

Nondeterministic Direct Product Reductions and the Success Probability of SAT Solvers [Preliminary version]

Andrew Drucker*

May 6, 2013

Last updated: May 16, 2013

Abstract

We give two nondeterministic reductions which yield new direct product theorems (DPTs) for Boolean circuits. In these theorems one assumes that a target function f is mildly hard against *nondeterministic* circuits, and concludes that the direct product $f^{\otimes t}$ is extremely hard against (only polynomially smaller) probabilistic circuits. The main advantage of these results compared with previous DPTs is the strength of the size bound in our conclusion.

As an application, we show that if NP is not in coNP/poly then, for every PPT algorithm attempting to produce satisfying assignments to Boolean formulas, there are infinitely many instances where the algorithm's success probability is nearly-exponentially small. This furthers a project of Paturi and Pudlák [STOC'10].

Contents

1	Introduction	2
1.1	Hardness amplification and direct product theorems	2
1.2	Our new direct product theorems	4
1.3	Application to the success probability of PPT SAT solvers	6
1.4	Our techniques	7
2	Preliminaries	9
2.1	Deterministic and probabilistic Boolean circuits	10
2.2	Nondeterministic circuits and nondeterministic mappings	10
2.3	Direct-product solvers	12
2.4	Hash families and their properties	13
2.5	A general nondeterministic circuit construction	14

*School of Mathematics, IAS, Princeton, NJ. Email: andy.drucker@gmail.com. This material is based upon work supported by the National Science Foundation under agreements Princeton University Prime Award No. CCF-0832797 and Sub-contract No. 00001583. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

3	Stage-based analysis of direct-product computations	15
3.1	Valid outputs and the α, β sequences	15
3.2	A lemma on random sequences	16
3.3	Behavior of the α, β sequences	19
3.4	Analysis of worst-case direct-product solvers: the key experiment	21
4	The main DPT for worst-case direct-product solvers	23
5	The success probability of SAT solvers	29
6	The DPT for sampleable distributions	31
6.1	A “confidence-building” lemma	32
6.2	A DPT for input-confident direct-product solvers	36
6.3	Proof of Theorem 6.1	39

1 Introduction

This work contributes to two central areas of study in complexity theory: *hardness amplification* on the one hand, and *the complexity of NP search problems* on the other.

1.1 Hardness amplification and direct product theorems

In the general hardness amplification project, we assume that we have identified a function f that is “mildly hard” to compute, for some class \mathcal{C} of resource-bounded algorithms. Our goal is to derive a second function f' that is “extremely hard” to compute, for some possibly-different class \mathcal{C}' . In our initial discussion we will focus on the case where $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$, $f' : \{0, 1\}^{n'} \rightarrow \{0, 1\}^{d'}$ are finite functions, and $\mathcal{C}, \mathcal{C}'$ are two sets of probabilistic Boolean circuits, but we note that the project can be studied in other models as well.

The notion of difficulty suggested above can be formalized in two ways (both relevant to our work). Let $p \in [0, 1]$. In the *average-case* setting, let us say that f is *p-hard* for \mathcal{C} with respect to an input distribution \mathcal{D} if, on an input $\mathbf{x} \sim \mathcal{D}$, every algorithm in \mathcal{C} computes $f(\mathbf{x})$ correctly with probability at most p . In the *worst-case* setting, we say that f is *worst-case p-hard* for a class \mathcal{C} of randomized algorithms if, for every algorithm $C \in \mathcal{C}$, there is an input x such that $\Pr_C[C(x) = f(x)] \leq p$. In either setting, from a “mild” hardness guarantee $p = 1 - \varepsilon$ for \mathcal{C} in computing f , we want to obtain a much stronger bound $p' \ll 1$ for the second class \mathcal{C}' in computing f' .

There are several motivations to pursue hardness amplification. First, the security of most modern cryptographic primitives, such as one-way functions and public-key cryptosystems, inherently requires the existence of computational problems, solvable in NP, which possess a strong average-case hardness guarantee. While the mere existence of hard-on-average problems in NP is not known to be *sufficient* for doing cryptography, a better understanding of the sources of average-case hardness seems necessary for making progress in the foundations of cryptography. Moreover, hardness-amplification techniques in complexity theory have also helped to inspire tools for the *security-amplification* of cryptographic primitives. (See, e.g., [Gol07] for background on the complexity-theoretic underpinnings of modern cryptography and on ideas of security amplification.)

Second, average-case hardness is also inherent in the fundamental concept of *pseudorandomness*: a pseudorandom source is information-theoretically distinguishable from random bits, yet the distinguishing task must be computationally hard-on-average. Techniques of hardness amplification have played a key role in important constructions of pseudorandom generators [HILL99, IW97].

Third, hardness amplification, and in particular the *direct product* approach to amplifying hardness, is interesting in its own right, and helps us to critically examine some of our most basic intuitions about computation, as we will describe.

Given a function f as above and $t \in \mathbb{N}^+$, let $f^{\otimes t} : \{0, 1\}^{n \times t} \rightarrow \{0, 1\}^{d \times t}$, the t -fold *direct product* of f , be the function which takes t length- n inputs (x^1, \dots, x^t) and outputs $(f(x^1), \dots, f(x^t))$. A *direct product theorem (DPT)* is any result upper-bounding the success bound p' for computing $f^{\otimes t}$ in terms of an assumed success bound p for f (and, possibly, other parameters).

When f is Boolean, the direct product construction can be considered along with the related “ t -fold XOR” $f^{\oplus t}(x^1, \dots, x^t) := \bigoplus f(x^j)$ (i.e., the sum mod 2). An “XOR lemma” is any result upper-bounding the success bound p' for computing $f^{\oplus t}$ in terms of an assumed success bound p for f . The original XOR lemma was stated by Yao in unpublished work, with the first published proof due to Levin [Lev87]; see [GNW11] for more information on the lemma’s history and several proofs. Unlike the direct product, the t -fold XOR $f^{\oplus t}$ is itself a Boolean function, which can be advantageous in applications of hardness amplification, but which can also be a disadvantage since this limits its average-case hardness (an algorithm may compute $f^{\oplus t}$ with success probability $1/2$ by guessing a random bit). Several works show how in some settings XOR lemmas may be obtained from DPTs [GL89, GNW95] or vice versa [NRS99, VW08]. We will not prove or use XOR lemmas in this work; we merely point out that their study is often intimately linked with the study of direct product theorems (indeed, XOR lemmas are even referred to as “direct product theorems” in some papers).

The motivation for the direct product construction is as follows. Let $\mathcal{C}_{\leq s}$ be the class of probabilistic Boolean circuits of size at most s . It would *appear* that, if any circuit $C \in \mathcal{C}_{\leq s}$ has success probability at most $p < 1$ in computing f , then any circuit $C' \in \mathcal{C}_{\leq s}$ should have success probability not much more than p^t in computing $f^{\otimes t}$. The intuition here is that combining the t “unrelated” computations should not help much, and simply trying to solve each instance separately would be a nearly-optimal strategy.

This hypothesis can be considered in both the worst-case and the average-case setting; in the latter, if f is p -hard with respect to inputs drawn from \mathcal{D} , then it is natural to study the difficulty of computing $f^{\otimes t}$ with t inputs drawn independently from \mathcal{D} .¹

One might even be tempted to make the bolder conjecture that $f^{\otimes t}$ is p^t -hard against circuits in $\mathcal{C}_{\leq t \cdot s}$, but this was shown by Shaltiel to fail badly in the average-case setting [Sha03]. So what kind of DPT is known to hold in the circuit model? A version of the following theorem, with slightly weaker parameters, was proved in [GNW11, first version in '95]; a tighter argument leading to the bounds given below is described in [Wig97]. In the form of the DPT stated below, the focus is on getting a success probability bound of at most some $\varepsilon > 0$ for $f^{\otimes t}$, where t is chosen accordingly.

Theorem 1.1 (see [GNW11], Lemma 7, and [Wig97], Theorems 2.9, 2.10). *Suppose the Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is p -hard for circuits of size s with respect to input distribution \mathcal{D} , for*

¹“Derandomized” variants of the scenario, in which the t inputs are not fully independent, have also been studied, and powerful “derandomized” DPTs were obtained, notably in [IW97, IJKW10]. We will not consider these in the present paper. Neither will we consider “XOR lemmas”, which analyze a variant of the direct product construction in the case where f is Boolean [GNW11].

some $p < 1$. For an appropriate $t = \Theta\left(\frac{\ln(1/\varepsilon)}{1-p}\right)$, the function $f^{\otimes t}$ is ε -hard with respect to $\mathcal{D}^{\otimes t}$ for circuits with size bounded by s' , where

$$s' = \Theta\left(\frac{\varepsilon \cdot s}{\ln(1/\varepsilon)}\right).$$

This is a powerful and elegant result, but one whose parameters can be disappointing in many situations. The size bound s' degrades quickly as $\varepsilon \rightarrow 0$; if $s \leq \text{poly}(n)$, i.e., if our initial hardness assumption is against circuits of some fixed-polynomial size (and $p = 2/3$, say), then Theorem 1.1 cannot give a success bound of $n^{-\omega(1)}$ against any nontrivial class of circuits. In unpublished work, Steven Rudich has observed that this barrier is inescapable for a certain class of “black-box,” relativizing, *deterministic* (or probabilistic) reductions (see [GNW11] for more discussion). The limitations of more general black-box hardness-amplification reductions have been extensively studied, particularly for the case of XOR lemmas and other reductions that produce a Boolean function; Shaltiel and Viola’s work [SV10] is one notable contribution which also provides a good guide to this literature.

1.2 Our new direct product theorems

In this work we show that, if we are willing to assume that our starting function f is somewhat hard to compute by *nondeterministic* circuits, then we obtain very strong hardness results for $f^{\otimes t}$ against the class of probabilistic circuits. The direct product theorems we prove have quantitative parameters that are far superior to those in Theorem 1.1.

Our first DPT holds for the worst-case setting. We show:

Theorem 1.2. *Let $f = \{f_n\}$ be a family of Boolean functions on n input bits, and suppose that f is “hard for nondeterministic circuits” in the sense that $f \notin \text{NP/poly} \cap \text{coNP/poly}$.²*

Now let $100 \leq t(n) \leq \text{poly}(n)$ be a parameter, and let $\{C_n\}_{n>0}$ be any family of polynomial-size, probabilistic circuits outputting $t(n)$ bits. Then for infinitely many choices of n and $x \in \{0, 1\}^{n \times t(n)}$,

$$\Pr[C_n(x) = f_n^{\otimes t(n)}(x)] < \exp\left(-\frac{t(n)}{3 \cdot 10^5}\right).$$

Like all known direct product theorems in the circuit setting, this result is proved in its contrapositive form; it follows straightforwardly from our Theorem 4.1. The latter is a *nondeterministic direct product reduction*—a method for transforming a probabilistic circuit that weakly approximates $f_n^{\otimes t(n)}$ into a *nondeterministic* circuit that computes f_n with much greater confidence. In the present reduction, we get a nondeterministic circuit that computes f_n exactly. This transformation incurs only a polynomial blowup in circuit size. The reduction is fully black-box, but is not subject to the limitations identified by Rudich due to its use of nondeterminism.

Our next DPT holds for the average-case setting, and applies to input distributions that are efficiently sampleable.

Theorem 1.3. *Let $\{f_n\}$ be a family of Boolean functions on n input bits. Let $\mathcal{D} = \{\mathcal{D}_n\}$ be a family of input distributions, sampleable by a polynomial-size family of circuits. Let $\delta_n \in [.5, 1]$ be*

²Here, strictly speaking we mean that the language $L_f = \{x : f(x) = 1\}$ is not in $\text{NP/poly} \cap \text{coNP/poly}$.

a parameter. Suppose that f is “hard for nondeterministic circuits with respect to \mathcal{D} ” in the sense that, for all families $\{G_n\}$ of polynomial-size nondeterministic circuits, if n is sufficiently large then

$$\Pr_{\mathbf{x} \sim \mathcal{D}_n} [G_n(\mathbf{x}) = f_n(\mathbf{x})] < 1 - \varepsilon_n . \quad 3$$

Now let $100 \leq t(n) \leq \text{poly}(n)$ be a parameter, and let $\{C_n\}_{n>0}$ be any family of polynomial-size, probabilistic circuits outputting $t(n)$ bits. Then for sufficiently large n and $\mathbf{x} \sim \mathcal{D}_n^{\otimes t(n)}$, we have

$$\Pr[C_n(\mathbf{x}) = f_n^{\otimes t(n)}(\mathbf{x})] < \exp\left(-\frac{\varepsilon_n^{3/2} \sqrt{t(n)}}{10^8}\right) .$$

This result follows easily from our Theorem 6.1, a nondeterministic direct product reduction for sampleable distributions. Our two DPTs above, stated for Boolean functions, are special cases of our general results, in which the functions f_n need not be Boolean; there is no adverse dependence in these results on the range size.

Our nondeterministic direct product reductions are not the first use of nondeterministic reductions in the study of average-case complexity. In particular, previous authors have exploited nondeterminism to give worst-case to average-case reductions for computational problems. In [FL97, first version in '92], Feige and Lund gave a nondeterministic worst-case to average-case reduction for computing the Permanent over large fields. We state one of the resulting theorems, which is interesting to compare with ours:

Theorem 1.4 ([FL97], Thm. 3.6). *Suppose there is a $\gamma \in (0, 1)$ and a probabilistic polynomial-time algorithm P such that, for any $n \times n$ 0-1 input matrix X , we have*

$$\Pr[P(X) = \text{perm}_{\mathbb{Z}}(X)] \geq 2^{-n^\gamma} .$$

Then, $\text{P}^{\#P} = \text{PH} = \text{AM}$.

Earlier Amir, Beigel, and Gasarch [ABG03, first version in '90] had shown somewhat weaker results for #SAT, the problem of computing the number of satisfying assignments to a Boolean formula. Thus we have long had a good understanding of the worst-case difficulty against PPT algorithms of #P-complete counting problems. We also have quite strong average-case hardness results for #P, assuming $\text{P}^{\#P} \neq \text{AM}$ [FL97]. However, #P functions are not believed to be computable with the power of NP, or even PH (this would collapse PH, since $\text{PH} \subseteq \text{P}^{\#P}$ [Tod91]). This seems to limit the potential applications of the strong hardness of #P functions, e.g., in cryptography.

As another example of the power of nondeterminism for average-case complexity, we mention a work of Trevisan and Vadhan; in the course of constructing deterministic extractors for efficiently sampleable distributions, they gave a nondeterministic worst-case to average-case reduction for the task of computing an arbitrary low-degree polynomial over a finite field [TV00, Lem. 4.1].

None of these reductions from previous work proceed through a direct product reduction in our sense. However, [ABG03] considers t -fold direct products $f^{\otimes t}$ and gives complexity upper bounds for functions f such that $f^{\otimes t}$ can be computed using few queries to an oracle; some of these reductions use nondeterminism.

³In the most common definition, a nondeterministic circuit $G_n(x)$ is said to *compute* a function $g(x)$ if the following condition holds: $g(x) = 1$ iff there exists some setting of the nondeterministic gates of G_n causing it to accept input x . Our Theorem 1.3 is valid if our initial hardness assumption on f_n is with respect to this definition. However, for this result we actually only require hardness with respect to a more restrictive model of nondeterministic computation: the model of *nondeterministic mapping circuits*, described in Section 2.2.

1.3 Application to the success probability of PPT SAT solvers

In a recent work, Paturi and Pudlák [PP10] asked about the achievable worst-case success probability of PPT heuristics for circuit satisfiability. They argue for the importance of this question by observing that many of the known exact algorithms for k -SAT and other NP search problems, which achieve exponential runtimes with an improved exponent over naive search, can be converted to *polynomial-time* search heuristics with a success probability attaining nontrivial advantage over random guessing. Thus, exploring the limitations of PPT search heuristics also illuminates the limitations of a very natural paradigm for exponential-time computation.

Paturi and Pudlák show the following powerful result (a special case of a more general theorem): Suppose there is a $\gamma < 1$ and a PPT algorithm P_{solver} that, when given as input a description of any satisfiable Circuit-SAT instance Ψ with r variables, outputs a satisfying assignment to Ψ with probability at least $q(r) := 2^{-\gamma r}$. Then, there is a deterministic circuit family $\{C_r\}_r$ that succeeds at the same task on r -variable instances of size $n \leq \text{poly}(r)$ with probability 1, and C_r is of size at most 2^{r^μ} for some $\mu < 1$.

In this work, we exploit a simple connection between direct product computations for the satisfiability decision problem for Boolean formulas, and the task of producing satisfying assignments to such formulas. The connection is simple enough to describe in a single sentence:

- (\star) If ψ^1, \dots, ψ^t are formulas, and $s \in [0, t]$ is a correct guess for the number of satisfiable ψ^j s, then the 0/1 values $[\psi^1 \in \text{SAT}], \dots, [\psi^t \in \text{SAT}]$ can all be inferred given any satisfying assignment to the formula $\Psi^{(s)}$ which asks for satisfying assignments to at least s of the ψ^j s.

This observation has been used before in [BH88] for quite different purposes, as we will discuss shortly. Using this connection together with our worst-case DPT (Theorem 4.1), we prove the following result, bounding the achievable success probability of polynomial-time heuristics for SAT solvers under the standard complexity-theoretic assumption that NP is not contained in coNP/poly .

Theorem 1.5. *Let $\gamma \in (0, 1)$. Suppose there is a PPT algorithm P_{solver} that, when given as input a description of a satisfiable 3-CNF formula Φ , of description length $|\langle \Phi \rangle| = N$, outputs a satisfying assignment to Φ with probability at least $q(N) := 2^{-N^\gamma}$.*

Then, $\text{NP} \subseteq \text{coNP}/\text{poly}$ (and the Polynomial Hierarchy collapses to Σ_3^P).

In fact, the conclusion would hold even if P_{solver} were a non-uniform polynomial-size circuit family rather than a uniform PPT algorithm.

This theorem is incomparable to the result of Paturi and Pudlák. Their result implies a stronger upper bound for the success probability of SAT solvers, at the cost of assuming that NP has nearly-exponential circuit complexity. While we admire this past work (and it gave the original motivation for our project, although our methods are almost completely different), we also believe that it is valuable to understand how much one can infer about the difficulty of NP problems, starting from assumptions about *polynomial-time computation alone*. The hypothesis that the Polynomial Hierarchy does not collapse to Σ_3^P can be equivalently restated as one such assumption: namely, that there is no polynomial-time procedure that reduces a 4-round game, described by a Boolean circuit, to an equivalent 3-round game with the “existential” player taking the first move (see [Pap94, Chap. 17]).

Essentially the same observation (\star) above was used earlier by Buss and Hay [BH88] (see also [Pap94, Chap. 17]) to show the equality of complexity classes $\text{P}_{||}^{\text{NP}} = \text{P}^{\text{NP}[\log]}$; that is, $\text{poly}(n)$ non-adaptive queries to an NP oracle are equivalent in power to $O(\log(n))$ adaptive queries, at least

for solving decision problems in polynomial time.⁴ We are not aware of previous work on hardness amplification using (\star) . However, many past works have used various connections between direct product theorems and the difficulty of NP search problems. Indeed, cryptography’s demand for various types of hard-on-average NP search problems has been a prime motivation for the study of hardness amplification as far back as Yao’s seminal work [Yao82], and direct product theorems for NP search problems such as preimage-finding have been an central ingredient in this project (see [Gol07, Chap. 2]). In the modern “hardness-versus-randomness” paradigm [NW94, IW97] for obtaining pseudorandom generators (PRGs) under hardness assumptions about $E = \text{DTIME}[2^{O(n)}]$, one important step is to apply a direct product theorem to certain *decision problems* lying in E .⁵ Hardness amplification has also been previously applied to NP *decision problems* to obtain hard-on-average NP *search problems*, e.g., in a previous work of ours [Dru11, Thm. 5]; the connection used there was more complicated than in our present work, and proved a much stronger hardness for NP search problems under a rather strong assumption on the exponential-time complexity of NP. Beyond this, a great deal of work has been devoted to amplifying average-case hardness *within* the class of NP decision problems (see [O’D02, BT06]); such works avoid both the direct product construction (which produces a non-Boolean function) and the XOR construction (which may produce a decision problem outside of NP).

1.4 Our techniques

We prove two main nondeterministic direct product reductions in this work. While there are major technical differences between them, both begin with the same intuition, which we will describe at a high level here.

Let f be a Boolean function on n input bits, and \mathcal{D} a distribution over inputs to f . Suppose C is a probabilistic circuit that computes $f^{\otimes t}$ with some success probability $q > 2^{-ct}$ on inputs $\bar{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^t) \sim \mathcal{D}^{\otimes t}$, for some small constant $0 < c \ll 1$. We would like to obtain from C a nondeterministic circuit computing f with high success probability with respect to \mathcal{D} . (To prove our worst-case direct product reduction, Theorem 4.1, it will suffice to show how to do this for every \mathcal{D} . We can then use the minimax theorem to build a nondeterministic circuit that is correct on every input.)

Say that an execution of C is j -*valid*, for $0 \leq j \leq t$, if it correctly computes f on its first j inputs. We obviously have

$$\Pr [C(\bar{\mathbf{x}}) = f^{\otimes t}(\bar{\mathbf{x}})] = \prod_{j \in [t]} \Pr[j\text{-valid} | (j-1)\text{-valid}] > 2^{-ct}.$$

Thus, for a *typical* choice of index j , $\Pr[j\text{-valid} | (j-1)\text{-valid}] \approx 2^{-c} \approx 1$. This motivates us to choose such an index $j = j^*$, and to *fix* some settings y^1, \dots, y^{j^*-1} to the first $(j^* - 1)$ inputs. Then, by storing the values $f(y^1), \dots, f(y^{j^*-1})$, we can easily *recognize* a $(j^* - 1)$ -valid execution (as specified by a setting to the remaining inputs and to the random bits used by C). Now, given a single input $\mathbf{x} \sim \mathcal{D}$ on which we wish to compute f , we will try to obtain a $(j^* - 1)$ -valid execution of C on an input-tuple whose first j^* elements are $y^1, \dots, y^{j^*-1}, \mathbf{x}$. The idea is that, by our choice

⁴In the reduction used in Buss and Hay’s work, the value s is determined by binary search, whereas our reduction simply makes a random guess for s . Our goal is to output the values $[\psi^1 \in \text{SAT}], \dots, [\psi^t \in \text{SAT}]$, whereas in [BH88] the goal is to compute some Boolean function of these values.

⁵This is relevant because every PRG must be a one-way function, and the associated preimage-finding task is an NP search problem which must be hard on average.

of j^* , a “typical” such execution (in which the remaining inputs are drawn from $\mathcal{D}^{\otimes(t-j^*)}$) should also be j^* -valid, and give us the value $f(\mathbf{x})$. This basic strategy can be seen in previous direct product reductions; see, e.g., [GNW11]. In our reduction, however, nondeterminism will allow us to obtain a $(j^* - 1)$ -valid execution of C very efficiently, even when such executions are extremely rare; we simply need to “guess and check.”

This approach requires care, however, because an execution obtained by nondeterministic guessing need not be a representative sample of the population from which it was drawn. Thus, even if we successfully fix values of y^1, \dots, y^{j^*-1} and receive an input $\mathbf{x} \sim \mathcal{D}$ such that most $(j^* - 1)$ -valid executions are also j^* -valid, we need a way to “kill off” the “atypical” $(j^* - 1)$ -valid executions which fail to be j^* -valid.

For this task, a natural idea is to try to apply *random hashing*, a well-established tool for reducing the size of a set and culling atypical elements. The use of randomly selected hash functions for such purposes, in conjunction with nondeterministic guessing, was pioneered by Goldwasser and Sipser [GS86].⁶ This technique has an important requirement, however: to kill all the “atypical” $(j^* - 1)$ -valid executions while simultaneously leaving at least one j^* -valid execution alive, we need to know a good approximation to the *probability* of a $(j^* - 1)$ -valid execution, conditioned on $y^1, \dots, y^{j^*-1}, \mathbf{x}$. We want this probability to be predictable with high accuracy based on y^1, \dots, y^{j^*-1} alone, without any foreknowledge of the input \mathbf{x} , so that a good approximation can be encoded into our final nondeterministic circuit as helpful advice. To summarize, we hope to find and fix inputs y^1, \dots, y^{j^*-1} , such that with high probability over $\mathbf{x} \sim \mathcal{D}$, we have:

- (i) $\Pr[j^*\text{-valid}|y^1, \dots, y^{j^*-1}, \mathbf{x}] \approx \Pr[(j^* - 1)\text{-valid}|y^1, \dots, y^{j^*-1}, \mathbf{x}];$
- (ii) $\Pr[(j^* - 1)\text{-valid}|y^1, \dots, y^{j^*-1}, \mathbf{x}] \approx \Pr[(j^* - 1)\text{-valid}|y^1, \dots, y^{j^*-1}].$

In each condition above, we will tolerate a $(1 \pm .01)$ multiplicative error in our approximations.

So how do we choose the values y^1, \dots, y^{j^*-1} ? The obvious idea is to choose them as independent samples from \mathcal{D} , after selecting j^* uniformly at random. However, this approach may *fail* to guarantee condition (i) above, if the successful direct-product computations of C are “concentrated” in a pathological way. For example, it may be that $C(x^1, \dots, x^t)$ always outputs $f^{\otimes t}(x^1, \dots, x^t)$ provided the first input x^1 lies in some “good” set $G \subset \{0, 1\}^n$ of probability mass $\approx 2^{-ct}$, while if $x^1 \notin G$, then C simply outputs random guesses. In this case, condition (i) fails for a typical setting $x^1 := y^1$. Also, condition (ii) is only useful if the conditional probabilities involved are nonzero, and this may also fail with high probability.

We address these difficulties in two distinct ways in our two direct product reductions:

- In our worst-case reduction (Theorem 4.1), we assume from the start that our C has a probability $\geq q$ of computing $f^{\otimes t}$ under *every* input. Then, the problem of zero probabilities doesn’t arise, and the worst-case success guarantee of C turns out to very usefully constrain the effects of conditioning on $y^1, \dots, y^{j^*-1} \sim \mathcal{D}$.
- In our average-case reduction, we show how to transform C to a second probabilistic circuit \tilde{C} taking some number $t' \ll t$ of n -bit inputs, and whose success probability is $\geq .5q$ on *almost*

⁶We are aiming to build a nondeterministic circuit, not a probabilistic one; but the eventual plan will be to fix a polynomial number of representative hash functions as non-uniform advice. Let us also mention that hash families were used in Paturi and Pudlák’s work [PP10] as well, but in a very different way. Those authors used hash functions to *reduce the amount of randomness* used by a SAT solver having a worst-case success probability guarantee, as a step toward transforming a Circuit-SAT instance into an equivalent instance with *fewer variables*.

every t' -tuple from $\mathcal{D}^{\otimes t'}$. We call this key attribute the *input-confidence property*; it turns out to be nearly as useful as the worst-case guarantee assumed in the previous case.

To endow C with the input-confidence property, we regard most of the t input positions as auxiliary sources of randomness, and concentrate on computing f on the remaining t' coordinates. This has the effect of “smoothing out” the success probability over the different possible t' -tuples. In this step and its analysis (which uses Kullback-Leibler divergence) we are strongly influenced by Raz’s proof of the Parallel Repetition Theorem [Raz98]. Also, in this transformation, we crucially use that \mathcal{D} is sampleable; we need our transformed circuit to generate its own samples from \mathcal{D} in order to be useful.

Let us rename t' as t in the average-case setting, for uniformity of discussion. In either setting, we are then able to show that choosing $j^*, y^1, \dots, y^{j^*-1}$ randomly as before ensures conditions (i) and (ii) with high probability. This requires careful work, but ultimately derives from elementary facts about the behavior of random sequences (see Lemma 3.4).

Now, in our average-case direct product reduction for sampleable distributions, we can perform hashing over all possible outcomes to x^{j^*+1}, \dots, x^t , weighted by their likelihood under \mathcal{D} . In our *worst-case* direct product reduction, \mathcal{D} may not be efficiently sampleable, which causes an additional challenge. In this setting we will show that in fact, it is adequate to draw x^{j^*+1}, \dots, x^t independently at random from multisets S_{j^*+1}, \dots, S_t , each obtained by sampling a reasonable (polynomial) number of times from \mathcal{D} . These “sparsified” versions of \mathcal{D} can be coded into our nondeterministic circuit. The idea of this sparsification and its analysis are somewhat similar to (and inspired by) a sparsification step from our previous work on instance compression [Dru12, Lem. 6.3].

2 Preliminaries

When we speak of *probabilistic polynomial-time (PPT) algorithms*, we refer to any uniform model of classical computation, e.g., multi-tape Turing machines. We assume familiarity with the basic uniform complexity classes P , NP and $coNP$ and the higher levels Σ_k^P, Π_k^P of the Polynomial Hierarchy PH. (For background on complexity classes, consult [AB09, Pap94].)

For a language $L \subseteq \{0,1\}^*$, we let χ_L denote the characteristic function of L . Let $\chi_L|_{=n} : \{0,1\}^n \rightarrow \{0,1\}$ denote its restriction to inputs of length n . If $f : \{0,1\}^n \rightarrow \{0,1\}^d$ is a function and $t \in \mathbb{N}^+$, we let $f^{\otimes t} : \{0,1\}^{n \times t} \rightarrow \{0,1\}^{d \times t}$, the t -fold direct product of f , be defined by

$$f^{\otimes t}(x^1, \dots, x^t) := (f(x^1), \dots, f(x^t)) .$$

All probability spaces we consider will be finite. For distributions $\mathcal{D}, \mathcal{D}'$ over a finite set U , we let $\mathcal{D}(u) := \Pr_{\mathbf{u} \sim \mathcal{D}}[\mathbf{u} = u]$, and for $U' \subseteq U$ let $\mathcal{D}(U') := \Pr_{\mathbf{u} \sim \mathcal{D}}[\mathbf{u} \in U']$. Let $\text{supp}(\mathcal{D})$, the *support* of \mathcal{D} , be defined as $\text{supp}(\mathcal{D}) := \{u : \mathcal{D}(u) > 0\}$. We let $\|\mathcal{D} - \mathcal{D}'\|_{\text{stat}} := \frac{1}{2} \sum_{u \in U} |\mathcal{D}(u) - \mathcal{D}'(u)|$ denote the statistical distance between \mathcal{D} and \mathcal{D}' . Equivalently,

$$\|\mathcal{D} - \mathcal{D}'\|_{\text{stat}} = \max_{U' \subseteq U} |\mathcal{D}(U') - \mathcal{D}'(U')| .$$

The statistical distance between two random variables X, X' is defined as the statistical distance between their respective distributions.

We use $\mathcal{D} \otimes \mathcal{D}'$ to denote the distribution over pairs $(\mathbf{u}, \mathbf{u}')$ which independently samples $\mathbf{u} \sim \mathcal{D}$ and $\mathbf{u}' \sim \mathcal{D}'$. We let $\mathcal{D}^{\otimes k}$ denote k independent samples from \mathcal{D} . For a finite multiset S , we write $\mathbf{a} \in_r S$ to indicate that random variable \mathbf{a} is sampled uniformly from S (with elements weighted by their multiplicity in S). We let $\mathbf{1}[E]$ denote the indicator random variable for an event E .

If S, S' are multisets, we use $S \cup S'$ to indicate multiset union, and similarly for other multiset operations.⁷ We let $S^{\times k}$ denote the k -fold Cartesian product of S .

We will appeal to the following form of the Chernoff-Hoeffding bound:

Lemma 2.1. *Suppose X_1, \dots, X_k are independent random variables in the range $[0, 1]$, each satisfying $\mathbb{E}[X_i] = \mu$. Let $X_{\text{avg}} := \frac{1}{k} \sum_{i \in [k]} X_i$. Then for $\varepsilon > 0$,*

$$\Pr[X_{\text{avg}} \leq \mu - \varepsilon] \leq \exp(-2\varepsilon^2 k) \quad \text{and} \quad \Pr[X_{\text{avg}} \geq \mu + \varepsilon] \leq \exp(-2\varepsilon^2 k).$$

2.1 Deterministic and probabilistic Boolean circuits

When we refer to *Boolean circuits*, we will consider the model of Boolean circuits with gates from $\{\wedge, \vee, \neg\}$. The \wedge, \vee gates are of fanin two. We let $\text{size}(C)$ denote the total number of gates in a (deterministic, probabilistic, or nondeterministic) Boolean circuit C , including input gates and all other types of gates. Note that $\text{size}(C)$ is within a factor 2 of the number of wires in C .

Boolean circuits may compute multiple-output functions, as follows. We consider every Boolean circuit C to be equipped with an ordering on some nonempty subset of its gates, call these g_1^*, \dots, g_d^* . We then say that C computes the function which on input x , outputs $(g_1^*(x), \dots, g_d^*(x))$. We denote this function simply as $C(x)$.

A *probabilistic Boolean circuit* is a circuit C which, in addition to having gates for its regular input variables $x = (x_1, \dots, x_n)$, has an additional designated set of (fanin-zero) “random” gates $r = (r_1, \dots, r_m)$. We write $C(x)$ to denote the random variable giving the output of C on input x , when the random gates are set to uniform random values. We write C^{det} to indicate the underlying deterministic circuit with input gates (x, r) . For a deterministic or probabilistic Boolean circuit C , we write

$$C : \{0, 1\}^m \rightarrow \{0, 1\}^d$$

to denote that C takes m regular input bits (not including any random gates), and produces d bits of output.

For a deterministic or probabilistic circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$, a value $\delta \in [0, 1]$, a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$, and a distribution \mathcal{D} over $\{0, 1\}^n$, we say that C $(1 - \delta)$ -computes f with respect to \mathcal{D} if $\Pr_{x \sim \mathcal{D}}[C(x) = f(x)] \geq 1 - \delta$.

A *sampling circuit* is a circuit C with no input gates, but with designated “random” gates r and output gates g_1^*, \dots, g_d^* . We say that C *samples* a distribution \mathcal{D} on $\{0, 1\}^d$ if a uniformly random setting to r induces output distribution \mathcal{D} . We let $C(r)$ denote the output on the setting r to the random gates.

2.2 Nondeterministic circuits and nondeterministic mappings

An *ordinary nondeterministic circuit* is a Boolean circuit C , accompanied with a specification of three sets of special gates in C :

⁷The multiplicity of (s, s') in the Cartesian product $S \times S'$ is equal to the product of the multiplicity of s in S with the multiplicity of s' in S' .

- a set of “ordinary” input gates $x = (x_1, \dots, x_n)$;
- a set of fanin-zero “nondeterministic” gates $w = (w_1, \dots, w_m)$;
- a single “accept/reject” gate g_0^* .

We require that the ordinary-input and nondeterministic gates be disjoint, but otherwise these sets are allowed to overlap. We say that $C(x, w)$ *accepts* (x, w) if $g_0^*(x, w) = 1$, otherwise we say C *rejects* (x, w) . Now if $\mathfrak{C} = \{C_n(x, w)\}_{n>0}$ is a family of nondeterministic circuits, with C_n having n input gates, we say that \mathfrak{C} *recognizes* a language L , if for all n and all $x \in \{0, 1\}^n$,

$$x \in L \iff \exists w : C_n(x, w) \text{ accepts.}$$

The complexity class NP/poly is defined as $\{L : \text{there is an ordinary nondeterministic circuit family } \{C_n\} \text{ recognizing } L, \text{ with } \text{size}(C_n) \leq \text{poly}(n)\}$. The class coNP/poly is defined as $\{L : \bar{L} \in \text{NP/poly}\}$.⁸ It appears unlikely that $\text{NP} \subseteq \text{coNP/poly}$, as this would imply a collapse of the Polynomial Hierarchy:

Theorem 2.2 ([Yap83]). *If $\text{NP} \subseteq \text{coNP/poly}$, then $\text{PH} = \Sigma_3^P = \Pi_3^P$.*

We will also use a fairly standard model of nondeterministic computation for functions with one or more bits of output (see [BLS84, BLS85], and [HO02, Sec. 3.3 and App. A.14] for background); here we specialize the model to the circuit setting and use somewhat different terminology. A *nondeterministic mapping circuit* is a circuit C which, in addition to the three classes of designated gates described above, also has a set of output gates g_1^*, \dots, g_d^* for some $d > 0$ (these gates may overlap with the other gate-classes). The *mapping defined by C* , denoted $F_C : \{0, 1\}^n \rightarrow \mathcal{P}(\{0, 1\}^d)$, is defined as

$$F_C(x) := \{v \in \{0, 1\}^d : \exists w \in \{0, 1\}^m \text{ such that } C \text{ accepts } (x, w) \text{ and } (g_1^*(x, w), \dots, g_d^*(x, w)) = v\}.$$

We will use the following obvious fact:

Proposition 2.3. *Let L be a language. Suppose there is a family $\{C_n\}_{n>0}$ of nondeterministic mapping circuits, satisfying*

1. *For every n and $x \in \{0, 1\}^n$, $F_{C_n}(x) = \{\chi_L(x)\}$;*
2. *$\text{size}(C_n) \leq \text{poly}(n)$.*

Then, $L \in \text{NP/poly} \cap \text{coNP/poly}$.

We say that a nondeterministic mapping circuit C $(1-\delta)$ -*defines* a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$ with respect to input distribution \mathcal{D} if $\Pr_{x \sim \mathcal{D}}[F_C(x) = \{f(x)\}] \geq 1 - \delta$.

⁸(Here, \bar{L} denotes the complement of L .)

2.3 Direct-product solvers

Definition 2.4 (Direct-product solvers). *Let $n, d, t \in \mathbb{N}^+$ be fixed, and let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$ be a function. Let $C : \{0, 1\}^{n \times t} \rightarrow \{0, 1\}^{d \times t}$ be a probabilistic circuit. Let $q, c \in [0, 1]$.*

1. *Let $\bar{\mathcal{D}}$ be a distribution over $\{0, 1\}^{n \times t}$. We say that C is a q -direct-product solver for $f^{\otimes t}$ with respect to $\bar{\mathcal{D}}$ if the following holds: if we sample $\bar{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^t) \sim \bar{\mathcal{D}}$, then*

$$\Pr [C(\bar{\mathbf{x}}) = (f(\mathbf{x}^1), \dots, f(\mathbf{x}^t))] \geq q ,$$

where the probability is taken over $\bar{\mathbf{x}}$ and over the randomness used by C .

2. *We say that the circuit C is a worst-case q -direct-product solver for $f^{\otimes t}$ if it satisfies item 1 above for every input distribution $\bar{\mathcal{D}}$ over $\{0, 1\}^{n \times t}$. Equivalently, for all $(x^1, \dots, x^t) \in \{0, 1\}^{n \times t}$,*

$$\Pr [C(x^1, \dots, x^t) = (f(x^1), \dots, f(x^t))] \geq q .$$

3. *For any $(x^1, \dots, x^t) \in \{0, 1\}^{n \times t}$, define the confidence*

$$\text{conf}(C; x^1, \dots, x^t) := \Pr [C(x^1, \dots, x^t) = (f(x^1), \dots, f(x^t))] .$$

We say that C is a (c, q) -input-confident direct-product solver for $f^{\otimes t}$ with respect to $\bar{\mathcal{D}}$ if the following holds: if we sample $\bar{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^t) \sim \bar{\mathcal{D}}$, then

$$\Pr [\text{conf}(C; \bar{\mathbf{x}}) \geq q] \geq c .$$

(Note that any such C is also a q' -direct-product solver for $f^{\otimes t}$ with respect to $\bar{\mathcal{D}}$, where $q' := qc$.)

4. *We extend each of the above definitions to families of circuits and input distributions, in the natural way. Let $d(n), t(n) : \mathbb{N}^+ \rightarrow \mathbb{N}^+$, $q(n) : \mathbb{N}^+ \rightarrow [0, 1]$ be functions. Let $\mathfrak{C} = \{C_n\}_{n>0}$ be a family of probabilistic circuits $C_n : \{0, 1\}^{n \times t(n)} \rightarrow \{0, 1\}^{d(n) \times t(n)}$. Let $\bar{\mathcal{D}} = \{\bar{\mathcal{D}}_n\}_{n>0}$ be a family of distributions over $\{0, 1\}^{n \times t(n)}$.*

We say that \mathfrak{C} is a $q(n)$ -direct-product solver for $f^{\otimes t(n)}$ with respect to $\bar{\mathcal{D}}$ if for any $n > 0$, C_n is a $q(n)$ -direct-product solver for $f^{\otimes t(n)}$ with respect to $\bar{\mathcal{D}}_n$. The definitions in items 2 and 3 are extended to circuit and input-distribution families in an analogous fashion.

If C is a q -direct-product solver for $f^{\otimes t}$ with respect to some input distribution, then we may non-uniformly fix the random gates of C to obtain a deterministic q -direct-product solver for $f^{\otimes t}$ with respect to the same input distribution. This is a standard observation in the study of average-case complexity. Note, however, that there is no obvious analogue of this transformation for worst-case direct-product solvers, or for input-confident direct product solvers, that would allow us to derandomize while preserving the quality parameters.

2.4 Hash families and their properties

We now introduce a widely-used class of function families called strongly universal hash families. Let U be a finite set, and let \mathcal{H} be a finite family of “hash functions” $h : U \rightarrow \mathbb{F}_2^k$. We say that \mathcal{H} is a *strongly universal hash family* if, for all $u, u' \in U$ with $u \neq v$, and for all $z, z' \in \mathbb{F}_2^k$, we have

$$\Pr_{h \in \mathcal{H}} [h(u) = z \wedge h(u') = z'] = 2^{-2k}.$$

We use the following standard, explicit construction:

Proposition 2.5. *Given $k, \ell > 0$, consider $U := \mathbb{F}_2^\ell$. The family of functions*

$$\mathcal{H}^{\ell, k} := \left\{ h_{A, v} : \mathbb{F}_2^\ell \rightarrow \mathbb{F}_2^k \right\}_{A \in \mathbb{F}_2^{k \times \ell}, v \in \mathbb{F}_2^k}$$

given by

$$h_{A, v}(x) := Ax + v \quad (\text{with addition over } \mathbb{F}_2^k)$$

is a strongly universal hash family.

Lemma 2.6. *Suppose that $k, \ell > 0$, that $U' \subseteq U = \mathbb{F}_2^\ell$, and that $|U'| = \theta \cdot 2^k$, for some $\theta > 0$. Say we select $(A, v) \in_r \mathbb{F}_2^{k \times \ell} \times \mathbb{F}_2^k$; then,*

$$1 - \theta^{-1} < \Pr[0^k \in h_{A, v}(U')] \leq \theta. \quad (1)$$

Proof. Let $U' = \{u_1, \dots, u_m\}$. For $i \in [m]$, define the indicator variable $X_i := [h_{A, v}(u_i) = 0^k]$. Note that $\mathbb{E}[X_i] = 2^{-k}$. Let $S := \sum_{i \in [m]} X_i$. Then $\Pr[0^k \in h_{A, v}(U')] = \Pr[S \geq 1]$, and by Markov’s inequality this is at most $\mathbb{E}[S] = \sum_{i \in [m]} 2^{-k} = \theta$. This proves the upper bound.

For the lower bound, note that $[0^k \notin h_{A, v}(U')]$ implies $S = 0$, so that in this case $|S - \mathbb{E}[S]| = \theta$. By Chebyshev’s inequality, this occurs with probability at most $\frac{\mathbb{E}[(S - \mathbb{E}[S])^2]}{\theta^2}$. The strongly universal property implies that the random variables $X_i, X_{i'}$ are pairwise independent, so we have

$$\Pr[0^k \notin h_{A, v}(U')] \leq \frac{\sum_{i \in [m]} \mathbb{E}[(X_i - 2^{-k})^2]}{\theta^2} = \frac{\theta \cdot 2^k}{\theta^2} \cdot (2^{-k}(1 - 2^{-k})) < \frac{1}{\theta}.$$

This gives the lower bound in the Lemma’s statement. \square

Corollary 2.7. *Suppose that $k, \ell > 0$ and that $U = \mathbb{F}_2^\ell$. Consider two subsets $U_{(i)}, U_{(ii)} \subseteq U$, with $|U_{(i)}| = \theta \cdot 2^k$ and $|U_{(ii)}| = \theta' \cdot 2^k$. If $(A, v) \in_r \mathbb{F}_2^{k \times \ell} \times \mathbb{F}_2^k$, then we have*

$$\Pr[0^k \in (h_{A, v}(U_{(i)}) \setminus h_{A, v}(U_{(ii)}))] \geq 1 - \theta^{-1} - \theta',$$

where $h_{A, v}$ is as in Proposition 2.5.

Proof. First we apply the lower bound of Lemma 2.6 with $U' := U_{(i)}$ to find that

$$\Pr[0^k \notin h_{A, v}(U_{(i)})] \leq \theta^{-1}.$$

Next we apply the upper bound of Lemma 2.6 with $U' := U_{(ii)}$ to find that

$$\Pr[0^k \in h_{A, v}(U_{(ii)})] \leq \theta'.$$

Taking a union bound completes the proof. \square

2.5 A general nondeterministic circuit construction

The following technical lemma, which is slightly complicated to state but conceptually simple, will be useful in our proofs of Theorems 4.1 and 6.1. Readers may wish to defer studying the lemma until it is used in its proper context.

Lemma 2.8. *Let $n, d \in \mathbb{N}^+$, and let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$ be a function. Let $K, N, T \in \mathbb{N}^+$ be additional parameters, and suppose we have identified several sets and mappings, as follows.*

- *For every $u \in \{0, 1\}^n$, there is a set $V_u \subseteq \{0, 1\}^N$, partitioned into disjoint subsets $\{V_u^z\}_{z \in \{0, 1\}^d}$;*
- *For $i \in [T]$, there is a mapping $h_i^* : \{0, 1\}^N \rightarrow \{0, 1\}^K$;*
- *There is a “favorable” set $\text{Fav} \subseteq \{0, 1\}^n$.*

Suppose that there is a deterministic circuit $C_{\text{Vtest}}(w, u) : \{0, 1\}^{N+u} \rightarrow \{0, 1\}^{d+1}$ that determines whether $w \in V_u$ and, if so, which subset V_u^z contains w . Suppose too that, for each $i \in [T]$, there is a circuit $C_{h_i^}$ of size $O(KN)$ that computes h_i^* .*

Say that h_i^ is good for $u \in \{0, 1\}^n$ if*

$$0^K \in h_i^* \left(V_u^{f(u)} \right) \setminus \left(\bigcup_{z \neq f(u)} h_i^* (V_u^z) \right).$$

Finally, assume the following condition:

- *For every $u \in \text{Fav}$, there are at least $.6T$ indices $i \in [T]$ such that h_i^* is good for u .*

Then, there exists a nondeterministic mapping circuit C^\dagger taking n input bits, such that for all $u \in \text{Fav}$,

$$F_{C^\dagger}(u) = \{f(u)\};$$

also, we have $\text{size}(C^\dagger) \leq O((\text{size}(C_{\text{Vtest}})KN) \cdot T)$.

Proof. C^\dagger takes input $u \in \{0, 1\}^n$ and has nondeterministic variables

$$\overline{w} = (w^1, \dots, w^T) \in \{0, 1\}^{N \times T}.$$

The circuit $C^\dagger(u, \overline{w})$ acts as follows:

1. For each $i \in [T]$, define $z^i \in \{0, 1\}^d \cup \{\perp\}$ by

$$z^i := \begin{cases} z & \text{if } w^i \in V_u^z \text{ and } h_i^*(w^i) = 0^K, \\ \perp & \text{otherwise (i.e., if } w^i \notin V_u \text{ or } h_i^*(w^i) \neq 0^K). \end{cases}$$

2. If there exists a $z \in \{0, 1\}^d$ such that at least $.6T$ of the strings z^i are equal to z , C' accepts and outputs z (which is necessarily unique). If there is no such z , C' rejects.

That C^\dagger can be implemented in the desired resource bounds is immediate from our assumptions. To prove correctness of the construction, consider any $u \in \text{Fav}$. By our choice of h_1^*, \dots, h_T^* , there is a set $I \subseteq [T]$ of size at least $.6T$ indices i such that h_i^* is good for u . For each such i , there is a string $\hat{w}^i \in V_u^{f(u)}$ for which $h_i^*(\hat{w}^i) = 0^K$. Define an assignment $\bar{w} = (w^1, \dots, w^T)$ to the nondeterministic gates of C' by letting

$$w^i := \begin{cases} \hat{w}^i & \text{if } i \in I, \\ 0^N & \text{otherwise.} \end{cases}$$

(Our use of the value 0^N here is arbitrary.) Observe that in the execution of $C^\dagger(u, \bar{w})$, we have $z^i = f(u)$ for each $i \in I$. Thus, C^\dagger accepts (u, \bar{w}) and outputs $f(u)$. This shows $f(u) \in F_{C'}(u)$.

To see that $F_{C'}(u) \subseteq \{f(u)\}$, consider any $z \neq f(u)$. Note that, for any $i \in I$ (for which h_i^* is good for u), we cannot have $z^i = z$ under any assignment to \bar{w} . Thus the number of i for which $z^i = z$ is at most $.4T$, so $C^\dagger(u, \bar{w})$ cannot accept with output z . This completes the proof of Lemma 2.8. \square

3 Stage-based analysis of direct-product computations

In this section we analyze the behavior of direct-product solvers on inputs (x^1, \dots, x^t) drawn from a known probability distribution \bar{D} over $\{0, 1\}^{n \times t}$. The case where \bar{D} is a t -fold product distribution will be of primary interest to us, although some of our lemmas will apply to non-product input distributions.

3.1 Valid outputs and the α, β sequences

Here we make some definitions that will be of central importance.

Definition 3.1 (j -valid outputs). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$ be a function, let $(x^1, \dots, x^t) \in \{0, 1\}^{n \times t}$, and let $z = (z_1, \dots, z_t) \in \{0, 1\}^{d \times t}$. For $j \in [t]$, say that z is j -valid for (x^1, \dots, x^t) , with respect to $f^{\otimes t}$, if*

$$z_\ell = f(x^\ell), \quad \text{for all } \ell \leq j.$$

When the target function f and the reference strings x^1, \dots, x^t are clear from the context, we will simply say that z is j -valid. Note that if z is j -valid for (x^1, \dots, x^t) then it is also j' -valid for $j' < j$. By convention, every $z \in \{0, 1\}^{d \times t}$ is said to be 0-valid.

If C is a probabilistic circuit taking a tuple of strings (x^1, \dots, x^t) as input (with each x^j of equal, predetermined length) and outputting a t -bit string, we say that a particular execution of C on input (x^1, \dots, x^t) is j -valid if it outputs some bitstring that is j -valid with respect to the inputs. We denote this event simply as $[C(x^1, \dots, x^t) \text{ is } j\text{-valid}]$.

Definition 3.2 (α, β sequences). *Let $n, t \in \mathbb{N}^+$ and let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$. Let $C : \{0, 1\}^{n \times t} \rightarrow \{0, 1\}^{d \times t}$ be a probabilistic circuit. Let \bar{D} be a distribution over $\{0, 1\}^{n \times t}$, and let $(\mathbf{x}^1, \dots, \mathbf{x}^t) \sim \bar{D}$. Define two sequences of random variables*

$$\alpha_0, \alpha_1, \dots, \alpha_t, \quad \beta_0, \beta_1, \dots, \beta_{t-1},$$

as follows. For $j \in [0, t]$, we let

$$\alpha_j := \Pr \left[C(\mathbf{x}^1, \dots, \mathbf{x}^t) \text{ is } j\text{-valid} \mid \mathbf{x}^1, \dots, \mathbf{x}^j \right] ,$$

with validity defined with respect to f . In other words, if $(\mathbf{x}^1, \dots, \mathbf{x}^j) = (x^1, \dots, x^j)$, then α_j takes on the value

$$\alpha_j = \Pr \left[C(\mathbf{x}^1, \dots, \mathbf{x}^t) \text{ is } j\text{-valid} \mid (\mathbf{x}^1, \dots, \mathbf{x}^j) = (x^1, \dots, x^j) \right] ,$$

where the probability is taken over the randomness in $\mathbf{x}^1, \dots, \mathbf{x}^t$ and in the randomness used by C . Similarly, for $j \in [0, t-1]$, define

$$\beta_j := \Pr \left[C(\mathbf{x}^1, \dots, \mathbf{x}^t) \text{ is } j\text{-valid} \mid \mathbf{x}^1, \dots, \mathbf{x}^{j+1} \right] .$$

Observe that $\alpha_0 = \beta_0 = 1$. Also, we have $\alpha_j, \beta_j \leq 1$ and $\alpha_{j+1} \leq \beta_j$.

3.2 A lemma on random sequences

Claim 3.3. 1. The function $\ln(1+x)$, defined on $(-1, \infty)$, satisfies

$$\ln(1+x) \leq \begin{cases} x - \frac{x^2}{6} & \text{if } x \in (-1, 1) , \\ (\ln 2)x & \text{if } x \geq 1 . \end{cases}$$

2. Consequently, for all $x \in (-1, \infty)$,

$$\ln(1+x) \leq x - \min \left\{ \frac{x^2}{6}, .3 \right\} .$$

Proof. (1) We have the Taylor series representation

$$\ln(1+x) = \sum_{n=1}^{\infty} (-1)^{n+1} \frac{x^n}{n} , \quad \text{valid for } x \in (-1, 1) .$$

This is an alternating series with terms monotonically decreasing in absolute value, so the Leibniz rule for alternating series tells us that $\ln(1+x) \leq x - \frac{x^2}{2} + \frac{x^3}{3}$ for $x \in (-1, 1)$. For such x we have $\frac{x^3}{3} \leq \frac{x^2}{3}$, which establishes the first half of item 1.

For the other half, observe that the function $f(x) = \ln(1+x)$ satisfies

$$f'(x) = \frac{1}{1+x} , \quad f''(x) = -\frac{1}{(1+x)^2} < 0 ,$$

with these expressions valid over its entire domain $(-1, \infty)$. Define the linear function $g(x) := (\ln 2)x$. We have

$$f(1) = \ln 2 = g(1) \quad \text{and} \quad f'(1) = \frac{1}{2} < \ln 2 = g'(1) ;$$

combining this with the fact that $f''(x) < 0$ everywhere implies that $f(x) \leq g(x)$ for $x \geq 1$. This completes the proof of item 1.

(2) From item 1 it immediately follows that for $x \in (-1, \infty)$,

$$\ln(1+x) \leq x - \min \left\{ \frac{x^2}{6}, (1 - \ln 2)|x| \right\}.$$

The first term in the min is smaller for $x \in (-1, 1]$, and for $x \geq 1$ we have $(1 - \ln 2)|x| > .3$. This proves item 2. \square

Lemma 3.4. *Let Y_1, Y_2, \dots, Y_T be (possibly dependent) nonnegative random variables satisfying $\mathbb{E}[Y_i] \leq 1$ for $i \in [T]$. Let $Y_{\text{prod}} := \prod_{i \in [T]} Y_i$. Let $\mathbf{i} \in_r [T]$ be chosen independently of Y_1, \dots, Y_t .*

1. *Suppose that there is some $q \in (0, 1]$ such that $Y_{\text{prod}} \geq q$ with probability 1. Then,*

$$\Pr[Y_{\mathbf{i}} \in [.99, 1.01]] \geq 1 - \frac{2^{16} \ln(1/q)}{T}.$$

2. *Suppose instead that for some $q \in (0, 1]$ and $\zeta \in [0, 1]$, we have $\Pr[Y_{\text{prod}} \geq q] \geq 1 - \zeta$. Then*

$$\Pr[Y_{\mathbf{i}} \in [.99, 1.012]] \geq 1 - \frac{2^{16} \ln(1/q)}{T} - 1.1 \cdot 2^{16} \zeta.$$

Proof. (1) Let Z_1, \dots, Z_T be defined by $Z_i := \ln(Y_i)$. Note that Z_i is well-defined since $Y_i > 0$, under our assumption $Y_{\text{prod}} \geq q$ in part 1. Letting $Z_{\text{sum}} := \sum_{i \in [T]} Z_i$, note that $Z_{\text{sum}} = \ln(Y_{\text{prod}})$. Our assumption in part 1 implies that $Z_{\text{sum}} \geq \ln q$. Thus,

$$\sum_{i \in [T]} \mathbb{E}[Z_i] = \mathbb{E}[Z_{\text{sum}}] \geq \ln q. \quad (2)$$

On the other hand, by applying Claim 3.3, item 2, we find that

$$Z_i \leq (Y_i - 1) - \min \left\{ \frac{(Y_i - 1)^2}{6}, .3 \right\}.$$

Taking expectations and using that $\mathbb{E}[Y_i] \leq 1$,

$$\mathbb{E}[Z_i] \leq -\mathbb{E} \left[\min \left\{ \frac{(Y_i - 1)^2}{6}, .3 \right\} \right].$$

Let $p_i := \Pr[Y_i \notin [.99, 1.01]]$. Then

$$\mathbb{E} \left[\min \left\{ \frac{(Y_i - 1)^2}{6}, .3 \right\} \right] \geq p_i \cdot \frac{(.01)^2}{6},$$

so that

$$\mathbb{E}[Z_i] \leq -\frac{p_i}{2^{16}}.$$

Combining this with Eq. (2) gives

$$\sum_{i \in [T]} p_i \leq -2^{16} \ln q = 2^{16} \ln(1/q),$$

which implies that

$$\Pr_{\mathbf{i} \in_r [T]} [Y_{\mathbf{i}} \notin [.99, 1.01]] = \frac{1}{T} \sum_{i \in [T]} p_i \leq \frac{2^{16} \ln(1/q)}{T},$$

as was to be shown.

(2) Note that if $\zeta > .001$, the assertion being made is trivial, since probabilities are always nonnegative. So let us assume that $\zeta \in [0, .001]$. Let (Y'_1, \dots, Y'_T) be random variables jointly distributed as (Y_1, \dots, Y_T) conditioned on the event $[Y_{\text{prod}} \geq q]$. That is, for any $(y_1, \dots, y_T) \in \mathbb{R}^T$, $\Pr[(Y'_1, \dots, Y'_T) = (y_1, \dots, y_T)] = \Pr[(Y_1, \dots, Y_T) = (y_1, \dots, y_T) | Y_{\text{prod}} \geq q]$. Let $c := 1 - \zeta$; our assumption in part 2 implies that $\mathbb{E}[Y_i] \geq c \cdot \mathbb{E}[Y'_i]$, which gives $\mathbb{E}[Y'_i] \leq 1/c$.

Now let $Y''_i := cY'_i$; we have $Y''_i > 0$ and $\mathbb{E}[Y''_i] \leq 1$. Let $Y''_{\text{prod}} := \prod_{i \in [T]} Y''_i$. We have $Y''_{\text{prod}} \geq c^T q$ always. Thus we may apply item 1 to the sequence Y''_1, \dots, Y''_T to find that

$$\Pr_{\mathbf{i} \in_r [T]} \left[Y'_{\mathbf{i}} \in \left[\frac{.99}{c}, \frac{1.01}{c} \right] \right] = \Pr_{\mathbf{i} \in_r [T]} [Y''_{\mathbf{i}} \in [.99, 1.01]] \geq 1 - \frac{2^{16} \ln \left(\frac{1}{c^T q} \right)}{T}. \quad (3)$$

Then it follows from our definitions that

$$\Pr_{\mathbf{i} \in_r [T]} \left[Y_{\mathbf{i}} \in \left[\frac{.99}{c}, \frac{1.01}{c} \right] \right] \geq c \cdot \Pr_{\mathbf{i} \in_r [T]} \left[Y'_{\mathbf{i}} \in \left[\frac{.99}{c}, \frac{1.01}{c} \right] \right].$$

Next, note that our assumption $c \in [.999, 1]$ implies that $[\frac{.99}{c}, \frac{1.01}{c}] \subseteq [.99, 1.012]$, and also that $1/c = 1/(1 - \zeta) \leq 1 + 1.01\zeta$. Using these facts and Eq. (3), we derive the bounds

$$\begin{aligned} \Pr_{\mathbf{i} \in_r [T]} [Y_{\mathbf{i}} \in [.99, 1.012]] &\geq \Pr_{\mathbf{i} \in_r [T]} \left[Y_{\mathbf{i}} \in \left[\frac{.99}{c}, \frac{1.01}{c} \right] \right] \\ &\geq c \cdot \left(1 - \frac{2^{16} \ln \left(\frac{1}{c^T q} \right)}{T} \right) \\ &\geq 1 - \zeta - \frac{2^{16} \ln \left(\frac{1}{c^T q} \right)}{T} \\ &\geq 1 - \zeta - \frac{2^{16} \ln \left(\frac{(1+1.01\zeta)^T}{q} \right)}{T} \\ &= 1 - \zeta - \frac{2^{16} T \cdot \ln(1 + 1.01\zeta)}{T} - \frac{2^{16} \ln(1/q)}{T} \\ &\geq 1 - 1.1 \cdot 2^{16} \zeta - \frac{2^{16} \ln(1/q)}{T}, \end{aligned}$$

where in the last step we used the bound $\ln(1 + x) \leq x$. This proves part 2. \square

3.3 Behavior of the α, β sequences

Lemma 3.5. Fix $n, d, t \in \mathbb{N}^+$. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$ be a function, and let $C : \{0, 1\}^{n \times t} \rightarrow \{0, 1\}^{d \times t}$ be a probabilistic circuit. Let $q \in (0, 1]$ and $\zeta \in [0, 1]$ be given.

1. Suppose that C is a $(1 - \zeta, q)$ -input-confident direct-product solver for $f^{\otimes t}$ with respect to an input distribution $\overline{\mathcal{D}}$ over $\{0, 1\}^{n \times t}$.

Let $\overline{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^t) \sim \overline{\mathcal{D}}$. Let $\alpha_0, \dots, \alpha_t, \beta_0, \dots, \beta_{t-1}$ be as in Definition 3.2, defined with respect to C and $\overline{\mathcal{D}}$. Let $\mathbf{j} \in_r [t]$ be sampled independently of $\overline{\mathbf{x}}$. Then with probability at least

$$1 - \frac{2^{16} \ln(1/q)}{t} - 2^{18} \zeta ,$$

we have

$$\alpha_{\mathbf{j}-1}, \beta_{\mathbf{j}-1} > 0 , \quad \left(\frac{\beta_{\mathbf{j}-1}}{\alpha_{\mathbf{j}-1}} \right) \in [.99, 1.012] , \quad \text{and} \quad \left(\frac{\alpha_{\mathbf{j}}}{\beta_{\mathbf{j}-1}} \right) \in [.99, 1] . \quad (4)$$

2. Now suppose instead that C is a worst-case q -direct-product solver for $f^{\otimes t}$. Let V be a finite set and let $\{\overline{\mathcal{D}}_v\}_{v \in V}$ be a set of distributions indexed by V , with each distribution over $\{0, 1\}^{n \times t}$. Let $\mathfrak{D}, \mathfrak{D}'$ be two distributions over $V \times [t]$, with the following properties:

- (a) If $(\mathbf{v}, \mathbf{j}) \sim \mathfrak{D}$, then \mathbf{v}, \mathbf{j} are independent and \mathbf{j} is uniform over $[t]$;
- (b) $\|\mathfrak{D} - \mathfrak{D}'\| \leq \gamma$, for some $\gamma \in [0, 1)$.

Consider the following experiment **Expt**(\mathfrak{D}'):

- (i) Sample $(\mathbf{v}', \mathbf{j}') \sim \mathfrak{D}'$;
- (ii) Sample $\overline{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^t) \sim \overline{\mathcal{D}}_{\mathbf{v}'}$;
- (iii) Let the sequence $\alpha_0, \dots, \alpha_t, \beta_0, \dots, \beta_{t-1}$ as in Definition 3.2 be defined with respect to $\overline{\mathcal{D}}_{\mathbf{v}'}$ and $\overline{\mathbf{x}}$.

Then with probability at least

$$1 - \frac{2^{16} \ln(1/q)}{t} - \gamma$$

over **Expt**(\mathfrak{D}'), we have

$$\left(\frac{\beta_{\mathbf{j}-1}}{\alpha_{\mathbf{j}-1}} \right) \in [.99, 1.02] \quad \text{and} \quad \left(\frac{\alpha_{\mathbf{j}}}{\beta_{\mathbf{j}-1}} \right) \in [.99, 1] . \quad (5)$$

Part 1 of Lemma will be applied to prove Theorem 6.1, our DPT for sampleable distributions, and also to prove part 2. Part 2 of our Lemma will be applied in the proof of Theorem 4.1, our DPT for worst-case direct-product solvers.

Proof. (1) Let $T := 2t$, and define random variables Y_1, \dots, Y_T as follows: for each $\ell \in [T]$, if $\ell = 2k + 1$ then let

$$Y_\ell := \begin{cases} \beta_k / \alpha_k & \text{if } \alpha_k \neq 0 , \\ 1 & \text{otherwise.} \end{cases}$$

If $\ell = 2k$, let

$$Y_\ell := \begin{cases} \alpha_k/\beta_{k-1} & \text{if } \beta_{k-1} \neq 0, \\ 1 & \text{otherwise.} \end{cases}$$

To illustrate the pattern, if all α_k, β_k are nonzero in some outcome then we have

$$Y_1 = \left(\frac{\beta_0}{\alpha_0}\right) = 1, Y_2 = \left(\frac{\alpha_1}{\beta_0}\right), Y_3 = \left(\frac{\beta_1}{\alpha_1}\right), \dots, Y_T = \left(\frac{\alpha_t}{\beta_{t-1}}\right).$$

Observe that if $\alpha_t > 0$, then all of $\alpha_0, \dots, \alpha_{t-1}, \beta_0, \dots, \beta_{t-1}$ are also positive, so that the product $Y_{\text{prod}} := \prod_{\ell \in [T]} Y_\ell$ equals $\alpha_t/\alpha_0 = \alpha_t$. By the definition of α_t and the input-confidence property of C , it follows that

$$\Pr[Y_{\text{prod}} \geq q] \geq 1 - \zeta.$$

For the even indices, we have $Y_{2k} \leq 1$ always. Also, we claim that $\mathbb{E}[Y_{2k+1}] = 1$. To see this, just observe that for $k \in [t(n)]$ we have the identity $\mathbb{E}[\beta_k|\alpha_k] = \alpha_k$, which tells us that Y_{2k+1} has expected value 1 conditioned on any nonzero value of α_k ; and if we condition on $[\alpha_k = 0]$, then $Y_{2k+1} = 1$ with certainty.

We have verified that the assumptions of Lemma 3.4, part 2 are satisfied by (Y_1, \dots, Y_T) ; we infer that if $\mathbf{i} \in_r [T]$,

$$\Pr[Y_{\mathbf{i}} \notin [.99, 1.012]] \leq \frac{2^{16} \ln(1/q)}{T} + 1.1 \cdot 2^{16} \zeta = \frac{2^{15} \ln(1/q)}{t} + 1.1 \cdot 2^{16} \zeta.$$

Recall that $T = 2t$ is even. If we instead choose $\hat{\mathbf{i}} \in_r \{1, 3, 5, \dots, T-1\}$, and then select $\hat{\mathbf{i}}' \in_r \{\hat{\mathbf{i}}, \hat{\mathbf{i}}+1\}$, then $\hat{\mathbf{i}}'$ is uniform over $[T]$ and we get the same bound for $\Pr[Y_{\hat{\mathbf{i}}'} \notin [.99, 1.012]]$. It follows that

$$\Pr[Y_{\hat{\mathbf{i}}} \notin [.99, 1.012] \vee Y_{\hat{\mathbf{i}}+1} \notin [.99, 1.012]] \leq \frac{2^{16} \ln(1/q)}{t} + 1.1 \cdot 2^{17} \zeta.$$

Now, note that $\mathbf{j} := \hat{\mathbf{i}}/2$ is distributed as a uniform element in $[t]$, and that the relations

$$\alpha_{\mathbf{j}-1}, \beta_{\mathbf{j}-1} > 0, \quad Y_{\hat{\mathbf{i}}} = \left(\frac{\beta_{\mathbf{j}-1}}{\alpha_{\mathbf{j}-1}}\right), \quad Y_{\hat{\mathbf{i}}+1} = \left(\frac{\alpha_{\mathbf{j}}}{\beta_{\mathbf{j}-1}}\right)$$

hold whenever $\alpha_t > 0$, which happens with probability at least $1 - \zeta$. Finally, if $\alpha_t > 0$ then $\beta_{\mathbf{j}-1} > 0$ and $\left(\frac{\alpha_{\mathbf{j}}}{\beta_{\mathbf{j}-1}}\right) \leq 1$. So Eq. (4) holds with probability at least

$$1 - \frac{2^{16} \ln(1/q)}{t} - 2^{18} \zeta.$$

This proves part 1 of the Lemma.

(2) First, suppose we run the alternative experiment **Expt**(\mathfrak{D}), which samples $(\mathbf{v}', \mathbf{j}')$ according to \mathfrak{D} rather than \mathfrak{D}' . Now, after conditioning upon any outcome $[\mathbf{v}' = v]$ of the first component, the index \mathbf{j} remains uniform over $[t]$ (by property (a) of \mathfrak{D}). Also, C is a $(1, q)$ -input-confident direct-product solver for $f^{\otimes t}$ with respect to the input distribution $\overline{\mathcal{D}}_v$. Thus we may set $\overline{\mathcal{D}} := \overline{\mathcal{D}}_v$ and $\zeta := 0$, and apply Lemma 3.5, part 1 to find that the probability that Eq. (5) holds is at least $1 - \frac{2^{16} \ln(1/q)}{t}$. Thus in **Expt**(\mathfrak{D}), Eq. (5) holds with probability at least $1 - \frac{2^{16} \ln(1/q)}{t}$.

Now let \mathbf{I}, \mathbf{I}' be the indicator variables for the events that Eq. (5) holds in **Expt**(\mathfrak{D}), **Expt**(\mathfrak{D}') respectively. Note that the two experiments are identically defined, except that the first draws a single sample from \mathfrak{D} while the second draws a sample from \mathfrak{D}' . Thus, $\|\mathbf{I} - \mathbf{I}'\|_{\text{stat}} \leq \|\mathfrak{D} - \mathfrak{D}'\|_{\text{stat}} \leq \gamma$, using property (b). This proves part 2 of the Lemma. \square

3.4 Analysis of worst-case direct-product solvers: the key experiment

Lemma 3.6. Fix $n, d, t, M \in \mathbb{N}^+$ and $q \in (0, 1], \zeta \in [0, 1]$. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$ be a function, and let \mathcal{D} be a distribution over $\{0, 1\}^n$. Suppose the probabilistic circuit C is a worst-case q -direct-product solver for $f^{\otimes t}$ with respect to the product input distribution $\mathcal{D}^{\otimes t(n)}$ over $\{0, 1\}^{n \times t}$. Consider the following experiment $\mathbf{Expt}^*(\mathcal{D})$:

1. Let $\mathbf{u} \in \{0, 1\}^n$ be sampled according to \mathcal{D} , and let $\mathbf{j} \in_r [t]$;
2. For each $j \in [t]$:
 - (i) let $s_j \in_r \{2, 3, \dots, M+1\}$;
 - (ii) Define a multiset S_j over $\{0, 1\}^n$, obtained by drawing s_j independent samples from \mathcal{D} ;
 - (iii) Let $\hat{S}_j := S_j \cup \{\mathbf{u}\}$ if $j = \mathbf{j}$; otherwise let $\hat{S}_j := S_j$;
 - (iv) Let $\mathbf{y}^j \in_r S_j$;
 - (v) Let $\hat{\mathbf{y}}^j := \mathbf{u}$ if $j = \mathbf{j}$; otherwise let $\hat{\mathbf{y}}^j := \mathbf{y}^j$.
3. Define the random variables

$$\alpha_0, \alpha_1, \dots, \alpha_t, \quad \beta_0, \beta_1, \dots, \beta_{t-1}, \quad \hat{\alpha}_0, \hat{\alpha}_1, \dots, \hat{\alpha}_t, \quad \hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_{t-1},$$

by the rules

$$\begin{aligned} \alpha_j &:= \Pr \left[C(\mathbf{y}^1, \dots, \mathbf{y}^t) \text{ is } j\text{-valid} \mid S_1, S_2, \dots, S_t, \mathbf{y}^1, \dots, \mathbf{y}^j \right], \\ \hat{\alpha}_j &:= \Pr \left[C(\hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^t) \text{ is } j\text{-valid} \mid \hat{S}_1, \hat{S}_2, \dots, \hat{S}_t, \hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^j \right], \\ \beta_j &:= \Pr \left[C(\mathbf{y}^1, \dots, \mathbf{y}^t) \text{ is } j\text{-valid} \mid S_1, S_2, \dots, S_t, \mathbf{y}^1, \dots, \mathbf{y}^{j+1} \right], \\ \hat{\beta}_j &:= \Pr \left[C(\hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^t) \text{ is } j\text{-valid} \mid \hat{S}_1, \hat{S}_2, \dots, \hat{S}_t, \hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^{j+1} \right]. \end{aligned}$$

Then with probability at least $1 - \frac{2^{16} \ln(1/q)}{t} - \frac{1}{M}$, we have

$$\left(\frac{\hat{\beta}_{\mathbf{j}-1}}{\alpha_{\mathbf{j}-1}} \right) \in [.98, 1.02] \quad \text{and} \quad \left(\frac{\hat{\alpha}_{\mathbf{j}}}{\hat{\beta}_{\mathbf{j}-1}} \right) \in [.99, 1]. \quad (6)$$

Proof. We aim to apply part 2 of Lemma 3.5 to the sequences $\hat{\alpha}_0, \dots, \hat{\alpha}_t, \hat{\beta}_0, \dots, \hat{\beta}_t$. Let $\mathbb{S}, \hat{\mathbb{S}}$ denote the random tuples (S^1, \dots, S^t) and $(\hat{S}^1, \dots, \hat{S}^t)$ respectively. Let \mathfrak{D} denote the distribution governing the tuple (\mathbb{S}, \mathbf{j}) , and let \mathfrak{D}' denote the distribution governing $(\hat{\mathbb{S}}, \mathbf{j})$. Each t -tuple $\bar{B} = (B_1, \dots, B_t)$ of multisets over $\{0, 1\}^n$ naturally defines a distribution $\bar{\mathcal{D}}_{\bar{B}}$ over elements $(\mathbf{z}^1, \dots, \mathbf{z}^t) \in \{0, 1\}^{n \times t}$, namely, the product distribution that independently chooses $\mathbf{z}^j \in_r B_j$ for each j .

The random variable \mathbf{j} is uniform over $[t]$ and independent of \mathbb{S} , so condition (a) of Lemma 3.5, part 2 is satisfied by \mathfrak{D} . For condition (b), note that $\hat{\mathbb{S}}$ is not fully independent of \mathbf{j} . However, we will show that $(\hat{\mathbb{S}}, \mathbf{j})$ is quite close in distribution to (\mathbb{S}, \mathbf{j}) :

Claim 3.7. $\left\| (\hat{S}_1, \dots, \hat{S}_t, \mathbf{j}) - (S_1, \dots, S_t, \mathbf{j}) \right\|_{\text{stat}} \leq \frac{1}{M}.$

Proof. We use a coupling argument. First, we generate a random multiset S_0 consisting of 2 independent samples from \mathcal{D} . Now define a random multiset \tilde{S} by letting

$$\tilde{S} := \begin{cases} \hat{S}_{\mathbf{j}} & \text{if } s_{\mathbf{j}} < M + 1, \\ S_0 & \text{if } s_{\mathbf{j}} = M + 1. \end{cases}$$

Note that $|\tilde{S}|$ is uniform over $\{2, 3, \dots, M + 1\}$, and its elements are distributed as independent samples from \mathcal{D} . Thus, if we form the random tuple

$$(S_1, \dots, S_{\mathbf{j}-1}, \tilde{S}, S_{\mathbf{j}+1}, \dots, S_t, \mathbf{j}),$$

we see that it is identically distributed to $(S_1, \dots, S_{\mathbf{j}}, \dots, S_t, \mathbf{j})$. Also, we have $\tilde{S} = \hat{S}_{\mathbf{j}}$ unless $s_{\mathbf{j}} = M + 1$, which happens only with probability $\frac{1}{M}$; and we always have $\hat{S}_j = S_j$ for all $j \neq \mathbf{j}$. This proves our Claim. \square

Thus condition (b) is satisfied by $\mathfrak{D}, \mathfrak{D}'$ with $\gamma := \frac{1}{M}$.

Observe that, after conditioning on \mathbb{S} , the sequence $(\mathbf{y}^1, \dots, \mathbf{y}^t)$ sampled in $\mathbf{Expt}^*(\mathcal{D})$ is distributed precisely according to $\overline{\mathcal{D}}_{\mathbb{S}}$. Similarly, after conditioning on $\hat{\mathbb{S}}$, the sequence $(\hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^t)$ is distributed according to $\overline{\mathcal{D}}_{\hat{\mathbb{S}}}$. We can therefore combine part 2 of Lemma 3.5 with Claim 3.7 to find that, with probability at least $1 - \frac{2^{16} \ln(1/q)}{t} - \frac{1}{M}$ over $\mathbf{Expt}^*(\mathcal{D})$, we have

$$\left(\frac{\hat{\beta}_{\mathbf{j}-1}}{\hat{\alpha}_{\mathbf{j}-1}} \right) \in [.99, 1.012] \quad \text{and} \quad \left(\frac{\hat{\alpha}_{\mathbf{j}}}{\hat{\beta}_{\mathbf{j}-1}} \right) \in [.99, 1].$$

Next we will need the following simple but important claim:

Claim 3.8. *With probability 1 we have*

$$\hat{\alpha}_{\mathbf{j}-1} = \frac{1}{s_{\mathbf{j}} + 1} \cdot \hat{\beta}_{\mathbf{j}-1} + \left(\frac{s_{\mathbf{j}}}{s_{\mathbf{j}} + 1} \right) \cdot \alpha_{\mathbf{j}-1}.$$

Proof. Consider any outcome of $\mathbf{Expt}^*(\mathcal{D})$, which is fully determined by the values of the random variables

$$(S_1, \dots, S_t, \mathbf{y}^1, \dots, \mathbf{y}^t, \mathbf{j}, \mathbf{u}) = (S_1^*, \dots, S_t^*, y^1, \dots, y^t, j, u).$$

Under these conditionings, observe that $\hat{\alpha}_{\mathbf{j}-1}$ equals the probability that the output of the computation $C(y^1, \dots, y^{j-1}, \mathbf{z}^j, \mathbf{z}^{j+1}, \dots, \mathbf{z}^t)$ is $(j-1)$ -valid, where

$$(\mathbf{z}^j, \mathbf{z}^{j+1}, \dots, \mathbf{z}^t) \in_r ((S_j^* \cup \{u\}) \times S_{j+1}^* \times \dots \times S_t^*).$$

This distribution on $(\mathbf{z}^j, \mathbf{z}^{j+1}, \dots, \mathbf{z}^t)$ can be equivalently realized as follows:

1. First, let $\mathbf{a} \in_r [s_j + 1]$ (noting here that $s_j = |S_j^*|$);
2. If $\mathbf{a} = s_j + 1$, let $\mathbf{z}^j := u$ and sample $(\mathbf{z}^{j+1}, \dots, \mathbf{z}^t) \in_r (S_{j+1}^* \times \dots \times S_t^*)$; otherwise choose $(\mathbf{z}^j, \mathbf{z}^{j+1}, \dots, \mathbf{z}^t) \in_r (S_j^* \times S_{j+1}^* \times \dots \times S_t^*)$.

On the other hand, $\hat{\beta}_{j-1}$ equals the probability that $C(y^1, \dots, y^{j-1}, u, \mathbf{z}^{j+1}, \mathbf{z}^{j+2}, \dots, \mathbf{z}^t)$ is $(j-1)$ -valid, where $(\mathbf{z}^{j+1}, \dots, \mathbf{z}^t) \in_r (S_{j+1}^* \times \dots \times S_t^*)$; and α_{j-1} equals the probability that $C(y^1, \dots, y^{j-1}, \mathbf{z}^j, \mathbf{z}^{j+1}, \dots, \mathbf{z}^t)$ is $(j-1)$ -valid, where $(\mathbf{z}^j, \dots, \mathbf{z}^t) \in_r (S_j^* \times \dots \times S_t^*)$. Combining these facts with our observations about $\hat{\alpha}_{j-1}$ yields the Claim. \square

Now, suppose that

$$\hat{\beta}_{j-1} \geq .99 \cdot \hat{\alpha}_{j-1} ,$$

which, as we have seen, occurs with high probability. Using Claim 3.8, this implies that

$$\frac{\hat{\beta}_{j-1}}{.99} \geq \frac{1}{s_j + 1} \cdot \hat{\beta}_{j-1} + \left(\frac{s_j}{s_j + 1} \right) \cdot \alpha_{j-1} ,$$

so that

$$\left(\frac{1}{.99} - \frac{1}{s_j + 1} \right) \cdot \frac{\hat{\beta}_{j-1}}{\alpha_{j-1}} \geq \left(\frac{s_j}{s_j + 1} \right) ,$$

which simplifies to

$$\frac{\hat{\beta}_{j-1}}{\alpha_{j-1}} \geq \left(\frac{99s_j}{100s_j + 1} \right) .$$

This is greater than .98. By a similar calculation, if $\hat{\beta}_{j-1} \leq 1.012 \cdot \hat{\alpha}_{j-1}$ then

$$\frac{\hat{\beta}_{j-1}}{\alpha_{j-1}} \leq \left(\frac{1012 \cdot s_j}{1000s_j - 12} \right) ,$$

which is less than 1.02 since $s_j \geq 2$. Combining our work, we conclude that with probability at least $1 - \frac{2^{16} \ln(1/q)}{t} - \frac{1}{M}$, Eq. (6) is satisfied. This completes the proof of Lemma 3.6. \square

4 The main DPT for worst-case direct-product solvers

In this section, we prove:

Theorem 4.1. *Let $n, d, t \in \mathbb{N}^+$, with $t \geq 100$. Let $q \in (0, 1]$, and let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$ be a function. Suppose there is a probabilistic circuit $C : \{0, 1\}^{n \times t} \rightarrow \{0, 1\}^{d \times t}$, such that C is a worst-case q -direct-product solver for $f^{\otimes t}$. Assume that*

$$q \geq \exp\left(-\frac{t}{3 \cdot 10^5}\right) .$$

Then, there is a nondeterministic mapping circuit C^ with n input bits and d output bits, such that $F_{C^*}(u) = \{f(u)\}$ for every $u \in \{0, 1\}^n$, and for which*

$$\text{size}(C^*) \leq \text{poly}(\text{size}(C)) .$$

Theorem 1.2 from the Introduction follows from this more general result:

Proof of Theorem 1.2. Let $f = \{f_n\}$ satisfy the hypothesis of Theorem 1.2. Suppose for contradiction's sake that there is a family of probabilistic circuits $\{C_n\}$ of size at most n^k for some $k > 0$, such that for all sufficiently large n and all $x \in \{0, 1\}^{n \times t(n)}$,

$$q := \Pr[C_n(x) = f_n^{\otimes t(n)}(x)] \geq \exp\left(-\frac{t(n)}{3 \cdot 10^5}\right).$$

Then by applying Theorem 4.1, to sufficiently large values of n , we find that there is a family $\{C_n^*\}$ of nondeterministic mapping circuits on n input bits, each with a single output bit, such that for all $u \in \{0, 1\}^n$,

$$F_{C_n^*}(u) = \{f_n(u)\},$$

and satisfying

$$\text{size}(C_n^*) \leq \text{poly}(\text{size}(C_n)) = \text{poly}(n^k) = \text{poly}(n).$$

By Proposition 2.3, we conclude that $f \in \text{NP}/\text{poly} \cap \text{coNP}/\text{poly}$. This contradicts our initial hypothesis, proving the Theorem. \square

Proof of Theorem 4.1. Let $\delta := \frac{10^5 \ln(1/q)}{t}$. Note that $\delta < .49$ by our assumption on q . As our main technical effort in this section, we will prove the following lemma:

Claim 4.2. *For any distribution \mathcal{D} over $\{0, 1\}^n$, there is a nondeterministic mapping circuit $C' = C'_\mathcal{D}$ that $(1 - \delta)$ -defines f with respect to \mathcal{D} , and for which $\text{size}(C') \leq \text{size}(C)^a$ for some $a > 0$.*

We defer the proof of Claim 4.2. Our Theorem follows from this Claim by a standard type of minimax argument (following [Yao77]), which we give next. Consider the following two-player, simultaneous-move, zero-sum game G :

- **Player 1:** Chooses an input $u \in \{0, 1\}^n$;
- **Player 2:** Chooses a nondeterministic mapping circuit C' with n input bits and d output bits, for which $\text{size}(C') \leq \text{size}(C)^a$;
- **Payoff to Player 2:** A reward of 1 if $F_{C'}(u) = \{f(u)\}$, or 0 otherwise.

By our Claim, for every mixed strategy $u \sim \mathcal{D}$ Player 1 may use, there is a pure strategy $C'_\mathcal{D}$ for Player 2 that gives expected payoff greater than .51 for Player 2. By the minimax theorem, there exists a single mixed strategy \mathcal{C} for Player 2 that achieves expected payoff at least .51 against *every* pure strategy u for Player 1. This is a distribution over circuits of size at most $\text{size}(C)^a$. Fix any $u \in \{0, 1\}^n$. By Lemma 2.1, if we take $H = O(n)$ large enough, and sample C_1, \dots, C_H independently from \mathcal{C} , then with probability greater than $1 - 2^{-n}$, we have $F_{C_i}(u) = \{f(u)\}$ for at least $.505 \cdot H$ indices $i \in [H]$. By a union bound, there exists a possible outcome C_1^*, \dots, C_H^* for which this occurs for every $u \in \{0, 1\}^n$.⁹

Let w^1, \dots, w^H denote the nondeterministic gate-sets for C_1^*, \dots, C_H^* respectively (their lengths may differ). We define C^* , with input $u \in \{0, 1\}^n$ and nondeterministic gates $\bar{w} = (w^1, \dots, w^H)$, as follows. C^* first evaluates $C_i^*(u, w^i)$ for each $i \in [H]$. If there is some $z \in \{0, 1\}^d$ such that at least $.505 \cdot H$ of the circuits C_i^* accept with output value z , then C^* accepts and outputs z ; otherwise, C^* rejects. It is immediate from our construction that $F_{C^*}(u) = \{f(u)\}$ for each $u \in \{0, 1\}^n$, and clearly we can achieve $\text{size}(C^*) \leq O(H \cdot \text{size}(C)^a) \leq \text{poly}(\text{size}(C))$.¹⁰ \square

⁹We remark that this step is an instance of the “strategy sparsification” technique of [LY94, Alt94].

¹⁰Recall here that $\text{size}(C) \geq \max\{nt, dt\}$, since we count input and output gates towards the circuit size.

Proof of Claim 4.2. Fix any distribution \mathcal{D} over $\{0, 1\}^n$. Our construction of the circuit $C' = C'_{\mathcal{D}}$ will require some preliminary work. First, refer to the experiment $\mathbf{Expt}^*(\mathcal{D})$ of Lemma 3.6, defined with respect to $\mathcal{D}, C, n, d, t, f$, and with $M := \lceil t / \ln(1/q) \rceil$. With the random variables $\alpha_j, \hat{\alpha}_j, \beta_j, \hat{\beta}_j$ as defined in that experiment, let us fix outcomes to the random variables

$$s_1, \dots, s_t, S_1, \dots, S_t, \mathbf{j}, \mathbf{y}^1, \dots, \mathbf{y}^t$$

that maximize the probability that Eq. (6) holds, where the probability is now taken over $\mathbf{u} \sim \mathcal{D}$. Let Λ denote the collection of random variables whose values we are fixing, and let $[\Lambda = \lambda]$ denote the particular setting we are making. When in $\mathbf{Expt}^*(\mathcal{D})$ we condition on $[\Lambda = \lambda]$, Eq. (6) holds with probability at least $1 - \frac{2^{16} \ln(1/q)}{t} - \frac{1}{M} \geq \frac{(2^{16}+1) \ln(1/q)}{t}$, which is $> 1 - \delta$ by our setting. Also, under our conditioning, \mathbf{u} remains undetermined and is distributed according to \mathcal{D} .

Note that our settings determine outcomes to $\alpha_0, \dots, \alpha_t, \beta_0, \dots, \beta_{t-1}$. The value $\alpha_{j^*-1} > 0$ in particular will be useful to us in defining our circuit. Abusing notation somewhat, we now let $s_1, \dots, s_t, S_1, \dots, S_t, \alpha_0, \dots, \alpha_t, \beta_0, \dots, \beta_{t-1}$ denote the fixed outcomes to these variables. Let $j^* \in [t]$ denote the fixed outcome to \mathbf{j} . For each $j \in [t]$, let

$$S_j = \{ y^{j,\ell} \}_{\ell \in [s_j]}$$

be an indexing of S_j (with some elements possibly appearing multiple times, according to their multiplicity in S_j). We define a mapping $\ell^* : [t] \rightarrow \mathbb{N}^+$ by the relation that, for our outcomes to $\mathbf{y}^1, \dots, \mathbf{y}^t$, we have

$$\mathbf{y}^j = y^{j,\ell^*(j)}.$$

As another important piece of non-uniform data in our construction, we will need to know the values taken by f on $y^{1,\ell^*(1)}, \dots, y^{j^*-1,\ell^*(j^*-1)}$. For $j \in [j^* - 1]$, we let

$$\hat{z}^j := f(y^{j,\ell^*(j)}).$$

Next we make some further definitions and observations. Suppose the probabilistic circuit C uses R bits of randomness (we may assume $R > 0$, or else we would have $q = 1$ and the Claim would be a triviality). Recall that $C^{\det}(x^1, \dots, x^t; r) : \{0, 1\}^{n \times t + R} \rightarrow \{0, 1\}$ denotes C considered as a deterministic circuit with random string r as part of its input. For any string $u \in \{0, 1\}^n$, we define a *viable certificate* for u as a tuple

$$w = (m_{j^*+1}, m_{j^*+2}, \dots, m_t, r) \in [s_{j^*+1}] \times \dots \times [s_t] \times \{0, 1\}^R$$

for which the first $(j^* - 1)$ length- d output blocks of the computation

$$C^{\det} \left(y^{1,\ell^*(1)}, \dots, y^{j^*-1,\ell^*(j^*-1)}, u, y^{j^*+1,m_{j^*+1}}, \dots, y^{t,m_t}; r \right) \quad (7)$$

equal $(\hat{z}^1, \dots, \hat{z}^{j^*-1})$. For such a w and $z \in \{0, 1\}^d$, we say that w is a *viable z -certificate* for u if the $(j^*)^{\text{th}}$ output block of the computation in Eq. (7) equals z , i.e., if in this computation C makes the “guess” that $f(u) \stackrel{?}{=} z$.

We fix some natural encoding of $[s_{j^*+1}] \times \dots \times [s_t] \times \{0, 1\}^R$ in which each element w has a unique representation as a binary string in $\{0, 1\}^N$; here, we may take $N \leq R + O(t \log_2 M) \leq \text{poly}(\text{size}(C))$. Let $V_u \subseteq \{0, 1\}^N$ denote the set of viable certificates for u , and for $z \in \{0, 1\}^d$, let $V_u^z \subseteq V_u$ denote the viable z -certificates for u . The sets V_u^z form a partition of V_u .

Claim 4.3. *Let us condition on $[\Lambda = \lambda]$ as above in $\mathbf{Expt}^*(\mathcal{D})$. Then,*

1. *For the random variable \mathbf{u} over $\{0,1\}^n$, the equality*

$$|V_{\mathbf{u}}| = 2^R \left(\prod_{j=j^*+1}^t s_j \right) \cdot \hat{\beta}_{j^*-1} \quad (8)$$

holds with probability 1.

2. *Also, we have the equality*

$$|V_{\mathbf{u}}^{f(\mathbf{u})}| = 2^R \left(\prod_{j=j^*+1}^t s_j \right) \cdot \hat{\alpha}_{j^*} , \quad (9)$$

and therefore

$$\frac{|V_{\mathbf{u}}^{f(\mathbf{u})}|}{|V_{\mathbf{u}}|} = \frac{\hat{\alpha}_{j^*}}{\hat{\beta}_{j^*-1}} . \quad (10)$$

Proof. (1) Condition further on any possible outcome $[\mathbf{u} = u]$ in $\mathbf{Expt}^*(\mathcal{D})$. Together with our prior conditioning $[\Lambda = \lambda]$, this determines the values of $\hat{S}_1, \dots, \hat{S}_t, \hat{\alpha}_1, \dots, \hat{\alpha}_t, \hat{\beta}_1, \dots, \hat{\beta}_t$. Under this conditioning, we see from the definition that

$$\hat{\beta}_{j^*-1} = \Pr \left[C(\hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^t) \text{ is } (j^* - 1)\text{-valid} \mid \hat{S}_1, \hat{S}_2, \dots, \hat{S}_t, \hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^{j^*} \right] \quad (11)$$

$$= \Pr \left[C(y^{1,\ell^*(1)}, \dots, y^{j^*-1,\ell^*(j^*-1)}, u, \mathbf{v}^{j^*+1}, \dots, \mathbf{v}^t) \text{ is } (j^* - 1)\text{-valid} \right] , \quad (12)$$

where we sample $(\mathbf{v}^{j^*+1}, \dots, \mathbf{v}^t) \in_r S_{j^*+1} \times \dots \times S_t$ (and where validity is with respect to f). Here, S_{j^*+1}, \dots, S_t are our fixed values under $[\Lambda = \lambda]$, and the probability in Eq. (12) is taken over $\mathbf{v}^{j^*+1}, \dots, \mathbf{v}^t$ and over the random bits r used by C .

Let us calculate the probability in Eq. (12). The selection of $(\mathbf{v}^{j^*+1}, \dots, \mathbf{v}^t)$ may be equivalently performed by choosing $(m_{j^*+1}, \dots, m_t) \in_r [s_{j^*+1}] \times \dots \times [s_t]$, and setting $\mathbf{v}^j := y^{j,m_j}$ for $j \in \{j^* + 1, \dots, t\}$. There are $\left(\prod_{j=j^*+1}^t s_j \right) \cdot 2^R$ possible outcomes to $(m_{j^*+1}, \dots, m_t, r)$, each one equally likely. The outcomes that cause the computation indicated in Eq. (12) to be $(j^* - 1)$ -valid are, under our definition, precisely those for which

$$(m_{j^*+1}, \dots, m_t, r) \in V_u .$$

Thus, under our conditioning $[\mathbf{u} = u]$ we have

$$\hat{\beta}_{j^*-1} = \frac{|V_u|}{2^R \left(\prod_{j=j^*+1}^t s_j \right)} , \quad \text{i.e.,} \quad |V_u| = 2^R \left(\prod_{j=j^*+1}^t s_j \right) \hat{\beta}_{j^*-1} .$$

As u was an arbitrary outcome to \mathbf{u} , we have proved part 1 of the Claim.

(2) Condition again on any possible outcome $[\mathbf{u} = u]$ in $\mathbf{Expt}^*(\mathcal{D})$. Then we have

$$\hat{\alpha}_{j^*} = \Pr \left[C(\hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^t) \text{ is } j^*\text{-valid} \mid \hat{S}_1, \hat{S}_2, \dots, \hat{S}_t, \hat{\mathbf{y}}^1, \dots, \hat{\mathbf{y}}^{j^*} \right] \quad (13)$$

$$= \Pr \left[C(y^{1,\ell^*(1)}, \dots, y^{j^*-1,\ell^*(j^*-1)}, u, \mathbf{v}^{j^*+1}, \dots, \mathbf{v}^t) \text{ is } (j^* - 1)\text{-valid} \right], \quad (14)$$

where again $(\mathbf{v}^{j^*+1}, \dots, \mathbf{v}^t) \in_r S_{j^*+1} \times \dots \times S_t$. Let these \mathbf{v}^j be generated by (m_{j^*+1}, \dots, m_t) just as in part 1. The outcomes that cause the computation in Eq. (14) to be j^* -valid are exactly those for which the following two conditions hold:

1. $(m_{j^*+1}, \dots, m_t, r) \in V_u$;
2. The $(j^*)^{th}$ output block of $C^{\det}(y^{1,\ell^*(1)}, \dots, y^{j^*-1,\ell^*(j^*-1)}, u, \mathbf{v}^{j^*+1}, \dots, \mathbf{v}^t; r)$ equals $f(u)$.

These outcomes are exactly those for which $(m_{j^*+1}, \dots, m_t, r) \in V_u^{f(u)}$. Then by a calculation following that in part 1, we can verify that Eq. (9) holds under $[\mathbf{u} = u]$. As u was arbitrary, Eq. (9) holds identically. Combining this with part 1 gives Eq. (10). \square

We have chosen the settings $[\Lambda = \lambda]$ so that Eq. (6) holds with high probability over $\mathbf{u} \sim \mathcal{D}$. Using Claim 4.3, this implies that $|V_{\mathbf{u}}^{f(u)}| \approx |V_{\mathbf{u}}|$ with high probability, i.e., almost all viable certificates for u correctly guess $f(u)$. Furthermore, by both parts of Claim 4.3 and the first condition of Eq. (6), both of $|V_{\mathbf{u}}^{f(u)}|, |V_{\mathbf{u}}|$ are (with high probability) approximately equal to $\rho := 2^R \left(\prod_{j=j^*+1}^t s_j \right) \alpha_{j^*-1}$; this latter quantity is determined by the setting $[\Lambda = \lambda]$, and does not depend on \mathbf{u} .

This motivates our strategy for building a nondeterministic mapping circuit to compute f on an input u sampled from \mathcal{D} . First, we choose a random hash function h from a strongly universal hash family with domain $U := \{0, 1\}^N$ (identified with \mathbb{F}_2^N), and with a range space of size determined by ρ . We consider an element of U “dead” unless it maps to the all-0 vector under h ; our aim is to “kill off” all of the viable certificates making incorrect guesses for $f(u)$, while leaving alive some viable certificate that makes a correct guess. Corollary 2.7, combined with our control on the sizes of $|V_{\mathbf{u}}|, |V_{\mathbf{u}}^{f(u)}|$, makes it possible to ensure this outcome with reasonable success probability. Nondeterminism will then allow us to guess and verify a live, viable certificate.

To make this strategy succeed with the required probability, we will perform repeated trials, selecting multiple hash functions and taking a majority vote over the trials. In the end we will remove the randomness in these trials and apply Lemma 2.8 to obtain our final nondeterministic circuit.

Let

$$K := \lceil \log_2 \rho - 4 \rceil,$$

and let

$$\mathcal{H}^{N,K} := \left\{ h_{A,v} : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^K \right\}_{A \in \mathbb{F}_2^{K \times N}, v \in \mathbb{F}_2^K}$$

be the strongly universal hash family given by Proposition 2.5. Note that $K = O(R + Mt) = O(R + t^2) \leq \text{poly}(\text{size}(C))$. Under our setting $[\Lambda = \lambda]$, say that $h_{A,v} \in \mathcal{H}^{N,K}$ is *good* for $u \in \{0, 1\}^n$ if

$$0^K \in h_{A,v} \left(V_u^{f(u)} \right) \setminus \left(\bigcup_{z \neq f(u)} h_{A,v} \left(V_u^z \right) \right).$$

Say that a string $u \in \{0, 1\}^n$, contained in the support of \mathcal{D} , is *favorable*, and write $u \in \text{Fav}$, if Eq. (6) holds under $[\Lambda = \lambda, \mathbf{u} = u]$. By our choice of λ we have $\Pr_{\mathbf{u} \sim \mathcal{D}}[\mathbf{u} \in \text{Fav}] \geq 1 - \delta$.

Claim 4.4. *Suppose u is favorable. If $(A, v) \in_r \mathbb{F}_2^{K \times N} \times \mathbb{F}_2^K$, then*

$$\Pr[h_{A,v} \text{ is good for } u] > .75 .$$

Proof. Define the set

$$V_u^{\text{wrong}} := \bigcup_{z \neq f(u)} V_u^z .$$

Under the conditioning $[\Lambda = \lambda, \mathbf{u} = u]$ in $\mathbf{Expt}^*(\mathcal{D})$, which defines outcomes to $\hat{\beta}_{j^*-1}, \hat{\alpha}_{j^*}$, part 2 of Claim 4.3 tells us that

$$\frac{|V_u^{f(u)}|}{|V_u|} = \frac{\hat{\alpha}_{j^*}}{\hat{\beta}_{j^*-1}} .$$

As u is favorable, this implies

$$\frac{|V_u^{f(u)}|}{|V_u|} \in [.99, 1] , \quad \text{and so} \quad \frac{|V_u^{\text{wrong}}|}{|V_u|} \leq .01 . \quad (15)$$

Next, part 1 of Claim 4.3 combined with Eq. (6) tells us that

$$\frac{|V_u|}{\rho} = \frac{|V_u|}{2^R \left(\prod_{j=j^*+1}^t s_j \right) \alpha_{j^*-1}} = \frac{\hat{\beta}_{j^*-1}}{\alpha_{j^*-1}} \in [.98, 1.02] . \quad (16)$$

Define real numbers θ, θ' by the relations

$$|V_u| = \theta \cdot 2^K, \quad |V_u^{\text{wrong}}| = \theta' \cdot 2^K .$$

By Eqs. (15)-(16) and our setting $K = \lceil \log_2 \rho - 2 \rceil$, we have the bounds

$$\theta \geq 8 \cdot .98 > 7 , \quad \theta \leq 16 \cdot 1.02 < 17 , \quad \theta' \leq .01\theta .$$

We apply Corollary 2.7, with $U := \mathbb{F}_2^N$, $U_{(i)} := V_u$, and $U_{(ii)} := V_u^{\text{wrong}}$ (there is no requirement that $U_{(i)}, U_{(ii)}$ be disjoint) to find that

$$\Pr_{(A,v)} [0^K \in h_{A,v}(V_u) \setminus h_{A,v}(V_u^{\text{wrong}})] \geq 1 - \theta^{-1} - .1\theta > .75 ,$$

where we used the fact that $1 - x^{-1} - .01x > .75$ for $x \in [7, 17]$. Thus $h_{A,v}$ is good for u with probability greater than .75. This proves the Claim. \square

Let $T := 40n$. We consider T -tuples

$$\bar{h} = (h_{A^1, v^1}, \dots, h_{A^T, v^T}) \in (\mathcal{H}^{N,K})^{\times T} .$$

Suppose we select the functions (A^i, v^i) independently as $(A^i, v^i) \in_r \mathbb{F}_2^{K \times N} \times \mathbb{F}_2^K$. For any $u \in \{0, 1\}^n$, let

$$X^+(u) := \sum_{i \in [T]} \mathbf{1} [h_{A^i, v^i} \text{ is good for } u]$$

Fix attention to any favorable u . By using Claim 4.4 and applying Lemma 2.1 to $X^+(u)$, we find that with probability at least $1 - \exp(-2(.15)^2 \cdot T) > 1 - 2^{-2n}$,

$$X_+(u) \geq .6T. \quad (17)$$

Taking a union bound, we conclude that with probability at least $1 - 2^n \cdot 2^{-2n} > 0$, Eq. (17) holds for *every* favorable $u \in \{0, 1\}^n$. Thus there exists a choice

$$\bar{h} = (h_1^*, \dots, h_T^*) = (h_{A^1, v^1}, \dots, h_{A^T, v^T})$$

such that Eq. (17) holds for every $u \in \text{Fav}$.

We are now ready to construct our circuit C' promised in the statement of Claim 4.2, by appealing to Lemma 2.8. First, each hash function h_i^* is determined by the values (N, K) and the pair (A^i, v^i) , which can be specified by $O(KN)$ bits. Using this description, we can directly compute $h_i^*(w)$ for a given $w \in \{0, 1\}^N$ using $O(KN)$ gates, as needed.

Let us next consider the problem of testing membership in the sets V_u, V_u^z . Our encoding of elements of $[s_{j^*+1}] \times \dots \times [s_t] \times \{0, 1\}^R$ as strings in $\{0, 1\}^N$ can be explicitly defined given the values s_{j^*+1}, \dots, s_t , and these are specifiable with $O(t \log_2 M) = O(t \log_2 t)$ bits. The multisets S_1, \dots, S_t we used, and the distinguished elements $y^{1, \ell^*(1)}, \dots, y^{t, \ell^*(t)}$, are specifiable with $O(tMn) = O(t^2n)$ bits; given these sets, we can efficiently determine whether w is a well-formed representative of an element of $[s_{j^*+1}] \times \dots \times [s_t] \times \{0, 1\}^R$ and, if so, map it to the input to C^{det} it represents under the correspondence presented in Eq. (7). By performing this mapping, evaluating C^{det} on the resulting input, and comparing its output to the hard-coded values $\hat{z}^1, \dots, \hat{z}^{j^*-1}$, we can determine which set V_u^z (if any) contains a given $w \in \{0, 1\}^N$. This can be performed by a circuit $C_{\text{Vtest}}(w, u)$ with at most $\text{size}(C) + \text{poly}(n + d + t) \leq \text{poly}(\text{size}(C))$ gates.

Thus, we may apply Lemma 2.8 to obtain a nondeterministic mapping circuit C^\dagger on n input gates, of size $\text{size}(C^\dagger) \leq \text{poly}(\text{size}(C))$, such that $F_{C^\dagger}(u) = \{f(u)\}$ for every $u \in \text{Fav}$. Recall that $\Pr_{\mathbf{u} \sim \mathcal{D}}[\mathbf{u} \in \text{Fav}] \geq 1 - \delta$. Thus C^\dagger $(1 - \delta)$ -defines f with respect to \mathcal{D} . So we may choose C^\dagger as our desired circuit C' , finding that $\text{size}(C^\dagger) \leq \text{size}(C)^a$ for an appropriately large $a > 0$. This proves Claim 4.2, completing the proof of Theorem 4.1. \square

5 The success probability of SAT solvers

For $k \in \mathbb{N}^+$, a k -CNF is a Boolean formula ψ of the form $\psi = \bigwedge_i D_i$, where each D_i is a disjunction of exactly k terms; here, a term is a literal or negated literal.

In this section we prove:

Theorem 5.1. [Theorem 1.5, restated] *Let $\gamma \in (0, 1)$. Suppose there is a PPT algorithm P_{solver} that, when given as input a description of a satisfiable 3-CNF formula Φ , of description length $|\langle \Phi \rangle| = N$, outputs a satisfying assignment to Φ with probability at least $q(N) := 2^{-N^\gamma}$.*

Then, $\text{NP} \subseteq \text{coNP}/\text{poly}$.

We will need the following straightforward lemma.

Lemma 5.2. *There is a deterministic polynomial-time algorithm M that takes as inputs descriptions of 3-CNF formulas $\psi^1, \psi^2, \dots, \psi^t$ and a value $s \geq 0$, and outputs a 3-CNF formula $\Psi^{(s)}$ that is satisfiable if and only if at least s of the ψ_j 's are satisfiable. Moreover, $\Psi^{(s)}$ contains designated variables $y = (y_1, \dots, y_t)$ such that for each $j \in [t]$, if $\Psi^{(s)}$ has a satisfying assignment with y_j set to 1, then ψ^j is satisfiable. If ψ^j has m_j clauses, then $\Psi^{(s)}$ has $O\left(t + \sum_{j \in [t]} m_j\right)$ clauses.*

Proof. Consider Boolean strings y of length t . It is a standard fact that for any $s \geq 0$ we can explicitly (in polynomial time) construct a fanin-two Boolean circuit $G^{(s)}(y)$ over $\{\wedge, \vee, \neg\}$ with $O(t)$ gates, that accepts exactly if the Hamming weight $\|y\|_1$ is at least s .

Next we rename the variables of each ψ^j if necessary to ensure that for $j \neq j'$, ψ^j and $\psi^{j'}$ contain disjoint sets of variables. Let x^j denote the input variables to ψ^j . Then for each j we construct a circuit $H^j(x^j, y_j)$ that accepts exactly if $[y_j = 0 \vee \psi^j(x^j) = 1]$. This H^j can be implemented with $O(m_j)$ gates.

Define a circuit $F^{(s)}(x^1, \dots, x^t, y)$ which outputs $G^{(s)}(y) \wedge \left(\bigwedge_{j \in [t]} H^j(x^j, y_j) \right)$. Using Cook's reduction again, we derive from $F^{(s)}$ a 3-CNF formula $\Psi^{(s)}(y_1, \dots, y_t, z_1, \dots, z_T)$, such that $\Phi^{(s)}(y, \cdot)$ is satisfiable exactly if there are settings to x^1, \dots, x^t such that $F^{(s)}(x^1, \dots, x^t, y) = 1$. This is the case exactly if $[\|y\|_1 \geq s] \wedge [\forall j \text{ with } y_j = 1, \psi^j \text{ is satisfiable}]$. We have established the correctness of our reduction. The number of clauses in $\Psi^{(s)}$ is bounded by a constant factor times the number of gates in $F^{(s)}$, which is $O\left(t + \sum_{j \in [t]} m_j\right)$, so $\Psi^{(s)}$ obeys the required size bound; and $\Psi^{(s)}$ can be constructed in polynomial time as needed. \square

The next lemma connects the task of satisfying 3CNFs to direct-product computations for 3SAT.

Lemma 5.3. *Let P_{solver} be a PPT algorithm that takes as input a description of a 3-CNF formula Φ . Assume that there is a nonincreasing function*

$$q : \mathbb{N} \rightarrow (0, 1] ,$$

such that, whenever Φ is a satisfiable 3-CNF of description length N , $P_{\text{solver}}(\langle \Psi \rangle)$ outputs a satisfying assignment to Φ with probability at least $q(N)$.

Then there is a PPT algorithm P_{decider} that takes as input an arbitrary-size list $\langle \langle \psi^1 \rangle, \dots, \langle \psi^t \rangle \rangle$ of descriptions of 3-CNFs, and outputs a length- t Boolean string $v = (v_1, \dots, v_t)$, giving guesses for the satisfiability status of each ψ^i , i.e., for the values $(\chi_{3\text{SAT}}(\langle \psi^1 \rangle), \dots, \chi_{3\text{SAT}}(\langle \psi^t \rangle))$.

P_{decider} has the following property: if each ψ^j has description length $|\langle \psi^j \rangle| = n$, then for some $N = O(nt \cdot \log_2(nt))$, we have

$$\Pr \left[(v_1, \dots, v_t) = (\chi_{3\text{SAT}}(\langle \psi^1 \rangle), \dots, \chi_{3\text{SAT}}(\langle \psi^t \rangle)) \right] \geq \frac{q(N)}{t+1} .$$

Proof. P_{decider} first guesses a value $s \in_r \{0, 1, \dots, t\}$. It then constructs the formula $\Psi^{(s)}$ defined by $(\psi^1, \dots, \psi^t, s)$ as in Lemma 5.2, and applies P_{solver} to $\langle \Psi^{(s)} \rangle$. If P_{solver} returns a satisfying assignment to $\Psi^{(s)}$, with setting (y_1, \dots, y_t) to the y -variables, then P_{decider} outputs $(v_1, \dots, v_t) := (y_1, \dots, y_t)$. Otherwise P_{decider} sets v_1, \dots, v_t arbitrarily.

P_{decider} is clearly polynomial-time. We analyze its success probability when given as input descriptions of 3-CNFs ψ^1, \dots, ψ^t , each of description length n . Let $S \subseteq [t]$ denote the set of indices j for which ψ^j is satisfiable. With probability at least $\frac{1}{t+1}$, P_{decider} sets $s := |S|$. Let us condition on this event. Now $\Psi^{(s)}$ is satisfiable, and (by our guarantee on the y -variables from Lemma 5.2) is satisfiable *only* by assignments for which $\|y\|_1 \geq s$ while $y_j = 0$ for all $j \notin S$. For such assignments, we have $(y_1, \dots, y_t) = (\chi_{3\text{SAT}}(\langle \psi^1 \rangle), \dots, \chi_{3\text{SAT}}(\langle \psi^t \rangle))$.

Thus, P_{decider} succeeds provided that P_{solver} finds a satisfying assignment to $\Psi^{(s)}$; by the success property of P_{solver} occurs with probability at least $q(N)$, where we define N as an upper bound

on the description length $|\langle \Psi^{(s)} \rangle|$. (Here we are using that $q(\cdot)$ is nonincreasing.) We may take $N = O(nt \cdot \log_2(nt))$, since $\Psi^{(s)}$ has $O(nt)$ clauses. Overall, our success probability is at least $q(N)/(t+1)$. This proves Lemma 5.3. \square

Proof of Theorem 5.1. We apply Lemma 5.3 to $P_{\text{solver}, q(\cdot)}$ to obtain a second PPT algorithm P_{decider} as described in that Lemma. Fixing any values of n, t , we can construct a probabilistic Boolean circuit $C_{n,t} : \{0,1\}^{n \times t} \rightarrow \{0,1\}^t$ such that on input $(\langle \psi^1 \rangle, \dots, \langle \psi^t \rangle) \in \{0,1\}^{n \times t}$, the circuit $C_{n,t}$ simulates an execution of P_{decider} on input $\langle \langle \psi^1 \rangle, \dots, \langle \psi^t \rangle \rangle$. Here, we may ensure $\text{size}(C_{n,t}) \leq \text{poly}(n+t)$ by using any standard translation of algorithms to circuits. From the property of P_{decider} guaranteed by Lemma 5.3, we find that $C_{n,t}$ is a worst-case q' -direct-product solver for $(\chi_{\text{3SAT}}|_{=n})^{\otimes t}$, where

$$q' = \frac{q(\kappa n t \cdot \log_2(nt))}{t+1} = 2^{-(\kappa n t \log_2(nt)^\gamma)^\gamma - \log_2(t+1)}, \quad (18)$$

for some absolute constant $\kappa > 0$.

Let $\rho := 2\gamma/(1-\gamma)$; we now fix $t := \lceil n^\rho \rceil$. Then $(nt \cdot \log_2(nt))^\gamma = O(n^{(1+\rho)\gamma} \cdot \log_2(nt)^\gamma) = O(n^{\gamma(1+\gamma)/(1-\gamma)} \cdot \log_2(nt)^\gamma) = o(n^\rho)$. Also, $\log_2(t+1) = O(\log_2 n) = o(n^\rho)$. Thus, for sufficiently large n we have

$$(\kappa n t \cdot \log_2(nt))^\gamma + \log_2(t+1) < \frac{t}{3 \cdot 10^5},$$

so that, using Eq. (18), we have

$$q' \geq 2^{-t/(3 \cdot 10^5)} \geq \exp\left(\frac{-t}{3 \cdot 10^5}\right).$$

We can therefore apply Theorem 4.1 to $C_{n,t}$, with $f := \chi_{\text{3SAT}}|_{=n}$, to obtain a nondeterministic mapping circuit C_n^* on n input bits, such that $F_{C_n^*}(\langle \psi \rangle) = \{\chi_{\text{3SAT}}(\langle \psi \rangle)\}$ for all formulas ψ of description length n , and for which $\text{size}(C_n^*) \leq \text{poly}(\text{size}(C_{n,t})) \leq \text{poly}(n+t) \leq \text{poly}_\gamma(n)$. It follows from Proposition 2.3 that $\text{3SAT} \in \text{coNP/poly}$. As 3SAT is NP-complete, this proves the Theorem. \square

6 The DPT for sampleable distributions

In this section we prove:

Theorem 6.1. *Let $n, d, t \in \mathbb{N}^+$ with $t \geq 100$. Let $f : \{0,1\}^n \rightarrow \{0,1\}^d$, and let $q \in (0,1]$. Let \mathcal{D} be a distribution over $\{0,1\}^n$ sampleable by a probabilistic circuit $C_{\mathcal{D}\text{-samp}}$. Suppose there is a probabilistic circuit $C : \{0,1\}^{n \times t} \rightarrow \{0,1\}^t$, such that C is a q -direct-product solver for $f^{\otimes t}$ with respect to $\mathcal{D}^{\otimes t}$.*

Let $\varepsilon \in (0,1)$, and assume that

$$q \geq \exp\left(-\frac{\varepsilon^{3/2}\sqrt{t}}{10^8}\right); \quad (19)$$

then, there is a nondeterministic mapping circuit C_ε^ on n input bits that $(1-\varepsilon)$ -defines f with respect to \mathcal{D} , and such that*

$$\text{size}(C_\varepsilon^*) \leq \text{poly}(\text{size}(C) + \text{size}(C_{\mathcal{D}\text{-samp}})).$$

In the size bound on C_ε^* above, there is no hidden dependence on ε . This result readily implies Theorem 1.3:

Proof of Theorem 1.3. Let $f = \{f_n\}$ satisfy the hypothesis of Theorem 1.3. Suppose for contradiction's sake that there is an infinite set $S \subseteq \mathbb{N}$, and a family of probabilistic circuits $\{C_n\}$ of size at most n^k for some $k > 0$, such that for all $n \in S$ and all $\mathbf{x} \sim \mathcal{D}_n^{\otimes t(n)}$,

$$q := \Pr[C_n(\mathbf{x}) = f_n^{\otimes t(n)}(\mathbf{x})] \geq \exp\left(-\frac{\varepsilon_n^{3/2} \sqrt{t(n)}}{10^8}\right).$$

Let $\{C_{\text{samp},n}\}$ be the polynomial-sized sampling circuit family, assumed to exist, for which $C_{\text{samp},n}$ samples \mathcal{D}_n . Suppose $\text{size}(C_{\text{samp},n}) \leq n^{k'}$. Then by applying Theorem 6.1 to each $n \in S$, we find that there is a family $\{C_{n,\varepsilon_n}^*\}_{n \in S}$ of nondeterministic mapping circuits on n input bits, each with a single output bit; for each $n \in S$, C_{n,ε_n}^* is a circuit that $(1 - \varepsilon_n)$ -defines f_n with respect to \mathcal{D}_n , and satisfies

$$\text{size}(C_{n,\varepsilon_n}^*) \leq \text{poly}(\text{size}(C_n) + \text{size}(C_{\text{samp},n})) = \text{poly}(n^k + n^{k'}) = \text{poly}(n).$$

This contradiction to our hypothesis proves the Theorem. \square

6.1 A “confidence-building” lemma

Our first step toward proving Theorem 6.1 will be to give a reduction which converts a direct-product solver into an input-confident direct-product solver. This will occupy us in Section 6.1, where our goal is to prove the following lemma.

Lemma 6.2. *Fix $n, d, t \in \mathbb{N}^+$ and $q \in (0, 1]$. Let \mathcal{D} be a distribution over $\{0, 1\}^n$, sampleable by a probabilistic circuit $C_{\mathcal{D}\text{-smp}}$ with s non-output gates. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$, and suppose the deterministic circuit C is a q -direct-product solver for $f^{\otimes t}$ with respect to the input distribution $\mathcal{D}^{\otimes t}$ over $\{0, 1\}^{n \times t}$.*

Fix any $t' \in [t]$. Let $k := \lfloor t/t' \rfloor$, and let $\zeta := \sqrt{(2 \ln 2) \log_2(1/q)/k}$. Assume that $\zeta < 1$. Then there is a probabilistic circuit

$$\tilde{C} : \{0, 1\}^{n \times t'} \rightarrow \{0, 1\}^{t'},$$

satisfying

$$\text{size}(\tilde{C}) \leq \text{size}(C) + st,$$

such that \tilde{C} is a $(1 - \zeta, .5q)$ -input-confident direct product solver for $f^{\otimes t'}$ with respect to the input distribution $\mathcal{D}^{\otimes t'}$.

Note the requirement that the reference input distribution be efficiently sampleable.

If \mathcal{D} is the uniform distribution on $\{0, 1\}^n$, then for $C_{\mathcal{D}\text{-smp}}$ we may take a circuit all of whose n gates are designated both as random gates and as output gates. Then $s = 0$, and for the circuit \tilde{C} obtained from Lemma 6.2 we have $\text{size}(\tilde{C}) = \text{size}(C)$. The added efficiency of the reduction in such cases is not important for our main results; we merely point it out.

For our work in proving Lemma 6.2 we will need some additional definitions and facts. Suppose B is a finite set, $k > 0$ is an integer, and $\bar{\nu}$ is a (possibly non-product) distribution over $B^{\times k}$. For $i \in [k]$, let $\bar{\nu}^{(i)}$ denote the distribution over B governing the i^{th} coordinate of $\bar{\nu}$. Also, if $A \subseteq B^{\times k}$

is assigned nonzero probability by ν , we let $\bar{\nu}_A$ denote the distribution of a random variable $\mathbf{u} \sim \bar{\nu}$ after conditioning on $[\mathbf{u} \in A]$. If $\bar{\nu}(A) = 0$ then $\bar{\nu}_A$ is left undefined.

We will use the following well-known, non-symmetric measure of difference between random variables (see [CT06] for more information).

Definition 6.3 (KL divergence). *The (binary) Kullback-Leibler divergence, or KL divergence between two distributions $\mathcal{R}, \mathcal{R}'$, denoted $D_{\text{KL}}(\mathcal{R}||\mathcal{R}')$, is defined as follows. If $\text{supp}(\mathcal{R}) \subseteq \text{supp}(\mathcal{R}')$, set*

$$D_{\text{KL}}(\mathcal{R}||\mathcal{R}') := \sum_{x \in \text{supp}(\mathcal{R})} \mathcal{R}(x) \cdot \log_2 \left(\frac{\mathcal{R}(x)}{\mathcal{R}'(x)} \right) ;$$

otherwise set $D_{\text{KL}}(\mathcal{R}||\mathcal{R}') := +\infty$.

Fact 6.4. [Raz98, Special case of Lemma 3.3] *Suppose μ is a distribution over a finite set B . Let $k > 0$, and let $\bar{\nu}$ be any distribution over $B^{\times k}$. Then,*

$$D_{\text{KL}}(\bar{\nu}||\mu^{\otimes k}) \geq \sum_{i \in [k]} D_{\text{KL}}(\bar{\nu}^{(i)}||\mu) .$$

Fact 6.5. [Raz98, Lemma 3.5] *Let B be a finite set, $k > 0$, and suppose $\bar{\pi}$ is any distribution over $B^{\times k}$. Suppose that $V \subseteq B^{\times k}$ satisfies $\bar{\pi}(V) > 0$. Then,*

$$D_{\text{KL}}(\bar{\pi}_V||\bar{\pi}) = -\log_2(\bar{\pi}(V)) .$$

The next corollary follows a similar use of Lemmas 6.4 and 6.5 in [Raz98].

Corollary 6.6. *Let μ be a distribution over the finite set B , and $k > 0$. Suppose that $A \subseteq B^{\times k}$ satisfies $\mu^{\otimes k}(A) \geq q > 0$. Let us use the notation $\bar{\nu} := (\mu^{\otimes k})_A$. If we let*

$$\gamma_i := D_{\text{KL}}(\bar{\nu}^{(i)}||\mu) ,$$

then

$$\frac{1}{k} \sum_{i \in [k]} \gamma_i \leq \frac{\log_2(1/q)}{k} .$$

Proof. First, Lemma 6.5, applied to $\bar{\pi} := \mu^{\otimes k}$, tells us that

$$D_{\text{KL}}(\bar{\nu}||\mu^{\otimes k}) = D_{\text{KL}}(\bar{\pi}_A||\bar{\pi}) = -\log_2(\mu^{\otimes k}(A)) \leq \log_2(1/q) .$$

Next, Lemma 6.5 tells us that $\sum_{i \in [k]} \gamma_i \leq D_{\text{KL}}(\bar{\nu}||\mu^{\otimes k})$. Combining these facts gives the result. \square

We will also use the following important relation between KL divergence and statistical distance [CT06] (see Lemma 11.6.1, p. 370).

Lemma 6.7 (Pinsker's inequality, stated for binary KL divergence). *For any random variables Z, Z' ,*

$$D(Z||Z') \geq \frac{2}{\ln 2} \cdot \|Z - Z'\|_{\text{stat}}^2 .$$

Lemma 6.8. *Let B be a finite set, and μ a distribution over B . Let $k \geq 2$ be an integer and let $A \subseteq B^{\times k}$. Suppose that A is “large” according to the product measure $\mu^{\otimes k}$; that is, $\mu^{\otimes k}(A) \geq q$ for some $q \in (0, 1]$.*

Consider the following experiment $\mathbf{Expt}^\dagger(\mu)$:

1. *Sample $\mathbf{b} = (\mathbf{b}^1, \dots, \mathbf{b}^k) \sim \mu^{\otimes k}$;*
2. *For $i \in [k]$, define τ_i , a random variable in $[0, 1]$ determined by \mathbf{b}^i , by*

$$\tau_i := \Pr \left[(\mathbf{b}^1, \dots, \mathbf{b}^k) \in A \mid \mathbf{b}^i \right] .$$

For $i \in [k]$, let $\zeta_i := \Pr[\tau_i < .5q]$. Then,

$$\frac{1}{k} \sum_{i \in [k]} \zeta_i \leq \sqrt{(2 \ln 2) \frac{\log_2(1/q)}{k}} .$$

In particular, there exists an $i \in [k]$ with $\zeta_i \leq \sqrt{(2 \ln 2) \log_2(1/q)/k}$.

Proof. We will prove that for any $i \in [k]$,

$$\zeta_i \leq \sqrt{(2 \ln 2) D_{\text{KL}}(\bar{\nu}^{(i)} \parallel \mu)} .$$

It will follow that

$$\frac{1}{k} \sum_{i \in [k]} \zeta_i \leq \frac{1}{k} \sum_{i \in [k]} \sqrt{(2 \ln 2) D_{\text{KL}}(\bar{\nu}^{(i)} \parallel \mu)} \leq \sqrt{\frac{2 \ln 2}{k} \sum_{i \in [k]} D_{\text{KL}}(\bar{\nu}^{(i)} \parallel \mu)} \leq \sqrt{(2 \ln 2) \frac{\log_2(1/q)}{k}} ,$$

where in the last two steps we used Jensen’s inequality and Corollary 6.6. This will establish the Lemma.

So fix attention to some $i \in [k]$. For notational simplicity, let us assume $i = 1$; the other cases are handled identically. Define

$$B_{(\text{ii})} := \left\{ b \in B : \Pr_{(\mathbf{b}^2, \dots, \mathbf{b}^k) \sim \mu^{\otimes(k-1)}} \left[(b, \mathbf{b}^2, \dots, \mathbf{b}^k) \in A \right] < .5q \right\} , \quad A' := A \cap \{(b^1, \dots, b^k) : b^1 \in B_{(\text{ii})}\} ,$$

and observe that we have

$$\zeta_1 = \mu(B_{(\text{ii})}) .$$

Also, it follows from the definition of $\bar{\nu} = (\mu^{\otimes k})_A$ that we have

$$\bar{\nu}^{(1)}(B_{(\text{ii})}) = \frac{\mu^{\otimes k}(A')}{\mu^{\otimes k}(A)} . \tag{20}$$

The numerator on the right-hand side of Eq. (20) is at most $\mu(B_{(\text{ii})}) \cdot (.5q)$ by definition of $B_{(\text{ii})}$, while the denominator is at least q , by our assumption on A . Thus,

$$\bar{\nu}^{(1)}(B_{(\text{ii})}) \leq .5\mu(B_{(\text{ii})}) ,$$

which implies that

$$\|\bar{\nu}^{(1)} - \mu\|_{\text{stat}} \geq .5\mu(B_{(\text{ii})}) . \quad (21)$$

On the other hand, Lemma 6.7 tells us that

$$\|\bar{\nu}^{(1)} - \mu\|_{\text{stat}} \leq \sqrt{\frac{\ln 2 \cdot D_{\text{KL}}(\bar{\nu}^{(1)} \parallel \mu)}{2}} .$$

Combining this with Eq. (21) gives

$$\zeta_1 = \mu(B_{(\text{ii})}) \leq \sqrt{(2 \ln 2) D_{\text{KL}}(\bar{\nu}^{(1)} \parallel \mu)} ,$$

as claimed. This proves Lemma 6.8. \square

Proof of Lemma 6.2. First we make a definition. Let (x^1, \dots, x^t) denote the t blocks of input variables to C , with each x^j consisting of n bits. Given a pair U, V of disjoint subsets of $[t]$, with U nonempty, and given a mapping $\Phi : V \rightarrow \{0, 1\}^n$ describing an assignment to the blocks $(x^j)_{j \in V}$, define a probabilistic circuit

$$C_{U,V,\Phi} : \{0, 1\}^{n \times |U|} \rightarrow \{0, 1\}^{|U|} ,$$

obtained from C , as follows. First, each input block x^j with $j \in V$ is fixed to the assignment $\Phi(j)$. Next, if the designated output gates of C are

$$\{g_{j,e}^*\}_{j \in [t], e \in [d]} ,$$

where $g_{j,e}^*$ tries to compute the e^{th} bit of $f(x^j)$, then the output gates of $C_{U,V,\Phi}$ are $\{g_{j,e}^*\}_{j \in U, e \in [d]}$. Finally, for each of the blocks $(x^j)_{j \in [t] \setminus (U \cup V)}$, we construct a separate copy of $C_{\mathcal{D}\text{-samp}}$, and identify x^j with the output gates of the corresponding copy of $C_{\mathcal{D}\text{-samp}}$. The inputs to each such copy of $C_{\mathcal{D}\text{-samp}}$ are regarded as random gates in $C_{U,V,\Phi}$; the blocks $(x^j)_{j \in U}$ are left as the designated input gates of $C_{U,V,\Phi}$.

Our plan is to take $\tilde{C} := C_{U^*, V^*, \Phi^*}$, for suitably chosen (U^*, V^*, Φ^*) . As a first step, note that we may write $t = kt' + \ell$, for some $0 \leq \ell < t'$. Set $V^* := \{j : kt' < j < t\}$, a possibly-empty set. For any $\Phi : V \rightarrow \{0, 1\}^n$, let $q_\Phi \in [0, 1]$ be the maximal value such that $C_{[t] \setminus V^*, V^*, \Phi}$ is a q_Φ -direct-product solver for $f^{\otimes(t-\ell)}$ with respect to the input distribution $\mathcal{D}^{\otimes(t-\ell)}$. By averaging over settings to Φ induced by outcomes of $\mathcal{D}^{\otimes \ell}$ to $(x^j)_{j \in V^*}$ and applying our assumption on C , we conclude there exists a Φ such that $q_\Phi \geq q$. Let us fix Φ^* as any such mapping.

Next, for $i \in [k]$, let U_i be the interval $\{(i-1)k+1, \dots, (i-1)k+t'\}$, so that $|U_i| = t'$ and the sets U_1, \dots, U_k, V^* form a partition of $[t]$. We will choose U^* as one of our sets U_i , and will apply Lemma 6.8 to find a choice of i with good properties. Define the finite set $B := \{0, 1\}^{n \times t'}$, and the distribution $\mu := \mathcal{D}^{\otimes t'}$ over B . Define

$$A := \{\bar{x} = (x^1, \dots, x^{kt'}) \in \{0, 1\}^{n \times kt'} : C_{[t] \setminus V^*, V^*, \Phi^*}(\bar{x}) = (f(x^1), \dots, f(x^{kt'}))\} .$$

Note that, as $([t] \setminus V^*) \cup V^* = [t]$, the circuit $C_{[t] \setminus V^*, V^*, \Phi^*}$ is deterministic (no copies of $C_{\mathcal{D}\text{-samp}}$ are introduced), so that A is well-defined. We may regard A as a subset of B^k ; by our choice of Φ^* we have

$$\mu^{\otimes k}(A) \geq q .$$

Lemma 6.8 tells us that there is an $i^* \in [k]$ for which, if we choose $(\mathbf{b}^1, \dots, \mathbf{b}^k) \sim \mu^{\otimes k}$, then with probability at least $1 - \sqrt{(2 \ln 2) \log_2(1/q)/k} = 1 - \zeta$, we have

$$\Pr \left[(\mathbf{b}^1, \dots, \mathbf{b}^k) \in A \mid \mathbf{b}^{i^*} \right] \geq .5q .$$

Thus, if we modify $C_{[t] \setminus V^*, V^*, \Phi^*}$ by supplying all blocks $(x^j)_{j \notin U_{i^*} \cup V^*}$ with an independent random sample from \mathcal{D} (not considered as part of the input), and by using $\{g_{j,e}^*\}_{j \in U_{i^*}, e \in [d]}$ as the designated output gates, we obtain a $(1 - \zeta, .5q)$ -input-confident direct-product solver for $f^{\otimes t'}$ with respect to the input distribution $\mathcal{D}^{\otimes t'}$ over the remaining input blocks $(x^j)_{j \in U_{i^*}}$. Now observe that the modification we have just described is precisely implemented by $C_{U_{i^*}, V^*, \Phi^*}$. This circuit has at most $\text{size}(C) + st$ gates, since we have obtained it from C by fixing some input gate values, adding fewer than t copies of $C_{\mathcal{D}\text{-smp}}$ (with output gates of these copies identified with gates in C), and changing the designated type of other gates. This proves Lemma 6.2. \square

6.2 A DPT for input-confident direct-product solvers

In the previous section we showed how to endow direct-product solvers with the input-confidence property; in this section we will show how to *exploit* this property. The DPT we show here for input-confident direct-product solvers may also be of interest in its own right. We prove:

Theorem 6.9. *Fix $n, d, t' \in \mathbb{N}^+$ and $\zeta \in [0, 1), q' \in (0, 1]$. Let \mathcal{D} be a distribution over $\{0, 1\}^n$, sampled by the circuit $C_{\mathcal{D}\text{-smp}}$. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^d$. Suppose the probabilistic circuit C is a $(1 - \zeta, q')$ -input-confident direct-product solver for $f^{\otimes t'}$ with respect to the input distribution $\mathcal{D}^{\otimes t'}$ over $\{0, 1\}^{n \times t'}$.*

Let $\delta := \frac{2^{16} \ln(1/q')}{t'} + 2^{18} \zeta$. Then there is a nondeterministic mapping circuit $C' : \{0, 1\}^n \rightarrow \{0, 1\}^d$ that $(1 - \delta)$ -defines f with respect to input distribution \mathcal{D} , and such that

$$\text{size}(C') \leq \text{poly}(\text{size}(C) + \text{size}(C_{\mathcal{D}\text{-smp}})) .$$

Proof. Let $\bar{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^{t'}) \sim \mathcal{D}^{\otimes t'}$. Let $\alpha_0, \dots, \alpha_t, \beta_0, \dots, \beta_{t'-1}$ be as in Definition 3.2, defined with respect to C, f , and $\bar{\mathbf{x}}$. Let $\mathbf{j} \in_r [t']$ be sampled independently of $\bar{\mathbf{x}}$. Now let us fix a $j^* \in [t']$ and settings

$$[\mathbf{j} = j^*, \mathbf{x}^1 = y^1, \dots, \mathbf{x}^{j^*} = y^{j^*}] \tag{22}$$

that maximize the probability that Eq. (4) holds. Let Λ denote the collection of variables whose values we are setting, and let $[\Lambda = \lambda]$ denote the event described in Eq. (22).¹¹ By Part 1 of Lemma 3.5, Eq. (4) holds with probability at least $1 - \delta$ after conditioning on $[\Lambda = \lambda]$. Our setting $[\Lambda = \lambda]$ determines the value of $\alpha_{j^*-1} > 0$, but (in general) do not determine $\beta_{j^*-1}, \alpha_{j^*}$. Also, under our conditioning, $\mathbf{x}^{j^*}, \dots, \mathbf{x}^{t'}$ remain undetermined and are distributed as independent samples from \mathcal{D} .

For $j \in [j^* - 1]$, let

$$\hat{z}^j := f(y^j) .$$

¹¹Note, here and in what follows, that our notation indicates a similarity with our work in the proof of Claim 4.2, but that significant differences are also present. For example, our setting $[\Lambda = \lambda]$ here plays an analogous role to the setting $[\Lambda = \lambda]$ in that proof, but the random variables being fixed are different.

Suppose that the circuit C has R random gates, and that the sampling circuit $C_{\mathcal{D}\text{-smp}}$ uses R_{smp} random gates; we may assume $R, R_{\text{smp}} > 0$. For any $u \in \{0, 1\}^n$, we define a *viable certificate* for u as a tuple

$$w = (r^{j^*+1}, \dots, r^{t'}, r) \in \{0, 1\}^{R_{\text{smp}} \times (t' - j^*) + R}$$

for which the first $(j^* - 1)$ length- d output blocks of the computation

$$C^{\text{det}}(y^1, \dots, y^{j^*-1}, u, C_{\mathcal{D}\text{-smp}}(r^{j^*+1}), \dots, C_{\mathcal{D}\text{-smp}}(r^{t'}); r) \quad (23)$$

equal $(\hat{z}^1, \dots, \hat{z}^{j^*-1})$. Set

$$N := R_{\text{smp}} \cdot (t' - j^*) + R.$$

We let $V_u \subseteq \{0, 1\}^N$ denote the viable certificates for u . For $z \in \{0, 1\}^d$, we say that $w \in V_u$ is a *viable z -certificate* for u if the $(j^*)^{\text{th}}$ output block of the computation in Eq. (23) equals z . Let V_u^z denote the viable z -certificates for u .

Claim 6.10. *Let us condition on $[\Lambda = \lambda]$ as above. Then,*

1. *Defining the random variable $\mathbf{u} := \mathbf{x}^j$ over $\{0, 1\}^n$, the equality*

$$|V_{\mathbf{u}}| = 2^N \cdot \beta_{j^*-1} \quad (24)$$

holds with probability 1.

2. *Also, we have the equality*

$$|V_{\mathbf{u}}^{f(\mathbf{u})}| = 2^N \cdot \alpha_{j^*}, \quad (25)$$

and therefore

$$\frac{|V_{\mathbf{u}}^{f(\mathbf{u})}|}{|V_{\mathbf{u}}|} = \frac{\alpha_{j^*}}{\beta_{j^*-1}}. \quad (26)$$

Proof. (1) Condition further on any possible outcome $[\mathbf{u} = u]$. Under this conditioning, we see that

$$\beta_{j^*-1} = \Pr \left[C(\mathbf{x}^1, \dots, \mathbf{x}^{t'}) \text{ is } (j^* - 1)\text{-valid} \mid \mathbf{x}^1, \dots, \mathbf{x}^{j^*} \right] \quad (27)$$

$$= \Pr \left[C(y^1, \dots, y^{j^*-1}, u, \mathbf{v}^{j^*+1}, \dots, \mathbf{v}^{t'}) \text{ is } (j^* - 1)\text{-valid} \right], \quad (28)$$

where validity is defined with respect to f , and where we sample $\mathbf{v}^{j^*+1}, \dots, \mathbf{v}^{t'}$ independently from \mathcal{D} . Equivalently, in Eq. (28) we may regard $\mathbf{v}^{j^*+1}, \dots, \mathbf{v}^{t'}$ as being sampled by independent copies of $C_{\mathcal{D}\text{-smp}}$. Now we may regard the probability in Eq. (28) as being taken over uniform random seeds $r^{j^*+1}, \dots, r^{t'}$ to $(t' - j^*)$ copies of $C_{\mathcal{D}\text{-smp}}$, and over a uniform setting $r \in \{0, 1\}^R$ to the random gates in C .

There are $2^{R_{\text{smp}} \times (t' - j^*) + R}$ equally-likely outcomes to $(r^{j^*+1}, \dots, r^{t'}, r)$. The outcomes that cause the computation indicated in Eq. (12) to be $(j^* - 1)$ -valid are, under our definition, precisely those for which $(r^{j^*+1}, \dots, r^{t'}, r) \in V_u$. Thus, under our conditioning $[\mathbf{u} = u]$ we have

$$\beta_{j^*-1} = \frac{|V_u|}{2^{R_{\text{smp}} \times (t' - j^*) + R}} = \frac{|V_u|}{2^N}.$$

This proves part 1 of the Claim.

(2) Condition again on any outcome $[\mathbf{u} = u]$. Under this conditioning, we see that

$$\alpha_{j^*} = \Pr \left[C(\mathbf{x}^1, \dots, \mathbf{x}^{t'}) \text{ is } j^* \text{-valid} \mid \mathbf{x}^1, \dots, \mathbf{x}^{j^*} \right] \quad (29)$$

$$= \Pr \left[C(y^1, \dots, y^{j^*-1}, u, \mathbf{v}^{j^*+1}, \dots, \mathbf{v}^{t'}) \text{ is } j^* \text{-valid} \right], \quad (30)$$

where we sample $\mathbf{v}^{j^*+1}, \dots, \mathbf{v}^{t'}$ independently from \mathcal{D} ; these may again be regarded as sampled by independent copies of $C_{\mathcal{D}\text{-smp}}$, with random gate-sets $r^{j^*+1}, \dots, r^{t'}$. Again let r denote the random gates of C . The settings to $(r^{j^*+1}, \dots, r^{t'}, r)$ that cause the computation in Eq. (12) to be j^* -valid are exactly those for which $(r^{j^*+1}, \dots, r^{t'}, r) \in V_u^{f(u)}$. Thus, conditioned on $[\mathbf{u} = u]$ we have $\alpha_{j^*} = |V_u|/2^N$. This proves part 2. \square

Next, let $K := \lceil (N - 4) + \log_2(\alpha_{j^*-1}) \rceil$, and let

$$\mathcal{H}^{N,K} := \left\{ h_{A,v} : \mathbb{F}_2^N \rightarrow \mathbb{F}_2^K \right\}_{A \in \mathbb{F}_2^{K \times N}, v \in \mathbb{F}_2^K}$$

be the hash family given by Proposition 2.5. Under the setting $[\Lambda = \lambda]$, say that $h_{A,v} \in \mathcal{H}^{N,K}$ is *good for* $u \in \{0,1\}^n$ if

$$0^K \in h_{A,v} \left(V_u^{f(u)} \right) \setminus \left(\bigcup_{z \neq f(u)} h_{A,v} (V_u^z) \right).$$

Say that a string $u \in \{0,1\}^n$, satisfying $\mathcal{D}(u) > 0$, is *favorable*, and write $u \in \text{Fav}$, if Eq. (4) holds under $[\Lambda = \lambda, \mathbf{x}^{j^*} = u]$. By our choice of λ , we have $\Pr_{\mathbf{u} \sim \mathcal{D}}[\mathbf{u} \in \text{Fav}] \geq 1 - \delta$.

Claim 6.11. *Suppose u is favorable. If $(A, v) \in_r \mathbb{F}_2^{K \times N} \times \mathbb{F}_2^K$, then*

$$\Pr[h_{A,v} \text{ is good for } u] > .75.$$

Proof. Define

$$V_u^{\text{wrong}} := \bigcup_{z \neq f(u)} V_u^z.$$

Under the conditioning $[\Lambda = \lambda, \mathbf{x}^{j^*} = u]$, which defines outcomes to $\beta_{j^*-1}, \alpha_{j^*}$, part 2 of Claim 4.3 tells us that

$$\frac{|V_u^{f(u)}|}{|V_u|} = \frac{\alpha_{j^*}}{\beta_{j^*-1}}.$$

As u is favorable, this gives

$$\frac{|V_u^{f(u)}|}{|V_u|} \in [.99, 1], \quad \text{and so} \quad \frac{|V_u^{\text{wrong}}|}{|V_u|} \leq .01.$$

Next we combine part 1 of Claim 6.10 with Eq. (4) to find that

$$\frac{|V_u|}{2^N \cdot \alpha_{j^*-1}} = \frac{\beta_{j^*-1}}{\alpha_{j^*-1}} \in [.99, 1.012];$$

we also have $2^K \in [2^{-4} \cdot 2^N \alpha_{j^*-1}, 2^{-3} \cdot 2^N \alpha_{j^*-1}]$, so that

$$\theta := \frac{|V_u|}{2^K} \in [8 \cdot .99, 16 \cdot 1.012] \subset [7, 17] \quad \text{and} \quad \theta' := \frac{|V_u^{\text{wrong}}|}{2^K} \leq .01\theta .$$

We apply Corollary 2.7 with $U := \mathbb{F}_2^N$, $U_{(i)} := V_u$, and $U_{(ii)} := V_u^{\text{wrong}}$; by the same calculations used to prove Claim 4.4, we find that

$$\Pr_{(A,v)} [0^K \in h_{A,v}(V_u) \setminus h_{A,v}(V_u^{\text{wrong}})] > .75 .$$

This proves Claim 6.11. \square

Let $T := 40n$. Suppose we select T hash functions $(h_{A^1,v^1}, \dots, h_{A^T,v^T})$ independently, with $h_{A^i,v^i} \in_r \mathcal{H}^{N,K}$. For any $u \in \{0,1\}^n$, let

$$X^+(u) := \sum_{i \in [T]} \mathbf{1} [h_{A^i,v^i} \text{ is good for } u] .$$

Consider any favorable u . By using Claim 4.4 and applying Lemma 2.1 to $X^+(u)$, we find that with probability at least $1 - \exp(-2(.15)^2 \cdot T) > 1 - 2^{-2n}$,

$$X^+(u) \geq .6T . \tag{31}$$

Then with positive probability, Eq. (31) holds for *every* favorable $u \in \{0,1\}^n$. Thus there exists some choice

$$\bar{h} = (h_1^*, \dots, h_T^*) = (h_{A^1,v^1}, \dots, h_{A^T,v^T})$$

such that Eq. (31) holds for every favorable u .

We can now apply Lemma 2.8, just as in the proof of Claim 4.2. We note that under our current settings, $N, K, T \leq \text{poly}(\text{size}(C) + \text{size}(C_{\mathcal{D}\text{-samp}}))$, and each hash function h_i^* can be evaluated using $O(KN)$ gates. The one novel element is that to check whether a string $w \in \{0,1\}^N$ lies in the set V_u (as defined here), and if so to determine which V_u^z contains it, now requires the evaluation of $C_{\mathcal{D}\text{-samp}}$ on $(t - j^*)$ given random seeds, as in Eq. (23). The total number of gates needed for $C_{V\text{test}}$ is polynomial in $\text{size}(C) + \text{size}(C_{\mathcal{D}\text{-samp}})$. We obtain a circuit C^\dagger of size $\leq \text{poly}(\text{size}(C) + \text{size}(C_{\mathcal{D}\text{-samp}}))$, such that $F_{C^\dagger}(u) = \{f(u)\}$ for every $u \in \text{Fav}$. As $\Pr_{\mathbf{u} \sim \mathcal{D}}[\mathbf{u} \in \text{Fav}] \geq 1 - \delta$, we conclude that C^\dagger $(1 - \delta)$ -defines f with respect to \mathcal{D} , and is a suitable choice for C' . This proves Theorem 6.9. \square

6.3 Proof of Theorem 6.1

Theorem 6.12. [Theorem 6.1, restated] Let $n, d, t \in \mathbb{N}^+$, with $t \geq 100$. Let $f : \{0,1\}^n \rightarrow \{0,1\}^d$, and let $q \in (0, 1]$. Let \mathcal{D} be a distribution over $\{0,1\}^n$ sampleable by a probabilistic circuit $C_{\mathcal{D}\text{-samp}}$. Suppose there is a probabilistic circuit $C : \{0,1\}^{n \times t} \rightarrow \{0,1\}^t$, such that C is a q -direct-product solver for $f^{\otimes t}$ with respect to $\mathcal{D}^{\otimes t}$.

Let $\varepsilon \in (0, 1)$, and assume that

$$q \geq \exp\left(-\frac{\varepsilon^{3/2}\sqrt{t}}{6 \cdot 10^8}\right) ; \tag{32}$$

then, there is a nondeterministic mapping circuit C_ε^* on n input bits that $(1 - \varepsilon)$ -defines f with respect to \mathcal{D} , and such that

$$\text{size}(C_\varepsilon^*) \leq \text{poly}(\text{size}(C) + \text{size}(C_{\mathcal{D}\text{-smp}})) .$$

Proof. Let us set

$$t' := \left\lfloor (t \cdot \log_2(1/q))^{1/3} \right\rfloor .$$

From Eq. (32), it easily follows that we have $1 \leq t' < t$. Thus we may apply Lemma 6.2 with our choice of t' , to obtain a probabilistic circuit

$$\tilde{C} : \{0, 1\}^{n \times t'} \rightarrow \{0, 1\}^{t'} ,$$

satisfying

$$\text{size}(\tilde{C}) \leq \text{size}(C) + \text{size}(C_{\mathcal{D}\text{-smp}}) \cdot t ,$$

such that \tilde{C} is a $(1 - \zeta, .5q)$ -input-confident direct product solver for $f^{\otimes t'}$ with respect to the input distribution $\mathcal{D}^{\otimes t'}$; here,

$$\zeta = \sqrt{\frac{(2 \ln 2) \log_2(1/q)}{\lfloor t/t' \rfloor}} < \sqrt{\frac{2 \log_2(1/q) t'}{t}} \leq \sqrt{\frac{2 \log_2(1/q)^{4/3} \cdot t^{1/3}}{t}} = \frac{\sqrt{2} \cdot \log_2(1/q)^{2/3}}{t^{1/3}} .$$

Next, we apply Theorem 6.9 to \tilde{C} , with $q' := .5q$; we obtain a nondeterministic mapping circuit $C' : \{0, 1\}^n \rightarrow \{0, 1\}^d$ that $(1 - \delta)$ -defines f with respect to input distribution \mathcal{D} , where

$$\delta = \frac{2^{16} \ln(2/q)}{t'} + 2^{18} \zeta < \frac{5 \cdot 10^5 \log_2(1/q)^{2/3}}{t^{1/3}}$$

(we have chosen t' so that the two terms in the expression for δ are equal up to constant factors). From Eq. (32) it follows that $\delta < \varepsilon$.

Also, as guaranteed by Theorem 6.9, we have

$$\text{size}(C') \leq \text{poly}(\text{size}(\tilde{C}) + \text{size}(C_{\mathcal{D}\text{-smp}})) \leq \text{poly}(\text{size}(C) + \text{size}(C_{\mathcal{D}\text{-smp}})) .$$

Thus we may take C' as our desired circuit C_ε^* . This proves Theorem 6.1. \square

Acknowledgments

I thank Avi Wigderson for helpful comments, and Robin Moser for a discussion of related questions.

References

- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.
- [ABG03] Amihoud Amir, Richard Beigel, and William I. Gasarch. Some connections between bounded query classes and non-uniform complexity. *Inf. Comput.*, 186(1):104–139, 2003.

- [Alt94] Ingo Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199, Supplement 1(0):339 – 355, 1994.
- [BH88] Samuel R. Buss and Louise Hay. On truth-table reducibility to sat and the difference hierarchy over np. In *3rd Structure in Complexity Theory Conference*, pages 224–233, 1988.
- [BLS84] Ronald V. Book, Timothy J. Long, and Alan L. Selman. Quantitative relativizations of complexity classes. *SIAM J. Comput.*, 13(3):461–487, 1984.
- [BLS85] Ronald V. Book, Timothy J. Long, and Alan L. Selman. Qualitative relativizations of complexity classes. *J. Comput. Syst. Sci.*, 30(3):395–413, 1985.
- [BT06] Andrej Bogdanov and Luca Trevisan. Average-case complexity. *Foundations and Trends in Theoretical Computer Science*, 2(1), 2006.
- [CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006.
- [Dru11] Andrew Drucker. A PCP characterization of AM. In *ICALP*, pages 581–592, 2011. Full version at <http://eccc.hpi-web.de/report/2010/019/>.
- [Dru12] Andrew Drucker. New limits to classical and quantum instance compression. In *53rd IEEE FOCS*, pages 609–618, 2012. Full version at <http://eccc.hpi-web.de/report/2012/112/>.
- [FL97] Uriel Feige and Carsten Lund. On the hardness of computing the permanent of random matrices. *Computational Complexity*, 6(2):101–132, 1997.
- [GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *21st ACM STOC*, pages 25–32, 1989.
- [GNW95] Oded Goldreich, Noam Nisan, and Avi Wigderson. On Yao’s XOR-lemma. *Electronic Colloquium on Computational Complexity (ECCC)*, TR95-050, 1995.
- [GNW11] Oded Goldreich, Noam Nisan, and Avi Wigderson. On Yao’s XOR-lemma. In Oded Goldreich, editor, *Studies in Complexity and Cryptography*, volume 6650 of *Lecture Notes in Computer Science*, pages 273–301. Springer, 2011.
- [Gol07] Oded Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 1 edition, 2007.
- [GS86] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *18th ACM STOC*, pages 59–68, 1986.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [HO02] Lane A. Hemaspaandra and Mitsunori Ogihara. *The Complexity Theory Companion*. Texts in Theoretical Computer Science. Springer, 2002.

- [IJKW10] Russell Impagliazzo, Ragesh Jaiswal, Valentine Kabanets, and Avi Wigderson. Uniform direct product theorems: Simplified, optimized, and derandomized. *SIAM J. Comput.*, 39(4):1637–1665, 2010. Earlier version in STOC ’08.
- [IW97] Russell Impagliazzo and Avi Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *29th ACM STOC*, pages 220–229, 1997.
- [Lev87] Leonid A. Levin. One-way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987.
- [LY94] Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *26th ACM STOC*, pages 734–740, 1994.
- [NRS99] Noam Nisan, Steven Rudich, and Michael E. Saks. Products and help bits in decision trees. *SIAM J. Comput.*, 28(3):1035–1050, 1999. Earlier version in FOCS ’94.
- [NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167, 1994. Earlier version in FOCS ’88.
- [O’D02] Ryan O’Donnell. Hardness amplification within NP. In *34th ACM STOC*, pages 751–760, 2002.
- [Pap94] Christos H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.
- [PP10] Ramamohan Paturi and Pavel Pudlák. On the complexity of circuit satisfiability. In Leonard J. Schulman, editor, *STOC*, pages 241–250. ACM, 2010.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. Earlier version in STOC ’95.
- [Sha03] Ronen Shaltiel. Towards proving strong direct product theorems. *Computational Complexity*, 12(1-2):1–22, 2003. Earlier version in CCC ’01.
- [SV10] Ronen Shaltiel and Emanuele Viola. Hardness amplification proofs require majority. *SIAM J. Comput.*, 39(7):3122–3154, 2010.
- [Tod91] Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991.
- [TV00] Luca Trevisan and Salil P. Vadhan. Extracting randomness from samplable distributions. In *FOCS*, pages 32–42. IEEE Computer Society, 2000.
- [VW08] Emanuele Viola and Avi Wigderson. Norms, XOR lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4(1):137–168, 2008. Earlier version in CCC ’07.
- [Wig97] Avi Wigderson. Derandomizing BPP, 1997. Lecture notes prepared by Ronen Shaltiel. <http://www.math.ias.edu/~avi/BOOKS/rand.pdf>.
- [Yao77] Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *18th IEEE FOCS*, pages 222–227, 1977.

- [Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *FOCS*, pages 80–91. IEEE Computer Society, 1982.
- [Yap83] Chee-Keng Yap. Some consequences of non-uniform conditions on uniform classes. *Theor. Comput. Sci.*, 26:287–300, 1983.