

Nondeterministic Circuit Lower Bounds from Mildly Derandomizing Arthur-Merlin Games

Barış Aydınlioğlu * Dieter van Melkebeek †

June 18, 2012

Abstract

In several settings derandomization is known to follow from circuit lower bounds that themselves are equivalent to the existence of pseudorandom generators. This leaves open the question whether derandomization implies the circuit lower bounds that are known to imply it, i.e., whether the ability to derandomize in *any* way implies the ability to do so in the canonical way through pseudorandom generators.

For the setting of decision problems, Impagliazzo et al. implicitly showed the following equivalence: Randomized polynomial-time decision procedures can be simulated in NSUBEXP (the subexponential version of NP) with subpolynomial advice on infinitely many input lengths if *and only if* $\text{NEXP} \not\subseteq \text{P/poly}$. We establish a full analogue in the setting of verification procedures: Arthur-Merlin games can be simulated in $\Sigma_2\text{SUBEXP}$ (the subexponential version of $\Sigma_2\text{P}$) with subpolynomial advice on infinitely many input lengths if *and only if* $\Sigma_2\text{EXP} \not\subseteq \text{NP/poly}$.

A key ingredient in our proofs is improved Karp-Lipton style collapse results for nondeterministic circuits. The following are two instantiations that may be of independent interest: Assuming that Arthur-Merlin games can be derandomized in $\Sigma_2\text{P}$, we show that (i) $\text{PSPACE} \subseteq \text{NP/poly}$ implies $\text{PSPACE} \subseteq \Sigma_2\text{P}$, and (ii) $\text{coNP} \subseteq \text{NP/poly}$ implies $\text{PH} \subseteq \text{P}^{\Sigma_2\text{P}}$.

In proving our result we provide a general framework that also captures earlier conditional circuit lower bound results similar to ours.

1 Introduction

The power of randomness constitutes a central topic in complexity theory. In the context of randomized decision procedures the question is whether the class BPP , or its promise version prBPP , can be computed deterministically without much overhead – in subexponential or maybe even polynomial time. Similarly, in the context of randomized verification procedures one seeks for efficient nondeterministic computations of Arthur-Merlin games: the class AM , or its promise version prAM .

A major development in the area are hardness versus randomness tradeoffs [Yao82, NW94, BFNW93, IW97], which state that either nonuniformity speeds up computations significantly or else nontrivial derandomization is possible. More precisely, these results show how to use a language

*Partially supported by NSF grant 1017597.

†Partially supported by NSF grant 1017597.

in some complexity class \mathcal{C} that is *assumed* to require “large” circuits, to construct a pseudorandom generator (PRG) computable in \mathcal{C} with “small” seed length. The PRG transforms its seed into a longer string, say of length s , such that the average behavior of any circuit C of size s is almost the same when the input to C is provided from the uniform distribution or from the output distribution of the PRG on a uniform seed. We say that the PRG fools the circuit C . If \mathcal{C} requires large circuits of a certain type τ , then the resulting PRG fools circuits of type τ , and can be used to derandomize any procedure that can be modeled by small circuits of type τ [KvM02]. See the table below for some examples from the above papers and [SU06], where $\mathbf{E} \doteq \mathbf{DTIME}(2^{O(n)})$, $\mathbf{NE} \doteq \mathbf{NTIME}(2^{O(n)})$, and $\tau \in \{\mathbf{d}, \mathbf{n}\}$ denotes deterministic (d) or nondeterministic (n) circuits.

τ	class \mathcal{C}	randomized class	derandomization
d	\mathbf{E}	\mathbf{prBPP}	$\mathbf{DTIME}(t)$
n	$\mathbf{NE} \cap \mathbf{coNE}$	\mathbf{prAM}	$\mathbf{NTIME}(t)$

Once we have such a PRG, the derandomization is obtained by cycling over all seeds and simulating the randomized procedure on the output of the PRG for each seed. Stronger circuit lower bounds for \mathcal{C} imply smaller seed lengths for the generator, yielding more efficient derandomizations. At the “low end”, superpolynomial circuit lower bounds yield subpolynomial seed length and derandomizations that run in subexponential time t . At the “high end”, linear-exponential circuit lower bounds yield logarithmic seed length and derandomizations that run in polynomial time t .

As the circuit lower bounds seem plausible, even at the high end, the hardness versus randomness tradeoffs have fueled the conjecture that \mathbf{prBPP} can be fully derandomized to \mathbf{P} , and \mathbf{prAM} to \mathbf{NP} . However, even the low-end hardness conditions remain open to date. This raises the question whether there are means of derandomizing that do not need any of the above hardness conditions, or whether derandomization is *equivalent* to hardness.

Proving an equivalence would establish a “canonical form” of derandomization, namely through PRGs. It would show that if we can individually derandomize each procedure of the type considered, then we can derandomize them “all at once” – we do not need to know the particulars of a procedure in order to derandomize it. On the other hand, proving a non-equivalence would establish that the hardness-based PRG approach to derandomization is incomplete, i.e., that there are better avenues to derandomization. In fact, since the existence of PRGs is known to imply the hardness conditions that yield them, this would render incomplete *any* PRG-based approach – using hardness or not. Therefore another way of phrasing the question is whether or not the ability to derandomize implies the ability to do so through PRGs.

In recent years we have seen a number of result for $(\mathbf{pr})\mathbf{BPP}$ showing that derandomization of some sort implies hardness of some sort, although typically not strong enough so as to imply back the derandomization (e.g., [IKW02, KI04, KvMS12]). One exceptional example is an equivalence for a mild notion of derandomization for \mathbf{prBPP} , implicit in the work of Impagliazzo et al. [IKW02]. They show that $\mathbf{NE} \not\subseteq \mathbf{P}/\text{poly}$ if and only if every \mathbf{prBPP} -problem can be decided in $\mathbf{NTIME}(2^{n^\epsilon})$ with an advice of length n^ϵ for infinitely-many input lengths, for arbitrarily small $\epsilon > 0$. Note that this notion of derandomization is indeed mild compared to what is conjectured to hold (namely a derandomization of \mathbf{prBPP} in \mathbf{P}): it uses nondeterminism, a subexponential amount of time, a subpolynomial amount of nonuniformity, and it is only required to work infinitely often. Mild though it may be, it is a nontrivial derandomization nonetheless, and one that is not known to

hold. If it can be done at all, the [IKW02] result shows that it can be done in a canonical way – through PRGs.

In contrast to the class prBPP , not much is known regarding derandomization-to-hardness connections for Arthur-Merlin games. The only result in this direction is a “hybrid” connection that shows an implication from derandomizing prAM to *deterministic* circuit lower bounds [AGHK11], whereas the hardness-to-derandomization implication for prAM involves *nondeterministic* circuits. In particular, what kind of nondeterministic circuit lower bounds, if any, are implied by derandomizing prAM is an open question.

Our results. In this paper we take up this last question and obtain an analogue of the equivalence of hardness and derandomization from [IKW02] for the class prAM instead of prBPP .

Theorem 1 (Equivalence for Arthur-Merlin games). *The following are equivalent:*

- Every prAM -problem can be decided in $\Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$ at infinitely many input lengths, for every $\epsilon > 0$.
- $\Sigma_2\text{E} \not\subseteq \text{NP}/\text{poly}$.

Recall that prAM can be simulated in $\Pi_2\text{P}$. Although plausible hardness assumptions imply simulations in NP , it is open whether a simulation in $\Sigma_2\text{TIME}(t)$ is possible for subexponential t – even with subpolynomial advice, and even if the simulation is only required to succeed infinitely often. In this sense, Theorem 1 is a full analogue of the equivalence result by Impagliazzo et al. [IKW02] for prBPP . In both their equivalence and ours the derandomizations are mild in the same way: compared to the standard notion they use extra time, extra advice, and an extra level of nondeterminism (where each “extra” is quantified identically in both results), and they are only required to work infinitely often.

As it is currently open whether $\Sigma_2\text{E} \not\subseteq \text{NP}/\text{poly}$, Theorem 1 implies that mildly derandomizing Arthur-Merlin games would yield new circuit lower bounds. In fact, it is even open whether $\text{E}^{\Sigma_2\text{P}} \not\subseteq \text{NP}/\text{poly}$. The situation for nondeterministic circuits mimics the one for deterministic circuits “one level down” in the exponential-time hierarchy, where it is open whether $\text{E}^{\text{NP}} \not\subseteq \text{P}/\text{poly}$.

More interestingly, Theorem 1 implies that mildly derandomizing prAM *in any way* implies that the same derandomization can be done *canonically*.

Corollary 1 (Derandomization-to-PRG connection for Arthur-Merlin games). *If prAM can be derandomized as in Theorem 1, then there exist pseudorandom generators that yield the same derandomization.*

A key step in the proof of Theorem 1 shows that mild derandomizations of prAM imply improved Karp-Lipton style collapse results for nondeterministic circuits. The following are two instantiations that may be of independent interest.

Theorem 2 (High-end collapse results). *Suppose that every prAM -problem can be decided in $\Sigma_2\text{P}$. Then*

- (i) $\text{PSPACE} \subseteq \text{NP}/\text{poly} \implies \text{PSPACE} \subseteq \Sigma_2\text{P}$, and
- (ii) $\text{coNP} \subseteq \text{NP}/\text{poly} \implies \text{PH} \subseteq \text{P}^{\Sigma_2\text{P}}$.

Karp and Lipton [KL82] showed that if $\text{NP} \subseteq \text{P/poly}$ then $\text{PH} \subseteq \Sigma_2\text{P}$. Yap’s adaptation [Yap83] of the Karp-Lipton result gives that if $\text{coNP} \subseteq \text{NP/poly}$ then $\text{PH} \subseteq \Sigma_3\text{P}$. Modulo the derandomization assumption, the second item of Theorem 2 improves Yap’s result by “half a level”. Another variant of the Karp-Lipton argument (attributed to Meyer) states that if $\text{PSPACE} \subseteq \text{P/poly}$ then $\text{PSPACE} \subseteq \Sigma_2\text{P}$. Relativizing the latter result yields the strongest collapse consequence of $\text{PSPACE} \subseteq \text{NP/poly}$ known unconditionally, namely $\text{PSPACE} \subseteq \Sigma_3\text{P}$. The first item of Theorem 2 improves this collapse by one level, modulo the derandomization assumption. We refer to Section 2.4 for related work on more refined collapses.

We use a low-end version of the first collapse result in Theorem 2 to establish the forward direction of Theorem 1. The proof relies on interactive proofs for PSPACE , and as such does not relativize. If we use the second collapse result instead of the first one, we obtain a relativizing proof of a weaker result, namely that the same derandomization assumption implies $\text{E}^{\Sigma_2\text{P}} \not\subseteq \text{NP/poly}$.

A lower bound framework. Once we have established our collapse results, we develop the derandomization-to-hardness part of Theorem 1 as an instantiation of a generic framework that captures our results as well as the corresponding one from [IKW02]. Also captured is the result by Kabanets and Impagliazzo [KI04] that either $\text{NE} \not\subseteq \text{P/poly}$ or that PERM (the permanent over the integers) requires superpolynomial-size arithmetic circuits, under the same mild derandomization assumption as in [IKW02] but for the class BPP instead of prBPP .

Organization. In Section 2 we sketch the ideas behind our results, set up the lower bound framework, and discuss the relationship with earlier work. In Section 3 we introduce our notation for the formal development. In Section 4 we present the collapse results, and in Section 5 the main result.

2 Outline of the Arguments and Related Work

We now outline the proofs of Theorem 1 and Theorem 2. We focus on the most novel part of the proof of Theorem 1, which is the direction from derandomization to hardness, as well as the weaker but relativizing variant mentioned in the introduction. Both can be cast as instantiations of a general framework that also captures several known results of this ilk. The framework is based on Kannan’s theorem [Kan82] that some level of the polynomial-time hierarchy cannot be decided by circuits of fixed-polynomial size, and critically relies on collapse results in order to bring down the required level of the polynomial-time hierarchy.

We first explain the collapse results we need and then present the framework. For ease of exposition, in this overview we use the assumption that prAM can be simulated in $\Sigma_2\text{P}$, yielding the high-end collapse results of Theorem 2, although for the proof of Theorem 1 it suffices to have the weaker assumption that prAM can be simulated in $\bigcap_{\epsilon > 0} \text{i.o.} - \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$.

2.1 First high-end collapse result

Our first collapse result is an adaptation to the nondeterministic setting of the classical result that if PSPACE has polynomial-size deterministic circuits then $\text{PSPACE} \subseteq \text{MA}$ [LFKN92, Sha92]. Let us first recall the proof of that classical result.

Assuming that PSPACE has polynomial-size deterministic circuits, we want to compute some PSPACE-complete language L in MA. The proof hinges on the existence of an interactive proof system for L in which the honest prover's responses are computable in PSPACE. By assumption, there is a polynomial-size deterministic circuit D_{prover} that encodes the honest prover's strategy; i.e., given the transcript of a message history, the circuit computes the next bit the honest prover sends. Now the MA-protocol for L is as follows: Merlin sends a polynomial-size circuit D' to Arthur, who then carries out the interactive protocol for L by himself, evaluating the circuit D' to determine the prover's responses. If the input is in L , Merlin can send the circuit D_{prover} , which makes Arthur accept with high probability. If the input is not in L , the soundness property of the interactive proof system guarantees that no deterministic circuit can make Arthur accept with significant probability. This proves that $L \in \text{MA}$.

Now let us turn to our setting and try to achieve a similar collapse under the assumption that PSPACE has *nondeterministic* circuits of polynomial size. Since Arthur may need Merlin's help in evaluating nondeterministic circuits, we allow for one more round of interaction between Arthur and Merlin, and aim for a collapse to MAM=AM rather than MA. By assumption, there exists a polynomial-size nondeterministic circuit C_{prover} implementing the honest prover's strategy; i.e., given the transcript of a message history and a bit b , C_{prover} accepts if the honest prover's next message bit is b , and rejects otherwise. Now consider the following attempt at a protocol for the PSPACE-complete language L : Merlin sends a polynomial-size nondeterministic circuit C' , purported to encode the strategy of the honest prover. Upon receiving the circuit C' , Arthur reveals his coin flips ρ to Merlin. Merlin then provides the certificates for the circuit C' that allow Arthur to construct and verify every bit of the transcript of the interactive proof corresponding to the coin flips ρ . Finally, Arthur accepts if the resulting transcript is accepting.

This protocol is complete but not necessarily sound. Indeed, Merlin can send a nondeterministic circuit C' that has accepting and rejecting computation paths on every input, which allows Merlin to adapt his strategy to the coin flips ρ in whatever way he wants, by revealing accepting computation paths only if he wishes to. We somehow need to force Merlin to commit to a fixed strategy in advance.

In order to do so, we use our derandomization assumption and aim for a collapse to $\Sigma_2\text{P}$ instead of AM. First, in an existential phase we "guess" a nondeterministic circuit C' supposed to implement the honest prover's strategy for L . We provide C' as additional input to the AM-protocol, and require Merlin to convince Arthur using the specific circuit C' provided. Moreover, in *parallel* to the AM-protocol, we make sure C' commits Merlin to a fixed strategy. More precisely, in a universal phase we check that on every partial transcript for at least one of the choices of the bit b , C' rejects on all computation paths. This fixes the soundness problem while maintaining completeness.

Note that the above procedure can be implemented within $\Sigma_2\text{P}$: We use two alternations outside the AM-protocol, where the second alternation for checking the circuit is executed in parallel to the protocol. Since by our derandomization assumption the AM-protocol can be simulated in $\Sigma_2\text{P}$, a collapse of PSPACE to $\Sigma_2\text{P}$ follows.

2.2 Second high-end collapse result

To outline the proof of our second collapse result, let us first recall the proof of the classical result by Karp and Lipton for the case of NP and deterministic circuits. Assuming that satisfiability (SAT) has polynomial-size circuits, we consider any $\Pi_2\text{P}$ -predicate of the form $(\forall u)(\exists v)\varphi(u, v)$, where φ

is a Boolean formula, and translate it into an equivalent $\Sigma_2\text{P}$ -predicate.

One way to construct the $\Sigma_2\text{P}$ -predicate goes as follows. Use an existential quantifier to “guess” a deterministic circuit D , verify that D correctly decides SAT, and then use D to transform the inner existential phase of the original $\Pi_2\text{P}$ -predicate into a deterministic one, effectively eliminating one quantifier alternation. Hence, we obtain an equivalent predicate that reads as

$$(\exists D) [\text{correct}(D) \wedge (\forall u)D(“(\exists v)\varphi(u, v)” = 1)]. \quad (1)$$

Exploiting the self-reducibility of SAT, $\text{correct}(D)$ can be expressed as a coNP -predicate. This way (1) becomes a $\Sigma_2\text{P}$ -predicate.

Let us now turn to our setting and try to achieve the same collapse under the assumption that $\overline{\text{SAT}}$ has *nondeterministic* circuits of polynomial size. Mimicking the above proof, we use an existential quantifier to guess a nondeterministic circuit C , check its correctness for $\overline{\text{SAT}}$, and then feed the existential phase of the $\Pi_2\text{P}$ -predicate into C . The latter transforms the existential phase into an equivalent universal phase, which can be merged with the initial universal phase of the original $\Pi_2\text{P}$ -predicate. This complementation gives us a new equivalent predicate of the form

$$(\exists C) [\text{correct}(C) \wedge (\forall u)C(“(\exists v)\varphi(u, v)” = 0)]. \quad (2)$$

In fact, it suffices that the predicate $\text{correct}(C)$ checks the *completeness* of C for $\overline{\text{SAT}}$ (that C accepts every unsatisfiable formula) without explicitly checking the *soundness* of C for $\overline{\text{SAT}}$ (that C only accepts unsatisfiable formulas). However, while soundness can be tested in coNP , completeness seems to require $\Pi_2\text{P}$. This turns (2) into a $\Sigma_3\text{P}$ -predicate rather than a $\Sigma_2\text{P}$ -predicate. Note that obtaining an equivalent $\Sigma_3\text{P}$ -predicate is trivial since we started from a $\Pi_2\text{P}$ -predicate. If we start from a Π_3^{P} -predicate instead, an analogous transformation yields an equivalent $\Sigma_3\text{P}$ -predicate, implying a collapse of the polynomial-time hierarchy to the third level (this is Yap’s theorem [Yap83]). If we could check for completeness in $\Sigma_2\text{P}$, then the collapse would deepen to the second level and we would be done, but we do not know how to do this, even under our derandomization assumption. What we *can* do under our assumption, is to construct a correct nondeterministic circuit for $\overline{\text{SAT}}$ “half a level” up from $\Sigma_2\text{P}$, namely in $\text{P}^{\Sigma_2\text{P}}$ (just like we could if we knew how to check completeness in $\Sigma_2\text{P}$). This collapses the polynomial-time hierarchy down to $\text{P}^{\Sigma_2\text{P}}$, giving our second collapse result.

The particular problem in prAM that we assume can be derandomized is the following approximate lower bound problem: Given a circuit C and an integer a , decide whether C accepts at least a inputs or noticeably less than a , say, less than $a(1 - \epsilon)$ where $\epsilon = 1/\text{poly}(n)$. Goldwasser and Sipser [GS86] showed that this problem lies in prAM , even when C is nondeterministic.

The key difference our derandomization assumption makes is the added ability to guarantee, within $\Sigma_2\text{P}$, that a nondeterministic circuit C is “almost” complete, i.e., that C accepts at least a $(1 - \epsilon)$ -fraction of all unsatisfiable formulas, and even more generally, that a nondeterministic circuit C accepts at least a $(1 - \epsilon)$ -fraction of all unsatisfiable formulas that are rejected by some other given nondeterministic circuit. This enables us to construct in $\text{P}^{\Sigma_2\text{P}}$, out of a sound nondeterministic circuit C_i for $\overline{\text{SAT}}$, another sound nondeterministic circuit C_{i+1} for $\overline{\text{SAT}}$ that misses at most half as many unsatisfiable formulas as C_i does. Starting from the trivial sound circuit C_0 that rejects everything, this process yields a sound and complete nondeterministic circuit for $\overline{\text{SAT}}$ within n iterations.

To explain the role of the derandomization hypothesis in more detail, we first sketch how to find C_1 because that case is easier. In constructing C_1 we make oracle queries of the form: Is there

a sound nondeterministic circuit of size $s(n)$ for $\overline{\text{SAT}}$ that accepts a inputs. These queries can be decided approximately by a $\Sigma_2\text{P}$ -oracle because of our derandomization assumption and because soundness can be checked in coNP . Assuming $\overline{\text{SAT}}$ has nondeterministic circuits of size $s(n)$, this enables us to approximate the number \bar{a} of unsatisfiable formulas of length n through a binary search using a $\Sigma_2\text{P}$ -oracle. Moreover, by self-reduction we get a sound circuit C_1 that accepts at least $(1 - \epsilon)\bar{a}$ inputs, with the factor $(1 - \epsilon)$ being due to the gap between the yes and no instances of the approximate lower bound problem. Setting $\epsilon = 1/2$, we thus find a sound circuit C_1 for $\overline{\text{SAT}}$ that accepts at least half of all unsatisfiable formulas of length n .

To construct C_2 we want to employ a similar strategy as in the construction of C_1 , namely to find a sound circuit \tilde{C}_1 for $\overline{\text{SAT}}$ that seems to accept as many inputs as possible, and then set $C_2 = C_1 \vee \tilde{C}_1$. The difference, however, is that this time we want to maximize not over all inputs, but just over those inputs that C_1 rejects. This causes a problem because the set of inputs that C_1 rejects is in coNP , whereas the approximate lower bound problem allows us to estimate the size of NP sets only.

We overcome this obstacle by using the complementation idea again. By assumption, $\overline{\text{SAT}}$ has small nondeterministic circuits not only at input length n , but also at larger input lengths. In particular, there is a nondeterministic circuit C' of size $s(n')$ for $\overline{\text{SAT}}$ at input length n' , where n' is large enough that we can express the computation of an n -input size- $s(n)$ nondeterministic circuit – in particular C_1 – with a Boolean formula of length n' . If we can get a hold of such a circuit C' , then we can express the coNP -set $\{x \in \{0, 1\}^n : C_1 \text{ rejects } x\}$ alternately as the NP-set $\{x \in \{0, 1\}^n : C' \text{ accepts } \phi_{C_1}(x, \cdot)\}$, where $\phi_{C_1}(x, y)$ is a Boolean formula of size n' that expresses that y is a valid accepting computation of C_1 on input x . Since the latter set is in NP, it can be provided as input to the approximate lower bound problem. Of course, getting a hold of a circuit C' for $\overline{\text{SAT}}$ at length n' is the very problem we are trying to solve – only harder since $n' > n$. We observe, however, that we do not need to explicitly check the completeness of C' for $\overline{\text{SAT}}$; it suffices to check the soundness of C' for $\overline{\text{SAT}}$. Since the latter can be done in coNP , we can guess and check the circuit C' in $\Sigma_2\text{P}$.

To recapitulate, we want to find a sound nondeterministic circuit \tilde{C}_1 for $\overline{\text{SAT}}$ that misses at most half of the unsatisfiable formulas that C_1 misses. We accomplish this by first encoding the computation of C_1 on a generic input x as a Boolean formula $\phi_{C_1}(x, y)$ of length n' , such that $\phi_{C_1}(x, \cdot)$ is satisfiable for a particular x iff C_1 accepts x . Then we make oracle queries that ask: Is there a nondeterministic circuit C' of size $s(n')$ on n' inputs, and a nondeterministic circuit \tilde{C}_1 of size $s(n)$ on n inputs, such that (i) the set $\{x \in \{0, 1\}^n : \tilde{C}_1 \text{ accepts } x \text{ and } C' \text{ accepts } \phi_{C_1}(x, \cdot)\}$ is of size at least a , and (ii) C' and \tilde{C}_1 are both sound for $\overline{\text{SAT}}$. By our derandomization assumption that $\text{prAM} \subseteq \Sigma_2\text{P}$, these queries can be made to a $\Sigma_2\text{P}$ -oracle, which allows us to construct \tilde{C}_1 in $\text{P}^{\Sigma_2\text{P}}$. We then set $C_2 = C_1 \vee \tilde{C}_1$.

As a side remark we point out that, although we do not explicitly require the circuit C' to be complete for $\overline{\text{SAT}}$, the maximization of a forces C' to be complete (on the relevant instances). This is how we can avoid checking completeness explicitly (which seems to require the power of $\Pi_2\text{P}$) although we rely on it.

Having found C_2 , we then iterate to get a third nondeterministic circuit C_3 that misses at most half as many unsatisfiable formulas as C_2 does, and so on until we reach perfect completeness. This way we construct in $\text{P}^{\Sigma_2\text{P}}$ a nondeterministic circuit of size $O(n \cdot s(n))$ for $\overline{\text{SAT}}$ at length n . The collapse of the polynomial-time hierarchy to $\text{P}^{\Sigma_2\text{P}}$ follows.

2.3 The lower bound framework

In both our main equivalence for prAM and the one for prBPP due to Impagliazzo et al. [IKW02], the proof of the forward direction – from derandomization to hardness – can be cast as an instantiation of a generic framework. We now describe that framework.

Our goal is to show, under some derandomization assumption, that some class \mathcal{C} does not have polynomial size circuits of type τ , where \mathcal{C} is a linear-exponential class such as E or NE , and τ could be deterministic, nondeterministic, or even arithmetic¹.

The proof goes by contradiction and consists of the following ingredients. Assume that \mathcal{C} has τ -circuits of polynomial size.

1. Collapsing the polynomial-time hierarchy. Use the hypothesis that \mathcal{C} has τ -circuits of polynomial size to show that the polynomial-time hierarchy can be simulated in some randomized class \mathcal{R} .
2. Derandomization. Use the derandomization assumption to show that every language in \mathcal{R} is decidable in \mathcal{C} with a fixed-polynomial amount of advice for infinitely many input lengths.

Note that classes \mathcal{C} such as E and NE contain complete languages under fixed-polynomial time reductions (in fact linear-time reductions), and therefore \mathcal{C} having polynomial size circuits is equivalent to \mathcal{C} having size- n^c circuits for a *fixed* constant c . By combining the above two ingredients, we can thus conclude that all of the polynomial-time hierarchy can be decided by τ -circuits of size less than n^d for some fixed constant d and infinitely many input lengths. This contradicts Kannan’s result that for any fixed polynomial n^d , there exists a language in some level of the polynomial-time hierarchy that requires τ -circuits of size at least n^d for all but finitely many input lengths [Kan82].²

Figure 2.3 lists the instantiations of the above framework that yield derandomization-to-hardness results implied by the works of Impagliazzo et al. [IKW02] and of Kabanets and Impagliazzo [KI04] mentioned in the introduction. The last two lines represent our main result (Theorem 1) obtained via our first collapse involving PSPACE , and our weaker but relativizing result obtained via our second collapse involving coNP .

	\mathcal{C}	τ	\mathcal{R}	derandomization hypothesis
[IKW02]	NE	d	MA	$\text{prBPP} \subseteq \cap_{\epsilon>0} \text{i.o.} - \text{NTIME}(2^{n^\epsilon})/n^\epsilon$
[KI04]	$\text{NE} \cup \{\text{PERM}\}$	a	N-BPP	$\text{BPP} \subseteq \cap_{\epsilon>0} \text{i.o.} - \text{NTIME}(2^{n^\epsilon})/n^\epsilon$
main result	$\Sigma_2\text{E}$	n	$\text{M}(\text{AM} \text{coNP})$	$\text{prAM} \subseteq \cap_{\epsilon>0} \text{i.o.} - \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$
relativizing result	$\text{E}^{\Sigma_2\text{P}}$	n	$\text{PprM}(\text{AM} \text{coNP})$	

Figure 1: Framework to show that \mathcal{C} does not have polynomial-size circuits of type τ , assuming a derandomization hypothesis and using an intermediate randomized class \mathcal{R} .

The values “d,” “a,” and “n” for τ in Figure 2.3 stand for deterministic, arithmetic, and nondeterministic circuits, respectively. In the line for [KI04], the class \mathcal{R} is expressed using the

¹We say that an arithmetic circuit C on n inputs decides a language L at length n if C agrees with the characteristic function of L on $\{0, 1\}^n$.

²Kannan’s result is originally stated for deterministic circuits, but it is straightforward to generalize it to nondeterministic or arithmetic circuits.

operator “N”; this is the natural operator that yields the class NP when applied to the class P. The same operator can also be used to express the class MA as N-prBPP. The classes \mathcal{R} we use for our results in Figure 2.3 are technical augmentations of Arthur-Merlin protocols; they are implicit in the proof of Theorem 2 and are explicitly defined in Section 3.

As a side remark we point out that the first two lines of Figure 2.3 represent proofs that are somewhat simpler than the original ones. The original proofs use the relatively involved collapse $\text{NEXP} \subseteq \text{P/poly} \implies \text{NEXP} \subseteq \text{MA}$ [IKW02], which itself relies on the nondeterministic time hierarchy theorem, a hardness-based PRG construction, as well as multi-prover interactive proofs. The framework shows that the same result can be obtained by collapsing PH rather than NEXP, hence it suffices to use a more basic collapse technology using the assumption $\text{PSPACE} \subseteq \text{P/poly}$ or even the weaker assumption $\text{P}^{\#P} \subseteq \text{P/poly}$ (for [KI04] this was observed earlier in [AvM11]).

2.4 Related work

In the setting of decision problems, Goldreich [Gol11a, Gol11b] showed that the standard notion of derandomization for prBPP (deterministic simulations without advice that work for all but finitely many inputs) is equivalent to the existence of so-called “targeted PRGs,” which are PRGs that have access to the input or even to the circuit that models the randomized computation on the given input. For standard PRGs (which are oblivious to the input and only depend on the running time of the randomized computation) he showed an equivalence between their existence and derandomizations of prBPP in an average-case setting: deterministic simulations without advice that may err on some inputs but such that generating erroneous inputs is computationally difficult. In both cases the equivalence with circuit lower bounds remains open.

In the setting of Arthur-Merlin games, Gutfreund et al. [GST03] suggest an approach to prove that AM can be simulated in $\Sigma_2\text{SUBEXP}$. Theorem 1 implies that if the approach works for prAM then it would yield new circuit lower bounds.

Using Kannan’s argument to get circuit lower bounds from a derandomization assumption for prAM was carried out by Gutfreund and Kawachi [AGHK11]. The same paper also presents an alternate and simpler proof that does not use Kannan’s argument, but uses the power of prAM to directly diagonalize against deterministic circuits.

The technique of using a prAM-oracle to iteratively construct a sound circuit with rapidly increasing completeness, appears in the work of Chakaravathy and Roy [CR08]. Using this technique they show that PH collapses to P^{prAM} , under the classical Karp-Lipton assumption that $\text{NP} \subseteq \text{P/poly}$.

Be it for diagonalization as in [AGHK11], or for finding a circuit as in [CR08, AGHK11] and our work, the use of a prAM-oracle can be viewed as finding a witness \tilde{y} that approximately maximizes a “quality measure” f defined on the set of all strings. For diagonalization purposes this measure would be the number of circuits that a given string y eliminates when viewed as the characteristic string of a function. For finding a circuit for $\overline{\text{SAT}}$ at length n , y is viewed as a circuit and $f(y)$ measures the number of unsatisfiable formulas that y accepts provided that y is sound.

In [Gol11a, Gol11b] a prBPP-oracle is used to construct targeted PRGs. The constructions also involve approximately maximizing a quality measure f ; in this case $f(y)$ may be defined recursively as the average quality of the extensions of y , i.e., $f(y) = \frac{1}{2}(f(y0) + f(y1))$. The difference between the works mentioned in the previous paragraph and Goldreich’s is that in the latter f can be additively approximated using a prBPP-oracle, whereas in the former f is multiplicatively approximated using a prAM-oracle.

Regarding our high-end collapse result involving coNP and NP/poly , at a more refined level of granularity the strongest unconditional collapse consequence of the condition $\text{coNP} \subseteq \text{NP/poly}$ is that PH collapses to $\text{S}_2\text{P}^{\text{NP}}$ [CCHO05], a class that contains $\text{P}^{\Sigma_2\text{P}}$ but is not known to equal it. Similarly, the strongest unconditional collapse consequence of $\text{PSPACE} \subseteq \text{NP/poly}$ is that $\text{PSPACE} \subseteq \text{S}_2\text{P}^{\text{NP}}$.

3 Notation and Conventions

In this section we introduce our notation and conventions, including the notion of an augmented Arthur-Merlin protocol, which is a technical construct that naturally arises in our collapse arguments.

Promise problems and languages. A *promise problem* Π is a pair of disjoint sets (Π_Y, Π_N) of strings over the binary alphabet $\{0, 1\}$. A language L is a promise problem of the form (L, \bar{L}) , where $\bar{L} \doteq \{x \in \{0, 1\}^* : x \notin L\}$. A promise problem $\Pi' = (\Pi'_Y, \Pi'_N)$ is said to *agree with* a promise problem $\Pi = (\Pi_Y, \Pi_N)$ if $\Pi_Y \subseteq \Pi'_Y$ and $\Pi_N \subseteq \Pi'_N$. To *decide* a promise problem Π is to correctly determine, for all inputs $x \in \Pi_Y \cup \Pi_N$, which of Π_Y and Π_N contains x . For two classes of promise problems \mathcal{C} and \mathcal{C}' , we write $\mathcal{C} \subseteq \mathcal{C}'$ if for every $\Pi \in \mathcal{C}$ there exists $\Pi' \in \mathcal{C}'$ such that Π' agrees with Π . We say that Π reduces to Π' if there exists an oracle Turing machine M such that $M^{L'}$ agrees with Π for *every* language L' that agrees with Π' . In particular, a language L is in $\text{P}^{\Pi'}$ if there exists a polynomial-time oracle Turing machine M such that $M^{L'}$ decides L for *every* language L' that agrees with Π' . For more on promise problems, see the survey [Gol06].

prAM represents the class of promise problems Π for which there exists a constant c and a language $L \in \text{P}$ such that for every input x

$$\begin{aligned} [\text{completeness}] \quad x \in \Pi_Y &\Rightarrow \Pr_y[(\exists z)\langle x, y, z \rangle \in L] \geq 2/3, \text{ and} \\ [\text{soundness}] \quad x \in \Pi_N &\Rightarrow \Pr_y[(\exists z)\langle x, y, z \rangle \in L] \leq 1/3, \end{aligned}$$

where n denotes the length of x , the variables y and z range over $\{0, 1\}^{n^c}$, and the probabilities are with respect to the uniform distribution. AM denotes those problems in prAM that are languages. Underlying each problem in prAM there is a protocol between a randomized polynomial-time verifier (Arthur) and an all-powerful prover (Merlin); we refer to these protocols as Arthur-Merlin protocols or Arthur-Merlin games.

In our proofs the following technical augmentation of Arthur-Merlin protocols arises naturally. For lack of a better name, we refer to them as “augmented” Arthur-Merlin protocols.

Definition 1 (Augmented Arthur-Merlin protocol). *The class $\text{prM}(\text{AM}|\text{coNP})$ consists of all promise problems Π for which there exists a constant c , a promise problem $\Gamma \in \text{prAM}$, and a language $V \in \text{coNP}$ such that*

$$\begin{aligned} [\text{completeness}] \quad x \in \Pi_Y &\Rightarrow (\exists y)(\langle x, y \rangle \in \Gamma_Y \wedge \langle x, y \rangle \in V), \text{ and} \\ [\text{soundness}] \quad x \in \Pi_N &\Rightarrow (\forall y)(\langle x, y \rangle \in \Gamma_N \vee \langle x, y \rangle \notin V), \end{aligned}$$

where n denotes the length of x , and y ranges over $\{0, 1\}^{n^c}$. $\text{M}(\text{AM}|\text{coNP})$ denotes those problems in $\text{prM}(\text{AM}|\text{coNP})$ that are languages.

Similar to the class prAM , underlying each problem in $\text{prM}(\text{AM}||\text{coNP})$ there is a protocol between an all-powerful prover, Merlin, and – in this case – *two* verifiers, who cannot communicate with each other. One verifier is the usual randomized polynomial-time Arthur from the prAM -problem Γ ; the other one is the coNP -verifier V . Merlin goes first and sends a common message to both verifiers. At this point, V has to make a decision to accept/reject, whereas Arthur can interact with Merlin as in the Arthur-Merlin protocol for Γ before making a decision. The input is accepted by the protocol iff both verifiers accept.

Nondeterministic circuits. A nondeterministic Boolean circuit C consists of AND and OR gates of fan-in 2, NOT gates of fan-in 1, input gates of fan-in 0, and additionally, choice gates of fan-in 0. We say that the circuit accepts input x , or $C(x) = 1$ in short, if there is some assignment of Boolean values to the choice gates that makes the circuit evaluate to 1; otherwise we say that C rejects x , or $C(x) = 0$ in short. We measure the *size* of a circuit by the number of its connections. A circuit of size s can be described by a binary string of length $O(s \log s)$.

4 Collapse Results

In this section we establish our collapse result (Theorem 2), which uses the assumption that prAM can be mildly derandomized. In fact, we prove an unconditional collapse result involving the class of augmented Arthur-Merlin protocols introduced in Definition 1, from which Theorem 2 follows under the derandomization assumption. We first establish a collapse result assuming PSPACE has nondeterministic circuits of polynomial size (corresponding to the first part in Theorem 2) and then do the same for coNP instead of PSPACE (corresponding to the second part in Theorem 2).

4.1 Collapse result for PSPACE

The proof of the following theorem uses interactive proofs for PSPACE and as such does not relativize.

Theorem 3. *If $\text{PSPACE} \subseteq \text{NP}/\text{poly}$ then $\text{PSPACE} \subseteq \text{M}(\text{AM}||\text{coNP})$.*

Proof. Let L be in PSPACE , fix an interactive proof system for L , and consider the language L_{prover} consisting of all tuples $\langle x, y, b \rangle$ such that y is a prefix of the transcript of an interaction of the verifier with the honest prover on input x , and the next bit in the transcript is sent by the prover and equals b . Without loss of generality we can assume that L_{prover} is paddable such that given a random string ρ for the verifier, we can construct the entire transcript with the honest prover on an input $x \in \{0, 1\}^n$ by making queries to L_{prover} of a *single* length $\ell(n) = \text{poly}(n)$.

By the assumption that $\text{PSPACE} \subseteq \text{NP}/\text{poly}$, L_{prover} can be decided by some polynomial-size nondeterministic circuit C_{prover} . Now consider the following augmented Arthur-Merlin protocol for deciding L on $x \in \{0, 1\}^n$. Merlin sends to both verifiers a polynomial-size nondeterministic circuit C' , purported to compute L_{prover} at length $\ell(n)$. The coNP -verifier V checks that C' is “single-valued”, i.e., for all possible queries to the circuit C' , if for some query $\langle x_0, y_0, b_0 \rangle$ the circuit C' accepts then C' rejects the complementary query $\langle x_0, y_0, \neg b_0 \rangle$. Note that this check is indeed in coNP .

Arthur picks a random string ρ of the appropriate length (at most $\ell(n)$) and sends it to Merlin. Merlin sends the transcript for the interactive protocol for L on input x corresponding to the coin

flips ρ . Merlin also sends the certificates for C' that purportedly produce that transcript. Arthur accepts iff C' produces the transcript when given those certificates, and the transcript is accepting.

To argue completeness, consider $x \in L$. Then Merlin can just send $C' = C_{\text{prover}}$. That circuit passes the coNP -verifier V and also passes Arthur's verification with high probability.

For the soundness, consider $x \notin L$, and suppose that Merlin sends a circuit C' that passes the coNP -verifier. This means that C' is single-valued, and corresponds to a fixed prover strategy. Then Arthur rejects with high probability by the soundness of the original interactive proof system for L . \square

The proof of the first part of Theorem 2 follows immediately from Theorem 3:

Proof (of part (i) of Theorem 2). If prAM can be simulated in $\Sigma_2\text{P}$ then so can $\text{prM}(\text{AM}|\text{coNP})$. This follows because replacing Γ_Y in Definition 1 by a $\Sigma_2\text{P}$ -predicate and Γ_N by its complement, turns Π_Y into a $\Sigma_2\text{P}$ -predicate and Π_N into its complement. If in addition $\text{PSPACE} \subseteq \text{NP}/\text{poly}$, Theorem 3 implies that $\text{PSPACE} \subseteq \text{M}(\text{AM}|\text{coNP}) \subseteq \Sigma_2\text{P}$. \square

4.2 Collapse result for coNP

We proceed with a relativizable proof of the following unconditional collapse result assuming coNP has nondeterministic circuits of polynomial size.

Theorem 4. *If $\text{coNP} \subseteq \text{NP}/\text{poly}$ then $\Sigma_3\text{P} \subseteq \text{P}^{\text{prM}(\text{AM}|\text{coNP})}$.*

The second part of Theorem 2 follows immediately from Theorem 4 under the mild derandomization assumption for prAM , in a relativizable way.

Proof (of part (ii) of Theorem 2). As we argued in the proof of part (i) in Section 4.1, if prAM can be simulated in $\Sigma_2\text{P}$ then so can $\text{prM}(\text{AM}|\text{coNP})$. If in addition $\text{coNP} \subseteq \text{NP}/\text{poly}$, Yap's theorem [Yap83] and Theorem 4 imply that $\text{PH} \subseteq \Sigma_3\text{P} \subseteq \text{P}^{\text{prM}(\text{AM}|\text{coNP})} \subseteq \text{P}^{\Sigma_2\text{P}}$. \square

We now argue Theorem 4. Assume that $\overline{\text{SAT}}$ has nondeterministic circuits of size $s(n)$, where s is some polynomial. Following the outline of Section 2.2, with the aid of a $\text{prM}(\text{AM}|\text{coNP})$ -oracle, we construct a circuit of size $O(n \cdot s(n))$ that correctly decides $\overline{\text{SAT}}$ on all instances of size n . The circuit is obtained as the end of a sequence of sound circuits with rapidly improving completeness, starting from the trivial circuit that rejects everything.

To measure the improvement in each step, we consider the following function $f : \{0,1\}^* \times \{0,1\}^* \rightarrow \mathbb{N}$, which takes as arguments the current circuit C in the sequence, and a candidate circuit \tilde{C} to improve the completeness of C in the next step, while maintaining soundness.

$$f(C, \tilde{C}) = \begin{cases} |C^{-1}(0) \cap \tilde{C}^{-1}(1)| & \text{if } \tilde{C} \text{ is sound for } \overline{\text{SAT}} \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We map circuits \tilde{C} that are not sound to zero because their use would violate the soundness of the sequence. If \tilde{C} is sound, f counts the number of instances of $\overline{\text{SAT}}$ that are missed by C but caught by \tilde{C} .

For a given circuit C that is sound but not complete, our goal is to find a circuit \tilde{C} that approximately maximizes $f(C, \tilde{C})$. For this task we only need access to an approximation of f to within a constant multiplicative factor. In general, for a function $f : \{0,1\}^* \times \{0,1\}^* \rightarrow \mathbb{N}$,

approximating f to within a multiplicative factor is captured by the following promise problem $A_f = (Y_f, N_f)$, which additionally takes an integer a in binary and a positive integer $1/\epsilon$ in unary.

$$\begin{aligned} Y_f &= \{\langle x, y, a, \epsilon \rangle : f(x, y) \geq a\} \\ N_f &= \{\langle x, y, a, \epsilon \rangle : f(x, y) < (1 - \epsilon)a\}. \end{aligned} \quad (4)$$

The crux of our argument is the following lemma.

Lemma 1. *Let f be the function defined by (3), and A_f the promise problem given by (4). If $\text{coNP} \subseteq \text{NP/poly}$ then $A_f \in \text{prM}(\text{AM} \parallel \text{coNP})$.*

Proof. We follow the outline from Section 2.2, but cast the resulting algorithm for A_f in terms of an augmented Arthur-Merlin protocol with coNP -verifier V on input (C, \tilde{C}) .

Since V can check whether \tilde{C} is sound for $\overline{\text{SAT}}$, it suffices to construct an augmented Arthur-Merlin protocol for A_g , where g is the simplification of f defined by $g(C, \tilde{C}) \doteq |C^{-1}(0) \cap \tilde{C}^{-1}(1)|$.

Goldwasser and Sipser [GS86] showed that for every predicate $L \in \text{NP}$ and function

$$h : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \mathbb{N}, h : (u, v) \mapsto |\{w \in \{0, 1\}^* : \langle u, v, w \rangle \in L\}|, \quad (5)$$

the promise problem A_h is in prAM . Note that the function g is of the form (5), except that the underlying predicate $C(w) = 0 \wedge \tilde{C}(w) = 1$ is the difference of two NP sets rather than just an NP set. We remedy this issue by invoking the hypothesis $\text{coNP} \subseteq \text{NP/poly}$ as follows.

Let C and \tilde{C} have n inputs and be of size at most s . By the Cook-Levin Theorem, we can construct in time $\text{poly}(s)$ a Boolean formula $\phi_C(x, y)$ of size n' such that for all $x \in \{0, 1\}^n$, $C(x) = 0$ iff $\phi_C(x, \cdot) \in \overline{\text{SAT}}$. Let C' denote a circuit that takes inputs of size n' , and let L denote the predicate that on input $\langle u, v, w \rangle$ with $u = \langle C, C' \rangle$ and $v = \tilde{C}$, decides whether $C'(\phi_C(w, \cdot)) = 1 \wedge \tilde{C}(w) = 1$. Note that $L \in \text{NP}$, so $A_h \in \text{prAM}$ for h defined by (5).

Whenever C' is sound for $\overline{\text{SAT}}$ on instances of size n' , $\langle \langle C, C' \rangle, \tilde{C}, w \rangle \in L$ only if $C(w) = 0$ and $\tilde{C}(w) = 1$, which implies that $h(\langle \langle C, C' \rangle, \tilde{C} \rangle) \leq g(C, \tilde{C})$. Moreover, if C' correctly decides $\overline{\text{SAT}}$ at length n' , then $\langle \langle C, C' \rangle, \tilde{C}, w \rangle \in L$ iff $C(w) = 0 \wedge \tilde{C}(w) = 1$, and $h(\langle \langle C, C' \rangle, \tilde{C} \rangle) = g(C, \tilde{C})$. By the hypothesis that $\text{coNP} \subseteq \text{NP/poly}$, there exists a circuit C' of size $\text{poly}(n')$ that computes $\overline{\text{SAT}}$ on instances of length n' . This suggests the following augmented Arthur-Merlin protocol for A_g on input $\langle C, \tilde{C}, a, \epsilon \rangle$.

If $a \leq 0$ the protocol trivially accepts, as g is nonnegative. Otherwise, Merlin sends as his initial message a circuit C' of size $\text{poly}(n')$ on n' -inputs, purported to be a circuit for $\overline{\text{SAT}}$ at length n' . The coNP -verifier checks that C' is sound with respect to $\overline{\text{SAT}}$, and Arthur engages in a protocol with Merlin for the promise problem A_h on input $\langle \langle C, C' \rangle, \tilde{C}, a, \epsilon \rangle$.

To argue completeness of the protocol for A_g , suppose that $g(C, \tilde{C}) \geq a > 0$. In order to make both verifiers accept, Merlin can send as his initial message a polynomial-size circuit $C'_{\overline{\text{SAT}}}$ for $\overline{\text{SAT}}$ at length n' , which exists by the hypothesis that $\text{coNP} \subseteq \text{NP/poly}$. Since $C'_{\overline{\text{SAT}}}$ is sound for $\overline{\text{SAT}}$, the coNP -verifier accepts. As for the other verifier (Arthur), since $C'_{\overline{\text{SAT}}}$ is a correct circuit for $\overline{\text{SAT}}$, we have that $h(\langle \langle C, C'_{\overline{\text{SAT}}} \rangle, \tilde{C} \rangle) = g(C, \tilde{C}) \geq a$. The completeness of the Arthur-Merlin protocol for A_h then guarantees that Arthur can be convinced with high probability.

To argue soundness, suppose that $g(C, \tilde{C}) < a(1 - \epsilon)$. First, if the coNP -verifier accepts, then Merlin must have sent a sound circuit C' for $\overline{\text{SAT}}$ in the first round. In that case $h(\langle \langle C, C' \rangle, \tilde{C} \rangle) \leq g(C, \tilde{C})$, so $h(\langle \langle C, C' \rangle, \tilde{C} \rangle) < a(1 - \epsilon)$. By the soundness of the Arthur-Merlin protocol for A_h , this means that Arthur rejects with high probability. Thus, whenever the coNP -verifier accepts, Arthur rejects with high probability. This completes the proof. \square

Lemma 1 allows us to efficiently improve the completeness of a sound but incomplete circuit C for $\overline{\text{SAT}}$ when given oracle access to a language that agrees with A_f by finding a circuit \tilde{C} that approximately maximizes $f(C, \tilde{C})$, and outputting $C \vee \tilde{C}$. The approximate maximization can be done using the following generic lemma.

Lemma 2. *Let $f : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \mathbb{N}$ be such that $f(x, y) \leq 2^{|x|^c}$ for some constant c . If $A_f \in \text{prM}(\text{AM}|\text{coNP})$, where A_f denotes the promise problem defined by (4), then there exists a promise problem $\Pi \in \text{prM}(\text{AM}|\text{coNP})$ such that the following holds for any language L that agrees with Π . On input $x \in \{0, 1\}^*$, a nonnegative integer m in unary, and a positive integer $1/\tilde{\epsilon}$ in unary, we can find, in deterministic polynomial time with oracle access to L , a value $\tilde{y} \in \{0, 1\}^m$ such that*

$$f(x, \tilde{y}) \geq (1 - \tilde{\epsilon}) \cdot \max_{y \in \{0, 1\}^m} f(x, y).$$

Lemma 2 holds more generally when $\text{prM}(\text{AM}|\text{coNP})$ is replaced by any reasonable complexity class \mathcal{C} that is closed under the existential operator,³ such as Σ_k^P for any $k \geq 1$. Chakaravarthy and Roy [CR08] implicitly use it with $\mathcal{C} = \text{prAM}$.

Proof (of Lemma 2). We run a prefix search for \tilde{y} . In order to do so, we make use of the auxiliary function $g(x, y) \doteq \max_{y' \in \{0, 1\}^m} f(x, yy')$, where yy' denotes the concatenation of y and y' . The fact that $A_f \in \text{prM}(\text{AM}|\text{coNP})$ implies that $\Pi \doteq A_g \in \text{prM}(\text{AM}|\text{coNP})$. This is because on input $\langle x, y, a, \epsilon \rangle$, the augmented Arthur-Merlin protocol for A_g can have Merlin first guess $y' \in \{0, 1\}^{m-|y|}$ and then run the augmented Arthur-Merlin protocol for A_f on input $\langle x, yy', a, \epsilon \rangle$. Let L denote a language that agrees with the promise problem Π .

In a first phase we find an approximation \tilde{a} to $a^* \doteq \max_{y \in \{0, 1\}^m} f(x, y)$. To do so, we make use of the predicate $P(a) \doteq \langle x, \lambda, a, \epsilon \rangle \in L$, where λ denotes the empty string, ϵ will be set later, and the rest of the parameters are the inputs given in the statement of the lemma. Note that $P(0)$ holds because f is nonnegative. At the other end, $P(2^{|x|^c} + 1)$ fails by the assumption on the range of f . We run a binary search for an integer value $\tilde{a} \in [0, 2^{|x|^c}]$ such that (i) $P(\tilde{a})$ holds and (ii) $P(\tilde{a} + 1)$ fails. This guarantees that $(1 - \epsilon)\tilde{a} \leq a^* \leq \tilde{a}$, where the first inequality follows from (i) and the fact that $P(\tilde{a})$ implies that $g(x, \lambda) \geq (1 - \epsilon)\tilde{a}$, and the second inequality follows from (ii) and the fact that $\neg P(\tilde{a} + 1)$ implies that $g(x, \lambda) < \tilde{a} + 1$. This concludes the first phase.

In a second phase we run the actual prefix search for \tilde{y} . We maintain the invariant that

$$g(x, \tilde{y}_{1\dots i}) \geq \tilde{a}_i, \tag{6}$$

for $0 \leq i \leq m$, where $\tilde{y}_{1\dots i}$ denotes the prefix of length i of \tilde{y} , and the values \tilde{a}_i are chosen not too much smaller than \tilde{a} . More specially, we set $\tilde{a}_0 = (1 - \epsilon)\tilde{a}$, and for $i = 0, \dots, m - 1$ we extend the prefix of \tilde{y} of length i to length $i + 1$ as follows. By (6) we know that for at least one choice of $\tilde{y}_{i+1} \in \{0, 1\}$,

$$\langle x, \tilde{y}_{1\dots i+1}, \tilde{a}_i, \epsilon \rangle \in L, \tag{7}$$

and for any choice of $\tilde{y}_{i+1} \in \{0, 1\}$ satisfying (7), $g(x, \tilde{y}_{1\dots i+1}) \geq (1 - \epsilon)\tilde{a}_i$. Thus, we may pick \tilde{y}_{i+1} as the lesser value in $\{0, 1\}$ for which (7) holds, and set $\tilde{a}_{i+1} = (1 - \epsilon)\tilde{a}_i$.

In the end, we obtain $\tilde{y} \in \{0, 1\}^m$ satisfying

$$f(x, \tilde{y}) = g(x, \tilde{y}) \geq \tilde{a}_m = (1 - \epsilon)^m \tilde{a}_0 \geq (1 - \epsilon)^{m+1} \tilde{a} \geq (1 - \epsilon)^{m+1} a^*,$$

³More precisely, the property needed is that if $\Pi = (\Pi_Y, \Pi_N)$ is in \mathcal{C} then so is $\Pi' = (\Pi'_Y, \Pi'_N)$, where $\Pi'_Y = \{\langle x, y', 1^m \rangle : (\exists y'' \in \{0, 1\}^m) \langle x, y'y'' \rangle \in \Pi_Y\}$ and $\Pi'_N = \{\langle x, y', 1^m \rangle : (\forall y'' \in \{0, 1\}^m) \langle x, y'y'' \rangle \in \Pi_N\}$.

which is at least $(1 - \tilde{\epsilon})a^*$ provided we set $\epsilon = \tilde{\epsilon}/(m + 1)$.

Since both phases run in polynomial time with oracle access to L , the result follows. \square

Starting from the trivial sound circuit C_0 that rejects all inputs, we iteratively apply the improvement step based on Lemma 1 and Lemma 2 with $\tilde{\epsilon} = 1/2$. After no more than n iterations this yields a circuit of polynomial size that decides $\overline{\text{SAT}}$ on inputs of size n . We have proved the following theorem.

Theorem 5. *Suppose that $\text{coNP} \subseteq \text{NP}/\text{poly}$. There exists a promise problem $\Pi \in \text{prM}(\text{AM}||\text{coNP})$ such that the following holds for any language L that agrees with Π . Given n , we can construct a polynomial-size nondeterministic circuit for $\overline{\text{SAT}}$ at length n in deterministic polynomial time with oracle access to L .*

With Theorem 5 in hand, the nondeterministic variant of the Karp-Lipton argument yields the collapse stated in Theorem 4.

Proof (of Theorem 4). Let K denote the $\Sigma_3\text{P}$ -complete language consisting of all Boolean formulas $\varphi(x, y, z)$ on three sets of variables x, y, z such that $(\exists x)(\forall y) \varphi(x, y, \cdot) \in \text{SAT}$. We show that under the assumptions of the theorem, $K \in \text{PprM}(\text{AM}||\text{coNP})$.

Consider the related language K' consisting of all pairs $\langle \varphi, C \rangle$, where φ is a Boolean formula as above and C is a circuit such that $(\exists x)(\forall y) C(\varphi(x, y, \cdot)) = 0$. As the condition $C(\varphi(x, y, \cdot)) = 0$ can be decided in quasilinear time on a co-nondeterministic machine, $K' \in \Sigma_2\text{P} \subseteq \text{M}(\text{AM}||\text{coNP})$. Moreover, if C is a circuit that correctly decides $\overline{\text{SAT}}$ on inputs of the appropriate size, then $\varphi \in K$ iff $\langle \varphi, C \rangle \in K'$.

In order to decide L on an input φ of size n , we first run the algorithm from Theorem 5 on input n to obtain a circuit C of polynomial size for $\overline{\text{SAT}}$ on inputs of the required size, and then check whether $\langle \varphi, C \rangle \in K'$. Note that any language L that agrees with the promise problem $\Pi \in \text{prM}(\text{AM}||\text{coNP})$ from Theorem 5, suffices as oracle for the construction; the circuit C we construct may depend on the choice of L , but the final membership decision to K does not. The theorem follows. \square

5 Equivalence Result

In this section we establish our hardness-derandomization equivalence for Arthur-Merlin games (Theorem 1). We first argue the derandomization-to-hardness direction for $\Sigma_2\text{E}$, as well as a weaker but relativizing claim for $\text{E}^{\Sigma_2\text{P}}$. We finish with the hardness-to-derandomization direction for $\Sigma_2\text{E}$.

5.1 From derandomization to hardness

We use the lower bound framework introduced in Section 2.3. We assume that $\mathcal{C} \subseteq \text{NP}/\text{poly}$, where $\mathcal{C} = \Sigma_2\text{E}$ or $\mathcal{C} = \text{E}^{\Sigma_2\text{P}}$, and derive a contradiction with Kannan's result that the polynomial-time hierarchy does not have (nondeterministic) circuits of fixed-polynomial size. The derivation entails two key ingredients: (i) collapsing the polynomial-time hierarchy to some randomized class \mathcal{R} , and (ii) derandomizing \mathcal{R} assuming that $\text{prAM} \subseteq \bigcap_{\epsilon > 0} \text{i.o.} - \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$. We developed (i) in Section 4. We now discuss (ii).

The classes \mathcal{R} we consider involve augmented Arthur-Merlin protocols as introduced in Definition 1 of Section 3. More precisely, we consider $\mathcal{R} = \text{prM}(\text{AM}|\text{coNP})$ and $\mathcal{R} = \text{PprM}(\text{AM}|\text{coNP})$. Under the stronger derandomization assumption that $\text{prAM} \subseteq \Sigma_2\text{SUBEXP} \doteq \bigcap_{\epsilon > 0} \Sigma_2\text{TIME}(2^{n^\epsilon})$, those classes \mathcal{R} can trivially be simulated in $\Sigma_2\text{SUBEXP}$ or $\text{SUBEXP}^{\Sigma_2\text{P}}$, respectively. To carry over that argument to the i.o.-setting with small advice, we need to make sure that the derandomization of \mathcal{R} on inputs of length n only makes use of the derandomization of prAM on one of the infinitely many good input lengths m where the latter simulation is guaranteed to work. The following lemma shows how to do that for infinitely many input lengths n by exploiting the paddability of prAM and using an additional short advice string to point to a nearby good length m .

Lemma 3 (Derandomization Lemma). *If $\text{prAM} \subseteq \bigcap_{\epsilon > 0} \text{i.o.} - \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$ then*

- (i) $\text{prM}(\text{AM}|\text{coNP}) \subseteq \bigcap_{\epsilon > 0} \text{i.o.} - \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$, and
- (ii) $\text{PprM}(\text{AM}|\text{coNP}) \subseteq \bigcap_{\epsilon > 0} \text{i.o.} - \text{DTIME}^{\Sigma_2\text{P}}(2^{n^\epsilon})/n^\epsilon$.

Proof. Part (i): Let \mathcal{C} denote $\bigcap_{\epsilon > 0} \text{i.o.} - \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$, and let $\Pi \in \text{prM}(\text{AM}|\text{coNP})$. Let Γ denote the prAM -problem underlying Π , and V the coNP -verifier, as in Definition 1. By assumption, there is a language $Q \in \mathcal{C}$ that agrees with Γ . If we replace the predicate Γ by Q in the definition of Π , we obtain the language

$$R = \{x : (\exists y \in \{0, 1\}^{n^c}) (\langle x, y \rangle \in Q \wedge \langle x, y \rangle \in V)\}. \quad (8)$$

Observe that R agrees with Π . It remains to show that $R \in \mathcal{C}$.

We can assume without loss of generality that Q is paddable; by this we mean (a) $\langle x, y \rangle \in Q$ iff $\langle x, y, 0^{pad} \rangle \in Q$ for all $pad \in \mathbb{N}$, and (b) if $\langle x, y \rangle$ is of length m then there is a setting of pad such that $\langle x, y, 0^{pad} \rangle$ is of length m' for all $m' \geq m$.

Fix any $\epsilon > 0$. We want to exhibit a language $R_\epsilon \in \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$ that agrees with R on infinitely many lengths n . By assumption, for every $\delta > 0$ there is a language $Q_\delta \in \Sigma_2\text{TIME}(2^{m^\delta})/m^\delta$ and an infinite set of lengths M_δ such that Q_δ agrees with Q on all lengths in M_δ . We use Q_δ for a sufficiently small value of $\delta > 0$ to construct R_ϵ as follows.

Let $\ell(n)$ denote the maximum length of the (unpadded) queries issued to the language Q when deciding the language R on inputs x of length n . Suppose there exists a length $m \in M_\delta$ in the range $\ell(n) \leq m \leq \ell(n+1)$. Then we pick such a length m , give R_ϵ at length n as advice the value of m as well as the advice for Q_δ at length m , and let R_ϵ at length n be defined by (8) but with each query “ $\langle x, y \rangle \in Q$ ” replaced by the equivalent query to Q_δ padded to length m , i.e., by “ $\langle x, y, 0^{pad} \rangle \in Q_\delta$ ”, where pad is such that $|\langle x, y, 0^{pad} \rangle| = m$. Note that in this case R_ϵ agrees with R at length n . If there is no length $m \in M_\delta$ in the range $\ell(n) \leq m \leq \ell(n+1)$, we define R_ϵ in the same way but with m set to $\ell(n)$. In this case there is no guarantee that R_ϵ and R agree at length n .

Since the intervals $[\ell(n), \ell(n+1)]$ cover all but finitely many lengths m , and Q_δ agrees with Q for infinitely many lengths m , R_ϵ agrees with R on infinitely many lengths n .

All that remains is the complexity analysis of R_ϵ . The queries to Q_δ are padded up to length no more than $\ell(n+1)$, which is polynomially bounded in n . It follows that those queries can be decided in $\Sigma_2\text{TIME}(2^{n^{c\delta}})/n^{c\delta}$ for some fixed constant c . The advice for R_ϵ is of length at most $\log(\ell(n+1)) + n^{c\delta}$. Thus, if we set $\delta = \epsilon/(c+1)$ we have that $R_\epsilon \in \Sigma_2\text{TIME}(2^{n^\epsilon})/n^\epsilon$. This completes part (i).

Part (ii): Let $L \in \text{DTIME}^\Pi(n^c)$ where $\Pi \in \text{prM}(\text{AM}|\text{coNP})$. By padding all queries of the base machine up to an appropriate length m depending on n as in part (i), and providing the base machine with the advice we gave R_ϵ in part (i), we obtain that

$$L \in \cap_{\epsilon>0} \text{i.o.} - \text{DTIME}^{\Sigma_2\text{TIME}(2^{n^\epsilon})}(n^d)/n^\epsilon,$$

where d is a constant depending on c . Since $\Sigma_2\text{TIME}(2^{n^\epsilon}) \subseteq \text{DTIME}^{\Sigma_2\text{P}}(2^{n^\epsilon})$, the conclusion follows. \square

For our main result, and more generally for superpolynomial lower bounds, the framework exploits the fact that linear-exponential classes such as E , NE , etc. have polynomial-size circuits if and only if they have *fixed*-polynomial-size circuits; this follows because those classes contain complete languages under linear-time reductions. The next lemma formalizes this fact in a way that will be handy later.

Lemma 4. *Let $\mathcal{C} \in \{\Sigma_2\text{E}, \text{E}^{\Sigma_2\text{P}}\}$. Suppose $\mathcal{C} \subseteq \text{NP}/\text{poly}$. Then every language in \mathcal{C}/n has nondeterministic circuits of size n^d for all but finitely many input lengths, where d is a fixed constant.*

The framework derives a contradiction with the following nondeterministic version of a classical result of Kannan's.

Lemma 5 (implicit in [Kan82]). *For every constant $d > 0$ there exists a language in $\text{P}^{\Sigma_3\text{P}}$ that requires nondeterministic circuits of size n^d for all but finitely many input lengths n .*

We include a proof for completeness.

Proof. Let $s(n) = n^d$. For all but finitely many n , $s(n)$ is no more than the maximum circuit complexity at length n . Therefore, there exists a characteristic sequence of length 2^n that cannot be computed by circuits of size less than $s(n)$. Moreover, the length ℓ of the shortest prefix σ such that no circuit of size less than $s(n)$ agrees with σ satisfies $\ell = O(s(n) \log s(n))$. This follows because circuits of size $s(n)$ can be described by strings of length $O(s(n) \log s(n))$ and distinct prefixes need distinct circuits.

The lexicographically least such prefix σ , say σ^* , can be found through a binary search, by using a $\Sigma_3\text{P}$ -oracle, in time $\text{poly}(s(n))$. To see this, given a size- s circuit C and a length- ℓ string σ , consider the task of deciding whether the circuit C does not agree with σ . This task can be performed with an NP -oracle in time $\text{poly}(s(n))$, and hence in $\Pi_2\text{TIME}(\text{poly}(s(n)))$. To run the binary search, we need to answer queries as to whether a given string can be extended to a string σ such that for all circuits C of size less than $s(n)$, C does not agree with σ . As these queries can be decided in $\Sigma_3\text{TIME}(\text{poly}(s(n)))$ and $s(n)$ is polynomial, we can construct σ^* in polynomial time with oracle access to $\Sigma_3\text{P}$.

Once found, σ^* is viewed as a string σ' of length 2^n that is all zeroes beyond the first ℓ bits (and that equals σ^* in its first ℓ bits). It follows that σ' is the characteristic string of a language L that cannot be decided by a circuit of size less than $s(n)$ at length n . On input x of length n , whether $x \in L$ can be decided according to the x th bit of σ' . Thus, $L \in \text{P}^{\Sigma_3\text{P}}$. \square

We now have all the ingredients to instantiate the lower bound framework described in Section 2.3 and obtain the following derandomization-to-hardness results for Arthur-Merlin games.

Theorem 6. (i) If $\text{prAM} \subseteq \cap_{\epsilon > 0} \text{i.o.} - \Sigma_2 \text{TIME}(2^{n^\epsilon})/n^\epsilon$ then $\Sigma_2 \text{E} \not\subseteq \text{NP/poly}$.

(ii) Relative to any oracle, if $\text{prAM} \subseteq \cap_{\epsilon > 0} \text{i.o.} - \Sigma_2 \text{TIME}(2^{n^\epsilon})/n^\epsilon$ then $\text{E}^{\Sigma_2 \text{P}} \not\subseteq \text{NP/poly}$.

Proof. For part (i) we set $\mathcal{C} = \Sigma_2 \text{E}$ and $\mathcal{R} = \text{M}(\text{AM} \parallel \text{coNP})$, and for part (ii) we set $\mathcal{C} = \text{E}^{\Sigma_2 \text{P}}$ and $\mathcal{R} = \text{P}^{\text{prM}(\text{AM} \parallel \text{coNP})}$. The following derivation proves both parts by contradiction.

$$\begin{aligned}
& \mathcal{C} \subseteq \text{NP/poly} \\
\implies & \text{PH} \subseteq \mathcal{R} && \text{(by the collapse results from Section 4)} \\
\implies & \text{PH} \subseteq \text{i.o.} - \mathcal{C}/n && \text{(by Lemma 3)} \\
\implies & \text{PH} \subseteq \text{i.o.} - \text{NSIZE}(n^d) && \text{(by Lemma 4)} \\
\implies & \text{contradiction} && \text{(by Lemma 5),}
\end{aligned}$$

where d is some constant and $\text{NSIZE}(n^d)$ denotes the class of languages with nondeterministic circuits of size n^d .

The relativization claim in (ii) follows because all steps in that part relativize (whereas the collapse argument for (i) involves a nonrelativizing step). \square

Note that part (i) of Theorem 6 yields the forward direction of Theorem 1.

5.2 From hardness to derandomization

Finally, we argue the direction from hardness to derandomization in Theorem 1. This direction follows in a straightforward way from the known hardness versus randomness tradeoffs.

Theorem 7 (implicit in [KvM02, SU06]). Let s and σ denote monotone constructible functions from \mathbb{N} to \mathbb{N} . If

$$\Sigma_2 \text{E} \not\subseteq \text{NSIZE}(s(n))$$

then there exists a pseudorandom generator that yields the derandomization

$$\text{prAM} \subseteq \text{i.o.} - \Sigma_2 \text{TIME}(2^{\sigma(n^{O(1)})})/\sigma(n^{O(1)}),$$

provided that $\sigma(n) \geq (s^{-1}(n^c))^2/\log n$, where $c > 0$ is a universal constant and $s^{-1}(m)$ denotes $\min\{n \in \mathbb{N} : s(n) \geq m\}$.

Given the hardness hypothesis, for any $\epsilon > 0$, we can pick $s(n)$ in Theorem 7 to be n^k for a large enough k and get a PRG that yields the simulation of a polynomial-time Arthur-Merlin protocol in $\Sigma_2 \text{TIME}(2^{n^\epsilon})/n^\epsilon$ for infinitely many input lengths n . This establishes the backward direction of Theorem 1.

Proof (of Theorem 7). The proof follows by combining two lemmas that are implicit in the literature. First, in order to derandomize prAM it suffices to construct pseudorandom generators that fool nondeterministic circuits.

Lemma 6 (implicit in [KvM02]). If there is a pseudorandom generator G that on seed length $\sigma(n)$ is computable in $\Sigma_2 \text{TIME}(2^{O(\sigma(n))})$ with advice of length $\sigma(n)$, and fools nondeterministic circuits for infinitely many n , then prAM can infinitely often be simulated in $\Sigma_2 \text{TIME}(2^{O(\sigma(n^{O(1)})}))n^{O(1)}$ with advice of length $\sigma(n^{O(1)})$.

The pseudorandom generators needed in Lemma 6 follow from the given hardness assumption by the known hardness versus randomness tradeoffs for prAM.

Lemma 7 (implicit in [SU06]). *There exists a constant $c > 0$ such that the following holds for any constructible function $\ell : \mathbb{N} \rightarrow \mathbb{N}$. If there is a language in $\Sigma_2\text{E}$ that requires nondeterministic circuits of size n^c at length $\ell(n)$ for infinitely many n , then there exists a pseudorandom generator G that has constructible seed length $\sigma(n) = O(\ell^2(n)/\log n)$, is computable in $\Sigma_2\text{TIME}(2^{O(\sigma(n))})$ on seed length $\sigma(n)$ with advice of length $\sigma(n)$, and fools nondeterministic circuits at infinitely many lengths n .*

Using G in Lemma 6 yields the required derandomization of prAM. □

6 Concluding Remarks

Both our main result (Theorem 1) and the corresponding result for prBPP in [IKW02] establish an equivalence of derandomization and PRGs at just one point within the derandomization spectrum. A natural question to ask is whether these equivalences can be extended to other regions of the derandomization spectrum: simulations of prBPP/prAM using less time, less advice, or with one fewer alternation – if they can be done at all, then can they be done through PRGs? What about simulations that succeed on all-but-finitely-many input lengths rather than infinitely-many?

As for the amount of advice, in both the prBPP and the prAM setting we can show that an equivalence still holds if we reduce the advice from subpolynomial down to polylogarithmic. In the setting of prMA we can reduce the advice down to a constant. These results follow from parameterizations of the proofs in Sections 4 and 5, and will be developed in the final version of the paper.

References

- [AGHK11] Barış Aydınhoğlu, Dan Gutfreund, John M. Hitchcock, and Akinori Kawachi. Derandomizing Arthur-Merlin games and approximate counting implies exponential-size lower bounds. *Computational Complexity*, 20(2):329–366, 2011.
- [AvM11] Scott Aaronson and Dieter van Melkebeek. On circuit lower bounds from derandomization. *Theory of Computing*, 7:177–184, 2011.
- [BFNW93] László Babai, Lance Fortnow, Noam Nisan, and Avi Wigderson. BPP has subexponential time simulations unless EXPTIME has publishable proofs. *Computational Complexity*, 3:307–318, 1993.
- [CCHO05] Jin-Yi Cai, Venkatesan T. Chakaravarthy, Lane A. Hemaspaandra, and Mitsunori Ogi-hara. Competing provers yield improved Karp-Lipton collapse results. *Information and Computation*, 198(1):1–23, 2005.
- [CR08] Venkatesan T. Chakaravarthy and Sambuddha Roy. Finding irrefutable certificates for S_2^P via Arthur and Merlin. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 157–168, 2008.

- [Gol06] Oded Goldreich. On promise problems: A survey. In *Essays in Memory of Shimon Even*, pages 254–290, 2006.
- [Gol11a] Oded Goldreich. In a world of $P=BPP$. In *Studies in Complexity and Cryptography*, pages 191–232. 2011.
- [Gol11b] Oded Goldreich. Two comments on targeted canonical derandomizers. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:47, 2011.
- [GS86] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 59–68, 1986.
- [GST03] Dan Gutfreund, Ronen Shaltiel, and Amnon Ta-Shma. Uniform hardness versus randomness tradeoffs for Arthur-Merlin games. *Computational Complexity*, 12(3-4):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 *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 220–229, 1997.
- [Kan82] Ravi Kannan. Circuit-size lower bounds and nonreducibility to sparse sets. *Information and Control*, 55(1):40–56, 1982.
- [KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1/2):1–46, 2004.
- [KL82] Richard M. Karp and Richard J. Lipton. Turing machines that take advice. *L’Enseignement Mathématique*, 28(2):191–209, 1982.
- [KvM02] Adam 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.
- [KvMS12] Jeff Kinne, Dieter van Melkebeek, and Ronen Shaltiel. Pseudorandom generators, typically-correct derandomization, and circuit lower bounds. *Computational Complexity*, 21(1):3–61, 2012.
- [LFKN92] Carsten Lund, Lance Fortnow, Howard J. Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.
- [NW94] Noam Nisan and Avi Wigderson. Hardness vs. randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994.
- [Sha92] Adi Shamir. $IP = PSPACE$. *Journal of the ACM*, 39(4):869–877, 1992.

- [SU06] Ronen Shaltiel and Christopher Umans. Pseudorandomness for approximate counting and sampling. *Computational Complexity*, 15(4):298–341, 2006.
- [Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 80–91, 1982.
- [Yap83] Chee-Keng Yap. Some consequences of non-uniform conditions on uniform classes. *Theoretical Computer Science*, 26:287–300, 1983.