Electronic Colloquium on Computational Complexity, Report No. 40 (2002)

Three-Query PCPs with Perfect Completeness over non-Boolean Domains

Lars Engebretsen and Jonas Holmerin

Department of Numerical Analysis and Computer Science
Royal Institute of Technology
SE-100 44 Stockholm
SWEDEN
E-mail: {enge,joho}@kth.se

Abstract. We study non-Boolean PCPs that have perfect completeness and read three positions from the proof. For the case when the proof consists of values from a domain of size d for some integer constant $d \geq 2$, we construct a non-adaptive PCP with perfect completeness and soundness $d^{-1} + d^{-2} + \varepsilon$, for any constant $\varepsilon > 0$, and an adaptive PCP with perfect completeness and soundness $d^{-1} + \varepsilon$, for any constant $\varepsilon > 0$. These results match the best known constructions for the case $d = 2$ and our proofs also show that the particular predicates we use in our PCPs are non-approximable beyond the random assignment threshold.

1 Introduction

A language belongs to **NP** if it has the property that inputs in the language admit a proof that can be verified in polynomial time. More specifically, if L is an **NP** language, there exists a polynomial time verifier V_L with the following properties: For every $x \in L$, there exists a proof y with size polynomial in the size of x such that V_L accepts (x, y), but for $x \notin L$ there does not exist any polynomially sized proof y such that V_L accepts (x, y). In the above definition, the verifier V_L is *deterministic*, i.e., it always accepts a proof of a correct input and never accepts a proof of an incorrect input. One way to modify the definition of **NP** is to let the verifier be probabilistic and allow it to make mistakes. One can also change the notion of a proof and instead view the process of verifying a proof as an interaction between the verifier and several provers.

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 [1]. In the PCP model, the proof can be viewed as a table that the verifier has oracle access to. The verifier also has access to a specified amount of random bits. Based on the random bits and the input, the verifier decides which positions in the proof it should look at. Once it has examined the

ISSN 1433-8092

positions of its choice, it uses all available information to decide if the input should be accepted or rejected. The PCP theorem asserts the startling fact that any language in **NP** can be probabilistically checked by a verifier that uses logarithmic randomness, always accepts a correct proof of an input in the language, accepts proofs of inputs not in the language with probability at most $1/2$, and examines a *constant* number of bits of the proof.

PCPs using a logarithmic number of random bits can be used to prove approximation hardness results for many combinatorial optimization problems. In particular, PCPs querying a small number of bits, say k bits, are intimately connected with Boolean constraint satisfaction problems on k variables. In a constraint satisfaction problem, we are given a set of constraints over some variables and are asked to find an assignment satisfying as many constraints as possible. Indeed, a k-query PCP for **NP** with logarithmic randomness, completeness c, and soundness s immediately gives rise to a constraint satisfaction problem on k variables that is **NP**-hard to approximate within c/s, as long as the verifier is *non-adaptive*, that is, the queries may not depend on the values of previously queried bits. In this case, the problem of computing the proof which maximizes the acceptance probability is a constraint satisfaction problem corresponding to the constraint verified by the verifier. For instance, if the verifier reads three bits from the proof and accepts for seven out of the eight possible answers, the corresponding constraint satisfaction problem is maximum E3-satisfiability.

A natural question to ask is: What is the best soundness that can be achieved with a small number of queries? In his seminal paper [6] Håstad constructed a three-query PCP where s is arbitrarily close to $1/2$. To be more precise: for any $L \in \mathbf{NP}$ and any given constant $\varepsilon > 0$, Håstad's construction gives a proof system that accepts inputs in L with probability at least $1-\varepsilon$ and accepts inputs that are not in L with probability at least $1/2+\varepsilon$. Since Håstad's proof system accepts inputs in L only with probability $1-\varepsilon$, and not always, we say that it has *near-perfect completeness*. By contrast, a proof system is said to have *perfect completeness* if inputs in the language are always accepted.

Samorodnitsky and Trevisan [13] used a clever extension of Håstad's construction to prove that any language in **NP** can be recognized by a (k^2+2k)-query PCP with near-perfect completeness and soundness $2^{-k^2}+\varepsilon$ for any $\varepsilon > 0$. This result is known to be essentially optimal since it is not possible to get better soundness than 2^{1-q} for a non-adaptive q-query PCP for **NP** unless $\mathbf{P} = \mathbf{NP}$ [16].

For several reasons we would prefer the proof systems to have perfect completeness. Firstly, it is simply esthetically pleasing to have a proof system where a correct proof is always be accepted. Secondly, in several proofs of hardness of approximation, perfect completeness has been important. For example, when proving hardness of hypergraph coloring [4, 11] and vertex cover on hypergraphs [9] perfect completeness appears to be crucial. Thirdly, it is of theoretical interest to get an understanding of how near-perfect versus perfect completeness and adaptive versus non-adaptive verifiers affect our ability to construct proof systems with low soundness. For non-adaptive 3-query PCPs, Håstad's construction provides a PCP for **NP** with near-perfect completeness and soundness $1/2+\varepsilon$ for every constant $\varepsilon > 0$ and it is impossible to achieve soundness better than $5/8$ in proof systems with perfect completeness unless $\mathbf{P} = \mathbf{NP}$ [17, 18]. Using the tools from Håstad's paper, it is possible to construct a three query PCP with perfect completeness and soundness $3/4+\varepsilon$ for any constant $\varepsilon > 0$ and Guruswami *et al.* [5], building on Håstad's work, proved that there is an adaptive PCP for **NP** with perfect completeness and soundness $1/2+\varepsilon$ for any constant $\varepsilon > 0$. For more than 3 queries, Håstad and Khot [7] proved that a non-adaptive verifier can achieve perfect completeness and soundness $2^{-k^2}+\varepsilon$ with k^2+4k queries, while an adaptive verifier can achieve the same completeness and soundness with only k^2+2k queries.

In all PCPs mentioned so far, the proof has been a list of bits, i.e., the proof has been *binary*. Håstad also constructed a PCP where the positions in the proof contain elements from a larger domain in his paper [6]: For the case when each position in the proof contains an element from a domain of size d, where d is an arbitrary integer constant, he constructed a proof system with near-perfect completeness and soundness $d^{-1}+\varepsilon$ for any constant $\varepsilon > 0$. Engebretsen [2] adapted the construction of Samorodnitsky and Trevisan [13] to domains of size d, thereby constructing a proof system that queries k^2+2k positions in the proof, has near-perfect completeness and soundness $d^{-k^2}+\varepsilon$ for any $\varepsilon > 0$.

This brings up the question that we address in this paper: How does the soundness behave, as a function of the domain size d, for 3-query PCPs with perfect completeness? The result of Håstad and Khot for non-adaptive PCP generalizes to domains of size p where p is a prime to give soundness $p^{-k^2}+\varepsilon$ with $4k+k^2$ queries. However, since $4k+k^2=5$ when $k=1$, this doesn't give us anything for the case of three queries. It is easy to construct a proof system with perfect completeness and soundness $(d-1)/d$ and Holmerin [8] has constructed a proof system with soundness $2/d$ for $d \geq 3$. In this paper, we prove the following:

Theorem 1. *Any language in **NP** is decided by a non-adaptive PCP with answers from a domain of size d that queries three positions in the proof, has perfect completeness and soundness $d^{-1}+d^{-2}+\varepsilon$ for any constant $\varepsilon > 0$.*

Theorem 2. *Any language in **NP** is decided by an adaptive PCP with answers from a domain of size d that queries three positions in the proof, has perfect completeness and soundness $d^{-1}+\varepsilon$ for any constant $\varepsilon > 0$.*

This generalizes the constructions for the binary case [5, 6] and shows that we can achieve, as a function of the domain size, the same higher order term in the soundness for PCPs with perfect completeness as for the best currently known PCPs with near-perfect completeness. As a technical component in our construction, we use parts of new variant of Håstad's PCP with a simpler analysis, due to Khot [10]. We remark that the hardness results obtained are the best possible for the constraints we use in our verifiers since a random assignment to the proof makes the verifier accept with probabilities $d^{-1}+d^{-2}$ and d^{-1}, respectively. Hence, we establish in this paper that the predicates are what Håstad [6] calls "non-approximable beyond the random assignment threshold".

It appears difficult to get stronger results using current techniques. For $d > 2$ it is currently only known that it is impossible to get soundness d^{-2} in a 3-query PCP for **NP** [14] and it is at present unknown if it is possible to get better soundness with near-perfect completeness than with perfect completeness.

2 PCPs and hardness of approximation

In his paper [6], Håstad introduced a methodology for proving lower bounds 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 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.

2.1 The 2-prover 1-round protocol

We start with an instance of the **NP**-hard [1, 3] problem μ-gap E3-Sat(5).

Definition 1. 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 soundness by repeating u times in parallel as follows: The verifier picks u clauses $(C_1, \dots, C_u)$ 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 $(C_1, \dots, C_u)$ to P_1 and $(x_1, \dots, x_u)$ to P_2. It receives an assignment to the variables in $(C_1, \dots, C_u)$ from P_1 and an assignment to $(x_1, \dots, x_u)$ from P_2, and accepts if these assignments are consistent and satisfy $C_1 \wedge \dots \wedge C_u$. As above, the completeness of this proof system is 1, and it can be shown [12] 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.

2.2 A more balanced setting

In this paper, we use a version of the u-parallel repetition of the basic 2P1R protocol that was recently applied by Khot [10] to Håstad's PCP for E3-Sat [6]. It turns out that a problem in the analysis of the particular type of PCP verifier that Håstad had in his PCP for **NP** with perfect completeness and soundness 3/4—and that we have in our PCP in this paper—is that a large set of satisfying assignments to the clauses $(C_1, \dots, C_u)$ from § 2.1 may project down to a very small set of assignments to the variables $(x_1, \dots, x_u)$. Håstad solved this problem by making a very careful analysis of the PCP verifier. Khot recently obtained a simpler analysis by a modification of the basic 2P1R protocol from § 2.1.

The modified protocol is parameterized by both u and T. In this version, the verifier selects at random a multiset W consisting of $(T+1)u$ clauses. It then selects at random a multiset of T clauses and u variables by selecting, uniformly at random from W, a multiset of u clauses and then selecting, independently and uniformly at random, a variable from each of those clauses.

Definition 2. *Given a μ-gap E3-Sat(5) formula ϕ, a (T, u)*-block selected *from ϕ is a multiset of $(T+1)u$ clauses from ϕ.*

Definition 3. *Given a (T, u)-block W, a* random u-projection of W *is a multiset formed by first selecting u clauses at random from W and then replacing each of those clauses with a variable selected uniformly and independently at random from the respective clause.*

Having selected a (T, u)-block W uniformly at random and a random u-projection U from W, the verifier sends W to the first prover and U to the second prover. The verifier accepts if the assignments returned by the provers agree and satisfy both U and W.

Lemma 1. *The protocol described in this section has perfect completeness and soundness c_μ^u, where $c_\mu < 1$ is some constant depending on μ but not on u or the size of the instance.*

Proof. For a satisfiable formula, the verifier always accepts if the two provers answer according to the same satisfying assignment.

To prove that the verifier accepts unsatisfiable formulae with probability at most c_μ^u, we reduce from the protocol described in § 2.1. It is known [12] that this protocol has soundness c_μ^u, where $c_\mu < 1$ is some constant depending on μ but not on u or the size of the instance. Now suppose that there exists provers Q_1 and Q_2 for the protocol from this section such that the verifier accepts an unsatisfiable formula with probability $s > c_\mu^u$. Then the provers P_1 and P_2 from the protocol described in § 2.1 can use Q_1 and Q_2 to construct strategies that make the verifier in that protocol accept an unsatisfiable formula with probability s.

Given a multiset of clauses $(C_1, \dots, C_u)$, P_1 selects T more clauses uniformly at random. The thereby obtained (T, u)-block is sent to Q_1 and P_1 then returns the assignment to $(C_1, \dots, C_u)$ obtained from Q_1. Given a multiset of variables $(x_1, \dots, x_u)$, P_2 selects T more clauses uniformly at random. The thereby obtained u-projection is sent to Q_2 and P_2 then returns the assignment to $(x_1, \dots, x_u)$ obtained from Q_2.

The answers sent back by P_1 and P_2 make the verifier accept with probability s. But this is a contradiction since $s > c_\mu^u$. ∎

2.3 The long d-code and the Fourier transform

Our PCP construction is used to simulate the 2P1R game from § 2.2. The proof in our PCP should therefore contain answers to the queries for all possible choices of the verifier in the 2P1R game. Since we want to use our PCP construction to prove lower bounds for constraints over domain size d, we use a variant of the standard long code.

Definition 4. *Let V be a multiset of variables and clauses and denote by* SAT^V *the set of assignments to the variables in V and the clauses in V that satisfy all the clauses in V. The* long d-code of some $y \in \mathrm{SAT}^V$ *is a function* $A_{V,y}\colon \boldsymbol{Z}_d^{\mathrm{SAT}^V} \to \boldsymbol{Z}_d$ *defined by* $A_{V,y}(g) = g(y)$.

Definition 5. *A* standard written d-proof with parameters u and T *contains for each multiset U containing u variables and T clauses a string of length* $d^{2^u 7^{Tu}}$*, which we interpret as the table of a function* $A_U\colon \boldsymbol{Z}_d^{\mathrm{SAT}^U} \to G$. *It also contains for each multiset W of $(T+1)u$ clauses a string of length* $d^{7^{(T+1)u}}$ *which we interpret as the table of a function* $A_W\colon \boldsymbol{Z}_d^{\mathrm{SAT}^W} \to G$.

Definition 6. *A standard written d-proof with parameters u and T is a* correct proof for a formula ϕ *if there is an assignment x, satisfying ϕ, such* *that A_V is the long d-code of $x|_V$ for any multiset V containing u variables and Tu clauses and any multiset V of $(T+1)u$ clauses.*

The verifier in our PCP selects a random (T,u)-block W and then a random u-projection U of W. It then queries the tables A_U and A_W in the standard written d-proof at cleverly chosen positions. The analysis of the acceptance predicate of the verifier needs certain facts regarding the Fourier transform of the long d-code.

The Fourier series of a function from an Abelian group to $\boldsymbol{C}$ is a linear combination of the characters of that group, i.e., homomorphisms from the group to $\boldsymbol{C}$. For a general treatment of this theory, we refer the reader to Terras's book [15]. Since we only work with powers of $\boldsymbol{Z}_d$, we use the following conventions to simplify the framework: The group $\boldsymbol{Z}_d$ is represented by the powers of $\omega = e^{2\pi i/d}$ with multiplication as the group operator and $\hat{\boldsymbol{Z}}_d$, the characters of that group, is represented by the integers $\{0,1,2,\dots,d-1\}$ with addition modulo d as the group operator. For $g \in \boldsymbol{Z}_d$ and $\gamma \in \hat{\boldsymbol{Z}}_d$ we denote the action of γ on g by g^γ and this should be interpreted as normal exponentiation. Similarly, for functions A taking values in $\boldsymbol{Z}_d$ we write A^γ for the function $x \mapsto (A(x))^\gamma$.

Lemma 2. *Let $\gamma \in \hat{\boldsymbol{Z}}_d$ be arbitrary. Then*

$$\frac{1}{d}\sum_{g\in\boldsymbol{Z}_d} g^\gamma = \begin{cases} 1 & \textit{if } \gamma = 0,\\ 0 & \textit{otherwise.}\end{cases}$$

Lemma 3. *Let $g \in \boldsymbol{Z}_d$ be arbitrary. Then*

$$\frac{1}{d}\sum_{\gamma=0}^{d-1} g^\gamma = \begin{cases} 1 & \textit{if } g = 1,\\ 0 & \textit{otherwise.}\end{cases}$$

The spaces of functions from SAT^U to $\boldsymbol{Z}_d$ and from SAT^W to $\boldsymbol{Z}_d$ can be viewed as powers of $\boldsymbol{Z}_d$ since we can identify a function with the table of its values. We always use the shorthands

$$F = \boldsymbol{Z}_d^{|\,\mathrm{SAT}^U\,|}, \qquad H = \boldsymbol{Z}_d^{|\,\mathrm{SAT}^W\,|}$$

for given U and W. The dual groups $\hat{F}$ and $\hat{H}$ are then represented by vectors that contain elements of $\hat{\boldsymbol{Z}}_d$ and are indexed by elements in SAT^U and SAT^W, respectively, and given $f \in F$, $h \in H$, $\alpha \in \hat{F}$, and $\beta \in \hat{H}$, we define

$$f^\alpha = \prod_{x\in\mathrm{SAT}^U} (f(x))^{\alpha(x)}, \qquad h^\beta = \prod_{y\in\mathrm{SAT}^W} (h(y))^{\beta(y)}.$$

We can then expand arbitrary functions $A\colon F \to \boldsymbol{Z}_d$ and $B\colon H \to \boldsymbol{Z}_d$ in their respective Fourier series:

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

where the Fourier coefficients are computed as

$$\hat{A}_\alpha = \frac{1}{|F|}\sum_{f\in F} A(f)f^{-\alpha}, \qquad \hat{B}_\beta = \frac{1}{|H|}\sum_{h\in H} B(h)h^{-\beta}.$$

The Fourier coefficients satisfy Plancherel's equality:

$$\sum_{\alpha\in\hat{F}} |\hat{A}_\alpha|^2 = \frac{1}{|F|}\sum_{f\in F}|A(f)|^2 = 1 = \frac{1}{|H|}\sum_{h\in H}|B(h)|^2 = \sum_{\beta\in\hat{H}}|\hat{B}_\beta|^2.$$

2.3.1 Folding

As in many previous PCP constructions, we require the tables in the PCP to be *folded*. The lemmas in this section were originally proved by Håstad [6], we state them here for easy reference.

Definition 7. *A function A from $\boldsymbol{Z}_d^{\mathrm{SAT}^V}$ to $\boldsymbol{Z}_d$ is* folded *if $A(gf) = gA(f)$ for all $g \in \boldsymbol{Z}_d$ and it is γ-homogeneous for $\gamma \in \hat{\boldsymbol{Z}}_d$ if $A(gf) = g^\gamma A(f)$.*

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

Lemma 4. *If $A\colon \boldsymbol{Z}_d^{\mathrm{SAT}^V} \to \boldsymbol{Z}_d$ is folded, A^γ is γ-homogeneous for every $\gamma \in \hat{\boldsymbol{Z}}_d$.*

Lemma 5. *Suppose that the function $A\colon \boldsymbol{Z}_d^{\mathrm{SAT}^V} \to \boldsymbol{Z}_d$ is γ-homogeneous for some $\gamma \in \hat{\boldsymbol{Z}}_d$ and let $\hat{A}_\alpha$ be the Fourier coefficients of that function at α. Then $\hat{A}_\alpha = 0$ unless $\gamma = \sum_{x\in\mathrm{SAT}^V}\alpha(x)$.*

Now let W be a (T, u)-block and U be a u-projection of W. This u-projection defines a projection function $\pi\colon \mathrm{SAT}^W \to \mathrm{SAT}^U$. We define the projection function $\pi_d\colon \hat{H} \to \hat{F}$ as follows: $\alpha = \pi_d(\beta)$ if $\alpha(x) = \sum_{y\in\pi^{-1}(x)}\beta(y)$ for all $x \in \mathrm{SAT}^U$.

Lemma 6. *Let U, W, F, H and π be as above. Then, for any $\beta \in \hat{H}$, $(f\circ\pi)^\beta = f^{\pi_d(\beta)}$.*

For a given $\beta \in \hat{H}$ we define $|\beta| = |\{y \in \hat{H} : \beta(y) \neq 0\}|$. We also need the projection $\pi(\beta) = \{x \in F : \exists y \in \pi^{-1}(x)[\beta(y) \neq 0]\}$. Note that $\pi(\beta)$ is a subset of F while π_d is a function from $\hat{H}$ to $\hat{F}$.

2.3.2 Projections

As we mentioned in § 2.2, the 2P1R for μ-gap E3-Sat(5) was modified to handle certain difficulties in the analysis of the PCP verifier we present in this paper. In this section, we formalize the properties needed in the analysis. The facts from this section are straightforward generalizations of the corresponding definitions and lemmas from Khot's paper [10].

Lemma 7. *Let W be a (T, u)-block and consider two different assignments $y, y' \in \mathrm{SAT}^W$. Then*

$$\Pr[\pi(y) \neq \pi(y')] > 1 - \frac{1}{T}$$

where the probability is over the selection of a random u-projection of W and π is the projection induced by the u-projection of W.

Proof. Since $y \neq y'$, there is at least one variable that is assigned different values by y and y'. If this clause is present in the u-projection—which happens with probability $T/(T+1) \geq 1 - 1/T$—the projections are certainly different. ∎

Corollary 1. *Let W be a (T, u)-block and consider a $\beta \colon \mathrm{SAT}^W \to \hat{\boldsymbol{Z}}_d$. Then, with probability at least $1 - |\beta|/T$ over the choice of a random u-projection U of W, it holds for an arbitrary $y \in \mathrm{SAT}^W$ such that $\beta(y) \neq 0$ that there is no other y' such that $\beta(y') \neq 0$ and $y|_U = y'|_U$.*

Proof. Take an arbitrary $y \in \mathrm{SAT}^W$. By Lemma 7, for any $y' \in \mathrm{SAT}^W \setminus \{y\}$,

$$\Pr[\pi(y) = \pi(y')] < \frac{1}{T}$$

where the probability is over the selection of a random u-projection of W and π is the projection induced by the u-projection of W. By the union bound,

$$\Pr\Bigg[\bigcup_{\substack{y' \in \mathrm{SAT}^W \setminus \{y\} \\ \beta(y) \neq 0}} \{\pi(y) = \pi(y')\}\Bigg] < \frac{|\beta|}{T}$$

where the probability is over the same probability space. ∎

Lemma 8. *Let W be a (T, u)-block and β be a function from SAT^W to $\hat{\boldsymbol{Z}}_d$. Then*

$$\mathrm{E}\left[\frac{1}{|\pi(\beta)|}\right] \leq \frac{1}{|\beta|} + \frac{1}{T}$$

where the probability is over the selection of a random u-projection of W and π is the projection induced by the u-projection of W.

Proof. Since every $y \in \mathrm{SAT}^W$ projects down to at most one $x \in \mathrm{SAT}^U$,

$$|\beta| = \sum_{x \in \pi(\beta)} |\{y \in \mathrm{SAT}^W : \pi(y) = x \wedge \beta(y) \neq 0\}|.$$

Now apply the Cauchy-Schwartz inequality to the above equation. Then,

$$|\beta|^2 = \Big(\sum_{x \in \pi(\beta)} 1 \cdot |\{y \in \mathrm{SAT}^W : \pi(y) = x \wedge \beta(y) \neq 0\}|\Big)$$

$$\begin{aligned}
&\leq \Bigl(\sum_{x\in\pi(\beta)} 1^2\Bigr)\Bigl(\sum_{x\in\pi(\beta)} |\{y \in \mathrm{SAT}^W : \pi(y) = x \wedge \beta(y) \neq 0\}|^2\Bigr) \\
&= |\pi(\beta)| \sum_{x\in\pi(\beta)} |\{y \in \mathrm{SAT}^W : \pi(y) = x \wedge \beta(y) \neq 0\}|^2 \\
&= |\pi(\beta)| N_\pi(\beta),
\end{aligned}$$

where $N_\pi(\beta)$ is the numbers of pairs $(y_1, y_2) \in \mathrm{SAT}^W \times \mathrm{SAT}^W$ such that $\pi(y_1) = \pi(y_2)$, $\beta(y_1) \neq 0$ and $\beta(y_2) \neq 0$. Hence,

$$\frac{1}{|\pi(\beta)|} \leq \frac{N_\pi(\beta)}{|\beta|^2}.$$

Now introduce for each pair $(y_1, y_2) \in \mathrm{SAT}^W \times \mathrm{SAT}^W$ an indicator random variable $I_{(y_1,y_2)}$ that is one if $\pi(y_1) = \pi(y_2)$, $\beta(y_1) \neq 0$ and $\beta(y_2) \neq 0$. Then

$$\begin{aligned}
\mathrm{E}\Bigl[\frac{1}{|\pi(\beta)|}\Bigr] &= \frac{1}{|\beta|^2} \sum_{\substack{y_1\in\mathrm{SAT}^W\\ \beta(y_1)\neq 0}} \sum_{\substack{y_2\in\mathrm{SAT}^W\\ \beta(y_2)\neq 0}} \mathrm{E}[I_{(y_1,y_2)}] \\
&= \frac{1}{|\beta|^2}\Bigl(|\beta| + \sum_{\substack{y_1\in\mathrm{SAT}^W\\ \beta(y_1)\neq 0}} \sum_{\substack{y_2\in\mathrm{SAT}^W\\ \beta(y_2)\neq 0\\ y_1\neq y_2}} \mathrm{E}[I_{(y_1,y_2)}]\Bigr) \\
&= \frac{1}{|\beta|^2}\Bigl(|\beta| + \sum_{\substack{y_1\in\mathrm{SAT}^W\\ \beta(y_1)\neq 0}} \sum_{\substack{y_2\in\mathrm{SAT}^W\\ \beta(y_2)\neq 0\\ y_1\neq y_2}} \mathrm{E}[I_{(y_1,y_2)}]\Bigr) \\
&\leq \frac{1}{|\beta|^2}\Bigl(|\beta| + \sum_{\substack{y_1\in\mathrm{SAT}^W\\ \beta(y_1)\neq 0}} \sum_{\substack{y_2\in\mathrm{SAT}^W\\ \beta(y_2)\neq 0\\ y_1\neq y_2}} \frac{1}{T}\Bigr) \\
&= \frac{1}{|\beta|} + \frac{|\beta|-1}{|\beta|T} \leq \frac{1}{|\beta|} + \frac{1}{T},
\end{aligned}$$

where the first inequality follows from Lemma 7. ∎

Corollary 2. *Let W be a (T,u)-block and β be a function from SAT^W to $\hat{\boldsymbol{Z}}_d$. Then $|\pi(\beta)| \geq \delta \min\{|\beta|, T\}$ with probability $1-2\delta$, where the probability is over the selection of a random u-projection of W and π is the projection induced by the u-projection of W.*

Proof. Apply Markov's inequality to the conclusion of Lemma 8. ∎

2.4 Constructing strategies for the provers

The verifier in our PCP expects as proof encodings of the answers of the two provers in the balanced version of the 2P1R game from § 2.2. Specifically, the proof in our PCP consists of purported Long d-Codes of the assignments to the variables in U and W for each possible choice (U, W) of the 2P1R verifier. To prove that our PCP has a certain soundness, we use the Fourier expansion of the purported long codes to extract probabilistic strategies for the provers P_1 and P_2. In particular, we 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 PCP verifier accepts with large probability.

Lemma 9. *Suppose that $B\colon H \to \boldsymbol{Z}_d$ is γ_2-homogeneous for some $\gamma_2 \neq 0$ and known to the first prover in the 2P1R game once it has received a (T, u)-block W, that $A\colon F \to \boldsymbol{Z}_d$ is γ_1-homogeneous for some $\gamma_1 \neq 0$ and known to the first prover in the 2P1R game once it has received a u-projection U of W. Let $\hat{A}_\alpha$ be the Fourier coefficients of A and $\hat{B}_\beta$ be the Fourier coefficients of B and let $\alpha \preceq \beta$ mean that $\alpha(x) \neq 0 \implies x \in \pi(\beta)$. If*

$$\mathrm{E}\Big[\sum_{\beta \in \hat{H}} \sum_{\substack{\alpha \in \hat{F} \\ \alpha \preceq \beta}} |\hat{A}_\alpha|^2 |\hat{B}_\beta|^2 |\beta|^{-1}\Big] \geq \eta,$$

where the probability is over the selection of a (T, u)-block W uniformly at random, and a random u-projection U of W, there exists a strategy for the provers in the balanced version of the 2P1R game from § 2.2 that makes the verifier in that game accepts with probability at least η.

Proof. The strategy is as follows: On receiving a multiset 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 $\beta(y) \neq 0$ uniformly—by Lemma 5 such a y always exists—this y is returned to the verifier.

On receiving a multiset U of clauses and 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 $\alpha(x) \neq 0$ uniformly—by Lemma 5 such an x always exists—this x is returned to the verifier.

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

$$\mathrm{E}\Big[\sum_{\beta \in \hat{H}} \sum_{\substack{\alpha \in \hat{F} \\ \alpha \preceq \beta}} |\hat{A}_\alpha|^2 |\hat{B}_\beta|^2 |\beta|^{-1}\Big] \geq \eta. \qquad \blacksquare$$

3 Our PCP construction

The PCP is shown in Fig. 1. The intuition behind the protocol is to take a linearity test—which we know has soundness at most $d^{-1} + \varepsilon$ for any

```
The proof is a standard written $d$-proof with parameter $u$ and $T$:
The verifier acts as follows:
1. Let $W$ be a random $(T, u)$-block selected from $\Phi$.
2. Let $U$ be a random $u$-projection of $W$.
3. Let $\pi\colon \mathrm{SAT}^W \to \mathrm{SAT}^U$ be the function that creates
   an assignment in $\mathrm{SAT}^U$ from an assignment in $\mathrm{SAT}^W$.
4. Let $F = \boldsymbol{Z}_d^{|\mathrm{SAT}^U|}$ and $H = \boldsymbol{Z}_d^{|\mathrm{SAT}^W|}$.
5. Select $f \in F$ and $h \in H$ uniformly at random.
6. Select $e_f \in H$ by selecting, independently for every $y \in \mathrm{SAT}^W$,
   $e_f(y)$ such that:
   – $f(\pi(y)) \neq 1 \implies e_f(y) = 1$;
   – $f(\pi(y)) = 1 \implies (\Pr[e_f(y) = 1] = 1 - \delta) \wedge (\Pr[e_f(y) = \omega] = \delta)$.
7. Accept if $A_U(f) A_W(h) A_W(h^{-1}(f \circ \pi)^{-1} e_f) = 1$
   or if $A_U(f) = 1$ and $A_W(h) A_W(h^{-1}(f \circ \pi)^{-1} e_f) = \omega$;
   Reject otherwise.
```

Figure 1. The above PCP is parameterized by the positive integers d, u and T and the positive real δ and tests if a μ-gap E3-Sat(5) formula Φ is satisfiable by querying three positions in a Standard Written d-proof with parameter u. With suitable choices of the parameters u, T and δ as functions of ε and d, the above PCP has perfect completeness and soundness $d^{-1} + d^{-2} + \varepsilon$ for any constant $\varepsilon > 0$.

constant $\varepsilon > 0$ and near-perfect completeness—and modify it a slightly in such a way that we get at test with perfect completeness which is still close enough to a linearity test to have soundness $d^{-1} + d^{-2} + \varepsilon$ for any constant $\varepsilon > 0$. The following lemma is immediate:

Lemma 10. *The PCP in Fig. 1 has perfect completeness.*

To analyze the soundness of the protocol we follow the standard approach. We estimate the acceptance probability of the verifier using Fourier analysis and prove that if the acceptance probability is large there must exist pairs of correlated coefficients whose product is large. We then use these products to devise a strategy for the provers in the balanced version of the 2P1R game from § 2.2.

By Lemma 3, the expression $d^{-1} \sum_{\gamma=0}^{d-1} g^\gamma$ is an indicator for the event $g = 1$ when g assumes values in $\boldsymbol{Z}_d$. Hence, the test accepts with probability $\mathrm{E}[I_1 + I_2]$ where

$$I_1 = \frac{1}{d} \sum_{\gamma=0}^{d-1} \Bigl(A_U(f) A_W(h) A_W(h^{-1}(f \circ \pi)^{-1} e_f) \Bigr)^\gamma,$$

$$I_2 = \biggl(\frac{1}{d} \sum_{\gamma=0}^{d-1} \Bigl(A_U(f) \Bigr)^\gamma \biggr) \biggl(\frac{1}{d} \sum_{\gamma=0}^{d-1} \Bigl(A_W(h) A_W(h^{-1}(f \circ \pi)^{-1} e_f) \omega^{-1} \Bigr)^\gamma \biggr).$$

This expression can be rewritten as $d^{-1} + d^{-2} + \mathrm{E}[L + Q + C_1 + C_2]$ where

$$L = \frac{1}{d}\sum_{\gamma=1}^{d-1}(A_U(f))^\gamma, \tag{1}$$

$$Q = \frac{1}{d^2}\sum_{\gamma=1}^{d-1}\Big(A_W(h)A_W(h^{-1}(f \circ \pi)^{-1}e_f)\omega^{-1}\Big)^\gamma, \tag{2}$$

$$C_1 = \frac{1}{d}\sum_{\gamma=1}^{d-1}\Big(A_U(f)A_W(h)A_W(h^{-1}(f \circ \pi)^{-1}e_f)\Big)^\gamma, \tag{3}$$

$$C_2 = \frac{1}{d^2}\sum_{\gamma_1=1}^{d-1}\sum_{\gamma_2=1}^{d-1}\Big(A_U(f)\Big)^{\gamma_1}\Big(A_W(h)A_W(h^{-1}(f \circ \pi)^{-1}e_f)\omega^{-1}\Big)^{\gamma_2}. \tag{4}$$

We now prove that L and Q have small magnitude and then prove that most of the terms in C_1 and C_2 also have small magnitude. Finally, we prove that the remaining terms can be used to extract a strategy for the provers in the balanced version of the 2P1R game from § 2.2.

Lemma 11. $\mathrm{E}[L] = 0$

Proof. The tables in the proof are folded. ∎

Let us now analyze $\mathrm{E}[Q]$:

$$\mathrm{E}[Q] = \mathrm{E}\Big[\frac{1}{d^2}\sum_{\gamma=1}^{d-1}\omega^{-\gamma}\,\mathrm{E}\Big[(A_W(h)A_W(h^{-1}(f \circ \pi)^{-1}e_f))^\gamma \;\Big|\; U, W\Big]\Big].$$

Our aim is to show that the inner expectation

$$Q_\gamma = \mathrm{E}\Big[(A_W(h)A_W(h^{-1}(f \circ \pi)^{-1}e_f))^\gamma \;\Big|\; U, W\Big] \tag{5}$$

always has small magnitude, regardless of γ. To this end, we expand it in a Fourier series:

$$\begin{aligned} Q_\gamma &= \mathrm{E}\Big[\sum_{\beta_1 \in \hat{H}}\sum_{\beta_2 \in \hat{H}}\hat{B}_{\beta_1}\hat{B}_{\beta_1}h^{\beta_1}(h^{-1}(f \circ \pi)^{-1}e_f)^{\beta_2} \;\Big|\; U, W\Big] \\ &= \sum_{\beta_1 \in \hat{H}}\sum_{\beta_2 \in \hat{H}}\hat{B}_{\beta_1}\hat{B}_{\beta_1}\,\mathrm{E}[e_f^{\beta_2}f^{-\pi_d(\beta_2)} \mid U, W]\,\mathrm{E}[h^{\beta_1-\beta_2} \mid U, W], \end{aligned}$$

where $\hat{B}_\beta$ is the Fourier coefficient of A_W^γ at β. Since $\mathrm{E}[h^{\beta_1-\beta_2} \mid U, W]$ unless $\beta_1 = \beta_2$ by Lemma 2, the above expression can be simplified to

$$Q_\gamma = \sum_{\beta \in \hat{H}}\hat{B}_\beta^2\,\mathrm{E}[e_f^\beta f^{-\pi_d(\beta)} \mid U, W].$$

Let us now compute

$$\mathrm{E}[e_f^\beta f^{-\pi_d(\beta)} \mid U, W] = \mathrm{E}\Big[\prod_y (e_f(y))^{\beta(y)}(f(y|_U))^{-\beta(y)} \;\Big|\; U, W\Big]$$

$$= \prod_{x\in\pi(\beta)} \mathrm{E}\Big[\prod_{y\in\pi^{-1}(x)} (e_f(y)f^{-1}(x))^{\beta(y)}\Big].$$

Consider the factor corresponding to a fixed x in the above product. Since $e_f(y)$ depends on $f(x)$ we have to condition on $f(x)$:

$$\begin{aligned}
&\mathrm{E}\Big[\prod_{y\in\pi^{-1}(x)} (e_f(y)f^{-1}(x))^{\beta(y)}\Big] \\
&\quad = \frac{1}{d}\sum_{t=0}^{d-1} \mathrm{E}\Big[\prod_{y\in\pi^{-1}(x)} (e_f(y)f^{-1}(x))^{\beta(y)} \;\Big|\; f(x)=\omega^t\Big].
\end{aligned}$$

If $f(x) \neq 1$, $e_f(y)$ is always 1, otherwise $e_f(y)$ is selected according to a biased distribution on $\{1, \omega\}$. Hence

$$\begin{aligned}
&E\Big[\prod_{y\in\pi^{-1}(x)} (e_f(y)f^{-1}(x))^{\beta(y)}\Big] \\
&\quad = \frac{1}{d}\prod_{y\in\pi^{-1}(x)} (1-\delta+\omega^{\beta(y)}\delta) + \frac{1}{d}\sum_{t=1}^{d-1}\prod_{y\in\pi^{-1}(x)} \omega^{-t\beta(y)}
\end{aligned}$$

Recall the notation $\pi_d(\beta)(x) = \sum_{x\in\pi^{-1}(y)} \beta(y) \bmod d$. If $\pi_d(\beta)(x) \neq 0$, the last sum above is -1; otherwise it is $d-1$. If we define

$$a_{\beta,x} = \frac{1}{d}\Big(\prod_{y\in\pi^{-1}(x)} (1-\delta+\omega^{\beta(y)}\delta) + d - 1\Big), \tag{6}$$

$$b_{\beta,x} = \frac{1}{d}\Big(\prod_{y\in\pi^{-1}(x)} (1-\delta+\omega^{\beta(y)}\delta) - 1\Big), \tag{7}$$

we can write

$$\begin{aligned}
|Q_\gamma|^2 &= \Big|\sum_{\beta\in\hat{H}} \hat{B}_\beta^2 \prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)=0}} a_{\beta,x} \prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)\neq 0}} b_{\beta,x}\Big|^2 \\
&\leq \Big(\sum_{\beta\in\hat{H}} |\hat{B}_\beta|^2\Big)\Big(\sum_{\beta\in\hat{H}} |\hat{B}_\beta|^2 \prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)=0}} |a_{\beta,x}|^2 \prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)\neq 0}} |b_{\beta,x}|^2\Big) \\
&= \sum_{\beta\in\hat{H}} |\hat{B}_\beta|^2 \prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)=0}} |a_{\beta,x}|^2 \prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)\neq 0}} |b_{\beta,x}|^2,
\end{aligned}$$

where the inequality is the Cauchy-Schwartz inequality and the last equality follows from Plancherel's equality.

Lemma 12. *For any integer $b \neq 0 \bmod d$ and any real $\delta \in [0,1]$, $|1-\delta+\omega^b\delta|^2 \leq 1-\delta d^{-2}$ and, consequently, $|1-\delta+\omega^b\delta| \leq 1-\delta d^{-2}/2$.*

Proof. Since $\omega = e^{2\pi i/d}$,

$$\begin{aligned} |1-\delta+\omega^b\delta|^2 &\le |1-\delta+\omega\delta|^2 \\ &= (1-\delta(1-\cos(2\pi/d)))^2 + \delta^2 \sin^2(2\pi/d) \\ &= 1 - 2(\delta-\delta^2)(1-\cos(2\pi/d)). \end{aligned}$$

For $d \in \{2,3,4\}$ this expression is at most $1-2\delta(1-\delta) \le 1-\delta$ while for larger d it is at most

$$1 - 2(\delta-\delta^2)(2\pi^2 d^{-2} - 2\pi^4 d^{-4}/3) \le 1 - 2\pi^2 d^{-2}\delta(1-\delta) \le 1 - \delta d^{-2}.$$

The conclusion of the lemma now follows since $(1-\eta)^{1/2} \le e^{-\eta/2} \le 1-\eta/2$ for all $\eta \in [0,1]$. ∎

Using separate arguments for terms where $|\beta| = |\{y \in \hat{H} : \beta(y) \ne 0\}|$ is large and terms where $|\beta|$ is small, we now bound the factor multiplying $|\hat{B}_\beta|^2$ by 2δ. This is the first place where we use the new construction due to Khot [10] mentioned in § 2.2: For the case when $|\beta|$ is small, we use Corollary 1 and when $|\beta|$ is large, we use Corollary 2.

Lemma 13. *Let Q_γ be defined as in equation (5) and let $\gamma \in \hat{\boldsymbol{Z}}_d \setminus \{0\}$ and $\delta \in (0,1)$ be arbitrary. Then* $\mathrm{E}[|Q_\gamma|^2] \le 2\delta$ *provided that* $T \ge \delta^{-2} d^3 \ln \delta^{-1}$.

Proof. By the above reasoning, we have established that

$$\mathrm{E}[|Q_\gamma|^2] \le \mathrm{E}\Bigg[\sum_{\beta\in\hat{H}} |\hat{B}_\beta|^2 \,\mathrm{E}\Big[\prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)=0}} |a_{\beta,x}|^2 \prod_{\substack{x\in\pi(\beta)\\ \pi_d(\beta)(x)\ne 0}} |b_{\beta,x}|^2 \,\Big|\, W\Big]\Bigg] \tag{8}$$

We now bound the inner expectation above separately for every β and then use Plancherel's equality to bound the entire sum.

Suppose that $|\beta| < \delta T$ and $\hat{B}_\beta \ne 0$. It follows by Lemma 5, that $\hat{B}_\beta = 0$ if $\beta(y) = 0$ for all y; hence there exists some y such that $\beta(y) \ne 0$. By Corollary 1, with probability at least $1 - |\beta|/T \ge 1-\delta$ over the choice of U there is no other y' such that $y|_U = y'|_U$. Given that this event happens, the factor corresponding to $x = y|_U$ in the product has magnitude

$$\frac{|1-\delta+\omega^{\beta(y)}\delta - 1|}{d^2} = \frac{\delta|1-\omega^{\beta(y)}|}{d^2} \le \frac{2\delta}{d^2}.$$

Since the other factors have at most unit magnitude, the magnitude of the inner expectation is bounded by $2\delta d^{-2}$ when $|\beta| < \delta T$. Given that $|\pi^{-1}(x)| \ne 1$, which happens with probability at most δ, the entire product has at most unit magnitude. Therefore, the factor multiplying $|\hat{B}_\beta|^2$ in (8) is at most $(1-\delta)\cdot 2\delta d^{-2} + \delta \cdot 1 \le 3\delta/2$ when $|\beta| \le \delta T$.

Next, we consider the case when $|\beta| \ge \delta T$. Since $|a_{\beta,x}| \ge |b_{\beta,x}|$, we can upper bound the product inside the inner expectation in (8) by $|a_{\beta,x}|^{2|\pi(\beta)|} \le$

$(1-\delta d^{-3}/2)^{2|\pi(\beta)|}$. By Corollary 2, $|\pi(\beta)| \geq \delta T = \delta^{-1} d^3 \ln \delta^{-1}$ with probability at least $1 - 2\delta$ over the choice of U. When this holds,

$$(1-\delta d^{-3}/2)^{2|\pi(\beta)|} \leq e^{-\delta d^{-3}|\pi(\beta)|} = \delta.$$

When $|\pi(\beta)|$ fails to be large, which happens with probability at most 2δ, the product has at most unit magnitude. Therefore, the factor multiplying $|\hat{B}_\beta|^2$ in (8) is at most $(1-2\delta)\cdot\delta + \delta \cdot 1 \leq 2\delta$ when $|\beta| \geq \delta T$.

To conclude, the term corresponding to an arbitrary β in (8) has magnitude at most $2\delta|\hat{B}_\beta|^2$, hence

$$\mathrm{E}\Big[|Q_\gamma|^2\Big] \leq 2\delta\,\mathrm{E}\Big[\sum_{\beta\in\hat{H}} |\hat{B}_\beta|^2\Big] = 2\delta. \qquad \blacksquare$$

Corollary 3. *Let Q be defined as in equation (2) and $\delta \in (0,1)$ be arbitrary. Then $|\mathrm{E}[Q]|^2 \leq \delta$ provided that $T \geq \delta^{-2} d^3 \ln \delta^{-1}$.*

Proof. Since $\mathrm{E}[Q] = \mathrm{E}[d^{-2}\sum_{\gamma=1}^{d-1}\omega^{-\gamma}Q_\gamma]$, and the function $z \mapsto |z|^2$ is convex, it follows from Jensen's inequality that

$$\begin{aligned}|\mathrm{E}[Q]|^2 &= \frac{(d-1)^2}{d^4}\left|\frac{1}{d-1}\sum_{\gamma=1}^{d-1}\omega^{-\gamma}\,\mathrm{E}[Q_\gamma]\right|^2 \leq \frac{d-1}{d^4}\sum_{\gamma=1}^{d-1}|\mathrm{E}[Q_\gamma]|^2 \\ &\leq \frac{d-1}{d^4}\sum_{\gamma=1}^{d-1}\mathrm{E}[|Q_\gamma|^2] \leq \frac{2\delta(d-1)^2}{d^4},\end{aligned}$$

where the last inequality follows from Lemma 13 and the first two inequalities follow from Jensen's inequality. $\blacksquare$

Let us now look at the terms of degree three, i.e., the expectations (3) and (4). They can be written as

$$\mathrm{E}[C_1] = \mathrm{E}\Big[\frac{1}{d}\sum_{\gamma=1}^{d-1} C_{\gamma,\gamma}\Big], \qquad \mathrm{E}[C_2] = \mathrm{E}\Big[\frac{1}{d^2}\sum_{\gamma_1=1}^{d-1}\sum_{\gamma_2=1}^{d-1}\omega^{-\gamma_2}C_{\gamma_1,\gamma_2}\Big]$$

where

$$C_{\gamma_1,\gamma_2} = \mathrm{E}\Big[(A_U(f))^{\gamma_1}(A_W(h)A_W(h^{-1}(f\circ\pi)^{-1}e_f))^{\gamma_2} \;\Big|\; U, W\Big]. \tag{9}$$

To prove that $\mathrm{E}[C_1]$ and $\mathrm{E}[C_2]$ have small magnitude, it is enough to bound $|C_{\gamma_1,\gamma_2}|$ for arbitrary $\gamma_1,\gamma_2 \neq 0$. To this end, we expand C_{γ_1,γ_2} in a Fourier series:

$$C_{\gamma_1,\gamma_2} = \sum_{\alpha\in F}\sum_{\beta\in\hat{H}} \hat{A}_\alpha \hat{B}_\beta^2\,\mathrm{E}[e_f^\beta f^{\alpha-\pi_d(\beta)} \mid U, W]$$

where $\hat{A}_\alpha$ is the Fourier coefficient of $A_U^{\gamma_1}$ at α and $\hat{B}_\beta$ is the Fourier coefficient of $A_W^{\gamma_2}$ at β. Since e_f and f are dependent, we cannot simply discard all terms where $\alpha \neq \pi_d(\beta)$. However, we *can* use Lemma 2 to discard terms for which there exists an x such that $\alpha(x) \neq 0$ but $\beta(y) = 0$ for *all* $y \in \pi^{-1}(x)$. Recall that $\alpha \preceq \beta$ denotes precisely that this does *not* hold. Thus

$$C_{\gamma_1,\gamma_2} = \sum_{\beta \in \hat{H}} \sum_{\substack{\alpha \in F \\ \alpha \preceq \beta}} \hat{A}_\alpha \hat{B}_\beta^2 \operatorname{E}[e_f^\beta f^{\alpha - \pi_d(\beta)} \mid U, W].$$

Let us now compute

$$\operatorname{E}[e_f^\beta f^{\alpha - \pi_d(\beta)} \mid U, W] = \prod_{x \in \pi(\beta)} \operatorname{E}\Big[(f(x))^{\alpha(x)} \prod_{y \in \pi^{-1}(x)} (e_f(y) f^{-1}(x))^{\beta(y)}\Big].$$

We consider each factor separately. The factor corresponding to an arbitrary x can be written

$$\begin{aligned} &\operatorname{E}\Big[(f(x))^{\alpha(x)} \prod_{y \in \pi^{-1}(x)} (e_f(y) f^{-1}(x))^{\beta(y)}\Big] \\ &= \frac{1}{d} \prod_{y \in \pi^{-1}(x)} (1 - \delta + \omega^{\beta(y)}\delta) + \frac{1}{d} \sum_{t=1}^{d-1} \omega^{t\alpha(x)} \prod_{y \in \pi^{-1}(x)} \omega^{-t\beta(y)} \\ &= \frac{1}{d} \prod_{y \in \pi^{-1}(x)} (1 - \delta + \omega^{\beta(y)}\delta) + \frac{1}{d} \sum_{t=1}^{d-1} \omega^{t(\alpha(x) - \pi_d(\beta)(x))}, \end{aligned}$$

where $\pi_d(\beta)$ is defined as above. If $\alpha(x) \neq \pi_d(\beta)(x)$, the second term evaluates to $-d^{-1}$; otherwise it evaluates to $(d-1)/d$. Therefore,

$$C_{\gamma_1,\gamma_2} = \sum_{\beta \in \hat{H}} \sum_{\substack{\alpha \in F \\ \alpha \preceq \beta}} \hat{A}_\alpha \hat{B}_\beta^2 p(\alpha, \beta) \tag{10}$$

where

$$p(\alpha, \beta) = \prod_{\substack{x \in \pi(\beta) \\ \alpha(x) = \pi_d(\beta)(x)}} a_{\beta,x} \prod_{\substack{x \in \pi(\beta) \\ \alpha(x) \neq \pi_d(\beta)(x)}} b_{\beta,x} \tag{11}$$

and $a_{\beta,x}$ and $b_{\beta,x}$ are defined in equations (6) and (7) above. We now proceed to bound the terms corresponding to terms with $|\beta| > T$.

Lemma 14. *For any $\beta \in \hat{H}$,*

$$\sum_{\substack{\alpha \in F \\ \alpha \preceq \beta}} |p(\alpha, \beta)|^2 = \prod_{x \in \pi(\beta)} (|a_{\beta,x}|^2 + (d-1)|b_{\beta,x}|^2),$$

where $a_{\beta,x}$, $b_{\beta,x}$ and $p(\alpha, \beta)$ are defined in equations (6), (7) and (11).

Proof. There are $d^{|\pi(\beta)|}$ different α such that $\alpha \preceq \beta$. Of those, exactly one α, namely $\alpha = \pi_d(\beta)$ gives a product with only factors of type a. More generally, let $A \subseteq \pi(\beta)$ be the set of x such that $\alpha(x) = \pi_d(\beta)(x)$. Then there are $(d-1)^{|\pi(\beta)\setminus A|}$ different α such that α and $\pi_d(\beta)$ agree exactly on A and for all those α,

$$|p(\alpha,\beta)|^2 = \prod_{x\in A} |a_{\beta,x}|^2 \prod_{x\in\pi(\beta)\setminus A} |b_{\beta,x}|^2.$$

Hence,

$$\begin{aligned}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |p(\alpha,\beta)|^2 &= \sum_{A\subseteq\pi(\beta)} (d-1)^{|\pi(\beta)\setminus A|} \prod_{x\in A} |a_{\beta,x}|^2 \prod_{x\in\pi(\beta)\setminus A} |b_{\beta,x}|^2 \\ &= \prod_{x\in\pi(\beta)} (|a_{\beta,x}|^2 + (d-1)|b_{\beta,x}|^2).\end{aligned}$$

■

Lemma 15. *For all integers $d \geq 2$,*

$$\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |p(\alpha,\beta)|^2 \leq \left(1 - \frac{\delta}{d^3}\right)^{|\pi(\beta)|}$$

Proof. By Lemma 14,

$$\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |p(\alpha,\beta)|^2 \leq \prod_{x\in\pi(\beta)} (|a_{\beta,x}|^2 + (d-1)|b_{\beta,x}|^2).$$

We now bound each factor in the above product using the observation that

$$\begin{aligned}|a_{\beta,x}|^2 + (d-1)|b_{\beta,x}|^2 &= (a_{\beta,x} + (d-1)b_{\beta,x})(a_{\beta,x} + (d-1)b_{\beta,x})^* \\ &\quad - (d-1)(a_{\beta,x}b^*_{\beta,x} + a^*_{\beta,x}b_{\beta,x}) - (d-1)(d-2)|b_{\beta,x}|^2, \qquad (12)\end{aligned}$$

where z^* denotes the complex conjugate of z. Using the definitions of $a_{\beta,x}$ and $b_{\beta,x}$ together with the shorthand

$$c_{\beta,x} = \prod_{y\in\pi^{-1}(x)} (1 - \delta + \omega^{\beta(y)}\delta)$$

and the notation $\Re\{z\}$ for the real part of z, we now expand each of the expressions in the right-hand side of (12):

$$\begin{aligned}&a_{\beta,x} + (d-1)b_{\beta,x} = c_{\beta,x},\\ &(a_{\beta,x} + (d-1)b_{\beta,x})(a_{\beta,x} + (d-1)b_{\beta,x})^* = |c_{\beta,x}|^2,\\ &a_{\beta,x}b^*_{\beta,x} = d^{-2}(|c_{\beta,x}|^2 + (d-1)c_{\beta,x} - c^*_{\beta,x} - (d-1)),\\ &a_{\beta,x}b^*_{\beta,x} + a^*_{\beta,x}b_{\beta,x} = 2d^{-2}(|c_{\beta,x}|^2 + (d-2)\Re\{c_{\beta,x}\} - (d-1)),\end{aligned}$$

$$|b_{\beta,x}|^2 = d^{-2}(|c_{\beta,x}|^2 + 2\Re\{c_{\beta,x}\} + 1),$$

Hence $a_{\beta,x}b^*_{\beta,x} + a^*_{\beta,x}b_{\beta,x} + (d-2)|b_{\beta,x}|^2 = d^{-1}|c_{\beta,x}|^2 - d^{-1}$ and

$$|a_{\beta,x}|^2 + (d-1)|b_{\beta,x}|^2 = |c_{\beta,x}|^2 - \frac{d-1}{d}\Big(|c_{\beta,x}|^2 - 1\Big) = 1 - \frac{1 - |c_{\beta,x}|^2}{d}.$$

By Lemma 12,

$$|c_{\beta,x}|^2 = \prod_{y\in\pi^{-1}(x)} |1 - \delta + \omega^{\beta(y)}\delta|^2 \le (1 - \delta d^{-2})^{|\pi^{-1}(x)|} \le 1 - \delta d^{-2}$$

and we can therefore bound each $|a_{\beta,x}|^2 + (d-1)|b_{\beta,x}|^2$ from above by $1 - \delta d^{-3}$. Since there are $|\pi(\beta)|$ factors in the product, the proof is complete. ■

As before, the terms in the Fourier series are bounded in separate ways depending on the size of $|\beta|$. We first prove that the sum of all terms corresponding to large $|\beta|$ must have a magnitude that is upper bounded by an expression linear in $\delta^{1/2}$. Intuitively, this follows since $p(\alpha,\beta)$ has small magnitude with very high probability when $|\beta|$ is large. Here we, again, use Corollary 2 from § 2.2. We then prove that the terms corresponding to small $|\beta|$ must also have a magnitude that is upper bounded by an expression linear in $\delta^{1/2}$, or else there exists a strategy for the provers in the 2P1R game from § 2.2 that makes the verifier accept with probability larger than c^u_μ.

Lemma 16. *Let C_{γ_1,γ_2} be defined as in equation (9) and let $\gamma_1,\gamma_2 \in \hat{\boldsymbol{Z}}_d\setminus\{0\}$ be arbitrary. Then, provided that $T \ge \delta^{-2}d^3\ln\delta^{-1}$,*

$$|C_{\gamma_1,\gamma_2}|^2 \le 2\left|\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}} \hat{A}_\alpha\hat{B}^2_\beta p(\alpha,\beta)\Big]\right|^2 + 6\delta,$$

where $\hat{A}_\alpha$ is the Fourier coefficient of $A^{\gamma_1}_U$ at α and $\hat{B}_\beta$ is the Fourier coefficient of $A^{\gamma_2}_W$ at β.

Proof. By equation (10),

$$\frac{|C_{\gamma_1,\gamma_2}|^2}{2} = \frac{1}{2}\left|\mathrm{E}\Big[\sum_{\beta\in\hat{H}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}} \hat{A}_\alpha\hat{B}^2_\beta p(\alpha,\beta)\Big]\right|^2$$

$$\leq \left|\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\leq T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}\hat{A}_\alpha\hat{B}_\beta^2 p(\alpha,\beta)\Big]\right|^2 + \left|\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}\hat{A}_\alpha\hat{B}_\beta^2 p(\alpha,\beta)\Big]\right|^2,$$

where the inequality follows from Jensen's inequality. We now bound the latter sum above. By the Cauchy-Schwartz inequality and Parseval's inequality applied to $\hat{B}_\beta$,

$$\begin{aligned}
&\left|\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}\hat{A}_\alpha\hat{B}_\beta^2 p(\alpha,\beta)\right|^2\\
&\quad\leq \Big(\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}|\hat{A}_\alpha|^2|\hat{B}_\beta|^2\Big)\Big(\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}|\hat{B}_\beta|^2|p(\alpha,\beta)|^2\Big)\\
&\quad\leq \sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}|\hat{B}_\beta|^2\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}|p(\alpha,\beta)|^2,
\end{aligned}$$

hence

$$\left|\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}\hat{A}_\alpha\hat{B}_\beta^2 p(\alpha,\beta)\Big]\right|^2 \leq \mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}|\hat{B}_\beta|^2\,\mathrm{E}\Big[\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}|p(\alpha,\beta)|^2\;\Big|\;W\Big]\Big].$$

Consider the expression multiplying $|\hat{B}_\beta|^2$. Since $|\beta|\geq T\geq\delta^{-2}d^3\ln\delta^{-1}$, Corollary 2 implies that $|\pi(\beta)|\geq\delta^{-1}d^3\ln\delta^{-1}$ with probability at least $1-2\delta$ over the choice of U. When this holds,

$$\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}|p(\alpha,\beta)|^2\leq\Big(1-\frac{\delta}{d^3}\Big)^{|\pi(\beta)|}\leq\exp(\delta d^{-3}|\pi(\beta)|)\leq\delta.$$

In the other case, when $|\pi(\beta)|$ fails to be large,

$$\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}|p(\alpha,\beta)|^2\leq 1.$$

To conclude,

$$\mathrm{E}\Big[\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}|p(\alpha,\beta)|^2\;\Big|\;W\Big]\leq(1-2\delta)\cdot\delta+2\delta\cdot 1\leq 3\delta$$

and hence

$$\mathrm{E}\Big[\Big|\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}\sum_{\substack{\alpha\in F\\\alpha\preceq\beta}}\hat{A}_\alpha\hat{B}_\beta^2 p(\alpha,\beta)\Big|^2\Big]\leq 3\delta\,\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\geq T}}|\hat{B}_\beta|^2\Big].$$

Since the latter expectation is at most 1 due to Plancherel's equality, the conclusion of the lemma follows. ■

Lemma 17. *Let C_{γ_1,γ_2} be defined as in equation (9) and let $\gamma_1, \gamma_2 \in \hat{\boldsymbol{Z}}_d \setminus \{0\}$ and $\delta \in (0,1)$ be arbitrary. Then* $\mathrm{E}[|C_{\gamma_1,\gamma_2}|^2] \le 8\delta$ *provided that* $T \ge \delta^{-2} d^3 \ln \delta^{-1}$ *and* $u > (\log \delta^{-1} + \log T + T \log d)/\log c_\mu^{-1}$.

Proof. By Lemma 16,

$$|C_{\gamma_1,\gamma_2}|^2 \le 2\left|\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} \hat{A}_\alpha \hat{B}_\beta^2 p(\alpha,\beta)\Big]\right|^2 + 6\delta,$$

where $\hat{A}_\alpha$ is the Fourier coefficient of $A_U^{\gamma_1}$ at α and $\hat{B}_\beta$ is the Fourier coefficient of $A_W^{\gamma_2}$ at β. Now suppose that $|C_{\gamma_1,\gamma_2}|^2 > 8\delta$. Then

$$\begin{aligned}
\delta &< \left|\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} \hat{A}_\alpha \hat{B}_\beta^2 p(\alpha,\beta)\Big]\right|^2 \\
&\le \mathrm{E}\left[\Big|\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} \hat{A}_\alpha \hat{B}_\beta^2 p(\alpha,\beta)\Big|^2\right] \\
&\le \mathrm{E}\left[\Big(\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |\hat{A}_\alpha|^2|\hat{B}_\beta|^2\Big)\Big(\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |\hat{B}_\beta|^2|p(\alpha,\beta)|^2\Big)\right] \\
&\le \mathrm{E}\left[\Big(\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |\hat{A}_\alpha|^2|\hat{B}_\beta|^2\Big)\Big(d^T\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}} |\hat{B}_\beta|^2\Big)\right] \\
&\le d^T\,\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |\hat{A}_\alpha|^2|\hat{B}_\beta|^2\Big] \\
&\le Td^T\,\mathrm{E}\Big[\sum_{\substack{\beta\in\hat{H}\\|\beta|\le T}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |\hat{A}_\alpha|^2|\hat{B}_\beta|^2|\beta|^{-1}\Big] \\
&\le Td^T\,\mathrm{E}\Big[\sum_{\beta\in\hat{H}}\sum_{\substack{\alpha\in F\\ \alpha\preceq\beta}} |\hat{A}_\alpha|^2|\hat{B}_\beta|^2|\beta|^{-1}\Big].
\end{aligned}$$

We can now apply Lemma 9 with $\eta = \delta T^{-1} d^{-T}$ to see that there exists a strategy for the provers in the balanced version of the 2P1R game from § 2.2 with success rate η. Note that the functions $A_U^{\gamma_1}$ and $A_W^{\gamma_2}$ are known to the provers in that game since the provers can simply try all possible combinations of γ_1, γ_2, A_U, and A_W and select the combination that gives the largest expectation. But since u has been selected in such a way that $c_\mu^u < \eta$, we obtain a contradiction, and hence $|C_{\gamma_1,\gamma_2}|^2 \le 8\delta$. ■

Corollary 4. *Let C_1 and C_2 be defined as in equations (3)–(4) and $\delta \in (0,1)$ be arbitrary. Then $|\operatorname{E}[C_1]|^2 \le 8(d-1)\delta$ and $|\operatorname{E}[C_2]|^2 \le 8\delta$ provided that $T \ge \delta^{-2}d^3 \ln \delta^{-1}$ and $u > (\log \delta^{-1} + \log T + T \log d)/\log c_\mu^{-1}$.*

Proof. By the definitions of C_1 and C_2 from equations (3)–(4) it follows that

$$\begin{aligned} |\operatorname{E}[C_1]|^2 &= \frac{(d-1)^2}{d^2}\left|\frac{1}{d-1}\sum_{\gamma=1}^{d-1}\operatorname{E}[C_{\gamma,\gamma}]\right|^2 \le \frac{(d-1)}{d^2}\sum_{\gamma=1}^{d-1}|\operatorname{E}[C_{\gamma,\gamma}]|^2 \\ &\le \frac{(d-1)}{d^2}\sum_{\gamma=1}^{d-1}\operatorname{E}[|C_{\gamma,\gamma}|^2] \le \frac{8\delta(d-1)^2}{d^2} \end{aligned}$$

and similarly that

$$|\operatorname{E}[C_2]|^2 \le \frac{(d-1)^2}{d^4}\sum_{\gamma_1=1}^{d-1}\sum_{\gamma_2=1}^{d-1}\operatorname{E}[|C_{\gamma_1,\gamma_2}|^2] \le \frac{8\delta(d-1)^4}{d^4}.$$

Hence $|\operatorname{E}[C_1]|^2 \le 8\delta$ and $|\operatorname{E}[C_2]|^2 \le 8\delta$. ■

Putting the pieces together, it turns out that if we first select δ as a function of ε, then select T as a function of ε and δ and finally select u as a function of ε, δ and T, we can prove that all the terms above sum up to something having magnitude strictly less than ε under the assumption that the verifier in the PCP accepts with probability $d^{-1} + d^{-2} + \varepsilon$.

Lemma 18. *For any integer $d \ge 2$ and any constant $\varepsilon > 0$, there are choices of the parameters δ, T and u such that the verifier in Fig. 1 with these parameter choices has soundness at most $d^{-1} + d^{-2} + \varepsilon$.*

Proof. Given ε, first select $\delta \le \varepsilon^2/52$ and then select $T \ge \delta^{-2}d^3 \ln \delta^{-1}$ and $u > (\log \delta^{-1} + \log T + T \log d)/\log c_\mu^{-1}$. Now suppose that the verifier in Fig. 1 with these parameter choices accepts an incorrect input with probability $d^{-1} + d^{-2} + \varepsilon$. Then

$$\varepsilon = \operatorname{E}[L] + \operatorname{E}[Q] + \operatorname{E}[C_1] + \operatorname{E}[C_2] = \operatorname{E}[Q] + \operatorname{E}[C_1] + \operatorname{E}[C_2],$$

where L, Q, C_1, and C_2 are defined as in (1)–(4) and the last equality follows since $\operatorname{E}[L] = 0$ by Lemma 11. Jensen's inequality now implies that

$$\varepsilon^2 \le 3|\operatorname{E}[Q]|^2 + 3|\operatorname{E}[C_1]|^2 + 3|\operatorname{E}[C_2]|^2$$

Since $|\operatorname{E}[Q]|^2 \le \delta$ by Corollary 3 and $|\operatorname{E}[C_1]|^2 + |\operatorname{E}[C_2]|^2 \le 16\delta$ by Corollary 4, we obtain a contradiction. ■

Combing Lemma 10 and Lemma 18, we get Theorem 2.

The proof is a standard written d-proof with parameter u.

The verifier acts as follows:
Steps 1–6 are as in Fig. 1.

7. Select $e_{\omega f} \in H$ by selecting, independently for every $y \in \mathrm{SAT}^W$, $e_{\omega f}(y)$ such that:
   - $\omega f(\pi(y)) \neq 1 \implies e_{\omega f}(y) = 1$;
   - $\omega f(\pi(y)) = 1 \implies (\Pr[e_{\omega f}(y) = 1] = 1 - \delta) \wedge (\Pr[e_{\omega f}(y) = \omega] = \delta)$.
8. If $A_U(f) \neq 1$, accept if $A_U(f)A_W(h)A_W(h^{-1}(f \circ \pi)^{-1}e_f) = 1$;
   If $A_U(f) = 1$, accept if $A_W(h)A_W(h^{-1}(\omega^{-1}f \circ \pi)^{-1}e_{\omega^{-1}f}) = \omega$;
   Reject otherwise.

Figure 2. The above PCP is parameterized by the positive integers d, u and T and the positive real δ and tests if a μ-gap E3-Sat(5) formula Φ is satisfiable by adaptively querying three positions in a Standard Written d-proof with parameter u. With suitable choices of the parameters u, T and δ as functions of ε and d, the above PCP has perfect completeness and soundness $d^{-1} + \varepsilon$ for any constant $\varepsilon > 0$.

4 The adaptive PCP

In this section we give an adaptive version of the PCP construction in the previous section. The PCP is shown in Figure 2. The verifier proceeds as in the non-adaptive case when the first queried value, $A_U(f)$, is not 1. When $A_U(f) = 1$, the verifier essentially makes the same test as it would have if $\omega^{-1}f$ would have taken the place of f. Since $A_U(\omega^{-1}f) \neq 1$, the verifier may thus accept for only one value of $A_W(\cdot)A_W(\cdot)$, as in the case $A_U(f) \neq 1$. Due to folding, $A_U(\omega^{-1}f) = \omega^{-1}A_U(f)$, so the test can be performed without reading any extra values. Again we have perfect completeness:

Lemma 19. *The PCP in Fig. 2 has perfect completeness.*

Turning to the soundness, we note that the acceptance probability of the verifier can be written as $\mathrm{E}[I_1 + I_2]$ where

$$I_1 = \Big(1 - \frac{1}{d}\sum_{\gamma=0}^{d-1}\big(A_U(f)\big)^\gamma\Big) \times \Big(\frac{1}{d}\sum_{\gamma=0}^{d-1}\big(A_U(f)A_W(h)A_W(h^{-1}(f \circ \pi)^{-1}e_f)\big)^\gamma\Big)$$

and

$$I_2 = \Big(\frac{1}{d}\sum_{\gamma=0}^{d-1}\big(A_U(f)\big)^\gamma\Big) \times \Big(\frac{1}{d}\sum_{\gamma=0}^{d-1}\big(\omega^{-1}A_W(h)A_W(h^{-1}(\omega f \circ \pi)^{-1}e_{\omega f})\big)^\gamma\Big)$$

$$= \Big(\frac{1}{d}\sum_{\gamma=0}^{d-1}\Big(\omega A_U(\omega^{-1}f)\Big)^{\gamma}\Big)$$
$$\times \Big(\frac{1}{d}\sum_{\gamma=0}^{d-1}\Big(\omega^{-1}A_W(h^{-1}(\omega f\circ\pi)^{-1}e_{\omega^{-1}f})\Big)^{\gamma}\Big),$$

where the second equality follows since $A_U(\omega^{-1}f)=\omega^{-1}A_U(f)$. Since f is selected uniformly at random in the protocol,

$$\mathrm{E}\Big[\Big(A_W(h)A_W(h^{-1}(f\circ\pi)^{-1}e_f)\Big)^{\gamma}\Big]=\mathrm{E}\Big[\Big(A_W(h^{-1}(\omega f\circ\pi)^{-1}e_{\omega^{-1}f})\Big)^{\gamma}\Big]$$

and therefore $\mathrm{E}[I_1+I_2]=d^{-1}-\mathrm{E}[L]-\mathrm{E}[Q]+\mathrm{E}[C_1]-\mathrm{E}[C_2]+\mathrm{E}[L']+\mathrm{E}[Q']+\mathrm{E}[C']$, where L, Q, C_1 and C_2 are defined as in (1)–(4) and

$$L'=\frac{1}{d}\sum_{\gamma=1}^{d-1}\omega^{-\gamma}(A_U(f))^{\gamma},$$
$$Q'=\frac{1}{d^2}\sum_{\gamma=1}^{d-1}\omega^{\gamma}\Big(A_W(h)A_W(h^{-1}(f\circ\pi)^{-1}e_f)\omega^{-1}\Big)^{\gamma},$$
$$C'=\frac{1}{d^2}\sum_{\gamma_1=1}^{d-1}\sum_{\gamma_2=1}^{d-1}\omega^{\gamma_1-\gamma_2}\Big(A_U(f)\Big)^{\gamma_1}\Big(A_W(h)A_W(h^{-1}(f\circ\pi)^{-1}e_f)\omega^{-1}\Big)^{\gamma_2}.$$

From the previous section, we already know how to bound all these terms.

Lemma 20. *Let $\delta\in(0,1)$ be arbitrary and suppose that $T\geq\delta^{-2}d^3\ln\delta^{-1}$ and $u>(\log\delta^{-1}+\log T+T\log d)/\log c_\mu^{-1}$. Then $\mathrm{E}[L']=0$, $|\,\mathrm{E}[Q']|^2\leq\delta$, and $|\,\mathrm{E}[C']|^2\leq 8\delta$*

Proof. $E[L']=0$ since the tables in the proof are folded. The two inequalities follow from Lemmas 13 and 17 in the same way as Corollaries 3 and 4. ■

To prove that the adaptive protocol has soundness $d^{-1}+\varepsilon$ for any constant $\varepsilon>0$ is now straightforward.

Lemma 21. *For any integer $d\geq 2$ and any constant $\varepsilon>0$, there are choices of the parameters δ, T and u such that the verifier in Fig. 2 with these parameter choices has soundness at most $d^{-1}+\varepsilon$.*

Proof. Given ε, first select $\delta\leq\varepsilon^2/131$ and then select $T\geq\delta^{-2}d^3\ln\delta^{-1}$ and $u>(\log\delta^{-1}+\log T+T\log d)/\log c_\mu^{-1}$. Now suppose that the verifier in Fig. 1 with these parameter choices accepts an incorrect input with probability $d^{-1}+d^{-2}+\varepsilon$. Then

$$\varepsilon=-\mathrm{E}[L]-\mathrm{E}[Q]+\mathrm{E}[C_1]-\mathrm{E}[C_2]+\mathrm{E}[L']+\mathrm{E}[Q']+\mathrm{E}[C']$$

$$= -\operatorname{E}[Q] + \operatorname{E}[C_1] - \operatorname{E}[C_2] + \operatorname{E}[Q'] + \operatorname{E}[C'],$$

where L, Q, C_1, and C_2 are defined as in (1)–(4), Q' and C' are defined above, and the last equality follows since $\operatorname{E}[L] = \operatorname{E}[L'] = 0$ by Lemma 11. Hence Jensen's inequality implies that

$$\varepsilon^2 \leq 5|\operatorname{E}[Q]|^2 + 5|\operatorname{E}[C_1]|^2 + 5|\operatorname{E}[C_2]|^2 + 5|\operatorname{E}[Q']|^2 + 5|\operatorname{E}[C']|^2.$$

Since $|\operatorname{E}[Q]|^2 \leq \delta$ by Corollary 3, $|\operatorname{E}[C_1]|^2 + |\operatorname{E}[C_2]|^2 \leq 16\delta$ and $|\operatorname{E}[Q']|^2 + |\operatorname{E}[C']|^2 \leq 9\delta$ by Lemma 20, we obtain a contradiction. ∎

Combing Lemma 19 and Lemma 21, we get Theorem 2.

5 Conclusions and open problems

We have established that there exists, for every constant $\varepsilon > 0$, a non-adaptive PCP for **NP** that reads three values from a domain of size d, has perfect completeness and soundness $d^{-1} + d^{-2} + \varepsilon$. Moreover, we have presented an adaptive version of that PCP where the soundness is improved to $d^{-1}+\varepsilon$. As mentioned in the introduction, for $d > 2$ it is currently an open question whether these results are optimal. However, improvements would probably require new techniques for constructing PCPs. We believe that a search for better approximation algorithms for 3-ary constraint satisfaction over domains of size d could be very fruitful. In particular, an algorithm with approximation ratio better than $d^{-1} + d^{-2}$ for satisfiable instances would prove that adaptive verifiers are strictly more powerful than non-adaptive ones also for non-Boolean PCPs.

References

1. 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.
2. Lars Engebretsen. The non-approximability of non-Boolean predicates. In Michel Goemans, Klaus Jansen, José D. P. Rolim, and Luca Trevisan, editors, *Proceedings of 5th International Workshop on Randomization and Approximation Techniques in Computer Science*, volume 2129 of *Lecture Notes in Computer Science*, pages 241–248. Springer-Verlag, Berkeley, California, USA, 18–20 August 2001.
3. Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, July 1998.
4. Venkatesan Guruswami, Johan Håstad, and Mahdu Sudan. Hardness of approximate hypergraph coloring. In *41st Annual Symposium on Foundations of Computer Science*, pages 149–158. IEEE, Redondo Beach, California, 12–14 November 2000.
5. Venkatesan Guruswami, Daniel Lewin, Madhu Sudan, and Luca Trevisan. A tight characterization of NP with 3-query PCPs. In *39th Annual Symposium on Foundations of Computer Science*, pages 8–17. IEEE, Palo Alto, California, 8–11 November 1998.

6. Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, July 2001.

7. Johan Håstad and Subhash Khot. Query efficient PCPs with perfect completeness. In *42nd Annual Symposium on Foundations of Computer Science*, pages 610–619. IEEE, Las Vegas, Nevada, 14–17 November 2001.

8. Jonas Holmerin. Improved inapproximability results for vertex cover on k-uniform hypergraphs. ICALP 2002, to appear.

9. Jonas Holmerin. Vertex cover on 4-regular hyper-graphs is hard to approximate within $2 - \varepsilon$. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 544–552. Montréal, Québec, Canada, 19–21 May 2002.

10. Subhash Khot. Hardness of coloring 3-uniform hypergraphs and hardness of the not-all-equal predicate, April 2002. Manuscript.

11. Subhash Khot. Hardness results for approximate hypergraph coloring. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 351–359. Montréal, Québec, Canada, 19–21 May 2002.

12. Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.

13. Alex Samorodnitsky and Luca Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 191–199. Portland, Oregon, 21–23 May 2000.

14. Maria Serna, Luca Trevisan, and Fatos Xhafa. The (parallel) approximability of non-Boolean satisfiability problems and restricted integer programming. In Michel Morvan, Christoph Meinel, and Daniel Krob, editors, *Proceedings of the 15th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1373 of *Lecture Notes in Computer Science*, pages 488–498. Springer-Verlag, Paris, 25–27 February 1998.

15. Audrey Terras. *Fourier Analysis on Finite Groups and Applications*, volume 43 of *London Mathematical Society student texts.* Cambridge University Press, Cambridge, 1999.

16. Luca Trevisan. Parallel approximation algorithms by positive linear programming. *Algorithmica*, 21(1):72–88, May 1998.

17. Luca Trevisan. Approximating satisfiable satsifiabiliy problems. *Algorithmica*, 28(1):145–172, September 2000.

18. 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.