Electronic Colloquium on Computational Complexity, Revision 2 of Report No. 120 (2009)

Almost Optimal Bounds for Direct Product Threshold Theorem

Charanjit S. Jutla
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598

Abstract

We consider weakly-verifiable puzzles which are challenge-response puzzles such that the responder may not be able to verify for itself whether it answered the challenge correctly. We consider k-wise direct product of such puzzles, where now the responder has to solve k puzzles chosen independently in parallel. Canetti et al have earlier shown that such direct product puzzles have a hardness which rises exponentially with k. In the threshold case addressed in Impagliazzo et al, the responder is required to answer correctly a fraction of challenges above a threshold. The bound on hardness of this threshold parallel version was shown to be similar to Chernoff bound, but the constants in the exponent are rather weak. Namely, Impagliazzo et al show that for a puzzle for which probability of failure is δ, the probability of failing on less than $(1-\gamma)\delta k$ out of k puzzles, for any parallel strategy, is at most $e^{-\gamma^2\delta k/40}$.

In this paper, we develop new techniques to bound this probability, and show that it is arbitrarily close to Chernoff bound, i.e. $e^{-\gamma^2\delta k/2}$. We show that given any responder that solves k parallel puzzles with a good threshold, there is a uniformized parallel solver who has the same threshold of solving k parallel puzzles, while being oblivious to the permutation of the puzzles. This enhances the analysis considerably, and may be of independent interest.

1 Introduction

Consider challenge-response puzzles where the responder may not be able to determine if its answer is a correct response or not, either because the challenge may have multiple correct responses (and the challenger seeks a particular one of those), or because the responder is computationally constrained, e.g. in CAPTCHA puzzles [8]. Such puzzles are called weakly-verifiable puzzles [2].

In cryptography, and other applications, the challenge-response puzzles are often used to distinguish between a real and fake responder, where the differentiation is obtained by the probability of their solving a randomly chosen challenge. For example, the authentic party may have a probability α of solving the challenge correctly, whereas non-authentic parties may have a probability only β ($< \alpha$) of

ISSN 1433-8092

solving the puzzles correctly. However, if the gap is not large then direct product, or (parallel) repetition of such puzzles may be sought. Ideally, one would like that if k puzzles are chosen independently in parallel, then the probability of the non-authentic party solving all puzzles correctly is at most β^k [2]. Unfortunately, this also makes the success probability of the authentic party go down (if $\alpha < 1$).

In [4], the authors observe that the authentic party is on average expected to solve αk puzzles, and if a Chernoff-like bound holds, then the probability of fake parties solving αk puzzles may go down exponentially. They show that their intuition is correct, and indeed give an exponential bound. However, the bound they obtain has a weak constant in the exponent. In particular they show that (setting $\delta = 1 - \beta$) the probability of the non-authentic party responding *incorrectly* to less than $(1 - \gamma)\delta k$ puzzles (out of k parallel puzzles) is at most $e^{-\gamma^2 \delta k/40}$. For real problems like CAPTCHA, the 1/40 factor in the exponent is debilitating, and the authors mention it as an open problem to improve this constant.

As is to be expected, the result in [4] is proved by reducing a single puzzle instance to a (simulated) direct product puzzle instance. However, multiple simulations are required to get a good reduction. The complication in analyzing the reduction then stems from the fact that the given single puzzle instance must be embedded in all simulated direct product puzzle instances, and hence they are not independent. In [4], the authors use a nice duality property of good (bi-partite graph based) samplers to analyze the dependent simulations.

In this paper we develop further new techniques to analyze this probability and show that one can indeed bound the probability arbitrarily close to as in Chernoff bound. In particular we upper bound the above probability by about $e^{-\gamma^2 \delta k/2}$.

We show that a uniformized parallel solver, which first permutes its given k-puzzles randomly, solves them as before, and permutes the results back, has the same probability of success as before. However, this uniformized solver is much easier to analyze. While this in itself, when plugged into the "trust reduction" strategy of [4] gives better bounds than before, to get the bounds similar to Chernoff bound we need further new techniques. In particular, while a count of other simulated puzzles being answered incorrectly gives a good guess of whether the given puzzle may be answered incorrectly, a linearly weighted metric we consider leads to more optimal bounds.

While the idea of uniformized parallel solver also applies to Raz's Theorem [7], in particular because of Holenstein's observation that the two provers can use shared randomness [3], it is to be seen if it leads to improved analysis.

The rest of the paper is organized as follows. In Section 2 we describe a result about samplers which we employ, as well as give definitions of threshold weakly-verifiable puzzles. In section 3 we consider uniformized parallel solvers and give the main technical lemmas. In section 4 we give the main theorem and its proof. In section 5 we describe the pre-processing phase.

2 Preliminaries

2.1 Basics

Lemma 1 *[Chernoff Bound [1]] Let $X = (X_1 + X_2 + ... + X_n)/n$, where the X_i are mutually independent indicator random variables, each with mean μ. Then, for $\beta \geq 0$,*

$$\begin{aligned} \Pr[X \geq (1+\beta)\mu] &< (e^{\beta}(1+\beta)^{-1-\beta})^{\mu n} \\ \Pr[X < (1-\beta)\mu] &< e^{-\beta^2 \mu n/2} \end{aligned}$$

2.2 Samplers

Consider bipartite graphs $F = G(L \cup R, E)$. We allow graphs with multiple edges. For a vertex v of G, we denote by $N_G(v)$ the multi-set of its neighbours in G. When the graph G is clear from context, we will drop the subscript G, and simply write $N(v)$. We say that G is *bi-regular* if the degree of each vertex in L is same, and the degree of each vertex in R is same.

Let $G = G(L \cup R, E)$ be any bi-regular bipartite graph. For a function $\lambda : [0,1] \times [0,1] \to [0,1]$, we say that G is a λ-*sampler* [4] if, for every function $F : L \to [0,1]$ with the average value $\mathbf{E}_{x \in L}[F(x)] \geq \mu$ and any $0 < \nu < 1$, there are at most $\lambda(\mu, \nu) \cdot |R|$ vertices $r \in R$ such that $\mathbf{E}_{y \in N(r)}[F(y)] \leq (1-\nu)\mu$.

We will employ the following lemma from [5, 4]. It says that for any two large vertex subsets W and F of a sampler, the fraction of edges between W and F is close to the product of the densities of W and F.

Lemma 2 *[5, 4] Suppose $G = G(L \cup R, E)$ is a λ-sampler. Let $W \subseteq R$ be any set of measure at least τ, and let $V \subseteq L$ be any set of measure at least β. Then, for all $0 < \nu < 1$ and $\lambda_0 = \lambda(\beta, \nu)$, we have*

$$\Pr_{x \in L, y \in N(x)}[x \in V \ \& \ y \in W] \ \geq \ \beta(1-\nu)(\tau - \lambda_0)$$

where the probability is over first picking x uniformly from L, and then picking y uniformly from $N(x)$.

We will also need the following observation from [4], which shows that the direct product is an extremely good sampler. Consider the following bipartite graph $G = G(L \cup R, L)$: the set of left vertices is the set of n-bit strings $\{0,1\}^n$; the right vertices R are a pair $\langle r, c \rangle$, where r range over all k-tuples of n-bit strings $\{0,1\}^{nk}$, and c range over m-bit strings $\{0,1\}^m$; for every $y = \langle (r_1, r_2, ..., r_k), c \rangle \in R$, there are k edges (y, r_1), (y, r_2),...,(y, r_k) in E.

Lemma 3 *[4] The graph G defined above is a λ-sampler for $\lambda(\mu, \nu) = e^{-\nu^2 \mu k/2}$.*

2.3 Weakly-Verifiable Puzzles

Definition 1. [2] A weakly-verifiable puzzle $\mathcal{P} = (C, R, d(n))$, with security parameter n, consists of a polynomial time computable function C, a polynomial time computable predicate R, and a polynomial $d(n)$. For any functions $t(n)$ and $c(n)$, the $(t(n), c(n))$-value (*failure value*) of the puzzle is

$$\mathrm{val}(\mathcal{P}, t, c) \ := \ \min_{X} \Pr_{r \in \mathcal{U}_{d(n)}, s \in \mathcal{U}_{c(n)}} [\neg R(r, X(s, C(r)))]$$

where the minimization is over $t(n)$-computable randomized algorithms X using $c(n)$ bits of randomness.

For a parameter δ, $0 \leq \delta \leq 1$, we say that a puzzle $\mathcal{P}$ is $(\delta, t(n), c(n))$-hard if the $(t(n), c(n))$-value of $\mathcal{P}$ is at least δ. In other words, every algorithm X running in time $t(n)$, and using $c(n)$ bits of randomness, has probability at least δ of answering the puzzle wrong.

Definition 2. The k-wise direct product $\mathcal{P}^k$ of a weakly-verifiable puzzle $\mathcal{P} = (C, R, d(n))$ is the weakly-verifiable puzzle $(C^k, R^k, kd(n))$, where $C^k(\langle r_1, ..., r_k \rangle)$ is defined to be $(C(r_1), ..., C(r_k))$, and

$$R^k(\langle r_1, ..., r_k \rangle, \langle x_1, ..., x_k \rangle) \ := \ \bigwedge_{i=1}^{k} R(r_i, x_i)$$

For any parameters ν and δ, $0 \leq \nu, \delta \leq 1$, and any functions $t(n), c(n)$, the puzzle $\mathcal{P}^k$ is said to be ν-approximate $(\delta, t(n), c(n))$-hard if the following minimum probability

$$\min_{X} \Pr_{\vec{r} \in \mathcal{U}^k_{d(n)}, s \in \mathcal{U}_{c(n)}} \left[\, |\{i \in [1..k] \ : \ \neg R(r_i, X_i(s, C^k(\vec{r})))\}| \ > \ \nu k \right]$$

is at least δ, where the minimization is over all randomized algorithms X running in time $t(n)$ and using $c(n)$ bits of randomness. Note that X here takes k puzzles and returns k answers, $\langle X_1, ..., X_k \rangle$. Such an algorithm X will be referred to as a **k-parallel solver**.

3 Uniformized Parallel Solvers

Given a k-parallel solver X, we consider its *uniformized* version $\overline{X}$, which first randomly permutes its given k puzzles, solves them using X, and permutes back the results. In other words, for all $i = 1..k$,

$$\overline{X}_i(\langle s, \pi \rangle, \langle C(r_1), ..., C(r_k) \rangle) \ := \ X_{\pi^{-1}(i)}(s, \langle C(r_{\pi(1)}), ..., C(r_{\pi(k)}) \rangle)$$

where π is any permutation of $[1..k]$.

It is easy to see that the "failure value" of the uniformized parallel solver remains the same, as the following shows.

$$\begin{aligned}
&\Pr_{r_1,...,r_k}\Pr_{s,\pi}\big[\,|\{i\in[1..k]\ :\ \neg R(r_i,\overline{X}_i(\langle s,\pi\rangle,\langle C(r_1),...,C(r_k)\rangle))\}|\ >\ \nu k\big]\\
&=\ \Pr_{r_1,...,r_k}\Pr_{s,\pi}\big[\,|\{i:\ \neg R(r_i,X_{\pi^{-1}(i)}(s,\langle C(r_{\pi(1)}),...,C(r_{\pi(k)})\rangle))\}|\ >\ \nu k\big]\\
&=\ \Pr_{r_1,...,r_k}\Pr_{s,\pi}\big[\,|\{j=\pi^{-1}(i):\ \neg R(r_{\pi(j)},X_j(s,\langle C(r_{\pi(1)}),...,C(r_{\pi(k)})\rangle))\}|\ >\nu k\big]\\
&=\ \Pr_{r_1,...,r_k}\Pr_{s}\big[\,|\{j:\ \neg R(r_j,X_j(s,\langle C(r_1),...,C(r_k)\rangle))\}|\ >\ \nu k\big]
\end{aligned}$$

where the last equality follows because r_1,...,r_k are chosen independently and identically. Thus, without loss of generality, we can consider only uniformized parallel solvers.

Notation.

Let us fix a parallel solver X, and its uniformized parallel solver $\overline{X}$. We will use the following shorthands to denote some useful quantities and predicates. Let $C^k(\vec{r},\pi)$ denote $\langle C(r_{\pi(1)}),...,C(r_{\pi(k)})\rangle$. Thus, $C^k(\vec{r},\mathbf{1})$ (where $\mathbf{1}$ is the identity permutation) just denotes $\langle C(r_1),...,C(r_k)\rangle$. Given the randomness $r_1,...r_k$ to generate the k puzzles, and the randomness $\langle s,\pi\rangle$ used by $\overline{X}$, **define random variables**

- total($\overline{X}$) := $|\{i\in[1..k]\ :\ \neg R(r_i,\overline{X}_i(\langle s,\pi\rangle,C^k(\vec{r},\mathbf{1})))\}|$
- F($\overline{X}$) (short for first) := $\neg R(r_1,\overline{X}_1(\langle s,\pi\rangle,C^k(\vec{r},\mathbf{1})))$
- others($\overline{X}$) := $|\{i\in[2..k]\ :\ \neg R(r_i,\overline{X}_i(\langle s,\pi\rangle,C^k(\vec{r},\mathbf{1})))\}|$
- others($\overline{X},j$) := $|\{i\in[1..j-1,j+1..k]\ :\ \neg R(r_i,\overline{X}_i(\langle s,\pi\rangle,C^k(\vec{r},\mathbf{1})))\}|$
- for $\Gamma\subseteq[1..k]$, (failure-) pattern($\overline{X}$, Γ) denotes

$$\bigwedge_{i\in\Gamma}R\big(r_i,\overline{X}_i(\langle s,\pi\rangle,C^k(\vec{r},\mathbf{1}))\big)\ \wedge\ \bigwedge_{i\notin\Gamma}\neg R\big(r_i,\overline{X}_i(\langle s,\pi\rangle,C^k(\vec{r},\mathbf{1}))\big)$$

From now on, unless otherwise stated, all probabilities will be over $r_1,...,r_k$ each chosen uniformly and independently from $\mathcal{U}_{d(n)}$, s chosen uniformly (and independently) from $\mathcal{U}_{c(n)}$, and π chosen uniformly (and independently) from all permutations of $[1..k]$. Further, define

- For any t, $0\le t\le k$, let p_t denote $\Pr[\text{total}(\overline{X})=t]$.
- Let $\tau=(1-\gamma)\delta k$.
- Define $P=\sum_{t\le\tau}p_t$.
- In the following, we use a positive integer d which we will fix later. This d is clearly unrelated to the polynomial $d(n)$ above in the definition of weakly-verifiable puzzles.

- For $j \geq 0$, let $\psi_j = \gamma\delta(1-\gamma^d) + j \cdot (\gamma^d/k)$.

Lemma 4 *For any integer t, $0 \leq t \leq k$,*

$$\Pr[F(\overline{X}) \mid total(\overline{X}) = t] \;=\; \frac{t}{k}$$

Proof: We first show that for any t, $0 \leq t \leq k$, and any subset Γ of $[1..k]$ of size t, the probability of pattern($\overline{X}$, Γ) is a function only of t, and is independent of the subset Γ.

Indeed, consider Γ, and another subset Γ' of size t, and let σ be any permutation of $[1..k]$, such that $\Gamma' = \sigma(\Gamma)$ (a permutation applied to a subset Γ just yields the set which is the range of the permutation with domain Γ). It is clear that such a permutation exists. Then,

$$\begin{aligned}
&\Pr[\text{pattern}(\overline{X},\, \Gamma')] \\
&= \Pr\Big[\bigwedge_{i\in\Gamma'} R\big(r_i, X_{\pi^{-1}(i)}(s, C^k(\vec{r},\pi))\big) \bigwedge_{i\notin\Gamma'} \neg R\big(r_i, X_{\pi^{-1}(i)}(s, C^k(\vec{r},\pi))\big)\Big] \\
&= \Pr\Big[\bigwedge_{i\in\sigma(\Gamma)} R\big(r_i, X_{\pi^{-1}(i)}(s, C^k(\vec{r},\pi))\big) \bigwedge_{i\notin\sigma(\Gamma)} \neg R\big(r_i, X_{\pi^{-1}(i)}(s, C^k(\vec{r},\pi))\big)\Big] \\
&= \Pr\Big[\bigwedge_{j\in\Gamma} R\big(r_{\sigma(j)}, X_{\pi^{-1}(\sigma(j))}(s, C^k(\vec{r},\pi))\big) \;\wedge \\
&\qquad \bigwedge_{j\notin\Gamma} \neg R\big(r_{\sigma(j)}, X_{\pi^{-1}(\sigma(j))}(s, C^k(\vec{r},\pi))\big)\Big]
\end{aligned}$$

Now, denote $r_{\sigma(j)}$ by w_j. Then, the above becomes (with probability now over $w_{\sigma^{-1}(1)}$,..., $w_{\sigma^{-1}(k)}, s, \pi$)

$$\begin{aligned}
\Pr\Big[&\bigwedge_{j\in\Gamma} R\big(w_j, X_{\pi^{-1}(\sigma(j))}(s, C^k(\vec{w},\sigma^{-1}\pi))\big) \wedge \\
&\bigwedge_{j\notin\Gamma} \neg R\big(w_j, X_{\pi^{-1}(\sigma(j))}(s, C^k(\vec{w},\sigma^{-1}\pi))\big)\Big]
\end{aligned}$$

Now, $\pi^{-1}\sigma = (\sigma^{-1}\pi)^{-1}$. Denote $\sigma^{-1}\pi$ by $\hat{\pi}$. Since permutations form a group, $\hat{\pi}$ is independent of σ, with π chosen uniformly and independently of σ. Then, the above probability can be written as (with probability now over $w_{\sigma^{-1}(1)}$, ...,$w_{\sigma^{-1}(k)}$, s, $\hat{\pi}$)

$$\Pr\Big[\bigwedge_{j\in\Gamma} R\big(w_j, X_{\hat{\pi}^{-1}(j)}(s, C^k(\vec{w},\hat{\pi}))\big) \bigwedge_{j\notin\Gamma} \neg R\big(w_j, X_{\hat{\pi}^{-1}((j)}(s, C^k(\vec{w},\hat{\pi}))\big)\Big]$$

Since, $w_1, ..., w_k$ are chosen identically and independently, the above remains same even when the probability is considered over $w_1, ...w_k, s, \hat{\pi}$. This proves that the

above probability is a function only of t, and independent of the particular subset Γ. Now,

$$\begin{aligned}
&\Pr[\mathrm{F}(\overline{X}) \mid \mathrm{total}(\overline{X}) = t] \\
&= \sum_{\Gamma:|\Gamma|=t} \Pr[\mathrm{F}(\overline{X}) \wedge \mathrm{pattern}(\overline{X}, \Gamma) \mid \mathrm{total}(\overline{X}) = t] \\
&= \sum_{\Gamma:|\Gamma|=t, 1\in\Gamma} \Pr[\mathrm{pattern}(\overline{X}, \Gamma) \mid \mathrm{total}(\overline{X}) = t] \\
&= \frac{\sum_{\Gamma:|\Gamma|=t,1\in\Gamma} \Pr[\mathrm{pattern}(\overline{X}, \Gamma)]}{\Pr[\mathrm{total}(\overline{X}) = t]} \\
&= \frac{\sum_{\Gamma:|\Gamma|=t,1\in\Gamma} \Pr[\mathrm{pattern}(\overline{X}, \Gamma)]}{\sum_{\Gamma:|\Gamma|=t} \Pr[\mathrm{pattern}(\overline{X}, \Gamma)]} \\
&= \binom{k-1}{t-1} / \binom{k}{t} \\
&= t/k
\end{aligned}$$

□

Lemma 5 *For $t < k$,*

$$\Pr[F(\overline{X}) \mid others(\overline{X}) \ \leq t] \ = \ \frac{\sum_{t' \leq t+1} (t'/k) p_{t'}}{\sum_{t' \leq t} p_{t'} \ + \ ((t+1)/k) p_{t+1}}$$

Proof: First note that, for $t < k$,

$$\begin{aligned}
&\Pr[\mathrm{others}(\overline{X}) \ = t] \\
&= \Pr[\mathrm{F}(\overline{X}) \wedge \mathrm{others}(\overline{X}) \ = t] + \Pr[\neg\mathrm{F}(\overline{X}) \wedge \mathrm{others}(\overline{X}) \ = t] \\
&= \Pr[\mathrm{F}(\overline{X}) \wedge \mathrm{total}(\overline{X}) \ = t+1] + \Pr[\neg\mathrm{F}(\overline{X}) \wedge \mathrm{total}(\overline{X}) \ = t] \\
&= \frac{t+1}{k} p_{t+1} + \frac{k-t}{k} p_t \quad \text{(by Lemma 4)}
\end{aligned}$$

The lemma follows easily from this observation. □

Lemma 6 *For any $i > 0$, suppose for all j, $0 \leq j < i$*

$$\Pr[F(\overline{X}) \mid others(\overline{X}) \ \leq \tau + j] \ > \ \frac{\tau}{k} + \psi_j$$

then

$$p_{\tau+i} \ > \ \psi_0 P \cdot \frac{k}{\tau+i} \cdot \prod_{0<j<i} \left(1 + (\psi_j - \frac{j}{k})(\frac{k}{\tau+j})\right) \tag{1}$$

Proof: For any j, $j < i$, we first note that Lemma 5, along with the hypothesis of the lemma for j, yields (by simple manipulation)

$$p_{\tau+j+1} \;>\; \frac{k}{\tau+j+1} \cdot \big(\psi_j P + \sum_{0<j'\leq j} p_{\tau+j'}(\psi_j - \frac{j'}{k})\big) \tag{2}$$

The base case, i.e. $i = 1$, follows immediately from this by considering $j = 0$. Now suppose the induction hypothesis holds for i, and we will prove the lemma for $i+1$. The antecedent for $i+1$ completely yields the antecedent for $i' < i+1$. Thus, inequality (1) holds for all such i'.

Let $\Psi(j)$ stand for $(\psi_j - \frac{j}{k})(\frac{k}{\tau+j})$.

Then by inequality (2), and plugging in inequality (1) for each $p_{\tau+j}$ $(j < i+1)$, while noting that ψ_j is an increasing function of j, we get that $p_{\tau+i+1}$ is greater than

$$\begin{aligned}
& \frac{\psi_0 Pk}{\tau+i+1} \cdot \Big(1 + \sum_{0<j\leq i} \Psi(j) \prod_{0<j'<j} \big(1+\Psi(j')\big)\Big) \\
= \;& \frac{\psi_0 Pk}{\tau+i+1} \cdot \Big(1 + \sum_{0<j\leq i} \big(1+\Psi(j)-1\big) \prod_{0<j'<j} \big(1+\Psi(j')\big)\Big) \\
= \;& \frac{\psi_0 Pk}{\tau+i+1} \cdot \Big(1 + \sum_{0<j\leq i} \prod_{0<j'\leq j} \big(1+\Psi(j')\big) \;-\; \sum_{0<j\leq i} \prod_{0<j'<j} \big(1+\Psi(j')\big)\Big) \\
= \;& \frac{\psi_0 Pk}{\tau+i+1} \cdot \Big(1 + \prod_{0<j'\leq i} \big(1+\Psi(j')\big) - \prod_{0<j'<1} \big(1+\Psi(j')\big)\Big) \\
= \;& \frac{\psi_0 Pk}{\tau+i+1} \cdot \prod_{0<j'<i+1} \big(1+\Psi(j')\big)
\end{aligned}$$

□

Lemma 7 *For $0 < \gamma < 1$, and for any positive integer $t < \gamma\delta k \,(= M)$,*

$$\prod_{j=1}^{t} \frac{\tau+\psi_j k}{\tau+j} \;\geq\; e^{\gamma(1-\gamma^d)(1-\frac{t+1}{2M})t}$$

Proof: From the definition of ψ_j, the product can be written as

$$\prod_{j=1}^{t} \frac{\delta k - \gamma^{d+1}\delta k + \gamma^d j}{\delta k - \gamma\delta k + j} \;=\; \prod_{j=1}^{t} \frac{1-\gamma^{d+1}(1-j/M)}{1-\gamma(1-j/M)}$$

Let θ_j denote $(1-j/M)$. It is an easy exercise to show that $e^{-z} \geq 1-z$, for $0 \leq z \leq 1$. Since $0 < \gamma < 1$, and $d > 0$, as well as $0 < \theta_j < 1$ for each $j < M$, it follows that

$$\frac{1-\gamma^{d+1}\theta_j}{1-\gamma\theta_j} \geq e^{(\gamma-\gamma^{d+1})\theta_j}.$$

The lemma follows easily from this. □

Finally, we need the following simple calculation. Define

$$q_{\tau+i} = \psi_0 \cdot \frac{k}{\tau+i} \cdot \prod_{j=1}^{i-1} \frac{\tau+\psi_j k}{\tau+j}$$

Lemma 8 *Let $M = \lceil \gamma\delta k \rceil$, and suppose $\delta k \geq 1$.*

1. *For any i, $0 \leq i < M$, and for any $\chi \geq 1$,*

$$q_{\tau+i} \cdot \chi \cdot \frac{4}{\gamma(1-\gamma^d)} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2} - e^{-(1-\gamma^d)(\delta-\tau/k-\psi_i)k/2} \;>\; \frac{\chi}{e} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2 + \gamma(1-\gamma^d)i/2}$$

2. $q_{\tau+M} \cdot \frac{4}{\gamma(1-\gamma^d)} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2} \;>\; 1$

The detailed calculations can be found in Appendix A.

4 The Main Theorem

Theorem 9 *Let $\mathcal{P} = (C, R, d(n))$ be a weakly-verifiable puzzle that is (δ, $t(n)$, $c(n)$)-hard. Let k be any positive integer such that $\delta k \geq 1$, and γ ($1 > \gamma > 0$) be arbitrary. Further, let ϵ_0 be any arbitrary positive real, d be an arbitrary positive integer, and let*

$$\epsilon \geq \frac{8}{\gamma} \cdot e^{(-\gamma^2\delta+\epsilon_0)k/2}.$$

Then the direct product puzzle $\mathcal{P}^k$ is $(1-\gamma)\delta$-approximate $((1-\epsilon), t'(n), c'(n))$-hard with $t'(n) = t(n) \cdot \epsilon/(\gamma\delta k \ln(1/\epsilon_0)) \cdot poly(1/n)$, and $c'(n) = c(n) \cdot \epsilon/(\gamma\delta k \ln(1/\epsilon_0)) \cdot poly(1/n)$.

In the following let $\epsilon_1 = \epsilon_2 = \epsilon_3 = \epsilon_0/6$. Recall the definitions of τ, P, and ψ_j from Section 3. Let $d > 0$ be such that $\gamma^{d+1} < \min(1/10, 1/(\gamma\delta k))$.

Consider, for contradiction sake, a k-parallel solver X which for the k-wise direct product $\mathcal{P}^k$ has $(1-\gamma)\delta$-approximate (failure) value less than $1-\epsilon$, i.e

$$P \;=\; \Pr_{\vec{r}\in\mathcal{U}^k_{d(n)}, s\in\mathcal{U}_{c(n)}} \left[\, |\{i \in [1..k] \;:\; \neg R(r_i, X_i(s, C^k(\vec{r})))\}| \;\leq\; \tau \right] \;>\; \epsilon$$

As explained earlier in Section 3, we consider its uniformized version $\overline{X}$, which has the same failure value $(1-P)$. Using $\overline{X}$ as an oracle, we will give an algorithm Y to solve the underlying puzzle $\mathcal{P}$ with failure value less than δ, leading to a contradiction.

The algorithm Y will have a pre-processing phase (i.e. independent of the given target puzzle instance x, and function of security parameter n), where it runs some statistical tests using X to determine the appropriate algorithm $\mathcal{C}[i]$ to run, where $\mathcal{C}[0], ..., \mathcal{C}[M-1]$ are M $(= \lceil \gamma\delta k \rceil)$ algorithms as follows:

$\mathcal{C}[i]$: On input x, run $\mathcal{C}'[i]$ below on x. If the value returned is different from $\perp$, then return that value; otherwise repeat by calling $\mathcal{C}'[i]$ on x again, for a total of at most T iterations ($T = \frac{8e}{\epsilon\gamma^{d+1}(1-\gamma)} \ln(1/\epsilon_1)$).

If no output is produced in these T iterations, return $\perp$.

$\mathcal{C}'[i]$: On input x, choose $k-1$ random tapes $\alpha_2, ..., \alpha_k$ uniformly and independently from $\{0,1\}^{d(n)}$. Let $x_2, ..., x_k$ be the corresponding puzzles, i.e. $x_l = C(\alpha_l)$, for $l = 2..k$. Set $\bar{x} = \langle x, x_2, ..., x_k \rangle$. Run $\overline{X}$ on $\bar{x}$. Check if **others** $\leq \tau + i$, and if so return $\overline{X}_1(\bar{x})$; otherwise return $\perp$.

The pre-processing phase η returns $\eta(\overline{X}, n, \delta, \gamma, k)$, a value between 0 and $M-1$. When it is clear from context, we just call the value η. Thus, $0 \leq \eta \leq M-1$. As mentioned above, Y runs $\mathcal{C}[\eta]$ on x.

The event **valid** stands for the following being satisfied by the returned η:

1. For all $i < \eta$, $\Pr[\mathrm{F}(\overline{X}) \mid \text{others}(\overline{X}) \ \leq \tau + i] \ > \ \frac{\tau}{k} + \psi_i$, and
2. $\Pr[\mathrm{F}(\overline{X}) \mid \text{others}(\overline{X}) \ \leq \tau + \eta] \ \leq \ \frac{\tau}{k} + \psi_\eta + \epsilon_2$.

We will later bound the probability of *valid* not happening by ϵ_3 (lemma 12); i.e. after we describe how the pre-processing works. In rest of this section, we condition on the event *valid* being true, and we will not mention it explicitly in the probabilities.

We first need to bound the probability of $\mathcal{C}[\eta]$ timing out, i.e. returning $\perp$. Note that, $\mathcal{C}'[\eta]$ returns something other than $\perp$ if (others $\leq \tau + \eta$). As in Lemma 5, it is easy to see that the probability of this happening is at least P which is at least ϵ (by hypothesis of the theorem). However, multiple calls to $\mathcal{C}'[\eta]$ are *not* independent, as they all include the query x. However, as shown in [4], the corresponding graph is a good sampler, and that helps us analyze the probability of $\mathcal{C}[\eta]$ timing out. Of course, we require Lemma 6, and the idea therein of a linearly increasing ψ_j, to obtain better bounds.

To this end, we consider a (k-colored) bipartite graph $G = G(L \cup R, E)$; the set of left vertices is the set of $d(n)$-bit strings $\{0,1\}^{d(n)}$; the right vertices are triples $\langle \bar{\alpha}, s, \pi \rangle$, where $\bar{\alpha}$ ranges over all k-tuples of $d(n)$-bit strings, and s ranges over $c(n)$ bit strings, and π ranges over permutations of $[k]$; for every $y = \langle (\alpha_1, ..., \alpha_k), s, \pi \rangle \in R$ there are k edges (y, α_1), ..., (y, α_k) in E, *colored* 1..k respectively.

By lemma 3, this graph is a λ-sampler for $\lambda(\mu, \nu) = e^{-\nu^2 \mu k/2}$.

Corresponding to each $(\alpha_1, ..., \alpha_k)$ are puzzles $(x_1, ..., x_k)$. Now, *define* Good_η to be the subset of R (the right vertices) such that $\overline{X}$ when run on input $(x_1, ..., x_k)$, with randomness s and π, has the following property

$$|\{i \in [1..k] \ : \ \neg R(\alpha_i, \overline{X}_i(\langle s, \pi \rangle, \langle x_1, ..., x_k \rangle))\}| \ \leq \ \tau + \eta$$

In other words, $\text{total}(\overline{X}) \leq \tau + \eta$. Let the density of Good_η in R be g_η. We now *define* $H_\eta \subseteq L$ to be all those vertices α such that α has less than $(\epsilon \cdot \frac{\gamma^{d+1}(1-\gamma)}{8e})$

fraction of its neighbours in the set Good_η. We will later see in Lemma 11 how H_η is relevant, even though $\mathcal{C}'[\eta]$ embeds α (or it's x) only in the first position. We can bound the size of H_η, just as in [4], by employing Lemma 2.

Lemma 10 *H_η has density at most $\delta - \tau/k - \psi_\eta - \epsilon_0$.*

Proof: Suppose to the contrary, the density of H_η is greater than $\beta = \delta - \tau/k - \psi_\eta - \epsilon_0$. Let $H' \subseteq H_\eta$ be any subset of density exactly β. Now, by definition of H_η, we have $\Pr_{\alpha \in L, w \in N(\alpha)}[\alpha \in H' \,\&\, w \in \text{Good}_\eta] < \beta\epsilon\gamma^{d+1}(1-\gamma)/8e$. On the other hand, by Lemma 2, we get that the same probability is at least $\beta(g_\eta - \lambda_0)(1-\bar{\nu})$ for $\lambda_0 = \lambda(\beta, \bar{\nu})$, for any $0 \leq \bar{\nu} \leq 1$. We set $\bar{\nu} = \sqrt{1-\gamma^d}$.

Now, note that $g_\eta = \Pr[\text{total}(\overline{X}) \leq \tau + \eta]$. If $\eta = 0$, then $g_\eta = P > \epsilon$. Otherwise, since event *valid* is true, we can use Lemma 6 and 7 to get an explicit lower bound for $p_{\tau+\eta}$, and hence for g_η.

Then, using Lemma 8.1, and noting that $1 - \bar{\nu} > \gamma^d/2$, it can be seen by a simple calculation that $\beta(g_\eta - \lambda_0)(1-\bar{\nu})$ is more than $\beta\epsilon\gamma^{d+1}(1-\gamma)/8e$, a contradiction. □

Lemma 11 *For every $\alpha \notin H_\eta$ and the puzzle x corresponding to that random α, we have $\Pr[\mathcal{C}[\eta](x) = \perp] \leq \epsilon_1$, where the probability is over the random coins of $\mathcal{C}[\eta]$ (including those of $\overline{X}$ and X).*

Proof: We consider a variation of $\mathcal{C}[i]$, where instead of calling $C'[i]$, it calls the following $\mathcal{C}''[i]$ instead.

$\mathcal{C}''[i]$: On input x, choose $k-1$ random tapes $\alpha_1, ..., \alpha_{k-1}$ uniformly and independently from $\{0,1\}^{d(n)}$. Let $x_1, ..., x_{k-1}$ be the corresponding puzzles, i.e. $x_l = C(\alpha_l)$, for $l = 1..k-1$. Pick $j \in [1..k]$ at random and set $\bar{x} = \langle x_1, ..., x_{j-1}, x, x_j, ..., x_{k-1} \rangle$. Run $\overline{X}$ on $\bar{x}$. Check if $\textbf{others}(\overline{X}, j) \leq \tau + i$, and if so return $\overline{X}_j(\bar{x})$; otherwise return $\perp$.

For each fixed α, the behaviour of $\mathcal{C}'[i]$ and $\mathcal{C}''[i]$ is statistically identical, because placing x in the random j-th place is just a permutation of placing x in the first place, and that the permutations form a group.

Further, picking a color $j \in [1..k]$ at random, and then picking $\alpha_1, ..., \alpha_{k-1}$ at random and placing α in the j-th place to form $\bar{\alpha}$ is same as picking a random neighbour of α (random element of $N_G(\alpha)$, and note that $N_G(\alpha)$ is defined to be a multi-set)[1].

Now, $\mathcal{C}''[\eta]$ returns something other than $\perp$ if the neighbour satisfies **others** $(\overline{X}, j) \leq \tau + \eta$, which is implied by total $\leq \tau + \eta$. But, for $\alpha \notin H_\eta$, the density of neighbours satisfying total $\leq \tau + \eta$ is more than $\epsilon\gamma^{d+1}(1-\gamma)/8e$. Hence for such α, the probability of $\mathcal{C}''[\eta]$ returning something other than $\perp$ is more than $\epsilon\gamma^{d+1}(1-\gamma)/8e$.

[1]This follows formally by noting that $\sum_{j=1}^{k} j\binom{k}{j}(A-1)^{k-j} = kA^{k-1}$, for any A

But, the probability of $\mathcal{C}[\eta](x)$, using $\mathcal{C}'$, returning $\perp$ is same as probability of $\mathcal{C}$, using $\mathcal{C}''$, returning $\perp$, which is at most $(1-(\epsilon \cdot \frac{\gamma^{d+1}(1-\gamma)}{8e}))^T \leq \epsilon_1$. $\square$

Proof of Main Theorem: Now we are ready to prove the main theorem. Since there are a potential T attempts by $\mathcal{C}[\eta]$ on x, we call the values returned in the q-th attempt by $\mathcal{C}'[\eta]^q$ $(1 \leq q \leq T)$. Now, the input x was set by choosing α uniformly from $\{0,1\}^{d(n)}$. Thus,

$$\Pr_{\alpha}[\,\mathcal{C}[\eta](x) \text{ is wrong}] \leq \Pr_{\alpha}[\,\alpha \in H_\eta] + \Pr_{\alpha}[\,\mathcal{C}[\eta](x) \text{ is wrong } \&\ \alpha \notin H_\eta]$$

The first term on the right-hand side is at most $\delta - \tau/k - \psi_\eta - \epsilon_0$ by Lemma 10. We now focus on the second term.

$$\begin{aligned}
&\Pr[\,\mathcal{C}[\eta](x) \text{ is wrong } \&\ \alpha \notin H_\eta] \\
&\leq \Pr[\,\mathcal{C}[\eta](x) = \perp \,\&\, \alpha \notin H_\eta] + \Pr[\,\mathcal{C}[\eta](x) \text{ is wrong } \&\ \mathcal{C}[\eta](x) \neq \perp \,\&\, \alpha \notin H_\eta] \\
&\leq \epsilon_1 + \Pr[\,\mathcal{C}[\eta](x) \text{ is wrong } \&\ \mathcal{C}[\eta](x) \neq \perp] \qquad \text{(by Lemma 11)} \\
&\leq \epsilon_1 + \Pr[\,\mathcal{C}[\eta](x) \text{ is wrong} \mid \mathcal{C}[\eta](x) \neq \perp] \\
&= \epsilon_1 + \Pr[\,\mathcal{C}'[\eta]^q(x) \text{ is wrong} \mid \exists q : \mathcal{C}'[\eta]^q(x) \neq \perp] \\
&= \epsilon_1 + \Pr[F(\overline{X}) \mid others(\overline{X}) \leq \tau + \eta] \\
&\leq \epsilon_1 + \frac{\tau}{k} + \psi_\eta + \epsilon_2 \qquad \text{(by event } \textbf{valid}\text{)}
\end{aligned}$$

Thus,

$$\Pr_{\alpha}[\,\mathcal{C}[\eta](x) \text{ is wrong}] = \delta + \epsilon_1 - \epsilon_0 + \epsilon_2.$$

Finally, by Lemma 12 (of the following section), the probability of η being not valid is at most ϵ_3, and this leads to a contradiction as $\epsilon_0 > \epsilon_1 + \epsilon_2 + \epsilon_3$. $\square$

5 Pre-Processing and Hypothesis Testing

As mentioned in Section 4, the algorithm Y first does some pre-processing using algorithm η, and using $\overline{X}$ as an oracle. The inputs to η are the security parameter n, k, as well as γ and δ. It returns a value η, $0 \leq \eta \leq M-1$ $(M = \lceil \gamma\delta k \rceil)$. As before, let d be such that $\gamma^d < 1(\gamma\delta k)$.

Before we describe this pre-processing algorithm, we remark that it is intended to compute the smallest $j < M$, such that $\Pr[\mathrm{F}(\overline{X}) \mid \text{others}(\overline{X}) \ \leq \tau + j] \leq \frac{\tau}{k} + \psi_j$. Now, we had assumed that $P > \epsilon$, and the hypothesis of Theorem 9 assumes a lower bound on ϵ. Then, by Lemmas 6, 7 and 8.2, it follows that if such a j does not exist, then $p_{\tau+M} > 1$, an impossibility. So, let $\bar{\eta}$ be that smallest $0 \leq j < M$.

The algorithm $\eta(\overline{X}, n, \delta, \gamma, k)$ does the following:

η : For each $i = 1..M-1$, compute the following statistics

$$t_i = \frac{\#(F(\overline{X}) \,\&\, \text{others}(\overline{X}) \ \leq \tau + i)}{1 + \#(\text{others}(\overline{X}) \ \leq \tau + i)}$$

where the count is over running $\overline{X}$ on random and independent $\bar{\alpha}$ (each in $\{0,1\}^{d(n)k}$), for a total of N times (N to be determined below). Set η to be the smallest i such that $t_i < \tau/k + \psi_i + \epsilon_2/2$. If no such η exists then set $\eta = M - 1$. Return η.

Lemma 12 *There is a polynomial ϕ, independent of n, such that with $N = \phi(\gamma\delta k, \ln(1/\epsilon_2), \ln(1/\epsilon_3))$,*

$$\Pr[\textit{ not } \textbf{valid }] < \epsilon_3$$

Proof: Clearly, for $i = \bar{\eta}$, the actual conditional probability of F is is no more than $\tau/k + \psi_i$. Hence, $t_i > \tau/k + \psi_i + \epsilon_2/2$ is an exponentially low probability event by Chernoff bound. Now, for some smaller i, if conditional probability of F is greater than $\tau/k + \psi_i + \epsilon_2$, then again t_i being less than $\tau/k + \psi_i + \epsilon_2/2$ is an exponentially low probability event. □

References

[1] N. Alon, J. Spencer, "The Probabilistic Method", John Wiley and Sons, 1992.

[2] R. Canetti, S. Halevi, M. Steiner, "Hardness Amplification of Weakly-Verifiable Puzzles", Proc. TCC 2005, pp 17-33.

[3] T. Holenstein, "Parallel Repetition: Simplifications and the No-Signalling Case", Proc. ACM STOC 2007.

[4] R. Impagliazzo, R. Jaiswal, V. Kabanets, "Chernoff-Type Direct Product Theorem", J. Cryptology, (2009) 22:75-93.

[5] R. Impagliazzo, R. Jaiswal, V. Kabanets, A. Wigderson, "Uniform direct-product theorems: Simplified, optimized, and de-randomized", Proc. ACM STOC 2008.

[6] D. Knuth , "Art of Computer Programming, Vol 1.", Addison Wesley 1973.

[7] R. Raz, "A Parallel Repetition Theorem". SIAM J. of Computing, 27(3):763-803.

[8] L. von Ahn, M. Blum, N.J. Hopper, J. Langford, " CAPTCHA: Using hard AI problems for security", in Proc. Eurocrypt 2003.

Appendix A.

Recall,

$$q_{\tau+i} = \psi_0 \cdot \frac{k}{\tau + i} \cdot \prod_{j=1}^{i-1} \frac{\tau + \psi_j k}{\tau + j}$$

Lemma 8: Let $M = \lceil \gamma\delta k \rceil$, and suppose $\delta k \geq 1$.

1. For any i, $0 \leq i < M$, and for any $\chi \geq 1$,

$$q_{\tau+i} \cdot \chi \cdot \frac{4}{\gamma(1-\gamma^d)} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2} - e^{-(1-\gamma^d)(\delta-\tau/k-\psi_i)k/2} > \frac{\chi}{e} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2+\gamma(1-\gamma^d)i/2}$$

2. $q_{\tau+M} \cdot \frac{4}{\gamma(1-\gamma^d)} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2} > 1$

Proof: For the first item in the lemma, we have

$$\begin{aligned} \delta - \tau/k - \psi_i &= \delta - (1-\gamma)\delta - \gamma(1-\gamma)\delta - i \cdot (\gamma/k) \\ &= \gamma^2\delta - \gamma \cdot (i/k) \end{aligned}$$

Now, by Lemma 7,

$$q_{\tau+i} > \frac{\gamma\delta(1-\gamma^d)k}{\tau+i} \cdot e^{\gamma(1-\gamma^d)(1-\frac{i}{2M})(i-1)}$$

But, $\tau + i < \delta k$, and further, $e^{-\gamma(1-\gamma^d)} > 1/e$. Thus,

$$\begin{aligned} & q_{\tau+i} \cdot \chi \cdot \frac{4}{\gamma(1-\gamma^d)} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2} - e^{-(1-\gamma^d)(\delta-\tau/k-\psi_i)k/2} \\ > \; & \frac{4}{e} \cdot \chi \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2+\gamma(1-\gamma^d)(1-\frac{i}{2M})i} - e^{\gamma(1-\gamma^d)(i/2)-\gamma^2(1-\gamma^d)\delta k/2} \\ = \; & e^{-\gamma^2(1-\gamma^d)\delta k/2+\gamma(1-\gamma^d)(i/2)} \cdot (\frac{4}{e} \cdot \chi \cdot e^{\gamma(1-\gamma^d)(\frac{1}{2}-\frac{i}{2M})i} - 1) \\ > \; & e^{-\gamma^2(1-\gamma^d)\delta k/2} \cdot \frac{\chi}{e} \cdot e^{\gamma(1-\gamma^d)i/2} \end{aligned}$$

For the second item in the lemma, again using Lemma 7 we have

$$\begin{aligned} & q_{\tau+M} \cdot \frac{4}{\gamma} \cdot e^{-\gamma^2(1-\gamma^d)\delta k/2} \\ > \; & \frac{4}{e} \cdot e^{-\gamma^2(1-\gamma)\delta k/2+\gamma(1-\gamma^d)(1-\frac{M-1}{2M})M} \\ > \; & \frac{4}{e} \cdot e^{-\gamma^2(1-\gamma)\delta k/2+\gamma(1-\gamma^d)\gamma\delta k/2} \end{aligned}$$

□

ECCC ISSN 1433-8092

http://eccc.hpi-web.de