Electronic Colloquium on Computational Complexity, Report No. 53 (2002)

ECCC

Is Constraint Satisfaction Over Two Variables Always Easy?*

Lars Engebretsen[1,**] and Venkatesan Guruswami[2,***]

[1] Department of Numerical Analysis and Computer Science
Royal Institute of Technology
SE-100 44 Stockholm
SWEDEN
[2] University of California at Berkeley
Miller Institute for Basic Research in Science
Berkeley, CA 94720
USA

July, 2002

Abstract. By the breakthrough work of Håstad, several constraint satisfaction problems are now known to have the following *approximation resistance* property: satisfying more clauses than what picking a random assignment would achieve is **NP**-hard. This is the case for example for Max E3-Sat, Max E3-Lin and Max E4-Set Splitting. A notable exception to this extreme hardness is constraint satisfaction over two variables (2-CSP); as a corollary of the celebrated Goemans-Williamson algorithm, we know that every Boolean 2-CSP has a non-trivial approximation algorithm whose performance ratio is better than that obtained by picking a random assignment to the variables. An intriguing question then is whether this is also the case for 2-CSPs over larger, non-Boolean domains. This question is still open, and is equivalent to whether the generalization of Max 2-SAT to domains of size d, can be approximated to a factor better than $(1 - 1/d^2)$.

In an attempt to make progress towards this question, in this paper we prove, firstly, that a slight restriction of this problem, namely a generalization of linear inequations with two variables per constraint, *is not* approximation resistant, and, secondly, that the Not-All-Equal Sat problem over domain size d with three variables per constraint, *is* approximation resistant, for every $d \geq 3$. In the Boolean case, Not-All-Equal Sat with three variables per constraint is equivalent to Max 2-SAT and thus has a non-trivial approximation algorithm; for larger domain sizes, Max 2-SAT can be reduced to Not-All-Equal Sat with three variables per constraint. Our approximation algorithm implies that a wide class of 2-CSPs called *regular 2-CSPs* can all be approximated beyond their random assignment threshold.

*A preliminary version of this work appears as an extended abstract in the *Proceedings of the 6th International Workshop on Randomization and Approximation Techniques in Computer Science* (RANDOM'02), 13–15 September 2002, Cambridge, Massachusetts.

**Research partly performed while the author was visiting MIT with support from the Marcus Wallenberg Foundation and the Royal Swedish Academy of Sciences.

***Supported by a Miller Research Fellowship.

ISSN 1433-8092

1 Introduction

In a breakthrough paper, Håstad [10] studied the problem of giving approximate solutions to maximization versions of several constraint satisfaction problems. An instance of a such a problem is given as a set of variables and a collection of constraints, i.e., functions from some domain to $\{0,1\}$, on certain subsets of variables, and the objective is to find an assignment to the variables that satisfies as many constraints as possible. An approximate solution of a constraint satisfaction program is simply an assignment that satisfies roughly as many constraints as possible. In this setting we are interested in proving either that there exists a polynomial time algorithm producing approximate solutions, i.e., solutions that are at most some constant factor worse compared to the optimum, or that no such algorithms exist.

In this paper we will study the common special case where each individual constraint depends on a fixed number k of the variables—this case is usually called the Max k-CSP problem and the size of the instance is given as the total number of variables that appear in the constraints. The complexity of the constraint satisfaction problem (CSP) is determined by the precise set of constraints that may be posed on subsets of k variables, and accordingly we get various families of Max k-CSP problems. For each such CSP, there exists a very naive algorithm that approximates the optimum within a constant factor: The algorithm that just guesses a solution at random. In his paper, Håstad [10] proved the very surprising fact that this algorithm is essentially the best possible efficient algorithm for several constraint satisfaction problems, unless $\mathbf{P} = \mathbf{NP}$. The proofs unify constructions from interactive proof systems with harmonic analysis over finite groups and give a general framework for proving strong impossibility results regarding the approximation of constraint satisfaction programs.

Håstad [10] suggests that predicates with the property that the naive randomized algorithm is the best possible polynomial time approximation algorithm should be called *non-approximable beyond the random assignment threshold*; we also use the phrase *approximation resistant* to refer to the same phenomenon.

Definition 1. *A solution to a maximization problem is* α-approximate *if it is feasible and has weight at least* α *times the optimum. An approximation algorithm has* performance ratio α *if it delivers* α*-approximate solutions in polynomial time.*

Definition 2. *A CSP is said to be* approximation resistant *or* non-approximable beyond the random assignment threshold *if, for any constant* $\varepsilon > 0$*, it is* ***NP****-hard to compute a* $(\rho+\varepsilon)$*-approximate solution, where* ρ *is the expected fraction of constraints satisfied by a solution guessed uniformly at random.*

Clearly, understanding which predicates are approximation resistant is an important pursuit. The current knowledge is that for Boolean CSPs, which understandably have received the most attention so far, there is a precise understanding of which CSPs on *exactly three* variables are approximation resistant: All predicates that are implied by parity have this property [10, 17]. It is known that *no* Boolean CSP over two variables is approximation resistant; this is a corollary of the breakthrough Goemans-Williamson algorithm [8]. For the case of four or more variables, very little is known; therefore it seems to be a good approach to first understand the situation for two and three variables.

Accordingly, we are interested in the situation for CSPs with two and three variables over larger, non-Boolean, domains. In particular, it is a really intriguing question whether every CSP over two variables can be approximated better than random, no matter what the domain size is. The central aim of this paper is to study this question. We are not able to resolve it completely, but we conjecture that the answer to the question is yes.

1.1 Formal definitions of some CSPs

Before discussing our results, we will need to define some of the CSPs that we will be concerned with in this paper. A specific k-CSP problem is defined by the family of constraints that may be imposed on subsets of k variables. Allowing arbitrary constraints gives the most general problem, which we call Max Ek-CSP(d). In this paper, d refers to the domain size from which the variables may take values, with $d = 2$ corresponding to the Boolean case. Over domain size d, a constraint is simply a function $f\colon [d]^k \to \{0,1\}$, where $[d] = \{0, 1, \ldots, d-1\}$. Equivalently, a constraint f can be viewed as a subset of $[d]^k$ consisting of all inputs which it maps to 1.

The Max Ek-Sat(d) problem is defined by the constraint family $\{f \subseteq [d]^k : |f| = d^k - 1\}$, i.e., the family of all constraints having just one non-satisfying assignment. Max Ek-NAE-Sat(d) is the problem where the constraints assert that the specific variables are not all equal, except that we also allow translates of variables, e.g., for the two variable case, a constraint can assert $x_1+1 \neq x_2+3$ (the addition being done modulo d); this is the analog of complementation of Boolean variables. In the Max Ek-Lin(d) problem, the constraint family is given by all linear constraints: $\{\mathrm{Lin}(\alpha_1, \ldots, \alpha_k, c) : \alpha_i, c \in [d]\}$ where $\mathrm{Lin}(\alpha_1, \ldots, \alpha_k, c) = \{(x_1, \ldots, x_k) : \sum_i \alpha_i x_i = c \bmod d\}$. The Max E$k$-LinInEq($d$) problem is defined by the family of all linear inequations: $\{f \subseteq [d]^k : [d]^k \setminus f \text{ is a linear constraint}\}$.

For the two variable case, we define the constraint satisfaction problems Max BIJ(d) and Max Co-BIJ(d) which are generalizations of Max E2-Lin(d) and Max E2-LinInEq(d) respectively. Let S_d be the set of all bijections from $[d]$ to $[d]$. For each $\pi \in S_d$, define the 2-ary constraint $f_{\pi,d} = \{(a, b) \in [d]^2 : b = \pi(a)\}$. Now define the family $\mathrm{BIJ}(d) = \{f_{\pi,d} : \pi \in S_d\}$; we call the CSP associated with this family Max BIJ(d). The problem Max Co-BIJ(d) is obtained by constraints which are complements of those in BIJ(d), i.e., a constraint is of the form $\pi(x_1) \neq x_2$ for some bijection π defined over $[d]$. It is clear that these problems generalize Max E2-Lin(d) and Max E2-LinInEq(d) respectively.

For the three variable case, we define the problem Max E3-NAE-Sat(G) for finite Abelian groups G. For each triple $(g_1, g_2, g_3) \in G^3$ define the constraint $N_{g_1,g_2,g_3} = \{(x_1, x_2, x_3) \in G^3 : \neg(g_1 x_1 = g_2 x_2 = g_3 x_3)\}$. Now define the family $\mathrm{NAE}(G) = \{N_{g_1,g_2,g_3} : (g_1, g_2, g_3) \in G^3\}$; we denote by Max E3-NAE-Sat(G) the CSP associated with this family of constraints. Note that the group structure is indeed present in the problem since the constraints involve multiplication by elements from G. In fact, we are able to prove in this paper that Max E3-NAE-Sat($\boldsymbol{Z}_4$) is approximation resistant while we are unable to determine the approximability of Max E3-NAE-Sat($\boldsymbol{Z}_2 \times \boldsymbol{Z}_2$).

It is an interesting open question to determine what kind of hardness holds for the restricted version of Max E3-NAE-Sat(G) where group multipliers are not allowed; for this problem the group structure is, of course, not present at all. Recently, Khot [13] has shown that Max E3-NAE-Sat($\boldsymbol{Z}_3$) is approximation resistant even without group multipliers.

1.2 Our results

Preliminaries: First, we make explicit the easily seen result that an approximation algorithm for Max E2-Sat(d) with performance ratio better than $1 - 1/d^2$, i.e., better than the random assignment threshold, implies that any CSP over 2 variables can be approximated to within better than *its respective* random assignment threshold. In other words, Max E2-Sat(d) is the hardest problem in this class, and if there is some Max E2-CSP(d) which is approximation resistant, then Max E2-Sat(d) has to be approximation resistant.

Consequently, our interest is in the approximability of Max E2-Sat(d), specifically to either find a polynomial time approximation algorithm with performance ratio greater than $1 - 1/d^2$ or to prove a tight hardness result that the trivial $1 - 1/d^2$ is the best one can hope for. While we are unable to

resolve this question, we consider and prove results for two predicates whose difficulty sandwiches that of solving Max E2-Sat(d): namely Max Co-BIJ(d) and Max E3-NAE-Sat($\boldsymbol{Z}_d$). The former problem is (in a loose sense) the natural 2-CSP which is next in "easiness" after Max E2-Sat(d) as far as approximating better than the random assignment threshold is concerned. There is an approximation preserving reduction from Max E2-Sat(d) to Max E3-NAE-Sat($\boldsymbol{Z}_d$), implying that Max E3-NAE-Sat($\boldsymbol{Z}_d$) is a harder problem than Max E2-Sat(d).

Algorithms: For the Max Co-BIJ(d) problem, we prove that it is *not* approximation resistant by presenting a polynomial time approximation algorithm with performance ratio $1 - d^{-1} + 0.07d^{-4}$. This result implies that a large class of 2-CSPs, called *regular 2-CSPs* (defined below), are *not* approximation resistant. Viewing a 2-ary constraint C over domain size d as a subset of $[d] \times [d]$, the constraint is said to be *r-regular* if for each $a \in [d]$, $|\{x : (x,a) \in C\}| = |\{y : (a,y) \in C\}| = r$ (the term regular comes from the fact that the bipartite graph defined by C is regular). The constraint is *regular* if it is r-regular for some $1 \leq r < d$. A 2-CSP is regular if all the constraints in the CSP are regular and it is r-regular if all the constraints are r-regular.

Our result for regular 2-CSPs includes as a special case the result of Frieze and Jerrum [7] that Max d-Cut can be approximated to better than its random threshold. Our performance ratio is weaker, but our analysis is simpler and gives a more general result. Another special case is the result for Max E2-Lin(d) where our result actually improves the approximation ratio of Andersson *et al* [3]. Recently, Khot [12] gave a simpler algorithm that beats the random assignment threshold for Max E2-Lin(d) as well as the more general Max BIJ(d) problems—his result is actually more general and can find a near-satisfying assignment given a near-satisfiable instance, i.e., an instance where the optimum solution satisfies a fraction $1 - \varepsilon$ of constraints. Our approximation algorithm for Max Co-BIJ(d) is based on a semidefinite programming relaxation, similar to that used for Max BIJ(d) by Khot [12], combined with a rounding scheme used by Andersson [1, 2] to construct an approximation algorithm for Max d-Section, the generalization of Max Bisection to domains of size d. Technically, we view this algorithmic result as the main contribution of this paper.

Inapproximability results: For the Boolean case, $d = 2$, it is known that Max E3-NAE-Sat can be approximated to better than random. The GW-algorithm [8] for Max E2-Sat essentially gives such an algorithm, and the performance ratio was later improved by Zwick [18]. We prove that for larger domains, the problem becomes approximation resistant; in other words, it is **NP**-hard to approximate Max E3-NAE-Sat($\boldsymbol{Z}_d$) to better than $(1 - 1/d^2 + \varepsilon)$ for any $d \geq 3$ and any $\varepsilon > 0$. This result rules out the possibility of a non-trivial algorithm for Max E2-Sat(d) that works by reducing it to Max E3-NAE-Sat($\boldsymbol{Z}_d$). In fact, we prove that for any finite group G which is not of the form $\boldsymbol{Z}_2 \times \boldsymbol{Z}_2 \times \cdots \times \boldsymbol{Z}_2$, Max E3-NAE-Sat($G$) is hard to approximate within a factor $(1 - 1/|G|^2 + \varepsilon)$.

We also obtain tight hardness results for CSPs such as Max E3-Sat(d), and Max E3-LinInEq(d). They are known to be approximation resistant over the Boolean domain [10], and hence one expects them to be hard for larger domains too. We verify that this is indeed the case. For example, we prove that for $k, d \geq 3$, Max Ek-LinInEq(d) is approximation resistant and thus cannot be approximated within a factor better than $(1 - 1/d + \varepsilon)$, for any constant $\varepsilon > 0$, in polynomial time unless $\mathbf{P} = \mathbf{NP}$. A simple gadget then gives the tight hardness result of $(1 - 1/d^k + \varepsilon)$ for Max Ek-Sat(d).

We remark that the above hardness results hold with *perfect completeness*; in other words, the stated approximation factors are hard to obtain even on satisfiable instances of the concerned constraint satisfaction problems.

$$
\begin{array}{lll}
\text{maximize} & \displaystyle\sum_{x,x',\pi} w_{x,x',\pi}\left(1 - \sum_{j=0}^{d-1}\langle u_j^x, u_{\pi(j)}^{x'}\rangle\right) & \\
\text{such that} & \langle u_j^x, u_{j'}^{x'}\rangle \geq 0 & \forall x \in X, x' \in X, j \in \boldsymbol{Z}_d, j' \in \boldsymbol{Z}_d \\
& \langle u_j^x, u_{j'}^x\rangle = 0 & \forall x \in X, (j,j') \in \boldsymbol{Z}_d^2 : j \neq j' \\
& \displaystyle\sum_{j=0}^{d-1}\langle u_j^x, u_j^x\rangle = 1 & \forall x \in X \\
& \displaystyle\sum_{j=0}^{d-1}\sum_{j'=0}^{d-1}\langle u_j^x, u_{j'}^{x'}\rangle = 1 & \forall (x,x') \in X^2 : x \neq x'
\end{array}
$$

Figure 1. Semidefinite relaxation of Max Co-BIJ(d) with variable set X. A clause in the Max Co-BIJ(d) instance is denoted by (x, x', π) where $x \in X$ and $x' \in X$ are variables and $\pi: \boldsymbol{Z}_d \to \boldsymbol{Z}_d$ is a permutation. The clause is satisfied unless $x = j$ and $x' = \pi(j)$ for some $j \in \boldsymbol{Z}_d$. Each clause (x, x', π) has a non-negative weight $w_{x,x',\pi}$ associated with it.

Conclusions: We are not able to completely resolve the status of the 2-CSP problem over larger domains. Using reductions, we prove a hardness result of $1 - \Omega(1/d^2)$ for Max E2-Sat(d), which compares reasonably well with the $(1-1/d^2)$ random assignment threshold. For satisfiable instances of Max E2-Sat(d), we prove a hardness result of $1-\Omega(1/d^3)$. Nevertheless, we conjecture that there is an approximation algorithm beating the random assignment threshold for Max E2-Sat(d), and hence for all instances of 2-CSP.

Organization: We begin with a brief Section 2 highlighting why Max E2-Sat(d) is the hardest Max E2-CSP(d) problem in terms of beating the random assignment threshold. Next, in Section 3 we prove that every Max Co-BIJ(d) problem admits an algorithm that beats the random assignment threshold, and record some of its consequences. In Section 7, we prove that Max E3-NAE-Sat(G) is approximation resistant for most groups, including $G = \boldsymbol{Z}_d$ (the case of most interest in the context of Max E2-Sat(d)). Finally, we record results that directly apply to Max E2-Sat(d) in Section 8.

2 The "Universality" of Max E2-Sat(d)

We note that the existence of an approximation algorithm that beats the random threshold for every 2-CSP is equivalent to the existence of such an algorithm for Max E2-Sat(d). Thus, an algorithm for Max E2-Sat(d) with performance ratio better than $(1 - 1/d^2)$ will imply that no 2-CSP is approximation resistant, thus resolving our conjecture that every 2-CSP is "easy".

This claim is seen by a "gadget" reducing an arbitrary CSP(d) to Max E2-Sat(d). Given an instance of any 2-CSP, construct an instance of Max E2-Sat(d) by repeating the following for every constraint C in the original 2-CSP: For every non-satisfying assignment to C, add one 2SAT(d) constraint which has precisely this non-satisfying assignment. If an assignment satisfies C then it also satisfies all the 2SAT(d) constraints in the gadget, and otherwise it satisfies precisely all but one of the 2SAT(d) constraints. Using this fact, it is straightforward to show that if Max E2-Sat(d) can be approximated beyond the random threshold, the above procedure gives an approximation algorithm that approximates an arbitrary 2-CSP beyond the random threshold. Conversely, if any 2-CSP at all is approximation resistant, then Max E2-Sat(d) must be approximation resistant.

1. Solve the semidefinite program in Fig. 1.
2. Denote by u_j^x the vectors obtained from the solution.
3. For every $(x, j) \in X \times \boldsymbol{Z}_d$, let $v_j^x = u_j^x - \frac{1}{d}\sum_{j=0}^{d-1} u_j^x$.
4. Select r from a dn-dimensional $\mathrm{N}(0, I)$ Gaussian distribution.
5. Set $q_{xj} = \frac{1}{d} + K\langle r, v_j^x\rangle$ for all $(x, j) \in X \times \boldsymbol{Z}_d$.
6. For each $x \in X$,
   - set $p_{xj} = q_{xj}$ if $q_{xj} \in [0, 2/d]$ for all $j \in \boldsymbol{Z}_d$;
   - set $p_{xj} = 1/d$ for all $j \in \boldsymbol{Z}_d$ otherwise.
7. For each $x \in X$, let $x = j$ with probability p_{xj}.

Figure 2. Approximation algorithm for Max Co-BIJ(d) with variable set X. The algorithm is parameterized by the positive constant K.

3 Approximation algorithm for Max Co-BIJ(d)

To construct an approximation algorithm for Max Co-BIJ(d) we combine a modification of the semidefinite relaxation used by Khot [12] for the Max BIJ(d) problem with a modification of the randomized rounding used by Andersson [1, 2] for the Max d-Section problem. Recall that a specific clause in the Max Co-BIJ(d) problem is of the form (x, x', π), where x and x' are variables in the Max Co-BIJ(d) instance and π is a permutation, and that the clause is satisfied unless $x = j$ and $x' = \pi(j)$ for some j. In our semidefinite relaxation of Max Co-BIJ(d) there are d vectors $\{u_0^x, \ldots, u_{d-1}^x\}$ for every variable x in the Max Co-BIJ(d) instance. Intuitively, the vector u_j^x sets the value of the variable x to j. To prove that our algorithm beats the random assignment threshold, we first establish that the semidefinite program in Fig. 1 is a relaxation of Max Co-BIJ(d), then prove that the rounding scheme proposed in Fig. 2 is well-defined, and finally analyze the expected performance of the rounding scheme using local analysis.

Lemma 1. *The semidefinite program in Fig. 1 is a relaxation of Max Co-BIJ(d).*

Proof. Suppose that we have an instance of Max Co-BIJ(d) with variable set X. Given an assignment ρ to the variables in X consider the following configuration of vectors: Let a be a vector of unit norm and define

$$u_j^x = \begin{cases} a & \text{if } j = \rho(x), \\ 0 & \text{otherwise.} \end{cases}$$

This configuration of vectors is feasible for the semidefinite program in Fig. 1 and the corresponding value of the objective function is exactly the weight of the satisfied equations in the Max Co-BIJ(d) instance. ∎

Lemma 2. *Suppose that $\{u_j^x : (x, j) \in X \times \boldsymbol{Z}_d\}$ is a feasible solution to the semidefinite program in Fig. 1. Then the barycenter*

$$b_x = \frac{1}{d}\sum_{j=0}^{d-1} u_j^x$$

of the vectors $\{u_j^x : j \in \boldsymbol{Z}_d\}_{j=0}^{d-1}$ is independent of x.

Proof. The constraints in the semidefinite program imply that $|b_x|^2 = |b_{x'}|^2 = \langle b_x, b_{x'}\rangle = d^{-2}$. ∎

Lemma 3. *For any clause (x, x', π) in the Max Co-BIJ(d) instance, the algorithm in Fig. 2 satisfies (x, x', π) with probability at least*

$$\left(1 - \sum_{j=0}^{d-1} \langle u_j^x, u_{\pi(j)}^{x'} \rangle\right) \int_B \left(1 - \frac{1}{d} + \frac{K^2 r_1^2}{d}\right) dP(r)$$

where K is any positive constant, the vectors u_j^x and $u_{j'}^{x'}$ are as described in the algorithm, $B = \{r \in \boldsymbol{R}^{2d} : |r| \leq 1/Kd\}$, and P is the probability distribution of a $2d$-dimensional Gaussian with mean zero and identity covariance matrix.

Proof. Consider an arbitrary clause (i, i', π) and the corresponding values q_{xj} computed by the algorithm. Let $B = \{r \in \boldsymbol{R}^{2d} : |r| \leq 1/Kd\}$. When $r \in B$, both q_{xj} and $q_{x'j}$ are in the interval $[0, 2/d]$; hence the clause (i, i', π) is satisfied with probability

$$\begin{aligned} 1 - \sum_{j=0}^{d-1} p_{xj} p_{x',\pi(j)} &= 1 - \sum_{j=0}^{d-1} \left(\frac{1}{d} + K\langle r, v_j^x \rangle\right)\left(\frac{1}{d} + K\langle r, v_{\pi(j)}^{x'} \rangle\right) \\ &= 1 - \frac{1}{d} - K^2 \sum_{j=0}^{d-1} \langle r, v_j^x \rangle \langle r, v_{\pi(j)}^{x'} \rangle \end{aligned}$$

given r in this case. By the definition of v_j^x from the algorithm, $\langle r, v_j^x \rangle \langle r, v_{\pi(j)}^{x'} \rangle = \langle r, u_j^x - b \rangle \langle r, u_{\pi(j)}^{x'} - b \rangle = \langle r, u_j^x \rangle \langle r, u_{\pi(j)}^{x'} \rangle - \langle r, u_j^x \rangle \langle r, b \rangle - \langle r, u_{\pi(j)}^{x'} \rangle \langle r, b \rangle + \langle r, b \rangle \langle r, b \rangle$ where b is the barycenter of the vectors $\{u_j^x : j \in \boldsymbol{Z}_d\}$, which is independent of x according to Lemma 2. Therefore

$$-K^2 \sum_{j=0}^{d-1} \langle r, v_j^x \rangle \langle r, v_{\pi(j)}^{x'} \rangle = K^2 \left(d\langle r, b \rangle \langle r, b \rangle - \sum_{j=0}^{d-1} \langle r, u_j^x \rangle \langle r, u_{\pi(j)}^{x'} \rangle\right).$$

To conclude, the probability that the clause is satisfied given r is

$$1 - \frac{1}{d} + K^2 \left(d\langle r, b \rangle \langle r, b \rangle - \sum_{j=0}^{d-1} \langle r, u_j^x \rangle \langle r, u_{\pi(j)}^{x'} \rangle\right) \tag{1}$$

when $r \in B$. Integrating over B, we can thus bound the probability that the clause is satisfied from below by

$$\int_B \left(1 - \frac{1}{d} + K^2 \left(d\langle r, b \rangle \langle r, b \rangle - \sum_{j=0}^{d-1} \langle r, u_j^x \rangle \langle r, u_{\pi(j)}^{x'} \rangle\right)\right) dP(r).$$

To compute the integral of $\langle r, b \rangle \langle r, b \rangle$, introduce an orthonormal basis $\{e_k\}$ such that $b = e_1/d$ and write $r = \sum_k r_k e_k$ in this basis. Then

$$\int_B \langle r, b \rangle \langle r, b \rangle \, dP(r) = \frac{1}{d^2} \int_B r_1^2 \, dP(r)$$

where the last integral is actually independent of the basis since both P and B are spherically symmetric. To compute the integral of $\langle r, u_j^x \rangle \langle r, u_{\pi(j)}^{x'} \rangle$, we proceed similarly: Introduce an orthonormal basis $\{e_k\}$ such that $u_j^x = x_1 e_1$ and $u_{\pi(j)}^{x'} = y_1 e_1 + y_2 e_2$. Then

$$\int_B \langle r, u_j^x \rangle \langle r, u_{\pi(j)}^{x'} \rangle \, dP(r) = \int_B (r_1^2 x_1 y_1 + r_1 r_2 x_1 y_2) \, dP(r).$$

The integral of the second term vanishes since the integrand is odd and the interval is symmetric around the origin. Therefore

$$\int_B \langle r, u_j^x \rangle \langle r, u_{\pi(j)}^{x'} \rangle \, dP(r) = x_1 y_1 \int_B r_1^2 \, dP(r) = \langle u_j^x, u_{\pi(j)}^{x'} \rangle \int_B r_1^2 \, dP(r).$$

To conclude, we can write the probability that the clause is satisfied as $a - cx$ where

$$\begin{aligned} a &= \int_B \left(1 - \frac{1}{d} + \frac{K^2 r_1^2}{d}\right) dP(r), \\ c &= K^2 \int_B r_1^2 \, dP(r), \\ x &= \sum_{j=0}^{d-1} \langle u_j^x, u_{\pi(j)}^{x'} \rangle. \end{aligned}$$

We now find an α such that $a - cx \geq \alpha(1 - x)$. Since the expression (1) is a probability it is non-negative for every r. Hence $a \geq c \geq 0$ and, since the constraints in the semidefinite program imply that $x \geq 0$,

$$a - cx \geq a - ax = \left(1 - \sum_{j=0}^{d-1} \langle u_j^x, u_{\pi(j)}^{x'} \rangle\right) \int_B \left(1 - \frac{1}{d} + \frac{K^2 r_1^2}{d}\right) dP(r). \qquad \blacksquare$$

Theorem 1. *The algorithm in Fig. 2 with $K = 1/\sqrt{13d^3}$ is a randomized polynomial time approximation algorithm for Max Co-BIJ(d) with expected performance ratio $1 - d^{-1} + 0.07d^{-4}$ and thus better than the random assignment threshold.*

Proof. Consider the algorithm in Fig. 2 applied to an instance of Max Co-BIJ(d). By Lemma 1, the semidefinite program in Fig. 1 is a relaxation of Max Co-BIJ(d); therefore the optimum value of the Max Co-BIJ(d) instance can be bounded by

$$\sum_{x,x',\pi} w_{x,x',\pi} \left(1 - \sum_{j=0}^{d-1} \langle u_j^x, u_{\pi(j)}^{x'} \rangle\right)$$

where the vectors u_j^x are the solution to the program. Now consider the rounding. By Lemma 3, the clause (x, x', π) is satisfied with probability

$$\left(1 - \sum_{j=0}^{d-1} \langle u_j^x, u_{\pi(j)}^{x'} \rangle\right) \int_B \left(1 - \frac{1}{d} + \frac{K^2 r_1^2}{d}\right) dP(r)$$

By Lemmas 21 and 22,

$$\int_B \left(1 - \frac{1}{d} + \frac{K^2 r_1^2}{d}\right) dP(r) \geq 1 - \frac{1}{d} + \frac{0.07}{d^4}$$

with the parameter choice $K^2 = d^{-3}/13$. So with this parameter choice, it now follows by linearity of expectation and the above bound on the optimum value of the Max Co-BIJ(d) instance that the expected weight of the solution produced by the algorithm is at least a factor $(1 - d^{-1} + 0.07d^{-4})$ times the optimum value of the Max Co-BIJ(d) instance. $\blacksquare$

1. Select a random assignment ρ to the variables in the instance.
2. Replace each equation $ax + by = c$ with the $d - 1$ inequations $ax + by \neq c_i$ for all $c_i \neq c$. Obtain an assignment τ by running the algorithm in Fig. 2 on this instance.
3. Pick the best of the assignments ρ and τ.

Figure 3. Approximation algorithm for Max E2-Lin(d).

3.1 An approximation algorithm for Max E2-Lin(d)

We can use the above algorithm for Max Co-BIJ(d) to construct an algorithm also for Max E2-Lin(d): Simply replace an equation $ax + by = c$ with the $d - 1$ inequations $ax + by \neq c_i$ for all $c_i \neq c$. Then an assignment that satisfies a linear equation satisfies all of the corresponding linear inequations and an assignment that does not satisfy a linear equation satisfies $d - 2$ of the $d - 1$ corresponding linear equations. This algorithm gives a performance ratio of $1/d + \Omega(1/d^4)$ which improves significantly on the previously best known ratio of $1/d + \Omega(1/d^{14})$ [3].

Theorem 2. *For all $d \geq 4$, the algorithm in Fig. 3 is a randomized polynomial time approximation algorithm for Max E2-Lin(d) with expected performance ratio $d^{-1} + 0.05d^{-4}$.*

Proof. If the optimum of the instance is smaller than a fraction $1 - 0.05d^{-3}$ of all equations, the random assignment—which satisfies an expected fraction $1/d$ of all equations—satisfies an expected fraction

$$\frac{1/d}{1 - 0.05d^{-3}} \geq \frac{1}{d} + \frac{0.05}{d^4}$$

of the optimum.

If the optimum is larger than a fraction $1-0.05d^{-3}$ of all equations, the optimum of the instance of inequations is at least $1-0.05d^{-3}(d-1)^{-1}$. By Theorem 1, our approximation for Max Co-BIJ(d) from Fig. 2 therefore finds a solution with expected weight at least a fraction

$$\left(1 - \frac{0.05}{d^3(d-1)}\right)\left(\frac{d-1}{d} + \frac{0.07}{d^4}\right) = \frac{d-1}{d} + \frac{0.02}{d^4} - \frac{0.0035}{d^7(d-1)}$$

of all inequations. An assignment satisfying a fraction $1/d + \alpha$ of the equations satisfies a fraction

$$\left(\frac{1}{d} + \alpha\right) + \frac{d-2}{d-1}\left(1 - \frac{1}{d} - \alpha\right) = \frac{d-1}{d} + \frac{\alpha}{d-1}$$

of the inequations and vice versa; therefore the assignment constructed above satisfies at least an expected fraction

$$\frac{1}{d} + \frac{0.02(d-1)}{d^4} - \frac{0.0035}{d^7} > \frac{1}{d} + \frac{0.02(d - 1 - 0.0035d^{-3})}{d^4}$$

of the equations. The function $d-1-0.0035d^{-3}$ is increasing in d and for $d = 4$ it is $3-0.0035/64 > 2.5$. Therefore, at least an expected fraction $d^{-1} + 0.05d^{-4}$ of all equations are satisfied when $d \geq 4$. ∎

Corollary 1. *For all $d \geq 2$, there is a polynomial time approximation algorithm for Max E2-Lin(d) with expected performance ratio $d^{-1}+0.05d^{-4}$ and thus better than the random assignment threshold.*

Repeat the following for every relation in the Max Co-BIJ(d) instance:

1. Let R be the bipartite graph defined by the relation.
2. Let r be the degree of R.
3. Let R^c be the complement of R.
4. Decompose the edges of R^c into perfect matchings $R_1, \ldots, R_{d-r}$.
5. Define the relations π_i such that $d_1 \sim_{\pi_i} d_2$ if $(d_1, d_2) \in R_i$.
6. Add the relations $\pi_1, \pi_2, \ldots, \pi_{d-r}$ to the Max Co-BIJ(d) instance.

Figure 4. Construction of an instance of Max Co-BIJ(d) from an instance of any regular CSP.

1. Select a random assignment ρ to the variables in the instance.
2. Create an instance of Max Co-BIJ(d) as described in Fig. 4. Obtain an assignment τ by running the algorithm in Fig. 2 on this instance.
3. Pick the best of the assignments ρ and τ.

Figure 5. Approximation algorithm for any regular CSP.

Proof. Algorithms for $d = 2$ and $d = 3$ have been provided by Goemans and Williamson [8, 9], for $d \geq 4$ the result follows by Theorem 2 ∎

3.2 An approximation algorithm for regular 2-CSPs

We can obtain an approximation algorithm for all regular CSPs by a straightforward generalization of the ideas from the previous section. Given an r-regular 2-CSP, we proceed as follows for every relation R defining the CSP: Decompose R^c, the "bipartite complement" of the graph defined by R, into $(d-r)$ perfect matchings $\pi_R^1, \pi_R^2, \ldots, \pi_R^{d-r}$. Then let these matchings define the Max Co-BIJ(d) instance. An assignment that satisfies R satisfies all of the $d - r$ matchings while an assignment that does not satisfy R satisfies $d - r - 1$ of them. Run the following two algorithms and take the assignment producing the largest weight as the result: The first algorithms selects a random assignment to the variables; the second algorithm runs the above algorithm for Max Co-BIJ(d).

Theorem 3. *For all $d \geq 2$ and all $1 \leq r \leq d-1$, the algorithm in Fig. 5 is a randomized polynomial time approximation algorithm for r-regular CSPs with expected performance ratio $r/d + \Omega(d^{-4})$.*

Proof. Consider the result of the algorithm on an arbitrary instance. If the optimum of the instance is smaller than a fraction $1-0.05(d-r)d^{-3}(d-1)^{-1}$ of all equations, the random assignment—which satisfies an expected fraction r/d of all equations—satisfies an expected fraction

$$\frac{r/d}{1 - 0.05(d-r)d^{-3}(d-1)^{-1}} \geq \frac{r}{d} + \frac{0.05r(d-r)}{d^4(d-1)}$$

of the optimum.

If the optimum is larger than a fraction $1 - 0.05(d-r)d^{-3}(d-1)^{-1}$ of all constraints, the optimum of the corresponding instance of Max Co-BIJ(d) is at least a fraction

$$1 - \frac{0.05(d-r)}{d^3(d-1)}\left(1 - \frac{d-r-1}{d-r}\right) = 1 - \frac{0.05}{d^3(d-1)}$$

By Theorem 1, our approximation for Max Co-BIJ(d) from Fig. 2 therefore finds a solution with expected weight at least a fraction

$$\left(1 - \frac{0.05}{d^3(d-1)}\right)\left(\frac{d-1}{d} + \frac{0.07}{d^4}\right) = \frac{d-1}{d} + \frac{0.02}{d^4} - \frac{0.0035}{d^7(d-1)}$$

of all constraints in the Max Co-BIJ(d) instance. An assignment satisfying a fraction $r/d + \alpha$ of the constraints in the r-regular CSP satisfies a fraction

$$\frac{d-r}{d-r}\left(\frac{r}{d} + \alpha\right) + \frac{d-r-1}{d-r}\left(1 - \frac{r}{d} - \alpha\right) = \frac{d-1}{d} + \frac{\alpha}{d-r}$$

of the constraints in the Max Co-BIJ(d) instance and vice versa; therefore the assignment constructed above satisfies at least a fraction

$$\frac{r}{d} + \frac{0.02(d-r)}{d^4} - \frac{0.0035(d-r)}{d^7(d-1)}$$

of the constraints in the r-regular CSP. ∎

It is not necessary that the various constraints be r-regular for the same r, and a similar argument also shows that every regular 2-CSP can be approximated in polynomial time beyond its random assignment threshold.

4 PCPs and inapproximability results: Background

In his seminal paper [10], Håstad introduced a methodology for proving inapproximability results for constraint satisfaction problems. On a high level, the method can be viewed as a simulation of the well-known two-prover one-round (2P1R) protocol for E3-Sat where the verifier sends a variable to one prover and a clause containing that variable to the other prover, accepting if the returned assignments are consistent and satisfy the clause.

4.1 The 2-prover 1-round protocol

We start with an instance of the **NP**-hard [4, 6] problem μ-gap E3-Sat(5).

Definition 3. μ-gap E3-Sat(5) *is the following decision problem: We are given a Boolean formula ϕ in conjunctive normal form, where each clause contains exactly three literals and each literal occurs exactly five times. We know that either ϕ is satisfiable or at most a fraction $\mu < 1$ of the clauses in ϕ are satisfiable and are supposed to decide if the formula is satisfiable.*

There is a well-known two-prover one-round (2P1R) interactive proof system that can be applied to μ-gap E3-Sat(5). It consists of two provers, P_1 and P_2, and one verifier. Given an instance, i.e., an E3-Sat formula ϕ, the verifier picks a clause C and variable x in C uniformly at random from the instance and sends C to P_1 and x to P_2. It then receives an assignment to the variables in C from P_1 and an assignment to x from P_2, and accepts if these assignments are consistent and satisfy C. If the provers are honest, the verifier always accepts with probability 1 when ϕ is satisfiable, i.e., the proof system has *completeness* 1, or *perfect completeness*. It can be shown that the provers can fool

the verifier with probability at most $(2+\mu)/3$ when ϕ is not satisfiable, i.e., that the above proof system has *soundness* $(2+\mu)/3$.

The soundness can be lowered to $((2+\mu)/3)^u$ by repeating the protocol u times independently, but it is also possible to construct a one-round proof system with lower the soundness by repeating u times in parallel as follows: The verifier picks u clauses $(C_1, \ldots, C_u)$ uniformly at random from the instance. For each C_i, it also picks a variable x_i from C_i uniformly at random. The verifier then sends $(C_1, \ldots, C_u)$ to P_1 and $(x_1, \ldots, x_u)$ to P_2. It receives an assignment to the variables in $(C_1, \ldots, C_u)$ from P_1 and an assignment to $(x_1, \ldots, x_u)$ from P_2, and accepts if these assignments are consistent and satisfy $C_1 \wedge \cdots \wedge C_u$. As above, the completeness of this proof system is 1, and it can be shown [15] that the soundness is at most c_μ^u, where $c_\mu < 1$ is some constant depending on μ but not on u or the size of the instance.

4.2 Constructing strategies for the provers

In the above setting, the proof is simply an assignment to all the variables. In that case, the verifier can just compare the assignments it receives from the provers and check if they are consistent and satisfying. The construction we use to prove that several non-Boolean constraint satisfaction programs are non-approximable beyond the random assignment threshold can be viewed as a simulation of the u-parallel repetition of the above 2P1R interactive proof system for μ-gap E3-Sat(5). We use a probabilistically checkable proof system (PCP) with a verifier closely related to the particular constraint we want to analyze. To find predicates that depend on variables from some domain of size d and are non-approximable beyond the random assignment threshold, we work with an Abelian group of order d. That enables us to use representation theory for Abelian groups to analyze our protocols.

The final verifier expects as proof encodings of the answers of P_1 and P_2 in the Raz 2P1R, and then checks very efficiently, by making very few queries, that the proof is close to valid encodings of answers that would have made the 2P1R verifier accept with good probability. To get hardness results for CSPs over domain size d, the specific encoding used is the so called *long G-code*, which will be defined in Sec. 6, for some Abelian group G of order d.

The proof expected by the PCP verifier consists of purported Long G-Codes of the assignments to the variables in U and W for each possible choice (U, W) of the 2P1R verifier. The PCP design task thus reduces to designing an "inner" verifier to check if two purported Long G-Codes encode assignments which are consistent answers for P_1 and P_2 in the 2P1R. One designs such a verifier with an acceptance predicate closely tied to the problem at hand, and its performance is analyzed using Fourier analysis. The basic strategy here is to show how proofs that make the "inner" verifier accept with high probability can be used to extract good strategies for P_1 and P_2 in the 2P1R protocol. On a high level, the Fourier expansion of the purported long codes are used to extract probabilistic strategies for the provers P_1 and P_2. We are able to express the acceptance probability of the verifier in the 2P1R protocol as a sum of certain pairwise products of Fourier coefficients and these products turn out to be large whenever the "inner" verifier accepts with large probability.

On a slightly more detailed level, let w be the probability that the considered constraint is satisfied by a random assignment. The aim of our analysis is to prove that it is **NP**-hard to satisfy more than a fraction w of the constraints. We do this by proving the contrapositive: If we can satisfy a fraction $w+\varepsilon$ of the constraints, for any constant $\varepsilon > 0$, we can decide any language in **NP** in polynomial time. This follows from the connection between our PCP and the 2P1R interactive proof system for μ-gap E3-Sat(5): We assume that it is possible to satisfy a fraction $w+\varepsilon$ for some constant $\varepsilon > 0$ and prove that this implies that there is a correlation between the

tables queried by the verifier in our PCP. We can then use this correlation to explicitly construct strategies for the provers in the 2P1R proof system for μ-gap E3-Sat(5) such that the verifier in that proof system accepts with probability that is independent of u. By selecting u large enough, we can then reach a contradiction. The final link in the chain is the observation that since our verifier uses only logarithmic randomness, we can form a CSP with polynomial size by enumerating the checked constraints for every possible outcome of the random bits. If the resulting constraint satisfaction program is approximable beyond the random assignment threshold, we can use it to decide the **NP**-hard language μ-gap E3-Sat(5) in polynomial time.

5 Representation theory and the Fourier transform

In this section we give a brief account of the representation theory for Abelian groups and the associated Fourier transform. For more details, we refer the reader to Terras's book [16]. In this paper, we always denote groups by the letters F, G and H, members in those groups by lowercase Roman letters and members in the corresponding dual groups $\hat{F}$, $\hat{G}$ and $\hat{H}$ by lowercase Greek letters.

5.1 Representation theory for Abelian groups

Let G be an Abelian group and $\hat{G}$ be its dual group, i.e., the group of all homomorphisms from G to $\boldsymbol{C}$. A central concept in the representation theory is the action of a member of $\hat{G}$ on a member of G. Since the dual groups consists of functions from G to $\boldsymbol{C}$, such an action is simply the evaluation of a member of $\hat{G}$ at a member of G.

Example 1. The group $\boldsymbol{Z}_d$ for some positive integer d is isomorphic to the group consisting of powers of $e^{2\pi i/d}$ with multiplication as the group operation and 1 as identity. The dual group then consists of all functions $x \mapsto x^n$ for the integers n from 0 to $d-1$. The product of the functions $x \mapsto x^n$ and $x \mapsto x^m$ is the function $x \mapsto x^{m+n}$, where addition is defined modulo d. This implies that the trivial homomorphism, that maps every group element to 1 and thus is the function $x \mapsto x^0$, is the identity in the dual group.

The above example motivates the following notation in the general case: For $g \in G$ and $\gamma \in \hat{G}$ we let g^γ denote $\gamma(g)$. We let multiplication be the group operation in G and addition be the group operation in $\hat{G}$ and denote the identities by $\mathbf{1}$ and $\mathbf{0}$, respectively. Then we can write the following identities: $g^{\gamma_1} g^{\gamma_2} = g^{\gamma_1+\gamma_2}$, $g_1^\gamma g_2^\gamma = (g_1 g_2)^\gamma$, $(g^2)^\gamma = g^{2\gamma} = g^\gamma g^\gamma$, $g^{-\gamma} = (g^{-1})^\gamma$, $g^{\mathbf{0}} = 1$, $\mathbf{1}^\gamma = 1$. We also need the following symmetry relations:

$$\sum_{g\in G} g^\gamma = \begin{cases} |G| & \text{if } \gamma = \mathbf{0}, \\ 0 & \text{otherwise,} \end{cases} \qquad \sum_{\gamma\in\hat{G}} g^\gamma = \begin{cases} |G| & \text{if } g = \mathbf{1}, \\ 0 & \text{otherwise.} \end{cases}$$

5.2 The Fourier transform

Now let f be a function from G to $\boldsymbol{C}$. Then we can write f as a Fourier series

$$f(g) = \sum_{\gamma\in\hat{G}} \hat{f}_\gamma g^\gamma \tag{2}$$

where

$$\hat{f}_\gamma = \frac{1}{|G|} \sum_{g\in G} f(g) g^{-\gamma}. \tag{3}$$

Moreover the following version of Plancherel's equality holds in this case:

$$\sum_{\gamma \in \hat{G}} |\hat{f}_\gamma|^2 = \frac{1}{|G|} \sum_{g \in G} |f(g)|^2. \tag{4}$$

6 The long G-code and its Fourier transform

Recall that our construction should simulate the 2P1R game for 3-Sat where the verifier sends u clauses to one prover and one variable from each clause to the other prover. Our proof, encoded by the long G-code, should contain these queries for all possible choices of the verifier in the 2P1R game. Since the verifier in the 2P1R games always rejects if it receives an answer which does not satisfy the u clauses we can in fact assume that the clause-prover always returns a satisfying assignment. Our encoding of the proof also reflects this in a way that will now be made explicit.

Definition 4. *Let U be a set of variables and denote by $\{-1,1\}^U$ the set of assignments to the variables in U. The* long G-code of some $x \in \{-1,1\}^U$ *is a function $A_{U,x}\colon \{-1,1\}^U \to G$ defined by $A_{U,x}(f) = f(x)$.*

Definition 5. *Let W be a set of clauses and denote by SAT^W the set of satisfying assignments to the clauses in W. The* long G-code of some $y \in \mathrm{SAT}^W$ *is a function $A_{W,y}\colon \mathrm{SAT}^W \to G$ defined by $A_{W,y}(h) = h(y)$.*

Definition 6. *A* standard written G-proof with parameter u *contains for each sequence U of u variables a string of length $|G|^{2^u}$, which we interpret as the table of a function $A_U\colon G^{\{-1,1\}^U} \to G$. It also contains for each set W of u clauses a string of length $|G|^{7^u}$ which we interpret as the table of a function $A_W\colon G^{\mathrm{SAT}^W} \to G$.*

Definition 7. *A standard written G-proof with parameter u is a* correct proof for a formula ϕ *if there is an assignment x, satisfying ϕ, such that A_V is the long G-code of $x|_V$ for any sequence V of u variables or any sequence V of u clauses.*

The verifier in our PCPs will typically select a random sequence W of clauses and then form U by selecting a variable at random from each clause. It will then query the tables A_U and A_W in the standard written G-proof at cleverly chosen positions. We analyze the verifier using the Fourier expansion of the alleged long codes; we therefore need to understand the Fourier transform of functions from G^I to $\boldsymbol{C}$ for some finite set I. Since the Fourier coefficients are used to devise a strategy for the provers in the 2P1R game, certain Fourier coefficients must be identically zero. The analysis also needs certain facts regarding the Fourier transform of two functions $A\colon G^I \to \boldsymbol{C}$ and $B\colon G^J \to \boldsymbol{C}$ where there is a projection from J to I. All these identities have already been obtained by Håstad [10, § 2.6], we only state them here for easy reference.

Let I be a finite set and consider the space G^I, i.e., the space of all functions from I to G. This space can be identified with the group $F = G^{|I|}$ since a function is simply a table of all its values for all of the $|I|$ possible inputs. Then the products of two functions is just the products of the corresponding group elements, and so on. For an element $f \in F$ and an element $x \in I$, we let $f(x)$ denote the coordinate in f corresponding to x. Now let A be a function from F to $\boldsymbol{C}$. We can write this function as a Fourier series

$$A(f) = \sum_{\alpha \in \hat{F}} \hat{A}_\alpha f^\alpha.$$

To describe the elements of $\hat{F}$, it is convenient to view them as functions from I to $\hat{G}$; then we can write

$$f^\alpha = \prod_{x \in I} (f(x))^{\alpha(x)}.$$

The latter expression is well-defined since $f(x) \in G$ and $\alpha(x) \in \hat{G}$.

Definition 8. *A function $A: F \to \boldsymbol{C}$ is* γ-homogeneous *for $\gamma \in \hat{G}$ if $A(gf) = g^\gamma A(f)$ for all $g \in G$.*

Lemma 4. *Let $F = G^I$ and suppose that $A: F \to \boldsymbol{C}$ is γ-homogeneous. Let $\hat{A}_\alpha$ be the Fourier coefficients of A. Then $\hat{A}_\alpha = 0$ unless $\gamma = \sum_{i \in I} \alpha(i)$. In particular, $\hat{A}_{\mathbf{0}} = 0$ if $\gamma \neq \mathbf{0}$.*

Now let J be a finite set satisfying $|I| \leq |J|$ and $\pi: J \to I$ be an onto mapping. Consider the space $H = G^J$. Any $f \in F$, i.e., any function from I to G can then be transformed in a canonical way to a function from J to G, i.e., to a member of H, by composing it with π; we denote this new function by $f \circ \pi$. Similarly, a $\beta \in \hat{H}$, i.e., a function from H to $\boldsymbol{C}$ can be associated in a natural way with a function from F to H, i.e., with a member of $\hat{F}$. We denote this new function by $\pi_G(\beta)$ and it is defined in such a way that $\pi_G(\beta)$, viewed as a homomorphism from F to $\boldsymbol{C}$, is the map $f \mapsto (f \circ \pi)^\beta$. It is easy to see that if $\alpha = \pi_G(\beta)$ then it holds that $\alpha(i) = \sum_{j \in \pi^{-1}(i)} \beta(j)$ for all $i \in I$.

Lemma 5. *Given an Abelian group G, two finite sets I and J such that $|I| \leq |J|$ and an onto mapping $\pi: J \to I$, let $F = G^I$ and $H = G^J$. Then, for any $\beta \in \hat{H}$, $(f \circ \pi)^\beta = f^{\pi_G(\beta)}$.*

As we already mentioned, in our case $I = \{-1, 1\}^U$ and $J = \mathrm{SAT}^W$ for some sequence W of u variables and some sequence U formed by selecting a variable from each clause in W. Our analysis involves the Fourier transform of functions $\gamma \circ A_U$ and $\gamma \circ A_W$ and it turns out that we need those functions to be γ-homogeneous. This can be enforced by certain access conventions in the verifier.

Definition 9. *A function A from F to G is* folded over G *if $A(gf) = gA(f)$ for all $g \in G$.*

Lemma 6. *If A is folded over G, then $\gamma \circ A$ is γ homogeneous.*

We can assume that all tables in the proof are folded since this can be simulated with the following access convention in the verifier: Partition G^F into equivalence classes by the relation $\equiv$, where $f \equiv h$ if there is $g \in G$ such that $f = gh$, i.e., $\forall w(f(w) = gh(w))$. Write $[f]$ for the equivalence class of f. Then, whenever the position corresponding to some function h is queried, return $gA(f)$ where g and f are such that $h = gf$ and f is the chosen representative for $[h]$.

7 Inapproximability results for Max E3-NAE-Sat(G)

In this section, our aim is to prove that unlike the Boolean case, for every $d \geq 3$, Max E3-NAE-Sat($\boldsymbol{Z}_d$) is approximation resistant. We will actually prove that Max E3-NAE-Sat(G) is approximation resistant for pretty much every finite Abelian group.

Theorem 4 (Main hardness result). *For every constant $\varepsilon > 0$ and every finite Abelian group G that is not isomorphic to $\boldsymbol{Z}_2^m$ for any positive integer m, it is **NP**-hard to distinguish instances of Max E3-NAE-Sat(G) that are satisfiable from instances where at most a fraction $(1 - |G|^{-2} + \varepsilon)$ of the constraints are simultaneously satisfiable.*

As a warm-up, we first prove the easier result that Max E3-LinInEq(G) is approximation resistant, even with perfect completeness, for every finite Abelian group of order at least three. This result can be used together with a simple gadget to prove that Max E3-NAE-Sat($\boldsymbol{Z}_3$) is approximation resistant. To prove the same result for all finite Abelian groups, we proceed in three stages: We first treat all groups of odd order, then $\boldsymbol{Z}_{2^m}$ for $m \geq 2$, and finally treat all groups by combining the proofs for the first two cases. In all our protocols we use the following notation:

Convention 1. When analyzing a PCP that tests if some μ-gap E3-Sat(5) formula Φ is satisfiable, expectations over U and W range over all sequences W of u clauses from Φ, with uniform measure, and all sequences U formed by selecting, uniformly and independently, one variable from each clause in W. Given U and W, we define the shorthands $F = G^{\{-1,1\}^U}$, $H = G^{\mathrm{SAT}^W}$ and let π be the projection that constructs an assignment in $\{-1,1\}^U$ from an assignment in SAT^W.

7.1 Intuition behind our PCP constructions

A verifier in a PCP typically first selects sets U and W uniformly at random and then checks a small number of positions in tables corresponding to U and W. Specifically, the standard way to get a PCP with three queries is to query one position in a table corresponding to U and two positions in a table corresponding to W. The three values obtained are then tested to see if they satisfy some given constraint—such a construction gives a hardness result for the CSP corresponding to the type of constraint checked. To get a hardness result for Max E3-NAE-Sat(G) the constraint checked by the verifier therefore has to be a not-all-equal constraint. Moreover, we want the verifier to have perfect completeness, i.e., to always verify a correct proof. We accomplish this by querying the positions $A_U(f)$, $A_W(h)$ and $A_W(f^{-1}h^2e)$ where f and h are selected uniformly at random and e is selected such that $e(y)$ is selected independently and uniformly at random from $G \setminus \{\mathbf{1}\}$. Here, the function $f^{-1}h^2e$ is the map $y \mapsto (f(y|_U))^{-1}(h(y))^2e(y)$. For a correct proof of a satisfying assignment, the answers to these queries will be $f(y|_U)$, $h(y)$ and $(f(y|_U))^{-1}(h(y))^2e(y)$ where y is a satisfying assignment to the clauses in W. These three values can never be all equal, since $f(y|_U) = h(y)$ implies that $(f(y|_U))^{-1}(h(y))^2e(y) = h(y)e(y) \neq h(y)$. Therefore the verifier always accepts a correct proof and to prove that it accepts a proof corresponding to an unsatisfying assignment with probability at most $1 - |G|^{-2} + \varepsilon$, where $\varepsilon > 0$ is an arbitrary constant, we proceed as follows: The assumption that the verifier accepts with probability $1 - |G|^{-2} + \varepsilon$ implies that a sum of certain pairs of related Fourier coefficients is large. Those coefficients can be used to devise strategies for the provers in the u-parallel version of the 2P1R game for μ-gap E3-Sat(5), strategies that make the verifier of that game accept with probability independent of u. This leads to a contradiction since it is known [15] that this protocol as soundness c_μ^u. The precise coupling can be expressed as follows:

Lemma 7. *Given a finite Abelian group G, suppose that $\gamma \in \hat{G} \setminus \{\mathbf{0}\}$ is arbitrary and that A_U and A_W are the folded tables of a standard written G-proof with parameter u that corresponds to an unsatisfiable instance of μ-gap E3-Sat(5). Then*

$$\mathrm{E}_{U,W}\left[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\beta|^{-1}\right] < \eta,$$

where U, W, F, H, and π are as in Convention 1, $A = \gamma \circ A_U$, and $B = \gamma \circ A_W$, provided that $u > \log \eta / \log c_\mu$.

The proof is a standard written G-proof with parameter u:

The verifier acts as follows:

1. Select a sequence W of clauses, each clause uniformly and independently at random from Φ.
2. Select a sequence U of variables by selecting one variable from each clause in W, uniformly and independently.
3. Let $\pi\colon \mathrm{SAT}^W \to \{-1,1\}^U$ be the function that creates an assignment in $\{-1,1\}^U$ from an assignment in SAT^W.
4. Let $F = G^{\{-1,1\}^U}$ and $H = G^{\mathrm{SAT}^W}$.
5. Select $f \in F$ and $h \in H$ uniformly at random.
6. Select $e \in H$ such that independently for every $y \in \mathrm{SAT}^W$, $e(y)$ is uniformly distributed in $G \setminus \{\mathbf{1}\}$.
7. Accept if $A_U(f)A_W(h)A_W((f \circ \pi)^{-1}h^{-1}e) \neq \mathbf{1}$; Reject otherwise.

Figure 6. The PCP used to prove optimal approximation hardness for Max E3-LinInEq(G) for any finite Abelian group G such that $|G| \geq 3$. The PCP is parameterized by the constant u and tests if a μ-gap E3-Sat(5) formula Φ is satisfiable.

Proof. We use the tables to construct strategies for the provers in the 2P1R game as follows: The provers first compute the $\gamma \in \hat{G} \setminus \{\mathbf{0}\}$ that maximizes

$$\mathrm{E}_{U,W}\left[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\beta|^{-1}\right]$$

and keep this γ fixed.

Given a sequence U of variables, the second prover computes the Fourier coefficients $\hat{A}_\alpha$ selects an α according to probability distribution given by $|\hat{A}_\alpha|^2$ and then an x such that $\alpha(x) \neq \mathbf{0}$ uniformly; this x is returned to the verifier.

Given a sequence W of clauses, the first prover computes the Fourier coefficients $\hat{B}_\beta$ selects a β according to probability distribution given by $|\hat{B}_\beta|^2$ and then a y such that $\beta(y) \neq \mathbf{0}$ uniformly; this y is returned to the verifier.

The assignment y always satisfies the clauses in W and it is guaranteed to be consistent if $\alpha = \pi_G(\beta)$ and the second prover happens to select precisely the y that projects onto the x selected by the first prover. Therefore, the success probability of the above strategy is at least

$$\mathrm{E}_{U,W}\left[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\beta|^{-1}\right].$$

Since it is known that the soundness of the 2P1R game is at most c_μ^u and the above expression is independent of u, the conclusion follows by selecting $u > \log \eta / \log c_\mu$. ∎

7.2 Warm-up: Linear inequations

The PCP is described in Fig. 6. The setting is as usual; the verifier first picks u clauses of the 3SAT instance at random and then a variable from each clause at random. It then queries three positions in tables in the proof that should correspond to an assignment to the clauses and variables and accepts if a certain inequation is satisfied. Perfect completeness follows easily:

Lemma 8. *The verifier in Fig. 6 has perfect completeness.*

Proof. Suppose that the proof is the correct encoding of some satisfying assignment y. Then $A_U(f) = f(y|_U)$ and $A_W(h) = h(y|_W)$, hence

$$A_U(f)A_W(h)A_W((f \circ \pi)^{-1}h^{-1}e) = e(y|_W) \neq \mathbf{1}$$

and the verifier accepts. ∎

Let us now study the error function e selected by the verifier.

Lemma 9. *Let e be picked as in Fig. 6 applied to a finite Abelian group of order at least three. Then $|\operatorname{E}[e^\beta]| \leq 2^{-|\beta|}$, where $|\beta|$ is defined as the size of the support of β, i.e., the number of y such that $\beta(y) \neq \mathbf{0}$.*

Proof. Since e is selected in such a way that the $e(y)$ are independent,

$$|\operatorname{E}[e^\beta]| = \prod_{y \in H} \left|\operatorname{E}\Big[(e(y))^{\beta(y)}\Big]\right|.$$

Since $g^{\mathbf{0}} = 1$ for all $g \in G$, the only factors that contribute are those where $\beta(y) \neq \mathbf{0}$. For those factors

$$\operatorname{E}\Big[(e(y))^{\beta(y)}\Big] = \frac{1}{|G|-1} \sum_{g \in G \setminus \{\mathbf{1}\}} g^{\beta(y)} = \frac{-1}{|G|-1},$$

where the last equality follows since $\sum_{g \in G} g^\gamma = 0$ for all $\gamma \neq \mathbf{0}$. Hence

$$|\operatorname{E}[e^{\beta_2}]| = (|G|-1)^{-|\beta|} \leq 2^{-|\beta|}.$$ ∎

The soundness is straightforward to analyze using the now standard methodology introduced by Håstad [10]. As usual, we assume that the test accepts with probability $1 - |G|^{-1} + \delta$ and prove that the tables in the proof must then be correlated. We then use this correlation to extract strategies for the provers in the 2P1R game.

Lemma 10. *Suppose that the verifier in Fig. 6 accepts with probability $1 - |G|^{-1} + \delta$ for some $\delta > 0$. Then there exists some $\gamma \in \hat{G} \setminus \{\mathbf{0}\}$ such that*

$$\operatorname{E}_{U,W}\left[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 4^{-|\beta|}\right] \geq \delta^2,$$

where U, W, F, H, and π are as in Convention 1, $\hat{A}_\alpha$ are the Fourier coefficients of $\gamma \circ A_U$, and $\hat{B}_\beta$ are the Fourier coefficients of $\gamma \circ A_W$.

Proof. The suggested test accepts unless $A_U(f)A_W(h)A_W((f \circ \pi)^{-1}h^{-1}e) = \mathbf{1}$, therefore we can write

$$\Pr[\text{accept}] = 1 - \frac{1}{|G|} \operatorname{E}\Big[\sum_{\gamma \in \hat{G}} (A_U(f)A_W(h)A_W((f \circ \pi)^{-1}h^{-1}e))^\gamma\Big]$$

$$= 1 - \frac{1}{|G|} - \frac{1}{|G|}\,\mathrm{E}\Big[\sum_{\gamma\in\hat{G}\setminus\{\mathbf{0}\}} (A_U(f)A_W(h)A_W((f\circ\pi)^{-1}h^{-1}e))^{\gamma}\Big].$$

Now suppose that $\Pr[\text{accept}] \geq 1 - |G|^{-1} + \delta$; then there exists a $\gamma \in \hat{G}$ such that

$$\left|\mathrm{E}\Big[(A_U(f)A_W(h)A_W((f\circ\pi)^{-1}h^{-1}e))^{\gamma}\Big]\right| \geq \delta.$$

Now expand $\gamma \circ A_U$ and $\gamma \circ A_W$ in their Fourier series for this value of γ:

$$(\gamma\circ A_U)(f) = \sum_{\alpha\in\hat{F}} \hat{A}_\alpha f^\alpha, \qquad (\gamma\circ A_W)(h) = \sum_{\beta\in\hat{H}} \hat{B}_\beta h^\beta.$$

This gives

$$\begin{aligned}
\delta &\leq \left|\mathrm{E}\Big[\sum_{\alpha\in\hat{F}}\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}} \hat{A}_\alpha\hat{B}_{\beta_1}\hat{B}_{\beta_2}\,\mathrm{E}\Big[f^\alpha h^{\beta_1}((f\circ\pi)^{-1}h^{-1}e)^{\beta_2}\Big]\Big]\right| \\
&= \left|\mathrm{E}\Big[\sum_{\alpha\in\hat{F}}\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}} \hat{A}_\alpha\hat{B}_{\beta_1}\hat{B}_{\beta_2}\,\mathrm{E}\Big[f^{\alpha-\pi_G(\beta_2)}h^{\beta_1-\beta_2}e^{\beta_2}\Big]\Big]\right| \\
&= \left|\mathrm{E}\Big[\sum_{\alpha\in\hat{F}}\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}} \hat{A}_\alpha\hat{B}_{\beta_1}\hat{B}_{\beta_2}\,\mathrm{E}\Big[f^{\alpha-\pi_G(\beta_2)}\Big]\,\mathrm{E}\Big[h^{\beta_1-\beta_2}\Big]\,\mathrm{E}\Big[e^{\beta_2}\Big]\Big]\right|.
\end{aligned}$$

The first of the inner expectations is zero unless $\alpha = \pi_G(\beta_2)$, the second of the inner expectations is zero unless $\beta_1 - \beta_2 = \mathbf{0}$. Putting this together, we get the bound

$$\begin{aligned}
\delta^2 &\leq \left|\mathrm{E}\Big[\sum_{\beta\in\hat{H}} \hat{A}_{\pi_G(\beta)}\hat{B}_\beta^2\,\mathrm{E}[e^\beta]\Big]\right|^2 \\
&\leq \mathrm{E}\left[\left|\sum_{\beta\in\hat{H}} \hat{A}_{\pi_G(\beta)}\hat{B}_\beta^2\,\mathrm{E}[e^\beta]\right|^2\right] \\
&\leq \mathrm{E}\left[\Big(\sum_{\beta\in\hat{H}} |\hat{A}_{\pi_G(\beta)}|^2|\hat{B}_\beta|^2|\,\mathrm{E}[e^\beta]|^2\Big)\Big(\sum_{\beta\in\hat{H}}|\hat{B}_\beta|^2\Big)\right] \\
&= \mathrm{E}\Big[\sum_{\beta\in\hat{H}} |\hat{A}_{\pi_G(\beta)}|^2|\hat{B}_\beta|^2|\,\mathrm{E}[e^\beta]|^2\Big] \\
&\leq \mathrm{E}\Big[\sum_{\beta\in\hat{H}} |\hat{A}_{\pi_G(\beta)}|^2|\hat{B}_\beta|^2 4^{-|\beta|}\Big],
\end{aligned}$$

where the equality follows since

$$\sum_{\beta\in\hat{H}} |\hat{B}_\beta|^2 = \frac{1}{|H|}\sum_{h\in H}|(A_W(h))^\gamma|^2 = 1$$

by Plancherel's equality. ∎

Theorem 5. *For every constant $\varepsilon > 0$ and every finite Abelian group G such that $|G| \geq 3$, it is **NP**-hard to distinguish instances of Max E3-LinInEq(G) that are satisfiable from instances where at most a fraction $(1 - |G|^{-1} + \varepsilon)$ of the inequations are simultaneously satisfiable.*

The proof is a standard written $\boldsymbol{Z}_d$-proof with parameter u:
The verifier acts as follows:
Steps 1–6 are as in Fig. 6 applied to $G = \boldsymbol{Z}_d$.
7. Accept if $A_U(f)$, $A_W(h)$ and $A_W((f \circ \pi)^{-1} h^2 e)$ are not all equal; Reject otherwise.

Figure 7. The PCP used to prove optimal approximation hardness for Max E3-NAE-Sat($\boldsymbol{Z}_d$) for odd d. The PCP is parameterized by the constant u and tests if a μ-gap E3-Sat(5) formula Φ is satisfiable.

Proof. Given ε, select $u > 2 \log \varepsilon / \log c_\mu$. Then Lemmas 10 and 7 together imply that the soundness of the PCP in Fig. 6 has soundness at most $1 - |G|^{-1} + \varepsilon$. ∎

Corollary 2. *Max E3-NAE-Sat($\boldsymbol{Z}_3$) is hard to approximate within $8/9+\varepsilon$ with perfect completeness for any $\varepsilon > 0$.*

Proof. We reduce Max E3-LinInEq(3), which we know to be hard to approximate within $2/3$ with perfect completeness from Theorem 5 applied to the group $\boldsymbol{Z}_3$, to Max E3-NAE-Sat($\boldsymbol{Z}_3$). A clause $x+y+z \neq 0 \bmod 3$ is replaced with the clauses NAE(x, y, z), NAE$(x+2, y+1, z)$, NAE$(x+1, y+2, z)$. If all three NAE clauses are satisfied, then $x + y + z \neq 0 \bmod 3$; if $x + y + z = 0$ two of the NAE clauses are satisfied. Therefore it is hard to distinguish the case when all the constraints are satisfied from the case when a fraction

$$\frac{3(2/3 + 3\varepsilon) + 2(1/3 - 3\varepsilon)}{3} = 8/9 + \varepsilon$$

of the constraints are satisfied. ∎

7.3 Case I: Groups of odd order

The PCP is described in Fig. 7. The setting is the same as in Sec. 7.2 but the acceptance predicate of the verifier is different.

Lemma 11. *The verifier in Fig. 7 has perfect completeness.*

Proof. Suppose that the proof is the correct encoding of a satisfying assignment. We get two cases: If $A_U(f) = f(y|_U)$ and $A_W(h) = h(y|_W)$ are not equal, the verifier accepts; if $f(y|_U) = h(y|_W)$, then $A_W((f \circ \pi)^{-1} h^2 e) = h(y|_W) e(y|_W) \neq h(y|_W) = A_W(h)$ and the verifier accepts. ∎

The soundness is a bit more complicated to analyze. As usual, we assume that the test accepts with probability $1 - |G|^{-2} + \delta$ and prove that the tables in the proof must then be correlated. We then use this correlation to extract strategies for the provers in the 2P1R game.

Lemma 12. *Suppose that the verifier in Fig. 7 accepts with probability $1-|G|^{-2}+\delta$ for some $\delta > 0$. Then there exists $\gamma_1 \in \hat{G} \setminus \{\mathbf{0}\}$ and $\gamma_2 \in \hat{G} \setminus \{\mathbf{0}\}$ such that*

$$\mathrm{E}_{U,W}\left[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 4^{-|\beta|}\right] \geq \delta^2,$$

where U, W, F, H, and π are as in Convention 1, $\hat{A}_\alpha$ are the Fourier coefficients of $\gamma_1 \circ A_U$, and $\hat{B}_\beta$ are the Fourier coefficients of $\gamma_2 \circ A_W$.

Proof. The test accepts unless all three queries return the same value, therefore we can write the acceptance probability as $1 - \sum_{g \in G} \mathrm{E}[I_g]$ where I_g denotes the indicator for the event that all three queries return g:

$$I_g = \frac{1}{|G|^3}\Big(1 + \sum_{\gamma_1 \neq \mathbf{0}} (g^{-1} A_U(f))^{\gamma_1}\Big)\Big(1 + \sum_{\gamma_2 \neq \mathbf{0}} (g^{-1} A_W(h))^{\gamma_2}\Big) \times \\ \Big(1 + \sum_{\gamma_3 \neq \mathbf{0}} (g^{-1} A_W((f \circ \pi)^{-1} h^2 e))^{\gamma_3}\Big).$$

We now expand the products and get

$$\begin{aligned} I_g = \frac{1}{|G|^3}\Big(1 &+ \sum_{\gamma_1 \neq \mathbf{0}} (g^{-1} A_U(f))^{\gamma_1} \\ &+ \sum_{\gamma_2 \neq \mathbf{0}} (g^{-1} A_W(h))^{\gamma_2} + \sum_{\gamma_3 \neq \mathbf{0}} (g^{-1} A_W((f \circ \pi)^{-1} h^2 e))^{\gamma_3} \\ &+ \sum_{\gamma_1 \neq \mathbf{0}} \sum_{\gamma_2 \neq \mathbf{0}} (g^{-1} A_U(f))^{\gamma_1} (g^{-1} A_W(h))^{\gamma_2} \\ &+ \sum_{\gamma_1 \neq \mathbf{0}} \sum_{\gamma_3 \neq \mathbf{0}} (g^{-1} A_U(f))^{\gamma_1} (g^{-1} A_W((f \circ \pi)^{-1} h^2 e))^{\gamma_3} \\ &+ \sum_{\gamma_2 \neq \mathbf{0}} \sum_{\gamma_3 \neq \mathbf{0}} (g^{-1} A_W(h))^{\gamma_2} (g^{-1} A_W((f \circ \pi)^{-1} h^2 e))^{\gamma_3} \\ &+ \sum_{\gamma_1 \neq \mathbf{0}} \sum_{\gamma_2 \neq \mathbf{0}} \sum_{\gamma_3 \neq \mathbf{0}} (g^{-1} A_U(f))^{\gamma_1} (g^{-1} A_W(h))^{\gamma_2} (g^{-1} A_W((f \circ \pi)^{-1} h^2 e))^{\gamma_3}\Big). \end{aligned}$$

Let us now consider the contributions of the above terms when we sum I_g over all $g \in G$. The constant term contributes $|G|/|G|^3 = 1/|G|^2$. The linear terms vanish since

$$\sum_{g \in G} \sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} (gh)^{\gamma}$$

where h is independent of g can be rewritten as

$$\sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} h^{\gamma} \sum_{g \in G} g^{\gamma} = 0$$

where the equality follows since $\sum_{g \in G} g^{\gamma} = 0$ for every $\gamma \neq \mathbf{0}$. The quadratic terms also disappear, but the argument is somewhat more involved. Consider terms of the first form:

$$\begin{aligned} &\sum_{g \in G} \sum_{\gamma_1 \neq \mathbf{0}} \sum_{\gamma_2 \neq \mathbf{0}} (g^{-1} A_U(f))^{\gamma_1} (g^{-1} A_W(h))^{\gamma_2} \\ &\qquad = \sum_{\gamma_1 \neq \mathbf{0}} \sum_{\gamma_2 \neq \mathbf{0}} (A_U(f))^{\gamma_1} (A_W(h))^{\gamma_2} \sum_{g \in G} g^{-\gamma_1 - \gamma_2}. \end{aligned}$$

For fixed, γ_1 and γ_2 such that $\gamma_1 + \gamma_2 \neq \mathbf{0}$, the inner sum is definitely 0. Therefore, we only have to care about

$$\sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} \mathrm{E}\Big[(A_U(f))^{\gamma} (A_W(h))^{-\gamma}\Big] = \sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} \sum_{\alpha \in \hat{F}} \sum_{\beta \in \hat{H}} \mathrm{E}\Big[\hat{A}_{\alpha} \hat{B}_{\beta} \, \mathrm{E}[f^{\alpha} h^{\beta}]\Big].$$

The inner expectation above is always zero unless $\alpha = \mathbf{0}$ and $\beta = \mathbf{0}$. Thanks to folding, we always have $\hat{B}_\beta = 0$ when $\beta = \mathbf{0}$, therefore all terms of the above form vanish. The terms of the second form need an additional argument. Expanding them as above, we get

$$\sum_{\gamma\in\hat{G}\setminus\{\mathbf{0}\}}\sum_{\alpha\in\hat{F}}\sum_{\beta\in\hat{H}}\mathrm{E}\Big[\hat{A}_\alpha\hat{B}_\beta\,\mathrm{E}[f^\alpha((f\circ\pi)^{-1}h^2e)^\beta]\Big]$$
$$=\sum_{\gamma\in\hat{G}\setminus\{\mathbf{0}\}}\sum_{\beta\in\hat{H}}\mathrm{E}\Big[\hat{A}_{\pi_G(\beta)}\hat{B}_\beta\,\mathrm{E}[h^{2\beta}]\,\mathrm{E}[e^\beta]\Big]. \tag{5}$$

The inner expectation $\mathrm{E}[h^{2\beta}]$ is zero as soon as $2\beta \neq \mathbf{0}$. It is true for groups of odd order that $\beta \neq \mathbf{0}$ implies $2\beta \neq \mathbf{0}$, therefore all terms of the above form vanish. For terms of the third form we get

$$\sum_{\gamma\in\hat{G}\setminus\{\mathbf{0}\}}\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}}\mathrm{E}\Big[\hat{B}_{\beta_1}\hat{B}_{\beta_2}\,\mathrm{E}[h^{\beta_1}((f\circ\pi)^{-1}h^2e)^{\beta_2}]\Big]$$
$$=\sum_{\gamma\in\hat{G}\setminus\{\mathbf{0}\}}\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}}\mathrm{E}\Big[\hat{B}_{\beta_1}\hat{B}_{\beta_2}\,\mathrm{E}[f^{-\pi_G(\beta_2)}]\,\mathrm{E}[h^{\beta_1-2\beta_2}]\,\mathrm{E}[e^{\beta_2}]\Big]$$
$$=\sum_{\gamma\in\hat{G}\setminus\{\mathbf{0}\}}\sum_{\beta\in\hat{H}}\mathrm{E}\Big[\hat{B}_{-2\beta}\hat{B}_\beta\,\mathrm{E}[f^{-\pi_G(\beta)}]\,\mathrm{E}[e^\beta]\Big].$$

The inner expectation $\mathrm{E}[f^{-\pi_G(\beta)}]$ is zero as soon as $\pi_G(\beta) \neq \mathbf{0}$. Since the tables are folded, $\hat{B}_\beta = 0$ for all β such that $\pi_G(\beta) = \mathbf{0}$, therefore also the terms of the above form vanish. Since we have killed all unwanted terms, the acceptance probability can then be written

$$1-\frac{1}{|G|^2}-\frac{1}{|G|^2}\sum_{\substack{\gamma_1,\gamma_2,\gamma_2\\ \gamma_1\neq\mathbf{0},\gamma_2\neq\mathbf{0},\gamma_3\neq\mathbf{0}\\ \gamma_1+\gamma_2+\gamma_3=\mathbf{0}}}\mathrm{E}\Big[(A_U(f))^{\gamma_1}(A_W((f\circ\pi)^{-1}h^2e))^{\gamma_2}(A_W(h))^{\gamma_3}\Big].$$

Suppose that the acceptance probability is at least $1-|G|^{-2}+\delta$ for some $\delta > 0$. Then there exists γ_1, γ_2 and γ_3 such that

$$\Big|\mathrm{E}\Big[(A_U(f))^{\gamma_1}(A_W((f\circ\pi)^{-1}h^2e))^{\gamma_2}(A_W(h))^{\gamma_3}\Big]\Big| \geq \delta.$$

Now expand the expression inside $\mathrm{E}[\cdot]$ in a Fourier series for those γ_1, γ_2, γ_3. We get

$$\delta \leq \Big|\mathrm{E}\Big[\sum_{\alpha\in\hat{F}}\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}}\hat{A}_\alpha\hat{B}_{\beta_1}\hat{C}_{\beta_2}f^\alpha((f\circ\pi)^{-1}h^2e)^{\beta_1}g^{\beta_2}\Big]\Big|$$
$$=\Big|\mathrm{E}\Big[\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}}\hat{A}_{\pi_G(\beta_1)}\hat{B}_{\beta_1}\hat{C}_{\beta_2}h^{2\beta_1+\beta_2}\,\mathrm{E}[e^{\beta_1}]\Big]\Big|$$
$$=\Big|\mathrm{E}\Big[\sum_{\beta\in\hat{H}}\hat{A}_{\pi_G(\beta)}\hat{B}_\beta\hat{C}_{-2\beta}\,\mathrm{E}[e^\beta]\Big]\Big|.$$

Putting the pieces together, we have shown that

$$\delta^2 \leq \Big|\mathrm{E}\Big[(A_U(f))^{\gamma_1}(A_W(h))^{\gamma_2}(A_W((f\circ\pi)^{-1}h^2e))^{\gamma_3}\Big]\Big|^2$$

The proof is a standard written $\boldsymbol{Z}_d$-proof with parameter u:
The verifier acts as follows:
Steps 1–5 are as in Fig. 6 applied to $G = \boldsymbol{Z}_d$.
6. Select $e_1 \in H$ such that independently for every $y \in \mathrm{SAT}^W$, $e_1(y)$ is uniformly distributed in $\{\omega^{4i}, \omega^{4i+1}\}_{i=0}^{d/4-1}$.
   Select $e_2 \in H$ such that independently for every $y \in \mathrm{SAT}^W$, $e_2(y)$ is uniformly distributed in $\{\omega^{4i+1}, \omega^{4i+2}\}_{i=0}^{d/4-1}$.
   Let $e = e_1 e_2$.
7. Accept if $A_U(f)$, $A_W(h)$ and $A_W((f \circ \pi)^{-1} h^2 e)$ are not all equal; Reject otherwise.

Figure 8. The PCP used to prove optimal approximation hardness for Max E3-NAE-Sat($\boldsymbol{Z}_d$) where $d = 2^m$ for integers $m \geq 2$. The group $\boldsymbol{Z}_d$ is represented by $\{\omega^i\}_{i=0}^{d-1}$ where $\omega = e^{2\pi i/d}$ and multiplication is the group operator. The PCP is parameterized by the constant u and tests if a μ-gap E3-Sat(5) formula Φ is satisfiable.

$$= \left| \mathrm{E}\Big[\sum_{\beta \in \hat{H}} \hat{A}_{\pi_G(\beta)} \hat{B}_\beta \hat{C}_{-2\beta} \, \mathrm{E}[e^\beta] \Big] \right|^2$$

$$\leq \mathrm{E}\left[\left| \sum_{\beta \in \hat{H}} \hat{A}_{\pi_G(\beta)} \hat{B}_\beta \hat{C}_{-2\beta} \, \mathrm{E}[e^\beta] \right|^2 \right]$$

We now apply Cauchy-Schwartz to rewrite the above bound as

$$\delta^2 \leq \mathrm{E}\left[\Big(\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\mathrm{E}[e^\beta]|^2 \Big) \Big(\sum_{\beta \in \hat{H}} |\hat{C}_{-2\beta}|^2 \Big) \right].$$

Since G has odd order, the second factor above can be bounded by

$$\sum_{\beta \in \hat{H}} |\hat{C}_{-2\beta}|^2 = \sum_{\beta \in \hat{H}} |\hat{C}_\beta|^2 = 1, \tag{6}$$

where the second equality is Plancherel's equality. Therefore,

$$\delta^2 \leq \mathrm{E}\left[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\mathrm{E}[e^\beta]|^2 \right]$$

$$\leq \mathrm{E}\left[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 4^{-|\beta|} \right]. \qquad \blacksquare$$

Theorem 6. *For every constant $\varepsilon > 0$ and every finite Abelian group G of odd order, it is **NP**-hard to distinguish instances of Max E3-NAE-Sat(G) that are satisfiable from instances where at most a fraction $(1 - |G|^{-2} + \varepsilon)$ of the constraints are simultaneously satisfiable.*

Proof. Given ε, select $u > 2 \log \varepsilon / \log c_\mu$. Then Lemmas 12 and 7 together imply that the PCP in Fig. 6 has soundness at most $1 - |G|^{-1} + \varepsilon$ $\blacksquare$

7.4 Case II: Z_d where $d = 2^m$ for integers $m \geq 2$

To handle this case we modify the protocol from Sec. 7.3 slightly. We represent $\boldsymbol{Z}_d$ as $\{\omega^i\}_{i=0}^{d-1}$ where $\omega = e^{2\pi i/d}$ and multiplication is the group operator. By representing the elements of the dual group—recall that they are actually the homomorphisms $x \mapsto x^n$ for integers n—as $\{0, 1, 2, \ldots, d-1\}$ with addition mod d as the group operator, we can still use the syntax g^γ to denote an element $\gamma \in \hat{\boldsymbol{Z}}_d$ acting on some $g \in \boldsymbol{Z}_d$. The only change in the protocol is the way we select error function. Recall that the fact that the underlying group has odd order was needed at two places in the proof. It was first needed for the quadratic terms from expression (5) in the proof of Lemma 12 to always be zero, and it was the necessary to for the first equality in the bound (6) to hold. In the protocol from Fig. 8, we select the error function in such a way that the proof can be made to work in the above two places also when G has order 2^m for integers $m \geq 2$. Let us first note some properties of the error functions

Lemma 13. *Let e_1 and e_2 be selected as in Fig. 8. Then*

$$e_1(y)e_2(y) \neq 1 \quad \textit{with probability } 1,$$

$$\mathrm{E}[(e_1(y))^\gamma] = \begin{cases} 1 & \textit{if } \gamma = 0, \\ (1+i)/2 & \textit{if } \gamma = d/4, \\ (1-i)/2 & \textit{if } \gamma = 3d/4, \\ 0 & \textit{otherwise.} \end{cases}$$

$$\mathrm{E}[(e_2(y))^\gamma] = \omega^\gamma \, \mathrm{E}[(e_2(y))^\gamma],$$

$$|\mathrm{E}[(e_1(y)e_2(y))^\gamma]| = \begin{cases} 1 & \textit{if } \gamma = 0, \\ 1/2 & \textit{if } \gamma = d/4 \textit{ or } \gamma = 3d/4, \\ 0 & \textit{otherwise.} \end{cases}$$

$$|\,\mathrm{E}[(e_1 e_2)^\beta]| = \begin{cases} 2^{-|\beta|} & \textit{if } \beta(y) \in \{0, d/4, 3d/4\} \textit{ for all } y, \\ 0 & \textit{otherwise.} \end{cases}$$

Proof. For e_1 and e_2 selected as in Fig. 8, $e_1(y)e_2(y)$ can assume the values

$$\omega^{4i}\omega^{4i'+1}, \quad \omega^{4i+1}\omega^{4i'+1}, \quad \omega^{4i+1}\omega^{4i'+1}, \quad \omega^{4i+1}\omega^{4i'+2},$$

or, expressed differently,

$$\omega^{4(i+i')+1}, \quad \omega^{4(i+i')+2}, \quad \omega^{4(i+i')+3},$$

for integers $0 \leq i, i' \leq d/4 - 1$. Since $\omega = e^{2\pi i/d}$, the expressions above are different from 1 if $4(i + i') \leq d - 3$. To prove the second equality, notice that

$$\mathrm{E}[(e_1(y))^\gamma] = \frac{1+\omega^\gamma}{d/2} \sum_{k=0}^{d/4-1} \omega^{4\gamma k} = \frac{(1+\omega^\gamma)(1-\omega^{\gamma d})}{d(1-\omega^{4\gamma})/2} = 0,$$

where the second equality is valid when $\omega^{4\gamma} \neq 1$, i.e., when γ is not an integer multiple of $d/4$. For the other cases,

$$\mathrm{E}[(e_1(y))^0] = \mathrm{E}[1] = 1,$$

$$\mathrm{E}[(e_1(y))^{d/4}] = \frac{1+\omega^{d/4}}{d/2} \sum_{k=0}^{d/4-1} 1 = \frac{1+i}{2},$$

$$\mathrm{E}[(e_1(y))^{d/2}] = \frac{1+\omega^{d/2}}{d/2} \sum_{k=0}^{d/4-1} 1 = 0,$$

$$\mathrm{E}[(e_1(y))^{3d/4}] = \frac{1+\omega^{3d/4}}{d/2} \sum_{k=0}^{d/4-1} 1 = \frac{1-i}{2}.$$

The third equality follows since $e_2(y)$ is distributed as $\omega e_1(y)$ and the fourth equality follows since e_1 and e_2 are selected independently. The fifth equality is an immediate consequence of the fourth.

■

Lemma 14. *Suppose that the verifier in Fig. 8 accepts with probability $1 - d^{-2} + \delta$ for some $\delta > 0$. Then there exists $\gamma_1 \in \hat{\boldsymbol{Z}}_d \setminus \{\mathbf{0}\}$ and $\gamma_2 \in \hat{\boldsymbol{Z}}_d \setminus \{\mathbf{0}\}$ such that*

$$\mathrm{E}_{U,W}\Bigg[\sum_{\substack{\beta \in \hat{H} \\ \beta(y) \in \{0, d/4, 3d/4\} \forall y}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 2^{-|\beta|} \Bigg] \geq \delta^2,$$

where U, W, F, H, and π are as in Convention 1 with $G = \boldsymbol{Z}_d$, $\hat{A}_\alpha$ are the Fourier coefficients of $\gamma_1 \circ A_U$, and $\hat{B}_\beta$ are the Fourier coefficients of $\gamma_2 \circ A_W$.

Proof. If we let $G = \boldsymbol{Z}_d$ and denote the product $e_1 e_2$ by e, the proof proceeds exactly as the proof of Lemma 12 up to the expression (5):

$$\sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} \sum_{\beta \in \hat{H}} \mathrm{E}\Big[\hat{A}_{\pi_G(\beta)} \hat{B}_\beta \, \mathrm{E}[h^{2\beta}] \, \mathrm{E}[e^\beta] \Big].$$

The argument used in the proof of Lemma 12 to prove that this expression vanishes is not valid in this case since it relies on the fact that the domain size is odd. Instead, we use our modified error function together with some other observations. Consider the term

$$\hat{A}_{\pi_G(\beta)} \hat{B}_\beta \, \mathrm{E}[h^{2\beta}] \, \mathrm{E}[e^\beta] = \hat{A}_{\pi_G(\beta)} \hat{B}_\beta \, \mathrm{E}[h^{2\beta}] \, \mathrm{E}[e_1^\beta] \, \mathrm{E}[e_2^\beta]$$

for a fixed β. We now argue that this term vanishes for every β. Since $\mathrm{E}[h^{2\beta}] = 0$ as soon as $2\beta \neq \mathbf{0}$, we only need to consider the case when $\beta(y) \in \{0, d/2\}$ for all $y \in \mathrm{SAT}^W$. If there exists a y_0 such that $\beta(y_0) = d/2$,

$$\mathrm{E}[(e_1(y_0))^{\beta(y_0)}] = \mathrm{E}[(e_1(y_0))^{d/2}] = 0$$

by Lemma 13, and therefore

$$\mathrm{E}[e_1^\beta] = \prod_{y \in \mathrm{SAT}^W} \mathrm{E}[(e_1(y))^{\beta(y)}] = 0.$$

Finally, $\hat{B}_\beta = 0$ as soon as $\beta(y) = 0$ for all y since the tables are folded, therefore the term vanishes also in this case. To conclude, all terms of the above form always vanish, for every β.

The argument used to kill the quadratic forms of the third form in the proof of Lemma 12 is valid also in this case since it does not rely on any assumption regarding the domain size. Therefore, the acceptance probability can then be written

$$1 - \frac{1}{d^2} - \frac{1}{d^2} \sum_{\substack{\gamma_1,\gamma_2,\gamma_2 \\ \gamma_1 \neq \mathbf{0}, \gamma_2 \neq \mathbf{0}, \gamma_3 \neq \mathbf{0} \\ \gamma_1+\gamma_2+\gamma_3=\mathbf{0}}} \mathrm{E}\Big[(A_U(f))^{\gamma_1}(A_W((f\circ\pi)^{-1}h^2e))^{\gamma_2}(A_W(h))^{\gamma_3}\Big].$$

Suppose that the acceptance probability is at least $1 - d^{-2} + \delta$ for some $\delta > 0$. Then there exists γ_1, γ_2 and γ_3 such that

$$\Big|\mathrm{E}\Big[(A_U(f))^{\gamma_1}(A_W((f\circ\pi)^{-1}h^2e))^{\gamma_2}(A_W(h))^{\gamma_3}\Big]\Big| \geq \delta.$$

Now expand the expression inside $\mathrm{E}[\cdot]$ in a Fourier series for those γ_1, γ_2, γ_3. We get

$$\begin{aligned}
\delta^2 &\leq \left|\mathrm{E}\Big[\sum_{\alpha\in\hat{F}}\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}} \hat{A}_\alpha\hat{B}_{\beta_1}\hat{C}_{\beta_2}f^\alpha((f\circ\pi)^{-1}h^2e)^{\beta_1}g^{\beta_2}\Big]\right|^2 \\
&= \left|\mathrm{E}\Big[\sum_{\beta_1\in\hat{H}}\sum_{\beta_2\in\hat{H}} \hat{A}_{\pi_G(\beta)}\hat{B}_{\beta_1}\hat{C}_{\beta_2}h^{2\beta_1+\beta_2}\,\mathrm{E}[e^{\beta_1}]\Big]\right|^2 \\
&= \left|\mathrm{E}\Big[\sum_{\beta\in\hat{H}} \hat{A}_{\pi_G(\beta)}\hat{B}_{\beta}\hat{C}_{-2\beta}\,\mathrm{E}[e^{\beta}]\Big]\right|^2 \\
&= \left|\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H} \\ \beta(y)\in\{0,d/4,3d/4\}\forall y}} \hat{A}_{\pi_G(\beta)}\hat{B}_{\beta}\hat{C}_{-2\beta}\,\mathrm{E}[e^{\beta}]\Big]\right|^2.
\end{aligned}$$

where the last equality follows from Lemma 13. We now apply the Cauchy-Schwartz inequality to the above bound:

$$\begin{aligned}
\delta^2 \leq \mathrm{E}\Big[&\Big(\sum_{\substack{\beta\in\hat{H} \\ \beta(y)\in\{0,d/4,3d/4\}\forall y}} |\hat{A}_{\pi_G(\beta)}|^2|\hat{B}_\beta|^2|\,\mathrm{E}[e^\beta]|\Big)\times \\
&\Big(\sum_{\substack{\beta\in\hat{H} \\ \beta(y)\in\{0,d/4,3d/4\}\forall y}} |\hat{C}_{-2\beta}|^2|\,\mathrm{E}[e^\beta]|\Big)\Big].
\end{aligned} \tag{7}$$

To bound the second factor above, we collect terms containing the same Fourier coefficient $\hat{C}_{-2\beta}$. Notice that each β such that $\beta(y) \in \{0, d/4, 3d/4\}$ for all y maps onto a β' such that $\beta(y) = 0$ implies that $\beta'(y) = 0$ and $\beta(y) \in \{d/4, 3d/4\}$ implies that $\beta'(y) = d/2$. Therefore, $|\beta| = |\beta'|$ and there are $2^{|\beta'|} = 2^{|\beta|}$ different β that map onto each β' with the property that $\beta' \in \{0, d/2\}$ for all y. To sum up, the second factor above is

$$\sum_{\substack{\beta\in\hat{H} \\ \beta(y)\in\{0,d/4,3d/4\}\forall y}} |\hat{C}_{-2\beta}|^2 2^{-|\beta|} = \sum_{\substack{\beta'\in\hat{H} \\ \beta'(y)\in\{0,d/2\}\forall y}} |\hat{C}_{\beta'}|^2 \leq 1.$$

The proof is a standard written G-proof with parameter u:

The verifier acts as follows:

Steps 1–5 are as in Fig. 6.

6. Select e by selecting independently the components of $e(y)$ according to Step 6 in Figs. 7 and 8, respectively.
7. Accept if $A_U(f)$, $A_W(h)$ and $A_W((f \circ \pi)^{-1} h^2 e)$ are not all equal; Reject otherwise.

Figure 9. The PCP used to prove optimal approximation hardness for Max E3-NAE-Sat(G) for any finite Abelian group $G \cong G_0 \times \boldsymbol{Z}_{2^{\alpha_1}} \times \boldsymbol{Z}_{2^{\alpha_2}} \times \cdots \times \boldsymbol{Z}_{2^{\alpha_s}}$ where G_0 is a finite Abelian group of odd order and $\alpha_i > 1$ for all i. The PCP is parameterized by the constant u and tests if a μ-gap E3-Sat(5) formula Φ is satisfiable.

Inserting the above bound into (7) gives

$$\begin{aligned}\delta^2 &\le \mathrm{E}\Bigg[\sum_{\substack{\beta \in \hat{H} \\ \beta(y) \in \{0, d/4, 3d/4\} \forall y}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\,\mathrm{E}[e^\beta]| \Bigg] \\ &\le \mathrm{E}\Bigg[\sum_{\substack{\beta \in \hat{H} \\ \beta(y) \in \{0, d/4, 3d/4\} \forall y}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 2^{-|\beta|} \Bigg].\end{aligned}$$

∎

Theorem 7. *For every constant $\varepsilon > 0$ and every $d = 2^m$ where $m \ge 2$ is an arbitrary integer, it is **NP**-hard to distinguish instances of Max E3-NAE-Sat($\boldsymbol{Z}_d$) that are satisfiable from instances where at most a fraction $(d^{-2} + \varepsilon)$ of the constraints are simultaneously satisfiable.*

Proof. Given ε, select $u > 2 \log \varepsilon / \log c_\mu$. Then Lemmas 14 and 7 together imply that the PCP in Fig. 6 has soundness at most $1 - |G|^{-1} + \varepsilon$ ∎

7.5 Proof of Theorem 4

Remember that every finite Abelian group is isomorphic to some group of the form

$$\boldsymbol{Z}_{p_1^{\alpha_1}} \times \boldsymbol{Z}_{p_2^{\alpha_2}} \times \cdots \times \boldsymbol{Z}_{p_s^{\alpha_s}} \tag{8}$$

where the p_i are not necessarily distinct primes and $|G| = p_1^{\alpha_1} p_2^{\alpha_2} \cdots p_s^{\alpha_s}$. To prove Theorem 4 we first combine the protocols of the previous sections to prove the result for groups such that $p_i^{\alpha_i} \ne 2$ for all i in the expansion (8). This result is then combined with a simple gadget reduction to complete the proof.

Lemma 15. *Suppose that the verifier in Fig. 9 accepts with probability $1 - |G|^{-2} + \delta$ for some $\delta > 0$. Then there exists $\gamma_1 \in \hat{G} \setminus \{\mathbf{0}\}$ and $\gamma_2 \in \hat{G}_d \setminus \{\mathbf{0}\}$ such that*

$$\mathrm{E}_{U,W}\Bigg[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 2^{-|\beta|} \Bigg] \ge \delta^2,$$

where G is a finite Abelian group such that $p_i^{\alpha_i} \ne 2$ for all i in the expansion (8), U, W, F, H, and π are as in Convention 1, $\hat{A}_\alpha$ are the Fourier coefficients of $\gamma_1 \circ A_U$, and $\hat{B}_\beta$ are the Fourier coefficients of $\gamma_2 \circ A_W$.

Proof. Since the proof is very similar to the proofs of Lemmas 12 and 14 we only give the most important details. From the requirements on G from the formulation of the lemma, we can deduce that G is isomorphic to a product of the form $G_0 \times \boldsymbol{Z}_{2^{\alpha_1}} \times \boldsymbol{Z}_{2^{\alpha_2}} \times \cdots \times \boldsymbol{Z}_{2^{\alpha_s}}$ where G_0 is a finite Abelian group of odd order and $\alpha_i > 1$ for all i. Any $g \in G$ can thus be written as a tuple $(g_0, g_1, g_2, \ldots, g_s)$ where $g_0 \in G_0$ and $g_i \in \boldsymbol{Z}_{2^{\alpha_i}}$ for $i = 1, 2, \ldots, s$. Similarly, any $\gamma \in \hat{G}$ can be decomposed as $(\gamma_0, \gamma_1, \gamma_2, \ldots, \gamma_s)$ and it then holds that $g^\gamma = g_0^{\gamma_0} g_1^{\gamma_1} g_2^{\gamma_2} \cdots g_s^{\gamma_s}$.

The two critical points of the previous proofs are the quadratic terms of the form (5) and the application of Plancherel's equality. Let us first consider the quadratic terms:

$$\sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} \sum_{\beta \in \hat{H}} \mathrm{E}\Big[\hat{A}_{\pi_G(\beta)} \hat{B}_\beta \, \mathrm{E}[h^{2\beta}] \, \mathrm{E}[e^\beta]\Big].$$

Consider an arbitrary term in the sum. Since h is selected uniformly and e is selected componentwise, the inner expectations can be written as

$$\mathrm{E}[h^{2\beta}] \, \mathrm{E}[e^\beta] = \mathrm{E}[h_0^{2\beta_0}] \, \mathrm{E}[e_0^{\beta_0}] \prod_{i=1}^{s} \mathrm{E}[h_i^{2\beta_i}] \, \mathrm{E}[e_i^{\beta_i}]$$

Now we know from the proof of Lemma 12 that $\mathrm{E}[h_0^{2\beta_0}] \, \mathrm{E}[e_0^{\beta_0}] = 0$ as soon as $\beta_0 \neq \mathbf{0}$. Similarly, we now from the proof of Lemma 14 that $\mathrm{E}[h_i^{2\beta_i}] \, \mathrm{E}[e_i^{\beta_i}] = 0$ as soon as $\beta_i \neq 0$. And, finally, if $\beta_0 = \mathbf{0}$ and $\beta_1 = \beta_2 = \cdots = \beta_s = 0$, $\hat{B}_\beta = 0$ thanks to folding. Therefore,

$$\sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} \sum_{\beta \in \hat{H}} \mathrm{E}\Big[\hat{A}_{\pi_G(\beta)} \hat{B}_\beta \, \mathrm{E}[h^{2\beta}] \, \mathrm{E}[e^\beta]\Big] = 0.$$

We now turn to the application of the Cauchy-Schartz inequality and Parseval's equality. With arguments similar to those in the proofs of Lemmas 12 and 14 it follows that

$$\begin{aligned}
\delta^2 &\leq \left| \mathrm{E}\Big[\sum_{\beta \in \hat{H}} \hat{A}_{\pi_G(\beta)} \hat{B}_\beta \hat{C}_{-2\beta} \, \mathrm{E}[e^\beta] \Big] \right|^2 \\
&\leq \mathrm{E}\Big[\Big(\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\, \mathrm{E}[e^\beta]| \Big) \Big(\sum_{\beta \in \hat{H}} |\hat{C}_{-2\beta}|^2 |\, \mathrm{E}[e^\beta]| \Big) \Big]
\end{aligned}$$

under the assumption that the test in Fig. 9 accepts with probability $1 - |G|^{-2} + \delta$. In the above expression, all sums over β are over all $\beta \in \hat{H}$ such that for all $i \in \{1, 2, \ldots, s\}$ it holds that $\beta_i(y) \in \{0, d/4, 3d/4\}$ for all y. We now want to bound the last sum using Parseval's equality. As in the proof of Lemma 14, we gather terms containing the same $\hat{C}_{\beta'}$. Since the original sum is only over $\beta \in \hat{H}$ such that for all $i \in \{1, 2, \ldots, s\}$ it holds that $\beta_i(y) \in \{0, d/4, 3d/4\}$ for all y, there are, for each β', at most $2^{|\beta'|} = 2^{|\beta|}$ different β with the property that $-2\beta = \beta'$. Therefore, the last sum can be bounded by

$$\sum_{\beta \in \hat{H}} |\hat{C}_{-2\beta}|^2 |\, \mathrm{E}[e^\beta]| = \sum_{\beta \in \hat{H}} |\hat{C}_{-2\beta}|^2 2^{-|\beta|} \leq \sum_{\beta' \in \hat{H}} |\hat{C}_{\beta'}|^2 = 1.$$

To conclude, we have shown that

$$\delta^2 \leq \mathrm{E}\Big[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\, \mathrm{E}[e^\beta]| \Big] \leq \mathrm{E}\Big[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 2^{-|\beta|} \Big]. \qquad \blacksquare$$

Lemma 16. *Suppose that, for every constant $\varepsilon > 0$, it is **NP**-hard to distinguish instances of Max E3-NAE-Sat(G) that are satisfiable from instances where at most a fraction $(1 - |G|^{-2} + \varepsilon)$ of the constraints are simultaneously satisfiable for some finite Abelian group G. Then, for every constant $\varepsilon > 0$, it is **NP**-hard to distinguish instances of Max E3-NAE-Sat($\boldsymbol{Z}_2 \times G$) that are satisfiable from instances where at most a fraction $(1-(2|G|)^{-2}+\varepsilon)$ of the constraints are simultaneously satisfiable.*

Proof. Given an arbitrary clause $\mathrm{NAE}(g_i x_i, g_j x_j, g_k x_k)$ from the Max E3-NAE-Sat(G) instance construct four clauses

$$\mathrm{NAE}((1, g_i)y_i, (1, g_j)y_j, (1, g_k)y_k),$$
$$\mathrm{NAE}((-1, g_j)y_j, (1, g_j)y_j, (1, g_k)y_k),$$
$$\mathrm{NAE}((1, g_i)y_i, (-1, g_j)y_j, (1, g_k)y_k),$$
$$\mathrm{NAE}((1, g_i)y_i, (1, g_j)y_j, (-1, g_k)y_k)$$

in the Max E3-NAE-Sat($\boldsymbol{Z}_2 \times G$) instance for the group $\boldsymbol{Z}_2 \times G$. Above we use multiplication as the group operator and represent $\boldsymbol{Z}_2$ as $\{-1, 1\}$.

Given a solution to the Max E3-NAE-Sat($\boldsymbol{Z}_2 \times G$) instance, construct a solution to the Max E3-NAE-Sat(G) in the obvious way. If x, y, and z are not all equal, all four clauses will be satisfied—if they are all equal, the last three clauses will be satisfied. Therefore, it is hard to distinguish the case when all constraints are satisfied from the case when a fraction

$$\frac{4(1 - d^{-2} + 4\varepsilon) + 3(d^{-2} - 4\varepsilon)}{4} = \frac{4 - d^{-2} + 4\varepsilon}{4} = 1 - (2d)^{-2} + \varepsilon$$

of the constraints are satisfied. ∎

8 The status of Max E2-Sat(d)

Although we were not able to completely resolve the status of Max E2-Sat(d) as far as approximability is concerned, we have obtained some evidence regarding its hardness. In this paper, we resolved the status of the easier Max Co-BIJ(d) problem—it is *not* approximation resistant according to Theorem 1—and the harder Max E3-NAE-Sat($\boldsymbol{Z}_d$) problem—it *is* approximation resistant according to Theorem 4.

For the Max E2-Sat(d) problem itself, we present some hardness results below. We first prove a result for domain size 3 and then prove a result for general domains.

Lemma 17. *For every constant $\varepsilon > 0$, the predicate $(x \neq a) \vee (y = b)$ over domains of size 3 is hard to approximate within $(23/24 + \varepsilon)$ with perfect completeness.*

Proof. Consider the following 2P1R interactive proof system for 3SAT: The first prover is given a variable and returns an assignment to that variable, the second prover is given a clause and returns an index of a literal that makes the clause satisfied. The verifier selects a clause at random, then a variable in the clause at random, sends the variable to P_1, the clause to P_2 and accepts unless P_2 returns the index of the variable sent to P_1 and P_1 returned an assignment that does not satisfy the literal. It is known that there are satisfiable instances of 3SAT such that it is **NP**-hard to satisfy more than $7/8 + \varepsilon$ of the clauses, for any constant $\varepsilon > 0$ [10]. When the above protocol is applied to such an instance of 3SAT, the test has perfect completeness and soundness $(1 - 1/24 + \varepsilon)$. To obtain the hardness for the claimed constraint satisfaction problem, we just use the following reduction: x specifies the name of a clause, y specifies a variable in this clause, a specifies the location of y in

x (encoded as 0,1 or 2), and b specifies a Boolean assignment to y (encoded over 0,1,2, where 2 is meaningless). ∎

Theorem 8. *For every constant $\varepsilon > 0$, it is **NP**-hard to approximate Max E2-Sat(3) within $47/48+\varepsilon$ with perfect completeness.*

Proof. Follows from Lemma 17 since $(x \neq a) \vee (y = b)$ can be written as a two E2-Sat(3) clauses. ∎

Theorem 9. *For every $d \geq 3$ and every constant $\varepsilon > 0$, Max E2-Sat(d) is hard to approximate within a factor of $(1 - d^{-4} + \varepsilon)$ with perfect completeness.*

Proof. We reduce Max E3-Sat(d), which is known to be hard to approximate within $(1 - d^{-3} + d\varepsilon)$ with perfect completeness to Max E2-Sat(d). A constraint SAT(x, y, z), which requires that at least one of x, y, z does not equal 0, is replaced with the constraints SAT(x, t), SAT($x, t+1$), SAT($x, t+2$), ..., SAT($x, t + d - 3$), SAT($y, t + d - 2$), SAT($z, t + d - 1$), where t is an auxiliary variable specific to this constraint and the additions are done modulo d. If all d 2SAT clauses are satisfied, the 3SAT clause has to be satisfied; if the 3SAT clause is not satisfied we can satisfy $d - 1$ of the 2SAT clauses. Therefore it is hard to distinguish the case when all the constraints are satisfied from the case when a fraction $\frac{1}{d}(d(1 - d^{-3} + d\varepsilon) + (d - 1)(d^{-3} - d\varepsilon)) = (1 - d^{-4} + \varepsilon)$ of the constraints are satisfied. ∎

Theorem 10. *Max E2-Sat(d) is hard to approximate within a factor $1 - \Omega(d^{-2})$ with non-perfect completeness. It is also hard to approximate within a factor $1 - \Omega(d^{-3})$ with perfect completeness for all $d \geq 3$.*

Proof. We reduce d-CUT, which is hard to approximate within $1 - 1/34d + \varepsilon$ with non-perfect completeness [11], to Max E2-Sat(d). A clause CUT(x, y) is replaced with the clauses SAT($x+i, y+i$) for all i from 0 to $d - 1$. A d-CUT instance with n constraints corresponds to a 2SAT(d) instance with dn constraints and an assignment satisfying all but k 2SAT(d) constraints satisfies all but k d-CUT constraints.

The hardness result with perfect completeness follows using the above reduction together with the result of Lemma 18 below which shows a factor $1 - \Omega(1/d^2)$ inapproximability result for Max d-CUT with perfect completeness, or in other words Max d-CUT on d-colorable graphs. Combining with the above gadget that reduces d-CUT to 2SAT(d), we get the claimed $1 - \Omega(1/d^3)$ hardness for satisfiable instances of Max E2-Sat(d) for all $d \geq 3$. ∎

Lemma 18. *There is a constant $\gamma > 0$ such that for all $d \geq 3$, given as input a d-colorable graph, it is NP-hard to find a d-cut that cuts a fraction $1 - \gamma/d^2$ of the edges.*

Proof. The reduction is from Max 3-CUT on 3-colorable graphs (i.e., Max 3-CUT with perfect completeness), which is known to be hard to approximate within a factor $(1 - \delta)$ for some absolute constant $\delta > 0$ [14]. This already gives the result for $d = 3$, so assume $d \geq 4$. Given a graph $G = (V, E)$ construct a new *weighted* graph H as follows. H will contain G as an induced subgraph. It will further contain $(d - 3)$ new vertices $u_1, u_2, \ldots, u_{d-3}$, and each u_i is connected to each $v \in V$ by an edge of weight $d(v)$, the degree of v in G. Also, between u_i and u_j for $1 \leq i < j \leq d - 3$, we add an edge of weight $2m$ where m is the number of edges in G.*

*For the reduction, it is possible to be more frugal in the size of the weights we assign, but since we are not optimizing for the value of γ in the statement of the lemma, we chose large enough weights to make the proof very simple.

The proof is a standard written G-proof with parameter u:

The verifier acts as follows:

Steps 1–4 are as in Fig. 6.

5. Select $f \in F$ and $h_1, h_2, \ldots, h_{k-2} \in H$ uniformly at random.
6. Select $e \in H$ such that independently for every $y \in \mathrm{SAT}^W$, $e(y)$ is uniformly distributed in $G \setminus \{\mathbf{1}\}$.
7. Define $h_{k-1} = (f \circ \pi)^{-1}(h_1 h_2 \cdots h_{k-2})^{-1} e$.
8. Accept if $A_U(f) \prod_{i=1}^{k-1} A_W(h_i) \neq \mathbf{1}$; Reject otherwise.

Figure 10. The PCP used to prove optimal approximation hardness for Max Ek-LinInEq(G) for any $k \geq 3$ and any finite Abelian group G such that $|G| \geq 3$. The PCP is parameterized by the constant u and tests if a μ-gap E3-Sat(5) formula Φ is satisfiable.

Clearly, H is d-colorable if G is 3-colorable as one can give $(d-3)$ new colors to the u_i's. Furthermore, given a d-coloring of H that fails to cut a set of edges of total weight W, we can use it to find a 3-coloring of G that fails to cut at most W edges. Indeed, given a d-coloring of H, we can first modify it so that the u_i's receive distinct colors without increasing the weight of the uncut edges. We can do so since if, say, u_1 and u_2 receive the same color, then we can recolor u_2 by a color that is not used by any of u_i's; under the new coloring the edge (u_1, u_2) of weight $2m$ is cut, and the worst that could happen is that all edges (u_2, v) become uncut and these have a total weight of only $2m$. Once the u_i's all get distinct colors, in the next step, if some $v \in V$ and u_i have the same color, we recolor v using one of the 3 colors not used by $u_1, u_2, \ldots, u_{d-3}$, until no edge of the form (v, u_i) is uncut. These steps do not increase the weight of the uncut edges since the weight of an edge (v, u_i) is greater than the total weight of all edges incident upon v in G.

We now have a d-coloring of H in which the vertices in V receive three colors distinct from the $(d-3)$ colors used to color the u_i's and thus the only uncut edges are those of G. Since the weight of uncut edges was originally W and could have only decreased in the process of recoloring H, it follows that there is a 3-coloring of G that leaves at most W uncut (i.e., miscolored) edges.

The total weight of edges of H, say M, is at most $d^2 m$. By the above discussion, if the maximum 3-cut of G cuts at most $(1-\delta)m$ edges, then the maximum d-cut of H cuts edges of total weight at most $M - \delta m \leq (1 - \delta/d^2)M$.

Hence, it is **NP**-hard to approximate Max d-cut on *weighted* graphs within a factor of $(1 - \Omega(1/d^2))$ even with perfect completeness. This is almost the result we want except that we only showed hardness for weighted graphs. However, one can deduce inapproximability within the same factor also for unweighted graphs by appealing to a general procedure for getting rid of weights due to Crescenzi, Silvestri, and Trevisan [5]. ∎

The result of Theorem 9 is not entirely subsumed by the result of Theorem 10 for satisfiable instances, since the constant in front of $1/d^3$ implies that the result of Theorem 9 will actually be stronger for small values of d. An interesting question is whether a factor $(1 - \Omega(d^{-2}))$ hardness can be shown for satisfiable instances of Max E2-Sat(d).

Conjecture. Although we did not completely resolve the status of Max E2-Sat(d), we conjecture at this point that the problem is *not* approximation resistant.

9 The status of Max Ek-Sat(d) for $k \geq 3$

Theorem 5 generalizes to Max Ek-LinInEq(G) for any $k > 3$ by a small modification to the protocol. The modified PCP is described in Fig. 10.

Lemma 19. *The verifier in Fig. 10 has perfect completeness.*

Proof. Suppose that the proof is the correct encoding of some satisfying assignment y. Then

$$\begin{aligned} A_U(f) \prod_{i=1}^{k-1} A_W(h_i) &= f(\pi(y)) \Big(\prod_{i=1}^{k-2} h_i(y) \Big) (f \circ \pi)^{-1}(\pi(y)) \Big(\prod_{i=1}^{k-2} h_i^{-1}(y) \Big) e(y) \\ &= e(y) \neq \mathbf{1} \end{aligned}$$

and the verifier accepts. ∎

Lemma 20. *Suppose that the verifier in Fig. 6 accepts with probability $1 - |G|^{-1} + \delta$ for some $\delta > 0$. Then there exists some $\gamma \in \hat{G} \setminus \{\mathbf{0}\}$ such that*

$$\mathrm{E}_{U,W} \Big[\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\beta|^{-1} \Big] \geq \delta^2,$$

where U, W, F, H, and π are as in Convention 1, $\hat{A}_\alpha$ are the Fourier coefficients of $\gamma \circ A_U$, and $\hat{B}_\beta$ are the Fourier coefficients of $\gamma \circ A_W$.

Proof. The suggested test accepts unless $A_U(f) \prod_{i=1}^{k-1} A_W(h_i) = \mathbf{1}$, therefore

$$\Pr[\text{accept}] = 1 - \frac{1}{|G|} - \frac{1}{|G|} \mathrm{E} \Big[\sum_{\gamma \in \hat{G} \setminus \{\mathbf{0}\}} \Big(A_U(f) \prod_{i=1}^{k-1} A_W(h_i) \Big)^\gamma \Big].$$

If $\Pr[\text{accept}] \geq 1 - |G|^{-1} + \delta$ there exists a $\gamma \in \hat{G}$ such that

$$\Big| \mathrm{E} \Big[\Big(A_U(f) \prod_{i=1}^{k-1} A_W(h_i) \Big)^\gamma \Big] \Big| \geq \delta.$$

Now expand $\gamma \circ A_U$ and $\gamma \circ A_W$ in their Fourier series for this value of γ:

$$\begin{aligned} \delta^2 &\leq \Bigg| \sum_{\alpha \in \hat{F}} \sum_{\beta_1 \in \hat{H}} \sum_{\beta_2 \in \hat{H}} \cdots \sum_{\beta_{k-1} \in \hat{H}} \hat{A}_\alpha \hat{B}_{\beta_1} \hat{B}_{\beta_2} \cdots \hat{B}_{\beta_{k-1}} \\ &\qquad \mathrm{E} \Big[f^\alpha \prod_{i=1}^{k-2} h_1^{\beta_i} \Big((f \circ \pi)^{-1} \prod_{i=1}^{k-2} h_i^{-1} e \Big)^{\beta_{k-1}} \Big] \Bigg|^2 \\ &= \Bigg| \mathrm{E} \Big[\sum_{\beta \in \hat{H}} \hat{A}_{\pi_G(\beta)} \hat{B}_\beta^{k-1} (1 - |G|)^{-|\beta|} \Big] \Bigg|^2 \\ &\leq \mathrm{E} \Big[\Big(\sum_{\beta \in \hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 (|G| - 1)^{-2|\beta|} \Big) \Big(\sum_{\beta \in \hat{H}} |\hat{B}_\beta|^{2(k-2)} \Big) \Big] \end{aligned}$$

$$\leq \mathrm{E}\Big[\sum_{\beta\in\hat{H}} |\hat{A}_{\pi_G(\beta)}|^2 |\hat{B}_\beta|^2 |\beta|^{-1}\Big],$$

where the last inequality follows since

$$\sum_{\beta\in\hat{H}} |\hat{B}_\beta|^{2(k-2)} \leq \sum_{\beta\in\hat{H}} |\hat{B}_\beta|^2 = 1$$

by Plancherel's equality. ∎

The strategy from the case $k = 3$ now works without modifications.

Theorem 11. *For every constant $\varepsilon > 0$, every integer $k \geq 3$ and every finite Abelian group of order at least three, it is **NP**-hard to distinguish instances of Max Ek-LinInEq(G) that are satisfiable from instances where at most a fraction $(1 - |G|^{-1} + \varepsilon)$ of the inequations are simultaneously satisfiable.*

As a corollary, using a reduction similar to that described in Section 2, we also obtain that Max Ek-Sat(d) is non-approximable beyond the random assignment threshold.

Corollary 3. *For all $k \geq 3$ and all $d \geq 2$, Max Ek-Sat(d) is hard to approximate within $(1-d^{-k}+\varepsilon)$ with perfect completeness for any constant $\varepsilon > 0$.*

Proof. The result for $d = 2$ follows from a difficult proof due to Håstad [10]. For $d \geq 3$, we reduce Max Ek-LinInEq(G) for some group of order d, which we know is hard to approximate within $1 - d^{-1}$ with perfect completeness by Theorem 11, to Max Ek-Sat(d). An equation $x_1 \cdots x_k \neq g$ is replaced with $d^{k-1}(d-1)$ Max Ek-Sat(d) clauses such that there is one clause for every non-satisfying assignment to the constraint $x_1 \cdots x_k \neq g$. If all Max Ek-Sat(d) clauses are satisfied, the Max Ek-LinInEq(G) clause is satisfied; if $x_1 \cdots x_k = g$ there is one unsatisfied Max Ek-Sat(d) clause. Therefore it is hard to distinguish the case when all the constraints are satisfied from the case when a fraction

$$\frac{d^{k-1}(d-1) - (1 - d^{-1} - d^{k-1}(d-1)\varepsilon)}{d^{k-1}(d-1)} = 1 - d^{-k} + \varepsilon$$

of the constraints are satisfied. ∎

Acknowledgments

We would like to thank Subhash Khot for sending us a copy of [12]; his algorithm for Max BIJ(d) directly inspired our algorithm for Max Co-BIJ(d).

References

1. Gunnar Andersson. An approximation algorithm for max p-section. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1563 of *Lecture Notes in Computer Science*, pages 237–247. Trier, 4–6 March 1999.
2. Gunnar Andersson. *Some New Randomized Approximation Algorithms.* Doctoral dissertation, Department of Numerical Analysis and Computer Science, Royal Institute of Technology, May 2000.
3. Gunnar Andersson, Lars Engebretsen, and Johan Håstad. A new way of using semidefinite programming with applications to linear equations mod p. *Journal of Algorithms*, 39(2):162–204, May 2001.
4. Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Márió Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, May 1998.

5. Pierluigi Crescenzi, Riccardo Silvestri, and Luca Trevisan. On weighted vs unweighted versions of combinatorial optimization problems. *Information and Computation*, 167(1):10-26, May 2001.
6. Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, July 1998.
7. Alan Frieze and Mark Jerrum. Improved approximation algorithms for MAX k-CUT and MAX BISECTION. *Algorithmica*, 18:67–81, 1997.
8. Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, November 1995.
9. Michel X. Goemans and David P. Williamson. Approximation algorithms for Max-3-Cut and other problems via complex semidefinite programming. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 443–452. Hersonissos, Crete, Grece, 6–8 July 2001.
10. Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, July 2001.
11. Viggo Kann, Sanjeev Khanna, Jens Lagergren, and Alessandro Panconesi. On the hardness of approximating Max k-Cut and its dual. *Chicago Journal of Theoretical Computer Science*, 1997(2), June 1997.
12. Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 767–775. Montréal, Québec, Canada, 19–21 May 2002.
13. Subhash Khot. Hardness results for coloring 3-colorable 3-uniform hypergraphs. To appear in *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science*. Vancouver, Canada, 16–19 November 2002.
14. Erez Petrank. The hardness of approximation: Gap location. *Computational Complexity*, 4(2):133–157, 1994.
15. Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.
16. Audrey Terras. *Fourier Analysis on Finite Groups and Applications*, volume 43 of *London Mathematical Society student texts*. Cambridge University Press, Cambridge, 1999.
17. Uri Zwick. Approximation algorithms for constraint satisfaction programs involving at most three variables per constraint. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 201–210. San Francisco, California, 25–27 January 1998.
18. Uri Zwick. Outward rotations: a tool for rounding solutions of semidefinite programming relaxations, with applications to MAX CUT and other problems. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 679–687. Atlanta, Georgia, 1–4 May 1999.

A Analysis of the algorithm for Max Co-BIJ(d)

Lemma 21. *Let r be a $2d$-dimensional normal variable with mean zero and identity covariance matrix. Let $B = \{r : |r| \leq 1/Kd\}$. Then,*

$$\int_B \left(1 - \frac{1}{d} + \frac{K^2 r_1^2}{d}\right) dP(r)$$
$$= 1 - \frac{1}{d} + \frac{K^2}{d} - e^{-1/2(Kd)^2}\left(\left(1 - \frac{1}{d}\right)\sum_{k=0}^{d-1} \frac{1}{k!2^k(Kd)^{2k}} + \frac{K^2}{d}\sum_{k=0}^{d} \frac{1}{k!2^k(Kd)^{2k}}\right).$$

Proof. We will show that

$$\int_B dP(r) = 1 - e^{-1/2(Kd)^2}\sum_{k=0}^{d-1} \frac{1}{k!2^k(Kd)^{2k}},$$

$$\int_B r_1^2\, dP(r) = 1 - e^{-1/2(Kd)^2} \sum_{k=0}^{d} \frac{1}{k! 2^k (Kd)^{2k}}\,.$$

Let us start by computing the first integral. The probability distribution for a $2d$-dimensional normally distributed random variable with zero mean and identity covariance matrix is $P(r) = (2\pi)^{-d} e^{-|r|^2/2}$, therefore

$$\int_B dP(r) = \frac{1}{(2\pi)^d} \int_0^{1/(Kd)} e^{-\rho^2/2} \int_{r:|r|=\rho} dS(r)\, d\rho.$$

The volume of a $2d$-dimensional hypersphere with radius ρ is $\pi^d \rho^{2d}/d!$ and the area of this sphere is $2d/\rho$ times the volume; thus $\int_{r:|r|=\rho} dS(r) = 2d\pi^d \rho^{2d-1}/d!$ and

$$\int_B dP(r) = \frac{1}{2^{d-1}(d-1)!} \int_0^{1/Kd} \rho^{2d-1} e^{-\rho^2/2}\, d\rho$$

With the substitution $\rho^2/2 = t$, $\rho\, d\rho = dt$,

$$\int_B dP(r) = \frac{1}{(d-1)!} \int_0^{1/(2K^2d^2)} t^{d-1} e^{-t}\, dt.$$

Using the identity

$$\int t^d e^{-t}\, dt = -e^{-t} \sum_{k=0}^{d} \frac{d!}{k!} t^k,$$

the latter integral can be easily evaluated:

$$\frac{1}{(d-1)!} \int_0^{1/Kd} t^{d-1} e^{-t}\, dt = 1 - e^{-1/2(Kd)^2} \sum_{k=0}^{d-1} \frac{1}{k! 2^k (Kd)^{2k}}\,.$$

The second integral can be computed with similar techniques. Since B is spherically symmetric,

$$\int_B r_1^2\, dP(r) = \frac{1}{2d} \int_B |r|^2\, dP(r).$$

Using the fact that the probability distribution for a $2d$-dimensional normally distributed random variable with zero mean and identity covariance matrix is $P(r) = (2\pi)^{-d} e^{-|r|^2/2}$, the integral of interest is

$$\begin{aligned}
\int_B |r|^2\, dP(r) &= \frac{1}{(2\pi)^d} \int_B |r|^2 e^{-|r|^2/2}\, dV(r) \\
&= \frac{1}{(2\pi)^d} \int_0^{1/(Kd)} \rho^2 e^{-\rho^2/2} \int_{r:|r|=\rho} dS(r)\, d\rho \\
&= \frac{1}{d!} \int_0^{1/2(Kd)^2} t^d e^{-t}\, dt \\
&= \left[-e^{-t} \sum_{k=0}^{d} \frac{t^k}{k!} \right]_0^{t=1/2(Kd)^2}
\end{aligned}$$

To conclude,

$$\int_B r_1^2\, dP(r) = 1 - e^{-1/2(Kd)^2} \sum_{k=0}^{d} \frac{1}{k! 2^k (Kd)^{2k}}\,. \qquad \blacksquare$$

Lemma 22. *Let*

$$S_n = \sum_{k=0}^{n} \frac{1}{k! 2^k (Kd)^{2k}}$$

With the parameter choice $K^2 = d^{-3}/13$,

$$\frac{K^2}{d} - e^{-1/2(Kd)^2}\left(\left(1-\frac{1}{d}\right)S_{n-1} + \frac{K^2}{d}S_n\right) \geq 0.07d^{-4}.$$

for all integers $d \geq 2$.

Proof. With the choice $K^2 = d^{-3}/2\alpha$, $e^{-1/2(Kd)^2} = e^{-\alpha d}$ and $S_n = \sum_{k=0}^{d} \frac{(\alpha d)^k}{k!}$. Let $a_k = (\alpha d)^k/k!$. Note that all a_k are positive and that $a_{k+1}/a_k = \alpha d/k \geq \alpha$ for all $k \leq d$. Therefore, provided that $\alpha \geq 2$ and $n \leq d$, $S_n \leq a_{n+1} \leq 2a_n$, which implies that

$$\left(1-\frac{1}{d}\right)S_{d-1} + \frac{K^2}{d}S_d \leq \frac{d-1}{d}\frac{(\alpha d)^d}{d!} + \frac{1}{\alpha d^4}\frac{(\alpha d)^d}{d!} \leq \frac{(\alpha d)^d}{d!} \leq \frac{(\alpha e)^d}{\sqrt{2\pi d}}$$

where the last inequality follows since $d! \geq d^d e^{-d}\sqrt{2\pi d}$ by Stirling's formula. Therefore

$$\frac{K^2}{d} - e^{-1/2(Kd)^2}\left(\left(1-\frac{1}{d}\right)S_{n-1} + S_n\right) \geq \frac{K^2}{d} - e^{-\alpha d}\frac{(\alpha e)^d}{\sqrt{2\pi d}} = \beta d^{-4}$$

where

$$\beta = \frac{1}{2\alpha} - (2\pi)^{-1/2}\exp((1-\alpha+\ln\alpha)d + 3.5\ln d).$$

All that remains is to select a good α, more precisely, an α such that the exponential function is decreasing in d. A simple calculation shows that $\alpha = 6.5$ suffices for all integers $d \geq 2$ and gives a lower bound of $\beta \geq 0.07$. ∎

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'