Instance-Wise Hardness versus Randomness Tradeoffs for Arthur-Merlin Protocols

Dieter van Melkebeek   Nicollas Mocelin Sdroievski

March 17, 2023

Abstract

A fundamental question in computational complexity asks whether probabilistic polynomial-time algorithms can be simulated deterministically with a small overhead in time (the BPP vs. P problem). A corresponding question in the realm of interactive proofs asks whether Arthur-Merlin protocols can be simulated nondeterministically with a small overhead in time (the AM vs. NP problem). Both questions are intricately tied to lower bounds. Prominently, in both settings *blackbox* derandomization, i.e., derandomization through pseudo-random generators, has been shown equivalent to lower bounds for decision problems against circuits.

Recently, Chen and Tell (FOCS'21) established near-equivalences in the BPP setting between *whitebox* derandomization and lower bounds for multi-bit functions against algorithms on almost-all inputs. The key ingredient is a technique to translate hardness into targeted hitting-sets in an instance-wise fashion based on a layered arithmetization of the evaluation of a uniform circuit computing the hard function f on the given instance.

In this paper we develop a corresponding technique for Arthur-Merlin protocols and establish similar near-equivalences in the AM setting. As an example of our results in the hardness to derandomization direction, consider a length-preserving function f computable by a nondeterministic algorithm that runs in time n^a. We show that if every Arthur-Merlin protocol that runs in time n^c for $c = O(\log^2 a)$ can only compute f correctly on finitely many inputs, then AM is in NP. Our main technical contribution is the construction of suitable targeted hitting-set generators based on probabilistically checkable proofs for nondeterministic computations.

As a byproduct of our constructions, we obtain the first result indicating that whitebox derandomization of AM may be equivalent to the existence of targeted hitting-set generators for AM, an issue raised by Goldreich (LNCS, 2011). Byproducts in the average-case setting include the first uniform hardness vs. randomness tradeoffs for AM, as well as an unconditional mild derandomization result for AM.

1 Introduction

The power of randomness constitutes a central theme in the theory of computing. In some computational settings, randomness is indispensable for any algorithmic solution. In others, it is provably needed for attaining efficiency. In yet others, the use of randomness leads to algorithms that run faster than all known deterministic ones, but the question remains open: Does an efficient deterministic algorithm exist?

In the context of decision problems, the key question is whether probabilistic polynomial-time algorithms with bounded error (the class BPP) can be simulated deterministically with a small overhead in time. In the realm of interactive verification protocols, the corresponding question asks whether Arthur-Merlin protocols (the class AM) can be simulated nondeterministically with a small overhead in time. In both settings, polynomial overhead is conjectured to suffice but even subexponential overhead remains open. Both settings have intricate connections to the quest for lower bounds, referred to as hardness vs. randomness tradeoffs. In some cases equivalences are known. We first describe the situation for BPP and then the one for AM, the focal point of this paper.

BPP setting. The first hardness vs. randomness tradeoffs were developed for *blackbox* derandomization, where a pseudo-random generator (PRG) produces, in an input-oblivious way, a small set of strings that "look random" to the process under consideration on every input of a given length. A long line of research established tight equivalences between *blackbox* derandomization of prBPP (the promise version of the class BPP) and *nonuniform* lower bounds for exponential-time classes. At the low end of the derandomization spectrum, subexponential-time blackbox derandomizations of prBPP are equivalent to super-polynomial circuit lower bounds for $\text{EXP} \doteq \text{DTIME}[2^{\text{poly}(n)}]$ [BFNW93]. At the high end, polynomial-time blackbox derandomizations of prBPP are equivalent to linear-exponential circuit lower bounds for $\text{E} \doteq \text{DTIME}[2^{O(n)}]$ [IW97]. A smooth interpolation between the two extremes exists and yields tight equivalences over the entire derandomization spectrum [Uma03]. The results are also robust in the sense that if the circuit lower bound holds at infinitely many input lengths (equivalent to the separation $\text{EXP} \not\subseteq \text{P/poly}$ at the low end), then the derandomization works at infinitely many input lengths, and if the circuit lower bound holds at almost-all input lengths, then the derandomization works at almost-all input lengths.

A uniformization of the underlying arguments led to equivalences between derandomizations that work on *most* inputs of a given length, and *uniform* lower bounds, i.e., lower bounds against algorithms. This derandomization setting is often referred to as the *average-case* setting.[1] At the low end, there exist subexponential-time simulations of BPP that work on all but a negligible fraction of the inputs of infinitely many lengths if and only if $\text{EXP} \not\subseteq \text{BPP}$ [IW01]. Unfortunately, the known construction does not scale well (see [TV07, CRTY20, CRT22] for progress toward an equivalence at the high end) and is not robust (a version for almost-all input lengths remains open). On the other hand, the result holds for blackbox derandomization as well as for general, "whitebox" derandomization, and implies an equivalence between blackbox and whitebox derandomization in this setting: If derandomization is possible at all, it can be done through pseudo-random generators.

This left open the setting of *whitebox* derandomizations that work for *almost all* inputs. For prBPP, such derandomizations are equivalent to the construction of *targeted* pseudo-random generators, which take an input x for the underlying randomized process, and produce a small set of

---

[1]The underlying distribution may be the uniform one or any other polynomial-time sampleable distribution.

strings that "look random" on that specific input x [Gol11]. Recently, Chen and Tell [CT21] raised the question of an equivalent lower bound condition, and proposed a candidate: *uniform* lower bounds for *multi-bit* functions (rather than usual decision problems) that hold on *almost-all inputs* in the following sense.

Definition 1 (Hardness on almost-all inputs). *A computational problem f is hard on almost-all inputs against a class of algorithms if for every algorithm A in the class there is at most a finite number of inputs on which A computes f correctly.*

Chen and Tell started from the following observation about derandomization to hardness at the high end of the spectrum.

Proposition 2 (Chen and Tell [CT21]). *If $\mathrm{prBPP} \subseteq \mathrm{P}$, then for every constant c there exists a length-preserving function f that is computable in deterministic polynomial time and is hard on almost-all inputs against $\mathrm{prBPTIME}[n^c]$.*

Remarkably, they also established a converse, albeit with an additional uniform-circuit depth restriction on the hard function f. Their approach naturally yields a targeted *hitting-set generator* (HSG), the counterpart of a pseudo-random generator for randomized decision processes with one-sided error (the class RP and its promise version prRP).

Theorem 3 (Chen and Tell [CT21]). *Let f be a length-preserving function f computable by logspace-uniform circuits of polynomial size and depth n^b for some constant b. If f is hard on almost-all inputs against $\mathrm{prBPTIME}[n^{b+O(1)}]$, where $O(1)$ denotes some universal constant, then $\mathrm{prRP} \subseteq \mathrm{P}$.*

Note that the hardness hypothesis of Theorem 3 necessitates the depth n^b of the uniform circuits computing the function f to be significantly less than their size. Otherwise, there exists even a deterministic algorithm that computes f in time $n^{b+O(1)}$.

The proof of Theorem 3 constructs a polynomial-time targeted hitting-set generator for prRP, which generically implies a polynomial-time targeted pseudo-random generator for prBPP, and thus that $\mathrm{prBPP} \subseteq \mathrm{P}$. Theorem 3 scales smoothly over the entire derandomization spectrum for prRP. Due to losses in the generic conversion from hitting-sets to derandomizations for two-sided error, the corresponding result for prBPP does not scale that well. In particular, a low-end variant of Theorem 3 for prBPP remains open. That said, the results are robust in a similar sense as above with respect to input lengths. In fact, the approach inherently yields a much higher degree of robustness because it effectuates a hardness vs. randomness tradeoff on an input-by-input basis, as we explain further in the paragraph below about our techniques.

Table 1 provides a qualitative overview of the lower bound equivalences for each of the three types of derandomization considered. We point out that, in the new setting of whitebox derandomizations that work on almost-all inputs, an actual equivalence along the lines of Chen and Tell [CT21] remains open due to the additional uniform-circuit depth requirement that is needed in the direction from hardness to derandomization. We refer to such results as *near-equivalences*. Follow-up works managed to obtain full-fledged equivalences in terms of other types of hardness, namely hardness of a computational problem related to Levin-Kolmogorov complexity [LP22a] and hardness in the presence of efficiently-computable leakage [LP22b].

Derandomization	Lower bound
blackbox, almost-all inputs	non-uniform
most inputs	uniform
whitebox, almost-all inputs	uniform, almost-all inputs

Table 1: Equivalences between various types of derandomization and lower bounds

AM setting. An equivalence corresponding to the first line of Table 1 is known throughout the entire spectrum [KvM02, MV05, SU05]. The role of EXP is now taken over by NEXP ∩ coNEXP, and the circuits are nondeterministic (or single-valued nondeterministic, or deterministic with oracle access to an NP-complete problem like SAT). The simulations use hitting-set generators for AM that are efficiently computable nondeterministically. Hitting-set generators are the natural constructs in the setting of AM because every Arthur-Merlin protocol can be efficiently transformed into an equivalent one with perfect completeness. As in the BPP setting, the lower bound equivalences for blackbox derandomization of prAM scale smoothly and are robust with respect to input lengths.

Regarding derandomizations that work on all but a negligible fraction of the inputs of a given length (the second line in Table 1), no hardness vs. randomness tradeoffs for AM were known prior to our work. What was known, are high-end results on derandomizations where no efficient nondeterministic algorithm can locate inputs on which the simulation is guaranteed to be incorrect [GSTS03, SU09]. Indeed, the authors of [GSTS03] explicitly mention the average-case setting and why their approach fails to yield average-case simulations that are correct on a large fraction of the inputs. The setting corresponding to the third line in Table 1 was not studied before.

Main results. As our main results, we obtain near-equivalences in this third setting, i.e., between whitebox derandomizations of Arthur-Merlin protocols that work on almost-all inputs, on the one hand, and hardness on almost-all inputs against Arthur-Merlin protocols, on the other hand.

We start from a similar observation in the derandomization to hardness direction as the one Chen and Tell made for BPP at the high end of the spectrum.

Proposition 4. *If* prAM ⊆ NP, *then for every constant c there exists a length-preserving function f that is computable in nondeterministic polynomial time with "a few" bits of advice, and is hard on almost-all inputs against* AMTIME[n^c].

We refer to Section 5.1 for the quantification of "a few."

Importantly, we are able to establish an almost-converse of Proposition 4. Under a slightly stronger hardness assumption, we construct a targeted hitting-set generator for prAM that is computable in nondeterministic polynomial time, yielding the following derandomization result.

Theorem 5. *Let f be a length-preserving function computable in nondeterministic time n^a for some constant a. If f is hard on almost-all inputs against* prAMTIME[n^c] *for $c = O((\log a)^2)$, where $O(\cdot)$ hides some universal constant, then*

$$\mathrm{prAM} \subseteq \mathrm{NP}.$$

Note that, in contrast to Theorem 3 in the BPP setting, Theorem 5 in the AM setting has no uniform-circuit depth restriction on the function f. Together with Proposition 4, Theorem 5

4

represents a near-equivalence between $\mathrm{prAM} \subseteq \mathrm{NP}$ and hardness on almost-all inputs of length-preserving[2] functions against Arthur-Merlin protocols. Whereas in the BPP setting, the remaining gap relates to uniform-circuit depth, in the AM setting the remaining gap relates to the advice and the technical distinction between AM and prAM protocols. We point out that the approaches in [LP22a] and [LP22b], which yield full-fledged equivalences in the BPP setting, do not seem compatible with the AM setting [Liu22].

Both Proposition 4 and Theorem 5 scale quite smoothly across the derandomization spectrum. The generalization of Theorem 5 has the following form: Let f be a length-preserving function computable in nondeterministic time $T(n)$. If f is hard on almost-all inputs against $\mathrm{prAMTIME}[t(n)]$, then $\mathrm{prAM} \subseteq \mathrm{NTIME}[\mathrm{poly}(T(n))]$. Intuitively, we may think of $t(n)$ as only slightly smaller than $T(n)$ for high-end results and much smaller for low-end results. Pushing our techniques as far as possible toward the low end, we obtain the following variant of Theorem 5.

Theorem 6. *Let f be a length-preserving function computable in nondeterministic exponential time. If f is hard on almost-all inputs against $\mathrm{prAMTIME}[n^{b(\log n)^2}]$ for all constants b, then for some constant c*

$$\mathrm{prAM} \subseteq \mathrm{NTIME}[2^{n^c}]. \tag{1}$$

As $\mathrm{prAM} \subseteq \mathrm{NEXP}$ trivially holds, the conclusion (1) of Theorem 6 represents the very low end of the derandomization spectrum. Note that a perfectly smooth scaling of Theorem 5 would only need a polynomial lower bound to arrive at the conclusion of Theorem 6, but the hypothesis of Theorem 6 requires a lower bound of $n^{\omega((\log n)^2)}$. We remark that the same discrepancy shows up in the current best-scaling uniform hardness vs. randomness tradeoffs for AM [SU09]. We refer to Theorem 27 in Section 4 for the full scaling and to Table 2 in the same section for other interesting instantiations.

Byproducts. Using our targeted hitting-set generators we are able to make progress on a number of related topics. We mention three representative ones here; more are described in the body of the paper.

First, there is the relationship between whitebox derandomization of prAM and the existence of targeted hitting-set generators for prAM. In the paper [Gol11] where Goldreich introduced targeted pseudo-random generators for prBPP and showed that their existence is equivalent to whitebox derandomization of prBPP, he asked about analogous results for prAM. To the best of our knowledge, there have been no prior results along those lines. We take a first step toward an equivalence in this setting.

Theorem 7. *If $\mathrm{prAMTIME}[2^{\mathrm{polylog}(n)}] \subseteq \mathrm{io\text{-}NEXP}$, then there exists a targeted hitting-set generator for prAM that yields the simulation $\mathrm{prAM} \subseteq \mathrm{io\text{-}NTIME}[2^{n^c}]/n^\epsilon$ for some constant c and all $\epsilon > 0$.*

Second, we establish the first hardness vs. randomness tradeoffs for Arthur-Merlin protocols in the average-case setting. Informally, under a high-end worst-case hardness assumption, we obtain

---

[2]The focus on length-preserving functions f in Proposition 4 and Theorem 5 is for concreteness. For Proposition 4 to hold, the number of output bits needs to grow with n in an efficiently computable fashion. For Theorem 5 any number of output bits suffices as long as there are not so many that the function f becomes trivially hard for Arthur-Merlin protocols running in time n^c.

nondeterministic polynomial-time simulations of prAM that are correct on all but a negligible fraction of the inputs.

Theorem 8. *If* $\mathrm{NTIME}[2^{an}] \cap \mathrm{coNTIME}[2^{an}] \not\subseteq \mathrm{BPTIME}[2^{(\log(a+1))^2 n}]_{||}^{\mathrm{SAT}}$ *for some constant* $a > 0$, *then for every problem in* prAM *and all* $e > 0$ *there exists a simulation of the problem in* NP *that is correct on all but a fraction* $1/n^e$ *of the inputs of length* n *for infinitely many lengths* n.

The class $\mathrm{BPTIME}[t(n)]_{||}^{\mathrm{SAT}}$ denotes probabilistic algorithms with bounded error that run in time $t(n)$ and can make parallel (i.e., non-adaptive) queries to an oracle for SAT. Theorem 8 answers a question in [GSTS03], which presents results in the different but related "pseudo" setting, where the simulation may err on many inputs of any given length, but no polynomial-time nondeterministic algorithm can pinpoint an error at that length. We remark that our technique also leads to identical results in the "pseudo" setting by replacing the hardness assumption with hardness against $\mathrm{AMTIME}[t(n)]$.

The model $\mathrm{prBPP}_{||}^{\mathrm{SAT}}$ was used as a proxy for prAM in the initial derandomization results for Arthur-Merlin protocols [KvM02] and is seemingly more powerful. However, derandomization results for prAM typically translate into similar derandomization results for $\mathrm{prBPP}_{||}^{\mathrm{SAT}}$. In particular, the conclusion $\mathrm{prAM} \subseteq \mathrm{NP}$ of Theorem 5 implies that $\mathrm{prBPP}_{||}^{\mathrm{SAT}} \subseteq \mathrm{P}_{||}^{\mathrm{SAT}}$, and the conclusion $\mathrm{prAM} \subseteq \mathrm{NTIME}[2^{n^c}]$ for some constant c in Theorem 6 implies that $\mathrm{prBPP}_{||}^{\mathrm{SAT}} \subseteq \mathrm{DTIME}[2^{n^c}]_{||}^{\mathrm{SAT}}$ for some constant c. In the case of Theorem 8, we argue that the hardness assumption implies simulations of $\mathrm{prBPP}_{||}^{\mathrm{SAT}}$ in $\mathrm{P}_{||}^{\mathrm{SAT}}$ of the same strength as the simulations of prAM in NP. This way, we obtain a hardness vs. randomness tradeoff in which the hardness model and the model to-be-derandomized match, namely probabilistic algorithms with bounded error and non-adaptive access to an oracle for SAT.

As our third byproduct, we present an unconditional mild derandomization result for AM in the average-case setting. By a *mild* derandomization of AM we mean a nontrivial simulation on Σ_2-machines. Recall that $\mathrm{AM} \subseteq \Pi_2 \mathrm{P}$, and proving that $\mathrm{AM} \subseteq \Sigma_2 \mathrm{P}$ is a required step if we hope to show that $\mathrm{AM} \subseteq \mathrm{NP}$. It is known that AM can be simulated (at infinitely many input lengths n) on Σ_2-machines that run in subexponential time and take n^c bits of advice for some constant c [Wil16]. It remains open whether AM can be simulated on Σ_2-machines in subexponential time with subpolynomial advice. Indeed, such a simulation for prAM would imply lower bounds against nondeterministic circuits that are still open [AvM17]. We show an unconditional subexponential-time and subpolynomial-advice Σ_2-simulation for prAM in the average-case setting.

Theorem 9. *For every problem in* prAM *and every constant* $\epsilon > 0$ *there exists a simulation of the problem in* $\Sigma_2 \mathrm{TIME}[2^{n^\epsilon}]/n^\epsilon$ *that is correct on all but a fraction* $1/n^e$ *of the inputs of length* n, *for all constants* e *and infinitely many lengths* n.

In fact, we can extend Theorem 9 to $\mathrm{prBPP}_{||}^{\mathrm{SAT}}$ in lieu of prAM.

Techniques. For our main result, we develop an instance-wise transformation of hardness into targeted hitting-sets tailored for AM. In the setting of BPP, Chen and Tell combine the Nisan-Wigderson pseudo-random generator construction [NW94] with the doubly-efficient proof systems of Goldwasser, Kalai, and Rothblum [GKR15] (as simplified in [Gol18]). The latter allows them to capture the computation of a uniform circuit of size T and depth d for f on a given input x by a downward self-reducible sequence of polynomials, which they use to instantiate the NW generator.

In case the derandomization of a one-sided error algorithm on a given input x fails, a bootstrapping strategy à la [IW01], based on a learning property of the NW generator, allows them to retrieve the value of $f(x)$ in time $O(d \cdot \text{polylog}(T))$. Thus, provided the depth d is small compared to the size T, either the derandomization on input x works or else the computation of $f(x)$ can be sped up.

A similar approach based on [GKR15] applies to the AM setting by replacing the NW construction with a hitting-set generator construction for AM that also has the learning property. Like in the BPP setting, the construction is only of interest when the circuits for f have relatively small depth. Moreover, the construction can only handle a limited amount of nondeterminism in the computation for f, whereas the direction from derandomization to hardness seems to require more.

In order to remedy both shortcomings, we develop a new method to extract hardness from a nondeterministic computation on a given input x, based on probabilistically checkable proofs rather than [GKR15]. The soundness of our method presupposes some type of resilience of the underlying regular pseudo-random generator. The required property was first identified and used by Gutfreund, Shaltiel and Ta-Shma [GSTS03] for the Miltersen-Vinochandran generator MV [MV05], and later by Shaltiel and Umans [SU09] for their recursive variant of the MV generator, RMV. We combine RMV with the probabilistically checkable proofs of Ben-Sasson, Goldreich, Harsha, Sudan, and Vadhan [BGH$^+$06] to transform hardness into pseudo-randomness for AM in an instance-wise fashion, without any uniform-circuit depth restriction or limitation on the amount of nondeterminism.

We highlight one strong feature of *all* instance-wise approaches. If the hardness condition holds on almost-all inputs, then the derandomization works on almost-all inputs. This is the setting in which we stated the results of Chen and Tell and our main results. Similarly, if the hardness condition holds on all inputs of a given length, then the derandomization works on all inputs of that length. This is the robustness property that we alluded to earlier. However, an instance-wise approach yields much more, including average-case derandomization results: To obtain a nondeterministic simulation for some prAM problem that works with high probability over any given distribution, it suffices to assume that every prAM protocol can only compute the hard function f with low probability over that same distribution.

Our derandomization-to-hardness result follows by diagonalization, as does the one by Chen and Tell. To obtain our byproducts, we combine our targeted hitting-set generator with several other ingredients, including diagonalization, the "easy-witness" method and traditional hardness vs. randomness tradeoffs. Our average-case derandomization results require a modification of our targeted HSG so that it respects a stronger resilience property. Along the way to our unconditional mild derandomization result, we establish an "easy witness lemma" for Σ_2 computations, which may be of independent interest.

Organization. In Section 2, we develop the ideas behind our results and relate them to existing techniques. We start the formal treatment in Section 3 with definitions, notation, and other preliminaries. In Section 4, we construct our targeted HSG and establish our hardness-to-derandomization results that make use of it (Theorems 5 and 6). Section 5 presents the derandomization-to-hardness side of our near-equivalence, as well as a proof of our byproduct on derandomization to targeted hitting-set generators (Theorem 7). In Section 6, we derive our derandomization byproducts under uniform worst-case hardness (the average-case simulation of Theorem 8 as well as a simulation that works on all inputs of infinitely many lengths). Section 7 contains our unconditional mild

derandomization result for AM (Theorem 9).

2 Technical overview

In this section, we start with an overview of techniques used in prior hardness vs. randomness tradeoffs for BPP and AM in a way that facilitates a high-level exposition of our main hardness-to-derandomization result for AM. We also provide the intuition for our derandomization-to-hardness result and for our byproducts.

2.1 Main results

We start with an overview of the techniques used for hardness-to-derandomization results in the traditional setting for BPP (lines 1 and 2 in Table 1), followed by those in the new setting (line 3 in Table 1). We then transition to AM, discuss the additional challenges, the known techniques in the traditional setting and, finally, our results in the new setting.

Traditional setting for BPP. The key ingredient in all known hardness vs. randomness tradeoffs is a pseudo-random generator construction G that takes a function h as an oracle and produces a pseudo-random distribution G^h with the following property: Any statistical test D that distinguishes G^h from uniform suffices as an oracle to efficiently learn h approximately from a small number of queries. Thus, if G^h does not "look random" to an efficient randomized process A on an input x, an approximation to h can be reconstructed efficiently when provided with x and the values of h on a small number of points, as well as oracle access to the distinguisher $D(r) = A(x, r)$, where $A(x, r)$ denotes the output of A on input x and random-bit string r. If the function h can be self-corrected (e.g., by being random self-reducible or by its truth-table being a codeword in a locally-correctable error correcting code), then the exact function h can be reconstructed efficiently.

In order to obtain hardness vs. randomness tradeoffs from pseudo-random generator constructions with the learning property, two questions need to be addressed:

1. How to obtain the distinguishers D?

2. How to obtain the answers to the learning queries?

The first question asks how to find inputs x on which the process A is not fooled by G^h. In the non-uniform setting such an input can be included in the advice. In the uniform setting for BPP, such inputs can be found by sampling x at random and testing for a difference in behavior of $D \doteq A(x, \cdot)$ between the uniform and the pseudo-random distributions, which can be done in prBPP.

Regarding the second question, in the non-uniform setting, the answers to the learning queries can also be provided as advice. In the uniform setting, [IW01] employs a function h that is not only random self-reducible but also downward self-reducible, and uses the downward self-reduction to answer the learning queries for length n by evaluating the circuit that resulted from the reconstruction for length $n-1$. This bootstrapping strategy presupposes that the reconstruction works at almost-all inputs lengths. This is why we only know how to obtain simulations that are correct at infinitely many input lengths in the uniform setting for BPP.

8

New setting for BPP. In the setting of line 3 in Table 1, the role of pseudo-random generators is taken over by *targeted* pseudo-random generators. Whereas PRGs are oblivious to x (beyond its length), targeted PRGs take x as an input and are only supposed to fool the randomized process on that particular x. This approach obviates the problem of obtaining the distinguisher D (question 1 above) as we can use $D = A(x, \cdot)$ for the given x. Targeted PRGs can be constructed from a PRG G by instantiating G with an oracle $h = h_x$ that depends on x. This raises a third question in the application of a PRG for hardness vs. randomness tradeoffs:

3. How to obtain the function h_x from x?

Chen and Tell [CT21] use the doubly-efficient proof systems of Goldwasser, Kalai, and Rothblum [GKR15] (as simplified in [Gol18]) to obtain h_x from x and combine it with the Nisan-Wigderson pseudo-random generator construction [NW94]. The GKR proof system takes a logspace-uniform family of circuits of size $T(n)$ and depth $d(n)$ computing a (multi-bit) Boolean function f, and transforms the circuit for f on a given input x into a downward self-reducible sequence of multivariate low-degree polynomials $\hat{g}_{x,0} \ldots, \hat{g}_{x,d'(n)}$ where $d'(n) = O(d(n) \log(T(n)))$. The polynomial $\hat{g}_{x,0}$ is efficiently computable at any point given input x, and the value of $f(x)$ can be extracted efficiently from $\hat{g}_{x,d'(n)}$. We refer to the sequence of polynomials as a *layered arithmetization* of the circuit for f on input x.

Chen and Tell instantiate the NW generator with the Hadamard encoding of each of the polynomials $\hat{g}_{x,i}$ as the function $h = h_{x,i}$, and follow a bootstrapping strategy similar to [IW01] to construct $\hat{g}_{x,d'(n)}$ from $\hat{g}_{x,0}$. For the strategy to work, the NW reconstructor needs to succeed at every level. This is the reason why Chen and Tell only end up with a (targeted) *hitting-set* generator rather than a *pseudo-random* generator. The time required by the bootstrapping process is proportional to the number of layers and thus to the depth $d(n)$ of the circuit computing f. By setting the parameters of the arithmetization appropriately, the dependency on the size $T(n)$ is only polylogarithmic. This is what enables the reconstruction to compute $f(x)$ very quickly as long as the depth $d(n)$ is not too large.

Liu and Pass [LP22b] also use the NW generator but obtain h_x as an encoding of the value of $f(x)$ itself, where f is an almost-all inputs leakage-resilient hard function (a function that remains hard even if some efficiently-computable information about $f(x)$ is leaked to an attacker). The answers to the learning queries are provided as part of the information about $f(x)$ that is leaked, which allows them to reconstruct $f(x)$ directly and efficiently. This approach leads to a (targeted) *pseudo-random* generator since it only involves a single instantiation of the NW generator. Reversing the hardness-to-derandomization direction yields an equivalence between derandomization of prBPP and the existence of almost-all inputs leakage-resilient hard functions.

Transition to AM. A number of changes are in order in terms of the requirements for similar results for AM. First, we need to handle *co-nondeterministic* distinguisher circuits D instead of deterministic ones. Co-nondeterministic circuits suffice because Arthur-Merlin protocols can be assumed to have perfect completeness. The only requirement for a correct derandomization is in the case of negative instances, in which case we want to hit the set of Arthur's random-bit strings for which Merlin cannot produce a witness. By the soundness property of the Arthur-Merlin protocol, the set contains at least half of the random-bit strings.

Second, we need to accommodate *nondeterministic* algorithms computing the function f. This is because the direction from derandomization to hardness seems to need them (see Proposition 4).

On each input x, such an algorithm needs to have at least one successful computation path, and on every successful computation path, the output should equal $f(x)$.

Third, the algorithm for the targeted hitting-set generator can also be nondeterministic, which is natural when the algorithm for f is nondeterministic. In the case of a generator, the nondeterministic algorithm should still have at least one successful computation path on every input, but it is fine to produce different outputs on different successful computation paths. For any given x and D, on every successful computation path, the output should be a hitting-set for D. This allows us to nondeterministically simulate a promise Arthur-Merlin protocol on input x as follows: Guess a computation path of the targeted HSG; if it succeeds, say with output S, guess a computation path for the Arthur-Merlin protocol on input x using each of the elements in S as the random-bit string, and accept if all of them accept; otherwise, reject.

Finally, we need to be able to run the reconstruction procedure as a (promise) *Arthur-Merlin protocol*. This is because we want the model in which we can compute $f(x)$ in case of a failed derandomization on input x, to match the class we are trying to derandomize. There are two requirements for the protocol to compute $f(x)$ on input x:

- *Completeness* demands that there exists a strategy for Merlin that leads Arthur to succeed with output $f(x)$ with high probability.

- *Soundness* requires that, no matter what strategy Merlin uses, the probability for Arthur to succeed with an output other than $f(x)$ is small.

The reconstructor naturally needs the power of nondeterminism in order to simulate the distinguisher D. Making sure the reconstructor is sound and needs no more power than prAM is the challenge.

Traditional setting for AM. In reference to the first two questions above, the answer to the one about obtaining a distinguisher D is similar as for BPP, except that in the uniform setting we do not know how to check in prAM for a difference in behavior of $D \doteq A(x, \cdot)$ between the uniform and the pseudo-random distributions. This is why average-case results remain open for AM. Instead, one assumes that some nondeterministic algorithm produces, on every successful computation path on input 1^n, an input x of length n on which the difference in behavior is guaranteed.

As for obtaining answers to the learning queries in the uniform setting for AM, we can make use of the nondeterminism allowed during the reconstruction and ask Merlin to provide the answers to the learning queries. However, we need to guard against a cheating Merlin. A strategy proposed by Gutfreund, Shaltiel and Ta-Shma in [GSTS03] consists of employing a function h that has a length-preserving instance checker. After Merlin has provided the supposed answers to the learning queries, to compute $h(z)$ for a given input z, we run the instance checker on input z and answer the queries y of the instance checker by running the evaluator part of the reconstruction process on input y. All the runs of the evaluator can be executed in parallel, ensuring a bounded number of rounds overall, which can be reduced to two in the standard way at the cost of a polynomial blowup in the running time [BM88].

To guarantee soundness, the reconstruction process needs to have an additional *resilience* property, namely that it remains partial single-valued even when the learning queries are answered incorrectly. Two hitting-set generators tailored for AM are known to have the property: the Miltersen-Vinodchandran generator MV [MV05], which is geared toward the high end, and a re-

cursive version, RMV, developed by Shaltiel and Umans [SU09] to cover a broader range. MV is used for the high end in [GSTS03], and RMV for the rest of the spectrum in [SU09].

New setting for AM. We build a targeted hitting-set generator for AM based on the RMV hitting-set generator. To obtain h_x from x, we make use of Probabilistically Checkable Proofs (PCPs) for the nondeterministic computation of the string $f(x)$ from x. Let V denote the verifier for such a PCP system that uses $O(\log(T(n)))$ random bits and poly $\log(T(n))$ queries for nonde-terministic computations that run in time $T(n)$. On input x, our targeted HSG guesses the value of $f(x)$ and a candidate PCP witness y_i for the i-th bit of $f(x)$ for each i, and runs all the checks of the verifier V on y_i (by cycling through all random-bit strings for V). If all checks pass, our targeted HSG instantiates RMV with y_i for each i as (the truth-table of) the oracle h_x, and outputs the union of all the instantiations as the hitting-set, provided those nondeterministic computations all accept; otherwise, the targeted HSG fails.

For the reconstruction of the i-th bit of $f(x)$, Arthur generates the learning queries of the RMV reconstructor for the oracle y_i, and Merlin provides the purported answers as well as the value of the i-th bit of $f(x)$. Arthur then runs some random checks of the verifier V on input x, answering the verifier queries by executing the evaluator of the RMV reconstructor. All the executions of the evaluator can be performed in parallel, ensuring a bounded number of rounds overall. The resilient partial single-valuedness property of the RMV reconstructor guarantees that the verifier queries are all consistent with some candidate proof $\tilde{y}_i$. The completeness and soundness of the PCP then imply the completeness and soundness of the reconstruction process for our targeted HSG. As V makes few queries and is very efficient, the running time of the process is dominated by the running time of the RMV reconstructor.

Abstracting out the details of our construction and how the distinguisher D is obtained, the result can be captured in two procedures: a nondeterministic one, H, which has at least one successful computation path for every input and plays the role of a targeted hitting-set generator, and a promise Arthur-Merlin protocol, R, which plays the role of a reconstructor for the targeted hitting-set generator. H and R have access to the input x and a co-nondeterministic circuit D, and have the following property.[3]

Property 10. *For every $x \in \{0,1\}^*$ and for every co-nondeterministic circuit D that accepts at least half of its inputs, at least one of the following holds:*

1. *$H(x, D)$ outputs a hitting-set for D on every successful computation path.*

2. *$R(x, D)$ computes $f(x)$ in a complete and sound fashion.*

Theorem 5 follows by considering nondeterministic running time $T(n) = n^a$ and co-nondeter-ministic circuits D of size n^c for some $c > 1$. In this regime, H runs in time $n^{O(a+c)}$ and R in time $n^{O(c(\log a)^2)}$. Under the hypothesis of Theorem 5, the second item in Property 10 cannot happen except for finitely many x of length n, so the first item needs to hold. For any constant $c' < c$, this yields a polynomial-time targeted hitting-set generator for prAMTIME$[n^{c'}]$, which can be used for

---

[3]The dependency of H on D is only through the number of input bits of D. For R, blackbox access to D suffices (in addition to the input x). However, we may as well give both H and R full access to the input x and the circuit D. In the intended application, the co-nondeterministic circuit D is obtained by hardwiring the input x into the Arthur-Merlin protocol being derandomized, but this is not essential for the construction.

all of prAM by padding. Theorem 6 follows along the same lines; the running time is dictated by the RMV reconstructor.

We point out that the approach of Chen and Tell can be ported to the AM setting by replacing NW with a generator for AM that has the learning property and a reconstructor running in prAM. The nondeterminism allows us to run the bootstrapping process in parallel, so the number of rounds of Arthur and Merlin remains bounded, but the overall running time remains proportional to the depth of the circuits for f. This means that, like in the setting of BPP, this approach only yields meaningful results when the depth is small compared to the size. Nondeterministic circuits for f can be accommodated in this approach by treating them as deterministic circuits with nondeterministic guess bits as additional inputs. However, this limits the amount of nondeterminism that can be handled. Our approach based on PCPs remedies the limitations on depth as well as nondeterminism.

Derandomization to hardness. Our derandomization-to-hardness result is proven by diagonalization. Under the prAM $\subseteq$ NP assumption, every fixed-polynomial time AM protocol computing a length-preserving function can be simulated in nondeterministic fixed-polynomial time. We would like to diagonalize against these simulating nondeterministic machines to construct our hard function. Due to the lack of an almost-everywhere hierarchy result for NTIME, we do not know how to do this efficiently for generic nondeterministic machines. This is where the advice comes to rescue: We use advice to indicate which nondeterministic machines are *single-valued* at a particular input length. We only need to consider single-valued machines, and diagonalizing against them is easy for a nondeterministic machine with a little more running time, but figuring out which nondeterministic machines are single-valued at a given input length is hard.

2.2 Byproducts

In this section, we develop the intuition for our byproducts.

Targeted hitting-set generators from derandomization (Theorem 7). To obtain a targeted HSG from derandomization of prAM, we employ our targeted hitting-set generator in a win-win argument. Either a complexity class separation holds, in which case a result of [IKW02] guarantees the existence of a regular (oblivious) hitting-set generator that yields the derandomization result, or we get a strong complexity class collapse. The collapse allows us to bypass some of the difficulties in diagonalizing against prAM protocols on almost-all inputs (one of the reasons we require advice in the derandomization-to-hardness direction of our near-equivalence), thus allowing us to do so efficiently and uniformly, and then instantiate our targeted hitting-set generator construction.

Average-case derandomization (Theorem 8). Our average-case derandomization results under worst-case hardness assumptions also make use of our targeted hitting-set generator construction, but in a different way. They do not exploit the potential of the hitting-sets to depend on the input x. In fact, they set $f(x)$ to the truth-table of the worst-case hard language L from the hypothesis at an input length determined by $|x|$. Instead, they hinge on the strong resilient soundness properties of the reconstructor.

As we are considering the average-case derandomization setting, the problem of obtaining the distinguisher D for the reconstruction resurfaces. Our approach is similar to the one for the

traditional average-case derandomization setting for BPP. If the simulation fails for protocol A with noticeable probability over a random input, then we can sample multiple inputs $x_1, x_2, \ldots$ and construct a list of "candidate distinguishers" $D_{x_1} \doteq A(x_1, \cdot), D_{x_2} \doteq A(x_2, \cdot), \ldots$ such that the list contains, with high probability, at least one "true" distinguisher. Whereas in the BPP setting one can test each candidate and discard, with high probability, the ones that are not distinguishers, we do not know how to do that in the AM setting. Instead, we employ a different approach: We run the reconstructor with each distinguisher with the hope that every execution either fails or outputs the correct value.

This approach necessitates a stronger form of resilience than the one provided by the RMV generator: That its reconstruction is sound when given as input *any* co-nondeterministic circuit D, not just those that accept at least half of their inputs (as in Property 10). We don't know how to guarantee this with our prAM reconstruction, but we are able to do so in $\mathrm{prBPP}_{||}^{\mathrm{SAT}}$ by approximating the fraction of inputs that D accepts and outright failing if the fraction is too low.

We point out that earlier works [GSTS03, SU09] also manage to guarantee soundness of the reconstructor for co-nondeterministic circuits D that accept at least half of their inputs, based on the resilient partial single-valuedness of the reconstructor for MV or RMV. They do so by running an instance checker, which limits the hard function f to classes for which instance checkers are known to exist, such as complete problems for E and EXP. Instead, we achieve soundness of the reconstructor based on the soundness of a PCP. As PCPs exist for all nondeterministic computations, this makes our approach more suitable in this setting. In particular, we do not know how to obtain Theorem 8 along the lines of [GSTS03, SU09].

Unconditional mild derandomization (Theorem 9).   Our unconditional mild derandomization result relies on a similar win-win argument as in the proof of Theorem 7: Either some hardness assumption/class separation holds, in which case we get derandomization right away, or we get a complexity collapse that we use to construct, by diagonalization, a hard function f that has the efficiency requirements we need to obtain the derandomization result using our targeted hitting-set generator.

Since our result is unconditional, we cannot use derandomization assumptions to make diagonalizing against prAM protocols easier. Instead, we rely on the inclusion $\mathrm{prAM} \subseteq \Pi_2 P$, which allows for diagonalizing against such protocols in $\Sigma_2 \mathrm{TIME}[n^{\omega(1)}]$. Our generator, however, requires the hard function to be computable by efficient nondeterministic algorithms. To help bridge the gap, we prove an "easy witness lemma" for Σ_2 computations that guarantees a strong collapse in case the aforementioned hardness assumption does not hold. The collapse then allows us to instantiate our targeted hitting-set generator construction with the diagonalizing function.

3   Preliminaries

We assume familiarity with standard complexity classes such as NP, AM, and prAM. We often consider inputs and outputs from non-Boolean domains, such as $\mathbb{F}^r$ for a field $\mathbb{F}$ and $r \in \mathbb{N}$. In such cases, we implicitly assume an efficient binary encoding for the elements of these domains. Finally, as is customary, all time bounds considered are implicitly assumed to be time-constructible.

3.1 Nondeterministic, co-nondeterministic and single-valued computation

We make use of nondeterministic, co-nondeterministic, and single-valued circuits in our results. A nondeterministic circuit is a Boolean circuit C with two sets of inputs, x and y. We say that C accepts x if there exists some y such that $C(x, y) = 1$, and that C rejects x otherwise. A co-nondeterministic circuit has a symmetric acceptance criterion: It accepts x if for all y it holds that $C(x, y) = 1$, and rejects x otherwise. A partial single-valued circuit also has two inputs, x and y; on input (x, y) it either fails (which we represent by $C(x, y) = \bot$) or succeeds and outputs a bit $b = C(x, y)$. Moreover, we require that for all y, y' such that both $C(x, y)$ and $C(x, y')$ succeed, $C(x, y) = C(x, y')$, i.e., the circuit computes a partial function on its first input. If, furthermore, for all x there exists a y such that $C(x, y)$ succeeds, we call the circuit total single-valued or just single-valued.

We are also interested in nondeterministic algorithms that compute multi-bit functions $f : \{0, 1\}^* \to \{0, 1\}^*$. Let $T(n)$ be a time bound. We say that $f \in \mathrm{NTIME}[T(n)]$ if there exists a nondeterministic algorithm N running in time $O(T(n))$ such that for all $x \in \{0, 1\}^*$, there exists at least one computation path on which $N(x)$ succeeds, and $N(x)$ outputs $f(x)$ on all successful computation paths. Note, in particular, that if $f \in \mathrm{NTIME}[T(n)]$, then the language $L_f = \{(x, i, b) \mid f(x)_i = b\}$ is in $\mathrm{NTIME}[T(n)]$.

3.2 Arthur-Merlin protocols

A promise Arthur-Merlin protocol P is a computational process in which Arthur and Merlin receive a common input x and operate as follows in alternate rounds for a bounded number of rounds. Arthur samples a random string and sends it to Merlin. Merlin sends a string that depends on the input x and all prior communication from Arthur; the underlying function is referred to as Merlin's strategy, which is computationally unrestricted. At the end of the process, a deterministic computation on the input x and all communication determines acceptance. The running time of the process is the running time of the final deterministic computation.

Any promise Arthur-Merlin protocol can be transformed into an equivalent one with just two rounds and Arthur going first, at the cost of a polynomial blow-up in running time, where the degree of the polynomial depends on the number of rounds [BM88]. As such, we often use the notation prAM to refer to promise Arthur-Merlin protocols with any bounded number of rounds, even though, strictly speaking, the notation refers to a two-round protocols with Arthur going first.

Promise Arthur-Merlin protocols can be simulated by probabilistic algorithms with oracle access to SAT: Instead of interacting with Merlin, Arthur asks the SAT oracle whether there exists a response of Merlin that would lead to acceptance. Similarly, $\mathrm{P}^{\mathrm{prAM}}_{\|}$ can be simulated in $\mathrm{BPP}^{\mathrm{SAT}}_{\|}$, the class of problems decidable by probabilistic polynomial-time algorithms with bounded error and non-adaptive oracle access to SAT. In fact, a converse also holds and helps to extend some of our results for prAM to the class $\mathrm{prBPP}^{\mathrm{SAT}}_{\|}$.

Lemma 11 ([CR11]). $\mathrm{prBPP}^{\mathrm{SAT}}_{\|} \subseteq \mathrm{P}^{\mathrm{prAM}}_{\|}$.

In Lemma 11, the deterministic machines with oracle access to prAM on the right-hand side are guaranteed to work correctly irrespective of how the queries outside of the promise are answered, even if those queries are answered inconsistently, i.e., different answers may be given when the same query is made multiple times.

Arthur-Merlin protocols that output values. A promise Arthur-Merlin protocol P may also output a value. In this case, at the end of the interaction, the deterministic computation determines success/failure and, in case of success, an output value. We denote this value by $P(x, M)$, which is a random variable defined relative to a strategy M for Merlin. Similar to the setting of circuits, we indicate failure by setting $P(x, M) = \perp$, a symbol disjoint from the set of intended output values. Our choice of using success and failure for protocols that output values is to avoid confusion with the decisional notions of acceptance and rejection.

Definition 12 (Arthur-Merlin protocol with output). *Let P be a promise Arthur-Merlin protocol. We say that on a given input $x \in \{0,1\}^*$:*

- $\circ$ *P outputs v with completeness c if there exists a Merlin strategy such that the probability that P succeeds and outputs v is at least c. In symbols: $(\exists M) \Pr[P(x, M) = v] \geq c$.*

- $\circ$ *P outputs v with soundness s if, no matter what strategy Merlin uses, the probability that P succeeds and outputs a value other than v is at most s. In symbols: $(\forall M) \Pr[P(x, M) \notin \{v, \perp\}] \leq s$.*

- $\circ$ *P has partial single-valuedness s if there exists a value v such that P outputs v with soundness s. In symbols: $(\exists v)(\forall M) \Pr[P(x, M) \notin \{v, \perp\}] \leq s$.*

Note that if P on input x outputs v with completeness c and has partial single-valuedness s, then it outputs v with soundness s, provided $s > 1 - c$. If we omit c and s, then they take their default values of $c = 1$ (perfect completeness) and $s = 1/3$.

For a given function $f : X \to \{0,1\}^*$ where $X \subseteq \{0,1\}^*$, we say that P computes f with completeness $c(n)$ and soundness $s(n)$ if on every input $x \in X$, P outputs $f(x)$ with completeness $c(|x|)$ and soundness $s(|x|)$. Note that P may behave arbitrarily on inputs that are not in X. In contrast, an AM protocol computing f still computes some value in a complete and sound fashion on inputs $x \notin X$.

3.3 Learn-and-evaluate and commit-and-evaluate protocols

The reconstruction processes for hardness-based hitting-set generators for prAM are typically special types of promise Arthur-Merlin protocols. We distinguish between two types.

A *learn-and-evaluate protocol* is composed of two phases: A *learning* phase followed by an *evaluation* phase. In the learning phase, a probabilistic algorithm makes queries to a function f and produces an output (which we call a sketch). The evaluation phase then consists of a promise Arthur-Merlin protocol that computes $f(x)$ correctly on every input x when given the sketch as additional input.

Definition 13 (Learn-and-evaluate protocol). *A learn-and-evaluate protocol P consists of a probabilistic oracle algorithm A_{learn} and a promise Arthur-Merlin protocol P_{eval}. Let $f : X \to \{0,1\}^*$ where $X \subseteq \{0,1\}^*$. We say that P computes f with error $e(n)$ for completeness $c(n)$ and soundness $s(n)$ if on every input $x \in X$ of length n the following hold: The probability over the randomness of A_{learn} that P_{eval} with input x and additional input $\pi = A_{learn}^f(1^n)$ outputs $f(x)$ with completeness $c(n)$ and soundness $s(n)$ is at least $1 - e(n)$.*

The learning phase of a learn-and-evaluate protocol can be simulated by an Arthur-Merlin protocol with output, where Merlin guesses the queries that A_{learn} makes on a given random-bit

string and answers them in parallel, and the output is a sketch of f. In this view, a learn-and-evaluate protocol becomes a pair of promise Arthur-Merlin protocols: one for the learning phase, and one for the evaluation phase. Note that the quality of the evaluation phase is only guaranteed when the learning queries are answered correctly, i.e., when Merlin is honest in the learning phase.

A *commit-and-evaluate* protocol [SU09] has the syntactic structure of a pair of promise Arthur-Merlin protocols without the restriction that Merlin in the first phase only answers queries about f. Semantically, a commit-and-evaluate protocol is more constrained than a learn-and-evaluate protocol. The first protocol of the pair now represents a commitment phase instead of a learning phase. In this phase, Arthur and Merlin interact and produce an output π, which we call a commitment. Similar to a learn-and-evaluate protocol, the commitment is given as input to the protocol of the evaluation phase. Whereas in a learn-and-evaluate protocol there are no guarantees whatsoever when Merlin is dishonest in the first phase, in a commit-and-evaluate protocol there is a strong guarantee: With high probability over Arthur's randomness in the commitment phase, the evaluation protocol is partial single-valued, meaning that Merlin cannot make Arthur output different values for the same input x with high probability. The guarantee is referred to as resilient partial single-valuedness.

Definition 14 (Commit-and-evaluate protocol). *A commit-and-evaluate protocol is a pair of promise Arthur-Merlin protocols $P = (P_{commit}, P_{eval})$. P has resilience $r(n)$ for partial single-valuedness $s(n)$ on domain $X \subseteq \{0,1\}^*$ if for all n, no matter what strategy Merlin uses during the commit phase, the probability that in the commitment phase, on input 1^n, P_{commit} succeeds and outputs a commitment π that fails to have the following property (2) is at most $r(n)$:*

For every x of length n in the domain X, $P_{eval}(x, \pi)$ has partial single-valuedness $s(n)$. (2)

In symbols:

$$(\forall n)(\forall M_{commit}) \Pr[(\forall x \in X \cap \{0,1\}^n) P_{eval}(x, \pi) \text{ has partial single-valuedness } s(n)] \geq 1 - r(n),$$

where $\pi = P_{commit}(1^n, M_{commit})$.

A commit-and-evaluate protocol naturally induces a promise Arthur-Merlin protocol: On input x, run P_{commit} on input $1^{|x|}$. If this process succeeds, let π denote its output and run P_{eval} on input (x, π).

3.4 Hitting-set generators and targeted hitting-set generators

In the setting of prBPP, Goldreich [Gol20] discusses two equivalent definitions of targeted pseudo-random generators: one for deterministic linear-time machines that take both the input x and the random-bit string r as inputs, and one based on circuits D that only take the random-bit string r as input. The circuit D can be obtained by first constructing a circuit C that simulates the machine on inputs of length $|x|$, and then hardwiring the input x. The difference between a regular and targeted pseudo-random generator lies in the dependency of the output on x (in the first definition) or the circuit D (in the second definition): For a regular PRG the output can only depend on $|x|$ or the size of D, whereas for a targeted PRG it can depend on x and D proper.

In the setting of prAM, without loss of generality, we can assume that promise Arthur-Merlin protocols have perfect completeness [FGM$^+$89]. Therefore, we only need to consider targeted hitting-set generators, the variant of targeted PRGs for one-sided error. Similar to the BPP setting,

there are two equivalent definitions of targeted HSGs for prAM. We propose a third, hybrid, and also equivalent definition, where the targeted generator is given access to both x and the circuit C. For prAM with perfect completeness the circuit C (as well as D) is co-nondeterministic. For regular HSGs, the output can only depend on the size of C. Our definition highlights that, in principle, there are two types of obliviousness that regular PRGs/HSGs exhibit: With respect to the input (where only dependencies on its size are allowed) and with respect to the algorithm being derandomized (where only dependencies on its running time are allowed). Since the algorithm description can be incorporated as part of the input, the dependency on C can be avoided. This is essentially why all three definitions are equivalent. In our targeted hitting-set generator constructions the dependency will only be through x and the size of C.

We start by defining hitting-sets for co-nondeterministic circuits.

Definition 15 (Hitting-set for co-nondeterministic circuits). *Let D be a co-nondeterministic circuit of size m. A set S of strings of length m is a hitting-set for D if there exists at least one $z \in S$ such that $D(z) = 1$ (where D might take a prefix of z as input if necessary). In that case, we say that S hits D.*

The notion allows us to define targeted hitting-set generators for prAM as follows, where we assume, without loss of generality, perfect completeness and soundness $1/2$. Regular hitting-set generators are viewed as a special case.

Definition 16 (Regular and targeted hitting-set generator for prAM). *A targeted hitting-set generator for prAM is a nondeterministic algorithm that, on input $x \in \{0,1\}^*$ and a co-nondeterministic circuit C, has at least one successful computation path, and if $\Pr_r[C(x,r) = 1] \geq 1/2$, outputs a hitting set for $D(r) \doteq C(x,r)$ on every successful computation path. A regular hitting-set generator for prAM is a targeted hitting-set generator where the output only depends on the size of C.*

For completeness, we state the standard way of obtaining the co-nondeterministic circuits C and D capturing promise Arthur-Merlin protocols.

Proposition 17. *There exists an algorithm that, on input 1^n and the description of a (Boolean output, two-round) prAMTIME$[t(n)]$ protocol P, runs in time $O(t(n)^2)$ and outputs a co-nondeterministic circuit C of size $m = O(t(n)^2)$ that simulates and negates the computation of P for input length n, i.e., its is comprised of $x \in \{0,1\}^n$ and Arthur's random-bit string r, and it co-nondeterministically verifies that there is no Merlin message that would lead to acceptance. In particular:*

- *If P with input x accepts all random inputs, then $D_x(r) \doteq C(x,r)$ rejects every input.*

- *If P with input x rejects at least a fraction $1/2$ of its random-bit strings, then $D_x(r) \doteq C(x,r)$ accepts at least a fraction $1/2$ of its inputs.*

3.5   PCPs and low-degree extensions

We use the following construction that follows from the PCP of proximity of Ben-Sasson, Goldreich, Harsha, Sudan, and Vadhan [BGH$^+$06].

Lemma 18 ([BGH$^+$06]). *Let T be a time bound. For every $s = s(n) : \mathbb{N} \to (0, 1]$ and every language $L \in \mathrm{NTIME}[T(n)]$ there exists a PCP verifier V with perfect completeness, soundness s, randomness complexity $\log(1/s) \cdot (\log T(n) + O(\log \log T(n)))$, non-adaptive query complexity $\log(1/s) \cdot \mathrm{polylog}(T(n))$, and verification time $\log(1/s) \cdot \mathrm{poly}(n, \log T(n))$. More precisely,*

○ *V has oracle access to a proof of length $T(n) \cdot \mathrm{polylog}(T(n))$, uses $\log(1/s) \cdot (\log T(n) + O(\log \log T(n)))$ random bits in any execution, makes $\log(1/s) \cdot \mathrm{polylog}(T(n))$ non-adaptive queries to the proof and runs in time $\log(1/s) \cdot \mathrm{poly}(n, \log T(n))$.*

○ *If $x \in L$, $|x| = n$, then there exists y of length $T(n) \cdot \mathrm{polylog}(T(n))$ such that $\Pr[V^y(x) = 1] = 1$.*

○ *If $x \notin L$, $|x| = n$, then for all y' of length $T(n) \cdot \mathrm{polylog}(T(n))$, $\Pr[V^{y'}(x) = 1] \leq s$.*

We also need standard low-degree extensions. Let $g : \{0,1\}^\ell \to \{0,1\}$ be a function, $\mathbb{F} = \mathbb{F}_p$ be the field with p elements (for prime p) and h and r integers such that $h^r \geq 2^\ell$. The low-degree extension of g with respect to p, h, r is the unique r-variate polynomial $\hat{g} : \mathbb{F}^r \to \mathbb{F}$ with degree $h - 1$ in each variable, for which $\hat{g}(\vec{v}) = g(y)$ for all $\vec{v} \in [h]^r$ representing a $y \in \{0,1\}^\ell$ and $\hat{g}(\vec{v}) = 0$ for the $\vec{v} \in [h]^r$ that do not represent a string y. The total degree of $\hat{g}$ is $\Delta = hr$ and $\hat{g}$ is computable in time $\mathrm{poly}(h^r, \log p, r)$ given oracle access to g.

3.6 Average-case simulation

The instance-wise nature of our technique allows us to conclude derandomization on average with respect to arbitrary distributions by assuming hardness with respect to that same distribution. The notion of average-case simulation that we use is the one where the simulation works correctly with high probability over inputs drawn from the distribution. We typically want good simulations to exist with respect to every efficiently sampleable distribution (where the simulation may depend on the distribution). This is usually referred to as the "heuristic" setting.

Definition 19 (Heuristic). *Let Π be a promise-problem, $\mu : \mathbb{N} \to [0, 1)$, $\mathcal{C}$ a complexity class and $\mathbf{x} = \{\mathbf{x}_n\}_{n \in \mathbb{N}}$ an ensemble of distributions where $\mathbf{x}_n$ is supported on $\{0,1\}^n$ and such that for all n, every x in the support of $\mathbf{x}_n$ satisfies the promise of Π. We write*

$$\Pi \in \mathrm{Heur}_{\mathbf{x}, \mu} \mathcal{C}$$

if there exists a language $L \in \mathcal{C}$ such that for all sufficiently large n, $\Pr_{x \in \mathbf{x}_n}[L(x) \neq \Pi(x)] \leq \mu(n)$. We write

$$\Pi \in \mathrm{Heur}_\mu \mathcal{C}$$

if the above property holds for every polynomial-time sampleable ensemble of distributions with the above support restriction.

The notions of average-case simulation extend to the infinitely-often setting in the natural way.

4   Targeted hitting-set generator construction

In this section, we develop our targeted HSG construction, which leads to our instance-wise hardness vs. randomness tradeoffs for Arthur-Merlin protocols.

Our construction builds on the RMV generator due to Shaltiel and Umans [SU09], which is a recursive variant of the MV generator that shares the desired resilience property with MV. We start with the definition of the RMV generator in Section 4.1 and state its reconstruction properties in terms of a commit-and-evaluate protocol. We present our construction and analysis in Section 4.2 and the derandomization consequences in Section 4.3.

4.1   Recursive Miltersen-Vinodchandran generator

We need a couple of ingredients to describe how the RMV generator works. The first one is a local extractor for the Reed-Müller code. A local extractor is a randomness extractor that only needs to know a few bits of the sample. In the following definition the sample is provided as an oracle, and the structured domain from which the sample is drawn is given as an additional parameter.

Definition 20 (Local extractor). *Let S be a set. A (k, ϵ) local S-extractor is an oracle function $E : \{0,1\}^s \to \{0,1\}^t$ that is computable in time $\mathrm{poly}(s,t)$ and has the following property: For every random variable X distributed on S with min-entropy at least k, $E^X(U_s)$ is ϵ-close to uniform.*

We make use of the following local extractor for Reed-Müller codes.

Lemma 21 (Implicit in [SU05]). *Fix parameters $r < \Delta$, and let S be the set of polynomials $\hat{g} : \mathbb{F}^r \to \mathbb{F}$ having total degree at most Δ, where $\mathbb{F} = \mathbb{F}_p$ denotes the field with p elements. There is a $(k, 1/k)$ local S-extractor for $k = \Delta^5$ with seed length $s = O(r \log p)$ and output length $t = \Delta$.*

Note that for every subcube with sides of size $\frac{\Delta}{r}$ and choice of values at its points, there exists an interpolating polynomial $\hat{g}$ with the parameters of Lemma 21. It takes $(\Delta/r)^r \log p$ bits to describe these polynomials, but the local extractor only accesses $\mathrm{poly}(\Delta, r, \log p)$ bits.

When instantiated with a polynomial $\hat{g} : \mathbb{F}^r \to \mathbb{F}$, the RMV generator groups variables and operates over axis-parallel (combinatorial) lines over the grouped variables.[4] Shaltiel and Umans call these MV-lines, which we define next.

Definition 22 (MV line). *Let $\mathbb{F} = \mathbb{F}_p$ for a prime p. Given a function $\hat{g} : \mathbb{F}^r \to \mathbb{F}$ where r is an even integer, we define $B = \mathbb{F}^{r/2}$ and identify $\hat{g}$ with a function from B^2 to $\mathbb{F}$. Given a point $\vec{a} = (\vec{a}_1, \vec{a}_2) \in B^2$ and $i \in \{1,2\}$, we define the line passing through $\vec{a}$ in direction i to be the function $L : B \to B^2$ given by $L(\vec{z}) = (\vec{z}, \vec{a}_2)$ if $i = 1$ and $L(\vec{z}) = (\vec{a}_1, \vec{z})$ if $i = 2$. This is an axis-parallel, combinatorial line, and we call it an MV-line. Given a function $\hat{g} : \mathbb{F}^r \to \mathbb{F}$ and an MV-line L we define the function $\hat{g}_L : B \to \mathbb{F}$ by $\hat{g}_L(z) = \hat{g}(L(z))$.*

The input for the RMV construction is a multivariate polynomial $\hat{g} : \mathbb{F}^r \to \mathbb{F}$ of total degree at most Δ, and the output is a set of m-bit strings for $m = \Delta^{1/100}$. The construction is recursive and requires that r is a power of 2 and that p is a prime larger than Δ^{100} (say, between Δ^{100} and $2\Delta^{100}$). Let E be the $(k, 1/k)$-local extractor from Lemma 21 for polynomials of degree Δ in $(r/2)$

---

[4]In the original construction [SU09], the RMV generator is defined with the number d of groups of variables as an additional parameter. Eventually, d is set to 2, which is the value we use for our results as well.

variables over $\mathbb{F}$. Remember that $k = \Delta^5$ and that the extractor uses seed length $O(r \log p)$ and output length $t = \Delta \geq m$. By using only a prefix of the output, we have it output exactly m bits.

The operation of the RMV generator on input $\hat{g}$ is as follows: Set $B = \mathbb{F}^{r/2}$. For every $\vec{a} \in B^2$ and $i \in \{1, 2\}$, let $L : B \to B^2$ be the MV-line passing through $\vec{a}$ in direction i. Compute $E^{\hat{g}_L}(y)$ for all seeds y. For $r = 2$, output the set of all strings of length m obtained over all $\vec{a} \in B^2$, MV lines L through $\vec{a}$, and seeds y. For $r > 2$, output the union of this set and the sets output by the recursive calls $\mathrm{RMV}(\hat{g}_L)$ for each of the aforementioned MV-line L.

The construction runs in time $p^{O(r)}$ and therefore outputs at most that many strings. If the set output by the procedure fails as a hitting-set for a co-nondeterministic circuit D of size m, then there exists an efficient commit-and-evaluate protocol P for $\hat{g}$ with additional input D. This is the main technical result of [SU09], which we present in a format that is suitable for obtaining our results.[5]

Lemma 23 ([SU09]). *Let Δ, m, r, p be such that $m = \Delta^{1/100}$, r is a power of 2 and p is a prime between Δ^{100} and $2\Delta^{100}$. Let also $\mathbb{F} = \mathbb{F}_p$ and $s \in (0, 1]$. There exists a commit-and-evaluate protocol $P = (P_{commit}, P_{eval})$ with additional input D, where D is a co-nondeterministic circuit of size m, such that the following holds for any polynomial $\hat{g} : \mathbb{F}^r \to \mathbb{F}$ of total degree at most Δ.*

- *Completeness: If D rejects every element output by $\mathrm{RMV}(\hat{g})$ then there exists a strategy M_{commit} for Merlin in the commit phase such that P_{eval} on input (z, D, π) outputs $\hat{g}(z)$ with completeness 1 for every $z \in \mathbb{F}^r$, where $\pi \doteq P_{commit}(1^n, M_{commit})$.*

- *Resilience: If D accepts at least a fraction $1/2$ of its inputs then P has resilience s for partial single-valuedness s on domain $\mathbb{F}^r$.*

- *Efficiency: Both P_{commit} and P_{eval} have two rounds. P_{commit} runs in time $\log{(1/s)} \cdot \mathrm{poly}(\Delta, r)$ and P_{eval} runs in time $(\log{(1/s)})^2 \cdot \Delta^{O((\log r)^2)}$.*

P only needs blackbox access to the deterministic predicate that underlies D.

4.2 Targeted generator and reconstruction

In this section, we present our targeted HSG construction, which works as follows: On input x and a co-nondeterministic circuit D of size m, it guesses a PCP (as in Lemma 18) for each bit of $f(x)$ and verifies each PCP deterministically by enumerating over the PCP verifier's randomness. It encodes each PCP as a low-degree polynomial (as in Section 3.5), instantiates the RMV generator with each of the polynomials and outputs the union of the outputs for each instantiation. For the reconstruction, we have Merlin send a bit b and commit to the low-degree extension of a proof that the i-th bit of $f(x)$ equals b. Arthur then runs the PCP verifier using the evaluation protocol to answer proof queries. The protocol succeeds and outputs b if and only if the PCP verifier accepts. Here is the formal statement of the result.

Theorem 24. *Let $T(n)$ be a time bound and $f \in \mathrm{NTIME}[T(n)]$. There exists a nondeterministic algorithm H (the generator) that always has at least one successful computation path per input, and a promise Arthur-Merlin protocol R (the reconstructor) such that for every $x \in \{0, 1\}^*$ and every*

---

[5]Shaltiel and Umans present the evaluation protocol as a multi-round protocol (with $\log r$ rounds). We collapse it into a two-round protocol by standard amplification (which also amplifies the crucial resilience property) [BM88, SU09].

co-nondeterministic circuit D that accepts at least half of its inputs, at least one of the following holds.

1. *$H(x, D)$ outputs a hitting-set for D on every successful computation path.*

2. *$R(x, D)$ computes $f(x)$ with completeness 1 and soundness 1/3.*

The construction also has the following properties:

- Resilient soundness: *In either case, the probability that $R(x, D)$ outputs a value other than $f(x)$ is at most 1/3.*

- Efficiency: *On inputs x of length n and D of size m, H runs in time $\mathrm{poly}(T(n), m)$, and R, given an additional index i, computes the i-th bit of $f(x)$ in time $\mathrm{poly}(n) \cdot (m \cdot \log T(n))^{O((\log r)^2)}$ for $r = O(\log(T(n))/\log m)$.*

Moreover, $H(x, D)$ only depends on x and the size of D, and $R(x, D)$ only needs blackbox access to the deterministic predicate that underlies D.

Proof. Let $f \in \mathrm{NTIME}[T(n)]$, consider the language $L_f = \{(x, i, b) \mid f(x)_i = b\}$ and note that $L_f \in \mathrm{NTIME}[T(n)]$. Let V be the PCP verifier of Lemma 18 for L_f with soundness $s = s(n) = (100T(n)^{-1})$. Let also $h = h(m) = m^{100}$, $p = p(m)$ be the smallest prime in the interval $[h^{100}, 2h^{100}]$ and $r = r(n, m)$ be the smallest power of 2 such that h^r is greater than the proof length of V on input length n. Note, in particular, that $h^r = \mathrm{poly}(T(n), m)$ and $r = O(\log(T(n))/\log m)$.

Generator. The generator H, on input x and a co-nondeterministic circuit D of size m, first guesses the value of $z = f(x)$ and a proof y_i of the correct length $T(n) \cdot \mathrm{polylog}(T(n))$ for the i-th bit of z for each i. Then it verifies that $\Pr[V^{y_i}(x, i, z_i) = 1] = 1$ for all i by deterministically enumerating over the $\mathrm{poly}(T(n))$ random-bit strings for V. If any of the verifications fail, it fails. Otherwise, it views each y_i as a function $g_i : \{0, 1\}^\ell \to \{0, 1\}$ for $\ell = \log|y_i|$ and outputs $\mathrm{RMV}(\hat{g}_i)$, where $\hat{g}_i$ is the low-degree extension of g_i with parameters p, h and r. The initial verification step takes time $\mathrm{poly}(T(n))$, and executing $\mathrm{RMV}(\hat{g})$ takes time $p^{O(r)} = \mathrm{poly}(T(n), m)$ and outputs strings of length m. This culminates in a running time of $\mathrm{poly}(T(n), m)$. Finally, since for the correct output $z = f(x)$ there always exist proofs y_i that are accepted with probability 1 for each i, there always exists a nondeterministic guess that leads the generator to succeed.

Reconstructor. We describe and analyze the prAM protocol R, which uses the commit-and-evaluate protocol $P = (P_{\mathrm{commit}}, P_{\mathrm{eval}})$ of Lemma 23 with soundness parameter $s' = s'(n) = (100T(n) \cdot q)^{-1}$, where $q = q(n) = \mathrm{polylog}(T(n))$ denotes the query complexity of V at input length n. On inputs x, D and an index i, Arthur and Merlin play the commit phase P_{commit}, which produces a commitment π_i to be fed into the evaluation phase. In parallel, Merlin also sends a bit b to Arthur. The idea is for an honest Merlin to send $b = f(x)_i$ and commit to the low-degree extension $\hat{g}_i$ of a proof y_i that witnesses $(x, i, b) \in L_f$ (or $f(x)_i = b$), though a dishonest Merlin may send a different bit and/or commit to some different function. Let γ_i denote the function that Merlin committed to via P_{commit}, which may be accessed with high probability by executing the evaluation protocol P_{eval} with input π_i. The restriction of γ_i to $[h]^r$ defines a candidate PCP proof $\tilde{y}_i$. Arthur then runs the verifier $V^{\tilde{y}_i}(x, i, b)$, employing Merlin's help to evaluate $\tilde{y}_i$ whenever V makes a query to it (where binary queries are first converted into the respective $\vec{v} \in \mathbb{F}_p^r$ and all queries are evaluated in parallel). If $V^{\tilde{y}_i}(x, i, b)$ accepts, then R succeeds and outputs b, otherwise it fails.

Completeness. If D is not hit by $H(x, D)$, then for all indices i there exists at least one proof y_i that witnesses $(x, i, f(x)_i) \in L_f$ and such that $\mathrm{RMV}(\hat{g}_i)$ fails to hit D, where $\hat{g}_i$ is the low-degree extension of y_i with parameters p, h and r. In that case, an honest Merlin can commit to such a $\hat{g}_i$ with probability 1 by the completeness property of Lemma 23 as well as send the correct value of $f(x)_i$ during the first phase. Then perfect completeness of V and P_{eval} guarantee that R succeeds and outputs $f(x)_i$ with probability 1.

(Resilient) soundness. If D accepts at least half of its inputs, then for a fixed index i the resilience property of P in Lemma 23 guarantees that with probability at least $1 - s'$, the commit phase is successful and thus the evaluation protocol with input π_i has partial single-valuedness s'. In that case, by a union bound over the at most q queries that V makes, with probability at least $1 - (100T(n))^{-1} = 1 - s$, every execution of the evaluation protocol results in the evaluation of a fixed function $\gamma_i : \mathbb{F}^r \to \mathbb{F}$. If Merlin sends the incorrect value of $b \neq f(x)_i$ in the first round (the only way he could try to have Arthur output the wrong value), the soundness property of V in Lemma 18 guarantees that R fails with probability at least $1 - s$ since $(x, i, b) \notin L_f$. By a union bound over these three "bad" events, all of which have probability at most s since $s \geq s'$, for any fixed index i, $R(x, D)$ with additional input i either fails or outputs $f(x)_i$ with probability at least $1 - 3s$. Finally, a union bound over the at most $T(n)$ possible indices i guarantees that R either fails or outputs $f(x)$ with soundness $1/3$. In particular, if completeness also holds then $R(x, D)$ computes $f(x)$ with completeness 1 and soundness $1/3$.

Efficiency. The commit phase takes time $\log(1/s') \cdot \mathrm{poly}(\Delta, r) = \mathrm{poly}(m, \log T(n))$ and two rounds of communication. Afterwards, evaluating each query made by $V(x, i, b)$ with P_{eval} takes time $(\log(1/s'))^2 \cdot \Delta^{O((\log r)^2)} = (m \cdot \mathrm{polylog}(T(n)))^{O((\log r)^2)}$. The verification step for V takes time $\log(1/s) \cdot \mathrm{poly}(n, \log T(n)) = \mathrm{poly}(n, \log T(n))$, and it makes at most $\log(1/s) \cdot \mathrm{polylog}(T(n)) = \mathrm{polylog}(T(n))$ queries, resulting in a total running time of $\mathrm{poly}(n) + (m \cdot \log T(n))^{O((\log r)^2)}$. Moreover, because V is non-adaptive, each execution of the evaluation protocol can be carried out in parallel, and thus the total number of rounds is four. Collapsing this protocol into a two-round one [BM88] leads to a prAM protocol with running time $\mathrm{poly}(n) \cdot (m \cdot \log T(n))^{O((\log r)^2)}$.

For the moreover part, we observe that computing $\mathrm{RMV}(\hat{g}_i)$ for each i only requires knowledge of m, the size of circuit D (instead of the circuit itself) and thus the generator H also only requires knowledge of m. Similarly, the commit-and-evaluate protocol in Lemma 23 only requires blackbox access to the deterministic predicate that underlies the circuit D, and thus so does our reconstructor R since it just gives D as input to P. $\qquad\square$

We remark that we can amplify the resilient soundness property for the reconstructor so that the probability that it outputs a value outside of $\{f(x)_i, \bot\}$ is at most 2^{-t} by running it $\Theta(t)$ times in parallel and outputting $\bot$ as soon as at least one of the answers is $\bot$ or the answers are inconsistent, and outputting the consistent answer bit otherwise.

We also present a version of the generator with a stronger resilient soundness property at the expense of increasing the complexity of the reconstructor from a promise Arthur-Merlin protocol to a probabilistic algorithm with parallel access to SAT. This version is useful for obtaining our byproducts in the average-case setting.

Corollary 25. *Let $T(n)$ be a time bound and $f \in \mathrm{NTIME}[T(n)]$. There exists a nondeterministic algorithm H (the generator) and a probabilistic algorithm R (the reconstructor) with parallel access to SAT that have the same properties as in Theorem 24 but such that the resilient soundness property holds for* every *co-nondeterministic circuit.*

In the setting of Corollary 25, item 2 of Theorem 24 should be interpreted as saying that $R(x, D)$ outputs $f(x)$ with probability at least $2/3$. We refer to the stronger resilience property in Corollary 25 as *strong resilient soundness*.

The idea behind Corollary 25 is for the reconstructor to first check whether the co-nondeterministic circuit D accepts at least somewhat less than half of its inputs. This is where the parallel access to an oracle for SAT comes in; it allows us to distinguish with high probability between the cases where the fraction of accepted inputs is, say, at most $1/3$ and at least $1/2$. In the former case, the new reconstructor indicates failure with high probability. Otherwise, we boost the fraction of accepted inputs to at least $1/2$ by trying D on two independent inputs, and then run the old reconstructor on the corresponding co-nondeterministic circuit D'.

Proof of Corollary 25. Let H' be the generator and R' the reconstructor of Theorem 24 instantiated with function f and amplified to have (resilient) soundness $1/6$.

Generator. The generator H, on input x and D of size m, first constructs the circuit D' of size $2m$ as $D'(r_1 r_2) = D(r_1) \vee D(r_2)$. We then define $H(x, D)$ as $\text{Left}(H'(x, D')) \cup \text{Right}(H'(x, D'))$, where $\text{Left}(S)$ and $\text{Right}(S)$ output the set of the left and right halves of every string in S, respectively.

Reconstructor. On input (x, D) and an index i, the reconstructor R estimates up to error $1/12$ and with probability of failure $1/6$ the fraction of inputs accepted by D by evaluating circuit D on $O(1)$ random inputs of length m, which can be done in probabilistic time $\text{poly}(m)$ with $O(1)$ parallel queries to a SAT oracle. If the estimated fraction is less than $5/12$ (the midpoint between $1/3$ and $1/2$), then R declares failure. In parallel, R builds the circuit D' in the same way as H, samples Arthur's randomness for protocol R' with inputs (x, D') and i and makes three queries to the SAT oracle to obtain the protocol's output: Whether there is a Merlin response that leads to success and whether there are Merlin responses that lead to outputting 0 and 1. If the first query is answered negatively, or the last two queries give inconsistent answers, then R declares failure. Otherwise, R outputs whatever R' does.

Strong resilient soundness. Consider two cases in relation to circuit D: Either D accepts fewer than $1/3$ of its inputs, or it accepts at least a $1/3$ of its inputs. In the first case, the initial verification fails with probability at least $5/6$. In the second case, D' accepts at least $2/3 - 1/9 = 5/9 > 1/2$ of its inputs. The resilient soundness property of protocol R' guarantees that with probability at least $5/6$, R either fails or outputs $f(x)$ correctly. In either case, it follows that R outputs an incorrect value for $f(x)$ with probability at most $1/6 \leq 2/3$.

Correctness. If a co-nondeterministic circuit D accepts at least half of its inputs, so does the circuit D'. Moreover, if $H(x, D)$ fails to hit D, then $H'(x, D')$ fails to hit D'. The correctness of protocol R' then guarantees that there exists a strategy for Merlin that makes R' output $f(x)$ with probability 1, and no strategy can make R' output an incorrect value for $f(x)$ with probability at least $1/6$. It follows that the second parallel phase of R yields $f(x)$ with probability at least $5/6$. Accounting for the error probability of $1/6$ in the initial verification, we conclude that R outputs $f(x)$ with probability at least $2/3$.

Efficiency. The running time of H is asymptotically identical to that of H', and the running time of R is polynomial in the running time of R'.

Finally, the moreover part follows right away from the moreover part of Theorem 24. □

Similar to the case of Theorem 24, we can amplify the strong resilient soundness property for the reconstructor so that the probability that it outputs a value outside of $\{f(x)_i, \bot\}$ (or different from $f(x)_i$ in case D is not hit by the generator) is at most 2^{-t} by running it $\Theta(t)$ times in parallel and outputting the majority answer.

4.3 Derandomization consequences

First, we present a generic derandomization result for prAM that works under hardness against arbitrary distributions.

Theorem 26. *There exists a constant c such that the following holds. Let $t, T : \mathbb{N} \to \mathbb{N}$ be time bounds such that $t(n) \geq n$, $\Pi \in \mathrm{prAMTIME}[t(n)]$ and $\{\mathbf{x}_n\}_{n \in \mathbb{N}}$ be an ensemble of distributions such that $\mathbf{x}_n$ is supported over $\{0,1\}^n$ and such that for all n, every x in the support of $\mathbf{x}_n$ satisfies the promise of Π. Assume that for $\mu : \mathbb{N} \to [0, 1)$ there exists a length-preserving function $f \in \mathrm{NTIME}[T(n)]$ such that for every $\mathrm{prAMTIME}[t(n)^{O((\log r)^2)}]$ protocol P for $r = O(\log(T(n))/\log(t(n)))$, it holds that the probability over $x \sim \mathbf{x}_n$ that $P(x) = f(x)$ is at most $\mu(n)$ for all but finitely many n. Then, it holds that*

$$\Pi \in \mathrm{Heur}_{\mathbf{x},\mu}\mathrm{NTIME}[T(n)^c].$$

Proof. First, notice that if $t(n) \leq \log T(n)$, then the conclusion is trivial and if $t(n) \geq T(n)$ then the premise is impossible, so we focus on the case that $\log T(n) \leq t(n) \leq T(n)$. Let $\Pi \in \mathrm{prAMTIME}[t(n)]$ and let P be a two-round protocol for Π running in time $O(t(n))$ on inputs of length n. On input $x \in \{0,1\}^n$, compute the circuit D_x of Proposition 17 with protocol P, and note that D_x has size $O(t(n)^2)$. Then, instantiate the HSG of Theorem 24 with f. Feed H inputs x and D_x and run the usual derandomization procedure for protocol P with the set output by $H(x, D_x)$: For each string $\rho \in H(x, D_x)$, nondeterministically guess Merlin's message y_ρ and compute the output of P with randomness ρ and message y_ρ, accepting if and only if P accepts for every $\rho \in H(x, D_x)$. The entire procedure runs in nondeterministic time $\mathrm{poly}(T(n), t(n)) = O(T(n)^c)$ for some constant c, since $T(n) \geq t(n)$.

Assume, with the intent of deriving a contradiction, that with probability at least $\mu(n)$ over $x \sim \mathbf{x}_n$, this derandomization fails for input x. First, notice that by the perfect completeness of P it must be the case that such an x lies in Π_N and that P with input x accepts every string in $H(x, D_x)$. Therefore, D_x acts as a distinguisher for $H(x, D_x)$, i.e., it rejects every string output by D_x while accepting at least half of its inputs. By computing D_x and feeding it to the prAM protocol R of Theorem 24, we obtain a prAM protocol that computes individual bits of $f(x)$ correctly for every x for which the derandomization fails, i.e., with probability at least $\mu(n)$ over $x \sim \mathbf{x}_n$. By running this protocol n times in parallel to compute every bit of $f(x)$, we obtain a prAM protocol that runs in time

$$\mathrm{poly}(n) \cdot (t(n) \cdot \log T(n))^{O((\log r)^2)} = t(n)^{O((\log r)^2)}$$

since $t(n) \geq \log T(n)$ and $t(n) \geq n$. This is a contradiction to the hardness of f so we are done. $\square$

We remark that we require hardness not just against AM protocols but against prAM protocols, which may not respect the completeness and/or soundness conditions on some inputs. However, an input of length n only contributes to the success fraction $\mu(n)$ provided the completeness and soundness conditions are met on that input.

As a consequence of Theorem 26, if the hardness assumption holds for almost-all inputs, then we obtain full derandomization of prAM.

Theorem 27. *There exists a constant c such that the following holds. Let $t, T : \mathbb{N} \to \mathbb{N}$ be time bounds such that $t(n) \geq n$. If there is a length-preserving function $f \in \mathrm{NTIME}[T(n)]$ that is hard on almost-all inputs against $\mathrm{prAMTIME}[t(n)^{O((\log r)^2)}]$ for $r = O(\log (T(n))/\log (t(n)))$, then*

$$\mathrm{prAMTIME}[t(n)] \subseteq \mathrm{NTIME}[T(n)^c].$$

Moreover, there exists a targeted hitting-set generator that achieves this derandomization result.

Proof. The statement follows from Theorem 26 by noting that the assumption that f is hard on almost-all inputs implies that f is hard for all possible distributions $\mathbf{x}_n$ with success probability $\mu(n) = 0$. In particular, the following nondeterministic algorithm is a hitting-set generator for prAM: On input $x \in \{0,1\}^*$ and a co-nondeterministic circuit C of size m, output $H(x, D)$ where H is the generator of Theorem 24 and $D \doteq C(x, \cdot)$. This algorithm has a successful computation path for any input and, on every successful computation path on inputs where D accepts at least half of its inputs, it outputs a set that hits D. The running time of the generator is $\mathrm{poly}(T(n), m)$. $\square$

Setting	$T(n)$	Hard for	Derandomization
high end	n^a	$n^{O((\log a)^2)}$	$\mathrm{prAM} \subseteq \mathrm{NP}$
middle-of-the-road	$2^{\mathrm{polylog}(n)}$	$n^{O((\log \log n)^2)}$	$\mathrm{prAM} \subseteq \mathrm{NTIME}[2^{\mathrm{polylog}(n)}]$
low end	$2^{n^{o(1)}}$	$n^{o((\log n)^2)}$	$\mathrm{prAM} \subseteq \mathrm{NTIME}[2^{n^{o(1)}}]$
very low end	$2^{\mathrm{poly}(n)}$	$n^{b(\log n)^2} \ \forall b$	$\exists c \ \mathrm{prAM} \subseteq \mathrm{NTIME}[2^{n^c}]$

Table 2: Derandomization consequences that follow from different instantiations of Theorem 27.

By setting parameters in Theorem 27, we obtain the derandomization results listed on Table 2. In particular, the first line of Table 2 establishes Theorem 5 and the last line establishes Theorem 6. We now provide more details on how to obtain each line of Table 2:

- For the high end, set $t(n) = n$, in which case $r = O(a)$. Then, $\mathrm{prAMTIME}[n] \subseteq \mathrm{NP}$ follows as long as f is hard on almost-all inputs against $\mathrm{prAMTIME}[n^{O((\log a)^2)}]$. The result for prAM follows by padding.

- For the middle-of-the-road result, set $t(n) = n$, in which case $r = \mathrm{polylog}(n)$. Then, $\mathrm{prAMTIME}[n] \subseteq \mathrm{NTIME}[2^{\mathrm{polylog}(n)}]$ follows as long as f is hard on almost-all inputs against $\mathrm{prAMTIME}[n^{O((\log \log n)^2)}]$. The result for prAM follows by padding.

- For the low end, let $\nu = \nu(n) = o(1)$ be such that $T(n) = 2^{n^\nu}$ and set $t(n) = n$. In this case, $r \leq n^\nu$. Then, $\mathrm{prAMTIME}[n] \subseteq \mathrm{NTIME}[\mathrm{poly}(n, 2^{n^\nu})]$ follows as long as f is hard on almost-all inputs against $\mathrm{prAMTIME}[n^{O((\nu \log n)^2)}]$. Since $\mathrm{poly}(n, 2^{n^\nu}) = 2^{n^{o(1)}}$ and $n^{O((\nu \log n)^2)} = n^{o((\log n)^2)}$, the result for prAM follows by padding.

- For the very low end, set $t(n) = n^b$ for a constant b, in which case $r = \mathrm{poly}(n)$. Then, $\mathrm{prAMTIME}[n^b] \subseteq \mathrm{NTIME}[2^{n^c}]$ for some constant c follows as long as f is hard on almost-all inputs against $\mathrm{prAMTIME}[n^{O(b(\log n)^2)}]$. To get the result for prAM, it suffices for hardness to hold for all b.

25

5 Consequences of derandomization

In this section, we prove the derandomization-to-hardness and derandomization-to-targeted HSGs directions of our near-equivalences.

5.1 Hardness on almost-all inputs

We start with our derandomization-to-hardness implication: If $\mathrm{prAM} \subseteq \mathrm{NP}$ then for all constants c there is a length-preserving function f computable in nondeterministic polynomial time (with a few bits of advice) that is hard on almost-all inputs against $\mathrm{AMTIME}[n^c]$. The basic idea is that, under the derandomization hypothesis, every (single-bit) AM protocol that runs in time n^c can be simulated by a single-valued nondeterministic machine without too much time overhead. If we have as advice whether a particular nondeterministic machine is single-valued or not at input length n, we can negate its input efficiently, obtaining a function f computable in nondeterministic time $\mathrm{poly}(n)$ that is almost-all inputs hard against AM protocols that run in time n^c. We now state Proposition 4 formally.

Proposition 28 (Formal version of Proposition 4). *If $\mathrm{prAM} \subseteq \mathrm{NP}$, then for every constant c and increasing function $\alpha : \mathbb{N} \to \mathbb{N}$ there exists a length-preserving function $f \in \mathrm{NP}/\alpha(n)$ that is hard on almost-all inputs against $\mathrm{AMTIME}[n^c]$.*

Proof. Assume that $\mathrm{prAM} \subseteq \mathrm{NP}$ and let c' be a constant to be defined later (which depends on c). The basic idea for the function f is as follows: On an input x of length n, we set its i-th output bit (for $1 \leq i \leq \min(n, \alpha(n))$) to the opposite of the i-th bit output by the i-th nondeterministic Turing machine N_i on input x (if N_i is single-valued and halts in at most $n^{c'+2}$ steps at input length n), and otherwise we set it to 0. Formally, on input x of length n and for $1 \leq i \leq n$

$$f(x)_i = \begin{cases} 1 - N_i(x)_i & \text{if } i \leq \alpha(n), N_i \text{ is single-valued and halts in at most } n^{c'+2} \text{ steps,} \\ 0 & \text{otherwise.} \end{cases}$$

Note that f is computable by a single-valued nondeterministic machine running in time $O(n^{c'+3})$ with $\alpha(n)$ bits of advice (indicating whether N_i is single-valued and halts in at most $n^{c'+2}$ steps at input length n for $1 \leq i \leq \alpha(n)$).[6] This holds because, when N_i is single-valued, computing $1 - N_i(x)_i$ can be done by guessing a path on which N_i succeeds, which must result in the unique value $N_i(x)$, and then outputting the opposite of the i-th bit of that. Assume, with the intent of deriving a contradiction, that there exists an AM protocol P that runs in time $O(n^c)$ and computes f on an infinite set of inputs $X \subseteq \{0,1\}^*$. Consider the protocol P' that takes as regular input a triple (x, i, b) and accepts iff the i-th bit of the output of protocol P with input x equals b (if $i > |x|$ then P' rejects). Note that P' induces a language L in $\mathrm{AMTIME}[n^c]$. Since $\mathrm{prAM} \subseteq \mathrm{NP}$ and $\mathrm{prAMTIME}[n^c]$ has a complete problem under linear-time reductions, it follows that there exists a constant c' such that $\mathrm{AMTIME}[n^c] \subseteq \mathrm{NTIME}[n^{c'}]$.[7] Let N be a nondeterministic machine that

---

[6] The nondeterministic machine computing f is only guaranteed to be single-valued when given the correct advice string.

[7] While our argument only requires that there exists a constant c' such that $\mathrm{AMTIME}[n^c] \subseteq \mathrm{NTIME}[n^{c'}]$, we use the assumption $\mathrm{prAM} \subseteq \mathrm{NP}$ instead of $\mathrm{AM} \subseteq \mathrm{NP}$ since it is unknown whether $\mathrm{AMTIME}[n^c]$ contains a complete problem under linear-time reductions.

runs in time $n^{c'}$ and computes L. Note that for every $x \in \{0,1\}^*$ and $1 \le i \le |x|$, $N(x,i,b) = 1$ for exactly one $b \in \{0,1\}$, and when $x \in X$, $N(x,i,b) = 1$ if and only if $f(x)_i = b$.

Now consider the following procedure N': On input $x \in \{0,1\}^n$, guess a value b_i and a witness y_i for each $1 \le i \le n$ and run $N(x,i,b_i;y_i)$. If for all i, $N(x,i,b_i;y_i)$ accepts, N' succeeds and prints the concatenation of the guessed b_i's, otherwise N' fails. Note that N' is a nondeterministic machine that runs in time $O(n^{c'+1})$. Moreover, by our assumption that P is an AM protocol and that prAM $\subseteq$ NP, N' is single-valued on *every* input. By construction, the single value equals $f(x)$ for all $x \in X$.

Let i be the index of N' in our enumeration, i.e., $N_i = N'$. By definition of f, for every input $x \in \{0,1\}^*$ of sufficiently large length $n \ge \alpha^{-1}(i)$ (so that it has a chance to negate the output of N_i), and in particular for all sufficiently large $x \in X$, we have that $f(x)_i = 1 - N'(x)_i = 1 - f(x)_i$, which is a contradiction. $\square$

This result extends to other parameter settings. As an example, we state a version of Proposition 28 at the very low end.

Proposition 29. *If there exists a constant c such that that* AM $\subseteq$ NTIME$[2^{n^c}]$, *then for every increasing function $\alpha : \mathbb{N} \to \mathbb{N}$ there exists a function $f \in$ NEXP$/\alpha(n)$ that is hard on almost-all inputs against* AM *protocols running in polynomial time.*

Proof (Sketch). The proof is essentially identical to that of Proposition 28, but with a different time bound. Since AM $\subseteq$ NTIME$[2^{n^c}]$, the diagonalizing machine N needs to diagonalize against single-valued nondeterministic algorithms running in time $2^{n^{c'}}$ for some fixed constant $c' > c$, and thus we get a nondeterministic algorithm that runs in time $O(2^{n^k})$ for any constant $k > c'$. $\square$

We conclude this section by noting in more detail where the gaps between our hardness-to-derandomization and derandomization-to-hardness results lie. The first gap lies in the fact that in the derandomization-to-hardness direction, the hard function f we construct requires a few bits of advice that we don't know how to handle in the other direction. There is, however, a subtler difference — In the hardness-to-derandomization direction, we require hardness against prAM protocols, which may not obey the AM promise on all inputs (though we only consider the protocol as computing $f(x)$ on input x if it obeys the promise and respects both completeness and soundness on input x). In the derandomization-to-hardness direction, we can only guarantee hardness against AM protocols, which necessarily obey the AM promise on all inputs. We remark that a similar problem shows up in other hardness vs. randomness tradeoffs for AM [GSTS03, SU09]. For example, to conclude almost-everywhere derandomization of AM, the authors of [GSTS03] require hardness of EXP against AM protocols for which completeness only holds infinitely-often. Finally, we also note that, while Chen and Tell only state their derandomization-to-hardness result for BPP [CT21], in that setting one can actually achieve hardness against prBPP (where the probabilistic algorithm might not have a high-probability output for every input).

5.2 Targeted hitting-set generator

In this section, we prove Theorem 7 along the lines of the intuition provided in Section 2.2. We make use of a win-win argument: Either the EXP $\ne$ NEXP hardness assumption holds, in which case there is a regular (oblivious) HSG that guarantees the derandomization result [IKW02]. Or else we may assume that EXP = NEXP, which allows us to construct a function f that is hard

against prAM protocols by diagonalization, with which we then instantiate Theorem 24 to obtain the targeted HSG.

We need the following result that follows from the "easy-witness" method.

Lemma 30 ([IKW02]). *If* NEXP $\neq$ EXP *then* prAM $\subseteq$ io-NTIME$[2^{n^\epsilon}]/n^\epsilon$ *for every* $\epsilon > 0$. *Moreover, there exists a (regular) HSG that achieves this derandomization.*

We now prove Theorem 7, which we restate here for convenience.

Theorem 7 (Restated). *If* prAMTIME$[2^{\text{polylog}(n)}]$ $\subseteq$ io-NEXP, *then there exists a targeted hitting-set generator for* prAM *that yields the simulation* prAM $\subseteq$ io-NTIME$[2^{n^c}]/n^\epsilon$ *for some constant c and all $\epsilon > 0$.*

Proof. If EXP $\neq$ NEXP, we are done by Lemma 30. Otherwise, it holds that NEXP = EXP. We use this collapse to construct a length-preserving multi-bit function $f \in$ EXP that is hard against prAMTIME$[n^{(\log n)^3}]$. We then instantiate Theorem 24 with f to obtain the targeted HSG. Hardness against protocols running in this time bound suffices along the lines of Theorem 6.

Before constructing f, we make an observation: Due to the instance-wise nature of our construction, to obtain an infinitely-often derandomization result using Theorem 24 it suffices to have an *infinitely-often all-inputs hardness assumption*. More precisely, we require the following: For every prAMTIME$[n^{(\log n)^3}]$ protocol P, there exist infinitely many input lengths n such that P fails to compute f for every x of length n. Thus, we construct a function f with this requirement in mind.

Under the hypothesized derandomization assumption and because prAMTIME$[n^{(\log n)^3}]$ has a complete problem under linear-time reductions, it follows that there exists a constant k such that prAMTIME$[n^{(\log n)^3}] \subseteq$ io-NTIME$[2^{n^k}]$. Since NTIME$[2^{n^k}]$ also has a complete problem under linear-time reductions, under the assumption EXP = NEXP, there exists a constant k' such that prAMTIME$[n^{(\log n)^3}] \subseteq$ io-DTIME$[2^{n^{k'}}]$. In that case, it suffices to diagonalize against fixed-exponential time machines to construct f. Similar to Proposition 28, we define the i-th bit of $f(x)$ to be the opposite of the i-th bit output by $M_i(x)$ when it runs for at most $2^{|x|^{k'+1}}$ steps, where M_i is the i-th deterministic Turing machine. Formally, on input x of length n and for $1 \leq i \leq n$,

$$f(x)_i = \begin{cases} 1 - M_i(x) & \text{if } M_i(x) \text{ halts in at most } 2^{n^{k'+1}} \text{ steps,} \\ 0 & \text{otherwise.} \end{cases}$$

Note that f is computable by a deterministic machine running in time $O(n \cdot 2^{n^{k'+1}})$ and thus $f \in$ EXP.

Assume, with the intent of deriving a contradiction, that there exists a prAMTIME$[n^{(\log n)^3}]$ protocol P such that for almost-all input lengths n, P computes f on at least one input $x \in \{0,1\}^n$, and call the set of inputs where P computes f correctly X. Again, similar to the proof of Proposition 28, P induces a problem Π in prAMTIME$[n^{(\log n)^3}]$, and by our assumptions, there is a language $L \in$ DTIME$[2^{n^{k'}}]$ such that L and Π agree on infinitely many input lengths. Let M be a deterministic Turing machine running in time $O(2^{n^{k'}})$ that decides L. Recall that yes-instances of Π are triples (x, i, b) such that $x \in X$ and $f(x)_i = b$ while no-instances have $x \in X$ and $f(x)_i \neq b$. Let M' be the deterministic Turing machine that, on input x of length n, outputs $M'(x)$ of length n such that $M'(x)_i = 1$ if and only if M accepts $(x, i, 1)$ for $1 \leq i \leq n$. Note that M' runs in time

$2^{n^{k'+1}}$. By construction and our assumption on P, for infinitely many input lengths n there exists at least one $x \in X \cap \{0,1\}^n$ such that $M'(x) = f(x)$.

Let i be the index of M' in our enumeration. By definition of f, for every input $x \in \{0,1\}^*$ of sufficiently large length $n \geq i$ (so that it has a chance to negate the output of M'), and in particular for all sufficiently large inputs $x \in X$, we have that $f(x)_i = 1 - M'(x)_i = 1 - f(x)_i$, a contradiction. Finally, we instantiate Theorem 26 with f to obtain a targeted HSG for prAM that runs in exponential time, which suffices to obtain the conclusion. $\qquad\square$

6 Derandomization under uniform worst-case hardness

Our technique also leads to new results in the traditional uniform worst-case setting. Under worst-case hardness against probabilistic algorithms with non-adaptive oracle access to SAT, we obtain average-case derandomization results for prAM. Moreover, by further strengthening the hardness assumption, we may also conclude full (infinitely-often) derandomization of prAM. As previously mentioned, these results extend to average-case derandomization of $\text{prBPP}_{||}^{\text{SAT}}$.

6.1 Average-case simulation

In this section, we develop our average-case derandomization results for prAM under worst-case uniform hardness assumptions (where hardness is against $\text{BPTIME}_{||}^{\text{SAT}}$). Our results in this setting work as follows: Assume there exists a hard language $L \in \text{NTIME}[T(n)] \cap \text{coNTIME}[T(n)]$. To derandomize some prAM protocol P on input length n, we first consider the hard language L at some suitable input length ℓ, which depends on the hardness of L (for Theorem 8, for example, we take $\ell = \Theta(\log n)$). Then we let f be the function that maps any input $x \in \{0,1\}^n$ to the truth-table of L at input length ℓ, and it follows from the complexity of L that $f \in \text{NTIME}[2^\ell \cdot T(\ell)]$. Finally, we instantiate our targeted HSG construction H with f and use it to derandomize P.

For the reconstruction, we make use of the strong resilient soundness property of Corollary 25. If the average-case derandomization fails, to decide whether z of length ℓ is in L, we first sample multiple candidate "good" strings x that hopefully lead to a distinguisher D_x for the generator (enough so that we expect at least one "good" x with high probability). Then, we run the reconstruction for all of them, accepting if and only if at least one of those output 1. By the strong resilient soundness property and amplification, with high probability every execution either fails or outputs $f(x)_z = L(z)$, and in the high probability case that we sample at least one "good" x, some execution outputs $L(z)$, meaning we can compute L efficiently on input length ℓ.

First, we present such a result at the high end of the derandomization spectrum.

Theorem 31 (Strengthening of Theorem 8). *If* $\text{NTIME}[2^{an}] \cap \text{coNTIME}[2^{an}]$ *is not included in* $\text{BPTIME}[2^{(\log(a+1))^2 n}]_{||}^{\text{SAT}}$ *for some constant* $a > 0$, *then for all* $e > 0$ *it holds that*

$$\text{prAM} \subseteq \text{io-Heur}_{1/n^e}\text{NP}$$
$$\text{prBPP}_{||}^{\text{SAT}} \subseteq \text{io-Heur}_{1/n^e}\text{P}_{||}^{\text{SAT}}.$$

Proof. We first argue the result for prAM. Consider derandomizing a prAM protocol P for a problem Π running in time $O(n^k)$ for some constant k. Let S be an $O(n^s)$-time sampler for a distribution in $\{0,1\}^n$ and e be a constant such that we want to "fool" S with probability at least $1 - 1/n^e$. Let f be a function mapping every $x \in \{0,1\}^n$ to the truth-table of the hard language

$L \in \text{NTIME}[2^{an}] \cap \text{coNTIME}[2^{an}]$ at input length $\ell = \ell(n) = \Theta(\log n)$ to be set precisely later. Note that $f \in \text{NTIME}[T(n)]$ for $T(n) = 2^{(a+1)\ell}$. Instantiate the generator H of Corollary 25 with f, run H on input $x = 0^n$ (recall f maps every string in $\{0,1\}^n$ to the same truth-table) and co-nondeterministic circuit size $m = O(n^{2k})$, and use it to attempt to derandomize P in nondeterministic time $\text{poly}(T(n), n^{2k}) = \text{poly}(n)$.

If the derandomization fails for almost-all input lengths, even heuristically, then for almost-all input lengths n, $S(1^n)$ outputs with probability at least $1/n^e$ a string $x \in \{0,1\}^n$ such that the simulation errs on x, i.e., the circuit D_x obtained from x and P using Proposition 17 is a distinguisher for $H(0^n, D_x)$. To compute L at input length ℓ, it then suffices to do the following: On input $z \in \{0,1\}^\ell$, first use S to sample $t = \Theta(n^e)$ inputs $x_1, \ldots, x_t$ and use these to construct a list $D_{x_1}, \ldots, D_{x_t}$ of candidate distinguishers for $H(0^n, D_x)$. With high probability, this list contains an actual distinguisher for the generator. Let R be the algorithm of Corollary 25, amplified by parallel repetition to have negligible soundness 2^{-n}, i.e., with probability at least $1 - 2^n$, the algorithm outputs either $f(x)$ or $\bot$. Finally, run R with inputs 0^n, index z (recall $f(0^n)$ equals the truth-table of L at input length ℓ) and D_{x_i} for every sampled input x_i, and accept if and only if some execution outputs 1. To see that this is correct, note that by a union bound, with high probability every execution of R is successful in the sense that it either outputs $f(0^n)_z = L(z)$ or $\bot$. Conditioned on there being a distinguisher in the list, we are guaranteed to output the correct value of $L(z)$ with high probability.

The running time for the reconstruction is $O(n^{e+s})$ for generating the $t = \Theta(n^e)$ samples, and $O(n^{2k})^{O((\log r)^2)}$ per sample for running R, where $r = O(((a+1)\ell)/(k \log n))$, for a total of $O(n^e(n^s + n^{O(k(\log r)^2)}))$. By setting $\ell = dk \log n$, we have that $r = O(d(a+1))$ and we can upper bound the total running time by $n^{O(e+s+k(\log(d(a+1)))^2)}$. In terms of the input length ℓ, this is $2^{(\log(a+1))^2 \ell}$ when d is a sufficiently large constant depending on a, e, s. This concludes the argument for prAM.

Now, we argue the result for $\text{prBPP}_{\parallel}^{\text{SAT}}$. To do so, we use the containment $\text{prBPP}_{\parallel}^{\text{SAT}} \subseteq \text{P}_{\parallel}^{\text{prAM}}$ [CR11]. It suffices to show that every deterministic polynomial-time algorithm with non-adaptive oracle access to a paddable prAM-complete problem $\Gamma \in \text{prAMTIME}[n]$ can be simulated by deterministic polynomial-time algorithms with non-adaptive oracle access to SAT. Let M be a deterministic algorithm with non-adaptive oracle access to Γ running in time $O(n^b)$ and S be an $O(n^s)$-time sampler that we want to "fool" with probability at least $1 - 1/n^e$. Since Γ is paddable, we may assume that every query made by M on inputs of length n is of length $O(n^b)$ (at the expense of increasing its running time to $O(n^{2b})$). To simulate M on input x, let f be a function mapping every $x \in \{0,1\}^n$ to the truth-table of L at input length $\ell = \ell(n) = \Theta(\log n)$. As before, $f \in \text{NTIME}[2^{(a+1)\ell}]$. Instantiate the generator H of Corollary 25 with f and use it to derandomize Γ at input length $O(n^b)$ in order to obtain a $\text{P}_{\parallel}^{\text{SAT}}$ simulation for M. Whenever M with input x queries Γ, we instead query the SAT oracle whether the nondeterministic simulation of Γ using H with input 0^n and co-nondeterministic circuit size $m = O(n^{2b})$ accepts. This simulation runs in $\text{P}_{\parallel}^{\text{SAT}}$ since M is non-adaptive.

If this derandomization fails on almost-all input lengths n, then as before we can use S to sample $t = \Theta(n^e)$ inputs $x_1, \ldots, x_t$ such that with high probability the simulation fails on some x_i. Let $Q(M, x)$ be the set of queries to Γ made by M on input x. If the simulation fails on x_i, it must be the case that some query q in $Q(M, x_i)$ (and also in the promise of Γ) was answered incorrectly. Since the protocol for Γ has perfect completeness, it must be the case that $q \in \Pi_N$ and that D_q is a distinguisher for $H(0^n, D_q)$. The reconstruction is as before though we use the sets $Q(M, x_i)$ for

$i \in [t]$ to obtain the list of candidate generators, and correctness follows by the same argument as in the prAM case. The running time analysis is similar to the one for the case of prAM. □

At the low end, we are able to obtain a slightly stronger average-case derandomization result. Instead of having a different simulation for each sampler, we obtain a single simulation (depending on the problem in $\text{prAM}/\text{prBPP}^{\text{SAT}}_{||}$ and the constant ϵ) that "fools" every polynomial-time sampler.

Theorem 32. *If* $\text{NEXP} \cap \text{coNEXP} \not\subseteq \text{BPTIME}[n^{b((\log n)^2)}]^{\text{SAT}}_{||}$ *for all* $b > 0$, *then for every* $\epsilon > 0$ *and all* $e > 0$

$$\text{prAM} \subseteq \text{io-Heur}_{1/n^e}\text{NTIME}[2^{n^\epsilon}]$$
$$\text{prBPP}^{\text{SAT}}_{||} \subseteq \text{io-Heur}_{1/n^e}\text{DTIME}[2^{n^\epsilon}]^{\text{SAT}}_{||}.$$

Moreover, for any Π *in* prAM *or* $\text{prBPP}^{\text{SAT}}_{||}$ *and* $\epsilon > 0$, *there is a single simulation that works for all* $e > 0$.

Proof. We begin with the argument for prAM. Let L be a hard language in $\text{NTIME}[2^{n^a}] \cap \text{coNTIME}[2^{n^a}]$ for some constant $a \geq 1$. Consider derandomizing a protocol P for a problem $\Pi \in \text{prAMTIME}[n^k]$ for constant k. Let $\epsilon > 0$ and f be the function mapping every $x \in \{0,1\}^n$ to the truth-table of L at input length $\ell = n^\epsilon$. Note that $f \in \text{NTIME}[T(n)]$ for $T(n) = 2^{n^{a\epsilon}}$. Instantiate the generator H of Corollary 25 with f, run H on input $x = 0^n$ and co-nondeterministic circuit size $m = O(n^{2k})$, and use it to derandomize P. The simulation runs in nondeterministic time $\text{poly}(T(n), n^{2k})$, which is at most $2^{n^{\epsilon'}}$ for any $\epsilon' > 0$ by taking a sufficiently small $\epsilon > 0$.

The reconstruction is identical to that of Theorem 31 but with $\ell = n^\epsilon$. The running time is $O(n^{e+s})$ to generate the samples and $(n^{2k})^{O((\log r)^2)}$ per sample for running R, where $r = O(\log (T(n))/\log n)$, for a total of $O(n^e(n^s + n^{O((\log r)^2)}))$. Given our parameter choices, $r = O(n^{a\epsilon})$, and the expression is upper bounded by $O(n^e(n^s + n^{O((a\epsilon \log n)^2)}))$. As the input length is $\ell = n^\epsilon$ for constant ϵ, there exists a constant b (depending on a, e, s, ϵ) such that the running time is upper bounded by $\ell^{b(\log n)^2}$. If hardness holds for all $b > 0$, then the same simulation works for any constant value of s and e, i.e., for any polynomial-time sampler and any inverse-polynomial error probability.

The proof for $\text{prBPP}^{\text{SAT}}_{||}$ is also almost identical to that of Theorem 31, where we derandomize the "oracle" Γ using the generator H from Corollary 25 instantiated with the function f that maps every $x \in \{0,1\}^n$ to the truth-table of L at input length $\ell = n^\epsilon$ and use a set of queries instead of a set of inputs to obtain the list of candidate distinguishers for the reconstruction. This approach naturally leads to a simulation in $\text{P}^{\text{NTIME}[2^{n^\epsilon}]}_{||}$, and we obtain the $\text{DTIME}[2^{n^\epsilon}]^{\text{SAT}}_{||}$ simulation by replacing the original queries with padded SAT queries.

□

6.2 Infinitely-often all-input simulation

By introducing nondeterminism in the algorithms we require hardness for, we are able to extend Theorem 8 to conclude full (infinitely-often) derandomization of prAM. We have shown that, if the HSG construction of Theorem 8 fails to obtain average-case derandomization of prAM, then we are able to efficiently sample candidate distinguishers with the hope that at least one is "good". However, if the HSG fails in the worst case, it is harder to pinpoint exactly where it does so as to

obtain a distinguisher. To solve this, we have Merlin send a "good" input x. This necessitates a lower bound against $\text{MATIME}^{\text{SAT}}_{\|}$, but allows for concluding full (infinitely-often) derandomization of prAM and $\text{prBPP}^{\text{SAT}}_{\|}$.

Theorem 33. *If* $\text{NTIME}[2^{an}] \cap \text{coNTIME}[2^{an}] \not\subseteq \text{MATIME}[2^{(\log(a+1))^2 \ell}]^{\text{SAT}}_{\|}$ *for some constant* $a > 0$, *then*

$$\text{prAM} \subseteq \text{io-NP}$$
$$\text{prBPP}^{\text{SAT}}_{\|} \subseteq \text{io-P}^{\text{SAT}}_{\|}.$$

Proof. We argue the result for prAM first. Let $\Pi \in \text{prAMTIME}[n^k]$ for some constant k and let L be a hard language in $\text{NTIME}[2^{an}] \cap \text{coNTIME}[2^{an}]$. Let f be a function mapping every string in $\{0,1\}^n$ to the truth-table of L at input length $\ell = \Theta(\log n)$ to be set precisely later. Note that $f \in \text{NTIME}[T(n)]$ for $T(n) = 2^{(a+1)\ell}$. Instantiate the generator H of Corollary 25 with f, run H on input 0^n and co-nondeterministic circuit size $m = O(n^{2k})$, and use it to derandomize P in time $\text{poly}(T(n), n) = \text{poly}(n)$.

If the simulation fails for some input of almost-all input lengths, then for almost-all input lengths n there exists an $x \in \Pi_N$ of length n such that the simulation errs on x, i.e., the circuit D_x of Proposition 17 instantiated with the protocol for Π and x is a distinguisher for $H(0^n, D_x)$. Let R be the reconstructor of Corollary 25 and consider the following Merlin-Arthur protocol for L, where the protocol has parallel oracle access to SAT: On input $z \in \{0,1\}^\ell$, Merlin sends x, and Arthur runs $R(0^n, D_x)$ to compute the z-th bit of $f(0^n)$ (which equals $L(z)$). If R outputs $\bot$, then the protocol rejects, otherwise, it accepts if and only if R outputs 1. Because R is a probabilistic algorithm with parallel access to an oracle for SAT, Arthur can sample the randomness required for it and then run the underlying deterministic parallel-SAT-oracle computation, meaning this is indeed a $\text{MA}^{\text{SAT}}_{\|}$ protocol. Completeness follows since Merlin can send a correct value of x, and soundness follows from the strong resilience property of R: Even if Merlin sends a "bad" x', R is still guaranteed to either fail or output $L(z)$ with high probability.

To finish the argument for prAM, note that the running time of the protocol is just the running time of R, which is $\text{poly}(n) \cdot (m \cdot \log T(n))^{O((\log r)^2)}$ for $r = O(\log(T(n))/\log m)$. Since $m = O(n^{2k})$ and setting $\ell = dk \log n$, we have $r = O(d(a+1))$ and the running time for the protocol is upper bounded by $n^{O(k(\log(d(a+1)))^2)}$. In terms of the input length ℓ, this is $2^{(\log(a+1))^2 \ell}$ when d is a sufficiently large constant depending on a.

The simulation for $\text{prBPP}^{\text{SAT}}_{\|}$ is similar to before and the reconstruction is identical to the prAM case: If the simulation fails, then there is a query q of length $O(n^k)$ (which results in a distinguisher of size $O(n^{2k})$) that Merlin can send Arthur to make Arthur output $L(z)$ with high probability. Soundness also follows exactly as in the prAM case and the running time is again $2^{(\log(a+1))^2 \ell}$. $\qquad\square$

We only state the previous result for the high-end parameter setting because stronger results are already known for the low end. For example, to conclude a subexponential derandomization of prAM, it suffices for there to exist a language in $\text{NEXP} \cap \text{coNEXP}$ that is hard for a subclass of $\text{MA}^{\text{SAT}}_{\|}$ [AvM17]. In comparison with ours, other results that conclude the same derandomization either require hardness of nondeterministic algorithms against much larger deterministic time bounds, e.g., $\text{NE} \cap \text{coNE} \not\subseteq \text{DTIME}[2^{2^{n^\epsilon}}]$ for some $\epsilon > 0$ [IKW02] or hardness of deterministic algorithms against slightly less space, e.g., $\text{E} \not\subseteq \text{SPACE}[2^{\epsilon n}]$ for some $\epsilon > 0$ [Lu00].

7 Unconditional mild derandomization

In this section, we establish our unconditional mild derandomization result for prAM and extend it to $\text{prBPP}_{||}^{\text{SAT}}$. We employ a similar win-win argument to that of the proof of Theorem 7: Either some hardness assumption/class separation holds (here, $\Sigma_2\text{EXP} \not\subseteq \text{NP/poly}$), in which case we get derandomization right away. Or else we get a complexity collapse which we can use to construct a hard function f that has the efficiency requirements we need to apply one of our targeted hitting-set constructions (in this case Theorem 32, which requires hardness against $\text{BPTIME}[2^{\text{polylog}(n)}]_{||}^{\text{SAT}}$).

As a first step toward the win-win argument, we prove an "easy-witness lemma" for $\Sigma_2\text{EXP}$, which allows for the collapse $\text{P}^{\Sigma_2\text{EXP}} \subseteq \text{EXP}$ from the assumption that $\Sigma_2\text{EXP} \subseteq \text{NP/poly}$. Then we consider two cases:

- $\Sigma_2\text{EXP} \not\subseteq \text{NP/poly}$. In this case, the derandomization result follows from standard hardness vs. randomness tradeoffs.

- $\Sigma_2\text{EXP} \subseteq \text{NP/poly}$. In this case, we diagonalize against $\text{BPTIME}[2^{\text{polylog}(n)}]_{||}^{\text{SAT}}$ in $\text{P}^{\Sigma_2\text{EXP}} = \text{EXP}$, and then instantiate Theorem 32 to conclude the proof.

To diagonalize against $\text{BPTIME}[2^{\text{polylog}(n)}]_{||}^{\text{SAT}}$, we again make use of the inclusion $\text{prBPP}_{||}^{\text{SAT}} \subseteq \text{P}_{||}^{\text{prAM}}$ and diagonalize against deterministic algorithms with non-adaptive oracle access to prAM instead.

7.1 Nondeterministic easy witnesses

In this section, we prove our "easy witness lemma" for $\Sigma_2\text{EXP}$. One way of thinking of Σ_2 computations is as follows: On input x, guess a string y and then run a co-nondeterministic verifier on input (x, y). This view allows us to abstract the co-nondeterministic verification and think of y as a witness for x. In this section, we show that if $\Sigma_2\text{EXP} \subseteq \text{NP/poly}$, then every language in $\Sigma_2\text{EXP}$ has witnesses that are the truth-tables of functions computed by polynomial-size single-valued circuits. To do so, we use the following result to convert hardness against single-valued circuits into hitting sets for co-nondeterministic circuits.

Lemma 34 ([Uma03]). *There is a universal constant b and and a deterministic polynomial-time algorithm that, on input 1^m and a truth-table y of a function with single-valued circuit complexity at least m^b, outputs a set S of size $O(|y|^b)$ that hits co-nondeterministic circuits of size m that accept at least half of their inputs.*

We also need the following equivalence from [AvM17].

Lemma 35 ([AvM17]). $\Sigma_2\text{EXP} \not\subseteq \text{NP/poly}$ *if and only if* $\text{prAM} \subseteq \text{io-}\Sigma_2\text{TIME}[2^{n^\epsilon}]/n^\epsilon$ *for all* $\epsilon > 0$.

We are now ready to prove our easy witness result for $\Sigma_2\text{EXP}$.

Theorem 36. *Assume $\Sigma_2\text{EXP} \subseteq \text{NP/poly}$. Then $\Sigma_2\text{EXP}$ has single-valued witnesses of polynomial size, i.e., for every $L \in \Sigma_2\text{EXP}$ and linear-time (in its input length) co-nondeterministic verifier H for L, the following holds: For every $x \in L$, there exists a single-valued circuit C_x of size $\text{poly}(|x|)$ such that $H(x, \cdot)$ accepts the exponential-length truth-table of C_x.*

Proof. We show that $\Sigma_2\mathrm{E}$ has single-valued witness circuits of size n^c for some constant c. The result for $\Sigma_2\mathrm{EXP}$ then follows by padding.

Assume that $\Sigma_2\mathrm{E}$ does not have single-valued witness circuits of size n^c for any constant c. This implies that for all $c \geq 1$, there is a co-nondeterministic verifier H_c that takes as input a string x and a string y of length $2^{O(|x|)}$, runs in time $2^{O(|x|)}$, and has the following property: For infinitely many n, there is a input x' of length n such that $H_c(x', y')$ accepts for some y', but *every* y accepted by $H_c(x', \cdot)$ has single-valued circuit complexity at least n^c. Thus, there are infinitely many n such that, if we give x' as n bits of advice, guess a string y of length $2^{O(n)}$, and verify that $H_c(x', y)$ accepts (using co-nondeterminism), we are guaranteed that y encodes the truth-table of a function with single-valued circuit complexity at least n^c. This gives us a Σ_2-procedure for obtaining hard functions, which we use to derandomize prAM and obtain a contradiction to Lemma 35.

Let $\Pi \in \mathrm{prAM}$ and let P be a protocol for Π that runs in time $O(\ell^a)$ on input length ℓ. By Proposition 17, to derandomize P it suffices to have a set S that hits any co-nondeterministic circuit of size $O(\ell^{2a})$ that accepts at least half of its inputs. To obtain such a set using Lemma 34, we need to first obtain a truth-table of single-valued circuit complexity at least $\Omega(\ell^{2ab})$, where b is the constant from the lemma. Recall that our objective is to obtain a subexponential (time 2^{n^ϵ} for all $\epsilon > 0$) simulation. To this end, let $\epsilon > 0$ be sufficiently small and consider the verifier H_c for $c = 3ab/\epsilon$ on inputs of length $n = \ell^\epsilon$. If n is one of the infinitely many input lengths for which there exists x' such that every string accepted by $H_c(x', \cdot)$ has single-valued circuit complexity at least $n^c = \ell^{3ab}$, then we can obtain such a hard string by having x' as advice, guessing $y \in \{0,1\}^{2^{O(\ell^\epsilon)}}$ and verifying that $H_c(x', y)$ accepts.

In parallel, apply Lemma 34 to y to obtain a set S of size $2^{O(\ell^\epsilon)}$, and use S to derandomize the prAM computation (guessing a Merlin response for each string in S). Finally, accept if and only if both $H_c(x', y)$ and the prAM simulation accept. All of this can be carried out in $\Sigma_2\mathrm{TIME}[2^{O(\ell^\epsilon)}]/\ell^\epsilon$. Since ϵ is an arbitrarily small constant and the simulation works for infinitely many input lengths ℓ, we obtain a contradiction to Lemma 35. $\square$

Theorem 36 allows us to establish the following complexity class collapse in case $\Sigma_2\mathrm{EXP} \subseteq \mathrm{NP}/\mathrm{poly}$. The corollary represents the role our easy witness result plays in the proof of Theorem 9.

Corollary 37. *If* $\Sigma_2\mathrm{EXP} \subseteq \mathrm{NP}/\mathrm{poly}$, *then* $\mathrm{P}^{\Sigma_2\mathrm{EXP}} = \mathrm{EXP}$.

Proof. Under the hypothesis from the statement, we show that $\Sigma_2\mathrm{EXP} = \mathrm{coNEXP}$, which suffices by combining Lemma 35 and Lemma 30. The hypothesis and Lemma 35 guarantee the negation of $\mathrm{prAM} \subseteq \mathrm{io}\text{-}\Sigma_2\mathrm{TIME}[2^{n^\epsilon}]/n^\epsilon$ for all ϵ, which in turn implies the negation of $\mathrm{prAM} \subseteq \mathrm{io}\text{-}\mathrm{NTIME}[2^{n^\epsilon}]/n^\epsilon$ for all ϵ, and thus the contrapositive of Lemma 30 implies $\mathrm{EXP} = \mathrm{NEXP}$ and finally $\Sigma_2\mathrm{EXP} = \mathrm{coNEXP} = \mathrm{EXP}$.

To show that $\Sigma_2\mathrm{EXP} = \mathrm{coNEXP}$, by padding, it suffices to show that every $L \in \Sigma_2\mathrm{E}$ is in coNEXP. Fix $L \in \Sigma_2\mathrm{E}$. By Theorem 36, L has single-valued witnesses of size n^c for some constant c. On input $x \in \{0,1\}^n$, we cycle through all nondeterministic circuits C of size n^c and compute their truth-tables in time $O(2^{n^c})$. For each truth-table T, we then run $V(x, T)$ (where V is a co-nondeterministic verifier for L), accepting if and only if some verification accepts. All of this runs in exponential co-nondeterministic time, so we are done. $\square$

7.2 Simulation

We now execute our win-win strategy and establish Theorem 9 and its strengthening for $\mathrm{prBPP}_{||}^{\mathrm{SAT}}$ in lieu of prAM. We first consider the case where $\Sigma_2\mathrm{EXP} \not\subseteq \mathrm{NP/poly}$. In this case simulations of the required type that work on all inputs of a given length are provided by Lemma 35 for prAM. We argue the same simulations follow for $\mathrm{prBPP}_{||}^{\mathrm{SAT}}$.

Lemma 38. *If $\Sigma_2\mathrm{EXP} \not\subseteq \mathrm{NP/poly}$, then for every $\epsilon > 0$*

$$\mathrm{prBPP}_{||}^{\mathrm{SAT}} \subseteq \text{io-}\Sigma_2\mathrm{TIME}[2^{n^\epsilon}]/n^\epsilon.$$

Proof. We use the inclusion $\mathrm{prBPP}_{||}^{\mathrm{SAT}} \subseteq \mathrm{P}_{||}^{\mathrm{prAM}}$. Let k be a constant and M be an $O(n^k)$-time deterministic machine with non-adaptive oracle access to a paddable prAM-complete problem $\Gamma \in \mathrm{prAMTIME}[n]$. We assume that all queries made by M on inputs of length n are of length $O(n^k)$ at the expense of increasing M's running time to $O(n^{2k})$.

Our approach is to use Lemma 34 to derandomize the queries made to Γ while making sure that the overall simulation of M can be carried out in subexponential Σ_2. To derandomize Γ at input length $O(n^k)$ using the lemma, we need to obtain a truth-table of single-valued circuit complexity at least $\Omega(n^{2bk})$, where b is the constant from the lemma. Let $\epsilon > 0$ and $L \in \Sigma_2\mathrm{E}$ be a language that has nondeterministic circuit complexity at least $n^{3bk/\epsilon}$ for infinitely many input lengths (which is guaranteed to exist by the hypothesis of the theorem). The simulation of M on inputs x goes as follows: Given as advice the number of strings of length n^ϵ that are in L, the Σ_2 algorithm guesses the truth-table of L at input length n^ϵ, verifies it, and uses it as the string y in Lemma 34. More precisely, after guessing the truth-table, the algorithm performs the following operations in parallel:

○ It uses an existential and an universal guess to verify that the guessed truth-table for L is correct. This is possible because the algorithm has as advice the number of strings of length n^ϵ that are in L, and thus it can existentially guess which strings are in L and only verify those, with the guarantee that the others are not in L.

○ It guesses which of the queries to Γ that M makes on input x are answered positively and which are answered negatively. For each query that is guessed to be answered positively, it uses the set S from Lemma 34 and the existential phase to verify that there is a random-bit string in S for which Merlin can provide a witness. Similarly, it uses S and the universal phase to verify each query that is guessed to be answered negatively.

We note that the only existential computation paths that survive the computation are the ones where the truth-table of L at input length n^ϵ was guessed correctly. In this case, and in the case that n^ϵ is one of the infinitely many input lengths where L has nondeterministic circuit complexity at least $n^{3bk/\epsilon}$, it holds that the guessed truth-table has high enough (single-valued) nondeterministic circuit complexity such that S hits the co-nondeterministic circuits given by Proposition 17 for negative instances of Γ at input length $O(n^k)$. This further guarantees that the surviving existential computation paths are those that correctly guess the answers to all queries M makes on input x that are in the promise of Γ. This suffices to obtain a simulation of M that is correct on infinitely many input lengths since M is insensitive to variations in answers to queries that are outside the promise (even when the same query is answered differently on different occasions). Finally, we note that the entire procedure runs in time $2^{O(n^\epsilon)}$, which can be made smaller than $2^{n^{\epsilon'}}$ for any $\epsilon' > 0$ by taking ϵ to be sufficiently small. □

The other case of the win-win analysis is when $\Sigma_2\mathrm{EXP} \subseteq \mathrm{NP}/\mathrm{poly}$. In this case, we use the collapse $\mathrm{P}^{\Sigma_2\mathrm{EXP}} = \mathrm{EXP}$ given by Corollary 37 and our targeted hitting-generator construction to obtain the desired simulation. We conclude:

Theorem 39 (Strengthening of Theorem 9). *For every $\epsilon > 0$ and every $e > 0$*

$$\mathrm{prBPP}^{\mathrm{SAT}}_{||} \subseteq \text{io-Heur}_{1/n^e}\Sigma_2\mathrm{TIME}[2^{n^\epsilon}]/n^\epsilon.$$

Proof. If $\Sigma_2\mathrm{EXP} \not\subseteq \mathrm{NP}/\mathrm{poly}$, then it follows that $\mathrm{prBPP}^{\mathrm{SAT}}_{||} \subseteq \Sigma_2\mathrm{TIME}[2^{n^\epsilon}]/n^\epsilon$ for all $\epsilon > 0$ by Lemma 38. Otherwise, by Corollary 37, we have that $\mathrm{P}^{\Sigma_2\mathrm{EXP}} = \mathrm{EXP}$. By Theorem 31, all we need to show is that $\mathrm{P}^{\Sigma_2\mathrm{EXP}} \not\subseteq \bigcup_{b\in\mathbb{N}}\mathrm{BPTIME}[n^{b((\log n)^2)}]^{\mathrm{SAT}}_{||}$. Given the containment $\mathrm{prBPP}^{\mathrm{SAT}}_{||} \subseteq \mathrm{P}^{\mathrm{prAM}}_{||}$ and a padding argument, it follows that $\bigcup_{b\in\mathbb{N}}\mathrm{BPTIME}[n^{b((\log n)^2)}]^{\mathrm{SAT}}_{||} \subseteq \mathrm{DTIME}[2^{\mathrm{polylog}(n)}]^{\mathrm{prAM}}_{||}$. It remains to show that $\mathrm{P}^{\Sigma_2\mathrm{EXP}} \not\subseteq \mathrm{DTIME}[2^{\mathrm{polylog}(n)}]^{\mathrm{prAM}}_{||}$, which we do by diagonalization.

Fix a prAM-complete problem Γ and note that if $L \in \mathrm{DTIME}[2^{\mathrm{polylog}(n)}]^{\mathrm{prAM}}_{||}$, then there exists a Turing machine M running in time $2^{\mathrm{polylog}(n)}$ with non-adaptive oracle access to Γ that computes L. Thus, it suffices to diagonalize against such machines with Γ as an oracle. Let S be the following $\Sigma_2\mathrm{EXP}$-oracle machine: On input $x \in \{0,1\}^n$, interpret x as a non-adaptive oracle Turing machine M_x with an oracle for Γ. Then, using binary search and the $\Sigma_2\mathrm{EXP}$ oracle, compute the number q of queries that M_x on input x makes that are answered negatively, where we let M_x run for at most 2^n steps. This is possible in $\mathrm{P}^{\Sigma_2\mathrm{EXP}}$ because $\mathrm{prAM} \subseteq \Pi_2\mathrm{P}$, so we can verify negative instances in $\Sigma_2\mathrm{EXP}$. Once we know q, we can simulate $M_x(x)$ for at most 2^n steps in $\Sigma_2\mathrm{EXP}$ as follows: Guess which q queries are negative and verify them in $\Sigma_2\mathrm{EXP}$ (again using the fact that $\mathrm{prAM} \subseteq \Pi_2\mathrm{P}$); then assume that the other queries are answered positively and simulate $M_x(x)$ directly with these answers. By querying the $\Sigma_2\mathrm{EXP}$ oracle S then outputs the opposite of this simulation. By construction, the language of S is in $\mathrm{P}^{\Sigma_2\mathrm{EXP}} \setminus \mathrm{DTIME}[2^{\mathrm{polylog}(n)}]^{\mathrm{prAM}}_{||}$. $\square$

This concludes our discussion of the byproducts of our main results.

Acknowledgments

References

[AvM17] Barış Aydınlıoğlu and Dieter van Melkebeek. Nondeterministic circuit lower bounds from mildly derandomizing Arthur-Merlin games. *Computational Complexity*, 26(1):79–118, 2017.

[BFNW93] László Babai, Lance Fortnow, Noam Nisan, and Avi Wigderson. BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3(4):307–318, 1993.

[BGH+06] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Robust PCPs of proximity, shorter PCPs, and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006.

[BM88] László Babai and Shlomo Moran. Arthur-Merlin games: A randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36(2):254–276, 1988.

[CR11] Venkatesan T. Chakaravarthy and Sambuddha Roy. Arthur and Merlin as oracles. *Computational Complexity*, 20(3):505–558, 2011.

[CRT22] Lijie Chen, Ron D. Rothblum, and Roei Tell. Unstructured hardness to average-case randomness. Technical Report TR22-097, Electronic Colloquium on Computational Complexity (ECCC), 2022.

[CRTY20] Lijie Chen, Ron D. Rothblum, Roei Tell, and Eylon Yogev. On exponential-time hypotheses, derandomization, and circuit lower bounds: Extended abstract. In *Symposium on Foundations of Computer Science (FOCS)*, pages 13–23, 2020.

[CT21] Lijie Chen and Roei Tell. Hardness vs randomness, revised: Uniform, non-black-box, and instance-wise. In *Symposium on Foundations of Computer Science (FOCS)*, 2021. Full version available as ECCC Technical Report TR21-80.

[FGM+89] Martin Furer, Oded Goldreich, Yishay Mansour, Michael Sipser, and Stahis Zachos. On completeness and soundness in interactive proof systems. *Advances in Computing Research.*, 5:429–442, 1989.

[GKR15] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. *Journal of the ACM*, 62(4), 2015.

[Gol11] Oded Goldreich. In a world of P=BPP. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 191–232. Springer, 2011. Part of the Lecture Notes in Computer Science book series (LNCS, volume 6650).

[Gol18] Oded Goldreich. On doubly-efficient interactive proof systems. *Foundations and Trends in Theoretical Computer Science*, 13:157–246, 2018.

[Gol20] Oded Goldreich. Two comments on targeted canonical derandomizers. In *Computational Complexity and Property Testing: On the Interplay Between Randomness and Computation*, pages 24–35. Springer, 2020.

[GSTS03] Dan Gutfreund, Ronen Shaltiel, and Amnon Ta-Shma. Uniform hardness versus randomness tradeoffs for Arthur-Merlin games. *Computational Complexity*, 12(3):85–130, 2003.

[IKW02] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. In search of an easy witness: Exponential time vs. probabilistic polynomial time. *Journal of Computer and System Sciences*, 65(4):672 – 694, 2002.

[IW97] Russell Impagliazzo and Avi Wigderson. P = BPP if E requires exponential circuits: Derandomizing the XOR lemma. In *Symposium on Theory of Computing (STOC)*, page 220–229, 1997.

[IW01] Russell Impagliazzo and Avi Wigderson. Randomness vs time: Derandomization under a uniform assumption. *Journal of Computer and System Sciences*, 63(4):672–688, 2001.

[KvM02] Adam R. Klivans and Dieter van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM Journal on Computing*, 31(5):1501–1526, 2002.

[Liu22] Yanyi Liu. Personal communication, October 2022.

[LP22a] Yanyi Liu and Rafael Pass. Characterizing derandomization through hardness of Levin-Kolmogorov complexity. In *Computational Complexity Conference (CCC)*, volume 234, pages 35:1–35:17, 2022.

[LP22b] Yanyi Liu and Rafael Pass. Leakage-resilient hardness v.s. randomness. Technical Report TR22-113, Electronic Colloquium on Computational Complexity (ECCC), 2022.

[Lu00] Chi-Jen Lu. Derandomizing Arthur-Merlin games under uniform assumptions. In *Algorithms and Computation*, pages 302–312. Springer, 2000.

[MV05] Peter Bro Miltersen and N. V. Vinodchandran. Derandomizing Arthur–Merlin games using hitting sets. *Computational Complexity*, 14(3):256–279, 2005.

[NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994.

[SU05] Ronen Shaltiel and Christopher Umans. Simple extractors for all min-entropies and a new pseudorandom generator. *Journal of the ACM*, 52(2):172–216, 2005.

[SU09] Ronen Shaltiel and Christopher Umans. Low-end uniform hardness versus randomness tradeoffs for AM. *SIAM Journal on Computing*, 39(3):1006–1037, 2009.

[TV07] Luca Trevisan and Salil Vadhan. Pseudorandomness and average-case complexity via uniform reductions. *Computational Complexity*, 16(4):331–364, 2007.

[Uma03] Christopher Umans. Pseudo-random generators for all hardnesses. *Journal of Computer and System Sciences*, 67(2):419 – 440, 2003.

[Wil16] R. Ryan Williams. Natural proofs versus derandomization. *SIAM Journal on Computing*, 45(2):497–529, 2016.