

Parallel Repetition From Fortification

Dana Moshkovitz *

April 16, 2015

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. In this work we give a simple transformation on games, which we call “fortification”, and show that for fortified games \mathcal{G} , the value of the 2-repeated game is approximately $val(\mathcal{G})^2$. Our proof is extremely short and simple. As a corollary, starting with the combinatorial PCP of Dinur that has soundness error close to 1, we get a simple, combinatorial, construction of a PCP with soundness error close to 0. The latter can be used for hardness of approximation as in the work of Håstad.

*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 [13, 10, 12]. In a breakthrough result, Raz [20] 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 [20]). *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 [12]. 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 [15], as well as for games with a special structure, most notably *projection games* [19, 9] and *expanding projection games* [22, 9]. 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. *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 $\pi_e(a) = b$. Definition 1.1 is a symmetric version of this definition, and as useful to hardness of approximation (or more). 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 first inequality follows from convexity, while the second inequality follows from a probabilistic assignment).*

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 $val(\mathcal{G}) \leq \epsilon$. 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 [4, 3, 2, 1], 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 [14].

Understanding the significance of projection games, we now turn to review what is known about their repetition. Interestingly, in the projection case $val(\mathcal{G}^{\otimes k})$ does not depend on the number of possible answers of the provers [19]. The state of the art results are as follows:

Theorem 2 (Parallel Repetition Theorem for projection games [19, 9]). *For any projection game \mathcal{G} as in Remark¹ 1.1,*

1. *If $val(\mathcal{G}) = 1 - \epsilon$, then $val(\mathcal{G}^{\otimes k}) \leq (1 - \epsilon/2)^{\Omega(\epsilon k)}$ [19].*
2. *If $val(\mathcal{G}) = 1 - \epsilon$ and $\epsilon \ll 1/\sqrt{k}$, then $val(\mathcal{G}^{\otimes k}) \leq 1 - \Omega(\sqrt{k} \cdot \epsilon)$ [9].*
3. *$val(\mathcal{G}^{\otimes k}) \leq 2^k \cdot val(\mathcal{G})^{k/2}$ [9].*

The first result is best when $val(\mathcal{G})$ is a constant close to 1; the second result is best for $val(\mathcal{G})$ very close to 1; while the last result is best for the case of small $val(\mathcal{G})$ (note that in the first result the base of the exponent is about $\frac{1}{2}$ rather than $val(\mathcal{G})$ when $val(\mathcal{G})$ is very small). The second result is tight when it applies, as Raz [21] showed a unique game \mathcal{G} with $val(\mathcal{G}) = 1 - \epsilon$ for which $val(\mathcal{G}^{\otimes k}) \geq 1 - O(\sqrt{k} \cdot \epsilon)$. More generally, Barak et al [5] analyze the behavior of general unique games under parallel repetition.

For projection games on expanders, Dinur and Steurer’s proof is somewhat simpler than its general case [9]. More than that, Raz and Rosen [22] prove a stronger result in the expander case: if $val(\mathcal{G}) = 1 - \epsilon$ for $\epsilon < 1/2$, then $val(\mathcal{G}^{\otimes k}) \leq (1 - \epsilon)^{\Omega(k)}$.

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 $val(\mathcal{G}^{\otimes 2}) = val(\mathcal{G})$ [10].

1.2 Our Contribution

Instead of exploring the subtle behavior of general projection games under repetition, in this work we engineer the games so they behave well under repetition. We present a simple combinatorial transformation on projection games, which we call “fortification”. Fortification endows the game with extractor structure and ensures that certain sub-games of the game have (approximately) the same value as the global game. Fortification preserves a projection structure, while increasing $|X|$ and $|\Sigma_X|$ in a controlled way. We show that for fortified projection games \mathcal{G} , the value of the k -repeated game is, approximately, $val(\mathcal{G})^k$, i.e.,

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

¹Using 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.

where the small additive error err can be made arbitrarily small by fortification. While we can handle general k , our focus will be on getting the simplest proof possible so we concentrate on the case $k = 2$, which allows amplification of the value via repeated squaring.

In the fortified game, rather than sending the first prover a question x and the second prover a question x' , the verifier sends the first prover a set of correlated questions $\{x_1, \dots, x_t\} \ni x$, and it sends the second prover a set of correlated questions $\{x'_1, \dots, x'_t\} \ni x'$. The provers are asked to provide answers for all t questions they got. The verifier then uses their answers to perform the test involving x and x' (note that other questions among the $2t$ typically induce no tests). The choice of the correlated questions is done using an extractor or a random walk on an expander, in a manner that was inspired by ideas in combinatorial construction of error correcting codes.

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 a heavy use of linear algebra and Cheeger's inequality as in the recent analysis of Dinur and Steurer for projection games.

As a corollary, starting from a PCP Theorem with soundness error bounded away from 1 [4, 3, 2, 1], we get a PCP with arbitrarily small constant soundness error. In particular, starting with the combinatorial PCP of Dinur, we get a combinatorial PCP with low error whose analysis is combinatorial. The latter can be used for hardness of approximation as in the work of Håstad.

Our proof evolved from a previous work of the author [17] about soundness amplification for low degree testing. As happened several times in the past in PCP, we could transform some of the ideas from the algebraic analysis into a purely combinatorial setting.

1.3 Simplification, Erratum, Relation to [6]

There are two major differences between this version and the conference version of this work. First, the proof of parallel repetition is considerably simplified and extremely short. Second, the proof of the fortification lemma is corrected. Specifically, the conference version used (δ, ε) -extractors where (roughly speaking) $\delta \approx val(\mathcal{G})^k$ and $\varepsilon \approx val(\mathcal{G})$, whereas the correct parameters are $\delta, \varepsilon \approx val(\mathcal{G})^k$. The latter parameters weaken the result in one respect: the alphabet becomes exponential in inverse the value, rather than polynomial (in contrast, the size of the repeated game remains nearly as efficient as it was before, and improves upon previous works on parallel repetition). The increase in the alphabet size is inconsequential in the most important application of parallel repetition where the value and alphabet are constants, however it rules out the use of our technique as is for proving optimal hardness of SET-COVER. The latter hardness result was already known from Dinur and Steurer's work on parallel repetition [9].

Amey Bhangale, Ramprasad Saptharishi, Girish Varma and Rakesh Venkat gave an example where the previous parameter setting was insufficient. Then, following a correspondence with the author in which she suggested to consider stronger parameters, they were able to come up with an analysis that they wrote up as [6]. Their analysis is different and independent from the analysis that the author came up with and is detailed in the appendix.

1.4 Previous Work on Combinatorial Analysis of Parallel Repetition

Feige and Kilian [11], as well as Impagliazzo, Kabanets and Wigderson [16] already gave combinatorial analyses of parallel repetition. Crucially, those parallel repetition theorems were *weaker* than what was known via other techniques, while our theorem is *stronger* than what is known

via other techniques. As in the current paper, Feige and Kilian, as well as Impagliazzo, Kabanets and Wigderson, first apply a transformation on the game, and then repeat the game in parallel. The transformation differs from our fortification, and is (up to variants) as follows: The verifier picks uniformly at random either (i) “compare”: edges with a common endpoint $e = (x, y), e' = (x', y) \in E$; or (ii) “confuse”: independent edges $e = (x, y), e' = (x', y') \in E$. One prover is sent x and the other prover is sent x' . The provers reply $a, a' \in \Sigma_X$, respectively; In case the two edges have a common endpoint y , the verifier checks that $\pi_e(a) = \pi_{e'}(a')$. The intuition of this transformation is that in some of the rounds the provers are compared, hence for the verifier to accept with good probability, the provers are forced to a consistent strategy. The confuse rounds ensure that the consistent strategy is pervasive.

Feige and Kilian [11] show that for games \mathcal{G} transformed as we described, for any $\delta > 0$ such that k is a sufficiently large polynomial in $1/\delta$ and $1/(1 - \text{val}(\mathcal{G}))$, it holds that $\text{val}(\mathcal{G}^{\otimes k}) \leq \delta$. In this theorem, the decay of the value of the game with repetition is polynomial in k , rather than exponential in k . Impagliazzo, Kabanets and Wigderson [16] prove that $\text{val}(\mathcal{G}^{\otimes k}) \leq 2^{-\Omega(\sqrt{k}/(1 - \text{val}(\mathcal{G})))}$. Here the decay is exponential in \sqrt{k} instead of in k .

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} . If a distribution is uniform over a set $S \subseteq X$ (a “flat” distribution; we’ll also refer to S as an event), then it has min-entropy $\log |S|$. It is known that any distribution with min-entropy k can be viewed as a convex combination $\{S_i\}_i$ where for every i , it holds that $\Pr[S_i] = |S_i|/|X| \geq 2^{-k}$.

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 in statistical distance.

Lemma 2.1 (Extractor construction; follows from expander random walk [23]). *For any $\delta, \varepsilon > 0$, 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 $\text{val}(\mathcal{G})$, is the maximum of this quantity over all prover strategies. The size of \mathcal{G} , denoted $\text{size}(\mathcal{G})$, 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 question to the first prover, and T is an event depending on the question 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. That is, the verifier picks at random $y \in Y$, $x, x' \in X$ such that $(x, y), (x', y) \in E$, conditioned on $x \in S, x' \in T$. It then performs the test as before. We say that the rectangular game is δ -large if $\Pr[S], \Pr[T] \geq \delta$.

We define the fortified value of a game as follows:

Definition 3.1. *The δ -fortified value of a game \mathcal{G} , denoted $val_\delta(\mathcal{G})$, is the maximum of $val(\mathcal{G}_{|S \times T})$ over all δ -large rectangular games $\mathcal{G}_{|S \times T}$.*

We say that a projection game \mathcal{G} is *δ -fortified* if the underlying graph is a (δ, δ^2) -extractor and δ -large rectangular sub-games have value at most $val(\mathcal{G}) + \delta$. We show that every projection game can be fortified easily. Fortification increases the size and the alphabets of the game by factor polynomial in $1/\delta$. 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.

Our fortification lemma assumes that the bipartite graph underlying the projection game is bi-regular. Projection games on general graphs can be transformed to bi-regular using transformations of [18]. The first transformation regularizes the Y side, so each Y vertex has a small degree:

Lemma 3.1 (*Y-degree reduction* [18]). *For any $\eta > 0$, any projection game \mathcal{G} can be efficiently transformed to a new projection game \mathcal{G}' on a graph (X, Y, E) that is Y -regular with degree $\text{poly}(1/\eta)$, where $size(\mathcal{G}') \leq size(\mathcal{G}) \cdot \text{poly}(1/\eta)$ and $val(\mathcal{G}') \leq val(\mathcal{G}) + \eta$ (the alphabets are unchanged).*

The second transformation switches between the Y and the X side. The idea is that each assignment to a vertex $y \in Y$ now contains assignments to all the neighbors of y , such that the assignments to the neighbors agree on their projection to y :

Lemma 3.2 (*Switching sides* [18]). *Any projection game \mathcal{G} on a graph $G = (X, Y, E)$ with Y -degree D and alphabets Σ_X, Σ_Y can be transformed into a projection game \mathcal{G}' on a graph $G' = (Y, X, E)$ and alphabets Σ_X^D, Σ_Y , where $val(\mathcal{G}') = val(\mathcal{G})$.*

By applying Y -degree reduction, switching sides, and Y -degree reduction again, we obtain a projection game on a bi-regular graph that has approximately the same value as the original game.

Having gotten bi-regularity out of the way, let us describe the fortification transformation:

Lemma 3.3 (*Fortification*). *For any $\varepsilon, \delta > 0$, 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, δ^2) -extractor.
2. $D = D(\delta, \delta^2)$, where D is as in Lemma 2.1.
3. The size of G^* is linear in the size of G and $\text{poly}(1/\delta)$.
4. $\text{val}(\mathcal{G}^*) = \text{val}(\mathcal{G})$.
5. $\text{val}_\delta(\mathcal{G}^*) \leq \text{val}(\mathcal{G}) + O(\delta)$.

Proof. Let $H = (X^*, X, E_H)$ be a (δ, δ^2) -extractor. By Lemma 2.1, such can be constructed so $|X^*|$ is linear in $|X|$ and $\text{poly}(1/\delta)$, and each vertex in X^* has $D = D(\delta, \delta^2)$ 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, δ^2) -extractor, and that $\text{size}(\mathcal{G}^*)$ is $\text{size}(\mathcal{G}) \cdot \text{poly}(1/\delta)$. 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 $\text{val}(\mathcal{G}) \leq \text{val}(\mathcal{G}^*)$, 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 $S, T \subseteq X^*$, $|S|, |T| \geq \delta |X^*|$. We'd like to prove that $\text{val}(\mathcal{G}^*)_{|S \times T} \leq \text{val}(\mathcal{G}) + O(\delta)$. Lemma A.4 shows that conditioned on the events S and T , when picking $y \in Y$ and two edges $(x^*, y), ((x^*)', y) \in E_H$, $x^* \in S, (x^*)' \in T$ – which involves picking $(x, y), (x', y) \in E$, as well as $(x^*, x), ((x^*)', x') \in E_H$ – we get that (x, x', y) is $O(\delta)$ -close to (x, x', y) picked by sampling uniformly $y \in Y, (x, y), (x', y) \in E$. Since the latter induces a uniform test in \mathcal{G} , our claim follows. \square

Fortification preserves projection, but does not preserve uniqueness. Indeed, due to the works [20, 5], we do not expect to prove a strong parallel repetition for unique games.

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 $\text{poly}(1/\delta)$, where we fortify against sub-games of fraction δ . When repeating the game for k rounds, the size increases by a factor $\text{poly}(1/\delta)^k$. However, due to fortification, $\text{val}(\mathcal{G}^{\otimes k})$ decreases exponentially with k , rather than with $k/2$. Hence, to reach a target $\text{val}(\mathcal{G}^{\otimes k})$ 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 $\approx (\text{size}(\mathcal{G}) \text{poly}(1/\delta))^k$ for us. Since typically $\text{size}(\mathcal{G})$ is much larger than $\text{poly}(1/\delta)$ (which is polynomial in inverse the target value), our method yields better size than before. Fortification also raises the size of the alphabet Σ_X to a power $D = D(\delta, \delta) = \text{poly}(1/\delta)$. This is a disadvantage of our method, since standard parallel repetition has alphabet size $\text{poly}(1/\delta)$ where δ

is the final value, whereas our alphabet size is exponentially worse. When parallel repetition is applied to achieve a (small) constant soundness error (as is typically the case, since repetition raises the size to the k), the increase in the alphabet size is inconsequential.

4 A Parallel Repetition Theorem

In this work we prove a parallel repetition theorem assuming that the underlying game is fortified. While we could handle a large number of repetitions at once, we present a proof for only two repetitions, since this case is exceptionally short and simple. Using repeated squaring we can then decrease the value of the game further.

Theorem 3 (Parallel repetition). *There exists a constant $c \geq 1$, such that for any $\delta, \varepsilon > 0$, if \mathcal{G} is δ -fortified where $\delta \leq \varepsilon/(c|\Sigma_Y|)$, then*

$$\text{val}(\mathcal{G}^{\otimes 2}) \leq \text{val}(\mathcal{G}) \cdot (\text{val}(\mathcal{G}) + \varepsilon) + \varepsilon.$$

Proof. Suppose on way of contradiction that there exists a strategy for the two players $a : X \times X \rightarrow \Sigma_A \times \Sigma_A$, $b : X \times X \rightarrow \Sigma_A \times \Sigma_A$, that makes the verifier accept with probability larger than $\text{val}(\mathcal{G}) \cdot (\text{val}(\mathcal{G}) + \varepsilon) + \varepsilon$. Let $(x_1, y_1), (x'_1, y_1) \in E$, be the random choices of the verifier in the first repetition, and let $(x_2, y_2), (x'_2, y_2) \in E$, be those of the second repetition. Denote the replies of the provers by $(a_1, a_2) = a(x_1, x_2)$ and $(a'_1, a'_2) = a'(x'_1, x'_2)$.

Fix $(x_1, y_1), (x'_1, y_1) \in E$. Then a_1, a'_1 become functions of x_2, x'_2 , respectively. For every $b \in \Sigma_Y$, let $S_b = \{x_2 \in X \mid \pi_{(x_1, y_1)}(a_1(x_2)) = b\}$ and $T_b = \{x'_2 \in X \mid \pi_{(x'_1, y_1)}(a'_1(x'_2)) = b\}$. There are two cases:

- $|S_b| < \delta|X|$ or $|T_b| < \delta|X|$: such b 's contribute to $\text{val}(\mathcal{G}^{\otimes 2})$ at most $2|\Sigma_Y| \cdot \delta \leq \varepsilon$ (for sufficiently large c).

Our assumption implies that conditioned on the players agreeing in the first repetition, namely, $\pi_{(x_1, y_1)}(a_1) = \pi_{(x'_1, y_1)}(a'_1)$, the probability p that the players agree in the second repetition, namely, $\pi_{(x_2, y_2)}(a_2) = \pi_{(x'_2, y_2)}(a'_2)$, and $|S_b|, |T_b| \geq \delta|X|$, is more than $\text{val}(\mathcal{G}) + \varepsilon$.

- $|S_b|, |T_b| \geq \delta|X|$: by fortification, such b 's contribute to p at most $\text{val}(\mathcal{G}) + O(\delta)$. We pick c so the $O(\delta)$ term is at most ε .

Overall, we get a contradiction. □

Theorem 3 assumes that $\delta \leq c\varepsilon/|\Sigma_Y|$, where ε should be taken to be sufficiently smaller than $\text{val}(\mathcal{G})^2$. Existing constructions of projection games \mathcal{G} have $|\Sigma_Y| = \text{poly}(1/\text{val}(\mathcal{G}))$ [18]. Moreover, there is a simple transformation of Dinur and Harsha [8] based on code concatenation that decreases $|\Sigma_Y|$ to $\text{poly}(1/\text{val}(\mathcal{G}))$ for any projection game \mathcal{G} . Overall, this means that δ should be taken to be polynomially small in $\text{val}(\mathcal{G})$.

Lemma 4.1 (Σ_Y reduction [8]). *For any $\eta > 0$, any projection game \mathcal{G} with $\text{val}(\mathcal{G}) \leq \varepsilon$ and alphabet Σ_Y 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 $\text{poly}(1/\eta)$.*

5 Repeated Squaring

Using our parallel repetition theorem, in this section we develop a new combinatorial soundness amplification technique for two prover games that allows one to decrease the soundness error of projection games from close to 1 to close to 0. The latter is required for optimal hardness of approximation results [14].

Theorem 4 (Combinatorial PCP with low error). *For any $\alpha > 0$, it is NP-hard to distinguish, given a projection game \mathcal{G} , between the case where $\text{val}(\mathcal{G}) = 1$ and the case where $\text{val}(\mathcal{G}) \leq \alpha$.*

Proof. We start with the PCP theorem of Dinur [7]. This theorem shows an efficient reduction from an instance H of the NP-hard 3COLORING problem to a regular constraint graph $G = (V, E)$, an alphabet Σ for the vertices, and predicates $p_e \subseteq \Sigma \times \Sigma$ on the edges $e \in E$. If H is 3-colorable, then G has an assignment $f : V \rightarrow \Sigma$ that satisfies all the predicates on the edges. If H is not 3-colorable, then all assignments to G 's vertices satisfy at most $1 - \beta$ fraction of the predicates, where $\beta > 0$ is some global constant.

We can convert G into a projection game \mathcal{G} as follows: the set of questions is E ; the set of answers Σ_X consists of all satisfying assignments to the endpoints of an edge; $Y = V$; and there are edges between $e = (u, v) \in E$ and each of u, v . The projection on those edges maps an assignment $(a_u, a_v) \in \Sigma \times \Sigma$ to the endpoints to either a_u or a_v depending on whether the endpoint is u or v . Thus, Dinur's theorem shows that it is NP-hard to distinguish the case that a given projection game \mathcal{G} with sufficiently large constant alphabets $|\Sigma_X|, |\Sigma_Y|$ satisfies $\text{val}(\mathcal{G}) = 1$ and the case that it satisfies $\text{val}(\mathcal{G}) \leq 1 - \beta'$, where β' is a constant that depends on β .

We repeat the following process until $\text{val}(\mathcal{G}) \leq \alpha$:

- Set $\delta = \alpha / (c |\Sigma_Y|)$ for a sufficiently large constant c .
- Fortify \mathcal{G} to get a δ -fortified \mathcal{G}^* . We have $\text{size}(\mathcal{G}^*) \leq \text{size}(\mathcal{G}) \cdot \text{poly}(1/\delta)$. The X alphabet grows to $|\Sigma_X|^{\text{poly}(1/\delta)}$ and the Y alphabet remains the same.
- Consider $(\mathcal{G}^*)^{\otimes 2}$. By Theorem 3, the soundness error goes down to $\text{val}(\mathcal{G})^2 \cdot (1 + \epsilon)$ (for a suitable constant ϵ and assuming c is large enough). The size of the game, as well as the sizes of the alphabets, are squared.
- Set \mathcal{G} to the new game we constructed, and repeat the process.

After k iterations, the size is $\text{size}(\mathcal{G})^{2^k} \cdot (|\Sigma_Y|/\alpha)^{2^{O(k)}}$, the X alphabet is of size $|\Sigma_X|^{(|\Sigma_Y|/\alpha)^{2^{O(k)}}}$, and the Y alphabet is of size $|\Sigma_Y|^{2^k}$. If $\text{val}(\mathcal{G}) = 1$, then the value remains 1, and in any case $\text{val}(\mathcal{G}) \leq (1 + \epsilon)^k \cdot \text{val}(\mathcal{G})^{2^k}$. For $k \approx \log((1/\beta') \log(1/\alpha))$, the value becomes α . \square

Acknowledgements

I am thankful to Ran Raz, Henry Yuen and an anonymous reviewer. Special thanks to Amey Bhangale, Ramprasad Saptharishi, Girish Varma and Rakesh Venkat for discussions.

References

- [1] 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.
- [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] 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.
- [4] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.
- [5] 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.
- [6] A. Bhangale, R. Saptharishi, G. Varma, and R. Venkat. On fortification of projection games. 2015.
- [7] I. Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):12, 2007.
- [8] 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.
- [9] I. Dinur and D. Steurer. Analytical approach to parallel repetition. In *Proc. 46th ACM Symp. on Theory of Computing*, 2014.
- [10] 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.
- [11] U. Feige and J. Kilian. Two-prover protocols - low error at affordable rates. *SIAM Journal on Computing*, 30(1):324–346, 2000.
- [12] U. Feige and O. Verbitsky. Error reduction by parallel repetition - a negative result. *Combinatorica*, 22(4):461–478, 2002.
- [13] 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.
- [14] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [15] T. Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009.
- [16] R. Impagliazzo, V. Kabanets, and A. Wigderson. New direct-product testers and 2-query PCPs. *SIAM Journal on Computing*, 41(6):1722–1768, 2012.
- [17] D. Moshkovitz. An approach to the sliding scale conjecture via parallel repetition for low degree testing. Technical Report 30, ECCO, 2014.

- [18] D. Moshkovitz and R. Raz. Two query PCP with sub-constant error. *Journal of the ACM*, 57(5), 2010.
- [19] A. Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing*, 40(6):1871–1891, 2011.
- [20] R. Raz. A parallel repetition theorem. In *SIAM Journal on Computing*, volume 27, pages 763–803, 1998.
- [21] R. Raz. A counterexample to strong parallel repetition. *SIAM Journal on Computing*, 40(3):771–777, 2011.
- [22] 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.
- [23] 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.

A Auxiliary Lemmas

Lemma A.1. *Suppose that for every y we have $0 \leq a_y \leq 1$, $\sum_y a_y = 1$. Additionally, $\max_y |c_y - 1| \leq \varepsilon$. Then, $\sum_y |c_y a_y - b_y| \leq \sum_y |a_y - b_y| + \varepsilon$.*

Proof. Let $\varepsilon_y = |c_y - 1|$. First consider the case where $\max_y \varepsilon_y \leq \varepsilon$.

$$\begin{aligned} \sum_y |c_y a_y - b_y| &= \sum_y |a_y - b_y \pm \varepsilon_y a_y| \\ &\leq \sum_y |a_y - b_y| + |\varepsilon_y a_y| \\ &\leq \sum_y |a_y - b_y| + \varepsilon \sum_y a_y \\ &\leq \sum_y |a_y - b_y| + \varepsilon. \end{aligned}$$

□

Lemma A.2. *Let $H = (X, Y, E)$ be a $(\delta, \varepsilon\delta)$ -extractor, and let $S, T \subseteq X$, $|S|, |T| \geq \delta |X|$. When picking uniformly at random $y \in Y$ and $(x, y), (x', y) \in E$,*

$$\left| \frac{\Pr[x \in S, x' \in T]}{\Pr[x \in S] \cdot \Pr[x' \in T]} - 1 \right| \leq 3\varepsilon.$$

Proof.

$$\begin{aligned} \Pr[x \in S, x' \in T] &= \frac{1}{|Y|} \cdot \sum_{y \in Y} \Pr[x \in S|y] \cdot \Pr[x' \in T|y] \\ &= \frac{1}{|Y|} \cdot \sum_{y \in Y} \frac{\Pr[y|x \in S] \cdot \Pr[x \in S]}{\Pr[y]} \cdot \frac{\Pr[y|x' \in T] \cdot \Pr[x' \in T]}{\Pr[y]} \\ &= \Pr[x \in S] \cdot \Pr[x' \in T] \cdot \frac{1}{|Y|} \cdot \sum_{y \in Y} \frac{\Pr[y|x \in S]}{\Pr[y]} \cdot \frac{\Pr[y|x' \in T]}{\Pr[y]} \end{aligned}$$

Let

$$\begin{aligned} \varepsilon_y &= \left| \frac{\Pr[y|x \in S]}{\Pr[y]} - 1 \right|, \\ \varepsilon'_y &= \left| \frac{\Pr[y|x' \in T]}{\Pr[y]} - 1 \right|. \end{aligned}$$

By the extractor property, $\sum \varepsilon_y \leq \varepsilon\delta |Y|$ and $\sum \varepsilon'_y \leq \varepsilon\delta |Y|$. Moreover, we have

$$\Pr[y|x \in S] = \frac{\Pr[x \in S|y] \cdot \Pr[y]}{\Pr[x \in S]} \leq \frac{1}{\delta |Y|}$$

and similarly, $\Pr[y|x \in T] \leq \frac{1}{\delta |Y|}$.

Hence,

$$\begin{aligned}
\frac{1}{|Y|} \cdot \sum_{y \in Y} \frac{\Pr[y|x \in S]}{\Pr[y]} \cdot \frac{\Pr[y|x' \in T]}{\Pr[y]} &= \frac{1}{|Y|} \cdot \sum_{y \in Y} (1 \pm \varepsilon_y) \cdot (1 \pm \varepsilon'_y) \\
&= \frac{1}{|Y|} \cdot \sum_{y \in Y} (1 \pm \varepsilon_y \pm \varepsilon'_y \pm \varepsilon_y \varepsilon'_y) \\
&= 1 \pm 2\varepsilon\delta \pm \varepsilon.
\end{aligned}$$

□

Lemma A.3. *Let $H = (X, Y, E)$ be a $(\delta, \varepsilon\delta)$ -extractor. Let $S, T \subseteq X$, $|S|, |T| \geq \delta|X|$. Consider the distribution over Y where $y \in Y$ is picked uniformly conditioned on two uniformly random neighbors $x, x' \in X$ of y in H satisfying $x \in S$, $x' \in T$. Then, the distribution of y is $O(\varepsilon)$ -close to uniform.*

Proof. Denote

$$A \doteq \frac{\Pr[x \in S] \cdot \Pr[x' \in T]}{\Pr[x \in S, x' \in T]}.$$

By Lemma A.2, $|A - 1| \leq 3\varepsilon$. First we bound:

$$\begin{aligned}
\sum_{y \in Y} \left| \Pr[y|x \in S, x' \in T] - A \cdot \Pr[y|x \in S] \right| &= \sum_{y \in Y} \left| \frac{\Pr[x \in S, x' \in T|y]}{|Y| \cdot \Pr[x \in S, x' \in T]} - A \cdot \Pr[y|x \in S] \right| \\
&= \sum_{y \in Y} \left| \frac{\Pr[x \in S|y] \cdot \Pr[x' \in T|y]}{|Y| \cdot \Pr[x \in S, x' \in T]} - A \cdot \Pr[y|x \in S] \right| \\
&= \sum_{y \in Y} \left| \frac{A \cdot \Pr[y|x \in S] \cdot \Pr[y|x' \in T]}{1/|Y|} - A \cdot \Pr[y|x \in S] \right|
\end{aligned}$$

We wish to bound:

$$\sum_{y \in Y} \frac{A \cdot \Pr[y|x \in S]}{(1/|Y|)} \cdot \left| \Pr[y|x' \in T] - \frac{1}{|Y|} \right|$$

By the extractor property,

$$\sum_{y \in Y} \left| \Pr[y|x' \in T] - \frac{1}{|Y|} \right| \leq O(\varepsilon\delta).$$

We have

$$\Pr[y|x \in S] = \frac{\Pr[x \in S|y] \cdot \Pr[y]}{\Pr[x \in S]} \leq \frac{1}{\delta|Y|}$$

Therefore, we get the bound

$$\sum_{y \in Y} \left| \Pr[y|x \in S, x' \in T] - A \cdot \Pr[y|x \in S] \right| \leq O(\varepsilon).$$

The lemma follows from the triangle inequality and using Lemma A.1 and the extractor property to bound $\sum_{y \in Y} |A \cdot \Pr[y|x \in S] - 1/|Y||$. □

Lemma A.4. *Let $H = (X^*, X, E_H)$ be a $(\delta, \varepsilon\delta)$ -extractor, and let $G = (X, Y, E)$. Let $S, T \subseteq X^*$, $|S|, |T| \geq \delta |X^*|$. Consider the distribution over $Y \times X \times X$ where $y \in Y$, two neighbors $x, x' \in X$ of y in G , and $(x^*, x), ((x^*)', x') \in E_H$, are picked uniformly conditioned on $x^* \in S, (x^*)' \in T$. Then, the distribution of y, x, x' is $O(\varepsilon)$ -close to uniform.*

Proof. The graph $G^* = (X^*, Y, E^*)$ where there is an edge $(x^*, y) \in E^*$ if there are edges $(x^*, x) \in E_H$ and $(x, y) \in G$, is a $(\delta, \varepsilon\delta)$ -extractor. By Lemma A.3 applied on G^* , we know that y is $O(\varepsilon)$ -close to uniform. We'll first bound the following expression

$$\begin{aligned}
& \sum_{x, x', y} \left| \Pr [x, x', y | x^* \in S, (x^*)' \in T] - \Pr [x, x' | y, x^* \in S] \Pr [y | x^* \in S, (x^*)' \in T] \right| \\
&= \sum_{x, x', y} \Pr [y | x^* \in S, (x^*)' \in T] \cdot \left| \Pr [x, x' | y, x^* \in S, (x^*)' \in T] - \Pr [x, x' | y, x^* \in S] \right| \\
&= \sum_{x, x', y} \Pr [y | x^* \in S, (x^*)' \in T] \cdot \left| \Pr [x | y, x^* \in S] \Pr [x' | y, (x^*)' \in T] - \Pr [x | y, x^* \in S] \Pr [x' | y] \right| \\
&= \sum_{x, x', y} \Pr [y | x^* \in S, (x^*)' \in T] \cdot \Pr [x | y, x^* \in S] \cdot \left| \Pr [x' | y, (x^*)' \in T] - \Pr [x' | y] \right| \\
&\leq \sum_{x, x', y} \Pr [y | x^* \in S, (x^*)' \in T] \cdot \left| \Pr [x' | y, (x^*)' \in T] - \Pr [x' | y] \right|
\end{aligned}$$

Let us denote $\varepsilon_y = |\Pr [y | x^* \in S, (x^*)' \in T] - \Pr [y]|$, so $\sum_y \varepsilon_y \leq O(\varepsilon)$, and we can continue bounding as follows:

$$\begin{aligned}
&\leq \sum_{x, x', y} \left(\Pr [y] + \varepsilon_y \right) \cdot \left| \Pr [x' | y, (x^*)' \in T] - \Pr [x' | y] \right| \\
&\leq O(\varepsilon) + \sum_{x, x', y} \Pr [y] \cdot \left| \frac{\Pr [x', y | (x^*)' \in T]}{\Pr [y | (x^*)' \in T]} - \frac{\Pr [x', y]}{\Pr [y]} \right| \\
&= O(\varepsilon) + \sum_{x, x', y} \left| \frac{\Pr [y]}{\Pr [y | (x^*)' \in T]} \cdot \Pr [y | x'] \Pr [x' | (x^*)' \in T] - \Pr [y | x'] \Pr [x'] \right| \\
&\leq O(\varepsilon) + \sum_{x, x', y} \left| \frac{\Pr [y]}{\Pr [y | (x^*)' \in T]} \cdot \Pr [x' | (x^*)' \in T] - \Pr [x'] \right| \\
&\leq O(\varepsilon)
\end{aligned}$$

In the last inequality we use the extractor property and Lemma A.1. To complete the proof we bound:

$$\begin{aligned}
& \sum_{x, x', y} \left| \Pr [x, x' | y, x^* \in S] \Pr [y | x^* \in S, (x^*)' \in T] - \Pr [x, x', y] \right| \\
&= \sum_{x, x', y} \left| \Pr [x' | y] \Pr [x | y, x^* \in S] \Pr [y | x^* \in S, (x^*)' \in T] - \Pr [x' | y] \Pr [x | y] \Pr [y] \right| \\
&= \sum_{x, x', y} \Pr [x' | y] \Pr [y] \cdot \left| \Pr [x | y, x^* \in S] \frac{\Pr [y | x^* \in S, (x^*)' \in T]}{\Pr [y]} - \Pr [x | y] \right| \\
&\leq O(\varepsilon)
\end{aligned}$$

The last inequality follows from Lemmas A.3 and A.1, plus the bound on $\sum |\Pr [x'|y, (x^*)' \in T] - \Pr [x'|y]|$.
The lemma follows from the triangle inequality. \square