A 3-Query Non-Adaptive PCP with Perfect Completeness

Subhash Khot
khot@cc.gatech.edu

Rishi Saket
saket@cc.gatech.edu

December 3, 2005

Abstract

We study a very basic open problem regarding the PCP characterization of NP, namely, the power of PCPs with 3 non-adaptive queries and perfect completeness. Optimal results are known if one sacrifices either non-adaptiveness or perfect completeness. Håstad [11] constructs a 3-query non-adaptive PCP with soundness $\frac{1}{2} + \epsilon$ but it loses prefect completeness ($\epsilon > 0$ is an arbitrarily small constant). Guruswami *et al.* [9] construct a 3-query PCP with perfect completeness and soundness $\frac{1}{2} + \epsilon$ but the queries are adaptive. In a sharp contrast, Zwick [15] shows that a 3-query non-adaptive PCP with perfect completeness cannot achieve soundness below $\frac{5}{8}$. The lowest soundness known till now for such a PCP is $\frac{6}{8} + \epsilon$ given by a construction of Håstad [11].

In this paper, we construct a 3-query non-adaptive PCP with perfect completeness and soundness $\frac{20}{27} + \epsilon$, which improves upon the previous best soundness of $\frac{6}{8} + \epsilon$. A standard reduction from PCPs to constraint satisfaction problems (CSPs) implies that it is NP-hard to tell if a boolean CSP on 3-variables has a satisfying assignment or no assignment satisfies more than $\frac{20}{27} + \epsilon$ fraction of the constraints.

Our construction uses "biased Long Codes" introduced by Dinur and Safra [6]. We develope new 3-query tests to check consistency between such codes. These tests are analyzed by extending Håstad's Fourier methods [11] to the biased case.

1 Introduction

The celebrated PCP Theorem ([2], [1]) states that NP has probabilistic proof systems where the verifier is extremely efficient in terms of the number of random bits used and the number of queries made to the proof. A probabilistic polynomial-time verifier is said to be $(r(n), q(n))$-restricted, if on an input x of length n, the verifier uses at most $r(n)$ random bits and queries at most $q(n)$ bits from the proof. For $0 < s < c \leq 1$, let $\mathrm{PCP}_{c,s}[r(n), q(n)]$ denote the class of languages L which have a proof system where the verifier is $(r(n), q(n))$-restricted and satisfies the following properties :

- Completeness : If input $x \in L$, there exists a proof that the verifier accepts with probability c.
- Soundness : If $x \notin L$, no proof is accepted with probability more than s.

The parameters c and s are called the completeness and the soundness parameter respectively. The queries made by the verifier could be adaptive or non-adaptive. To make this point explicit, we will denote the corresponding classes by $\mathrm{aPCP}_{c,s}[r(n), q(n)]$ (adaptive queries) and $\mathrm{naPCP}_{c,s}[r(n), q(n)]$ (non-adaptive queries) respectively. With this notation, the PCP Theorem can be stated as

Theorem 1.1 (The PCP Theorem [1], [2]) $\quad \mathrm{NP} \subseteq \mathrm{naPCP}_{1,1/2}[O(\log n), O(1)]$

In this statement we have $c = 1$, i.e. when $x \in L$, there exists a proof that the verifier always accepts. Such a verifier is said to have perfect completeness, which is a natural property one may desire of a proof system. After the discovery of the PCP Theorem, a series of papers ([3], [4], [11], [9], [14]) led to constructions of verifiers which achieve better and better trade-off between the number of queries and the soundness parameter. Such constructions have direct implications for hardness of approximating optimization problems, for example Max-3SAT, Max-Cut and Vertex Cover. In this paper, we study the power of PCPs when the verifier is allowed to make only 3 non-adaptive queries to the proof and required to have perfect completeness. The question we address is :

$$\text{What is the smallest value of } s \text{ s.t. } \mathrm{NP} \subseteq \mathrm{naPCP}_{1,s}[O(\log n), 3] \text{ ?}$$

This question has been well-studied before and optimal results are known if we sacrifice either the perfect completeness or non-adaptiveness. Håstad's [11] famous 3-bit PCP construction shows that

Theorem 1.2 ([11]) $\quad \forall\, \epsilon > 0, \quad \mathrm{NP} \subseteq \mathrm{naPCP}_{1-\epsilon, \frac{1}{2}+\epsilon}[O(\log n), 3]$

Håstad's verifier loses perfect completeness and the analysis of this verifier makes an essential use of this feature. Guruswami *et al.* [9] consider adaptive verifiers and prove that

Theorem 1.3 ([9]) $\quad \forall\, \epsilon > 0, \quad \mathrm{NP} \subseteq \mathrm{aPCP}_{1, \frac{1}{2}+\epsilon}[O(\log n), 3]$

However when we require both perfect completeness and non-adaptiveness, the situation changes dramatically. Zwick [15] gives a polynomial-time randomized algorithm which given a satisfiable instance of a boolean 3-CSP, finds an assignment satisfying $\frac{5}{8}$ fraction of the constraints. This result implies that

Theorem 1.4 ([15]) $\quad \mathrm{naPCP}_{1,5/8}[O(\log n), 3] \subseteq \mathrm{BPP}$

Therefore a 3-query non-adaptive verifier with perfect completeness cannot achieve soundness below $\frac{5}{8}$ unless NP $\subseteq$ BPP, in sharp contrast with the adaptive or imperfect completeness case where the verifiers can achieve soundness $\frac{1}{2} + \epsilon$. On the other hand, such a verifier can achieve soundness $\frac{6}{8} + \epsilon$ as shown by Håstad [11] and this result is the best known result till date.

In this paper we partially bridge the gap ($\frac{6}{8}$ vs $\frac{5}{8}$) between Håstad's result and Zwick's result, by constructing a 3 query non-adaptive PCP with perfect completeness and soundness of $\frac{20}{27} + \epsilon$. The following theorem states the main result in this paper :

Theorem 1.5 $\forall\, \epsilon > 0, \quad \mathrm{NP} \subseteq \mathrm{naPCP}_{1, \frac{20}{27}+\epsilon}[O(\log n), 3]$

A standard reduction from PCPs to CSPs (taking the bits in the proof as variables of a CSP and the tests of the verifier as the constraints of the CSP) gives the following theorem :

Theorem 1.6 *For any constant $\epsilon > 0$, it is NP-hard to tell if a boolean CSP on 3-variables has a satisfying assignment or no assignment satisfies more than $\frac{20}{27} + \epsilon$ fraction of the constraints.*

Main techniques : The main technique in this paper is to use "biased Long Codes" introduced by Dinur and Safra [6]. Their paper uses Long Codes in a very combinatorial way whereas we use Fourier analysis of biased Long Codes, extending Håstad's Fourier methods to the biased case. We build new PCP tests where the verifier uses 3 non-adaptive queries, has perfect completeness and reasonable soundness.

The main test in the paper relies on the following observation. For $p = \frac{1}{2} + \epsilon$ and $q = 1 - p$, let μ_p denote the distribution on a bit x where one sets $x = 1$ with probability p and $x = 0$ with probability q. Consider the following distribution on 3 bits (x, y, z) :

$$(x,y,z) = \begin{cases} (0,0,0) & \text{with probability } q^2 \\ (0,1,1) & \text{with probability } pq \\ (1,0,1) & \text{with probability } pq \\ (1,1,0) & \text{with probability } pq \\ (1,1,1) & \text{with probability } p^2 - pq \end{cases}$$

This distribution satisfies the following properties (which turn out to be crucial for analysis) :

- Each of the bits x, y, z is distributed according to μ_p.
- The bits (x, y, z) are pairwise independent.
- $\Pr[x \oplus y \oplus z = 0] < 1$.

This observation leads (in a straightforward manner) to a 3-bit non-adaptive PCP test with perfect completeness. The fact that the triple (x, y, z) takes only 5 different settings corresponds to the fact that the PCP verifier's test has only 5 satisfying assignments. One may ideally expect soundness $\frac{5}{8}$, however the test breaks down when the proof consists solely of 0s or of 1s since the verifier accepts when the bits read are $(0, 0, 0)$ and $(1, 1, 1)$. We handle this problem by combining the test with three more tests. We are then able to show that the soundness is at most $\frac{20}{27} + \epsilon$.

We would like to point out that Håstad's 3-bit PCP (Theorem 1.2) is based on a distribution on 3 bits with $p = 1/2$. This distribution gives non-zero probability mass to all the 8 settings of bits (x, y, z). However the verifier rejects the settings for which $x \oplus y \oplus z = 1$. These 4 settings have a total probability mass of ϵ and therefore the completeness is only $1 - \epsilon$.

Overview of the paper : Section 2 introduces the tools used in the paper including the 2-Prover Games, the biased Long Codes and the Fourier Analysis. Section 3 describes the 3-bit PCP tests of the verifier which is the crux of the paper. We construct the final PCP verifier in Section 4. We conclude in Section 5 suggesting ways to improve the results in this paper.

2 Preliminaries

2.1 Standard Framework for PCP Constructions

An equivalent statement of the PCP Theorem is the following :

Theorem 2.1 *For some constant $c < 1$, it is NP-hard to distinguish whether a 3-SAT formula ψ is satisfiable (the YES instance) or no assignment satisfies more than a fraction c of the clauses (the NO instance).*

One can assume that the formula ψ in Theorem 2.1 has a regular structure, meaning every clause contains exactly 3 variables and every variable appears in exactly 5 clauses. We call such a formula an instance of 3-SAT-5.

We follow the standard framework for PCP constructions developed by Bellare et al [4] and Håstad [11]. In this framework, we first construct a 2-Prover-1-Round Game from the 3-SAT-5 instance ψ given by Theorem 2.1. The PCP verifier then expects as a proof the encodings of provers' answers in the 2-Prover Game. The specific encoding used is the Long Code introduced by Bellare et al. The test of the verifier consists of reading a few bits from the proof and performing a local consistency check.

2.2 The 2-Prover-1-Round Game

We will use the 2-Prover Game constructed by Khot [12] which is a slight modification of the 2-Prover Games used earlier (see [4], [11]). The verifier in this game will be denoted by V_{2p1r} (to distinguish it from the PCP verifier we want to construct).

Let $\{x_1, x_2, \ldots\}$ be the variables and $\{C_1, C_2, \ldots\}$ be the clauses of the 3-SAT-5 instance ψ. The game is parameterized by two integers T and u. Think of $T, u \gg 1$ and these parameters can be made as large as one wants independent of each other. The verifier V_{2p1r} picks a set of Tu clauses at random, say $W = \{C_1, C_2, \ldots, C_{Tu}\}$. W will be the question to the Prover 1 who is required to give as an answer, a satisfying assignment to the clauses in W. Denoting the set of satisfying assignments to W by $\mathcal{M}_W$, the answer of Prover 1 is some $\sigma \in \mathcal{M}_W$. Now the verifier picks a random subset of W with size u, say $S = \{C_{i_1}, C_{i_2}, \ldots, C_{i_u}\}$ where $1 \leq i_1 < i_2 \ldots < i_u \leq Tu$. Each clause C_{i_j} contains 3 variables and the verifier picks one of these variables at random, say x_{i_j}. By abuse of notation, let $U = \{x_{i_1}, x_{i_2}, \ldots, x_{i_u}\} \cup (W \setminus S)$. Note that U is a set of u variables and $(T-1)u$ clauses. U will be the question to the Prover 2 who is required to give as an answer, an assignment to U satisfying all the clauses in U. Denoting the set of all such assignments by $\mathcal{M}_U$, the answer of Prover 2 is some $\tau \in \mathcal{M}_U$. Note that every assignment to W can be restricted to an assignment to U and this is precisely the consistency check the verifier performs. The verifier V_{2p1r} accepts iff τ is a restriction of σ. Defining a map $\pi^{W,U} : \mathcal{M}_W \mapsto \mathcal{M}_U$ which maps an assignment to W to its restriction to U, the verifier accepts iff $\pi^{W,U}(\sigma) = \tau$.

We will denote by $\mathcal{W}$, the set of all questions asked to the Prover 1 and by $\mathcal{U}$, the set of all questions asked to the Prover 2. Clearly, if the formula ψ is a YES instance (i.e. satisfiable), the provers in this game have a strategy that makes the verifier accept with probability 1. The strategy is to fix one satisfying assignment to ψ and give answers consistent with this assignment.

If the formula ψ is a NO instance (i.e. no assignment satisfies more than a fraction c of the clauses), Raz's Parallel Repetition Theorem [13] implies the following :

Theorem 2.2 *If ψ is a NO instance, no strategy of the provers can make the verifier accept with probability more than c_0^u where $c_0 < 1$ is an absolute constant.*

We need the following *smoothness property* of this 2-Prover Game which is proved in Appendix B.

Lemma 2.3 *For a fixed $W \in \mathcal{W}$ and $\beta \subseteq \mathcal{M}_W, \beta \neq \emptyset$, we have*

$$E_U\left[\frac{1}{|\pi(\beta)|}\right] \leq \frac{1}{|\beta|} + \frac{1}{T} \qquad (\pi = \pi^{W,U})$$

where the expectation is taken over the choice of the question U to Prover-2 conditional on the question to Prover-1 being W. In particular, if $T \geq \frac{1}{\epsilon^4}$ and $|\beta| \geq \frac{1}{\epsilon^3}$, then except with probability 2ϵ, $|\pi(\beta)| \geq 1/\epsilon^2$.

As mentioned before, the PCP verifier expects as a proof, the Long Codes of provers' answers in the 2-Prover Game. We define the Long Code and the biased version of the Long Code in the next section.

2.3 Biased Long Code and Fourier Analysis

It is convenient to change the $\{0,1\}$-notation to $\{1,-1\}$-notation. Henceforth, we will assume that the encodings/proofs will consist of 1 and -1 instead of bits 0 and 1 respectively.

The Long Code on a set $\mathcal{M}$ is indexed by all functions $g : \mathcal{M} \mapsto \{-1,1\}$. We denote

$$\mathcal{G} := \{ \, g \mid g : \mathcal{M} \mapsto \{-1,1\} \, \}$$

The Long Code B of $b \in \mathcal{M}$ is defined as

$$B(g) = g(b) \quad \forall \, g \in \mathcal{G}$$

A "biased" Long Code with bias $0 < p < 1$ has a probability distribution on the indices g of the code, where an index is selected by picking a function g with $g(x) = -1$ with probability p and $g(x) = 1$ with probability $1-p$ independently for all $x \in \mathcal{M}$. We denote this as $g \in_R \mu_p(\mathcal{M})$. Let $q = 1-p$.

We briefly explain the Fourier analysis of biased Long Codes. It is well-known how to extend the Fourier methods to the biased case (e.g. see [8]). One needs to identify the right orthonormal basis.

The space of all "tables" $B : \mathcal{G} \mapsto \mathbf{R}$ forms a real vector space with dimension $2^{|\mathcal{M}|}$ where addition of two tables is defined as pointwise addition. For example, a long code is one such table. We define an inner product on this space as

$$< B_1, B_2 > := \; E_{g \in_R \mu_p(\mathcal{M})}[B_1(g) B_2(g)]$$

For every $x \in \mathcal{M}$, define a function $\phi_x : \mathcal{G} \mapsto \mathbf{R}$ as

$$\phi_x(g) = \begin{cases} -\sqrt{q/p} & \text{if } g(x) = -1 \\ \sqrt{p/q} & \text{if } g(x) = 1 \end{cases}$$

Now we identify an orthonormal basis for the vector space. For every subset $\beta \subseteq \mathcal{M}$, the character $\chi_\beta : \mathcal{G} \mapsto \mathbf{R}$ is defined as

$$\chi_\beta := \prod_{x \in \beta} \phi_x$$

With this definition, $\chi_\emptyset \equiv 1$ and for any $x \in \mathcal{M}$, $\chi_{\{x\}} = \phi_x$. It is instructive to verify that the characters χ_β are orthonormal and we do this in Appendix A. It follows that any table B can be expressed as

$$B = \sum_{\beta \subseteq \mathcal{M}} \widehat{B}_\beta \, \chi_\beta$$

where $\widehat{B}_\beta$ are real numbers called Fourier Coefficients. When the range of B is $\{-1,1\}$, we have Parseval's identity, i.e. $\sum_\beta \widehat{B}_\beta^2 = 1$.

3 The Tests of the Verifier

The test of the verifier will be a combination of 4 tests, T_1, T_2, T_3 and T_4. The verifier will perform the 4 tests with probabilities to be decided later. Let us fix $p = \frac{1}{2} + \epsilon$ and let $q = 1-p$. The parameters ϵ, T, u are chosen in the following way : Given ϵ arbitrarily small, we choose $T = \frac{1}{\epsilon^4}$ and then choose u sufficiently large. This particular order is necessary in the analysis.

3.1 The Test T_1

The test T_1 is based on the 3-bit distribution described in the introduction.

Test T_1

1. Pick a random set $W \in \mathcal{W}$ and its random sub-set $U \in \mathcal{U}$ as the verifier V_{2p1r} would do. Let B, A be the supposed Long Codes of assignments to W and U respectively. Let $\pi = \pi^{W,U}$ be the projection between W and U.

2. Pick functions $f \in_R \mu_p(\mathcal{M}_U)$ and $g \in_R \mu_p(\mathcal{M}_W)$ independently.

3. Define a function $h : \mathcal{M}_W \mapsto \{-1, 1\}$ as follows : For every $y \in \mathcal{M}_W$,
   - If $f(\pi(y)) = 1$ and $g(y) = 1$, define $h(y) = 1$
   - If $f(\pi(y)) = 1$ and $g(y) = -1$, define $h(y) = -1$
   - If $f(\pi(y)) = -1$ and $g(y) = 1$, define $h(y) = -1$
   - If $f(\pi(y)) = -1$ and $g(y) = -1$, define $h(y) = 1$ with probability q/p and define $h(y) = -1$ with probability $1 - q/p$.

4. Accept iff the triple $(A(f), B(g), B(h)) \in S_1$ where

$$S_1 := \{(1,1,1), (1,-1,-1), (-1,1,-1), (-1,-1,1), (-1,-1,-1)\}$$

3.2 Completeness

Note that when the 3-SAT-5 formula ψ is satisfiable, the provers in the 2-Prover Game have a strategy that makes the verifier V_{2p1r} always accept. Consider a proof obtained by encoding the provers' answers by correct long codes. For this proof, A, B are Long Code of some $x \in \mathcal{M}_U$ and $y \in \mathcal{M}_W$ with $\pi(y) = x$. Now by definition of Long Code, $(A(f), B(g), B(h)) = (f(x), g(y), h(y)) = (f(\pi(y)), g(y), h(y))$ and we note that the test always chooses (f, g, h) so that

$$(f(\pi(y)), g(y), h(y)) \in S_1 \quad \forall\, y \in \mathcal{M}_W$$

Therefore the PCP verifier always accepts, i.e. it has perfect completeness.

3.3 Soundness

Now consider the case when the formula ψ is a NO instance. We will bound the acceptance probability of the verifier using Fourier analysis methods.

Lemma 3.1 *When $x, y, z \in \{-1, 1\}$, the expression*

$$\frac{5 - x - y - z + xy + xz + yz + 3xyz}{8}$$

takes value 1 *when* $(x, y, z) \in S_1$ *and* 0 *otherwise.*

Using this lemma, the acceptance probability of the verifier can be written as the expression

$$\Pr[\text{Acc}] = E_{W,U,f,g,h}\Big[\frac{5 - A(f) - B(g) - B(h) + A(f)B(g) + A(f)B(h) + B(g)B(h) + 3A(f)B(g)B(h)}{8}\Big] \tag{1}$$

Lemma 3.2 *In Test T_1, the function h is identically distributed as the function g (i.e. both are distributed according to $\mu_p(\mathcal{M}_W)$) and the functions (f, g) and (f, h) are pairwise independent.*

It follows from this lemma that

$$E_f[A(f)] = \widehat{A}_\emptyset, \quad E_g[B(g)] = E_h[B(h)] = \widehat{B}_\emptyset$$
$$E_{f,g}[A(f)B(g)] = E_{f,h}[A(f)B(h)] = \widehat{A}_\emptyset \widehat{B}_\emptyset$$

We postpone the analysis of the term $E_{g,h}[B(g)B(h)]$ and look at the "interesting term"

$$E_{W,U,f,g,h}[A(f)B(g)B(h)]$$

Using Fourier expansions of A and B, we get

$$E_{W,U,f,g,h}[\sum_{\alpha,\beta,\gamma} \widehat{A}_\alpha \widehat{B}_\beta \widehat{B}_\gamma \cdot \chi_\alpha(f)\chi_\beta(g)\chi_\gamma(h)]$$

Because of pairwise independence of (f, g) and (f, h), it follows that the expectation is non-zero only if $\beta = \gamma$ and $\alpha \subseteq \pi(\beta)$. Thus the expression reduces to

$$E_{W,U,f,g,h}[\sum_{\alpha,\beta:\alpha\subseteq\pi(\beta)} \widehat{A}_\alpha \widehat{B}_\beta^2 \cdot \chi_\alpha(f)\chi_\beta(g)\chi_\beta(h)] \tag{2}$$

We need the following technical lemma which is proved in Appendix C.

Lemma 3.3 *If $\beta \subseteq \mathcal{M}_W$, $\alpha \subseteq \pi(\beta)$ and for every $x \in \pi(\beta)$*

$$c_x := |\{\ y \mid y \in \beta,\ \pi(y) = x\}| \tag{3}$$

then in Test T_1 we have

$$E_{f,g,h}[\chi_\alpha(f)\chi_\beta(g)\chi_\beta(h)] = \prod_{x\in\pi(\beta)\setminus\alpha} \left(q + p(-\frac{q}{p})^{c_x}\right) \cdot \prod_{x\in\alpha} \sqrt{pq}\left(1 - (-\frac{q}{p})^{c_x}\right)$$

In particular, this expectation has magnitude at most 1.

$$\text{Define} \qquad R(\beta) = \sum_{\alpha\subseteq\pi(\beta),\alpha\neq\emptyset} \widehat{A}_\alpha\ E_{f,g,h}[\chi_\alpha(f)\chi_\beta(g)\chi_\beta(h)]$$

Thus expression (2) can be written as (with expectation over W, U implicit),

$$\widehat{A}_\emptyset \widehat{B}_\emptyset^2 + \widehat{A}_\emptyset \sum_{\beta\neq\emptyset} \widehat{B}_\beta^2 \prod_{x\in\pi(\beta)} \left(q + p(-\frac{q}{p})^{c_x}\right) + \sum_{\beta\neq\emptyset,|\beta|>1/\epsilon^3} \widehat{B}_\beta^2\ R(\beta) + \sum_{\beta\neq\emptyset,|\beta|\leq 1/\epsilon^3} \widehat{B}_\beta^2\ R(\beta) \tag{4}$$

We upper bound the second, third and the fourth terms separately. Note that

$$\begin{aligned} |R(\beta)| &\leq \sum_{\alpha\subseteq\pi(\beta)} |\widehat{A}_\alpha| \prod_{x\in\pi(\beta)\setminus\alpha} \left|q + p(-\frac{q}{p})^{c_x}\right| \cdot \prod_{x\in\alpha} \left|\sqrt{pq}\left(1 - (-\frac{q}{p})^{c_x}\right)\right| \\ &\leq \sqrt{\sum_{\alpha\subseteq\pi(\beta)} \widehat{A}_\alpha^2} \sqrt{\sum_{\alpha\subseteq\pi(\beta)} \prod_{x\in\pi(\beta)\setminus\alpha} \left|q + p(-\frac{q}{p})^{c_x}\right|^2 \cdot \prod_{x\in\alpha} \left|\sqrt{pq}\left(1 - (-\frac{q}{p})^{c_x}\right)\right|^2} \end{aligned}$$

$$\begin{aligned}
&\le \sqrt{\prod_{x\in\pi(\beta)}\left(\left|q+p(-\frac{q}{p})^{c_x}\right|^2+\left|\sqrt{pq}\Big(1-(-\frac{q}{p})^{c_x}\Big)\right|^2\right)}\\
&= \sqrt{\prod_{x\in\pi(\beta)}\Big(q+p(\frac{q}{p})^{2c_x}\Big)}\\
&\le (1-2\epsilon)^{|\pi(\beta)|/2}
\end{aligned}$$

Consider the expectation over the choice of U. When $|\beta| > 1/\epsilon^3$, using Lemma 2.3, except with probability 2ϵ, we have $|\pi(\beta)| \ge 1/\epsilon^2$ and consequently $|R(\beta)| \le (1-2\epsilon)^{1/(2\epsilon^2)} \le \epsilon$. This upper bounds the third term in expression (4). The fourth term is

$$\begin{aligned}
\left|\sum_{\substack{\beta\neq\emptyset,|\beta|\le 1/\epsilon^3,\\ \alpha\neq\emptyset,\alpha\subseteq\pi(\beta)}} \widehat{B}_\beta^2\widehat{A}_\alpha E_{f,g,h}[\chi_\alpha(f)\chi_\beta(g)\chi_\beta(h)]\right| &\le \sum_{\beta\neq\emptyset,|\beta|\le 1/\epsilon^3,\alpha\subseteq\pi(\beta),\alpha\neq\emptyset} \widehat{B}_\beta^2\,|\widehat{A}_\alpha|\\
&\le \sqrt{\sum_{|\beta|\le 1/\epsilon^3,\alpha\subseteq\pi(\beta)}\widehat{B}_\beta^2}\sqrt{\sum_{\beta\neq\emptyset,|\beta|\le 1/\epsilon^3,\alpha\subseteq\pi(\beta),\alpha\neq\emptyset}\widehat{A}_\alpha^2\widehat{B}_\beta^2}\\
&\le \sqrt{2^{1/\epsilon^3}}\sqrt{\sum_{\beta\neq\emptyset,|\beta|\le 1/\epsilon^3,\alpha\subseteq\pi(\beta),\alpha\neq\emptyset}\widehat{A}_\alpha^2\widehat{B}_\beta^2} \qquad (5)
\end{aligned}$$

Now note that expression (5) (with the outer expectation over W, U) gives a way of defining provers' strategy in the 2-Prover Game. On question W, Prover-1 picks $\beta \subseteq \mathcal{M}_W$ with probability $\widehat{B}_\beta^2$, picks a random $y \in \beta$ and gives it as an answer. On question U, Prover-2 picks $\alpha \subseteq \mathcal{M}_U$ with probability $\widehat{A}_\alpha^2$, picks a random $x \in \alpha$ and gives it as an answer. Expression (5) gives the probability that $\alpha \subseteq \pi(\beta)$, $|\beta| \le 1/\epsilon^3$ and $\alpha \neq \emptyset$. With a further ϵ^3 probability, we have $\pi(y) = x$ and the verifier in the 2-Prover Game accepts. However by Theorem 2.2 the acceptance probability of the 2-Prover Game can be assumed to be arbitrarily small (by choosing u large enough) and hence expression (5) can be upper bounded by ϵ.

Now we are left with the second term in equation (4) and the term $E_{g,h}[B(g)B(h)]$. We observe that both these terms look alike and can be bounded simultaneously.

$$\begin{aligned}
E_{g,h}[B(g)B(h)] &= E_{g,h}[\sum_{\beta,\gamma}\widehat{B}_\beta\widehat{B}_\gamma\,\chi_\beta(g)\chi_\gamma(g)]\\
&= \sum_\beta \widehat{B}_\beta^2 E_{g,h}[\chi_\beta(g)\chi_\beta(h)]\\
&= \widehat{B}_\emptyset^2 + \sum_{\beta\neq\emptyset}\widehat{B}_\beta^2\prod_{x\in\pi(\beta)}(q+p(-\frac{q}{p})^{c_x}) \qquad (6)\\
&\le \widehat{B}_\emptyset^2 + \sum_{1\le|\beta|\le 1/\epsilon^3}\widehat{B}_\beta^2\prod_{x\in\pi(\beta)}(q+p(-\frac{q}{p})^{c_x}) + \sum_{|\beta|>1/\epsilon^3}\widehat{B}_\beta^2\,(1-2\epsilon)^{|\pi(\beta)|} \qquad (7)
\end{aligned}$$

where computing the expectation in equation (6) is a special case of Lemma 3.3 with $\alpha = \emptyset$. When $|\beta| \le 1/\epsilon^3$, Corollary B.2 implies that except with probability $|\beta|/T \le \epsilon$, there exists $x \in \pi(\beta)$ such that $c_x = 1$. When $c_x = 1$, the product in (6) vanishes. When $|\beta| > 1/\epsilon^3$, Lemma 2.3 implies that except with probability 2ϵ, $|\pi(\beta)| \ge 1/\epsilon^2$ and consequently $(1-2\epsilon)^{|\pi(\beta)|}$ is a negligible quantity.

Combining all the bounds proved so far, the acceptance probability of the test T_1 can be bounded by,

$$\Pr[\mathrm{Acc}] \le E_{W,U}\Big[\frac{5-\widehat{A}_\emptyset-2\widehat{B}_\emptyset+2\widehat{A}_\emptyset\widehat{B}_\emptyset+\widehat{B}_\emptyset^2+3\widehat{A}_\emptyset\widehat{B}_\emptyset^2}{8}\Big]+O(\epsilon) \qquad (8)$$

3.4 The Test T_2

The Test T_2 is a variation of the Test T_1. After picking functions f, g, the verifier picks h in a different way and the acceptance condition is also different.

Test T_2

1. Pick a random set $W \in \mathcal{W}$ and its random sub-set $U \in \mathcal{U}$ as the verifier V_{2par} would do. Let B, A be the supposed Long Codes of assignments to W and U respectively. Let $\pi = \pi^{W,U}$ be the projection between W and U.

2. Pick functions $f \in_R \mu_p(\mathcal{M}_U)$ and $g \in_R \mu_p(\mathcal{M}_W)$ independently.

3. Define a function $h : \mathcal{M}_W \mapsto \{-1, 1\}$ as follows : For every $y \in \mathcal{M}_W$,

   - If $f(\pi(y)) = 1$ and $g(y) = 1$, define $h(y) = -1$
   - If $f(\pi(y)) = 1$ and $g(y) = -1$, define $h(y) = 1$ with probability q/p and define $h(y) = -1$ with probability $1 - q/p$.
   - If $f(\pi(y)) = -1$ and $g(y) = 1$, define $h(y) = 1$
   - If $f(\pi(y)) = -1$ and $g(y) = -1$, define $h(y) = -1$

4. Accept iff the triple $(A(f), B(g), B(h)) \in S_2$ where

$$S_2 := \{(1, 1, -1), (1, -1, 1), (1, -1, -1), (-1, 1, 1), (-1, -1, -1)\}$$

It is clear that the test has perfect completeness. To analyze the soundness, we arithmetize the acceptance condition as,

$$\Pr[\text{Acc}] = E_{W,U,f,g,h}\left[\frac{5 + A(f) - B(g) - B(h) - A(f)B(g) - A(f)B(h) + B(g)B(h) - 3A(f)B(g)B(h)}{8}\right] \tag{9}$$

Note that in Test T_2, h is distributed identically as g. Also, (f, g) and (f, h) are pairwise independent. The test T_2 can be analyzed along more or less the same lines as the test T_1 and it can be shown that

$$\Pr[\text{Acc}] \leq E_{W,U}\left[\frac{5 + \widehat{A}_\emptyset - 2\widehat{B}_\emptyset - 2\widehat{A}_\emptyset\widehat{B}_\emptyset + \widehat{B}_\emptyset^2 - 3\widehat{A}_\emptyset\widehat{B}_\emptyset^2}{8}\right] + O(\epsilon) \tag{10}$$

We omit the proof. However we state equivalent of Lemma 3.3 for the test T_2 as Lemma D.1.

3.5 The Test T_3

The test T_3 is a Not-All-Equal test on a supposed Long Code B. Note that tests T_1, T_2 check consistency between two tables A and B. However, test T_3 is a test on a single table B.

Test T_3

1. Pick a random set $W \in \mathcal{W}$ and let B be the supposed Long Code of the assignment to W.

2. Pick 3 functions $g_1, g_2, g_3 : \mathcal{M}_W \mapsto \{-1, 1\}$, where for every $y \in \mathcal{M}_W$, we set

$$(g_1(y), g_2(y), g_3(y)) = \begin{cases} (-1, 1, 1), (1, -1, 1), (1, 1, -1) & \text{with probability } \frac{2}{3} - p \text{ each} \\ (1, -1, -1), (-1, 1, -1), (-1, -1, 1) & \text{with probability } p - \frac{1}{3} \text{ each} \end{cases}$$

3. Accept iff Not-All-Equal$(B(g_1), B(g_2), B(g_3))$

Clearly, the test always accepts a correct Long Code and has perfect completeness. We note that each $g_i \in \mu_p(\mathcal{M}_W)$, however there is no independence between any pair of them. Arithmetizing the Not-All-Equal predicate, we get

$$\Pr[\text{Acc}] = E_{W,g_1,g_2,g_3}\Big[\frac{3 - B(g_1)B(g_2) - B(g_2)B(g_3) - B(g_3)B(g_1)}{4}\Big] \tag{11}$$

This expression can be shown to be (see Appendix E)

$$\begin{aligned}
\Pr[\text{Acc}] &= E_W\Big[\frac{3 - 3\sum_\beta \widehat{B}_\beta^2\Big(-(\frac{1}{3}\frac{1/4+3\epsilon^2}{1/4-\epsilon^2})\Big)^{|\beta|}}{4}\Big] \\
&\leq E_W\Big[\frac{3 - 3\widehat{B}_\emptyset^2 + 3\cdot(\frac{1}{3}+\epsilon)\sum_{\beta\neq\emptyset}\widehat{B}_\beta^2}{4}\Big] \\
&= E_W\Big[\frac{3 - 3\widehat{B}_\emptyset^2 + 3(\frac{1}{3}+\epsilon)(1-\widehat{B}_\emptyset^2)}{4}\Big] \\
&\leq E_W[1 - \widehat{B}_\emptyset^2] + \epsilon
\end{aligned} \tag{12}$$

3.6 The Test T_4

This test is also on a supposed long code B.

Test T_4

1. Let B be the supposed Long Code of the assignment to a random W in $\mathcal{W}$.
2. Pick 3 functions $g_1, g_2, g_3 : \mathcal{M}_W \mapsto \{-1, 1\}$, where for every $y \in \mathcal{M}_W$, we set

$$(g_1(y), g_2(y), g_3(y)) = \begin{cases} (-1,-1,1) & \text{with probability } 2p-1 \\ (-1,1,1),(1,-1,1) & \text{with probability } 1-\frac{3p}{2} \text{ each} \\ (-1,1,-1),(1,-1,-1) & \text{with probability } \frac{p}{2} \text{ each} \end{cases}$$

3. Accept iff $(B(g_1), B(g_2), B(g_3)) \neq (-1,-1,-1)$

We can see that the test accepts a correct Long Code and has perfect completeness. In the No case, we split the probability of acceptance in the following manner,

$$\Pr[\text{Acc}] \leq \Pr[(B(g_1), B(g_3)) \neq (-1,-1)] + \Pr[B(g_2) = 1 \;\&\; B(g_3) = -1] \tag{13}$$

Arithmetizing the terms we get (see Appendix F for details),

$$\begin{aligned}
\Pr[(B(g_1), B(g_3)) \neq (-1,-1)] &= E_{W,g_1,g_3}\Big[\frac{3 + B(g_1) + B(g_3) - B(g_1)B(g_3)}{4}\Big] \\
&\leq E_W\Big[\frac{3 + 2\widehat{B}_\emptyset - \widehat{B}_\emptyset^2}{4}\Big] + O(\epsilon)
\end{aligned} \tag{14}$$

and

$$\begin{aligned}
\Pr[B(g_2) = 1 \;\&\; B(g_3) = -1] &= E_{W,g_2,g_3}\Big[\frac{1 + B(g_2) - B(g_3) - B(g_2)B(g_3)}{4}\Big] \\
&\leq E_W\Big[\frac{1 - \widehat{B}_\emptyset^2}{4}\Big] + O(\epsilon)
\end{aligned} \tag{15}$$

Combining (13), (14) and (15) we get,

$$\Pr[\text{Acc}] \le E_W\Big[1 + \frac{\widehat{B}_\emptyset}{2} - \frac{\widehat{B}_\emptyset^2}{2}\Big] + O(\epsilon) \tag{16}$$

4 The Final PCP and Proof of Theorem 1.5

Now we are ready to construct the final PCP verifier. Let $\eta \geq 0$ be a parameter to be chosen later. The verifier perform the tests with the following probabilities.

$$\text{Verifier performs test}\begin{cases} T_1 \text{ with probability } \dfrac{4\eta+4}{12+9\eta} \\[2ex] T_2 \text{ with probability } \dfrac{4\eta+4}{12+9\eta} \\[2ex] T_3 \text{with probability } \dfrac{\eta}{12+9\eta} \\[2ex] T_4 \text{ with probability } \dfrac{4}{12+9\eta} \end{cases}$$

We have,

$$\begin{aligned}\Pr[\text{Acc}] \le E_{W,U}\Big[&\Big(\frac{4\eta+4}{12+9\eta}\Big)\frac{5 - \widehat{A}_\emptyset - 2\widehat{B}_\emptyset + 2\widehat{A}_\emptyset\widehat{B}_\emptyset + \widehat{B}_\emptyset^2 + 3\widehat{A}_\emptyset\widehat{B}_\emptyset^2}{8} + \\ &\Big(\frac{4\eta+4}{12+9\eta}\Big)\frac{5 + \widehat{A}_\emptyset - 2\widehat{B}_\emptyset - 2\widehat{A}_\emptyset\widehat{B}_\emptyset + \widehat{B}_\emptyset^2 - 3\widehat{A}_\emptyset\widehat{B}_\emptyset^2}{8} \\ &+\Big(\frac{\eta}{12+9\eta}\Big)(1-\widehat{B}_\emptyset^2) + \Big(\frac{4}{12+9\eta}\Big)\Big(1 + \frac{\widehat{B}_\emptyset}{2} - \frac{\widehat{B}_\emptyset^2}{2}\Big)\Big] + O(\epsilon)\end{aligned}$$

Simplifying and omitting the expectation over U, W we get,

$$\begin{aligned}\Pr[\text{Acc}] &\le \frac{6\eta + 9 - (\widehat{B}_\emptyset^2 + 2\eta\widehat{B}_\emptyset)}{12+9\eta} + O(\epsilon) \\ &= \frac{6\eta + 9 + \eta^2 - (\widehat{B}_\emptyset + \eta)^2}{12+9\eta} + O(\epsilon) \\ &\le \frac{6\eta+9+\eta^2}{12+9\eta} + O(\epsilon) \end{aligned} \tag{17}$$

Now the above expression is minimized for $\eta = \frac{1}{3}$. At this value we get $\Pr[\text{Acc}] \le \frac{20}{27} + O(\epsilon) < \frac{3}{4}$, for ϵ small enough. This gives us an improvement over the previous bound. This proves the main result in the paper, i.e. Theorem 1.5.

5 Conclusion

We constructed a 3-query non-adaptive PCP with perfect completeness and soundness below $\frac{6}{8}$. This makes a partial progress on a fairly long-standing problem. However, it seems difficult to construct a PCP with soundness $\frac{5}{8} + \epsilon$ and new ideas seem to be needed. Or perhaps there are algorithms that do better than $\frac{5}{8}$ on satisfiable 3-CSPs.

We hope that the techniques in this paper would be useful to settle some other open problems regarding the power of PCPs with a small number of queries. For example, can a 4-query PCP achieve soundness below $\frac{1}{2}$ even with imperfect completeness ?

6 Acknowledgments

We sincerely thank Johan Håstad for reading an earlier version of the paper and pointing out a mistake in the analysis.

References

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

[2] S. Arora and S. Safra. Probabilistic checking of proofs : A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[3] M. Bellare and M. Sudan. Improved non-approximability results. *In Proc. of the 26th Annual ACM Symposium on Theory of Computing*, pages 184-193, 1994

[4] M. Bellare, O. Goldreich, and M. Sudan. Free bits, pcps and non-approximability. *Electronic Colloquium on Computational Complexity, Technical Report TR95-024*, 1995.

[5] I. Dinur, V. Guruswami, S. Khot and O. Regev. Vertex cover in k-uniform hypergraphs is hard to approximate within $k - 1 - \epsilon$. *Manuscript*, 2002.

[6] I. Dinur and S. Safra. Importance of being biased. *In Proc. of the 34th Annual ACM Symposium on Theory of Computing*, 2002.

[7] I. Dinur, O. Regev and C. Smyth. On the hardness of coloring 3-uniform hyper-graphs. In *Proc. of the 43rd IEEE Symposium on Foundations of Computer Science*, 2002.

[8] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, Vol 18 (1), 1998, pages 27-36.

[9] V. Guruswami, D. Lewin, M. Sudan, and L. Trevisan. A new characterization of NP with 3 query PCPs. *In Proc. of 39th Annual IEEE Symposium of Foundations of Computer Science*, 1998.

[10] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proc. of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 627–636, 1996.

[11] J. Håstad. Some optimal inapproximability results. In *Proc. of the 29th Annual ACM Symposium on Theory of Computing*, pages 1–10, 1997.

[12] S. Khot. Hardness of coloring 3-colorable 3-uniform hypergraphs. In *Proc. of the 43rd IEEE Symposium on Foundations of Computer Science*, 2002.

[13] R. Raz. A parallel repetition theorem. *SIAM J. of Computing*, 27(3):763–803, 1998.

[14] A. Samorodnitsky, and Luca Trevisan. A PCP characterization of NP with optimal amortized query complexity. *In Proc. of the 32nd Annual ACM Symposium on Theory of Computing*, pages 191-199, 2000.

[15] U. Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *Proc. of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1998.

A Orthonormality of Characters χ_β

We have

$$< \chi_\beta, \chi_\beta > = E_g\Big[\chi_\beta(g)^2\Big] = E_g\Big[\prod_{x\in\beta}\phi_x(g)^2\Big] = \prod_{x\in\beta} E_g\Big[\phi_x(g)^2\Big] = 1$$

where in the last step we note that since $g \in_R \mu_p(\mathcal{M})$, we have $\phi_x(g) = -\sqrt{q/p}$ with probability p and $\phi_x(g) = \sqrt{p/q}$ with probability q. Therefore $E_g\Big[\phi_x(g)^2\Big] = p\cdot q/p + q\cdot p/q = 1$

When $\beta \neq \gamma$, assume w.l.o.g. that $x_0 \in \beta\setminus\gamma$. Then

$$\begin{aligned} < \chi_\beta, \chi_\gamma > &= E_g\Big[\phi_{x_0}(g)\prod_{x\in\beta\setminus\{x_0\}}\phi_x(g)\prod_{x'\in\gamma}\phi_{x'}(g)\Big] \\ &= E_g\Big[\phi_{x_0}(g)\Big]\cdot E_g\Big[\prod_{x\in\beta\setminus\{x_0\}}\phi_x(g)\prod_{x'\in\gamma}\phi_{x'}(g)\Big] \\ &= 0 \end{aligned}$$

where we used the fact that $E_g\Big[\phi_{x_0}(g)\Big] = p\cdot(-\sqrt{q/p}) + q\cdot(\sqrt{p/q}) = 0.$

B Proof of Lemma 2.3

First we prove the following lemma.

Lemma B.1 *For any $x, y \in \mathcal{M}_W$, $x \neq y$,* $\quad \Pr_U[\pi(x) = \pi(y)] \leq \frac{1}{T}$

Proof: Assume that $W = \{C_1, C_2, \ldots, C_{Tu}\}$ and assume w.l.o.g. that assignments x and y differ on the clause C_1. Now consider the process of picking a random U. One picks a subset $S \subseteq W$ of size u at random and replaces the clauses in S by variables. If X denotes this set of variables, then $U = X \cup (W\setminus S)$. With probability $1 - \frac{1}{T}$, we have $C_1 \notin S$, therefore $C_1 \in U$ and consequently, the restrictions of assignments x and y to the set U are distinct. This proves the claim. ∎

Corollary B.2 *For any fixed $W \in \mathcal{W}$ and any $\beta \subseteq \mathcal{M}_W, \beta \neq \emptyset$, over the choice of U, except with probability $\frac{|\beta|}{T}$, there exists $y \in \beta$ such that*

$$\forall\, y' \in \beta, y' \neq y, \qquad \pi(y') \neq \pi(y)$$

Proof: Fix any $y \in \beta$. Apply Lemma B.1 to (y, y') for every $y' \in \beta, y' \neq y$ and then take a union bound. ∎

Now we prove Lemma 2.3. According to Lemma B.1, for any $x \neq y$, their "collision probability" via the map π is small. Therefore, π must map "large" sets to "large" sets with high probability. To be precise, for any $\beta \subseteq \mathcal{M}_W, \beta \neq \emptyset$, we have

$$E_U\Big[\frac{1}{|\pi(\beta)|}\Big] \leq \frac{1}{|\beta|} + \frac{1}{T}$$

Proof:

$$E_U\Big[\frac{1}{|\pi(\beta)|}\Big] \quad\leq\quad E_U\Big[\Pr_{x,y\in\beta}[\pi(x) = \pi(y)]\Big]$$

$$\begin{aligned}
&\leq \quad E_U\left[\frac{1}{|\beta|} + \Pr_{\substack{x,y\in\beta,\\ x\neq y}}[\pi(x)=\pi(y)]\right] \\
&= \quad \frac{1}{|\beta|} + E_{\substack{x,y\in\beta,\\ x\neq y}}\left[\Pr{}_U[\pi(x)=\pi(y)]\right] \\
&\leq \quad \frac{1}{|\beta|} + \frac{1}{T}
\end{aligned}$$

■

C Proof of Lemma 3.3

It suffices to consider the case when $|\pi(\beta)| = 1$, i.e. all elements of $\beta \subseteq \mathcal{M}_W$ map to the same element of $\mathcal{M}_U$. The general case follows by considering every $x \in \pi(\beta)$ separately. In the special case when $|\pi(\beta)| = 1$, let $\pi(\beta) = \{x_0\}$ and we have $c_{x_0} = |\beta|$. We will compute the desired expectation by carefully looking at the way $f(x_0), g(y), h(y)$ are defined in the Test T_1. We consider the two possibilities $\alpha = \{x_0\} = \pi(\beta)$ and $\alpha = \emptyset$ separately.

C.1 Calculation for $\alpha = \pi(\beta)$

We want to compute

$$E_{f,g,h}[\chi_\alpha(f)\chi_\beta(g)\chi_\beta(h)] = E_{f,g,h}[\phi_{x_0}(f)\prod_{y\in\beta}\phi_y(g)\phi_y(h)] \tag{18}$$

Case i : $f(x_0) = 1$ which happens with probability q. In this case, $\phi_{x_0}(f) = \sqrt{p/q}$. After fixing $f(x_0)$, the values $(g(y), h(y))$ are picked independently for different $y \in \beta$. Thus

$$E_{g,h}[\prod_{y\in\beta}\phi_y(g)\phi_y(h)] = \prod_{y\in\beta}E_{g,h}[\phi_y(g)\phi_y(h)] = \Big(E_{g,h}[\phi_{y_0}(g)\phi_{y_0}(h)]\Big)^{|\beta|} \tag{19}$$

where $y_0 \in \beta$ is any fixed element. To compute this expectation, note that after fixing $f(x_0) = 1$, one sets

$$(g(y_0), h(y_0)) = \begin{cases} (1,1) \text{ with probability } q \\ (-1,-1) \text{ with probability } p \end{cases}$$

Thus the last expectation in (19) is $\quad q\cdot\sqrt{p/q}\sqrt{p/q} + p\cdot(-\sqrt{q/p})\cdot(-\sqrt{q/p}) = 1$

Case ii : $f(x_0) = -1$ which happens with probability p. In this case $\phi_{x_0}(f) = -\sqrt{q/p}$ and equation (19) holds again. To compute this expecctation, we note that after fixing $f(x_0) = -1$, we set

$$(g(y_0), h(y_0)) = \begin{cases} (1,-1) \text{ with probability } q \\ (-1,1) \text{ with probability } q \\ (-1,-1) \text{ with probability } p-q \end{cases}$$

Thus the last expectation in (19) is

$$q\cdot\sqrt{p/q}\cdot(-\sqrt{q/p}) + q\cdot(-\sqrt{q/p})\cdot\sqrt{p/q} + (p-q)\cdot(-\sqrt{q/p})\cdot(-\sqrt{q/p}) = -q/p$$

Combining the two cases, the expectation in equation (18) is

$$q\cdot\sqrt{p/q}\cdot(1)^{|\beta|} + p\cdot(-\sqrt{q/p})\cdot(-q/p)^{|\beta|} = \sqrt{pq}\cdot\Big(1-(-q/p)^{|\beta|}\Big) \tag{20}$$

as desired.

C.2 Calculation for $\alpha = \emptyset$

We want to compute

$$E_{f,g,h}[\chi_\emptyset(f)\chi_\beta(g)\chi_\beta(h)] = E_{f,g,h}[\prod_{y\in\beta}\phi_y(g)\phi_y(h)] \tag{21}$$

This calculation is very similar to the calculation for the expectation (18), except that we do not multiply by $\phi_{x_0}(f)$. From equation (20), the desired expectation can be written as

$$q\cdot(1)^{|\beta|} + p\cdot(-q/p)^{|\beta|} = q + p(-q/p)^{|\beta|}$$

D Main Technical Lemma for Test T_2

We state a technical lemma needed for analysis of Test T_2.

Lemma D.1 *If $\beta \subseteq \mathcal{M}_W$, $\alpha \subseteq \pi(\beta)$ and for every $x \in \pi(\beta)$*

$$c_x := |\{\ y \mid y\in\beta,\ \pi(y) = x\}| \tag{22}$$

then in Test T_2 we have

$$E_{f,g,h}[\chi_\alpha(f)\chi_\beta(g)\chi_\beta(h)] = \prod_{x\in\pi(\beta)\setminus\alpha}\Big(p + q(-\frac{q}{p})^{c_x}\Big)\ \cdot\prod_{x\in\alpha}\sqrt{pq}\Big(-1 + (-\frac{q}{p})^{c_x}\Big)$$

In particular, this expectation has magnitude at most 1.

Using this Lemma, the test T_2 can be analyzed in almost the same way as the test T_1.

E Soundness Analysis for Test T_3

We will analyze $E[B(g_1)B(g_2)]$, the remaining two terms are identical.

$$\begin{aligned}
E_{g_1,g_2}[B(g_1)B(g_2)] &= E_{g_1,g_2}[\sum_{\beta,\gamma}\widehat{B}_\beta\widehat{B}_\gamma\chi_\beta(g_1)\chi_\gamma(g_2)] \\
&= \sum_\beta \widehat{B}^2_\beta E_{g_1,g_2}[\chi_\beta(g_1)\chi_\beta(g_2)] \\
&= \sum_\beta \widehat{B}^2_\beta\Big(-\frac{1}{3}\frac{1/4+3\epsilon^2}{1/4-\epsilon^2}\Big)^{|\beta|}
\end{aligned}$$

where we used the fact that $g_1 \in_R \mu_p(\mathcal{M}_W)$, $g_2 \in_R \mu_p(\mathcal{M}_W)$ and therefore expectation is non-zero only when $\beta = \gamma$. The expectation on the last line can be compted explicitly. This proves equation (12).

F Soundness analysis of Test T_4

We want to analyse the term,

$$\begin{aligned}
&\Pr[(B(g_1), B(g_2)) \neq (-1,-1)] \\
=\ & E_{W,g_1,g_2}\Big[\frac{3 + B(g_1) + B(g_3) - B(g_1)B(g_3)}{4}\Big] \\
=\ & \frac{1}{4}\Big(3 + E_{W,g_1}[B(g_1)] + E_{W,g_3}[B(g_3)] - E_{W,g_1,g_3}[B(g_1)B(g_3)]\Big)
\end{aligned} \tag{23}$$

Now, $E_{W,g_1}[B(g_1)] = E_{W,g_3}[B(g_3)] = E_W[\widehat{B}_\emptyset]$. Also,

$$\begin{aligned}
E_{g_1,g_3}[B(g_1)B(g_3)] &= E_{g_1,g_3}[\sum_{\beta,\gamma} \widehat{B}_\beta \widehat{B}_\gamma \chi_\beta(g_1)\chi_\gamma(g_3)] && (24)\\
&= \sum_\beta \widehat{B}_\beta^2 E_{g_1,g_3}[\chi_\beta(g_1)\chi_\beta(g_3)]\\
&= \sum_\beta \widehat{B}_\beta^2 E_{g_1,g_3}[\prod_{y\in\beta} \phi_y(g_1)\phi_y(g_3)]\\
&= \sum_\beta \widehat{B}_\beta^2 [\prod_{y\in\beta} E_{g_1,g_3}[\phi_y(g_1)\phi_y(g_3)]]\\
&= \sum_\beta \widehat{B}_\beta^2 (E_{g_1,g_3}[\phi_{y_0}(g_1)\phi_{y_0}(g_3)])^{|\beta|} && (25)
\end{aligned}$$

The distribution of $(g_1(y_0), g_3(y_0))$ is as follows,

$$(g_1(y_0), g_3(y_0)) = \begin{cases} (-1,1) \text{ with probability } p/2\\ (1,-1) \text{ with probability } p/2\\ (1,1) \text{ with probability } 1-(3p/2)\\ (-1,-1) \text{ with probability } p/2 \end{cases}$$

Therefore, we get,

$$\begin{aligned}
E_{g_1,g_3}[\phi_{y_0}(g_1)\phi_{y_0}(g_3)] &= 2\cdot\frac{p}{2}\cdot\sqrt{\frac{p}{q}}\cdot\Big(-\sqrt{\frac{q}{p}}\Big) + \Big(1-\frac{3p}{2}\Big)\cdot\frac{p}{q} + \frac{p}{2}\cdot\frac{q}{p} && (26)\\
&= -p + \frac{p}{q} - \frac{3p^2}{q} + \frac{q}{2}\\
&= \frac{-2pq + 2p - 6p^2q + q^2}{2q}\\
&= \frac{-2p(1-p) + 2p - 6p^2(1-p) + (1-p)^2}{2q}\\
&= \frac{6p^3 - 3p^2 - 2p + 1}{2q}\\
&= \frac{1}{2q}\Big[6\Big(\frac{1}{2}+\epsilon\Big)^3 - 3\Big(\frac{1}{2}+\epsilon\Big)^2 - 2\Big(\frac{1}{2}+\epsilon\Big) + 1\Big]\\
&= u(\epsilon) && (27)
\end{aligned}$$

where, $u(\epsilon)$ is $O(\epsilon)$. Using this we get,

$$\begin{aligned}
E_{g_1,g_3}[B(g_1)B(g_3)] &= \sum_\beta \widehat{B}_\beta^2 (E_{g_1,g_3}[\phi_{y_0}(g_1)\phi_{y_0}(g_3)])^{|\beta|}\\
&= \sum_\beta \widehat{B}_\beta^2 (u(\epsilon))^{|\beta|}\\
&= \widehat{B}_\emptyset^2 + O(\epsilon) && (28)
\end{aligned}$$

which gives us that,

$$\Pr[(B(g_1), B(g_2)) \neq (-1,-1)] \leq E_W\Big[\frac{3 + 2\widehat{B}_\emptyset - \widehat{B}_\emptyset^2}{4}\Big] + O(\epsilon) \quad (29)$$

Now we analyse the term,

$$\begin{aligned}
& \Pr[B(g_2) = 1 \ \&\ B(g_3) = -1] \\
= \;& E_{W,g_2,g_3}\Big[\frac{1 + B(g_2) - B(g_3) - B(g_2)B(g_3)}{4}\Big] \\
= \;& \Big(\frac{1 + E_{W,g_2}[B(g_2)] - E_{W,g_3}[B(g_3)] - E_{W,g_2,g_3}[B(g_2)B(g_3)]}{4}\Big) \qquad (30)
\end{aligned}$$

Now, $E_{W,g_2}[B(g_2)] = E_{W,g_3}[B(g_3)] = E_W[\widehat{B}_\emptyset]$. We also notice that the distribution of $(g_2(y_0), g_3(y_0))$ and $(g_1(y_0), g_3(y_0))$ is identical and therefore we get that,

$$E_{g_2,g_3}[B(g_2)B(g_3)] = \widehat{B}_\emptyset^2 + O(\epsilon) \qquad (31)$$

which gives us that,

$$\Pr[B(g_2) = 1 \ \&\ B(g_3) = -1] \leq E_W\Big[\frac{1 - \widehat{B}_\emptyset^2}{4}\Big] + O(\epsilon) \qquad (32)$$