Electronic Colloquium on Computational Complexity, Report No. 109 (2009)

Tight Parallel Repetition Theorems for Public-coin Arguments

Kai-Min Chung*(kmchung@fas.harvard.edu), Feng-Hao Liu (fenghao@cs.brown.edu)

November 3, 2009

Abstract

Following Hastad et al. [HPPW08], we study parallel repetition theorems for public-coin interactive arguments and their generalizations. We obtain the following results:

1. We show that the reduction of Hastad et al. [HPPW08] actually gives a *tight* direct product theorem for public-coin interactive arguments. That is, n-fold parallel repetition reduces 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 to the previous results.
2. We give a new reduction to strengthen the direct product theorem of Hastad et al. for arguments with extendable and simulatable verifiers. We show that n-fold parallel repetition reduces 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 give a simple and generic reduction which shows that tight direct product theorems imply almost-tight Chernoff-type theorems. The reduction extends our results to Chernoff-type theorems, and gives an alternative proof to the Chernoff-type theorem of Impagliazzo et al. [IJK07] for weakly-verifiable puzzles.
4. As an additional contribution, we observe that the reduction of Pass and Venkitasubramaniam [PV07] for constant-round public-coin arguments gives tight parallel repetition theorems for threshold verifiers, who accept when more than a certain number of repetition accepts.

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

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

ISSN 1433-8092

1 Introduction

In an interactive protocol, a prover P wants to convince the verifier V the validity of some statement x. Two desired 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 reduce it is by repetition. That is, a prover and a verifier run n copies of a 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 outcomes are acceptance, and more generally a *threshold verifier* $\mathsf{V}^{n,k}$ accepts if at least k outcomes are acceptance. Repetitions can be either sequential or parallel. Sequential repetition reduces 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 reduces soundness error for interactive proofs. However, for interactive arguments, whether parallel repetition reduces soundness error is a subtle question.

For three-message arguments, a sequence of works [BIN97, CHS05, IJK07, CLLY09, HS09] shows that parallel repetition reduces the soundness error for threshold verifiers $\mathsf{V}^{n,k}$ at the optimal, information-theoretic rate, namely, the probability that n independent Bernoulli random variables with expectation δ have sum at least k. However, Bellare, Impagliazzo [BIN97], and Naor, and Pietrzak and Wikström [PW07] construct four-message protocols where the soundness error does not reduce *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 reduces the error; see Section 1.4.) On the other hand, for the class of public-coin arguments, recent study shows that the soundness error is reduced even for protocols with an arbitrary (polynomial) number of messages. In this paper, we continue the study of parallel repetition theorems for public-coin arguments.

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 reduction that converts a parallel prover (for a direct product verifier) with success probability δ^n to a single-copy prover with success probability essentially[1] δ, which is optimal. However, the reduction is only efficient for constant-round protocols.

Håstad, Pass, Pietrzak, and Wikström [HPPW08] give a more efficient reduction that proves parallel repetition theorem for public-coin arguments with an arbitrary number of rounds. The reduction also gives a more general Chernoff-type theorem where a parallel prover P^{n*} for $\mathsf{V}^{n,(1-\gamma)n}$ with success probability ε is converted to a prover P^* for V with success probability $1-\gamma-O(m\sqrt{\log(1/\varepsilon)/n})$, where $\gamma \in [0,1)$ and m is the number of rounds. In particular, when $\gamma = 0$ (i.e., the direct product case), the success probability is $1-O(m\sqrt{\log(1/\varepsilon)/n})$, which is suboptimal in compare to $\varepsilon^{1/n} \approx 1-O(\log(1/\varepsilon)/n)$. Their analysis uses Raz's Sampling Lemma [Raz98] in every

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

round, which is the reason for the factor $O(m\sqrt{\log(1/\varepsilon)/n})$ in the bound.[2] 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 of every round. He shows that the reduction of [HPPW08] actually achieves success probability $1-\gamma-O(\sqrt{\log(1/\varepsilon)/n})$ for Cheonoff-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. However, it remained open whether n-fold parallel repetition for public-coin arguments with a super-constant number of rounds reduces soundness error in an optimal, information-theoretic rate from δ to δ^n.

Our Result. In this paper, we answer the question affirmatively. We show that the reduction of Håstad et al. [HPPW08] reduces soundness error from δ to δ^n for direct product case for public-coin arguments with an arbitrary number of rounds. The crux of our improvement is a way to avoid using Raz's Sampling Lemma.

Techniques. The reductions of efficient parallel repetition theorems mentioned above share the following structure. Let P^{n*} be a deterministic parallel prover. The reduced 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 (wlog), 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 uses *recursive sampling* to select a good coordinate $i\in[n]$ and good messages for V_{-i} such that P^{n*} could convince V_i with highest probability. However, since P^* recursively takes many samples in each round, the number of samples grows exponentially in the number of rounds. Thus, the reduction 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 $(\vec{v}_1,\vec{p}_1,\ldots,\vec{v}_m,\vec{p}_m)$ denote the messages of $\langle\mathsf{P}^{n*},\mathsf{V}^{n,n}\rangle$, where $\vec{v}_j=(v_{j,1},\ldots,v_{j,n})$ and $\vec{p}_j=(p_{j,1},\ldots,p_{j,n})$ 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 $\vec{v}_{j,-i}=(v_{j,1},\ldots,v_{j,i-1},v_{j,i+1},\ldots,v_{j,n})$, and $\vec{v}_{j+1},\ldots,\vec{v}_m$ uniformly at random to complete the interaction. Once the continuation is *successful*, i.e., $\mathsf{V}^{n,n}$ accepts, P^* selects the $\vec{v}_{j,-i}$ 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 a 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 distance $O(\sqrt{\log(1/\varepsilon)/n})$.

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

Thus, the real prover P^* can succeed with probability at least $1 - O(m\sqrt{\log(1/\varepsilon)/n})$. 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 reduction algorithm directly, avoiding the use of any form of Raz's Lemma. We lower bound the success probability of the reduction algorithm 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[3] show by induction on the round $j \in [m]$ that

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

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 Extendable and Simulatable Verifiers

The results of Håstad et al. [HPPW08] extend to arguments with *extendable* and *simulatable* verifiers defined in [HPPW08]. The model generalizes both three-message arguments and public-coin arguments, and contains other natural protocols. Roughly speaking, extendability and simulatability refer to the ability of efficiently simulating a random continuation of internal and external verifiers, respectively. Extendability means that given a partial transcript and the coins of a verifier, one can efficiently sample from the distribution of the possible next messages conditioning on the partial transcript, including the final decision of the verifier. Simulatability means that given a partial transcript (without the verifier's coins), one can efficiently generate a next message. However, since the verifier's coins are not given, one may not know the decision of the verifier in the end of the interaction.

The argument of Håstad et al. [HPPW08] extends to this model, and gives parallel repetition theorems with the same parameters. That is, the reduction achieves success probability $1 - \gamma - O(m\sqrt{\log(1/\varepsilon)/n})$ for Cheonoff-type case, and $1 - O(m\sqrt{\log(1/\varepsilon)/n})$ for direct product case, where m is the number of rounds. Unfortunately, the analysis of Wikström [Wik09] does not extend to this model. Thus the best known bound remains dependent on m.

Our Result. We give a new reduction 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

[3] 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.

tight. In particular, removing the dependency on m extends the "concurrent" repetition theorem of Wikström [Wik09] (for public-coin arguments) to arguments with extendable and simulatable verifiers.

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^* cannot check whether a random continuation is successful. To handle this issue, [HPPW08] ignores the decision of the external verifier and uses rejection sampling with "soft decision": the more the number of accepting internal verifiers, the higher the probability that the prover selects a random continuation. [HPPW08] shows that in the ideal version of this procedure, the success probability of $\tilde{\mathsf{P}}^*$ is still close to 1. Thus, the success probability of P^* is close to $1 - O(m\sqrt{\log(1/\varepsilon)/n})$ by Raz's Lemma.

As we do not analyze the ideal procedure, we do not use the soft decision. Rather, our reduction P^* simply selects a random continuation where all the internal verifiers accept. We are able to lower bound the success probability of P^* by a similar induction. Let η_i be the success probability of P^* when the external verifier V is embedded in the i-th coordinate. Let γ be the success probability of P^{n*}, and γ_i the probability that P^{n*} convinces verifiers $\mathsf{V}_1, \ldots, \mathsf{V}_{i-1}, \mathsf{V}_{i+1}, \ldots, \mathsf{V}_n$. We show by induction on the round $j \in [m]$ that

$$\prod_{i=1}^{n} \eta_i \geq \left(\frac{\gamma^{n+1}}{\prod_{i=1}^{n} \gamma_i} \right), \text{ when conditioning on any partial interaction } (\vec{v}_1, \vec{p}_1, \ldots, \vec{v}_j, \vec{p}_j).$$

By assumption on P^{n*}, we have $\gamma \geq \varepsilon$. If P^{n*} satisfies the additional property that $\gamma_i \leq \varepsilon^{\frac{n-1}{n}}$, then the success probability of P^* is

$$\frac{1}{n} \cdot \sum_{i=1}^{n} \eta_i \geq \left(\prod_{i=1}^{n} \eta_i \right)^{1/n} \geq \left(\frac{\gamma^{n+1}}{\prod_{i=1}^{n} \gamma_i} \right)^{1/n} \geq \varepsilon^{2/n}.$$

Of course, the additional property may not hold. Fortunately, we can enforce this property by the following observation. If there exists some coordinate $i \in [n]$ such that $\gamma_i \geq \varepsilon^{(n-1)/n}$, then we can obtain a prover $\mathsf{P}^{(n-1)*}$ that convinces $\mathsf{V}^{n-1,n-1}$ with probability at least $\gamma_i \geq \varepsilon^{(n-1)/n}$. For simplicity, let us assume that we can compute γ_i's exactly.[4] The prover $\mathsf{P}^{(n-1)*}$ simply finds such a coordinate i and interacts with $\mathsf{V}^{n-1,n-1}$ by simulating the interaction of P^{n*} and $\mathsf{V}^{n,n}$ with the i-th coordinate played by an internal verifier and the rest coordinates played by $\mathsf{V}^{n-1,n-1}$. Clearly, $\mathsf{P}^{(n-1)*}$ convinces $\mathsf{V}^{n-1,n-1}$ if and only if P^{n*} convinces verifiers $\mathsf{V}_1, \ldots, \mathsf{V}_{i-1}, \mathsf{V}_{i+1}, \ldots, \mathsf{V}_n$ of $\mathsf{V}^{n,n}$, and the probability is γ_i by definition. Applying the observation iteratively, we obtain a prover $\mathsf{P}^{n'*}$ such that (i) $\mathsf{P}^{n'*}$ convinces $\mathsf{V}^{n',n'}$ with probability at least $\varepsilon^{n'/n}$, and (ii) either $n' = 1$ or the additional property holds. If $n' = 1$, then we are done. Otherwise, applying the reduction on $\mathsf{P}^{n'*}$, we obtain a prover P^* with success probability at least $(\varepsilon^{n'/n})^{2/n'} = \varepsilon^{2/n}$.

1.3 Extension to Chernoff-type Theorems

We give a simple and generic reduction which shows that tight direct product theorems imply almost tight Chernoff-type theorems, and thus extend our results to Chernoff-type Theorems. Our reduction applies to various models such as weakly-verifiable puzzles, and gives an alternative proof

[4] It is easy to estimate γ_i by sampling, and a more careful argument can handle the estimation error.

to the Chernoff-type theorem of Impagliazzo et al. [IJK07] as a consequence of the tight direct product theorem of Canetti et al. [CHS05].

The reduction converts a parallel prover P^{n*} for $\mathsf{V}^{n,k}$ to a parallel prover P^{t*} for $\mathsf{V}^{t,t}$ for any $t \leq k$. 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 $k = (1-\gamma)n$, 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 $(\varepsilon \cdot \binom{k}{t} / \binom{n}{t})^{1/t} \approx 1 - \gamma - O(\sqrt{\log(1/\varepsilon)/n})$.[5]

For public-coin arguments, the reduction extends our direct product theorem to a Chernoff-type theorem with similar parameter to [Wik09]. For arguments with extendable and simulatable verifiers, the reduction and our improved direct product theorem gives a prover P^* with success probability $(1-\gamma)^2 - O(\sqrt{\log(1/\varepsilon)/n})$. The bound is incomparable to the previous bound of $1 - \gamma - O(m\sqrt{\log(1/\varepsilon)/n})$ of [HPPW08] in that our bound does not depend on m, but has a slightly worse dependency on γ.

As an additional contribution, we also prove that the reduction 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 $\Pr[X_i = 1] = \delta$.

1.4 Related Work

An important work that is not mentioned in the above discussion is the work of Haitner [Hai09]. Haitner proves that any interactive arguments can be modified slightly so that parallel repetition does reduce the soundness error. However, Haitner's result holds only for sufficiently large n and is far from the information-theoretic bound: the error probability is $1 - O(mn^{-1/10})$. Haitner's analysis also seems to use Raz's Lemma in an essential way.

1.5 Organization of the Paper

We first present some basic notation. Then we prove the direct product theorem for public-coin arguments in Section 3. We extend our results to arguments with extendable and simulatable verifiers in Section 4, and Chernoff-type theorems in Section 5. Finally we discuss constant-round public-coin arguments in Section 6.

2 Preliminary and Notation

We use s to denote the security parameter, and 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

[5] 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 $1 - \gamma - \max\{\alpha, O(\sqrt{\log(1/\varepsilon)/n})\}$ for an arbitrarily small constant α, which suffices for most applications.

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 $(\vec{v}_1, \vec{p}_1, \ldots, \vec{v}_m, \vec{p}_m) = \langle \mathsf{P}^n, \mathsf{V}^{n,k}\rangle(x)$, where $\vec{v}_j = (v_{j,1}, \ldots, v_{j,n})$ and $\vec{p}_j = (p_{j,1}, \ldots, p_{j,n})$.

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 $(\vec{v}_1, \ldots, \vec{v}_m)$. Thus, we can skip prover's messages and describe an interaction by $(\vec{v}_1, \ldots, \vec{v}_m)$. We refers to a partial transcript as a *history* $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_j)$.

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 2 *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 can assume without loss of generality 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 accepting probability.

Let us first recall the common approach of such a reduction. On input x, the reduced 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 $(\vec{v}_1, \ldots, \vec{v}_m)$. Let $T_i(\cdot)$ denote whether V_i accepts. That is, $T_i(\vec{v}_1, \ldots, \vec{v}_m) = 1$ iff V_i accepts in $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x)$ with history $(\vec{v}_1, \ldots, \vec{v}_m)$.

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 $\vec{v}_{j,-i} = (v_{j,1}, \ldots, v_{j,i-1}, v_{j,i+1}, \ldots, v_{j,n})$ alternately. At the end, P^* *succeeds* if $T_i(\vec{v}_1,\ldots,\vec{v}_m)=1$. Note that a node of the game tree is of the form either $(i;\vec{v}_1,\ldots,\vec{v}_j)$, in which case it is V's turn to move, or of the form $(i;\vec{v}_1,\ldots,\vec{v}_{j-1},v_{j,i})$, 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 3 (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 $u=(i;\vec{v}_1,\ldots,\vec{v}_{j-1},v_{j,i})$, 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 $\vec{v}_{j,-i},\ldots,\vec{v}_{m,i},\vec{v}_{m,-i}$. A continuation is* successful *if all verifiers accept, i.e., $T_\ell(\vec{v}_1,\ldots,\vec{v}_m)=1$ for all $\ell\in[n]$. The first time a successful continuation is found, P^*_{rej} plays the corresponding move $\vec{v}_{j,-i}$. If no successful continuations are found, P^*_{rej} aborts.*

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

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 subsequent section.

3.1 Analysis of 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 $(\vec{v}_1,\ldots,\vec{v}_j)$ or $(\vec{v}_1,\ldots,\vec{v}_{j-1},v_{j,i})$. 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 4 *For every $i \in [n]$ and full history $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_m)$, we have $\eta_i(\bar{h}) = T_i(\bar{h})$. For every $i \in [n]$, $j \in [m]$, and history $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_{j-1})$, we have*[6]

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

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 $\vec{v}_{j,-i}$ with probability $\Pr[\vec{v}_{j,-i}] \cdot (\gamma(\bar{h}, \vec{v}_j)/\gamma(\bar{h}, v_{j,i}))$, which corresponds to the expectation operator over $\vec{v}_{j,-i}$ with factor $\gamma(\bar{h}, \vec{v}_j)/\gamma(\bar{h}, v_{j,i})$ 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 $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_j)$,

$$\prod_{i=1}^{n} \eta_i(\bar{h}) \geq \gamma(\bar{h}).$$

The base case $j = m$ is trivial. For every full history $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_m)$, $\gamma(\bar{h}) = 1$ 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 $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_j)$, we want to prove the inductive hypothesis for $j - 1$ and every $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_{j-1})$. More precisely, for every $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_{j-1})$, we want to show that

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

provided that for every $\vec{v}_j$,

$$\prod_{i=1}^{n} \eta_i(\bar{h}, \vec{v}_j) \geq \gamma(\bar{h}, \vec{v}_j).$$

For notational simplicity, we abstract the above statement as the following lemma.

Lemma 5 *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 $\prod_i \eta_i(\vec{q}) \geq \gamma(\vec{q})$ for every $\vec{q} = (q_1, \ldots, q_n) \in \Omega^n$. Let $\gamma = \mathrm{E}_{\vec{q}}[\gamma(\vec{q})]$. For every $i \in [n]$, let*

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

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

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

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

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.

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

We now explain the application of Hölder's Inequalities.

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

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

- 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

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

■

Remark 6 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

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

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

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

3.2 Analysis of P^*_{rej}

In this section, we modify the analysis in Section 3.1 to analyze the success probability of strategy P^*_{rej}. 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. Nevertheless, we can still lower bound the success probability by induction on a modified inductive hypothesis.

Again, we start by expressing the success probability of P^*_{rej}. Recall γ and η_i defined in the previous subsection. We use the same γ and modify the definition of η_i to be the conditional success probability of P^*_{rej}. That is, for every $i \in [n]$ and history $\bar{h}$, we define

$$\eta_i(\bar{h}) \stackrel{\text{def}}{=} \Pr[\mathsf{P}^*_{rej} \text{ succeeds } |u = (i; \bar{h})].$$

The formula for $\eta_i(\cdot)$ is given by the following claim.

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

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

where $f(\alpha) = (1 - (1-\alpha)^M)$, and $M = O(mn/\varepsilon\xi)$.

Proof. Observing that P^*_{rej} can find a successful continuation with probability exactly $f(\gamma(\bar{h}, v_{j,i}))$, and that conditioning on a successful continuation is found, P^*_{rej} plays $\vec{v}_{j,-i}$ with the same probability as P^*_{ideal}, we obtain the above formula for η_i. ■

Our goal is to show that if the accept probability of $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x)$ is $\gamma \geq \varepsilon$, then the success probability of P^*_{rej} is at least $\delta \cdot (1-\xi)$. Let $\nu = 1/M$. We use the following inductive hypothesis: for every $j \in \{0, \ldots, m\}$ and history $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_j)$,

$$\prod_{i=1}^{n} \eta_i(\bar{h}) \geq \left(\frac{(\gamma(\bar{h}) - (m-j)\cdot\nu)_+^{n+1}}{(\gamma(\bar{h}) + \nu)^n}\right),$$

where $(\alpha)_+ \stackrel{\text{def}}{=} \max\{\alpha, 0\}$. Again, the base case is trivial to check. We prove the following lemma for the inductive step.

Lemma 8 *Let $\nu \in (0,1)$ and $t, M \geq 0$ such that $M \cdot \nu \geq 1$. Let $\gamma, \eta_1, \ldots, \eta_n : \Omega^n \to [0,1]$ be $[0,1]$-valued functions over Ω^n such that*

$$\prod_i \eta_i(\vec{q}) \geq \left(\frac{(\gamma(\vec{q}) - t\cdot\nu)_+^{n+1}}{(\gamma(\vec{q}) + \nu)^n}\right)$$

for every $\vec{q} = (q_1, \ldots, q_n) \in \Omega^n$. Let $\gamma = \mathrm{E}_{\vec{q}}[\gamma(\vec{q})]$. For every $i \in [n]$, let

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

where $f(\alpha) = (1-(1-\alpha)^M)$, and the above expectation is over uniform distribution over Ω^n. We have

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

Proof. The proof is similar to that of Lemma 16 but a bit more technical. Again, we first write down the whole computation, and then justify the inequalities.

$$\begin{aligned}
&\prod_{i=1}^{n} \mathop{\mathrm{E}}_{\vec{q}}\left[\left(\frac{\gamma(\vec{q})\cdot\eta_i(\vec{q})}{\gamma(q_i)}\cdot f(\gamma(q_i))\right)\right] \\
&\quad= \prod_{i=1}^{n} \mathop{\mathrm{E}}_{\vec{q}}\left[\left(\frac{\gamma(\vec{q})\cdot\eta_i(\vec{q})}{\gamma(q_i)/f(\gamma(q_i))}\right)\right] \\
&\quad\geq \mathop{\mathrm{E}}_{\vec{q}}\left[\left(\frac{\gamma(\vec{q})^n\cdot\prod_{i=1}^{n}\eta_i(\vec{q})}{\prod_{i=1}^{n}(\gamma(q_i)/f(\gamma(q_i)))}\right)^{1/n}\right]^n \qquad \text{(by Hölder's Inequality)} \\
&\quad\geq \mathop{\mathrm{E}}_{\vec{q}}\left[\left(\frac{\gamma(\vec{q})^n\cdot(\gamma(\vec{q})-t\cdot\nu)_+^{n+1}/(\gamma(\vec{q})+\nu)^n}{\prod_{i=1}^{n}(\gamma(q_i)/f(\gamma(q_i)))}\right)^{1/n}\right]^n \qquad \text{(by inductive hypothesis)} \\
&\quad\geq \mathop{\mathrm{E}}_{\vec{q}}\left[\left(\frac{(\gamma(\vec{q})-(t+1)\cdot\nu)_+^{n+1}}{\prod_{i=1}^{n}(\gamma(q_i)/f(\gamma(q_i)))}\right)^{1/n}\right]^n \\
&\quad\geq \left[\left(\frac{\mathrm{E}_{\vec{q}}[(\gamma(\vec{q})-(t+1)\cdot\nu)_+]^{n+1}}{E_{\vec{q}}[\prod_{i=1}^{n}(\gamma(q_i)/f(\gamma(q_i)))]}\right)^{1/n}\right]^n \qquad \text{(by Hölder's Inequality)} \\
&\quad\geq \left(\frac{(\gamma-(t+1)\cdot\nu)_+^{n+1}}{(\gamma+\nu)^n}\right)
\end{aligned}$$

In the first equality, observing that $\alpha/f(\alpha)\to 1/M$ as $\alpha\to 0$, we can take the convention that $0/f(0)=1/M$. This gives us a correct formula for η_i's, and gets around the zero denominator issue. The application of Hölder's Inequalities are the same as the proof in Lemma 16. We check the third and last inequality below.

- Third inequality: we check that the inequality holds pointwisely for every $\vec{q}$. The denominator is the same. For the numerator, we need to check that

  $$\left(\frac{\gamma(\vec{q})\cdot(\gamma(\vec{q})-t\cdot\nu)_+^{n+1}}{(\gamma(\vec{q})+\nu)^n}\right)\geq(\gamma(\vec{q})-(t+1)\cdot\nu)_+^{n+1},$$

  which follows by inequality $\alpha^n\cdot(\alpha-t\nu)_+^{n+1}\geq(\alpha+\nu)^n\cdot(\alpha-(t+1)\nu)_+^{n+1}$ for every $t,\alpha,\nu\geq 0$ (Claim 9 below).

- Last inequality: We have $\mathrm{E}_{\vec{q}}[(\gamma(\vec{q})-(t+1)\cdot\nu)_+]\geq(\gamma-(t+1)\cdot\nu)_+$ for the numerator by Jensen's inequality. For the denominator, we check that for every $i\in[n]$,

  $$E_{q_i}\left[\frac{\gamma(q_i)}{1-(1-\gamma(q_i))^M}\right]\leq\gamma+\nu.$$

  This holds since if $M\nu\geq 1$, then $\alpha/(1-(1-\alpha)^M)\leq\alpha+\nu$ for every $\alpha\in[0,1]$ (Claim 10 below).

■

Claim 9 *The inequality $\alpha^n\cdot(\alpha-t\nu)_+^{n+1}\geq(\alpha+\nu)^n\cdot(\alpha-(t+1)\nu)_+^{n+1}$ holds for every $t,\alpha,\nu\geq 0$.*

Proof. Fix arbitrary $t, \nu \geq 0$, the inequality is trivial for $\alpha \leq (t+1)\nu$. For $\alpha \geq (t+1)\nu$, let us consider $h(x) \stackrel{\text{def}}{=} (\alpha + x)^n \cdot (\alpha - t\nu - x)^{n+1}$. Clearly, we have $h(0) = \alpha^n \cdot (\alpha - t\nu)_+^{n+1}$, and $h(\nu) = (\alpha+\nu)^n \cdot (\alpha - (t+1)\nu)_+^{n+1}$. Furthermore, it is easy to verify that $h'(x) \leq 0$ for every $x \in [0, \nu]$. Therefore, we have $h(0) \geq h(\nu)$, which proves the claim. ■

Claim 10 *Let $M \in \mathbb{R}$, $\nu \in (0,1]$ be two numbers with $M\nu \geq 1$. Let $g(\alpha) = \begin{cases} \frac{\alpha}{1-(1-\alpha)^M} & \alpha \in (0,1] \\ 1/M & \alpha = 0 \end{cases}$. Then $g(\alpha) \leq \alpha + \nu$ for $\alpha \in [0,1]$.*

Proof. If $\alpha = 0$, then the inequality holds trivially. For the case $\alpha \in (0,1]$, first we consider the function $h(\alpha) = (1 - (1-\alpha)^M)(\alpha + \nu) - \alpha$, and prove that $h(\alpha) \geq 0$. By the fact $h'(\alpha) = (1-\alpha)^M(M\alpha + M\nu - 1) \geq 0$ for $\alpha \in [0,1]$, which tells that h is a non-decreasing function in $[0,1]$, we have $h(0) = 0$ implies $h(\alpha) \geq h(0) \geq 0$ for $\alpha \in [0,1]$. Then we observe that for $\alpha \in (0,1]$, $h(\alpha) \geq 0$ implies $g(\alpha) \leq \alpha + \nu$ (since $(1 - (1-\alpha)^M) > 0$). Thus the claim holds for $\alpha \in [0,1]$. ■

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

$$\frac{1}{n} \cdot \sum_{i=1}^{n} \eta_i \geq \left(\prod_{i=1}^{n} \eta_i \right)^{1/n} \geq \left(\frac{(\gamma - m \cdot \nu)_+^{n+1}}{(\gamma + \nu)^n} \right)^{1/n} \geq \gamma^{1/n} \cdot (1 - O(mn\nu/\gamma)) \geq \varepsilon^{1/n} \cdot (1 - \xi).$$

4 Protocols with Extendable and Simulatable Verifier

In this section, we present a new reduction that extends our results to interactive protocols with extendable and simulatable verifiers. As mentioned in the introduction, extendability and simulatability refer to the ability of efficiently simulating a random continuation of internal and external verifiers, respectively. Extendability means that given a partial transcript and the coins of a verifier, one can efficiently sample from the distribution of the possible next messages conditioning on the partial transcript, including the final decision of the verifier. Simulatability means that given a partial transcript (without the verifier's coins), one can efficiently generate a next message. However, since the verifier's coins are not given, one may not know the decision of the verifier in the end of the interaction. We will not define the properties formally and refer the reader to [HPPW08] for more details.

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 extendable and simulatable verifier V with success probability $\delta^2 = \varepsilon^{2/n} \approx 1 - O(\log(1/\varepsilon)/n)$.[7] As discussed in the introduction, when the verifier is extendable and simulatable, in the game $\mathcal{G}(\mathsf{P}^{n*}, x)$, one can still simulate a random continuation from any P^*-move node u, but no longer be able to know the decision of the external verifier. Thus, we first apply a preprocessing algorithm and then use a modified rejection sampling strategy.

Let P^{n*} be a parallel prover for $\mathsf{V}^{n,n}$ and for some input x, $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x) = 1] \geq \delta^n$. We first apply a preprocessing algorithm to turn P^{n*} to a parallel prover $\mathsf{P}^{n'*}$ for $\mathsf{V}^{n',n'}$ for some $1 \leq n' \leq n$ such that (i) $\Pr[\langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'} \rangle(x) = 1] \gtrsim \delta^{n'}$, and (ii) for all $i \in [n']$, $\Pr[\ \mathsf{V}_j$ accepts $\forall j \neq i$ in $\langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'}(x) \rangle] \lesssim \delta^{n'-1}$. If $n' = 1$, then we are done. Otherwise, we use a rejection sampling strategy where a random continuation is selected when all internal verifiers accept. By a similar induction argument as in the

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

proof of Theorem 2, we show that this reduction algorithm can break soundness with probability close to δ^2. Formally, we obtain the following theorem.

Theorem 11 *Let* $\mathsf{V} \in \mathrm{PPT}$ *be extendable and simulatable. There exists a prover strategy* P^* *such that for every common input* x*, every* $n \in \mathbb{N}$*, every* $\varepsilon, \xi \in (0,1)$*, 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. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x) = 1] \geq \varepsilon \Rightarrow$

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

We present the preprocessing algorithm and its analysis in the following section, and the rejection sampling strategy and its analysis in the subsequent section.

4.1 The Preprocessing Algorithm

Let $\langle \mathsf{P}, \mathsf{V}\rangle$ be an arbitrary interactive argument. Let P^{n*} be a parallel solver and x a common input such that $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x) = 1] \geq \delta^n$. In this subsection, we present a transformation $\mathcal{T}$ that runs in time $\mathrm{poly}(s, \delta^n, \xi)$ and converts P^{n*} to another parallel prover $\mathsf{P}^{n'*}$ for some $1 \leq n' \leq n$ with $\Pr[\langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'}\rangle(x) = 1] \geq \delta^{n'}(1 - \xi)$ and $\Pr[\ \mathsf{V}_j \text{ accepts } \forall j \neq i \text{ in } \langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'}(x)\rangle] \leq \delta^{n'-1}(1 + \xi)$ for every $i \in [n']$. Formally, we prove the following lemma.

Lemma 12 *Let* $\mathsf{V} \in \mathrm{PPT}$*. There exists a* PPT *transformation* $\mathcal{T}$ *such that for every common input* x*, every* $n \in \mathbb{N}$*, every* $\delta, \xi \in (0,1)$*, and every efficient parallel prover strategy* P^{n*}*,*

1. $\mathcal{T}(x, n, \delta, \xi)$ *runs in time* $\mathrm{poly}(|x|, n, \delta^{-n}, \xi^{-1})$ *given oracle access to* $\mathsf{P}^{n*}(x)$*.*
2. *If* $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x) = 1] \geq \delta^n$*, then with overwhelming probability,* $\mathcal{T}(x, n, \delta, \xi)$ *outputs an integer* $n' \in [n]$ *and a parallel prover strategy* $\mathsf{P}^{n'*}$ *such that giving* $\mathsf{P}^{n'*}$ *oracle access to* P^{n*}*,* $\langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'}\rangle(x)$ *runs in time* $\mathrm{poly}(|x|, n)$*. Furthermore,*

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

and for every $i \in [n']$*,*

$$\Pr[\ \mathsf{V}_j \textit{ accepts } \forall j \neq i \textit{ in } \langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'}\rangle(x)] \leq \delta^{n'-1}(1 + \xi).$$

Proof. We present the transformation in a slightly different way to that in the introduction. The transformation $\mathcal{T}$ selects a subset $S \subset [n]$ of coordinates by sampling and outputs the following $\mathsf{P}_S^{n'*}$ with $n' = n - |S|$: to interact with $\mathsf{V}^{n',n'}$, $\mathsf{P}_S^{n'*}$ simulates the interaction of $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle$ with internal verifiers plays coordinate S and external verifier $\mathsf{V}^{n',n'}$ plays the rest coordinates. Note that $\mathsf{P}_S^{n'}$ runs in time $\mathrm{poly}(|x|, n)$ given oracle access to P^{n*}. $\mathcal{T}$ selects $S \subset [n]$ as follows:

1. Initially, set $S = \phi$.
2. Simulate the interaction $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x)$ for $M = O(n/(\delta^{2n} \cdot \xi^2))$ times. For a coordinate set $T \subset [n]$, let $p(T) = (\ \#\text{ of simulation that } \mathsf{V}_i \text{ accepts } \forall i \in T)/M$ be an estimation of the probability $\Pr[\forall i \in T, \mathsf{V}_i \text{ accepts in } \langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x)]$.

3. Iteratively, if there exists $i \notin S$ such that $p([n]\backslash(S \cup \{i\})) \geq \delta^{n-|S|-1}$, then add i to S until no such coordinate i exists or $|S| = n - 1$.

4. Outputs $n' = n - |S|$ and $\mathsf{P}_S^{n'}$.

Clearly, the process terminates in n iterations and in each iteration, we check at most n estimation $p(\cdot)$'s. By standard Chernoff bounds and union bounds, all estimations have error at most δ^n/ξ with overwhelming probability. Thus, $\mathcal{T}$ runs in time $\mathrm{poly}(|x|, \delta^n, \xi)$ and outputs a number $n' \in [n]$ and a prover $\mathsf{P}_S^{n'}$ such that either $n' = 1$, or $\Pr[\langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'} \rangle(x) = 1] \geq \delta^{n'}(1-\xi)$ and $\Pr[\ \mathsf{V}_j \text{ accepts } \forall j \neq i \text{ in } \langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'}(x) \rangle] \leq \delta^{n'-1}(1+\xi)$ for every $i \in [n']$. ■

4.2 The Rejection Sampling Strategy

The preprocessing algorithm in the above section converts a parallel prover P^{n*} with $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x) = 1] \geq \delta^n$ to a parallel prover $\mathsf{P}^{n'*}$ for some $1 \leq n' \leq n$ with $\Pr[\langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'} \rangle(x) = 1] \geq \delta^{n'}(1-\xi)$, and in addition, $\Pr[\ \mathsf{V}_j \text{ accepts } \forall j \neq i \text{ in } \langle \mathsf{P}^{n'*}, \mathsf{V}^{n',n'}(x) \rangle] \leq \delta^{n'-1}(1+\xi)$ for all $i \in [n']$. In this section, we present a modified rejection sampling strategy of P^* in the game $\mathcal{G}(\mathsf{P}^{n'*}, x)$, which can be implemented efficiently when the verifier is extendable and simulatable, and can succeed with probability close to δ^2. Formally we prove the following lemma.

Lemma 13 *Let $\mathsf{V} \in$ PPT be extendable and simulatable. There exists a prover strategy P^* such that for every common input x, every $n \in \mathbb{N}$, every $\delta, \xi \in (0,1)$, and every parallel prover strategy P^{n*},*

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

$$\Pr[\ \mathsf{V}_j \textit{ accepts } \forall j \neq i \textit{ in } \langle \mathsf{P}^{n*}, \mathsf{V}^{n,n} \rangle(x)] \leq \delta^{n-1}(1+\xi),$$

then

$$\Pr[\langle \mathsf{P}^*, \mathsf{V} \rangle(x) = 1] \geq \delta^2 \cdot (1 - 4\xi).$$

We start by defining the modified rejection sampling strategy, which is the same as Definition 3 except that the definition of a successful continuation is modified.

Definition 14 (Strategy $\tilde{\mathsf{P}}^*_{rej}$) *We define strategy $\tilde{\mathsf{P}}^*_{rej}$ as follows. Let P^{n*} be a deterministic parallel prover, x a common input, and $\mathcal{G}(\tilde{\mathsf{P}}^{n*}, x)$ the corresponding game.*

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

For simplicity, we again consider the ideal version $\tilde{\mathsf{P}}^*_{ideal}$ of $\tilde{\mathsf{P}}^*_{rej}$, and prove the success probability of $\tilde{\mathsf{P}}^*_{ideal}$ is at least $\delta^2(1-4\xi)$. The proof can be generalized to handle $\tilde{\mathsf{P}}^*_{rej}$ as in the public-coin case.

We introduce similar notations to express the success probability of $\tilde{\mathsf{P}}^*_{ideal}$. For every history $\bar{h}$, let $\gamma(\bar{h}) \stackrel{\text{def}}{=} \Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x) = 1|\bar{h}]$, $\eta_i(\bar{h}) \stackrel{\text{def}}{=} \Pr[\tilde{\mathsf{P}}^*_{ideal}$ succeeds $|u = (i; \bar{h})]$, and

$$\gamma_i(\bar{h}) \stackrel{\text{def}}{=} \Pr[\ \mathsf{V}_j \text{ accepts } \forall j \neq i \text{ in } \langle \mathsf{P}^{n*}, \mathsf{V}^{n,n}\rangle(x)|\bar{h}],$$

for every $i \in [n]$. Note that by assumption, we have $\gamma \geq \delta^n$ and $\gamma_i \leq \delta^{n-1}$ for every $i \in [n]$. We have the following claim for $\eta_i(\cdot)$ similar to Claim 4.

Claim 15 *For every $i \in [n]$ and full history $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_m)$, we have $\eta_i(\bar{h}) = T_i(\bar{h})$. For every $i \in [n]$, $j \in [m]$, and history $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_{j-1})$, we have*

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

Proof. The proof is similar to that of Claim 4. The only difference is that now rejection sampling is accepted when all internal verifiers accept, which corresponds to γ_i instead of γ. ■

This time, we induct on the following inductive hypothesis: for every $j \in \{0, \ldots, m\}$ and history $\bar{h} = (\vec{v}_1, \ldots, \vec{v}_j)$,

$$\prod_{i=1}^{n} \eta_i(\bar{h}) \geq \frac{\gamma^{n+1}(\bar{h})}{\prod_{i=1}^{n} \gamma_i(\bar{h})}.$$

The base case $j = m$ is trivial. We prove the following lemma for the inductive step.

Lemma 16 *Let $\gamma, \gamma_1, \ldots, \gamma_n, \eta_1, \ldots, \eta_n : \Omega^n \to [0,1]$ be $[0,1]$-valued functions over a product space Ω^n such that*

$$\prod_{i} \eta_i(\vec{q}) \geq \frac{\gamma^{n+1}(\vec{q})}{\prod_{i=1}^{n} \gamma_i(\vec{q})}$$

for every $\vec{q} = (q_1, \ldots, q_n) \in \Omega^n$. Let $\gamma = \mathrm{E}_{\vec{q}}[\gamma(\vec{q})]$. For every $i \in [n]$, let

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

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

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

Proof. The proof is similar.

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

■

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

$$\frac{1}{n}\cdot\sum_{i=1}^{n}\eta_i \geq \left(\prod_{i=1}^{n}\eta_i\right)^{1/n} \geq \left(\frac{\gamma^{n+1}}{\prod_{i=1}^{n}\gamma_i}\right)^{1/n} \geq \left(\frac{(\delta^n(1-\xi))^{(n+1)}}{(\delta^{n-1}(1+\xi))^n}\right)^{1/n} \geq \delta^2(1-4\xi).$$

Theorem 11 follows straightforwardly by applying Lemma 12 and 13 with $\delta = \varepsilon^{1/n}$ and $\xi/4$.

5 Extension to Chernoff-type Theorems

In this section, we present a generic reduction 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 reduction can be used to show that tight direct product theorems implies Chernoff-type theorems. For example, using our reduction with the direct product theorem of Canetti et al. [CHS05] yields an alternative proof of the Chernoff-type theorem of Impagliazzo et al. [IJK07] for weakly-verifiable puzzles. The reduction also extends our direct product theorems to Chernoff-type theorems.

The reduction 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 17 *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*

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

Proof. By definition, we have

$$\begin{aligned}
&\Pr[\langle \mathsf{P}^{t*}, \mathsf{V}^{t,t}\rangle(x) = 1] \\
&\quad\geq \Pr[\mathsf{V}_i \text{ accepts } \forall i \in S \text{ in } \langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) \wedge \langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \\
&\quad= \Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \cdot \Pr\left[\mathsf{V}_i \text{ accepts } \forall i \in S \text{ in } \langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) \,\middle|\, \langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1\right] \\
&\quad\geq \Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \cdot \frac{\binom{k}{t}}{\binom{n}{t}}.
\end{aligned}$$

■

When V is public-coin, the above lemma and Theorem 2 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[8] 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 18 *Let $\alpha, \gamma \in (0,1)$ be any* constants *such that $\alpha + \gamma < 1$. 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. *$\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,(1-\gamma)n}\rangle(x) = 1] \geq \varepsilon \Rightarrow$*

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

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

Proof. For simplicity, let us assume $(1-\gamma)n$ is an integer (if not, we can replace $(1-\gamma)n$ by $\lceil (1-\gamma)n \rceil$.) Applying Lemma 17 and Theorem 2 with $t = 2 \cdot \log(1/\varepsilon)/\alpha$ and $\xi = \alpha/2$, we obtain a prover P^* such that for every common input x with $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,(1-\gamma)n}\rangle(x) = 1] \geq \varepsilon$,

$$\Pr[\langle \mathsf{P}^*, \mathsf{V}\rangle(x) = 1] \geq \left(\varepsilon \cdot \frac{\binom{(1-\gamma)n}{t}}{\binom{n}{t}}\right)^{1/t} \cdot (1 - \xi) \geq 1 - \gamma - \alpha,$$

since by the setting of parameters, we have $\varepsilon^{1/t} \geq (1 - \alpha/2)$, $(1 - \xi) \geq (1 - \alpha/2)$, and

$$\left(\frac{\binom{(1-\gamma)n}{t}}{\binom{n}{t}}\right)^{1/t} \geq \left(\frac{(1-\gamma)n - t}{n - t}\right) \geq \left(\frac{1 - \gamma - \alpha/2}{1 - \alpha/2}\right).$$

Furthermore, P^* runs in time $\mathrm{poly}(|x|, n, (\binom{n}{t}/\binom{(1-\gamma)n}{t}), \varepsilon^{-1}, \xi^{-1})$, which is efficient since

$$\frac{\binom{n}{t}}{\binom{(1-\gamma)n}{t}} \leq \left(\frac{n - t}{(1-\gamma)n - t}\right)^t \leq \left(\frac{1 - \alpha/2}{1 - \gamma - \alpha/2}\right)^t = \mathrm{poly}(1/\varepsilon).$$

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

■

Similarly, when V is extendable and simulatable, we can extend Theorem 11 to the following Chernoff-type theorem.

Theorem 19 *Let $\alpha, \gamma \in (0,1)$ be any* constants *such that $\alpha + \gamma < 1$. Let $\mathsf{V} \in \mathrm{PPT}$ 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. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,(1-\gamma)n}\rangle(x) = 1] \geq \varepsilon \Rightarrow$

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

Proof. Applying Lemma 17 and Theorem 11 with $t = 4 \cdot \log(1/\varepsilon)/\alpha$ and $\xi = \alpha/4$, we obtain a prover P^* such that for every common input x with $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,(1-\gamma)n}\rangle(x) = 1] \geq \varepsilon$,

$$\Pr[\langle \mathsf{P}^*, \mathsf{V}\rangle(x) = 1] \geq \left(\varepsilon \cdot \frac{\binom{(1-\gamma)n}{t}}{\binom{n}{t}}\right)^{2/t} \cdot (1-\xi) \geq (1-\gamma)^2 - \alpha,$$

since by the setting of parameters, we have $\varepsilon^{2/t} \geq (1 - \alpha/2)$, $(1-\xi) \geq (1-\alpha/4)$, and

$$\left(\frac{\binom{(1-\gamma)n}{t}}{\binom{n}{t}}\right)^{2/t} \geq \left(\frac{(1-\gamma)n - t}{n-t}\right)^2 \geq \left(\frac{1-\gamma-\alpha/4}{1-\alpha/4}\right)^2.$$

Furthermore, P^* runs in time $\mathrm{poly}(|x|, n, (\binom{n}{t}/\binom{(1-\gamma)n}{t}), \varepsilon^{-1}, \xi^{-1})$, which is efficient since

$$\frac{\binom{n}{t}}{\binom{(1-\gamma)n}{t}} \leq \left(\frac{n-t}{(1-\gamma)n - t}\right)^t \leq \left(\frac{1-\alpha/4}{1-\gamma-\alpha/4}\right)^t = \mathrm{poly}(1/\varepsilon).$$

■

6 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].

Theorem 20 *Let $m \in \mathbb{N}$ be an arbitrary constant, and $\mathsf{V} \in \mathrm{PPT}$ 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. $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \geq P(n,k,\delta) \Rightarrow$

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

Let P^{n*} be a parallel prover with $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle] > P(n,k,\delta)$. Without loss of generality, we assume that P^{n*} is deterministic. Thus, the outcome of $\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle$ depends only on $\vec{v}_1, \vec{v}_2 \ldots, \vec{v}_m$, all the verifiers' messages. Our goal is to design a prover P^* that can make V wrongly accept with probability greater equal than $\delta \cdot (1-\xi)$.

The framework here is similar to the one in previous sections. Here we give a brief review. Recall that we have defined predicates $T_i(\vec{v}_1, \ldots, \vec{v}_m) = 1$ if on the history the i-th verifier $\mathsf{V}_i(\vec{v}_1, \ldots, \vec{v}_m)$ accepts, and otherwise 0. Now P^* and V is playing a game induced by P^{n*}, x, denoted as $\mathcal{G}(\mathsf{P}^{n*}, x)$, where in the i-th round V makes a move $v_{i,j}$ and P^* makes a move $\vec{v}_{i,-j}$.

Pass and Venkitasubramaniam[PV07] used this approach for proving a Direct Product Theorem for AM argument systems. They showed that, (i) if $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \geq \delta^n$, then the prover P^*_{opt} using *optimal strategy* in this approach can succeed with probability at least δ, and (ii) by recursive sampling, there exists an efficient prover P^* and a modified $\tilde{\mathsf{P}}^{n*}$ such that (a) $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \approx \Pr[\langle \tilde{\mathsf{P}}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1]$, and (b) the efficient prover P^* achieves almost the same success probability as the optimal solver $\tilde{\mathsf{P}}^*_{opt}$, who plays optimal strategy on the game $\mathcal{G}(\tilde{\mathsf{P}}^{n*}, x)$ (instead of $\mathcal{G}(\mathsf{P}^{n*}, x)$). Thus, by (i), $\Pr[\langle \mathsf{P}^*, \mathsf{V}\rangle(x) = 1] \approx \Pr[\langle \tilde{\mathsf{P}}^*_{opt}, \mathsf{V}\rangle(x) = 1] \approx \delta$, and the Direct Product Theorem follows. Note that the technical subtlety of introducing the modified prover $\tilde{\mathsf{P}}^{n*}$ is necessary, as it is possible that $\delta \approx \Pr[\langle \mathsf{P}^*, \mathsf{V}\rangle(x) = 1] \ll \Pr[\langle \mathsf{P}^*_{opt}, \mathsf{V}\rangle(x) = 1]$.

We generalize Pass and Venkitasubramaniam's[PV07] argument to prove Theorem 2. Our key observation is that for every $k \in \{1, \ldots, n\}$, if $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \geq P(n,k,\delta)$, then the optimal strategy prover in the aforementioned approach can success with probability at least δ, i.e., $\Pr[\langle \mathsf{P}^*, \mathsf{V}\rangle(x) = 1] \geq \delta$. Therefore, Theorem 2 follows by using exactly the same recursive sampling algorithm of [PV07].

Now we analyze the success probability of the optimal strategy prover P^*_{opt}.

Definition 21 *We define functions γ_i's and η_i's, which represent the success probability of the optimal P^*_{opt} at each node of the game tree of $\mathcal{G}(\mathsf{P}^{n*}, x)$ as follows. For every leaf $(\vec{v}_1, \ldots, \vec{v}_m)$, we define*

$$\gamma_i(\vec{v}_1, \ldots, \vec{v}_m) = T_i(\vec{v}_1, \ldots, \vec{v}_m).$$

For $j \in \{1, \ldots, m-1\}$ and $\vec{v}_1, \ldots, \vec{v}_{j-1}$, we inductively define $\gamma_i(\vec{v}_1, \ldots, \vec{v}_{j-1})$ and $\eta_i(\vec{v}_1, \ldots, \vec{v}_{j-1}, v_{j,i})$ for every $i = 1, \ldots, n$ as follows.

$$\eta_i(\vec{v}_1, \ldots, \vec{v}_{j-1}, v_{j,i}) = \max_{\vec{v}_{j,-i}} \gamma_i(\vec{v}_1, \ldots, \vec{v}_j), \ \ \textit{and} \ \ \gamma_i(\vec{v}_1, \ldots, \vec{v}_{j-1}) = \mathop{\mathrm{E}}_{v_{j,i}} [\eta_i(\vec{v}_1, \ldots, \vec{v}_{j-1}, v_{j,i})]$$

Finally, we define $\gamma = \max_i \gamma_i$.

Observe that γ is the success probability of the optimal P^*_{opt} in the game $\mathcal{G}(\mathsf{P}^{n*}, x)$, because the uniformly random V corresponds to expectation operator and the optimal P^*_{opt} corresponds to the maximal operator. Our goal is to prove $\Pr[\langle \mathsf{P}^*_{opt}, \mathsf{V}\rangle(x) = 1] = \gamma \geq \delta$.

The idea is to use a coupling argument. We can view T_i's as binary random variables with randomness $\vec{v}_1, \ldots, \vec{v}_m \in_R \{0,1\}^{n \times t}$. In this notation,

$$\Pr\left[\sum_{i=1}^{n} T_i \geq k\right] = \Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \geq P(n,k,\delta).$$

We show that there exists a coupling of $(T_1, \ldots, T_n)$ and random variables $(R_1, \ldots, R_n)$ such that (i) $T_i \leq R_i$ for every $i = 1, \ldots, n$ with probability 1 and (ii) R_i's are mutually independent and $\Pr[R_i = 1] = \gamma_i$. By (i), $\Pr[\sum_i R_i \geq k] \geq \Pr[\sum_i T_i \geq k] \geq P(n, k, \delta)$. Since R_i's are independent bits with bias γ_i, it is not hard to show that there exists some $\gamma_i \geq \delta$. Therefore, $\Pr[\langle \mathsf{P}^*_{opt}, \mathsf{V}\rangle(x) = 1] = \gamma = \max_i \gamma_i \geq \delta$. Formally, we prove the following two lemmas.

Lemma 22 *Let $\vec{T} = (T_1, \ldots, T_n)$ be binary random variables with randomness $\vec{p}_1, \ldots, \vec{p}_m \in_R \{0,1\}^{n\times t}$, and γ_i's and η_i's defined as in Definition 21. There exists a coupling of $(T_1, \ldots, T_n)$ and binary random variables $(R_1, \ldots, R_n)$ such that (i) $T_i \leq R_i$ for every $i = 1, \ldots, n$ with probability 1 and (ii) R_i's are mutually independent with bias $\Pr[R_i = 1] = \gamma_i$ for $i = 1, \ldots, n$.*

Proof. We define the desired random variables $\vec{R} = \vec{R}_0 = (R_{0,1}, \ldots, R_{0,n})$ by induction. We start by defining $\vec{R}_m = \vec{T}$ and inductively define $\vec{R}_j$ for $j = m-1, \ldots, 0$ using the following two inductive hypotheses.

1. $T_i \leq R_{j,i}$ for every $i = 1, \ldots, n$ with probability 1.
2. For every $\vec{p}_1, \ldots, \vec{p}_j$, the conditional random variables $(R_{j,1}, \ldots, R_{j,n})|_{\vec{p}_1,\ldots,\vec{p}_j}$ are mutually independent with $\Pr[R_{j,i} = 1] = \gamma_i(\vec{p}_1, \ldots, \vec{p}_j)$.

It follows that $\vec{R}_0 = (R_{0,1}, \ldots, R_{0,n})$ are the desired random variables. It is easy to verify that the inductive hypotheses hold for the base case $j = m$ trivially, as $\vec{R}_m = \vec{T}$ and there is no randomness after conditioning on $\vec{p}_1, \ldots, \vec{p}_m$. It remains to define $\vec{R}_{j-1}$ from $\vec{R}_j$ as follows.

We define $\vec{R}_{j-1}$ by defining its conditional distribution $\vec{R}_{j-1}|_{\vec{p}_1,\ldots,\vec{p}_j}$ for every $\vec{p}_1, \ldots, \vec{p}_j$. By the inductive hypothesis, $(R_{j,1}, \ldots, R_{j,n})|_{\vec{p}_1,\ldots,\vec{p}_j}$ are independent bits with bias $\gamma_i(\vec{p}_1, \ldots, \vec{p}_j)$. Since $\eta_i(\vec{p}_1, \ldots, \vec{p}_{j-1}, p_{j,i}) \geq \gamma_i(\vec{p}_1, \ldots, \vec{p}_j)$ for every i, we can define $(R_{j,1}, \ldots, R_{j,n})|_{\vec{p}_1,\ldots,\vec{p}_j}$ with the following two properties easily.[9]

- $R_{j,i}|_{\vec{p}_1,\ldots,\vec{p}_j} \leq R_{j-1,i}|_{\vec{p}_1,\ldots,\vec{p}_j}$ for $i = 1, \ldots, n$ with probability 1.
- $(R_{j-1,1}, \ldots, R_{j-1,n})|_{\vec{p}_1,\ldots,\vec{p}_j}$ are independent bits with $\Pr[R_{j-1,i}|_{\vec{p}_1,\ldots,\vec{p}_j} = 1] = \eta_i(\vec{p}_1, \ldots, \vec{p}_{j-1}, p_{j,i})$.

This completes the definition of $\vec{R}_{j-1}$. We now to check that $\vec{R}_{j-1}$ satisfies the inductive hypotheses. The first condition holds because $T_i \leq R_{j,i} \leq R_{j-1,i}$ for every $i = 1, \ldots, n$ with probability 1. The second condition holds because once we fix $\vec{p}_1, \ldots, \vec{p}_{j-1}$, the bias of $R_{j-1,i}$ depends only on the $p_{j,i}$ component of $\vec{p}_j$, and the $p_{j,i}$'s are independent. More formally, for every $\vec{p}_1, \ldots, \vec{p}_{j-1}$, every $i = 1, \ldots, n$ and every $\vec{r}_{-i} = (r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n) \in \{0,1\}^{n-1}$, we have

$$\begin{aligned}
&\Pr[R_{j-1,i} = 1|\vec{p}_1, \ldots, \vec{p}_{j-1}, \vec{R}_{j-1,-i} = \vec{r}_{-i}] \\
&= \mathop{\mathrm{E}}_{\vec{p}_j}[\Pr[R_{j-1,i} = 1|\vec{p}_1, \ldots, \vec{p}_j, \vec{R}_{j-1,-i} = \vec{r}_{-i}]] \\
&= \mathop{\mathrm{E}}_{\vec{p}_j}[\eta_i(\vec{p}_1, \ldots, \vec{p}_{j-1}, p_{j,i})] \quad \text{(because } R_{j-1,i} \text{ and } R_{j-1,-i} \text{ are indep. given } \vec{p}_1, \ldots, \vec{p}_j) \\
&= \mathop{\mathrm{E}}_{p_{j,i}}[\eta_i(\vec{p}_1, \ldots, \vec{p}_{j-1}, p_{j,i})] \\
&= \gamma_i(\vec{p}_1, \ldots, \vec{p}_{j-1})
\end{aligned}$$

■

[9] For example, when $R_{j,i} = 1$, we set $R_{j-1,i} = 1$, and when $R_{j,i} = 0$, we toss independent coins and set $R_{j-1,i} = 1$ with certain probability to make $\Pr[R_{j-1,i} = 1] = \eta_i(\vec{p}_1, \ldots, \vec{p}_{j-1}, p_{j,i})$.

Lemma 23 *Let* $(R_1, \ldots, R_n)$ *be independent binary random variables with bias* $\Pr[R_i = 1] = \gamma_i$ *for every* $i = 1, \ldots, n$. *If* $\Pr[\sum_i R_i \geq k] \geq P(n,k,\delta)$*, then there exists some* $\gamma_i \geq \delta$.

Proof. Let $f(\alpha_1, \ldots, \alpha_n)$ be $\Pr[\sum_i S_i \geq k]$, where the S_i's are independent binary random variables with bias $\Pr[S_i = 1] = \alpha_i$. Clearly, f is strictly increasing in every coordinate and $f(\delta, \ldots, \delta) = P(n,k,\delta)$. Therefore, if $\gamma_i < \delta$ for every i, then $\Pr[\sum_i R_i \geq k] = f(\gamma_1, \ldots, \gamma_n) < f(\delta, \ldots, \delta) = P(n,k,\delta)$, a contradiction. ■

Combining Lemma 22 and 23, we obtain the following lemma.

Lemma 24 *Let* V *be an* AM *verifier. Suppose there exists a parallel prover* P^{n*} *with* $\Pr[\langle \mathsf{P}^{n*}, \mathsf{V}^{n,k}\rangle(x) = 1] \geq P(n,k,\delta)$*, then the prover* P^*_{opt}*, who plays the optimal strategy in the game* $\mathcal{G}(\mathsf{P}^{n*}, x)$ *can successfully convince* V *with probability greater equal than* δ*, i.e.* $\Pr[\langle \mathsf{P}^*_{opt}, \mathsf{V}\rangle(x) = 1] \geq \delta$.

Followed by the argument of Pass and Venkitasubramaniam [PV07], one can relate the optimal prover P^*_{opt} to an efficient prover P^* and then this completes the proof of Theorem 2. We omit the argument and refer to [PV07] for the curious readers.

Acknowledgments

We thank Salil Vadhan for very helpful discussions throughout this work.

References

[BIN97] Mihir Bellare, Russell Impagliazzo, and Moni Naor. Does parallel repetition lower the error in computationally sound protocols? In *FOCS*, pages 374–383, 1997.

[CHS05] Ran Canetti, Shai Halevi, and Michael Steiner. Hardness amplification of weakly verifiable puzzles. In *TCC*, pages 17–33, 2005.

[CLLY09] Kai-Min Chung, Feng-Hao Liu, Chi-Jen Lu, and Bo-Yin Yang. Efficient string-commitment from weak bit-commitment and full-spectrum theorem for puzzles. Unpublished manuscript, 2009.

[Dur04] Richard Durrett. *Probability: Theorey and Examples. Third Edition.* Duxbury, 2004.

[Hai09] Iftach Haitner. A parallel repetition theorem for any interactive argument. In *FOCS*, 2009.

[HPPW08] Johan Håstad, Rafael Pass, Krzysztof Pietrzak, and Douglas Wikström. An efficient parallel repetition theorem. Unpublished manuscript, 2008.

[HS09] Thomas Holenstein and Grant Schoenebeck. General hardness amplification of predicates and puzzles. Unpublished manuscript, 2009.

[IJK07] Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Chernoff-type direct product theorems. In *CRYPTO*, pages 500–516, 2007.

[PV07] Rafael Pass and Muthuramakrishnan Venkitasubramaniam. An efficient parallel repetition theorem for arthur-merlin games. In *STOC*, pages 420–429, 2007.

[PW07] Krzysztof Pietrzak and Douglas Wikström. Parallel repetition of computationally sound protocols revisited. In *TCC*, pages 86–102, 2007.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.

[Wik09] Douglas Wikström. An efficient concurrent repetition theorem. Cryptology ePrint Archive, Report 2009/347, 2009.