Parallel Repetition Theorems for Interactive Arguments*

Kai-Min Chung[1,**] and Feng-Hao Liu[2]

[1] School of Engineering & Applied Sciences, Harvard University, Cambridge MA, USA
kmchung@fas.harvard.edu
[2] Department of Computer Science, Brown University, Providence RI, USA
fenghao@cs.brown.edu

Abstract. We study efficient parallel repetition theorems for several classes of interactive arguments and obtain the following results:

1. We show a *tight* parallel repetition theorem for public-coin interactive arguments by giving a tight analysis for a reduction algorithm of Håstad et al. [HPPW08]. That is, n-fold parallel repetition decreases the soundness error from δ to δ^n. The crux of our improvement is a new analysis that avoid using Raz's Sampling Lemma, which is the key ingredient to the previous results.
2. We give a new security analysis to strengthen a parallel repetition theorem of Håstad et al. for a more general class of arguments. We show that n-fold parallel repetition decreases the soundness error from δ to $\delta^{n/2}$, which is almost tight. In particular, we remove the dependency on the number of rounds in the bound, and as a consequence, extend the "concurrent" repetition theorem of Wikström [Wik09] to this model.
3. We obtain a way to turn *any* interactive argument to one in the class above using fully homomorphic encryption schemes. This gives a way to amplify the soundness of any interactive argument without increasing the round complexity.
4. We give a simple and generic transformation which shows that tight direct product theorems imply almost-tight Chernoff-type theorems. This extends our results to Chernoff-type theorems, and gives an alternative proof to the Chernoff-type theorem of Impagliazzo et al. [IJK09] for weakly-verifiable puzzles.

Keywords: Parallel repetition, interactive argument, public-coin, Arthur-Merlin, direct product theorem.

1 Introduction

In an interactive protocol $\langle \mathsf{P}, \mathsf{V} \rangle$, the prover P wants to convince the verifier V of the validity of some statement (e.g., $x \in L$ for some language L). Two desired

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-3-642-11799-2_36

* A full version of this paper can be found on [CL09].

** Supported by US-Israel BSF grant 2006060 and NSF grant CNS-0831289.

D. Micciancio (Ed.): TCC 2010, LNCS 5978, pp. 19–36, 2010.

properties are *completeness*: for a valid statement, the honest prover can always convince the honest verifier; and *soundness*: for an invalid statement, an honest verifier, even when interacting with an adversarial prover, should accept with bounded probability, namely at most some δ, where δ is called the *soundness error* or error probability of the protocol. A protocol is called an interactive *proof* if the soundness holds against computationally unbounded provers, and an interactive *argument* if the soundness only holds against efficient provers.

When the soundness error of a protocol is too high, a natural way to decrease it is by repetition. That is, a prover and a verifier run n copies of the protocol, and the verifier decides whether to accept or not based on the outcomes of the n executions. For example, a *direct product verifier* $\mathsf{V}^{n,n}$ accepts if all constituent verifiers accept, and more generally the *threshold verifier* $\mathsf{V}^{n,k}$ accepts if at least k constituent verifiers accept. Repetitions can be either sequential or parallel. Sequential repetition decreases soundness error for all known settings, but increases the round complexity, which is usually undesirable. Parallel repetition does not increase the number of rounds and decreases soundness error for interactive proofs. However, for interactive arguments, whether parallel repetition decreases soundness error is a subtle question.

For three-message arguments, a sequence of works [BIN97, CHS05, IJK09, CLLY09, HS09] shows that parallel repetition decreases the soundness error for the threshold verifier $\mathsf{V}^{n,k}$ at the optimal, information-theoretic rate, namely, the probability that n independent Bernoulli random variables with expectation δ have sum at least k. In contrast, Bellare, Impagliazzo, and Naor [BIN97], and Pietrzak and Wikström [PW07] construct some protocols where the soundness error does not decrease *at all* under parallel repetition. Thus, parallel repetition theorems for general arguments have been ruled out. (However, Haitner [Hai09] recently showed that any interactive arguments can be slightly modified so that parallel repetition decreases the error.) On the other hand, for public-coin arguments, recent study shows that the soundness error decreases even for protocols with an arbitrary (polynomial) number of messages.

1.1 Parallel Repetition for Public-Coin Arguments

The first parallel repetition theorem for public-coin arguments is by Pass and Venkitasubramaniam [PV07] for constant-round protocols. They give an efficient transformation that converts a (cheating) parallel prover P^{n*} who interacts with a direct product verifier $\mathsf{V}^{n,n}$ with success probability δ^n to a (cheating) prover P^* who interacts with V with success probability essentially[1] δ, where the success probability refers to the probability that P^{n*} (resp., P^*) successfully convinces the verifier $\mathsf{V}^{n,n}$ (resp., V). This is essentially optimal since one can easily turn a single-copy prover strategy P^* with success probability δ to a parallel prover strategy P^{n*} with success probability δ^n by applying P^* independently to each copy. However, their analysis is only efficient for constant-round protocols.

[1] Throughout the introduction, we ignore the required negligible slackness for such reductions in the discussion.

Håstad, Pass, Pietrzak, and Wikström [HPPW08] give a more efficient reduction algorithm that allows them to prove parallel repetition theorem for public-coin arguments with an arbitrary number of rounds. They actually proves a more general threshold theorem which says that a (cheating) prover P^{n*} interacting with a threshold verifier[2] $\mathsf{V}^{n,(1-\rho)n}$ with success probability ε can be converted to a (cheating) prover P^* interacting with V with success probability $1-\rho-O(m\sqrt{\log(1/\varepsilon)/n})$, where $\rho \in [0,1)$ and m is the number of rounds. In the literature (e.g., [IJK09]), this type of theorems is often referred as *Chernoff-type* theorems. In particular, when $\rho = 0$ (i.e., the direct product case), the success probability is $1 - O(m\sqrt{\log(1/\varepsilon)/n})$, which is suboptimal in comparison to $\varepsilon^{1/n} \approx 1 - O(\log(1/\varepsilon)/n)$. Their analysis uses Raz's Sampling Lemma [Raz98] in every round, which is the reason for the factor $O(m\sqrt{\log(1/\varepsilon)/n})$ in the bound.[3] An immediate question is whether the sub-optimality is inherent for super-constant round protocols.

Recently, Wikström [Wik09] strengthened the bound of Håstad et al. [HPPW08] by generalizing Raz's Lemma and applying it only once instead in every round. He improves the analysis of [HPPW08] and shows that the construction in [HPPW08] actually achieves success probability $1 - \rho - O(\sqrt{\log(1/\varepsilon)/n})$ for Chernoff-type case, and $1 - O(\sqrt{\log(1/\varepsilon)/n})$ for direct product case. Removing the dependency on m allows him to prove a more general "concurrent repetition" theorem. The previous works give bounds on the rate at which the soundness error decreases, but it remains open whether the bounds are tight for the parallel repetition of public-coin arguments.

Our Result. In this paper, we prove a *tight* parallel repetition theorem for public-coin interactive arguments. We show that n-fold parallel repetition decreases the soundness error of public-coin arguments from δ to δ^n. We use the same reduction algorithm as [HPPW08], and the crux of our improvement is a way to avoid using Raz's Sampling Lemma.

Techniques. The constructions of P^* from P^{n*} mentioned above share the following structure. Without loss of generality, let P^{n*} be a deterministic parallel prover. The constructed prover P^* simulates internally an interaction between P^{n*} (given as a black-box) and n verifiers $\mathsf{V}_1, \ldots, \mathsf{V}_n$, where one coordinate V_i for some $i \in [n]$ chosen by P^* is played by the external verifier V. That is, throughout the interaction, P^* forwards the message that P^{n*} sends to V_i to the external V, and forwards V's message to P^{n*} as V_i's message. Since P^{n*} is deterministic, the interaction of P^{n*} and $\mathsf{V}^{n,n}$ is determined by the verifiers' messages. In each round, V selects a uniformly random message for V_i, and the task of P^* is to select good messages for the rest verifiers (denoted by V_{-i}) that maximize the probability of $\mathsf{V} = \mathsf{V}_i$ accepting at the end of interaction.

For example, the prover P^* of Pass and Venkitasubramaniam [PV07] uses *recursive sampling* to select a good coordinate $i \in [n]$ and good messages for V_{-i}

[2] Recall that a threshold verifier $\mathsf{V}^{n,(1-\rho)n}$ accepts iff at least $(1-\rho)n$ constituent verifiers accept.

[3] We elaborate more detail in the Techniques paragraph below.

such that P^{n*} could convince V_i with the highest probability among the samples he sees. However, since P^* recursively takes many samples in each round, the number of samples grows exponentially in the number of rounds. Thus, this is only efficient for constant-round protocols.

To cope with the inefficiency, the prover P^* of Håstad et al. [HPPW08] selects coordinate $i \in [n]$ uniformly at random, and uses *rejection sampling* to select good messages for V_{-i}. More precisely, let $(\boldsymbol{v}_1, \boldsymbol{p}_1, \ldots, \boldsymbol{v}_m, \boldsymbol{p}_m)$ denote the messages of $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle$, where $\boldsymbol{v}_j = (v_{j,1}, \ldots, v_{j,n})$ and $\boldsymbol{p}_j = (p_{j,1}, \ldots, p_{j,n})$ are messages of $\mathsf{V}^{n,n}$ and P^{n*} in round $j \in [m]$, respectively. In the j-th round, when P^* receives V's message, P^* considers the message as $v_{j,i}$, and repeatedly samples *random continuations* from the current partial interaction of P^{n*} and $\mathsf{V}^{n,n}$ for a polynomial number of times. That is, P^* samples messages $\boldsymbol{v}_{j,-i} = (v_{j,1}, \ldots, v_{j,i-1}, v_{j,i+1}, \ldots, v_{j,n})$, and $\boldsymbol{v}_{j+1}, \ldots, \boldsymbol{v}_m$ uniformly at random to complete the interaction. Once the continuation is *successful*, i.e., $\mathsf{V}^{n,n}$ accepts, P^* selects the $\boldsymbol{v}_{j,-i}$ of this continuation as V_{-i}'s messages, and forwards P^{n*}'s response $p_{j,i}$ to the external verifier V. If no successful continuations are found in polynomially many samples, P^* simply aborts.

To analyze the success probability, Håstad et al. [HPPW08] consider an "ideal" version of the procedure, where there is no external verifier, and the prover $\tilde{\mathsf{P}}^*$ simulates the interaction of P^{n*} and $\mathsf{V}^{n,n}$ alone by selecting each round of *all* internal verifiers' messages by rejection sampling, i.e., conditioning on a successful random continuation. Since successful continuation always exists by construction, $\tilde{\mathsf{P}}^*$ can always complete a successful interaction (i.e., $\mathsf{V}^{n,n}$ accepts) with probability 1. They then apply Raz's Lemma [Raz98] for every round to upper bound the statistical distance between the two experiments. Each application of Raz's Lemma incurs statistical distance $O(\sqrt{\log(1/\varepsilon)/n})$. Thus, the constructed prover P^* can succeed with probability at least $1 - O(m\sqrt{\log(1/\varepsilon)/n})$, where m is the number of the round. The analysis of Wikström [Wik09] follows the same structure as Håstad et al. [HPPW08]. He generalizes Raz's Lemma to a "multi-round" setting which allows him to bound the statistical distance by one application of the generalized lemma, and hence remove the dependency on m. However, to get a tight direct product theorem, we cannot afford the $O(\sqrt{\log(1/\varepsilon)/n})$ loss of applying the Raz's Lemma. It is also not clear whether the bound on the statistical distance of two experiments can be improved to $1 - \varepsilon^{1/n}$.

We instead analyze the construction directly, avoiding the use of any form of Raz's Lemma. We lower bound the success probability of the constructed prover P^* by induction. Let η_i be the success probability of P^* (i.e., the probability that P^* convinces V) when the external verifier V is embedded in the i-th coordinate, and γ the success probability of P^{n*} (i.e., the probability that P^{n*} convinces $\mathsf{V}^{n,n}$). We essentially[4] show by induction on the round $j \in [m]$ that

[4] Technically, this is for a stronger prover who can sample random continuation for unbounded number of times. For the real prover, we need to modify the inductive hypothesis to take into account the fact that the prover may fail to find a successful continuation and abort.

$\prod_i^n \eta_i \geq \gamma$, when conditioning on any partial interaction $(\boldsymbol{v}_1, \boldsymbol{p}_1, \ldots, \boldsymbol{v}_j, \boldsymbol{p}_j)$.

The base case where $j = m$ is trivial. The inductive step is proved by two applications of Hölder's Inequality. It follows that the success probability of P^* when $j = 0$ is

$$\frac{1}{n} \cdot \sum_{i=1}^{n} \eta_i \geq \left(\prod_{i=1}^{n} \eta_i \right)^{1/n} \geq \gamma^{1/n},$$

which is at least $\varepsilon^{1/n}$ by assumption.

1.2 Extension to Arguments with Simulatable Verifiers without Verdict

The results of Håstad et al. [HPPW08, HPWP10][5] extend to arguments with *simulatable* verifiers *without verdict* defined in [HPWP10]. The model generalizes both three-message arguments and public-coin arguments, and contains other natural protocols. Roughly speaking, simulatability of a verifier means that given only the prover's view of any partial interaction (which thus excludes the verifier's internal state) one can efficiently simulate the verifier in the rest of the interaction. However, since the verifier's coins are not given, one may not know the decision of the verifier in the end of the interaction. In such cases, it is referred as *simulatable verifiers without verdict*.

The argument of Håstad et al. [HPPW08] extends to this model, and gives parallel repetition theorems with the same parameters. That is, the constructed prover P^* achieves success probability $1-\rho-O(m\sqrt{\log(1/\varepsilon)/n})$ for Chernoff-type case, and $1-O(m\sqrt{\log(1/\varepsilon)/n})$ for direct product case, where m is the number of rounds. The bounds are further improved to $1-\rho-O(\sqrt{m}\sqrt{\log(1/\varepsilon)/n})$ and $1-O(\sqrt{m}\sqrt{\log(1/\varepsilon)/n})$, respectively in the new version of Håstad et al. [HPWP10], which remain dependent on m.

Our Result. We give a new reduction algorithm that converts a parallel prover P^{n*} for $\mathsf{V}^{n,n}$ with success probability ε to a prover P^* for V with success probability $\varepsilon^{2/n} \approx 1 - O(\log(1/\varepsilon)/n)$, which is almost tight.

Techniques. Recall that the prover P^* of Håstad et al. [HPPW08] selects good messages of V_{-i} by sampling and selecting a "successful" random continuation. However, since the decision of the external verifier is not known, P^* needs to select a "successful" random continuation based only on the decisions of the internal verifiers. A naive approach for P^* is to choose a continuation where all the internal verifiers accept. However, such naive P^* cannot succeed with good probability if the "success pattern" has certain *bad correlations*, as illustrated by the following example.

[5] [HPWP10] is a new version of [HPPW08] that merges the paper [Wik09] and contains additional results.

Consider a two-message protocol $\langle \mathsf{P}, \mathsf{V} \rangle$, and a (deterministic) parallel prover P^{n*} such that when interacting with $\mathsf{V}^{n,n}$, (i) P^{n*} can convince the parallel verifier $\mathsf{V}^{n,n}$ with probability ε, and (ii) for every $i \in [n]$, P^{n*} can convince all except the i-th verifier with probability $(1-\varepsilon)/n$. As there are only two messages, the naive prover P^* receives a message v from the external verifier V, and selects a response as follows. P^* randomly selects $i \in [n]$, views v as v_i, selects v_{-i} such that P^{n*} convinces V_{-i}, and forwards the corresponding p_i to V. Observe that P^* will select continuations in both cases but can only successfully convince the external verifier V in case (i). It is possible that P^* may succeed with probability $\text{(i)}/(\text{(i)}+\text{(ii)}) = \varepsilon/(\varepsilon + (1-\varepsilon)/n)$ for every external verifier V's message v. Thus, the success probability of P^* may be only $\text{(i)}/(\text{(i)}+\text{(ii)}) \approx n\varepsilon \ll \varepsilon^{1/n}$.

Two techniques have been developed to handle this bad correlation issue since the study of three-message arguments. Bellare et al. [BIN97] use the idea of *soft decision*, namely, the more the number of accepting internal verifiers, the higher the probability that the prover selects a random continuation. This approach is taken in both Impagliazzo et al. [IJK09] and Håstad et al. [HPWP10]. All these results used Raz's Sampling Lemma in their analysis.

To avoid the use of Raz's Sampling Lemma, we adopt another technique developed by Canetti et al. [CHS05] who prove a tight parallel repetition theorem for three-message arguments. The key observation is that one can exploit such a bad correlation to decrease the problem size: they present a transformation that turns a badly correlated parallel prover P^{n*} (interacting with $\mathsf{V}^{n,n}$) to a parallel prover $\mathsf{P}^{(n-1)*}$ (interacting with $\mathsf{V}^{n-1,n-1}$) that still has good success probability. To illustrate the idea using the above example, a such $\mathsf{P}^{(n-1)*}$ can simply interact with $\mathsf{V}^{n-1,n-1}$ by simulating the interaction of P^{n*} and $\mathsf{V}^{n,n}$ with the first coordinate played by an internal verifier and the rest coordinates played by $\mathsf{V}^{n-1,n-1}$. It is not hard to see that such $\mathsf{P}^{(n-1)*}$ can succeed with probability $\varepsilon + (1-\varepsilon)/n$.

It follows that one can iteratively apply the transformation until either (i) $n = 1$ or (ii) no such bad correlations exist. In case (i), we trivially obtain a single-copy prover P^* with good success probability, while in case (ii), Canetti et al. manage to show that the naive approach works for three-message arguments. We observe that this idea is generic and applicable to our setting, where our analysis technique described in Section 1.1 can be generalized to prove that, in out setting, the naive approach works in case (ii) as well, which leads to an almost tight bound.

1.3 Reducing Soundness Error for *Any* Interactive Arguments

We obtain a way to turn *any* interactive argument $\langle \mathsf{P}, \mathsf{V} \rangle$ to an interactive argument $\langle \mathsf{P}', \mathsf{V}' \rangle$ with simulatable verifier without verdict that preserves the completeness and soundness of the original protocol. As a consequence of the above section, parallel repetition decreases soundness error of the modified protocol $\langle \mathsf{P}', \mathsf{V}' \rangle$ in a nearly optimal rate.

The idea is to run the protocol $\langle \mathsf{P}, \mathsf{V} \rangle$ with all messages under the encryption of a *fully homomorphic encryption scheme* [Gen09] using verifier's public key,

which still allows the prover to simulate the original protocol with messages under encryption. Intuitively, completeness and soundness are preserved since the two parties effectively run the same protocol. Furthermore, since all the messages are encrypted under the verifier's key, they look random to the prover. Therefore, the verifier is easy to simulate – simply generate the verifier's message by encrypting some junks.

We remark that our result in this section is incomparable to the result of Haitner [Hai09] (subsequently improved by Håstad et al. [HPWP10]), who also gives a simple transformation that turns any interactive argument to one with comparable soundness where parallel repetition decreases the soundness error. Our result achieves nearly optimal rate, while the result of Haitner [Hai09] and Håstad et al. [HPWP10] has the undesirable dependency on the number of rounds m. In particular, the number of repetition is required to be at least $\Omega(m^4)$ for the soundness error to decrease. On the other hand, we use a relatively strong cryptographic assumption of fully homomorphic encryption schemes while their result holds unconditionally.

1.4 Extension to Chernoff-Type Theorems

We give a simple and generic transformation which shows that tight direct product theorems imply almost tight Chernoff-type theorems, and thus extend our results to Chernoff-type theorems. Our transformation applies to various models such as weakly-verifiable puzzles, and gives an alternative proof to the Chernoff-type theorem of Impagliazzo et al. [IJK09] as a consequence of the tight direct product theorem of Canetti et al. [CHS05].

The transformation converts a parallel prover P^{n*} for $\mathsf{V}^{n,k}$ to a parallel prover P^{t*} for $\mathsf{V}^{t,t}$ for any given $t \leq k$. The prover P^{t*} simply samples a random set of coordinate $S \subset [n]$ of size t, and interacts with $\mathsf{V}^{t,t}$ by simulating the interaction of P^{n*} and $\mathsf{V}^{n,k}$ with coordinates S played by $\mathsf{V}^{t,t}$ and the remaining coordinates played by internal verifiers. Clearly, P^{t*} convinces $\mathsf{V}^{t,t}$ if and only if P^{n*} convinces verifiers V_i's for $i \in S$ of $\mathsf{V}^{n,k}$. Let ε be the success probability of P^{n*}. It is not hard to show that P^{t*} has success probability at least $\varepsilon \cdot \binom{k}{t}/\binom{n}{t}$ by an averaging argument. Let $k = (1-\rho)n$, and suppose a tight direct theorem holds, then applying the reduction on P^{t*} with properly chosen t gives a prover P^* with success probability $(\varepsilon \cdot \binom{k}{t}/\binom{n}{t})^{1/t} \approx 1 - \rho - O(\sqrt{\log(1/\varepsilon)/n})$.[6]

For public-coin arguments, the transformation extends our direct product theorem to a Chernoff-type theorem with similar parameter to [Wik09]. For arguments with simulatable verifiers without verdict, the transformation and our improved direct product theorem yield a prover P^* with success probability $(1-\rho)^2 - O(\sqrt{\log(1/\varepsilon)/n})$. This bound is incomparable to the bound $1 - \rho - O(\sqrt{m}\sqrt{\log(1/\varepsilon)/n})$ of [HPWP10] in that our bound does not depend on m, but has a slightly worse dependency on ρ.

[6] Technically, for the reduction to be efficient, we cannot set the parameter t to be too large. Thus, the reduction P^* can only success with probability $1 - \rho - \max\{\alpha, O(\sqrt{\log(1/\varepsilon)/n})\}$ for an arbitrarily small constant α, which suffices for most applications.

As an additional contribution, we also prove that the reduction algorithm of Pass and Venkitasubramaniam [PV07] for *constant-round* public-coin arguments gives tight parallel repetition theorems for any threshold verifiers, i.e., if V has soundness error δ, then $\mathsf{V}^{n,k}$ has soundness error essentially $P(n,k,\delta)$, where $P(n,k,\delta) = \Pr[\sum_{i=1}^{n} X_i \geq k]$ with X_i's being i.i.d. binary random variables and $\Pr[X_i = 1] = \delta$.

2 Preliminary and Notation

We introduce the following notation for an interactive protocol $\langle \mathsf{P}, \mathsf{V} \rangle$. Let x denote the common input. We assume the verifier speaks first. One round contains two message exchanges – from the verifier to the prover and back. Let m denote the number of rounds. A transcript of an interaction is denoted by $(v_1, p_1, \ldots, v_m, p_m) = \langle \mathsf{P}, \mathsf{V} \rangle(x)$. When V is public-coin, verifier's messages $v_1, \ldots, v_m$ are independent uniformly random strings.

Consider parallel execution of a protocol. We use $\langle \mathsf{P}^n, \mathsf{V}^{n,k} \rangle$ to denote a n-fold parallel repetition of $\langle \mathsf{P}, \mathsf{V} \rangle$, where n copies of verifiers are denoted by $\mathsf{V}_1, \ldots, \mathsf{V}_n$, and $\mathsf{V}^{n,k}$ accepts iff at least k copies of V_i's accept. A transcript of an interaction is denoted by $(\boldsymbol{v}_1, \boldsymbol{p}_1, \ldots, \boldsymbol{v}_m, \boldsymbol{p}_m) = \langle \mathsf{P}^n, \mathsf{V}^{n,k} \rangle(x)$, where $\boldsymbol{v}_j = (v_{j,1}, \ldots, v_{j,n})$ and $\boldsymbol{p}_j = (p_{j,1}, \ldots, p_{j,n})$.

When a parallel prover P^{n*} is deterministic, the interaction $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k} \rangle$ is determined by the verifier's messages $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m)$. Thus, we can skip prover's messages and describe an interaction by $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m)$. We refers to a partial transcript as a *history* $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_j)$.

The main tool used in our analysis is Hölder's Inequality.

Lemma 1 (Hölder's Inequality [Dur04])

- *Let F, G be two non-negative functions from Ω to $\mathbb{R}$, and $a, b > 0$ satisfying $1/a + 1/b = 1$. Let q be a uniformly random variable over Ω. We have*

$$\mathop{\mathrm{E}}_{q}[F(q) \cdot G(q)] \leq \mathop{\mathrm{E}}_{q}[F(q)^a]^{1/a} \cdot \mathop{\mathrm{E}}_{q}[G(q)^b]^{1/b}.$$

- *In general, let $F_1, \ldots, F_n$ be non-negative functions from Ω to $\mathbb{R}$, and $a_1, \ldots a_n > 0$ satisfying $1/a_1 + \ldots 1/a_n = 1$. We have*

$$\mathop{\mathrm{E}}_{q}[F_1(q) \cdots F_n(q)] \leq \mathop{\mathrm{E}}_{q}[F_1(q)^{a_1}]^{1/a_1} \cdots \mathop{\mathrm{E}}_{q}[F_n(q)^{a_n}]^{1/a_n}.$$

3 Tight Direct Product Theorem for Public-Coin Arguments

In this section, we prove a tight direct product theorem for public-coin interactive arguments.

Theorem 1. *Let $\mathsf{V} \in \mathrm{PPT}$ be public-coin. There exists a prover strategy P^* such that for every common input x, every $n \in \mathbb{N}$, every $\varepsilon, \xi \in (0,1)$, and every parallel prover strategy P^{n*},*

1. $\mathsf{P}^*(x, n, \varepsilon, \xi)$ *runs in time* $\mathrm{poly}(|x|, n, \varepsilon^{-1}, \xi^{-1})$ *given oracle access to* $\mathsf{P}^{n*}(x)$.
2. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x) = 1] \geq \varepsilon \Rightarrow$

$$\Pr[\langle \mathsf{P}^*(n, \varepsilon, \xi), \mathsf{V} \rangle(x) = 1] \geq \varepsilon^{1/n} \cdot (1 - \xi).$$

We remark that the theorem also holds for interactive arguments with simulatable verifier with verdict defined in Section 4.

Without loss of generality we assume that P^{n*} is deterministic, since by sampling, we can find a fixing of the coin tosses of P^{n*} with only a small loss in the success probability. Let us first recall the common approach of such a reduction algorithm. On input x, the constructed prover P^* simulates the interaction of $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x)$ internally, where P^* simulates $n-1$ internal verifiers by himself, and lets the external verifier V play V_i for some coordinate $i \in [n]$ by forwarding the messages accordingly. Since P^{n*} is deterministic, the interaction is determined by $\mathsf{V}^{n,n}$'s message $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m)$. Let $T_i(\cdot)$ denote whether V_i accepts. That is, $T_i(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m) = 1$ iff V_i accepts in $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x)$ with history $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m)$.

This can be viewed as a *game* $\mathcal{G}(\mathsf{P}^{n*}, x)$ played between P^* and V as follows. At beginning, P^* plays a *move* $i \in [n]$. Then for each round $j \in [m]$, V plays a random move $v_{j,i}$, and P^* plays a (carefully chosen) move $\boldsymbol{v}_{j,-i} = (v_{j,1}, \ldots, v_{j,i-1}, v_{j,i+1}, \ldots, v_{j,n})$ alternately. At the end, P^* *succeeds* if $T_i(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m) = 1$. Note that a node of the game tree is of the form either $(i; \boldsymbol{v}_1, \ldots, \boldsymbol{v}_j)$, in which case it is V's turn to move, or of the form $(i; \boldsymbol{v}_1, \ldots, \boldsymbol{v}_{j-1}, v_{j,i})$, in which case it if P^*'s turn to move. Phrased in this way, the task is to design a strategy for P^* such that if $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x)$ accepts with probability at least ε, then P^* can succeed with probability close to $\varepsilon^{1/n}$ in game $\mathcal{G}(\mathsf{P}^{n*}, x)$. We present the "rejection sampling" reduction algorithm of Hastad et al. [HPPW08] as a strategy of P^* in this game:

Definition 1 (Strategy P^*_{rej}). *We define strategy* P^*_{rej} *as follows. Let* P^{n*} *be a deterministic parallel prover,* x *a common input, and* $\mathcal{G}(\mathsf{P}^{n*}, x)$ *the corresponding game defined as above.*

- *In the first* P^**-move,* P^*_{rej} *selects a coordinate* $i \in_R [n]$ *uniformly at random.*
- *On* P^**-move node* $u = (i; \boldsymbol{v}_1, \ldots, \boldsymbol{v}_{j-1}, v_{j,i})$, P^*_{rej} *simulates a* random continuation *of* $\mathcal{G}(\mathsf{P}^{n*}, x)$ *(i.e., the interaction of* $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x)$*) at most* $M \stackrel{\text{def}}{=} O(mn/\varepsilon\xi)$ *times. That is,* P^*_{rej} *simulates the game from* u *with both parties playing random moves* $\boldsymbol{v}_{j,-i}, \ldots, \boldsymbol{v}_{m,i}, \boldsymbol{v}_{m,-i}$*. A continuation is* successful *if all verifiers accept, i.e.,* $T_\ell(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m) = 1$ *for all* $\ell \in [n]$*. The first time a successful continuation is found,* P^*_{rej} *plays the corresponding move* $\boldsymbol{v}_{j,-i}$*. If no successful continuations are found,* P^*_{rej} *aborts.*

Note that if P^*_{rej} *does not abort,* P^*_{rej} *plays move* $\boldsymbol{v}_{j,-i}$ *with the probability proportional to the conditional success probability of* P^{n*} *given on the history* $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_j)$.

Clearly, strategy P^*_{rej} can be implemented in time $\mathrm{poly}(|x|, n, \varepsilon^{-1}, \xi^{-1})$. We next analyze the success probability of P^*_{rej} by induction on the round $j \in [m]$. For

the sake of clarity, below we first present the analysis of an ideal version P^*_{ideal} of P^*_{rej}, where P^*_{ideal} can simulate random continuations for unbounded number of times. The analysis of P^*_{rej} is presented in the full version of this paper [CL09].

3.1 Analysis of $\mathbf{P}^*_{ideal}$

In this subsection, we analyze the success probability of an ideal version P^*_{ideal} of strategy P^*_{rej}, which is the same as P^*_{rej} except that P^*_{ideal} can simulate the random continuations an unbounded number of times. Thus, P^*_{ideal} will never abort whenever there is a successful continuation from the current P^*-move node. We will show that if $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x) = 1] \geq \varepsilon$, then P^*_{ideal} can succeed with probability at least $\varepsilon^{1/n}$ in game $\mathcal{G}(\mathsf{P}^{n*}, x)$.

We first introduce the following notation to express the success probability of P^*_{ideal}. We define $\gamma(\bar{h}) \stackrel{\text{def}}{=} \Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^n\rangle(x) = 1|\bar{h}]$, where $\bar{h}$ is a history of the form either $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_j)$ or $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{j-1}, v_{j,i})$. That is, $\gamma(\bar{h})$ is the accepting probability of $\langle \mathsf{P}^{n*}, \mathsf{V}^n\rangle$ conditioning on the history $\bar{h}$. Note that $\gamma = \Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^n\rangle(x) = 1] \geq \varepsilon$ by assumption. Next, for every $i \in [n]$, we define $\eta_i(\bar{h}) \stackrel{\text{def}}{=} \Pr[\mathsf{P}^*_{ideal} \text{ succeeds } |u = (i; \bar{h})]$ to be the success probability of P^*_{ideal} conditioning on node $u = (i; \bar{h})$ of the game tree. Note that the success probability of P^*_{ideal} is $(1/n) \cdot \sum_{i=1}^{n} \eta_i$.

Claim. For every $i \in [n]$ and full history $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m)$, we have $\eta_i(\bar{h}) = T_i(\bar{h})$. For every $i \in [n]$, $j \in [m]$, and history $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{j-1})$, we have[7]

$$\eta_i(\bar{h}) = \mathop{\mathrm{E}}_{\boldsymbol{v}_j}\left[\frac{\gamma(\bar{h}, \boldsymbol{v}_j) \cdot \eta_i(\bar{h}, \boldsymbol{v}_j)}{\gamma(\bar{h}, v_{j,i})}\right].$$

Proof. The first part follows by definition. For the second part, recall that V plays the random strategy and P^*_{ideal} plays the rejection sampling strategy. V plays each $v_{j,i}$ with probability $\Pr[v_{j,i}]$, which corresponds to the expectation operator over $v_{j,i}$. P^*_{ideal} plays each $\boldsymbol{v}_{j,-i}$ with probability $\Pr[\boldsymbol{v}_{j,-i}] \cdot (\gamma(\bar{h}, \boldsymbol{v}_j)/\gamma(\bar{h}, v_{j,i}))$, which corresponds to the expectation operator over $\boldsymbol{v}_{j,-i}$ with factor $\gamma(\bar{h}, \boldsymbol{v}_j)/\gamma(\bar{h}, v_{j,i})$ in the expectation.

We now prove that the success probability of P^*_{ideal} is at least $\varepsilon^{1/n}$ by induction. In fact, we induct on a slightly stronger inductive hypothesis: for every $j \in \{0, \ldots, m\}$ and history $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_j)$, $\prod_{i=1}^{n} \eta_i(\bar{h}) \geq \gamma(\bar{h})$.

The base case $j = m$ is trivial. For every full history $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m)$, $\gamma(\bar{h}) = 1$ iff $\eta_i(\bar{h}) = T_i(\bar{h}) = 1$ for every $i \in [n]$. Assuming that the inductive hypothesis holds for j and every $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_j)$, we want to prove the inductive hypothesis for $j - 1$ and every $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{j-1})$. More precisely, for every $\bar{h} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{j-1})$, we want to show that

[7] We use the convention that if $\gamma(\bar{h}, v_{j,i}) = 0$ (which implies $\gamma(\bar{h}, \boldsymbol{v}_j) = 0$), then the ratio is 0.

$$\prod_{i=1}^{n} \eta_i(\bar{h}) = \prod_{i=1}^{n} \mathop{\mathrm{E}}_{\boldsymbol{v}_j} \left[\frac{\gamma(\bar{h}, \boldsymbol{v}_j) \cdot \eta_i(\bar{h}, \boldsymbol{v}_j)}{\gamma(\bar{h}, v_{j,i})} \right] \geq \gamma(\bar{h}),$$

provided that for every $\boldsymbol{v}_j$, $\prod_{i=1}^{n} \eta_i(\bar{h}, \boldsymbol{v}_j) \geq \gamma(\bar{h}, \boldsymbol{v}_j)$. For notational simplicity, we abstract the above statement as the following lemma.

Lemma 2. *Let $\gamma, \eta_1, \ldots, \eta_n : \Omega^n \to [0,1]$ be $[0,1]$-valued functions over a product space Ω^n such that $\prod_i \eta_i(\boldsymbol{q}) \geq \gamma(\boldsymbol{q})$ for every $\boldsymbol{q} = (q_1, \ldots, q_n) \in \Omega^n$. Let $\gamma = \mathrm{E}_{\boldsymbol{q}}[\gamma(\boldsymbol{q})]$. For every $i \in [n]$, let*

$$\gamma(q_i) = \mathop{\mathrm{E}}_{\boldsymbol{q}_{-i}} [\gamma(\boldsymbol{q})] \quad \text{and} \quad \eta_i = \mathop{\mathrm{E}}_{\boldsymbol{q}} \left[\frac{\gamma(\boldsymbol{q}) \cdot \eta_i(\boldsymbol{q})}{\gamma(q_i)} \right],$$

where the above expectation is over uniform distribution over Ω^n. We have

$$\prod_{i=1}^{n} \eta_i = \prod_{i=1}^{n} \mathop{\mathrm{E}}_{\boldsymbol{q}} \left[\left(\frac{\gamma(\boldsymbol{q}) \cdot \eta_i(\boldsymbol{q})}{\gamma(q_i)} \right) \right] \geq \gamma.$$

Proof. The trick is to apply Hölder's Inequality to "swap the operators". We present the whole computation first, and then explain how Hölder's Inequality is applied.

$$\begin{aligned}
&\prod_{i=1}^{n} \mathop{\mathrm{E}}_{\boldsymbol{q}} \left[\left(\frac{\gamma(\boldsymbol{q}) \cdot \eta_i(\boldsymbol{q})}{\gamma(q_i)} \right) \right] \\
&\geq \mathop{\mathrm{E}}_{\boldsymbol{q}} \left[\left(\frac{\gamma(\boldsymbol{q})^n \cdot \prod_{i=1}^{n} \eta_i(\boldsymbol{q})}{\prod_{i=1}^{n} \gamma(q_i)} \right)^{1/n} \right]^n \qquad \text{(by Hölder's Inequality)} \\
&\geq \mathop{\mathrm{E}}_{\boldsymbol{q}} \left[\left(\frac{\gamma(\boldsymbol{q})^{n+1}}{\prod_{i=1}^{n} \gamma(q_i)} \right)^{1/n} \right]^n \qquad \text{(by inductive hypothesis)} \\
&\geq \left[\left(\frac{\mathrm{E}_{\boldsymbol{q}}[\gamma(\boldsymbol{q})]^{n+1}}{\mathrm{E}_{\boldsymbol{q}}[\prod_{i=1}^{n} \gamma(q_i)]} \right)^{1/n} \right]^n \qquad \text{(by Hölder's Inequality)} \\
&= (\gamma^{n+1}/\gamma^n) = \gamma.
\end{aligned}$$

We now explain the application of Hölder's Inequalities.

- The first inequality uses $\mathrm{E}[X_1^n]^{1/n} \cdot \cdots \cdot \mathrm{E}[X_n^n]^{1/n} \geq \mathrm{E}[X_1 \cdot \cdots \cdot X_n]$ with

$$X_i = \left(\frac{\gamma(\boldsymbol{q}) \cdot \eta_i(\boldsymbol{q})}{\gamma(q_i)} \right)^{1/n}.$$

- The third inequality uses $\mathrm{E}\left[B^{n+1}\right]^{1/(n+1)} \cdot \mathrm{E}\left[(A/B)^{(n+1)/n}\right]^{n/(n+1)} \geq \mathrm{E}[A]$, or equivalently,

$$\mathrm{E}\left[\left(\frac{A^{n+1}}{B^{n+1}} \right)^{1/n} \right] \geq \left(\frac{\mathrm{E}[A]^{n+1}}{\mathrm{E}[B^{n+1}]} \right)^{1/n}$$

with

$$\begin{cases} A = \gamma(\boldsymbol{q}), \\ B^{n+1} = \prod_{i=1}^{n} \gamma(q_i). \end{cases}$$

Remark 1. One might worry about the legitimacy of the manipulation when the denominators are zeros. One way to justify it is by adding some μ in the denominators before the manipulation. Formally, we have

$$\prod_{i=1}^{n} \mathop{\mathrm{E}}_{\boldsymbol{q}} \left[\left(\frac{\gamma(\boldsymbol{q}) \cdot \eta_i(\boldsymbol{q})}{\gamma(q_i)} \right) \right] \geq \prod_{i=1}^{n} \mathop{\mathrm{E}}_{\boldsymbol{q}} \left[\left(\frac{\gamma(\boldsymbol{q}) \cdot \eta_i(\boldsymbol{q})}{\gamma(q_i) + \mu} \right) \right] \geq \cdots \geq (\gamma^{n+1}/(\gamma + \mu)^n),$$

which is valid for arbitrary $\mu > 0$. Taking $\mu \to 0$, we obtain the desired result.

Applying the above lemma directly completes the proof of the induction. It follows that the success probability of P^*_{ideal} is

$$\frac{1}{n} \cdot \sum_{i=1}^{n} \eta_i \geq \left(\prod_{i=1}^{n} \eta_i \right)^{1/n} \geq \gamma^{1/n} \geq \varepsilon^{1/n}.$$

The next step is to analyze P^*_{rej} in a similar way as above. The challenge is that P^*_{rej} may abort due to the failure of finding a successful continuation in M trials, which makes the success probability a more complicated formula. Details of the analysis of P^*_{rej} can be found in the full version of this paper [CL09].

4 Arguments with Simulatable Verifier without Verdict

In this section, we present a new reduction algorithm that extends our results to interactive arguments with simulatable verifiers defined by Håstad et al. [HPWP10]. Roughly speaking, a verifier is simulatable if given only the prover's view of any partial interaction (which thus excludes the verifier's internal state), one can efficiently simulate verifier in the rest of the interaction. In terms of the terminology in [HPWP10], our results holds for arguments with "1-simulatable verifiers without verdict," which we refer to as just simulatable verifiers below for simplicity. For the sake of completeness, we repeat their definition in this special case. For a more general definition of simulatability, we refer the reader to [HPWP10].

The definition requires the following notation. Recall that we use p_j and v_j to denote the prover and verifier's j-th messages, respectively. We let s_j and t_j be the states of the prover and verifier after computing the j-th messages, respectively. We think of the verifier as using independent random tape R_j for computing j-th message. Namely, V computes message v_j from its previous state t_{j-1}, prover's message p_j, and fresh randomness R_j. Note that the verifier's state t_j implicitly contains the content of the random tapes $r_1, \ldots, r_j$ (generated in the previous rounds) of V. For convenience, we use $p_{[j]}$ to denote $p_1, \ldots, p_j$, and the same rule applies to other variables.

Definition 2 (Simulatable Verifier [HPWP10]). *A verifier* V *is said to be* simulatable without verdict*, or just* simulatable*, if for every* PPT *prover strategy* P^* *there exists a* PPT *simulator* S *such that for every partial interaction* $(s_{[j]}, t_{[j]}, x, p_{[j]}, v_{[j]})$*, the distribution of* P^**'s view of an interaction with* V *(not including the decision bit of* V*), starting from states* s_j *and* t_j *and message* p_j*, is* computationally indistinguishable *to the distribution of* P^**'s view of an interaction with* S *starting from states* s_j *and* $[s_{[j]}, x, p_{[j]}, v_{[j]}]$ *and message* p_j*. When the decision bit of* V *is included in the consideration, we say that* V *is* simulatable with verdict*.*

Remark 2. In the above definition, we only require the distributions to be computational indistinguishable, as opposed to the statistical closeness defined in [HPWP10]. Håstad et al. requires statistical closeness since they need to handle a general notion of "δ-simulatability." On the other hand, for the case of 1-simulatability, it can be shown (e.g., in the old version of Håstad el al. [HPPW08]) that the requirement can be relaxed to computational indistinguishability. The relaxation to computational indistinguishability is essential to our application of fully-homomorphic encryption in Section 5.

Remark 3. Another deviation from [HPWP10] is that, in the above definition, our simulator S interacts with P^*, as opposed to generate the view by himself in [HPWP10]. This difference is not essential. We adopt to the above definition since it makes the simulation of the random continuation described below more intuitive.

We observe that for arguments with simulatable verifier, in the corresponding game $\mathcal{G}(\mathsf{P}^{n*}, x)$, P^* can still simulate a random continuation from any P^*-move node u. Each internal verifier's next message is easy to generate since the message depends only on the verifier's state, the prover's message, and fresh randomness. For the external verifier V, although P^* does not know V's state, P^* can invoke the simulator to generate the verifier's message. However, P^* is not able to know the decision of the external verifier. Thus, P^* needs to select a "successful" random continuation based only on the decisions of the internal verifiers. As illustrated in the example in Section 1.2, there is an issue of "bad correlations." We resolve this issue in the spirit of Canetti et al. [CHS05], where we iteratively exploit bad correlations to decrease the problem size in a preprocssessing stage, and use a modified rejection sampling strategy when no such bad correlations exist. Our reduction turns a parallel prover P^{n*} for $\mathsf{V}^{n,n}$ with success probability $\delta^n \stackrel{\text{def}}{=} \varepsilon$ to a prover P^* for a single simulatable verifier V with success probability $\delta^2 = \varepsilon^{2/n} \approx 1 - O(\log(1/\varepsilon)/n)$.[8] Formally, we obtain the following theorem.

Theorem 2. *Let* $\mathsf{V} \in$ PPT *be simulatable without verdict. There exists a prover strategy* P^* *such that for every common input* x*, every* $n \in \mathbb{N}$*, every* $\varepsilon, \xi \in (0,1)$*, and every parallel prover strategy* P^{n*}*,*

[8] It is more convenient to present our proof using parameter δ^n instead of ε in this section.

1. $\mathsf{P}^*(x, n, \varepsilon, \xi)$ *runs in time* $\mathrm{poly}(|x|, n, \varepsilon^{-1}, \xi^{-1})$ *given oracle access to* $\mathsf{P}^{n*}(x)$.
2. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x) = 1] \geq \varepsilon \Rightarrow$

$$\Pr[\langle \mathsf{P}^*, \mathsf{V} \rangle(x) = 1] \geq \varepsilon^{2/n} \cdot (1 - \xi).$$

Detailed description of the reduction algorithms and analysis can be found in the full version of this paper [CL09].

5 Reducing Soundness Error for Any Arguments

(The ideas in this section were obtained in discussions with Boaz Barak, Yael Tauman Kalai, and Salil Vadhan.)

In this section, we present a way to turn *any* interactive argument $\langle \mathsf{P}, \mathsf{V} \rangle$ to an interactive argument $\langle \mathsf{P}', \mathsf{V}' \rangle$ with simulatable verifier that preserves the completeness and soundness of the original protocol. It follows that parallel repetition reduces soundness error of the modified protocol $\langle \mathsf{P}', \mathsf{V}' \rangle$ in a nearly optimal rate by Theorem 2.

Recall that the idea is to run the protocol $\langle \mathsf{P}, \mathsf{V} \rangle$ with all messages under the encryption of a *fully homomorphic encryption scheme*. Roughly speaking, a fully homomorphic encryption scheme is a public key encryption scheme with the additional property that given a public key pk, and an encryption $\mathsf{Enc}_{\mathsf{pk}}(m)$, one can homomorphically evaluate any function f (described by a poly-size circuit C) on the underlying message to obtain an encrypted function value $\mathsf{Enc}_{\mathsf{pk}}(f(m))$ without knowing the message m. That is, in addition to the standard functions $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ in public key encryption schemes, a fully homomorphic encryption scheme has an additional efficient function Eval that on inputs a public key pk, a description of a poly-size circuit $C(\cdot)$, and a cipher text c that is a valid encryption of m, outputs a cipher text c' which is a valid encryption of $C(m)$.

Recently in a breakthrough, Gentry [Gen09] showed the first construction of a fully homomorphic encryption scheme under reasonable hardness assumptions on ideal lattice problems and sparse subset sum problems. We refer the reader to [Gen09] for the formal definitions and constructions.

Let $\langle \mathsf{P}, \mathsf{V} \rangle$ be any interactive argument. Recall our notation, P and V receive some common input x and alternately send to each other messages denoted as $(v_1, p_1, v_2, p_2, \ldots, v_m, p_m)$ where m is the number of the rounds. We define a modified protocol $\langle \mathsf{P}', \mathsf{V}' \rangle$ that executes the protocol $\langle \mathsf{P}, \mathsf{V} \rangle$ under a fully homomorphic encryption of the verifier's key as follows. For simplicity, we assume that V always makes his decision in the end of the protocol, and all messages of $\langle \mathsf{P}, \mathsf{V} \rangle$ have some fixed length. We also assume that the encryption scheme has perfect correctness and the decryption algorithm Dec always outputs some messages (perhaps junks).

– In the first round, the verifier V' generates $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}()$,[9] prepares V's first message v_1, and sends the public key pk and the encrypted message $v_1' = \mathsf{Enc}_{\mathsf{pk}}(v_1)$ to P'.

[9] For simplicity, we omit the security parameter throughout this section.

- The prover P' on the received message v'_1, homomorphically computes p'_1, a valid encryption of the first message p_1 of P. Namely, let $C_1(x, v_1)$ be the next-message function of P. The prover P' uses Eval to compute $p'_1 = \mathsf{Eval}_{\mathsf{pk}}(v'_1, C_1(x, \cdot))$.
- In general, in the ℓ-th round, the verifier V' receives message $p'_{\ell-1}$. V' first decrypts the message $p'_{\ell-1}$ to obtain $p_{\ell-1} = \mathsf{Dec}_{\mathsf{sk}}(p'_{\ell-1})$. V' simulates V to generate the next message v_ℓ, and sends the encrypted message $v'_\ell = \mathsf{Enc}_{\mathsf{pk}}(v_\ell)$ to P'.
- The prover P' on the received message v'_ℓ, homomorphically computes p'_ℓ, a valid encryption of the first message p_ℓ of P. Namely, let $C_\ell(x, v_{[\ell]}, p_{[\ell-1]})$ be the next-message function of P. The prover P' uses Eval to compute $p'_\ell = \mathsf{Eval}_{\mathsf{pk}}((v'_{[\ell]}, p'_{[\ell-1]}), C_\ell(x, \cdot))$.
- At the end, V' decrypts the last message p'_m. V' accepts iff V accepts.

We first observe that $\langle \mathsf{P}', \mathsf{V}' \rangle$ has exactly the same completeness and soundness as $\langle \mathsf{P}, \mathsf{V} \rangle$ suppose the homomorphic encryption scheme has perfect correctness. The completeness is trivially the same, since $\langle \mathsf{P}', \mathsf{V}' \rangle$ simply simulates $\langle \mathsf{P}, \mathsf{V} \rangle$ under a fully homomorphic encryption. For the soundness, note that for every (cheating) prover strategy P'^* for $\langle \mathsf{P}', \mathsf{V}' \rangle$, we can construct a (cheating) prover strategy P^* that interacts with V by simulating the interaction of P'^* and V' as follows. P^* first generates $(\mathsf{pk}, \mathsf{sk})$ by himself and forwards pk to P'^*. P^* then simulates the interaction of P'^* and V' by (i) encrypting the messages of V and forwarding them to P'^*, and (ii) decrypting the messages of P'^* and forwarding them to V. It follows that P^* can convince V with the same probability as P'^* convincing V'. Similarly, for every P^*, there is a P'^* that applies the same strategy as P^* (homomorphically) and convinces V' with the same probability as P^* convincing V.

It remains to show that V' is simulatable. To argue this, we need to specify the random tape used by V' in each round, since this affects the states t_ℓ's of the verifier. For convenience, we define $m + 1$ random tapes $R_0, R_1, \ldots, R_m$ for V', where both R_0 and R_1 are generated in the first round. We let R_0 be the random tape that contains all the randomness used in V. For $\ell \in [m]$, we let R_ℓ be the randomness that V' uses to encrypt the ℓ-th round message. Note that defined in this way, given the state $t'_{\ell-1}$ of V' and prover P''s message p'_ℓ, the underlying verifier V's message v_i is *deterministic*, and the randomness of V''s message v'_i comes only from the encryption. Now it is trivial to simulate V'. A simulator S simply ignores the prover's message, and sends a fresh encryption of junks in each round. By the semantic security of the encryption scheme, the prover's view when interacting with V' is computationally indistinguishable from that when interacting with S.

We summarize the above discussion in the following theorem.

Theorem 3. *Let $\langle \mathsf{P}, \mathsf{V} \rangle$ be any interactive argument with soundness error δ. Suppose there exists a fully homomorphic encryption scheme with perfect correctness, then the modified interactive argument $\langle \mathsf{P}', \mathsf{V}' \rangle$ defined above satisfies the following properties.*

– $\langle \mathsf{P}', \mathsf{V}' \rangle$ *has exactly the same completeness and soundness as* $\langle \mathsf{P}, \mathsf{V} \rangle$.
– V' *is simulatable without verdict, and thus* n*-fold parallel repetition reduces soundness error from* δ *to* $\delta^{n/2} + \mathsf{ngl}$.

6 Extension to Chernoff-Type Theorems

In this section, we present a generic transformation that converts a parallel prover P^{n*} that has good success probability against a threshold verifier to a parallel prover P^{t*} that has good success probability against a direct product verifier for some $t \leq n$. The transformation can be used to show that tight direct product theorems implies Chernoff-type theorems. For example, using our transformation with the direct product theorem of Canetti et al. [CHS05] yields an alternative proof of the Chernoff-type theorem of Impagliazzo et al. [IJK09] for weakly-verifiable puzzles. The transformation also extends our direct product theorems to Chernoff-type theorems.

The transformation is defined as follows. P^{t*} first selects a set $S \subset [n]$ of size t uniformly at random, and then interacts with $\mathsf{V}^{t,t}$ by simulating the interaction of $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k} \rangle$ with $\mathsf{V}^{t,t}$ playing the coordinates of $\mathsf{V}^{n,k}$ in S and the remaining $n-t$ coordinates played by internal verifiers. The following simple lemma easily follows by the definition.

Lemma 3. *Let* $\langle \mathsf{P}, \mathsf{V} \rangle$ *be an interactive protocol, and* $t, k, n \in \mathbb{N}$ *such that* $1 \leq t \leq k \leq n$. *Let* P^{n*} *be a parallel prover strategy, and* P^{t*} *the induced parallel prover strategy defined as above. For every common input* x, *we have*

$$\Pr[\langle \mathsf{P}^{t*}, \mathsf{V}^{t,t} \rangle(x) = 1] \geq \Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k} \rangle(x) = 1] \cdot \frac{\binom{k}{t}}{\binom{n}{t}}.$$

When V is public-coin, the above lemma and Theorem 1 implies that for every parallel prover P^{n*}, every $t \leq k$ and $\xi \in (0,1)$, there exists a prover P^* such that for every x with $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k} \rangle(x) = 1] \geq \varepsilon$, we have $\Pr[\langle \mathsf{P}^*, \mathsf{V} \rangle(x) = 1] \geq \left(\varepsilon \cdot \binom{k}{t} / \binom{n}{t}\right)^{1/t} \cdot (1 - \xi)$. However, P^* runs in time $\mathrm{poly}(|x|, n, \binom{n}{t}/\binom{k}{t}, \varepsilon^{-1}, \xi^{-1})$, which may not be efficient[10] for large t. Nevertheless, we can obtain the following Chernoff-type theorem by setting the parameters properly. We state the theorem in a similar form to [HPPW08] and [Wik09].

Theorem 4. *Let* $\alpha, \rho \in (0,1)$ *be any* constants *such that* $\alpha + \rho < 1$. *Let* $\mathsf{V} \in \mathrm{PPT}$ *be public-coin. There exists a prover strategy* P^* *such that for every common input* x, *every* $n \in \mathbb{N}$, *every* $\varepsilon, \xi \in (0,1)$ *with* $n \geq 4\log(1/\varepsilon)/\alpha^2$, *and every parallel prover strategy* P^{n*},

1. $\mathsf{P}^*(x, n, \varepsilon, \xi)$ *runs in time* $\mathrm{poly}(|x|, n, \varepsilon^{-1}, \xi^{-1})$ *given oracle access to* $\mathsf{P}^{n*}(x)$.
2. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,(1-\rho)n} \rangle(x) = 1] \geq \varepsilon \Rightarrow$

$$\Pr[\langle \mathsf{P}^*(n, \varepsilon, \xi), \mathsf{V} \rangle(x) = 1] \geq 1 - \rho - \alpha.$$

[10] Here, by efficient we mean the running time is polynomial in $|x|, n, \varepsilon^{-1}, \xi^{-1}$.

In comparison, the simple reduction and tight direct product theorem yields a Chernoff-type theorem with a slightly restricted parameter range where α and ρ are constants. Nevertheless, it suffices for conceivable applications and achieves almost tight bound $1 - \rho - 2\sqrt{\log(1/\varepsilon)/n}$ in this regime.

Similarly, when V is simulatable, we can extend Theorem 2 to the following Chernoff-type theorem.

Theorem 5. *Let* $\alpha, \rho \in (0,1)$ *be any* constants *such that* $\alpha+\rho < 1$. *Let* $\mathsf{V} \in$ PPT *be exteandable and simulatable. There exists a prover strategy* P^* *such that for every common input* x, *every* $n \in \mathbb{N}$, *every* $\varepsilon, \xi \in (0,1)$ *with* $n \geq 16\log(1/\varepsilon)/\alpha^2$, *and every parallel prover strategy* P^{n*},

1. $\mathsf{P}^*(x, n, \varepsilon, \xi)$ *runs in time* $\mathrm{poly}(|x|, n, \varepsilon^{-1}, \xi^{-1})$ *given oracle access to* $\mathsf{P}^{n*}(x)$.
2. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,(1-\rho)n}\rangle(x) = 1] \geq \varepsilon \Rightarrow$

$$\Pr[\langle \mathsf{P}^*(n, \varepsilon, \xi), \mathsf{V}\rangle(x) = 1] \geq (1-\rho)^2 - \alpha.$$

Detailed proofs of Lemma 3, Theorem 4, 5 can be found in the full version of this paper [CL09].

7 Constant-Round AM Arguments Systems

In this section, we prove a tight parallel repetition theorem for threshold verifiers $\mathsf{V}^{n,k}$ for *constant-round* public-coin arguments, which generalizes the direct product theorem of Pass and Venkitasubramaniam [PV07]. We state the theorem and further details of the proofs can be found in the full version of this paper [CL09].

Theorem 6. *Let* $m \in \mathbb{N}$ *be an arbitrary constant, and* $\mathsf{V} \in$ PPT *be* m*-round and public coin. There exists a prover strategy* P^* *such that for every common input* x, *every* $n, k \in \mathbb{N}$ *with* $k \in [n]$, *every* $\delta, \xi \in (0,1)$, *and every parallel prover strategy* P^{n*},

1. $\mathsf{P}^*(x, n, k, \delta, \xi)$ *runs in time* $\mathrm{poly}(|x|, n, \delta^{-m}, P(n,k,\delta)^{-m}, \xi^{-m})$ *given oracle access to* $\mathsf{P}^{n*}(x)$.
2. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \geq P(n,k,\delta) \Rightarrow$

$$\Pr[\langle \mathsf{P}^*(n, k, \delta, \xi), \mathsf{V}\rangle(x) = 1] \geq \delta \cdot (1-\xi).$$

Acknowledgments

We thank Boaz Barak, Yael Tauman Kalai, and Salil Vadhan for the useful discussion that leads to the results in Section 5. We also thank Salil Vadhan for very helpful discussions throughout this work.

References

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

[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)

[CL09] Chung, K.-M., Liu, F.-H.: Parallel repetition theorems for interactive arguments. Electronic Colloquium on Computational Complexity (ECCC) (109) (2009), http://eccc.uni-trier.de/report/2009/109/

[CLLY09] Chung, K.-M., Liu, F.-H., Lu, C.-J., Yang, B.-Y.: Efficient string-commitment from weak bit-commitment and full-spectrum theorem for puzzles (2009) (unpublished manuscript)

[Dur04] Durrett, R.: Probability: Theorey and Examples, 3rd edn. Duxbury (2004)

[Gen09] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC, pp. 169–178 (2009)

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

[HPPW08] Håstad, J., Pass, R., Pietrzak, K., Wikström, D.: 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)

[HS09] Holenstein, T., Schoenebeck, G.: General hardness amplification of predicates and puzzles (2009) (unpublished manuscript)

[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)

[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)

[Wik09] Wikström, D.: An efficient concurrent repetition theorem. Cryptology ePrint Archive, Report 2009/347 (2009)