

Parallel Repetition of Entangled Games*

Julia Kempe

Thomas Vidick

Abstract

We consider one-round games between a classical verifier and two provers. One of the main questions in this area is the *parallel repetition question*: Is there a way to decrease the maximum winning probability of a game without increasing the number of rounds or the number of provers? Classically, this question, open for many years, has culminated in Raz's celebrated parallel repetition theorem on one hand, and in efficient product testers for PCPs on the other.

In the case where provers share entanglement, the only previously known results are for special cases of games, and are based on techniques that seem inherently limited. Here we show for the first time that the maximum success probability of entangled games can be reduced through parallel repetition, provided it was not initially 1. Our proof is inspired by a seminal result of Feige and Kilian in the context of classical two-prover one-round interactive proofs. One of the main components in our proof is an orthogonalization lemma for operators, which might be of independent interest.

Two-prover games play a major role both in theoretical computer science, where they led to many breakthroughs such as the discovery of tight inapproximability results, and in quantum physics, where they first arose in the context of Bell inequalities. In such games, a referee chooses a pair of questions from some distribution and sends one question to each of two non-communicating players, who then respond with answers taken from a finite set. The referee, based on the questions and answers, decides whether to accept (i.e., whether the players win). The main question of interest is the following: given a game, what is the player's maximum winning probability? Somewhat surprisingly, the answer to this depends on whether the players behave classically, or are allowed to use quantum mechanics. In the former case, the players' answers are simply deterministic functions of their inputs¹, and the maximum probability of winning is known as the (*classical*) *value* of the game. In the latter case the players, though still not allowed to communicate, may share an arbitrary entangled state. The maximum winning probability in this case is known as the *entangled value* of the game.

The parallel repetition question. One of the most important and interesting questions in this context is that of parallel repetition. While it is well known that the (entangled) value of a game can be decreased by repeating it either sequentially or in parallel with several independent pairs of players, for many applications (like hardness of approximation results or amplifications preserving zero-knowledge) one needs a way to decrease the winning probability without increasing the number of rounds or the number of players. Parallel repetition is designed to do just that: in its most basic form, the referee simply chooses ℓ pairs of questions independently and sends to each player his corresponding ℓ -tuple of questions. Each player then replies with an ℓ -tuple of answers, which are accepted if and only if each of the ℓ answer pairs would have been accepted in the original game. Clearly the value of an ℓ -parallel repeated game is *at least* the ℓ -th power of the value of the original game, however it is known that parallel repetition does *not* necessarily decrease the value of a game in a straightforward exponential manner². The challenge of the parallel repetition question is that of showing *upper bounds* on the value of a repeated game, and for a long time no such upper bound, even very weak, could be proved for classical games. Initial results date to Verbitsky [Ver94], who showed that indeed the value goes to zero with the number of repetitions, and Feige and Kilian [FK00], who showed that it decreases polynomially with the number of repetitions for the special case of *projection games*³. They used a mod-

*See the full paper at arXiv:1012.4728.

¹One can allow the players to use randomness, but this does not change their maximum winning probability.

²See [Fei91] for a classical example, and [CSUU08] for an example using entangled players due to Watrous.

³A projection game is one in which the second player's answer is uniquely determined by the first player's. Projection games form a wide class of games that captures most of the games one typically encounters in the classical literature

ified parallel repetition procedure in which a large fraction of the repetitions are made of *dummy* questions, that is, questions which are chosen independently at random independently for both players, and to which any answer is accepted. In this paper we deviate somewhat from the common terminology, and use the term “parallel repetition” even when referring to such more general procedures. Finally, in a breakthrough result, Raz [Raz98] showed that the value of a game repeated in parallel indeed decreases exponentially with the number of repetitions. There is still very active research in the area, mostly on simplifying the analysis, which, over a decade later, remains quite involved, and improving it for certain special cases of games [Hol07, Rao08, FKO07, Raz08, BHH⁺08, BRR⁺09, AKK⁺08, RR10].

Previous work. In this paper we focus on parallel repetition of *games with entangled players*. The only two previous results in this area are for two special classes of games. First, Cleve et al. showed that for the class of *XOR games*⁴, *perfect* parallel repetition holds [CSUU08]. Second, Kempe et al. show that parallel repetition (albeit not perfect) also holds for the more general (but still quite restricted) class of *unique games*⁵ [KRT08]. It is important to note that in all these results, the entangled value of the parallel repeated game is never analyzed directly; instead, one uses a “proxy” such as a semidefinite program [CSUU08, KRT08] or the no-signaling value [Hol07], whose behavior under parallel repetition is well understood. Moreover, in all these cases, the proxy’s value is efficiently computable. This unfortunately gives a very strong indication that such techniques cannot be extended to deal with general games.⁶ This raises the following question: *Can parallel repetition decrease the entangled value of games? If so, can we bound the rate of decrease?*

In parallel to the work on parallel repetition theorems, the related question of *product testing* arose in the context of error amplification for PCPs [DR06, DG08, IJKW08, IKW09]. Roughly speaking, the question here is to design tests by which a referee can check that the players play according to a *product strategy*, i.e., answer each question independently of the other questions — which would in particular imply that their maximum winning probability must necessarily go down exponentially. The result of [FK00] mentioned above in fact also shows that the strategy of the players must have some product structure, and recently there has been lots of interest in this question [DM10]. In the case of entangled players, however, absolutely nothing was known. We ask: *is there a way to test if the strategy of entangled players is in some sense close to a product strategy?*

Our results

In this work we answer these questions affirmatively and show that the value of entangled games can be decreased through parallel repetition, albeit at a rate polynomial in the inverse of the target value.

Theorem 1 (informal). *For any $s < 1$, $\delta > 0$, and game G , there is a corresponding ℓ -parallel repeated game G' , where $\ell = \text{poly}((1 - s)^{-1}, \delta^{-1})$, such that if the value of G is less than s then the value of G' is at most δ , whereas if the value of G is 1 then this is also true⁷ for the repeated game.*

In the course of our proof of the theorem we also establish that the prover’s strategies have a certain (weak) “serial” or “product” structure (see proof ideas and techniques below). We emphasize that, even though the repeated game obtained through this theorem is not the direct parallel repetition of G , it comes pretty close, and we describe it more precisely below. We also elaborate on the precise conditions under which the existence of a perfect strategy for G implies the same for G' . The kind of parallel repetition we perform depends on the structure of the game G :

Repetition for projection games. If G is a projection game, then the repeated game is obtained by independently playing G on a subset of the repetitions, and playing dummy rounds in the other

⁴XOR games are games with binary answers in which the referee’s decision is based solely on the XOR of the two answers.

⁵Unique games are ones in which the referee applies some permutation to the answers of the second player and accepts if and only they match those from the first player.

⁶Indeed, it is known that it is NP-hard to tell if the entangled value of a given game is 1 or not [KKM⁺08, IKM09]; hence, unless P=NP, for any efficiently computable upper bound on the entangled value, there are necessarily games whose entangled value is strictly less than 1 yet for which that upper bound is 1.

⁷See the discussion following the theorem for some caveats.

repetitions. If, in addition, the game happens to be a *free* game (i.e., a game in which the referee’s distribution on question pairs is a product distribution), then the dummy questions are no longer needed and hence our analysis applies to the *standard* ℓ -fold repetition.

Repetition for general games. If the game G does not have the projection property, then it is necessary to add a number of *consistency* rounds to the repetition. In those rounds the referee sends identical questions to the players, and expects identical answers. As before, the other rounds of the repetition are either the game G or dummy rounds. The consistency questions are added to play the role of the projection constraints. However, since it is not obvious that honest entangled players can answer the consistency questions correctly, even if G had value 1 it is not guaranteed that G' has value 1 anymore. This may or may not be an issue depending on where the original game comes from. In many cases it is known that, if there is a perfect strategy, it does not require any entanglement at all, or maybe it can be achieved using the maximally entangled state. In both cases it is not hard to see that players will be able to answer consistency questions perfectly, and hence our result holds.

Proof idea and techniques. We focus on the case of projection games, as the proof of the other cases does not present additional challenges. The starting point of our proof is the work of Feige and Kilian [FK00], for which the following intuition can be given. Our goal as the referee is to force the players to use a product strategy: we want to make sure that the player chooses his answer to the i th question based only on that question and not on any of the other $\ell - 1$ questions. Towards this end, the referee chooses a (typically large) fraction of the ℓ question pairs to be independently distributed *dummy questions*, the answers to which are ignored. These dummy questions are meant to confuse the players: if they were indeed trying to carefully choose their answer to a certain question by looking at many other questions, now most of those will be completely random and uncorrelated with the other player’s questions, so that such a strategy cannot possibly be helpful. In more detail, Feige and Kilian prove the following dichotomy theorem on the structure of single-player repeated strategies: either the strategy looks rather *random* (in which case the players cannot win the game with good probability — this is where the projection property is used) or it is almost a *serial* or *product* strategy (i.e. the player is playing the rounds independently, and his success probability will suffer accordingly).

Our proof follows a similar structure. However, an important challenge immediately surfaces: the proof in [FK00], and indeed *all* proofs of parallel repetition theorems or direct product tests, make the important initial step of assuming that the player’s strategies are deterministic. And indeed, it is not at all trivial to extend those proofs to even the randomized setting without making this initial simplifying assumption. To give a simple example, an important notion in Feige and Kilian’s proof is that of a *dead* question — simply put, a question to which the player does not give any majority answer, when one goes over all possible ways of completing that specific question into a tuple of questions for the repeated game. It is easily seen that, in the case of a deterministic strategy, dead questions are harmful, as the players are unlikely to satisfy the projection property on them. However, it is just as easily seen that for most randomized strategies, good or bad, *all* questions are dead.

This illustrates the kinds of difficulties that one encounters while trying to show parallel repetition in the case of entangled players, when one cannot simply “fix the randomness”. The issue we just raised is not too hard to solve, but others are more challenging. Indeed the main difficulty is to define a proper notion of *almost serial* for operators, which would in particular incorporate the inherent randomness of quantum strategies. It turns out that the right notion is that of consecutive measurements (rather than tensor products of independent measurements, a tempting but excessively strong possibility). Based on a quantum analogue of Feige and Kilian’s dichotomy theorem, we are able to show that the almost serial condition induces a condition of *almost orthogonality* on the player’s operators. At this point we need to prove a genuinely quantum lemma, which lets us extract a *product* strategy from the almost-orthogonal condition. This novel *orthogonalization lemma* is at the heart of our proof. We obtain that the players approximately perform a series of consecutive measurements, each depending only on the current question. An upper bound on the value of the repeated game then follows.

References

- [AKK⁺08] S. Arora, S. A. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. K. Vishnoi. Unique games on expanding constraint graphs are easy: extended abstract. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 21–28. New York, NY, USA, 2008.
- [BHH⁺08] B. Barak, M. Hardt, I. Haviv, A. Rao, O. Regev, and D. Steurer. Rounding Parallel Repetitions of Unique Games. In *Proc. 49th Annual IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 374–383. 2008.
- [BRR⁺09] B. Barak, A. Rao, R. Raz, R. Rosen, and R. Shaltiel. Strong Parallel Repetition Theorem for Free Projection Games. In *Proc. 13th RANDOM*, pages 352–365. 2009.
- [CSUU08] R. Cleve, W. Slofstra, F. Unger, and S. Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity*, 17:282–299, 2008.
- [DG08] I. Dinur and E. Goldenberg. Locally Testing Direct Product in the Low Error Range. *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 613–622, October 2008.
- [DM10] I. Dinur and O. Meir. Derandomized Parallel Repetition of Structured PCPs. *2010 IEEE 25th Annual Conference on Computational Complexity*, pages 16–27, June 2010.
- [DR06] I. Dinur and O. Reingold. Assignment Testers: Towards a Combinatorial Proof of the PCP Theorem. *SIAM Journal on Computing*, 36(4):975–1024, January 2006.
- [Fei91] U. Feige. On the success probability of two provers in one-round proof systems. In *Proc. 6th IEEE Structure in Complexity Theory*, pages 116–123. 1991.
- [FK00] U. Feige and J. Kilian. Two-Prover Protocols—Low Error at Affordable Rates. *SIAM Journal on Computing*, 30(1):324, 2000.
- [FKO07] U. Feige, G. Kindler, and R. O’Donnell. Understanding Parallel Repetition Requires Understanding Foams. In *IEEE Conference on Computational Complexity*, pages 179–192. 2007.
- [Hol07] T. Holenstein. Parallel repetition: simplifications and no-signaling case. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of Computing*. ACM, 2007.
- [IJKW08] R. Impagliazzo, R. Jaiswal, V. Kabanets, and A. Wigderson. Uniform direct product theorems:simplified, optimized, and derandomized. *Annual ACM Symposium on Theory of Computing*, pages 579–588, 2008.
- [IKM09] T. Ito, H. Kobayashi, and K. Matsumoto. Oracularization and Two-Prover One-Round Interactive Proofs against Nonlocal Strategies. In *Proc. 24th IEEE Conference on Computational Complexity*, pages 217–228. 2009.
- [IKW09] R. Impagliazzo, V. Kabanets, and A. Wigderson. New direct-product testers and 2-query PCPs. *Annual ACM Symposium on Theory of Computing*, pages 131–140, 2009.
- [KKM⁺08] J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick. Entangled Games are Hard to Approximate. *49th Annual IEEE Symposium on Foundations of Computer Science*, pages 447–456, October 2008.
- [KRT08] J. Kempe, O. Regev, and B. Toner. Unique Games with Entangled Provers are Easy. In *Proc. 49th IEEE Symp. on Foundations of Computer Science*, pages 457–466. 2008.

- [Rao08] A. Rao. Parallel Repetition in Projection Games and a Concentration Bound. In *Proc. 40th ACM Symp. on Theory of Computing*, pages 1–10. 2008.
- [Raz98] R. Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27:763–803, 1998. ISSN 0097-5397.
- [Raz08] R. Raz. A Counterexample to Strong Parallel Repetition. In *49th Annual IEEE Symposium on Foundations of Computer Science*, pages 369–373. 2008.
- [RR10] R. Raz and R. Rosen. A Strong Parallel Repetition Theorem for Projection Games on Expanders. *Technical report ECCC TR10-142*, 2010.
- [Ver94] O. Verbitsky. Towards the parallel repetition conjecture. *Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory*, pages 304–307, 1994.