

Direct Product Hardness Amplification*

David Lanzenberger and Ueli Maurer

Department of Computer Science
ETH Zurich
8092 Zurich, Switzerland
{landavid,maurer}@inf.ethz.ch

Abstract. We revisit one of the most fundamental hardness amplification constructions, originally proposed by Yao (FOCS 1982). We present a hardness amplification theorem for the direct product of certain games that is simpler, more general, and stronger than previously known hardness amplification theorems of the same kind. Our focus is two-fold. First, we aim to provide close-to-optimal concrete bounds, as opposed to asymptotic ones. Second, in the spirit of abstraction and reusability, our goal is to capture the essence of direct product hardness amplification as generally as possible. Furthermore, we demonstrate how our amplification theorem can be applied to obtain hardness amplification results for non-trivial interactive cryptographic games such as MAC forgery or signature forgery games.

1 Introduction

1.1 Security Amplification

Security amplification is a central theme of cryptography. Turning weak objects into strong objects is useful as it allows to weaken the required assumptions.

Almost all cryptographic constructions rely on the hardness of a certain problem, often modeled as games. As such, hardness amplification is of fundamental importance. Direct product theorems are one of the most natural and intuitive ways to amplify hardness: If a game can be won with probability at most δ , one would expect that n parallel instances of the game can be won with probability at most δ^n . While intuitive and usually trivial in an information-theoretic setting, these results are often surprisingly difficult to establish in a typical computational setting.

The main challenge of computational direct product hardness amplification statements is that they are essentially always based on a reduction, trying to turn *any* winner (or solver) for the direct product with some small winning probability into a winner for a single instance with much larger winning probability. Even though the instances of the direct product are all independent, a winner is not restricted to solving these instances independently. The main difficulty is usually to work around such potential dependencies.

* A preliminary version of this paper appears in the proceedings of TCC 2021. This is the full version.

1.2 Hardness of the Direct Product of two Games

Consider two probabilistic games G and H , i.e., probability distributions over deterministic instances of games from sets \mathcal{G} and \mathcal{H} , respectively. Let $[G, H]^\wedge$ denote the independent parallel composition of the two games that is won exactly if *both* G and H are won. Consider a winner (or player) W winning the two games $[G, H]^\wedge$ in parallel with probability δ .

Intuitively, we would like to argue that if W wins $[G, H]^\wedge$ with probability δ , then we can (by a simple reduction) use W to win at least one of the games G or H with much higher probability, say, $\sqrt{\delta}$. Note that this is trivially possible if W played both games G and H completely independently.

If W was known to play deterministically¹, certain instances $\mathcal{S} \subseteq \mathcal{G} \times \mathcal{H}$ are (always) solved successfully, while none of the other instances $\bar{\mathcal{S}}$ are ever solved. How does the set \mathcal{S} look like? Suppose that $\mathcal{S} = \mathcal{S}_G \times \mathcal{S}_H$ for some $\mathcal{S}_G \subseteq \mathcal{G}$ and $\mathcal{S}_H \subseteq \mathcal{H}$. This would, for example, be the case if W solved both given instances independently. Visually, this means that \mathcal{S} forms a rectangle as depicted in Figure 1.

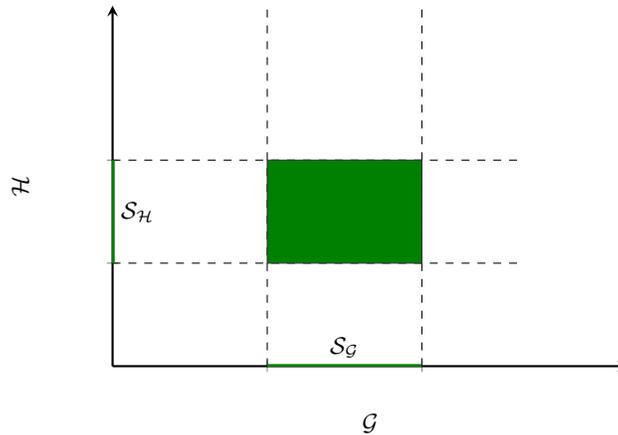


Fig. 1. The considered winner W is deterministic and wins exactly the instances $\mathcal{S}_G \times \mathcal{S}_H$ (marked in green) of the game $[G, H]^\wedge$.

If we want to use W to win, say, the game G , we need to simulate an instance of H towards W . In general, the only easy way to do this is simply by sampling an independent instance from the distribution of H , resulting in a winner we denote by $W_{(\cdot, H)}$. However, it is easy to see (in our example) that it is necessary that we hit into the set \mathcal{S}_H in order to win G . This means that the winner $W_{(\cdot, H)}$ for G might have the exact same winning probability that W has for $[G, H]^\wedge$.

¹ Of course, this is not without loss of generality.

For many types of games² such as one-way function inversion or collision-finding for hash functions, we can overcome this problem by repeating the given winner, such that we are overall successful if only a single one of our attempts has been successful. It is important to note that this property for itself does *not* allow to amplify the winning probability of *any* winner. In particular, if the winning probability on any instance is always either 0 or 1 (i.e., never in-between), no amplification is achieved. However, one can argue that in the given scenario, at least one of $W_{(\cdot, H)}$ or $W_{(G, \cdot)}$ must be amplifiable to a certain degree. A typical analysis such as in [Gol01, Mau17] would achieve this as follows.

1. First, it is proved³ that if W wins $[G, H]^\wedge$ with probability δ , the following is true for any $\epsilon > 0$:

With probability at least $\sqrt{\delta - 2\epsilon}$ over G , the winner $W_{(\cdot, H)}$ wins the sampled instance of G with probability at least ϵ , or otherwise the analogue is true for $W_{(G, \cdot)}$ on H .

2. Second, it is argued that repeating $W_{(\cdot, H)}$ for some q number of times, we obtain a winning probability of at least

$$\sqrt{\delta - 2\epsilon} \cdot (1 - (1 - \epsilon)^q),$$

approaching $\sqrt{\delta}$ as desired.

For example, to amplify a winning probability of $\delta = 0.01$ to close to $\sqrt{\delta} = 0.1$, say to 0.099, we need about $q \approx 76'600$ repetitions (with the optimal choice of $\epsilon \approx 8.65 \cdot 10^{-5}$). Even for a much less ambitious amplification to 0.09 only, we need $q = 4'981$ repetitions (choosing $\epsilon \approx 7.56 \cdot 10^{-4}$).

In the above two-step analysis, it seems that both steps are (essentially) optimal. Yet, their composition is, at least in certain regimes, quite far from optimal. We present a *combined* analysis that takes into account how the winning set \mathcal{S} behaves under the amplification, proving the very same reduction to be more efficient. In the above example, the desired amplification is achieved with only $q = 90$ and $q = 45$ repetitions, respectively (instead of $q \approx 76'600$ and $q = 4'981$).

It is easy to verify that if the winning set \mathcal{S} was a perfect square (similar to the rectangle in Figure 1), we would need $q = 44$ and $q = 22$ repetitions. A consequence of our results is that a rectangle as in Figure 1, even though it may seem to be a naively optimistic perspective, is actually close to the *worst* that can happen for the amplification.

1.3 Contributions and Outline

We briefly state our main contributions in a simplified manner. In Section 3, we present amplification theorems at the level of probability theory. We start

² Such games are called *clonable* in [Mau17].

³ See our proof of Theorem 1.

by showing a basic amplification theorem (Theorem 2) that yields an amplification similar to the known results [Gol01,CRS+07,Mau17]. Then, we show an improved analysis of the same type of statement, obtaining stronger amplification (Theorem 4).

In Section 4, we discuss that the proved amplification result is close to optimal, though still not perfect. We state a conjecture for a perfectly optimal amplification bound.

Finally, in Section 5, we demonstrate how the presented type of amplification theorem can be applied to non-trivial interactive games. We prove hardness amplification results for a general type of game (which includes the MAC forgery and the signature forgery game, and the simpler one-way function inversion as well as the hash function collision finding game), and give a comparison to related results of [DIJK09].

1.4 Related work

There exists a vast amount of literature on hardness amplification. We just mention some of them. Yao [Yao82] originally proposed the direct product construction for one-way functions. Goldreich [Gol01] showed an asymptotic hardness amplification result, stating that the direct product of weak one-way functions is a strong one-way function. Canetti et. al [CRS+07] studied the amplification of hash function collision resistance. They analyze a direct product construction similar to [Gol01], mainly to provide a baseline to compare against other constructions. [Mau17] introduced the notion of clonable games, and proved a bound similar to [Gol01] but for concrete parameters (non-asymptotic).

A related line of research [CHS05,IJK07,Jut10,DIJK09,HS11] studies hardness amplification via the direct product for games that are weakly-verifiable, i.e., where a solver may not be able to verify itself (efficiently) whether a given answer is correct. Some of these results are based on (a variant of) the XOR-Lemma.

In [DJMW12], it is shown that direct product hardness amplification “beyond negligible” is in general impossible (under certain plausible assumptions), meaning that for any negligible function $\epsilon(n)$, there exist cryptographic games such that their direct product can always be won with probability $\epsilon(n)$, no matter how many copies one takes.

2 Preliminaries

Notation. For $n \in \mathbb{N}$, we let $[n]$ denote the set $\{1, \dots, n\}$ with the convention $[0] = \emptyset$. The set of sequences (or strings) of length n over the alphabet \mathcal{A} is denoted by \mathcal{A}^n . An element of \mathcal{A}^n is denoted by $a^n = (a_1, \dots, a_n)$ for $a_i \in \mathcal{A}$.

In this paper, we assume all probability distributions to be over a finite set (or at least to have finite support). We let $\text{supp}(X)$ denote the *support* of a probability distribution X . Moreover, for two probability distributions X and Y , we let XY denote the independent joint distribution of X and Y . For example, we have $\mathbb{E}_{XY}[f(X, Y)] = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \Pr^X(X = x) \cdot \Pr^Y(Y = y) \cdot f(x, y)$.

We will need the following lemma.

Lemma 1. *For some $\gamma \in \mathbb{R}_+$, let $\psi : [0, \gamma] \rightarrow \mathbb{R}_+$ be a concave function. Then, we have for any $0 \leq a \leq b \leq \gamma$*

$$a\psi(b) \leq b\psi(a).$$

Proof. Assume without loss of generality $b \neq 0$ (since otherwise $a = 0$, so the inequality holds trivially). We have

$$\begin{aligned} a\psi(b) &= b \cdot \frac{a}{b} \cdot \psi(b) \\ &\leq b \cdot \left(\frac{a}{b} \psi(b) + \left(1 - \frac{a}{b}\right) \psi(0) \right) \\ &\leq b \cdot \psi\left(\frac{a}{b} \cdot b + \left(1 - \frac{a}{b}\right) \cdot 0\right) \\ &= b\psi(a). \end{aligned}$$

In the third step, we have used that ψ is concave. □

3 The Amplification Theorem

3.1 The Setting

In order to justify the type of amplification theorems we will prove (and in order to provide some intuition), we briefly explain the typical way they can be used.

We assume two finite sets \mathcal{G} and \mathcal{H} , representing the deterministic instances of games⁴. Since the actual games are probabilistic, they are (not necessarily uniform) probability distributions G and H over the sets \mathcal{G} and \mathcal{H} . Wherever a joint distribution of G and H is needed, we mean the *independent* joint distribution (i.e., the product distribution).

We further consider a winner W for the product game $[G, H]^\wedge$, and let the function $\mu : \mathcal{G} \times \mathcal{H} \rightarrow [0, 1]$ denote the winning probability of W . This means that for each pair of instances $(g, h) \in \mathcal{G} \times \mathcal{H}$, the probability that W wins *both* g and h is $\mu(g, h)$. Hence, the probability that W wins the game $[G, H]^\wedge$ is the expected value

$$\mathbb{E}_{GH}[\mu(G, H)].$$

In order to use W as a winner for G , we simulate (or absorb) an instance of H towards W to obtain a winner $W_{(\cdot, H)}$. On a sampled instance $g \in \text{supp}(G)$ we want to win, we then apply an *amplification* to our winner $W_{(\cdot, H)}$, such that if its original success probability is⁵ ϵ on this fixed instance g , we obtain an amplified

⁴ For the amplification theorem itself, it will not be important what exact (type of) object the games (i.e., the elements of \mathcal{G} and \mathcal{H}) are. For example, they may be Turing machines (of a certain kind).

⁵ Of course, this probability will depend on the sampled instance $g \in \text{supp}(G)$, so we will not actually know the value of ϵ in general.

success probability of $\psi(\epsilon)$ (for some amplification function $\psi : [0, 1] \rightarrow [0, 1]$). This means that our winning probability on G is (at least) the nested expectation

$$\mathbb{E}_G[\psi(\mathbb{E}_H[\mu(G, H)])].$$

In the most straightforward applications, the amplification is achieved by repeating the winner q times independently, such that we are successful exactly if one repetition has been successful, resulting in the amplification function $\psi(x) = 1 - (1-x)^q$. This works for example for one-way function inversion [Gol01] or for hash function collision finding [CRS⁺07], where the winner needs to provide a solution (such as a pre-image of a given value) and we can efficiently verify whether an obtained solution is correct or not.

Loosely speaking, the amplification statements we will prove are of the following type:

If $\mathbb{E}_G[\psi(\mathbb{E}_H[\mu(G, H)])]$ and $\mathbb{E}_H[\psi(\mathbb{E}_G[\mu(G, H)])]$ are both “somewhat small”, then $\mathbb{E}_{GH}[\mu(G, H)]$ must be “much smaller”.

Turned around this means that

If $\mathbb{E}_{GH}[\mu(G, H)]$ is at least “somewhat large”, then at least one of $\mathbb{E}_G[\psi(\mathbb{E}_H[\mu(G, H)])]$ or $\mathbb{E}_H[\psi(\mathbb{E}_G[\mu(G, H)])]$ is “much larger”.

3.2 Amplification for Monotonic ψ

We first present a basic amplification theorem that works whenever the amplification function ψ is monotonically increasing. Technically, the proof is a simplified version of the main idea in the amplification theorems of [Gol01, Mau17].

Theorem 1. *Let $\mu : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ be any function, and let X and Y be probability distributions over \mathcal{X} and \mathcal{Y} , respectively. Moreover, let ψ and ψ' be monotonically increasing on $[0, 1]$, and assume that*

$$\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])] \leq \epsilon\psi(\delta) \quad \text{and} \quad \mathbb{E}_Y[\psi'(\mathbb{E}_X[\mu(X, Y)])] \leq \epsilon'\psi'(\delta')$$

for some $\epsilon, \delta, \epsilon', \delta' \in [0, 1]$. Then we have

$$\mathbb{E}_{XY}[\mu(X, Y)] \leq \epsilon\epsilon' + \delta + \delta'.$$

Proof. We first define the two sets

$$\mathcal{X}_{\geq \delta} := \{x \in \mathcal{X} \mid \mathbb{E}_Y[\mu(x, Y)] \geq \delta\} \quad \text{and} \quad \mathcal{Y}_{\geq \delta'} := \{y \in \mathcal{Y} \mid \mathbb{E}_X[\mu(X, y)] \geq \delta'\}.$$

The assumption implies that

$$\Pr^X(X \in \mathcal{X}_{\geq \delta}) \leq \epsilon \quad \text{and} \quad \Pr^Y(Y \in \mathcal{Y}_{\geq \delta'}) \leq \epsilon'.$$

Now, observe that

$$\begin{aligned}
& \mathbb{E}_{XY}[\mu(X, Y)] \\
& \leq \Pr^{XY}((X, Y) \in \mathcal{X}_{\geq \delta} \times \mathcal{Y}_{\geq \delta'}) \cdot \mathbb{E}_{XY}[\mu(X, Y) \mid (X, Y) \in \mathcal{X}_{\geq \delta} \times \mathcal{Y}_{\geq \delta'}] \\
& \quad + \Pr^X(X \notin \mathcal{X}_{\geq \delta}) \cdot \mathbb{E}_X[\mathbb{E}_Y[\mu(X, Y) \mid X \notin \mathcal{X}_{\geq \delta}]] \\
& \quad + \Pr^Y(Y \notin \mathcal{Y}_{\geq \delta'}) \cdot \mathbb{E}_Y[\mathbb{E}_X[\mu(X, Y) \mid Y \notin \mathcal{Y}_{\geq \delta'}]] \\
& \leq \Pr^X(X \in \mathcal{X}_{\geq \delta}) \cdot \Pr^Y(Y \in \mathcal{Y}_{\geq \delta'}) \\
& \quad + \mathbb{E}_X[\mathbb{E}_Y[\mu(X, Y) \mid X \notin \mathcal{X}_{\geq \delta}]] + \mathbb{E}_Y[\mathbb{E}_X[\mu(X, Y) \mid Y \notin \mathcal{Y}_{\geq \delta'}]] \\
& \leq \epsilon \epsilon' + \delta + \delta'.
\end{aligned}$$

This concludes the proof. \square

A generalized n -fold version of [Theorem 1](#) is as follows.

Theorem 2. *Let $\mu : \mathcal{X}^n \rightarrow [0, 1]$ be any function, and let $\{X_i\}_{i \in [n]}$ be probability distributions over \mathcal{X} . Moreover, let $\{\psi_i\}_{i \in [n]}$ be a family of monotonically increasing functions on $[0, 1]$, and assume that for all $i \in [n]$ we have*

$$\mathbb{E}_{X_i}[\psi_i(\mathbb{E}_{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n}[\mu(X_1, \dots, X_n)])] \leq \epsilon_i \cdot \psi_i(\delta_i)$$

for some $\epsilon_i, \delta_i \in [0, 1]$. Then we have

$$\mathbb{E}_{X_1 \dots X_n}[\mu(X_1, \dots, X_n)] \leq \prod_{i \in [n]} \epsilon_i + \sum_{i \in [n]} \delta_i.$$

As mentioned in the introduction, the typical amplification function is of the form $\psi(x) = 1 - (1 - x)^q$ for some $q \in \mathbb{N}$. This motivates the following corollary that is proved in [Appendix A](#).

Corollary 1. *For arbitrary $\epsilon \in (0, 1)$ and $\delta_i \in (0, 1)$, let $\psi(x) = 1 - (1 - x)^q$ for q such that*

$$q \geq n \cdot \nu_{n, \epsilon} \cdot \prod_{i \in [n]} \delta_i^{-1},$$

where $\nu_{n, \epsilon} := \inf_{c \in (0, 1)} \frac{-\ln(1 - (1 - \epsilon)^{1 - c})}{1 - (1 - \epsilon)^{cn}} \in \left[\frac{\ln(1/\epsilon)}{1 - (1 - \epsilon)^n}, \frac{\ln(2/\epsilon)}{1 - (1 - \epsilon/2)^n} \right]$.

Assume that for all $i \in [n]$

$$\mathbb{E}_{X_i}[\psi(\mathbb{E}_{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n}[\mu(X_1, \dots, X_n)])] \leq (1 - \epsilon)\delta_i.$$

Then, we have

$$\mathbb{E}_{X_1 \dots X_n}[\mu(X_1, \dots, X_n)] \leq \prod_{i \in [n]} \delta_i.$$

3.3 Amplification for Monotonic and Concave ψ

As mentioned in [Section 3.1](#), the standard amplification function for such theorems is $\psi(x) = 1 - (1-x)^q$, which is concave. In the following, we exploit the concavity of ψ to obtain a stronger amplification.

Theorem 3. *Let $\mu : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ be any function, and let X and Y be probability distributions over \mathcal{X} and \mathcal{Y} , respectively. Moreover, let ψ and ψ' be monotonically increasing and concave on $[0, 1]$, and assume that*

$$\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])] \leq \epsilon\psi(\delta) \quad \text{and} \quad \mathbb{E}_Y[\psi'(\mathbb{E}_X[\mu(X, Y)])] \leq \epsilon'\psi'(\delta')$$

for some $\epsilon, \delta, \epsilon', \delta' \in [0, 1]$. Then we have

$$\mathbb{E}_{XY}[\mu(X, Y)] \leq \max(\epsilon\epsilon', \epsilon\delta + \epsilon'\delta').$$

Before proving the theorem, we remark that at first glance, one might expect the proof to rely on Jensen's inequality. For concave ψ , Jensen's inequality would give us

$$\begin{aligned} \mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])] &\leq \psi(\mathbb{E}_{XY}[\mu(X, Y)]) \\ \iff \psi^{-1}(\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])]) &\leq \mathbb{E}_{XY}[\mu(X, Y)]. \end{aligned}$$

However, our goal is to *upper* bound $\mathbb{E}_{XY}[\mu(X, Y)]$. Observe that by considering one dimension only, say $\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])]$, no non-trivial bound on $\mathbb{E}_{XY}[\mu(X, Y)]$ can be obtained, as we might have

$$\mathbb{E}_{XY}[\mu(X, Y)] = \mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])].$$

To consider *both* dimensions $\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])]$ and $\mathbb{E}_Y[\psi(\mathbb{E}_X[\mu(X, Y)])]$ is what will enable us to obtain a good upper bound on $\mathbb{E}_{XY}[\mu(X, Y)]$.

Proof (of [Theorem 3](#)). Just as in the proof of [Theorem 1](#), we first define the two sets

$$\mathcal{X}_{\geq\delta} := \{x \in \mathcal{X} \mid \mathbb{E}_Y[\mu(x, Y)] \geq \delta\} \quad \text{and} \quad \mathcal{Y}_{\geq\delta'} := \{y \in \mathcal{Y} \mid \mathbb{E}_X[\mu(X, y)] \geq \delta'\}.$$

We derive

$$\begin{aligned} \mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)]) \mid X \notin \mathcal{X}_{\geq\delta}] &= \mathbb{E}_X\left[\frac{\delta}{\delta} \cdot \psi(\mathbb{E}_Y[\mu(X, Y)]) \mid X \notin \mathcal{X}_{\geq\delta}\right] \\ &\geq \mathbb{E}_X\left[\frac{\psi(\delta)}{\delta} \cdot \mathbb{E}_Y[\mu(X, Y)] \mid X \notin \mathcal{X}_{\geq\delta}\right] \\ &= \frac{\psi(\delta)}{\delta} \mathbb{E}_X[\mathbb{E}_Y[\mu(X, Y)] \mid X \notin \mathcal{X}_{\geq\delta}]. \end{aligned}$$

The second step is due to [Lemma 1](#). This implies that

$$\begin{aligned} \mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])] &= \Pr^X(X \in \mathcal{X}_{\geq\delta}) \cdot \mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)]) \mid X \in \mathcal{X}_{\geq\delta}] \\ &\quad + \Pr^X(X \notin \mathcal{X}_{\geq\delta}) \cdot \mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)]) \mid X \notin \mathcal{X}_{\geq\delta}] \\ &\geq \Pr^X(X \in \mathcal{X}_{\geq\delta}) \cdot \psi(\delta) \\ &\quad + \Pr^X(X \notin \mathcal{X}_{\geq\delta}) \cdot \frac{\psi(\delta)}{\delta} \mathbb{E}_X[\mathbb{E}_Y[\mu(X, Y)] \mid X \notin \mathcal{X}_{\geq\delta}]. \end{aligned}$$

Since we have $\epsilon\psi(\delta) \geq \mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])]$ by assumption, we obtain

$$\Pr^X(X \notin \mathcal{X}_{\geq \delta}) \cdot \mathbb{E}_X[\mathbb{E}_Y[\mu(X, Y)] | X \notin \mathcal{X}_{\geq \delta}] \leq \delta \cdot (\epsilon - \Pr^X(X \in \mathcal{X}_{\geq \delta})).$$

Analogously, we obtain

$$\Pr^Y(Y \notin \mathcal{Y}_{\geq \delta'}) \cdot \mathbb{E}_Y[\mathbb{E}_X[\mu(X, Y)] | Y \notin \mathcal{Y}_{\geq \delta'}] \leq \delta' \cdot (\epsilon' - \Pr^Y(Y \in \mathcal{Y}_{\geq \delta'})).$$

Now, observe that

$$\begin{aligned} \mathbb{E}_{XY}[\mu(X, Y)] &\leq \Pr^{XY}((X, Y) \in \mathcal{X}_{\geq \delta} \times \mathcal{Y}_{\geq \delta'}) \cdot \mathbb{E}_{XY}[\mu(X, Y) | (X, Y) \in \mathcal{X}_{\geq \delta} \times \mathcal{Y}_{\geq \delta'}] \\ &\quad + \Pr^X(X \notin \mathcal{X}_{\geq \delta}) \cdot \mathbb{E}_X[\mathbb{E}_Y[\mu(X, Y)] | X \notin \mathcal{X}_{\geq \delta}] \\ &\quad + \Pr^Y(Y \notin \mathcal{Y}_{\geq \delta'}) \cdot \mathbb{E}_Y[\mathbb{E}_X[\mu(X, Y)] | Y \notin \mathcal{Y}_{\geq \delta'}] \\ &\leq \Pr^X(X \in \mathcal{X}_{\geq \delta}) \cdot \Pr^Y(Y \in \mathcal{Y}_{\geq \delta'}) \\ &\quad + \delta \cdot (\epsilon - \Pr^X(X \in \mathcal{X}_{\geq \delta})) + \delta' \cdot (\epsilon' - \Pr^Y(Y \in \mathcal{Y}_{\geq \delta'})). \end{aligned}$$

By assumption we must have $\Pr^X(X \in \mathcal{X}_{\geq \delta}) \leq \epsilon$ and $\Pr^Y(Y \in \mathcal{Y}_{\geq \delta'}) \leq \epsilon'$, so let $\Pr^X(X \in \mathcal{X}_{\geq \delta}) = \gamma\epsilon$ and $\Pr^Y(Y \in \mathcal{Y}_{\geq \delta'}) = \omega\epsilon'$ for $\gamma, \omega \in [0, 1]$. Then we get

$$\begin{aligned} \mathbb{E}_{XY}[\mu(X, Y)] &\leq \gamma\omega\epsilon\epsilon' + \epsilon\delta(1 - \gamma) + \epsilon'\delta'(1 - \omega) \\ &\leq \gamma\omega\epsilon\epsilon' + \epsilon\delta(1 - \gamma\omega) + \epsilon'\delta'(1 - \gamma\omega) \\ &= \gamma\omega\epsilon\epsilon' + (1 - \gamma\omega)(\epsilon\delta + \epsilon'\delta') \\ &\leq \max(\epsilon\epsilon', \epsilon\delta + \epsilon'\delta'). \end{aligned}$$

□

In the symmetric case, the optimal choice is $\epsilon = \epsilon' = 2\delta = 2\delta'$. This gives the following bound.

Corollary 2. *For any $\mu : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ and any monotonically increasing and concave function $\psi : [0, 1] \rightarrow [0, 1]$, let $\xi(x) := x \cdot \psi(x)$. We have*

$$\mathbb{E}_{XY}[\mu(X, Y)] \leq 4 \cdot \xi^{-1} \left(\frac{\max(\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])], \mathbb{E}_Y[\psi(\mathbb{E}_X[\mu(X, Y)])])}{2} \right)^2.$$

Equivalently,

$$\max(\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])], \mathbb{E}_Y[\psi(\mathbb{E}_X[\mu(X, Y)])]) \geq 2\xi(\sqrt{\mathbb{E}_{XY}[\mu(X, Y)]}/2).$$

A generalized n -fold version of [Theorem 3](#) is as follows.

Theorem 4. Let $\mu : \mathcal{X}^n \rightarrow [0, 1]$ be any function, and let $\{X_i\}_{i \in [n]}$ be probability distributions over \mathcal{X} . Moreover, let $\{\psi_i\}_{i \in [n]}$ be a family of monotonically increasing and concave functions on $[0, 1]$, and assume that for all $i \in [n]$ we have

$$\mathbb{E}_{X_i}[\psi_i(\mathbb{E}_{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n}[\mu(X_1, \dots, X_n)])] \leq \epsilon_i \cdot \psi_i(\delta_i)$$

for some $\epsilon_i, \delta_i \in [0, 1]$. Then we have

$$\mathbb{E}_{X_1 \dots X_n}[\mu(X_1, \dots, X_n)] \leq \max \left(\prod_{i \in [n]} \epsilon_i, \sum_{i \in [n]} \epsilon_i \delta_i \right).$$

The following corollary is proved in [Appendix A](#).

Corollary 3. For any $i \in [n]$, let $\ell_i \geq 1$, $\epsilon_i \in (0, 1)$, $\delta_i \in (0, 1)$, and

$$q_i \geq n \cdot \ell_i \cdot \ln(1/\epsilon_i) \cdot \prod_{j \in [n], j \neq i} \delta_j^{-1}$$

be arbitrary, and assume that for $\psi_i(x) = 1 - (1 - x/\ell_i)^{q_i}$ we have

$$\mathbb{E}_{X_i}[\psi_i(\mathbb{E}_{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n}[\mu(X_1, \dots, X_n)])] \leq (1 - \epsilon_i)\delta_i.$$

Then,

$$\mathbb{E}_{X_1 \dots X_n}[\mu(X_1, \dots, X_n)] \leq \prod_{i \in [n]} \delta_i.$$

Note that [Corollary 3](#) is a strictly stronger version of [Corollary 1](#), where (assuming all ϵ_i are equal to ϵ , and $\ell_i = 1$) we needed

$$q \geq n \cdot \nu_{n, \epsilon} \cdot \prod_{i \in [n]} \delta_i^{-1}.$$

The improvements of the new bound are two-fold:

1. First, the weaker version has a factor of (at least)

$$\nu_{n, \epsilon} \geq \frac{\ln(1/\epsilon)}{1 - (1 - \epsilon)^n} \geq \frac{\ln(1/\epsilon)}{n\epsilon} \quad \text{instead of just } \ln(1/\epsilon).$$

For fixed n , this means that q is proportional to $(1/\epsilon) \ln(1/\epsilon)$ instead of just $\ln(1/\epsilon)$. It is easy to see that, at least in certain regimes, the value $\nu_{n, \epsilon}$ is significantly larger than $\ln(1/\epsilon)$. For example, for $n = 2$ and $\epsilon = 0.001$ we have $\nu_{n, \epsilon} \approx 5118.5$, and $\ln(1/\epsilon) \approx 6.9$.

Moreover, this means that how close one can efficiently amplify δ^n to δ does not depend any more on n .

2. Second, the weaker version has a factor of

$$\prod_{j \in [n]} \delta_j^{-1} \quad \text{instead of just} \quad \prod_{j \in [n], j \neq i} \delta_j^{-1}.$$

Technically, the difference may seem to be small (in particular for large n and all δ_j close to 1). Conceptually, however, the new term is exactly what one would naturally expect, and the best one can hope for in an amplification theorem of a very general type: If we want to boost the winning probability of a winner W from δ^n to δ by running W q times with a success probability of at most δ^n in each run, we need $q \cdot \delta^n \geq \delta \iff q \geq (1/\delta)^{n-1}$. Put differently: When amplifying the hardness from δ to δ^n , the cost of the reduction is inversely proportional to the hardness *increase* (which is unavoidable), as opposed to the *target* hardness δ^n .

4 The Square is not (Always) Optimal

How tight are the bounds shown in the previous section, in particular those for concave amplification function ([Theorem 3](#) and [Theorem 4](#))? It is easy to see that the rectangle (or, in the symmetric case, the square) is optimal within a factor of at most 2 (see the discussion in [Section 1.2](#)).

Corollary 4. *Let $\mu : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ be any function, let ψ and ψ' be monotonically increasing and concave on $[0, 1]$, and assume that*

$$\mathbb{E}_X[\psi(\mathbb{E}_Y)] \leq \epsilon \psi(\epsilon') \quad \text{and} \quad \mathbb{E}_Y[\psi'(\mathbb{E}_X)] \leq \epsilon' \psi'(\epsilon).$$

Then we have

$$\mathbb{E}_{XY}[\mu(X, Y)] \leq 2\epsilon\epsilon'.$$

More generally, an n -dimensional orthotope (or hyperrectangle) is optimal within a factor of at most n .

One might conjecture that the rectangle is always optimal, i.e., that the factor of 2 in the above corollary can be removed. In the following, we show that this is not true.

Proposition 1. *There exist $\mu : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$, monotonically increasing and concave functions ψ and ψ' on $[0, 1]$, as well as distributions X and Y over \mathcal{X} and \mathcal{Y} , and $\epsilon, \epsilon' \in [0, 1]$, such that*

$$\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])] \leq \epsilon \cdot \psi(\epsilon') \quad \text{and} \quad \mathbb{E}_Y[\psi'(\mathbb{E}_X[\mu(X, Y)])] \leq \epsilon' \cdot \psi'(\epsilon),$$

but

$$\mathbb{E}_{XY}[\mu(X, Y)] > \epsilon\epsilon'.$$

Proof. Consider the following function $\mu : \{x_1, x_2\} \times \{y_1, y_2\} \rightarrow [0, 1]$:

y_2	1	0
y_1	1	1
	x_1	x_2

Moreover, let $\Pr^X(x_1) = \Pr^Y(y_1) = \frac{1}{4}$, and $\psi(x) = \psi'(x) = 1 - (1 - x)^2$. For $\epsilon = \epsilon' = 0.65582$ we have

$$\mathbb{E}_X[\psi(\mathbb{E}_Y[\mu(X, Y)])] = \mathbb{E}_Y[\psi'(\mathbb{E}_X[\mu(X, Y)])] = \frac{37}{64} \leq \epsilon' \cdot \psi'(\epsilon) = \epsilon \cdot \psi(\epsilon').$$

However,

$$\mathbb{E}_{XY}[\mu(X, Y)] = \frac{7}{16} = .4375 > \epsilon\epsilon' = \epsilon^2 \approx .431.$$

□

The choice of μ in the above example seems to work for any function of the form $\psi(x) = 1 - (1 - x)^q$, though for larger q we need $\Pr^X(x_1)$ and $\Pr^Y(y_1)$ to be closer to 0.

Even though the square is not optimal in general, we believe that whenever it is not optimal, the “opposite square” is optimal. By “opposite square” we mean that there is a square $\mathcal{S} = \mathcal{X}' \times \mathcal{Y}' \subseteq \mathcal{X} \times \mathcal{Y}$ such that $\mu(x, y) = 0$ if $(x, y) \in \mathcal{S}$ and $\mu(x, y) = 1$ otherwise. Loosely speaking, this means that the worst that can happen in terms of amplification is that either the *success* probability of a winner is maximally concentrated (into a square), or the *failure* probability is maximally concentrated. The following makes this mathematically rigorous.

Conjecture 1. Let $\mu : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ be any function, and X and Y arbitrary distributions over \mathcal{X} and \mathcal{Y} , respectively. Moreover, let $\psi(x) = 1 - (1 - x)^q$ for some $q \in \mathbb{N}$, and assume that

$$\max(\mathbb{E}_X(\psi(\mathbb{E}_Y[\mu(X, Y)])), \mathbb{E}_Y(\psi(\mathbb{E}_X[\mu(X, Y)]))) \leq \epsilon\psi(\epsilon)$$

for some $\epsilon \in [0, 1]$. Let $\delta \in [0, 1]$ be the (unique) value such that

$$\epsilon\psi(\epsilon) = (1 - \delta) + \delta\psi(1 - \delta).$$

Then, we have

$$\mathbb{E}_{XY}[\mu(X, Y)] \leq \max(\epsilon^2, 1 - \delta^2).$$

The above conjecture is stated for the two-dimensional symmetric case and only for the function $\psi(x) = 1 - (1 - x)^q$, but we conjecture natural generalizations to be true as well.

5 Applying the Amplification Theorem

As mentioned in the introduction, it is easy to obtain concrete hardness amplification results for certain games such as one-way function inversion or hash collision finding. Known asymptotic results such as “weak one-way functions imply strong one-way functions” are straightforward to derive from [Corollary 3](#) (with a more efficient reduction). Such games have been called *clonable* in [\[Mau17\]](#).

In the following, we demonstrate how the presented amplification theorems can be applied to more involved games that are not clonable, such as MAC forgery or signature forgery games.

We first give a redefinition (with some minor changes) of the type of game that has been introduced as “Dynamic Weakly Verifiable Puzzle” (or DWVP) in [\[DIJK09\]](#). A DWVP is an abstraction that captures certain cryptographic games such as the MAC forgery game or the signature forgery game (but includes the simpler one-way function inversion game as well as the hash function collision finding game).

Definition 1. *A deterministic DWVP is characterized by a function $h : \mathcal{M} \rightarrow \mathcal{H}$ and a relation $\sigma \subseteq \mathcal{M} \times \mathcal{S}$. The game supports the following query types:*

- *HINT-query: A query of the form $m \in \mathcal{M}$ that is answered with $h(m)$.*
- *VERIFICATION-query: A query of the form $(m', s) \in \mathcal{M} \times \mathcal{S}$. This query is always answered with a fixed symbol (say, \perp).*

The game is won when a VERIFICATION-query (m', s) is made such that

$$(m', s) \in \sigma \text{ and } m' \text{ was not asked before as HINT-query.}$$

Moreover, the game may support arbitrary additional query types.

Definition 2. *A probabilistic DWVP is a probability distribution over (compatible) deterministic DWVPs.*

For the MAC forgery game, for example, the HINT-queries would enable the winner to ask for tags of chosen messages, and the VERIFICATION-queries would correspond to forgery attempts ($\mathcal{H} = \mathcal{S}$ would correspond to the set of tags).

For certain games, an additional query type may be used to inform the winner about the instance to be solved (in a way that does not count as a hint). For example, a signature forgery game may use this to output the generated public key. Or, a one-way function inversion game would use this to output the function image y that is supposed to be inverted.

Now, we define the direct product of DVWPs. In contrast to [\[DIJK09\]](#), we give a more general definition, taking the direct product of arbitrary (potentially different) DVWPs, and define the direct product in a way such that the resulting game is not necessarily a DVWP anymore.

Definition 3. The direct product of deterministic DWVPs $\{g_i\}_{i \in [n]}$, denoted by

$$[g_1, \dots, g_n]^\wedge,$$

is the game which answers queries of the form (i, q) , where $i \in [n]$ and q is a query for the subgame g_i . It is won exactly when all subgames g_i are won.

The direct product $[G_1, \dots, G_n]^\wedge$ of probabilistic DWVPs is defined by lifting the deterministic definition via the independent joint distribution.

Notation 1. Analogous to the above games, we model (compatible) winners (or solvers) as probability distributions over deterministic winners.

We assume a predicate ω that describes whether a deterministic winner w wins a game g or not, i.e., $\omega(g, w) = 1$ if w wins g (and $\omega(g, w) = 0$ otherwise). For a given probabilistic winner W for a game G , we let

$$\omega(G, W) = \mathbb{E}_{GW}[\omega(G, W)] \in [0, 1]$$

denote the winning (or success) probability of W playing G .

In the following, we present a direct product hardness amplification theorem for arbitrary DWVPs.

Theorem 5. Let $\{G_i\}_{i \in [n]}$ be a family of probabilistic DWVPs. Let W be a winner for the direct product $[G_1 \dots G_n]^\wedge$, and asking up to h_i HINT-queries to G_i .

For any $\{\delta_j\}_{j \in [n]}$ and $\{\epsilon_j\}_{j \in [n]}$ with $\delta_j, \epsilon_j \in (0, 1]$, there are uniform reductions $\{\rho_i\}_{i \in [n]}$ and non-uniform reductions $\{\rho'_i\}_{i \in [n]}$, such that if W has a winning probability of

$$\prod_{j \in [n]} \delta_j,$$

then,

(i) For some $i \in [n]$, the winner $\rho_i(W)$ for G_i has winning probability at least

$$\omega(G_i, \rho_i(W)) \geq \frac{(1 - \epsilon_i)\delta_i}{e(h_i + 1)},$$

where ρ_i runs the winner W for the direct product $\lceil q_i \rceil$ times for

$$q_i = n \cdot \ln(1/\epsilon_i) \cdot \prod_{j \in [n], j \neq i} \delta_j^{-1}.$$

(ii) For some $i \in [n]$, the winner $\rho'_i(W)$ for G_i has winning probability at least

$$\omega(G_i, \rho'_i(W)) \geq (1 - \epsilon_i)\delta_i,$$

where ρ'_i runs the winner W for the direct product $\lceil q'_i \rceil$ times for

$$q'_i = n \cdot e(h_i + 1) \cdot \ln(1/\epsilon_i) \cdot \prod_{j \in [n], j \neq i} \delta_j^{-1}.$$

We provide some intuition before proving the theorem. We would like to amplify the winning probability of W by repeating it multiple times. The problem is that this might not increase our winning probability, since it can happen that W makes a successful VERIFICATION-query on a message which was asked as a HINT-query in an earlier repetition. To overcome this problem, the natural idea, originating in [DIJK09], is to disallow certain messages to be asked as HINT-query. We show two possibilities of achieving this: In the first version, we simply pick messages randomly (ad-hoc) to disallow as HINT-query. This enables a uniform and efficient reduction, but comes at the cost of introducing a loss factor of $(h_i + 1)$ in the obtained winning probability. This is why we present a second version, in which we provide the reduction with some non-uniform advice (that depends on the winner W). The advice essentially describes a fixed set of messages that are supposed to be disallowed as HINT-query, such that the loss in winning probability of the first (uniform) version can be overcome just by repeating more often (by a factor of $(h_i + 1)$). Our non-uniform version can be made uniform in a similar way as in [DIJK09], at the cost of introducing a similarly expensive precomputation.

We stress that the following proof is almost entirely concerned with analyzing the loss when certain messages are disallowed as HINT-queries, whereas the actual direct product amplification simply follows from Corollary 3.

Proof (of Theorem 5). The main idea, originating in [DIJK09], is to prevent certain messages to be asked as HINT-query. This is why we introduce a filter system F , acting as a proxy between a winner W and a game G_i that does the following:

1. First, for each⁶ message $m \in \mathcal{M}$, F decides independently with probability $1/(h_i + 1)$ that m is disallowed to ask as a HINT-query.
2. Then, all queries from the connected winner are proxied and the response is forwarded back, unless a HINT-query m is asked for a disallowed message m , in which case F just returns an error symbol, say \perp , as response.

For any filter $f \in \text{supp}(F)$ and any winner W_i for an instance $g \in \text{supp}(G_i)$, we let

$$\hat{\omega}(g, W_i, f)$$

denote the f -restricted winning probability of W_i playing g through the filter f , where only VERIFICATION-queries that are *disallowed* as HINT-queries by the filter f are taken into account⁷. This gives us the following useful property: When

⁶ Of course, this is most efficiently done by sampling lazily as we go, only for the messages that actually appear.

⁷ Note that W_i 's *actual* winning probability through the filter may be larger than $\hat{\omega}(g, W_i, f)$, since we allow to ask VERIFICATION-queries that are allowed as HINT-queries as well. We do not want to disallow those with the filter, since it may happen that W_i first asks some VERIFICATION-queries that are allowed as HINT-queries and then still makes a successful VERIFICATION-query that is disallowed as HINT-query.

a winner W_i for G_i is run q times independently through any fixed (deterministic) filter $f \in \text{supp}(F)$, the obtained success probability is at least

$$1 - (1 - \hat{\omega}(G_i, W_i, f))^q.$$

Observe that for any deterministic instance g , we have

$$\omega(g, W_i) \leq (e \cdot (h_i + 1)) \cdot \hat{\omega}(g, W_i, F).$$

This is because if we have h_i (distinct) hint queries M_1, \dots, M_{h_i} and the first successful attack query is M_{h_i+1} , the probability that the attack is also successful through the filter F and M_{h_i+1} is disallowed as a HINT-query is at least

$$\left(1 - \frac{1}{h_i + 1}\right)^{h_i} \cdot \frac{1}{h_i + 1} \geq \frac{1}{e \cdot (h_i + 1)}.$$

In the following, let $W_{i\sim}$ denote the winner for G_i that is obtained from W by simulating independent instances of $(G_1, \dots, G_{i-1}, G_{i+1}, \dots, G_n)$ towards W .

For claim (i), consider the following reduction: ρ_i maps a winner W for the game $[G_1 \dots G_n]^\wedge$ to a winner W_i for G_i which simply runs $W_{i\sim}$ q_i times independently through the filter F (without resetting the filter). Let $\chi_i(x) := 1 - (1 - x)^{q_i}$. For any $g \in \text{supp}(G_i)$ we have

$$\begin{aligned} \frac{\chi_i(\omega([G_1 \dots G_{i-1} \ g \ G_{i+1} \dots G_n]^\wedge, W))}{e \cdot (h_i + 1)} &\leq \frac{\chi_i(\omega(g, W_{i\sim}))}{e \cdot (h_i + 1)} \\ &\leq \frac{\chi_i(\omega(g, W_{i\sim}))}{\omega(g, W_{i\sim})} \cdot \hat{\omega}(g, W_{i\sim}, F) \\ &= \frac{\chi_i(\omega(g, W_{i\sim}))}{\omega(g, W_{i\sim})} \cdot \mathbb{E}_F[\hat{\omega}(g, W_{i\sim}, F)] \\ &= \mathbb{E}_F \left[\frac{\hat{\omega}(g, W_{i\sim}, F)}{\omega(g, W_{i\sim})} \chi_i(\omega(g, W_{i\sim})) \right] \\ &\leq \mathbb{E}_F \left[\frac{\omega(g, W_{i\sim})}{\omega(g, W_{i\sim})} \chi_i(\hat{\omega}(g, W_{i\sim}, F)) \right] \\ &= \mathbb{E}_F[\chi_i(\hat{\omega}(g, W_{i\sim}, F))] \\ &\leq \omega(g, \rho_i(W)). \end{aligned}$$

In the first step, we have used that χ_i is monotonically increasing and that

$$\omega([G_1 \dots G_{i-1} \ g \ G_{i+1} \dots G_n]^\wedge, W) \leq \omega(g, W_{i\sim}).$$

The second step is due to the inequality $\omega(g, W_{i\sim}) \leq (e \cdot (h_i + 1)) \cdot \hat{\omega}(g, W_{i\sim}, F)$. The fifth step is due to χ_i being concave, [Lemma 1](#), and the above inequality $\hat{\omega}(g, W_{i\sim}, f) \leq \omega(g, W_{i\sim})$ for any $f \in \text{supp}(F)$. Since the shown inequality holds for any $g \in \text{supp}(G_i)$, it also holds in expectation:

$$\begin{aligned} e(h_i + 1) \cdot \omega(G_i, \rho_i(W)) &= e(h_i + 1) \cdot \mathbb{E}_{G_i}[\omega(G_i, \rho_i(W))] \\ &\geq \mathbb{E}_{G_i}[\chi_i(\omega([G_1 \dots G_{i-1} \ G_i \ G_{i+1} \dots G_n]^\wedge, W))]. \end{aligned}$$

By [Corollary 3](#) we must have

$$\mathbb{E}_{G_i}[\chi_i(\omega([G_1 \dots G_{(i-1)} \ G_i \ G_{(i+1)} \dots G_n]^\wedge, W))] \geq (1 - \epsilon_i)\delta_i$$

for some $i \in [n]$, implying the first claim.

Now, we consider claim (ii). Recall that we have

$$\omega(g, W_{i\sim}) \leq (e \cdot (h_i + 1)) \cdot \hat{\omega}(g, W_{i\sim}, F)$$

for any $g \in \text{supp}(G_i)$, so the same is true in expectation:

$$\omega(G_i, W_{i\sim}) \leq (e \cdot (h_i + 1)) \cdot \hat{\omega}(G_i, W_{i\sim}, F) = (e \cdot (h_i + 1)) \cdot \mathbb{E}_F[\hat{\omega}(G_i, W_{i\sim}, F)].$$

Thus, there exists $f' \in \text{supp}(F)$ such that

$$\omega(G_i, W_{i\sim}) \leq (e \cdot (h_i + 1)) \cdot \hat{\omega}(G_i, W_{i\sim}, f').$$

Now, let the reduction ρ'_i map a winner W for $[G_1 \dots G_{(i-1)} \ G_i \ G_{(i+1)} \dots G_n]^\wedge$ to the winner W'_i that simply runs $W_{i\sim}$ q'_i times independently through the filter⁸ f' . Note that since we only use that the events \mathcal{E}_i of message m_i being disallowed as a HINT-query are $(h_i + 1)$ -wise independent, an appropriate f' with short description always exists (one can take, for example, a $(h_i + 1)$ -universal hash function).

For $\chi_i(x) = 1 - (1 - x)^{q'_i}$ we have

$$\begin{aligned} \mathbb{E}_{G_i} \left[\chi_i \left(\frac{\omega([G_1 \dots G_{(i-1)} \ G_i \ G_{(i+1)} \dots G_n]^\wedge, W)}{e \cdot (h_i + 1)} \right) \right] &\leq \mathbb{E}_{G_i} \left[\chi_i \left(\frac{\omega(G_i, W_{i\sim})}{e \cdot (h_i + 1)} \right) \right] \\ &\leq \chi_i \left(\frac{\mathbb{E}_{G_i}[\omega(G_i, W_{i\sim})]}{e \cdot (h_i + 1)} \right) \\ &= \chi_i \left(\frac{\omega(G_i, W_{i\sim})}{e \cdot (h_i + 1)} \right) \\ &\leq \chi_i(\hat{\omega}(G_i, W_{i\sim}, f')) \\ &\leq \omega(G_i, \rho'_i(W)). \end{aligned}$$

The second step is due to Jensen's inequality (χ_i is concave). By instantiating [Corollary 3](#) with $\ell_i = e(h_i + 1)$ we obtain

$$\mathbb{E}_{G_i} \left[\chi_i \left(\frac{\omega([G_1 \dots G_{(i-1)} \ G_i \ G_{(i+1)} \dots G_n]^\wedge, W)}{e \cdot (h_i + 1)} \right) \right] \geq (1 - \epsilon_i)\delta_i$$

for some $i \in [n]$, implying the second claim. \square

⁸ If the filter answers a HINT-query with \perp , the current repetition can be aborted.

We point out some differences between our non-uniform amplification statement from [Theorem 5](#) and the non-uniform DWVP amplification Theorem 4 of [\[DIJK09\]](#).

- The reduction in [\[DIJK09\]](#) guarantees that only a *single* VERIFICATION-query is asked. This makes their analysis very complicated, and comes at the cost of an increased number of asked HINT-queries (by an additional factor of h compared to our bounds, where h is the total number of HINT-queries asked by the considered winner W). We describe a reduction that executes the winner multiple times and submits *all* VERIFICATION-queries.

It is important to note that it depends on the concrete game whether the number of VERIFICATION-queries asked is important or not. For example, in the case of the signature forgery game, it is trivial to reduce the number of submitted VERIFICATION-queries to one, since one can efficiently check whether a forgery attempt will be accepted or not. The same is true for any game where one can verify a VERIFICATION-query efficiently before submitting it.

For the MAC forgery game it will in general not be possible to verify a forgery attempt efficiently. However, it is still meaningful (and quite natural) to consider the case where the adversary is allowed to make multiple forgery attempts. Note, however, that our amplification statement is not applicable for games that allow only very few (or even just one) VERIFICATION-queries. This may be the case, for example, if the goal of the game is to guess a value from a small set (say, a bit).

- The statement in [\[DIJK09\]](#) is a Chernoff-type amplification result that covers the threshold case, i.e., it is in the more general setting where a winner does not solve *all* n independent instances, but only a fraction of them. It seems though that for MAC forgery and signature forgery games, the basic non-threshold case (which we cover) is of most interest.
- Our [Theorem 5](#) provides concrete (non-asymptotic) bounds with very small constant factors. In contrast, the statements of [\[DIJK09\]](#) hide large constants in asymptotic bounds. Moreover, we have a loss of a factor $(h_i + 1)$, that is independent of the number of VERIFICATION-query asked, whereas the loss in [\[DIJK09\]](#) is $\mathcal{O}(h + v)$, where h and v are the total number of HINT- and VERIFICATION-queries asked.
- We consider the direct product of n arbitrary DWVPs, i.e., the individual games are not required to be the same. In contrast, [\[DIJK09\]](#) studies the case where all n games are equal, and uses a restricted direct product definition that requires to ask the same query m to all instances in parallel.

Note that because our games $\{G_i\}_{i \in [n]}$ may be all different, we obtain an amplifying reduction for *some* G_i . If the games $\{G_i\}_{i \in [n]}$ are all the same, and one is aiming for a uniform reduction, it is a standard technique to embed the given instance g at a uniform random position $I \in [n]$. It may seem that one would lose a factor of n in winning probability when this is done. However, we note that by the AM–GM inequality, the conclusion

of [Corollary 3](#) can be extended to

$$\mathbb{E}_{X_1 \dots X_n}[\mu(X_1, \dots, X_n)] \leq \prod_{i \in [n]} \delta_i \leq \left(\sum_{i \in [n]} \frac{\delta_i}{n} \right)^n.$$

This prevents losing a factor of n when embedding the given instance at a uniform position.

6 Conclusions and Open Problems

We presented an abstract direct product hardness amplification theorem at the level of probability theory. Our hope is that phrasing it at this level enables reusability and leads to a more modular analysis of hardness amplification statements, similar as in our proof of hardness amplification for DVWPs ([Theorem 5](#)). The theorem assumes an arbitrary concave amplification function ψ , simply because the proof does not require further assumptions. This leads to the question of whether natural examples of games with corresponding reductions exist, where the function ψ is something totally different than $1 - (1 - x)^q$ or $1 - (1 - x/\ell)^q$.

Moreover, the shown bounds are close to optimal, but still not perfectly tight. We phrased a conjecture for a perfectly tight bound, which states that the worst case in terms of amplification is that either the *success* probability or the *failure* probability of the considered winner is maximally concentrated. Independently of this conjecture, it seems that the factor of n in the number of repetitions q (see [Corollary 3](#)) can be significantly reduced.

Finally, we leave it for future work to generalize the amplification statements beyond the “product” setting, for example to the “threshold” setting.

References

- CHS05. Ran Canetti, Shai Halevi, and Michael Steiner. Hardness amplification of weakly verifiable puzzles. In Joe Kilian, editor, *Theory of Cryptography*, pages 17–33, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- CRS⁺07. Ran Canetti, Ron Rivest, Madhu Sudan, Luca Trevisan, Salil Vadhan, and Hoeteck Wee. Amplifying collision resistance: A complexity-theoretic treatment. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007*, pages 264–283, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- DIJK09. Yevgeniy Dodis, Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Security amplification for interactive cryptographic primitives. In Omer Reingold, editor, *Theory of Cryptography*, pages 128–145, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- DJMW12. Yevgeniy Dodis, Abhishek Jain, Tal Moran, and Daniel Wichs. Counterexamples to hardness amplification beyond negligible. In Ronald Cramer, editor, *Theory of Cryptography*, pages 476–493, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- Gol01. Oded Goldreich. *Foundations of Cryptography*, volume 1. Cambridge University Press, 2001.
- HS11. Thomas Holenstein and Grant Schoenebeck. General hardness amplification of predicates and puzzles. In Yuval Ishai, editor, *Theory of Cryptography*, pages 19–36, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- IJK07. Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Chernoff-type direct product theorems. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007*, pages 500–516, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- Jut10. Charanjit S. Jutla. Almost optimal bounds for direct product threshold theorem. In Daniele Micciancio, editor, *Theory of Cryptography*, pages 37–51, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- Mau17. Ueli Maurer. An information-theoretic approach to hardness amplification. In *2017 IEEE International Symposium on Information Theory (ISIT)*, June 2017.
- Yao82. Andrew C. Yao. Theory and application of trapdoor functions. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, pages 80–91, 1982.

Appendix

A Proofs

Proof (of Corollary 1). Let $c \in (0, 1)$ be arbitrary. Moreover, we let $a_i = (1 - \epsilon)^c \delta_i$ and $b_i = -\ln(1 - (1 - \epsilon)^{1-c})/q$.

$$\begin{aligned}
 a_i \cdot \psi(b_i) &= (1 - \epsilon)^c \delta_i \cdot \psi_i(-\ln(1 - (1 - \epsilon)^{1-c})/q) \\
 &= (1 - \epsilon)^c \delta_i \cdot (1 - (1 - (-\ln(1 - (1 - \epsilon)^{1-c})/q)^q)) \\
 &\geq (1 - \epsilon)^c \delta_i \cdot (1 - (\exp(\ln(1 - (1 - \epsilon)^{1-c})/q)^q)) \\
 &= (1 - \epsilon)^c \delta_i \cdot (1 - \epsilon)^{1-c} \\
 &= (1 - \epsilon) \delta_i.
 \end{aligned}$$

From [Theorem 2](#) we obtain

$$\begin{aligned}
 \mathbb{E}_{X_1, \dots, X_n}[\mu(X_1, \dots, X_n)] &\leq \prod_{i \in [n]} a_i + \sum_{i \in [n]} b_i \\
 &= (1 - \epsilon)^{cn} \prod_{i \in [n]} \delta_i + \sum_{i \in [n]} b_i \\
 &= (1 - \epsilon)^{cn} \prod_{i \in [n]} \delta_i + n \cdot -\ln(1 - (1 - \epsilon)^{1-c})/q \\
 &\leq (1 - \epsilon)^{cn} \prod_{i \in [n]} \delta_i + n \cdot \frac{-\ln(1 - (1 - \epsilon)^{1-c})}{n \cdot \nu_{n, \epsilon} \cdot \prod_{i \in [n]} \delta_i^{-1}} \\
 &\leq (1 - \epsilon)^{cn} \prod_{i \in [n]} \delta_i + (1 - (1 - \epsilon)^{cn}) \cdot \prod_{i \in [n]} \delta_i \\
 &= \prod_{i \in [n]} \delta_i.
 \end{aligned}$$

Finally, we show that

$$\nu_{n, \epsilon} \in \left[\frac{\ln(1/\epsilon)}{1 - (1 - \epsilon)^n}, \frac{\ln(2/\epsilon)}{1 - (1 - \epsilon/2)^n} \right].$$

First, observe that

$$\begin{aligned}
 \nu_{n, \epsilon} &= \inf_{c \in (0, 1)} \frac{-\ln(1 - (1 - \epsilon)^{1-c})}{1 - (1 - \epsilon)^{nc}} \geq \frac{\inf_{c \in (0, 1)} -\ln(1 - (1 - \epsilon)^{1-c})}{\sup_{c \in (0, 1)} 1 - (1 - \epsilon)^{nc}} \\
 &= \frac{\ln(1/\epsilon)}{1 - (1 - \epsilon)^n}.
 \end{aligned}$$

The upper bound is shown as follows.

$$\begin{aligned}
\nu_{n,\epsilon} &= \inf_{c \in (0,1)} \frac{-\ln(1 - (1 - \epsilon)^{1-c})}{1 - (1 - \epsilon)^{nc}} \leq \inf_{c \in (0,1)} \frac{-\ln(1 - (1 - (1 - c)\epsilon))}{1 - (1 - c \cdot \epsilon)^n} \\
&= \inf_{c \in (0,1)} \frac{-\ln((1 - c)\epsilon)}{1 - (1 - c \cdot \epsilon)^n} \\
&\leq \frac{\ln(2/\epsilon)}{1 - (1 - \epsilon/2)^n}.
\end{aligned}$$

This concludes the proof. □

Proof (of [Corollary 3](#)). Observe that for any $i \in [n]$ we have

$$\begin{aligned}
\delta_i \cdot \psi_i(\ell_i \ln(1/\epsilon_i)/q_i) &= \delta_i \cdot (1 - (1 - \ln(1/\epsilon_i)/q_i)^{q_i}) \\
&\geq \delta_i \cdot (1 - (e^{-\ln(1/\epsilon_i)/q_i})^{q_i}) \\
&= (1 - \epsilon_i)\delta_i.
\end{aligned}$$

Thus, we have by [Theorem 4](#)

$$\begin{aligned}
\mathbb{E}_{X_1 \dots X_n}[\mu(X_1, \dots, X_n)] &\leq \max \left(\prod_{i \in [n]} \delta_i, \sum_{i \in [n]} \delta_i \ell_i \ln(1/\epsilon_i)/q_i \right) \\
&\leq \max \left(\prod_{i \in [n]} \delta_i, \prod_{i \in [n]} \delta_i \right) \\
&= \prod_{i \in [n]} \delta_i.
\end{aligned}$$

This concludes the proof. □