

Parallel Repetition of Fortified Games

Dana Moshkovitz *

April 26, 2014

Abstract

The Parallel Repetition Theorem upper-bounds the value of a repeated (tensored) two prover game $\mathcal{G}^{\otimes k}$ in terms of the value of the game \mathcal{G} and the number of repetitions k . Contrary to what one might have guessed, the value of $\mathcal{G}^{\otimes k}$ is not $val(\mathcal{G})^k$, but rather a more complicated expression depending on properties of \mathcal{G} . Indeed, there are numerous works aiming to understand the value of repeated games, both for general games \mathcal{G} and for special cases. A case of particular interest is that of projection games, where the answer of the first prover determines at most one acceptable answer for the second prover.

In this work we give a simple transformation, which we call “fortification”, that can be applied to any projection game. We show that for fortified games \mathcal{G} , the value of the k -repeated game is approximately $val(\mathcal{G})^k$. This results in nearly a quadratic improvement in the size of projection PCP with the lowest error known today. Unlike previous proofs of the parallel repetition theorem that relied on information theory or linear algebra, our proof is purely combinatorial and quite short.

We then discuss the problem of derandomizing parallel repetition, and the limitations of the fortification idea in this setting. We point out a connection between the problem of derandomizing parallel repetition and the problem of composition. This connection could shed light on the so-called Projection Games Conjecture, which asks for projection PCP with minimal error.

*dmoshkov@csail.mit.edu. Department of Electrical Engineering and Computer Science, MIT. This material is based upon work supported by the National Science Foundation under Grant Number 1218547.

1 Introduction

1.1 The Parallel Repetition Theorem

In a two prover game \mathcal{G} , a verifier picks at random a pair of questions (x, y) from a specified set of possible questions, sends x to the first prover, and sends y to the second prover; the first prover replies with an answer a , and the second prover replies with an answer b ; the verifier, knowing x and y , and having inspected both a and b , decides whether to accept or reject. The value the prover strategies achieve is the probability that the verifier accepts. The *value* of \mathcal{G} , denoted $val(\mathcal{G})$, is the maximum of this quantity over all prover strategies.

A k -repetition (tensor) of a game \mathcal{G} is the game $\mathcal{G}^{\otimes k}$, in which the verifier picks at random k question pairs $(x_1, y_1), \dots, (x_k, y_k)$; sends one prover x_1, \dots, x_k , and sends the other prover y_1, \dots, y_k ; the first prover replies with a_1, \dots, a_k , and the second prover replies with b_1, \dots, b_k ; the verifier checks that it would have accepted in all k tests.

A long line of work analyzes how $val(\mathcal{G}^{\otimes k})$ depends on $val(\mathcal{G})$ and k . Clearly, $val(\mathcal{G}^{\otimes k}) \geq val(\mathcal{G})^k$, since the provers can follow the same strategy in each one of the k rounds. One might guess that $val(\mathcal{G}^{\otimes k}) = val(\mathcal{G})^k$, but this turns out to be false [16, 13, 29, 15]. In a breakthrough result, Raz [24] showed that $val(\mathcal{G}^{\otimes k})$ does exhibit an exponential decay with k when $val(\mathcal{G}) < 1$ (below Σ_X is the set of possible answers a of the first prover, while Σ_Y is the set of possible answers b of the second prover):

Theorem 1 (Raz’s Parallel Repetition Theorem [24]). *There exists $W : [0, 1] \rightarrow [0, 1]$ such that $W(x) < 1$ for $x < 1$, and*

$$val(\mathcal{G}^{\otimes k}) \leq (W(val(\mathcal{G})))^{k/\log(|\Sigma_X||\Sigma_Y|)}.$$

Interestingly, the dependence of the exponent in $|\Sigma_X|$ and $|\Sigma_Y|$ is inherent [15]. Disappointingly, the base of the exponent is quite far from $val(\mathcal{G})$. In fact, in Raz’s theorem, $W(val(\mathcal{G}))$ is close to 1 even when $val(\mathcal{G})$ is close to 0! Many works simplified and improved the parameters of the Parallel Repetition Theorem for general games [18], as well as for games with a special structure, most notably *projection games* [23, 12] and *expanding projection games* [26, 12]. Before we describe the main results of those papers, let us discuss projection games and their importance.

Arguably, the most important application of the Parallel Repetition Theorem is soundness amplification for projection games. In this paper it will be convenient for us to consider the following definition of a projection game:

Definition 1.1 (Projection game). *A projection game is defined by a bipartite graph $G = (X, Y, E)$, alphabets Σ_X and Σ_Y and functions $\{\pi_e : \Sigma_X \rightarrow \Sigma_Y\}_{e \in E}$, called “projections”. In the game, the verifier picks uniformly at random $y \in Y$, and two edges $e = (x, y), e' = (x', y) \in E$, sends x to the first prover, and sends x' to the second prover; the first prover replies with $a \in \Sigma_X$, and the second prover replies with $a' \in \Sigma_X$; the verifier accepts if $\pi_e(a) = \pi_{e'}(a')$.*

Remark 1.1. *While Definition 1.1 is often useful (e.g., [19]), the more standard definition of projection games is as follows: the verifier picks uniformly at random an edge $(x, y) \in E$, sends x to the first prover, and sends y to the second prover; the first prover replies with $a \in \Sigma_X$, and the second prover replies with $b \in \Sigma_Y$; the verifier accepts if $\pi_e(a) = b$. If \mathcal{G}' is the game in Definition 1.1, and \mathcal{G} is the game we defined here, then $val(\mathcal{G})^2 \leq val(\mathcal{G}') \leq val(\mathcal{G})$.*

The PCP Theorem, in the form that is most useful for hardness of approximation, states that it is NP-hard, given a projection game \mathcal{G} , to distinguish the case where $val(\mathcal{G}) = 1$ from the case where $val(\mathcal{G}) \leq \varepsilon$. The parameter ε is called the *soundness error* of the PCP. Since parallel repetition of a projection game is itself a projection game, the Parallel Repetition Theorem, when applied on the basic PCP Theorem [5, 4, 3, 2], yields a projection PCP theorem with arbitrarily small soundness error. Projection PCP with low soundness error is the basis of most of the best NP-hardness of approximation results we have today. In particular, it is the basis of the hardness results in Håstad’s seminal paper [17].

An unfortunate aspect of parallel repetition is that it raises the size of the game (i.e., the number of possible questions and question pair; we use the notation $size(\mathcal{G})$) to the k ’th power. In particular, if k is super-constant, one gets a super-polynomial reduction from SAT to the repeated game, rather than an NP-hardness result. When one assumes that solving SAT on inputs of size n requires time $exp(n)$ (“The Exponential Time Hypothesis”), the reductions obtained using parallel repetition only yield time lower bounds of the form $exp(N^{1/k})$ for input size N . Due to this state of affairs, parallel repetition is used mostly for constant k .

One of the most important open problems in approximability is to construct projection games with error that is inverse polynomial in the size of the game. The author named this problem “The Projection Games Conjecture” in [20]. One of the most notable applications of this conjecture is an NP-hardness result for approximating CLOSEST-VECTOR-PROBLEM in lattices to within polynomial factors (see [20] for a discussion of more applications). The lowest soundness error known today is $\varepsilon = 1/(\log N)^c$, for any constant $c > 0$, when N is the size of the game. This is by a reduction of the author and Raz [21] from SAT on input of size n to projection games of size $N = n^{1+o(1)}$, where the soundness error is $\varepsilon = 1/(\log N)^\beta$ for some $\beta > 0$. When this game is repeated in parallel, the soundness error can be $\varepsilon = 1/(\log N)^c$ for any constant $c > 0$, while the size is raised to $O(c/\beta)$ as shown by Dinur and Steurer [12].

Understanding the significance of projection games, we now turn to review what is known about their repetition. Rao [23] showed that in the projection case $val(\mathcal{G}^{\otimes k})$ does not depend on the number of possible answers of the provers. The state of the art results for projection games are from the beautiful recent paper by Dinur and Steurer [12]:

Theorem 2 (Parallel Repetition Theorem for projection games [12]). *For any projection game \mathcal{G} as in Remark 1.1,*

$$val(\mathcal{G}^{\otimes k}) \leq 2^k \cdot val(\mathcal{G})^{k/2}.$$

Using the relation between the projection games of Definition 1.1 and the projection games of Remark 1.1 (explained in Remark 1.1), Theorem 2 yields a (weaker) parallel repetition theorem for projection games as in Definition 1.1. It is quite possible that the techniques of Dinur and Steurer yield bounds as in Theorem 2 for games as in Definition 1.1 too.

The case of $val(\mathcal{G}) = 1 - \epsilon$ for small $\epsilon > 0$ is extremely well-studied due to its importance to soundness amplification of unique games (unique games are projection games where the projections are one-to-one). In this case, for projection games as in Remark 1.1, the works of Rao [23] and Dinur-Steurer [12] yield

$$val(\mathcal{G}^{\otimes k}) \leq \min \left\{ 1 - \Omega(\sqrt{k} \cdot \epsilon), (1 - \epsilon/2)^{\Omega(\epsilon k)} \right\}.$$

For unique games of interest, the $1 - \Omega(\sqrt{k} \cdot \epsilon)$ is tight¹. For projection games on expanders,

¹Up to the constants in the $\Omega(\cdot)$, the upper bound in the case of $1 - \Omega(\sqrt{k} \cdot \epsilon)$ is tight, as Raz [25] showed a unique game \mathcal{G} for which $val(\mathcal{G}^{\otimes k}) \geq 1 - O(\sqrt{k} \cdot \epsilon)$. More generally, Barak et al [6] analyze the behavior of general unique games under parallel repetition.

Raz and Rosen [26] prove $\text{val}(\mathcal{G}^{\otimes k}) \leq (1 - \epsilon)^{\Omega(k)}$. We note that in all the aforementioned results, either explicitly or hiding in $\Omega(\cdot)$, is the fact that *not all repetitions count*. That is, in many of the k repetitions, the provers may win with probability 1 conditioned on winning other rounds. This phenomenon is known to actually occur – there are unique games with $\text{val}(\mathcal{G}^{\otimes 2}) = \text{val}(\mathcal{G})$ [13].

1.2 Our Contribution

In this paper we suggest a new approach to parallel repetition – rather than try to explore the subtle behavior of general projection (or unique) games under repetition, make the games into ones that behave well under repetition. We present a simple combinatorial transformation on projection games, which we call “fortification”, and which was inspired by ideas in combinatorial construction of error correcting codes. Fortification preserves the projection structure, while increasing $|X|$ and $|\Sigma_X|$ in a controlled way. We then show that for fortified projection games \mathcal{G} , the value of the k -repeated game is, approximately, $\text{val}(\mathcal{G})^k$, i.e.,

$$\text{val}(\mathcal{G})^k \leq \text{val}(\mathcal{G}^{\otimes k}) \leq \text{val}(\mathcal{G})^k + \text{err},$$

where the small additive error err can be made arbitrarily small by fortification. Notably, our analysis of parallel repetition is much simpler than all existing analyses. Unlike Raz’s proof, our analysis does not require information theory, or clever choices of sub-games à la Razborov, nor does it require linear algebra as in the recent analysis of Dinur and Steurer for projection games. Using our paradigm we are able to reprove the state of the art projection PCP Theorem, achieving soundness error $1/(\log n)^c$ for any $c > 0$ [21, 12]. In fact, due to our tighter bound on $\text{val}(\mathcal{G}^{\otimes k})$, for any desirable $c > 0$, our construction saves roughly a quadratic factor in the size compared to the result of [12].

We then turn to explore the possibility of obtaining stronger projection PCP theorems using our ideas. The bottleneck here is the large size blow-up introduced by parallel repetition, and hence the question is whether parallel repetition could be “derandomized” for appropriately fortified games. While we do not know how to extend our fortification ideas to this case (we explain the difficulty in Section 6), we are able to point out an intriguing connection between the problem of derandomizing parallel repetition and the well-studied problem of *composition* of two prover games. The connection – which holds for general two prover games – is that both problems share a combinatorial hard core. Since repetition and composition constitute the two existing approaches to the Projection Games Conjecture (error reduction and alphabet reduction, respectively), the connection sheds light on the difficulty of proving the conjecture.

1.3 Previous Work on Derandomizing Parallel Repetition

For simplicity, let us continue to denote the repeated game $\mathcal{G}^{\otimes k}$, with the understanding that the k rounds may be correlated. Feige and Kilian [14] showed that in the derandomized case, for $\text{val}(\mathcal{G}^{\otimes k}) \leq \delta$, it must be the case that the degrees in \mathcal{G} ’s graph are at least $\approx 1/\delta$ (under an assumption on \mathcal{G} they call *softness*, which indeed holds in the cases of interest). The degrees in the graph correspond to the uncertainty each prover has with respect to the questions of the other prover. For any two prover game in which each of the verifier’s tests can be satisfied by itself, if the graph is bi-regular, and one of the sides has degree D , then the value of the game is at least $1/D$. The interesting feature of Feige and Kilian’s result is that they relate the value of $\mathcal{G}^{\otimes k}$ to the degree in \mathcal{G} . Taking this restriction into account, one might hope for

a derandomized parallel repetition whose size is $size(\mathcal{G}) \cdot (1/\delta)^{O(k)}$. If such a derandomization had been available, it would have given projection PCP with soundness error $\delta = 2^{-(\log n)^\beta}$ for some constant $\beta > 0$.

However, so far there has been little progress even on suggesting candidate games \mathcal{G} with a derandomization $\mathcal{G}^{\otimes k}$. The two exceptions have been *free games* [28] and *linear games* [11]. A free game is a game that in which the questions of the two provers are independent. A linear game is a game in which the questions correspond to points in a linear space, and the verifier's tests correspond to linear sub-spaces. For free games, Shaltiel [28] analyzed repetition where the dependence between the randomness the verifier needs in order to reach a given target $val(\mathcal{G}^{\otimes k}) \leq \delta$, and the number of possible answers of the provers, is improved (recall that for general two prover games $val(\mathcal{G}^{\otimes k})$ depends on the number of possible answers of the provers). The size of the game in Shaltiel's theorem is still $(size(\mathcal{G}))^{\Omega(\log(1/\delta))}$. For linear games, Dinur and Meir [11] analyzed derandomized repetition, but where the soundness error does not decrease exponentially. For both types of games, known transformations from general games incur a large blow-up in the size (for free games [1]) or in the soundness error (for linear games [11]). In fact, for free games it was proved that the size blow-up is inherent [1]. Hence, neither free games nor linear games seem useful for making further progress toward the Projection Games Conjecture.

2 Preliminaries

Let $Dist$ be a distribution over a space X . The *entropy* in the distribution, denoted $H(Dist)$, is $\sum_{x \in X} Dist(x) \log(1/Dist(x))$. We say that $Dist$ has *min-entropy* at least k , and denote $H_\infty(Dist) \geq k$, if no $x \in X$ has probability higher than 2^{-k} . A (δ, ε) -extractor is a bi-regular bipartite graph $H = (X, Y, E)$, such that for every distribution $Dist$ over X with min-entropy at least $\log(\delta|X|)$, the distribution on Y obtained by picking x according to $Dist$ and picking a uniformly random neighbor $y \in Y$ of x is ε -close to uniform over Y .

Lemma 2.1 (Extractor constructions, [27]). *There exist (δ, ε) -extractors $G = (X, Y, E)$ such that $|X| = O(|Y|/\delta)$ and each vertex in X has degree $D = O(\log(1/\delta) \cdot (1/\varepsilon)^2)$. Moreover, there exist explicit constructions achieving $|X| = O(|Y|/\delta)$ and $D = D(\delta, \varepsilon) = \exp(\text{poly log log}(1/\delta)) \cdot (1/\varepsilon)^2$.*

A two prover game \mathcal{G} is defined by a set X of questions to the first prover, a set Y of questions to the second prover, an alphabet Σ_X for the answers of the first prover, and an alphabet Σ_Y for the answers of the second prover. In addition, there is a distribution μ over question pairs $X \times Y$, and a predicate $V \subseteq X \times Y \times \Sigma_X \times \Sigma_Y$. The verifier picks (x, y) from μ sends x to the first prover and y to the second prover; receives $a \in \Sigma_X$ from the first prover and $b \in \Sigma_Y$ from the second prover; then accepts or reject based on $V(x, y, a, b)$.

One often considers the bipartite graph associated with \mathcal{G} . This is the graph $G = (X, Y, E)$ on vertex sets X and Y , where the edges are the question pairs (x, y) with non-zero probability in μ . When one refers to degrees in \mathcal{G} , the intention is degrees in G . Typically, and by default in this paper, μ is uniform over E , and a question pair (x, y) from μ is such that x is uniform over X , while y is uniform over Y . The value achieved by certain prover strategies is the probability that the verifier accepts. The value of \mathcal{G} , denoted $val(\mathcal{G})$, is the maximum of this quantity over all prover strategies. The size of \mathcal{G} , denoted $size(\mathcal{G})$, is $|X| + |Y| + |E|$. The randomness of the verifier is $\log |E|$.

3 Fortification

If \mathcal{G}' is a sub-game of a game \mathcal{G} obtained by picking only a subset of the possible question pairs of the verifier, then the value of \mathcal{G}' can be much higher than the value of the original game \mathcal{G} . Fortified games \mathcal{G} are such that certain large sub-games \mathcal{G}' of \mathcal{G} have $val(\mathcal{G}') \approx val(\mathcal{G})$. The largeness is with respect to the upper bound $val(\mathcal{G}^{\otimes k}) \leq \delta$ we wish to obtain. Note that the requirement that *every* sub-game \mathcal{G}' of fraction δ in \mathcal{G} has $val(\mathcal{G}') < 1$ is equivalent to saying that $val(\mathcal{G}) < \delta$. Hence, it is important that we focus on a family of large sub-games, rather than on all large sub-games.

Specifically, we focus on *rectangular* sub-games, defined as follows: If S is an event depending on the questions to the first prover, and T is an event depending on the questions to the second prover, then the rectangular sub-game $\mathcal{G}_{|S \times T}$ is the game \mathcal{G} conditioned on the questions to the provers satisfying S and T , respectively. We also extend this definition to convex combinations over events $\{S_i \times T_i\}_i$. Here we first pick $S_i \times T_i$ from the combination, then consider the relevant sub-game. The value of the game is the convex combination of the values of the sub-games.

We say that a projection game on extractors is *fortified* if large rectangular sub-games have low value:

Definition 3.1. *The δ -fortified value of a game \mathcal{G} , denoted $val_\delta(\mathcal{G})$, is the maximum of $val(\mathcal{G}_{\{S_i \times T_i\}_i})$ over all convex combinations $\{S_i \times T_i\}_i$ such that for every i , $\Pr[S_i], \Pr[T_i] \geq \delta$.*

Every game can be fortified easily, without increasing the size or the alphabets of the game too much. Fortification does not change the value of the game, only makes sure that the value of large rectangular sub-games is similar to the value of the overall game:

Lemma 3.1 (Fortification). *A projection game \mathcal{G} on a bi-regular graph $G = (X, Y, E)$, with alphabets Σ_X, Σ_Y , and projections $\{\pi_e\}_{e \in E}$ can be efficiently converted to a game \mathcal{G}^* on a graph $G^* = (X^*, Y, E^*)$ with alphabets Σ_X^D, Σ_Y , and projections $\{\pi_{e^*}\}_{e^* \in E^*}$, such that*

1. G^* is a (δ, ε) -extractor.
2. $D = D(\delta, \varepsilon)$ as in Lemma 2.1.
3. The size of G^* is polynomial in the size of G , $1/\delta$ and $1/\varepsilon$.
4. $val(\mathcal{G}^*) = val(\mathcal{G})$.
5. $val_\delta(\mathcal{G}^*) \leq val(\mathcal{G}) + 2\varepsilon$.

Proof. Let $H = (X^*, X, E_H)$ be a (δ, ε) -extractor. By Lemma 2.1, such can be constructed so $|X^*| = \text{poly}(|X|, 1/\delta)$ and each vertex in X^* has $D = D(\delta, \varepsilon)$ neighbors in X . Let E^* contain an edge $e^* = (x^*, y)$ for every pair $(x^*, x) \in E_H$ and $e = (x, y) \in E$. An assignment \vec{a} to x^* consists of assignments to all D neighbors of x^* in H , and in particular some $a(x)$ to x . The projection on the edge e^* is $\pi_{e^*}(\vec{a}) = \pi_e(a(x))$. Note that G^* is a (δ, ε) -extractor, and that $size(\mathcal{G}^*)$ is $O(size(\mathcal{G})/\delta)$. Consider the game \mathcal{G}^* associated with the graph G^* , alphabets Σ_X^D, Σ_Y and projections $\{\pi_{e^*}\}_{e^* \in E^*}$. In this game, the verifier picks uniformly at random $y \in Y$ and $x^*, (x^*)' \in X^*$ such that $e^* = (x^*, y) \in E^*$ and $(e^*)' = ((x^*)', y) \in E^*$. Upon receipt of answers $\vec{a}, (\vec{a})' \in \Sigma_X^D$, the verifier checks that $\pi_{e^*}(\vec{a}) = \pi_{(e^*)'}((\vec{a})')$.

We have $val(\mathcal{G}) \leq val(\mathcal{G}^*)$, since any strategy $a : X \rightarrow \Sigma_X$ for \mathcal{G} induces a strategy for \mathcal{G}^* achieving the same value: given $x^* \in X^*$, the answer \vec{a} of the prover assigns every neighbor

$x \in X$ of x^* in H the answer $a(x)$. Moreover, $\text{val}(\mathcal{G}^*) \leq \text{val}(\mathcal{G})$, since every strategy in \mathcal{G}^* induces a randomized strategy in \mathcal{G} achieving the same value in expectation (and hence there exists a strategy for \mathcal{G} achieving this value): given $x \in X$, the prover picks at random a neighbor $x^* \in X^*$ of x in H , and responds according to the strategy for x^* .

Let $\{S_i \times T_i\}_i$ be a convex combination of events, where for all i , the event S_i depends only on x^* , the event T_i depends only on $(x^*)'$ and $\Pr[S_i], \Pr[T_i] \geq \delta$. We'd like to prove that $\text{val}(\mathcal{G}^*)_{|\{S_i \times T_i\}_i} \leq \text{val}(\mathcal{G}) + 2\varepsilon$. Select at random i , and $y \in Y$, $x^*, (x^*)' \in X^*$, conditioned on the events S_i and T_i . Let $x, x' \in X$ be the vertices for which $(x, y), (x', y) \in E$, while $(x^*, x), ((x^*)', x') \in E_H$. By the extractor property of H , the vertices x and x' are each ε -close to uniform over X . The claim that $\text{val}_\delta(\mathcal{G}^*) \leq \text{val}(\mathcal{G}) + 2\varepsilon$ follows from the definition of \mathcal{G} and \mathcal{G}^* . \square

We wish to emphasize that *not every projection game on extractors is fortified*. Indeed, if we take any projection game on extractors and change the projections on edges touching δ fraction of the vertices so they are trivially satisfied, we hardly change the value of the game, but we make sure that the game is not fortified.

Fortification increases the size by a factor $O(1/\delta)$, where we fortify against sub-games of fraction δ . When repeating the game for k rounds, the size increases by a factor $\approx (1/\delta)^k$. However, due to fortification, $\text{val}(\mathcal{G}^{\otimes k})$ decreases exponentially with k , rather than with $k/2$. Hence, to reach a target $\text{val}(\mathcal{G}^{\otimes k}) \leq \delta$ previous methods required twice as many rounds k as we do, and thus the right comparison is between size $\gg (\text{size}(\mathcal{G}))^{2k}$ for previous methods and size $\approx (\text{size}(\mathcal{G})/\delta)^k$ for us. Since typically $\text{size}(\mathcal{G})$ is much larger than $1/\delta$, our method yields better size than before.

Fortification also raises the size of the alphabet Σ_X to a power $D = D(\delta, \varepsilon)$. This price is quite tolerable since in order to reach a target $\text{val}(\mathcal{G}^{\otimes k}) \leq \delta$, we take $k = \Theta(\log(1/\delta))$, and in repetition, Σ_X is raised to a power k anyway ($1/\varepsilon$ is typically smaller than, or comparable to, $\log(1/\delta)$). Moreover, there is a hope that the large alphabet due to fortification could be re-used for the repeated tests.

4 A Parallel Repetition Theorem

In this work we suggest to prove parallel repetition theorems assuming that the underlying game is fortified:

Theorem 3 (Parallel repetition). *If \mathcal{G} is a projection game on a (δ, ε) -extractor where $\delta \leq \varepsilon^4 |\Sigma_Y|^{-(k-1)}$, then there exists $\text{err} = O(k\varepsilon)$ for which*

$$\text{val}(\mathcal{G}^{\otimes k}) \leq (\text{val}_{\delta/\varepsilon^3}(\mathcal{G}) + \text{err})^k.$$

For our parallel repetition theorem we need fortification against sub-games of fraction $\delta \approx |\Sigma_Y|^k$. Crucially, existing constructions of projection games \mathcal{G} with $\text{val}(\mathcal{G}) \leq \varepsilon$ have $|\Sigma_Y| \approx 1/\varepsilon$. Moreover, there is a simple transformation of Dinur and Harsha [10] based on code concatenation that decreases Σ_Y for any projection game:

Lemma 4.1 (Σ_Y reduction [10]). *Any projection game \mathcal{G} with $\text{val}(\mathcal{G}) \leq \varepsilon$ as before can be efficiently transformed to a new projection game \mathcal{G}' where $\text{size}(\mathcal{G}') \leq \text{size}(\mathcal{G}) \cdot \log |\Sigma_Y| \cdot \text{poly}(1/\eta)$, $\text{val}(\mathcal{G}') \leq \varepsilon + O(\eta)$ and the new $|\Sigma_Y|$ is $\text{poly}(1/\eta)$.*

The proof of Theorem 3 is in Section 5. Using fortification as in Lemma 3.1 and our parallel repetition theorem in Theorem 3, we obtain a projection PCP theorem with soundness error $1/(\log n)^c$ for any constant $c \geq 1$, matching the recent result of Dinur and Steurer. It is convenient to phrase this theorem in terms of the computational problem LABEL-COVER, where given a bipartite graph $G = (X, Y, E)$, alphabets Σ_X, Σ_Y , and projections on the edges $\{\pi_e\}_{e \in E}$, the goal is to find assignments $a : X \rightarrow \Sigma_X, b : Y \rightarrow \Sigma_Y$ that satisfy as many of the projections π_e as possible, i.e., $\pi_e(a(x)) = b(y)$ for as many edges $e = (x, y) \in E$ as possible.

Corollary 4.2 (Projection PCP). *For any constant $c \geq 1$, it is NP-hard to distinguish, given a LABEL-COVER instance of size n , between the case where all edges can be satisfied and the case where only $1/(\log n)^c$ fraction of the edges could be satisfied.*

Proof. Start with the PCP theorem of the author and Raz [21], showing that for some small constants $\beta, \beta' > 0$, it is NP-hard to distinguish the case that in a given LABEL-COVER instance \mathcal{G} on a graph of size n and alphabets $|\Sigma_X| \leq \exp((\log n)^\beta), |\Sigma_Y| = O(\log n)$, all edges could be satisfied from the case where only $1/(\log n)^{\beta'}$ fraction of the edges can be satisfied. Apply fortification (Lemma 3.1) on \mathcal{G} with parameters $\delta = 1/\text{poly log } n$ for a sufficiently large poly log n and $\varepsilon = 1/(\log n)^{\beta'}$, to get a LABEL-COVER instance \mathcal{G}^* of size $O(n/\delta)$ and alphabets $|\Sigma_Y| \leq O(\log n)$ and $|\Sigma_X^D| = \exp(D(\log n)^\beta) \leq \text{poly } n$. Consider the repeated game $(\mathcal{G}^*)^{\otimes k}$, for k that is a sufficiently large constant.

The alphabets of $(\mathcal{G}^*)^{\otimes k}$ are of sizes $|\Sigma_Y|^k \leq \text{poly log } n$ and $|\Sigma_X|^{Dk} \leq \text{poly } n$. The size of $(\mathcal{G}^*)^{\otimes k}$ is polynomial in n . If $\text{val}(\mathcal{G}) = 1$, then $\text{val}((\mathcal{G}^*)^{\otimes k}) = 1$. Meanwhile, by our parallel repetition theorem (Theorem 3), provided that $\delta = 1/\text{poly log } n$ for a sufficiently large poly log n , if $\text{val}(\mathcal{G}) \leq 1/(\log n)^{\beta'}$, then $\text{val}((\mathcal{G}^*)^{\otimes k}) \leq O((\log n)^{-\beta'k})$. \square

Corollary 4.2 implies that approximating SET-COVER on inputs of size n better than $(1 - \alpha)\ln n$ requires time $2^{n^{\Omega(\alpha)}}$ assuming the Exponential Time Hypothesis. This matches the best known approximation algorithms for SET-COVER (See [20] for details).

5 Proof of Parallel Repetition Theorem

Set $\hat{\varepsilon} = \text{val}_{\delta/\varepsilon^3}(\mathcal{G}) + ck\varepsilon$ where c is a sufficiently large constant. We assume that the value of $\mathcal{G}^{\otimes k}$ is larger than $\hat{\varepsilon}^k$, and wish to arrive at a contraction. If for each round $i = 1, \dots, k$ the provers fix strategies $a_i : X \rightarrow \Sigma_X$ and $a'_i : X \rightarrow \Sigma_X$ that depend only on the questions of the i 'th round, then the value they achieve is at most $\text{val}(\mathcal{G})^k$ by definition. However, the provers may answer the questions in each round based also on the questions to the other rounds. For example, suppose that the provers win in the first round iff there is a certain relation between their questions in the second round. Then, a-priori, it is possible that they win the first round with probability $\text{val}(\mathcal{G})$, and conditioned on winning the first round, win the second round with probability much larger than $\text{val}(\mathcal{G})$, since the second round is effectively played in a sub-game of the base game that potentially can be won with higher probability. We show that thanks to fortification this cannot happen.

We first identify a list $I \subset \{1, \dots, k\}$ of *influential rounds*, where the provers win with probability at most $\hat{\varepsilon}$. The intuition is that these are the rounds where the provers try to make gains that will help them win other rounds better than expected. Note that the list cannot contain all rounds – as otherwise the total probability of winning all k rounds would have been too small. For $i \in \{1, \dots, k\}$, Let W_i be the event that the provers win the i 'th round. For $I \subseteq \{1, \dots, k\}$ let W_I be the event that the provers win all the rounds in I .

Lemma 5.1 (Influential rounds). *There exists $I \subseteq \{1, \dots, k\}$, $l \doteq |I| < k$, such that for every $i \in \{1, \dots, k\} - I$, it holds that $\Pr [W_i | W_I] > \hat{\varepsilon}$.*

Proof. Construct I as follows: start with $I = \phi$, and while there is still $i \in \{1, \dots, k\}$ such that $\Pr [W_i | W_I] \leq \hat{\varepsilon}$, add i to I .

By construction, for every $i \in \{1, \dots, k\} - I$, it holds that $\Pr [W_i | W_I] > \hat{\varepsilon}$. We claim that at each step $\Pr [W_I] \leq \hat{\varepsilon}^{|I|}$. This is certainly true when $|I| = 0$. Moreover, if it is true for I , it continues to be true if we decide to insert i to I , as

$$\Pr [W_{I \cup \{i\}}] = \Pr [W_i | W_I] \cdot \Pr [W_I] \leq \hat{\varepsilon} \cdot \hat{\varepsilon}^{|I|} = \hat{\varepsilon}^{|I \cup \{i\}|}.$$

Since $\hat{\varepsilon}^{|I|} \geq \Pr [W_I] \geq \Pr [W_{1..k}] > \hat{\varepsilon}^k$, necessarily $|I| < k$. \square

Let $W = W_I$ be the event that the provers win all l influential rounds. By Lemma 5.1, conditioned on W , the provers win each of the other rounds with probability larger than $\hat{\varepsilon}$. We will argue that this cannot happen.

Consider a fixing of the questions to the provers in the influential rounds,

$$\{y_j\}_{j \in I} \subseteq Y, \{x_j\}_{j \in I} \subseteq X, \{x'_j\}_{j \in I} \subseteq X.$$

Let $(\mathcal{G}^{\otimes k})'$ be the sub-game associated with this fixing. Let W' be the event of winning all l influential rounds in $(\mathcal{G}^{\otimes k})'$. Note that even though this event depends only on the answers in l rounds, it may in fact depend on all the questions, including those in the remaining $k - l$ rounds. We further partition $(\mathcal{G}^{\otimes k})'$ into sub-games: There is a sub-game per choice of l labels from Σ_Y for $\{y_j\}_{j \in I}$ of the influential rounds. If the labels are denoted $\{\sigma_j\}_{j \in I} \subseteq \Sigma_Y$, then the event $S_{\vec{\sigma}}$ is that the first prover agrees with the choice, i.e., for all $j \in I$, we have $\pi_{e_j}(a_j) = \sigma_j$, while the event $T_{\vec{\sigma}}$ is that the second prover agrees with the choice, i.e., for all $j \in I$, we have $\pi_{e'_j}(a'_j) = \sigma_j$. Note that there are only $|\Sigma_Y|^l$ sub-games $S_{\vec{\sigma}} \times T_{\vec{\sigma}}$. Hence, for any $0 < \delta' < 1$, the contribution to W' from sub-games $S_{\vec{\sigma}} \times T_{\vec{\sigma}}$ where $S_{\vec{\sigma}}$ or $T_{\vec{\sigma}}$ have probability at most δ' is at most $\delta' |\Sigma_Y|^l$.

Set $\delta' = \delta/\varepsilon^3$. For the remainder of the analysis, we focus on a choice of $S_{\vec{\sigma}}$ and $T_{\vec{\sigma}}$ whose probabilities are at least δ' . Let $(\mathcal{G}^{\otimes k})''$ be the sub-game after the additional conditioning in $S_{\vec{\sigma}}$ and $T_{\vec{\sigma}}$. Let $\hat{k} = k - l$, and denote $[\hat{k}] = \{1, \dots, k\} - I$. Note that effectively $(\mathcal{G}^{\otimes k})''$ has only \hat{k} rounds. For every $i \in [\hat{k}]$, define the game $\hat{\mathcal{G}}_i$ as the restriction of the game $(\mathcal{G}^{\otimes k})''$ to the i 'th round, where the provers are given their questions in all \hat{k} rounds, but are tested only on their answers in the i 'th round. Let $\{S_{i,j}\}_j$ be the marginal of $S_{\vec{\sigma}}$ corresponding to the i 'th question (this is a convex combination of events over X , where $\Pr [S_{i,j}] \geq \delta'$ for all j). Similarly, let $\{T_{i,j}\}_j$ be the marginal of $T_{\vec{\sigma}}$ corresponding to the i 'th question (this is a convex combination of events over X , where $\Pr [T_{i,j}] \geq \delta'$ for all j). Let \mathcal{G}_i denote the sub-game $\mathcal{G}_{|\{S_{i,j} \times T_{i,j}\}_j}$.

In Lemma 5.2 we use the independence between the rounds and the extractor structure of G to argue that, no matter what was the fixing of questions and Σ_Y -labels for the influential rounds, a strategy for $\hat{\mathcal{G}}_i$ can be used to derive a strategy for \mathcal{G}_i whose value is at least $\text{val}(\hat{\mathcal{G}}_i) - O(k\varepsilon)$. By the fortification, we have $\text{val}(\mathcal{G}_i) \leq \text{val}_{\delta'}(\mathcal{G})$, and hence $\text{val}(\hat{\mathcal{G}}_i) \leq \text{val}_{\delta'}(\mathcal{G}) + O(k\varepsilon)$. On the other hand, from Lemma 5.1, if we take expectation of $\text{val}(\hat{\mathcal{G}}_i)$ over all fixing of questions and Σ_Y -labels for the influential rounds,

$$\mathbf{E} \left[\text{val}(\hat{\mathcal{G}}_i) \right] > \hat{\varepsilon} - \delta' |\Sigma_Y|^l \geq \hat{\varepsilon} - \varepsilon.$$

Since the left hand side is upper bounded by $\text{val}_{\delta'}(\mathcal{G}) + O(k\varepsilon)$, we get a contradiction (recall the definition of $\hat{\varepsilon}$). It remains to prove:

Lemma 5.2 (One round approximation). *There is $err = O(k\varepsilon)$, such that for every fixing of questions and Σ_Y -labels to the influential rounds, for every $i \in [\hat{k}]$,*

$$val(\hat{\mathcal{G}}_i) \leq val(\mathcal{G}_i) + err.$$

Proof. We consider the event $y_i = y$ for each $y \in Y$, and relate the provers winning in $\hat{\mathcal{G}}_i$ to their winning in \mathcal{G}_i .

For every $y \in Y$, consider the bipartite graph $(G^{\otimes k})_{y,i}$ whose vertices consist of all $\vec{x} = (x_1, \dots, x_k) \in X^k$ such that $\{x_j\}_{j \in I}$ is as fixed and $(x_i, y) \in E$, and all $\vec{y} = (y_1, \dots, y_k) \in Y^k$ such that $\{y_j\}_{j \in I}$ is as fixed and $y_i = y$. There is an edge between \vec{x} and \vec{y} if $e_j = (x_j, y_j) \in E$ for all $1 \leq j \leq k$. Denote $(G^{\otimes k})_{y,i} = ((X^k)_{y,i}, (Y^k)_{y,i}, (E^k)_{y,i})$. Since G is a (δ, ε) extractor, the product graph $(G^{\otimes k})_{y,i}$ is a $(\delta'' = \delta/\varepsilon, \varepsilon'' = 2k\varepsilon)$ -extractor [7].

Let $S_y \subseteq (X^k)_{y,i}$ be those vertices \vec{x} where event $S_{\vec{\sigma}}$ happens. Let $T_y \subseteq (X^k)_{y,i}$ be those vertices \vec{x} where event $T_{\vec{\sigma}}$ happens. Partition the vertices $\vec{x} \in S_y$ according to the assignments to y : For every $b \in \Sigma_Y$, let $S_{y,b} \subseteq S_y$ consist of those \vec{x} for which prover 1 assigns label $a_i \in \Sigma_X$ to x_i and $\pi_{e_i}(a_i) = b$. Partition the vertices $\vec{x} \in T_y$ according to the assignments to y : For every $b \in \Sigma_Y$, let $T_{y,b} \subseteq T_y$ consist of those \vec{x} for which prover 2 assigns label $a_i \in \Sigma_X$ to x_i and $\pi_{e_i}(a_i) = b$.

Focus on $y \in Y$ such that $|S_y| \geq \varepsilon \sum_y |S_y|$ and $|T_y| \geq \varepsilon \sum_y |T_y|$. The probability other y 's are selected as y_i in $\hat{\mathcal{G}}_i$ is at most ε . Focus on $b \in \Sigma_Y$ such that $|S_{y,b}| \geq \varepsilon |S_y|$ and $|T_{y,b}| \geq \varepsilon |T_y|$. The contribution to winning $\hat{\mathcal{G}}_i$ from b 's that do not satisfy this is at most ε . We have $|S_{y,b}| \geq \varepsilon |S_y| \geq \varepsilon \cdot \varepsilon \sum_y |S_y| \geq \varepsilon^2 \delta' |(X^k)_{y,i}|$ and similarly, $|T_{y,b}| \geq \varepsilon^2 \delta' |(X^k)_{y,i}|$. Since $\varepsilon^2 \delta' \geq \delta''$, by the extractor property, conditioned on $y_i = y$, the probability that the provers win $\hat{\mathcal{G}}_i$ by agreeing on the label b for y , is at most

$$\frac{|S_{y,b}|}{|S_y|} \cdot \frac{|T_{y,b}|}{|T_y|} + O(\varepsilon'').$$

The overall probability of winning $\hat{\mathcal{G}}_i$ is at most

$$\begin{aligned} & \mathbf{E}_{y \in Y} \left[\sum_{b \in \Sigma_Y} \frac{|S_{y,b}|}{|S_y|} \cdot \frac{|T_{y,b}|}{|T_y|} \right] + O(\varepsilon'' + \varepsilon) \\ & \leq val(\mathcal{G}_i) + err. \end{aligned}$$

The last inequality holds since the strategy in $\hat{\mathcal{G}}_i$ naturally induces a strategy for \mathcal{G}_i , in which given $x \in X$, each prover picks $\vec{x} \in X^k$ by placing x in the i 'th coordinate and picking the rest of the questions at random conditioned on the prior fixing and the event $S_{\vec{\sigma}}$ or $T_{\vec{\sigma}}$. The probability that this strategy succeeds is (up to the error term) on the left hand side. \square

6 Derandomized Parallel Repetition, Two Rounds and Composition

A natural question is whether it is possible to apply our parallel repetition and fortification ideas in order to obtain a projection PCP with soundness error lower than $1/\text{poly log } n$. To obtain such a low error we can no longer apply parallel repetition with k independent rounds. The reason is that this requires a super-constant k , for which parallel repetition blows-up the size

to n^k . A natural idea is to use k correlated rounds; an idea often referred to as “derandomizing parallel repetition”. In this section we explain the difficulty in “fortifying” in the derandomized setting. Moreover, we relate the problem of derandomizing parallel repetition to a different well-studied problem in PCP; that of *composition*. While we continue to use our notation from the previous part of the paper, everything in this part of the paper holds for general two prover games.

6.1 Correlation and Fortification

We start with explaining what breaks down in the analysis in Section 5 when considering the correlated case. In Section 5 we fix questions in the influential rounds and relate the game in the remaining rounds to \mathcal{G} . This approach fails in the correlated case, as the questions in the remaining rounds are likely to be extremely far from uniform in \mathcal{G} after such a fixing. The fixing was used in order to prevent conditioning in W from introducing dependencies between the questions of the provers beyond those captured by the graph of $\mathcal{G}^{\otimes k}$. The latter is what allowed us to fortify only against rectangular sub-games.

When conditioning on an event W that arbitrarily depends on the questions to both provers, fortification with respect to rectangular sub-games is no longer sufficient, nor is fortifying a single round without taking others into account. A natural generalization of fortification is with respect to general large sub-games of $\mathcal{G}^{\otimes k}$. However, the condition that *any* (non-rectangular) sub-game of fraction at least δ in $\mathcal{G}^{\otimes k}$ has value smaller than 1 is equivalent to the statement that $\text{val}(\mathcal{G}^{\otimes k}) < \delta$, which is precisely what we try to prove! Interestingly, the value of a *random* large sub-game of $\mathcal{G}^{\otimes k}$ (indeed, of any game with value sufficiently smaller than 1) does have value smaller than 1 with high probability [8]. However, since the provers are adversarial, this does not constitute a useful fortification. An intriguing open problem following this work is to define fortification that is both easy to analyze and useful for the correlated case.

6.2 Degrees and Two Rounds

The degrees in the graph associated with a two prover game \mathcal{G} correspond to the uncertainty each prover has with respect to the questions of the other prover. We know that the degrees have to be at least $1/\delta$ to allow for $\text{val}(\mathcal{G}) \leq \delta$ (assuming that each test of the verifier can be satisfied by itself). In fact, Feige and Kilian [14] show that in the randomness-efficient case, for $\text{val}(\mathcal{G}^{\otimes k}) \leq \delta$ it must be the case that the degrees in \mathcal{G} 's graph are at least $\approx 1/\delta$ (under an assumption on \mathcal{G} they call *softness*, which indeed holds in the cases of interest). In this section we relax the problem of derandomizing parallel repetition to a corresponding combinatorial problem about degrees in graphs. We call the combinatorial problem the Two Rounds problem. We then relate the Two Rounds problem to the well-studied problem of composition of two prover games, and show that any efficient composition scheme yields a solution to the problem.

We remark that large degrees are a necessary, but not a sufficient, condition for small value. In fact, in general two prover games one can increase the degree artificially, and without decreasing the value, by duplicating questions. Interestingly, for projection games on (δ, ε) -extractors $G = (X, Y, E)$, the degree in Y is necessarily large $\approx 1/\delta$, while the degree in X cannot be artificially increased by duplicating Y vertices due to the extractor property. This observation supports the intuition that large degrees are “morally” a sufficient condition for low value, at least in cases of interest.

Definition 6.1 (Two Rounds). *Given two projection games \mathcal{G}, \mathcal{H} on bi-regular graphs $G = (X, Y, E)$ and $H = (X', Y', E')$, respectively, where the degrees are at least d , we say that a distribution over pairs (\vec{x}, \vec{y}) , where $\vec{x} = (x, x') \in X \times X'$ and $\vec{y} = (y, y') \in Y \times Y'$, yields “two rounds” of \mathcal{G} and \mathcal{H} for parameter d , if:*

- (x, y) is a uniformly distributed edge in E .
- (x', y') is a uniformly distributed edge in E' .
- For any fixing of x, x' , we have $H(y|y'), H(y'|y) \geq \log d$.
- For any fixing of y, y' , we have $H(x|x'), H(x'|x) \geq \log d$.

Note that picking two *independent* uniform edges $(x, y) \in E$ and $(x', y') \in E'$ yields two rounds. The challenge is to pick two rounds using less randomness. Ideally, one could hope to use $\log |E| + O(\log D)$ randomness, when D is the maximal degree in the graph, since given (x, x') (similarly, given (y, y')), there are at most D^2 alternatives for (y, y') (respectively, (x, x')).

A randomness-efficient solution to the Two Rounds problem for games \mathcal{G} and \mathcal{H} yields a candidate construction for a derandomized 2-round parallel repetition, where the tensored games are \mathcal{G} and \mathcal{H} : the first prover gets questions x, x' , while the second prover gets questions y, y' ; the first prover answers $a, a' \in \Sigma_X$, while the second prover answers $b, b' \in \Sigma_Y$; the verifier checks that $\pi_e(a) = b$ and $\pi_{e'}(a') = b'$ (for simplicity, in this part of the paper we consider the more standard definition of projection games; see Remark 1.1). The definition of two rounds guarantees that a prover who knows both x and x' , even if it has information on y (e.g., by virtue of conditioning on an event W), has a lot of uncertainty about y' . The same goes for the other question and the other prover.

The hope is that there are randomness-efficient two rounds – and, more generally, k rounds – for “interesting” games \mathcal{G} and \mathcal{H} , namely, ones whose value is NP-hard to approximate. Ideally, such a derandomized parallel repetition scheme would yield $\text{val}(\mathcal{G}^{\otimes k}) \leq \delta$ when the size of $\mathcal{G}^{\otimes k}$ is $\text{size}(\mathcal{G}) \cdot D^{O(k)}$ for $D, d = \Theta(1/\delta)$. In other words, this would give a projection PCP with soundness error $\delta = 2^{-(\log n)^\beta}$ for some constant $\beta > 0$, i.e., exponentially smaller than what we know today.

6.3 Composition

We wish to relate the Two Rounds problem to the problem of composition for PCP. The goal in the composition problem is to take an *outer* game with large alphabets, as well as small *inner* games with small alphabets, and compose them into a single game with small alphabets. This is similar to concatenation for codes, where one combines an outer code with large alphabet and inner codes over a small alphabet to get a single code over the small alphabet. The idea of composition is to simulate a test of the outer game using a test of an inner game, since the latter only requires small alphabet. We will argue that the Two Rounds problem is tightly connected to the composition problem, as it is to the derandomized parallel repetition problem. The difference between composition and repetition is that in composition the second round comes to replace the first round, while in repetition the second round is in addition to the first round.

More formally, one is given a projection game \mathcal{G} , referred to as the *outer game*, on a graph $G = (X, Y, E)$, alphabets Σ_X, Σ_Y and projections $\{\pi_e\}_{e \in E}$ as before. One is also given a collection of $|E|$ projection games $\{\mathcal{G}_e\}_{e \in E}$ that are referred to as *inner games*. All of the games

\mathcal{G}_e in the collection are on vertices X' and Y' , and the alphabets are $\Sigma_{X'}$ and $\Sigma_{Y'}$. The graphs are bi-regular, the set of edges in \mathcal{G}_e is denoted E_e , and $|E_e|$ is the same for all $e \in E$. We use d to denote the minimal degree in E_e .

The inner games satisfy the following property:

Definition 6.2 (Inner game). *Let $e = (x, y) \in E$. We say that \mathcal{G}_e is an inner game of value ε associated with e , if:*

- *For every label $a \in \Sigma_X$ to x there is a prover strategy $a' : X' \rightarrow \Sigma_{X'}$ to X' , and for every label $b \in \Sigma_Y$ to y there is a prover strategy $b' : Y' \rightarrow \Sigma_{Y'}$. If $\pi_e(a) = b$, then the value that the prover strategies a' and b' achieve in \mathcal{G}_e is 1.*
- *Every prover strategy $a' : X' \rightarrow \Sigma_{X'}$ corresponds to a label $a \in \Sigma_X$ (or a small set of labels), and every prover strategy $b' : Y' \rightarrow \Sigma_{Y'}$ corresponds to a label $b \in \Sigma_Y$ (or a small set of labels). If a' and b' achieve value at least ε in \mathcal{G}_e , then there are corresponding labels $a \in \Sigma_X$ and $b \in \Sigma_Y$ such that $\pi_e(a) = b$.*

The game $\mathcal{G} \diamond \{\mathcal{G}_e\}$ is as follows: one picks uniformly at random an edge $e = (x, y) \in E$ and an edge $e' = (x', y') \in E_e$; the first prover is sent (x, x') , while the second prover is sent (y, y') ; the first prover is supposed to reply with the assignment $a'(x')$, where a' is the strategy associated with a label a for x ; the second prover is supposed to reply with the assignment $b'(y')$ where b' is the strategy associated with a label b for y ; it is verified that $\pi_{e'}(a'(x')) = b'(y')$.

The hope is that there are games \mathcal{G} whose value is NP-hard to approximate, as well as games $\{\mathcal{G}_e\}$ satisfying the inner property for \mathcal{G} , where the size of the inner games is small (roughly $\text{poly log} |\Sigma_X|$) and whose alphabets $\Sigma_{X'}, \Sigma_{Y'}$ are small as well. As it stands now, there are constructions of inner games based on the Hadamard code and based on the long code (these are variants of standard constructions as in [9] and [22]). Hadamard-based constructions have $|X'|, |Y'| = \text{poly}(|\Sigma_X|)$, while the long code-based constructions have $|X'|, |Y'| = \exp(|\Sigma_X|)$. Both have alphabets that are of size polynomial in $1/\varepsilon$. There are also constructions that have size polynomial in $\log |\Sigma_X|$ [21, 11], alas, they have a large alphabet $|\Sigma_{X'}| = \exp(\text{poly}(1/\varepsilon))$. To improve on the current state of the art in PCP one has to design inner PCPs for the case of outer alphabet Σ_X that is larger than $\exp(\text{poly log } n)$.

The connection between composition and the Two Rounds problem is as follows. Let \mathcal{H} be the projection game on vertex sets X' and Y' defined by picking $e \in E$ uniformly at random and playing \mathcal{G}_e . Denote by $H = (X', Y', E')$ the bi-partite graph associated with \mathcal{H} . It holds:

Claim 6.1. *If the degrees in G are at least D , and the degrees in the \mathcal{G}_e 's are at least d , then the distribution over questions in $\mathcal{G} \diamond \{\mathcal{G}_e\}$ yields two rounds of the games \mathcal{G} and \mathcal{H} , where the parameter is $\min\{d, D/|X'|, D/|Y'|\}$.*

Proof. Let (x, x') and (y, y') be as picked in $\mathcal{G} \diamond \{\mathcal{G}_e\}$. Note that (x, y) is uniformly distributed over E , while (x', y') is uniformly distributed over E' . Fix x and x' . For any fixed y , we have that y' is uniform over d possibilities. We also have that $H(y|y') \geq H(y) - H(y') \geq \log D - \log |Y'|$. The argument is similar when one fixes y and y' . \square

Acknowledgements

I am thankful to Ran Raz for discussions.

References

- [1] S. Aaronson, D. Moshkovitz, and R. Impagliazzo. AM with multiple merlins. In *Computational Complexity Conference*, 2014.
- [2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
- [3] S. Arora and S. Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.
- [4] L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy. Checking computations in polylogarithmic time. In *Proc. 23rd ACM Symp. on Theory of Computing*, pages 21–32, 1991.
- [5] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.
- [6] B. Barak, M. Hardt, I. Haviv, A. Rao, O. Regev, and D. Steurer. Rounding parallel repetitions of unique games. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, pages 374–383, 2008.
- [7] M. R. Capalbo, O. Reingold, S. P. Vadhan, and A. Wigderson. Randomness conductors and constant-degree lossless expanders. In *IEEE Conference on Computational Complexity*, page 15, 2002.
- [8] M. Dinitz, G. Kortsarz, and R. Raz. Label cover instances with large girth and the hardness of approximating basic k-spanner. In *ICALP*, pages 290–301, 2012.
- [9] I. Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):12, 2007.
- [10] I. Dinur and P. Harsha. Composition of low-error 2-query PCPs using decodable PCPs. In *Proc. 50th IEEE Symp. on Foundations of Computer Science*, pages 472–481, 2009.
- [11] I. Dinur and O. Meir. Derandomized parallel repetition via structured PCPs. *Computational Complexity*, 20(2):207–327, 2011.
- [12] I. Dinur and D. Steurer. Analytical approach to parallel repetition. In *Proc. 46th ACM Symp. on Theory of Computing*, 2014.
- [13] U. Feige. On the success probability of the two provers in one round proof systems. In *Proc. of 6th IEEE Symposium on Structure in Complexity Theory*, pages 116–123, 1991.
- [14] U. Feige and J. Kilian. Impossibility results for recycling random bits in two-prover proof systems. In *Proc. 27th ACM Symp. on Theory of Computing*, pages 457–468, 1995.
- [15] U. Feige and O. Verbitsky. Error reduction by parallel repetition - a negative result. *Combinatorica*, 22(4):461–478, 2002.
- [16] L. Fortnow, J. Rompel, and M. Sipser. Errata for on the power of multi-prover interactive protocols. In *Structure in Complexity Theory Conference*, pages 318–319, 1990.
- [17] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.

- [18] T. Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009.
- [19] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for MAX-CUT and other two-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [20] D. Moshkovitz. The projection games conjecture and the NP-hardness of $\ln n$ -approximating set-cover. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX 2012*, volume 7408, pages 276–287, 2012.
- [21] D. Moshkovitz and R. Raz. Two query PCP with sub-constant error. *Journal of the ACM*, 57(5), 2010.
- [22] J. Radhakrishnan and M. Sudan. On Dinur’s proof of the PCP theorem. *Bulletin of the AMS*, 44(1):19–61, 2007.
- [23] A. Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011.
- [24] R. Raz. A parallel repetition theorem. In *SIAM Journal on Computing*, volume 27, pages 763–803, 1998.
- [25] R. Raz. A counterexample to strong parallel repetition. *SIAM Journal on Computing*, 40(3):771–777, 2011.
- [26] R. Raz and R. Rosen. A strong parallel repetition theorem for projection games on expanders. In *IEEE Conference on Computational Complexity*, pages 247–257, 2012.
- [27] O. Reingold, S. P. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. *Annals of Mathematics*, 155(1):157–187, 2002.
- [28] R. Shaltiel. Derandomized parallel repetition theorems for free games. *computational complexity*, 22(3):565–594, 2013.
- [29] O. Verbitsky. Towards the parallel repetition conjecture. In *Structure in Complexity Theory Conference*, pages 304–307, 1994.