Electronic Colloquium on Computational Complexity, Revision 01 of Report No. 030 (2002)

ECCC

Inapproximability Results for Equations over Finite Groups

Lars Engebretsen [a,1] Jonas Holmerin [a] Alexander Russell [b,2]

[a] *Department of Numerical Analysis and Computer Science, Royal Institute of Technology, SE-100 44 Stockholm, SWEDEN*

[b] *Department of Computer Science and Engineering, University of Connecticut, Storrs, CT 06269, USA*

Abstract

An *equation* over a finite group G is an expression of form $w_1 w_2 \dots w_k = 1_G$, where each w_i is a variable, an inverted variable, or a constant from G; such an equation is *satisfiable* if there is a setting of the variables to values in G so that the equality is realized. We study the problem of simultaneously satisfying a family of equations over a finite group G and show that it is **NP**-hard to approximate the number of simultaneously satisfiable equations to within $|G| - \epsilon$ for any $\epsilon > 0$. This generalizes results of Håstad (2001, J. ACM, 48 (4)), who established similar bounds under the added condition that the group G is Abelian.

Key words: Optimization, Approximation, Groups, Finite groups, Probabilistically Checkable Proofs, NP-hardness

1 Introduction

Many fundamental computational problems can be naturally posed as questions concerning the simultaneous solvability of families of equations over finite groups. This connection has been exploited to achieve a variety of strong inapproximability results for problems such as Max Cut, Max Di-Cut, Exact

Email addresses: enge@kth.se (Lars Engebretsen), joho@kth.se (Jonas Holmerin), acr@cse.uconn.edu (Alexander Russell).

[1] Research partly performed at MIT with support from the Marcus Wallenberg Foundation and the Royal Swedish Academy of Sciences.

[2] This is research is partially supported by NSF CAREER award CCR-0093065 and NSF grants CCR-0220264 and EIA-0218443.

ISSN 1433-8092

Satisfiability, and Vertex Cover [8,10,11,15–17,21,27]. A chief technical ingredient in these hardness results is a tight lower bound on the approximability of the problem of simultaneously satisfying equations over a finite *Abelian* group; in this article we extend these results to cover all finite groups.

An *equation* in variables $x_1, \ldots, x_n$ over a group G is an expression of form $w_1 \ldots w_k = 1_G$, where each w_i is either a variable, an inverted variable, or a group constant and 1_G denotes the identity element. A *solution* is an assignment of the variables to values in G that realizes the equality. A collection of equations $\mathcal{E}$ over the same variables induces a natural optimization problem, the problem of determining the maximum number of simultaneously satisfiable equations in $\mathcal{E}$. We let EQ_G denote this optimization problem. The special case where a variable may only appear *once* in each equation is denoted EQ^1_G; when each equation has single occurrences of exactly k variables, the problem is denoted $\mathrm{EQ}^1_G[k]$. Our main theorem asserts that for any finite group G it is **NP**-hard to approximate $\mathrm{EQ}^1_G[3]$ (and hence EQ^1_G and EQ_G) to within $|G| - \epsilon$ for any $\epsilon > 0$; this is tight.

As mentioned above, EQ_G is tightly related to a variety of familiar optimization problems. When $G = \boldsymbol{Z}_2$, for example, instances of EQ_G where exactly two variables occur in each clause, i.e., $\mathrm{EQ}^1_{\boldsymbol{Z}_2}[2]$, correspond precisely to the familiar optimization problem Max Cut, the problem of determining the largest number of edges which cross some bipartition of an undirected graph. If, for example, $G = S_3$, the (non-Abelian) symmetric group on three letters, then the problem of maximizing the number of bichromatic edges in a three coloring of a given graph can be reduced to EQ_G [14]; other examples are described by Håstad [16] and Zwick [27]. The general problem has also been studied due to applications to the fine structure of $\mathbf{NC}^1$ [3,14] specializing the framework of Barrington *et al.* [4,5]. Finally, the problem naturally gives rise to a number of well-studied combinatorial enumeration problems: see, e.g., [6,13,25] and [24, pp. 110ff.].

If G is Abelian and $\mathcal{E}$ is a collection of equations over G, each of which can individually be satisfied, the trivial randomized approximation algorithm which independently assigns to each variable a uniformly selected value in G satisfies an expected fraction $|G|^{-1}$ of the equations. This algorithm can be efficiently derandomized by the method of conditional expectation [1, §15.1] and it in fact also applies to EQ^1_G for any finite group G. In 1997, Håstad [16] showed that if $\mathbf{P} \neq \mathbf{NP}$ and G is Abelian, then no polynomial time approximation algorithm can approximate $\mathrm{EQ}^1_G[3]$ to within $|G| - \epsilon$ for any $\epsilon > 0$. The main theorem of this paper shows that this same lower bound holds for all finite groups.

Theorem 1 *For any finite group G and any constant $\epsilon > 0$, it is **NP**-hard to approximate $\mathrm{EQ}^1_G[3]$ to within $|G| - \epsilon$.*

The paper is organized as follows: After an overview of our contribution in Sec. 2 we briefly describe the representation theory of finite groups and the generalization of the so called *long code* to non-Abelian groups in Sections 3 and 4. The main theorem then appears in Section 5.

2 Overview of our results

A burst of activity focusing on the power of various types of interactive proof systems in the 80s and early 90s culminated in the so called *PCP theorem*, described briefly below. A *probabilistically checkable proof system* (PCP) for a language L consists of a probabilistic polynomial time verifier which, given an input x and oracle access to a purported proof that $x \in L$, probabilistically verifies the validity of the proof. In this paper, we only consider PCPs where the number of random bits used by the verifier is *logarithmic in the input size* and the number of "positions" of the proof examined by the verifier is a *constant*. The verifier is also *nonadaptive* in the sense that the queries may not depend on the values of previously queried positions in the proof. A PCP is said to have *completeness* c and *soundness* s if a correct proof that $x \in L$ is accepted with probability at least c and, when $x \notin L$, no proof is accepted with probability exceeding s.

The PCP theorem [2] asserts the startling fact that any **NP**-language has a PCP with completeness 1 and soundness 1/2 where the verifier uses logarithmic randomness and examines a constant number of bits of the proof. To prove our inapproximability results for certain families of equations over finite groups we use the PCP theorem to construct, for any finite group G and any positive constants ϵ and δ, a PCP with completeness $1-\epsilon$ and soundness $|G|^{-1}+\delta$ where the verifier uses logarithmic randomness and examines *three positions* in the proof. Each "position" in our setting holds a value from the group G; this corresponds to reading $\lceil \log |G| \rceil$ adjacent bits if the proof is written in binary.

There is an approximation-preserving reduction from conjunctive normal form Boolean formulas containing exactly three literals per clause (E3-Sat) to E3-Sat formulas where each variable occurs in exactly five clauses [12,19]. Coupling the PCP theorem and this reduction shows that for every language L in **NP**, an arbitrary instance x can be transformed in polynomial time into an E3-Sat formula $\phi_{x,L}$ with the following property: if $x \in L$, then $\phi_{x,L}$ is satisfiable, and if $x \notin L$ then at most a fraction $\mu < 1$ of the clauses can be satisfied. Here μ is a universal constant, independent of the language and the instance.

In his seminal paper [16], Håstad introduced a methodology for proving lower

bounds for constraint satisfaction problems. At 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. It follows from Raz's parallel repetition theorem [20] that if the 2P1R protocol is repeated u times in parallel and applied to the formula $\phi_{x,L}$ above then the verifier always accepts an unsatisfiable formula with probability at most c_μ^u where $c_\mu < 1$ is independent of u.

To prove his inapproximability result for equations over finite Abelian groups, Håstad constructed a PCP where the verifier tests a given assignment of variables $x_1, \ldots, x_n$ to group values to determine if it satisfies an equation selected at random from a certain family of equations. As each such equation involves three variables, this can be tested with three oracle queries. He then, in essence, reduced the problem of finding a strategy for the 2P1R protocol for E3-Sat to the problem of finding an assignment which satisfies many of the group equations by showing that if the verifier in the PCP accepts with high probability, there is a strategy for the provers in the 2P1R protocol that makes the verifier of that protocol accept with high probability. The inapproximability result follows since it is known that the verifier in the latter protocol cannot accept an unsatisfiable instance of E3-Sat with high probability.

To adapt Håstad's method to equations over arbitrary finite groups we need to overcome a couple of technical difficulties. The first one regards the coding of the proof in Håstad's proof system. Our second two contributions regard the analysis of the probability that the verifier accepts an incorrect proof. These are surveyed briefly in §2.1 and §2.2, below.

2.1 The non-Abelian long code

To encode the proof, Håstad used a proof with several different tables, each coded with the so called *long code*. For any finite group G, the long G-code of a binary string x of length n consists of the values of all functions from n-bit strings to G evaluated on x. In his proofs, Håstad has to assume that the alleged proofs that the verifier examines have a certain structure. For instance, the positions corresponding to some function f and the function gf for any $g \in G$ must be consistent. This can be enforced by employing certain *access conventions* in the verifier, which we describe in detail later. Our first technical contribution in this paper is to formulate the Fourier transform of the long G-code for all finite groups G and to prove that certain access conventions, slightly different from those used by Håstad, imply that we can assume that the Fourier coefficients of the alleged proofs examined by our verifier have

certain desirable properties.

2.2 Analysis of the verifier

Our main technical contributions are from the part of the analysis where we establish the connection between the proof system that tests a group equation and the 2P1R protocol for E3-Sat. The first step in this analysis is to "arithmetize" the acceptance probability of the former protocol. For the case of an Abelian finite group G, this is straightforward: The acceptance probability can be written as a sum of $|G|$ terms. If the acceptance probability is large, there has to be at least one large term in the sum. Håstad then proceeds by expanding this allegedly large term in its Fourier expansion and then uses the Fourier coefficients to devise a strategy for the provers in the 2P1R game for E3-Sat. Specifically, the probability distribution induced by the Fourier coefficients is used to construct a probabilistic strategy for the provers in the 2P1R game. Roughly speaking, the acceptance probability of the verifier in the 2P1R game is large because some pair of related Fourier coefficients is large.

For non-Abelian groups, the way to arithmetize the test turns out to require a sum of the traces of products of certain matrices given by the representation theory of the group in question. As in Håstad's case, we find that if the acceptance probability of the linear test is large, there has to be one product of matrices with a large trace. Our next step is to expand this matrix product in its Fourier series. Unfortunately, the Fourier expansion of each entry in those matrices contains matrices that could be very large; consequently, the Fourier expansion of the entire trace contains a product of matrices with potentially huge dimension. Thus, the fact that this trace is large does not necessarily mean that the individual entries in the matrices are large and directly using the entries in the matrices to construct the probabilistic strategy for the provers in the 2P1R game does not appear to work. Instead, and this is our first main technical contribution, we prove that the terms in the Fourier expansion corresponding to matrices with large dimension cannot contribute much to the value of the trace. Having done that, we know that the terms corresponding to matrices with reasonably small dimension actually sum up to a significantly large value and we use those terms to construct a strategy for the provers in the 2P1R game; this is our second main technical contribution.

3 Representation theory and the Fourier transform

In this section, we give a short account of the representation theory needed to state and prove our results. For more details, we refer the reader to the

excellent accounts by Serre [23] and Terras [26].

The traditional Fourier transform, as appearing in, say, signal processing [9], algorithm design [22], or PCPs [16], focuses on decomposing functions $f: G \to \boldsymbol{C}$ defined over an Abelian group G. This "decomposition" proceeds by writing f as a linear combination of *characters* of the group G. Unfortunately, this same procedure cannot work over a non-Abelian group since in this case there are not enough characters to span the space of all functions from G into $\boldsymbol{C}$; the theory of group representations fills this gap, being the natural framework for Fourier analysis over non-Abelian groups and shall be the primary tool utilized in the analysis the "non-Abelian PCPs" introduced in Section 4.

Group representation theory studies realizations of groups as collections of matrices: specifically, a *representation* of a group G associates a matrix with each group element so that the group multiplication operation corresponds to normal matrix multiplication. Such an association gives an embedding of the group into $\mathrm{GL}(V)$, the group of invertible linear operators on a finite dimensional $\boldsymbol{C}$-vector space V. (Note that if V is one-dimensional, then this is exactly the familiar notion of character used in the Fourier analysis over Abelian groups.)

Definition 2 *Let G be a finite group. A* representation *of G is a homomorphism $\gamma: G \to \mathrm{GL}(V)$; the dimension of V is denoted by d_γ and called the* dimension of the representation.

Two representations are immediate: the *trivial representation* has dimension 1 and maps everything to 1. The permutation action of a group on itself gives rise to the *left regular representation*. Concretely, let V be a $|G|$-dimensional vector space with an orthogonal basis $B = \{e_g : g \in G\}$ indexed by elements of G. Then the *left regular representation* $\mathrm{reg}: G \to \mathrm{GL}(V)$ is given by $\mathrm{reg}(g): e_h \mapsto e_{gh}$; the matrix associated with $\mathrm{reg}(g)$ is simply the permutation matrix given by mapping each element h of G to gh.

If γ is a representation, then for each group element g, $\gamma(g)$ is a linear operator and, as mentioned above, can be identified with a matrix. We denote by $(\gamma(g)_{ij})$ the matrix corresponding to $\gamma(g)$. Two representations γ and θ of G are *isomorphic* if they have the same dimension and there is a change of basis U so that $U\gamma(g)U^{-1} = \theta(g)$ for all g. A representation non-isomorphic to the trivial representation is said to be *nontrivial*.

If $\gamma: G \to \mathrm{GL}(V)$ is a representation and $W \subseteq V$ is a subspace of V, we say that W is *invariant* if $\gamma(g)(W) \subseteq W$ for all g. If the only invariant subspaces are $\{0\}$ and V, we say that γ is *irreducible*. Otherwise, γ does have a nontrivial invariant subspace W_0 and notice that by restricting each $\gamma(g)$ to W_0 we obtain a new representation. When this happens, it turns out that there is always another invariant subspace W_1 so that $V = W_0 \oplus W_1$ and in this case we write

$\gamma = \gamma_0 \oplus \gamma_1$, where γ_0 and γ_1 are the representations obtained by restricting to W_0 and W_1. This is equivalent to the existence of a basis in which the $\gamma(g)$ are all block diagonal, where the matrix of $\gamma(g)$ consists of $\gamma_0(g)$ on the first block and $\gamma_1(g)$ on the second block. In this way, any representation can be decomposed into a sum of irreducible representations. The matrix entries of irreducible representations of a finite group G are "orthogonal" with respect to the pairing function

$$\langle f_1 \mid f_2 \rangle_G = \frac{1}{|G|} \sum_{g \in G} f_1(g) f_2(g^{-1}) \tag{1}$$

for functions from G to $\boldsymbol{C}$:

Proposition 3 *Let γ and θ be two non-isomorphic irreducible representations of G. Suppose that they are represented by the matrices (γ_{ij}) and $(\theta_{k\ell})$, respectively. Then $\langle \gamma_{ij} \mid \theta_{k\ell} \rangle_G = 0$ for all i, j, k, l and $d_\gamma \langle \gamma_{ij} \mid \gamma_{k\ell} \rangle_G = \delta_{i\ell}\delta_{jk}$.*

Corollary 4 *Let $\gamma\colon G \to \mathrm{GL}(V)$ be a nontrivial irreducible representation of G. Then $\sum_{g \in G} \gamma(g) = 0$.*

For a finite group G, there are only a finite number of irreducible representations up to isomorphism; we let $\hat{G}$ denote the set of distinct irreducible representations of G. It is not hard to show that any irreducible representation is isomorphic to a representation where each $\gamma(g)$ is unitary, and we will always work under this assumption.

We remark that if $\mathcal{R}(G)$ denotes the collection of *all* representations of a finite group G upto isomorphism, then the transformation $G \mapsto \mathcal{R}(G)$ is "functorial" in the sense that if $\phi\colon G \to H$ is a group homomorphism, then there is a natural map

$$\phi^*\colon \mathcal{R}(H) \to \mathcal{R}(G) \tag{2}$$

given by $\phi^*(\rho) = \rho \circ \phi$. Note that $\phi^*(\rho)$ need not be irreducible even if ρ is.

There is a natural product of representations, the *tensor product*. We define the tensor product $A \otimes B$ of two matrices $A = (a_{ij})$ and $B = (b_{k\ell})$ to be the matrix indexed by pairs $(i,k);(j,\ell)$ so that $(A \otimes B)_{(i,k);(j,\ell)} = a_{ij}b_{k\ell}$. We will use the so called *inner trace* of a tensor product: For a matrix M indexed by pairs $(i,k);(j,\ell)$ the inner trace, denoted by $\operatorname{Tr} M$, is defined by $(\operatorname{Tr} M)_{ij} = \sum_k M_{(i,k),(j,k)}$. We let tr denote the normal trace. The inner trace is the "opposite" of the tensor product in the sense that $\operatorname{Tr}(A \otimes B) = (\operatorname{tr} B)A$. If $\gamma\colon G \to \mathrm{GL}(V)$ and $\theta\colon H \to \mathrm{GL}(W)$ are representations of G and H, respectively, we define $\gamma \otimes \theta\colon G \times H \to \mathrm{GL}(V \otimes W)$ to be the representation of $G \times H$ given by $(\gamma \otimes \theta)(g,h) = \gamma(g) \otimes \theta(h)$.

Proposition 5 *Let G and H be finite groups. Then the irreducible representations of $G \times H$ are precisely $\{\gamma \otimes \theta : \gamma \in \hat{G}, \theta \in \hat{H}\}$. Furthermore, each of*

these representations is distinct up to isomorphism.

For a representation γ, the function $g \mapsto \operatorname{tr} \gamma(g)$ is called the *character corresponding to* γ and is denoted by χ_γ. Note that χ_γ takes values in $\boldsymbol{C}$ even if γ has dimension larger than one. Our principal use of the character relies on the following fact:

Proposition 6 *Let G be a finite group. Then*

$$\sum_{\gamma \in \hat{G}} d_\gamma \chi_\gamma(g) = \begin{cases} |G| & \text{if } g = 1_G, \\ 0 & \text{otherwise.} \end{cases}$$

As $\gamma(1_G)$ is always an identity matrix, $\chi_\gamma(1_G) = d_\gamma$ and we conclude the following:

Corollary 7 $\sum_{\gamma \in \hat{G}} d_\gamma^2 = |G|$.

Note that the pairing function defined in (1) is not an inner product on the space of functions from G to $\boldsymbol{C}$, as it is not semilinear in f_2. It does, however, coincide with the usual inner product

$$\langle f_1 \mid f_2 \rangle_2 = \frac{1}{|G|} \sum_g f_1(g) f_2(g)^*$$

on the class of functions for which $f(g^{-1}) = f(g)^*$, which will play a distinguished role in the proofs below. In fact, the functions corresponding to the entries in the irreducible representations of G are indeed orthogonal with respect to the inner product $\langle \cdot \mid \cdot \rangle_2$. Since there are $|G|$ such functions by Corollary 7, there are sufficiently many to span the space of functions from G to $\boldsymbol{C}$.

3.1 The Fourier transform

We now proceed to describe the Fourier transform of functions from an arbitrary finite group G to $\boldsymbol{C}$. Let f be a function from G to $\boldsymbol{C}$ and γ be an irreducible representation of G. Then

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

is the *Fourier coefficient of f at γ*. Moreover, f can be written as a Fourier series

$$f(g) = \sum_{\gamma \in \hat{G}} d_\gamma \operatorname{tr}\left(\hat{f}_\gamma \gamma(g^{-1}) \right) \tag{4}$$

In our analysis, we need the following version of Plancherel's equality:

Lemma 8 *Suppose that f is a function from G to $\boldsymbol{C}$. Then*

$$\sum_{\gamma\in\hat{G}}\sum_{1\le i\le d_\gamma}\sum_{1\le j\le d_\gamma} d_\gamma\Big|\langle f \mid \gamma_{ij}\rangle_G\Big|^2 = \frac{1}{|G|}\sum_{g\in G}|f(g)|^2 \tag{5}$$

if the representations $\gamma\in\hat{G}$ are represented in unitary bases.

PROOF. We expand the expression above using the definition of $\langle\cdot\mid\cdot\rangle_G$:

$$\Big|\langle f \mid \gamma_{ij}\rangle_G\Big|^2 = \frac{1}{|G|^2}\sum_{g\in G}\sum_{h\in G} f(g)\gamma_{ij}(g^{-1})f^*(h)\Big(\gamma_{ij}(h^{-1})\Big)^*$$

Since γ is a unitary representation, $\gamma(h^{-1}) = \gamma^{-1}(h) = \gamma^*(h)$, and hence $\gamma_{ij}(h^{-1}) = (\gamma_{ji}(h))^*$. Therefore,

$$\begin{aligned}
&\sum_{\gamma\in\hat{G}}\sum_{1\le i\le d_\gamma}\sum_{1\le j\le d_\gamma} d_\gamma\Big|\langle f \mid \gamma_{ij}\rangle_G\Big|^2\\
&\qquad= \frac{1}{|G|^2}\sum_{\gamma\in\hat{G}} d_\gamma \sum_{1\le i\le d_\gamma}\sum_{1\le j\le d_\gamma}\sum_{g\in G}\sum_{h\in G} f(g)f^*(h)\gamma_{ij}(g^{-1})\gamma_{ji}(h)\\
&\qquad\qquad= \frac{1}{|G|^2}\sum_{g\in G}\sum_{h\in G} f(g)f^*(h)\sum_{\gamma\in\hat{G}} d_\gamma\,\mathrm{tr}\Big(\gamma(g^{-1})\gamma(h)\Big)\\
&\qquad\qquad= \frac{1}{|G|^2}\sum_{g\in G}\sum_{h\in G} f(g)f^*(h)\sum_{\gamma\in\hat{G}} d_\gamma\,\mathrm{tr}\Big(\gamma(g^{-1}h)\Big) = \frac{1}{|G|}\sum_{g\in G}|f(g)|^2,
\end{aligned}$$

where the last equality follows from Proposition 6. (See also Serre's account [23, §6.2, Exercise 6.2], which discusses this in different language.)

3.2 The Fourier transform of matrix-valued functions

We also need to use the Fourier transform on functions $f\colon G\to \mathrm{End}(V)$, where $\mathrm{End}(V)$ is the set of linear maps from the vector space V to itself; here we identify $\mathrm{End}(V)$ with the space of all $\dim V\times\dim V$ matrices over $\boldsymbol{C}$. Although we have not found any treatment of such transforms in the literature, it is straightforward to generalize the concepts from the previous section to matrix-valued functions. For a representation γ of G, we define

$$\hat{f}_\gamma = \frac{1}{|G|}\sum_{g\in G} f(g)\otimes\gamma(g). \tag{6}$$

Treating the $f(g)$ as matrices, this is nothing more than the component-wise Fourier transform of the function f. The reason for grouping them together

into these tensor products is the following: Let $f, h\colon G \to \mathrm{End}(V)$ be two such functions, and define their convolution as

$$(f * h)(g) = \frac{1}{|G|} \sum_{t \in G} f(t) h(t^{-1} g),$$

this product being the ring product in $\mathrm{End}(V)$ (that is, function composition). Then it turns out that $\widehat{(f * h)}_\gamma = \hat{f}_\gamma \hat{h}_\gamma$, this product being matrix multiplication:

$$\begin{aligned}
\widehat{(f * h)}_\gamma &= \frac{1}{|G|} \sum_{g \in G} (f * h)(g) \otimes \gamma(g) \\
&= \frac{1}{|G|^2} \sum_{g \in G} \sum_{t \in G} \Big(f(t) h(t^{-1} g) \Big) \otimes \gamma(t t^{-1} g) \\
&= \frac{1}{|G|^2} \sum_{g \in G} \sum_{t \in G} \Big(f(t) \otimes \gamma(t) \Big) \Big(h(t^{-1} g) \otimes \gamma(t^{-1} g) \Big) \\
&= \hat{f}_\gamma \hat{h}_\gamma.
\end{aligned}$$

In this case, the Fourier series is

$$f(g) = \sum_{\gamma \in \hat{G}} d_\gamma \operatorname{Tr}\Big(\hat{f}_\gamma \big(\boldsymbol{I} \otimes \gamma(g^{-1}) \big) \Big) \tag{7}$$

where $\operatorname{Tr} M$ is the inner trace. This also gives rise to a Plancherel equality: for two functions $f, h\colon G \to \mathrm{End}(V)$,

$$\frac{1}{|G|} \sum_{g \in G} f(g) h(g^{-1}) = (f * h)(1_G) = \sum_{\gamma \in \hat{G}} d_\gamma \operatorname{Tr}(\hat{f}_\gamma \hat{h}_\gamma). \tag{8}$$

As we noted above, the representations of a finite group can always be expressed in a unitary basis. When a function from a finite group G to $\mathrm{End}(V)$ behaves like a unitary matrix in a certain sense and the representations of G are expressed in a unitary basis, the Fourier coefficients are Hermitian. This turns out to be important in our analysis.

Definition 9 *Let G be a finite group, V be a finite-dimensional vector space, and f be a function from G to* $\mathrm{End}(V)$. *Then f is* skew-symmetric *if $f(g^{-1}) = f^*(g)$ for all $g \in G$.*

Lemma 10 *Let G be a finite group, V be a finite-dimensional vector space, and f be a skew-symmetric function from G to* $\mathrm{End}(V)$. *Then $\hat{f}_\gamma$ is Hermitian if γ is expressed in a unitary basis.*

PROOF. Recall that a matrix M is Hermitian if $M = M^*$. By equation (6),

$$\hat{f}_\gamma = \frac{1}{|G|}\sum_{g\in G} f(g)\otimes\gamma(g) = \frac{1}{2|G|}\sum_{g\in G}\Big(f(g)\otimes\gamma(g) + f(g^{-1})\otimes\gamma(g^{-1})\Big)$$
$$= \frac{1}{2|G|}\sum_{g\in G}\Big(f(g)\otimes\gamma(g) + f^*(g)\otimes\gamma^*(g)\Big),$$

where the last equality follows since f is skew-symmetric and γ is expressed in a unitary basis. Now, $f(g)\otimes\gamma(g) + f^*(g)\otimes\gamma^*(g)$ is clearly Hermitian, and since a sum of Hermitian matrices is Hermitian, $\hat{f}_\gamma$ is Hermitian.

4 The non-Abelian long code and its Fourier transform

The long code was introduced by Bellare, Goldreich and Sudan [7] and adapted by Håstad [16] to prove approximability bounds for linear equations over Abelian groups. In this section, we once more generalize the long code for use in our proof system, that must work for all finite groups.

Let K be a finite set and denote by G^K the set of all functions from K to G. The *long G-code* of some $x \in K$ is the function A_x from G^K to G such that $A_x(f) = f(x)$. The proof in our PCP consists of several separate tables, each of which is a purported long code. In the analysis of the soundness of the verifier, we study the Fourier transform of such purported long codes composed with a representation of G, i.e., the Fourier transform of functions from G^K to $\mathrm{End}(V)$, where V is the underlying vector space of a representation ρ.

4.1 Folding

We first note that the concept of *folding* that has been used extensively for ordinary long codes extends to the long G-code.

Definition 11 *Let G be a finite group, $\gamma \in \hat{G}$ be arbitrary, V be the space corresponding to γ, K be a finite set, and A be a function from G^K to* $\mathrm{End}(V)$. *Then A is* γ-homogeneous *if $A(gf) = \gamma(g)A(f)$ for all $g \in G$ and all $f \in G^K$.*

In the above definition, gf is interpreted in the natural way: it is the function defined by $x \mapsto gf(x)$.

Lemma 12 *Let G be a finite group, γ be an arbitrary nontrivial representation of G, V be the space corresponding to γ, K be a finite set, and A be a γ-homogeneous function from G^K to* $\mathrm{End}(V)$. *Then $\hat{A}_\rho = 0$ when ρ is the trivial representation of G^K.*

PROOF. Since $\rho(f) = 1$ for all f when ρ is the trivial representation, (6) immediately yields

$$\hat{A}_\rho = \frac{1}{|G|^{|K|}} \sum_{f \in G^K} A(f) = \frac{1}{|G|^{|K|+1}} \sum_{g \in G} \sum_{f \in G^K} A(gf)$$
$$= \frac{1}{|G|^{|K|+1}} \sum_{f \in G^K} \left(\sum_{g \in G} \gamma(g) \right) A(f) = 0,$$

where the last equality follows from Corollary 4.

By employing a certain access convention in the verifier, we can ensure that tables correspond to γ-homogeneous functions.

Definition 13 *Let G be a finite group, K be a finite set, and A be a function from G^K to G. Partition G^K into equivalence classes by the relation $\equiv$, where $f \equiv h$ if there is $g \in G$ such that $f = gh$. Write $[f]$ for the equivalence class of f. Define A_G, A* left-folded over *G by choosing a representative for each equivalence class and defining $A_G(h) = gA(f)$, if $h = gf$ and f is the chosen representative for $[h]$.*

Lemma 14 *Let G be a finite group, K be a finite set, and A be a function from G^K to G. Then $\gamma \circ A_G$ is γ-homogeneous for every $\gamma \in \hat{G}$.*

PROOF. Note that $A_G(gf) = gA_G(f)$ for all $g \in G$ and all $f \in G^K$. Hence $(\gamma \circ A_G)(gf) = \gamma(g)(\gamma \circ A_G)(f)$.

It turns out that our analysis only requires some of the tables in the proof to be folded, while the other tables must correspond to skew-symmetric functions as per Definition 9. Again, this can be accomplished by proper access conventions in the verifier. In this case, we achieve our goal by, when accessing $A(f)$, with probability $1/2$ using the value $A(f^{-1})^{-1}$ instead.

Lemma 15 *Let G be a finite group, K be a finite set, A be a function from G^K to G, $\gamma \in \hat{G}$ be unitary, and $B(f) = \mathrm{E}_{b \in \{-1,1\}}[((\gamma \circ A)(f^b))^b]$. Then B is skew-symmetric.*

PROOF. By the above definition of B,

$$\begin{aligned} B(f) &= \mathrm{E}_{b\in\{-1,1\}}\Big[\Big((\gamma\circ A)(f^b)\Big)^b\Big] \\ &= \frac{1}{2}(\gamma\circ A)(f) + \frac{1}{2}\Big((\gamma\circ A)(f^{-1})\Big)^{-1} \\ &= \frac{1}{2}(\gamma\circ A)(f) + \frac{1}{2}\Big((\gamma\circ A)(f^{-1})\Big)^{*}, \end{aligned}$$

where the last equality holds since γ is unitary. Hence $B(f^{-1}) = \frac{1}{2}(\gamma\circ A)(f^{-1}) + \frac{1}{2}((\gamma\circ A)(f))^* = (B(f))^*$, and B is skew-symmetric by Definition 9.

We remark, that if A is a long G-code, A_G is identical to A, and $A(f^{-1})^{-1} = A(f)$.

4.2 Projection

To state the final long G-code property that we need, we have to develop a more precise and detailed description of the Fourier transform. Since we can represent a function $f\colon K \to G$ by a table containing $f(x)$ for every $x \in K$, we can identify G^K with $G^{|K|}$. In order to reason about the Fourier transform of a function from G^K to $\mathrm{End}(V)$, we need an understanding of the irreducible representations of powers of G. It follows from Proposition 5 that the irreducible representations of G^K are precisely those representations obtained by taking tensor products of $|K|$ irreducible representations of G: when $\rho_x \in \hat{G}$ for each $x \in K$ this is the representation given by

$$\rho = \bigotimes_{x\in K} \rho_x \qquad \text{where} \qquad \rho(f) = \bigotimes_{x\in K} \rho_x(f(x)).$$

We treat the tensor product of two matrices as a matrix indexed by pairs. Analogously, we treat the tensor product of $|K|$ matrices as a matrix indexed by $|K|$-tuples. In order to reason about single entries in the tensor product that forms a representation $\rho = \bigotimes_{x\in K} \rho_x$ we define the *set of indices* $\iota(\rho)$. An element $i \in \iota(\rho)$ is a vector indexed by elements of K so that for all $x \in K, 1 \le i_x \le d_{\rho_x}$; we refer to such an element i as an *index*. Then for two indices $i, j \in \iota(\rho)$ we define

$$\rho_{ij}(f) = \prod_{x\in K} \Big(\rho_x\Big(f(x)\Big)\Big)_{i_x,j_x}.$$

We also define the *weight* $|\rho|$ of an irreducible representation ρ of G^K to be the number of $x \in K$ such that ρ_x is nontrivial.

The verifier in our PCP checks positions in tables corresponding to two related long codes. The precise details of how these tables are related is described

below; for now it is enough to know that the tables correspond to functions from $F = G^K$ to G and from $H = G^L$ to G, respectively, where there is an onto function $\pi\colon L \to K$. Such a function $\pi\colon L \to K$ gives rise to the "dual" function $\pi^*\colon F \to H$ given by $\pi^*(f) = f \circ \pi$; the π^* defined in this way is in fact a homomorphism. Applying now the functorial property noted in Section 3 (equation 2) to the homomorphism π^*, a representation $\rho \in \hat{H}$ may be transformed into the representation ρ^π of F given by $\rho^\pi(f) = \rho(f \circ \pi)$. In particular, this transforms the components of the representation $\rho \in \hat{H}$ into functions on F. Recall that for $i, j \in \iota(\rho)$, the components ρ_{ij} are functions from H to $\boldsymbol{C}$. We denote the new, associated, functions by $\rho^\pi_{ij}\colon F \to \boldsymbol{C}$; they are given by the rule $f \mapsto \rho_{ij}(f \circ \pi)$. Using our definition of the index sets,

$$\rho^\pi_{ij}(f) = \rho_{ij}(f \circ \pi) = \prod_{x \in K} \prod_{y \in \pi^{-1}(x)} \Big(\rho_y\big(f(x)\big)\Big)_{i_y, j_y}.$$

We are now ready to formulate the following projection lemma:

Lemma 16 *Let K and L be finite sets and $\pi\colon L \to K$ be an onto function. Let $F = G^K$ and $H = G^L$. Define the relation $\sim$ on $\hat{F} \times \hat{H}$ so that for $\tau \in \hat{F}$ and $\rho \in \hat{H}$, $\tau \sim \rho$ if for all $x \in K$ such that τ_x is nontrivial, there is some $y \in \pi^{-1}(x)$ such that ρ_y is nontrivial. Then*

(1) $\tau \sim \rho \implies |\tau| \le |\rho|$.
(2) $\tau \not\sim \rho \implies \forall i, j \in \iota(\rho), \forall k, \ell \in \iota(\tau), \Big(\langle \rho^\pi_{ij} \mid \tau_{k\ell} \rangle_F = 0\Big)$.

PROOF. The first implication follows directly from the definition of the relation $\sim$. To prove the second implication, assume that $\tau \not\sim \rho$; then there is some $x' \in K$ such that $\tau_{x'}$ is nontrivial but ρ_y is trivial for all $y \in \pi^{-1}(x')$. Recall that we can write

$$\rho^\pi_{ij}(f) = \prod_{x \in K} \prod_{y \in \pi^{-1}(x)} \Big(\rho_y\big(f(x)\big)\Big)_{i_y, j_y} \qquad \text{and} \qquad \tau_{k\ell}(f) = \prod_{x \in K} \Big(\tau_x\big(f(x)\big)\Big)_{k_x, \ell_x}$$

by our definition of the index sets; hence

$$\begin{aligned}
\langle \rho^\pi_{ij} \mid \tau_{k\ell} \rangle_F &= \frac{1}{|F|} \sum_{f \in F} \Big(\prod_{x \in K} \prod_{y \in \pi^{-1}(x)} \Big(\rho_y\big(f(x)\big)\Big)_{i_y, j_y} \Big) \Big(\prod_{x \in K} \Big(\tau_x\big(f^{-1}(x)\big)\Big)_{k_x, \ell_x} \Big) \\
&= \frac{1}{|F|} \sum_{f \in F} \prod_{x \in K} \Big(\tau_x\big(f^{-1}(x)\big)\Big)_{k_x, \ell_x} \prod_{y \in \pi^{-1}(x)} \Big(\rho_y\big(f(x)\big)\Big)_{i_y, j_y} \\
&= \prod_{x \in K} \Big(\frac{1}{|G|} \sum_{g \in G} \big(\tau_x(g^{-1})\big)_{k_x, \ell_x} \prod_{y \in \pi^{-1}(x)} \big(\rho_y(g)\big)_{i_y, j_y} \Big)
\end{aligned}$$

where the last equality holds since a sum over all functions $f \in F$ can be viewed as $|K|$ nested sums over the possible values of $f(x)$ for $x \in K$. We can

then change the sums of a product into a product of $|K|$ sums, i.e., a product of sums over all $g \in G$.

Since ρ_y is trivial for all $y \in \pi^{-1}(x')$, the factor corresponding to x' in the above product is

$$\frac{1}{|G|} \sum_{g \in G} \Big(\tau_{x'}(g^{-1})\Big)_{k_{x'},\ell_{x'}} = 0,$$

where the equality follows from Corollary 4.

5 The main result

In his paper [16], Håstad introduced a methodology for proving lower bounds for constraint satisfaction problems. At 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 clause to one prover and a variable contained in that clause to the other prover, accepting if the returned assignments are consistent and satisfy the clause.

5.1 The two-prover one-round protocol

The starting point for our PCPs will be the standard 2P1R protocol for **NP** which we will now describe. We begin by discussing the decision problem *μ-gap E3-Sat(5)*.

Definition 17 *A Boolean formula ϕ in conjunctive normal form is* μ-promised *if either ϕ is satisfiable or no more than a μ-fraction of the clauses of ϕ are simultaneously satisfiable.* μ-gap E3-Sat(5) *is the problem of determining satisfiability of a μ-promised Boolean formula, where each clause contains exactly three literals and each literal occurs exactly five times.*

Recall that it is possible to reduce any problem in **NP** to an instance of μ-gap E3-Sat(5) [2,12,19]. This gives rise to a natural 2P1R protocol consisting 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 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 [20] 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.

5.2 The protocol

The proof in our PCP contains a purported encoding of a pair of strategies for the provers in the above u-parallel game. For a multiset U of variables, we denote by $\{-1,1\}^U$ the set of all assignments to the variables in U. For a multiset W of clauses, we denote by SAT^W the set of all satisfying assignments to the clauses in W. A satisfying assignment to the clauses in W can be viewed as a string of length u consisting of the numbers 1 to 7. Each number represents one of the satisfying assignments to an E3-SAT clause according to some arbitrary, but fixed, convention. Of course, it may happen that U contains the same variable more than once or that W contains clauses with common variables. For technical reasons, we do not require the assignments in $\{-1,1\}^U$ and SAT^W to be internally consistent. When x is an assignment to all the variables in an instance and V is a multiset of variables or a multiset of clauses, we denote by $x|_V$ the assignment to the variables in V. If V is a multiset of clauses $x|_V$ is therefore an assignment to the variables that constitute the clauses in V.

Definition 18 *A* Standard Written G-Proof with parameter u *for a formula ϕ consists of a table $A_U \colon G^{\{-1,1\}^U} \to G$ for each multiset U of u variables from ϕ and a table $A_W \colon G^{\mathrm{SAT}^W} \to G$ for each multiset W of u clauses from ϕ.*

Definition 19 *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 multiset V containing either u variables from ϕ or u clauses from ϕ.*

The protocol itself is similar to that used by Håstad [16] to prove inapproximability of equations over Abelian groups; the only difference is in the coding of the proof. The tables corresponding to sets of variables are left-folded over G and the tables corresponding to sets of clauses are folded over inverse. The

Input: A μ-gap E3-Sat(5) formula ϕ and oracle access to a Standard Written G-Proof with parameter u.

(1) Select uniformly at random a multiset $W = \{C_{i_1}, \ldots, C_{i_u}\}$ of u clauses.
(2) Construct a multiset U by choosing a variable uniformly at random from each C_{i_k}.
(3) Let π be the function that creates an assignment in $\{-1,1\}^U$ from an assignment in SAT^W.
(4) Select uniformly at random $f\colon \{-1,1\}^U \to G$.
(5) Select uniformly at random $h\colon \mathrm{SAT}^W \to G$.
(6) Choose $e\colon \mathrm{SAT}^W \to G$, such that, independently for each $y \in \mathrm{SAT}^W$,
    (a) With probability $1-\epsilon$, $e(y) = 1_G$.
    (b) With probability ϵ, $e(y)$ is chosen uniformly at random from G.
(7) Choose b_1 and b_2 independently and uniformly at random from $\{-1,1\}$.
(8) If $A_{U,G}(f)(A_W(h^{b_1}))^{b_1}(A_W((h^{-1}(f\circ\pi)^{-1}e)^{b_2}))^{b_2} = 1_G$ then accept, else reject.

Fig. 1. The above PCP is parameterized by the positive integer u and the positive real ϵ and tests if a μ-gap E3-Sat(5) formula ϕ is satisfiable by querying three positions in a Standard Written G-Proof with parameter u.

verifier is given in Figure 1. It is straightforward to bound the number of random bits used by the verifier and the completeness of the PCP:

Lemma 20 *The verifier needs at most* $u\log(5n)+2^u\log|G|+7^u\log(|G|^2/\epsilon)+2$ *random bits.*

PROOF. Since every variable occurs exactly five times in the μ-gap E3-Sat(5) formula ϕ, at most $u\log(5n/3)$ random bits are needed to sample the set W. Once W has been selected, $u\log 3$ bits suffice to select U. It is enough to use $(2^u+7^u)\log|G|$ random bits to select the functions f and h. To sample the error function e, we need to use $\log(|G|/\epsilon)$ random bits for every possible assignment to the variables it depends on. Thus, $7^u\log(|G|/\epsilon)$ random bits suffice to sample the entire error function. Finally, the sampling of b_1 and b_2 requires 2 bits.

Lemma 21 *The verifier in Figure 1 has completeness at least* $1-(1-|G|^{-1})\epsilon$.

PROOF. Let x be the assignment corresponding to a correct Standard Written G-Proof with parameter u for a formula ϕ. Then, by the definition of the

long G-code, $A_{U,G}(f) = f(x|_U)$ for all f and $(A_W(h^b))^b = h(x|_W)$ for all h; hence

$$A_{U,G}(f)(A_W(h^{b_1}))^{b_1}(A_W((h^{-1}(f\circ\pi)^{-1}e)^{b_2}))^{b_2} \\ = f(x|_U)h(x|_W)h^{-1}(x|_W)f^{-1}(x|_U)e(x|_W) = e(x|_W).$$

Considering how e is selected by the verifier, $e(x|_W) = 1_G$ with probability $1-(1-|G|^{-1})\epsilon$ and hence the verifier accepts a correctly encoded proof of a satisfying assignment with probability $1-(1-|G|^{-1})\epsilon$.

5.3 Analysis of the soundness

The analysis follows the now standard approach. We assume that the verifier accepts a proof corresponding to an unsatisfiable formula with probability $|G|^{-1}+\delta$ and prove that it is then possible to construct strategies for the provers in the 2P1R game that make the verifier of that game accept with high probability. Since it is known that this cannot be the case, there cannot exist a proof corresponding to an unsatisfiable formula that the PCP verifier accepts with probability $|G|^{-1}+\delta$.

To this end, we first apply Proposition 6 to arrive at an expression for the acceptance probability. Since

$$|G|^{-1}\sum_{\gamma\in\hat{G}} d_\gamma\chi_\gamma\Big(A_{U,G}(f)\Big(A_W(h^{b_1})\Big)^{b_1}\Big(A_W((h^{-1}(f\circ\pi)^{-1}e)^{b_2})\Big)^{b_2}\Big)$$

is an indicator of the event that the verifier accepts, the acceptance probability can be written as:

$$|G|^{-1}\sum_{\gamma\in\hat{G}} d_\gamma \,\mathrm{E}\Big[\chi_\gamma\Big(A_{U,G}(f)\Big(A_W(h^{b_1})\Big)^{b_1}\Big(A_W((h^{-1}(f\circ\pi)^{-1}e)^{b_2})\Big)^{b_2}\Big)\Big] = \\ |G|^{-1}+|G|^{-1}\sum_{\gamma\in\hat{G}\setminus\{\mathbf{1}\}} d_\gamma \,\mathrm{E}\Big[\chi_\gamma\Big(A_{U,G}(f)\Big(A_W(h^{b_1})\Big)^{b_1}\Big(A_W((h^{-1}(f\circ\pi)^{-1}e)^{b_2})\Big)^{b_2}\Big)\Big]$$

where the expectations are over the choice of U, W, f, h, e, b_1, and b_2. With the aid of Corollary 7, we deduce that if the verifier in Figure 1 accepts with probability $|G|^{-1}+\delta$, there must be some nontrivial irreducible representation γ of G such that

$$\Big|\mathrm{E}\Big[\chi_\gamma\Big(A_{U,G}(f)\Big(A_W(h^{b_1})\Big)^{b_1}\Big(A_W(h^{-1}(f\circ\pi)^{-1}e)^{b_2}\Big)^{b_2}\Big)\Big]\Big| > d_\gamma\delta. \qquad (9)$$

We now proceed by applying Fourier-inversion to $\gamma\circ A_{U,G}$ and $\gamma\circ A_W$. More precisely, we first apply Fourier-inversion to $\gamma\circ A_W$, resulting in:

Lemma 22 *Suppose that the verifier in Figure 1 accepts with probability* $|G|^{-1}+\delta$. *Then there exists a nontrivial representation* γ *of* G *such that*

$$\left|\mathrm{E}_{f,U,W}\left[\mathrm{tr}\left(A(f)\sum_{\rho\in\hat{H}} d_\rho(1-\epsilon)^{|\rho|}\,\mathrm{Tr}\Big(\hat{B}_\rho^2\Big(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\Big)\Big)\right)\right]\right| > d_\gamma\delta.$$

where $A=\gamma\circ A_{U,G}$, $H=G^{\mathrm{SAT}^W}$ *and* $B(h)=\mathrm{E}_{b\in\{-1,1\}}[((\gamma\circ A_W)(h^b))^b]$.

PROOF. Since the verifier in Figure 1 accepts with probability $|G|^{-1}+\delta$ there exists a nontrivial representation γ of G such that the inequality (9) holds; we now fix that γ and select a basis such that it is unitary. We proceed by expanding the expectation in (9) in a Fourier series. Since γ is a homomorphism, the expectation in (9) is equal to

$$\mathrm{tr}\,\mathrm{E}_{f,h,e,U,W}\Bigg[(\gamma\circ A_{U,G})(f)\,\mathrm{E}_{b_1}\Big[\Big((\gamma\circ A_W)(h^{b_1})\Big)^{b_1}\Big]\\ \mathrm{E}_{b_2}\Big[\Big((\gamma\circ A_W)((h^{-1}(f\circ\pi)^{-1}e)^{b_2})\Big)^{b_2}\Big]\Bigg].$$

To shorten the notation, we introduce the shorthands $A=\gamma\circ A_{U,G}$ and $B(h)=\mathrm{E}_{b\in\{-1,1\}}[((\gamma\circ A_W)(h^b))^b]$. With these shorthands the above expectation is equal to

$$\mathrm{tr}\,\mathrm{E}_{f,h,e,U,W}\Big[A(f)B(h)B(h^{-1}(f\circ\pi)^{-1}e)\Big]\\ =\mathrm{tr}\,\mathrm{E}_{f,U,W}\Big[A(f)\,\mathrm{E}_{h,e}\Big[B(h)B(h^{-1}(f\circ\pi)^{-1}e)\Big]\Big] \quad (10)$$

We now expand the inner expectation in its Fourier series. Since $\mathrm{E}_{h,e}[B(h)B(h^{-1}(f\circ\pi)^{-1}e)]=\mathrm{E}_e[(B*B)((f\circ\pi)^{-1}e)]$, this is immediate:

$$\mathrm{E}_e[(B*B)((f\circ\pi)^{-1}e)]=\sum_{\rho\in\hat{H}} d_\rho\,\mathrm{Tr}\Big(\hat{B}_\rho^2\Big(\boldsymbol{I}_{d_\gamma}\otimes\Big(\rho(f\circ\pi)\,\mathrm{E}_e[\rho(e^{-1})]\Big)\Big)\Big).$$

To compute $\mathrm{E}_e[\rho(e^{-1})]$, note that

$$\mathrm{E}_e[\rho(e^{-1})]=\mathrm{E}_e\Bigg[\bigotimes_{y\in\mathrm{SAT}^W}\Big(\rho_y\Big(e(y)^{-1}\Big)\Big)\Bigg]=\bigotimes_{y\in\mathrm{SAT}^W}\mathrm{E}_{e(y)}\Big[\rho_y\Big(e(y)^{-1}\Big)\Big],$$

where the second equality follows since $e(y)$ is selected independently for every y. Now, $\mathrm{E}_{e(y)}[\rho_y(e(y)^{-1})]=\boldsymbol{I}_{d_{\rho_y}}$ if ρ_y is trivial; otherwise

$$\mathrm{E}_{e(y)}\Big[\rho_y\Big(e(y)^{-1}\Big)\Big]=(1-\epsilon)\rho_y(1_G)+\epsilon\,\mathrm{E}_{g\in G}[\rho_y(g)]=(1-\epsilon)\boldsymbol{I}_{d_{\rho_y}},$$

where the last equality follows from Corollary 4. Hence $\mathrm{E}_e[\rho(e^{-1})] = (1-\epsilon)^{|\rho|}\boldsymbol{I}_{d_\rho}$; when this is substituted into the above expression, (10) becomes

$$\operatorname{tr}\mathrm{E}\Big[A(f)B(h)B(h^{-1}(f\circ\pi)^{-1}e)\Big] = \mathrm{E}_f\left[\operatorname{tr}\left(A(f)\sum_{\rho\in\hat{H}} d_\rho(1-\epsilon)^{|\rho|}\operatorname{Tr}\left(\hat{B}_\rho^2\Big(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\Big)\right)\right)\right]. \quad (11)$$

We will now see that for any fixed f, the terms in the resulting sum corresponding to ρ with $|\rho| \geq c$ contribute very little.

Lemma 23 *Let G be a finite group, V be a d_γ-dimensional vector space, $A(f) \in \mathrm{End}(V)$ be unitary, H be a power of G, and $B\colon H \to \mathrm{End}(V)$ be a skew-symmetric function such that for all $h \in H$, $B(h)$ is a convex combination of unitary matrices. Then*

$$\left|\operatorname{tr}\left(A(f)\sum_{\substack{\rho\in\hat{H}\\|\rho|\geq c}} d_\rho(1-\epsilon)^{|\rho|}\operatorname{Tr}\left(\hat{B}_\rho^2\Big(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\Big)\right)\right)\right| \leq d_\gamma(1-\epsilon)^c \quad (12)$$

for any positive real ϵ and any positive integer $c > 0$.

Corollary 24 *Suppose that the verifier in Figure 1 accepts with probability $|G|^{-1}+\delta$. Then, for any unitary $\gamma \in \hat{G}$ and any $c > \lceil(\log\delta - 1)/\log(1-\epsilon)\rceil$, where logs are taken base 2,*

$$\left|\mathrm{E}_{f,U,W}\left[\operatorname{tr}\left(A(f)\sum_{\substack{\rho\in\hat{H}\\|\rho|\geq c}} d_\rho(1-\epsilon)^{|\rho|}\operatorname{Tr}\left(\hat{B}_\rho^2\Big(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\Big)\right)\right)\right]\right| < \frac{d_\gamma\delta}{2}$$

where $A = \gamma\circ A_{U,G}$, $H = G^{\mathrm{SAT}^W}$ and $B(h) = \mathrm{E}_{b\in\{-1,1\}}[((\gamma\circ A_W)(h^b))^b]$

PROOF of Lemma 23. Note that, since B is skew-symmetric, $\hat{B}_\rho$ is Hermitian by Lemma 10 and thus $\hat{B}_\rho^2$ is a positive semidefinite matrix. Hence, a direct application of Lemma 34 bounds the left-hand side of (12) from above by

$$\sum_{\substack{\rho\in\hat{H}\\|\rho|\geq c}} d_\rho(1-\epsilon)^{|\rho|}\operatorname{tr}\hat{B}_\rho^2.$$

By Lemma 33, the above expression can be bounded from above by

$$(1-\epsilon)^c\sum_{\rho\in\hat{H}} d_\rho\operatorname{tr}\hat{B}_\rho^2 = (1-\epsilon)^c\operatorname{tr}\left(\sum_{\rho\in\hat{H}} d_\rho\operatorname{Tr}\hat{B}_\rho^2\right).$$

By Plancherel's equality (8), $\sum_{\rho\in\hat{H}} d_\rho \operatorname{Tr} \hat{B}_\rho^2 = (B * B)(1_H)$. Since the product of unitary matrices is also a unitary matrix, $B(h)B(h^{-1})$ is a convex combination of unitary matrices for every $h \in H$. Hence

$$(B * B)(1_H) = \frac{1}{|H|} \sum_{h\in H} B(h)B(h^{-1})$$

is itself a convex combination of unitary matrices and thus has elements with at most unit magnitude by Corollary 29. Consequently, $\operatorname{tr}((B * B)(1_H))$ is at most d_γ and therefore

$$(1-\epsilon)^c \sum_{\rho\in\hat{H}} d_\rho \operatorname{tr} \hat{B}_\rho^2 \leq d_\gamma (1-\epsilon)^c,$$

which completes the proof.

While we bound the terms corresponding to ρ with $|\rho| \geq c$ by a purely algebraic argument, we bound the terms corresponding to ρ with $|\rho| < c$ by using them to devise a strategy for the provers in the 2P1R game for μ-gap E3-Sat(5). Since this strategy has a success probability that is independent of u, the number of repetitions in the 2P1R game, we can then select u in such a way that also the terms corresponding to ρ with $|\rho| < c$ have to be upper bounded by $d_\gamma \delta/2$.

Lemma 25 *Suppose that for any nontrivial $\gamma \in \hat{G}$,*

$$\left| \mathrm{E}_{f,U,W} \left[\operatorname{tr}\left(A(f) \sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} d_\rho (1-\epsilon)^{|\rho|} \operatorname{Tr}\left(\hat{B}_\rho^2 \left(\boldsymbol{I}_{d_\gamma} \otimes \rho(f \circ \pi) \right) \right) \right) \right] \right| \geq \eta \qquad (13)$$

where $A\colon G^{\{-1,1\}^U} \to \operatorname{End}(V)$ is unitary and $B\colon G^{\mathrm{SAT}^W} \to \operatorname{End}(V)$ is a convex combination of unitary matrices, both A and B are known to both provers in the 2P1R game from §5.1, V is the vector space corresponding to γ, and $H = G^{\mathrm{SAT}^W}$. Then there is a strategy for the provers in the 2P1R protocol with success probability at least $\eta^2 c^{-1} |G|^{-c} d_\gamma^{-6}$.

Corollary 26 *Let $A_{U,G}$ and A_W be the tables in a Standard Written G-Proof with parameter u corresponding to an unsatisfiable formula. Then, for any unitary $\gamma \in \hat{G}$,*

$$\left| \mathrm{E}_{f,U,W} \left[\operatorname{tr}\left(A(f) \sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} d_\rho (1-\epsilon)^{|\rho|} \operatorname{Tr}\left(\hat{B}_\rho^2 \left(\boldsymbol{I}_{d_\gamma} \otimes \rho(f \circ \pi) \right) \right) \right) \right] \right| < \frac{d_\gamma \delta}{2} \qquad (14)$$

where $A = \gamma \circ A_{U,G}$, $H = G^{\mathrm{SAT}^W}$ and $B(h) = \mathrm{E}_{b\in\{-1,1\}}[((\gamma \circ A_W)(h^b))^b]$; provided that $u > \lceil (2\log\delta^{-1} + \log c + c\log|G| + 4\log d_\gamma + 2)/\log c_\mu^{-1} \rceil$ where

c_μ is the constant from §5.1.

PROOF of Lemma 25. Expand $A(f)$ using Fourier inversion (7). Then the left hand side of (13) becomes

$$\left|\mathrm{E}_{U,W,f}\left[\mathrm{tr}\left(\sum_{\tau\in\hat{F}}\sum_{\substack{\rho\in\hat{H}\\|\rho|<c}} d_\tau d_\rho (1-\epsilon)^{|\rho|}\,\mathrm{Tr}\Big(\hat{A}_\tau\big(\boldsymbol{I}_{d_\gamma}\otimes\tau(f^{-1})\big)\Big)\,\mathrm{Tr}\Big(\hat{B}_\rho^2\big(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\big)\Big)\right)\right]\right|$$

where $F = G^{\{-1,1\}^U}$. If this expression is larger than η, then there must be some index t such that

$$\frac{\eta}{d_\gamma} \le \left|\mathrm{E}_{U,W,f}\left[\left(\sum_{\tau\in\hat{F}}\sum_{\substack{\rho\in\hat{H}\\|\rho|<c}} d_\tau d_\rho\,\mathrm{Tr}\Big(\hat{A}_\tau\big(\boldsymbol{I}_{d_\gamma}\otimes\tau(f^{-1})\big)\Big)\,\mathrm{Tr}\Big(\hat{B}_\rho^2\big(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\big)\Big)\right)_{tt}\right]\right| \tag{15}$$

We now fix this value of t. By our notation for the index sets $\iota(\tau)$ and $\iota(\rho)$ and the "projected" representation ρ^π_{op} from §4.2,

$$\begin{aligned}\left(\mathrm{Tr}\Big(\hat{A}_\tau\big(\boldsymbol{I}_{d_\gamma}\otimes\tau(f^{-1})\big)\Big)\right)_{tk} &= \sum_{m,n\in\iota(\tau)}\big(\hat{A}_\tau\big)_{tn,km}\tau_{mn}(f^{-1}),\ \text{and}\\ \left(\mathrm{Tr}\Big(\hat{B}_\rho^2\big(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\big)\Big)\right)_{kt} &= \sum_{o,p\in\iota(\rho)}\big(\hat{B}_\rho^2\big)_{ko,tp}\rho^\pi_{op}(f)\\ &= \sum_{o,p\in\iota(\rho)}\sum_{\substack{q\in\iota(\rho)\\1\le r\le d_\gamma}}(\hat{B}_\rho)_{ko,rq}(\hat{B}_\rho)_{rq,tp}\rho^\pi_{op}(f).\end{aligned}$$

Inserting these expressions into the right hand side of (15), we get

$$\frac{\eta}{d_\gamma} \le \left|\sum_{\substack{1\le k\le d_\gamma\\1\le r\le d_\gamma}}\mathrm{E}_{U,W}\left[\sum_{\tau\in\hat{F}}\sum_{\substack{\rho\in\hat{H}\\|\rho|<c}}\sum_{\substack{m,n\in\iota(\tau)\\o,p,q\in\iota(\rho)}} d_\tau d_\rho(\hat{A}_\tau)_{tn,km}(\hat{B}_\rho)_{ko,rq}(\hat{B}_\rho)_{rq,tp}\,\mathrm{E}_f\Big[\tau_{mn}(f^{-1})\rho^\pi_{po}(f)\ \Big|\ U,W\Big]\right]\right|. \tag{16}$$

Focus now on the innermost expectation

$$\mathrm{E}_f\Big[\tau_{mn}(f^{-1})\rho^\pi_{po}(f)\ \Big|\ U,W\Big] = \langle\rho^\pi_{po}\mid\tau_{mn}\rangle_F.$$

By Lemma 16, this is zero unless $\tau\sim\rho$, where $\sim$ is the relation defined in

Lemma 16. Hence (16) becomes

$$\frac{\eta}{d_\gamma} \leq \left| \sum_{\substack{1\leq k\leq d_\gamma \\ 1\leq r\leq d_\gamma}} \mathrm{E}_{U,W}\left[\sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} \sum_{\substack{\tau\in\hat{F} \\ \tau\sim\rho}} \sum_{\substack{m,n\in\iota(\tau) \\ o,p,q\in\iota(\rho)}} d_\tau d_\rho (\hat{A}_\tau)_{tn,km} (\hat{B}_\rho)_{ko,rq} (\hat{B}_\rho)_{rq,tp} \langle \rho^\pi_{po} \mid \tau_{mn}\rangle_F \right] \right|. \quad (17)$$

We now apply Cauchy-Schwartz twice, first to the sum over k, r and then to the remaining sums, to simplify the above expression further:

$$\begin{aligned} \frac{\eta^2}{d_\gamma^4} &\leq \sum_{\substack{1\leq k\leq d_\gamma \\ 1\leq r\leq d_\gamma}} \mathrm{E}_{U,W}\left[\left| \sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} \sum_{\substack{\tau\in\hat{F} \\ \tau\sim\rho}} \sum_{\substack{m,n\in\iota(\tau) \\ o,p,q\in\iota(\rho)}} d_\tau d_\rho (\hat{A}_\tau)_{tn,km} (\hat{B}_\rho)_{ko,rq} (\hat{B}_\rho)_{rq,tp} \langle \rho^\pi_{po} \mid \tau_{mn}\rangle_F \right|^2 \right] \\ &\leq \sum_{\substack{1\leq k\leq d_\gamma \\ 1\leq r\leq d_\gamma}} \mathrm{E}_{U,W}\left[\left(\sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} \sum_{\substack{\tau\in\hat{F} \\ \tau\sim\rho}} \sum_{\substack{m,n\in\iota(\tau) \\ o,p,q\in\iota(\rho)}} d_\tau d_\rho \left|(\hat{A}_\tau)_{tn,km}\right|^2 \left|(\hat{B}_\rho)_{ko,rq}\right|^2 \right) \right. \\ &\qquad \left. \left(\sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} \sum_{\substack{\tau\in\hat{F} \\ \tau\sim\rho}} \sum_{\substack{m,n\in\iota(\tau) \\ o,p,q\in\iota(\rho)}} d_\tau d_\rho \left|(\hat{B}_\rho)_{rq,tp}\right|^2 \left|\langle \rho^\pi_{po} \mid \tau_{mn}\rangle_F\right|^2 \right) \right] \end{aligned} \quad (18)$$

We proceed by bounding the second factor above, i.e.,

$$\begin{aligned} &\sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} \sum_{\substack{\tau\in\hat{F} \\ \tau\sim\rho}} \sum_{\substack{m,n\in\iota(\tau) \\ o,p,q\in\iota(\rho)}} d_\tau d_\rho \left|(\hat{B}_\rho)_{rq,tp}\right|^2 \left|\langle \rho^\pi_{po} \mid \tau_{mn}\rangle_F\right|^2 \\ &\qquad = \sum_{\substack{\rho\in\hat{H} \\ |\rho|<c}} \sum_{p,q\in\iota(\rho)} d_\rho \left|(\hat{B}_\rho)_{rq,tp}\right|^2 \left(\sum_{\substack{\tau\in\hat{F} \\ \tau\sim\rho}} \sum_{\substack{m,n\in\iota(\tau) \\ o\in\iota(\rho)}} d_\tau \left|\langle \rho^\pi_{po} \mid \tau_{mn}\rangle_F\right|^2 \right). \end{aligned} \quad (19)$$

By equation (5) in Lemma 8,

$$\sum_{\tau\in\hat{F}} \sum_{\substack{m,n\in\iota(\tau) \\ o\in\iota(\rho)}} d_\tau \left|\langle \rho^\pi_{po} \mid \tau_{nm}\rangle_F\right|^2 = \frac{1}{|F|} \sum_{f\in F} \sum_{o\in\iota(\rho)} |\rho_{po}(f\circ\pi)|^2 = 1,$$

where the last equality follows since ρ is written in a unitary basis and the inner sum is therefore exactly one for every f by Lemma 28. Regarding the rest of (19), note that $(\hat{B}_\rho)_{rq,tp} = |H|^{-1} \sum_{h\in H} B_{rt}(h)\rho_{qp}(h) = \langle B_{rt} \mid \rho_{qp}^{-1}\rangle_H$; another application of Lemma 8 therefore shows that

$$\sum_{\rho\in\hat{H}} \sum_{p,q\in\iota(\rho)} d_\rho \left|(\hat{B}_\rho)_{rq,tp}\right|^2 = \sum_{\rho\in\hat{H}} \sum_{p,q\in\iota(\rho)} d_\rho \left|\langle B_{rt} \mid \rho_{qp}^{-1}\rangle_H\right|^2 = \frac{1}{|H|} \sum_{h\in H} |B_{rt}(h)|^2 \leq 1,$$

where the inequality follows since B is a convex combination of unitary matrices and therefore has entries with at most unit magnitude by Corollary 29. Using the above bounds in (18) transforms that bound into

$$\begin{aligned}\frac{\eta^2}{d_\gamma^4} &\leq \sum_{\substack{1\leq k\leq d_\gamma\\1\leq r\leq d_\gamma}} \mathrm{E}_{U,W}\left[\sum_{\substack{\tau\in\hat{F}\\ \tau\sim\rho}}\sum_{\substack{\rho\in\hat{H}\\ |\rho|<c}}\sum_{\substack{m,n\in\iota(\tau)\\ o,p,q\in\iota(\rho)}} d_\tau d_\rho \left|(\hat{A}_\tau)_{tn,km}\right|^2 \left|(\hat{B}_\rho)_{ko,rq}\right|^2\right] \\ &\leq |G|^c \sum_{\substack{1\leq k\leq d_\gamma\\1\leq r\leq d_\gamma}} \mathrm{E}_{U,W}\left[\sum_{\substack{\tau\in\hat{F}\\ \tau\sim\rho}}\sum_{\substack{\rho\in\hat{H}\\ |\rho|<c}}\sum_{\substack{m,n\in\iota(\tau)\\ o,q\in\iota(\rho)}} d_\tau d_\rho \left|(\hat{A}_\tau)_{tn,km}\right|^2 \left|(\hat{B}_\rho)_{ko,rq}\right|^2\right]\end{aligned}$$

where the second inequality follows by summing over p in the innermost sum. To conclude, there must be some $k, r \in \{1, \ldots, d_\gamma\}$ such that

$$\mathrm{E}_{U,W}\left[\sum_{\substack{\rho\in\hat{H}\\ |\rho|<c}}\sum_{\substack{\tau\in\hat{F}\\ \tau\sim\rho}}\sum_{\substack{m,n\in\iota(\tau)\\ o,q\in\iota(\rho)}} d_\tau d_\rho \left|(\hat{A}_\tau)_{tn,km}\right|^2 \left|(\hat{B}_\rho)_{ko,rq}\right|^2\right] \geq \frac{\eta^2}{|G|^c d_\gamma^6}. \tag{20}$$

We now describe the strategies for the provers in the 2P1R protocol. The index t is independent of U and W and can be calculated by the provers in advance. Also the values of k and r mentioned above can be calculated in advance.

Upon receiving W, P_1 first picks $\rho \in \hat{H}$ with probability proportional to $\sum_{o,q\in\iota(\rho)} d_\rho |(\hat{B}_\rho)_{ko,rq}|^2$. This is a well-defined procedure since

$$\sum_{\rho\in\hat{H}}\sum_{o,q\in\iota(\rho)} d_\rho \left|(\hat{B}_\rho)_{ko,rq}\right|^2 = \frac{1}{|H|}\sum_{h\in H}|B_{kr}(h)|^2 \leq 1.$$

Having selected ρ, P_1 then returns a random y such that ρ_y is nontrivial. If no such y exists—this happens only if ρ is trivial—P_1 gives up.

Upon receiving U, P_2 picks $\tau \in \hat{F}$ with probability $\sum_{n,m\in\iota(\tau)} d_\tau |(\hat{A}_\tau)_{tn,km}|^2$. This is a well-defined procedure since

$$\sum_{\tau\in\hat{F}}\sum_{n,m\in\iota(\tau)} d_\tau \left|(\hat{A}_\tau)_{tn,km}\right|^2 = \frac{1}{|F|}\sum_{f\in F}|A_{tk}(f)|^2 \leq 1.$$

Then P_2 picks a random x such that τ_x is nontrivial and returns this x as its answers. This is always possible since $\hat{A}_\tau$ is nonzero only for nontrivial τ by Lemma 12.

To give a lower bound on the success rate of this strategy, we argue that there are many choices of the provers that make the verifier accept: Specifically, suppose that P_1 picks ρ and P_2 picks τ such that $\tau \sim \rho$. If P_2 returns x', then there must be some $y' \in \pi^{-1}(x')$ such that $\rho_{y'}$ is nontrivial. The probability of

P_1 picking this y' is at least $|\rho|^{-1}$. Summing over all τ and ρ such that $\tau \sim \rho$, we get that the probability of success is at least

$$\mathrm{E}_{U,W}\left[\sum_{\substack{\rho\in\hat{H}\\|\rho|<c}}\sum_{\substack{\tau\in\hat{F}\\\tau\sim\rho}}\sum_{\substack{m,n\in\iota(\tau)\\o,q\in\iota(\rho)}}\frac{d_\tau d_\rho\left|(\hat{A}_\tau)_{tn,km}\right|^2\left|(\hat{B}_\rho)_{ko,rq}\right|^2}{|\rho|}\right] \geq \frac{\eta^2}{c|G|^c d_\gamma^6}$$

where the inequality follows from the bound (20).

Finally, we put together these two parts and establish the soundness of the verifier.

Lemma 27 *For any constants $\delta > 0$ and $0 < \epsilon < 1$, there is a choice of the parameters c and u such that the soundness of the PCP in Fig. 1 is at most $|G|^{-1} + \delta$.*

PROOF. Suppose for contradiction that ϕ is not satisfiable and there is a proof which the verifier accepts with probability $|G|^{-1} + \delta$. By Lemma 22, for this proof, there is a nontrivial irreducible representation γ of G such that

$$\left|\mathrm{E}_{f,U,W}\left[\mathrm{tr}\left(A(f)\sum_{\rho\in\hat{H}} d_\rho(1-\epsilon)^{|\rho|}\,\mathrm{Tr}\left(\hat{B}_\rho^2\left(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\right)\right)\right)\right]\right| > d_\gamma\delta.$$

where $A = \gamma \circ A_{U,G}$, $H = G^{\mathrm{SAT}^W}$ and $B(h) = \mathrm{E}_{b\in\{-1,1\}}[((\gamma\circ A_W)(h^b))^b]$. However, by selecting constants $c > \lceil(\log\delta - 1)/\log(1-\epsilon)\rceil$ and $u > \lceil(2\log\delta^{-1} + \log c + c\log|G| + 4\log d_\gamma + 2)/\log c_\mu^{-1}\rceil$, Corollaries 24 and 26 show that

$$\left|\mathrm{E}_{f,U,W}\left[\mathrm{tr}\left(A(f)\sum_{\substack{\rho\in\hat{H}\\|\rho|\geq c}} d_\rho(1-\epsilon)^{|\rho|}\,\mathrm{Tr}\left(\hat{B}_\rho^2\left(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\right)\right)\right)\right]\right| < \frac{d_\gamma\delta}{2},$$

$$\left|\mathrm{E}_{f,U,W}\left[\mathrm{tr}\left(A(f)\sum_{\substack{\rho\in\hat{H}\\|\rho|<c}} d_\rho(1-\epsilon)^{|\rho|}\,\mathrm{Tr}\left(\hat{B}_\rho^2\left(\boldsymbol{I}_{d_\gamma}\otimes\rho(f\circ\pi)\right)\right)\right)\right]\right| < \frac{d_\gamma\delta}{2},$$

which is a contradiction.

5.4 Hardness of approximating $\mathrm{EQ}_G^1[3]$

We now apply this PCP to obtain hardness results for approximating systems of equations over G.

PROOF of Theorem 1. Let G be a finite group and let $0 < \epsilon < 1$ and $0 < \delta < 1$ be two constants satisfying the inequality $|G|^{-1} + \delta < 1 - \epsilon$. By Lemma 21 and Lemma 27 it is possible to choose the parameters of the verifier in Figure 1 such that

(1) the constant u is chosen so that $|G|^{-2^u} < \delta/6$, and
(2) it is **NP**-hard to distinguish between the case that there is a proof which the verifier accepts with probability $1 - \epsilon/2$, and the case that there is no proof which is accepted with probability more than $|G|^{-1} + \delta/2$.

Now we create a system of equations $\mathcal{E}$ in the natural way: the variables correspond to the positions in the proofs, and an equation is added for each random string corresponding to the test made for this random string. By a discussion similar to that in Lemma 20, it can be shown that the instance so obtained contains $m = 4(5n)^u|G|^{2^u}(2|G|^2/\epsilon)^{7^u}$ equations, which is polynomial in n as u, $|G|$, and ϵ are constants. One may think that the equations would always be of the form $xy^iz^j = 1_G$, but this is not the case due to folding over G; in general an equation is of the form $gxy^iz^j = 1_G$, where g is a group constant and $i, j \in \{1, -1\}$.

There is a technicality in that when $h^{b_1} = (h^{-1}(f \circ \pi)^{-1}e)^{b_2}$ in the protocol, the resulting equation contains two occurrences of the same variable. Observe, however, that $h^{b_1} = (h^{-1}(f \circ \pi)^{-1}e)^{b_2} \iff (f \circ \pi) = eh^{-b_1 b_2 - 1}$ and, as π is onto, the probability η that $(f \circ \pi)$ takes this particular value is no more than $|G|^{-2^u} < \delta/6$. Thus, removing ηm equations from $\mathcal{E}$ results in a new family of equations $\mathcal{E}'$ of size $m' = (1 - \eta)m \geq (1 - \delta/6)m$ which is indeed a proper instance of $\mathrm{EQ}^1_G[3]$. Moreover, if it is possible to satisfy at least $(1 - \epsilon/2)m$ equations in $\mathcal{E}$, then it is possible to satisfy at least

$$(1 - \epsilon/2)m - \eta m = \frac{(1 - \epsilon/2 - \eta)}{(1 - \eta)}m' = \left(1 - \frac{\epsilon}{2(1 - \eta)}\right)m' > (1 - \epsilon)m'$$

equations of $\mathcal{E}'$, where the last inequality follows since $\eta < \delta/6 < 1/2$. Similarly, if it possible to satisfy no more than $(|G|^{-1} + \delta/2)m$ equations of $\mathcal{E}$, then it is possible to satisfy no more than

$$\begin{aligned}\left(|G|^{-1} + \frac{\delta}{2}\right)m = \frac{|G|^{-1} + \delta/2}{1 - \eta}m' &< \left(|G|^{-1} + \frac{\delta}{2}\right)(1 + 2\eta)m' \\ &< \left(|G|^{-1} + \frac{\delta}{2}\right)\left(1 + \frac{\delta}{3}\right)m' \leq \left(|G|^{-1} + \delta\right)m'\end{aligned}$$

equations of $\mathcal{E}'$, where the first inequality follows since $(1 - \eta)^{-1} < 1 + 2\eta$. Furthermore, by appealing to condition (2) above, distinguishing these two cases is **NP**-hard, as desired.

6 Open questions

An interesting question is that of *satisfiable instances*. Some problems, such as E3-Sat, retain their inapproximability properties even when restricted to satisfiable instances. This is not the case for $\mathrm{EQ}^1_G[k]$ when G is a finite Abelian group, since if such a system is satisfiable a solution may be found essentially by Gaussian elimination. However, when G is non-Abelian, deciding whether a system of equations over G is satisfiable is **NP**-complete [14], so it seems reasonable that the problem over non-Abelian groups retains some hardness of approximation for satisfiable instances. However, the following simple argument shows that we can not hope, even for the non-Abelian groups, for a lower bound of $|G|^{-1} + \delta$: Given an instance σ of $\mathrm{EQ}^1_G[k]$ over some non-Abelian group G, we construct an instance σ' over $\mathrm{EQ}^1_H[k]$, where $H = G/G'$ and G' is the commutator subgroup of G, i.e., the subgroup generated by the elements $\{g^{-1}h^{-1}gh : g, h \in G\}$. The instance σ' is the same as σ, except that all group constants are replaced by their equivalence class in G/G'. Now since H is an Abelian group, we can solve over H. The solution is an assignment of cosets to the variables. We then construct a random solution of x by for each variable choosing a random element in the corresponding coset. Now the value of the left hand side of each equation will be uniformly distributed in the coset of the right hand side, and thus we will satisfy an expected fraction $|G'|^{-1}$ of all equations.

7 Acknowledgments

Volker Diekert gave many valuable comments on a preliminary version of this paper and simplified the proof of Lemma 23. We would also like to thank Johan Håstad for useful discussions and the anonymous referees for constructive comments that helped improve the presentation of our results.

A Identities from linear and multilinear algebra

This appendix contains the identities and bounds that are needed in the proof of Lemma 23. They all follow in a straightforward manner from standard linear and multilinear algebra. As a service to the reader, we also include a very short summary of the less known background results from linear algebra that we use in this paper; for more information on linear algebra, the reader is referred to Lang's book [18].

We first recall that a complex matrix $A = (a_{ij})$ is *unitary* if $A^{-1} = A^*$, where the matrix A^* has a^*_{ji} at position (i, j) and the latter asterisk denotes complex conjugation: For a complex number $z = x + iy$, $z^* = x - iy$ and $|z|^2 = x^2 + y^2 = zz^*$. Then recall that a matrix is *Hermitian* if $A = A^*$ and that Hermitian matrices have only real eigenvalues. Since the eigenvalues of A^2 are the squares of the eigenvalues of A, the square of a Hermitian matrix has only non-negative real eigenvalues, i.e., it is *positive semidefinite.*

Lemma 28 *Let $A = (a_{ij})$ be a unitary $n \times n$ matrix. Then $\sum_j |a_{ij}|^2 = 1$ for all $1 \le i \le n$.*

PROOF. Let A be unitary. Then AA^* is the identity matrix. Since for all $1 \le i \le n$, $(AA^*)_{ii} = \sum_j a_{ij}a^*_{ij} = \sum_j |a_{ij}|^2$, the lemma follows.

Corollary 29 *Let $\{A_k\}$ be a family of unitary $n \times n$ matrices and let $\{\lambda_k\}$ be a sequence of non-negative real numbers such that $\sum_k \lambda_k = 1$. Let $B = \sum_k \lambda_k A_k$. Then the elements of B have at most unit magnitude.*

PROOF. By Jensen's inequality $|(B)_{ij}| = |\sum_k \lambda_k (A_k)_{ij}| \le \sum_k \lambda_k |(A_k)_{ij}|$. Lemma 28 implies that $|(A_k)_{ij}| \le 1$; hence $|(B)_{ij}| \le \sum_k \lambda_k = 1$.

Lemma 30 *For any complex numbers $\{a_1, \ldots, a_n\}$ and $\{b_1, \ldots, b_n\}$,*

$$\left|\sum_{i=1}^n a_i b_i\right|^2 \le \left(\sum_{i=1}^n |a_i|^2\right)\left(\sum_{i=1}^n |b_i|^2\right).$$

As a special case, corresponding to $b_i = 1$,

$$\left|\sum_{i=1}^n a_i\right|^2 \le n \sum_{i=1}^n |a_i|^2.$$

PROOF. This is the Cauchy-Schwartz inequality; we provide a proof for the sake of completeness. If $a_i = 0$ for all i, the inequality clearly holds. Otherwise, let $a = \sum_{i=1}^n |a_i|^2$, $b = 2|\sum_{i=1}^n a_i b_i|$, and $c = \sum_{i=1}^n |b_i|^2$. Now,

$$\sum_{i=1}^n |ta_i + b^*_i|^2 = \sum_{i=1}^n \Big(t^2|a_i|^2 + t(a_i b_i + a^*_i b^*_i) + |b_i|^2\Big) = t^2 a + tb + c$$

for any real t. Since the above sum is non-negative, $t^2 a + tb + c \ge 0$, or, equivalently, $4ac \ge b^2$.

A.2 Tensor products and traces

Given two matrices $A = (a_{ij})$ and $B = (b_{k\ell})$, the *tensor product* $A \otimes B$ is the matrix indexed by pairs $(i,k);(j,\ell)$ so that $(A \otimes B)_{(i,k);(j,\ell)} = a_{ij}b_{k\ell}$. Note that $(A_1 \otimes A_2)(B_1 \otimes B_2) = (A_1B_1) \otimes (A_2B_2)$ and that the tensor product is bilinear. For a matrix M indexed by pairs $(i,k);(j,\ell)$ the inner trace, denoted by $\operatorname{Tr} M$, is defined by $(\operatorname{Tr} M)_{ij} = \sum_k M_{(i,k),(j,k)}$. We let tr denote the normal trace, i.e., the sum of the diagonal elements. The trace is invariant under similarity, i.e., $\operatorname{tr} A = \operatorname{tr}(U^*AU)$ for any unitary matrix U. Furthermore, the trace of a matrix is equal to the sum of its eigenvalues.

Lemma 31 *Let A be a tensor product of an $n \times n$ matrix and a $k \times k$ matrix and let B be an $n \times n$ matrix. Then $(\operatorname{Tr} A)B = \operatorname{Tr}(A(B \otimes \boldsymbol{I}_k))$.*

PROOF. Suppose that $A = A_1 \otimes A_2$. Using the identity $(A_1 \otimes A_2)(B_1 \otimes B_2) = (A_1B_1) \otimes (A_2B_2)$ we obtain that $A(B \otimes \boldsymbol{I}_k) = (A_1B) \otimes A_2$. Since $\operatorname{Tr}((A_1B) \otimes A_2) = (\operatorname{tr} A_2)A_1B$ it follows that $\operatorname{Tr}(A(B \otimes \boldsymbol{I}_k)) = (\operatorname{tr} A_2)A_1B = (\operatorname{Tr} A)B$.

Lemma 32 *Let A be a positive semidefinite matrix and U be a unitary matrix. Then $|\operatorname{tr}(AU)| \leq \operatorname{tr} A$.*

PROOF. As A is positive semidefinite, it may be written VDV^* where V is unitary and D is diagonal with non-negative entries on the diagonal. Since the trace is invariant under similarity, $\operatorname{tr} A = \sum_i D_{ii}$ and we may rewrite

$$\operatorname{tr}(AU) = \operatorname{tr}(VDV^*U) = \operatorname{tr}(VDV^*UVV^*) = \operatorname{tr}(DV^*UV) = \operatorname{tr}(DW),$$

where $W = V^*UV$ is a product of unitary matrices and therefore unitary. All entries of a unitary matrix have absolute value less than or equal to one; hence

$$|\operatorname{tr}(AU)| = \left|\sum_i D_{ii}W_{ii}\right| \leq \sum_i |D_{ii}W_{ii}| \leq \sum_i D_{ii} = \operatorname{tr} A,$$

as desired.

A.3 Bounds used in Lemma 23

Lemma 33 *Let $\epsilon \in [0,1]$, S_ρ be a family of positive semidefinite matrices and d_ρ and n_ρ be positive integers. Then*

$$\sum_{\rho:n_\rho>c} d_\rho(1-\epsilon)^{n_\rho} \operatorname{tr} S_\rho \leq (1-\epsilon)^c \sum_\rho d_\rho \operatorname{tr} S_\rho.$$

PROOF. Since the S_ρ is positive semidefinite $\sum_{\rho:n_\rho>c} d_\rho(1-\epsilon)^{n_\rho} \operatorname{tr} S_\rho$ is a sum of non-negative numbers and thus non-negative. Therefore

$$\sum_{\rho:n_\rho>c} d_\rho(1-\epsilon)^{n_\rho} \operatorname{tr} S_\rho \leq (1-\epsilon)^c \sum_{\rho:n_\rho>c} d_\rho \operatorname{tr} S_\rho \leq (1-\epsilon)^c \sum_\rho d_\rho \operatorname{tr} S_\rho.$$

Lemma 34 *Let U be an $n \times n$ unitary matrix, $\{S_\rho\}$ be a family of positive semidefinite matrices that are tensor products of $n \times n$ matrices and $k \times k$ matrices, $\{V_\rho\}$ be a family of unitary matrices that are tensor products of $n \times n$ matrices and $k \times k$ matrices, and $\{a_\rho\}$ be a family of non-negative real numbers. Then*

$$\left|\operatorname{tr}\left(U \sum_\rho a_\rho \operatorname{Tr}(S_\rho V_\rho)\right)\right| \leq \sum_\rho a_\rho \operatorname{tr} S_\rho,$$

where the inner trace is with respect to the tensor products forming S_ρ and V_ρ.

PROOF. Since $\operatorname{tr}(AB) = \operatorname{tr}(BA)$ for any matrices A and B,

$$\begin{aligned}\left|\operatorname{tr}\left(U \sum_\rho a_\rho \operatorname{Tr}(S_\rho V_\rho)\right)\right| &= \left|\operatorname{tr}\left(\sum_\rho a_\rho \operatorname{Tr}(S_\rho V_\rho) U\right)\right| \\ &= \left|\operatorname{tr}\left(\sum_\rho a_\rho \operatorname{Tr}\left(S_\rho V_\rho (U \otimes I_k)\right)\right)\right|\end{aligned}$$

where the last equality follows from Lemma 31. Since the a_ρ are non-negative,

$$\begin{aligned}\left|\operatorname{tr}\left(\sum_\rho a_\rho \operatorname{Tr}\left(S_\rho V_\rho (U \otimes I_k)\right)\right)\right| &\leq \sum_\rho a_\rho \left|\operatorname{tr} \operatorname{Tr}\left(S_\rho V_\rho (U \otimes I_k)\right)\right| \\ &= \sum_\rho a_\rho \left|\operatorname{tr}\left(S_\rho V_\rho (U \otimes I_k)\right)\right| \\ &\leq \sum_\rho a_\rho \operatorname{tr} S_\rho,\end{aligned}$$

where the last inequality follows from Lemma 32.

References

[1] N. Alon, A. Spencer, The Probabilistic Method, 2nd edition, John Wiley & Sons, 2000.

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, M. Szegedy, Proof verification and the hardness of approximation problems, Journal of the ACM 45 (3) (1998) 501–555.

[3] D. A. M. Barrington, P. McKenzie, C. Moore, P. Tesson, D. Thérien, Equation satisfiability and program satisfiability for finite monoids, in: Proceedings of the

25th Annual Symposium on Mathematical Foundations of Computer Science, Bratislava, Slovak Republic, 2000, pp. 172–181.

[4] D. A. M. Barrington, H. Straubing, D. Thérien, Non-uniform automata over groups, Information and Computation 89 (2) (1990) 109–132.

[5] D. A. M. Barrington, D. Thérien, Finite monoids and the fine structure of NC^1, Journal of the ACM 35 (4) (1988) 941–952.

[6] F. Bédard, A. Goupil, The poset of conjugacy classes and decomposition of products in the symmetric group, Canadian Mathematical Bulletin 35 (2) (1992) 152–160.

[7] M. Bellare, O. Goldreich, M. Sudan, Free bits, PCPs and non-approximability—towards tight results, SIAM Journal on Computing 27 (3) (1998) 804–915.

[8] I. Dinur, S. Safra, The importance of being biased, in: Proceedings of the 34th Annual ACM Symposium on Theory of Computing, Montréal, Quebéc, Canada, 2002, pp. 33–42.

[9] H. Dym, H. P. McKean, Fourier Series and Integrals, Vol. 14 of Probability and Mathematical Statistics, Academic Press, 1972.

[10] L. Engebretsen, The non-approximability of non-Boolean predicates, in: M. Goemans, K. Jansen, J. D. P. Rolim, L. Trevisan (Eds.), Proceedings of 5th International Workshop on Randomization and Approximation Techniques in Computer Science, Vol. 2129 of Lecture Notes in Computer Science, Springer-Verlag, Berkeley, California, USA, 2001, pp. 241–248.

[11] L. Engebretsen, J. Holmerin, Towards optimal lower bounds for clique and chromatic number, Theoretical Computer Science, to appear.

[12] U. Feige, A threshold of $\ln n$ for approximating set cover, Journal of the ACM 45 (4) (1998) 634–652.

[13] H. Finkelstein, K. I. Mandelberg, On solutions of "equations in symmetric groups", Journal of Combinatorial Theory, Series A 25 (2) (1978) 142–152.

[14] M. Goldmann, A. Russell, The computational complexity of solving systems of equations over finite groups, in: Proceedings of the Fourteenth Annual IEEE Conference on Computational Complexity, Atlanta, Georgia, 1999, pp. 80–86.

[15] V. Guruswami, J. Håstad, M. Sudan, Hardness of approximate hypergraph coloring, in: 41st Annual Symposium on Foundations of Computer Science, IEEE, Redondo Beach, California, 2000, pp. 149–158.

[16] J. Håstad, Some optimal inapproximability results, Journal of the ACM 48 (4) (2001) 798–859.

[17] S. Khot, Improved inapproximability results for maxclique, chromatic number and approximate graph coloring, in: 42nd Annual Symposium on Foundations of Computer Science, IEEE, Las Vegas, Nevada, 2001, pp. 600–609.

[18] S. Lang, Linear Algebra, Undergraduate Texts in Mathematics, Springer-Verlag, New York, 1987.

[19] C. H. Papadimitriou, M. Yannakakis, Optimization, approximation, and complexity classes, Journal of Computer and System Sciences 43 (3) (1991) 425–440.

[20] R. Raz, A parallel repetition theorem, SIAM Journal on Computing 27 (3) (1998) 763–803.

[21] A. Samorodnitsky, L. Trevisan, A PCP characterization of NP with optimal amortized query complexity, in: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, Portland, Oregon, 2000, pp. 191–199.

[22] A. Schönhage, V. Strassen, Schnelle Multiplikation großer Zahlen, Computing 7 (1971) 281–292.

[23] J.-P. Serre, Linear Representations of Finite Groups, Vol. 42 of Graduate Texts in Mathematics, Springer-Verlag, New York, 1977.

[24] R. P. Stanley, Enumerative Combinatorics, Volume 2, Vol. 62 of Cambridge Studies in Advances Mathematics, Cambridge University Press, Cambridge, 1999.

[25] S. P. Strunkov, K teorii uravnenij na konečnyh gruppah, Izvestiâ Rossijskoj Akademii Nauk, Seriâ Matematičeskaâ 59 (6) (1995) 171–180, English translation: On the theory of equations in finite groups. Izvestiya: Mathematics, 59(6):1273–1282, 1995.

[26] A. Terras, Fourier Analysis on Finite Groups and Applications, Vol. 43 of London Mathematical Society student texts, Cambridge University Press, Cambridge, 1999.

[27] U. 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, San Francisco, California, 1998, pp. 201–210.

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'