 3-XOR), the family of tractable cardinality CSPs is far less diverse than for “ordinary” CSPs.

The main focus of this work is to investigate CSPs with a much less strict global constraint, which we refer to as a *modular* (or *congruency*) constraint. We will restrict ourselves to the Boolean case in this work, and impose the requirement that the number of 1’s in the solution be congruent to ℓ modulo M , for some integers ℓ, M .^{*} We refer to this class of problems as (Boolean) Mod-CSPs.

Informally, it is easy to see that any such Mod-CSP is at least as hard as the corresponding local CSP, as we can take a local CSP instance and add M dummy variables not part of any clauses, so that the modular constraint is now trivially satisfiable. Conversely, these Mod-CSPs are at least as easy (up to polynomial factors) as the corresponding cardinality problem, because we can brute force the cardinality of 1s by trying all $c_1 = \ell \pmod M$.

By Schaefer’s Boolean CSP dichotomy theorem [Sch78], there are only three essentially different non-trivial tractable cases of Boolean CSP: 2-SAT, HORN-SAT, and LIN-2. We thus study each of these problems when we seek a solution of Hamming weight in some congruence class modulo M , for a fixed M . (When M can grow with the input, these problems become NP-hard as one can encode a global cardinality constraint [CSS10].) Our goal is to classify the cases which are polynomial time solvable, as a function of M . In order to better appreciate the difficulty of this endeavor, we encourage the reader to not peek ahead and write down a guess for each of the cases listed in the table below. (Note that 2-SAT-MOD $_M$ refers to 2-SAT with the global constraint modulo M , etc.)

Name	Constraint types	$M = 3$	$M = 4$	$M = 6$	$M = 15$
2-SAT-MOD $_M$	$x \vee y; x = \neg y$				
HORN-SAT-MOD $_M$	$x \wedge y \rightarrow z$				
LIN-2-MOD $_M$	$x \oplus y \oplus z = 0$ or 1				

Our work hinges on several connections which makes our investigation interesting beyond the specific application to modular CSPs. The complexity of the problems (except 2-SAT-MOD $_M$) are tied to the parameters of certain polynomial representations — lower bounds on such representations yields our algorithmic guarantees, and at the same the existence of efficient representations implies hardness results. Thus the work illustrates an interesting duality between algorithms and hardness as originating from the same object. The relevant complexity measures for polynomial representations are novel and deserve further study. As the particular choice of complexity measure for each problem seems closely linked to the underlying polymorphisms of the CSP, we hope that initiating such a study could help bring together computational complexity theorists and specialists in the algebraic theory of CSPs.

^{*}Our actual setup is a bit more general, associating a weight from an abelian group for each variable-value pair, and requiring that the sum of the weights equals some value.

The result for linear equations has interesting connections to coding theory, namely the extremal dimension of binary linear codes whose codewords have modular restrictions on their weights (which relates to concepts like doubly even codes that have been studied in coding theory), as well as to locally decodable codes, via relationship between polynomial representations and matching vector families and the Polynomial Freiman-Ruzsa (PFR) conjecture in additive combinatorics.

1.2 Our Results

We resolve the complexity of 2-SAT-MOD $_M$, HORN-SAT-MOD $_M$, and LIN-2-MOD $_M$, namely whether it belongs to or is unlikely to be polynomial time solvable, for all moduli.[†] We will denote by 2-SAT-MOD $_M(n)$ an instance of 2-SAT with n variables and a global constraint modulo M , HORN-SAT-MOD $_M(n)$ and LIN-2-MOD $_M(n)$ are similarly defined.

Theorem 1.1 (Informal statement of main results). *Suppose we have a single global modular constraint with a fixed modulus M .*

1. 2-SAT-MOD $_M(n)$ is polytime solvable for all moduli M .
2. HORN-SAT-MOD $_M(n)$ is polytime solvable when the modulus M equals a prime power, and assuming ETH, cannot be solved in $n^{o((\log n / \log \log n)^{r-1})}$ time when M has r distinct prime factors.
3. LIN-2-MOD $_M(n)$ is randomized polytime solvable when the modulus $M = 2^\ell p^s$ for some odd prime p and $\ell \leq 2$, and assuming ETH, cannot be solved in $n^{o((\log n / \log \log n)^{r-1})}$ when M has r distinct odd prime factors. Further, assuming the Polynomial Freiman-Ruzsa conjecture, there is an algorithm with runtime $\exp(O_M(n / \log n))$ for all M .[‡]

The polynomial runtime grows like $n^{O(M)}$, and assuming ETH this cannot be improved to $n^{O(M/\log M)}$ for the case of HORN-SAT-MOD $_M$ and LIN-2-MOD $_M$. Moreover, for the LIN-2-MOD $_M$ problem, the randomized polytime algorithms have quasipolynomial derandomizations.

For the one uncovered case of LIN-2-MOD $_M(n)$, when $M = 2^\ell p^s$ for an odd prime p and $\ell > 2$, we give a quasi-polynomial time algorithm with runtime $\exp(O_M((\log n)^{2^{\ell-1}-1}))$. We are not sure if there should be a polynomial time algorithm also in this case.

The following table summarizes our results. We will assume $M \leq cn$ for some sufficiently small constant $c > 0$. All the hardness results are assuming ETH. For simplicity, the results for LIN-2-MOD $_M$ are only stated for odd M and the running times are for randomized algorithms, the algorithm for general M is assuming PFR conjecture.

	M is a prime power		M has r distinct prime factors ($r \geq 2$)	
	Algorithm	Hardness	Algorithm	Hardness
2-SAT-MOD $_M(n)$	$n^{M+O(1)}$	-	$n^{M+O(1)}$	-
HORN-SAT-MOD $_M(n)$	$n^{M+O(1)}$	$n^{\Omega(\frac{M}{\log M})}$	-	$\exp\left(\Omega_M\left(\left(\frac{\log n}{\log \log n}\right)^r\right)\right)$
LIN-2-MOD $_M(n)$ [§]	$n^{M+O(1)}$	$n^{\Omega(\frac{M}{\log M})}$	$\exp\left(O_M\left(\frac{n}{\log n}\right)\right)$	$\exp\left(\Omega_M\left(\left(\frac{\log n}{\log \log n}\right)^r\right)\right)$

[†]Unfortunately, one uncovered case is LIN-2-MOD $_M$ for moduli $M = 2^\ell p^s$ for an odd prime p and $\ell \geq 3$.

[‡]We write $A = O_M(B)$ or $A \lesssim_M B$ to say that $A \leq C_M B$ for some constant $C_M > 0$ which depends only on M and not on any other parameters. $A = \Omega_M(B)$, $A \gtrsim_M B$, $A = o_M(B)$, $A = \omega_M(B)$ are similarly defined.

Extensions to multiple modular constraints. We also consider natural extensions of the three problems, where we allow a small number of mod M constraints or the more general version where we allow ℓ constraints with different moduli M_1, M_2, \dots, M_ℓ . Our algorithm our 2-SAT-MOD $_M$ is presented in this more general model. For HORN-SAT and LIN-2, we show that these general versions can be reduced to the basic version with a single mod M constraint, without increasing the size of the instance too much. For example we can reduce an instance of HORN-SAT over n variables with ℓ constraints modulo M to an instance of HORN-SAT-MOD $_M$ over $n^{O_M(\ell)}$ variables. Note that once ℓ becomes linear in n , these problems are NP-hard, so an exponential dependence in ℓ is necessary. A similar statement also holds for LIN-2 with multiple modular constraints.

Completing a classification of Boolean Mod-CSPs. Although HORN-SAT, LIN-2, and 2-SAT are the most important CSPs that need to be analyzed, some additional care is needed to extend these classifications to all Boolean CSPs. This work is done in Appendix A, where we classify the computational complexity. The main observation is to use a classification by Post [Pos41] of the polymorphisms of Boolean CSPs. Using this classification, we can show that when M is a prime power, the Mod-CSP problem is always either in RP or is NP-complete. On the other hand, when M is divisible by distinct primes, the classification becomes a bit more difficult. The main difficulty is that we have an upper bound for 2-SAT-MOD $_M$, but lower bounds on HORN-SAT-MOD $_M$ and LIN-2-MOD $_M$ (when M is divisible by distinct odd primes). Because of the confluence of lower and upper bounds, there are some additional tractable cases that show up, which are similar in structure to 2-SAT. See Section A.4 for more details.

1.3 Our Techniques and Connections

We now give brief overviews of our approach to establish Theorem 1.1. We discuss each of the three constraint types in turn.

1.3.1 2-SAT-MOD $_M$ and recursive methods

For 2-SAT-MOD $_M$, our algorithm is recursive. The one key idea is that we work with a more general form of modular constraint to make the recursion work, one that allows the Hamming weight to belong to a subset S of congruence classes (rather than a single value). Standard methods, with some care to update the modular constraint, allow us to reduce to the case when the “implication graph” on the literals of the 2-SAT instance is a DAG. We then select a literal y with no outgoing edge, and first set it to 1 (which doesn’t impact any other literal), and solve the 2-SAT instance on the remaining variables with an updated modular constraint (that takes into account the setting of y). If this succeeds, we can output this assignment and be done.

Otherwise we set $y = 0$, which forces all literals which have a path to y in the DAG also to 0. We can update the modular constraint accordingly, but note that in the end we are allowed to flip y to 1 and the 2-SAT instance will still be satisfied. While there is no need to do this for normal 2-SAT, this flip might allow us to satisfy the modular constraint. As a result, we allow the set S of congruence classes in the recursive call to also include this possibility. This is the reason why we need to work with the more general form of modular constraint. To implement this idea to run in polynomial time is a bit subtle, as naively we could reduce an instance with n variables to two instances with $n - O(1)$ variables leading to exponential runtime. To avoid this pitfall, we track the size of the allowed moduli S , and argue that if it doesn’t increase in the second recursive call

[§]For LIN-2-MOD $_M$, we are only stating the results for odd M and the algorithms are randomized. The running time of the algorithm for non-prime-power M is conditioned on Polynomial Freiman-Ruzsa conjecture.

(one where we set $y = 0$), we can truncate that call and return no solution for that branch. This is justified because any valid solution with $y = 0$ remains valid when $y = 1$, and the former is already ruled out in the first call. The increase in $|S|$ in one of the recursive calls implies a polynomially bounded solution to the recurrence for the runtime, with exponent at most $M - 1$.

1.3.2 HORN-SAT-MOD $_M$ and covering number of NAND mod M over $\{0, 1\}$ basis

Our algorithm and analysis for HORN-SAT-MOD $_M$, i.e., HORN-SAT with a single global linear constraint modulo M , is very different from the 2-SAT case. An important property of the HORN-SAT instance is that the set of solutions is intersection-closed. Given an instance Ψ of HORN-SAT and a subset S of its variables, one can efficiently find the minimal solution among all solutions of Ψ which set the variables of S to 1, this is called the FindMinimal(Ψ, S) routine. Now we run FindMinimal(Ψ, S) on all subsets S upto a give size R and check if any of the outputs satisfy the modular constraint. If none of them satisfy the modular constraint, we claim that Ψ has no solution which satisfies the modular constraint. The running time of this R -round algorithm is $n^{R+O(1)}$. If the R -round algorithm for HORN-SAT-MOD $_M$ fails, then we show that it is because of a special kind of obstruction. To describe these obstructions, we will need a few definitions.

Definition 1.2. A multilinear polynomial $p(x_1, \dots, x_d)$ is said to represent NAND $_d$ mod M over $\{0, 1\}^d$ if it has integer coefficients and

$$p(x) \begin{cases} = 0 \pmod M & \text{if } x = \mathbf{1} \\ \neq 0 \pmod M & \text{if } x \in \{0, 1\}^d \setminus \{\mathbf{1}\} \end{cases}$$

where $\mathbf{1}$ is the all ones vector.

Definition 1.3 (Covering number). The covering number of a multilinear polynomial $p(x)$, denoted by $\text{cov}(p)$, is the minimum number of monomials of p one can choose such that every variable that appears in p appears in one of them.

Note that the covering number is the minimum set cover of the family of subsets of variables given by the monomials. An obstruction for the R -round algorithm to solve HORN-SAT-MOD $_M(n)$ correctly is a polynomial p which represents NAND mod M with at most $n + 1$ monomials and $\text{cov}(p) > R$. Therefore we have the following proposition.

Proposition 1.4. If every polynomial p which represents NAND mod M over $\{0, 1\}$ basis with $n+1$ monomials has $\text{cov}(p) \leq R(n)$, then there exists an $n^{R+O(1)}$ -time algorithm for HORN-SAT-MOD $_M(n)$.

When M is a prime power. In this case, we can show that any polynomial which represents NAND mod M should have covering number at most $M - 1$. Note that this bound is independent of the number of monomials in the polynomial. This implies that our algorithm with $R = M - 1$ rounds solves HORN-SAT-MOD $_M(n)$ correctly.

When M has $r \geq 2$ distinct prime factors. In this case, it turns out that there are obstructions for any constant round algorithm. More precisely, there are polynomials which represent NAND mod M with n monomials, but their covering number is $\Omega((\log n / \log \log n)^{r-1})$. Such polynomials can be obtained from polynomials of degree $O(d^{1/r})$ which represent NAND $_d$ mod M constructed by Barrington, Beigel and Rudich [BBR94]. This implies that the R -round algorithm will not work if we choose $R = o((\log n / \log \log n)^{r-1})$. This by itself does not show hardness of HORN-SAT-MOD $_M$, it just shows that our algorithm doesn't work with constant rounds. But it turns out that we can

use low-degree polynomial representations of $\text{NAND} \bmod M$ directly as a gadget to reduce 3-SAT to HORN-SAT-MOD_M without blowing up the size too much. More precisely, we show the following hardness result.

Proposition 1.5. *Suppose there exists a polynomial which represents $\text{NAND}_d \bmod M$ over $\{0, 1\}^d$ with degree Δ . Assuming ETH, $\text{HORN-SAT-MOD}_M(n)$ requires $2^{\Omega(d)}$ time for some $n = (ed/\Delta)^{O(\Delta)}$.*

Using the upper bound $\Delta = O(d^{1/r})$ from [BBR94], we get $\exp(\Omega_M((\log n / \log \log n)^r))$ time ETH-hardness for $\text{HORN-SAT-MOD}_M(n)$. Better upper bounds on Δ lead to better hardness results. The best known lower bound on Δ is $\Omega_M((\log d)^{1/(r-1)})$ due to Barrington and Tardos [TB98]. If there is a polynomial whose degree matches this lower bound, then assuming ETH, we can get $\exp(\exp((\log n)^{1-1/r}))$ hardness for $\text{HORN-SAT-MOD}_M(n)$.

What about sub-exponential time algorithms? We conjecture that the covering number of any polynomial which represents $\text{NAND} \bmod M$ with n monomials should be $n^{o_M(1)}$ for any fixed M . If true, this would imply an $\exp(n^{o_M(1)})$ -time algorithm for $\text{HORN-SAT-MOD}_M(n)$. We give evidence towards our conjecture by showing that the fractional covering number (which is an LP relaxation of covering number) of any such polynomial is indeed $n^{o_M(1)}$.

1.3.3 LIN-2-MOD $_M$ and sparsity of the OR $\bmod M$ over $\{-1, 1\}$ basis

The algorithm and analysis for LIN-2-MOD_M has strong similarities to the analysis for HORN-SAT-MOD_M . The algorithm is quite simple to state: sample T uniformly random solutions to the linear system over \mathbb{F}_2 and check if any of them satisfy the modular constraint.

To analyze how many samples T to check so that the algorithm is correct with high probability, we need to prove that for any d -dimensional affine subspace of \mathbb{F}_2^n , there are either 0 points in this space satisfying the modular constraint or there are at least $2^d/f(n, M)$ such points. Then $T = O(f(n, M))$ samples would suffice. By a simple reduction, it suffices to bound the maximal dimension $\mathcal{D}(n, M)$ such that a \mathcal{D} -dimensional affine subspace of \mathbb{F}_2^n has *exactly one* element whose Hamming weight is $a \bmod M$ for some a . Quantitatively, we show that $f(n, M) \leq O(2^{\mathcal{D}(n, M)})$. In other words, the obstructions for our algorithm are large affine subspaces which have exactly one point which satisfies a linear constraint modulo M . When M is a power of 2, we prove that $\mathcal{D}(n, M) \leq M - 1$. For $M = 2^\ell M'$ for some odd $M' \geq 3$, we can get upper bounds on $\mathcal{D}(n, M)$ from upper bounds on $\mathcal{D}(n, M')$. So we can only focus on odd M . The obstructions for odd M can be represented using polynomials like we did in the HORN-SAT-MOD_M case.

Definition 1.6. *A polynomial $p(x_1, \dots, x_d)$ is said to represent $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ if it has integer coefficients and*

$$p(x) \begin{cases} = & 0 \pmod M \text{ if } x = \mathbf{1} \\ \neq & 0 \pmod M \text{ if } x \in \{-1, 1\}^d \setminus \{\mathbf{1}\} \end{cases}$$

where $\mathbf{1}$ is the all ones vector.

The existence of a d -dimensional affine subspace of \mathbb{F}_2^n with exactly one point satisfying the $\bmod M$ constraint is equivalent to an $(n + 1)$ -sparse polynomial representation of $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$. Thus, lower bounds on the sparsity of representations of $\text{OR} \bmod M$ in the $\{-1, 1\}$ basis imply upper bounds on $\mathcal{D}(n, M)$ and the runtime of our randomized algorithm. More precisely, we have the following proposition.

Proposition 1.7. *Let M be odd, if there is no polynomial which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ with at most $n + 1$ monomials, then there exists a $O(2^d n^{O(1)})$ time randomized algorithm for $\text{LIN-2-MOD}_M(n)$.*

Connections to circuit complexity and locally decodable codes. Note that an s -sparse polynomial which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ is equivalent to a two layer $\text{MOD}_M \circ \text{MOD}_2$ circuit which computes OR over $\{0, 1\}^d$ with s MOD_2 -gates at the bottom level connected to a MOD_M gate at the top.[¶] Circuits made up of such MOD gates are extensively studied in circuit lower bounds [CGPT06, CW09]. An s -sparse polynomial representation of $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ can be used to construct a *matching vector family* (MVF) of size 2^d over $(\mathbb{Z}/M\mathbb{Z})^s$. MVFs of large size and small dimension can be used to construct good *locally decodable codes* (LDCs) and private information retrieval (PIR) schemes [Yek08, Efr12, DGY11, DG16]. So lower bounds for LDCs or upper bounds on the size of MVFs imply lower bounds on the sparsity of polynomials which represent $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$. Concretely, superpolynomial lower bounds for the length of constant-query LDCs would imply $2^{o(n)}$ time algorithms for $\text{LIN-2-MOD}_M(n)$ when M is any constant; we find this connection quite surprising!

When M is an odd prime power. In this case, we can show that the sparsity of a polynomial which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ is at least $2^{d/(M-1)}$, which is nearly tight. In other words, $\mathcal{D}(n, M) \leq (M-1) \log n$. This immediately implies a $n^{M+O(1)}$ -time randomized algorithm for LIN-2-MOD_M .

When M has $r \geq 2$ odd prime factors. The results of [BBR94] imply that there exist degree $O(d^{1/r})$ polynomials which represent $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$, therefore they have at most $n = \exp(O(d^{1/r} \log d))$ monomials. In other words, $\mathcal{D}(n, M) \gtrsim (\log n / \log \log n)^r$. Thus there are obstructions for our algorithm to run in polynomial time. This by itself does not show hardness. But, similar to the HORN-SAT-MOD_M case, we can directly use low-degree polynomials which represent $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ to reduce 3-SAT to LIN-2-MOD_M without increasing the size too much.

Proposition 1.8. *Suppose there exists a polynomial of degree Δ which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$. Assuming ETH, $\text{LIN-2-MOD}_M(n)$ takes $2^{\Omega(d)}$ time for some $n = (ed/\Delta)^{O(\Delta)}$.*

Using the upper bound $\Delta = O(d^{1/r})$ from [BBR94], we get $\exp(\Omega_M((\log n / \log \log n)^r))$ time ETH-hardness for $\text{LIN-2-MOD}_M(n)$. Better upper bounds on Δ lead to better hardness results. The best known lower bound on Δ is $\Omega_M((\log d)^{1/(r-1)})$ due to Barrington and Tardos [TB98]. If there is a polynomial whose degree matches this lower bound, then assuming ETH, we can get $\exp(\exp((\log n)^{1-1/r}))$ -time hardness for $\text{LIN-2-MOD}_M(n)$.

What about subexponential time algorithms? Unfortunately, we do not know any unconditional superlinear (i.e. $\omega(d)$) lower bounds on the sparsity of polynomials which represent $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$. Such a lower bound would imply $2^{o(n)}$ -time algorithms for $\text{LIN-2-MOD}_M(n)$ for any constant M . But using the connection to MVFs, and an upper bound on the size of MVFs due to [BDL14] assuming Polynomial Freiman-Ruzsa (PFR) conjecture from additive combinatorics, we can conclude that the sparsity of a polynomial representing $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ is at least $\Omega_M(d \log d)$ under the same conjecture. This implies an $\exp(O_M(n / \log n))$ -time algorithm for $\text{LIN-2-MOD}_M(n)$ assuming the PFR conjecture.

[¶]Here a MOD_M gate outputs 0 if the number of 1-inputs is a multiple of M and outputs 1 otherwise.

1.3.4 Connections to submodular minimization and integer programming.

A classic result in combinatorial optimization is that integer linear programs which have a *totally unimodular* constraint matrix, i.e., every square submatrix has determinant in $\{-1, 0, 1\}$, can be solved in polynomial time. This is because the vertices of the feasible polytope of such a constraint matrix are integral. Recently this result has been extended to *totally bimodular* constraint matrices, where every square submatrix has determinant in $\{-2, -1, 0, 1, 2\}$ [AWZ17]. Such results have inspired a conjecture that any integer program for which every square submatrix has determinant bounded in absolute value by M , can be solved in $n^{O_M(1)}$ time.

A recent paper by Nägele, Sudakov and Zenklusen [NSZ18] tries to lay the groundwork for proving this conjecture by considering a special case. This special case is finding the minimum cut in a directed graph such that the number of vertices on one side of the cut satisfies a modular condition modulo M . As cut functions of directed graphs are submodular, they generalized this question to the following algorithmic problem, denoted by $\text{SUBMOD-MIN-MOD}_M(n)$: Given a, M , minimize a submodular function $f : \{0, 1\}^n \rightarrow \mathbb{Z}$ (given oracle access) over all x such that $\sum_{i=1}^n x_i = a \pmod M$.^{||}

Our algorithm for HORN-SAT-MOD_M is actually inspired by the algorithm from [NSZ18] for $\text{SUBMOD-MIN-MOD}_M(n)$. They show that the only obstructions to their R -round algorithm to work correctly are certain set families they called (M, R, d) -systems. They then showed that when M is a prime power, $(M, M - 1, d)$ -systems do not exist for any d . This implies an $n^{O(M)}$ -time algorithm for SUBMOD-MIN-MOD_M when M is a prime power. They asked if a polynomial time algorithm exists for SUBMOD-MIN-MOD_M when M has multiple prime factors.

We observe that (M, R, d) -systems are equivalent to polynomial representations of $\text{NAND} \pmod M$ over $\{0, 1\}$ basis with d monomials and covering number greater than R . Thus the obstructions for their algorithm for SUBMOD-MIN-MOD_M and our algorithm for HORN-SAT-MOD_M are exactly the same! Thus we can give a simpler proof of their result for prime power M , and also throw new light on SUBMOD-MIN-MOD_M when M has multiple prime factors. In particular, if M has r distinct prime factors, this implies that there are obstructions for their R -round algorithm to work for $\text{SUBMOD-MIN-MOD}_M(n)$ for $R = o_M((\log n / \log \log n)^{r-1})$. This answers an open question from their paper and explains why they couldn't extend their algorithm for any constant M . If true, our conjecture that the covering number of a polynomial which represents $\text{NAND} \pmod M$ with n monomials is $n^{o_M(1)}$, implies an $\exp(n^{o_M(1)})$ time algorithm for SUBMOD-MIN-MOD_M for any M . We also make a conjecture about the existence of certain submodular functions, which would allow us to prove superpolynomial ETH-hardness results for SUBMOD-MIN-MOD_M when M has multiple prime factors. It would be interesting if our methods have any implications for hardness of solving integer linear programs with bounded-minor constraint matrices, our results suggest that something different can happen at $M = 6$.

1.4 Future directions

Given that our work is the first to study the effect of global modular constraints on the tractability of CSPs, and our results unearth a rich picture rife with interesting connections to many central topics such as algebraic complexity measures of Boolean functions, coding theory, and combinatorial optimization, there are naturally many questions and directions for future work. We list a few below.

^{||}A function $f : \{0, 1\}^n \rightarrow \mathbb{Z}$ is called submodular if for all $x, y \in \{0, 1\}^n$, $f(x) + f(y) \geq f(x \vee y) + f(x \wedge y)$, where \vee and \wedge are bitwise OR and AND.

- Our work raises some intriguing new questions about polynomial representations modulo a composite number M . Can one prove an unconditional non-trivial (i.e. $\omega(d)$) lower bound on the sparsity of polynomial representing $\text{OR}_d \bmod M$ over the $\{1, -1\}$ -basis? Can one construct polynomial representations that are sparser than what is obtainable by simply appealing to the best known low-degree representations?

For the covering complexity measure for polynomials representing $\text{NAND} \bmod M$ over the $\{0, 1\}$ basis with n monomials, can we prove an upper bound of $n^{o_M(1)}$ (Conjecture 3.20)? This would imply a sub-exponential time algorithm for HORN-SAT-MOD_M and SUBMOD-MIN-MOD_M via our connections.

- The analysis of our algorithms for LIN-2-MOD_M and HORN-SAT-MOD_M use the sparsity and covering complexity of appropriate polynomial representations. On the other hand, our hardness results for these problems use the degree measure of these polynomials. Can one prove that sparsity and covering complexity directly dictate hardness as well? This would complete a very pleasing picture by giving matching algorithmic and hardness results and tying the complexity to some complexity parameter associated with polynomial representations.
- Can one show hardness of submodular minimization with a global mod 6 constraint? Constructing certain submodular functions as described in Conjecture 3.33 would imply such a hardness result. What about hardness of solving integer linear programs with bounded minors?
- Of the major cases for our classification of Boolean CSPs, 2-SAT-MOD_M is special in that its complexity is not tied to the complexity of certain polynomial representations. Given that the polynomials representations which play a role for HORN-SAT-MOD_M and LIN-2-MOD_M are very closely related to the polymorphisms of HORN-SAT and LIN-2 , AND and XOR respectively, how does the MAJ polymorphisms of 2-SAT play a role in the tractability of 2-SAT-MOD_M ? Understanding this would give us clues about how to generalize our results to CSPs over larger domains.
- Furthermore, 2-SAT-MOD_M is also the only case in which we lack a lower bound of the form $\Omega(n^{f(M)})$, for some nontrivial function of f . Could it be that 2-SAT-MOD_M is fixed-parameter tractable? For instance, does there exist a $2^M n^{O(1)}$ algorithm? We note that the dependence on M cannot be polynomial (unless $\text{P} = \text{NP}$), as setting M greater than n would solve 2-SAT with a global cardinality constraint, which is NP-complete.
- How do Mod-CSPs behave on a non-Boolean domain? Even for a domain of size three, the classification of ordinary CSPs is much more complex [Bul06]. As such, in order for such a program to be carried out, one needs to better understand the interplay between the global modular constraints and the polymorphisms of these CSPs. In particular, how do notions like cores, bounded width, identities, etc., interplay with the modular constraints? See [BKW17] for definitions of these terms.
- What other interesting global constraints can we impose on CSPs like 2-SAT , HORN-SAT , LIN-2 , while still keeping them tractable? What happens if we add a global constraint over a non-abelian group i.e. a global constraint of the form $\prod_{i=1}^n g_i^{x_i} = g_0$ for some $g_0, g_1, \dots, g_n \in G$ where G is a non-abelian group.

1.5 Organization

In Section 2, we formally define what a Mod-CSP is as well as state some basic facts about polynomials. In Section 3, we study HORN-SAT-MOD_M . In Section 4, we study LIN-2-MOD_M . In Section 5,

we prove our algorithmic results for 2-SAT-MOD_M. In Section A, we show how the results for these individual problems imply some dichotomy-type results for general Boolean Mod-CSPs.

2 Preliminaries

In this section, we give standard definitions of (local) CSPs as well as a formal definition of Mod-CSPs.

2.1 CSP basics

The type of CSP we study depends on a few parameters. The number of variables is often denoted by n and the variables are written as x_1, \dots, x_n . The number of constraints is denoted by m . Although the number of constraints in an instance tends to ∞ and $n \rightarrow \infty$, there are only finitely many types of constraints, which are specified by the *template* Γ . To specify the structure of the template, we have a *signature* $\sigma = (I, \text{arity})$, where I is an index set of the constraints, and $\text{arity} : I \rightarrow \mathbb{N}$ gives the *arity* of each constraint, how many variables it takes as arguments. We formally define the template as follows.

Definition 2.1. *A template with signature σ over domain D is an ordered tuple $\Gamma = (C_i \subset D^{\text{arity}_i} \mid i \in I)$.*

For this template, we say that a CSP over Γ is a formula Ψ over variables x_1, \dots, x_n which is a CNF of constraints from Γ . That is,

$$\Psi(x_1, \dots, x_n) = \bigwedge_{i=1}^m C_{f(i)}(x_{j_{i,1}}, \dots, x_{j_{i,\text{arity}_{f(i)}}})$$

for some $f : [m] \rightarrow I$. Note that variables need not be distinct. From this, we can define the corresponding computational problem.

Definition 2.2. *The decision problem $\text{CSP}(\Gamma)$ asks, given an instance Ψ does there exist a solution?*

The following are a few concrete examples of tractable Γ in the Boolean domain ($D = \{0, 1\}$).

- 2-SAT can be encoded by $\Gamma_{2\text{-SAT}} = \{C \subset \{0, 1\}^2 : |C| = 3\}$. There are several efficient algorithms for 2-SAT (e.g., [Kro67, BKW17]).
- LIN-2 can be encoded by $\Gamma_{\text{LIN-2}} = \{\{(x, y, z) : x \oplus y \oplus z = a\} : a \in \{0, 1\}\}$, where \oplus is addition modulo 2. LIN-2 can be solved in polynomial time using Gaussian elimination.
- HORN-SAT is encoded by $\Gamma_{\text{HORN-SAT}} = \{\{(x, y, z) : x \wedge y \rightarrow z\}, \{(0)\}, \{(1)\}\}$. HORN-SAT can be solved in polynomial time using a deduce-and-propagate style of algorithm. Despite the apparent simplicity of the constraint, $\Gamma_{\text{HORN-SAT}}$ is P-complete [CN10].
- DUAL-HORN-SAT is encoded by $\Gamma_{\text{DUAL-HORN-SAT}} = \{\{(x, y, z) : x \rightarrow y \vee z\}, \{(0)\}, \{(1)\}\}$. Note that $\Psi(x)$ is a DUAL-HORN-SAT formula if and only if $\Psi'(x) = \Psi(\neg x)$ is a HORN-SAT formula.

As previously mentioned, the computational complexity of any such Γ has been classified by the algebraic dichotomy theorem [Bul17, Zhu17]. In the Boolean case, this is known as Schaefer's theorem.

Theorem 2.3 ([Sch78]). *Let Γ be a Boolean CSP template. Then, either $\text{CSP}(\Gamma)$ is NP-complete or one of the following tractable cases. In the last four cases below, $\Psi(x)$ is true if and only if $\Psi'(x, y)$ is true some y and Ψ' can be constructed in polynomial time.*

1. *For all $\Psi(x) \in \text{CSP}(\Gamma)$, $x = (0, \dots, 0)$ is a solution.*
2. *For all $\Psi(x) \in \text{CSP}(\Gamma)$, $x = (1, \dots, 1)$ is a solution.*
3. *For all $\Psi(x) \in \text{CSP}(\Gamma)$, there is a formula $\Psi'(x, y) \in \text{CSP}(\Gamma_{2\text{-SAT}})$.*
4. *For all $\Psi(x) \in \text{CSP}(\Gamma)$, there is a formula $\Psi'(x, y) \in \text{CSP}(\Gamma_{\text{LIN-2}})$.*
5. *For all $\Psi(x) \in \text{CSP}(\Gamma)$, there is a formula $\Psi'(x, y) \in \text{CSP}(\Gamma_{\text{HORN-SAT}})$.*
6. *For all $\Psi(x) \in \text{CSP}(\Gamma)$, there is a formula $\Psi'(x, y) \in \text{CSP}(\Gamma_{\text{DUAL-HORN-SAT}})$.*

Note that the first two cases are rather superficial and case 6 (DUAL-HORN-SAT) is equivalent to HORN-SAT up to a global negation of the variables. As such, 2-SAT, LIN-2, and HORN-SAT will be the focus of our study.

2.2 Mod-CSPs

With CSPs formally defined, it is now easy to state what a Mod-CSP is. To assist the reader, we give multiple definitions of a Mod-CSP, each with an increasing level of generality.

Definition 2.4. *Let Γ be a template over the domain $\{0, 1\}$. Let $M \geq 2$ be an integer. An instance of $\text{MOD-CSP}(\Gamma, M)$ consists of an instance $\Psi(x_1, \dots, x_n)$ of $\text{CSP}(\Gamma)$ with an additional constraint,*

$$x_1 + x_2 + \dots + x_n = a \pmod{M}$$

for some integer a .

The only structure of arithmetic modulo M we are using is that $\mathbb{Z}/M\mathbb{Z}$ (the integers modulo M) is an abelian group and that $\{0, 1\}$ has a “natural” inclusion in $\mathbb{Z}/M\mathbb{Z}$. We generalize both of these aspects in the next definition.

Definition 2.5. *Let Γ be a template over a domain D . Let $(G, +, 0)$ be an abelian group. An instance of $\text{MOD-CSP}(\Gamma, G)$ consists of an instance $\Psi(x_1, \dots, x_n)$ of $\text{CSP}(\Gamma)$ and an additional constraint*

$$g_1(x_1) + \dots + g_n(x_n) = a$$

for some $a \in G$ and arbitrary maps $g_1, \dots, g_n : D \rightarrow G$.

Remark 2.6. *When Γ is a Boolean template, $\text{MOD-CSP}(\Gamma, M)$ and $\text{MOD-CSP}(\Gamma, \mathbb{Z}/M\mathbb{Z})$ are subtly different because $\text{MOD-CSP}(\Gamma, \mathbb{Z}/M\mathbb{Z})$ allows for weighted instances. Even so, these problems are polynomial-time equivalent as long as we make the mild, standard (e.g., [BKW17]) assumption that we allow equality of variables because we can simulate the term $i_j(x_j)$ by having $i_j(1) - i_j(0)$ copies of x_j all set equal to each other and subtracting $i_j(0)$ from a to account for the additive shift. Equality can be built as a gadget in 2-SAT, LIN-2, and HORN-SAT, so this assumption does not hurt our algorithmic or hardness results for those problems.*

In this paper, G is always finite, and unless otherwise specified the cardinality of G is independent of the number of variables in the instance. By the classification of finite abelian groups, G is a subgroup of $(\mathbb{Z}/M\mathbb{Z})^k$ for some positive integer M and k . Therefore $\text{MOD-CSP}(\Gamma, G)$ can then be written as an instance of $\text{CSP}(\Gamma)$ along with a system of k weighted equations modulo M . This is the perspective used in Sections 3 and 4.

For completeness, we also define a “list” version of Mod-CSPs, whether the modular constraint need not be an equality but rather an inclusion in a set of elements. Note that the list version can be reduced to the equality version by a brute force search over every element of the list. This version will only be used in the 2-SAT analysis in Section 5 and in the hardness analysis in Section A.

Definition 2.7. *Let Γ be a template over domain D . Let $(G, +)$ be an abelian group and let $S \subset G$ be nonempty. An instance of $\text{MOD-CSP}(\Gamma, G, S)$ consists of an instance Ψ of $\text{CSP}(\Gamma)$ over variables x_1, \dots, x_n with a single additional constraint*

$$g_1(x_1) + \dots + g_n(x_n) \in S,$$

for some choice of maps $g_1, \dots, g_n : D \rightarrow G$.

2.3 Lucas theorem and other polynomial preliminaries

We will collect some useful lemmas about polynomial representations of modular constraints. The reader may skip these and come back when needed.

Lemma 2.8 (Lucas’s theorem). *Let p be a prime and let a, b be non-negative integers with base p expansions given by $a = a_0 + a_1p + a_2p^2 + \dots$ and $b = b_0 + b_1p + b_2p^2 + \dots$ for some $0 \leq a_i, b_i \leq p-1$. Then*

$$\binom{a}{b} = \prod_{i \geq 0} \binom{a_i}{b_i} \pmod{p}$$

where $\binom{m}{n} = \frac{m(m-1)\dots(m-(n-1))}{n!}$ for $n > 0$ and $\binom{m}{0} = 1$ for all m .

Denote by $s_k(x)$ the k^{th} elementary symmetric polynomial given by $s_k(x) = \sum_{i_1 < i_2 < \dots < i_k} x_{i_1} x_{i_2} \dots x_{i_k}$.

Lemma 2.9. *Let p be some prime, $\ell \geq 1$ and $0 \leq a \leq p^\ell - 1$ be integers. Let $a = a_0 + a_1 2 + \dots + a_{\ell-1} p^\ell$ be the base p expansion of a . Then for every $x \in \{0, 1\}^n$,*

$$\text{Ham}(x) = a \pmod{p^\ell} \iff s_{p^t}(x) = a_t \forall 0 \leq t \leq \ell - 1.$$

Proof. Let $\text{Ham}(x) = b_0 + b_1 p + b_2 p^2 + \dots$ be the binary expansion of $\text{Ham}(x)$. $\text{Ham}(x) = a \pmod{p^\ell}$ iff $a_t = b_t$ for all $0 \leq t \leq \ell - 1$. By Lemma 2.8, $\binom{\text{Ham}(x)}{p^t} = \binom{b_t}{1} = b_t \pmod{p}$. Note that $\binom{\text{Ham}(x)}{p^t} = s_{2^t}(x)$. Thus $\text{Ham}(x) = a \pmod{p^\ell}$ iff $s_{p^t}(x) = a_t$ for every $0 \leq t \leq \ell - 1$. \square

Lemma 2.10. *Let p be some prime. For every $\ell \geq 1$ and $0 \leq a \leq p^\ell - 1$, there exists a degree $p^\ell - 1$ polynomial $\phi_{\ell, a} \in \mathbb{F}_p[x_1, \dots, x_n]$ such that for every $x \in \{0, 1\}^n$,*

$$\phi_{\ell, a}(x) \pmod{p} = \begin{cases} 0 & \text{if } \text{Ham}(x) = a \pmod{p^\ell} \\ 1 & \text{if } \text{Ham}(x) \neq a \pmod{p^\ell}. \end{cases}$$

Proof. Let $a = a_0 + a_1p + \dots + a_{\ell-1}p^{\ell-1}$ be the base p expansion of a . By Lemma 2.9 and Fermat's little theorem,

$$\phi_{\ell,a}(x) = 1 - \prod_{t=0}^{\ell-1} \left(1 - (a_t - s_{p^t}(x))^{p-1}\right)$$

is the required polynomial of degree $(p-1)(1+p+\dots+p^{\ell-1}) = p^\ell - 1$. \square

Lemma 2.11 (DeMillo-Lipton-Schwartz-Zippel Lemma). *Let $f \in \mathbb{F}_2[x_1, \dots, x_n]$ be a degree r polynomial and let $x_0 \in \mathbb{F}_2^n$ be some fixed point. Then*

$$\Pr_{x \in \mathbb{F}_2^n} [f(x) = f(x_0)] \geq \frac{1}{2^r}$$

where the probability is over a uniformly random point $x \in \mathbb{F}_2^n$.

3 HORN-SAT-MOD_M

Consider a Boolean CSP on variables x_1, \dots, x_n with three kinds of constraints: $x_i = 0$, $x_i = 1$, or $x_{i_1} \vee \neg x_{i_2} \vee \dots \vee \neg x_{i_k} = 1$ (with only the first variable not negated)**. Such a CSP is known as HORN-SAT. Fix an instance Ψ of HORN-SAT and let $\mathcal{F} \subset \{0, 1\}^n$ be the family of solutions. Think of each vector $v \in \{0, 1\}^n$ as the set $\{i : v_i = 1\}$, so \mathcal{F} is viewed as a family of subsets of $[n]$. The important property of HORN-SAT is that \mathcal{F} is an intersecting family: $A, B \in \mathcal{F}$ implies that $A \cap B \in \mathcal{F}$. †† In essence, HORN-SAT is a white-box model for studying intersecting families.

Definition 3.1. HORN-SAT-MOD_M(n) is the following algorithmic problem: Given an instance of HORN-SAT on n variables x_1, x_2, \dots, x_n along with a modular constraint $\sum_i a_i x_i = a_0 \pmod M$, decide if there is a solution in $x \in \{0, 1\}^n$.

In this section, we will present an algorithm for HORN-SAT-MOD_M and analyze its running time.

3.1 Algorithm for HORN-SAT-MOD_M

Fix an instance Ψ of HORN-SAT and let $\mathcal{F} \subset \{0, 1\}^n$ be the family of solutions. The algorithm FindMinimal shows that given any $A \subset [n]$, we can efficiently find the unique minimal $B \in \mathcal{F}$ such that $B \supset A$. Note that the uniqueness of B follows from the intersection-closed property of \mathcal{F} .

Note that FindMinimal will converge in n rounds, because in each round the Hamming weight of y will strictly increase. We are now ready to present our algorithm for HORN-SAT-MOD_M. Our algorithm uses this FindMinimal routine, and is directly inspired by the algorithm in Nagle, Sudakov, and Zenklusen [NSZ18] for submodular minimization with modular constraints. Wlog, we can assume that the coefficients $a_1, \dots, a_n = 1$ in the modular constraint. This is because, we can create a_i copies of variable x_i with equality constraints among these copies. And equality can be implemented using HORN-SAT clauses as: $a = b$ iff $a \rightarrow b$ and $b \rightarrow a$. This will increase the variables by a factor of M . For a bit vector $x \in \{0, 1\}^n$, Ham(x) denotes the Hamming weight of x i.e. number of 1's in x .

**This is equivalent to $x_{i_2} \wedge \dots \wedge x_{i_k} \rightarrow x_{i_1}$.

††This is equivalent to HORN-SAT having AND₂ as a polymorphism.

Algorithm 3.1 Algorithm FindMinimal

- Input: Instance Ψ of HORN-SAT, $x \in \{0, 1\}^n$.
 - Output: $y \in \{0, 1\}^n$ with $y \geq x$, y satisfies Ψ , and y minimal.
 - FindMinimal(x)
 1. Set $y \leftarrow x$.
 2. While y does not satisfy Ψ ,
 - (a) Find a clause $C(y_{i_1}, \dots, y_{i_k})$ which fails.
 - (b) Find the minimal $z \geq y$ which satisfies C .
 - (c) If z does not exist, return NO-SOLUTION.
 - (d) Set $y \leftarrow z$.
 3. Return y .
-

Algorithm 3.2 Algorithm for HORN-SAT-MOD $_M(n)$ with R rounds

- Input: instance Ψ of HORN-SAT with a global constraint $\sum x_i = r \pmod M$.
 - Output: either a solution $x \in \{0, 1\}^n$ or NO-SOLUTION.
 - Method:
 1. For all $x \in \{0, 1\}^n$ with $\text{Ham}(x) \leq R$, test if FindMinimal(x) returns a satisfying assignment.
 2. Else, return NO-SOLUTION.
-

We will relate the number of rounds needed in Algorithm 3.2 to certain combinatorial families called (M, r, d) -systems.

Definition 3.2 ([NSZ18]). *A collection of subsets $\mathcal{F} \subset 2^{[d]}$ is called an (M, r, d) -system if*

1. \mathcal{F} is closed under intersections i.e. if $F, G \in \mathcal{F}$ then $F \cap G \in \mathcal{F}$.
2. For every $F \in \mathcal{F}$, $|F| \not\equiv d \pmod M$.
3. For every subset $S \subset [d]$ of size at most r , there exists an $F \in \mathcal{F}$ which contains it.

Proposition 3.3. *If Algorithm 3.2 with R rounds fails on some instance of HORN-SAT-MOD $_M(n)$ then there exists an (M, R, d) -system for some $d \leq n$.*

Proof. Let Ψ be some HORN-SAT-MOD $_M(n)$ instance on which the algorithm with R rounds fails to find a solution even though one exists. Let \mathcal{F} be the set of solutions of the HORN-SAT instance in Ψ without the modulo M constraint, \mathcal{F} is an intersection-closed family. Let A be a minimal solution for Ψ that obeys the modular constraint. Therefore FindMinimal(B) $\not\subseteq A$ for all $B \subset A$ with $|B| \leq R$. Consider the family $\mathcal{F}^A = \{B : B \subset A, B \in \mathcal{F}\}$. We claim that $\mathcal{H} = \mathcal{F}^A \setminus \{A\}$ is a $(M, R, |A|)$ -system.

1. \mathcal{H} is an intersecting family since \mathcal{F}^A is an intersecting family whose universe is A .
2. $|H| \not\equiv |A| \pmod{M}$ for all $H \in \mathcal{H}$ as A is a minimal solution to Ψ .
3. Finally for all $S \subset A$ with $|S| \leq R$, there is $H \in \mathcal{H}$ with $S \subset H$ because $\text{FindMinimal}(S) \in \mathcal{H}$.

□

The following proposition from [NSZ18] gives bounds on (M, R, d) -systems when M is a prime power.

Proposition 3.4 ([NSZ18]). *Let M be a prime power. Then there does not exist an (M, R, d) -system with $R \geq M - 1$ for any d .*

We will give a simpler proof of the above proposition in Section 3.3 by a connection to polynomials representing $\text{NAND} \pmod{M}$ over $\{0, 1\}$ basis. Combining Propositions 3.3 and 3.4, we have the following corollary.

Corollary 3.5. *Let M be a prime power. Then Algorithm 3.2 with $R = M - 1$ rounds solves $\text{HORN-SAT-MOD}_M(n)$ correctly.*

When M has multiple prime factors, we show that HORN-SAT-MOD_M cannot be solved in polynomial time assuming ETH. But we believe that for any fixed M , Algorithm 3.2 with $R = n^{o(1)}$ rounds should solve HORN-SAT-MOD_M . We make a conjecture about polynomial representations of $\text{NAND} \pmod{M}$ in Section 3.3 (Conjecture 3.20) which would imply this.

Proposition 3.6. *Conjecture 3.20 implies that Algorithm 3.2 with $R = O_M(n^{o(1)})$ rounds solves $\text{HORN-SAT-MOD}_M(n)$ correctly for any M .*

3.2 HORN-SAT-MOD_M with multiple modular constraints

A natural extension of HORN-SAT-MOD_M is to allow k linear equations modulo M , which we will denote by $\text{HORN-SAT-MOD}_{M,k}$. We can show that our algorithm and its analysis can be naturally extended to show that $\text{HORN-SAT-MOD}_{M,k}(n)$ can be solved in time $n^{k(M-1)+O(1)}$ when M is a prime power. Note that once k becomes linear in n , $\text{HORN-SAT-MOD}_{M,k}$ becomes NP-hard for any $M \geq 2$, thus exponential dependence in k is necessary. A further generalization is to allow a bounded number of linear equations modulo different M 's. A clean way to capture all these generalizations is to look at HORN-SAT_G which is a HORN-SAT instance along with a linear equation with coefficients from some finite abelian group G . This is because having k equations modulo M_1, M_2, \dots, M_t is equivalent to a single equation over $G = \mathbb{Z}/M_1\mathbb{Z} \times \mathbb{Z}/M_2\mathbb{Z} \times \dots \times \mathbb{Z}/M_t\mathbb{Z}$. Our algorithm and its analysis for HORN-SAT-MOD_M can be extended easily to work for HORN-SAT_G for any finite abelian group. This is because any finite abelian group is a product of cyclic groups.

Fact 3.7 (Structure theorem of finite abelian groups). *Any finite abelian group G is a finite product of cyclic groups of prime power order i.e. $G = \prod_i \mathbb{Z}/p_i^{k_i}\mathbb{Z}$ for some prime numbers p_i (which may not be distinct) and $k_i \geq 1$.*

Instead of redoing everything for general groups, in the rest of this subsection we will reduce any instance of HORN-SAT_G for any finite abelian group to an instance of $\text{HORN-SAT}_{\mathbb{Z}/M\mathbb{Z}}$ which is the same as HORN-SAT-MOD_M without increasing the size too much.

We will need the following lemma which shows how to convert a linear equation modulo a prime power to a polynomial equation modulo a prime.

Lemma 3.8. *Let p be any prime and $k \geq 1$. Let $a_0, a_1, \dots, a_n \in \{0, 1, \dots, p^k - 1\}$ be some integers. Then there exists a polynomial $f(x_1, \dots, x_n)$ of degree $p^k - 1$ such that for every $x \in \{0, 1\}^n$,*

$$f(x) = \begin{cases} 0 \pmod p & \text{if } \sum_{i=1}^n a_i x_i = a_0 \pmod{p^k} \\ 1 \pmod p & \text{if } \sum_{i=1}^n a_i x_i \neq a_0 \pmod{p^k} \end{cases}.$$

Proof. Let \hat{x} be the vector where each x_i appears with multiplicity a_i . Then $\sum_{i=1}^n a_i x_i = \text{Ham}(\hat{x})$. By Lemma 2.10, there exists a polynomial $\phi(y)$ of degree $p^k - 1$ such that for every bit vector y ,

$$\phi(y) = \begin{cases} 0 \pmod p & \text{if } \text{Ham}(y) = a_0 \pmod{p^k} \\ 1 \pmod p & \text{if } \text{Ham}(y) \neq a_0 \pmod{p^k} \end{cases}.$$

Therefore $f(x) = \phi(\hat{x})$ is the required polynomial. \square

Proposition 3.9. *Let Ψ be an instance of $\text{HORN-SAT}_G(n)$ for some finite abelian group. Let $G = G_1 \times G_2 \times \dots \times G_t$ where each $G_i = \prod_j \mathbb{Z}/p_i^{k_{ij}} \mathbb{Z}$ for some distinct primes p_1, p_2, \dots, p_t . Let $M = \prod_{i=1}^t p_i$ and $d = \max_{i \in [t]} \left(\sum_j (p_i^{k_{ij}} - 1) \right)$. Then we can construct an instance Ψ' of $\text{HORN-SAT-MOD}_M(N)$ where $N = \binom{n}{\leq d}$ in $\text{poly}(N)$ time such that Ψ is satisfiable iff Ψ' is satisfiable.*

Proof. We have a constraint of the form $\sum_{\ell=1}^n g_\ell x_\ell = g_0$ for some $g_0, g_1, \dots, g_n \in G$. This is equivalent to a set of conditions of the form $\sum_{\ell=1}^n a_\ell x_\ell = a_0 \pmod{p_i^{k_{ij}}}$ for each cyclic component in G . By Lemma 3.8, there exists a polynomial $f_{ij}(x)$ of degree $p_i^{k_{ij}} - 1$ such that for every $x \in \{0, 1\}^n$,

$$f_{ij}(x) = \begin{cases} 0 \pmod{p_i} & \text{if } \sum_{i=1}^n a_i x_i = a_0 \pmod{p_i^{k_{ij}}} \\ 1 \pmod{p_i} & \text{if } \sum_{i=1}^n a_i x_i \neq a_0 \pmod{p_i^{k_{ij}}} \end{cases}.$$

Let $f_i(x) = 1 - \prod_j (1 - f_{ij}(x))$. The degree of f_i is $d_i = \sum_j (p_i^{k_{ij}} - 1)$. For every $x \in \{0, 1\}^n$, $\sum_i g_i x_i = g_0$ in G iff $\forall i \in [t] f_i(x) = 0 \pmod{p_i}$. We can combine the t polynomial conditions $f_i(x) = 0 \pmod{p_i}$ for $i \in [t]$ into a single polynomial condition $f(x) = 0 \pmod{M}$ by Chinese remainder theorem. The degree of f is $d = \max_{i \in [t]} d_i$.

Ψ' will have $N = \binom{n}{\leq d}$ variables corresponding to monomials of degree at most d in the variables x_1, \dots, x_n i.e. for every monomial $\prod_{i \in S} x_i$ with $|S| \leq d$, we create a variable y_S in Ψ' . Intuitively we would want $y_S = \prod_{i \in S} x_i$. To impose this, we will add HORN-SAT constraints of the form $y_S \implies y_i$ for each $i \in S$ and $\wedge_{i \in S} y_i \implies y_S$. We will set $y_\emptyset = 1$. We will also add all the original HORN-SAT constraints of Ψ into Ψ' by replacing x_i 's with y_i 's. Let $f(x) = \sum_{|S| \leq d} b_S \prod_{i \in S} x_i$ for some integer coefficients b_S . We will impose the modular constraint $\sum_{|S| \leq d} b_S y_S = 0 \pmod{M}$ to Ψ' . Now it is clear that Ψ' is satisfiable iff Ψ is satisfiable. Moreover the reduction only takes $\text{poly}(N)$ time. \square

We have the following immediate corollary.

Corollary 3.10. *Let $M = p^\ell$ be a prime power. Then $\text{HORN-SAT-MOD}_{M,k}$ can be solved in time $n^{k(M-1)(p-1)+O(1)}$.*

By analyzing Algorithm 3.2 directly, one can actually reduce the running time in the above corollary to $n^{k(M-1)+O(1)}$.

3.3 Covering number of $\text{NAND}_d \pmod M$ over $\{0, 1\}^d$

We will relate the existence of (M, r, d) -systems to polynomials which represent $\text{NAND} \pmod M$ over $\{0, 1\}$ basis.

Definition 3.11. *A polynomial $p(x_1, \dots, x_n)$ is said to represent $\text{NAND}_n \pmod M$ over $\{0, 1\}^n$ if it has integer coefficients and*

$$p(x) \begin{cases} = & 0 \pmod M \text{ if } x = \mathbf{1} \\ \neq & 0 \pmod M \text{ if } x \in \{0, 1\}^n \setminus \{\mathbf{1}\} \end{cases}$$

where $\mathbf{1}$ is the all ones vector.

We will now define the notion of covering number of polynomials.

Definition 3.12 (Covering number). *The covering number of a multilinear polynomial $p(x)$, denoted by $\text{cov}(p)$, is the minimum number of monomials of p one can choose such that every variable that appears in p appears in one of them.*

Monomials of a multilinear polynomial can be thought of subsets of variables. Thus $\text{cov}(p)$ is the usual covering number of the set system corresponding to the monomials of p . We are now ready to prove the relation between (M, r, d) -systems and polynomials which represent $\text{NAND} \pmod M$ over $\{0, 1\}$ basis. For a polynomial p , define $|p|$ as the sum of the absolute value of its coefficients.

Proposition 3.13. *Suppose there exists an (M, r, d) -system with n maximal sets, then there exists a polynomial $p(x_1, \dots, x_n)$ which represents $\text{NAND}_n \pmod M$ with $|p| \leq d + M - 1$ and $\text{cov}(p) > r$. Conversely given a polynomial $p(x_1, \dots, x_n)$ with non-negative coefficients which represents $\text{NAND}_n \pmod M$ with $|p| \leq d$ and $\text{cov}(p) > r$, there exists an (M, r, d) -system with at most n maximal sets.*

Proof. Let \mathcal{F} be an (M, r, d) -system with maximal sets F_1, \dots, F_n . Define the polynomial $\phi(x_1, x_2, \dots, x_n)$ as follows:

$$\phi(x) = \sum_{a \in [d]} \prod_{i: a \notin F_i} x_i.$$

From the definition, it is clear that $|\phi| \leq d$. For $x \in \{0, 1\}^n$,

$$|\cap_{i: x_i=0} F_i| = \sum_{a \in [d]} \prod_{i: x_i=0} \mathbf{1}(a \in F_i) = \sum_{a \in [d]} \prod_{i: a \notin F_i} x_i = \phi(x).$$

Therefore for every $x \neq \mathbf{1}$, we have $\phi(x) \neq d \pmod M$ and $\phi(\mathbf{1}) = d \pmod M$. Since any subset $S \subset [d]$ of size r is contained in some F_i , any r monomials in ϕ cannot cover all the n variables x_1, \dots, x_n i.e. $\text{cov}(\phi) > r$. Now $p(x) = \phi(x) - (d \pmod M)$ is the required polynomial which represents $\text{NAND}_n \pmod M$.

We will now show the converse. Let $p(x_1, \dots, x_n)$ be a multilinear polynomial with non-negative coefficients which represents $\text{NAND}_n \pmod M$ with $|p| \leq d$ and $\text{cov}(p) > r$. Define the map $\phi : \{0, 1\}^n \rightarrow \{0, 1\}^d$ as $\phi(x) = (\prod_{i \in T} x_i)_T$ where the coordinates of ϕ are monomials of p occurring with multiplicity equal to their coefficient in p . So,

$$\text{Ham}(\phi(x)) = \sum_{i=1}^d \phi_i(x) = p(x).$$

If $a \odot b$ is the coordinate-wise product, then $\phi(x) \odot \phi(y) = \phi(x \odot y)$. Define

$$\mathcal{F} = \{\phi(x) : x \in \{0, 1\}^n \setminus \{\mathbf{1}\}\}.$$

We claim that \mathcal{F} is an (M, r, d) -system.

1. Let $A, B \in \mathcal{F}$ and let $A = \phi(x), B = \phi(y)$ for some $x, y \in \{0, 1\}^n \setminus \{\mathbf{1}\}$. Then $A \cap B = \phi(x) \odot \phi(y) = \phi(x \odot y)$. Since $x \odot y \neq \mathbf{1}$, $A \cap B \in \mathcal{F}$.
2. $d = \text{Ham}(\phi(\mathbf{1})) = p(\mathbf{1}) = 0 \pmod{M}$. For $A \in \mathcal{F}$, $A = \phi(x)$ for some $x \in \{0, 1\}^n \setminus \{\mathbf{1}\}$. So $|A| = \text{Ham}(\phi(x)) = p(x) \neq 0 \pmod{M}$. Thus $|A| \neq d \pmod{M}$.
3. Let $S \subset [d]$ be of size at most r . Let $x \in \{0, 1\}^n$ be such that the variables appearing in the monomials $\{\phi_i(x) : i \in S\}$ are set to 1 and the rest of the variables are set to 0. Since $\text{cov}(p) > r$, $x \neq \mathbf{1}$. Then the set $\phi(x) \in \mathcal{F}$ contains S . \square

Thus proving upper bounds on the covering number of polynomials which represent $\text{NAND} \pmod{M}$ over $\{0, 1\}$ basis implies upper bounds on the number of rounds needed in Algorithm 3.2. In the next few subsections, we will focus on proving covering number upper bounds.

3.3.1 When M is a prime power

In this section, we will prove bounds on the covering number of polynomials which represent $\text{NAND} \pmod{M}$ over $\{0, 1\}$ basis when M is a prime power. We will collect some facts that we will need.

Fact 3.14. *If p is any prime, any function $f : \{0, 1\}^n \rightarrow \mathbb{F}_p$ has a unique representation as a multilinear polynomial.*

The following lemma explicitly gives the polynomial which calculates $\text{NAND}_n \pmod{p}$ exactly over $\{0, 1\}^n$.

Fact 3.15. *Suppose a multilinear polynomial $f(z_1, \dots, z_n)$ exactly represents $\text{NAND}_n \pmod{p}$ over $\{0, 1\}^n$ for some prime p i.e.*

$$f(z) \begin{cases} = 0 \pmod{p} & \text{if } z = \mathbf{1} \\ = 1 \pmod{p} & \text{if } z \in \{0, 1\}^n \setminus \{\mathbf{1}\} \end{cases}.$$

Then $f(z) = 1 - \prod_{i \in [n]} z_i$.

Proposition 3.16. *Let M be a prime power, then any polynomial $f(x_1, \dots, x_n)$ which represents $\text{NAND}_n \pmod{M}$ over $\{0, 1\}^n$ has $\text{cov}(f) \leq M - 1$.*

Proof. Let $M = p^k$ for some prime p . Wlog, we can assume that f has coefficients in $\{0, 1, \dots, M - 1\}$. Let $\Psi_f(x)$ be the vector of monomials of f where each monomial occurs with multiplicity equal to its coefficient in f . Therefore for $x \in \{0, 1\}^n$, $\text{Ham}(\Psi_f(x)) = f(x)$. By Lemma 2.10, there exists a polynomial ϕ of degree $p^k - 1 = M - 1$ such that

$$\phi(\Psi_f(x)) \begin{cases} = 0 \pmod{p} & \text{if } f(x) = 0 \pmod{p^k} \\ = 1 \pmod{p} & \text{if } f(x) \neq 0 \pmod{p^k} \end{cases}.$$

Therefore $\phi(\Psi_f(x))$ exactly represents $\text{NAND}_n \pmod{p}$ over $\{0, 1\}^n$. By Fact 3.15, $\phi(\Psi_f(x)) = 1 - \prod_{i \in [n]} x_i$. In particular, $\phi(\Psi_f(x))$ contains the monomial $\prod_{i \in [n]} x_i$. Since ϕ has degree $M - 1$, every monomial in $\phi(\Psi_f(x))$ is the product of at most $M - 1$ monomials in f . Thus there should be at most $M - 1$ monomials in f whose union contains all the variables and so $\text{cov}(f) \leq M - 1$. \square

Proof of Proposition 3.4. This follows immediately by combining Propositions 3.13 and 3.16. \square

3.3.2 When M has multiple prime factors

We will now focus on the case, when M has multiple prime factors. In this case, we will first show that there cannot be a constant bound on the covering number. For this, we need the following proposition by Barrington, Beigel and Rudich [BBR94], from the paper where they first introduced polynomial representations modulo composites. It shows that there are non-trivial low degree polynomials representing $\text{NAND} \pmod M$ if M has multiple prime factors.

Proposition 3.17 ([BBR94]). *Suppose M has r distinct prime factors. There exists an explicit degree $O_M(t^{1/r})$ polynomial which represents $\text{NAND}_t \pmod M$ over $\{0, 1\}^t$. Moreover it can be computed in time polynomial in its size.*

Proposition 3.18. *Let M be some fixed positive integer with $r \geq 2$ distinct prime factors. Then there exists a polynomial p with d monomials which represents $\text{NAND} \pmod M$ over $\{0, 1\}$ basis such that*

$$\text{cov}(f) \gtrsim_M \left(\frac{\log d}{\log \log d} \right)^{r-1}$$

for infinitely many d .

Proof. Let $p(x_1, \dots, x_t)$ be a polynomial of degree $O_M(t^{1/r})$ which represents $\text{NAND}_t \pmod M$ over $\{0, 1\}^t$ as given by Proposition 3.17. The number of monomials in p is $d \leq (\leq_{O_M} t^{1/r})$. Now note that

$$\text{cov}(p) \geq \frac{t}{\deg(p)} \gtrsim_M t^{1-1/r} \gtrsim_M (\log d / \log \log d)^{r-1}. \quad \square$$

Corollary 3.19. *Let M be some fixed positive integer with t distinct prime factors. Then there exists an (M, r, d) -system with $r \gtrsim_M \left(\frac{\log d}{\log \log d} \right)^{t-1}$ for infinitely many d .*

This addresses the open problem raised in [NSZ18], where they asked if (M, r, d) -systems exist with $r = \omega(1)$ when M has multiple prime factors. Though we show that the covering number can grow with d , we also conjecture that it shouldn't grow too quickly.

Conjecture 3.20. *Let M be some fixed constant. Any polynomial f with at most d monomials that represents $\text{NAND} \pmod M$ over $\{0, 1\}$ basis should have $\text{cov}(f) = d^{o_M(1)}$.*

We can prove a weaker form of Conjecture 3.20. We will show that the natural LP relaxation of covering number is indeed small. For this we need the following lower bound on the degree of polynomials which represent $\text{NAND} \pmod M$.

Proposition 3.21 ([TB98]). *Let M be a fixed constant with r distinct prime factors. Let $p(x_1, \dots, x_t)$ be a polynomial representing $\text{NAND}_t \pmod M$ over $\{0, 1\}^t$, then $\deg(p) \gtrsim_M (\log t)^{1/(r-1)}$.*

Fact 3.22 (Chernoff bound). *Let Z_1, Z_2, \dots, Z_n be independent random variables taking values in $[0, 1]$ and let $Z = Z_1 + Z_2 + \dots + Z_n$. Then for any $t > 1$,*

$$\Pr[Z \geq t] \leq \left(\frac{e\mathbb{E}[Z]}{t} \right)^t.$$

Proposition 3.23. *Let M be some fixed constant with r distinct prime factors and let $p(x_1, \dots, x_n)$ be a multilinear polynomial with d monomials representing $\text{NAND}_n \pmod M$ over $\{0, 1\}^n$. Then*

$$\text{cov}(p) \leq \log(n) \cdot \exp\left(C_M (\log d)^{1-1/r}\right)$$

where $C_M > 0$ is a constant depending only on M .

Proof. We will think of the monomials of p as sets $S \subset [n]$ and let \mathcal{F} be the collection of monomials in p . We are interested in the minimum set cover from \mathcal{F} which covers all of $[n]$. We can write the following linear programming relaxation for this problem.

$$\begin{aligned} \min \sum_{S \in \mathcal{F}} w_S \\ w_S \geq 0 \\ \forall i \in [n] \sum_{S \ni i} w_S \geq 1 \end{aligned} \tag{1}$$

Let L be the optimum value of the LP (1) attained for some $(w_S^*)_{S \in \mathcal{F}}$. Clearly L is a lower bound on the minimum set cover. By picking each subset $S \in \mathcal{F}$ in the cover with probability w_S^* , it is not hard to see that any fixed element in $[n]$ is covered with a constant probability and the number of sets picked is $O(L)$. By repeating this $O(\log n)$ times, with high probability, all the elements of $[n]$ will be covered. Therefore $\text{cov}(p) \leq O(L \log n)$. We will now prove an upper bound on L . We can write the dual of the LP (1) as follows:

$$\begin{aligned} \max \sum_{i \in [n]} p_i \\ p_i \geq 0 \\ \forall S \in \mathcal{F} \sum_{i \in S} p_i \leq 1 \end{aligned} \tag{2}$$

By LP duality the optimum value of the LP (2) is also L and is achieved for some $p_1^*, p_2^*, \dots, p_n^*$. Now let ρ be a random restriction the variables x_1, \dots, x_n where each x_i is set to 1 with probability $1 - p_i^*/L^\varepsilon$ where $\varepsilon > 0$ is a small constant that we will choose later. The restricted polynomial $p|_\rho$ represents $\text{NAND} \bmod M$ on the remaining variables. The expected number of remaining variables in $p|_\rho$ is $(\sum_{i \in [n]} p_i^*)/L^\varepsilon = L^{1-\varepsilon}$. So $p|_\rho$ has $\Omega(L^{1-\varepsilon})$ variables left with probability $1 - o(1)$.

Claim 3.24. $p|_\rho$ has degree $O(\frac{\log d}{\varepsilon \log L})$ with probability $1 - o(1)$.

Proof. Fix some $S \in \mathcal{F}$ and let $t = 10 \frac{\log d}{\varepsilon \log L}$. Let Z be the number of variables left in S after the random restriction. By Chernoff bound (Fact 3.22),

$$\Pr[Z \geq t] \leq \left(\frac{e\mathbb{E}[Z]}{t} \right)^t = \left(\frac{e(\sum_{i \in S} p_i^*)}{L^\varepsilon t} \right)^t \leq \left(\frac{e}{L^\varepsilon t} \right)^t = \exp(-t \log(L^\varepsilon t/e)) \leq \frac{1}{d^2}$$

By union bounding over all the d sets in \mathcal{F} , we can conclude that every monomial in $p|_\rho$ has degree at most t with probability at least $1 - 1/d$. \square

So there exists a restriction ρ such that $p|_\rho$ has degree $O(\frac{\log d}{\varepsilon \log L})$ and $\Omega(L^{1-\varepsilon})$ variables. So by Proposition 3.21, we get

$$(\log(L^{1-\varepsilon}))^{1/(r-1)} \lesssim_M \frac{\log d}{\varepsilon \log L}.$$

Choosing $\varepsilon = 1/2$, gives $L \leq \exp(C_m (\log d)^{1-1/r})$ for some constant $C_m > 0$ depending only on m . \square

So Proposition 3.23 proves Conjecture 3.20 for polynomials which represent $\text{NAND}_n \bmod M$ over $\{0, 1\}^n$ with d monomials if $n = 2^{d^{o(1)}}$. But there can be such polynomials where $n = 2^{\Omega(d)}$. Showing that the covering number is small for such polynomials is open.

3.4 Hardness of HORN-SAT-MOD_M

In this section, we will show hardness for HORN-SAT-MOD_M when M has multiple prime factors assuming exponential time hypothesis (ETH). We will show that if there are low degree polynomials representing $\text{OR}_d \bmod M$, then solving HORN-SAT-MOD_M is hard.

Conjecture 3.25 (Exponential time hypothesis (ETH) [IP01, IPZ01]). *There is no $2^{o(m)}$ time algorithm for 3-SAT with m clauses.*

Proposition 3.26. *Suppose $f(\cdot)$ is some function such that for every d , there exists a degree $f(d)$ polynomial which represents $\text{NAND}_d \bmod M$ over $\{0, 1\}^d$ which can be computed efficiently^{††}. Then assuming ETH, solving HORN-SAT-MOD_M(n) requires at least $2^{\Omega(m)} - \text{poly}(n)$ time for some m such that $f(m) \log(m/f(m)) \gtrsim \log n$.*

Proof. Choose the largest m such that $n \geq \binom{3m}{\leq 3f(m)}$, such an m will satisfy $f(m) \log(m/f(m)) = \Omega(\log n)$. Suppose $\phi(x_1, \dots, x_t) = C_1(x) \wedge C_2(x) \wedge \dots \wedge C_m(x)$ is some 3-SAT instance with m clauses and $t \leq 3m$ variables where each $C_i(x)$ depends on at most 3 variables. The variables x_1, \dots, x_t take $\{0, 1\}$ values and each $C_i(x)$ is a polynomial of degree at most 3 which takes these $\{0, 1\}$ values and outputs 1 if the i^{th} clause is satisfied and 0 if it is not. So ϕ is satisfiable iff there exists some $x \in \{0, 1\}^t$ such that $C_1(x) = \dots = C_m(x) = 1$. Now let $p(z_1, \dots, z_m)$ be a polynomial of degree $f(m)$ which represents $\text{NAND}_m \bmod M$. Then ϕ is satisfiable iff there exists some $x \in \{0, 1\}^t$ such that the polynomial $\Gamma(x) = p(C_1(x), \dots, C_m(x)) = 0 \bmod M$. The polynomial Γ has degree at most $3f(m)$ and t variables, so it has at most $N = \binom{t}{\leq 3f(m)}$ monomials. Let $\Gamma(x) = \sum_{S \subset [t]: |S| \leq 3f(m)} a_S \prod_{i \in S} x_i$. Wlog we can assume that $a_S \in 0, 1, 2, \dots, M-1$ because we only care about its values modulo M . We will now create an instance Ψ of HORN-SAT-MOD_M(N) on $N \leq n$ variables such that Ψ has a solution iff ϕ is satisfiable. The variables in Ψ will be indexed by subsets $S \subset [t]$ with $|S| \leq 3f(m)$, let us denote them by z_S . Intuitively, we would want $z_S = \prod_{i \in S} x_i$. To enforce this, we add the following HORN-SAT clauses to Ψ .

- For each variable z_S , add the clause $z_S \rightarrow z_i$ for every $i \in S$.
- For each variable z_S , add the clause $\bigwedge_{i \in S} z_i \rightarrow z_S$.

Finally, to Ψ we add the modular constraint $\sum_S a_S z_S = 0 \bmod M$. Ψ will have at most $O(tN)$ clauses and Ψ has a solution iff there exists an $x \in \{0, 1\}^n$ such that $\Gamma(x) = 0 \bmod M$. Therefore the HORN-SAT-MOD_M(N) instance Ψ has a solution iff the 3-SAT formula ϕ is satisfiable. The running time of the reduction is $\text{poly}(N)$. Therefore we can solve 3-SAT with m clauses in time $\text{poly}(N) + T(N)$ where $T(N)$ is the time it takes to solve HORN-SAT-MOD_M(N). This proves the required claim. \square

Remark 3.27. *Note that the instance of HORN-SAT produces by the ETH reduction does not have any constraints which force a variable to be a particular constant. This observation is needed for the dichotomy result in Appendix A.*

Remark 3.28. *While the obstructions to our algorithm are polynomials which represent $\text{NAND} \bmod M$ over $\{0, 1\}$ basis with high covering number, the gadgets used in the hardness proof are low degree polynomials which represent $\text{NAND} \bmod M$. Though these are clearly related, it is tempting to believe that the obstructions for the optimal algorithm should be the right gadgets to prove tight hardness results. Can we use polynomials which represent $\text{NAND} \bmod M$ over $\{0, 1\}$ basis with*

^{††}It should be computable in time which is polynomial in its size.

high covering number in the hardness reduction? Can we start with something else other than 3-SAT in the reduction?

Proposition 3.29. *For any integer $M \geq 2$, there exists a degree $\lceil d/(M-1) \rceil$ polynomial which represents $\text{NAND}_d \bmod M$ over $\{0, 1\}^d$.*

Proof. Partition the variables x_1, x_2, \dots, x_d into $M-1$ parts of size at most $d' = \lceil d/(M-1) \rceil$. We can compute the NAND of each part exactly with a degree d' polynomial of the form $1 - \prod_{i=1}^{d'} x_i$. Adding these polynomials which compute NAND on each part exactly, we get a polynomial which represents $\text{NAND}_d \bmod M$ over $\{0, 1\}^d$. \square

We have shown that Algorithm 3.2 runs in time $n^{M+O(1)}$ when M is a prime power. Combining Propositions 3.29 and 3.26 we have the following corollary, which shows that our algorithm is nearly tight assuming ETH when M is a prime power.

Corollary 3.30. *Suppose $M \leq n$. Assuming ETH, solving $\text{HORN-SAT-MOD}_M(n)$ requires at least $n^{\Omega(M/\log M)}$ time.*

We will now show hardness for M which is not a prime power.

Corollary 3.31. *Suppose M has r distinct prime factors. Assuming ETH, solving $\text{HORN-SAT-MOD}_M(n)$ requires at least $\exp(\Omega_M((\log n/\log \log n)^r))$ time.*

Proof. By Proposition 3.17, we can take $f(m) = O_M(m^{1/r})$ in Proposition 3.26. $m^{1/r} \log m = \Omega_M(n)$ implies that

$$m \gtrsim_M \left(\frac{\log n}{\log \log n} \right)^r$$

which implies the required bound. \square

When M has $r > 1$ prime factors, the lowest degree needed to represent $\text{OR}_d \bmod M$ is not well understood. The best upper known upper bound is $O_M(d^{1/r})$ as in Proposition 3.17 [BBR94]. The best lower bound on the degree is $\Omega_M((\log d)^{1/(r-1)})$ due to Barrington and Tardos [TB98]. If there is a polynomial whose degree matches this lower bound, then assuming ETH, we can get $\exp(\exp((\log n)^{1-1/r}))$ hardness for $\text{HORN-SAT-MOD}_M(n)$.

3.5 Submodular minimization with modular constraints

A function $f : 2^{[n]} \rightarrow \mathbb{R}$ is called submodular if for every $S, T \subset [n]$, $f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$. We will identify $2^{[n]}$ with $\{0, 1\}^n$ below, by identifying subsets with their indicator vectors.

Definition 3.32. $\text{SUBMOD-MIN-MOD}_M(n)$ denotes the following problem^{§§}. Given an evaluation oracle to a submodular function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ and integers $m \in \mathbb{Z}_{>0}$ and $0 \leq a_0, a_1, \dots, a_n \leq M-1$, find $\min\{f(x) : x \in \{0, 1\}^n, \sum_{i=1}^n a_i x_i = a \bmod M\}$.

[NSZ18] showed that $\text{SUBMOD-MIN-MOD}_M(n)$ can be solved in $n^{O(M)}$ time when M is a prime power. They asked if their methods can be extended to prove that $\text{SUBMOD-MIN-MOD}_M(n)$ can be solved in $n^{O_M(1)}$ for M which are not prime powers. Their algorithm and its analysis is

^{§§}[NSZ18] don't have coefficients in their original definition, they only look at $\text{Ham}(x)$. But by making copies of variables, one can reduce the more general problem with coefficients to their version.

closely related to the existence of (M, r, d) -systems and so to HORN-SAT-MOD_M . We showed that, assuming ETH, HORN-SAT-MOD_M cannot be solved in polynomial time if M is not a prime power. Can we show a similar hardness result for SUBMOD-MIN-MOD_M ? The following conjecture will imply such a hardness result.

Let $0 \leq r \leq d$ be integers and let $n = \binom{d}{\leq r}$. Let $\phi : \{0, 1\}^d \rightarrow \{0, 1\}^n$ given by $\phi(x) = (\prod_{i \in S} x_i)_{S \subseteq [d]; |S| \leq r}$. Let $\mathcal{F}_{r,d} = \phi(\{0, 1\}^d)$. Note that if we think of $\mathcal{F}_{r,d}$ as a collection of subsets of $[n]$, $\mathcal{F}_{r,d}$ is an intersection-closed family.

Conjecture 3.33. *For every $0 \leq r \leq d$ and $n = \binom{d}{\leq r}$, there exists a submodular function $f : \{0, 1\}^n \rightarrow \mathbb{Z}$ such that:*

1. $f(x)$ can be evaluated in $\text{poly}(n)$ time for every $x \in \{0, 1\}^n$.
2. For all $x \in \mathcal{F}_{r,d}$ (defined as above), $f(x) \leq -1$ and for all $x \notin \mathcal{F}_{r,d}$, $f(x) \geq 0$.

Proposition 3.34. *Suppose $\phi(\cdot)$ is some function such that for every d , there exists a degree $\phi(d)$ polynomial which represents $\text{NAND}_d \bmod M$ over $\{0, 1\}^d$ which can be efficiently computed. Then assuming ETH and Conjecture 3.33, solving $\text{SUBMOD-MIN-MOD}_M(n)$ requires at least $2^{\Omega(m)}/\text{poly}(n)$ time for some m such that $\phi(m) \log(m/\phi(m)) = \Omega(\log n)$.*

Proof. Note that the hard instances of $\text{HORN-SAT-MOD}_M(n)$ constructed in 3.26 all have the same HORN-SAT constraints, and only differ in the modular constraint. These instances need $2^{\Omega(m)} - \text{poly}(n)$ time to solve assuming ETH for some m such that $\phi(m) \log(m/\phi(m)) = \Omega(\log n)$. And the set of solutions to the HORN-SAT constraints in these hard instances is $\mathcal{F}_{3\phi(m),t}$ for some $t = O(m)$. Given an instance Ψ of $\text{HORN-SAT-MOD}_M(n)$ from these set of hard instances we can reduce it to a $\text{SUBMOD-MIN-MOD}_M(n)$ instance Ψ' where the modular constraint remains the same and the submodular function takes negative values on $\mathcal{F}_{3\phi(m),t}$ and non-negative values elsewhere, as given by Conjecture 3.33. This is a valid reduction because the value of Ψ' is negative iff Ψ is satisfiable. If $\text{SUBMOD-MIN-MOD}_M(n)$ can be solved in $T(n)$ time (assuming unit time evaluation oracle), then Ψ can be solved in $\text{poly}(n)T(n)$ time (because each evaluation oracle access now costs $\text{poly}(n)$ time). This implies that $T(n) \gtrsim 2^{\Omega(m)}/\text{poly}(n)$. \square

Corollary 3.35. *Suppose M has r distinct prime factors. Assuming ETH and Conjecture 3.33, solving $\text{SUBMOD-MIN-MOD}_M(n)$ requires at least $\exp(\Omega_M((\log n / \log \log n)^r))$ time.*

Proof. By Proposition 3.17, we can take $f(m) = O_M(m^{1/r})$ in Proposition 3.34. $m^{1/r} \log m = \Omega_M(n)$ implies that

$$m \gtrsim_M \left(\frac{\log n}{\log \log n} \right)^r$$

which implies the required bound. \square

Since the running time of the algorithm for $\text{SUBMOD-MIN-MOD}_M(n)$ from [NSZ18] depends on the existence of (M, R, d) -systems, Conjecture 3.20 will imply non-trivial algorithms for $\text{SUBMOD-MIN-MOD}_M(n)$ for any fixed M .

Proposition 3.36. *Conjecture 3.20 implies that for any fixed M , $\text{SUBMOD-MIN-MOD}_M(n)$ can be solved in $\exp(n^{o_M(1)})$ time.*

4 LIN-2-MOD_M

A LIN-2 instance is a system of linear equations modulo 2 in n variables $x_1, x_2, \dots, x_n \in \{0, 1\}$ i.e. each equation is of the form $\sum_i a_i x_i = a_0$ for some $a_0, a_1, \dots, a_n \in \{0, 1\}$. Satisfiability of a LIN-2 instance can be solved in polynomial time by Gaussian elimination.

Definition 4.1. LIN-2-MOD_M(n) is the following algorithmic problem: Given an instance of LIN-2 on n variables x_1, x_2, \dots, x_n along with a modular constraint $\sum_i a_i x_i = a_0 \pmod M$, decide if there is a solution in $x \in \{0, 1\}^n$.

In this section, we will present an algorithm for LIN-2-MOD_M and analyze its running time.

4.1 Algorithm for LIN-2-MOD_M

Wlog, we can assume that coefficients $a_1, \dots, a_n = 1$ in the modular constraint. This is because we can make a_i copies of x_i and add equality constraints among the copies. And equality is a LIN-2 constraint as $a = b$ iff $a \oplus b = 0$. Since we can assume that original coefficients $a_1, \dots, a_n \in \{0, 1, \dots, M-1\}$, this increases the number of variables by a factor of M . Consider the following algorithm for this problem which depends on the parameter R , the number of rounds. We can calculate a basis for the set of solutions of a LIN-2 instance in polynomial time, so we will start with such a basis.

Algorithm 4.1 Algorithm for LIN-2-MOD_M(n) with R rounds

- Input: An affine subspace V of \mathbb{F}_2^n given by $V = \text{span}\{v_1, v_2, \dots, v_d\} + b$ for some linearly independent vectors $v_1, \dots, v_d \in \mathbb{F}_2^n$ and some vector $b \in V$, and positive integers a, M .
 - Output: either a solution $x \in V$ with $\text{Ham}(x) = a \pmod M$ or NO-SOLUTION if no such x exists.
 - Method:
 1. If there exists a subset $S \subset [d]$ of size at most R such that $\text{Ham}(b + \sum_{i \in S} v_i) = a \pmod M$, output this solution.
 2. Else, return NO-SOLUTION.
-

We will now prove that if we choose the number of rounds R appropriately depending on n, M , then Algorithm 4.1 solves LIN-2-MOD_M(n) correctly. Since the running time of the algorithm is $O(n^R)$, the smaller the R the better. Surprisingly, the value of R required depends crucially on the prime factor decomposition of M ! Let us start with a simple proposition which shows that if Algorithm 4.1 fails, then there should be a special kind of obstruction.

Proposition 4.2. If Algorithm 4.1 with R rounds fails on an instance of LIN-2-MOD_M(n), then there exists an affine subspace U of \mathbb{F}_2^n with dimension greater than R with exactly one point $x^* \in U$ such that $\text{Ham}(x^*) = a \pmod M$.

Proof. Suppose Algorithm 4.1 failed to find a solution after R rounds. Therefore there exists a solution $x^* = b + \sum_{i \in S} v_i$ given a subset $S \subset [d]$ of size $|S| > R$ such that $\text{Ham}(x^*) = a \pmod M$. Wlog, we can assume that S has the minimum size among such sets. Now let U be the affine

subspace given by $U = b + \text{span}\{v_i : i \in S\}$. The dimension of U is $|S|$ which is greater than R . By minimality of S , every point in $y \in U$ other than x^* has $\text{Ham}(y) \not\equiv a \pmod{M}$. \square

Definition 4.3. Let n, M be some positive integers. $\mathcal{D}(n, M)$ denotes the largest dimension of an affine subspace C in \mathbb{F}_2^n such that for some $0 \leq a \leq M$, there exists exactly one point $x_0 \in C$ such that $\text{Ham}(x_0) = a \pmod{M}$.

Therefore Proposition 4.2 implies the following corollary.

Corollary 4.4. Algorithm 4.1 with R rounds solves $\text{LIN-2-MOD}_M(n)$ correctly if $R \geq \mathcal{D}(n, M)$.

The following proposition gives upper bounds on $\mathcal{D}(n, M)$ which in turn imply upper bounds on the number of rounds sufficient for our algorithm. The bounds depend crucially on the prime factor decomposition of M . Our bounds when M is a prime power are nearly tight. The bound for general M is conditional on a conjecture in additive combinatorics called the Polynomial Freiman-Ruzsa (PFR) conjecture, which we will define in Section 4.3.

Proposition 4.5. Let n, M be positive integers. Then:

1. $\mathcal{D}(n, M) \leq M - 1$ if M is a power of 2.
2. $\mathcal{D}(n, M) \leq (M - 1) \log_2(n + 1)$ if M is an odd prime power.
3. $\mathcal{D}(n, M) \leq 1 + (M' - 1) \log_2(n + 1)$ if $M = 2M'$ for some odd prime power M' .
4. $\mathcal{D}(n, M) \lesssim_\ell (M' \log n)^{2^{\ell-1}-1}$ if $M = 2^\ell M'$ for some odd prime power M' and $\ell \geq 2$ where $c > 0$ is some absolute constant.
5. $\mathcal{D}(n, M) \leq c_M n / \log n$ for some sufficiently large constant $c_M > 0$ depending only on M assuming the PFR conjecture.

We will prove Proposition 4.5 in Section 4.3 by a connection to sparsity of polynomials which represent $\text{OR} \pmod{M}$ over $\{-1, 1\}$ basis.

So when M is a power of 2, we have a polynomial time algorithm for $\text{LIN-2-MOD}_M(n)$. And when M is a product of a power of 2 and an odd prime power, we have a quasipolynomial time algorithm. For general M , we have a slightly non-trivial running time of $\exp(O(n \log \log n / \log n))$, whereas the trivial algorithm which checks every solution takes $\exp(\Omega(n))$ time. By using randomization, we can considerably speed up the above algorithms. For this we make use the following proposition, which uses an amplification trick. It allows us to conclude that if there is one solution, then there should be many solutions.

Proposition 4.6. Let V be an affine subspace of \mathbb{F}_2^n and let $N(V, a, M) = |\{x \in V : \text{Ham}(x) = a \pmod{M}\}|$. Then,

$$N(V, a, M) \neq 0 \Rightarrow N(V, a, M) \geq \frac{|V|}{2^{\mathcal{D}(n, M)+1}}.$$

Proof. Let $D = \mathcal{D}(n, M) + 1$. Wlog we can assume that the dimension of V is greater than D , otherwise the bound is trivially true. Since $N(V, a, M) \neq 0$, we can find some $x^* \in V$ such that $\text{Ham}(x^*) = a \pmod{M}$. Let $T \subset V \setminus \{x^*\}$ be the set of all points in $y \in V \setminus \{x^*\}$ such that $\text{Ham}(y) = a \pmod{M}$. Pick a random affine subspace A inside V of dimension D passing through

x^* . By the definition of $\mathcal{D}(n, M)$, there exists an other point $z \in A \setminus \{x^*\}$ such that $\text{Ham}(z) = a \pmod M$. Therefore $|T \cap A| \geq 1$. Therefore,

$$\begin{aligned} 1 &\leq \mathbb{E}_A[|T \cap A|] \\ &= \mathbb{E}_A\left[\sum_{x \in V \setminus \{x^*\}} \mathbf{1}_T(x) \mathbf{1}_A(x)\right] \\ &= \sum_{x \in V \setminus \{x^*\}} \mathbf{1}_T(x) \mathbb{E}_A[\mathbf{1}_A(x)] \\ &= \sum_{x \in V \setminus \{x^*\}} \mathbf{1}_T(x) \frac{2^D - 1}{|V| - 1} = |T| \frac{2^D - 1}{|V| - 1}. \end{aligned}$$

Therefore $N(V, a, M) = 1 + |T| \geq 1 + \frac{|V|-1}{2^D-1} \geq \frac{|V|}{2^D}$. \square

Algorithm 4.2 Randomized algorithm for $\text{LIN-2-MOD}_M(n)$ with T trials

- Input: An affine subspace V of \mathbb{F}_2^n and positive integers a, M .
 - Output: a solution $x \in V$ with $\text{Ham}(x) = a \pmod M$ or NO-SOLUTION.
 - Method:
 1. Pick a uniformly random subset of T points from V and output a solution if any of them satisfies the modular condition.
 2. Else, return NO-SOLUTION.
-

Combining Proposition 4.5 with Proposition 4.6 we get the following corollary.

Corollary 4.7. *Algorithm 4.2 with T trials outputs correctly given an instance of $\text{LIN-2-MOD}_M(n)$ with probability at least $2/3$ if $T \geq 4 \cdot 2^{\mathcal{D}(n, M)}$.*

Proof. By Proposition 4.6, if there exists a solution, then Algorithm 4.2 will find it with probability

$$1 - \left(1 - 1/2^{\mathcal{D}(n, M)+1}\right)^T \geq 1 - e^{-T/2^{\mathcal{D}(n, M)+1}} \geq 1 - e^{-2} \geq 2/3. \quad \square$$

4.2 LIN-2-MOD_M with multiple modular constraints

A natural extension of LIN-2-MOD_M is to allow k linear equations modulo M , which we will denote by $\text{LIN-2-MOD}_{M, k}$. As k becomes polynomial in n (i.e., $k \geq n^c$ for some constant $c > 0$), $\text{LIN-2-MOD}_{M, k}$ becomes \blacksquare NP-hard for any $M \geq 3$. We can show that our algorithm and its analysis can be naturally extended to show that $\text{LIN-2-MOD}_{M, k}(n)$ can be solved in time $n^{k(M-1)+O(1)}$ when M is an odd prime power. Similar to HORN-SAT_G , LIN-2_G which has a linear equation with coefficients from a finite abelian group G is the most general form of this problem. For example, $\text{LIN-2-MOD}_{M, k}$ is the same as $\text{LIN-2}_{(\mathbb{Z}/M\mathbb{Z})^k}$. Our algorithm and its analysis can be adapted for

\blacksquare This result follows as a consequence of Schaefer's theorem. If $M \geq 3$, a linear equation $\pmod M$ can simulate a 1-in-3-SAT constraint via $x_1 + x_2 + x_3 = 1 \pmod M$. Since polynomial-sized instances of 1-in-3-SAT are NP-hard, $\text{LIN-2-MOD}_{M, k}$ must be NP-hard for polynomial-sized k .

general groups. Instead of redoing everything for general groups, we will now present a reduction from LIN-2_G to LIN-2-MOD_M similar to our reduction from HORN-SAT_G to HORN-SAT-MOD_M.

Proposition 4.8. *Let Ψ be an instance of LIN-2_G(n) for some finite abelian group. Let $G = G_1 \times G_2 \times \dots \times G_t$ where each $G_i = \prod_j \mathbb{Z}/p_i^{k_{ij}} \mathbb{Z}$ for some distinct primes $p_1 < p_2 < \dots < p_t$. Let*

$$M = \begin{cases} \prod_{i=1}^t p_i & \text{if } p_1 \neq 2 \\ 2^{d_1+1} \prod_{i=2}^t p_i & \text{if } p_1 = 2 \end{cases}$$

where $d_1 = \sum_j (2^{k_{1j}} - 1)$ and $d = \max_{i \in [t]} \left(\sum_j (p_i^{k_{ij}} - 1) \right)$. Then we can construct an instance Ψ' of LIN-2-MOD_M(N) where $N = \binom{n}{\leq d}$ in poly(N) time such that Ψ is satisfiable iff Ψ' is satisfiable.

Proof. We have a constraint of the form $\sum_{\ell=1}^n g_\ell x_\ell = g_0$ for some $g_0, g_1, \dots, g_n \in G$. This is equivalent to a set of conditions of the form $\sum_{\ell=1}^n a_\ell x_\ell = a_0 \pmod{p_i^{k_{ij}}}$ for each cyclic component in G . By Lemma 3.8, there exists a polynomial $f_{ij}(x)$ of degree $p_i^{k_{ij}} - 1$ such that for every $x \in \{0, 1\}^n$,

$$f_{ij}(x) = \begin{cases} 0 \pmod{p_i} & \text{if } \sum_{i=1}^n a_i x_i = a_0 \pmod{p_i^{k_{ij}}} \\ 1 \pmod{p_i} & \text{if } \sum_{i=1}^n a_i x_i \neq a_0 \pmod{p_i^{k_{ij}}} \end{cases}$$

Let $f_i(x) = 1 - \prod_j (1 - f_{ij}(x))$. The degree of f_i is $d_i = \sum_j (p_i^{k_{ij}} - 1)$. For every $x \in \{0, 1\}^n$, $\sum_i g_i x_i = g_0$ in G iff $\forall i \in [t] f_i(x) = 0 \pmod{p_i}$. We now have two cases depending on $p_1 = 2$ or not.

Case 1: $p_1 \neq 2$

Now for $y \in \{-1, 1\}^n$, let

$$h_i(y_1, \dots, y_n) = f_i \left(\frac{1 - y_1}{2}, \dots, \frac{1 - y_n}{2} \right) \pmod{p_i}.$$

We can assume that the coefficients of h_i are in $\{0, 1, \dots, p_i - 1\}$ by inverting 2 mod p_i . We can combine the t polynomial conditions $h_i(y) = 0 \pmod{p_i}$ for $i \in [t]$ into a single polynomial condition $h(y) = 0 \pmod{M}$ for $M = p_1 p_2 \dots p_t$ by Chinese remainder theorem. The degree of h is $d = \max_{i \in [t]} d_i$.

Case 2: $p_1 = 2$

We will define h_i as before for $i \geq 2$. And define

$$h_1(y) = 2^{d_1} f_1 \left(\frac{1 - y_1}{2}, \dots, \frac{1 - y_n}{2} \right).$$

Since the degree of f_1 is d_1 , h_1 has integral coefficients. And $h_1 \left(\frac{1 - y_1}{2}, \dots, \frac{1 - y_n}{2} \right) = 0 \pmod{2}$ iff $h_1(y) = 0 \pmod{2^{d_1+1}}$. We can combine $h_1(y) = 0 \pmod{2^{d_1+1}}$ and remaining the $t - 1$ polynomial conditions $h_i(y) = 0 \pmod{p_i}$ for $2 \leq i \leq t$ into a single polynomial condition $h(x) = 0 \pmod{M}$ for $M = 2^{d_1+1} p_2 p_3 \dots p_t$ by Chinese remainder theorem. The degree of h is $d = \max_{i \in [t]} d_i$.

We now have a polynomial h of degree d such that for $x \in \{0, 1\}^n$, $\sum_{\ell=1}^n g_\ell x_\ell = g_0$ iff $h((-1)^x) = 0 \pmod{M}$ where $(-1)^x = ((-1)^{x_1}, \dots, (-1)^{x_n})$. Let

$$h((-1)^x) = \sum_{S \subset [n]: |S| \leq d} a_S (-1)^{\sum_{i \in S} x_i}.$$

We are now ready to create an instance Ψ' of LIN-2-MOD_M. Ψ' will have $N = \binom{n}{\leq d}$ variables. For every subset $S \in \binom{[n]}{\leq d}$, we create a variable z_S in Ψ' . Intuitively we would want $z_S = \sum_{i \in S} x_i$.

To impose this, we will add linear constraints of the form $z_S = \sum_{i \in S} z_i$ for every S . We will set $z_\phi = 0$. We will also add all the original linear constraints of Ψ into Ψ' by replacing x'_i 's with z'_i 's. Finally we will impose the modular constraint $\sum_{|S| \leq d} a_S(1 - 2z_S) = 0 \pmod{M}$ to Ψ' , note that $(-1)^{z_S} = (1 - 2z_S)$. Now it is clear that Ψ' is satisfiable iff Ψ is satisfiable. Moreover the reduction only takes $\text{poly}(N)$ time. \square

We have the following immediate corollary.

Corollary 4.9. *Let $M = p^\ell$ for some odd prime p . Then $\text{LIN-2-MOD}_{M,k}$ can be solved in time $n^{O(k(M-1)(p-1))}$ with high probability.*

By directly analyzing our Algorithm 4.2, one can reduce the running time to $n^{k(M-1)+O(1)}$ in the above corollary.

4.3 Sparsity of polynomials representing $\text{OR}_d \pmod{M}$ over $\{-1, 1\}^d$

In this section, we will prove the upper bounds on $\mathcal{D}(n, M)$ stated in Proposition 4.5 by a reduction to understanding the sparsity of polynomials which represent $\text{OR} \pmod{M}$ over $\{-1, 1\}$ basis (this is presented in Proposition 4.16). This reduction is only useful when M is odd. So we will start with a separate simple argument for the case when M is a power of 2, and then show how to reduce general M with an odd factor to the case when M itself is odd.

4.3.1 M is a power of 2

Proposition 4.10. $\mathcal{D}(n, 2^\ell) \leq 2^\ell - 1$.

Proof. Let C be an arbitrary affine subspace of \mathbb{F}_2^n of dimension $d \geq 2^\ell$ which contains a point $x_0 \in C$ such that $\text{Ham}(x_0) = a \pmod{2^\ell}$. By Lemma 2.10, there exists a polynomial $\phi_{\ell,a} \in \mathbb{F}_2[x_1, \dots, x_n]$ of degree $2^\ell - 1$ such that for every $x \in \mathbb{F}_2^n$, $\text{Ham}(x) = a \pmod{2^\ell}$ iff $\phi_{\ell,a}(x) = 0$. The restriction of $\phi_{\ell,a}$ to C is also a degree $2^\ell - 1$ polynomial in d variables. By Lemma 2.11,

$$\Pr_{x \in C}[\phi_{\ell,a}(x) = \phi_{\ell,a}(x_0)] \geq \frac{1}{2^{2^\ell - 1}}.$$

Since $\text{Ham}(x_0) = a \pmod{2^\ell}$, $\phi_{\ell,a}(x_0) = 0$. Therefore $|\{x \in C : \phi_{\ell,a}(x) = 0\}| \geq 2^{d - (2^\ell - 1)} \geq 2$. \square

This proves part (1) of Proposition 4.5.

4.3.2 When $M = 2^\ell M'$ for some odd M'

In this subsection, we will reduce the case when $M = 2^\ell M'$ to the case when M is odd. If $M = 2M'$ for some odd M' , then the reduction is easy.

Proposition 4.11. *Let $M = 2M'$ for some odd M' then $\mathcal{D}(n, M) \leq 1 + \mathcal{D}(n, M')$.*

Proof. Suppose C is an affine subspace in \mathbb{F}_2^n of dimension $d = \mathcal{D}(n, M)$ which contains exactly one point x_0 such that $\text{Ham}(x_0) = a \pmod{M}$. Let $C' = C \cap \{x : \text{Ham}(x) = a \pmod{2}\}$. Now C' is an affine subspace in \mathbb{F}_2^n of dimension $\geq d - 1$ such that there x_0 is the only point in C' with $\text{Ham}(x_0) = a \pmod{M'}$. Therefore $d - 1 \leq \mathcal{D}(n, M')$, which proves the claim. \square

To analyze the case when $M = 2^\ell M'$ for $\ell \geq 2$ and some odd M' , we will need the following lemma which states a low degree polynomial over \mathbb{F}_2 has a large subspace in which it is constant.

Lemma 4.12 ([CT15]). *Let $f_1, f_2, \dots, f_t \in \mathbb{F}_2[x_1, \dots, x_n]$ be polynomials of degree at most r and let $x_0 \in \mathbb{F}_2^n$ be some fixed point. Then there exists an affine subspace containing x_0 of dimension $\Omega((n/t)^{1/(r-1)})$ on which each of the f_i is constant.*

Proposition 4.13. *Let $M = 2^\ell M'$ for some $\ell \geq 2$ and odd M' . Then*

$$\mathcal{D}(n, M) \leq O(\ell) \cdot \mathcal{D}(n, M')^{2^{\ell-1}-1}.$$

Proof. Suppose C is an affine subspace in \mathbb{F}_2^n of dimension $d = \mathcal{D}(n, M)$ which contains exactly one point x_0 such that $\text{Ham}(x_0) = a \pmod{M}$. Let $a = a_0 + a_1 2 + \dots + a_{\ell-1} 2^{\ell-1}$ be the binary expansion of a . By Lemma 2.9, for every $x \in \mathbb{F}_2^n$, $\text{Ham}(x) = a \pmod{2^\ell}$ iff $s_{2^t}(x) = a_t \forall 0 \leq t \leq \ell - 1$. The restriction of s_{2^t} to C is still a degree 2^t polynomial. Therefore by Lemma 4.12, there exists an affine subspace A of C containing x_0 of dimension $\Omega((d/\ell)^{1/(2^{\ell-1}-1)})$ on which $s_1, s_2, \dots, s_{2^{\ell-1}}$ are constant. Therefore at every point $x \in A$, $s_{2^t}(x) = s_{2^t}(x_0) = a_t$ for every $0 \leq t \leq \ell - 1$. Therefore for every $x \in A$, $\text{Ham}(x) = a \pmod{2^\ell}$. Thus A contains exactly one point x such that $\text{Ham}(x) = a \pmod{M'}$, namely x_0 . Therefore $\dim(A) \leq \mathcal{D}(n, M')$ which implies that $d \leq O(\ell) \cdot \mathcal{D}(n, M')^{2^{\ell-1}-1}$. \square

4.3.3 Reduction of $\mathcal{D}(n, M)$ to sparsity of OR mod M over $\{-1, 1\}$ basis for odd M

Let's start with the definition of a polynomial representation of OR mod M over $\{-1, 1\}$ basis.

Definition 4.14. *A polynomial $p(x_1, \dots, x_d)$ is said to represent OR $_d$ mod M over $\{-1, 1\}^d$ if it has integer coefficients and*

$$p(x) \begin{cases} = 0 \pmod{M} & \text{if } x = \mathbf{1} \\ \neq 0 \pmod{M} & \text{if } x \in \{-1, 1\}^d \setminus \{\mathbf{1}\} \end{cases}$$

where $\mathbf{1}$ is the all ones vector.

To make the required reduction, we need the following proposition which relates the Hamming weights of points in a d -dimensional affine subspace to evaluations of a polynomial over $\{-1, 1\}^d$. For $\mathbf{a} = (a_1, \dots, a_\ell) \in \mathbb{F}_2^\ell$, let $(-1)^\mathbf{a}$ denote the vector $((-1)^{a_1}, \dots, (-1)^{a_\ell})$. Given a polynomial p with integer coefficients, define $|p|$ as sum of the absolute value of its coefficients. Note that the number of monomials in p is always at most $|p|$.

Proposition 4.15. *Let $\mathbf{b}, \mathbf{u}_1, \dots, \mathbf{u}_d \in \mathbb{F}_2^n$. Then there exists a multilinear polynomial $p(z_1, \dots, z_d)$ with integer coefficients and $|p| = n$ such that for every $\mathbf{y} \in \mathbb{F}_2^d$,*

$$n - 2\text{Ham}\left(\mathbf{b} + \sum_i y_i \mathbf{u}_i\right) = p((-1)^\mathbf{y}).$$

Conversely, given any multilinear polynomial $p(z_1, \dots, z_d)$ with integer coefficients and $|p| = n$, there exists $\mathbf{b}, \mathbf{u}_1, \dots, \mathbf{u}_d \in \mathbb{F}_2^n$ satisfying the above identity.

Proof. We will first start with $\mathbf{b}, \mathbf{u}_1, \dots, \mathbf{u}_d \in \mathbb{F}_2^n$ and construct such a polynomial p . Let $\mathbf{b} = (b_1, \dots, b_n)$ and $\mathbf{u}_i = (u_{i1}, \dots, u_{in})$. Define the map $\phi : \{-1, 1\}^d \rightarrow \{-1, 1\}^n$ as:

$$\phi_t(z_1, \dots, z_d) = (-1)^{b_t} \prod_{i \in [d]: u_{it}=1} z_i.$$

Then,

$$\phi_t((-1)^{\mathbf{y}}) = (-1)^{b_t} \prod_{i \in [d]: u_{it}=1} (-1)^{y_i} = (-1)^{b_t} \prod_{i \in [d]} (-1)^{y_i u_{it}} = (-1)^{b_t + \sum_{i \in [d]} y_i u_{it}}.$$

Therefore for any $\mathbf{y} \in \mathbb{F}_2^d$, $\phi((-1)^{\mathbf{y}}) = (-1)^{\mathbf{b} + \sum_{i=1}^d y_i \mathbf{u}_i}$. Define p as: $p(z) = \sum_{t=1}^n \phi_t(z)$. Note that p has integer coefficients and $|p| = n$. And finally,

$$\text{Ham}(\mathbf{b} + \sum_{i \in [d]} y_i \mathbf{u}_i) = \sum_{t=1}^n \frac{1}{2} (1 - \phi_t((-1)^{\mathbf{y}})) = \frac{1}{2} (n - p((-1)^{\mathbf{y}})).$$

To prove the converse, we just execute the steps of the above construction in reverse. Given a polynomial $p(z_1, \dots, z_d)$ with integer coefficients and $|p| = n$, let $\phi_1(z), \dots, \phi_n(z)$ be (signed, possibly repeated) monomials in z be such that $p(z) = \sum_{t=1}^n \phi_t(z)$. Now define $\mathbf{b}, \mathbf{u}_1, \dots, \mathbf{u}_d \in \mathbb{F}_2^n$ such that $\phi_t(z) = (-1)^{b_t} \prod_{i \in [d]: u_{it}=1} z_i$ is true, explicitly

$$u_{it} = \begin{cases} 1 & \text{if } \phi_t(z) \text{ contains } z_i \\ 0 & \text{else} \end{cases}.$$

Then by the same argument as above, the required identity is satisfied. \square

The following proposition shows the connection between affine subspaces which are obstructions for Algorithm 4.1 and polynomials representing $\text{OR} \pmod M$ over $\{-1, 1\}$ basis.

Proposition 4.16. *Suppose M is odd. Let $d = \mathcal{D}(n, M)$, then there exists a polynomial $f(x_1, \dots, x_d)$ with at most $n + 1$ monomials that represents $\text{OR}_d \pmod M$ over $\{-1, 1\}^d$. Conversely, given a polynomial with n monomials which represents $\text{OR}_d \pmod M$ over $\{-1, 1\}^d$, $\mathcal{D}(n', M) \geq d$ for some $n' \leq M \cdot n$.*

Proof. Suppose U is a d -dimensional affine subspace of \mathbb{F}_2^n which contains exactly one point $x^* \in U$ such that $\text{Ham}(x) = a \pmod M$. Let $\mathbf{b}, \mathbf{u}_1, \dots, \mathbf{u}_d \in \mathbb{F}_2^n$ be such that $U = \mathbf{b} + \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_d\}$. By Proposition 4.15, there exists a polynomial $p(z_1, \dots, z_d)$ with integer coefficients and at most n monomials such that for every $\mathbf{y} \in \mathbb{F}_2^d$, $n - 2\text{Ham}(\mathbf{b} + \sum_{i=1}^d y_i \mathbf{u}_i) = p((-1)^{\mathbf{y}})$. Suppose $x^* = \mathbf{b} + \sum_i y_i^* \mathbf{u}_i$, then $p(z) = (n - 2a) \pmod M$ for exactly one z in $\{-1, 1\}^d$ given by $z = z^* = (-1)^{y^*}$ (here we are using the fact that M is odd).

Define the polynomial $f(z) = p(z \odot z^*) - (n - 2a)$ where $z \odot z^*$ is the coordinate wise product. Note that $f(\mathbf{1}) = 0 \pmod M$ and $f(z) \neq 0 \pmod M$ for all $z \in \{-1, 1\}^d \setminus \{\mathbf{1}\}$. Thus $f(z)$ is a polynomial which represents $\text{OR}_d \pmod M$ over $\{-1, 1\}^d$ and f has at most $n + 1$ monomials.

To prove the converse, suppose $f(z)$ is a polynomial with n monomials which represents $\text{OR}_d \pmod M$ over $\{-1, 1\}^d$. Wlog, we can assume that the coefficients of f are in $\{0, 1, \dots, M - 1\}$. Let $n' = |f|$, by our assumption about coefficients of f , $n' \leq Mn$. By the converse part in Proposition 4.15, there exists $\mathbf{b}, \mathbf{u}_1, \dots, \mathbf{u}_d \in \mathbb{F}_2^{n'}$, such that $n' - 2\text{Ham}(\mathbf{b} + \sum_{i=1}^d y_i \mathbf{u}_i) = f((-1)^{\mathbf{y}})$ for every $\mathbf{y} \in \mathbb{F}_2^d$. Note that $\mathbf{u}_1, \dots, \mathbf{u}_d$ must be linearly independent. If not, there exists a $\mathbf{y} \neq 0$ such that $f((-1)^{\mathbf{y}}) = f(\mathbf{1}) = 0 \pmod M$ which is a contradiction.

Therefore in the d -dimensional affine subspace given by $V = \{\mathbf{b} + \sum_{i=1}^{n'} y_i \mathbf{u}_i : \mathbf{y} \in \mathbb{F}_2^d\}$, there exists exactly one point $\mathbf{x}^* \in V$ (given by $\mathbf{x}^* = \mathbf{b}$) such that $\text{Ham}(\mathbf{x}^*) = (n/2) \pmod M$. Thus $\mathcal{D}(n', M) \geq d$. \square

Because of the above proposition, if we prove sparsity lower bounds on polynomials which represent $OR \bmod M$ then we get good upper bounds on the number of rounds that will be enough in Algorithm 4.1.

4.3.4 When M is an odd prime power

Now we will show that when M is an odd prime power, a polynomial which represents $OR_d \bmod M$ over $\{-1, 1\}^d$ should have exponential number of monomials. We will collect some facts that we need to prove this.

Fact 4.17. *If p is an odd prime, any function $f : \{-1, 1\}^d \rightarrow \mathbb{F}_p$ has a unique representation as a multilinear polynomial.*

The following lemma explicitly gives the polynomial which calculates $OR_d \bmod p$ exactly.

Fact 4.18. *Suppose a multilinear polynomial $f(z_1, \dots, z_d)$ exactly represents $OR_d \bmod p$ over $\{-1, 1\}^d$ for some odd prime p i.e.*

$$f(z) \begin{cases} = 0 \pmod p & \text{if } z = \mathbf{1} \\ = 1 \pmod p & \text{if } z \in \{-1, 1\}^d \setminus \{\mathbf{1}\} \end{cases}.$$

Then $f(z)$ has 2^d monomials and explicitly given by,

$$f(z) = 1 - \prod_{i \in [d]} \frac{(1 + z_i)}{2}.$$

The following proposition provides a sparsity lower bound when M is an odd prime power.

Proposition 4.19. *Suppose M is an odd prime power. If a polynomial $f(z)$ represents $OR_d \bmod M$ over $\{-1, 1\}^d$, then f has at least $2^{d/(M-1)}$ monomials.*

Proof. Let $M = p^k$ for some odd prime p . Wlog, we can assume that the coefficients of f are in $\{0, 1, \dots, M-1\}$. Let $\Psi'_f(z)$ be the vector of monomials of f evaluated at z , where each monomial appears with multiplicity equal to its coefficient in f , and let N be its length. For $z \in \{-1, 1\}^d$, $\Psi'_f(z) \in \{-1, 1\}^N$ and $f(z) = \sum_{i \in [N]} (\Psi'_f(z))_i$. Let $\Psi_f(z)$ be the vector of the same length as $\Psi'_f(z)$ whose coordinates are given by

$$(\Psi_f(z))_i = \frac{1 + (\Psi'_f(z))_i}{2}.$$

For $z \in \{-1, 1\}^d$, $\Psi_f(z) \in \{0, 1\}^N$ and

$$\text{Ham}(\Psi_f(z)) = \sum_{i=1}^N \frac{1 + (\Psi'_f(z))_i}{2} = \frac{N}{2} + \frac{f(z)}{2}.$$

Therefore for $z \in \{-1, 1\}^d$, $\text{Ham}(\Psi_f(z)) = N/2 \pmod{p^k}$ iff $z = \mathbf{1}$.

By Lemma 2.10, there exists a polynomial ϕ of degree $p^k - 1 = M - 1$ such that

$$\phi(\Psi_f(z)) = \begin{cases} 0 \pmod p & \text{if } \text{Ham}(\Psi_f(z)) = N/2 \pmod{p^k} \\ 1 \pmod p & \text{if } \text{Ham}(\Psi_f(z)) \neq N/2 \pmod{p^k} \end{cases}.$$

Therefore $\phi(\Psi_f(z))$ exactly represents $OR_d \bmod p$ and therefore by Fact 4.18, it has 2^d monomials. Since ϕ has degree $M - 1$, the number of monomials in $\phi(\Psi_f(z))$ is at most $\binom{N'}{\leq M-1}$ where N' is the number of (distinct) monomials in f . Therefore, $N' \geq 2^{d/(M-1)}$. \square

Thus we have the following corollary which proves part (2) of Proposition 4.5.

Corollary 4.20. *Let M be an odd prime power. Then $\mathcal{D}(n, M) \leq (M - 1) \log_2(n + 1)$.*

Proof. Suppose C is an affine subspace of \mathbb{F}_2^n of dimension $d > (M - 1) \log_2(n + 1)$ which contains exactly one point x_0 such that $\text{Ham}(x_0) = a \pmod{M}$. Then by Proposition 4.15, there exists a polynomial $p(x_1, \dots, x_d)$ with at most $n + 1$ monomials such that $p(x)$ represents $OR_d \pmod{M}$ over $\{-1, 1\}^d$. Therefore by Proposition 4.19, $(n + 1) \geq 2^{d/(M-1)}$ which is a contradiction. \square

We remark that the above bound is nearly tight. Let $n = (M - 1)2^d$. The Hadamard code is a subspace of $\mathbb{F}_2^{2^d}$ of dimension d such that every non-zero point in the subspace has weight 2^{d-1} . Decompose $\mathbb{F}_2^n = \bigoplus \mathbb{F}_2^{2^d}$ where the copies of $\mathbb{F}_2^{2^d}$ are supported on mutually disjoint sets of variables. Let V be the subspace of \mathbb{F}_2^n which is the direct sum of Hadamard codes in each copy of $\mathbb{F}_2^{2^d}$. Then V has dimension $(M - 1)d = (M - 1) \log_2(n/(M - 1))$ and every non-zero point in V has weight in $\{1 \cdot 2^{d-1}, 2 \cdot 2^{d-1}, \dots, (M - 1) \cdot 2^{d-1}\}$ which is non-zero modulo M .

Combining Corollary 4.20 with Propositions 4.11 and 4.13 implies parts (3) and (4) of Proposition 4.5.

4.3.5 When M has multiple odd prime factors

We will now focus on the case when M is odd and has multiple prime factors. Unfortunately, in this case we do not know any unconditional super linear lower bounds on the sparsity of polynomials representing $OR_d \pmod{M}$ over $\{-1, 1\}^d$. But we can get a conditional super linear lower bound, assuming Polynomial Freiman-Ruzsa (PFR) conjecture which is a well-known conjecture in additive combinatorics. We achieve this by constructing matching vector families in $(\mathbb{Z}/M\mathbb{Z})$ starting from sparse representations of $OR_d \pmod{M}$ over $\{-1, 1\}^d$. We will first define matching vector families.

Definition 4.21. *A matching vector family (MVF) over $\mathbb{Z}/M\mathbb{Z}$ of rank r and size N is a collection of vectors $\mathbf{u}_1, \dots, \mathbf{u}_N \in (\mathbb{Z}/M\mathbb{Z})^r$ and $\mathbf{v}_1, \dots, \mathbf{v}_N \in (\mathbb{Z}/M\mathbb{Z})^r$ such that for every $i, j \in [N]$:*

$$\langle \mathbf{u}_i, \mathbf{v}_j \rangle \begin{cases} = 0 \pmod{M} & \text{if } i = j \\ \neq 0 \pmod{M} & \text{if } i \neq j. \end{cases}$$

MVFs over $\mathbb{Z}/M\mathbb{Z}$ of low rank and large size have found applications in many areas. They are used in the construction of constant query locally decodable codes [Yek08, Efr12, DGY11], Ramsey graphs [Gro00, Gop14], private information retrieval schemes [DG16] and secret sharing schemes [LV18]. In particular, this implies that lower bounds for constant query locally decodable codes give lower bounds on the rank of MVFs of a given size. For example, super polynomial lower bounds on the length of constant query locally decodable codes imply that the sparsity of a polynomial representing $OR_d \pmod{M}$ over $\{-1, 1\}^d$ should be $\omega(d)$. But only polynomial lower bounds on constant query locally decodable codes are known [KT00, KDW04]***. In fact, we do not even know any strong unconditional lower bounds on the rank of MVFs over $\mathbb{Z}/M\mathbb{Z}$ i.e. results of the form $N \leq \exp(o_M(r))$. But assuming the PFR conjecture, the following bound is known. We will not state the PFR conjecture here, for the precise statement see [BDL14].

Proposition 4.22 ([BDL14]). *Assuming the Polynomial Freiman-Ruzsa conjecture over $(\mathbb{Z}/M\mathbb{Z})^r$, any MVF over $\mathbb{Z}/M\mathbb{Z}$ of rank r should have size $N \leq \exp(O_M(r/\log r))$.*

***For 2-query locally decodable codes, it is known that the length of the encoding should be exponential in the message length [KDW04]. But for $q \geq 3$, the best lower bounds are only polynomial.

We are now ready to prove the super linear lower bound on the sparsity of polynomials representing $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ assuming PFR.

Proposition 4.23. *Assuming the Polynomial Freiman-Ruzsa (PFR) conjecture in $(\mathbb{Z}/M\mathbb{Z})^d$, any polynomial which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ needs to have $\Omega_M(d \log d)$ monomials.*

Proof. Given a sparse polynomial which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$, we will construct a MVF over $\mathbb{Z}/M\mathbb{Z}$ of small rank and large size. The construction is based on a similar construction due to Sudan which first appeared in [Gop09]. Suppose $p(z) = \sum_{t=1}^r a_t \prod_{i \in S_t} z_i$ is a polynomial with r monomials which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$. Define the $2^d \times 2^d$ matrix A with rows and columns indexed by $\{-1, 1\}^d$ as

$$A(z, z') = p(z \odot z') \bmod M$$

where $z \odot z'$ is the component-wise product. Note that the diagonal entries

$$A(z, z) = p(z \odot z) = p(\mathbf{1}) = 0 \bmod M$$

and the off-diagonal entries are

$$A(z, z') = p(z \odot z') \neq 0 \bmod M.$$

Moreover the rank of the matrix A is at most r since

$$A(z, z') = p(z \odot z') = \sum_t a_t \prod_{i \in S_t} z_i z'_i = \left\langle \left(a_t \prod_{i \in S_t} z_i \right)_t, \left(\prod_{i \in S_t} z'_i \right)_t \right\rangle.$$

Therefore the set of vectors $\mathbf{u}_z = (a_t \prod_{i \in S_t} z_i)_{t \in [r]} \in (\mathbb{Z}/M\mathbb{Z})^r$ and $\mathbf{v}_{z'} = (\prod_{i \in S_t} z'_i)_{t \in [r]} \in (\mathbb{Z}/M\mathbb{Z})^r$ for $z, z' \in \{-1, 1\}^d$ form a MVF over $(\mathbb{Z}/M\mathbb{Z})$ of size $N = 2^d$ of rank r . By Proposition 4.22,

$$N \leq \exp(O_M(r/\log r)) \Rightarrow d \leq O_M(r/\log r) \Rightarrow r \geq \Omega_M(d \log d). \quad \square$$

So we have the following corollary which proves part (5) of Proposition 4.5.

Corollary 4.24. *Assuming the PFR conjecture, for every positive integer M , there exists a constant c_M depending only on M such that, $\mathcal{D}(n, M) \leq c_M n / \log n$.*

Proof. Suppose C is an affine subspace of \mathbb{F}_2^n of dimension $d > c_M n / \log n$ which contains exactly one point x_0 such that $\text{Ham}(x_0) = a \bmod M$. Then by Proposition 4.15, there exists a polynomial $p(x_1, \dots, x_d)$ with at most $n + 1$ monomials such that $p(x)$ represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$. Therefore by Proposition 4.23, $(n + 1) \geq \Omega_M(d \log d)$ which is a contradiction if we choose c_M sufficiently large. \square

4.4 Hardness of LIN-2-MOD $_M$

We will now show hardness for LIN-2-MOD $_M$ when M has multiple odd prime factors assuming exponential time hypothesis (ETH). We will show that if there are low degree polynomials representing $\text{OR}_d \bmod M$, then solving LIN-2-MOD $_M$ is hard.

Proposition 4.25. *Let M be odd and suppose $f(\cdot)$ is some function such that for every d , there exists a degree $f(d)$ polynomial which represents $\text{OR}_d \bmod M$ over $\{-1, 1\}^d$ which is efficiently computable. Then assuming ETH, solving LIN-2-MOD $_M(n)$ requires at least $2^{\Omega(m)} - \text{poly}(n)$ time for some m such that $f(m) \log(m/f(m)) \gtrsim \log(n)$.*

Proof. Choose the largest m such that $\binom{3m}{\leq 3f(m)} \leq n$, such an m will satisfy $f(m) \log(m/f(m)) \gtrsim \log(n)$. Suppose $\phi(x) = C_1(x) \wedge C_2(x) \wedge \dots \wedge C_m(x)$ is some 3-SAT instance with m clauses and $t \leq 3m$ variables where each $C_i(x)$ depends on at most 3 variables. We can assume that the variables x_1, \dots, x_t take $\{-1, 1\}$ values and each $C_i(x)$ is a polynomial which takes these $\{-1, 1\}$ values and outputs 1 if the i^{th} clause is satisfied and -1 if it is not. So ϕ is satisfiable iff there exists some $x \in \{-1, 1\}^t$ such that $C_1(x) = \dots = C_m(x) = 1$. Now let $p(z_1, \dots, z_m)$ be a polynomial of degree $f(m)$ which represents $\text{OR}_m \pmod M$. Then ϕ is satisfiable iff there exists some $x \in \{-1, 1\}^t$ such that the polynomial $\Psi(x) = p(C_1(x), \dots, C_m(x)) = 0 \pmod M$. The polynomial Ψ has degree at most $3f(m)$ and so it has at most $\binom{t}{\leq 3f(m)} \leq n$ monomials. Wlog we can assume that Ψ has coefficients in $\{1, 2, \dots, M-1\}$ because we only care about its values modulo M , let us denote these coefficients by a_1, a_2, \dots, a_n . Emulating the proof of Proposition 4.15, there exists $\mathbf{u}_1, \dots, \mathbf{u}_t \in \mathbb{F}_2^n$ (which can be computed efficiently from Ψ) such that for every $\mathbf{y} \in \mathbb{F}_2^t$,

$$\Psi((-1)^{\mathbf{y}}) = \sum_{j=1}^n a_j \left(1 - 2 \left(\sum_{i=1}^t y_i \mathbf{u}_i \right)_j \right) = \sum_{j=1}^n a_j - 2 \sum_{j=1}^n a_j \left(\sum_{i=1}^t y_i \mathbf{u}_i \right)_j.$$

Therefore $\Psi(x) = 0 \pmod M$ for some $x \in \{-1, 1\}^t$ iff there exists some $\mathbf{x}' \in \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_t\}$ such that $\sum_{j=1}^n a_j (\mathbf{x}')_j = (\sum_{j=1}^n a_j)/2 \pmod M$. We can write the condition $\mathbf{x}' \in \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_t\}$ as a system of linear equations over \mathbb{F}_2 that \mathbf{x}' should satisfy, explicitly, $U^\perp \mathbf{x}' = \bar{0}$ where U^\perp is the matrix whose rows form a basis for the orthogonal complement of $\text{span}\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_t\}$.

Thus we reduced an instance of 3-SAT with m clauses to an instance of $\text{LIN-2-MOD}_M(n)$. The reduction itself takes $\text{poly}(n)$ time. By ETH, 3-SAT requires $2^{\Omega(m)}$ time. This proves that we need $2^{\Omega(m)} - \text{poly}(n)$ time to solve $\text{LIN-2-MOD}_M(n)$. \square

Remark 4.26. *Note that the gadgets we used in the hardness proof are low-degree polynomials which represent $\text{OR} \pmod M$ over $\{-1, 1\}$ basis. Whereas the obstructions to our algorithm are sparse polynomials which represent $\text{OR} \pmod M$ over $\{-1, 1\}$ basis. It is tempting to believe that the obstructions to the optimal algorithm should be the right gadgets that should be used in the hardness proof. Here is a different reduction. Start with a GAP-3LIN instance $\phi(x) = (E_1(x), \dots, E_m(x))$ over t variables where it is promised that either $1 - \varepsilon$ fraction of equations are satisfiable or less than $1/2 + \varepsilon$ fraction are satisfiable. GAP-3LIN is NP-hard and there are near-linear time reductions from 3-SAT to GAP-3LIN [MR10]. Suppose $p(z_1, \dots, z_m)$ is a polynomial over $\{-1, 1\}^m$ such that it weakly represents this (partial) threshold function modulo M . That is the values of $p(z) \pmod M$ when $\sum_i z_i \geq (1 - 2\varepsilon)m$ and when $\sum_i z_i \leq 2\varepsilon m$ are disjoint, say S_1 and S_0 respectively. Then ϕ is $(1 - \varepsilon)$ -satisfiable iff there exists some $x \in \{-1, 1\}^t$ such that $\Psi(x) = p(E_1(x), \dots, E_m(x)) \in S_1$. But note that in the $\{-1, 1\}$ basis, the sparsity of Ψ is the same as sparsity of p . Thus we get a good hardness reduction if there are sparse polynomials which weakly represent the $(1 - \varepsilon, 1/2 + \varepsilon)$ -threshold partial function. Another interesting question is, can we use the non-existence of such sparse polynomials in creating a good algorithm for LIN-2-MOD_M ?*

Proposition 4.27. *For any odd integer $M \geq 2$, there exists a degree $\lceil d/(M-1) \rceil$ polynomial which represents $\text{OR}_d \pmod M$ over $\{-1, 1\}^d$.*

Proof. Partition the variables x_1, x_2, \dots, x_d into $M-1$ parts of size at most $d' = \lceil d/(M-1) \rceil$. We can compute the OR of each part exactly with a degree d' polynomial of the form $1 - \prod_{i=1}^{d'} \left(\frac{1+x_i}{2} \right)$. Note that powers of 2 in the denominator can be inverted $\pmod M$ to get a polynomial with integer coefficients. Adding these polynomials which compute OR on each part exactly, we get a polynomial which represents $\text{OR}_d \pmod M$ over $\{-1, 1\}^d$. \square

We have shown that $\text{LIN-2-MOD}_M(n)$ can be solved in randomized time $n^{M+O(1)}$ when M is an odd prime power. Combining Propositions 4.27 and 4.25 we have the following corollary, which shows that our this running time is nearly tight assuming ETH when M is an odd prime power.

Corollary 4.28. *Suppose $3 \leq M \leq cn$ be an odd integer for some small enough constant $0 < c < 1$. Assuming ETH, solving $\text{LIN-2-MOD}_M(n)$ requires at least $n^{\Omega(M/\log M)}$ time.*

Corollary 4.29. *Suppose M has r distinct odd prime factors. Assuming ETH, solving $\text{LIN-2-MOD}_M(n)$ requires at least $\exp(\Omega_M((\log n / \log \log n)^r))$ time.*

Proof. Let $M = 2^\ell M'$ for some odd M' . Since we can reduce $\text{LIN-2-MOD}_{M'}(n)$ to $\text{LIN-2-MOD}_M(n)$ easily, it is enough to show hardness for $\text{LIN-2-MOD}_{M'}(n)$. So wlog, we can assume M is odd.

When M is odd, a polynomial which represents NAND over $\{0, 1\}$ basis can be converted into a polynomial with represents OR over $\{-1, 1\}$ basis by a linear basis change which preserves the degree. Therefore by Proposition 3.17, we can take $f(m) = O_M(m^{1/r})$ in Proposition 4.25. $m^{1/r} \log m = \Omega_M(\log n)$ implies that

$$m \gtrsim_M \left(\frac{\log n}{\log \log n} \right)^r$$

which implies the required bound. □

5 2-SAT-MOD $_M$

In this section, we present an algorithm for 2-SAT-MOD_M . The algorithm is recursive and for the recursion to work, we need to consider the more general list version of 2-SAT-MOD_M . Our algorithm works for any abelian group G in place of $\mathbb{Z}/M\mathbb{Z}$. So we will consider 2-SAT with a global modular constraint over a finite abelian group G , which we call $\text{MOD-CSP}(2\text{-SAT}, G, S)$. In this definition, S is the set of permitted values in the modular constraint.

The 2-SAT instance itself will be over variables $\{x_1, \dots, x_n\}$. We say that the set of literals are $V := \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$, where \bar{x}_i represents the negation of x_i . We let $E \subset V \times V$ be the constraints, where $(y, z) \in E$ implies that $(y \vee z)$ is a constraint.

As is standard, we can interpret a constraint $y \vee z$ as a pair of implications $\bar{y} \rightarrow z$ and $\bar{z} \rightarrow y$. As such, we consider a complementary implication digraph (V, F) where $F = \{(\bar{y}, z) : (y, z) \in E\} \cup \{(\bar{z}, y) : (y, z) \in E\}$.

For the global modular constraint, we have a constraint of the form

$$\sum_{j=1}^n g_j(x_j) \in S,$$

where $g_j : \{0, 1\} \rightarrow G$. To make it more symmetric in the literals, we can write each g_j as the sum of indicator functions $g_j(x) = g_{x_j}(x) + g_{\bar{x}_j}(\bar{x})$ such that

$$g_{x_j}(x) = \begin{cases} g_j(1) & x = 1 \\ 0 & x = 0 \end{cases} \quad g_{\bar{x}_j}(x) = \begin{cases} g_j(0) & x = 1 \\ 0 & x = 0 \end{cases}. \quad (3)$$

5.1 Preprocessing

The first step of our algorithm is to do standard preprocessing on the constraint graph. Note that if our implication graph has a cycle $y_1 \rightarrow y_2 \rightarrow \dots \rightarrow y_k \rightarrow y_1$, then we can deduce that $y_1 = y_2 = \dots = y_k$, and we can replace these variables by a single variable y and replace the corresponding g_y 's by their sum. By the duality of the implications, there must also be a component $\bar{y}_k \rightarrow \bar{y}_{k-1} \rightarrow \dots \rightarrow \bar{y}_1 \rightarrow \bar{y}_k$. Thus, whenever two vertices are merged in the constraint graph, their negations are also merged. The only exception is if a cycle contains both a y and its negation \bar{y} , in which case we can safely output NO SOLUTION. This preprocessing is described in Algorithm 5.1.

Algorithm 5.1 Preprocessing for MOD-CSP(2-SAT, G, S)

- Input: Instance Ψ of 2-SAT over x_1, \dots, x_n along with maps $g_j : \{0, 1\} \rightarrow G$ and a global constraint $\sum_{j=1}^n g_j(x_j) \in S$.
 - Output: Either NO SOLUTION or a new instance Ψ' on $\{x'_1, \dots, x'_k\}$ with an acyclic implication digraph $(V' = \{x'_1, \dots, x'_k, \bar{x}'_1, \dots, \bar{x}'_k\}, F')$ as well as new maps $g'_y : \{0, 1\} \rightarrow G, y \in V'$ such that Ψ is satisfiable iff Ψ' is satisfiable.
 - PREPROCESS
 1. Let $V := \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$, and construct $g_y, y \in V$ as described in (3).
 2. Construct the implication set F , and compute the strongly connected components in F .
 3. If any $y, \neg y$ appear in the same component, then output NO SOLUTION. Otherwise, label the components $\{C_1, \dots, C_k, \bar{C}_1, \dots, \bar{C}_k\}$, where \bar{C}_i has the complements of the literals in C_i .
 4. For $i \in \{1, \dots, k\}$
 - (a) Let x'_i and \bar{x}'_i be new variables representing C_i and \bar{C}_i .
 - (b) Let $g'_{x'_i} = \sum_{y \in C_i} g_y$ and $g'_{\bar{x}'_i} = \sum_{y \in \bar{C}_i} g_y$.
 5. For every pair of components for which there is at least one edge from one to the other, add an edge between the corresponding variables to F' .
-

It is clear that the run-time of the algorithm is $O(n + m)$. The correctness of this algorithm follows from the following claim.

Claim 5.1. *There is a bijection between solutions to Ψ and Ψ' which preserves the weights according to i and i' , respectively.*

Proof. Fix a solution $\mathbf{a} := (a_1, \dots, a_n)$ to Ψ . For each strongly connected component C_i (or \bar{C}_i), because there is a directed walk of implications from any pair of literals y, z in the same C_i , all those literals must have the same common value in that strongly connected component. Thus, $\mathbf{a}' = (a'_1, \dots, a'_k)$, where a'_i is defined to be the common value of C_i , is well-defined and the map $\mathbf{a} \mapsto \mathbf{a}'$ is injective. This \mathbf{a}' is a valid solution to Ψ' as any implication $y' \rightarrow z'$ in Ψ' is constructed from an implication $y \rightarrow z$ in Ψ . Also, as C_i and \bar{C}_i have complementary variables, x'_i and \bar{x}'_i are complementary.

By definition,

$$\begin{aligned}
& \sum_{i=1}^k g'_{x'_i}(a'_i) + \sum_{i=1}^k g'_{\bar{x}'_i}(\bar{a}'_i) \\
&= \sum_{i=1}^k \sum_{y \in C_i} g_y(a'_i) + \sum_{i=1}^k \sum_{\bar{y} \in \bar{C}_i} g_{\bar{y}}(\bar{a}'_i) \\
&= \sum_{i=1}^k \sum_{y \in C_i} g_y(a_y) + \sum_{i=1}^k \sum_{\bar{y} \in \bar{C}_i} g_{\bar{y}}(a_{\bar{y}}) \quad \left(a_y = \begin{cases} a_i & \text{if } y = x_i \\ \bar{a}_i & \text{if } y = \bar{x}_i \end{cases} \right) \\
&= \sum_{i=1}^n g_{x_i}(a_i) + \sum_{i=1}^n g_{\bar{x}_i}(\bar{a}_i).
\end{aligned}$$

For the other direction, consider a solution $\mathbf{a}' := (a'_1, \dots, a'_k)$ to Ψ' . We can lift this solution in the opposite manner by setting literal x_i equal to the literal of \mathbf{a}' corresponding to the strongly connected component to which x_i belongs. This lifting is also injective, and it preserves the property that all the constraints are satisfied, as whenever an implication $x_i \rightarrow x_j$ is needed, it exists between the corresponding components. By running the above equations in reverse, we have that this lifted solution also preserves the modular constraint.

Thus, we have a bijection between the solution sets of these problems. \square

It is also not hard to see that if there is an edge from C_i to C_j , there must also be an edge from \bar{C}_j to \bar{C}_i . This, along with logic from similar cases, shows that Ψ' does indeed encode a 2-SAT instance.

5.2 Acyclic case

With the preprocessing algorithm complete, we can now assume that the implication graph of the 2-SAT instance is a directed acyclic graph, DAG. The algorithm is presented in Algorithm 5.2.

Algorithm 5.2 DAG algorithm

- Input: Instance Ψ of 2-SAT as exhibited by an acyclic implication digraph on $V = \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$ and edge set F . Maps $g_j : \{0, 1\} \rightarrow G$ for all $j \in [n]$ and a global constraint $\sum_{j=1}^n g_j(x_j) \in S$.
- Output: Either NO SOLUTION or an assignment $x_1, \dots, x_n \in \{0, 1\}^n$ satisfying the global constraint.
- 2-SAT-DAG($V, F, \{g_v\}, S$)
 1. Define $g_v : \{0, 1\} \rightarrow G$ for all $v \in V$ as in (3).
 2. If $n = 0$, check if $0 \in S$. If so, output the empty assignment. Otherwise output NO SOLUTION.
 3. Otherwise, select a literal y with outdegree 0.
 4. Set $y = 1$; that is, let $V' = V \setminus \{y, \bar{y}\}$ and $F' = F \cap (V' \times V')$ and $S' = \{s - g_y(1) : s \in S\}$.
 5. Check if 2-SAT-DAG($V', F', \{g_v\}, S'$) has a solution, and if so, add “ $y = 1$ ” and “ $\bar{y} = 0$ ” to the assignment.
 6. Otherwise, set $y = 0$. This choice forces all literals z in the DAG for which there is a path from z to y to be assigned the value 0. Let W be the set of all such literals (including y itself), and let \bar{W} be the set of complements of literals in W . If $W \cap \bar{W} \neq \emptyset$ output NO SOLUTION.
 7. Otherwise, set $V'' = V \setminus (W \cup \bar{W})$, $F'' = F \cap (V'' \times V'')$, and $S'' = \{s - \sum_{z \in \bar{W}} g_z(1) : s \in S\} \cup \{s - g_y(1) - \sum_{z \in \bar{W} \setminus \{\bar{y}\}} g_z(1)\}$.
 8. If $|S''| = |S|$, then output NO SOLUTION.
 9. Check if 2-SAT-DAG($V'', F'', \{g_v\}, S''$) has a solution. If so, take the assignment and add “ $z = 0$ ” and “ $\bar{z} = 1$ ” for all $z \in W$, and return this solution.
 10. Otherwise, output NO SOLUTION.

Both the run-time and analysis of this algorithm take some work to analyze.

5.2.1 Correctness of Algorithm 5.2

We prove correctness by inducting on n . Clearly the case $n = 0$ is correct, as there are no variables and so the modular constraint is equivalent to $0 \in S$.

Assume the induction hypothesis is true for all $n < N$ for some positive integer N . Because the DAG is acyclic, a vertex of outdegree 0 must exist; call this literal y as in the algorithm. Note that \bar{y} must have indegree 0. Any valid assignment to the 2-SAT constraints (ignoring the modular constraint) must still be valid when y is set equal to 1 (and \bar{y} is set equal to 0). In particular, this means that setting $y = 1$, $\bar{y} = 0$ and removing both literals from the digraph will leave us with a valid DAG 2-SAT instance. By the induction hypothesis, this similar instance can be solved (with the adjusted set S') and then can be lifted back up to get a solution to Ψ and the global modular constraint. Thus, if the smaller instance has a valid solution, we are done.

Now what if the smaller instance fails to have a valid solution? Clearly then any valid solution must have $y = 0$ and $\bar{y} = 1$. Thus, any element of W (those literals which through a chain of

implications lead to y) must also have value 0, and those which belong to \bar{W} must have value 1 (since they are lead to by a chain of implications from \bar{y}). Thus, if \bar{W} and W intersect, the 2-SAT instance is inconsistent and we can safely reject.

Let $S_0'' = \{s - \sum_{z \in \bar{W}} g_z(1) : s \in S\}$. Clearly running 2-SAT-DAG on the digraph with $W \cup \bar{W}$ deleted will work using S_0'' . But recall we can flip the values of y and \bar{y} to get another valid solution. Thus, if a solution on the smaller DAG has weight in $S_1'' = \{s - g_y(1) - \sum_{z \in \bar{W} \setminus \{\bar{y}\}} g_z(1) : s \in S\}$, we can lift to a valid solution on the full DAG. If we run 2-SAT-DAG on the smaller digraph with $S'' = S_0'' \cup S_1''$, we can always lift back to a valid solution.

But, if $|S''| = |S|$, then $S_1'' = S_0'' = S''$. Thus, any valid solution with $y = 0$ is also a valid solution with $y = 1$. But, we have already ruled out in the first recursive case that no solutions with $y = 0$ exist. Thus, we can safely reject in this scenario.

When we do run on the smaller digraph, we know from the above logic that getting a modular value in S'' is equivalent to getting a modular value in $S'' \setminus S_1'' \subset S_0''$ (as any other element would imply a contradiction). Thus, as long as NO SOLUTION is not output from the recursive call, we can extend to a full solution.

5.2.2 Run-time of Algorithm 5.2

We claim that this algorithm runs in time $O((m+n) \cdot n^{|G|-|S|})$. Let $f(n, m, k)$ be an upper bound on the running time on instances with n variables, m implications, and $k = |G| - |S|$. If $k = 0$, then the modular constraint is trivial so $f(n, m, 0) = O(n+m)$ (the run-time of 2-SAT). For $k \geq 1$, note that

$$f(n, m, k) \leq f(n-1, m, k) + f(n-1, m, k-1) + O(m+n).$$

This is because, $f(n-1, m, k)$ is an upper bound on the work of the first recursive call, and $f(n-1, m, k-1)$ is an upper bound on the work of the second recursive call (if it is run). This recursion is consistent with a run-time of $O((m+n)n^k)$, as desired. This yields the following proposition.

Proposition 5.2. *Let G be an Abelian group. Then an instance of 2-SAT-MOD $_G$ on n variables and m constraints can be solved in $(n+m)^{|G|+O(1)}$ time.*

For special groups, this analysis can be improved. For example, when $G = \mathbb{F}_2^k$, then the algorithm in fact runs in $(n+m)^{O(k)}$ time.

A Establishing the Boolean Mod-CSP dichotomy

A.1 PP-reductions, Polymorphisms, and Galois correspondence

One family of simple gadget reductions from one CSP to another are known as pp-reductions. These are the gadget reductions used by Schaefer to prove the dichotomy for Boolean CSPs. They are formally defined as follows.

Definition A.1. *Let Γ and Γ' be templates over a domain D . We say that there is a primitive positive reduction from Γ' to Γ if for all $C' \in \Gamma'$ there exist $C_1, \dots, C_k \in \Gamma$ (perhaps with repetition) such that*

$$C'(x_1, \dots, x_\ell) = \exists y_1, \dots, y_{\ell'} \bigwedge_{i=1}^k C_i(z_{i,1}, \dots, z_{i,\text{arity}_i}),$$

where each $z_{i,j}$ is an $x_{i'}$ or a $y_{i'}$, allowing for repetition.

Informally a pp-reduction means that every constraint in Γ' can be expressed as a conjunction of constraints in Γ possibly with the addition of some auxiliary variables.

Note that if there is a pp-reduction from Γ' to Γ then there exists a polynomial-time reduction from $\text{CSP}(\Gamma')$ to $\text{CSP}(\Gamma)$ (in fact a logspace reduction). (e.g., [Che09, BKW17]).

As stated, it is rather difficult to determine whether there exists a pp-reduction between two templates Γ and Γ' . This issue can be resolved by looking at the *polymorphisms* of these constraint templates.

Definition A.2. *Let $C \subset D^k$ be a constraint. A polymorphism of C is a function $f : D^L \rightarrow D$ such that for all $x^1, \dots, x^L \in C$, we also have that $f(x^1, \dots, x^L) \in C$.^{†††} More pictorially (c.f., [BKW17]),*

$$\begin{array}{cccc} x_1^1 & \cdots & x_k^1 & \in C \\ x_1^2 & \cdots & x_k^2 & \in C \\ \vdots & \vdots & \vdots & \in C \\ x_1^L & \cdots & x_k^L & \in C \\ \hline f \Downarrow & y_1 & \cdots & y_k \in C \end{array}$$

The set of such polymorphisms is denoted by $\text{Pol}(C)$. For a general template Γ , the set of polymorphisms is

$$\text{Pol}(\Gamma) := \bigcap_{C \in \Gamma} \text{Pol}(C).$$

Here are a few examples (see [BKW17] for many more).

1. Let $\text{MAJ}_k : \{0, 1\}^k \rightarrow \{0, 1\}$ be the bitwise majority operator on k bits, then $\text{MAJ}_k \in \text{Pol}(2\text{-SAT})$ for all odd k .
2. Let $\text{XOR}_k : \{0, 1\}^k \rightarrow \{0, 1\}$ be the bitwise XOR on k bits, then $\text{XOR}_k \in \text{Pol}(3\text{-XOR})$ for all odd k .
3. Let $\text{AND}_k : \{0, 1\}^k \rightarrow \{0, 1\}$ be the bitwise AND operator on k bits, then $\text{AND}_k \in \text{Pol}(\text{HORN-SAT})$ for all k .

Intuitively, polymorphisms capture high-dimensional symmetries in the constraints. If the constraints have many symmetries (such as linear constraints are closed under affine operations), then the corresponding CSPs should be more likely to be tractable. This can be stated rigorously as a *Galois correspondence*.

Theorem A.3 (Galois correspondence for pp-reductions, [Jea98]). *Let Γ and Γ' be templates over a domain D . There exists a pp-reduction from Γ' to Γ if and only if $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma')$.*

Thus, to classify the computational complexity of CSPs, it suffices to classify sets of polymorphisms. Such an investigation was done by Post [Pos41] (in a slightly more general context) in the case of Boolean polymorphisms. This classification along with the Galois correspondence yields an elegant restatement of Schaefer's theorem.

^{†††}Here f acts coordinate-wise i.e. the j^{th} coordinate of $f(x^1, \dots, x^L)$ is obtained by applying f to the j^{th} coordinates of x^1, \dots, x^L .

Theorem A.4 (Schaefer’s theorem, polymorphism version [Sch78, Che09, BKW17]). *Let Γ be a Boolean template. Either $\text{CSP}(\Gamma)$ is NP-complete or it falls into one of the six following cases.*

1. $0 \in \text{Pol}(\Gamma)$, in which case “all zeros” is a solution to every instance.
2. $1 \in \text{Pol}(\Gamma)$, in which case “all ones” is a solution to every instance.
3. $\text{AND}_2 \in \text{Pol}(\Gamma)$, in which case Γ is pp-reducible to HORN-SAT.
4. $\text{OR}_2 \in \text{Pol}(\Gamma)$, in which case Γ is pp-reducible to DUAL-HORN-SAT.
5. $\text{MAJ}_3 \in \text{Pol}(\Gamma)$, in which case Γ is pp-reducible to 2-SAT.
6. $\text{XOR}_3 \in \text{Pol}(\Gamma)$, in which case Γ is pp-reducible to LIN-2.

A.2 Extension to Mod-CSPs

We now would like to take this theory of CSPs and port it to Mod-CSPs. To start, we show that the notion of pp-reduction is still meaningful for Mod-CSPs.

Claim A.5. *Fix an Abelian group G and $S \subset G$. Consider two Γ_1 and Γ_2 such that there is a pp-reduction from Γ_1 to Γ_2 , then there is a polynomial time reduction from $\text{MOD-CSP}(\Gamma_1, G, S)$ to $\text{MOD-CSP}(\Gamma_2, G, S)$.*

Proof. Consider an instance of $\text{MOD-CSP}(\Gamma_1, G, S)$ with local constraints $\Psi(x_1, \dots, x_n)$ and the global constraint

$$g_1(x_1) + \dots + g_n(x_n) \in S. \quad (4)$$

The pp-reduction says that $\Psi(x_1, \dots, x_n)$ is equivalent to $\Psi'(x_1, \dots, x_n; y_1, \dots, y_m)$, where Ψ' is a formula with constraints from Γ_2 .

Also observe that (4) is equivalent to

$$g_1(x_1) + \dots + g_n(x_n) + h_1(y_1) + \dots + h_m(y_m) \in S \quad (5)$$

where $h_1 = \dots = h_m = 0$.

Thus, Ψ' and (5) in $\text{MOD-CSP}(\Gamma_2, G, S)$ is equivalent to Ψ and (4) in $\text{MOD-CSP}(\Gamma_1, G, S)$. \square

Since Mod-CSPs are preserved under pp-reductions, Theorem A.3 tells us that Mod-CSPs with more polymorphisms are at least as tractable.

To understand the complexity of Mod-CSPs, we use a result of Post [Pos41], which is explicitly stated in [Che09]. First, we need a definition

Definition A.6 (c.f., [Che09]). *An operator $f : D^L \rightarrow D$ is essentially unary if there exists $i \in [L]$ such that $f(x_1, \dots, x_L) = f(y_1, \dots, y_L)$ whenever $x_i = y_i$.*

Note this definition says that constant functions are essentially unary. The other common example are *dictator* (or *projection*) functions: $f(x) = x_i$.

Theorem A.7 (Theorem 5.1 of [Che09]). *Let Γ be a Boolean template such that there exists $f(x_1, \dots, x_L) \in \text{Pol}(\Gamma)$ which is not essentially unary. Then, at least one of OR_2 , AND_2 , MAJ_3 , XOR_3 is in $\text{Pol}(\Gamma)$.*

Now, we show that Boolean Mod-CSPs whose polymorphisms only have essentially unary operators are NP-complete.

Lemma A.8. *Let Γ be a CSP template over a domain $D = \{0, 1\}$. Let G be an Abelian group, and let S be a nontrivial subset of G ($S \neq \emptyset$ and $S \neq G$). If $\text{Pol}(\Gamma)$ consists entirely of essentially unary operators, then $\text{MOD-CSP}(\Gamma, G, S)$ is NP-complete.*

By virtue of Schaefer’s dichotomy theorem, this result does not hold for CSPs as constant polymorphisms can lead to tractability. Thus, we need to use the global modular constraint to “break” these constant solutions.

Proof. This will be shown via a reduction from graph 3-coloring like in [BG16]. Let $C = \{1, 2, 3\}$ be the colors. Let $F = \{(1, 2), (1, 3), (2, 1), (2, 3), (3, 1), (3, 2)\}$ be all valid ways of coloring an edge.

Let (H, E) be a connected graph. For each vertex $v \in H$, construct a collection of variables $x_v(d_1, d_2, d_3)$ for all $d_1, d_2, d_3 \in D$. Likewise, for each edge $(u, v) \in E$ (think of the edge as directed so we can distinguish the two vertices), construct $y_{(u,v)}(d_{(i,j)})_{(i,j) \in F}$ for all $d_{(i,j)} \in D$ where $(i, j) \in F$.

To talk about assignments to the variables, we let $f_v : D^C \rightarrow D$ be such that $f_v(d_1, d_2, d_3)$ is the value assigned to $x_v(d_1, d_2, d_3)$. For each edge (u, v) we define $g_{u,v} : D^F \rightarrow D$ similarly.

Now, we constrain that the f ’s and g ’s are polymorphisms. Fix a $v \in H$. For any constraint $R \in \Gamma$ on k variables and for all $r^c \in R$ for all $c \in C$, we then specify that

$$(f_v(r_1^1, r_1^2, r_1^3), f_v(r_2^1, r_2^2, r_2^3), \dots, f_v(r_k^1, r_k^2, r_k^3)) \in R.$$

By definition of a polymorphism, the valid assignments to f_v are precisely the polymorphisms of Γ .

Likewise, for all $(u, v) \in E$ and $R \in \Gamma$ and for all $r^e \in R$ for $e \in F$ we specify

$$(g_{u,v}(r_1^e)_{e \in F}, \dots, g_{u,v}(r_k^e)_{e \in F}) \in R.$$

So far we haven’t linked these different polymorphisms to each other. To do that, we specify^{†††} that

$$f_u(d_1, d_2, d_3) = g_{u,v}(d_{(1,2)}, d_{(1,3)}, d_{(2,1)}, d_{(2,3)}, d_{(3,1)}, d_{(3,2)})$$

if for all $(i, j) \in F$, $d_i = d_{(i,j)}$. We likewise say that

$$f_v(d_1, d_2, d_3) = g_{u,v}(d_{(1,2)}, d_{(1,3)}, d_{(2,1)}, d_{(2,3)}, d_{(3,1)}, d_{(3,2)})$$

if for all $(i, j) \in F$, $d_j = d_{(i,j)}$. Formally, we are saying that f_u and f_v are *minors* (or *projections*) of $g_{u,v}$.

Let $0 \in G$ be the identity and pick $s_0 \in S$ (possibly 0). Let $s_1 \in G \setminus S$.

Fix $v_0 \in V$. Specify that

$$g_{x_{v_0}(0,0,0)}(d) = \begin{cases} 0 & d = 0 \\ s_1 - s_0 & d = 1 \end{cases}$$

$$g_{x_{v_0}(1,0,0)}(d) = \begin{cases} s_1 & d = 0 \\ s_0 & d = 1 \end{cases}$$

^{†††}Note that = can always be simulated by using a common variable for all the equal instances, so we do not need to add = to Γ .

Let all other i 's be 0. This completes the global constraint.

Now, we need to show that the reduction is complete and sound. For completeness, if (H, V) has a valid 3-coloring $c : H \rightarrow C$, there must be a permutation of the colors such that $c(v_0) = 1$. Consider the assignment

$$\begin{aligned} f_v(d_1, d_2, d_3) &= d_{c(v)}, & v \in H \\ g_{u,v}(d_{(i,j)}) &= d_{(c(u), c(v))} & (u, v) \in E. \end{aligned}$$

It is clear that f_v and $g_{u,v}$ are polymorphisms and that the minor constraints are satisfied. For the global constraint, observe that $g_{x_{v_0}(0,0,0)}(0) + g_{x_{v_0}(1,0,0)}(1) = 0 + s_0 \in S$.

For the soundness, imagine that the instance of MOD-CSP(Γ, G, S) has a solution. Thus, each f_v and $g_{u,v}$ either is constant or nontrivially depends on a single coordinate. It is apparent from the minor relations that if $g_{u,v}$ depends nontrivially on coordinate (i, j) then f_u must depend nontrivially on coordinate i and f_v must depend nontrivially on coordinate j . Conversely, if f_u depends nontrivially on coordinate i then $g_{u,v}$ depends nontrivially on coordinate (i, j) for some j . Similarly, if f_v depends nontrivially on coordinate j then $g_{u,v}$ depends nontrivially on (i, j) for some i .

Since (H, E) is connected, we must either have that all of the f 's and g 's are constant or they all nontrivially depend on some coordinate. If the latter case occurs, we can assign a color to each vertex $v \in H$ based on which coordinate f_v nontrivially depends on. The relations between these coordinates in the previous paragraph shows that this assignment is a valid 3-coloring.

Thus, (H, E) is 3-colorable as long as the assignment is not constant in each polymorphism. But, if the assignment is constant on each f_v , and in particular f_{v_0} , this would imply that the global constraint either satisfies

$$g_{x_{v_0}(0,0,0)}(0) + g_{x_{v_0}(1,0,0)}(0) = s_1,$$

or

$$g_{x_{v_0}(0,0,0)}(1) + g_{x_{v_0}(1,0,0)}(1) = s_1,$$

but $s_1 \notin S$, so we have a contradiction. \square

Remark A.9. Note that the size of the reduction is linear in the size of the original instance $|H| + |V|$. Since graph 3-coloring cannot be done in $2^{o(|H|+|V|)}$ time assuming ETH, we have that such MOD-CSP(Γ, G, S) cannot be done solve in $2^{o(n)} \text{poly}(m)$ time (where n is the number of variables and m is the number of constraints) assuming ETH.

Remark A.10. This result also holds for non-Boolean domains $D = \{1, \dots, k\}$. The reduction is essentially identical, except the global constraint is modified so that there are k nontrivial functions with

$$\begin{aligned} g_1(d) + \dots + g_k(d) &= s_1 \text{ for all } d \in D \\ g_1(1) + \dots + g_k(k) &= s_0, \end{aligned}$$

which is certainly possible as there are k^2 variables (k per function), but only $k + 1$ constraints.

Remark A.11. For non-Boolean CSPs, weaker conditions on $\text{Pol}(\Gamma)$ are known to imply that CSP(Γ) is NP-complete (e.g., [BKW17]). We leave as a challenge to the reader to find a suitable extension of such results to (non-Boolean) Mod-CSPs.

With these structural results for Boolean Mod-CSPs, we can now state a few dichotomy-like results.

A.3 Classification for prime powers

Theorem A.12. *Let Γ be a Boolean CSP template. Let G be a nontrivial Abelian group whose order is a prime power. Then, we have the following classification.*

1. *If one of $\text{MAJ}_3, \text{OR}_2, \text{AND}_2 \in \text{Pol}(\Gamma)$, then $\text{MOD-CSP}(\Gamma, G) \in \text{P}$.*
2. *Otherwise, if $\text{XOR}_3 \in \text{Pol}(\Gamma)$ and $|G|$ is a power of two, then $\text{MOD-CSP}(\Gamma, G) \in \text{P}$.*
3. *Otherwise, if $\text{XOR}_3 \in \text{Pol}(\Gamma)$, then $\text{MOD-CSP}(\Gamma, G) \in \text{RP} \cap \text{QP}$.*
4. *Otherwise, $\text{MOD-CSP}(\Gamma, G)$ is NP-complete.*

Proof. We prove the cases in order.

1. If $\text{MAJ}_3 \in \text{Pol}(\Gamma)$, then by Schaefer's theorem there is a pp-reduction from $\text{CSP}(\Gamma)$ to 2-SAT. Thus, there is a pp-reduction from $\text{MOD-CSP}(\Gamma, G)$ to 2-SAT-MOD $_G$, which can be solved in polynomial time by Section 5.

If $|G|$ is a prime power and $\text{AND}_2 \in \text{Pol}(\Gamma)$, then by Schaefer's theorem there is a pp-reduction from $\text{MOD-CSP}(\Gamma, G)$ to HORN-SAT-MOD $_G$, which can be solved in polynomial time by Corollary 3.5. Likewise, if $\text{OR}_2 \in \text{Pol}(\Gamma)$, then there is a pp-reduction from $\text{MOD-CSP}(\Gamma, G)$ to DUAL-HORN-SAT-MOD $_G$ which can also be solved in polynomial time by the same theorem, since solving HORN-SAT and DUAL-HORN-SAT instances are equivalent.

2. If $|G|$ is a power of two and $\text{XOR}_3 \in \text{Pol}(\Gamma)$, by Schaefer's theorem this is a pp-reduction from $\text{MOD-CSP}(\Gamma, G)$ to LIN-2-MOD $_G$, which can be solved in deterministic polynomial time by Corollary 4.4.
3. If $|G|$ is an odd prime power and $\text{XOR}_3 \in \text{Pol}(\Gamma)$, by Schaefer's theorem this is a pp-reduction from $\text{MOD-CSP}(\Gamma, G)$ to LIN-2-MOD $_G$, which can be solved in randomized polynomial time and deterministic quasi-polynomial time by Proposition 4.7 and Corollary 4.4.
4. By Theorem A.7, the only polymorphisms of Γ are essentially unary. Thus, by Lemma A.8, $\text{MOD-CSP}(\Gamma, G)$ is NP-complete. \square

A.4 Partial classification for non-prime powers

Recall, that when we motivated Mod-CSPs, we said in the Boolean case, there are essentially only three nontrivial cases: HORN-SAT-MOD $_G$, LIN-2-MOD $_G$, and 2-SAT-MOD $_G$. As shown in Theorem A.12, this view is correct when $|G|$ is a prime power, as all three problems admit polynomial time algorithms. When $|G|$ is a non-prime power, the general classification is a bit more complicated. In particular, although HORN-SAT-MOD $_G$ fails to have a polynomial-time algorithm, a special case of the problem does.

To define this special case, consider the operators

$$\begin{aligned} \text{ANDOR}(x, y, z) &= x \wedge (y \vee z) \\ \text{ORAND}(x, y, z) &= x \vee (y \wedge z). \end{aligned}$$

Note that $\text{ANDOR}(x, y, y) = \text{AND}_2(x, y)$ and $\text{ORAND}(x, y, y) = \text{OR}_2(x, y)$, so any CSP with one of these as a polymorphism is pp-reducible to either HORN-SAT or DUAL-HORN-SAT. But, in

the case where the group G does not have prime power order, HORN-SAT and DUAL-HORN-SAT have lower bounds away from P (assuming ETH). Thus, we need to study such problems separately. By the classification of Theorem A.17, understanding the case of ANDOR and ORAND will be the key to completing the complexity classification of Boolean Mod-CSPs in the case of non-prime power moduli.

The corresponding CSP for ANDOR is rather simple.

Proposition A.13 (e.g., [Pos41, CKZ08, ABI⁺09]). *Let Γ be a Boolean CSP template. If $\text{ANDOR} \in \text{Pol}(\Gamma)$, then Γ is pp-reducible to a template Λ_{ANDOR} with constraints of the form*

$$\Lambda_{\text{ANDOR}} = \left\{ \{(1)\}, \{(x, y) : x \rightarrow y\} \bigcup_{k=1}^{\infty} \{(x_1, \dots, x_k) : \neg x_1 \vee \neg x_2 \vee \dots \vee \neg x_k\} \right\}$$

Remark A.14. *This corresponds to the clone S_{00} and co-clone IS_{00} in Post's lattice.*

For ORAND, the corresponding template is the negation of the above

$$\Lambda_{\text{ORAND}} = \left\{ \{(0)\}, \{(x, y) : x \rightarrow y\} \bigcup_{k=1}^{\infty} \{(x_1, \dots, x_k) : x_1 \vee x_2 \vee \dots \vee x_k\} \right\}$$

We now show that both of these problems are tractable for any Abelian group G .

Lemma A.15. *For all finite Abelian groups G , $\text{MOD-CSP}(\Lambda_{\text{ANDOR}}, G)$ and $\text{MOD-CSP}(\Lambda_{\text{ORAND}}, G)$ are tractable in $O((n+m)^{|G|})$ time.*

Proof. First, note that these two problems are equivalent up to flipping 0 and 1. As a result, we restrict attention to $\text{MOD-CSP}(\Lambda_{\text{ORAND}}, G)$.

The overall algorithm is rather similar to the one for 2-SAT-MOD $_G$ with the auxiliary set $S \subset G$. As a result, we only state the major differences.

Consider the directed graph spanned by the implications. Like in the algorithm for 2-SAT-MOD $_G$ we can contract the strongly connected components to single variables. Note that this contract operation preserves $x_1 \vee \dots \vee x_k$ (although it may reduce to a smaller number of variables). Also, if any variables are forced in value, we can propagate that information through the digraph and update the modular constraint.

Now, take a vertex x of the digraph which has zero outdegree. Set $x = 1$, and solve the remaining instance. If a solution is found, then quit. Otherwise, set $x = 0$, but observe like in the 2-SAT algorithm, any solution with $x = 0$ yields another solution when x is set back to 1. Thus, we can either expand S in the $x = 0$ branch, or if S does not expand, we can skip the branch entirely. Thus, we get a $O((n+m)^{|G|})$ algorithm like for 2-SAT. \square

With this algorithmic result, we need another fact about Post's lattice. Stating this result requires us to define a couple of variants of HORN-SAT and LIN-2.

- HORN-SAT-NO-CONSTANTS are instances of HORN-SAT without any constraints of the form $x = 0$ or $x = 1$.
- DUAL-HORN-SAT-NO-CONSTANTS are instances of DUAL-HORN-SAT without any constraints of the form $x = 0$ or $x = 1$.

- LIN-2-EVEN-ZERO are instances of LIN-2 where every linear constraint is of the form $x_{i_1} \oplus \dots \oplus x_{i_k} = 0$, where k is even.

Note that the hardness (assuming ETH) of MOD-CSP(HORN-SAT-NO-CONSTANTS, G) and MOD-CSP(DUAL-HORN-SAT-NO-CONSTANTS, G) follow from the proof of Proposition 3.26, as no constants are specified in the constructed instance.

The hardness of MOD-CSP(LIN-2-EVEN-ZERO, G) is a bit more technical and requires a slight modification of the proof of hardness of LIN-2-MOD $_M$ in Proposition 4.25.

Claim A.16. *If $|G|$ is divisible by $r \geq 2$ distinct odd primes, then MOD-CSP(LIN-2-EVEN-ZERO, G) requires $\exp(\Omega_{|G|}((\log n / \log \log n)^r))$ time assuming ETH.*

Proof. Let M be a product of the r distinct odd primes dividing $|G|$. There is a subgroup of G isomorphic to $\mathbb{Z}/M\mathbb{Z}$, so it suffices to prove the hardness of MOD-CSP(LIN-2-EVEN-ZERO, $\mathbb{Z}/M\mathbb{Z}$). The proof is very similar to the proof of Proposition 4.25, so will only sketch the main differences. Let m be some positive integer to be chosen later. By Proposition 3.17, there exists a degree $O(m^{1/r})$ polynomial $q(z_1, z_2, \dots, z_m)$ which represents OR mod M over $\{-1, 1\}^m$ i.e. $q(z) = 0 \pmod M$ iff $z = \mathbf{1}$. Let a_0 be such that $2a_0$ is not a quadratic residue modulo M i.e. there doesn't exist any b such that $2a_0 = b^2 \pmod M$. Now define

$$p(z_0, z_1, \dots, z_m) = a_0 + z_0(q(z_1, \dots, z_m)^2 - a_0).$$

Now we claim that p represents OR mod M over $\{-1, 1\}^{m+1}$. If $z_0 = 1$, then $p(1, z_1, \dots, z_m) = q(z_1, \dots, z_m)^2 = 0 \pmod M$ iff $z_1, \dots, z_m = 1$. If $z_0 = -1$, then $p(-1, z_1, \dots, z_m) = 2a_0 - q(z_1, \dots, z_m)^2 \neq 0 \pmod M$ for any $z_1, \dots, z_m \in \{-1, 1\}$.

Now let Ψ be a 3-SAT instance with m clauses $C_1(x), \dots, C_m(x)$ and $t \leq 3m$ variables x_1, \dots, x_t . We can assume that the variables take $\{-1, 1\}$ values and $C_i(x)$ is a degree 3 polynomial which takes these $\{-1, 1\}$ values and outputs 1 if the clause is satisfied and -1 if not. Now consider the polynomial $\Gamma(x_0, x_1, \dots, x_t) = p(x_0, C_1(x), \dots, C_m(x))$. It is easy to see that Ψ is satisfiable iff there exists $x_0, x_1, \dots, x_t \in \{-1, 1\}$ such that $\Gamma(x_0, x_1, \dots, x_t) = 0 \pmod M$. Let the degree of Γ be $\Delta \leq O(m^{1/r})$. Let

$$\Gamma(x_0, x_1, \dots, x_t) = \gamma_0 + \gamma_S \sum_{S \subset \{0, 1, \dots, t\}, 1 \leq |S| \leq \Delta} \prod_{i \in S} x_i.$$

Note that every non-constant monomial of Γ has the variable x_0 in it. When we convert it to an instance of LIN-2-MOD $_M$ as in the proof of Proposition 4.25, then the subspace is spanned by $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_t$ where $\mathbf{u}_0 = \mathbf{1}$ as x_0 appears in every non-constant monomial. When we write $\mathbf{x}' \in \text{span}\{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_t\}$ as $U^\perp \mathbf{x}' = 0$, all the rows of U^\perp are orthogonal to $\mathbf{u}_0 = \mathbf{1}$ and therefore have even number of 1s. Thus we get an instance of LIN-2 where every equation has an even number of variables and the constant term is zero, which is precisely LIN-2-EVEN-ZERO. The final instance has $n \leq \binom{t+1}{\leq \Delta} \leq \binom{3m}{\leq O(m^{1/r})}$ variables. Therefore we can choose $m = \Omega_M((\log n / \log \log n)^r)$. \square

From these, we can classify a slice of Post's lattice.

Theorem A.17 (e.g., [Pos41, CKZ08, ABI⁺09]). *Let Γ be a Boolean CSP. Then, if none of ANDOR, ORAND, MAJ $_3 \in \text{Pol}(\Gamma)$, then one of the problems HORN-SAT-NO-CONSTANTS, DUAL-HORN-SAT-NO-CONSTANTS or LIN-2-EVEN-ZERO is pp-reducible to Γ .*

From this, we can now prove the classification for the non-prime-power case.

Theorem A.18. *Let Γ be a Boolean CSP template. Let G be a nontrivial Abelian group, such that $|G| = p_1^{e_1} \dots p_r^{e_r}$ where $p_1 < p_2 < \dots < p_r$ are primes, $e_1, \dots, e_r \geq 1$ and $r \geq 2$. Then, we have the following classification.*

1. *If one of MAJ₃, ANDOR, ORAND \in Pol(Γ), then MOD-CSP(Γ, G) \in P.*
2. *Otherwise, if $p_1 = 2$ and $r = 2$ and XOR₃ \in Pol(Γ). Then, MOD-CSP(Γ, G) \in QP.*
3. *Otherwise, if one of XOR₃, AND₂, OR₂ is in Pol(Γ), then MOD-CSP(Γ, G) cannot be solved in less than quasi-polynomial time, assuming the exponential time hypothesis.*
4. *Otherwise, MOD-CSP(Γ, G) is NP-complete.*

Proof. Again, we prove the results in order.

1. This follows from Section 5 and Lemma A.15.
2. This follows from Corollary 4.4.
3. By Theorem A.17, we have that one of HORN-SAT-NO-CONSTANTS, DUAL-HORN-SAT-NO-CONSTANTS, or LIN-2-EVEN-ZERO pp-reduces to Γ , so the quasi-polynomial lower bounds (assuming ETH) of Corollary 4.24 and Corollary 3.17 apply.
4. By Theorem A.7 and Lemma A.8, we have this NP-completeness result. □

References

- [ABI⁺09] Eric Allender, Michael Bauland, Neil Immerman, Henning Schnoor, and Heribert Vollmer. The complexity of satisfiability problems: Refining schaefer’s theorem. *Journal of Computer and System Sciences*, 75(4):245–254, 2009.
- [AWZ17] Stephan Artmann, Robert Weismantel, and Rico Zenklusen. A strongly polynomial algorithm for bimodular integer linear programming. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1206–1219. ACM, 2017.
- [BBR94] David A Mix Barrington, Richard Beigel, and Steven Rudich. Representing boolean functions as polynomials modulo composite numbers. *Computational Complexity*, 4(4):367–382, 1994.
- [BDL14] Abhishek Bhowmick, Zeev Dvir, and Shachar Lovett. New bounds for matching vector families. *SIAM Journal on Computing*, 43(5):1654–1683, 2014.
- [BG16] Joshua Brakensiek and Venkatesan Guruswami. New hardness results for graph and hypergraph colorings. In Ran Raz, editor, *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, volume 50 of *LIPICs*, pages 14:1–14:27. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- [BJK05] A. Bulatov, P. Jeavons, and A. Krokhin. Classifying the Complexity of Constraints Using Finite Algebras. *SIAM Journal on Computing*, 34(3):720–742, January 2005.

- [BKW17] Libor Barto, Andrei A. Krokhin, and Ross Willard. Polymorphisms, and how to use them. In Andrei A. Krokhin and Stanislav Zivny, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
- [BM10] Andrei A. Bulatov and Daniel Marx. The complexity of global cardinality constraints. *Logical Methods in Computer Science*, Volume 6, Issue 4, October 2010.
- [Bul06] Andrei A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, 2006.
- [Bul17] Andrei A. Bulatov. A dichotomy theorem for nonuniform csps. In Umans [Uma17], pages 319–330.
- [CGPT06] Arkadev Chattopadhyay, Navin Goyal, Pavel Pudlák, and Denis Thérien. Lower bounds for circuits with mod_m gates. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pages 709–718, 2006.
- [Che09] Hubie Chen. A Rendezvous of Logic, Complexity, and Algebra. *ACM Comput. Surv.*, 42(1):2:1–2:32, December 2009.
- [CKZ08] Nadia Creignou, Phokion G Kolaitis, and Bruno Zanuttini. Structure identification of boolean relations and plain bases for co-clones. *Journal of Computer and System Sciences*, 74(7):1103–1115, 2008.
- [CN10] Stephen Cook and Phuong Nguyen. *Logical foundations of proof complexity*, volume 11. Cambridge University Press Cambridge, 2010.
- [CSS10] Nadia Creignou, Henning Schnoor, and Ilka Schnoor. Nonuniform boolean constraint satisfaction problems with cardinality constraint. *ACM Trans. Comput. Logic*, 11(4):24:1–24:32, July 2010.
- [CT15] Gil Cohen and Avishay Tal. Two structural results for low degree polynomials and applications. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, page 680, 2015.
- [CW09] Arkadev Chattopadhyay and Avi Wigderson. Linear systems over composite moduli. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia USA*, pages 43–52, 2009.
- [DG16] Zeev Dvir and Sivakanth Gopi. 2-server pir with subpolynomial communication. *Journal of the ACM (JACM)*, 63(4):39, 2016.
- [DGY11] Zeev Dvir, Parikshit Gopalan, and Sergey Yekhanin. Matching vector codes. *SIAM Journal on Computing*, 40(4):1154–1178, 2011.
- [Efr12] Klim Efremenko. 3-query locally decodable codes of subexponential length. *SIAM Journal on Computing*, 41(6):1694–1703, 2012.
- [FV98] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1998.

- [GL16] Venkatesan Guruswami and Euiwoong Lee. Complexity of approximating csp with balance / hard constraints. *Theory of Computing Systems*, 59(1):76–98, Jul 2016.
- [Gop09] Parikshit Gopalan. A note on Efremenko’s locally decodable codes. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 16, 2009.
- [Gop14] Parikshit Gopalan. Constructing ramsey graphs from boolean function representations. *Combinatorica*, 34(2):173–206, 2014.
- [Gro00] Vince Grolmusz. Superpolynomial size set-systems with restricted intersections mod 6 and explicit ramsey graphs. *Combinatorica*, 20(1):71–86, 2000.
- [IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.
- [Jea98] Peter Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200(1–2):185 – 204, 1998.
- [KDW04] Iordanis Kerenidis and Ronald De Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *Journal of Computer and System Sciences*, 69(3):395–420, 2004.
- [Kro67] Melven R Krom. The decision problem for a class of first-order formulas in which all disjunctions are binary. *Mathematical Logic Quarterly*, 13(1-2):15–20, 1967.
- [KT00] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 80–86. ACM, 2000.
- [LV18] Tianren Liu and Vinod Vaikuntanathan. Breaking the circuit-size barrier in secret sharing. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 699–708, 2018.
- [MR10] Dana Moshkovitz and Ran Raz. Two-query pcp with subconstant error. *Journal of the ACM (JACM)*, 57(5):29, 2010.
- [NSZ18] Martin Nägele, Benny Sudakov, and Rico Zenklusen. Submodular minimization under congruency constraints. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 849–866. Society for Industrial and Applied Mathematics, 2018.
- [Pos41] Emil L Post. *The two-valued iterative systems of mathematical logic*. Number 5 in Annals of Mathematics Studies. Princeton University Press, 1941.
- [Sch78] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing, STOC ’78*, pages 216–226, New York, NY, USA, 1978. ACM.

- [TB98] Gábor Tardos and DA Mix Barrington. A lower bound on the mod 6 degree of the or function. *Computational Complexity*, 7(2):99–108, 1998.
- [Uma17] Chris Umans, editor. *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*. IEEE Computer Society, 2017.
- [Yek08] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. *Journal of the ACM (JACM)*, 55(1):1, 2008.
- [Zhu17] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In Umans [Uma17], pages 331–342.