

Improving Exhaustive Search Implies Superpolynomial Lower Bounds

Ryan Williams*
IBM Almaden Research Center

November 7, 2009

Abstract

The P vs NP problem arose from the question of whether exhaustive search is necessary for problems with short verifiable solutions. We do not know if even a slight algorithmic improvement over exhaustive search is universally possible for all NP problems, and to date no major consequences have been derived from the assumption that an improvement exists.

We show that there are natural NP and BPP problems for which minor algorithmic improvements over the trivial deterministic simulation already entail lower bounds such as $\text{NEXP} \not\subseteq \text{P/poly}$ and $\text{LOGSPACE} \neq \text{NP}$. These results are especially interesting given that similar improvements *have* been found for many other hard problems. Optimistically, one might hope our results suggest a new path to lower bounds; pessimistically, they show that carrying out the seemingly modest program of finding slightly better algorithms for all search problems may be extremely difficult (if not impossible).

We also prove unconditional superpolynomial time-space lower bounds for improving on exhaustive search: there is a problem verifiable with $k(n)$ length witnesses in $O(n^a)$ time (for some a and some function $k(n) \leq n$) that cannot be solved in $k(n)^c n^{a+o(1)}$ time and $k(n)^c n^{o(1)}$ space, *for every* $c \geq 1$. While such problems can always be solved by exhaustive search in $O(2^{k(n)} n^a)$ time and $O(k(n) + n^a)$ space, we can prove a *superpolynomial* lower bound in the parameter $k(n)$ when space usage is restricted.

*This work originated while the author was a member of the Institute for Advanced Study in Princeton, New Jersey, supported by NSF Grant CCF-0832797 (Expeditions in Computing) and NSF Grant DMS-0835373 (Pseudorandomness). At IBM, the author is supported by the Josef Raviv Memorial Fellowship.

1 Introduction

To what extent can we avoid exhaustive search for generic search problems? This is one of the many offshoots of the P versus NP problem and it is the central question addressed in the area of exact algorithms for NP-hard problems.¹ The general message of this research has been that the trivial enumeration of all possible solutions *can* be quantitatively improved upon for many problems (and their number is increasing every year). The broad success of this program leads one to wonder if *all* NP problems allow some modest improvement over brute-force search. More precisely, let $V(x, y)$ be a verifier that runs on witnesses y of length $|x|^c$ and takes $O(|x|^d)$ time. The runtime of exhaustive search is $O(2^{|x|^c} |x|^d)$. Can we always find a witness in $O(2^{.99|x|^c} |x|^d)$? What about $O(2^{|x|^c - \log^2 |x|} |x|^{100d})$? Can we approximate the fraction of witnesses that V accepts in this time? If V uses only $s(n)$ space, can we find a witness in $O(2^{.99|x|^c} |x|^d)$ time and $O(|x|^{100c} + s(n))$ space? These questions are the focus of the present paper.

Here we offer the first concrete evidence that the above minor improvements will be difficult to achieve for some problems. We do not rule out the possibility of these improvements; rather, we show that they would already imply superpolynomial lower bounds in complexity theory. (Optimists might say we have given a new approach to proving class separations.) We show that the task of finding improved algorithms is connected naturally to the task of proving superpolynomial lower bounds.

1.1 Main Results

1.1.1 Improved Algorithms Imply Circuit Lower Bounds

Let CIRCUIT SAT be the problem of finding a satisfying assignment to a Boolean circuit (over AND, OR, NOT gates). The CIRCUIT SAT problem is often the first NP-complete problem introduced in textbooks. Due to its applications in hardware/software verification and automated reasoning, the problem has been extensively studied in many areas of computer science. (Although CNF satisfiability gets all the fanfare, most formulas obtained in practice start as Boolean circuits.) Clearly, for circuits with n input variables and $\text{poly}(n)$ gates, the problem can be solved in $2^n \cdot \text{poly}(n)$ time on any robust model of computation.²

Theorem 1.1 *Suppose there is a superpolynomial function $s(n)$ such that CIRCUIT SAT on circuits with n variables and n^k gates can be solved in $2^n \cdot \text{poly}(n^k)/s(n)$ time by a (co-non)deterministic algorithm, for all k . Then $\text{NEXP} \not\subseteq \text{P/poly}$.*

That is, practically any nontrivial improvement over exhaustive search (or nontrivial proofs of unsatisfiability) implies superpolynomial circuit lower bounds for nondeterministic exponential time. This would be a breakthrough in the area of circuit complexity. The contrapositive is also interesting: it says that the small-circuit assumption for NEXP implies nearly tight lower bounds on CIRCUIT SAT.

The best known deterministic algorithm for CNF satisfiability takes $2^{n - \Omega(n/\ln(m/n))} \text{poly}(m)$, where m is the number of clauses and n is the number of variables [DH08, CIP06]. For more complex circuits, the current state of the art is a randomized SAT algorithm for AC^0 circuits which runs in $2^{n - n^{1-o(1)}}$ time on circuits with $n^{1+o(1)}$ gates [CIP09]. Both time bounds are noticeably smaller than the hypothesis of

¹In fact, in Gödel's famous letter to Von Neumann, he posed precisely this question:

“It would be interesting to know... in the case of finite combinatorial problems, how strongly *in general* the number of steps vis-à-vis *dem blossen Probieren* [the “bare trying” of solutions] can be reduced.” [Har89]

²We use the $\text{poly}(n)$ notation to denote $O(n^c)$ factors for a fixed $c > 0$ independent of n .

Theorem 1.1.³ Of course this does not mean that the hypothesis can be achieved, but it does not seem implausible at present.

Scaled-down versions of Theorem 1.1 are possible. For example, if satisfiability of Boolean formulas on n variables and n^c connectives can be solved in $2^n \cdot \text{poly}(n^c)/s(n)$, then E^{NP} does not have (non-uniform) polynomial size formulas.⁴ If satisfiability of $AC^0[6]$ circuits (constant depth circuits with AND, OR, NOT, and MOD6 gates) on n variables and n^c gates can be solved in $2^{n/2} \cdot \text{poly}(n^c)/s(n)$, then E^{NP} does not have polynomial size $AC^0[6]$ circuits.

Interestingly, it does not take a much stronger assumption to get a nearly optimal circuit lower bound for problems in nondeterministic $2^{O(n)}$ time:

Theorem 1.2 *If CIRCUIT SAT on n variables and m gates is in $O(2^{(1-\delta)n} \text{poly}(m))$ (co-non)deterministic time for some $\delta > 0$, then there is $\varepsilon > 0$ and a language in $NTIME[2^n]$ that does not have $2^{\varepsilon n}$ size circuits.*

One may be skeptical that these improvements are possible for Circuit Satisfiability. Similar results can also be established for problems efficiently solvable with randomness. The *Circuit Acceptance Probability Problem* (CAPP) is to approximate (within $\pm 1/6$) the fraction of satisfying assignments of a given Boolean circuit on n variables and n^c gates for some c [KC99, KRC00, For01, Bar02, IKW02]. CAPP can be solved in $O(n^c)$ time with a randomized algorithm. It is known that CAPP is in polynomial time if and only if $\text{PromiseBPP} = P$ [For01], so the problem is “complete” in some sense. We show that essentially any nontrivial derandomization for CAPP implies superpolynomial circuit lower bounds.

Theorem 1.3 *If there is an $O(2^n \cdot \text{poly}(n^c)/s(n))$ nondeterministic algorithm for CAPP (for any super-polynomial $s(n)$), then $NEXP \not\subseteq P/\text{poly}$.*

(Here, a *nondeterministic approximation algorithm* for CAPP always has at least one accepting computation, and always outputs a good approximation when it enters an accept state.) In fact the proof of Theorem 1.3 holds when we replace CAPP with the problem of producing a satisfying assignment to a circuit when we are guaranteed that at least half of the assignments are satisfying.

Theorem 1.3 also implies interesting “amplification” results:

Theorem 1.4 *If CAPP has a nondeterministic $O(2^n \cdot \text{poly}(n)/s(n))$ time algorithm that succeeds on all inputs, then for all $\varepsilon > 0$, CAPP has a nondeterministic $O(2^{n^\varepsilon})$ time algorithm with n^ε advice that succeeds on infinitely many inputs.*

Theorem 1.5 *If CAPP on n variables and m gates has a nondeterministic $O(2^{(1-\delta)n} \cdot \text{poly}(m))$ time algorithm that succeeds on all inputs, then CAPP has a nondeterministic polynomial time algorithm with $O(\log n)$ advice that succeeds on infinitely many inputs.*

The prospects seem better for achieving an improved CAPP algorithm. (Note we do not necessarily have to build a pseudorandom generator to achieve the simulations.) For CNF formulas of length n , Luby and Velikovic [LV96] have given a deterministic $n^{2^{O(\sqrt{\log \log n})}}$ algorithm for approximating the number of solutions. Note it was already known that $NEXP \neq MA$ iff $NEXP \not\subseteq P/\text{poly}$ [IKW02]. However it could still

³ AC^0 circuits have constant depth and are comprised of AND, OR, NOT gates, with unbounded fan-in for each gate.

⁴ E^{NP} is the class of languages recognized by an algorithm running in $2^{O(n)}$ time that can query an NP oracle (with queries of $2^{O(n)}$ length). Linear size circuit lower bounds are not known for this class, but a $n^{3-o(1)}$ lower bound on formula size follows from work of Hastad [Has98].

have been the case that $\text{NEXP} = \text{MA}$ and yet Circuit SAT and CAPP have slightly better algorithms. The above results show this conjunction is impossible.

1.1.2 Improved Algorithms Imply $\text{LOGSPACE} \neq \text{NP}$ and *Subexponential Algorithms*

One strength of the above results is that there is no space restriction needed for the improved algorithms: our hypotheses only require SAT algorithms running in $O(2^n/s(n))$ time and $O(2^n/s(n))$ space for sufficiently large $s(n)$. Next we show that if exhaustive search can be improved in a way that preserves the verifier's space bound, then very surprising consequences result.

Here we shall study problems with *limited nondeterminism*, which only need short witnesses (e.g., of $\text{poly}(\log n)$ length) although verification still takes polynomial time. In particular we look at the case where the amount of nondeterminism is *logarithmic* in n ("log-nondeterminism"), and hence exhaustive search is already in *polynomial time*. This case is of interest as there are many polynomial time problems which fall in this category, and for which researchers have tried to find faster algorithms than a canonical one. The 3SUM problem is a prominent example [GO95]; satisfiability of exponential size circuits is another.⁵

It is natural to think that it may be easier to universally improve on log-nondeterminism problems. After all, the search space is no longer exponential in the running time of the verifier — now they are both polynomials in n . From this perspective it appears more likely that a clever pruning of the search space can be done.

Consider a problem solvable with witnesses of $\log n$ length, $O(n^c)$ time, and $\text{poly}(\log n)$ space.⁶ The obvious deterministic simulation runs in $O(n^{c+1})$ time and $\text{poly}(\log n)$ space. We show that any universal improvement in the runtime exponent can be amplified into an *arbitrary polynomial* speedup with nondeterministic algorithms. Such a speedup would have dramatic consequences.

Theorem 1.6 *Suppose for all $c, d \geq 1$ that every problem Π solvable with $\log n$ nondeterministic bits, n^c time, and $(\log n)^d$ space can be solved by some deterministic algorithm in $O(n^{c+.99})$ time and $\text{poly}(\log n)^d$ space. Then every such Π can also be solved by some nondeterministic algorithm in $O(n^3)$ time.*

Of course, the .99 is not special, and can be substituted with any $\delta < 1$.

Corollary 1.1 *The hypothesis of Theorem 1.6 implies $\text{LOGSPACE} \neq \text{NP}$; in fact, $\text{SC} \neq \text{NP}$, where SC is the class of problems solvable in (simultaneous) polynomial time and polylogarithmic space.*

Corollary 1.2 *The hypothesis of Theorem 1.6 implies that the quantified Boolean formula problem has a proof system where every QBF of length n has proofs of $2^{\varepsilon n}$ length for all $\varepsilon > 0$.*

That is, either the exponent in the trivial algorithm is optimal for some constant c , or we separate complexity classes in a surprising way. Note if the trivial algorithm is optimal for $c = 1$, then SAT cannot be solved in subquadratic time and polylog space, a problem that remains open despite much effort.⁷

⁵In 3SUM, we are given a set A of n numbers and wish to find three numbers that sum to zero. The problem can be easily solved in $O(n^2)$ by sorting, and it is a key bottleneck in the solution of many problems in geometry and data structures. Finding an $O(n^{1.99})$ algorithm is a major challenge. Note 3SUM (on $\text{poly}(\log n)$ -bit numbers) can be reduced to CIRCUIT SAT with $\log n$ inputs and $O(n \cdot \text{poly}(\log n))$ gates. The idea is to sort A in increasing order, and on input $i \in [n]$ use two pointers (one starting at the beginning of A and one at the end) to sweep A for numbers a and b such that the i th number plus $a + b$ is zero. Thus any improvement on exhaustive search for exponential circuit satisfiability implies a 3SUM improvement.

⁶In place of $\text{poly}(\log n)$ space, one may substitute any constructible function $f(n)$ satisfying $\log n \leq f(n) \leq n^{o(1)}$ for the remaining results in this paper. We have chosen $\text{poly}(\log n)$ for concreteness.

⁷In fact, using the above footnote on 3SUM, one can show that if 3SUM on N numbers (expressible in $\text{poly}(\log N)$ bits) cannot be solved in $O(N^{2-\varepsilon})$ for all $\varepsilon > 0$, then SAT cannot be solved in $O(n^{2-\varepsilon})$, for all $\varepsilon > 0$.

Identifying an explicit natural problem in place of “every problem” in Theorem 1.6 is nontrivial. (We could always use a form of the “Bounded Halting Problem”, but this is undesirable.) The proof of Theorem 1.6 uses a delicate inductive argument that makes it difficult to extend to lower bounds on a natural problem.

1.1.3 Unconditional Lower Bounds

We want to understand the extent to which the exponential part of exhaustive search (namely, the exponential runtime in the witness length) can be reduced, without affecting the runtime of verification. We can prove unconditional lower bounds for improving on exhaustive search in this manner, based on Theorem 1.6 and its consequences. We first state a superpolynomial lower bound in terms of the witness length. To our knowledge no similar lower bounds have been reported.

Theorem 1.7 *There is a problem Π verifiable with $k(n)$ -length witnesses in $O(n^a)$ time (for some constant a and some $k(n) \leq n$) that cannot be solved in $k(n)^c n^a \cdot \text{poly}(\log n)$ time and $k(n)^c \cdot (\log n)^c$ space, for all c .*

Note if we could “only” change the “some $k(n)$ ” to “all $k(n)$ ” then we would separate P from NP. The theorem does rule out certain (daffy, but until now possible) strategies for trying to show LOGSPACE = NP. For example, it is not possible to transform an arbitrary $O(n^a)$ time verifier with witnesses of length k into an $O(n^{a+o(1)})$ time and $O(\text{poly}(\log n))$ space verifier with witnesses of length $10^{10} \log k$. The proof uses the framework of Theorem 1.6. Theorem 1.7 implies the following lower bound for parameterized problems (where the parameter is k):

Corollary 1.3 *There is a parameterized problem Π solvable in $O(2^k n^a)$ time that cannot be solved in $k^c \cdot n^{a+o(1)}$ time and $k^c \cdot (\log n)^c$ space, for all c .*

Extending this result to hold for circuit satisfiability on general computational models is an interesting challenge. We manage to show a related lower bound for CIRCUIT SAT:

Theorem 1.8 *For every $\varepsilon > 0$, CIRCUIT SAT on k variables and m gates cannot be solved in $k^{1-\varepsilon} m^{1+o(1)}$ time and $m^{o(1)}$ space.*

This is already interesting, because we do not know how to prove that SAT is not in $O(n^{2-\varepsilon})$ time and $n^{o(1)}$ space on general models (the case where $k = m$).⁸ It is known that SAT (and consequently, CIRCUIT SAT) cannot be solved in less than $n^{1.8}$ time and $n^{o(1)}$ space [Wil08]. It is also known (in a precise sense) that current techniques will not allow us to extend this bound to n^2 [Wil09]. This limitation seems related to our inability to extend Theorem 1.8 to arbitrary polynomials in k .

1.1.4 A Nontrivial Simulation

In Appendix B, we report a baby step towards improving brute force search. We observe a nontrivial deterministic simulation of nondeterministic multitape Turing machines, inspired by a nontrivial derandomization of randomized multitape Turing machines of Van Melkebeek and Santhanam [vMS05].

Theorem 1.9 *For every multitape Turing machine M running in time t that accesses a tape of nondeterministic bits, there is a $\delta > 0$ and a deterministic TM that simulates M in $O(2^{(1-\delta)t})$ time.*

⁸Here n is the total length of the input.

A group led by Lipton [Lip09a] has recently proved related results. Our simulation is not yet enough to imply lower bounds, as the simulation runtime does not scale with the amount of nondeterminism.

1.2 An Overview of Our Techniques

Our basic approach is simple: we assume the opposite of the desired lower bound (i.e., that we have decent uniform algorithms and very good non-uniform circuits), and construct an efficient simulation of the hard class, so tight that any nontrivial CIRCUIT SAT or CAPP algorithm contradicts a time hierarchy theorem.

To illustrate, let us sketch the result that faster circuit satisfiability implies NEXP circuit lower bounds. It is known that $\text{NEXP} \subseteq \text{P/poly}$ implies a simulation of $\text{NTIME}[2^n]$ where we *existentially* guess a polynomial size circuit encoding a witness, then run an exponential time verifier [IKW02]. We make the observation that $\text{NEXP} \subseteq \text{P/poly}$ implies *every* verifier for an NEXP language has small circuits that encode witnesses. Therefore we are free to construct any verifier we need to get a contradiction. We choose one that exploits efficient reductions from NP to SAT, translated up to NEXP. Using a small witness circuit, we can replace the exponential time verifier with a single call to CIRCUIT SAT on $n + O(\log n)$ variables and $\text{poly}(n)$ gates. It follows that an $O(2^n \cdot \text{poly}(n)/s(n))$ algorithm for CIRCUIT SAT implies $\text{NTIME}[2^n] \subseteq \text{NTIME}[2^n \cdot \text{poly}(n)/s(n)]$, a contradiction. For the lower bound that follows from approximating the solutions to a circuit, we use the highly efficient *PCP of Proximity* verifiers of Ben-Sasson *et al.* [BGHSV05] in our CIRCUIT SAT instance.

To prove that stronger improved algorithms yield stronger consequences such as $\text{LOGSPACE} \neq \text{NP}$, we use ideas from the existing time-space lower bounds for SAT, along with an inductive trick that (assuming a universal speedup over exhaustive search) lets us repeatedly reduce any exponent in the time bound of a complexity class until it becomes a fixed constant, leading to superpolynomial lower bounds. Further development of this trick yields unconditional lower bounds for parameterized problems.

2 Preliminaries

We assume the reader is familiar with basic concepts in complexity and algorithms. In the following paragraphs, we focus on a few particulars needed for this paper.

2.1 Notation and Background

We define $[n] := \{1, \dots, n\}$. As usual, unless otherwise specified a function has domain \mathbb{N} and co-domain \mathbb{N} , is assumed to be time (or space) constructible within the appropriate bounds, and is monotone nondecreasing. All logarithms are in base two. A circuit is Boolean (with AND, OR, NOT gates), unless indicated otherwise.

Computational Model and Complexity. In almost all our results one may assume any deterministic computational model in which $O(\text{poly}(\log n))$ bits of information may be processed in a single step. This includes all flavors of Turing machines (including random access), random access machines with logarithmic size registers, and so on [GS89]. We shall indicate clearly when this is not the case.

Fix a finite alphabet Σ . We usually assume $\Sigma = \{0, 1\}$. $\text{NTIME}[t(n)]$ ($\text{TIME}[t(n)]$) denote the classes of languages recognized in $O(t(n))$ time by a nondeterministic (deterministic) algorithm, respectively. $\text{TISP}[t(n), s(n)]$ denotes those languages recognized by some algorithm that runs in both $O(t(n))$ time and $O(s(n))$ space. Recall $\text{P} = \bigcup_{k \geq 1} \text{TIME}[n^k]$ and $\text{NP} = \bigcup_{k \geq 1} \text{NTIME}[n^k]$. Steve's Class (abbreviated SC) is $\bigcup_{k \geq 1} \text{DTISP}[n^k, (\log n)^k]$.

i.o. – \mathcal{C} is the class of languages $L \subset \Sigma^*$ such that there is a language $L' \in \mathcal{C}$ where $L \cap \Sigma^n = L' \cap \Sigma^n$ holds for infinitely many n .

$(\exists f(n))\mathcal{C}$ is the class of languages which are recognized by machines with the following behavior: first, the machine nondeterministically guesses $f(n)$ bits, then it runs a subroutine recognizing a language in \mathcal{C} . The class $(\forall f(n))\mathcal{C}$ is defined similarly. For example, a language in the class $(\exists n)\text{TISP}[n^2, \log n]$ is recognized by some machine that on input x guesses some $|x|$ -bit string y then runs a computation taking $O(|x|^2)$ time and $O(\log |x|)$ space on the input $\langle x, y \rangle$. (We assume the machine has a tape alphabet large enough to accommodate an efficient pairing function.)

We shall use the following translation lemmas, whose proofs follow by standard “padding arguments.”

Lemma 2.1 (Translation Lemma) *For all constructible $f(n)$, $t_1(n)$, $t_2(n)$, and $s(n)$,
If $\text{NTIME}[t_1(n)] \subseteq \text{TISP}[t_2(n), s(n)]$ then $\text{NTIME}[t_1(f(n))] \subseteq \text{TISP}[t_2(f(n)), s(f(n))]$.
If $\text{TISP}[t_2(n), s(n)] \subseteq \text{NTIME}[t_1(n)]$ then $\text{TISP}[t_2(f(n)), s(f(n))] \subseteq \text{NTIME}[t_1(f(n))]$.*

We say that L is a *unary language* if $L \subseteq \{1\}^*$. We shall need the following strong nondeterministic time hierarchy theorem.

Theorem 2.1 ([Zak83]) *Let t_1 and t_2 be time constructible functions such that $t_1(n+1) \leq o(t_2(n))$. There is a unary language in $\text{NTIME}[t_2(n)]$ that is not in $\text{NTIME}[t_1(n)]$.*

2.2 Related Work

Subexponential Complexity and Parameterized Complexity. A theory of (sub)exponential time complexity for NP-hard problems has been developing. Much work has gone into extending reducibility notions to subexponential time algorithms (e.g., [IP01, IPZ01, CJ03, CG07, CIP06, Tra08]). The theory has found interesting evidence for the *exponential time hypothesis* (ETH) that 3-SAT is not in $2^{o(m)}$ time, where m is the number of clauses. If ETH is false, then many other problems have subexponential algorithms as well.

In parameterized complexity, one studies paired problems with an input and parameter k written in unary. It is often convenient to define a parameterized problem when the computational problem of interest is hard, but only for certain inputs whose difficulty can be measured by the parameter. For example, the LONGEST PATH problem is NP-complete, but if one wants to find a longest path in a graph with treewidth at most k , there is a $2^{O(k \log k)} \text{poly}(n)$ algorithm [FL89]. Parameterized complexity is inherently tied to the study of exact algorithms for NP. A problem is *fixed-parameter tractable* (FPT) when it has an algorithm with $f(k)n^c$ runtime for some function f and a universal c .

If a problem does not seem to be FPT, one tries to find evidence for this, using reducibility notions similar to NP-completeness. There is a hierarchy of problems, all technically solvable in polynomial time when the parameter k is constant, which do not seem to be FPT. (In fact many hierarchies have been proposed; cf. the text of Flum and Grohe [FG06].) The most basic is the W -hierarchy, where the lowest level is the class FPT of problems which are FPT. The next level $W[1]$ already contains problems which are not in FPT , assuming ETH [DF99, FG06]. There is a rough analogy between the W -hierarchy and the polynomial time hierarchy: in the W -hierarchy, $f(k) \log n$ bits are guessed in each alternation (instead of $\text{poly}(n)$), and the deterministic verifier only reads some $f(k)$ bits of the input. FPT equals the entire W -hierarchy if and only if Circuit SAT (with n variables and m gates) is in $2^{n/s(n)} \text{poly}(m)$ time for some unbounded $s(n)$ [ADR95]. (Note this runtime is much faster than those in our hypotheses.) Indeed, there is a sense in which the aforementioned subexponential time theory is isomorphic to questions in parameterized complexity [CG07].

Limited Nondeterminism. Kintala and Fisher [KF77] initiated the study of classes with limited nondeterminism. Buss and Goldsmith [BG93] and Bloch, Buss, and Goldsmith [BBG98] considered classes with only logarithmic nondeterminism, as we study here. A related but more refined model (the guess-and-check model) was introduced by Cai and Chen [CC97]. Flum, Grohe, and Weyer [FGW06] gave a framework connecting limited nondeterminism with parameterized complexity.

Better Algorithms And Better Lower Bounds. The idea that better algorithms can lead to better lower bounds is a recurring theme in theoretical computer science. Here we cite a few examples which seem to be most relevant to our work.

- Work of F. Zane and his co-authors from the late 90’s [Zan98] alternated between finding faster SAT algorithms and proving new exponential lower bounds for depth-three circuits, both accomplished using new structure theorems for CNF formulas.
- In the DFA INTERSECTION problem, one is given k DFAs with at most n states each, and the task is to decide if all DFAs accept a common string. Karakostas, Lipton, and Viglas [KLV03] prove if DFA INTERSECTION is solvable in $n^{o(k)}$ time, then $\text{NTIME}[n] \subseteq \text{TIME}[2^{\varepsilon n}]$ for all $\varepsilon > 0$ and $\text{LOGSPACE} \neq \text{NP}$. However DFA INTERSECTION is PSPACE-complete [Koz77], while the problems we consider are either in BPP or NP, and we do not need to assume subexponential algorithms. Lipton has also advocated this type of approach for proving lower bounds in his popular blog [Lip09b].
- Impagliazzo, Kabanets, and Wigderson [IKW02] showed that if CAPP is solvable in nondeterministic $2^{n^{o(1)}}$ time infinitely often, then $\text{NEXP} \not\subseteq \text{P/poly}$. Kabanets and Impagliazzo [KI04] showed that a nondeterministic $2^{n^{o(1)}}$ time algorithm for *Polynomial Identity Testing* implies that either $\text{NEXP} \not\subseteq \text{P/poly}$ or the Permanent does not have arithmetic circuits of polynomial size. Our results appear to be incomparable: while identity testing looks easier than the problems we study (even CAPP), our running time assumptions are much weaker and there is no “or” in our conclusion.

If improved algorithms of the kind needed in this paper exist, their discovery probably will not be hindered by the existing barriers in complexity theory (but perhaps another unforeseen one). In Appendix C we briefly discuss this.

3 Improved Algorithms Imply Circuit Lower Bounds

We start by proving consequences of better CIRCUIT SAT algorithms. The notion of “universal witness circuits” shall be useful. A language L has universal witness circuits if every correct verifier for L has polynomial size circuits that encode a witness accepted by the verifier, on every input in L .

Definition 3.1 *Let $L \in \text{NTIME}[t(n)]$. A polynomial time algorithm V is a verifier for L if $x \in L \iff (\exists y : |y| \leq t(|x|))[V(x, y) = 1]$.*

Definition 3.2 *A language $L \in \text{NTIME}[t(n)]$ has universal witness circuits if for all polynomial time verifiers V for L , there is a k and an $O(n^k)$ -size Boolean circuit family $\{C_n\}$ with the property:*

$x \in L \iff V(x, w(x)) = 1$, where $w(x)$ is the concatenation of the outputs of $C_\ell(\langle x, z \rangle)$ evaluated over all $z \in \{0, 1\}^{\lceil \log_2 t(n) \rceil + 1}$ in lexicographical order, and $\ell \geq n$ is an integer of appropriate length.

The universal witness property may look strong, but all of NEXP has universal witness circuits if $\text{NEXP} \subseteq \text{P/poly}$. The key to this observation is that if a language does not have such circuits, then *some* correct verifier for NEXP accepts witnesses that *cannot* be encoded with small circuits, infinitely often. This verifier can be used to test strings for high circuit complexity, which is enough to obtain pseudorandom generators, leading to a contradiction. The following can be easily obtained from Impagliazzo, Kabanets, and Wigderson [IKW02]; we include a proof for completeness in Appendix A.

Lemma 3.1 (Follows from [IKW02]) *If $\text{NEXP} \subseteq \text{P/poly}$ then every language in NEXP has universal witness circuits.*

We also need a strong completeness result for satisfiability.

Theorem 3.1 (Tourelakis [Tou01], Fortnow *et al.* [FLvMV05]) *There is a fixed d such that for every $L \in \text{NTIME}[n]$, L reduces to 3SAT in $O(n(\log n)^d)$ time. Moreover there is an algorithm (with random access to its input) that, given an instance of L and an integer $i \in [cn(\log n)^d]$ in binary (for some c depending on L), outputs the i th clause of the resulting 3SAT formula in $O((\log n)^d)$ time.*

The value of d depends on the particular computational model chosen. For most models one can take d to be small, e.g. $d = 4$. Theorem 3.1 holds under any computational model that accesses up to $O(\text{poly}(\log n))$ bits per step. Most proofs of it essentially rely the fact that nondeterministic Turing machines can simulate most other nondeterministic models with polylog overhead, as well as the Hennie-Stearns two-tape simulation of multitape Turing machines (cf. [AB09], Section 1.7) which introduces another log overhead. The Hennie-Stearns simulation can be converted (with constant overhead) into a circuit [PF79], which can then be efficiently converted to 3-CNF using the Tseitin transformation [Tse68]. See Van Melkebeek ([vM07], Lemma 2.2) for an alternative proof based on sorting networks. By standard translation/padding arguments (substituting 2^n in place of n in the above), the theorem can be scaled up to exponential time bounds:

Corollary 3.1 *Every language L in $\text{NTIME}[2^n]$ can be reduced to 3SAT instances of $c2^n \cdot n^4$ size. Moreover there is an algorithm that, given an instance of L and an integer $i \in [c2^n \cdot n^4]$ in binary, outputs the i th clause of the resulting 3SAT formula in $O(n^4)$ time.*

We are now ready to prove the main result of this section: a generic theorem relating the solvability of CIRCUIT SAT to lower bounds for NEXP. We say that a function f is *superpolynomial* in n if for all k , $n^k = o(f(n))$.

Theorem 3.2 (Faster Circuit SAT Implies Circuit Lower Bounds) *Let $c \geq 1$. Let $a(n)$ be a monotone increasing and unbounded function. Let $S(n)$ and $T(n)$ be functions such that*

- $T(n)/(S(n) + n^8)^c \geq \Omega(n^4 \cdot a(n))$ and
- $n \leq S(n) \leq O(2^n/n \cdot 1/a(n))$.

Suppose CIRCUIT SAT on n variables and m gates can be solved in $O(2^n m^c / T(n))$ co-nondeterministic time. Then $\text{NTIME}[2^n]$ does not have circuits of size $S(n)$.

Proof. Suppose $\text{NTIME}[2^n]$ has $S(n)$ -size circuits. We show that a faster (co-non)deterministic algorithm for Circuit Satisfiability implies a contradiction to the nondeterministic time hierarchy.

Let $L \in \text{NTIME}[2^n]$ be arbitrary. By Corollary 3.1, L can be reduced to 3SAT instances of $c2^n \cdot n^4$ size, for some c . Let $V(x, y)$ be a verifier for L that reduces x to a 3-CNF formula ϕ_x of $c2^n \cdot n^4$ size, substitutes the i th bit of y for the i th variable of ϕ_x , then returns 1 if and only if the resulting formula evaluates to 1.

By Lemma 3.1, L has universal witness circuits. Therefore for all $x \in L$, there is some y of length at most $c2^n \cdot n^4$ such that $V(x, y) = 1$, and y can be encoded with a $S(|x|)$ size circuit. That is, for all x , there is a circuit C_y that takes inputs of length $\ell = \log(c2^n |x|^4)$ and has size at most $S(|x|)$, such that the witness y equals the concatenation of $C_y(z)$ over all $z \in \{0, 1\}^\ell$ in lexicographical order.

Consider the following nondeterministic algorithm N for L . On input x , existentially guess the circuit C_y , using $O(S(|x|) \log S(|x|)) \leq O(2^n/a(n))$ bits. Then construct a circuit D with ℓ input variables X , as follows. Given an integer $i \in [c2^n \cdot n^4]$, the i th clause of ϕ_x can be computed in $O(n^4)$ time (via Corollary 3.1). By the standard translation of algorithms into circuits, it follows that the i th clause can be computed with an $O(n^8)$ size circuit; call it E . Lead the input variables X of D into the inputs of E , whose $3n + O(\log n)$ output wires encode the X th clause of ϕ_x . These output wires encode the indices of three variables in ϕ_x , along with three “negation bits” indicating 1 for variables which are negated, if any. For convenience, call the variable indices z_1, z_2, z_3 . Evaluate $a_1 = C_y(z_1)$, $a_2 = C_y(z_2)$, and $a_3 = C_y(z_3)$. Letting b_1, b_2, b_3 be the negation bits of z_1, z_2, z_3 (respectively), output $\neg[(a_1 \oplus b_1) \vee (a_2 \oplus b_2) \vee (a_3 \oplus b_3)]$. That is, $D(X)$ outputs 1 if and only if the X th clause is not satisfied, i.e., D is unsatisfiable if and only if C_y encodes a satisfying assignment for ϕ_x . The circuit D has $O(n^8 + S(|x|))$ size and ℓ input variables.

Finally, N calls a fast algorithm for circuit satisfiability on D , and *accepts* if and only if D is unsatisfiable. N runs in time

$$O(2^n/a(n) + 2^\ell \cdot (n^8 + S(n))^c/T(n)) \leq O(2^n/a(n) + 2^n n^4 \cdot (n^8 + S(n))^c/T(n)).$$

By assumption on $T(n)$ and $S(n)$, this time bound is at most $O(2^n/a(n))$.

Recall that L was an arbitrary language from $\text{NTIME}[2^n]$, so we now have

$$\text{NTIME}[2^n] \subseteq \text{NTIME}[2^n/a(n)].$$

Since $2^{n+1}/a(n+1) = o(2^n)$, we have a contradiction to the strong nondeterministic time hierarchy theorem [SFM78, Zak83]. \square

The following are immediately implied:

Reminder of Theorem 1.1 *Suppose there is a superpolynomial function $s(n)$ such that CIRCUIT SAT on circuits with n variables and n^k gates can be solved in $2^n \cdot \text{poly}(n^k)/s(n)$ time by a (co-non)deterministic algorithm, for all k . Then $\text{NEXP} \not\subseteq \text{P/poly}$.*

Proof. Let $S(n) = n^k$ and $T(n) = s(n)$ in Theorem 3.2. Then a $2^n \cdot \text{poly}(n^k)/s(n)$ time algorithm implies that $\text{NTIME}[2^n]$ does not have n^k size circuits. Since k can be arbitrary, $\text{NTIME}[2^n]$ does not have polynomial size circuits, hence NEXP also does not. \square

Reminder of Theorem 1.2 *If CIRCUIT SAT on n variables and m gates is in $O(2^{(1-\delta)n} m^c)$ (co-non)deterministic time for some $\delta > 0$, then there is $\varepsilon > 0$ and a language in $\text{NTIME}[2^n]$ that does not have $2^{\varepsilon n}$ size circuits.*

Proof. Let $T(n) = 2^{\delta n}$ and $S(n) = 2^{\delta n/c}/n^5$. \square

There are also interesting consequences for finding small algorithmic improvements for FORMULA SAT, the problem of satisfying Boolean formulas. Unfortunately the derandomization results do not seem to apply to restricted circuits, so we need another way to obtain universal witness circuits for NEXP . This can be accomplished with the following lemma:

Lemma 3.2 (Folklore) *Let \mathcal{C} be any class of circuits. If E^{NP} has (non-uniform) circuits from class \mathcal{C} , then $NTIME[2^n]$ has universal witness circuits from class \mathcal{C} .*

Proof. Let L be a language in $NTIME[2^n]$. Let $V(x, y)$ be a nondeterministic verifier for L running in $d2^{|x|}$ time. Consider the following E^{NP} machine:

$N(x, i)$: *Binary search for the lexicographically smallest z such that $V(x, z)$ accepts, by querying: given (x, y) where $|y| = d2^{|x|}$, is there $z \leq y$ such that $V(x, z)$ accepts? Then output the i th bit of z .*

Note the queries can be computed in NP, and N needs at most $d2^n$ queries to the oracle. Since every such N has polysize circuits from class \mathcal{C} , $NTIME[2^n]$ has universal witness circuits from \mathcal{C} . \square

We also need another complete problem.

Theorem 3.3 ([BGHSV05], Section 4) *There is a constraint graph problem Π such that every language in $NTIME[2^n]$ can be reduced to graph instances Π of size $\ell = O(2^n \text{poly}(n))$, such that the graph can be succinctly represented by a $\text{poly}(n)$ size formula on $n + O(\log n)$ variables. In particular, on an input i the formula outputs the i th edge of the graph along with the relevant constraints for that edge.*

The proof goes through nearly the same process as the completeness theorem for SAT (Theorem 3.1).⁹ First use the fact that every $NTIME[2^n]$ language can be accepted by a nondeterministic multitape Turing machine in $O(2^n \text{poly}(n))$ time. (The polynomial factor depends on the computational model.) Then make this TM “oblivious” in that its head movements have extremely regular structure, running in $O(2^n \text{poly}(n)^2)$ time. Now [BGHSV05] show this TM computation can be represented by a (deBruijn) graph of $O(2^n \text{poly}(n)^4)$ size, where the edge function can be described by a $\text{poly}(n)$ size formula. The latter step uses the regular structure of the simulation. Carrying out the same argument as Theorem 3.2, and applying Lemma 3.2:

Theorem 3.4 *Let $s(n)$ be superpolynomial. If FORMULA SAT on n variables and n^c connectives can be solved in $2^n \cdot \text{poly}(n^c)/s(n)$ time, then E^{NP} does not have (non-uniform) polynomial size formulas.*

Similarly, let \mathcal{C} be any non-uniform circuit class that contains AC^0 .

Theorem 3.5 *If satisfiability of \mathcal{C} -circuits on n variables and n^c gates can be solved in $2^{n/2} \cdot \text{poly}(n^c)/s(n)$ time, then E^{NP} does not have polynomial size \mathcal{C} -circuits.*

Theorem 3.5 follows by replacing the reductions of Theorem 3.1 and Theorem 3.3 with the “textbook proof” of the Cook-Levin theorem (cf. [Sip05], Theorem 7.30), which reduces $NTIME[2^n]$ languages to 3SAT instances on $O(4^n)$ clauses and variables by constructing a tableau of $O(4^n)$ cells. Any clause of the resulting formula can be generated by an AC^0 circuit with $2n + O(1)$ input variables and $O(n)$ size.

Extending our approach to conclude EXP lower bounds is an interesting open problem. In order to get a lower bound for EXP, it seems that we need a much stronger hypothesis. The following result can be easily shown using the almost everywhere hierarchy for deterministic time [GHS91].

Proposition 1 *If CIRCUIT SAT (or even 3SAT) is in 2^{n^ε} time for all $\varepsilon > 0$, then $EXP \not\subseteq i.o. - P/\text{poly}$.*

⁹In fact, the reduction of Theorem 3.3 can probably be carried out with SAT, although we have not yet verified this.

3.1 Extremely Weak Derandomization Implies Circuit Lower Bounds

We now turn to the *Circuit Acceptance Probability Problem* (CAPP): given a Boolean circuit C with n input variables and $\text{poly}(n)$ gates, the goal is to compute an approximation to the number of assignments satisfying C , within a factor of $1/6$. More precisely, we wish to output a number v such that

$$\left| v - \frac{1}{2^n} \cdot \sum_{x \in \{0,1\}^n} C(x) \right| < 1/6.$$

CAPP has been studied extensively [KC99, KRC00, For01, Bar02, IKW02]. We prove:

Reminder of Theorem 1.3 *If there is an $O(2^n \cdot \text{poly}(n)/s(n))$ nondeterministic algorithm for CAPP (for any superpolynomial $s(n)$), then $\text{NEXP} \not\subseteq \text{P/poly}$.*

In the previous section, we saw that strong reductions to 3SAT imply an efficient NEXP verifier. The verifier universally tries all possible clauses in an exponentially long 3-CNF formula, and checks the values of three variables in each trial. This universal quantifier is replaced with a $\text{poly}(n)$ size circuit which (on a variable assignment) checks the clause indexed by the input to the circuit. Our idea is to replace this universal quantifier with a random choice, so that instead of testing satisfiability of the circuit, it suffices to approximate the number of solutions. Naturally, this suggests the use of *probabilistically checkable proofs* (PCPs). We use the *PCPs of Proximity* of Ben-Sasson, Goldreich, Harsha, Sudan, and Vadhan, which imply PCPs for NEXP with nice properties.¹⁰

Theorem 3.6 ([BGHSV05]) *Let $T : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be a non-decreasing function. Then for every $s > 0$ and every language $L \in \text{NTIME}[T(n)]$ there exists a PCP verifier $V(x, y)$ with soundness s , perfect completeness, randomness complexity $r = \log_2 T(|x|) + O(\log \log T(|x|))$, query complexity $q = \text{poly}(\log T(|x|))$, and verification time $t = \text{poly}(|x|, \log T)$. More precisely:*

- V has random access to x and y , uses at most r random bits in any execution, makes q queries to the candidate proof y , and runs in at most t steps.
- If $x \in L$ then there is a y of length $T(|x|)\text{poly}(\log T(|x|))$ such that $\Pr[V(x, y) \text{ accepts}] = 1$.
- If $x \notin L$ then for all y , $\Pr[V(x, y) \text{ accepts}] \leq s$.

We shall be interested in the case $T(n) = 2^n$. Then the above PCP verifier uses $n + O(\log n)$ bits of randomness, $\text{poly}(n)$ verification time, and $\text{poly}(n)$ query complexity.

Proof of Theorem 1.3. The proof proceeds similarly to Theorem 1.1. We start by assuming $\text{NEXP} \subseteq \text{P/poly}$, so that all languages in NEXP have universal witness circuits. Let L be a language in $\text{NTIME}[2^n] - \text{NTIME}[2^n \cdot \text{poly}(n)/s(n)]$.

Let $V(x, y)$ be the PCP verifier of Theorem 3.6 for L where $T(n) = 2^n$ and $s = 1/2$. For some $k > 0$, V tosses $n + k \log n$ coins, queries the $O(2^n)$ length string y in $O(n^k)$ places, and runs in $O(n^k)$ time. Let $c > 0$ be such that the minimum circuit complexity of any function encoding a $2^n n^k$ -bit witness of V on inputs of length n is at most $n^c + c$.

¹⁰Note that Or Meir has recently found a more combinatorial approach to these results, cf. [Mei09].

We now describe an alternative nondeterministic algorithm N for L . On input x of length n , existentially guess a circuit C with $n^c + c$ size. Construct a circuit D that simulates V as follows. The circuit D has $n + k \log n$ input wires r , corresponding to the random bits of V . Once r is fixed, the verifier V runs in $O(n^k)$ time and queries y in $O(n^k)$ positions. By the standard translation of oracle machines to circuits, V can be simulated by a circuit of size $O(n^{2k})$ with *oracle gates* for the string y . These oracle gates each have $n + O(1)$ input wires and output the bit of y indexed by the input wires. Replace each oracle gate with a copy of C . The resulting circuit D has $O(n^{2k+c})$ size. Run the presumed nondeterministic algorithm for CAPP on the D , and *accept* if and only if the fraction returned is greater than $3/4$.

Let y be the string obtained by evaluating C on all inputs in lexicographical order. By construction, $D(z) = 1$ if and only if $V(x, y)$ outputs 1 on the string of coin tosses z .

If there is a witness y for the input x , then there is a C of size $O(n^c)$ which encodes y . On such a C , the fraction of inputs accepted by D is 1. On the other hand, if there is no witness y for x , then certainly no circuit C encodes a witness for x , and hence on every circuit C , the fraction of inputs accepted by D is at most $1/2$. Hence the algorithm N is correct.

Assuming CAPP can be solved in $2^n \cdot \text{poly}(n)/s(n)$ time for some superpolynomial $s(n)$, the running time of the entire simulation is $2^n \cdot \text{poly}(n)/s(n)$. This contradicts the choice of L . \square

The above proof can be generalized along the lines of Theorem 3.2 in a straightforward way. A completely analogous argument to Theorem 1.2 shows the following.

Theorem 3.7 *If there is an $O(2^{(1-\delta)n} \cdot \text{poly}(m))$ nondeterministic algorithm for CAPP on circuits of n variables and m gates, then $\text{NTIME}[2^n]$ does not have $2^{\delta n/c}/n^5$ size circuits.*

Impagliazzo, Kabanets, Wigderson [IKW02] showed partial converses to Theorem 1.3 and Theorem 3.7:

- If $\text{NEXP} \not\subseteq \text{P}/\text{poly}$, then CAPP has a nondeterministic algorithm that works on infinitely many inputs in $O(2^{n^\varepsilon})$ time with n^ε advice, for every $\varepsilon > 0$.
- If $\text{NTIME}[2^{O(n)}] \not\subseteq \text{SIZE}[2^{\sigma(n)}]$ then CAPP has a nondeterministic polytime algorithm that works on infinitely many inputs with $O(\log n)$ advice.

Combining these results with the above, we have the following ‘‘amplifications’’ of CAPP algorithms:

Reminder of Theorem 1.4 *If CAPP has a nondeterministic $O(2^n \cdot \text{poly}(n)/s(n))$ time algorithm that succeeds on all inputs, then for all $\varepsilon > 0$, CAPP (in fact, all of MA) has a nondeterministic $O(2^{n^\varepsilon})$ time algorithm with n^ε advice that succeeds on infinitely many inputs.*

Reminder of Theorem 1.5 *If CAPP on n variables and m gates has a nondeterministic $O(2^{(1-\delta)n} \cdot \text{poly}(m))$ time algorithm that succeeds on all inputs, then CAPP has a nondeterministic polynomial time algorithm with $O(\log n)$ advice that succeeds on infinitely many inputs.*

The above can also be viewed as gap theorems: if CAPP can’t be solved infinitely often in NP with logarithmic advice, then CAPP can’t be in $2^{.99n} \cdot \text{poly}(m)$ nondeterministic time.

Evidence for Derandomization of CAPP(?) The hypotheses of Theorem 1.3 looks potentially provable to us. Van Melkebeek and Santhanam [vMS05] have shown that for every probabilistic k -tape Turing machine running in time t , there is a $\delta > 0$ such that the machine can be simulated deterministically in $O(2^{(1-\delta)t})$ time. Unfortunately this is not quite enough to achieve the circuit lower bound for NEXP. To do that,

we would need (for example) that every probabilistic k -tape Turing machine with $b(n)$ bits of randomness running in time $t(n)$ can be simulated deterministically in $\text{poly}(t) \cdot 2^b/s(b)$ time, for a superpolynomial function s . In our desired application, $t(n)$ is a polynomial factor larger than $b(n)$.

4 Further Improvements Imply LOGSPACE \neq NP and Subexponential proofs for QBF

We now show that improving the runtime of exhaustive search for problems with limited nondeterminism would have surprising consequences. Our results use a tool from prior work on time-space lower bounds for satisfiability. Namely, we use a lemma to “speed up” arbitrary space-bounded computations with alternating machines. (The notation below is discussed in the Preliminaries.¹¹)

Lemma 4.1 (Speedup Lemma [Kan84, FLvMV05]) *Let $a \geq 2$ and $\log n \leq s(n) \leq o(n)$. Then*

$$\text{TISP}[n^a, s(n)] \subseteq (\exists n)(\forall \log n - \log s(n))\text{TISP}[n^{a-1} \cdot s(n), s(n)].$$

Moreover, the TISP part of the RHS only reads $n + O(s(n))$ bits of its input, and erases the rest.

Proof. Let M be an algorithm using n^a time and $s(n)$ space. Note that a configuration of M on an input of length n (a description of the entire state of the algorithm at any moment in time) can be represented within $\ell = dS(n)$ bits, for some $d > 1$.

To simulate M in $(\exists n)(\forall \log n - \log s(n))\text{TISP}[n^{a-1} \cdot s(n), s(n)]$, the machine $N(x)$ existentially guesses a sequence of configurations $C_1, \dots, C_{n/\ell}$ of $M(x)$, using n bits. Then $N(x)$ appends the initial configuration C_0 to the beginning of the sequence and the (unique) accepting configuration $C_{(n/\ell)+1}$ to the end of the sequence. Then $N(x)$ universally guesses $i \in \{0, \dots, n/\ell\}$ using $\log(n/\ell) + 1 \leq \log n - \log s(n)$ bits, erases all configurations except C_i and C_{i+1} , then simulates $M(x)$ starting from C_i , accepting if and only if the C_{i+1} is reached within $n^{a-1} \cdot \ell$ steps. It is easy to see the simulation is correct. \square

The key idea behind our results is to exploit the fact that the universal quantifier in the Speedup Lemma is only *logarithmic*. Any improvement over exhaustive search of all $\log n - \log s(n)$ bit strings gives us a very slight advantage, which can be amplified by repeated applications of the assumed algorithm. We arrive at the main theorem of this section.

Reminder of Theorem 1.6 *Suppose there is $\delta > 0$ such that for all $c, d \geq 1$, every problem Π solvable with $\log n$ nondeterministic bits, n^c time, and $(\log n)^d$ space can be solved by a deterministic algorithm in $O(n^{c+(1-\delta)})$ time and $\text{poly}(\log n)^d$ space. Then every such Π can also be solved with a nondeterministic algorithm in $O(n^3)$ time, i.e., $\text{SC} \subseteq \text{NTIME}[n^3]$.*

In fact, the proof shows that $(\exists n)\text{TISP}[n^k, \text{poly}(\log n)] \subseteq (\exists O(n))\text{TISP}[n^3, \text{poly}(\log n)]$, given the hypothesis. (Interestingly, this containment of classes is not known to be false, but as we shall see, it would be surprising if it were true.) The containment is proven by applying the Speedup Lemma along with the hypothesis for a constant number of times, with different values of c in each application.

Proof of Theorem 1.6. Written in our complexity class notation, the hypothesis is:

$$\forall c, d \geq 1, \exists k \geq 1 (\exists \log n)\text{TISP}[n^c, (\log n)^d] \subseteq \text{TISP}[n^{c+1-\delta}, (\log n)^{dk}]. \quad (1)$$

¹¹Readers familiar with the notation of our prior work [Wil08, Wil09] should beware that the notation in this paper has very slightly different meaning. In this work, we must pay attention to the constant factors in alternations, whereas in prior work we did not bother even with $n^{o(1)}$ factors. Our notation here reflects these differences.

By taking the complement of both sides, we also have $(\forall \log n) \text{TISP}[n^c, (\log n)^d] \subseteq \text{TISP}[n^{c+1-\delta}, (\log n)^{dk}]$. That is, there is always a deterministic algorithm which can determine if there is a witness iff there is always one that can determine if every string is a witness.

Assume (1). First we prove the following containment by induction on ℓ , for all integers $\ell > 0$, $k > 2$, and $d > 0$ satisfying $k - \delta\ell \geq 2$:

$$(\exists n) \text{TISP}[n^k, (\log n)^d] \subseteq (\exists \ell n) \text{TISP}[n^{k-\delta\ell} \cdot \text{poly}(\log n), \text{poly}(\log n)], \quad (2)$$

(Note the poly factors depend only on k , d , and ℓ .)

When $\ell = 1$, we have

$$(\exists n) \text{TISP}[n^k, (\log n)^d] \subseteq (\exists n) (\forall \log n - d \log \log n) \text{TISP}[n^{k-1}(\log n)^d, (\log n)^d] \quad (3)$$

by the Speedup Lemma (Lemma 4.1). Consider the class $(\forall \log n - d \log \log n) \text{TISP}[n^{k-1}(\log n)^d, (\log n)^d]$ on the RHS of (3). An algorithm from the class receives an input of length $2n$ (the original input and the n bits guessed). Applying (1) with $c = k - 1$ and taking the complement, the universal quantifier can be removed, concluding that

$$(\forall \log n - d \log \log n) \text{TISP}[n^{k-1}(\log n)^d, (\log n)^d] \subseteq \text{TISP}[n^{k-1+(1-\delta)} \cdot \text{poly}(\log n), \text{poly}(\log n)],$$

provided that $k \geq 2$. Substituting the RHS of the above into (3), we obtain

$$(\exists n) \text{TISP}[n^k, (\log n)^d] \subseteq (\exists n) \text{TISP}[n^{k-\delta} \cdot \text{poly}(\log n), \text{poly}(\log n)].$$

This completes the base case of (2). For the inductive step, consider the following chain of containments. First, the induction hypothesis says

$$(\exists n) \text{TISP}[n^k, (\log n)^d] \subseteq (\exists \ell n) \text{TISP}[n^{k-\delta\ell} \cdot \text{poly}(\log n), \text{poly}(\log n)].$$

Note the $\text{TISP}[n^{k-\delta\ell} \cdot \text{poly}(\log n), \text{poly}(\log n)]$ part on the RHS receives an input of length $(\ell+1)n \leq O(n)$. Applying the Speedup Lemma to this part, the above class is then contained in

$$(\exists \ell n) (\exists n) (\forall \log n - \log \log n) \text{TISP}[n^{k-\delta\ell-1} \text{poly}(\log n), \text{poly}(\log n)].$$

Observe that \exists quantifiers can be merged, resulting in the class

$$(\exists (\ell+1)n) (\forall \log n - \log \log n) \text{TISP}[n^{k-\delta\ell-1} \text{poly}(\log n), \text{poly}(\log n)].$$

Applying (1) with $c = k - \delta\ell - 1$ (this is possible when $k - \delta\ell - 1 \geq 1$), the above is in

$$(\exists (\ell+1)n) \text{TISP}[n^{k-\delta\ell-1+(1-\delta)} \text{poly}(\log n), \text{poly}(\log n)].$$

Finally, note that the exponent $k - \delta\ell - 1 + (1 - \delta) = k - \delta(\ell + 1)$. This completes the proof of (2).

Now let $k > 3$ be arbitrary. Since $\delta > 0$, the quantity $\delta\ell$ is positive and can be made as large as desired, provided that $k - \delta\ell \geq 2$.

Set $\ell = \lfloor (k - 2)/\delta \rfloor$. We have $k - \delta\ell \geq k - \delta(k - 2)/\delta = 2$, and

$$k - \delta\ell = k - \delta \lfloor (k - 2)/\delta \rfloor < k - \delta((k - 2)/\delta - 1) = 2 + \delta < 3.$$

That is,

$$(\exists n)\text{TISP}[n^k, (\log n)^d] \subseteq (\exists \ell n)\text{TISP}[n^{k-\delta\ell} \cdot \text{poly}(\log n), \text{poly}(\log n)]$$

by (2), and

$$(\exists \ell n)\text{TISP}[n^{k-\delta\ell} \cdot \text{poly}(\log n), \text{poly}(\log n)] \subseteq (\exists \ell n)\text{TISP}[n^3, \text{poly}(\log n)] \subseteq \text{NTIME}[n^3],$$

by our choice of ℓ . □

Reminder of Corollary 1.1 *The hypothesis of Theorem 1.6 implies LOGSPACE \neq NP.*

Proof. If LOGSPACE = NP then there is a c such that $\text{NTIME}[n] \subseteq \text{TISP}[n^c, \log n]$. We have $\text{NTIME}[n^4] \subseteq \text{TISP}[n^{4c}, \log n]$ by the translation lemma (Lemma 2.1). By Theorem 1.6, $\text{TISP}[n^{4c}, \log n] \subseteq \text{NTIME}[n^3]$. This is a contradiction to the nondeterministic time hierarchy. □

Reminder of Corollary 1.2 *The hypothesis of Theorem 1.6 implies that the quantified Boolean formula problem has a proof system where every QBF of length n has proofs of $2^{\varepsilon n}$ length for all $\varepsilon > 0$.*

Proof. Theorem 1.6 says that for all k , $\text{TISP}[n^k, O(\log n)] \subseteq \text{NTIME}[n^3]$.

Let $\varepsilon > 0$, and let $f(n) = 2^{\varepsilon n/3}$. By the translation lemma (Lemma 2.1), $\text{TISP}[f(n)^k, O(\log f(n))] \subseteq \text{NTIME}[f(n)^3]$, which is $\text{TISP}[2^{\varepsilon kn/3}, O(n)] \subseteq \text{NTIME}[2^{\varepsilon n}]$. Since k can be arbitrarily large, it follows that $\text{SPACE}[O(n)] \subseteq \text{NTIME}[2^{\varepsilon n}]$. A quantified Boolean formula of length n can be easily solved using $O(n)$ space. Therefore quantified Boolean formulas of length n can be solved in $2^{\varepsilon n}$ time with a nondeterministic algorithm. The conclusion follows. □

5 Unconditional Lower Bounds

The ideas of the previous section lead to an *unconditional* lower bound.

Reminder of Theorem 1.7 *For some a and some $k(n) \leq n$, there is a problem Π solvable with $k(n)$ nondeterminism and $O(n^a)$ time that cannot be solved in $k(n)^c n^a \cdot \text{poly}(\log n)$ time and $k(n)^c \cdot \text{poly}(\log n)$ space, for all constants c .*

Proof. Assume the opposite: for all $a \geq 1$ and for every problem Π solvable with $k(n)$ nondeterminism and $O(n^a)$ time, there is a c_a such that Π can be solved in $k(n)^{c_a} n^a \cdot \text{poly}(\log n)$ time and $k(n)^{c_a} \cdot \text{poly}(\log n)$ space. For each $a \geq 1$, define

$$\Pi_a = \{(i, x) \mid \exists y \in \{0, 1\}^{\log |x|+1} \text{ machine } M_i(x, y) \text{ accepts within } |x|^a + a \text{ steps}\}.$$

Setting $k(n) = \log n$, we have by assumption that

$$\Pi_a \in \text{TISP}[n^a (\log n)^{c_a} \text{poly}(\log n), (\log n)^{c_a} \text{poly}(\log n) \text{poly}(\log n)].$$

By efficiently reducing all languages in $(\exists \log n)\text{TIME}[n^a]$ to Π_a , we have for all $a \geq 1$ that

$$\begin{aligned} (\exists \log n)\text{TIME}[n^a] &\subseteq \text{TISP}[n^a \cdot \text{poly}(\log n), \text{poly}(\log n)] \\ (\forall \log n)\text{TIME}[n^a] &\subseteq \text{TISP}[n^a \cdot \text{poly}(\log n), \text{poly}(\log n)]. \end{aligned} \tag{4}$$

Setting $k(n) = n$ and $a = 2$, we have a polynomial time algorithm for SAT. Let ℓ be an integer such that SAT is in $O(n^\ell)$ time. By Theorem 3.1, we have

$$\text{NTIME}[n] \subseteq \text{TIME}[n^\ell \text{poly}(\log n)]. \tag{5}$$

Similar to the proof of Theorem 1.6, we derive

$$\begin{aligned}
\text{TIME}[n^{\ell+1}] &\subseteq \text{TISP}[n^{\ell+1} \cdot \text{poly}(\log n), \text{poly}(\log n)] \quad (\text{applying (4) with } a = \ell + 1) \\
&\subseteq (\exists n)(\forall \log n) \text{TISP}[n^\ell \cdot \text{poly}(\log n), \text{poly}(\log n)] \quad (\text{Speedup Lemma}) \\
&\subseteq (\exists n) \text{TISP}[n^\ell \cdot \text{poly}(\log n), \text{poly}(\log n)] \quad (\text{applying (4)}) \\
&\subseteq (\exists 2n)(\forall \log n) \text{TISP}[n^{\ell-1} \cdot \text{poly}(\log n), \text{poly}(\log n)] \quad (\text{Speedup Lemma}) \\
&\subseteq (\exists 2n) \text{TISP}[n^{\ell-1} \cdot \text{poly}(\log n), \text{poly}(\log n)] \quad (\text{applying (4)}) \\
&\subseteq \dots \\
&\subseteq (\exists \ell n) \text{TISP}[n \cdot \text{poly}(\log n), \text{poly}(\log n)] \\
&\subseteq \text{NTIME}[n \cdot \text{poly}(\log n)] \quad (\text{trivial}) \\
&\subseteq \text{TIME}[n^\ell \cdot \text{poly}(\log n)], \quad (\text{applying (5)})
\end{aligned}$$

contradicting the deterministic time hierarchy. □

Corollary 5.1 *There is a parameterized problem Π solvable in $O(2^k n^a)$ time that cannot be solved in $k^c \cdot n^{a+o(1)}$ time and $k^c \cdot (\log n)^c$ space, for all c .*

In fact, a stronger statement holds: either polynomial nondeterminism cannot be simulated in polytime, or we have strong lower bounds on simulating log-nondeterminism in polylog space.

Theorem 5.1 *Either $P \neq NP$ or there is a problem solvable with $\log n$ bits of nondeterminism in $O(n^c)$ time that is not solvable in $O(n^{c+.99})$ time and $\text{poly}(\log n)$ space, for some $c \geq 1$.*

Proof. Assume the opposite, so that $P = NP$ and

$$(\exists \log n) \text{TIME}[n^c] \subseteq \text{TISP}[n^{c+.99}, \text{poly}(\log n)]. \quad (6)$$

for all c . Note that (6) implies the hypothesis of Theorem 1.6, so $\text{SC} \subseteq \text{NTIME}[n^3]$ by Theorem 1.6. Moreover, (6) implies that $P \subseteq \text{SC}$. Therefore $NP \subseteq \text{NTIME}[n^3]$, a contradiction. □

We cannot yet extend these lower bounds to problems like Circuit SAT on general computational models, due to the inefficiency of reductions from arbitrary languages to Circuit SAT. We could extend them to Circuit SAT on multitape Turing machines, but those lower bounds are easy: the opposite of the lower bound implies that Circuit Evaluation (the case where there are *no* input variables) is in $n \cdot \text{poly}(\log n)$ time and $\text{poly}(\log n)$ space, which is already known to be false. (In fact, on multitape Turing machines the set of palindromes requires nearly quadratic time in the polylog space setting.)

However we can extend the lower bound slightly to generic computational models, using the strong nondeterministic time hierarchy [Zak83] (Theorem 2.1).

Theorem 5.2 *For all $\varepsilon > 0$, CIRCUIT SAT on circuits with k input gates and m total gates cannot be solved in $m^{1+o(1)} k^{1-\varepsilon}$ time and $m^{o(1)}$ space.*

Recall we do not know how to prove that SAT can't be solved in $O(N^{2-\varepsilon})$ time and $N^{o(1)}$ space on general models (the particular case where $k = m$).

Proof. Let $\varepsilon > 0$ be arbitrarily small. By Theorem 2.1, there is a unary language L that is in $\text{NTIME}[n^2]$ but not in $\text{NTIME}[n^{2-\varepsilon+\varepsilon'}]$ for all $0 < \varepsilon' < \varepsilon$.

Assume that CIRCUIT SAT can always be solved in $k^{1-\varepsilon}m^{1+o(1)}$ time and $m^{o(1)}$ space. Letting $k = m = n^2$, we have that $L \in \text{NTIME}[n^2]$ implies $L \in \text{DTISP}[n^{4-2\varepsilon}, n^{o(1)}]$ by Theorem 3.1 (the strong completeness theorem for SAT).

We now describe a nondeterministic algorithm for L that runs in $n^{2-\varepsilon+o(1)}$ time (a contradiction). Let A be an algorithm for L running in $n^{4-2\varepsilon}$ time and $n^{o(1)}$ space. Start by rejecting if the input is not unary. Suppose the input is 1^n .

Similar to the Speedup Lemma (Lemma 4.1), nondeterministically guess a list of $n^{2-\varepsilon}$ configurations of $A(1^n)$ where the first configuration is the initial configuration and the last configuration is an accepting configuration. Build a circuit with $(2 - \varepsilon) \log n$ input variables, which “selects” an adjacent pair out of the $n^{2-\varepsilon}$ configurations. This pair is then fed to another circuit, which given a pair of configurations of $A(1^n)$, outputs 1 iff the first configuration does not lead to the second configuration within $n^{2-\varepsilon}$ steps of running A . To simulate A on the input 1^n , store the length of the input n in a batch of $O(\log n)$ wires W . Then on any simulated step, the circuit just needs to check if the current position of the input head (or, the current input register being read) is greater than the value in W . If it is, the input symbol is *blank*, otherwise it is 1. The circuit only has to carry the current configuration and these wires W throughout the simulation of $n^{2-\varepsilon}$ steps. Without loss of generality, we may assume that A simply iterates through its working storage cell by cell in a predictive way, adding a multiplicative $n^{o(1)}$ to the runtime. The upshot is that each step of $A(1^n)$ can be simulated by a circuit of $n^{o(1)}$ size. The circuit size overall is $n^{2-\varepsilon+o(1)}$ and has $(2 - \varepsilon) \log n$ inputs. Observe that the circuit is *unsatisfiable* if and only if the configuration list is part of a valid accepting computation history for $A(1^n)$.

Applying the assumed Circuit SAT algorithm to this circuit, we obtain a nondeterministic algorithm for L which guesses $n^{2-\varepsilon+o(1)}$ bits then runs deterministically in $n^{2-\varepsilon+o(1)} \cdot O(\log n)^c$ time and $n^{o(1)}$ space. That is, the nondeterministic algorithm for L runs in $n^{2-\varepsilon}$ time. This contradicts our choice of L . \square

6 Conclusion

We have seen that universal improvements over exhaustive search, even marginal ones, would have surprising consequences in complexity theory. This connection between improved exponential algorithms and superpolynomial lower bounds shows that two communities have been implicitly working towards similar goals. Our techniques prompt many questions, such as:

- Can we prove interesting non-uniform lower bounds assuming a $2^{.9n} \text{poly}(m)$ algorithm for CNF SAT on n variables and m clauses? Applying Theorem 3.5, one can show that if 3SAT is in $2^{\varepsilon n}$ time for all $\varepsilon > 0$, then E^{NP} does not have linear size circuits. That is, if the Exponential Time Hypothesis is false, then we obtain mild circuit lower bounds.
- If $\text{NEXP} \not\subseteq \text{P/poly}$, can this lower bound be used to obtain faster SAT algorithms? Note that we know how to partially derandomize CAPP assuming $\text{NEXP} \not\subseteq \text{P/poly}$ [IKW02].
- Can we replace CAPP with a form of *Polynomial Identity Testing* (PIT) in our results? The known proofs that “subexponential algorithms for PIT imply circuit lower bounds” go through Toda’s theorem [Tod91], which prevents us from getting a tight simulation.
- Is it possible to show (for example) that $\text{NTIME}[2^n] \not\subseteq i.o. - \text{NTIME}[2^{n/2}]$? This could be used to prove strong gap theorems for CAPP. The above separation could be used to show “if CAPP is

in $\text{NTIME}[2^{n/3}\text{poly}(n)]$ then $\text{NEXP} \not\subseteq i.o. - \text{P/poly}$,” which would in turn imply that CAPP is in $\text{NTIME}[2^{n^\epsilon}]$. Unfortunately the best known almost-everywhere time hierarchies for nondeterminism are somewhat weak [ABHH93].

- Can we prove unconditionally that there is a $k > 3$ such that

$$(\exists n)\text{TISP}[n^k, n^{o(1)}] \not\subseteq (\exists O(n))\text{TISP}[n^3, n^{o(1)}]?$$

Given our results, this separation would imply unconditional lower bounds for improving on exhaustive search with space-bounded verifiers. (Interestingly, such lower bounds may be achievable *without* proving class separations like $\text{LOGSPACE} \neq \text{NP}$.)

Acknowledgments

I am grateful to Mohan Paturi, Russell Impagliazzo, and Rahul Santhanam for useful discussions.

References

- [ADR95] K. Abrahamson, R. G. Downey, and M. R. Fellows. Fixed-parameter tractability and completeness IV: On completeness for W[P] and PSPACE analogs. *Annals of Pure and Applied Logic* 73:235–276, 1995.
- [ABHH93] E. Allender, R. Beigel, U. Hertrampf, and S. Homer. Almost everywhere complexity hierarchies for nondeterministic time. *Theor. Comp. Sci.*, 115(2):225–241, 1993.
- [AB09] S. Arora and B. Barak. *Computational Complexity – a modern approach*. Cambridge University Press, 2009.
- [BFNW93] L. Babai, L. Fortnow, N. Nisan, and A. Wigderson. BPP has subexponential simulations unless EXPTIME has publishable proofs. *Computational Complexity* 3:307–318, 1993.
- [Bar02] B. Barak. A Probabilistic-Time Hierarchy Theorem for Slightly Non-uniform Algorithms. In *Proc. of RANDOM*, 194–208, 2002.
- [BGHSV05] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. Vadhan. Succinct PCPs verifiable in polylogarithmic time. In *IEEE Conference on Computational Complexity*, 2005.
- [BBG98] S. Bloch, J. F. Buss, and J. Goldsmith. Sharply bounded alternation and quasilinear time. *Theory Comput. Syst.* 31(2):187–214, 1998.
- [BG93] J. F. Buss and J. Goldsmith. Nondeterminism within P. *SIAM J. Computing* 22:560–572, 1993.
- [CJ03] L. Cai and D. Juedes. On the existence of subexponential parameterized algorithms. *J. of Comp. and Syst. Sci.* 67(4):789–807, 2003.
- [CC97] L. Cai and J. Chen. On the amount of nondeterminism and the power of verifying. *SIAM J. Computing* 26(3):733–750, 1997.
- [CIP06] C. Calabro, R. Impagliazzo, and R. Paturi. A Duality between Clause Width and Clause Density for SAT. In *Proc. IEEE Conference on Computational Complexity*, 252–260, 2006.

- [CIP09] C. Calabro, R. Impagliazzo, and R. Paturi. The complexity of satisfiability of small depth circuits. In *Proc. International Workshop on Parameterized and Exact Computation*, 2009.
- [CG07] Y. Chen and M. Grohe. An isomorphism between subexponential and parameterized complexity theory. *SIAM Journal on Computing* 37:1228–1258, 2007.
- [DH08] E. Dantsin and E. A. Hirsch. Worst-Case Upper Bounds. In *Handbook of Satisfiability*, A. Biere, M. Heule, H. van Maaren and T. Walsh (eds.), 341–362, 2008.
- [DF99] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer-Verlag, 1999.
- [FL89] M. R. Fellows and M. A. Langston. On search, decision and the efficiency of polynomial-time algorithms. In *Proc. ACM Symposium on Theory of Computing*, 501–512, 1989.
- [FGW06] J. Flum, M. Grohe, and M. Weyer. Bounded fixed-parameter tractability and $\log^2 n$ nondeterministic bits. *J. Comput. Syst. Sci.* 72(1):34–71, 2006.
- [FG06] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
- [For01] L. Fortnow. Comparing notions of full derandomization. In *Proc. IEEE Conference on Computational Complexity*, 28–34, 2001.
- [FLvMV05] L. Fortnow, R. Lipton, D. van Melkebeek, and A. Viglas. Time-Space Lower Bounds for Satisfiability. *JACM* 52(6):835–865, 2005.
- [GO95] A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Computational Geometry: Theory and Applications* 5(3):165–185, 1995.
- [GHS91] J. G. Geske, D. T. Huynh, and J. I. Seiferas. A note on almost-everywhere complex sets and separating deterministic time complexity classes. *Information and Computation* 92(1):97–104, 1991.
- [GS89] Y. Gurevich and S. Shelah. Nearly linear time. *Logic at Botik’89*, Springer-Verlag LNCS 363, 1989.
- [Har89] J. Hartmanis. Gödel, von Neumann, and the $P \stackrel{?}{=} NP$ problem. *Bulletin of the European Association for Theoretical Computer Science* 101–107, June 1989. See also: M. Sipser. The history and status of the P versus NP question. In *Proceedings of ACM Symposium on Theory of Computing* 603–618, 1992.
- [Has98] J. Hastad. The Shrinkage Exponent of De Morgan Formulae is 2. *SIAM Journal on Computing*, 27:48–64, 1998.
- [IKW02] R. Impagliazzo, V. Kabanets, and A. Wigderson. In search of an easy witness: exponential time versus probabilistic polynomial time. *J. Comput. Syst. Sci.* 65(4):672–694, 2002.
- [IP01] R. Impagliazzo and R. Paturi. On the Complexity of k -SAT. *J. Comput. Syst. Sci.* 62(2):367–375, 2001.
- [IPZ01] R. Impagliazzo, R. Paturi, and F. Zane. Which Problems Have Strongly Exponential Complexity? *J. Comput. Syst. Sci.* 63(4):512–530, 2001.
- [KC99] V. Kabanets and J.-Y. Cai. Circuit minimization problem. In *Proc. ACM Symposium on Theory of Computing*, 73–79, 2000.

- [KRC00] V. Kabanets, C. Rackoff, and S. A. Cook. Efficiently approximable real-valued functions. *Electronic Colloquium on Computational Complexity*, TR00-034, 2000.
- [KI04] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity* 13(1-2):1–46, 2004.
- [Kan84] R. Kannan. Towards Separating Nondeterminism from Determinism. *Mathematical Systems Theory* 17(1):29–45, 1984.
- [KLV03] G. Karakostas, R. J. Lipton, and A. Viglas. On the complexity of intersecting finite state automata and NL versus NP. *Theor. Comput. Sci.* 302(1-3):257–274, 2003.
- [KF77] C. M. R. Kintala and P. C. Fisher. Computations with a Restricted Number of Nondeterministic Steps (Extended Abstract). In *ACM Symposium on Theory of Computing*, 178–185, 1977.
- [KvM99] A. Klivans and D. van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM Journal on Computing* 31(5):1501–1526, 2002.
- [Koz77] D. Kozen. Lower bounds for natural proof systems. *Proc. IEEE Symposium on Foundations of Computer Science*, 254–266, 1977.
- [Lip09a] R. J. Lipton. *Simulation of nondeterministic Machines*. Weblog post available at <http://rjlipton.wordpress.com/2009/05/11/simulation-of-nondeterministic-machines/>
- [Lip09b] R. J. Lipton. *An approach to the P=NP question?* Weblog post available at <http://rjlipton.wordpress.com/2009/10/02/an-approach-to-the-pnp-question/>
- [LV96] M. Luby and B. Velickovic. On Deterministic Approximation of DNF. *Algorithmica* 16(4/5):415–433, 1996.
- [vMS05] D. van Melkebeek and R. Santhanam. Holographic proofs and derandomization. *SIAM Journal on Computing*, 35:59–90, 2005.
- [vM07] D. van Melkebeek. A Survey of Lower Bounds for Satisfiability and Related Problems. *Foundations and Trends in Theoretical Computer Science*, 2007.
- [Mei09] O. Meir. Combinatorial PCPs with efficient verifiers. In *Proc. of IEEE Symposium on Foundations of Computer Science (FOCS)*, 2009.
- [PF79] N. Pippenger and M. J. Fischer. Relations Among Complexity Measures. *JACM* 26(2):361–381, 1979.
- [SFM78] J. Seiferas, M. J. Fischer, and A. Meyer. Separating Nondeterministic Time Complexity Classes. *JACM* 25:146–167, 1978.
- [Sip05] M. Sipser. Introduction to the Theory of Computation. *Course Technology*, 2005.
- [Tod91] S. Toda. PP is as Hard as the Polynomial-Time Hierarchy. *SIAM J. Computing* 20:865–877, 1991.
- [Tou01] I. Turlakis. Time-Space Tradeoffs for SAT on Nonuniform Machines. *J. Computer and System Sciences* 63(2):268–287, 2001.

- [Tra08] P. Traxler. The Time Complexity of Constraint Satisfaction. *International Workshop on Parameterized and Exact Computation*, 190–201, 2008.
- [Tse68] G. Tseitin. On the complexity of derivation in propositional calculus. *Studies in Constr. Math. and Math. Logic*, 1968.
- [Wil08] R. Williams. Time-space lower bounds for counting NP solutions modulo integers. *Computational Complexity* 17(2):179–219, 2008.
- [Wil09] R. Williams. Alternation-trading proofs, linear programming, and lower bounds. Manuscript, 2009.
- [Zak83] S. Zak. A Turing Machine Hierarchy. *Theor. Comput. Sci.* 26:327–333, 1983.
- [Zan98] F. Zane. Circuits, CNFs, and Satisfiability. Ph.D. Thesis, University of California at San Diego, 1998.

A Appendix: Proof of Lemma 3.1

Here we show that if $\text{NEXP} \subseteq \text{P/poly}$, then every language in NEXP has universal witness circuits. The argument is analogous to that used in the proof that $\text{NEXP} \in \text{P/poly} \implies \text{NEXP} = \text{MA}$ [IKW02].

Proof. First we show that the absence of witness-producing circuits for NEXP implies a faster nondeterministic simulation of MA . This is combined with the fact that $\text{NEXP} \subseteq \text{P/poly}$ implies an MA -simulation of NEXP , to yield a contradiction with a time hierarchy theorem.

Suppose there is an $L \in \text{NEXP}$ which does not have universal witness circuits. Let $c > 0$ be such that $L \in \text{NTIME}[2^{n^c}]$. Then there is *some* correct verifier V for L such that for all constants $d \geq 1$, there is an infinite sequence of inputs $S = \{x_{i_k}\}$ with the properties:

- for all k , $x_{i_k} \in L$ and
- for all sufficiently large k and all y of length $2^{|x_{i_k}|^c}$, $V(x_{i_k}, y) = 1$ implies that the circuit complexity of the function $f(x_{i_k}, i) = y_i$ is greater than $|x_{i_k}|^d$.

Work of Babai *et al.* [BFNW93] and Klivans-Van Melkebeek [KvM99] shows the following *hardness-randomness connection*: for every $\varepsilon > 0$ there is a $\delta < \varepsilon$ and integer e such that, given random access to a Boolean function on n^δ variables with circuit complexity at least $n^{\delta e}$, there is a pseudorandom generator $G : \{0, 1\}^{n^\varepsilon} \rightarrow \{0, 1\}^n$ computable in $2^{O(n^\varepsilon)}$ time which fools circuits of size n .

We can simulate MA infinitely often in nondeterministic time $O(2^n)$ with n bits of advice, as follows. Let P be an MA protocol, and suppose the computation of Arthur in P (given Merlin’s string) can be simulated by a circuit of size n^a . For each input length n , we set the advice for n -bit inputs to be the string $x_{i_\ell} \in S$, where x_{i_ℓ} has length $n^{\varepsilon a}$; if there is no such string, we set the advice to 0^n .

On an input x of length n with advice x_{i_k} , the MA simulation first nondeterministically guesses a witness y of length $2^{|x_{i_k}|^c}$ for the verifier V , and checks that $V(x_{i_k}, y) = 1$ (if this is not the case, the simulation rejects). Note by our choice of advice, this step takes $O(2^{n^{\varepsilon a c}})$ time. Next, we nondeterministically guess Merlin’s polynomial length string in the MA simulation. Finally, we simulate Arthur by evaluating G on all $n^{\varepsilon a}$ possible seeds (treating the string y as a hard function), evaluating the circuit for Arthur on the outputs of G , and taking the majority answer.

We claim the generator G fools Arthur. On the $n^{\varepsilon a}$ length input x_{i_k} , the string y can be treated as a Boolean function on $n^{\varepsilon a c}$ variables with circuit complexity at least $n^{\varepsilon a d}$. We can make $n^{\varepsilon a d} \geq n^{\varepsilon a \delta e}$ since we can set d to be arbitrarily large. Hence G armed with y can fool circuits of size $n^{\varepsilon a d / (\delta e)} \geq n^a$, by the hardness-randomness connection. That is, G fools circuits of size n^a , and therefore Arthur as well.

The total running time of the simulation is $O(2^{n^{\varepsilon a c}} + 2^{n^{\varepsilon a}})$. Setting $\varepsilon > 0$ to be arbitrarily small (note it is independent of c and a), we have established $\text{MA} \subseteq i.o. - \text{NTIME}[2^{n^{\varepsilon'}}]/n^{\varepsilon'}$ for all $\varepsilon' > 0$. Since $\text{NEXP} \subseteq \text{P/poly}$, there is a fixed constant q such that $\text{NTIME}[2^n]/n$ has circuits of size $O(n^q)$. So MA has infinitely many input lengths that can be computed by a circuit family of $O(n^q)$ size. (In the language of complexity classes, we have $\text{MA} \subseteq i.o. - \text{SIZE}[O(n^q)]$.)

Now, $\text{EXP} \subseteq \text{P/poly}$ implies that $\text{EXP} = \text{MA}$, by Babai, Fortnow, Nisan, and Wigderson [BFNW93]. Hence there is a $c' > 0$ such that

$$\text{EXP} = \text{MA} \subseteq i.o. - \text{NTIME}[2^n]/n \subseteq i.o. - \text{SIZE}[O(n^q)].$$

However, the above is false, by a simple diagonalization argument. □

B Appendix: Proof of Theorem 1.9

Reminder of Theorem 1.9 *For every multitape Turing machine M running in time t and accessing a tape of nondeterministic bits, there is a $\delta > 0$ and a deterministic Turing machine that can simulate M in $O(2^{(1-\delta)t})$ time.*

The idea is to perform a brute-force simulation of M up to $(1 - \delta)t$, for an appropriate $\delta > 0$. This generates $2^{(1-\delta)t}$ different configuration states of M . Observe that for a k -tape machine, there are only $2k\delta t$ different tape cells that could possibly be accessed in δt steps. Hence we can “compress” these configurations so that only $c^{2k\delta t}$ configurations remain for some $c > 0$, and the necessary information to complete the simulation is still preserved. Then we simulate the remaining configurations for δt steps to see if any accept. We can choose δ so that these procedures all take less than 2^t time.

Proof. Let $\delta > 0$ be a parameter. Let σ be the cardinality of the tape alphabet of M , and let k be the number of tapes of M . For all possible strings y of $(1 - \delta)t$ bits, simulate M for its first $(1 - \delta)t$ steps, assuming the nondeterministic tape has y written on it. For each y , save the *configuration* C_y of M after $(1 - \delta)t$ steps, where a configuration consists of all tape cells that are within δt cells of some tape head. For each of the k tapes, there are $2\delta t + 1$ such cells. Note that for the nondeterministic tape in C_y , some tape cells may be *undetermined*, as they refer to bits of the nondeterministic tape that have not yet been read by M . We denote those cells with a fresh symbol s that is not already in the alphabet of M .

Note the total number of such configurations is $(\sigma + 1)^{k(2\delta t + 1)}$. Hence in order to remember all possible configurations of M , it suffices to store a bit array A of only $(\sigma + 1)^{k(2\delta t + 1)}$ size.

Now for every configuration C marked in array A , and for all $2^{2\delta t}$ ways to fill in the undetermined symbols of the nondeterministic tape in C with $2\delta t$ bits, simulate M for δt steps. If any simulation of M results in acceptance, then accept.

To optimize the running time, we want to set δ to minimize $2^{(1-\delta)t} + 2^{2\delta t}(\sigma + 1)^{k(2\delta t + 1)}$. (Note that polynomial factors in t do not matter in the optimization, as they can be subsumed by simply adding an ε factor to the exponents.) Routine calculation shows that $\delta = 1/(3 + 2k \log(\sigma + 1))$ suffices. □

C A Remark About Barriers

Certainly, a slightly faster Circuit SAT or CAPP algorithm will need significantly new ideas. Let us remark why this strategy may be viable for proving lower bounds. Any potential approach to proving strong lower bounds must pass certain litmus tests that complexity theory has developed. The three primary barriers are relativization, natural proofs, and algebrization. We believe our work may provide a path that circumvents all three (provided the appropriate algorithms exist, of course). The natural proofs barrier does not apply due to our use of diagonalization, which avoids the “largeness” of natural proofs. We believe the relativization and algebrization barriers would be avoided, because *all nontrivial SAT algorithms that we know do not relativize or algebrize*. That is, the interesting SAT algorithms in the literature do not generalize to solve the SAT^A problem, when the algorithm is given oracle access to an arbitrary oracle A (or its algebraic extension \tilde{A}).¹² So it is not outrageous to think that we might design slightly better CIRCUIT SAT or CAPP algorithms with similar properties.

¹² SAT^A is the problem of satisfying CNF formulas with clauses of the form $(\ell_1 \vee \dots \vee \ell_k \vee A(\ell_{k+1}, \dots, \ell_{k+k'}))$, where ℓ_i are literals and the predicate $A(\ell_{k+1}, \dots, \ell_{k+k'})$ is true if and only if $\ell_{k+1} \dots \ell_{k+k'} \in A$. It is the natural variant on SAT obtained by applying the Cook-Levin theorem to a language in NP^A . We invite the reader to examine their favorite (nontrivial) SAT algorithm to understand why it does not extend to SAT^A . As a simple example, consider the branching algorithm which looks for clauses of length one, and sets the variable in that clause accordingly. (If there is no length-one clause, say it branches on an arbitrary variable.) In SAT^A , a length-one clause can have the form $A(\ell_{k+1}, \dots, \ell_{k+k'})$. Determining how to set the variables in this predicate requires at least an oracle for “Satisfiability of A .” If the A -predicate has more than one satisfying assignment, then at a stage where most variables have not been set, we cannot feasibly determine which is the correct assignment with an arbitrary A oracle or with \tilde{A} .