On the Complexity of Circuit Satisfiability (Extended Abstract)

Ramamohan Paturi [*]
University of California, San Diego
Email:paturi@cs.ucsd.edu
Pavel Pudlák
Mathematical Institute of the Czech Academy of Sciences
Email:pudlak@math.cas.cz

November 12, 2009

Abstract

In this paper, we are concerned with the exponential complexity of the *Circuit Satisfiability* (**CircuitSat**) problem and more generally with the exponential complexity of **NP**-complete problems. Over the past 15 years or so, researchers have obtained a number of exponential-time algorithms with improved running times for exactly solving a variety of **NP**-complete problems. The improvements are typically in the form of better exponents compared to exhaustive search. Our goal is to develop techniques to prove specific lower bounds on the exponents under plausible complexity assumptions. We consider natural, though restricted, algorithmic paradigms and prove lower bounds on the exponent of the success probability. Our approach has the advantage of clarifying the relative power of various algorithmic paradigms.

Our main technique is a a success probability amplification technique, called *the Exponential Amplification Lemma*, which shows that for any $f(n,m)$-size bounded probabilistic circuit family $\mathcal{A}$ that decides **CircuitSat** with success probability at least $2^{-\alpha n}$ for $\alpha < 1$ on inputs which are circuits of size m with n variables, there is another probabilistic circuit family $\mathcal{B}$ that decides **CircuitSat** with size roughly $f(\alpha n, f(m,n))$ and success probability about $2^{-\alpha^2 n}$. In contrast, the standard method for boosting success probability by repeated trials will improve it to $(1-(1-2^{-\alpha n})^t)$ ($\approx t2^{-\alpha n}$ for $t = O(2^{\alpha n})$) using circuits of size about $tf(n,m)$.

Using this lemma, we derive tight bounds on the exponent of the success probability for deciding the **CircuitSat** problem in a variety of probabilistic computational models under complexity assumptions. For example, we show that the success probability cannot be better than $2^{-n+o(n)}$ for deciding **CircuitSat** by probabilistic polynomial size circuits unless **CircuitSat** (thereby all of **NP**) on polynomial size instances can be decided by 2^{n^μ} size *deterministic* circuits for some $\mu < 1$, which is considered an unlikely event. As another example, we show that probabilistic quasilinear size circuits cannot achieve success probability better than $2^{-n+o(n)}$ unless **CircuitSat** (as well as **NP**) has $O(m^{\lg\lg m})$ size deterministic circuits, which is very close to the statement $\mathbf{NP} \in \mathbf{P/poly}$, a highly unlikely scenario.

[*] This research is supported by NSF grant CCF-0947262 from the Division of Computing and Communication Foundations. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

It is well-known that all **NP**-complete problems are equivalent as far as polynomial-time solvability is concerned. However, much less is known about the exact complexity of these problems. If we assume $\mathbf{NP} \neq \mathbf{P}$ or other appropriate complexity statement, what can we say about the exact worst-case complexity of **NP**-complete problems? An important context for this question is the development of a series of exact exponential-time algorithms with improved run times for a number of problems including IndependentSet, k-SAT, and k-colorability. The series of improvements are typically in the form of better exponents compared to exhaustive search. However, exact worst-case complexity of these problems seem to differ considerably. These improvements prompt several complexity questions, chief among them is whether we can expect continued improvements in the exponent. Is there a limit beyond which one should not expect improvement? How do these limits differ for different problems? Can we explain the differing limits in terms of the structural properties of the problems? What are the likely exact complexities of various **NP**-complete problems? Are the likely complexities of various problems related?

The current state of the art in complexity theory is far from being able to resolve these questions, especially the question of best exponents, even under reasonable complexity assumptions. We believe that it would be productive to approach these questions from the viewpoint of known algorithmic paradigms. Such an approach might be able to clarify the relative power of various algorithmic paradigms and might even be able to shed light on the best exponents in natural, though restricted, computational models under complexity assumptions. Furthermore, the study of the limitations of algorithmic paradigms might result in sharper versions of existing problems and suggest new directions of research. This paper presents an attempt in this direction and obtains nontrivial, interesting new results.

We propose to approach the randomized exact algorithms for **NP**-complete problems by studying the important subclass **OPP** of algorithms and its generalizations. **OPP** is the class of one-sided probabilistic polynomial-time algorithms. This class captures a common design paradigm for randomized exact exponential-time algorithms: to repeat sufficiently many times a one-sided error probabilistic polynomial-time algorithm that is correct with an exponentially small probability so that the overall algorithm finds a witness with constant probability. **OPP** includes Davis-Putnam-style backtracking algorithms developed in recent times to provide improved exponential-time upper bounds [BE95, Bei99, PPZ99, Epp01, DGH$^+$02, Epp03, Bys03, GHNR03, FGK06] for a variety of **NP**-hard problems. While the original versions of some of these algorithms are couched as exponential-time algorithms, one can observe from a formalization by Eppstein [Epp06] that these algorithms can be converted into probabilistic polynomial-time algorithms whose success probability is the reciprocal of the best exponential-time bound. **OPP** also includes local search algorithms such as Schöning's [Sch99]. **OPP** is interesting not just because of ubiquity, but because such algorithms are ideal from the point of view of space efficiency, parallelization, and speed-up by quantum computation. What are the limitations of such algorithms for deciding **NP**-complete problems? Could the best algorithm for a canonical **NP**-complete problem such as **CircuitSat** (the problem of deciding whether a circuit is satisfiable) be in **OPP**?

In addition to **OPP**, there are several other important algorithmic paradigms for designing exact algorithms for **NP**-hard problems, for example, exponential-time divide-and-conquer [Law76, Sak98], inclusion-exclusion/Möbius inversion [Kar82, Koi06, BH06, BKK07, BKK08], dynamic programming [HK61, Rob86], group algebra [Kou08, Wil09], and sieve algorithms [AKS01, MV09]. However, we argue that **OPP** and its generalizations could serve as an excellent starting point for the study of exponential-time algorithms for **NP**-complete problems in general.

Consider the problem of k-colorability for $k \geq 3$. The best-known algorithm [Koi06, BH06] for this problem applies the inclusion-exclusion principle to achieve an $\tilde{O}(2^n)$ algorithm where n in the number of vertices of the graph. This algorithm and the prior best-known algorithms [Sak98, Epp03, Bys05, BK06] do not belong to the class **OPP**. This raises a natural question whether we can expect an **OPP** algorithm for k-colorability whose success probability is at least 2^{-n}. Beyond this, can we expect **OPP**-style optimal algorithms for k-colorability? Does there exist any **OPP** algorithm for k-colorability whose success probability is at least c^{-n} where c is independent of k? Negative answers (or evidence to that effect) for these questions would provide convincing proof (or evidence) that exponential-time inclusion-exclusion and dynamic programming paradigms are strictly more powerful than that of **OPP**. On the other hand, algorithmic results that would place k-colorability in the class **OPP** with c^{-n} success probability would be exciting. Similar situation exists with respect to the Hamiltonian path problem where we know of no **OPP** algorithms that succeed with significantly better than $1/n!$ success probability whereas it is well-known that there are $O(n^2 2^n)$ algorithms [Bel62, Kar82] based on exponential-time dynamic programming and inclusion-exclusion techniques. It seems that resolution of the question, whether **OPP**-style optimal algorithms exist for k-colorability and Hamiltonian path is related to a fundamental issue regarding the trade-off between time and success probability which in turn may hold one of the keys (the others being **P** versus **NP** and derandomization) for the complexity theory of exact algorithms for **NP**-complete problems.

To study the exponents in the exact complexity of **NP**-complete problems, it is useful to parameterize **NP** problem instances with two parameters. Usually, **NP** problem instances are parameterized by the size of the input. However, for the purpose of capturing the exponential complexity, it is more natural to parameterize the instances in terms of an additional *complexity* parameter. For example, the **CircuitSat** problem instances are parameterized by the number of variables n and m, the length of the input, which is a description of a circuit. A graph problem such as k-colorability is parameterized by n, the number of vertices and m, the length of the representation of the graph. These parameters are natural and robust with respect to the representation of the input. We assume that problems are presented together with a complexity parameter. It is also useful to endow the class **NP** with a canonical complexity parameter so we can state properties of **NP** in terms of the complexity parameter and the size of the input. Consider an **NP** predicate in the form $\exists x \Phi(y, x)$ where y is the input instance, x the witness, and Φ is a polynomial-time (in the length of y) computable predicate. We canonically parameterize **NP** problem instances by n, the length of the witness as well as m, the length of the input. Focus on the length of the witness is natural when we consider exponential-time exact algorithms.

We sometimes use the notation $\mathbf{NP}(n, m)$ and $\mathbf{CircuitSat}(n, m)$ to be explicit about the parameterization. Complexity bounds are usually expressed as functions of both n and m. We would like to observe that any problem instance in $\mathbf{NP}(n, m)$ can be reduced to a $\mathbf{CircuitSat}(n, m')$ instance in time t preserving the witness length, where t and m' are polynomially bounded in m.

The results in the paper concern the more general class $\mathbf{OP}(T(n, m))$ of one-sided probabilistic algorithms that run in time $T(n, m)$. We usually want the success probability to be bounded by a small exponential function in n, independent of m. In addition to the case where $T(n, m)$ is polynomially bounded in m, we consider other natural cases. In fact, several of the branch-and-bound algorithms mentioned earlier run in quasilinear time where the success probability is inverse of the exponential run time. We will also consider more powerful models where the algorithm can check exponentially many candidates to find a witness. In particular, we consider the class $\mathbf{OP}(T(n, m))$ of algorithms for the following cases of $T(n, m)$:

- $T(n, m)$ is polynomially bounded in m,
- $T(n, m)$ is quasilinearly bounded in m (of the form $O(m \lg^k m)$ for some $k \geq 0$),

- $T(n, m)$ is subexponential in n and quasilinear in m, and
- $T(n, m)$ is a small exponential in n and quasilinear in m.

Our results include lower bounds on the success probability for deciding the **CircuitSat** problem in all of the models mentioned above subject to various complexity assumptions. These bounds are tight in the sense they achieve the best possible constant in the exponent. In particular, we show that the **CircuitSat** problem cannot be decided with success probability better than $2^{-n+o(n)}$ by probabilistic polynomial-time algorithms unless there are $2^{O(n^{\mu} \lg^{1-\mu} m)}$ size *deterministic* circuits with $\mu < 1$ for deciding $\textbf{CircuitSat}(n, m)$ (Theorem 4). In particular, the latter condition implies that $\textbf{CircuitSat}(n, m)$ (and consequently $\textbf{NP}(n, m)$) can be decided by deterministic circuits of size $2^{n^{\mu}}$ for $\mu < 1$ if m is polynomially bounded in n, which is considered an unlikely event. Assuming that this event does not happen, we get that there is no **OPP** algorithm for deciding **CircuitSat** with success probability better than $2^{-n+o(n)}$.

We will also prove similar lower bounds on the success probability for quasilinear probabilistic circuits. However, in this case, a much weaker assumption suffices. In particular, we show that the success probability cannot be better than $2^{-n+o(n)}$ in quasilinear probabilistic models unless $\textbf{CircuitSat}(n, m)$ (and consequently $\textbf{NP}(n, m)$) has $O(\text{poly}(m) n^{\lg \lg m})$ deterministic size circuits. The statement that $\textbf{CircuitSat}(n, m)$ has $O(\text{poly}(m) n^{\lg \lg m})$ size deterministic circuits is very close to the statement $\textbf{NP} \in \textbf{P}/\textbf{poly}$, which is a highly unlikely event.

We will further show that the success probability in the subexponential model cannot be better than $2^{-n+o(n)}$ unless $\textbf{CircuitSat}(n, m)$ (and consequently $\textbf{NP}(n, m)$) has $2^{o(n)}\text{poly}(m)$ deterministic circuits (Theorem 6). In particular, the latter condition violates **ETH**.

We will also show an optimal lower bound on the probability for small exponential time models of the form $2^{\alpha n} m \lg^k m$ for $\alpha < 1$. We show that the success probability in this model cannot be better than $2^{-(1-\alpha-\varepsilon)n+o(n)}$ for any $\varepsilon > 0$ unless $\textbf{CircuitSat}(n, m)$ (and consequently $\textbf{NP}(n, m)$) has $2^{\beta n}\text{poly}(m)$ size deterministic circuits where $\beta = 1/(1+\varepsilon/\alpha) < 1$. Although the latter statement is weaker, it still would be surprising if there is a constant factor reduction in the number of existential quantifiers for all of **NP**.

Our results essentially rule out any but exhaustive search algorithms for **CircuitSat** in the **OPP** and the corresponding quasilinear time models. We believe that our results are the first of their kind. We hope that these results would put a new emphasis on obtaining similar results for combinatorial problems such as k-colorability and Hamiltonian path as well as on the general question of time-success probability trade-offs. In the following we present some related prior work followed by an informal description of the key ideas.

1.1 Related Prior Work

While there is a large amount of literature on the topic of the satisfiability problem, we will focus in this section on previous research dealing with lower bounds on the exponential complexity of **CircuitSat** and related **NP**-complete problems.

Stearns and Hunt in their 1990 paper [SH90] considered the concept of *power index* to characterize the exponential complexities of basic **NP**-complete problems. Power index is defined to be the infimum of all θ for which the problem is in $\textbf{DTIME}(2^{m^{\theta}})$ where m is the length of the input. They have hypothesized that the power index of the CNF Satisfiability problem is 1 (*Satisfiability Hypothesis*) and using this assumption they have shown that the power index of the Clique problem is $\frac{1}{2}$ (in terms of the number of edges in the graph). They get their results mainly by analyzing how reductions change input lengths. The less a reduction blows up the input size, the tighter the connection between the power indices of the problems.

While their results are interesting, they are more sensitive to how the input is represented. We feel that it is more natural to parameterize **NP** in terms of witness size as well as input size, since the obvious exhaustive search algorithm is strongly exponential in the witness size, but may only be weakly exponential (as in the case of the Clique problem) in the input size. More importantly, while [SH90] assumed the Satisfiability Hypothesis to obtain lower bounds on the power indices for other **NP**-complete problems, in this paper we provide a justification for their assumption by showing that an analogous power index (for the success probability) for a closely related problem, **CircuitSat**, is 1 in **OPP**-style models under complexity assumptions that seem to be intimately related to the **P** versus **NP** question. Moreover, success probability is parameterized using the more natural and robust parameter, the witness length.

In an earlier paper, Schnorr [Sch78] considered the problems in the classes **NQL**, the nondeterministic quasilinear time and **QL**, the deterministic quasilinear time, under quasilinear time reductions. He showed that the CNF Satisfiability problem (and there by the **CircuitSat** problem) are complete for **NQL** under quasilinear time reductions. In fact, Stearns and Hunt cite this result to provide an indirect justification for the Satisfiability Hypothesis since the power index of the CNF Satisfiability problems is at least as large as the power index of any language in **NQL**. Schnorr in this paper comments that $\mathbf{P} \neq \mathbf{NP}$ implies $\mathbf{QL} \neq \mathbf{NQL}$ but the converse is not clear. Our Theorem 5 provides certain strengthening of the forward implication: the probabilistic version of **QL** cannot achieve better than $2^{-n+o(n)}$ success probability unless **NP** has almost polynomial size circuits.

More recently, Impagliazzo, Paturi and Zane [IPZ98] explored the question whether we can expect continued improvements in terms of better exponents for the exact complexity of problems such as k-SAT, k-colorability and Independent Set and showed that the possibility of arbitrarily small exponents for various **NP**-complete problems is one and the same. In particular, they defined the notion of subexponential time reduction families (SERF) and showed that several search problems including k-SAT, IndependentSet, k-Set Cover, Clique, Vertex Cover and k-colorability are SERF-equivalent. They also showed that some of these problems such as k-SAT and k-colorability are SERF-complete for the class of **SNP** of search problems expressible by second order existential formulas whose first order part is universal. If any of these problems can be solved subexponentially (in terms of witness length), then every problem in **SNP** can be solved in subexponential time (in terms of witness length). The key to the equivalence is a lemma called the *Sparsification Lemma* which shows one can achieve witness size-preserving reductions among these problems in subexponential time.

In a subsequent paper [IP01], Impagliazzo and Paturi considered the exact complexity s_k of k-SAT where $s_k = \inf\{\varepsilon | \exists \text{ a } 2^{\varepsilon n} \text{ randomized algorithm for deciding } k\text{-SAT}\}$. Under the assumption $s_3 > 0$, called the *Exponential Time Hypothesis* (**ETH**), they showed that the sequence is increasing infinitely often as k increases. Even under **ETH**, it is an open question to prove a specific lower bound for s_3, that is, to prove $s_3 > c$ for some specific $c > 0$. It is also open whether $s_\infty = 1$ where $s_\infty = \lim_{k\to\infty} s_k$. More generally, it is an open question to prove optimal exponential lower bounds for any **NP**-complete problem under plausible complexity assumptions.

Adopting **ETH** as an axiom casts light on the complexity of many other problems. Marx [Mar07a, Mar07b] used Sparsification Lemma to show that **ETH** implies that the complexity of database queries is determined by their treewidth. Very recently, Traxler [Tra08] has shown that **ETH** implies $(k, 2)$-CSP has k^{cn} complexity where c is an absolute constant, thus ruling out the possibility (under **ETH**) of a c^n time algorithm where c is independent of k. In contrast, k-coloring, a very important case of $(k, 2)$-CSP, has long been known to have such a c^n algorithm for c independent of k, and very recently, the c has been improved to 2 [BH06, Koi06].

There are a number of results that relate the complexity of **CircuitSat** in terms of the tractability of parameterized problems. Abrahamson, Downey and Fellows have shown that the existence of $2^{o(n)}\text{poly}(m)$ algorithms for **CircuitSat** problem for circuits of size m and n variables is equivalent to the problem of the tractability of the class of fixed parameter problems [ADF95]. Other interesting results regarding the connection between the possibilities of somewhat improved algorithms for parameterized problems and subexponential time algorithms of the form $2^{o(n)}\text{poly}(m)$ for **CircuitSat** can be found in [CHKX06].

While the work of [IPZ98] and the subsequent results based on **ETH** as well as the results connecting the complexity of **CircuitSat** with fixed parameter tractability are interesting and represent progress, we still do not have specific lower bounds on the exponents even under **ETH** as the constants in results based on **ETH** (such as those in [Tra08]) depend on the assumed constant s_3 in **ETH**. In contrast, our results do obtain a specific lower bound on the exponent of the success probability for **OPP** and other models under reasonable complexity assumptions. Interestingly, our work suggests the possibility that the question of specific lower bounds on the exponents is related to the questions of **P** versus **NP** and time-success probability trade-off for **NP**-complete problems.

We would also like to mention the recent time-space lower bounds (see [vM07] for a survey): the Formula Satisfiability problem cannot be decided by a deterministic random-access machine that runs in time $m^{1.801}$ and space $m^{o(1)}$ where m is the input length. Unlike our results, these results do not depend on any complexity assumptions. On the other hand, our results deal with exponential time/probability and related algorithmic paradigms.

1.2 Key Ideas

A key idea in our approach is to obtain a simultaneous trade-off between computational resources and instance parameters which in turn would lead to complexity relationships. Several basic conditions are needed to obtain such a trade-off. One of them is to parameterize the problem instances by two parameters, one of them is witness size and the other input size. Another is to parameterize the computational models with two resources or complexity parameters (as functions of the instance parameters). For example, size and success probability are the computational resources in the case of probabilistic circuit models. We then need a *non black-box* reduction technique where the computation itself (after hashing down) is the reduced instance. Identity of the space of instances and the space of computations seems to be crucially necessary for the reduction technique. These basic ideas can be brought together to obtain a simultaneous trade-off between computational resources and between instance parameters to prove our key lemma, *the Exponential Amplification Lemma* (Lemma 2). The Exponential Amplification Lemma is a success probability amplification technique which shows that for any $f(n,m)$-size bounded probabilistic circuit family $\mathcal{A}$ that decides **CircuitSat** with success probability at least $2^{-\alpha n}$ on inputs which are circuits of size m with n variables, there is another probabilistic circuit family $\mathcal{B}$ that decides **CircuitSat** with size roughly $f(\alpha n, f(m,n))$ and success probability about $2^{-\alpha^2 n}$. In contrast, the standard method for boosting success probability by repeated trials will improve it to $(1-(1-2^{-\alpha n})^t)$ ($\approx t2^{-\alpha n}$ for $t = O(2^{\alpha n})$) using circuits of size about $tf(n,m)$.

Elements of our technique are present in [IP01] where instances of CNF satisfiability are parameterized by the number of variables and the maximum width of the clauses and a subexponential time reduction was used to trade up the width for reducing the number of variables thereby obtaining relationships among the exponents. Traxler [Tra08] also obtains a similar trade-off between the size of the domain and the number of variables for constraint satisfaction problems. Whereas the techniques in [IP01, Tra08] only involve reductions among instances trading one parameter for another, our current technique for **CircuitSat** obtains simultaneous trade-offs between algorithmic resources and instance parameters by effecting them in terms of each other.

The proof of the lemma goes as follows: Circuits in the family $\mathcal{B}$ use *specializations* of circuits in the family $\mathcal{A}$ as instances. Specialization of the circuit $C(x,y)$ at $x = \bar{x}$ is the circuit $C^{\bar{x}}(y)$ obtained by plugging in a specific value $\bar{x}$ for the input x. $C^{\bar{x}}(y)$ is now a function of the random variables of the original probabilistic circuit. If the input $\bar{x}$ represents a satisfiable circuit on n variables, $C^{\bar{x}}(y)$ will have a large number of satisfying assignments if the family $\mathcal{A}$ has better than 2^{-n} success probability. Using a technique of Valiant and Vazirani [VV86], we can hash down the circuit $C^{\bar{x}}(y)$ to obtain another circuit H with reduced number of variables in such a way that $C^{\bar{x}}(y)$ and H are satisfiability equivalent. The description of H is then fed to an appropriate circuit in the family $\mathcal{A}$ to amplify the success probability. It turns out that the probabilistic circuit family $\mathcal{B}$ can be designed to implement this amplification.

2 Circuits and Circuit Satisfiability

In this paper, problem instances as well as computational objects are circuits with a single output over the standard, bounded fan-in basis AND, OR, and NOT. Let $\mathcal{C}$ denote the class of such circuits. For $C \in \mathcal{C}$, each source node in the directed acyclic graph of C is either labeled by a random variable or by an input variable or by a constant. Let $n = n(C)$ denote the total number of variables of C, the sum of the number of input variables and the number of random variables. Let $\mathbf{size}(C)$ denote the count of gates in the circuit where each source node is counted as a gate. For input variables y and random variables z, $C(y,z)$ denotes the output of the circuit. $C^y(z)$ denotes the *specialization* of the circuit when its first argument is fixed at the value y. Let $\mathbf{Pr}[C^y(z) = 1]$ denote the probability that the circuit C outputs 1 for the input y, where the random variables z of C take the values 0 or 1 with equal probability.

In our constructions, we also deal with circuits which multiple outputs. Earlier notions and notation extend naturally to such circuits.

Input instances for the **CircuitSat** problem are encodings of circuits and are parameterized by the number of variables and the length of the encoding. Encoding of a circuit contains two parts: the first part encodes the number of variables in unary notation and the second part is a standard encoding of the circuit as a binary string. For a circuit C, let $\mathbf{desc}(C)$ denote the encoding of the circuit and let $m(C) = |\mathbf{desc}(C)|$. $m(C)$ is at least $n(C)$ and is $O(\mathbf{size}(C)\lg(\mathbf{size}(C)))$.

2.1 Circuit Satisfiability and Circuit Families

We are primarily concerned with **CircuitSat**, the *Circuit Satisfiability* problem: given an encoding of $C \in \mathcal{C}$, does there exist a $x \in \{0,1\}^{n(C)}$, that is, a setting of the variables of C such that C with setting x outputs 1. In such a case, we say that the circuit C is *satisfiable*.

We consider probabilistic circuit families indexed by instance parameters as computational models. A family $\mathcal{F}$ of circuits is a collection $\{F_{n,m} | n, m \geq 1\}$ where $F_{n,m}$ is a probabilistic circuit whose inputs are encodings of circuits with n input variables and of encoding length m. For $f : \mathbb{N} \times \mathbb{N} \to \mathbb{R}$, we say that a circuit family $\{F_{n,m}\}$ is f-bounded if $\mathbf{size}(F_{n,m}) \leq O(f(n,m))$. We say that a circuit family $\{F_{n,m}\}$ *decides* **CircuitSat** with *success probability* $p(n)$ if for all inputs which are encodings of circuits with n variables and of encoding length m, $F_{n,m}$ outputs 1 with probability at least $p(n)$ for all satisfiable circuits and otherwise outputs 0 with probability 1. In other words, $p(n) = \inf_{m,y} \mathbf{Pr}[F^y_{n,m}(z) = 1]$, where y is a length m string which is an encoding of a satisfiable circuit with n variables and z denotes the string of random variables of $F_{n,m}$.

Let $\mathcal{F}$ be a circuit family $\{F_{n,m}\}$ deciding **CircuitSat** with success probability $p(n)$. We define the (exponential) *complexity* of $\mathcal{F}$ for deciding **CircuitSat** for inputs which are circuits with n

variables as

$$E_{\mathbf{CircuitSat}}(\mathcal{F}, n) = \lg(1/p(n))/n.$$

The idea is that $E_{\mathbf{CircuitSat}}(\mathcal{F}, n)$ captures the exponent c of the success probability $p(n)$ when expressed as $(2^{-n})^c$. The larger the exponent, the higher the complexity of $\mathcal{F}$. The complexity of $\mathcal{F}$ is defined as $E_{\mathbf{CircuitSat}}(\mathcal{F}) = \limsup E_{\mathbf{CircuitSat}}(\mathcal{F}, n)$.

We define the complexity $E_{\mathbf{CircuitSat}}(f)$ of deciding **CircuitSat** as the *best* exponent achievable by an f-bounded probabilistic circuit family. More precisely,

$$E_{\mathbf{CircuitSat}}(f) = \inf\{\varepsilon | \exists \text{ a } f\text{-bounded, } \mathbf{CircuitSat} \text{ deciding family } \mathcal{F} \text{ such that } E_{\mathbf{CircuitSat}}(\mathcal{F}) \leq \varepsilon\}.$$

We are interested in f-bounded circuit families where $f(n, m) = O(2^{\alpha(n)n} m^k \lg^l m)$ where $\alpha(n)$ is $0, o_n(1)$, or constant, $k \geq 1$, and $l \geq 0$. Such circuit families support computing paradigms where one can evaluate $2^{\alpha(n)n}$ witnesses to find a satisfying solution since a circuit of encoding length m can be evaluated at a given input by a circuit of size $m \lg^l m$ for some $l \geq 0$ [Pip77, PF79]. We use the notation $\tilde{O}(f)$ to suppress polylogarithmic factors in f to express circuit size bounds.

When $\alpha(n) = 0$ and $k \geq 1$, we refer to the circuit families as polynomially bounded. We single out the subclass of quasilinearly bounded circuit families when $\alpha(n) = 0$ and $k = 1$. Let $E_{\mathbf{CircuitSat}}(\mathbf{m^k})$ and $E_{\mathbf{CircuitSat}}(\tilde{\mathbf{O}}(\mathbf{m})))$ denote the complexity of deciding **CircuitSat** by polynomially and quasilinearly bounded circuit families respectively.

$E_{\mathbf{CircuitSat}}(\mathbf{2^{o(n)}m^k})$ ($E_{\mathbf{CircuitSat}}(\mathbf{2^{o(n)}}\tilde{\mathbf{O}}(\mathbf{m}))$) denote the complexity of deciding **CircuitSat** by f-bounded circuit families where $f(n, m) = O(2^{o(n)} m^k)$ with $k \geq 1$ ($k = 1$). $E_{\mathbf{CircuitSat}}(\mathbf{2^{\alpha n}}\tilde{\mathbf{O}}(\mathbf{m}))$ denotes the complexity of deciding **CircuitSat** by f-bounded circuit families where $f(n, m) = 2^{\alpha n}\tilde{O}(m)$ with $\alpha < 1$.

3 Complexity of Circuit Satisfiability

It is clear that $E_{\mathbf{CircuitSat}}(\mathbf{2^{o(n)}}\tilde{\mathbf{O}}(\mathbf{m})) \leq E_{\mathbf{CircuitSat}}(\mathbf{m^k}) \leq E_{\mathbf{CircuitSat}}(\tilde{\mathbf{O}}(\mathbf{m})) \leq 1$. Moreover, $E_{\mathbf{CircuitSat}}(\mathbf{2^{\alpha n}}\tilde{\mathbf{O}}(\mathbf{m})) \leq 1-\alpha$. It is open whether any of these complexities can be lower bounded. In this paper, we prove $E_{\mathbf{CircuitSat}}(\mathbf{2^{o(n)}}\tilde{\mathbf{O}}(\mathbf{m})) = E_{\mathbf{CircuitSat}}(\mathbf{m^k}) = E_{\mathbf{CircuitSat}}(\tilde{\mathbf{O}}(\mathbf{m})) = 1$ under complexity assumptions. We also prove that for any $\alpha < 1$ and $\varepsilon > 0$, $E_{\mathbf{CircuitSat}}(\mathbf{2^{\alpha n}}\tilde{\mathbf{O}}(\mathbf{m})) \geq 1 - \alpha - \varepsilon$ under a certain complexity assumption. Our key lemma, the Exponential Amplification Lemma, shows how to construct a circuit family for deciding **CircuitSat** with improved success probability from a given circuit family that decides **CircuitSat**. An important ingredient in the proof of the Exponential Amplification Lemma is the construction of a circuit **Sparse**(C), which is satisfiability-equivalent to the circuit C, but with fewer variables. This construction is closely related to that of the unique satisfiability construction of Valiant and Vazirani [VV86] and is captured in the following hash-down lemma, Lemma 1.

Let $C \in \mathcal{C}$ be a satisfiable circuit with n variables and let m denote its description length. Let $S \subseteq \{0, 1\}^n$ be the nonempty set of satisfying assignments to the variables of C. Let $s = \lfloor \lg |S^C| \rfloor - 2$; we will assume that $s > 0$.

The intuition for the construction of **Sparse**(C) is as follows. If the set S^C is intersected with a random subcube of dimension $(n - s)$, we expect to get a nonempty intersection. Therefore, $(n - s)$ bits are sufficient to locate a satisfying assignment in the intersection. This intuition is operationalized by restricting C to inputs from a random coset of a random linear transformation from $\{0, 1\}^n$ to $\{0, 1\}^s$. Such a restriction turns out to be satisfiability-equivalent since the random coset contains a satisfying assignment with sufficiently high probability. Moreover, members of the coset can be generated by a small-size circuit given their $(n - s)$-bit address in the coset. The details are provided below.

We follow the standard idea of using pairwise independent functions for hashing. However, rather than using random linear transformations we will use random Toeplitz matrices to achieve the desired pairwise independence. It requires only linear number of bits to specify a random Toeplitz matrix over GF(2). Linear randomness together with fast GCD and convolution algorithms [BGY80, Pan01] will only require a quasilinear computation to select a random coset and to address its members. While we can live with polynomial overhead for some of our theorems, quasilinear hashing is necessary when we work with quasilinear size circuit models.

Let n be fixed and let $t = (t_{-n+1}, \cdots, t_{-1}, t_0, t_1, \cdots, t_{n-1}) \in \{0,1\}^{2n-1}$. We will denote by T_t the Toeplitz matrix determined by t, which is the matrix defined by $T_t(i,j) = t_{(j-i)}$ for $0 \leq i, j \leq n-1$. For a column vector $z \in \{0,1\}^n$, let $(T_t(z))_s$ denote the column vector of the first s bits of $T_t(z)$. For $t \in \{0,1\}^{2n-1}$ and $w \in \{0,1\}^s$, define the affine linear transformation $h_{t,w}(z) = (T_t z)_s + w$. It is well-known that $\{h_{t,w}\}$ is a pairwise independent family of functions from $\{0,1\}^n$ to $\{0,1\}^s$.

We would like to parameterize the cosets $H_{t,w} = \{z | h_{t,w}(z) = 0\}$. For $w \in \{0,1\}^s$ and $x \in \{0,1\}^{n-s}$ we will define the column vector $(w; x)$ to be the concatenation of w and x. For $w \in \{0,1\}^s$ and t such that T_t is invertible, we define $J_{t,w}(x) := T_t^{-1}(w; x)$ to obtain the needed parameterization of the coset $H_{t,w}$. When T_t is nonsingular, it is easy to see that $\mathbf{Image}(J_{t,w}) = H_{t,w}$. Indeed, if $z \in H_{t,w}$, then by definition $h_{t,w}(z) = (T_t(z))_s + w = 0$. This implies that there exists an $x \in \{0,1\}^{n-s}$ such that $T_t(z) + (w; x) = 0$. Since we are working over GF(2) and since T_t is nonsingular, it follows that $z = T_t^{-1}(w; x) = J_{t,w}(x)$. Similarly, if $z \in \mathbf{Image}(J_{t,w})$, then $z \in H_{t,w}$.

The circuit $\mathbf{Sparse}^{t,w}(C)$ with $(n-s)$ variables is the composition of $J_{t,w}$ with C, *i.e.*, $\mathbf{Sparse}^{t,w}(C)(x) = C(J_{t,w}(x))$.

Lemma 1. *C and $\mathbf{Sparse}^{t,w}(C)$ are satisfiability-equivalent with probability at least $1/4$. Furthermore, $\mathbf{Sparse}^{t,w}(C)$ can be constructed in size $\tilde{O}(m)$ and its description can be computed by a circuit of size $\tilde{O}(m)$ given $\mathbf{desc}(C), t$ and w.*

Proof of Lemma 1 can be found in the Appendix. We now state and prove the Exponential Amplification Lemma.

Lemma 2. Exponential Amplification Lemma: *Let $\mathcal{F}$ be an f-bounded family for some $f : \mathbb{N} \times \mathbb{N} \to \mathbb{R}$ such that $E_{\mathbf{CircuitSat}}(\mathcal{F}) < \delta$ for $0 < \delta < 1$. Assume $f(n,m) \geq m$ for all n. Then there exists a g-bounded circuit family $\mathcal{G}$ such that, for all sufficiently large n, $E_{\mathbf{CircuitSat}}(\mathcal{G}) < \delta^2$ where*

$$g(n,m) = O(f(\lceil \delta n \rceil + 5, \tilde{O}(f(n,m)))).$$

Proof. Let $F_{n,m}$ be the probabilistic circuit from the family $\mathcal{F}$ that decides **CircuitSat** for circuits of description size m and n variables. $F_{n,m}$ itself has m input variables which encode the description of a circuit D of n variables. Also $F_{n,m}$ has r random variables y. By assumption $\mathbf{size}(F_{n,m}) = O(f(n,m))$, $r = O(f(n,m))$, and the success probability $p(n)$ of $F_{n,m}$ is greater than $2^{-\delta n}$ for all sufficiently large n.

Let $n' = \lceil \delta n \rceil + 5$ and $s = r - n'$. It follows that $p(n) > 2^{-n'}$. We will construct a g-bounded family $\mathcal{G}$ of probabilistic circuits that decides **CircuitSat** with success probability $p(r-s) = p(n')$ for circuits with n variables where $g(n,m) = O(f(n', \tilde{O}(f(n,m))))$. A key idea is, for a given circuit D, to view $F_{n,m}(\mathbf{desc}(D), y)$ as a (deterministic) boolean circuit $C(y) := F_{n,m}^{\mathbf{desc}(D)}(y)$ of r variables. When D is satisfiable, it follows that $C(y)$ has at least $2^{(r-\lceil \delta n \rceil)}$ solutions for all sufficiently large n. We apply the $\mathbf{Sparse}()$ function to C to obtain a circuit with n' variables which is satisfiability-equivalent to D. We then apply the description of $\mathbf{Sparse}(C)$ as input to an appropriate circuit from the family $\mathcal{F}$ to improve success probability. The details of this construction are presented in the following algorithm and figures.

Circuit $G_{n,m}$ (see Picture 1):

```
Input: $\mathbf{desc}(D)$ of $D \in \mathcal{C}$ with $n$ variables and description length $m$.
$C(y) := F_{n,m}^{\mathbf{desc}(D)}(y)$, the specialization of $F_{n,m}$ to the input $\mathbf{desc}(D)$.
If $r > n'$,
    Select random $t \in \{0,1\}^{2r-1}$ and $w \in \{0,1\}^s$.
    If $T_t$ is invertible, compute the description of a circuit $J_{t,w}(x)$ that computes $T_t^{-1}(w;x)$
for $x \in \{0,1\}^{n'}$.
    $\mathbf{Sparse}^{t,w}(C)(x) := C(J_{t,w}(x))$.
    $H(x) := \mathbf{Sparse}^{t,w}(C)(x)$ where $x \in \{0,1\}^{n'}$. (see Picture 2)
else
    $H(x) = C(x)$ where $x \in \{0,1\}^r$.
PrepCkt: Compute the description of the circuit $H$. Let $m' = |\mathbf{desc}(H)|$.
Apply $F_{n'',m'}$ to $\mathbf{desc}(H)$ where $n'' = n'$ if $r > n'$, otherwise $n'' = r$.
Output: $F_{n'',m'}(\mathbf{desc}(H))$.
```

We will first argue that D and H are satisfiability-equivalent. It is clear that if D is unsatisfiable, H is unsatisfiable. If D is satisfiable, then it follows from Lemma 1 that $\mathbf{Sparse}^{t,w}(C)$ is satisfiable with probability at least 1/4. Since $n'' \leq \lceil \delta n \rceil + 5$, by assumption $F_{n'',m'}$ outputs 1 on input $\mathbf{desc}(H)$ with probability at least $2^{-\delta n''}$, which implies that the success probability of $G_{n,m}$ is at least $2^{-\delta' \delta n - 6\delta' - 5}$ for some $\delta' < \delta$. It then follows that $E_{\mathbf{CircuitSat}}(\mathcal{G}, n) < \delta^2$ for all sufficiently large n.

We will now upper bound the size of $G_{n,m}$. By Lemma 1, $\mathbf{Sparse}^{t,w}(C)(x)$ can be described using at most $m' = \tilde{O}(f(n,m))$ bits and the description itself can be computed by a circuit of size $\tilde{O}(f(n,m)))$. It follows that the size of H is also bounded by $\tilde{O}(f(n,m))$ whether $r > n'$ or not.

Thus, the size of $G_{n,m}$ is upper bounded by

$$\begin{aligned} f(n'', m') + \tilde{O}(f(n,m)) &\leq f(\lceil \delta n \rceil + 5, \tilde{O}(f(n,m))) + \tilde{O}(f(n,m)) \\ &\leq O(f(\lceil \delta n \rceil + 5, \tilde{O}(f(n,m)))) \text{ since } f(n,m) \geq m. \end{aligned}$$

□

We think that the following lemma may be of independent interest, which uses the hash-down technique to boost the probability. This lemma presents an alternative technique to boosting the success probability by repeated independent trails. Its proof can be found in the Appendix.

Lemma 3. *If C is a probabilistic circuit of size M deciding* **CircuitSat** *with success probability $q > 0$ for a set of inputs, then there exists a deterministic circuit D of size $\tilde{O}(M^2)/q$ that decides* **CircuitSat** *on the same set of inputs. The same holds for any set in* **NP** *in place of* **CircuitSat**.

3.1 Results

For a variety of $f(n,m)$, we will argue it is implausible that the success probability could be as large as $2^{-\delta n}$ for $\delta < 1$ for any $f(n,m)$-bounded probabilistic circuit family deciding **CircuitSat**. Our technique is essentially the following: if the success probability is large, (by repeated applications of the Exponential Amplification Lemma followed by an application of Lemma 3) we can construct small size deterministic circuits for deciding $\mathbf{CircuitSat}(n,m)$ (or equivalently $\mathbf{NP}(n,m)$), which implies implausible events. The size of the deterministic circuit for deciding $\mathbf{CircuitSat}(n,m)$ depends on $f(n,m)$. The smaller the f, the smaller the size of the circuit for deciding $\mathbf{CircuitSat}(n,m)$.

However, the implausibility of the consequence for deciding $\mathbf{CircuitSat}(n, m)$ could also depend on the relative size of m with respect to n. Note that $\mathbf{CircuitSat}(n, m)$ can be decided by circuits of size $2^n\tilde{O}(m)$. If m is exponentially large in n (for example, $m = 2^n$) and if f is growing sufficiently fast with m (for example, $f(n, m) = \Omega(m^2\text{poly}(\lg m))$), then $f(n, m)$-bounded circuit can decide the **CircuitSat** problem with success probability 1.

In the following we select certain examples of f and state the resulting consequence for deciding $\mathbf{CircuitSat}(n, m)$ if the success probability is large enough. While our results hold for arbitrary parameters n and m, in the remarks following the theorems we point out the implausibility explicitly if it requires focusing on certain values of m as a function of n.

Theorem 4. *Either* $E_{\mathbf{CircuitSat}}(\mathbf{m^k}) = 1$ *or there exists a* $\mu < 1$ *such that* $\mathbf{CircuitSat}(n, m)$ *(and consequently* $\mathbf{NP}(n, m)$*) can be decided by deterministic circuits of size* $2^{O(n^\mu \lg^{1-\mu} m)}$.

Remarks: Consider the second clause in the disjunction in the statement of the theorem: there exists a $\mu < 1$ such that $\mathbf{CircuitSat}(n, m)$ can be decided by deterministic circuits of size $2^{O(n^\mu \lg^{1-\mu} m)}$. This statement implies that $\mathbf{CircuitSat}(n, m)$ can be decided by $O(2^{n^\mu})$ size deterministic circuits for $\mu < 1$ for the case where m is bounded by a polynomial in n. It is currently believed that such circuits are unlikely to exist. Even if $m = 2^{o(n)}$, we get that **CircuitSat** can be decided by deterministic circuits of size $2^{o(n)}$, which is also considered implausible. For example, this event contradicts **ETH**. Given these implausibilities, we can conclude that **OPP** algorithms cannot achieve success probability better than $2^{-n+o(n)}$. If $m \geq 2^{cn}$ for some constant $c > 0$, the theorem becomes a trivial statement.

Outline of the Proof: We assume that there exists a family $\mathcal{F}$ of probabilistic circuits of size $O(m^k)$ for some $k \geq 1$ achieving success probability $2^{-\delta n}$ for some $\delta < 1$. We apply the Exponential Amplification Lemma $d = (\lg \frac{n}{2\lg m+O(\lg\lg m)})/(\lg(\frac{1}{\delta} + \lg k)$ times followed by an application of Lemma 3 to obtain a deterministic circuit family of size $2^{O(n^\mu \lg^{1-\mu} m)}$ where $\mu < 1$ depends on k and δ.

Theorem 5. *Either* $E_{\mathbf{CircuitSat}}(\tilde{\mathbf{O}}(\mathbf{m})) = 1$ *or* $\mathbf{CircuitSat}(n, m)$ *(and consequently* $\mathbf{NP}(n, m)$*) can be decided by deterministic circuits of size* $O(\text{poly}(m)n^{O(\lg\lg m)})$.

Remarks: The consequence of better success probability is that $\mathbf{CircuitSat}(n, m)$ can be decided by quasi-polynomial size deterministic circuits which is very close to the statement $\mathbf{NP} \in \mathbf{P/poly}$. This highly improbable statement lets us conclude that we can only succeed with probability $2^{-n+o(n)}$ when we decide $\mathbf{CircuitSat}(n, m)$ using quasilinear probabilistic circuits. Theorem 5 can be proved by applying the Exponential Amplification Lemma about $O(\lg n)$ times.

We will also get tight lower bounds on the success probability when $f(n, m)$ is subexponential or exponential in n and quasilinear in m.

Theorem 6. *Either* $E_{\mathbf{CircuitSat}}(\mathbf{2^{o(n)}}\tilde{\mathbf{O}}(\mathbf{m})) = 1$ *or* $\mathbf{CircuitSat}(n, m)$ *(and consequently* $\mathbf{NP}(n, m)$*) can be decided by deterministic circuits of size* $2^{o(n)}\text{poly}(m)$.

Remarks: Theorem 6 is proved by applying the Exponential Amplification Lemma a number of times which grows with n. When we restrict m to be polynomial in n, the second statement in the theorem implies that $\mathbf{CircuitSat}(n, m)$ has subexponential size circuits which contradicts **ETH**.

Theorem 7. *For every* $\alpha, \varepsilon > 0$, *either* $E_{\mathbf{CircuitSat}}(\mathbf{2^{\alpha n}}\tilde{\mathbf{O}}(\mathbf{m})) \geq 1 - \alpha - \varepsilon$ *or* $\mathbf{CircuitSat}(n, m)$ *(and consequently* $\mathbf{NP}(n, m)$*) can be decided by circuits of size* $2^{n/(1+\varepsilon/\alpha)}\text{poly}(m)$.

Remark: In other words, if the success probability is better than $2^{-(1-\alpha)n+o(n)}$, we get that $\mathbf{CircuitSat}(n, m)$ can be decided by deterministic circuits of size $2^{cn}\text{poly}(m)$ where $c = 1/(1+\frac{\varepsilon}{\alpha}) <$

1. It should be noted that the standard success boosting technique (as opposed to the Exponential Amplification Lemma) would give deterministic circuits of size $2^{(1-\varepsilon)n}\text{poly}(m)$. It is easy to see that $c < (1-\varepsilon)$ as long as $\alpha < 1$.

It is useful to interpret our results as time-probability trade-offs. Consider $t := f(n, m)$-bounded circuit families that decide **CircuitSat**. As defined earlier, let $p := E_{\textbf{CircuitSat}}(f)$ denote as the best exponent achievable by an f-bounded probabilistic circuit family that decides **CircuitSat**. Consider p as a function of t. From the standard probability boosting technique, we conclude that as t increases that p must at least decrease at a rate of 1. In other words, the quantity $\lg t/n + p$ cannot increase. Our results say that if **CircuitSat** can be solved with probability $2^{-\delta n}$ for $\delta < 1$ at any t (ranging from quasilinear to small exponential function), probability will decrease at a rate higher than 1 as time increases. The rate of decline depends on δ as well as the earliest time at which such an advantage in guessing a satisfying solution can be achieved. If δ is already less than 1 already at quasilinear time, p will decline quite rapidly to yield a superpolynomial time algorithm for **NP**.

Acknowledgments: The authors would like to thank Boaz Barak, Chris Calabro, Russell Impagliazzo, Mike Saks, Avi Wigderson, and Ryan Williams for helpful discussions.

References

[ADF95] Karl R. Abrahamson, Rodney G. Downey, and Michael R. Fellows. Fixed-parameter tractability and completeness iv: On completeness for w[p] and pspace analogues. *Ann. Pure Appl. Logic*, 73(3):235–276, 1995.

[AKS01] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 601–610, New York, NY, USA, 2001. ACM.

[BE95] Richard Beigel and David Eppstein. 3-coloring in time $o(1.3446^n)$: A no-mis algorithm. In *FOCS*, pages 444–452, 1995.

[Bei99] Richard Beigel. Finding maximum independent sets in sparse and general graphs. In *SODA '99: Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms*, pages 856–857, Philadelphia, PA, USA, 1999. Society for Industrial and Applied Mathematics.

[Bel62] Richard Bellman. Dynamic programming treatment of the travelling salesman problem. *Journal of the ACM*, 9(1):61–63, 1962.

[BGY80] Richard P. Brent, Fred G. Gustavson, and David Y. Y. Yun. Fast solution of toeplitz systems of equations and computation of padé approximants. *J. Algorithms*, 1(3):259–295, 1980.

[BH06] Andreas Björklund and Thore Husfeldt. Inclusion–exclusion algorithms for counting set partitions. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 575–582, Washington, DC, USA, 2006. IEEE Computer Society.

[BK06] Hans Bodlaender and Dieter Kratsch. An exact algorithm for graph coloring with polynomial memory. Technical Report UU-CS-2006-015, Utrecht University, 2006.

[BKK07] Andreas Björklund, Thore Husfeldt Petteri Kaski, and Mikko Koivisto. Fourier meets möbius: Fast subset convolution. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 67–74, New York, NY, USA, 2007. ACM.

[BKK08] Andreas Björklund, Thore Husfeldt Petteri Kaski, and Mikko Koivisto. Computing the tutte polynomial in vertex-exponential time. In *FOCS '08: Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 677–686, Washington, DC, USA, 2008. IEEE Computer Society.

[Bys03] Jesper Makholm Byskov. Algorithms for k-colouring and finding maximal independent sets. In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 456–457, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.

[Bys05] J. M. Byskov. *Exact Algorithms for Graph Colouring and Exact Satisfiability.* PhD thesis, Aarhus University, Aarhus, Denmark, August 2005.

[CHKX06] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.*, 72(8):1346–1367, 2006.

[DGH$^+$02] Evgeny Dantsin, Andreas Goerdt, Edward A. Hirsch, Ravi Kannan, Jon Kleinberg, Christos Papadimitriou, Prabhakar Raghavan, and Uwe Schöning. A deterministic $(2-2/(k+1))^n$ algorithm for k-sat based on local search. *Theoretical Computer Science*, 289(1):69–83, 2002.

[Epp01] David Eppstein. Improved algorithms for 3-coloring, 3-edge-coloring, and constraint satisfaction. In *SODA '01: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 329–337, Philadelphia, PA, USA, 2001. Society for Industrial and Applied Mathematics.

[Epp03] David Eppstein. Small maximal independent sets and faster exact graph coloring. *Journal of Graph Algorithms and Applications*, 7:131–140, 2003.

[Epp06] David Eppstein. Quasiconvex analysis of multivariate recurrence equations for backtracking algorithms. *ACM Trans. Algorithms*, 2(4):492–509, 2006.

[FGK06] F. Fomin, F. Grandoni, and D. Kratsch. Measure and conquer : a simple $o(2^{0.288n})$ independent set algorithm. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 18–25, New York, NY, USA, 2006. ACM.

[GHNR03] Jens Gramm, Edward A. Hirsch, Rolf Niedermeier, and Peter Rossmanith. Worst-case upper bounds for max-2-sat with an application to max-cut. *Discrete Applied Mathematics*, 130(2):139–155, 2003.

[HK61] Michael Held and Richard M. Karp. A dynamic programming approach to sequencing problems. In *Proceedings of the 1961 16th ACM national meeting*, pages 71.201–71.204, New York, NY, USA, 1961. ACM.

[IP01] R. Impagliazzo and R. Paturi. The complexity of k-sat. *Journal of Computer and Systems Sciences*, 62(2):367–375, March 2001. Preliminary version in *14th Annual IEEE Conference on Computational Complexity*, pages 237–240, 1999.

[IPZ98] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63:512–530, 1998. Preliminary version in *39th Annual IEEE Symposium on Foundations of Computer Science*, pp 653-662, 1998.

[Kar82] Richard M. Karp. Dynamic programming meets the principle of inclusion and exclusion. *Operations Research Letters*, 1:49–51, April 1982.

[KL96] E. Kaltofen and A. Lobo. On rank properties of toeplitz matrices over finite fields. In *ISSAC '96: Proceedings of the 1996 International Symposium on Symbolic and Algebraic Computation*, pages 241–249, New York, NY, USA, 1996. ACM.

[Koi06] Mikko Koivisto. An $o^*(2^n)$ algorithm for graph coloring and other partitioning problems via inclusion-exclusion. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, pages 583–590, Washington, DC, USA, 2006. IEEE Computer Society.

[Kou08] Ioannis Koutis. Faster algebraic algorithms for path and packing problems. In *ICALP '08: Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, pages 575–586, Berlin, Heidelberg, 2008. Springer-Verlag.

[Law76] E. Lawler. A note on the complexity of the chromatic number problem. *Information Processing Letters*, 5(3):66–67, 1976.

[Mar07a] Dániel Marx. Can you beat treewidth? In *FOCS*, pages 169–179, 2007.

[Mar07b] Dániel Marx. On the optimality of planar and geometric approximation schemes. In *FOCS*, pages 338–348, 2007.

[MV09] Daniele Micciancio and Panagiotis Voulgaris. Faster exponential time algorithms for the shortest vector problem. manuscript, 2009.

[Pan01] Victor Pan. *Structured Matrices and Polynomials: Unified Superfast Algorithms.* Birkhäuser, Boston, 2001.

[PF79] Nicholas Pippenger and Michael J. Fischer. Relations among complexity measures. *J. ACM*, 26(2):361–381, 1979.

[Pip77] Nicholas Pippenger. Fast simulation of combinational logic circuits without random-access storage. In *Proceedings of the Fifteenth Allerton Conference on Communication, Control, and Computing*, pages 25–33, 1977.

[PPZ99] R. Paturi, P. Pudlák, and F. Zane. Satisfiability coding lemma. *Chicago Journal of Theoretical Computer Science*, December 1999. Preliminary version in *38th Annual Symposium on Foundations of Computer Science*, 566-574, 1997.

[Rob86] J. M. Robson. Algorithms for maximum independent sets. *Journal of Algorithms*, 7:425–440, September 1986.

[Sak98] Michael E. Saks. Space efficient algorithms for np-hard sequencing and partition problems. manuscript, November 1998.

[Sch78] Claus-Peter Schnorr. Satisfiability is quasilinear complete in nql. *J. ACM*, 25(1):136–145, 1978.

[Sch99] U. Schöning. A probabilistic algorithm for k-sat and constraint satisfaction problems. In *FOCS*, pages 410–414, 1999.

[SH90] Richard Edwin Stearns and Harry B. Hunt. Power indices and easier hard problems. *Mathematical Systems Theory*, 23(4):209–225, 1990.

[Tra08] Patrick Traxler. The time complexity of constraint satisfaction. In *2008 Dagstuhl Workshop on Moderately Exponential-time Algorithms*, 2008.

[vM07] Dieter van Melkebeek. A survey of lower bounds for satisfiability and related problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(099), 2007.

[VV86] L. Valiant and V. Vazirani. NP is as easy as detecting unique solutions. *Theoretical Computer Science*, 47:85–93, 1986.

[Wil09] Ryan Williams. Finding paths of length k in $o^*(2^k)$ time. *Information Processing Letters*, 109(6):315–318, 2009.

4 Appendix

Lemma 1: *C and $\mathbf{Sparse}^{t,w}(C)$ are satisfiability-equivalent with probability at least $1/4$. Furthermore, $\mathbf{Sparse}^{t,w}(C)$ can be constructed in size $\tilde{O}(m)$ and its description can be computed by a circuit of size $\tilde{O}(m)$ given $\mathbf{desc}(C), t$ and w.*

Proof. Our goal is to prove

$$\mathbf{Pr}_{t,w}[S^C \cap \mathbf{Image}(J_{t,w}) \neq \emptyset] \geq 1/4,$$

where we are assuming the uniform distribution of $t \in \{0,1\}^{2n-1}$ and $w \in \{0,1\}^s$. However, $\mathbf{Image}(J_{t,w}) = H_{t,w}$ if T_t is nonsingular. We will first argue that $\mathbf{Pr}_{t,w}[S^C \cap H_{t,w} \neq \emptyset] \geq 3/4$. We then use the result of Kaltofen and Lobo [KL96] which states $\mathbf{Pr}[T_t \text{ is nonsingular}] \geq 1/2$ to conclude

$$\begin{aligned} \mathbf{Pr}_{t,w}[S^C \cap \mathbf{Image}(J_{t,w}) \neq \emptyset] &\geq \mathbf{Pr}_{t,w}[S^C \cap H_{t,w} \neq \emptyset] - \mathbf{Pr}[T_t \text{ is singular}] \\ &\geq 1/4 \end{aligned}$$

For $z \in S^C$, let $q_{t,w}(z)$ denote the indicator function for the event $z \in H_{t,w}$, *i.e.*, $h_{t,w}(z) = 0$. Let $Q_{t,w} := \sum_{z \in S^C} q_{t,w}(z)$. We have $Q_{t,w} \neq 0$ iff $S^C \cap H_{t,w} \neq \emptyset$. We will upper bound the probability of the event $Q_{t,w} = 0$ using Chebyshev's inequality: $\mathbf{Pr}[Q_{t,w} = 0] \leq \frac{\mathbf{Var}(Q_{t,w})}{(\mathbf{E}[Q_{t,w}])^2}$. Using the property of pairwise independence of the functions $h_{t,w}$, we get

$$\mathbf{Var}[Q_{t,w}] = \mathbf{E}[Q_{t,w}^2] - \mathbf{E}[Q_{t,w}]^2 = (|S^C|^2 - |S^C|)2^{-2s} + |S^C|2^{-s} - |S^C|^2 2^{-2s} = |S^C|(2^{-s} - 2^{-2s}).$$

Hence

$$\frac{\mathbf{Var}(Q_{t,w})}{\mathbf{E}[Q_{t,w}]^2} = \frac{|S^C|(2^{-s} - 2^{-2s})}{|S^C|^2 2^{-2s}} = \frac{1 - 2^{-s}}{|S^C|2^{-s}} \leq \frac{1 - 2^{-s}}{4} < 1/4.$$

This completes the proof that C and $\mathbf{Sparse}^{t,w}(C)$ are satisfiability equivalent.

Using fast GCD and convolution algorithms as well as the Gohberg-Semencul formula for the inverse of a Toeplitz matrix [BGY80, Pan01], one can compute $J_{t,w}(x) = T_t^{-1}(w; x)$ with bit complexity $\tilde{O}(n)$. Since $m \geq n$, it turns out that the circuit $\mathbf{Sparse}^{t,w}(C)$ can be constructed in $\tilde{O}(m)$ size. It is also easy to check that the description of $\mathbf{Sparse}^{t,w}(C)$ can be computed by a circuit of size $\tilde{O}(m)$. □

Lemma 3: *If C is a probabilistic circuit of size M deciding* **CircuitSat** *with success probability $q > 0$ for a set of inputs, then there exists a deterministic circuit D of size $\tilde{O}(M^2)/q$ that decides* **CircuitSat** *on the same set of inputs. The same holds for any set in* **NP** *in place of* **CircuitSat**.

Proof. Let $C = C(x, y)$, where the input variables are x and the random variables are y. Let r the number of random variables. Let $J_{t,w} : \{0,1\}^{r-s} \to \{0,1\}^r$ be as in Lemma 1, where $s = r + \lfloor \lg q \rfloor - 2$. Consider the probabilistic circuit

$$\bigvee_{z \in \{0,1\}^{r-s}} C(x, J_{t,w}(z)),$$

with input variables x and random bits t and w. This circuit computes **CircuitSat** with success probability $1/4$ as shown in the proof of Lemma 1. The standard error reduction argument using the disjunction of about M independent (in terms of the selection of random t and w) copies of the probabilistic circuit obtains an exponentially small error probability. This ensures that there exists a setting of the random bits that always produces the correct answer. The size of the resulting circuit is $\tilde{O}(M^2/q)$. □

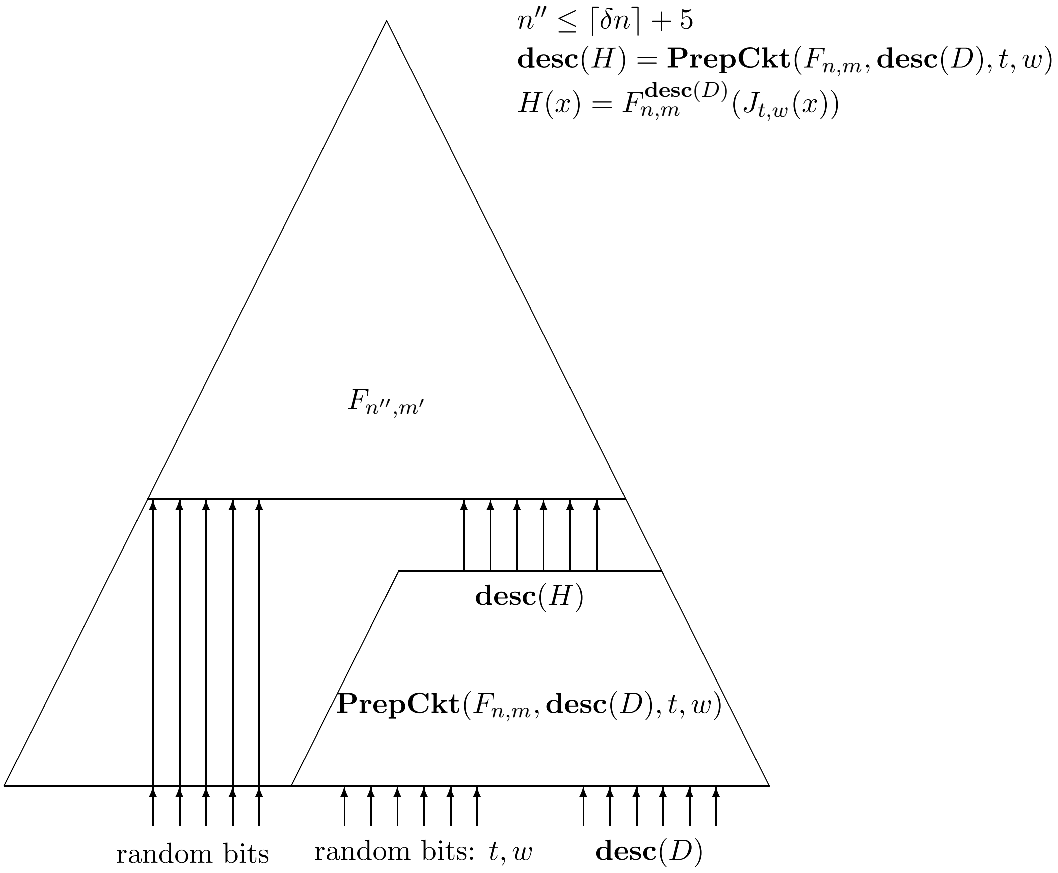
$n'' \leq \lceil \delta n \rceil + 5$
$\mathbf{desc}(H) = \mathbf{PrepCkt}(F_{n,m}, \mathbf{desc}(D), t, w)$
$H(x) = F_{n,m}^{\mathbf{desc}(D)}(J_{t,w}(x))$
$F_{n'',m'}$
$\mathbf{desc}(H)$
$\mathbf{PrepCkt}(F_{n,m}, \mathbf{desc}(D), t, w)$
random bits
random bits: t, w
$\mathbf{desc}(D)$

Figure 1: Circuit $G_{n,m}$

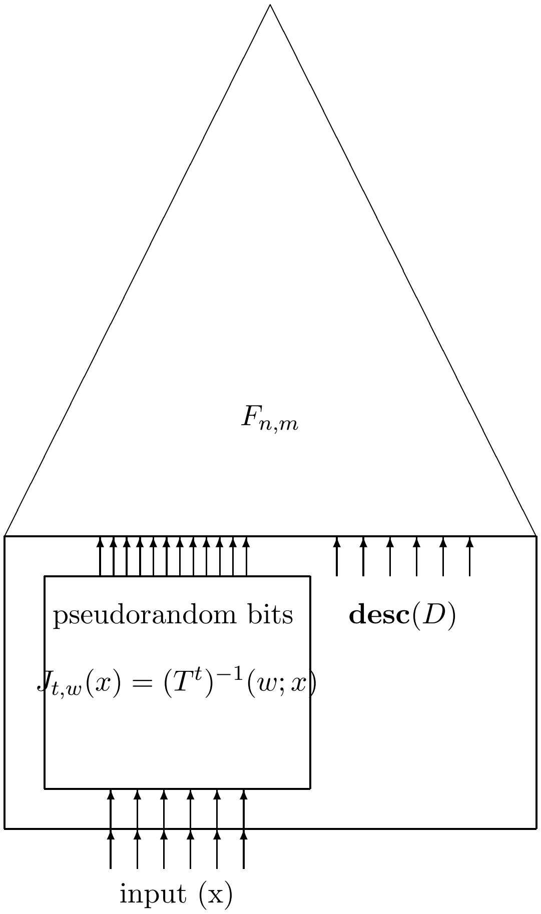

Figure 2: $H(x) = F_{n,m}^{\mathbf{desc}(D)}(J_{t,w}(x))$