Tight Parallel Repetition Theorems for Public-Coin Arguments Using KL-Divergence

Kai-Min Chung[1,⋆] and Rafael Pass[2,⋆⋆]

[1] Academia Sinica, Taiwan
kmchung@iis.sinica.edu.tw
[2] Cornell University, USA
rafael@cs.cornell.edu

Abstract. We present a new and conceptually simpler proof of a tight parallel-repetition theorem for public-coin arguments [Pass-Venkitasubramaniam, STOC'07], [Håstad et al, TCC'10], [Chung-Liu, TCC'10]. We follow the same proof framework as the previous non-tight parallel-repetition theorem of Håstad et al—which relied on *statistical distance* to measure the distance between experiments—and show that it can be made tight (and further simplified) if instead relying on *KL-divergence* as the distance between the experiments.

We then use this new proof to present the first tight "Chernoff-type" parallel repetition theorem for arbitrary public-coin arguments, demonstrating that parallel-repetition can be used to simultaneously decrease both the soundness and completeness error of *any* public-coin argument at a rate matching the standard Chernoff bound.

1 Introduction

Ideally, we would like the soundness error of an interactive proof [GMR89, BM88] or argument systems [BCC88] to be negligible. But, in many settings, our starting point is a protocol with somewhat large soundness error. For example, to design an interactive argument for a language L, it may be easier to first design a protocol with soundness error $1/2$. This leads to the question of *soundness amplification*: How can we to decrease the soundness error of a given protocol? (Ideally, we would like to simultaneously decrease both the soundness and completeness error; we return to this question shortly.) A natural approach to performing such soundness amplification is through a *direct product theorem*: Roughly speaking, a direct product theorem for a class of problems states that if an adversary can

⋆ Chung is supported in part by NSF Award CNS-1217821, NSF Award CCF-1214844, Pass' Sloan Fellowship, and Ministry of Science and Technology MOST 103-2221-E-001-022-MY3; part of this work was done while being at Cornell University.

⋆⋆ Pass is supported in part by an Alfred P. Sloan Fellowship, Microsoft New Faculty Fellowship, NSF CAREER Award CCF-0746990, NSF Award CCF-1214844, NSF Award CNS-1217821, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211.

Y. Dodis and J.B. Nielsen (Eds.): TCC 2015, Part II, LNCS 9015, pp. 229–246, 2015.

solve an instance of a problem with probability at most δ, then his chance of solving multiple independent instances should decrease exponentially, ideally to δ^k, if we have k independent instances. For the case of interactive proofs/arguments, the two most natural ways of running several instances are *sequential repetition* and *parallel repetition*. In the case of sequential repetitions, we run k instances of some underlying protocol sequentially, one after the other, and the verifier finally accepts if all instances were accepted. It is well-known that sequential repetition decreases the soundness error of both interactive proofs and arguments at an essentially optimal rate; see [BM88, Gol01, DP98]. However, sequential repetition increases the number of communication rounds of the original protocol. In the case of parallel repetition, we instead run the k protocols in parallel. It is known that parallel repetition decrease soundness error at an optimal rate for the case of interactive proofs (i.e., for the case of statistical soundness). For arguments (i.e., computational soundness), however, surprising things start happening: The seminal work of Bellare, Impagliazzo, and Naor [BIN97] demonstrate protocols for which parallel repetition fails to amplify soundness *at all*! These counter examples, however, are for *private-coin* protocols.

On the other hand, for the case of *public-coin* protocols, parallel repetition theorems have been established: Pass and Venkitasubramaniam [PV07] first showed a tight parallel repetition theorem for constant-round protocols. Håstad, Pass, Wikström and Pietrzak [HPWP10] next extended it to arbitrary (i.e., not necessarily constant-round) protocols; the rate at which the soundness decreases, however, was no longer optimal—roughly speaking, when $\delta = (1 - \mu)$, k repetitions decreases the error to $e^{-\Omega(\mu^2 k)}$ as opposed to $\delta^k = e^{-\Omega(\mu k)}$. Finally, Chung and Liu presented an optimal parallel repetition theorem—where k repetitions sufficed to decrease the error to δ^k. A more comprehensive survey of known parallel repetition theorems can be found in Section 1.3.

1.1 Our Results

A New Proof of Tight Parallel Repetition for Public-coin Protocols In this work, we revisit the result of Chung and Liu. Our central contribution is a new proof of their tight parallel repetition theorem. Our proof follows the same framework as the "simple" proof of the *non-tight* parallel-repetition theorem of Håstad et al—which relied on statistical distance to measure the distance between experiments—and shows that it can be made tight (and further simplified) if instead relying on *KL-divergence*[1] as the distance between the experiments. (KL-divergence was previously instrumental for proving tight parallel theorems for two-prover games [BRR+09] in an *information-theoretic* setting. Our new proof demonstrates that also in the computational setting, KL-divergence appears to be the right measure of distance when analyzing reductions through hybrid experiments.) As such, our proof significantly simplifies and "demystifies" the proof of [CL10], which directly analyzed the success probability through an intricate

[1] Recall that $\mathbf{KL}(X||Y) = \sum_{x \in \mathrm{supp}(X)} \Pr[X = x] \cdot \log \frac{\Pr[X=x]}{\Pr[Y=x]}$. For convenience, here we define KL-divergence with log base e. The choice is inconsequential.

inductive argument relying on Holder's inequality, providing little intuition for "why" the reduction works.

Additionally, as we now turn to discussing, our new proof enables considering more general scenarios (whereas the analysis in [CL10] is explicitly set up to analyze the particular direct-product case), and we believe this technique may be useful more broadly.

Tight Chernoff-type Parallel-repetition Theorem for Any Public-coin Protocols. So far we have only discussed *direct-product* parallel repetition theorems, where the parallel verifier accepts iff all parallel sessions are accepting. If the starting protocol also has a large completeness error, then parallel repetition with a "direct product" verifier also increases the completeness error. Ideally, we would like to have a way to simultaneously decrease both the completeness and the soundness error: just as for error reduction of the class **BPP**, the idea is to consider a *threshold verifier*, who accept whenever the fraction of accepting sessions is greater than a certain threshold γ (that is greater than the soundness error δ, or else there is no hope to reduce the soundness error). For error reduction of **BPP**, it follows by a standard Chernoff bound that such an approach works. For interactive arguments, such "Chernoff-type" parallel repetition theorems where first studied by Impagliazzo, Jaiswal, and Kabanets [IJK09] for the case of three-message protocols. Håstad et al. [HPWP10] extend the results of [IJK09] also to public-coin protocols and Chung and Liu [CL10] improved the error decrease rate obtaining "tight" Chernoff-type parallel repetition theorems (matching the parameters of the standard Chernoff bound)[2] in two setting:

- For the case of constant-round protocols (by relying on the direct product parallel repetition theorem of [PV07]).
- When the gap between the threshold γ and the soundness error δ is a constant (i.e., $\gamma - \delta = \Omega(1)$); this is done by relying on a generic reduction to the direct product case, which is only efficient when the gap is a constant.

In particular, for non-constant round protocols, previous result only enabled simultaneously decreasing the completeness and soundness error at a rate matching the standard Chernoff bound whenever the gap between the soundness error and "1-the completeness error" is a constant (as opposed to it being an inverse polynomial). We show that using our new proof technique, the analysis for the direct product case directly extends also to the case of threshold verifiers, yielding a Chernoff-type parallel repetition theorem for *any* public-coin protocol, and *any* threshold $\gamma > \delta$.

Specifically, we demonstrate the following theorem, which matches a "**KL**-version" Chernoff bound, and directly implies both tight direct product theorems and tight Chernoff-type theorems.

Theorem 1 (informal). *For a public-coin interactive argument with soundness error $\delta \in (0,1)$, k-fold parallel repetition with threshold $\gamma > \delta \in (0,1)$ decreases*

[2] Also the standard Chernoff bound is not "optimal" so we content ourselves with matching the parameters of the standard Chernoff bound.

soundness error to $e^{-k \cdot \mathbf{KL}(\gamma||\delta)} + \mathsf{ngl}$*, where* ngl *is a negligible function in the security parameter.*[3] *In particular,*[4]

- *For threshold* $\gamma = 1$ *(the direct product setting), the soundness error is decreased to* $\delta^k + \mathsf{ngl}$.
- *For threshold* $\gamma = (1+\mu)\delta$ *(the "multiplicative" Chernoff-bound setting), the soundness error is decreased to* $e^{-\Omega(\mu^2 \delta k)} + \mathsf{ngl}$ *for* $\mu \in (0,1)$ *and* $e^{-\Omega(\mu \delta k)} + \mathsf{ngl}$ *for* $\mu > 1$.

1.2 Proof Overview

We now explain our new proofs of a tight parallel repetition theorem for public-coin protocols. For simplicity of exposition, we start by focusing on the direct product case, and next extend the analysis to deal with threshold verifiers.

We first set-up some notation. Let us consider a public-coin protocol (P, V) with m rounds, where at each round $j \in [m]$, the verifier V sends a uniformly random message x_j to P, receives back a second-message y_j, and at the end *deterministically* decides to accept or reject based on the transcript $(x_1, y_1, \ldots, x_m, y_m)$. We denote by (P^k, V^k) the k-fold parallel repetition of (P, K); here V^k sends a message $\boldsymbol{x} = (x_{j,1}, \ldots, x_{j,k})$, receives back a message $\boldsymbol{y} = (y_{j,1}, \ldots, y_{j,k})$ at each round $j \in [m]$, and accepts iff $(x_{1,i}, y_{1,i}, \ldots, x_{m,i}, y_{m,i})$ is accepting for every instance $i \in [k]$. We refer to the different parallel executions of the protocol (P, V) inside (P^k, V^k) as the parallel *sessions*.

To prove that parallel repetition reduces the soundness error, we show how to transform any *parallel prover* P^{k*} that convinces V^k with probability $\epsilon \geq 1.1\delta^k$ to a *single-instance* prover P^* that convinces V with probability at least δ. This implies that parallel repetition reduces the soundness error at an essentially optimal rate (from δ to $1.1\delta^k$). We may without loss of generality assume that P^{k*} is deterministic—its optimal random coins can always be fixed non-uniformly.[5]

More precisely, P^* will internally emulate an execution of P^{k*} and use this execution in order to convince an external verifier V. On a high-level, the general strategy is quite straight forward. P^* picks one of the k sessions, i; this session will be externally forwarded (between P^{k*} and V), and all the other sessions, $-i$, will be appropriately emulated internally. In other words, the external verifier V is "embedded" in some session i of V^k, and P^* internally emulates P^{k*} and the remaining $k-1$ sessions $-i$ of V^k while forwarding P^{k*}'s messages $y_{j,i}$'s for session i to V; see Figure 1 for an illustration for the case of a one-round protocol.

Recall that since we have assumed that P^{k*} is deterministic, the interaction between P^{k*} and V^k is determined solely by V^k's message $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m$, where

[3] As shown by [DJMW12], under some cryptographic assumptions, the additive negligible term is necessary.

[4] For the direct product setting, it follows by the fact that $\mathbf{KL}(1||\delta) = \log(1/\delta)$. For the Chernoff-bound setting, it follows by the fact that $\mathbf{KL}((1+\mu)\delta||\delta) = \Theta(\mu^2\delta)$ for $\mu \in (0,1)$, and $\mathbf{KL}((1+\mu)\delta||\delta) = \Theta(\mu\delta)$ for $\mu > 1$.

[5] Alternatively, "close to optimal" coins can be uniformly fixed by sampling.

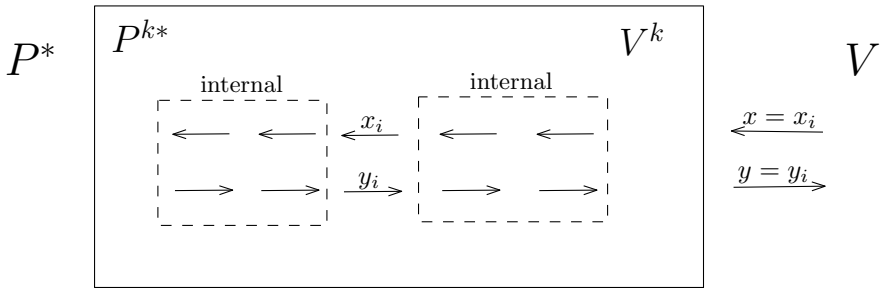

Fig. 1. Interaction between P^* and V for a one-round protocol: P^* embeds the external verifier V in session i of V^k and internally emulates P^{k*} and the remaining $k-1$ sessions $-i$ of V^k while forwarding P^{k*}'s message y for session i to V

each $\boldsymbol{x}_j = (x_{j,1}, \ldots, x_{j,k})$. Thus, P^* needs to decide *the session* i to embed V at beginning, and then at each round j, given an external message $x_{j,i}$, to *choose the remaining* $k-1$ *messages* $\boldsymbol{x}_{j,-i}$. We now recall the *rejection sampling* strategy of [HPWP10].

The Rejection Sampling Strategy. We consider a rejection sampling prover, P^*_{rej}, that selects the session $i \in [k]$ uniformly at random, and then at each round j, upon receiving the external verifier V's message $x_{j,i}$, P^*_{rej} selects $\boldsymbol{x}_{j,-i}$ using rejection sampling as follows: P^*_{rej} repeatedly samples a random continuation of (P^{k*}, V^k) (i.e., samples uniformly random $\boldsymbol{x}_{j,-i}, \boldsymbol{x}_{j+1} \ldots, \boldsymbol{x}_m$)[6] until it finds an *accepting continuation*, i.e., V^k accepts at the end of interaction (or a certain a-prior bound M on the number of samples is reached, in which case P^* aborts and fails). Then, P^* selects the corresponding messages in the accepting continuation as the messages of V_{-i} at round j.

To analyze the success probability of P^*_{rej}, let us first allow P^*_{rej} to make an unbounded number of samples (i.e., set $M = \infty$). Note that in this case, at each round j, P^*_{rej} simply selects $\boldsymbol{x}_{j,-i}$ conditioned on P^{k*} convincing V^k. See Figure 2 for an illustration. As we shall see, if P^{k*} convinces V^k with probability ϵ, then P^* convinces V with probability $\geq \epsilon^{1/k} > \delta$. We then deal with the bounded-sample case at the end (looking forward, as long as we make poly$(1/\epsilon)$ queries, having such a cut-off only slightly affects the success probability of P^*).

The main idea for analyzing (the unbounded sample version of) P^* is to consider an Ideal experiment, where P^* succeeds with probability 1 and next show that the actual execution of (P^*, V), referred to as the Real experiment, and the Ideal experiment are close (using an appropriate choice of a distance measure), from which we can conclude that P^* succeeds with high probability in the Real experiment. Let us start by formalizing the Real experiment.

[6] Note that here we use the fact that the protocol is public coin so that sampling a random continuation is simply sampling uniformly random $\boldsymbol{x}_{j,-i}, \boldsymbol{x}_{j+1} \ldots, \boldsymbol{x}_m$.

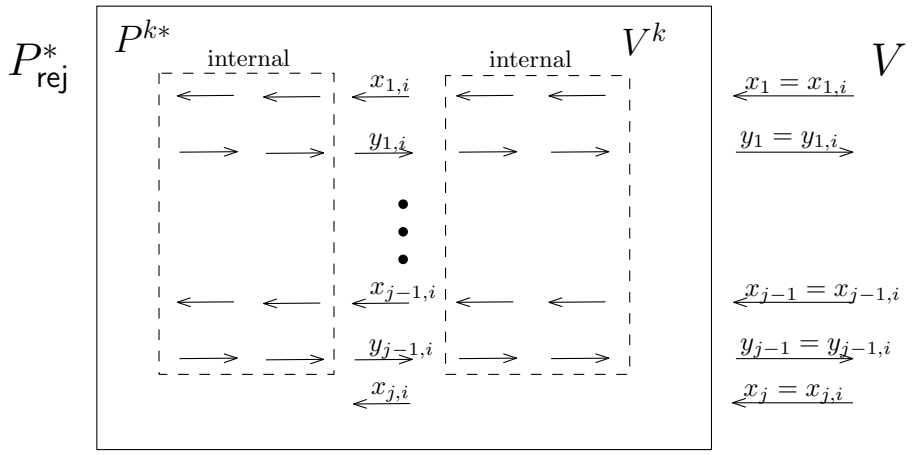

Fig. 2. Interaction between P^*_{rej} and V

The Real *Experiment.* Consider an execution of (P^*_{rej}, V) as follows. At beginning, P^*_{rej} selects a random coordinate $i \in [k]$. Then at each round $j \in [m]$, V selects a uniformly random $x_{j,i}$, and P^*_{rej} selects a random $\boldsymbol{x}_{j,-i}$ conditioned on W using rejection sampling (namely, repeatedly samples a random continuation of (P^{k*}, V^k) until it finds an *accepting continuation*, i.e., V^k accepts at the end of interaction, and selects the corresponding $\boldsymbol{x}_{j,-i}$). If no such $\boldsymbol{x}_{j,-i}$ exists, then P^*_{rej} *fails*. P^*_{rej} *succeeds* if it does not fail. The output of the experiment is defined to be $(i, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)$.

First, note that to prove that parallel repetition works (at an optimal rate) we need to show that P^* convinces V in the Real experiment with probability at least $\epsilon^{1/k}$. Secondly, observe that an equivalent way of defining the output $(i, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)$ of the experiment is as follows: uniformly sample $i \in [k]$, then for each $j \in [m]$, uniformly sample $x_{j,i} \in \{0,1\}^n$, and uniformly sample $\boldsymbol{x}_{-i} \in \{0,1\}^{(k-1)n}$ *conditioned* on P^{k*} convincing V^k.

The Ideal *Experiment.* Let us turn to defining the Ideal experiment. The experiment is defined identically to the Real experiment, except that now we additionally select $x_{j,i}$ conditioned on P^{k*} convincing V^k; that is, uniformly sample $i \in [k]$, then for each $j \in [m]$, uniformly sample $x_{j,i} \in \{0,1\}^n$ *conditioned* on P^{k*} convincing V^k, and uniformly sample $\boldsymbol{x}_{j,-i}$ *conditioned* on $P^{k*}(\boldsymbol{x})$ convincing V^k; again, the output of the experiment is defined to be $(i, \boldsymbol{x})$.

Note that an equivalent way of defining the Ideal experiment is to uniformly sampling $i \in [k]$, and then directly uniformly sample $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)$ conditioned on P^{k*} convincing V^k. Since P^{k*} convinces V^k with positive probability, it thus follows that in the Ideal experiment P^*_{rej} convinces V with probability 1.

Going from Ideal *to* Real. Observe that the only difference between the Real and the Ideal experiments is that at each round $j \in [m]$, in Real $x_{j,i}$'s are sampled

uniformly at random, and in Ideal they are sampled at random conditioned on P^{k*} convincing V^k. The following natural approach is taken in [HPWP10].

Consider a set of m "hybrid" experiments, where in H_j, the messages in the first j rounds are selected just as in Ideal (i.e., both $x_{j',i}$ and $\boldsymbol{x}_{j',-i}$ for $j' \leq j$ are sampled conditioned on P^{k*} convincing V^k), and the remaining $m-j$ rounds are selected just as in Real (i.e., for $j' > j$, only $\boldsymbol{x}_{j',-i}$ is sampled conditioned on P^{k*} convincing V^k, but $x_{j',i}$ is uniformly sampled without any conditioning). Clearly $H_0 = \mathsf{Real}$ and $H_m = \mathsf{Ideal}$. Furthermore, the only difference between two consecutive hybrids $j-1$ and j is whether $x_{j,i}$ is sampled conditioned on P^{k*} convincing V^k or not, where i is uniformly chosen. To bound the distance between the hybrids, [HPWP10] uses the following version of Raz' Lemma [Raz98].

Lemma 1 (Raz' Lemma [Raz98]). *Let $(H, \boldsymbol{X}) = (H, X_1, \ldots, X_k)$ be independent random variables and W be an event. Then,*

$$\frac{1}{k}\sum_{i=1}^{k} \mathbf{SD}((H, X_i)|_W, (H|_W, X_i)) \leq \sqrt{\frac{\log(1/\Pr[W])}{k}}.$$

Let W be the event that P^{k*} convinces V^k. The lemma directly implies that the statistical distance between any two consecutive hybrids H_{j-1} and H_j is at most $\sqrt{(\log(1/\Pr[W]))/k}$. Thus, by the triangle-inequality, the statistical distance between the Real and the Ideal experiments is at most $m \cdot \sqrt{(\log(1/\Pr[W]))/k}$, which yields a lower bound on the success probability of P^*_{rej} in the Real experiment that suffices to demonstrate that parallel repetition reduces the soundness error at an exponential rate.

However, the bound is not tight for two reasons. First, due to the "hybrid argument" we incur a linear loss in the number of rounds m (thus, to make the soundness error small we need the number of parallel repetitions to grow polynomially with the number of rounds in the protocol). [HPWP10] shows how to avoid the loss of m by proving a "multi-round" version of Raz' Lemma, which avoids the round-by-round hybrid argument. But the bound still does is not tight due to the use of statistical distance to measure the distance between Real and Ideal

KL Divergence as a Distance Measure. The crux of our new proof is to instead use Kullback-Leibler divergence (KL divergence, for short) as a distance measure between the Real and Ideal experiments. In fact, the proof of Raz' Lemma (and also the multi-round version in [HPWP10]) first provides a bound on the KL divergence between the random variables, and then arrives a bound on their statistical distance by relying on Pinsker inequality. The "translation" between KL divergence and statistical distance, however, incurs a quadratic loss. By directly working with KL divergence, we can avoid it.[7] By a calculation very similar to the proof of Raz' lemma (and essentially implicit in [HPWP10]), we directly show the following lemma:

[7] A similar phenomena occurred already in the context of parallel repetition for "free" two-prover games; see [BRR+09].

Lemma 2. $\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) \leq \frac{\log(1/\Pr[W])}{k}$.

Let us now show how to get a tight lower bound on the success probability of P^* in the Real experiment. Let $\mathsf{Suc}_{\mathsf{Real}}$ and $\mathsf{Suc}_{\mathsf{Ideal}}$ be indicator variables that indicate, respectively, whether P^* convinces V in the Real and the Ideal experiments.

$$\frac{\log(1/\Pr[W])}{k} \geq \mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) \geq \mathbf{KL}(\mathsf{Suc}_{\mathsf{Ideal}}||\mathsf{Suc}_{\mathsf{Real}}) = 1 \cdot \log \frac{1}{\Pr[\mathsf{Suc}_{\mathsf{Real}} = 1]}, \tag{1}$$

which implies that $\Pr[\mathsf{Suc}_{\mathsf{Real}} = 1] \geq \epsilon^{1/k}$ since $\Pr[W] = \epsilon$. The second inequality follows since applying the same function to two distributions can only decrease their KL divergence, whereas the last equality follows by the definition of KL divergence and the fact that $\Pr[\mathsf{Suc}_{\mathsf{Ideal}} = 1] = 1$. This concludes that P^* convinces V with probability at least $\epsilon^{1/k}$ in the Real experiment.

Dealing with Threshold Verifiers. Our analysis directly extends also to yield tight "Chernoff-type" parallel-repetition theorems where we consider a threshold verifier $V^{k,\gamma}$ that accepts iff more that $\gamma \cdot k$ sessions are accepting. Let us consider the same rejection sampling strategy P^*_{rej} that selects a uniform $i \in [k]$, and then at each round $j \in [m]$, samples $\boldsymbol{x}_{j,-i}$ conditioned on P^{k*} convinces $V^{k,\gamma}$ (note that we do not require $V^{k,\gamma}$ accepts the i-th session). Let us also consider the same Real and Ideal experiments as above. For the same reason, we have $\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) \leq \frac{\log(1/\Pr[W])}{k}$. The only difference is that in the Ideal experiment, we no longer have that the success probability is 1. However, since $V^{k,\gamma}$ accepts only when $\geq \gamma \cdot k$ sessions accepts, and $i \leftarrow [k]$ is uniform and independent of the transcript $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)$, P^*_{rej} convinces V with probability at least γ in Ideal. An analogous calculation shows that if P^{k*} convinces $V^{k,\gamma}$ with probability $\geq e^{-k \cdot \mathbf{KL}(\gamma||\delta)}$, then P^*_{rej} convinces V with probability at least δ in the Real experiment.

Handling The Bounded-Sample Case. In our analysis so far we have assumed that P^* can make an unbounded number of samples. Let us now show that its success probability is still high even if we impose a polynomial bound M on the number of samples it can make (and thus P^* becomes efficient). Let us first consider the Ideal experiment. The main observation is that, in the Ideal experiment, *in expectation*, P^* only needs to make $1/\epsilon$ samples to pick $\boldsymbol{x}_{j,-i}$ conditioned on $P^{k*}(\boldsymbol{x})$ convincing V^k for every $j \in [m]$ (since the prefix $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{j-1}, x_{j,i})$ is also picked conditioned on P^{k*} convincing V^k, and P^{k*} convinces V^k with probability ϵ). Thus, if the allowed number of samples M is sufficiently larger than $1/\epsilon$, then by the Markov inequality, P^* can successfully convince V^k with probability "almost" 1, even if we restrict P^* to use at most M samples.[8] Since the

[8] Since $M > 1/\epsilon$, we only get an efficient reduction as long as ϵ is an inverse polynomial. As a consequence, parallel repetition of arguments cannot decrease the soundness error beyond being "negligible". As shown by [DJMW12], under some cryptographic assumptions, this is inherent.

Ideal and the Real experiments are statistically close, this directly yields a lower bound on the success probability of P^* in the Real experiment. But as we saw, working with statistical distance does not give the tight bound. To obtain a tight bound, we again work with KL divergence. Here, the only difference is that $\Pr[\mathsf{Suc}_{\mathsf{Ideal}} = 1]$ is no longer 1, but can be made arbitrarily (inverse polynomially) close to 1 by increasing M. This is sufficient to conclude that $\Pr[\mathsf{Suc}_{\mathsf{Real}} = 1]$ can be made arbitrarily (inverse polynomially) close to $\epsilon^{1/k}$ as well (since the KL divergence of two binary random variables is a "smooth" function of the probabilities of both random variables).

1.3 Related Works: When Parallel Repetition Works

Let us briefly summarize the class of argument systems for which parallel repetitions is known to decrease the soundness error.

- The seminal work of Yao [Yao82] on hardness amplification of one-way functions can be viewed as showing that parallel repetition reduces the soundness error at an optimal rate for all two-message argument systems for which the verifier's decision to accept or reject is a public function of the transcript; that is, the verifier is not employing any secret randomness to decide whether to accept or reject—we refer to such protocols as being *publicly verifiable*. An important special case of publicly-verifiable protocol are *public-coin* protocols (a.k.a. *Arthur-Merlin* protocols [BM88]) where in each round the verifier simply tosses some random coins and directly sends them to the prover (that is, the verifier doesn't employ any secret randomness).
- The seminal work of Bellare, Impagliazzo and Naor [BIN97] was the first one to explicitly study parallel repetition of argument systems and demonstrated that parallel repetition reduces the soundness error for all three-message protocols (not just publicly-verifiable ones). The results of [BIN97] demonstrated that parallel repetition reduces the soundness error of such protocols at an exponential rate, but did not establish an optimal rate (i.e., reducing the soundness error from ϵ to ϵ^k). Nevertheless, the more recent work by Canetti, Halevi and Steiner [CHS05] shows that parallel repetition indeed reduces the soundness error at an optimal rate for this class of protocols.
- [BIN97, PW07] demonstrate 4-message protocols for which parallel repetition fails; in these protocols, the verifier uses "secret randomness". Pass and Venkitasubramaniam [PV12] turn to considering public-coin protocols, and demonstrate that parallel repetition decreases the soundness error for *constant-round* public-coin protocols at an optimal rate.
- Håstad, Pass, Wikström and Pierzak [HPWP08] show that parallel repetition, in fact, works for *all* (not necessarily constant-round) public-coin protocols, and decreases the soundness error at an exponential rate. Chung and Liu [CL10] demonstrate that it in fact decreases at an optimal rate.
- [HPWP08] consider a generalization of both public-coin and three-message protocol, called *simulatable* protocols—where roughly speaking the verifier's

messages can be predicted without knowing its randomness—and demonstrate that parallel repetition reduces the error at an exponential rate; an improved "nearly" optimal rate (reducing the soundness error from ϵ to $\epsilon^{k/2}$) is obtained by [CL10].
- [HPWP08], and more explicitly [CL10], also consider protocols satisfying a "computational" simulatability property and demonstrate that parallel repetition reduces the soundness error at a nearly optimal rate also for such protocols.
- The elegant work of Haitner [Hai09] considers a certain class of protocols with "random-terminating" verifiers and demonstrates that parallel repetition reduces the soundness error at an exponential rate for such protocols; random-terminating protocols are important since *any* argument can be turned into a random-terminating one, while only slightly increasing the soundness error. [HPWP10] provide a generalization, called δ-simulatable protocols—where, very roughly speaking, we only need to predict a δ-fraction of the verifier's messages—that encompasses both simulatable and random-terminating protocols, and demonstrate that parallel repetition decreases the soundness error at an exponential rate. Optimal, or even "nearly" optimal, parallel repetition theorems for δ-simulatable (or even random-terminating) protocols are not known.

2 Preliminaries

Throughout the paper, all log are base e.

2.1 Interactive Arguments

Definition 1 (Interactive Proofs/Arguments). *A pair of interactive algorithms (P, V) is an **interactive proof** for a* **NP** *language L with **completeness error** c and **soundness error** s if it satisfies the following properties:*

- *Completeness: For all $x \in L$ with* **NP** *witness w,*

$$\Pr[\langle P(w), V\rangle(x) = 1] = 1 - c(|x|).$$

- *Soundness: For all adversarial provers P^*, and for every all $x \notin L$,*

$$\Pr[\langle P^*, V\rangle(x) = 1] \leq s(|x|).$$

*where $\langle P, V\rangle(x)$ denotes the output of V after communicating with P if both players get x as a common input. (P, V) is an **interactive argument** for L if P runs in polynomial time and the soundness property holds only against all non-uniform polynomial-time adversarial provers P^*. (P, V) is public-coin if verifier's messages are uniformly random strings; otherwise, (P, V) is private-coin.*

Definition 2 (Parallel Repetition with Threshold Verifiers). *Let (P, V) be an interactive protocol. Let $k \in \mathbb{N}$ and $\gamma \in (0, 1)$. We use $(P^k, V^{k,\gamma})$ to denote k-fold parallel repetition of (P, V) with threshold γ, where P^k and $V^{k,\gamma}$ interact by executing k copies of (P, V) in parallel and at the end of execution, $V^{k,\gamma}$ accepts iff at least $\gamma \cdot k$ copies accept. For the direct product case with $\gamma = 1$, we use V^k to denote $V^{k,1}$.*

2.2 Kullback-Leibler Divergence

Here we review the definition and basic properties of Kullback-Leibler divergence.

Definition 3. *Let X and Y be discrete random variables over a finite support $[N]$. The Kullback-Leibler divergence (KL divergence for short) of X from Y is defined to be*

$$\mathbf{KL}(X||Y) = \sum_{x \in [N]} \Pr[X = x] \log \frac{\Pr[X = x]}{\Pr[Y = x]}.$$

For $p, q \in (0, 1)$, we use $\mathbf{KL}(p||q)$ to denote the KL divergence $\mathbf{KL}(X||Y)$ of two binary random variables X and Y with $\Pr[X = 1] = p$ and $\Pr[Y = 1] = q$.

The following properties of KL divergence can be found in any Information Theory textbook (e.g., [CT06]). We first recall the chain rule for KL divergence.

Lemma 3 (chain rule). *Let (X_1, X_2) and (Y_1, Y_2) be random variables. We have*

$$\mathbf{KL}((X_1, X_2)||(Y_1, Y_2)) = \mathbf{KL}(X_1||Y_1) + \mathop{\mathrm{E}}_{x \leftarrow X_1}[\mathbf{KL}(X_2|_{X_1=x}||Y_2|_{Y_1=x})].$$

The following lemma says that applying a deterministic function can only decrease the KL divergence.

Lemma 4. *Let X and Y be random variables and f a deterministic function. We have*

$$\mathbf{KL}(f(X)||f(Y)) \leq \mathbf{KL}(X||Y).$$

The following lemma allows us to decompose the KL divergence of two joint distributions by the sum of the KL divergence of their marginals.

Lemma 5 ([Hol09], Lemma 4.2). *Let $\boldsymbol{X} = (X_1, \ldots, X_k)$ be* independent *random variables, and $\boldsymbol{Y} = (Y_1, \ldots, Y_k)$ be random variables.*

$$\sum_{i=1}^{k} \mathbf{KL}(Y_i||X_i) \leq \mathbf{KL}(\boldsymbol{Y}||\boldsymbol{X})$$

The following lemma bounds how much conditioning can creates the KL divergence.

Lemma 6. *Let X be a random variable and W a (probabilistic) event.*

$$\mathbf{KL}(X|_W||X) \leq \log \frac{1}{\Pr[W]}.$$

The following simple lemma bounds the sensitivity of KL divergence of two binary random variables with respect to the first coordinate, which will be useful for us. The lemma follows by the fact that KL divergence is a smooth function, and is proved by a straightforward calculation. For the sake of completeness, we provide a proof in the appendix.

Lemma 7. *For every $p, q, \eta \in (0,1)$ such that $\eta \leq \min\{p/2, q\}$, we have*

$$\mathbf{KL}(p||q) - \mathbf{KL}(p-\eta||q) \leq \Theta\left(\eta \cdot \log(1/\eta)\right)$$

2.3 A Lemma on Sampling

The following simple lemma is taken from Håstad et al. [HPWP10]; for self-containment, we recall the proof.

Lemma 8 ([HPWP10]). *Let (X, Y) be a joint distribution over some finite domain. Let W be a deterministic event on (X, Y). Consider the following experiment:*

- *Sample $x \leftarrow X|_W$.*
- *Sample $y \leftarrow Y|_{W \wedge X=x}$ using rejection sampling; i.e., sample i.i.d. $y_1, y_2, \ldots \leftarrow Y|_{X=x}$ and outputs the first y_t such that $(x, y_t) \in W$.*

Let T be the number of sample used in the rejection sampling. We have $\mathrm{E}[T] = \frac{1}{\Pr[W]}$.

Proof. The lemma follows by the following calculation.

$$\begin{aligned}
\mathrm{E}[T] &= \sum_x \Pr[X = x|W] \cdot \mathrm{E}[T|X = x] \\
&= \sum_x \Pr[X = x|W] \cdot \frac{1}{\Pr[W|X = x]} \\
&= \sum_x \frac{\Pr[X = x \wedge W]}{\Pr[W]} \cdot \frac{\Pr[X = x]}{\Pr[W \wedge X = x]} \\
&= \sum_x \frac{\Pr[X = x]}{\Pr[W]} = \frac{1}{\Pr[W]}.
\end{aligned}$$

3 Proof of the Parallel Repetition Theorem

In this section, we present the formal of our tight Chernoff-type parallel repetition theorem for public-coin protocols.

Theorem 2. *Let (P, V) be a public-coin interactive argument for a language L. There exists an oracle adversarial prover $P^{(\cdot)*}$ such that for every $k \in \mathbb{N}$, input $z \in \{0,1\}^*$, every $\gamma, \delta, \xi \in (0,1)$ with $\gamma > \delta$, and every deterministic parallel adversarial prover P^{k*}, if*

$$\Pr[\langle P^{k*}, V^{k,\gamma}\rangle(z) = 1] \geq \epsilon \stackrel{\text{def}}{=} (1+\xi) \cdot e^{-k \cdot \mathbf{KL}(\gamma||\delta)},$$

then

$$\Pr[\langle P^{(P^{k*})*}(k, \gamma, \delta, \xi), V\rangle(z) = 1] \geq \delta.$$

Furthermore, $P^{(\cdot)*}$ runs in time $\mathrm{poly}(|z|, k, \epsilon^{-1}, \xi^{-1}, (\gamma-\delta)^{-1})$ given oracle access to P^{k*}.

Note that in the direct product setting with $\gamma = 1$, $\mathbf{KL}(1||\delta) = \log(1/\delta)$ so $e^{-k \cdot \mathbf{KL}(\gamma||\delta)} = \delta^k$. Thus, the above theorem implies a tight direct product theorem as a corollary. For the "multiplicative" Chernoff-bound setting with $\gamma = (1+\mu)\delta$, we have $\mathbf{KL}((1+\mu)\delta||\delta) = \Theta(\mu^2\delta)$ for $\mu \in (0,1)$, and $\mathbf{KL}((1+\mu)\delta||\delta) = \Theta(\mu\delta)$ for $\mu > 1$, which implies bounds $e^{-\Omega(\mu^2\delta k)}$ and $e^{-\Omega(\mu\delta k)}$, respectively. This matches the usual multiplicative Chernoff bounds.

Proof. Let m denote the round complexity of (P, V). Let us consider a $P_{\mathsf{rej}}^{(\cdot)*}$ that interacts with V by the aforementioned *rejection sampling* with $M = \Theta(\frac{k \cdot m}{\epsilon \cdot \xi \cdot (\gamma - \delta)} \cdot \log \frac{k}{\delta \cdot \xi})$. Specifically, P_{rej}^*, selects the session $i \in [k]$ uniformly at random, and then at each round j, upon receiving the external verifier V's message $x_{j,i}$, P_{rej}^* selects $\boldsymbol{x}_{j,-i}$ using rejection sampling as follows: P_{rej}^* repeatedly samples a random continuation of $(P^{k*}, V^{k,\gamma})$ until it finds an *accepting continuation*, i.e., $V^{k,\gamma}$ accepts at the end of interaction (note that we do not require $V^{k,\gamma}$ accepts the i-th coordinate), or $M = \Theta(\frac{k \cdot m}{\epsilon \cdot \xi \cdot (\gamma - \delta)} \cdot \log \frac{k}{\delta \cdot \xi})$ samples is reached, in which case P_{rej}^* aborts and fails. Then, P_{rej}^* selects the corresponding messages in the accepting continuation as the messages of V_{-i} at round j.

By inspection, $P_{\mathsf{rej}}^{(\cdot)*}$ runs in time $\mathrm{poly}(|z|, k, \epsilon^{-1}, \xi^{-1}, (\gamma-\delta)^{-1})$ on input $z, k, \gamma, \delta, \xi$. It remains to show that if P^{k*} convinces $V^{k,\gamma}$ with probability at least ϵ, then $P_{\mathsf{rej}}^{(\cdot)*}$ convinces V with probability at least δ. Let W denote the event that P^{k*} convinces $V^{k,\gamma}$ in the execution of $\langle P^{k*}, V^{k,\gamma}\rangle(z)$. We consider the following Real experiment, which is the same as the execution of $\langle P_{\mathsf{rej}}^{(P^{k*})*}(k, \gamma, \delta, \xi), V\rangle(z)$ except that P_{rej}^* takes an unbounded number of samples (i.e., set $M = \infty$).

The Real *Experiment.* Consider an execution of (P_{rej}^*, V) as follows. At beginning, P_{rej}^* selects a random coordinate $i \in [k]$. Then at each round $j \in [m]$, V selects a uniformly random $x_{j,i}$, and P_{rej}^* selects a random $\boldsymbol{x}_{j,-i}$ conditioned on W using rejection sampling (namely, repeatedly samples a random continuation of $(P^{k*}, V^{k,\gamma})$ until it finds an *accepting continuation*, i.e., $V^{k,\gamma}$ accepts at the end of interaction, and selects the corresponding $\boldsymbol{x}_{j,-i}$). Let T_j denotes the number of samples P_{rej}^* takes. If no such $\boldsymbol{x}_{j,-i}$ exists, then P_{rej}^* *fails*, and we set $T_j = \infty$ and all remaining $\boldsymbol{x}_{j,-i}, \boldsymbol{x}_{j+1}, \ldots, \boldsymbol{x}_m = \perp$. P_{rej}^* *succeeds* if it does not fail. The output of the experiment is defined to be $(i, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)$.

Note that the event that $P^{(\cdot)*}$ convinces V in $\langle P^{(P^{k*})*}(k,\gamma,\delta,\xi), V\rangle(z)$ corresponds to the event that in the Real experiment, P^* succeeds *and* $T_j \leq M$ for every $j \in [m]$. Let $\mathsf{Suc}_{\mathsf{Real}}$ be the indicator random variable of this event. Our goal is to lower bound

$$\Pr[\langle P^{(P^{k*})*}(k,\gamma,\delta,\xi), V\rangle(z) = 1] = \Pr[\mathsf{Suc}_{\mathsf{Real}} = 1].$$

We next compare it with an Ideal experiment, which is identical to the Real experiment, except that the messages $x_{1,i}, \ldots, x_{m,i}$ are also selected *conditioned* on W.

The Ideal *Experiment.* At beginning, P^*_{rej} selects a random coordinate $i \in [k]$. Then at each round $j \in [m]$, V selects a random $x_{j,i}$ conditioned on W, and P^*_{rej} selects a random $\boldsymbol{x}_{j,-i}$ conditioned on W using rejection sampling. Let T_j denotes the number of samples P^*_{rej} takes. The output of the experiment is defined to be $(i, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)$.

Note that sampling random $x_{1,i}, \boldsymbol{x}_{1,-i}, \ldots, x_{m,i}, \boldsymbol{x}_{m,-i}$ conditioned on W step by step is equivalent to sampling the whole $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m$ conditioned on W. Thus, the output distribution of the Ideal experiment is simply a uniformly random coordinate $i \in [k]$ and a uniformly random *accepting* transcript $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_m)$. Let $\mathsf{Suc}_{\mathsf{Ideal}}$ be the corresponding indicator random variable of $\mathsf{Suc}_{\mathsf{Real}}$ in the Ideal experiment; that is, $\mathsf{Suc}_{\mathsf{Ideal}}$ is the indicator random variable of the event that P^*_{rej} convinces V and $T_j \leq M$ for every $j \in [m]$.

In what follows, we will show that (i) $\Pr[\mathsf{Suc}_{\mathsf{Ideal}} = 1] \geq \gamma - (m/M\epsilon)$ and (ii) $\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) \leq (\log(1/\Pr[W]))/k$, and derive the desired lower bound on $\Pr[\mathsf{Suc}_{\mathsf{Real}} = 1]$ from them.

Lemma 9. $\Pr[\mathsf{Suc}_{\mathsf{Ideal}} = 1] \geq \gamma - (m/M\epsilon)$.

Proof. Note that in the Ideal experiment, for every $i \in [k]$ and $j \in [m]$, the prefix $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{j-1}, x_{j,i})$ is chosen randomly conditioned on W and then P^*_{rej} selects a random $\boldsymbol{x}_{j,-i}$ conditioned on W using rejection sampling. Applying Lemma 8 with $X = (\boldsymbol{X}_1, \ldots, \boldsymbol{X}_{j-1}, X_{j,i})$, $Y = X_{j,-i}$ and event W implies that $\mathrm{E}[T_j] = 1/\Pr[W] \leq 1/\epsilon$ for every $j \in [m]$. By the Markov inequality, we have $\Pr[T_j \leq M] \geq 1 - 1/(M\epsilon)$ for every $j \in [m]$. Also note that i is uniformly random and independent of $\boldsymbol{x}$ and T_j's so the probability that a random coordinate i is accepting is at least γ. Thus, it follows by an union bound that $\Pr[\mathsf{Suc}_{\mathsf{Ideal}} = 1] \geq \gamma - (m/M\epsilon)$.

Lemma 10. $\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) \leq (\log(1/\Pr[W]))/k$.

Proof. It is instructive to first prove the one-round case (i.e., $m = 1$), which is equivalent to the KL-version of Raz' Lemma. In this case by definition, $\mathsf{Ideal} = (I, \boldsymbol{X}_1|_W)$ and $\mathsf{Real} = (I, X_{1,I}, \boldsymbol{X}_{1,-I}|_{W,X_{1,I}})$. By applying the chain rule (Lemma 3), we have

$$\begin{aligned}\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) &= \mathbf{KL}(I||I) + \mathop{\mathrm{E}}_{I}\left[\mathbf{KL}\left(\boldsymbol{X}_1|_W||(X_{1,I}, \boldsymbol{X}_{1,-I}|_{W,X_{1,I}})\right)\right] \\ &= \frac{1}{k}\sum_{i=1}^{k}\mathbf{KL}\left(\boldsymbol{X}_1|_W||(X_{1,i}, \boldsymbol{X}_{1,-i}|_{W,X_{1,i}})\right).\end{aligned}$$

For each term $\mathbf{KL}(\boldsymbol{X}_1|_W||(X_{1,i}, \boldsymbol{X}_{1,-i}|_{W,X_{1,i}}))$, by applying the chain rule again, we have

$$\begin{aligned}
&\mathbf{KL}\left(\boldsymbol{X}_1|_W||(X_{1,i}, \boldsymbol{X}_{1,-i}|_{W,X_{1,i}})\right) \\
&= \mathbf{KL}(X_{1,i}|_W||X_{1,i}) + \mathop{\mathrm{E}}_{X_{1,i}|_W}[\mathbf{KL}(\boldsymbol{X}_{1,-i}|_{W,X_{1,i}}||\boldsymbol{X}_{1,-i}|_{W,X_{1,i}})] \\
&= \mathbf{KL}(X_{1,i}|_W||X_{1,i}).
\end{aligned}$$

Applying Lemma 5,

$$\begin{aligned}
\frac{1}{k}\sum_{i=1}^{k}\mathbf{KL}\left(\boldsymbol{X}_1|_W||(X_{1,i}, \boldsymbol{X}_{1,-i}|_{W,X_{1,i}})\right) &= \frac{1}{k}\sum_{i=1}^{k}\mathbf{KL}(X_{1,i}|_W||X_{1,i}) \\
&\leq \frac{1}{k}\mathbf{KL}(\boldsymbol{X}_1|_W||\boldsymbol{X}_1).
\end{aligned}$$

Therefore, by Lemma 6,

$$\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) \leq \frac{1}{k}\mathbf{KL}(\boldsymbol{X}_1|_W||\boldsymbol{X}_1) \leq \frac{\log(1/\Pr[W])}{k}.$$

We proceed to consider the general case, which is proved by the same calculation, except that we first apply an additional chain rule to break up terms corresponding to each round.

$$\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) = \sum_{j=1}^{m}\mathop{\mathrm{E}}_{I,\boldsymbol{X}_{<j}|_W}\left[\mathbf{KL}(\boldsymbol{X}_j|_{W,\boldsymbol{X}_{<j}}||(X_{j,I}|_{W,\boldsymbol{X}_{<j}}, \boldsymbol{X}_{j,-I}|_{W,\boldsymbol{X}_{<j},X_{j,I}}))\right].$$

Now, for each term, the same calculation as before using the chain rule and Lemma 5 shows that

$$\begin{aligned}
&\mathop{\mathrm{E}}_{I,\boldsymbol{X}_{<j}|_W}\left[\mathbf{KL}(\boldsymbol{X}_j|_{W,\boldsymbol{X}_{<j}}||(X_{j,I}|_{W,\boldsymbol{X}_{<j}}, \boldsymbol{X}_{j,-I}|_{W,\boldsymbol{X}_{<j},X_{j,I}}))\right] \\
&\leq \frac{1}{k}\mathop{\mathrm{E}}_{\boldsymbol{X}_{<j}|_W}\left[\mathbf{KL}(\boldsymbol{X}_j|_{W,\boldsymbol{X}_{<j}}||\boldsymbol{X}_j|_{\boldsymbol{X}_{<j}})\right].
\end{aligned}$$

Applying another chain rule and Lemma 6 gives,

$$\begin{aligned}
\mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) &\leq \frac{1}{k}\mathop{\mathrm{E}}_{\boldsymbol{X}_{<j}|_W}\left[\mathbf{KL}(\boldsymbol{X}_j|_{W,\boldsymbol{X}_{<j}}||\boldsymbol{X}_j|_{W,\boldsymbol{X}_{<j}})\right] \\
&= \frac{1}{k}\mathbf{KL}(\boldsymbol{X}_{\leq m}|_W||\boldsymbol{X}_{\leq m}) \\
&\leq \frac{\log(1/\Pr[W])}{k}
\end{aligned}$$

We now derive the desired lower bound on the probability $\Pr[\mathsf{Suc}_{\mathsf{Real}} = 1]$ using Lemma 9 and 10. Let $q = \Pr[\mathsf{Suc}_{\mathsf{Real}} = 1]$ and $\eta = m/M\epsilon$. Since our goal is to lower bound q by δ and $\gamma - \eta \geq \delta$, we can assume w.l.o.g., that $q \leq \gamma - \eta$. Lemma 10 implies that

$$\mathbf{KL}(\gamma - \eta||q) \leq \mathbf{KL}(\mathsf{Suc}_{\mathsf{Ideal}}||\mathsf{Suc}_{\mathsf{Real}}) \leq \mathbf{KL}(\mathsf{Ideal}||\mathsf{Real}) \leq (\log(1/\Pr[W]))/k,$$

where the second inequality follows since applying the same function to two distributions can only decrease their KL divergence. Now, note that the fact that $\Pr[W] \geq (1+\xi)e^{-k\cdot\mathbf{KL}(\gamma||\delta)}$ and Lemma 7 implies that

$$\begin{aligned}\Pr[W] &\geq (1+\xi)\cdot e^{-k\cdot\mathbf{KL}(\gamma||\delta)} \geq e^{-k\cdot\mathbf{KL}(\gamma||\delta)+\xi/2} \\ &\geq e^{-k\cdot(\mathbf{KL}(\gamma-\eta||\delta)+\Theta(\eta\cdot\log(1/\eta)))+\xi/2} \geq e^{-k\cdot\mathbf{KL}(\gamma-\eta||\delta)},\end{aligned}$$

where the last inequality follows by the fact that $k\cdot\Theta(\eta\cdot\log(1/\eta)) \leq \xi/2$ (which follows by the choice of M). Combining the above inequalities, we have $\mathbf{KL}(\gamma-\eta||q) \leq \mathbf{KL}(\gamma-\eta||\delta)$, which implies $q \geq \delta$.

References

[BCC88] Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. Journal of Computer and System Sciences 37(2), 156–189 (1988)

[BIN97] Bellare, M., Impagliazzo, R., Naor, M.: Does parallel repetition lower the error in computationally sound protocols? In: FOCS, pp. 374–383 (1997)

[BM88] Babai, L., Moran, S.: Arthur-Merlin games: A randomized proof system, and a hierarchy of complexity classes. J. Comput. Syst. Sci. 36(2), 254–276 (1988)

[BRR+09] Barak, B., Rao, A., Raz, R., Rosen, R., Shaltiel, R.: Strong parallel repetition theorem for free projection games. In: Dinur, I., Jansen, K., Naor, J., Rolim, J. (eds.) APPROX and RANDOM 2009. LNCS, vol. 5687, pp. 352–365. Springer, Heidelberg (2009)

[CHS05] Canetti, R., Halevi, S., Steiner, M.: Hardness amplification of weakly verifiable puzzles. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 17–33. Springer, Heidelberg (2005)

[CL10] Chung, K.-M., Liu, F.-H.: Parallel repetition theorems for interactive arguments. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 19–36. Springer, Heidelberg (2010)

[CT06] Cover, T.M., Thomas, J.A.: Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing). Wiley-Interscience (2006)

[DJMW12] Dodis, Y., Jain, A., Moran, T., Wichs, D.: Counterexamples to hardness amplification beyond negligible. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 476–493. Springer, Heidelberg (2012)

[DP98] Damgård, I., Pfitzmann, B.: Sequential iteration of interactive arguments and an efficient Zero-Knowledge argument for NP. In: Larsen, K.G., Skyum, S., Winskel, G. (eds.) ICALP 1998. LNCS, vol. 1443, pp. 772–783. Springer, Heidelberg (1998)

[GMR89] Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof-systems. SIAM Journal on Computing 18(1), 186–208 (1989)

[Gol01] Goldreich, O.: Foundations of Cryptography. Basic tools. Cambridge University Press (2001)

[Hai09] Haitner, I.: A parallel repetition theorem for any interactive argument. In: FOCS (2009)

[Hol09] Holenstein, T.: Parallel repetition: Simplification and the no-signaling case. Theory of Computing 5(1), 141–172 (2009)

[HPWP08] Håstad, J., Pass, R., Wikström, D., Pietrzak, K.: An efficient parallel repetition theorem (2008) (unpublished manuscript)
[HPWP10] Håstad, J., Pass, R., Wikström, D., Pietrzak, K.: An efficient parallel repetition theorem. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 1–18. Springer, Heidelberg (2010)
[IJK09] Impagliazzo, R., Jaiswal, R., Kabanets, V.: Chernoff-type direct product theorems. J. Cryptology 22(1), 75–92 (2009)
[PV07] Pass, R., Venkitasubramaniam, M.: An efficient parallel repetition theorem for arthur-merlin games. In: STOC, pp. 420–429 (2007)
[PV12] Pass, R., Venkitasubramaniam, M.: A parallel repetition theorem for constant-round arthur-merlin proofs. Transactions on Computation Theory 4(4), 10 (2012)
[PW07] Pietrzak, K., Wikström, D.: Parallel repetition of computationally sound protocols revisited. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 86–102. Springer, Heidelberg (2007)
[Raz98] Raz, R.: A parallel repetition theorem. SIAM J. Comput. 27(3), 763–803 (1998)
[Yao82] Yao, A.C.-C.: Theory and applications of trapdoor functions (extended abstract). In: FOCS, pp. 80–91 (1982)

Proof of Lemma 7

Lemma 11 (Lemma 7 Restated). *For every $p, q, \eta \in (0,1)$ such that $\eta \leq \min\{p/2, q\}$, we have*

$$\mathbf{KL}(p||q) - \mathbf{KL}(p-\eta||q) \leq \Theta\left(\eta \cdot \log(1/\eta)\right)$$

Proof. By definition,

$$\begin{aligned}
\mathbf{KL}(p||q) &= p\log\frac{p}{q} + (1-p)\log\frac{1-p}{1-q} \\
&= p\log p + p\log\frac{1}{q} + (1-p)\log(1-p) + (1-p)\log\frac{1}{1-q} \\
\mathbf{KL}(p-\eta||q) &= (p-\eta)\log\frac{p-\eta}{q} + (1-p+\eta)\log\frac{1-p+\eta}{1-q} \\
&= (p-\eta)\log(p-\eta) + (p-\eta)\log\frac{1}{q} + (1-p+\eta)\log(1-p+\eta) \\
&\quad + (1-p+\eta)\log\frac{1}{1-q}
\end{aligned}$$

By further expanding, we have

$$\begin{aligned}
(p-\eta)\log(p-\eta) &= p\log(p-\eta) - \eta\log(p-\eta) \\
&= p\log p + p\log(1-\frac{\eta}{p}) - \eta\log(p-\eta) \\
(p-\eta)\log\frac{1}{q} &= p\log\frac{1}{q} - \eta\log\frac{1}{q}
\end{aligned}$$

$$\begin{aligned}(1-p+\eta)\log(1-p+\eta) &= (1-p)\log(1-p+\eta) + \eta\log(1-p+\eta)\\ &= (1-p)\log(1-p) + (1-p)\log(1+\frac{\eta}{1-p})\\ &\quad + \eta\log(1-p+\eta)\\ (1-p+\eta)\log\frac{1}{1-q} &= (1-p)\log\frac{1}{1-q} + \eta\log\frac{1}{1-q}\end{aligned}$$

Therefore,

$$\begin{aligned}&\mathbf{KL}(p||q) - \mathbf{KL}(p-\eta||q)\\ &= -p\log(1-\frac{\eta}{p}) + \eta\log(p-\eta) + \eta\log\frac{1}{q} - (1-p)\log(1+\frac{\eta}{1-p})\\ &- \eta\log(1-p+\eta) - \eta\log\frac{1}{1-q}\\ &\leq -p\log(1-\frac{\eta}{p}) + \eta\log\frac{1}{q} - \eta\log(1-p+\eta)\\ &\leq 2\eta + \eta\log\frac{1}{q} - \eta\log\eta \leq \Theta(\eta\log(1/\eta)),\end{aligned}$$

where the first inequality follows by dropping negative terms, the second inequality follows by the monotonicity of logarithm and using Taylor expansion, and the last inequality uses $\eta < q$.