Strong ETH and Resolution via Games and the Multiplicity of Strategies

Ilario Bonacina[1] and Navid Talebanfard[2]

1 Computer Science Department, Sapienza University of Rome, via Salaria 113, 00198 Rome, Italy, bonacina@di.uniroma.it
2 Department of Mathematical and Computing Sciences, Tokyo Institute of Technology, Meguro-ku Ookayama 2-12-1, Japan 152-8552, navid@is.titech.ac.jp

── Abstract ──

We consider a restriction of the Resolution proof system in which at most a fixed number of variables can be resolved more than once along each refutation path. This system lies between regular Resolution, in which no variable can be resolved more than once along any path, and general Resolution where there is no restriction on the number of such variables. We show that when the number of re-resolved variables is not too large, this proof system is consistent with the Strong Exponential Time Hypothesis (SETH). More precisely for large n and k we show that there are unsatisfiable k-CNF formulas which require Resolution refutations of size $2^{(1-\epsilon_k)n}$, where n is the number of variables and $\epsilon_k = \widetilde{O}(k^{-1/5})$, whenever in each refutation path we only allow at most $\widetilde{O}(k^{-1/5})n$ variables to be resolved multiple times. However, these re-resolved variables along different paths do not need to be the same. Prior to this work, the strongest proof system shown to be consistent with SETH was regular Resolution [Beck and Impagliazzo, STOC'13]. This work strengthens that result and gives a different and conceptually simpler game-theoretic proof for the case of regular Resolution.

1 Introduction

The SAT problem is one of the most fundamental NP-complete problems. The theoretical significance of this problem has once again been demonstrated recently, after a series of results showing that SAT if not equivalent to circuit lower bounds but it is at least closely related. Paturi, Pudlák and Zane [19] proved tight depth-3 circuit lower bounds and from their technique they obtained a k-SAT algorithm which beats exhaustive search. Along similar lines, Santhanam [23] modified a lower bound argument to obtain improved satisfiability algorithms for De Morgan formulas of linear size. Employing stronger lower bound arguments, satisfiability algorithms were given for formulas of larger size in [7] and [8]. In a different direction, Williams [27] showed that even small improvements over exhaustive search for satisfiability on certain circuit classes implies a lower bound against that class. In fact he obtained his seminal NEXP $\not\subseteq$ ACC^0 result in [28] by giving a non-trivial ACC^0-SAT algorithm.

In this paper we will be focusing on the k-SAT problem. There are several non-trivial algorithms known for this problem (see e.g. [11, 19, 18, 24]). Despite this however, the exact complexity of k-SAT under suitable assumptions remains unknown. Formalising what this complexity could be, Impagliazzo and Paturi [15] formulated the following two hypotheses. The *Exponential Time Hypothesis* (ETH) which states that the are no sub-exponential time algorithms for the SAT problem, and the *Strong Exponential Time Hypothesis* (SETH) which states that the complexity of k-SAT grows as k increases and the running time of the best k-SAT algorithms approach that of exhaustive search. More formally, it says that k-SAT requires running time $2^{(1-\epsilon_k)n}$ where $\epsilon_k \to 0$ as $k \to \infty$.

Both ETH and SETH are stronger than $\mathsf{P} \neq \mathsf{NP}$ and hence we do not expect to be able to verify either of them in any new future. We can however ask whether known algorithms are consistent with these hypotheses, algorithms that work in a relatively intuitive way. For the PPSZ algorithm [18] strong lower bounds were proved in [9] supporting SETH. But one may ask for such a result that holds for a class of algorithms rather than for a specific one. Proof complexity provides a framework to do this. One can think of the run of a SAT algorithm on an unsatisfiable instance as a proof of unsatisfiability. If this proof is structured enough, we can employ tools from proof complexity and obtain lower bounds. For instance practical SAT-solvers are based on the *Davis-Putnam-Logemann-Loveland* algorithm (DPLL) that is a backtracking method introduced by [13, 12] to search for assignments satisfying a CNF formula. It is a well known result that DPLL is equivalent to a sub-system of the proof system Resolution where only proofs having a tree structure are allowed. Hence tree-like Resolution lower bounds transfer to lower bounds for the DPLL algorithm. In a series of works, [16, 25, 17] introduced the idea of *Conflict Driven Clause Learning* (CDCL) as a way for DPLL SAT-solvers to cut the search space and avoid duplicated work. This is done by performing a *conflict analysis* when the search for an assignments leads to a contradiction and then *learning* a clause encoding a reason for that failure. By definition Resolution simulates (polynomially) runs of CDCL solvers over unsatisfiable instances[1], hence lower bounds for Resolution transfer to lower bounds for CDCL solvers.

Exponential lower bounds consistent with ETH have long been known for natural proof systems such as Resolution, see e.g. [26]. These are $2^{\Omega(n)}$ lower bounds for k-CNF formulas on n variables and hence not strong enough to support SETH. Some thirteen years needed to be passed for the first SETH lower bounds. Pudlák and Impagliazzo [22] proved such lower bounds for tree-like Resolution via Prover-Delayer games. Another thirteen years later, Beck and Impagliazzo [4] obtained a very strong width lower bound which simplified and improved the result of [22] for tree-like Resolution and they were able to prove SETH lower bounds for regular Resolution. In this paper we prove another SETH lower bound for regular Resolution. One advantage of our proof is that it gives a SETH lower bound for a proof system which is more general than regular Resolution; we allow at most $\epsilon_k n$ variables to be re-queried along each path. We stress that these re-queried variables along different paths do not need to be the same.

Techniques

A standard technique to prove Resolution size lower bounds is due to Ben-Sasson and Wigderson [5]. They showed that if a formula requires refutations of large width, it also requires refutations with many clauses. More precisely they showed that if a k-CNF formula can only have Resolution refutations of width at least W, then it requires Resolution size at least $2^{(W-k)^2/16n}$, where n is the number of variables. Because of the $\frac{1}{16}$ in the exponent we do not immediately get $2^{(1-\epsilon_k)n}$ size lower bounds from strong width lower bounds. However, we note that if the formula is structured in some sense, for instance if it is a *xorification*, we can avoid this loss.

Beck and Impagliazzo in [4] showed that there are unsatisfiable k-CNF formulas in n variables requiring refutations of size at least $2^{n(1-\epsilon_k)}$ in *regular* Resolution, a sub-system of Resolution. Their proof is an adaptation of a probabilistic technique from [3] and, from an high level can be seen as a variation of the bottleneck counting of Haken in [14]. In their

---

[1] The converse also holds *under certain assumptions* on the behaviour of the CDCL solver, cf. [20] and [2].

argument a rule is given which maps assignments to particular clauses of the proof, at which a significant amount of 'work' is done.

We will be considering xorification of formulas where we replace each variable with the parity of a block of new variables. For such formulas we can strengthen the result of Ben-Sasson and Wigerson and show that the number of large clauses must be really large, and this gives us the desired lower bound. This result is achieved through *Pudlák games* that characterise Resolution size [21] applied to a structured formula, a xorification of some unsatisfiable CNF φ. This allows us to avoid the use of probabilistic arguments and it is the core of our main technical result, cf. Theorem 4. There we prove that if there is a width lower bound for refuting an unsatisfiable CNF φ in Resolution, then there exists a 'sufficiently strong' exponential size lower bound for refuting a xorification of φ. Our construction apply to a restriction of the Resolution proof system in which at most a fixed number of variables can be resolved more than once along each refutation path. For such system the SETH lower bound for size (Corollary 6) follows for our result on size of xorified formulas (Theorem 4) and from a strong width lower bound for some families of CNFs [4].

Informally, in the *Pudlák game* we have two players, Prover and Delayer, that play on some formula φ. Prover has the objective of showing that the formula φ is unsatisfiable by querying variables. Delayer on the other hand wants to play as long as possible before the formula is falsified while answering to the queries Prover asks her. The size of Resolution proofs of φ is then characterised as the minimal number of *records*, i.e. partial assignments, Prover has to consider in a winning strategy. Hence to prove a Resolution size lower bound we show that, in order to win, Prover must keep a large number of records and we can do that by producing a lot of sufficiently different strategies for Delayer. Prover must win against each of them, hence in his winning strategy he must have a lot of distinct records, since the strategies of Delayer are sufficiently different. In the literature this is done essentially by making Prover play against a Delayer that plays accordingly to a random strategy [21, 10]. Then the size lower bound, that is a lower bound on the number of records that Prover must have in a winning strategy, is obtained by probabilistic arguments. This may very likely lead to some loss in the constants that we need to avoid to prove a SETH lower bound for Resolution size. In the Pudlák game played on the xorification of a formula φ, we give a series of strategies for Delayer to which Prover has to answer in order to win. The construction of strategies relies on the characterisation of Resolution width as a game [1]. At a very high level, a winning strategy for Delayer in the width game on φ gives rise to a multitude of strategies for Delayer on the Pudlák game on the xorification of φ. The new strategies act differently from each other on the xorification of φ, but in a sense they all act the same as the original strategy σ on the original formula φ. This is done by exploiting the combinatorial properties of the xorified formula in such a way that the number of Delayer strategies, for the Pudlák game played on the xorified formula, does indeed hugely amplify. Then, the desired size lower bound follows from a counting argument.

2 Preliminaries

A *literal* is either a variable x or its negation $\neg x$. A *clause* C is a disjunction of literals and by its *width* we mean the number of literals appearing in C and we denote this by $|C|$. A *conjunctive normal form* formula (CNF) is a conjunction of a set of clauses.

Given a boolean function f on a set of variables X, a *partial assignment* is a function $\rho : X \to \{0, 1, *\}$. We call *domain* of ρ, $\mathrm{dom}(\rho)$ the set $\rho^{-1}(\{0, 1\})$. The restriction of f to ρ denoted by $f|_\rho$ is a function on $\rho^{-1}(*)$ obtained from f by fixing the value of all variables in

$\rho^{-1}(0) \cup \rho^{-1}(1)$ according to ρ. We write $\rho \subseteq \sigma$ if for all $x \in X$, $\rho(x) \neq *$ implies $\sigma(x) = \rho(x)$. For a partial assignment ρ for which $\rho(x) = *$, by $\rho \cup \{(x, b)\}$ we denote a partial assignment ρ' such that for all $y \neq x$, $\rho'(y) = \rho(y)$ and $\rho'(x) = b$. Given a (partial) assignment ρ and a subset $B \subseteq X$, $\rho|_B$ is a partial assignment defined only on the variables in B such that for all $x \in B$, $\rho|_B(x) = \rho(x)$.

Resolution [6] is a proof system for refuting unsatisfiable CNF formulas. The only inference rule in Resolution is given as follows

$$\frac{C \vee x, \quad D \vee \neg x}{C \vee D},$$

where C and D are clauses and we say that x is *resolved* and $C \vee D$ is called the *resolvant* of $C \vee x$ and $D \vee \neg x$.

A *Resolution derivation* of a clause D from a CNF φ is a sequence $\Pi = \langle C_1, \ldots, C_\tau \rangle$ of clauses such that $C_\tau = D$ and each C_i is either an *axiom*, i.e., a clause from φ, or it is derived by applying the Resolution rule on some clause C_j and $C_{j'}$ such that $j, j' < i$. We will denote this by $\Pi : \varphi \vdash D$. If φ is an unsatisfiable formula, a *Resolution refutation* of φ is a derivation of $\bot$, the empty clause, from φ. Resolution is *sound* and *complete*, that is we can derive $\bot$ from a CNF formula if and only if it is unsatisfiable.

A *δ-regular Resolution derivation* of a clause D from a formula φ in n variables is a Resolution derivation in which along any path at most δn variables are resolved multiple times. Hence a 0-regular Resolution refutation is just a standard regular refutation and a 1-regular Resolution refutation is one without any constraint.

The *size* of a Resolution derivation is the number of clauses appearing in it. We denote the minimum size of a derivation of D from φ by $\mathsf{size}(\varphi \vdash D)$. We also denote the minimum size of a δ-regular derivation of D from φ by $\mathsf{size}_\delta(\varphi \vdash D)$. Similarly we define the *width* of a derivation to be the width of the largest clause appearing in it. We denote the minimum width of a derivation of D from φ by $\mathsf{width}(\varphi \vdash D)$.

3   A game view of Resolution

In this section we present a common framework for the games described by Atserias and Dalmau [1] and Pudlák [21].

▶ **Definition 1** (Game$(\varphi, \mathcal{R})$)**.** Given an unsatisfiable CNF φ in n variables and a set of partial assignments $\mathcal{R}$ containing the empty assignment, we define a game, Game$(\varphi, \mathcal{R})$, between two players Prover (he) and Delayer (she).

At each step i of the game a partial assignment $\alpha_i \in \mathcal{R}$ is maintained (α_0 is the empty partial assignment), then at step $i + 1$ the following moves take place:

1. Prover picks some variable $x \notin \mathrm{dom}(\alpha_i)$.
2. Delayer then has to answer $x = b$ for some bit $b \in \{0, 1\}$.
3. Prover set $\alpha_{i+1} \in \mathcal{R}$ such that $\alpha_{i+1} \subseteq \alpha_i \cup \{(x, b)\}$.

If at any point in the game α_i falsify φ then Prover wins; otherwise Delayer wins. We say that Prover has a *winning strategy* for the game if for any strategy of Delayer, he can play so that he wins the game. Otherwise we say that Delayer has a *winning strategy*.

If in each run of the game Prover can query at most δn variables, we call the corresponding game Game$_\delta(\varphi, \mathcal{R})$.

For a suitable choice of $\mathcal{R}$ the Game$(\varphi, \mathcal{R})$ is exactly the one used by Atserias and Dalmau [1] to characterise the minimal width of Resolution refutations of φ. In particular in [1] the following result is shown (rephrased here with the notations we just set up).

▶ **Theorem 2** (Atserias and Dalmau [1])**.** *Let φ be an unsatisfiable CNF and $\mathcal{R}$ be the set of all possible partial assignments with a domain of size strictly less than w. The following are equivalent*
1. Prover *has a winning strategy for* Game$(\varphi, \mathcal{R})$*;*
2. width$(\varphi \vdash \bot) < w$.
Due to this equivalence, for this particular choice of $\mathcal{R}$, we will denote Game$(\varphi, \mathcal{R})$ *by* width-Game(φ, w).

The next result is essentially due to Pudlák [21]. He shows that we can also characterise the minimal size of Resolution refutations of φ in terms of these games. From a Resolution refutation Π we can construct a winning strategy for Prover with a set $\mathcal{R}$ of the same size of Π and vice versa. Moreover a play of the Game$_\delta(\varphi, \mathcal{R})$ corresponds to a path in Π and, if Π is δ-regular, in each run the set of variables Prover is going to query many times has size at most δn.

▶ **Theorem 3.** *Let φ be an unsatisfiable CNF and let δ be any real in the interval $[0, 1]$. The following are equivalent*
1. there exists a set of partial assignments $\mathcal{R}$ such that $|\mathcal{R}| \leq s$ for which Prover *has a winning strategy for* Game$_\delta(\varphi, \mathcal{R})$*;*
2. size$_\delta(\varphi \vdash \bot) \leq s$.

4 Games and Xorifications

Given a CNF φ on the variables $x_1, \ldots, x_n$, we define the *ℓ-xorification* of φ as follows: it is a formula on the new variables y_j^i, where $1 \leq i \leq n$ and $1 \leq j \leq \ell$ and it is obtained by replacing each x_i with $y_1^i \oplus \ldots \oplus y_\ell^i$. We denote this formula by $\varphi[\oplus^\ell]$ and note that if φ is a k-CNF, then $\varphi[\oplus^\ell]$ can be expanded to a $k\ell$-CNF. Due to this notation we will refer to the variables of φ as the *x-variables* and to the variables of $\varphi[\oplus^\ell]$ as the *y-variables*. Moreover we say that all the y-variables $y_1^i, \ldots, y_\ell^i$ form a *block* of variables corresponding to the x-variable x_i. We say that a partial assignment over the y-variables *fixes* a value for a x-variable x_i if it assigns all the y-variables in the block corresponding to x_i.

▶ **Theorem 4.** *Let φ an unsatisfiable CNF in n variables and w, δ and ℓ be parameters. If* width$(\varphi \vdash \bot) \geq w$ *then*

$$\text{size}_\delta(\varphi[\oplus^\ell] \vdash \bot) \geq 2^{w\ell(1-\epsilon)},$$

where $\epsilon = \frac{1}{\ell} \log(\frac{e^3 \ell n}{w}) + \frac{\delta n}{w} \log \frac{e^3 \ell}{\delta}$.

Proof. For each partial assignment α over the y-variables there is naturally associated a partial assignment α' over the x-variables, defined as follows

$$\alpha'(x_i) = \begin{cases} \alpha(y_i^1) \oplus \ldots \oplus \alpha_r(y_i^\ell) & \text{if } \forall j = 1, \ldots, \ell, \ y_i^j \in \text{dom}(\alpha), \\ * & \text{otherwise.} \end{cases}$$

By Theorem 3, it is enough to show that if Prover wins Game$_\delta(\varphi[\oplus^\ell], \mathcal{R})$ then

$$|\mathcal{R}| \geq 2^{w(\ell - \log(\frac{e^3 \ell n}{w}) - \frac{\delta \ell n}{w} \log \frac{e^3 \ell}{\delta})}.$$

So suppose Prover wins Game$_\delta(\varphi[\oplus^\ell], \mathcal{R})$ for some set of partial assignments $\mathcal{R}$. Since width$(\varphi \vdash \bot) \geq w$, by Theorem 2, there is a winning strategy σ for Delayer in the game width-Game(φ, w).

For each total assignment β on the y-variables, we consider a strategy σ_β for Delayer in the game $\mathsf{Game}_\delta(\varphi[\oplus^\ell], \mathcal{R})$ as follows. Let α_r be the partial assignment on y-variables at stage r of the game $\mathsf{Game}_\delta(\varphi[\oplus^\ell], \mathcal{R})$ and y_j^i the variable queried at stage $r + 1$. Then the strategy σ_β for Delayer goes as follows:

1. if there exists $j' \neq j$ such that $y_i^{j'} \notin \mathrm{dom}(\alpha_r)$, set y_i^j to $\beta(y_j^i)$;

2. otherwise, if for all $j' \neq j$, $y_i^{j'} \in \mathrm{dom}(\alpha_r)$, then look at the value $b \in \{0, 1\}$ the strategy σ sets the variable x_i when given the partial assignment α_r'. Then set y_i^j to $q \in \{0, 1\}$ such that

$$q \oplus \bigoplus_{j' \neq j} \alpha_r(y_i^j) = b.$$

This can be done since $x_i \equiv y_i^1 \oplus \ldots \oplus y_i^\ell$ and the value of x_i can be set freely to 0 or 1 appropriately even after all but one of $y_i^1, \ldots, y_i^\ell$ have been set.

Since we are assuming that Prover has a winning strategy for $\mathsf{Game}_\delta(\varphi[\oplus^\ell], \mathcal{R})$, in particular, this means that for any β he wins against the Delayer's strategy σ_β. It is immediate to see that for each total assignment β over the y-variables, σ_β is a winning strategy for Delayer in the game $\mathsf{width\text{-}Game}(\varphi[\oplus^\ell], w\ell)$. This means that for each total assignment β over the y-variables, $\mathcal{R}$ must contain some partial assignment, denoted by ρ_β, with domain of size at least $w\ell$ and such that at least w blocks of y-variables are completely fixed by ρ_β. Without loss of generality we assume that each ρ_β fixes exactly w blocks of y-variables, that is if ρ_β is setting more y-variables we simply ignore some of the variables and only consider w blocks. Our goal is to show that we have 'many distinct' such partial assignments ρ_β.

Let $B \subseteq [n]$ denote a generic set of size w and consider for each possible such B the set S_B of the total assignments βs such that ρ_β is fixing all the $y_1^i, \ldots, y_\ell^i$ corresponding to some i in B. There are $2^{n\ell}$ possible total assignments β and $\binom{n}{w}$ possible sets B, hence by the pigeonhole principle, there is a set $B^* \subseteq [n]$ of size w such that

$$|S_{B^*}| \geq \frac{2^{n\ell}}{\binom{n}{w}}. \tag{1}$$

Let S_{B^*}' be the set of partial assignments $\beta|_{B^*}$ where $\beta \in S_{B^*}$. We clearly have that

$$|S_{B^*}| \leq |S_{B^*}'| \cdot 2^{n\ell - \ell|B^*|} = |S_{B^*}'| \cdot 2^{n\ell - w\ell}.$$

By equation (1), we get

$$|S_{B^*}'| \geq \frac{2^{w\ell}}{\binom{n}{w}}. \tag{2}$$

We have now that S_{B^*}' and $\{\rho_\beta : \beta \in S_{B^*}\}$ both consist of assignments of domain $\{y_i^j : i \in B^* \wedge 1 \leq j \leq \ell\}$. We show that $|\{\rho_\beta : \beta \in S_{B^*}\}|$ cannot be too small compared to $|S_{B^*}'|$, this will be, intuitively, due to the fact that the βs we start with are very different.

Let Z^β be the set of variables that Prover re-queried when playing against σ_β and for any $i = 1, \ldots, n$ let $Z_i^\beta = Z^\beta \cap \{y_1^i, \ldots, y_\ell^i\}$. By hypothesis $|Z^\beta| \leq \delta\ell n$.

When Delayer follows the strategy σ_β and fixes all y-variables in a block corresponding to x_i, this assignment is within Hamming distance $|Z_i^\beta| + 1$ from β in this block. This means that for each $\beta \in S_{B^*}$ and for each i, $\rho_\beta|_{\{y_1^i, \ldots, y_\ell^i\}}$ has Hamming distance at most $|Z_i^\beta| + 1$ from some partial assignment in S_{B^*}' restricted to $\{y_1^i, \ldots, y_\ell^i\}$. This means that for each $\beta \in S_{B^*}$ and for each i, ρ_β restricted to the set $\{y_1^i, \ldots, y_\ell^i\}$ has Hamming distance at most

$|Z_i^\beta| + 1$ from some partial assignment in S'_{B^*} restricted to $\{y_1^i, \ldots, y_\ell^i\}$. Let $\mathcal{Z}$ be the set of all possible sets Z subsets of the y-variables of size $\delta\ell n$ such that there exists $\beta \in S_{B^*}$ with $Z^\beta \subseteq Z$. For any $i = 1, \ldots, n$ let $Z_i = Z \cap \{y_1^i, \ldots, y_\ell^i\}$. Then, by counting the variables where ρ_β and an assignment in S'_{B^*} could differ, we have that

$$|S'_{B^*}| \leq |\{\rho_\beta : \beta \in S_{B^*}\}| \cdot \sum_{Z \in \mathcal{Z}} \prod_{i \in B^*} 2^{|Z_i|+1} \binom{\ell}{|Z_i|+1}. \tag{3}$$

Hence we have the following chain of inequalities

$$|S'_{B^*}| \overset{eq.(3)}{\leq} |\{\rho_\beta : \beta \in S_{B^*}\}| \cdot \sum_{Z \in \mathcal{Z}} \prod_{i \in B^*} 2^{|Z_i|+1} \binom{\ell}{|Z_i|+1} \tag{4}$$

$$\leq |\{\rho_\beta : \beta \in S_{B^*}\}| \cdot \sum_{Z \in \mathcal{Z}} \prod_{i \in B^*} \left(\frac{e^2\ell}{|Z_i|+1}\right)^{|Z_i|+1} \tag{5}$$

$$\leq |\{\rho_\beta : \beta \in S_{B^*}\}| \cdot \sum_{Z \in \mathcal{Z}} \left(\frac{\sum_{i \in B^*} e^2\ell}{\sum_{i \in B^*}(|Z_i|+1)}\right)^{\sum_{i \in B^*}(|Z_i|+1)} \tag{6}$$

$$\leq |\{\rho_\beta : \beta \in S_{B^*}\}| \cdot \binom{\ell n}{\delta\ell n} \cdot \left(\frac{\sum_{i \in B^*} e^2\ell}{w}\right)^{\delta\ell n + w} \tag{7}$$

$$= |\{\rho_\beta : \beta \in S_{B^*}\}| \cdot \binom{\ell n}{\delta\ell n} \cdot \left(e^2\ell\right)^{\delta\ell n + w} \tag{8}$$

The inequality (6) follows from the weighted AM-GM inequality[2] and the inequality (7) follows from the fact that $w \leq \sum_{i \in B^*}(|Z_i|+1) \leq \delta\ell n + w$. Putting all together we have that

$$|\mathcal{R}| \overset{(\dagger\dagger)}{\geq} |\{\rho_\beta : \beta \in S_{B^*}\}| \geq \frac{|S'_{B^*}|}{\binom{n\ell}{\delta\ell n}\left(e^2\ell\right)^{\delta\ell n + w}} \overset{(eq.\ 2)}{\geq} \frac{2^{w\ell}}{\binom{n}{w}\binom{\ell n}{\delta\ell n}\left(e^2\ell\right)^{\delta\ell n + w}}$$

$$\geq \frac{2^{w\ell}}{\left(\frac{en}{w}\right)^w \left(\frac{e}{\delta}\right)^{\delta\ell n} \left(e^2\ell\right)^{\delta\ell n + w}}$$

$$= 2^{w\left(\ell - \log\left(\frac{e^3\ell n}{w}\right) - \frac{\delta\ell n}{w}\log\frac{e^3\ell}{\delta}\right)},$$

where the inequality $(\dagger\dagger)$ follows by the definition of ρ_β. ◀

The next step now is to obtain formulas which require very large Resolution width. Such a construction is given by Beck and Impagliazzo in [4].

▶ **Theorem 5** ([4]). *For any large n and k, there exist an unsatisfiable k-CNF formula φ on n variables and some $\zeta_k = \widetilde{O}(k^{-1/4})$ such that*

$$\mathsf{width}(\varphi \vdash \bot) \geq (1 - \zeta_k)n.$$

---

[2] The *weighted Arithmetic Mean - Geometric Mean inequality* says that given non-negative numbers $a_1, \ldots, a_n$ and non-negative weights $w_1, \ldots, w_n$ then

$$\prod_i a_i^{w_i} \leq \left(\frac{\sum_i w_i a_i}{w}\right)^w,$$

where $w = \sum_i w_i$. We applied this inequality with $a_i = e^2\ell$ and $w_i = |Z_i| + 1$.

Now, informally, our SETH lower bound for Resolution will follow from the existence of a CNF requiring very high Resolution width (Theorem 5) and the previous theorem about xorifications (Theorem 4).

▶ **Corollary 6.** *For any large n, k and $\ell = \widetilde{\Theta}(k^{1/4})$, there exists an unsatisfiable k-CNF formula φ on n variables such that*

$$\mathsf{size}_\delta(\varphi[\oplus^\ell] \vdash \bot) \geq 2^{(1-\epsilon_{k'})n\ell},$$

where $k' = k\ell$ is the initial width of the clauses of $\varphi[\oplus^\ell]$ and $\epsilon_{k'} = \delta = \widetilde{O}(k'^{-1/5})$.

Proof. Let φ be the k-CNF formula given by Theorem 5, in particular $\mathsf{width}(\varphi \vdash \bot) \geq (1 - \zeta_k)n$ where $\zeta_k = \widetilde{O}(k^{-1/4})$. Then $\varphi[\oplus^\ell]$ is a k'-CNF on $n\ell$ variables where $k' = k\ell$. By the choice of $\ell = \widetilde{\Theta}(k^{1/4})$, $\delta = \widetilde{O}(k^{-1/4})$ and by Theorem 4, it follows that

$$\mathsf{size}_\delta(\varphi[\oplus^\ell] \vdash \bot) \geq 2^{(1-\zeta_k)n(\ell - \log(\frac{e^3 \ell n}{w}) - \frac{\delta \ell n}{w}\log\frac{e^3 \ell}{\delta})}$$

$$\overset{(\dagger)}{=} 2^{(1-\zeta_k)n(\ell - O(\log k) - \ell\widetilde{O}(k^{-1/4}))} = 2^{(1-\widetilde{O}(k^{-1/4}))n\ell}$$

$$= 2^{(1-\epsilon_{k'})n\ell}.$$

In particular the equality (†) follows from the choice of $\ell = \widetilde{\Theta}(k^{1/4})$ and $\delta = \widetilde{O}(k^{-1/4})$. To obtain the asymptotic behaviour of $\epsilon_{k'}$ with respect to k', just observe that $k' = k\ell = \widetilde{\Theta}(k^{5/4})$ and $\epsilon_{k'} = \widetilde{O}(k^{-1/4})$, hence $\epsilon_{k'} = \widetilde{O}(k'^{-1/5})$. Similarly we get the asymptotic behaviour of δ as a function of k'.   ◀

5   Conclusion

We proved that there exist unsatisfiable k-CNF formulas in n variables that require δ-regular Resolution refutations of size at least $2^{(1-\epsilon)n}$, where $k = \widetilde{O}(\epsilon^{-5})$ and where $\delta = \widetilde{O}(\epsilon^{-5})$. A natural question is whether it is possible to improve the dependency of δ and k on ϵ.

More generally, we have some proof systems stronger than Resolution, such as Polynomial Calculus + Resolution, RES(k), Cutting Planes, for which we know that there are some unsatisfiable CNFs which require exponential size refutations. Are those proof systems consistent with SETH?

Acknowledgments

─── **References** ───

1   Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *J. Comput. Syst. Sci.*, 74(3):323–334, 2008.

2   Albert Atserias, Johannes Klaus Fichte, and Marc Thurley. Clause-learning algorithms with many restarts and bounded-width resolution. *J. Artif. Intell. Res. (JAIR)*, 40:353–373, 2011.

3   Paul Beame, Christopher Beck, and Russell Impagliazzo. Time-space tradeoffs in resolution: superpolynomial lower bounds for superlinear space. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 213–232. ACM, 2012.

4 Christopher Beck and Russell Impagliazzo. Strong ETH Holds for Regular Resolution. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 487–494. ACM, 2013.

5 Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *J. ACM*, 48(2):149–169, 2001.

6 Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, 1937. University of Chicago.

7 Ruiwen Chen, Valentine Kabanets, Antonina Kolokolova, Ronen Shaltiel, and David Zuckerman. Mining circuit lower bound proofs for meta-algorithms. In *IEEE 29th Conference on Computational Complexity, CCC*, pages 262–273, 2014.

8 Ruiwen Chen, Valentine Kabanets, and Nitin Saurabh. An improved deterministic #SAT algorithm for small de morgan formulas. In *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS*, pages 165–176, 2014.

9 Shiteng Chen, Dominik Scheder, Navid Talebanfard, and Bangsheng Tang. Exponential Lower Bounds for the PPSZ k-SAT Algorithm. In *SODA*, pages 1253–1263, 2013.

10 Stefan S. Dantchev. Relativisation provides natural separations for resolution-based proof systems. In *Computer Science - Theory and Applications, First International Computer Science Symposium in Russia, CSR 2006, St. Petersburg, Russia, June 8-12, 2006, Proceedings*, pages 147–158, 2006.

11 Evgeny Dantsin, Andreas Goerdt, Edward A. Hirsch, Ravi Kannan, Jon M. Kleinberg, Christos H. Papadimitriou, Prabhakar Raghavan, and Uwe Schöning. A deterministic $(2-2/(k+1))^n$ algorithm for k-SAT based on local search. *Theor. Comput. Sci.*, 289(1):69–83, 2002.

12 Martin Davis, George Logemann, and Donald W. Loveland. A machine program for theorem-proving. *Commun. ACM*, 5(7):394–397, 1962.

13 Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 1960.

14 Armin Haken. The intractability of resolution. *Theor. Comput. Sci.*, 39:297–308, 1985.

15 Russell Impagliazzo and Ramamohan Paturi. On the Complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.

16 Roberto J. Bayardo Jr. and Robert Schrag. Using CSP look-back techniques to solve real-world SAT instances. In Benjamin Kuipers and Bonnie L. Webber, editors, *Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Innovative Applications of Artificial Intelligence Conference, AAAI 97, IAAI 97, July 27-31, 1997, Providence, Rhode Island.*, pages 203–208. AAAI Press / The MIT Press, 1997.

17 Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. Chaff: Engineering an efficient SAT solver. In *Proceedings of the 38th Design Automation Conference, DAC 2001, Las Vegas, NV, USA, June 18-22, 2001*, pages 530–535. ACM, 2001.

18 Ramamohan Paturi, Pavel Pudlák, Michael E. Saks, and Francis Zane. An improved exponential-time algorithm for k-SAT. *J. ACM*, 52(3):337–364, 2005.

19 Ramamohan Paturi, Pavel Pudlák, and Francis Zane. Satisfiability coding lemma. In *38th Annual Symposium on Foundations of Computer Science, FOCS*, pages 566–574, 1997.

20 Knot Pipatsrisawat and Adnan Darwiche. On the power of clause-learning SAT solvers as resolution engines. *Artif. Intell.*, 175(2):512–525, 2011.

21 Pavel Pudlák. Proofs as games. *American Mathematical Monthly*, 107(6):541–550, 2000.

22 Pavel Pudlák and Russell Impagliazzo. A lower bound for DLL algorithms for k-SAT (preliminary version). In *SODA*, pages 128–136, 2000.

23 Rahul Santhanam. Fighting perebor: New and improved algorithms for formula and QBF satisfiability. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 183–192, 2010.

24 Uwe Schöning. A probabilistic algorithm for k-SAT and constraint satisfaction problems. In *40th Annual Symposium on Foundations of Computer Science, FOCS,*, pages 410–414, 1999.

25 João P. Marques Silva and Karem A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Trans. Computers*, 48(5):506–521, 1999.

26 Alasdair Urquhart. Hard examples for resolution. *J. ACM*, 34(1):209–219, 1987.

27 Ryan Williams. Improving exhaustive search implies superpolynomial lower bounds. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC*, pages 231–240, 2010.

28 Ryan Williams. Non-uniform ACC circuit lower bounds. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity, CCC*, pages 115–125, 2011.