

An Improved Deterministic #SAT Algorithm for Small De Morgan Formulas

Ruiwen Chen¹, Valentine Kabanets¹, and Nitin Saurabh²

¹ Simon Fraser University, Burnaby, Canada

ruiwenc@sfu.ca, kabanets@cs.sfu.ca

² Institute of Mathematical Sciences, Chennai, India

nitin@imsc.res.in

Abstract. We give a deterministic #SAT algorithm for de Morgan formulas of size up to $n^{2.63}$, which runs in time $2^{n-n^{\Omega(1)}}$. This improves upon the deterministic #SAT algorithm of [3], which has similar running time but works only for formulas of size less than $n^{2.5}$.

Our new algorithm is based on the shrinkage of de Morgan formulas under random restrictions, shown by Paterson and Zwick [12]. We prove a *concentrated* and *constructive* version of their shrinkage result. Namely, we give a deterministic polynomial-time algorithm that selects variables in a given de Morgan formula so that, *with high probability* over the random assignments to the chosen variables, the original formula shrinks in size, when simplified using a *deterministic polynomial-time* formula-simplification algorithm.

Keywords: de Morgan formulas, random restrictions, shrinkage, SAT algorithms.

1 Introduction

Subbotovskaya [16] introduced the method of *random restrictions* to prove that PARITY requires de Morgan formulas of size $\Omega(n^{1.5})$, where a de Morgan formula is a boolean formula over the basis $\{\vee, \wedge, \neg\}$. She showed that a random restriction of all but a fraction p of the input variables yields a new formula whose size is expected to reduce by at least the factor $p^{1.5}$. That is, the *shrinkage exponent* Γ for de Morgan formulas is at least 1.5, where the shrinkage exponent is defined as the least upper bound on γ such that the expected formula size shrinks by the factor p^γ under a random restriction leaving p fraction of variables free.

Impagliazzo and Nisan [9] argued that Subbotovskaya's bound $\Gamma \geq 1.5$ is not optimal, by showing that $\Gamma \geq 1.556$. Paterson and Zwick [12] improved upon [9], getting $\Gamma \geq (5 - \sqrt{3})/2 \approx 1.63$. Finally, Håstad [6] proved the tight bound $\Gamma = 2$; combined with Andreev's construction [1], this yields a function in P requiring de Morgan formulas of size $\Omega(n^{3-o(1)})$.

While the original motivation to study shrinkage in [16,9,12,6] was to prove formula lower bounds, the same results turn out to be useful also for designing non-trivial SAT algorithms for small de Morgan formulas. Santhanam [14] strengthened

Subbotovskaya's *expected* shrinkage result to *concentrated* shrinkage, i.e., shrinkage with high probability, and used this to get a deterministic #SAT algorithm (counting the number of satisfying assignments) for linear-size de Morgan formulas, with the running time $2^{n-\Omega(n)}$. Santhanam's algorithm deterministically selects a most frequent variable in the current formula, and recurses on the two subformulas obtained by restricting the chosen variable to 0 and 1; after $n - \Omega(n)$ recursive calls, almost all obtained formulas depend on fewer than the actual number of free variables remaining, which leads to nontrivial savings over the brute-force SAT algorithm for the original formula. A similar algorithm works also for formulas of size less than $n^{2.5}$, with the running time $2^{n-n^{\Omega(1)}}$ [3].

Motivated by average-case formula lower bounds, Komargodksi et al. [11] (building upon [8]) showed a concentrated-shrinkage version of Håstad's optimal result for the shrinkage exponent $\Gamma = 2$. Combined with the aforementioned algorithm of Chen et al. [3], this yields a nontrivial *randomized* zero-error #SAT algorithm for de Morgan formulas of size $n^{3-o(1)}$, running in time $2^{n-n^{\Omega(1)}}$.

The main question addressed by our paper is whether there is a *deterministic* #SAT algorithm, with similar running time, for formulas of size close to n^3 . This question is interesting since getting a deterministic algorithm often yields deeper understanding of the problem by revealing additional structural properties. It also provides better understanding of the role of randomness in efficient algorithms, as part of research on *derandomization*.

We give a deterministic #SAT algorithm for formulas of size up to $n^{2.63}$. In the process, we refine the results of Paterson and Zwick [12] on shrinkage of de Morgan formulas by making their results *constructive* in a certain precise sense. We provide more details next.

1.1 Our Main Results and Techniques

Our main result is a *deterministic* #SAT algorithm for de Morgan formulas of size up to $n^{2.63}$, running in time $2^{n-n^{\Omega(1)}}$.

Theorem 1 (Main). *There is a deterministic algorithm for counting the number of satisfying assignments in a given de Morgan formula on n variables of size at most $n^{2.63}$ which runs in time at most 2^{n-n^δ} , for some constant $0 < \delta < 1$.*

As in [14,3], we use a deterministic algorithm to choose a next variable to restrict, and then recurse on the two resulting restrictions of this variable to 0 and 1. Instead of Subbotovskaya-inspired selection procedure (choosing the most frequent variable), we use the weight function introduced by Paterson and Zwick [12], which measures the potential savings for each one-variable restriction, and selects a variable with the biggest savings. Since [12] gives the shrinkage exponent $\Gamma \approx 1.63$, rather than Subbotovskaya's 1.5, this could potentially lead to an improved #SAT algorithm for larger de Morgan formulas.

However, computing the savings, as defined by [12], is NP-hard, as it requires computing the size of a smallest logical formula equivalent to a given one-variable restriction. In fact, the shrinkage result of [12] is *nonconstructive* in the following

sense: the expected shrinkage in size is proved for the minimal logical formula computing the restricted boolean function, rather than for the formula obtained from the original formula using efficiently computable simplification rules. In contrast, the shrinkage results of [16,6] are constructive: the restricted formula is expected to shrink in size when simplified using a certain explicit set of logical rules, so that the new, simplified formula is computable in polynomial time from the original restricted formula.

While the constructiveness of shrinkage is unimportant for proving formula lower bounds, it is crucial for designing shrinkage-based #SAT algorithms for de Morgan formulas, such as those in [14,3,11]. Our main technical contribution is a proof of the *constructive* version of the result in [12]: we give deterministic polynomial-time algorithms for formula simplification and extend the analysis of [12] to show expected shrinkage of formulas with respect to this efficiently computable simplification procedure. The same simplification procedure allows us to choose, in deterministic polynomial-time, which variable should be restricted next. The merit of deterministic variable selection and concentrated and constructive shrinkage, for a shrinkage exponent Γ , is that they yield a deterministic satisfiability algorithm for de Morgan formulas up to size $n^{\Gamma+1-o(1)}$, using an approach of [3].

Namely, once we have this constructive shrinkage result, based on restricting one variable at a time, we apply the martingale-based analysis of [10,3] to derive a *concentrated* version of constructive shrinkage, showing that almost all random settings of the selected variables yield restricted formulas of reduced size, where the restricted formulas are simplified by our efficient procedure. The shrinkage exponent $\Gamma = (5 - \sqrt{3})/2 \approx 1.63$ is the same as in [12]. Using [3], we then get a deterministic #SAT algorithm, running in time $2^{n-n^{o(1)}}$, that works for de Morgan formulas of size up to $n^{\Gamma+1-o(1)} \approx n^{2.63}$.

1.2 Related Work

The deep interplay between lower bounds and satisfiability algorithms has been witnessed in several circuit models. For example, Paturi, Pudlak and Zane [13] give a randomized algorithm for k -SAT running in time $O(n^2 s 2^{n-n/k})$, where n is the number of variables and s is the formula size; they also show that PARITY requires depth-3 circuits of size $\Omega(n^{1/4} 2^{\sqrt{n}})$. More generally, Williams [18] shows that a “better-than-trivial” algorithm for Circuit Satisfiability, for a class \mathcal{C} of circuits, implies a super-polynomial lower bounds against the circuit class \mathcal{C} for some language in NEXP; using this approach, Williams [19] obtains a super-polynomial lower bound against ACC^0 circuits¹ by designing a nontrivial SAT algorithm for ACC^0 circuits.

Following [14], Seto and Tamaki [15] get a nontrivial #SAT algorithm for general linear-size formulas (over an arbitrary basis). Impagliazzo et al. [7] use a generalization of Håstad’s Switching Lemma [5], an analogue of shrinkage for

¹ constant-depth, unbounded fanin circuits, using AND, OR, NOT, and (MOD m) gates, for any integer m .

AC^0 circuits², to give a nontrivial randomized zero-error #SAT algorithm for depth- d AC^0 circuits on n inputs of size up to $2^{n^{1/(d-1)}}$. Beame et al. [2] give a nontrivial deterministic #SAT algorithm for AC^0 circuits, however, only for circuits of much smaller size than that of [7].

Recently, the method of (pseudo) random restrictions has also been used to get pseudorandom generators (yielding additive-approximation #SAT algorithms) for small de Morgan formulas [8] and AC^0 circuits [17].

Remainder of the Paper. We give basic definitions in Section 2. Section 3 contains our efficient formula-simplification procedures. We use these procedures in Section 4 to prove a constructive and concentrated shrinkage result for de Morgan formulas. This is then used in Section 5 to describe and analyze our #SAT algorithm from Theorem 1. Section 6 contains some open questions. Some proofs had to be omitted from this extended abstract due to space limitations; for a more complete version, please see [4].

2 Preliminaries

A (*de Morgan*) *formula* is a binary tree where each leaf is labeled by a literal (a variable x or its negation \bar{x}) or a constant (0 or 1), and each internal node is labeled by \wedge or \vee . A formula naturally computes a boolean function on its input variables.

Let F be a formula with no constant leaves. We define the *size* of F , denoted by $L(F)$, as number of leaves in F . Following [12], we define a *twig* to be a subtree with exactly two leaves. Let $T(F)$ be the number of twigs in F . We define the *weight* of F as $w(F) = L(F) + \alpha \cdot T(F)$, where $\alpha = \sqrt{3} - 1 \approx 0.732$. For convenience, if F is a constant, we define $L(F) = w(F) = 0$. We say F is *trivial* if it is a constant or a literal. Note that we define the size and weight only for formulas which are either constants or with no constant leaves; this is without loss of generality since constants can always be eliminated using a simplification procedure below.

It is easy to see that $L(F) + \alpha \leq w(F) \leq L(F)(1 + \alpha/2)$, since the number of twigs in a formula is at least one and at most half of the number of leaves.

We denote by $F|_{x=1}$ the formula obtained from F by substituting each appearance of x by 1 and \bar{x} by 0; $F|_{x=0}$ is similar. We say a formula \vee -*depends* (\wedge -*depends*) on a literal y if there is a path from the root to a leaf labeled by y such that every internal node on the path (including the root) is labeled by \vee (by \wedge).

3 Formula Simplification Procedures

3.1 Basic Simplification

We define a procedure **Simplify** to eliminate constants, redundant literals and redundant twigs in a formula. The procedure includes the standard constant

² constant-depth, unbounded fanin circuits, using AND, OR, and NOT gates.

simplification rules and a natural extension of the one-variable simplification rules from [6].

Simplify(F):

If F is trivial, done. Otherwise, apply the following transformations whenever applicable. We denote by y a literal and G a subformula.

1. **Constant elimination.**
 - (a) If a subformula is of the form $0 \wedge G$, replace it by 0.
 - (b) If a subformula is of the form $1 \vee G$, replace it by 1.
 - (c) If a subformula is of the form $1 \wedge G$ or $0 \vee G$, replace it by G .
2. **One-variable simplification.**
 - (a) If a subformula is of the form $y \vee G$ and y or \bar{y} appears in G , replace the subformula by $y \vee G|_{y=0}$.
 - (b) If a subformula is of the form $y \wedge G$ and y or \bar{y} appears in G , replace the subformula by $y \wedge G|_{y=1}$.
 - (c) If a subformula G is of the form $G_1 \vee G_2$ for non-trivial G_1 and G_2 , and G \vee -depends on a literal y , then replace G by $y \vee G|_{y=0}$.
 - (d) If a subformula G is of the form $G_1 \wedge G_2$ for non-trivial G_1 and G_2 , and G \wedge -depends on a literal y , then replace G by $y \wedge G|_{y=1}$.

We call a formula *simplified* if it is invariant under **Simplify**. Note that a simplified formula may not be a smallest logically equivalent formula; e.g., $(x \wedge y) \vee (\bar{x} \wedge y)$ is already simplified but it is logically equivalent to y .

The rules 1(a)–(c) and 2(a)–(b) are from [6,14]. Rules 2(c)–(d) are a natural generalization of the one-variable rule of [6], which allow us to eliminate more redundant literals and reduce the formula weight. For example, the formula $(x \vee y) \vee (x \wedge y)$ simplifies to $x \vee y$ under our rules but not the rules in [6,14]. For another example, the formula $(x \vee y) \vee (z \wedge w)$ with weight $4 + 2\alpha$ simplifies to $x \vee (y \vee (z \wedge w))$ with weight $4 + \alpha$.

The next lemma (proof omitted) shows **Simplify** is efficient.

Lemma 1. *Simplify runs in polynomial time.*

3.2 Simplification under All One-Variable Restrictions

Here we consider how a formula simplifies when one of its variables is restricted. Let F be a formula. We define a recursive procedure **RestrictSimplify** which produces a collection of formulas for F under all one-variable restrictions. We denote the output of the procedure by $\{F_y\}$, where y ranges over all literals. Note that each F_y is logically equivalent to $F|_{y=1}$.

The idea behind the transformations in **RestrictSimplify** is the following. When a formula simplifies to a literal under some one-variable restriction, then the formula must be logically equivalent to some special form. For example, if we know that $F|_{x=1}$ simplifies to a literal y , then F itself must be logically equivalent to $(x \wedge y) \vee (\bar{x} \wedge G)$ for some G . This logically equivalent form may help to simplify F under other one-variable restrictions.

RestrictSimplify(F):

If F is a constant c , then let $F_y := c$ for all y . If F is a literal, then let $F_y := F|_{y=1}$ for all y .

If F is $G \vee H$ or $G \wedge H$, recursively call **RestrictSimplify** to compute $\{G_y\}$ and $\{H_y\}$, and initialize each $F_y := \mathbf{Simplify}(G_y \vee H_y)$ or $F_y := \mathbf{Simplify}(G_y \wedge H_y)$, respectively. Then apply the following transformations whenever possible. We suppose there are two literals x and y over distinct variables such that $F_x = y$.

1. If $F_{\bar{x}} = y$, then let $F_w := y|_{w=1}$ for every literal w .
2. If $F_{\bar{x}} = z$ for some literal $z \notin \{x, \bar{x}, y\}$, then let $F_w := \mathbf{Simplify}((x \wedge y) \vee (\bar{x} \wedge z)|_{w=1})$ for every literal w .
3. (a) If neither x nor \bar{x} appears in F_y , then let $F_y := 1$; (b) otherwise, let $F_y := \mathbf{Simplify}(x \vee (F_y|_{x=0}))$.
4. (a) If neither x nor \bar{x} appears in $F_{\bar{y}}$, then let $F_{\bar{y}} := 0$; (b) otherwise, let $F_{\bar{y}} := \mathbf{Simplify}(\bar{x} \wedge (F_{\bar{y}}|_{x=0}))$.
5. For $z \notin \{x, \bar{x}, y, \bar{y}\}$, if neither x nor \bar{x} appears in F_z , then let $F_z := y$.

Correctness of RestrictSimplify. The above transformations are based on logical implications. In case 1, $F_x = F_{\bar{x}} = y$ implies that $F \equiv y$. In case 2, $F_x = y$ and $F_{\bar{x}} = z$ implies that $F \equiv (x \wedge y) \vee (\bar{x} \wedge z)$. Note that in this case z might be \bar{y} . In case 3, we have $F_y|_{x=1} \equiv F_x|_{y=1} = 1$; if neither x nor \bar{x} appears in F_y then $F_y = F_y|_{x=1} \equiv 1$, otherwise $F_y \equiv x \vee (F_y|_{x=0})$. Case 4 is dual to case 3. In case 5, if neither x nor \bar{x} appears in F_z then $F_z = F_z|_{x=1} \equiv F_x|_{z=1} = y$.

Remark 1. It is possible to introduce more simplifications rules in **RestrictSimplify**, e.g., when F_x is a constant for some literal x , or when, in case 5, x or \bar{x} appears in F_z ³. However, such simplifications are not needed for our proof of constructive shrinkage.

It is easy to show that **RestrictSimplify** is efficient.

Lemma 2. **RestrictSimplify** runs in polynomial time.

The *solo structure* of a formula F is the relation on literals defined by $x \Rightarrow y$ if $F_x = y$, where the collection of formulas $\{F_x\}$ is produced by the procedure **RestrictSimplify**. The following lemma gives all possible solo structures; it resembles the characterization of solo structures for boolean functions from [12].

Lemma 3. *The solo structure of a non-trivial formula F must be in one of the following forms:*

- (i) the empty relation,
- (ii) there exists y such that for all literals $x \notin \{y, \bar{y}\}$ we have $x \Rightarrow y$ in the relation,
- (iii) $\{x_1 \Rightarrow y, \dots, x_k \Rightarrow y\}$ for some $k \geq 1$ and x_i 's are over distinct variables,
- (iv) $\{x \Rightarrow y, y \Rightarrow x, \bar{x} \Rightarrow \bar{y}, \bar{y} \Rightarrow \bar{x}\}$,
- (v) $\{x \Rightarrow y, \bar{x} \Rightarrow z\}$,
- (vi) $\{x \Rightarrow y, y \Rightarrow x\}$,
- (vii) $\{x \Rightarrow y, \bar{y} \Rightarrow \bar{x}\}$.

³ then we could let $F_z := (x \wedge y) \vee (\bar{x} \wedge \mathbf{Simplify}(F_z|_{x=0}))$.

4 Constructive and Concentrated Shrinkage

Here we prove a constructive and concentrated version of the shrinkage result from [12]. For each literal y of a given formula F , we define the savings (reduction in weight of F) when we replace F by the new formula F_y , as computed by the procedure **RestrictSimplify**. We first prove that the lower bound on the average savings (over all variables of F) shown by [12] continues to hold with respect to our efficiently computable one-variable restrictions F_y .

4.1 Average Savings under One-Variable Restrictions

Assume a formula F is simplified; otherwise, let $F := \mathbf{Simplify}(F)$. For a formula F and a literal y , we define $\sigma_y(F) = w(F) - w(F_y)$, where F_y is produced by **RestrictSimplify**. Let $\sigma(F) = \sum_x (\sigma_x(F) + \sigma_{\bar{x}}(F))$, where the summation ranges over all variables of F . The quantity $\sigma(F)$ measures the total *savings* under all one-variable restrictions.

Theorem 2. *For any formula F , it holds that $\sigma(F)/w(F) \geq 2\gamma$, where $\gamma = (5 - \sqrt{3})/2 \approx 1.63$.*

The proof is by induction, as in [12]. The difficulty here is that we need to apply the “syntactic simplifications” defined by the procedure **RestrictSimplify**, instead of using the smallest logically equivalent formulas as in [12].

For the base case, the following lemma can be proved by enumerating all possible formulas of size at most 4 (the proof is omitted).

Lemma 4. *For any simplified F of size at most 4, we have $\sigma(F)/w(F) \geq 2\gamma$.*

For formulas of size larger than 4, we consider whether one child of the root is trivial. Without loss of generality, we assume the root is labeled by \vee ; the other case is dual. The following lemma considers if one child of the root is trivial. The proof is omitted here but it is similar to [12].

Lemma 5. *If F is a simplified formula of the form $x \vee G$ for some literal x and subformula G , and $L(F) \geq 5$, then $\sigma(F)/w(F) \geq 2\gamma$.*

Now we consider formulas where both children of the root are non-trivial.

Lemma 6. *Suppose F is of the form $G \vee H$ with $L(F) \geq 5$ and G, H are non-trivial. Then $\sigma(F)/w(F) \geq 2\gamma$.*

Intuitively, we need to take care of the cases where both G and H simplify to literals on distinct variables (thereby forming a new twig); otherwise the result holds by the induction hypothesis. Suppose $G_x \vee H_x$ is a twig for some literal x . Then $\sigma_x(F) = \sigma_x(G) + \sigma_x(H) - \alpha$, i.e., we get the savings from restricting x in G and H , but then need to pay the penalty α for the twig created. We will argue that there are “extra savings” from restricting other literals in the formula F that can be used to compensate for the penalty α at x .

Proof. We shall need the following basic property of **RestrictSimplify**.

Claim. For $F = G \vee H$ or $F = G \wedge H$, we have $w(F_y) \leq w(G_y) + w(H_y)$, for all literals y except those where G_y and H_y are literals over distinct variables.

Proof (of Claim). Let $F = G \vee H$; the other case is identical. For $F_y := \mathbf{Simplify}(G_y \vee H_y)$, the required inequality holds initially. All transformations, except 3(b) and 4(b), produce the smallest logically equivalent formula; rules 3(b) and 4(b) do not increase the weight of the formula. \square

We first prove that, for a literal x , if G_x and H_x are not literals over distinct variables, then $\sigma_x(F) \geq \sigma_x(G) + \sigma_x(H)$. Indeed, since $w(F) = w(G) + w(H)$, this follows from $w(F_x) \leq w(G_x) + w(H_x)$, which holds by the claim above.

Next, let k be the number of different literals x such that $G_x \vee H_x$ is a twig (i.e., G_x and H_x are literals over distinct variables). Thus there are k twigs created as we consider all possible one-variable restrictions. We will argue that, for different cases of k , the weight $k\alpha$ of these new twigs can be compensated from savings in other restrictions.

Case $k = 0$: We have $\sigma_y(F) \geq \sigma_y(G) + \sigma_y(H)$ for all literals y , and thus $\sigma(F) \geq \sigma(G) + \sigma(H)$. The result is by the induction hypothesis on G and H .

Case $1 \leq k \leq 2$: Let x be such that $G_x = y$ and $H_x = z$. Without loss of generality, assume x, y, z are distinct variables. Consider F under the restrictions $y = 1$ and $z = 1$. We will argue that the extra savings from applying **Simplify** on $G_y \vee H_y$ and $G_z \vee H_z$ are at least $2 > k\alpha$.

Since $G_x = y$, transformation 3(a)–(b) in **RestrictSimplify** guarantee that either G_y is constant 1 or it \vee -depends on x . Similarly either H_z is constant 1 or it \vee -depends on x . Since $H_y|_{x=1} \equiv H_x|_{y=1} = z$, we get that H_y is not a constant (it depends on z), and if it is a literal it must be z . Similarly G_z is not a constant (it depends on y), and if it is a literal it must be y .

We first consider the case that either G_y or H_z is constant 1. If $G_y = H_z = 1$, then there are at least 2 savings from simplifying $G_y \vee H_y$ and $G_z \vee H_z$ by eliminating constants. If $G_y = 1$ and H_y is not a literal, then there are at least 2 savings from simplifying $G_y \vee H_y$. If $G_y = 1$, $H_y = z$ and $H_z \neq 1$, we first have one saving from simplifying $G_y \vee H_y$; then since $H_y = z$ and $H_z \neq 1$, by the transformation 3(b) in **RestrictSimplify** H_z \vee -depends on y , and since G_z depends on y , we get another saving from simplifying $G_z \vee H_z$. The cases where $H_z = 1$ are similar.

Next we consider that both G_y and H_z \vee -depends on x . In the following we analyze different possibilities for H_y and G_z .

- If x appears in both H_y and G_z , then there are at least 2 savings from simplifying $G_y \vee H_y$ and $G_z \vee H_z$ by eliminating x .
- If x appears in H_y but not G_z , then by the transformation 5 in **RestrictSimplify** we have $G_z = y$, and thus G_y \vee -depends on both x and z . Then since H_y depends on both x and z , we have two savings from simplifying $G_y \vee H_y$ by eliminating both x and z from H_y .

- If x appears in G_z but not H_y , this is similar to the previous case.
- If x appears in neither H_y nor G_z , then by the transformation 5 in **RestrictSimplify** we have $G_z = y$ and $H_y = z$. Thus G_y \vee -depends on both x and z , and H_z \vee -depends on both x and y . Therefore we have at least 2 savings, one from simplifying $G_y \vee H_y$ by eliminating z , and another from simplifying $G_z \vee H_z$ by eliminating y .

Case $k \geq 3$: By Lemma 3, the solo structure of G and H must be one of cases (ii), (iii), or (iv).

First assume that either G or H is in case (ii) of Lemma 3. Without loss of generality, suppose G is in case (ii); then G is logically equivalent to a literal y but itself is non-trivial, which implies that $w(G) \geq 4 + \alpha$. (The smallest non-trivial, simplified formula equivalent to a literal has size at least 4). We have that $w(G_z) = 1$ for at least k literals $z \notin \{y, \bar{y}\}$, and $w(G_y) = w(G_{\bar{y}}) = 0$. Then by the fact that $w(F) = w(G) + w(H)$ and the induction hypothesis on H , we have

$$\begin{aligned} \sigma(F) &\geq k(w(G) - 1) + 2w(G) + \sigma(H) - k\alpha \\ &\geq 2\gamma \cdot w(F) + (2 + k - 2\gamma)w(G) - k(1 + \alpha) \geq 2\gamma \cdot w(F). \end{aligned}$$

If both G and H are in case (iv), then, under each restriction, they reduce to literals on the same variable. Since in case (iii) all x_i 's are over distinct variables, it is not possible that one of G and H is in case (iv) while the other is in case (iii). Thus, we now only need to analyze if both G and H are in case (iii).

Without loss of generality, suppose that x_1, \dots, x_k, y, z are distinct variables such that $G_{x_i} = y$ and $H_{x_i} = z$ for $i = 1, \dots, k$. By the transformation 3 in **RestrictSimplify**, either $G_y = 1$ or G_y \vee -depends on x_1, \dots, x_k ; and H_z is similar.

If every x_i appears in H_y , then there are k savings from simplifying $G_y \vee H_y$ by eliminating x_i 's. Similarly, if every x_i appears in G_z , there are also k savings from simplifying $G_z \vee H_z$.

If some x_i does not appear in H_y and some x_i does not appear in G_z . By the transformation 5 in **RestrictSimplify**, we have $H_y = z$ and $G_z = y$. Therefore,

$$\begin{aligned} \sigma_{x_i}(F) &= w(F) - (2 + \alpha), \quad i = 1, \dots, k \\ \sigma_y(F) &\geq 1 + (w(H) - 1) = w(H) \\ \sigma_z(F) &\geq 1 + (w(G) - 1) = w(G) \\ \sum_v \sigma_{\bar{v}}(F) &\geq L(F) \geq w(F)/(1 + \alpha/2), \quad v \text{ ranges over all variables of } F \end{aligned}$$

Summing the above cases together yields $\sigma(F) \geq 2\gamma \cdot w(F)$. □

Proof (Theorem 2). The proof is by combining the base case in Lemma 4 and the two inductive cases in Lemma 5 and Lemma 6. □

4.2 Concentrated Shrinkage

Theorem 2 characterizes the average shrinkage of the weight of a formula when a randomly chosen literal is restricted. Given a formula F on n variables, if we randomly pick one variable and randomly assign it 0 or 1, the weight of the restricted formula (produced by **RestrictSimplify**) reduces by at least $\gamma \cdot w(F)/n$ on average.

The procedure **RestrictSimplify** also allows us to deterministically pick the variable with the best savings in polynomial time. That is, given a formula F , we run **RestrictSimplify** to produce a collection of formulas $\{F_y\}$, and then pick a variable x such that $\sigma_x(F) + \sigma_{\bar{x}}(F)$ is maximized. We show that randomly restricting such a variable significantly reduces the expected weight of the simplified formula.

Lemma 7. *Let F be a formula on n variables. Let x be the variable such that $\sigma_x(F) + \sigma_{\bar{x}}(F)$ is maximized. Let F' be F_x or $F_{\bar{x}}$ with equal probability. Then we have $w(F') \leq w(F) - 1$ and $\mathbf{E}[w(F')] \leq (1 - \frac{1}{n})^\gamma \cdot w(F)$.*

Proof. Restricting one variable eliminates at least one leaf; therefore $w(F') \leq w(F) - 1$. By Theorem 2, $n(\sigma_x(F) + \sigma_{\bar{x}}(F)) \geq \sigma(F) \geq 2\gamma \cdot w(F)$. Then we have $\mathbf{E}[w(F')] = w(F) - \frac{1}{2}(\sigma_x(F) + \sigma_{\bar{x}}(F)) \leq (1 - \frac{\gamma}{n}) \cdot w(F) \leq (1 - \frac{1}{n})^\gamma \cdot w(F)$. \square

Next we use the martingale-based analysis from [10,3] to derive a “high-probability shrinkage” result from Lemma 7. Let $F_0 = F$ be a formula on n variables. For $1 \leq i \leq n$, let F_i be the (random) formula obtained from F_{i-1} by assigning the variable with the best savings with a random value $R_i \in \{0, 1\}$. The following Lemma shows the weight of a given de Morgan formula reduces with high probability under the restriction process. The proof, which is similar to [3], is omitted here due to space constraints.

Lemma 8 (Concentrated weight shrinkage). *For any de Morgan formula F on n variables and any $k > 10$, $\Pr \left[w(F_{n-k}) \geq 2 \cdot w(F) \cdot \left(\frac{k}{n}\right)^\gamma \right] < 2^{-k/10}$.*

Finally, by $w(F)/(1 + \alpha/2) \leq L(F) \leq w(F)$ for all F , we get from Lemma 8 the desired concentrated constructive shrinkage with respect to the restriction process defined above.

Corollary 1 (Concentrated constructive shrinkage). *Let F be an arbitrary de Morgan formula. There exist constants $c, d > 1$ such that, for any $k > 10$, $\Pr \left[L(F_{n-k}) \geq c \cdot L(F) \cdot \left(\frac{k}{n}\right)^\gamma \right] < 2^{-k/d}$.*

5 #SAT Algorithm for $n^{2.63}$ -size de Morgan Formulas

Here we prove our main result.

Theorem 3. *There is a deterministic algorithm for counting the number of satisfying assignments in a given formula on n variables of size at most $n^{2.63}$ which runs in time $t(n) \leq 2^{n-n^\delta}$, for some constant $0 < \delta < 1$.*

Proof. Suppose we have a formula F on n variables of size $n^{1+\gamma-\epsilon}$ for a small constant $\epsilon > 0$. Let $k = n^\alpha$ such that $\alpha < \epsilon/\gamma$. We build a restriction decision tree with 2^{n-k} branches as follows:

Starting with F at the root, run **RestrictSimplify** to produce a collection $\{F_y\}$, pick the variable x which will make the largest reduction in the weight of the current formula. Make the two formulas F_x and $F_{\bar{x}}$ the children of the current node. Continue recursively on F_x and $F_{\bar{x}}$ until get a full binary tree of depth exactly $n - k$.

Note that constructing this decision tree takes time $2^{n-k}\text{poly}(n)$, since the procedure **RestrictSimplify** runs in polynomial time. By Corollary 1, all but at most $2^{-k/d}$ fraction of the leaves have the formula size $L(F_{n-k}) < c \cdot L(F) \left(\frac{k}{n}\right)^\gamma = cn^{1-\epsilon+\gamma\alpha}$.

To solve #SAT for all “big” formulas (those that haven’t shrunk), we use brute-force enumeration over all possible assignments to the k free variables left. The running time is at most $2^{n-k} \cdot 2^{-k/d} \cdot 2^k \cdot \text{poly}(n) \leq 2^{n-k/d} \cdot \text{poly}(n)$.

For “small” formulas (those that shrunk to the size less than $cn^{1-\epsilon+\gamma\alpha}$), we use memoization. First, we enumerate all formulas of such size, and compute and store the number of satisfying assignments for each of them. Then, as we go over the leaves of the decision tree that correspond to small formulas, we simply look up the stored answers for these formulas.

There are at most $2^{O(n^{1-\epsilon+\gamma\alpha})}\text{poly}(n)$ such formulas, and counting the satisfying assignments for each one (with k inputs) takes time $2^k\text{poly}(n)$. Including pre-processing, computing #SAT for all small formulas takes time at most $2^{n-k} \cdot \text{poly}(n) + 2^{O(n^{1-\epsilon+\gamma\alpha})} \cdot 2^k \cdot \text{poly}(n) \leq 2^{n-k} \cdot \text{poly}(n)$.

The overall running time of our #SAT algorithm is bounded by 2^{n-n^δ} for some $\delta > 0$. □

6 Open Questions

The main open problem is to get a nontrivial deterministic #SAT algorithm for de Morgan formulas of size up to $n^{3-o(1)}$. Can one derandomize the zero-error algorithm of [11] that is based on Håstad’s shrinkage result [6]?

Can one improve the analysis of the shrinkage result of [12] (by considering more general patterns than just twigs), getting a better shrinkage exponent? If so, this could lead to a deterministic #SAT algorithm for larger de Morgan formulas.

References

1. Andreev, A.E.: On a method of obtaining more than quadratic effective lower bounds for the complexity of π -schemes. Vestnik Moskovskogo Universiteta. Matematika 42(1), 70–73 (1987); english translation in Moscow University Mathematics Bulletin

2. Beame, P., Impagliazzo, R., Srinivasan, S.: Approximating AC^0 by small height decision trees and a deterministic algorithm for $\#AC^0SAT$. In: Proceedings of the Twenty-Seventh Annual IEEE Conference on Computational Complexity, pp. 117–125 (2012)
3. Chen, R., Kabanets, V., Kolokolova, A., Shaltiel, R., Zuckerman, D.: Mining circuit lower bound proofs for meta-algorithms. In: Proceedings of the Twenty-Ninth Annual IEEE Conference on Computational Complexity (2014)
4. Chen, R., Kabanets, V., Saurabh, N.: An improved deterministic $\#SAT$ algorithm for small De Morgan formulas. *Electronic Colloquium on Computational Complexity (ECCC)* 20, 150 (2013)
5. Håstad, J.: Almost optimal lower bounds for small depth circuits. In: Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing, pp. 6–20 (1986)
6. Håstad, J.: The shrinkage exponent of de Morgan formulae is 2. *SIAM Journal on Computing* 27, 48–64 (1998)
7. Impagliazzo, R., Matthews, W., Paturi, R.: A satisfiability algorithm for AC^0 . In: Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 961–972 (2012)
8. Impagliazzo, R., Meka, R., Zuckerman, D.: Pseudorandomness from shrinkage. In: Proceedings of the Fifty-Third Annual IEEE Symposium on Foundations of Computer Science, pp. 111–119 (2012)
9. Impagliazzo, R., Nisan, N.: The effect of random restrictions on formula size. *Random Structures and Algorithms* 4(2), 121–134 (1993)
10. Komargodski, I., Raz, R.: Average-case lower bounds for formula size. In: Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, pp. 171–180 (2013)
11. Komargodski, I., Raz, R., Tal, A.: Improved average-case lower bounds for DeMorgan formula size. In: Proceedings of the Fifty-Fourth Annual IEEE Symposium on Foundations of Computer Science, pp. 588–597 (2013)
12. Paterson, M., Zwick, U.: Shrinkage of de Morgan formulae under restriction. *Random Structures and Algorithms* 4(2), 135–150 (1993)
13. Paturi, R., Pudlák, P., Zane, F.: Satisfiability coding lemma. *Chicago Journal of Theoretical Computer Science* (1999)
14. Santhanam, R.: Fighting peregbor: New and improved algorithms for formula and QBF satisfiability. In: Proceedings of the Fifty-First Annual IEEE Symposium on Foundations of Computer Science, pp. 183–192 (2010)
15. Seto, K., Tamaki, S.: A satisfiability algorithm and average-case hardness for formulas over the full binary basis. In: Proceedings of the Twenty-Seventh Annual IEEE Conference on Computational Complexity, pp. 107–116 (2012)
16. Subbotovskaya, B.: Realizations of linear function by formulas using \vee , $\&$, \neg . *Doklady Akademii Nauk SSSR* 136(3), 553–555 (1961); english translation in *Soviet Mathematics Doklady*
17. Trevisan, L., Xue, T.: A derandomized switching lemma and an improved derandomization of AC^0 . In: Proceedings of the Twenty-Eighth Annual IEEE Conference on Computational Complexity, pp. 242–247 (2013)
18. Williams, R.: Improving exhaustive search implies superpolynomial lower bounds. In: Proceedings of the Forty-Second Annual ACM Symposium on Theory of Computing, pp. 231–240 (2010)
19. Williams, R.: Non-uniform ACC circuit lower bounds. In: Proceedings of the Twenty-Sixth Annual IEEE Conference on Computational Complexity, pp. 115–125 (2011)