

A SATISFIABILITY ALGORITHM AND AVERAGE-CASE HARDNESS FOR FORMULAS OVER THE FULL BINARY BASIS

KAZUHISA SETO AND SUGURU TAMAKI

Abstract. We present a moderately exponential time algorithm for the satisfiability of Boolean formulas over the full binary basis. For formulas of size at most cn , our algorithm runs in time $2^{(1-\mu_c)n}$ for some constant $\mu_c > 0$. As a byproduct of the running time analysis of our algorithm, we obtain strong average-case hardness of affine extractors for linear-sized formulas over the full binary basis.

Keywords. Satisfiability, exact algorithm, formula, parity gate, average-case hardness

Subject classification. 68Q17, 68Q25.

1. Introduction

In this paper, we present a moderately exponential time algorithm for *the satisfiability of Boolean formulas over the full binary basis*, which is an interesting special case of the Circuit Satisfiability (Circuit SAT) problem. Circuit SAT is, given a Boolean circuit C with n input variables, to determine whether there exists a 0/1 assignment to the input variables such that C outputs 1. It is one of the most fundamental and important NP-complete problems and people have developed many efficient algorithms in both practical and theoretical sense. It is easy to see that one can solve the problem in time $\text{poly}(|C|)2^n$ by brute force search where $|C|$ denotes the size of C . An obvious question is whether there exist moderately exponential time algorithms, i.e., algorithms with the worst case running time of the form $\text{poly}(|C|)2^{(1-\mu)n}$ for some $\mu > 0$.

It is too difficult to answer the above question because of the generality of Circuit SAT, that is, many combinatorial problems can be represented as Circuit SAT, see, e.g., [Biere *et al.* (2009)]. Instead of considering Circuit SAT in the most general form, we may investigate the complexity of Circuit SAT over some restricted circuit class \mathcal{C} . We write such restricted Circuit SAT as \mathcal{C} -SAT. The most well studied restricted circuit class is k -CNF-formulas, which consist of a conjunction of clauses, where each clause is a disjunction of at most k literals. k -CNF-SAT is a central problem in the area of exact exponential algorithms and many efficient algorithms for it have been developed over the past 30 years, see, e.g., [Hertli (2011); Makino *et al.* (2011); Monien & Speckenmeyer (1985); Moser & Scheder (2011); Paturi *et al.* (2005); Schöning (1999)], and an excellent survey by [Dantsin & Hirsch (2009)]. The best running time upper bound is of the form $\text{poly}(|C|)2^{(1-\mu_k)n}$, where $\mu_k > 0$ is some constant only depending on k . Despite the success of exact algorithms for k -CNF-SAT, there are few works studying the exponential time complexity of Circuit SAT over more general circuit classes until recently.

Let us quickly review some results on \mathcal{C} -SAT for more general \mathcal{C} . We are aware of the works for

- CNF-formulas (without restriction on length of each clause) by [Arvind & Schuler (2003); Calabro *et al.* (2006); Dantsin *et al.* (2004, 2006); Pudlák (1998); Schuler (2005)] (see also [Hirsch (2008)]),
- \mathbf{AC}^0 circuits by [Calabro *et al.* (2009); Impagliazzo *et al.* (2012)],
- \mathbf{ACC}^0 circuits by [Williams (2011)], and
- U_2 -formulas (De Morgan formulas) by [Santhanam (2010)].

Here, \mathbf{AC}^0 circuits are constant depth circuits over the basis {and, or, not}, where the fan-in of each gate is unbounded, \mathbf{ACC}^0 circuits are the same as \mathbf{AC}^0 circuits except that the basis also contains arbitrary modulo gates of unbounded fan-in, and U_2 -formulas are formulas over the basis $U_2 = \{\text{and, or, not}\}$, where the fan-in of

{and, or} is two. The following summarizes the current best algorithms:

- For CNF-formulas with m clauses, [Calabro *et al.* (2006)] have shown that Circuit SAT can be solved in time $|C|2^{(1-1/\log(m/n))n}$ (see also [Dantsin & Hirsch (2009)]),
- for \mathbf{AC}^0 circuits of size cn and depth d , [Impagliazzo *et al.* (2012)] have shown that Circuit SAT can be solved in time $|C|2^{(1-1/O(\log c+d\log d))n}$,
- for \mathbf{ACC}^0 circuits of depth d , [Williams (2011)] has shown that Circuit SAT can be solved in time $|C|2^{n-\Omega(n^{2^{-O(d)}})}$, and
- for U_2 -formulas of size cn , [Santhanam (2010)] have shown that Circuit SAT can be solved in time $|C|2^{(1-1/c^{O(1)})n}$.

In this paper, we extend the result of [Santhanam (2010)] to the case of B_2 -formulas which are formulas over the full binary basis B_2 consisting of all two-variable functions. Our main result is the following:

THEOREM 1.1. *There is a deterministic algorithm for B_2 -formula-SAT which runs in time $2^{(1-\mu_c)n}$ on formulas of size at most cn . Here $\mu_c > 0$ is a constant only depending on c (roughly $\mu_c = 2^{-\Theta(c^3)}$).*

Santhanam's result has an application in proving strong average-case hardness of the parity function against linear-sized U_2 -formulas. From the proof of Theorem 1.1, we can show an analogous result, strong average-case hardness of affine extractors against linear-sized B_2 -formulas (affine extractors are formally defined in Section 4, see Definition 4.2).

THEOREM 1.2 (Informal). *For any constant $c > 0$, any sequence of B_2 -formulas of size at most cn must err in computing affine extractors on at least a $1/2 - 2^{-\Omega(n)}$ fraction of inputs of length n for each n .*

1.1. Background. In this section, we discuss the motivation of designing moderately exponential time algorithms for \mathcal{C} -SAT with more general \mathcal{C} .

Encoding practical instances. One of the motivations comes from practical applications. Because of its expressibility, Circuit SAT can represent many industrial problems such as software and hardware verification and testing, design automation, planning and automated reasoning in a natural way, see, e.g., [Biere *et al.* (2009)]. This motivates the development of faster SAT solvers for instances from practice and today we have very sophisticated SAT solvers which can treat instances of relatively large size. However, most SAT solvers require their input to be in CNF form although natural encoding of industrial problems such as hardware verification to Circuit SAT often results in instances represented by general circuits, e.g., formulas with no depth restriction, circuits with parity gates etc. To use fast SAT solvers, first we need to transform the original instances into CNF form. After such transformation, the size of instances must increase and in some case the size blow up can be exponential. For example, parity functions have linear size representation in formulas over the full binary basis, but requires quadratic size in De Morgan formulas and exponential size in CNF form. One can use the Tseitin transformation to avoid such huge blowups, but additional variables must be introduced in the transformation, which is very expensive in the context of exponential time algorithms. Thus, it is more desirable if one can develop efficient algorithms which can treat the original encoding of practical instances.

Proof Techniques. The analysis of running time savings of the best known SAT algorithms for certain circuit classes such as \mathbf{AC}^0 and De Morgan formulas follows from proof techniques for corresponding circuit lower bounds. That is, \mathbf{AC}^0 circuits and De Morgan formulas shrink their sizes significantly by “random restrictions”. Roughly speaking, random restriction chooses, say, $(1 - \rho)n$ variables for some $\rho > 0$ randomly and sets random 0/1 values to the chosen variables, and obtains a simplified circuit/formula over the remaining ρn variables. It is well known

that by appropriately choosing ρ , random restriction can collapse \mathbf{AC}^0 circuits and De Morgan formulas into a constant function with high probability, see, e.g., [Ajtai (1983); Furst *et al.* (1984); Håstad (1986); Yao (1985)], and [Andreev (1987); Håstad (1998); Impagliazzo & Nisan (1993); Subbotovskaya (1961)]. This implies that such circuit classes cannot compute parity functions since such functions remain non-constant functions after a random restriction is applied. The above lower bounds argument suggests that backtracking algorithms work well because expected depth of each path in backtracking tree is at most $(1 - \rho)n$ with high probability. Unfortunately, for our target class, formulas over the basis {and, or, xor}, random restriction cannot prove interesting lower bounds. It is easy to see that a formula consisting of only xor gates does not shrink into a constant unless we set values to all the variables. To achieve similar savings as Santhanam’s result for De Morgan formulas, our algorithm and its analysis require new ideas. We establish a structural result for B_2 -formulas of linear size, and use it to solve SAT.

SAT algorithm implies circuit lower bounds. As discussed in the previous paragraph, design and analysis of \mathcal{C} -SAT algorithm is often inspired by the corresponding circuit lower bound technique for \mathcal{C} . Interestingly, the connection also holds in the reverse order, that is, efficient SAT algorithms implies circuit lower bounds. One of such examples is the result by [Paturi *et al.* (1999)] showing tight lower bounds for depth three \mathbf{AC}^0 circuits computing parity functions. They exploited the connection between the success probability of their SAT algorithm for k -CNF formulas and bounds on the number of sub-circuits in depth three \mathbf{AC}^0 circuits. Another example is a recent breakthrough result due to Williams. In [Williams (2010)] he proves a generic result that SAT algorithm for the circuit class \mathcal{C} with “non-trivial” running time implies that \mathcal{C} does not contain NEXP, the class of languages computable in non-deterministic exponential time. Then in [Williams (2011)] he shows non-trivial SAT algorithms for \mathbf{ACC}^0 to conclude that \mathbf{ACC}^0 does not contain NEXP. Thus, developing SAT algorithms for richer circuit classes is tied to proving lower bounds for those classes.

Strong average-case hardness. We are often interested in proving strong average-case hardness for some circuit class \mathcal{C} rather than just proving worst-case hardness. Here, strong means that any circuit in \mathcal{C} must fail to compute the given function on at least $1/2 - 2^{-\Omega(n)}$ fraction of inputs. Such a hardness result can be used to construct very efficient pseudorandom generators for \mathcal{C} . For example, [Nisan & Wigderson (1994)] have shown that if there exists a family of functions computable in $E = \text{DTIME}(2^{O(n)})$ such that any subexponential size circuits must fail on at least $1/2 - 2^{-\Omega(n)}$ fraction of inputs, then $P = \text{BPP}$ holds. In general, proving strong (exponential) circuit lower bounds even in the worst-case is a very difficult task for even relatively weak circuit classes such as \mathbf{AC}^0 . However, if we can construct \mathcal{C} -SAT algorithms which run in time $2^{n-\Omega(n)}$, they often provide strong average-case lower bounds for \mathcal{C} and actually there are such (but) few results: [Calabro *et al.* (2009); Impagliazzo *et al.* (2012)] and [Santhanam (2010)], respectively show that the parity function is strongly hard for linear-sized \mathbf{AC}^0 circuits and U_2 formulas, respectively. (We remark that [Lu & Wu (2010)] also obtained strong average-case lower bounds for linear-sized \mathbf{AC}^0 circuits using different techniques.) Our result adds such rare hardness results in the case of linear-sized B_2 -formulas.

Paper organization. In the rest of our paper, we provide detailed algorithms and analysis to support our results. In section 2, we present some useful properties of B_2 -formulas, which play an important role in designing our satisfiability algorithm. In section 3, we give a high level idea, formal description and running time analysis of our algorithm. In section 4, we prove strong average-case hardness results.

2. Preliminaries

Let B_2 be the set of all Boolean functions of two variables. A B_2 -formula is a rooted binary tree in which each leaf is labeled by a literal from the set $\{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$ or a constant from $\{0, 1\}$ and each internal node is labeled by a function from B_2 . Given a B_2 -formula ϕ , a *subformula* of ϕ is a B_2 -formula which is a subtree in ϕ . By ϕ_v , we denote ϕ 's subformula whose root node is

v. Every B_2 -formula computes in a natural way a Boolean function from $\{0, 1\}^n$ to $\{0, 1\}$. The *size* of a B_2 -formula ϕ is defined to be the number of leaves in it, and it is denoted by $L(\phi)$. We denote by $\text{var}(\phi)$ the set of variables which appear as literals in ϕ . The *frequency* of a variable x in ϕ is defined to be the number of leaves labelled by x or \bar{x} , and it is denoted by $\text{freq}_\phi(x)$. We often omit the subscript ϕ when it is clear from the context. A $\{\wedge, \vee, \oplus\}$ -*formula* is a B_2 -formula in which each internal node is labeled by \wedge (“and”) or \vee (“or”) or \oplus (“xor”). It is easy to see that the following holds by using De Morgan’s laws and the fact that $\overline{\phi_1 \oplus \phi_2} = \bar{\phi}_1 \oplus \bar{\phi}_2$.

FACT 2.1. *For any B_2 -formula ϕ , there exists a $\{\wedge, \vee, \oplus\}$ -formula $\tilde{\phi}$ such that $\tilde{\phi}$ computes the same function as ϕ and $L(\tilde{\phi}) \leq L(\phi)$. Furthermore, we can obtain $\tilde{\phi}$ from ϕ in polynomial time in $L(\phi)$.*

PROOF. We modify each internal node u of ϕ from the root to leaves in breadth first search manner so that u does not have a label from $B_2 \setminus \{\wedge, \vee, \oplus\}$. Let $f(x, y)$ denote the label of u , let v, w denote the children of u . Since B_2 contains 16 different functions, we have to consider the following cases:

(i) $f(x, y) = a$ for $a \in \{0, 1\}$. In this case, we replace a subformula ϕ_u by a node u whose label is a .

(ii) $f(x, y) = x \oplus a$ ($f(x, y) = y \oplus a$, respectively) for $a \in \{0, 1\}$: In this case, we replace a subformula ϕ_u by a subformula ϕ_v (by a subformula ϕ_w , respectively) and then replace the label g of v (w , respectively) by $g \oplus a$.

(iii) $f(x, y) = (x \oplus a) \wedge (y \oplus b)$ ($f(x, y) = (x \oplus a) \vee (y \oplus b)$, respectively) for $a, b \in \{0, 1\}$: In this case, we replace the label of u by $x \wedge y$ ($x \vee y$, respectively) and the label g of v by $g \oplus a$ and the label h of w by $h \oplus b$.

(iv) $f(x, y) = ((x \oplus a) \wedge (y \oplus b)) \oplus 1$ ($f(x, y) = ((x \oplus a) \vee (y \oplus b)) \oplus 1$, respectively) for $a, b \in \{0, 1\}$: In this case, we replace the label of u by $x \vee y$ ($x \wedge y$, respectively) and the label g of v by $g \oplus \bar{a}$ and the label h of w by $h \oplus \bar{b}$.

(v) $f(x, y) = x \oplus y \oplus a$ for $a \in \{0, 1\}$: In this case, we replace the label of u by $x \oplus y$ and the label g of v by $g \oplus a$.

The resulting formula $\tilde{\phi}$ computes the same function as ϕ and $L(\tilde{\phi}) \leq L(\phi)$ as desired. \square

In our SAT algorithm for B_2 -formulas, we assume without loss of generality that input formulas are always $\{\wedge, \vee, \oplus\}$ -formulas. In what follows, we often write just formula instead of $\{\wedge, \vee, \oplus\}$ -formula for brevity. For convenience, we think of constant functions 0 and 1 as formulas.

2.1. Maximal linear nodes. Given a formula ϕ , we define the notion of *linear* node in ϕ inductively as follows. (i) Every leaf is linear. (ii) An internal node is linear if it is labeled by \oplus and both of its child nodes are linear. Then, a linear node is *maximal* if its parent node is not linear. Let v be a linear node in ϕ , ϕ_v be the subformula rooted by v and $S = \{y_1, \dots, y_k\}$ be the set of all leaves in ϕ_v . It is easy to see that ϕ_v computes a linear function $y_1 \oplus \dots \oplus y_k$. If there exists a variable x which appears in S more than once (as x or \bar{x}), ϕ_v is called *redundant*, otherwise it is called *irredundant*. By the commutativity of \oplus and the fact that $0 \oplus y = y$, $1 \oplus y = \bar{y}$, $y \oplus y = 0$ and $y \oplus \bar{y} = 1$ for any literal y , we have:

FACT 2.2. *A redundant subformula ϕ_v can be replaced by an irredundant formula $\tilde{\phi}_v$ which computes the same function as ϕ_v . Note that $L(\tilde{\phi}_v) < L(\phi_v)$ holds and we can obtain $\tilde{\phi}_v$ from ϕ_v in polynomial time in $L(\phi)$.*

The usefulness of the notion of maximal linear nodes is shown by the following lemma.

LEMMA 2.3. *Let ϕ be a formula which contains exactly m maximal linear nodes. Then, we can check the satisfiability of ϕ in time $2^m \cdot \text{poly}(L(\phi))$.*

PROOF. Let v_1, \dots, v_m be maximal linear nodes in ϕ . By the inductive definition of maximal linear node, the output of ϕ is fixed when all the outputs of $\phi_{v_1}, \dots, \phi_{v_m}$ are fixed to constants $a_1, \dots, a_m \in \{0, 1\}$. There are 2^m possible ways of fixing the outputs of $\phi_{v_1}, \dots, \phi_{v_m}$. For each possibility, we can check whether there exists a corresponding assignment to x_1, \dots, x_n or not, by solving a system of linear equations $\phi_{v_1} = a_1, \dots, \phi_{v_m} = a_m$ in polynomial time using Gaussian elimination. \square

The above lemma motivates us to reduce the number of maximal linear nodes in a formula ϕ . We need some definitions. Let u, v be maximal linear nodes in ϕ . Then u and v are *connected by a \oplus -path* if every node in the (unique) path from u to v is labeled by \oplus . Such a pair of nodes u, v is called *mergeable* for the following reason. Let s and t be the parent nodes of u and v respectively. Then we can write $\phi_s = \phi_u \oplus \phi_{u'}$ and $\phi_t = \phi_v \oplus \phi_{v'}$ where $\phi_{u'}$ and $\phi_{v'}$ are any formulas. Note that neither u' nor v' is maximal linear. We can assume $t \neq u'$ (by changing role of s and t if necessary), and consider the following transformation. First replace ϕ_s by $\tilde{\phi}_s = (\phi_u \oplus \phi_v) \oplus \phi_{u'}$, then replace ϕ_t by $\tilde{\phi}_t = \phi_{v'}$. The resulting formula $\tilde{\phi}$ obviously computes the same function as ϕ because $(\phi_u \oplus \phi_{u'}) \oplus \cdots \oplus (\phi_v \oplus \phi_{v'}) = ((\phi_u \oplus \phi_v) \oplus \phi_{u'}) \oplus \cdots \oplus \phi_{v'}$ by the commutativity of \oplus and we reduce the number of maximal linear nodes by one. Thus we have:

FACT 2.4. *Let ϕ be a formula which contains mergeable pairs of maximal linear nodes. Then there exists a formula $\tilde{\phi}$ which computes the same function as ϕ and does not contain mergeable pairs of maximal linear nodes. Furthermore, $L(\tilde{\phi}) \leq L(\phi)$ and we can obtain $\tilde{\phi}$ from ϕ in polynomial time in $L(\phi)$.*

2.2. Restrictions of $\{\wedge, \vee, \oplus\}$ -formulas. For any formula ϕ , any set of variables $\{x_{i_1}, \dots, x_{i_k}\}$ and any constants $a_1, \dots, a_k \in \{0, 1\}$, we denote by $\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]$ the formula obtained from ϕ by assigning to each x_{i_j}, \bar{x}_{i_j} the value a_j, \bar{a}_j and applying the following procedure **Simplify**.

The procedure **Simplify** reduces the size of a formula by applying rules to eliminate constants and redundant literals and gates. It also reduces the number of redundant subformulas and mergeable linear nodes in a formula. These are the same simplification rules used by [Håstad (1998)] and [Santhanam (2010)] with additional rules regarding \oplus gates.

Simplify (ϕ : formula)

Repeat the following until there is no decrease in size of ϕ , number of redundant subformulas and number of mergeable linear nodes.

- (a) If $0 \wedge \psi$ occurs as a subformula, where ψ is any formula, replace this subformula by 0.
- (b) If $0 \vee \psi$ occurs as a subformula, where ψ is any formula, replace this subformula by ψ .
- (c) If $1 \wedge \psi$ occurs as a subformula, where ψ is any formula, replace this subformula by ψ .
- (d) If $1 \vee \psi$ occurs as a subformula, where ψ is any formula, replace this subformula by 1.
- (e) If $y \vee \psi$ occurs as a subformula, where ψ is a formula and y is a literal, then replace all occurrences of y in ψ by 0 and all occurrence of \bar{y} by 1.
- (f) If $y \wedge \psi$ occurs as a subformula, where ψ is a formula and y is a literal, then replace all occurrences of y in ψ by 1 and all occurrence of \bar{y} by 0.
- (g) If $0 \oplus \psi$ occurs as a subformula, where ψ is any formula, replace this subformula by ψ .
- (h) If $1 \oplus \psi$ occurs as a subformula, where ψ is any formula, replace this subformula by $\bar{\psi}$, where $\bar{\psi}$ denotes a formula which computes the negation of ψ . ($L(\bar{\psi}) \leq L(\psi)$ by Fact 2.1.)
- (i) If ψ occurs as a redundant subformula, replace this subformula by an irredundant formula $\tilde{\psi}$ as Fact 2.2.
- (j) If ϕ contains mergeable maximal linear nodes, replace ϕ by $\tilde{\phi}$ without them using Fact 2.4.

It is easy to see that **Simplify** runs in time polynomial in the size of ϕ and the resulting formula computes the same function as ϕ . If **Simplify**(ϕ) returns ϕ itself, ϕ is called *irreducible*.

OBSERVATION 2.5. *Let ϕ be a formula, x be a variable and $a \in \{0, 1\}$ be a constant. Then*

$$L(\phi[x = a]) \leq L(\phi) - \text{freq}(x).$$

Furthermore, if $\text{freq}(x) \geq \frac{L(\phi)}{|\text{var}(\phi)|} + \gamma$ for some $\gamma \geq 0$, then

$$L(\phi[x = a]) \leq L(\phi) \left(1 - \frac{1}{|\text{var}(\phi)|}\right)^{1 + \frac{\gamma|\text{var}(\phi)|}{L(\phi)}}.$$

PROOF. The first inequality is obvious. The following calculation shows the second inequality.

$$\begin{aligned} L(\phi[x = a]) &\leq L(\phi) - \frac{L(\phi)}{|\text{var}(\phi)|} - \gamma \\ &= L(\phi) \left(1 - \frac{1 + \frac{\gamma|\text{var}(\phi)|}{L(\phi)}}{|\text{var}(\phi)|}\right) \\ &\leq L(\phi) \left(1 - \frac{1}{|\text{var}(\phi)|}\right)^{1 + \frac{\gamma|\text{var}(\phi)|}{L(\phi)}}, \end{aligned}$$

where the last inequality is by $(1 - bx) \leq (1 - x)^b$ for $0 < b, x < 1$. \square

OBSERVATION 2.6. *Let ϕ be a formula, v be a maximal linear node of ϕ where the parent node of v is labeled by \wedge (\vee , respectively), and u be a sibling of v . Assume $\text{var}(\phi_v) = \{x_{i_1}, \dots, x_{i_k}\}$ and there exists a variable x which is in $\text{var}(\phi_u)$ but not in $\text{var}(\phi_v)$. Then for any constants $a_1, \dots, a_k \in \{0, 1\}$ such that $\phi_v[x_{i_1} = a_1, \dots, x_{i_k} = a_k] = 0$ ($\phi_v[x_{i_1} = a_1, \dots, x_{i_k} = a_k] = 1$, respectively),*

$$L(\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]) \leq L(\phi) - \left(\sum_{x' \in \text{var}(\phi_v)} \text{freq}(x')\right) - 1.$$

Furthermore, if $\text{freq}(x') \geq \frac{L(\phi)}{|\text{var}(\phi)|}$ holds for every $x' \in \text{var}(\phi_v)$,

$$\begin{aligned} &L(\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]) \\ &\leq L(\phi) \left\{ \prod_{j=0}^{k-2} \left(1 - \frac{1}{|\text{var}(\phi)| - j}\right) \right\} \\ &\quad \times \left(1 - \frac{1}{|\text{var}(\phi)| - k + 1}\right)^{1 + \frac{|\text{var}(\phi)|}{L(\phi)}}. \end{aligned}$$

Note that a half fraction of assignments to x_{i_1}, \dots, x_{i_k} makes $\phi_v[x_{i_1} = a_1, \dots, x_{i_k} = a_k] = 0$ ($\phi_v[x_{i_1} = a_1, \dots, x_{i_k} = a_k] = 1$, respectively).

PROOF. If we assign 0/1 values to x_{i_1}, \dots, x_{i_k} , obviously $\sum_{x' \in \text{var}(\phi_v)} \text{freq}(x')$ of leaves become constants. Furthermore, by rule (a) of **Simplify** (by rule (d) of **Simplify**, respectively), v 's parent node becomes constant. That is, we can remove at least one leaf of ϕ_u whose label is x or \bar{x} . The second inequality follows from the first one and using Observation 2.5. \square

OBSERVATION 2.7. Let ϕ be a formula, v be a maximal linear node of ϕ where the parent node of v is labeled by \wedge or \vee , and u be a sibling of v . Assume $\text{var}(\phi_v) = \{x_{i_1}, \dots, x_{i_k}\}$ and $\text{var}(\phi_v) \supseteq \text{var}(\phi_u)$ and a variable x is in both $\text{var}(\phi_v)$ and $\text{var}(\phi_u)$. Assume $x = x_{i_1}$. Then for any constants $a_2, \dots, a_k \in \{0, 1\}$,

$$L(\phi[x_{i_2} = a_2, \dots, x_{i_k} = a_k]) \leq L(\phi) - \left(\sum_{x' \in \text{var}(\phi_v) \setminus \{x_{i_1}\}} \text{freq}(x') \right) - 1.$$

Furthermore, if $\text{freq}(x') \geq \frac{L(\phi)}{|\text{var}(\phi)|}$ holds for any $x' \in \text{var}(\phi_v) \setminus \{x_{i_1}\}$,

$$\begin{aligned} & L(\phi[x_{i_2} = a_2, \dots, x_{i_k} = a_k]) \\ & \leq L(\phi) \left\{ \prod_{j=0}^{k-3} \left(1 - \frac{1}{|\text{var}(\phi)| - j} \right) \right\} \\ & \quad \times \left(1 - \frac{1}{|\text{var}(\phi)| - k + 2} \right)^{1 + \frac{|\text{var}(\phi)|}{L(\phi)}}. \end{aligned}$$

PROOF. If we assign 0/1 values to x_{i_2}, \dots, x_{i_k} , obviously $\sum_{x' \in \text{var}(\phi_v) \setminus \{x_{i_1}\}} \text{freq}(x')$ of leaves become constants. Furthermore, ϕ_v becomes x or \bar{x} and ϕ_u becomes one of $0, 1, x, \bar{x}$. That is, by rules (a),(b),(c),(d),(e) and (f), we can remove at least one leaf whose label is x or \bar{x} . The second inequality follows from the first one and using Observation 2.5. \square

2.3. A structural lemma for $\{\wedge, \vee, \oplus\}$ -formulas. In this section, we present a structural lemma for $\{\wedge, \vee, \oplus\}$ -formulas, which is the main technical contribution of this paper. Let us explain the motivation of the lemma.

First recall that the SAT algorithm of Santhanam (2010) for U_2 -formula is based on the following observation:

OBSERVATION 2.8. *Let ϕ be an n -variable U_2 -formula and let x be a variable of maximum frequency in ϕ . Then $L(\phi[x = a]) \leq L(\phi) \left(1 - \frac{3}{2n}\right)$ and $L(\phi[x = \bar{a}]) \leq L(\phi) \left(1 - \frac{1}{n}\right)$ holds for some $a \in \{0, 1\}$.*

We can interpret the inequality $L(\phi[x = a]) \leq L(\phi) \left(1 - \frac{3}{2n}\right)$ as that ϕ shrinks non-trivially and this is the source of running time savings. However, the above observation does not hold for $\{\wedge, \vee, \oplus\}$ -formulas due to the existence of parity gates.

Our structural lemma guarantees that given a $\{\wedge, \vee, \oplus\}$ -formula ϕ , either (i) satisfiability of ϕ can be somewhat easily solved by brute force search, or (ii) there exists a good set of variables such that assigning values to them shrinks ϕ non-trivially. The following is the formal statement of the lemma.

LEMMA 2.9. *In any n -variable formula ϕ of size cn , $c \geq 3/4$, such that ϕ does not contain any pair of mergeable maximal linear nodes, at least one of the following holds.*

Case 1: *The number of maximal linear nodes in ϕ is less than $3n/4$.*

Case 2: *There exists a variable $x \in \text{var}(\phi)$, $\text{freq}(x) \geq c + \frac{1}{8c}$.*

Case 3: *There exists a maximal linear node v with $L(\phi_v) \leq 8c$ such that for any $x \in \text{var}(\phi_v)$, $\text{freq}(x) \geq c$ and the parent node of v is labelled by \wedge or \vee .*

PROOF. Assume that neither Case 1 nor Case 2 occurs. We need the following lemma and fact which are proven later.

LEMMA 2.10. *Let ϕ be a formula which contains at least one node labeled by \wedge or \vee but not any pair of mergeable maximal linear nodes. Then, the number of maximal linear nodes whose parent nodes are labeled by \wedge or \vee , denoted by $\#\text{MLin}_{\wedge,\vee}(\phi)$, is greater than the number of maximal linear nodes whose parent nodes are labeled by $\{\oplus\}$, denoted by $\#\text{MLin}_{\oplus}(\phi)$.*

FACT 2.11. (i) *The number of maximal linear nodes which have more than $8c$ leaves as descendants is at most $n/8$.* (ii) *The number of maximal linear nodes which have a variable x with $\text{freq}(x) < c$ is at most $n/8$.*

By Lemma 2.10, $\#\text{MLin}_{\wedge,\vee}(\phi) > (3n/4)/2 = 3n/8$. By Fact 2.11, there are at least $3n/8 - n/8 - n/8 = n/8$ maximal linear nodes satisfying the condition of Case 3.

PROOF (of Lemma 2.10). We will prove by induction on the size of ϕ .

If $L(\phi) = 1$, then ϕ does not contain a node labeled by \wedge or \vee . If $L(\phi) = 2$ or 3 , it is easy to check that $\#\text{MLin}_{\wedge,\vee}(\phi) > \#\text{MLin}_{\oplus}(\phi)$ holds.

Now we will show that $\#\text{MLin}_{\wedge,\vee}(\phi) > \#\text{MLin}_{\oplus}(\phi)$ for ϕ whose size is $\ell > 3$ and which contains at least one node labeled by \wedge or \vee . If the number of internal nodes which are labeled by \vee or \wedge is exactly one, it is easy to see that $\#\text{MLin}_{\wedge,\vee}(\phi) > \#\text{MLin}_{\oplus}(\phi)$ holds. Thus, assume otherwise. Consider the following two cases. (i) There is a maximal linear node of size at least two. (ii) Every maximal linear node is of size exactly one.

In case (i), pick any maximal linear node of size at least two, say v . Since $L(\phi_v) \geq 2$, ϕ_v contains an internal node w whose child nodes are leaves, say s and t . Let $\tilde{\phi}_v$ be a formula which is identical to ϕ_v except that the nodes s, t are removed from ϕ_v and the node w is replaced by s . Then let $\tilde{\phi}$ be the formula obtained from ϕ by replacing ϕ_v by $\tilde{\phi}_v$. Note that $\tilde{\phi}$ contains at least one node labeled by \wedge or \vee but does not contain any pairs of mergeable maximal linear nodes. Since $L(\tilde{\phi}) = \ell - 2$, it holds that $\#\text{MLin}_{\wedge,\vee}(\tilde{\phi}) > \#\text{MLin}_{\oplus}(\tilde{\phi})$ by the induction hypothesis. It is easy to see that

$\#\text{MLin}_{\wedge, \vee}(\phi) = \#\text{MLin}_{\wedge, \vee}(\tilde{\phi})$ and $\#\text{MLin}_{\oplus}(\phi) = \#\text{MLin}_{\oplus}(\tilde{\phi})$ by the construction of $\tilde{\phi}$, we are done.

In case (ii), there exists at least one internal node, say u , whose label is \vee or \wedge and child nodes are leaves, say p and q . Furthermore, there exists at least one more internal node whose label is \vee or \wedge . Let v be a parent node of u and w be a sibling of u .

If v is labeled by \vee or \wedge , consider a formula $\tilde{\phi}$ obtained from ϕ by replacing u by s . Since $L(\tilde{\phi}) = \ell - 1$, it holds that $\#\text{MLin}_{\wedge, \vee}(\tilde{\phi}) > \#\text{MLin}_{\oplus}(\tilde{\phi})$ by the induction hypothesis. It is easy to see that $\#\text{MLin}_{\wedge, \vee}(\tilde{\phi}) = \#\text{MLin}_{\wedge, \vee}(\phi) - 1$ and $\#\text{MLin}_{\oplus}(\tilde{\phi}) = \#\text{MLin}_{\oplus}(\phi)$, we are done.

If v is labeled by \oplus , consider a formula $\tilde{\phi}$ obtained from ϕ by replacing ϕ_v by ϕ_w . It is easy to see that $\#\text{MLin}_{\wedge, \vee}(\phi) - 2 = \#\text{MLin}_{\wedge, \vee}(\tilde{\phi})$ and $\#\text{MLin}_{\oplus}(\phi) = \#\text{MLin}_{\oplus}(\tilde{\phi})$ by the construction of $\tilde{\phi}$, we are done. \square

PROOF (of Fact 2.11). (i) is obvious by the averaging argument. We will show (ii). By the averaging argument, there exists a variable $x \in \text{var}(\phi)$ with $\text{freq}(x) \geq c$. Since there is no variable $x \in \text{var}(\phi)$ with $\text{freq}(x) \geq c + \frac{1}{8c}$ by the assumption that Case 2 does not hold, we have $c \leq \text{freq}(x) < c + \frac{1}{8c}$. Note that if there exists a variable $x' \in \text{var}(\phi)$ with $\text{freq}(x') < c$ then $\text{freq}(x') < c + \frac{1}{8c} - 1$ because $\text{freq}(x')$ is integer for any variable x' . By the averaging argument, there are at most $\frac{n}{8c}$ variables in $\text{var}(\phi)$ with frequency less than c . The total number of leaves labeled by such variables is at most $c \times \frac{n}{8c} \leq n/8$. \square

These proofs complete the proof of Lemma 2.9. \square

3. A Satisfiability Algorithm for B_2 -Formulas

Before describing our B_2 -formula-SAT algorithm and its running time analysis, let us give a basic idea behind them.

Let ϕ be an n -variable formula of size cn . If c is less than $3/4$ or the number of maximal linear nodes in ϕ is less than $3n/4$, then we can check the satisfiability of ϕ in time $2^{3n/4}$. Otherwise, Case 2 or 3 of Lemma 2.9 holds. In such a case, we can reduce the size of ϕ non-trivially by fixing some number of variables to be constants

as shown in Observation 2.5, Observation 2.6 and Observation 2.7. Note that if $\text{freq}(x_{i_j}) \geq c$ for any j , then

$$L(\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]) \leq L(\phi) \left\{ \prod_{j=0}^{k-1} \left(1 - \frac{1}{n-j} \right) \right\}$$

holds by repeatedly using Observation 2.5. This decrease of the size from ϕ to $\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]$ is called *trivial*. If for some $\gamma > 0$,

$$L(\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]) \leq L(\phi) \left\{ \prod_{j=0}^{k-1} \left(1 - \frac{1}{n-j} \right) \right\}^{\gamma} \left(1 - \frac{1}{n} \right)^{\gamma}$$

holds, then the decrease of the size from ϕ to $\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]$ is called *non-trivial*. If Observation 2.5 or Observation 2.7 applies to ϕ , $\phi[x = a]$ or $\phi[x_{i_2} = a_2, \dots, x_{i_k} = a_k]$ is non-trivially reduced for any $a, a_2, \dots, a_k \in \{0, 1\}$. However, if Observation 2.6 applies to ϕ , $\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]$ is non-trivially reduced for at least a half fraction of assignments of $a_1, \dots, a_k \in \{0, 1\}$, and $\phi[x_{i_1} = a_1, \dots, x_{i_k} = a_k]$ is at least trivially reduced for the remaining assignments of $a_1, \dots, a_k \in \{0, 1\}$. To summarize, if we choose certain number of variables appropriately and assign 0/1 values to them uniformly at random, then the formula size non-trivially reduces with probability at least $1/2$. We would like to estimate the expected size of the reduced formula after assigning values to $(1 - \alpha)n$ variables for some $\alpha > 0$.

The lemma described below captures the analysis of the above process. It is a generalization of the Lemma 5 shown in [Santhanam (2010)]. Let X_0, X_1, \dots be independent random variables which take 0/1 values uniformly at random. Let $\alpha \in (0, 1)$ and $\gamma > 0$ be real numbers and b, n be positive integers. We assume n is sufficiently larger than b . Let $Y_n(\alpha, b, \gamma)$ be a random variable defined as follows:

$$Y_n(\alpha, b, \gamma) := \left\{ \prod_{i=0}^{(1-\alpha)n} \left(1 - \frac{1}{n-i} \right) \right\} \left\{ \prod_{i=0}^{\frac{(1-\alpha)n}{b}} \left(1 - \frac{1}{n-bi} \right)^{\gamma X_i} \right\}.$$

LEMMA 3.1. *For any positive integer b , any real numbers $\delta \in (0, 1)$, $\gamma > 0$, there exist positive real numbers $\alpha \in (0, 1)$, $\epsilon > 0$ and a positive integer N , such that for any integer $n \geq N$,*

$$\Pr[Y_n(\alpha, b, \gamma) \leq \alpha\delta] \geq 1 - 2^{-\epsilon n}$$

holds.

PROOF. If what follows, we ignore integrality issues for simplicity, but the argument holds with slight modification. First note that $\left\{ \prod_{i=0}^{(1-\alpha)n} \left(1 - \frac{1}{n-i}\right) \right\} \leq \alpha$. Let $\zeta \in (0, 1)$ be a small positive real number chosen later. Let I_j be a set of consecutive integers defined as

$$I_j := \{(\zeta n)(j-1) + 1, \dots, (\zeta n)j\}$$

for $1 \leq j \leq \frac{1-\alpha}{b\zeta}$. It is easy to see by the Chernoff bound that

$$\Pr \left[\sum_{i \in I_j} X_i \leq \frac{\zeta n}{3} \right] = 2^{-\Omega_\zeta(n)}$$

for any j , $1 \leq j \leq \frac{1-\alpha}{b\zeta}$. Here Ω_ζ hides a constant factor determined by ζ . Thus, we have

$$\Pr \left[\sum_{i \in I_j} X_i > \frac{\zeta n}{3} \text{ for any } j, 1 \leq j \leq \frac{1-\alpha}{b\zeta} \right] = 1 - 2^{-\Omega_\zeta(n)}$$

by the union bound. We can show the following fact by an elementary calculation.

FACT 3.2. *If $\sum_{i \in I_j} X_i > \zeta n/3$ for any j , $1 \leq j \leq \frac{1-\alpha}{b\zeta}$, then it holds that*

$$\left\{ \prod_{i=0}^{(1-\alpha)n/b} \left(1 - \frac{1}{n-bi}\right)^{\gamma X_i} \right\} < (\alpha + \zeta)^{\gamma/3b}.$$

PROOF. We need the following two inequalities.

$$(3.3) \quad \prod_{i \in I_j} \left(1 - \frac{1}{n-bi}\right)^{X_i} \leq \prod_{i \in I_{j-1}} \left(1 - \frac{1}{n-bi}\right)^{\frac{1}{3}}$$

for $2 \leq j \leq \frac{1-\alpha}{b\zeta}$, and

$$(3.4) \quad \left(1 - \frac{1}{n - bi}\right) \leq \prod_{k=0}^{b-1} \left(1 - \frac{1}{n - bi + k}\right)^{\frac{1}{b}}.$$

Then, we can deduce the desired bound as follows.

$$\begin{aligned} & \left\{ \prod_{i=0}^{(1-\alpha)n/b} \left(1 - \frac{1}{n - bi}\right)^{\gamma X_i} \right\} \\ & \leq \left\{ \prod_{j=1}^{(1-\alpha)n/b\zeta} \prod_{i \in I_j} \left(1 - \frac{1}{n - bi}\right)^{\gamma X_i} \right\}, \end{aligned}$$

then, by (3.3),

$$\leq \left\{ \prod_{j=2}^{(1-\alpha)n/b\zeta} \prod_{i \in I_{j-1}} \left(1 - \frac{1}{n - bi}\right)^{\gamma/3} \right\},$$

then, by (3.4),

$$\begin{aligned} & \leq \left\{ \prod_{j=2}^{(1-\alpha)n/b\zeta} \prod_{i \in I_{j-1}} \prod_{k=0}^{b-1} \left(1 - \frac{1}{n - bi + k}\right)^{\gamma/3b} \right\} \\ & = \prod_{i=1}^{(\alpha+\zeta)n} \left(1 - \frac{1}{n - i}\right)^{\gamma/3b} \leq (\alpha + \zeta)^{\gamma/3b}. \end{aligned}$$

□

Set $\zeta = \alpha$, $2\alpha = \delta^{3b/\gamma}$ and choose $\epsilon = \epsilon(\zeta)$ appropriately, we have the desired bound. □

3.1. The Algorithm and Computation Tree. Our satisfiability algorithm for B_2 -formulas, **Evalformula** is described in Figure 3.1. Without loss of generality, we assume input formulas are irreducible $\{\wedge, \vee, \oplus\}$ -formulas. The correctness of **Evalformula** is guaranteed by Lemma 2.9.

We define a notion of “computation tree” corresponding to the execution of **EvalFormula** on a formula ϕ . A *computation tree* T_ϕ is a binary tree whose nodes are labeled by a triplet $\langle \psi, s, C \rangle$,

```

EvalFormula ( $\phi$ : Formula,  $n$ : integer)
01: /* Case 0 */
02: if  $L(\phi) = cn < 3n/4$ ,
03:   check the satisfiability of  $\phi$  by brute force search.
04:   if  $\phi$  is satisfiable, return “yes”, else return “no”.
05: /* Case 1 */
06: else if the number of maximal linear nodes is less than  $3n/4$ ,
07:   check the satisfiability of  $\phi$  by Lemma 2.3.
08:   if  $\phi$  is satisfiable, return “yes”, else return “no”.
09: /* Case 2 */
10: else if  $\exists x \in \text{var}(\phi)$ ,  $\text{freq}(x) \geq c + \frac{1}{8c}$ ,
11:   EvalFormula( $\phi[x = 0]$ ,  $n - 1$ ),
12:   EvalFormula( $\phi[x = 1]$ ,  $n - 1$ ).
13: /* Case 3 */
14: else if  $\exists$  maximal linear node  $v$  with  $L(\phi_v) \leq 8c$ 
:       such that  $\forall x \in \text{var}(\phi_v)$ ,  $\text{freq}(x) \geq c$  and
:       the parent node of  $v$  is labeled by  $\wedge$  or  $\vee$ ,
15:   assume  $\text{var}(\phi_v) = \{x_{i_1}, \dots, x_{i_k}\}$  and  $u$  is a sibling of  $v$ .
16:   /* Case 3a */
17:   if  $\exists x$  such that  $x \in \text{var}(\phi_u)$  and  $x \notin \text{var}(\phi_v)$ ,
18:     for each constants  $a_1, \dots, a_k \in \{0, 1\}$ ,
19:       EvalFormula( $\phi_i[x_{i_1} = a_1, \dots, x_{i_k} = a_k]$ ,  $n - k$ ).
20:   /* Case 3b */
21:   else if  $\text{var}(\phi_v) \supseteq \text{var}(\phi_u)$ ,
22:     assume  $x_{i_1} \in \phi_u$ .
23:     for each constants  $a_2, \dots, a_k \in \{0, 1\}$ ,
24:       EvalFormula( $\phi_i[x_{i_2} = a_2, \dots, x_{i_k} = a_k]$ ,  $n - k + 1$ ).

```

Figure 3.1: B_2 -formula SAT algorithm

where ψ is a formula, s is an integer and C is an element in $\{?, 0, 1, 2, 3a, 3a', 3b, 3b'\}$. ψ , s and C are called *formula label*, (*amortized-*)*size label* and *case label*, respectively. We construct T_ϕ recursively as follows.

Base step: The root node of T_ϕ is labeled by $\langle \phi, L(\phi), ? \rangle$.

Recursive Step: Let v be a leaf node of T_ϕ whose depth is d and label is $\langle \psi, s, ? \rangle$.

Case 0: If ψ satisfies the condition of Case 0 in **EvalFormula**, replace v 's label by $\langle \psi, 1, 0 \rangle$.

Case 1: If ψ satisfies the condition of Case 1 in **EvalFormula**, replace v 's label by $\langle \psi, 1, 1 \rangle$.

Case 2: If ψ satisfies the condition of Case 2 in **EvalFormula**, replace v 's label by $\langle \psi, s, 2 \rangle$. Add two nodes v_l, v_r as v 's children in the following way:

v_l is labeled by $\langle \psi[x = 0], L(\psi[x = 0]), ? \rangle$ and

v_r is labeled by $\langle \psi[x = 1], L(\psi[x = 1]), ? \rangle$.

Case 3a: If ψ satisfies the condition of Cases 3 and 3a in **EvalFormula**, replace v 's label by $\langle \psi, L(\psi), 3a \rangle$. Construct a complete binary tree T_k of height k starting from v as follows. If a node u is labeled by $\langle \psi', s, 3a \rangle$ or $\langle \psi', s, 3a' \rangle$ and at a distance $d' < k$ from v , u 's left child is labeled by $\langle \psi'[x_{i_{d'}} = 0], s - \frac{L(\psi')}{n-d'}, 3a' \rangle$ and u 's right child is labeled by $\langle \psi'[x_{i_{d'}} = 1], s - \frac{L(\psi')}{n-d'}, 3a \rangle$.

If a node u is labeled by $\langle \psi', s, 3a \rangle$ or $\langle \psi', s, 3a' \rangle$ and at a distance k from v , u 's left child is labeled by $\langle \psi'[x_{i_k} = 0], L(\psi'[x_{i_k} = 0]), ? \rangle$ and u 's right child is labeled by $\langle \psi'[x_{i_k} = 1], L(\psi'[x_{i_k} = 1]), ? \rangle$.

Case 3b: If ψ satisfies the condition of Cases 3 and 3b in **EvalFormula**, replace v 's label by $\langle \psi, L(\psi), 3b \rangle$. Construct a complete binary tree T_{k-1} of height $k - 1$ starting from v as follows.

If a node u is labeled by $\langle \psi', s, 3b \rangle$ or $\langle \psi', s, 3b' \rangle$ and

at a distance $d' < k - 1$ from v , u 's left child is labeled by $\langle \psi'[x_{i_{d'+1}} = 0], s - \frac{L(\psi)}{n-d}, 3b' \rangle$ and u 's right child is labeled by $\langle \psi'[x_{i_{d'+1}} = 1], s - \frac{L(\psi)}{n-d}, 3b' \rangle$.

If a node u is labeled by $\langle \psi', s, 3b \rangle$ or $\langle \psi', s, 3b' \rangle$ and at a distance $k - 1$ from v , u 's left child is labeled by $\langle \psi'[x_{i_k} = 0], L(\psi'[x_{i_k} = 0]), ? \rangle$ and u 's right child is labeled by $\langle \psi'[x_{i_k} = 1], L(\psi'[x_{i_k} = 1]), ? \rangle$.

We will assume that the computation tree is a complete binary tree of depth n by padding it - if there is a node v at depth less than n whose case label is 0 or 1, we add nodes whose labels are $\langle \text{null}, 1, 0 \rangle$ below v .

The following lemma is crucial in the running time analysis of **EvalFormula**. We use the notation $\tilde{L}(p)$ to denote the size label of p in T_ϕ and $d(p)$ to denote the depth of p in T_ϕ . Let p be a node in T_ϕ with $d(p) \leq n - 8c$ and $T_p(8c)$ denote the set of p 's descendants which are at distance $8c$ from p . Then, we have:

LEMMA 3.5. *There is a subset $\hat{T}_p(8c)$ of $T_p(8c)$ with $|\hat{T}_p(8c)| = \frac{|T_p(8c)|}{2}$ such that for any $q \in \hat{T}_p(8c)$,*

$$\tilde{L}(q) \leq \max \left[1, \tilde{L}(p) \left\{ \prod_{j=0}^{8c-1} \left(1 - \frac{1}{n - d(p) - j} \right) \right\} \times \left(1 - \frac{1}{n - d(p)} \right)^{\frac{1}{c^2}} \right]$$

and for any $q \in T_p(8c) \setminus \hat{T}_p(8c)$,

$$\tilde{L}(q) \leq \max \left\{ 1, \tilde{L}(p) \left\{ \prod_{j=0}^{8c-1} \left(1 - \frac{1}{n - d(p) - j} \right) \right\} \right\}.$$

PROOF. If the case label of p is 2, 3a or 3b, we obtain the desired bound from Observation 2.5, Observation 2.6 or Observation 2.7 respectively according to the case label of p . If the case label of p is 3a' or 3b', we can also use Observation 2.6 or Observation 2.7 respectively because of the definition of amortized-size. Note that $\frac{1}{c^2}$ in the exponent comes from Observation 2.5. If Observation 2.6 or Observation 2.7 apply, the exponent can be $\frac{1}{c}$. \square

3.2. Running time analysis. We begin with some definitions. Let $T_\phi(d)$ denote the set of depth d nodes in T_ϕ . Let p be a node in T_ϕ with $d(p) \leq n - 8c$ and let $\hat{T}_p(8c)$ denote the subset of $T_p(8c)$ such that the conditions in Lemma 3.5 hold. For $p \in T_\phi$, P_p denotes the path from the root of T_ϕ to p and $P_p(d)$ denotes the node of P_p with depth d . We define a function $X_i(p)$ from $T_\phi((1 - \alpha)n)$ to $\{0, 1\}$ as:

$$X_i(p) := \begin{cases} 1 & \text{if } P_p(8c(i + 1)) \in \hat{T}_{P_p(8ci)}(8c) \\ 0 & \text{otherwise} \end{cases}$$

for $0 \leq i \leq \frac{(1-\alpha)n}{8c}$.

The intuition behind these definitions is as follows: If a path P_p from the root of T_ϕ to $p \in T_\phi((1 - \alpha)n)$ goes through a node in $\hat{T}_{P_p(8ci)}(8c)$ for many i , then an assignment corresponding to P_p shrinks T_ϕ significantly.

The following lemma shows the independence of random variables $\{X_i\}$.

LEMMA 3.6. *Let p be drawn from $T_\phi((1 - \alpha)n)$ uniformly at random. Then*

$$\Pr \left[X_0(p) = a_0, \dots, X_{\frac{(1-\alpha)n}{8c}}(p) = a_{\frac{(1-\alpha)n}{8c}} \right] = \left(\frac{1}{2} \right)^{\frac{(1-\alpha)n}{8c} + 1}$$

for any $a_0, \dots, a_{\frac{(1-\alpha)n}{8c}} \in \{0, 1\}$. That is, $X_0, \dots, X_{\frac{(1-\alpha)n}{8c}}$ are independent random variables which take values 0 and 1 uniformly at random.

PROOF. It follows from that we define $\hat{T}_q(8c)$ as $|\hat{T}_q(8c)| = \frac{|T_q(8c)|}{2}$ in Lemma 3.5. \square

We can bound the size label of $p \in T_\phi((1 - \alpha)n)$ using $Y_n(\alpha, b, \gamma)$.

LEMMA 3.7. *Let p be a node of $T_\phi((1 - \alpha)n)$. Then*

$$\tilde{L}(p) \leq \max \left\{ 1, L(\phi) Y_n \left(\alpha, 8c, \frac{1}{c^2} \right) (p) \right\}.$$

where

$$Y_n(\alpha, b, \gamma)(p) := \left\{ \prod_{i=0}^{(1-\alpha)n} \left(1 - \frac{1}{n-i} \right) \right\} \\ \times \left\{ \prod_{i=0}^{(1-\alpha)n/b} \left(1 - \frac{1}{n-bi} \right)^{\gamma X_i(p)} \right\}.$$

PROOF. It follows from Lemma 3.5 and the definition of $X_i(p)$. \square

Combining Lemma 3.1 with $\delta = \frac{3}{4}$ and Lemma 3.7, we have:

LEMMA 3.8. *Let p be drawn from $T_\phi((1-\alpha)n)$ uniformly at random. Then*

$$\Pr \left[\tilde{L}(p) < \frac{3\alpha n}{4} \right] = 1 - 2^{-\Omega_\alpha(n)}$$

for $\alpha = \frac{1}{2} \left(\frac{3}{4} \right)^{24c^3}$.

Now we are ready to prove Theorem 1.1.

PROOF (of Theorem 1.1). Let ϕ be an n -variable formula and consider a corresponding computation tree T_ϕ . We upper bound the running time of **EvalFormula**(ϕ, n) by the sum of the following four values:

(i) Let $N_0(d)$ be the number of nodes whose depth is d and case label is 0. Define T_0 as

$$T_0 := \sum_{d=0}^{(1-\alpha)n} N_0(d) \cdot 2^{3(n-d)/4}.$$

(ii) Let $N_1(d)$ be the number of nodes whose depth is d and case label is 1. Define T_1 as

$$T_1 := \sum_{d=0}^{(1-\alpha)n} N_1(d) \cdot 2^{3(n-d)/4}.$$

(iii) Let $N_{2,3}$ be the number of nodes whose depth is $(1 - \alpha)n$ and size label is less than $3\alpha n/4$. Define $T_{2,3}$ as

$$T_{2,3} := N_{2,3} \cdot 2^{3\alpha n/4}.$$

(iv) Let $N'_{2,3}$ be the number of nodes whose depth is $(1 - \alpha)n$ and size label is at least $3\alpha n/4$. Define $T'_{2,3}$ as

$$T'_{2,3} := N'_{2,3} \cdot 2^{\alpha n}.$$

It is easy to see that $\text{poly}(L(\phi))(T_0 + T_1 + T_{2,3} + T'_{2,3})$ upper bounds the running time of **EvalFormula** (ϕ, n) . $N_0(d), N_1(d)$ is at most 2^d and $N_{2,3}$ is at most $2^{(1-\alpha)n}$. Lemma 3.8 shows that $N'_{2,3} \leq 2^{(1-\alpha)n} \times 2^{-\Omega(n)}$. Therefore, $T_0 + T_1 + T_{2,3} + T'_{2,3} = 2^{n-\Omega(n)}$. \square

4. Strong average-case hardness

We will show strong average-case hardness of affine extractors for linear-sized $\{\wedge, \vee, \oplus\}$ -formulas, as claimed in Theorem 1.2. At first, we give definitions of affine source and affine extractor.

DEFINITION 4.1. *Let F_2 be the finite field with 2 elements. Denote by F_2^n the n dimensional vector space over F_2 . A distribution X over F_2^n is an (n, k) -affine source if there exist linearly independent vectors $a_1, \dots, a_k \in F_2^n$ and another vector $b \in F_2^n$ such that X is sampled by choosing $x_1, \dots, x_k \in F_2$ uniformly and independently and computing*

$$X = \sum_{i=1}^k x_i a_i + b.$$

An affine extractor is a deterministic function such that given any affine source as the input, the output of the function is statistically close to the uniform distribution.

DEFINITION 4.2. A function $f_n : F_2^n \rightarrow F_2$ is a deterministic (k, ϵ) -affine extractor if for every (n, k) -affine source X ,

$$1/2 - \epsilon \leq \Pr[f_n(X) = 0] \leq 1/2 + \epsilon.$$

(And $1/2 - \epsilon \leq \Pr[f_n(X) = 1] \leq 1/2 + \epsilon$.)

We need the following Theorem due to [Bourgain (2007); Li (2011); Yehudayoff (2011)].

THEOREM 4.3. For every $\delta > 0$ there exists a polynomial time computable family of (k, ϵ) -extractors $\{\mathbf{AE}_\delta^n : F_2^n \rightarrow F_2\}$ with $k = \delta n$ and $\epsilon = 2^{-\Omega_\delta(n)}$.

THEOREM 4.4 (Formal restatement of Theorem 1.2). For any constant $c > 0$, there exists $\delta > 0$ such that any sequence of B_2 -formulas of size at most cn must err in computing affine extractors $\{\mathbf{AE}_\delta^n\}$ on at least a $1/2 - 2^{-\Omega(n)}$ fraction of inputs of length n for each n .

PROOF (of Theorem 4.4). Let ϕ be an n -variable formula of size at most cn and consider the computation tree T_ϕ constructed by **EvalFormula** on ϕ . Here we treat T_ϕ as the original one, i.e., we do not add any nodes to make T_ϕ a complete binary tree of height n . Let \tilde{T}_ϕ be the set of all leaf nodes in T_ϕ . Each node p of \tilde{T}_ϕ defines a subcube $C(p)$ of $\{0, 1\}^n$. Note that any node p in \tilde{T}_ϕ has case label 0 or 1.

We consider \mathbf{AE}_δ^n with $\delta = \frac{\alpha}{4}$ where $\alpha = \frac{1}{2} \left(\frac{3}{4}\right)^{24c^3}$ is chosen as Lemma 3.8 and prove the following facts.

FACT 4.5. If p has case label 0 and is at depth d , $C(p)$ can be partitioned into the set of subcubes $\{C_1, \dots, C_k\}$, $k \leq 2^{3(n-d)/4}$ such that each C_i has dimension at least $(n-d) - 3(n-d)/4 = (n-d)/4$ and ϕ becomes constant on C_i . Furthermore, if $d \leq (1 - \alpha)n$,

$$(4.6) \quad \Pr_{x \in C_i} [\phi(x) = \mathbf{AE}_\delta^n(x)] \leq 1/2 + 1/2^{\Omega(n)}.$$

PROOF. Each assignment to $\text{var}(\phi_p)$ determines a subcube C of dimension $n - d - |\text{var}(\phi_p)|$ in $C(p)$ where ϕ is constant over C . Since $|\text{var}(\phi_p)| \leq \frac{3(n-d)}{4}$, we have $k \leq 2^{3(n-d)/4}$. If $d \leq (1 - \alpha)n$, the dimension of each subcube C_i is at least $\frac{n-d}{4} \geq \frac{\alpha}{4}n = \delta n$. Thus, we have (4.6). \square

FACT 4.7. *If p has case label 1 and is at depth d , $C(p)$ can be partitioned into the set of affine subspaces $\{C_1, \dots, C_k\}$, $k \leq 2^{3(n-d)/4}$ such that each C_i has dimension at least $(n - d) - 3(n - d)/4 = (n - d)/4$ and ϕ becomes constant on C_i . Furthermore, if $d \leq (1 - \alpha)n$,*

$$(4.8) \quad \Pr_{x \in C_i} [\phi(x) = \mathbf{AE}_\delta^n(x)] \leq 1/2 + 1/2^{\Omega(n)}.$$

PROOF. Let v_1, v_2, \dots, v_m be maximal linear nodes of ϕ_p and $a_1, a_2, \dots, a_m \in \{0, 1\}$ be constants. If a system of linear equations $\phi_{v_1} = a_1, \phi_{v_2} = a_2, \dots, \phi_{v_m} = a_m$ is feasible, then it determines an affine subspace C' of dimension at least $n - d - m$ in $C(p)$ where ϕ is constant over C' . Since $m \leq \frac{3(n-d)}{4}$, we have $k \leq 2^{3(n-d)/4}$. If $d \leq (1 - \alpha)n$, the dimension of each affine subspace C_i is at least $\frac{n-d}{4} \geq \frac{\alpha}{4}n = \delta n$. Thus, we have (4.8). \square

In the proof of Theorem 1.1, we can see that

$$\Pr_{x \in \{0,1\}^n} \left[x \in \bigcup_{\substack{p \in \tilde{T}_\phi \\ d(p) > (1-\alpha)n}} C(p) \right] \leq 2^{-\Omega(n)}.$$

Therefore,

$$\begin{aligned} & \Pr_x [\phi(x) = \mathbf{AE}_\delta^n(x)] \\ &= \sum_{p \in \tilde{T}_\phi} \Pr_{x \in \{0,1\}^n} [x \in C(p)] \Pr_{x \in C(p)} [\phi(x) = \mathbf{AE}_\delta^n(x)] \\ &= \sum_{\substack{p \in \tilde{T}_\phi \\ d(p) \leq (1-\alpha)n}} \Pr_{x \in \{0,1\}^n} [x \in C(p)] \Pr_{x \in C(p)} [\phi(x) = \mathbf{AE}_\delta^n(x)] \\ & \quad + \sum_{\substack{p \in \tilde{T}_\phi \\ d(p) > (1-\alpha)n}} \Pr_{x \in \{0,1\}^n} [x \in C(p)] \Pr_{x \in C(p)} [\phi(x) = \mathbf{AE}_\delta^n(x)] \end{aligned}$$

and

$$\sum_{\substack{p \in \tilde{T}_\phi \\ d(p) \leq (1-\alpha)n}} \Pr_{x \in \{0,1\}^n} [x \in C(p)] \Pr_{x \in C(p)} [\phi(x) = \mathbf{AE}_\delta^n(x)] \leq \frac{1}{2} + 2^{-\Omega(n)}$$

by Fact 4.5 and Fact 4.7 , and

$$\begin{aligned} & \sum_{\substack{p \in \tilde{T}_\phi \\ d(p) > (1-\alpha)n}} \Pr_{x \in \{0,1\}^n} [x \in C(p)] \Pr_{x \in C(p)} [\phi(x) = \mathbf{AE}_\delta^n(x)] \\ & \leq \sum_{\substack{p \in \tilde{T}_\phi \\ d(p) > (1-\alpha)n}} \Pr_{x \in \{0,1\}^n} [x \in C(p)] \leq 2^{-\Omega(n)} \end{aligned}$$

where the last inequality follows from Lemma 3.8. This completes the proof. \square

Acknowledgment

The authors are deeply grateful to Alexander S. Kulikov for useful discussion and helpful comments on the manuscript. We also would like to thank two reviewers for their careful reading and comments on our paper. The comments helped us a lot in improving the paper.

References

- MIKLÓS AJTAI (1983). Σ_1^1 -formulae on finite structures. *Ann. Pure Appl. Logic* **24**(1), 1–48.
- A. E. ANDREEV (1987). On a method for obtaining more than quadratic effective lower bounds for the complexity of π -schemes. *Moscow Univ. Math. Bull.* **42**(1), 63–66.
- VIKRAMAN ARVIND & RAINER SCHULER (2003). The Quantum Query Complexity of 0-1 Knapsack and Associated Claw Problems. In *Proceedings of the 14th Annual International Symposium on Algorithm and Computation (ISAAC)*, 168–177.
- ARMIN BIERE, MARIJN HEULE, HANS VAN MAAREN & TOBY WALSH (editors) (2009). *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.

JEAN BOURGAIN (2007). On the construction of affine-source extractors. *Geometric and Functional Analysis* **17**(1), 33–57.

CHRIS CALABRO, RUSSELL IMPAGLIAZZO & RAMAMOHAN PATURI (2006). A Duality between Clause Width and Clause Density for SAT. In *Proceedings of the 21st Annual IEEE Conference Computational Complexity (CCC)*, 252–260.

CHRIS CALABRO, RUSSELL IMPAGLIAZZO & RAMAMOHAN PATURI (2009). The Complexity of Satisfiability of Small Depth Circuits. In *Proceedings of the 4th International Workshop on Parameterized and Exact Computation (IWPEC)*, 75–85.

EVGENY DANTSIN & EDWARD A. HIRSCH (2009). Worst-Case Upper Bounds. In *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, 403–424. IOS Press.

EVGENY DANTSIN, EDWARD A. HIRSCH & ALEXANDER WOLPERT (2004). Algorithms for SAT Based on Search in Hamming Balls. In *Proceedings of the 21st Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, 141–151.

EVGENY DANTSIN, EDWARD A. HIRSCH & ALEXANDER WOLPERT (2006). Clause Shortening Combined with Pruning Yields a New Upper Bound for Deterministic SAT Algorithms. In *Proceedings of the 6th Conference on Algorithms and Complexity (CIAC)*, 60–68.

MERRICK L. FURST, JAMES B. SAXE & MICHAEL SIPSER (1984). Parity, Circuits, and the Polynomial-Time Hierarchy. *Mathematical Systems Theory* **17**(1), 13–27.

JOHAN HÅSTAD (1986). Almost Optimal Lower Bounds for Small Depth Circuits. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC)*, 6–20.

JOHAN HÅSTAD (1998). The Shrinkage Exponent of de Morgan Formulas is 2. *SIAM J. Comput.* **27**(1), 48–64.

TIMON HERTLI (2011). 3-SAT Faster and Simpler - Unique-SAT Bounds for PPSZ Hold in General. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 277–284.

EDWARD A. HIRSCH (2008). Exact Algorithms for General CNF SAT. In *Encyclopedia of Algorithms*. Springer-Verlag.

RUSSELL IMPAGLIAZZO, WILLIAM MATTHEWS & RAMAMOHAN Paturi (2012). A satisfiability algorithm for AC^0 . In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 961–972.

RUSSELL IMPAGLIAZZO & NOAM NISAN (1993). The Effect of Random Restrictions on Formula Size. *Random Struct. Algorithms* **4**(2), 121–134.

XIN LI (2011). A New Approach to Affine Extractors and Dispersers. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity (CCC)*, 137–147.

CHI-JEN LU & HSIN-LUNG WU (2010). On the Hardness against Constant-Depth Linear-Size Circuits. In *Proceedings of the 16th Annual International Computings and Combinatorics Conference (COCOON)*, 13–22.

KAZUHISA MAKINO, SUGURU TAMAKI & MASAKI YAMAMOTO (2011). Derandomizing HSSW Algorithm for 3-SAT. In *Proceedings of the 17th Annual International Computings and Combinatorics Conference (COCOON)*, 1–12.

BURKHARD MONIEN & EWALD SPECKENMEYER (1985). Solving satisfiability in less than 2^n steps. *Discrete Applied Mathematics* **10**, 287–295.

ROBIN A. MOSER & DOMINIK SCHEDER (2011). A Full Derandomization of Schönning’s k -SAT Algorithm. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC)*, 245–252.

NOAM NISAN & AVI WIGDERSON (1994). Hardness vs Randomness. *J. Comput. Syst. Sci.* **49**(2), 149–167.

RAMAMOHAN Paturi, PAVEL PUDLÁK, MICHAEL E. SAKS & FRANCIS ZANE (2005). An improved exponential-time algorithm for k -SAT. *J. ACM* **52**(3), 337–364.

RAMAMOHAN Paturi, PAVEL PUDLÁK & FRANCIS ZANE (1999). Satisfiability Coding Lemma. *Chicago J. Theor. Comput. Sci.* **1999**, 19 pages.

PAVEL PUDLÁK (1998). Satisfiability - Algorithms and Logic. In *Proceedings of the 23rd International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 129–141.

RAHUL SANTHANAM (2010). Fighting Perebor: New and Improved Algorithms for Formula and QBF Satisfiability. In *Proceedings of the 51st International Symposium on Foundations of Computer Science (FOCS)*, 183–192.

UWE SCHÖNING (1999). A Probabilistic Algorithm for k -SAT and Constraint Satisfaction Problems. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, 410–414.

RAINER SCHULER (2005). An algorithm for the satisfiability problem of formulas in conjunctive normal form. *J. Algorithms* **54**(1), 40–44.

B. A. SUBBOTOVSKAYA (1961). Realizations of linear functions by formulas using and, or, not. *Soviet Mathematics Doklady* **2**, 110–112.

RYAN WILLIAMS (2010). Improving exhaustive search implies super-polynomial lower bounds. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, 231–240.

RYAN WILLIAMS (2011). Non-uniform ACC Circuit Lower Bounds. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity (CCC)*, 115–125.

ANDREW CHI-CHIH YAO (1985). Separating the Polynomial-Time Hierarchy by Oracles (Preliminary Version). In *Proceedings of the 26th International Symposium on Foundations of Computer Science (FOCS)*, 1–10.

AMIR YEHUDAYOFF (2011). Affine extractors over prime fields. *Combinatorica* **31**(2), 245–256.

KAZUHISA SETO
Graduate School of Informatics
Kyoto University
Kyoto, Japan
seto@kuis.kyoto-u.ac.jp
[http://www.lab2.kuis.kyoto-u.ac.jp/
~seto/](http://www.lab2.kuis.kyoto-u.ac.jp/~seto/)

SUGURU TAMAKI
Graduate School of Informatics
Kyoto University
Kyoto, Japan
tamak@kuis.kyoto-u.ac.jp
[http://www.lab2.kuis.kyoto-u.ac.jp/
~tamak/](http://www.lab2.kuis.kyoto-u.ac.jp/~tamak/)