



Towards a Complexity-Theoretic Understanding of Restarts in SAT Solvers

Chunxiao Li¹(✉), Noah Fleming², Marc Vinyals³, Toniann Pitassi²,
and Vijay Ganesh¹

¹ University of Waterloo, Waterloo, Canada
chunxiao.ian.li@gmail.com

² University of Toronto, Toronto, Canada

³ Technion, Haifa, Israel

Abstract. Restarts are a widely-used class of techniques integral to the efficiency of Conflict-Driven Clause Learning (CDCL) Boolean SAT solvers. While the utility of such policies has been well-established empirically, a theoretical understanding of whether restarts are indeed crucial to the power of CDCL solvers is missing.

In this paper, we prove a series of theoretical results that characterize the power of restarts for various models of SAT solvers. More precisely, we make the following contributions. First, we prove an exponential separation between a *drunk* randomized CDCL solver model with restarts and the same model without restarts using a family of satisfiable instances. Second, we show that the configuration of CDCL solver with VSIDS branching and restarts (with activities erased after restarts) is exponentially more powerful than the same configuration without restarts for a family of unsatisfiable instances. To the best of our knowledge, these are the first separation results involving restarts in the context of SAT solvers. Third, we show that restarts do not add any proof complexity-theoretic power vis-a-vis a number of models of CDCL and DPLL solvers with non-deterministic static variable and value selection.

1 Introduction

Over the last two decades, Conflict-Driven Clause Learning (CDCL) SAT solvers have had a revolutionary impact on many areas of software engineering, security and AI. This is primarily due to their ability to solve real-world instances containing millions of variables and clauses [2, 6, 15, 16, 18], despite the fact that the Boolean SAT problem is known to be an NP-complete problem and is believed to be intractable in the worst case.

This remarkable success has prompted complexity theorists to seek an explanation for the efficacy of CDCL solvers, with the aim of bridging the gap between theory and practice. Fortunately, a few results have already been established that lay the groundwork for a deeper understanding of SAT solvers viewed as proof systems [3, 8, 11]. Among them, the most important result is the one by Pipatsrisawat and Darwiche [18] and independently by Atserias et al. [2], that shows that

an idealized model of CDCL solvers with non-deterministic branching (variable selection and value selection), and restarts is *polynomially equivalent* to the general resolution proof system. However, an important question that remains open is whether this result holds even when restarts are disabled, i.e., whether configurations of CDCL solvers without restarts (when modeled as proof systems) are polynomial equivalent to the general resolution proof system. In practice there is significant evidence that restarts are crucial to solver performance.

This question of the “power of restarts” has prompted considerable theoretical work. For example, Bonet, Buss and Johannsen [7] showed that CDCL solvers with no restarts (but with non-deterministic variable and value selection) are strictly more powerful than regular resolution. Despite this progress, the central questions, such as whether restarts are integral to the efficient simulation of general resolution by CDCL solvers, remain open.

In addition to the aforementioned theoretical work, there have been many empirical attempts at understanding restarts given how important they are to solver performance. Many hypotheses have been proposed aimed at explaining the power of restarts. Examples include, the *heavy-tail* explanation [10], and the “restarts compact assignment trail and hence produce clauses with lower literal block distance (LBD)” perspective [14]. Having said that, the heavy-tailed distribution explanation of the power of restarts is not considered valid anymore in the CDCL setting [14].

1.1 Contributions

In this paper we make several contributions to the theoretical understanding of the power of restarts for several restricted models of CDCL solvers:

1. First, we show that CDCL solvers with backtracking, non-deterministic dynamic variable selection, randomized value selection, and restarts¹ are exponentially faster than the same model, but without restarts, with high probability (w.h.p)². A notable feature of our proof is that we obtain this separation on a family of satisfiable instances. (See Sect. 4 for details.)
2. Second, we prove that CDCL solvers with VSIDS variable selection, phase saving value selection and restarts (where activities of variables are reset to zero after restarts) are exponentially faster (w.h.p) than the same solver configuration but without restarts for a class of unsatisfiable formulas. This result holds irrespective of whether the solver uses backtracking or backjumping. (See Sect. 5 for details.)
3. Finally, we prove several smaller separation and equivalence results for various configurations of CDCL and DPLL solvers with and without restarts. For example, we show that CDCL solvers with non-deterministic static variable

¹ In keeping with the terminology from [1], we refer any CDCL solver with randomized value selection as a *drunk* solver.

² We say that an event occurs with high probability (w.h.p.) if the probability of that event happening goes to 1 as $n \rightarrow \infty$.

selection, non-deterministic static value selection, and with restarts, are polynomially equivalent to the same model but without restarts. Another result we show is that for DPLL solvers, restarts do not add proof theoretic power as long as the solver configuration has non-deterministic dynamic variable selection. (See Sect. 6 for details.)

2 Definitions and Preliminaries

Below we provide relevant definitions and concepts used in this paper. We refer the reader to the Handbook of Satisfiability [6] for literature on CDCL and DPLL solvers and to [4, 12] for literature on proof complexity.

We denote by $[c]$ the set of natural numbers $\{1, \dots, c\}$. We treat CDCL solvers as proof systems. For proof systems A and B , we use $A \sim_p B$ to denote that they are polynomially equivalent (p -equivalent). Throughout this paper it is convenient to think of the trail π of the solver during its run on a formula F as a restriction to that formula. We call a function $\pi: \{x_1, \dots, x_n\} \rightarrow \{0, 1, *\}$ a *restriction*, where $*$ denotes that the variable is unassigned by π . Additionally, we assume that our Boolean Constraint Propagation (BCP) scheme is greedy, i.e., BCP is performed till “saturation”.

Restarts in SAT Solvers. A restart policy is a method that erases part of the state of the solver at certain intervals during the run of a solver [10]. In most modern CDCL solvers, the restart policy erases the assignment trail upon invocation, but may choose not to erase the learnt clause database or variable activities. Throughout this paper, we assume that all restart policies are non-deterministic, i.e., the solver may (dynamically) non-deterministically choose its restart sequence. We refer the reader to a paper by Liang et al. [14] for a detailed discussion on modern restart policies.

3 Notation for Solver Configurations Considered

In this section, we precisely define the various heuristics used to define SAT solver configurations in this paper. By the term *solver configuration* we mean a solver parameterized with appropriate heuristic choices. For example, a CDCL solver with non-deterministic variable and value selection, as well as asserting learning scheme with restarts would be considered a solver configuration.

To keep track of these configurations, we denote solver configurations by the notation $M_{A,B}^{E,R}$, where M indicates the underlying *solver model* (we use C for CDCL and D for DPLL solvers); the subscript A denotes a *variable selection scheme*; the subscript B is a *value selection scheme*; the superscript E is a *backtracking scheme*, and finally the superscript R indicates whether the solver configuration comes equipped with a restart policy. That is, the presence of the superscript R indicates that the configuration has restarts, and its absence indicates that it does not. A $*$ in place of A, B or E denotes that the scheme is *arbitrary*, meaning that it works for any such scheme. See Table 1 for examples of solver configurations studied in this paper.

Table 1. Solver configurations in the order they appear in the paper. ND stands for non-deterministic dynamic.

	Model	Variable selection	Value selection	Backtracking	Restarts
$C_{ND,RD}^{T,R}$	CDCL	ND	Random dynamic	Backtracking	Yes
$C_{ND,RD}^T$	CDCL	ND	Random dynamic	Backtracking	No
$C_{VS,PS}^{J,R}$	CDCL	VSIDS	Phase saving	Backjumping	Yes
$C_{VS,PS}^J$	CDCL	VSIDS	Phase saving	Backjumping	No
$C_{S,S}^{J,R}$	CDCL	Static	Static	Backjumping	Yes
$C_{S,S}^J$	CDCL	Static	Static	Backjumping	No
$D_{ND,*}^T$	DPLL	ND	Arbitrary	Backtracking	No
$D_{ND,ND}^{T,R}$	DPLL	ND	ND	Backtracking	Yes
$D_{ND,ND}^T$	DPLL	ND	ND	Backtracking	No
$D_{ND,RD}^{T,R}$	DPLL	ND	Random dynamic	Backtracking	Yes
$D_{ND,RD}^T$	DPLL	ND	Random dynamic	Backtracking	No
$C_{ND,ND}^{J,R}$	CDCL	ND	ND	Backjumping	Yes
$C_{ND,ND}^J$	CDCL	ND	ND	Backjumping	No

3.1 Variable Selection Schemes

1. Static (S): Upon invocation, the S variable selection heuristic returns the unassigned variable with the highest rank according to some predetermined, fixed, total ordering of the variables.

2. Non-deterministic Dynamic (ND): The ND variable selection scheme non-deterministically selects and returns an unassigned variable.

3. VSIDS (VS) [16]: Each variable has an associated number, called its *activity*, initially set to 0. Each time a solver learns a conflict, the activities of variables appearing on the conflict side of the implication graph receive a constant bump. The activities of all variables are decayed by a constant c , where $0 < c < 1$, at regular intervals. The VSIDS variable selection heuristic returns the unassigned variable with highest activity, with ties broken randomly.

3.2 Value Selection Schemes

1. Static (S): Before execution, a 1-1 mapping of variables to values is fixed. The S value selection heuristic takes as input a variable and returns the value assigned to that variable according to the predetermined mapping.

2. Non-deterministic Dynamic (ND): The ND value selection scheme non-deterministically selects and returns a truth assignment.

3. Random Dynamic (RD): A randomized algorithm that takes as input a variable and returns a uniformly random truth assignment.

4. Phase Saving (PS): A heuristic that takes as input an unassigned variable and returns the previous truth value that was assigned to the variable. Typically

solver designers determine what value is returned when a variable has not been previously assigned. For simplicity, we use the phase saving heuristic that returns 0 if the variable has not been previously assigned.

3.3 Backtracking and Backjumping Schemes

To define different backtracking schemes we use the concept of *decision level* of a variable x , which is the number of decision variables on the trail prior to x . **Backtracking (T):** Upon deriving a conflict clause, the solver undoes the most recent decision variable on the assignment trail. **Backjumping (J):** Upon deriving a conflict clause, the solver undoes all decision variables with decision level higher than the variable with the second highest decision level in the conflict clause.

Note on Solver Heuristics. Most of our results hold irrespective of the choice of deterministic asserting clause learning schemes (except for Proposition 22). Additionally, it goes without saying that the questions we address in this paper make sense only when it is assumed that solver heuristics are polynomial time methods.

4 Separation for Drunk CDCL with and Without Restarts

Inspired by Alekhovich et al. [1], where the authors proved exponential lower bound for drunk DPLL solvers over a class of satisfiable instances, we studied the behavior of restarts in a drunken model of CDCL solver. We introduce a class of satisfiable formulas, $Ladder_n$, and use them to prove the separation between $C_{ND,RD}^{T,R}$ and $C_{ND,RD}^T$. At the core of these formulas is a formula which is hard for general resolution even after any small restriction (corresponding to the current trail of the solver). For this, we use the well-known Tseitin formulas.

Definition 1 (Tseitin Formulas). Let $G = (V, E)$ be a graph and $f: V \rightarrow \{0, 1\}$ a labelling of the vertices. The formula $Tseitin(G, f)$ has variables x_e for $e \in E$ and constraints $\bigoplus_{uv \in E} x_{uv} = f(v)$ for each $v \in V$.

For any graph G , $Tseitin(G, f)$ is unsatisfiable iff $\bigoplus_{v \in V} f(v) = 1$, in which case we call f an *odd labelling*. The specifics of the labelling are irrelevant for our applications, any odd labelling will do. Therefore, we often omit defining f , and simply assume that it is odd.

The family of satisfiable $Ladder_n$ formulas are built around the Tseitin formulas, unless the variables of the formula are set to be consistent to one of two satisfying assignments, the formula will become unsatisfiable. Furthermore, the solver will only be able to backtrack out of the unsatisfiable sub-formula by first refuting Tseitin, which is a provably hard task for any CDCL solver [20].

The $Ladder_n$ formulas contain two sets of variables, ℓ_j^i for $0 \leq i \leq n-2, j \in [\log n]$ and c_m for $m \in [\log n]$, where n is a power of two. We denote by ℓ^i

the block of variables $\{\ell_1^i, \dots, \ell_{\log n}^i\}$. These formulas are constructed using the following gadgets.

Ladder gadgets: $L^i := (\ell_1^i \vee \dots \vee \ell_{\log n}^i) \wedge (\neg \ell_1^i \vee \dots \vee \neg \ell_{\log n}^i)$.

Observe that L^i is falsified only by the all-1 and all-0 assignments.

Connecting gadgets: $C^i := (c_1^{\text{bin}(i,1)} \wedge \dots \wedge c_{\log n}^{\text{bin}(i,\log n)})$.

Here, $\text{bin}(i, m)$ returns the m th bit of the binary representation of i , and $c_m^1 := c_m$, while $c_m^0 := \neg c_m$. That is, C^i is the conjunction that is satisfied only by the assignment encoding i in binary.

Equivalence gadget: $EQ := \bigwedge_{i,j=0}^{n-2} \bigwedge_{m,k=1}^{\log n} (\ell_k^i \iff \ell_k^j)$.

These clauses enforce that every ℓ -variable must take the same value.

Definition 2 (Ladder formulas). For $G = (V, E)$ with $|E| = n - 1$ where n is a power of two, let $T\text{seitin}(G, f)$ be defined on the variables $\{\ell_1^0, \dots, \ell_1^{n-2}\}$. $Ladder_n(G, f)$ is the conjunction of the clauses representing

$$\begin{aligned} L^i &\Rightarrow C^i, & \forall 0 \leq i \leq n - 2 \\ C^i &\Rightarrow T\text{seitin}(G, f), & \forall 0 \leq i \leq n - 2 \\ C^{n-1} &\Rightarrow EQ. \end{aligned}$$

Observe that the $Ladder_n(G, f)$ formulas have polynomial size provided that the degree of G is $O(\log n)$. As well, this formula is satisfiable only by the assignments that sets $c_m = 1$ and $\ell_j^i = \ell_q^p$ for every $m, j, q \in [\log n]$ and $0 \leq i, p \leq n - 2$.

These formulas are constructed so that after setting only a few variables, any drunk solver will enter an unsatisfiable subformula w.h.p. and thus be forced to refute the Tseitin formula. Both the ladder gadgets and equivalence gadget act as trapdoors for the Tseitin formula. Indeed, if any c -variable is set to 0 then we have already entered an unsatisfiable instance. Similarly, setting $\ell_j^i = 1$ and $\ell_q^p = 0$ for any $0 \leq i, p \leq n - 2, j, q \in [\log n]$ causes us to enter an unsatisfiable instance. This is because setting all c -variables to 1 together with this assignment would falsify a clause of the equivalence gadget. Thus, after the second decision of the solver, the probability that it is in an unsatisfiable instance is already at least $1/2$. With these formulas in hand, we prove the following theorem, separating backtracking $C_{ND, RD}^T$ solvers with and without restarts.

Theorem 3. *There exists a family of $O(\log n)$ -degree graphs G such that*

1. $Ladder_n(G, f)$ can be decided in time $O(n^2)$ by $C_{ND, RD}^{T, R}$, except with exponentially small probability.
2. $C_{ND, RD}^T$ requires exponential time to decide $Ladder_n(G, f)$, except with probability $O(1/n)$.

The proof of the preceding theorem occupies the remainder of this section.

4.1 Upper Bound on Ladder Formulas via Restarts

We present the proof for part (1) of Theorem 3. The proof relies on the following lemma, stating that given the all-1 restriction to the c -variables, $C_{ND, RD}^T$ will find a satisfying assignment.

Lemma 4. *For any graph G , $C_{ND, RD}^T$ will find a satisfying assignment to $Ladder_n(G, f)[c_1 = 1, \dots, c_{\log n} = 1]$ in time $O(n \log n)$.*

Proof. When all c variables are 1, we have $C^{n-1} = 1$. By the construction of the connecting gadget, $C^i = 0$ for all $0 \leq i \leq n - 2$. Under this assignment, the remaining clauses belong to EQ , along with $\neg L^i$ for $0 \leq i \leq n - 2$. It is easy to see that, as soon as the solver sets an ℓ -variable, these clauses will propagate the remaining ℓ -variables to the same value. \square

Put differently, the set of c variables forms a *weak backdoor* [23, 24] for $Ladder_n$ formulas. Part (1) of Theorem 3 shows that, with probability at least $1/2$, $C_{ND, RD}^{T, R}$ can exploit this weak backdoor using only $O(n)$ number of restarts.

Proof (of Theorem 3 Part (1)). By Lemma 4, if $C_{ND, RD}^{T, R}$ is able to assign all c variables to 1 before assigning any other variables, then the solver will find a satisfying assignment in time $O(n \log n)$ with probability 1. We show that the solver can exploit restarts in order to find this assignment. The strategy the solver adopts is as follows: query each of the c -variables; if at least one of the c -variables was assigned to 0, restart. We argue that if the solver repeats this procedure $k = n^2$ times then it will find the all-1 assignment to the c -variables, except with exponentially small probability. Because each variable is assigned 0 and 1 with equal probability, the probability that a single round of this procedure finds the all-1 assignment is $2^{-\log n}$. Therefore, the probability that the solver has not found the all-1 assignment after k rounds is

$$(1 - 1/n)^k \leq e^{-k/n} = e^{-n}.$$

\square

4.2 Lower Bound on Ladder Formulas Without Restarts

We now prove part (2) of Theorem 3. The proof relies on the following three technical lemmas. The first claims that the solver is well-behaved (most importantly that it cannot learn any new clauses) while it has not made many decisions.

Lemma 5. *Let G be any graph of degree at least d . Suppose that $C_{ND, RD}^T$ has made $\delta < \min(d - 1, \log n - 1)$ decisions since its invocation on $Ladder_n(G, f)$. Let π_δ be the current trail, then*

1. *The solver has yet to enter a conflict, and thus has not learned any clauses.*
2. *The trail π_δ contains variables from at most δ different blocks ℓ^i .*

We defer the proof of this lemma to the arXiv version of the paper [13].

The following technical lemma states that if a solver with backtracking has caused the formula to become unsatisfiable, then it must *refute* that formula before it can backtrack out of it. For a restriction π and a formula F , we say that the solver has *produced a refutation* of an unsatisfiable formula $F[\pi]$ if it has learned a clause C such that C is falsified under π . Note that because general resolution p -simulates CDCL, any refutation of a formula $F[\pi]$ implies a general resolution refutation of $F[\pi]$ of size at most polynomial in the time that the solver took to produce that refutation.

Lemma 6. *Let F be any propositional formula, let π be the current trail of the solver, and let x be any literal in π . Then, $C_{ND,ND}^T$ backtracks x only after it has produced a refutation of $F[\pi]$.*

Proof. In order to backtrack x , the solver must have learned a clause C asserting the negation of some literal $z \in \pi$ that was set before x . Therefore, C must only contain the negation of literals in π . Hence, $C[\pi] = \emptyset$. \square

The third lemma reduces proving a lower bound on the runtime of $C_{ND,ND}^T$ on the $Ladder_n$ formulas under any well-behaved restriction to proving a general resolution lower bound on an associated Tseitin formula.

Definition 7. *For any unsatisfiable formula F , denote by $Res(F \vdash \emptyset)$ the minimal size of any general resolution refutation of F .*

We say that a restriction (thought of as the current trail of the solver) π to $Ladder_n(G, f)$ implies Tseitin if π either sets some c -variable to 0 or $\pi[\ell_j^i] = 1$ and $\pi[\ell_q^p] = 0$ for some $0 \leq i, q \leq n - 2, j, q \in [\log n]$. Observe that in both of these cases the formula $Ladder_n(G, f)[\pi]$ is unsatisfiable.

Lemma 8. *Let π be any restriction that implies Tseitin and such that each clause of $Ladder_n(G, f)[\pi]$ is either satisfied or contains at least two unassigned variables. Suppose that π sets variables from at most δ blocks ℓ^i . Then there is a restriction ρ_π^* that sets at most δ variables of Tseitin(G, f) such that*

$$Res(Ladder_n(G, f)[\pi] \vdash \emptyset) \geq Res(Tseitin(G, f)[\rho_\pi^*] \vdash \emptyset).$$

We defer the proof of this lemma to the arXiv version of the paper [13], and show how to use them to prove part (2) of Theorem 3. We prove this statement for any degree $O(\log n)$ graph G with sufficient expansion.

Definition 9. *The expansion of a graph $G = (V, E)$ is*

$$e(G) := \min_{V' \subseteq V, |V'| \leq |V|/2} \frac{|E[V', V \setminus V']|}{|V'|},$$

where $E[V', V \setminus V']$ is the set of edges in E with one endpoint in V' and the other in $V \setminus V'$.

For every $d \geq 3$, Ramanujan Graphs provide an infinite family of d -regular expander graphs G for which $e(G) \geq d/4$. The lower bound on solver runtime relies on the general resolution lower bounds for the Tseitin formulas [20]; we use the following lower bound criterion which follows immediately³ from [5].

Corollary 10 ([5]). *For any connected graph $G = (V, E)$ with maximum degree d and odd weight function f ,*

$$Res(Tseitin(G, f) \vdash \emptyset) = \exp \left(\Omega \left(\frac{(e(G)|V|/3 - d)^2}{|E|} \right) \right)$$

³ In particular, this follows from Theorem 4.4 and Corollary 3.6 in [5], noting that the definition of expansion used in their paper is lower bounded by $3e(G)/|V|$ as they restrict to sets of vertices of size between $|V|/3$ and $2|V|/3$.

We are now ready to prove the theorem.

Proof (of part (2) Theorem 3). Fix $G = (V, E)$ to be any degree- $(8 \log n)$ graph on $|E| = n - 1$ edges such that $e(G) \geq 2 \log n$. Ramanujan graphs satisfy these conditions.

First, we argue that within $\delta < \log n - 1$ decisions from the solver's invocation, the trail π_δ will imply Tseitin, except with probability $1 - 1/2^{\delta-1}$. By Lemma 5, the solver has yet to backtrack or learn any clauses, and it has set variables from at most δ blocks ℓ^i . Let x be the variable queried during the δ th decision. If x is a c variable, then with probability $1/2$ the solver sets $c_i = 0$. If x is a variable ℓ_j^i , then, unless this is the first time the solver sets an ℓ -variable, the probability that it sets ℓ_j^i to a different value than the previously set ℓ -variable is $1/2$.

Conditioning on the event that, within the first $\log n - 2$ decisions the trail of the solver implies Tseitin (which occurs with probability at least $(n - 8)/n$), we argue that the runtime of the solver is exponential in n . Let $\delta < \log n - 1$ be the first decision level such that the current trail π_δ implies Tseitin. By Lemma 6 the solver must have produced a refutation of $Ladder_n(G, f)[\pi_\delta]$ in order to backtrack out of the unsatisfying assignment. If the solver takes t steps to refute $Ladder_n(G, f)[\pi_\delta]$ then this implies a general resolution refutation of size $\text{poly}(t)$. Therefore, in order to lower bound the runtime of the solver, it is enough to lower bound the size of general resolution refutations of $Ladder_n(G, f)[\pi_\delta]$.

By Lemma 5, the solver has not learned any clauses, and has yet to enter into a conflict and therefore no clause in $Ladder_n(G, f)[\pi_\delta]$ is falsified. As well, π_δ sets variables from at most $\delta < \log n - 1$ blocks ℓ^i . By Lemma 8 there exists a restriction ρ_π^* such that $Res(Ladder_n(G, f)[\pi] \vdash \emptyset) \geq Res(Tseitin(G, f)[\rho_\pi^*] \vdash \emptyset)$. Furthermore, ρ_π^* sets at most $\delta < \log n - 1$ variables and therefore cannot falsify any constraint of $Tseitin(G, f)$, as each clause depends on $8 \log n$ variables. Observe that if we set a variable x_e of $Tseitin(G, f)$ then we obtain a new instance of $Tseitin(G_{\rho_\pi^*}, f')$ on a graph $G_{\rho_\pi^*} = (V, E \setminus \{e\})$. Therefore, we are able to apply Corollary 10 provided that we can show that $e(G_{\rho_\pi^*})$ is large enough.

Claim 11. *Let $G = (V, E)$ be a graph and let $G' = (V, E')$ be obtained from G by removing at most $e(G)/2$ edges. Then $e(G') \geq e(G)/2$.*

Proof. Let $V' \subseteq V$ with $|V'| \leq |V|/2$. Then, $E'[V', V \setminus V'] \geq e(G)|V'| - e(G)/2 \geq (e(G)/2)|V'|$. □

It follows that $e(G_{\rho_\pi^*}) \geq \log n$. Note that $|V| = n/8 \log n$. By Corollary 10,

$$Res(Ladder_n(G, f)[\pi] \vdash \emptyset) = \exp(\Omega(((n - 1)/24 - 8 \log n)^2/n)) = \exp(\Omega(n)).$$

Therefore, the runtime of $C_{ND, ND}^T$ is $\exp(\Omega(n))$ on $Ladder_n(G, F)$ w.h.p. □

5 CDCL + VSIDS Solvers with and Without Restarts

In this section, we prove that CDCL solvers with VSIDS variable selection, phase saving value selection and restarts (where activities of variables are reset to zero after restarts) are exponentially more powerful than the same solver configuration but without restarts, w.h.p.

Theorem 12. *There is a family of unsatisfiable formulas that can be decided in polynomial time with $C_{VS,PS}^{J,R}$ but requires exponential time with $C_{VS,PS}^J$, except with exponentially small probability.*

We show this separation using pitfall formulas $\Phi(G_n, f, n, k)$, designed to be hard for solvers using the VSIDS heuristic [22]. We assume that G_n is a constant-degree expander graph with n vertices and m edges, $f: V(G_n) \rightarrow \{0, 1\}$ is a function with odd support as with Tseitin formulas, we think of k as a constant and let n grow. We denote the indicator function of a Boolean expression B with $\llbracket B \rrbracket$. These formulas have k blocks of variables named X_j, Y_j, Z_j, P_j , and A_j , with $j \in [k]$, and the following clauses:

- $(\bigoplus_{e \ni v} x_{j,e} = f(v)) \vee \bigvee_{i=1}^n z_{j,i}$, expanded into CNF, for $v \in V(G_n)$ and $j \in [k]$;
- $y_{j,i_1} \vee y_{j,i_2} \vee \neg p_{j,i_3}$ for $i_1, i_2 \in [n], i_1 < i_2, i_3 \in [m+n]$, and $j \in [k]$;
- $y_{j,i_1} \vee \bigvee_{i \in [m+n] \setminus \{i_2\}} p_{j,i} \vee \bigvee_{i=1}^{i_2-1} x_{j,i} \vee \neg x_{j,i_2}$ for $i_1 \in [n], i_2 \in [m]$, and $j \in [k]$;
- $y_{j,i_1} \vee \bigvee_{i \in [m+n] \setminus \{m+i_2\}} p_{j,i} \vee \bigvee_{i=1}^m x_{j,i} \vee \bigvee_{i=1+\llbracket i_2=n \rrbracket}^{i_2-1} z_{j,i} \vee \neg z_{j,i_2}$ for $i_1, i_2 \in [n]$ and $j \in [k]$;
- $\neg a_{j,1} \vee a_{j,3} \vee \neg z_{j,i_1}, \neg a_{j,2} \vee \neg a_{j,3} \vee \neg z_{j,i_1}, a_{j,1} \vee \neg z_{j,i_1} \vee \neg y_{j,i_2}$, and $a_{j,2} \vee \neg z_{j,i_1} \vee \neg y_{j,i_2}$ for $i_1, i_2 \in [n]$ and $j \in [k]$; and
- $\bigvee_{j \in [k]} \neg y_{j,i} \vee \neg y_{j,i+1}$ for odd $i \in [n]$.

To give a brief overview, the first type of clauses are essentially a Tseitin formula and thus are hard to solve. The next four types form a pitfall gadget, which has the following easy-to-check property.

Claim 13. *Given any pair of variables y_{j,i_1} and y_{j,i_2} from the same block Y_j , assigning $y_{j,i_1} = 0$ and $y_{j,i_2} = 0$ yields a conflict.*

Furthermore, such a conflict involves all of the variables of a block X_j , which makes the solver prioritize these variables and it becomes stuck in a part of the search space where it must refute the first kind of clauses. Proving this formally requires a delicate argument, but we can use the end result as a black box.

Theorem 14 ([22, Theorem 3.6]). *For k fixed, $\Phi(G_n, f, n, k)$ requires time $\exp(\Omega(n))$ to decide with $C_{VS,PS}^J$, except with exponentially small probability.*

The last type of clauses, denoted by Γ_i , ensure that a short general resolution proof exists. Not only that, we can also prove that pitfall formulas have small backdoors [23, 24], which is enough for a formula to be easy for $C_{VS,PS}^{J,R}$.

Definition 15. *A set of variables V is a strong backdoor for unit-propagation if every assignment to all variables in V leads to a conflict, after unit propagation.*

Lemma 16. *If F has a strong backdoor for unit-propagation of size c , then $C_{VS,PS}^{J,R}$ can solve F in time $n^{O(c)}$, except with exponentially small probability.*

Proof. We say that the solver learns a beneficial clause if it only contains variables in V . Since there are 2^c possible assignments to variables in V and each beneficial clause forbids at least one assignment, it follows that learning 2^c beneficial clauses is enough to produce a conflict at level 0.

Therefore it is enough to prove that, after each restart, we learn a beneficial clause with large enough probability. Since all variables are tied, all decisions before the first conflict after a restart are random, and hence with probability at least n^{-c} the first variables to be decided before reaching the first conflict are (a subset of) V . If this is the case then, since V is a strong backdoor, no more decisions are needed to reach a conflict, and furthermore all decisions in the trail are variables in V , hence the learned clause is beneficial.

It follows that the probability of having a sequence of n^{2c} restarts without learning a beneficial clause is at most

$$(1 - n^{-c})^{n^{2c}} \leq \exp(-n^{-c} \cdot n^{2c}) = \exp(-n^c) \quad (1)$$

hence by a union bound the probability of the algorithm needing more than $2^c \cdot n^{2c}$ restarts is at most $2^c \cdot \exp(-n^c)$. \square

We prove Theorem 12 by showing that $\Phi(G_n, f, n, k)$ contains a backdoor of size $2k(k+1)$.

Proof (of Theorem 12). We claim that the set of variables $V = \{y_{j,i} \mid (j,i) \in [k] \times [2k+2]\}$ is a strong backdoor for unit-propagation. Consider any assignment to V . Each of the $k+1$ clauses $\Gamma_1, \Gamma_3, \dots, \Gamma_{2k+1}$ forces a different variable $y_{j,i}$ to 0, hence by the pigeonhole principle there is at least one block with two variables assigned to 0. But by Claim 13, this is enough to reach a conflict.

The upper bound follows from Lemma 16, while the lower bound follows from Theorem 14. \square

6 Minor Equivalences and Separations for CDCL/DPLL Solvers with and Without Restarts

In this section, we prove four smaller separation and equivalence results for various configurations of CDCL and DPLL solvers with and without restarts.

6.1 Equivalence Between CDCL Solvers with Static Configurations with and Without Restarts

First, we show that CDCL solvers with non-deterministic static variable and value selection without restarts ($C_{S,S}^J$) is as powerful as the same configuration with restarts ($C_{S,S}^{J,R}$) for both satisfiable and unsatisfiable formulas. We assume that the BCP subroutine for the solver configurations under consideration is “fixed” in the following sense: if there is more than one unit clause under a partial assignment, the BCP subroutine propagates the clause that is added to the clause database first.

Theorem 17. $C_{S,S}^J \sim_p C_{S,S}^{J,R}$ provided that they are given the same variable ordering and fixed mapping of variables to values for the variable selection and value selection schemes respectively.

We prove this theorem by arguing for any run of $C_{S,S}^{J,R}$, that restarts can be removed without increasing the run-time.

Proof. Consider a run of $C_{S,S}^{J,R}$ on some formula F , and suppose that the solver has made t restart calls. Consider the trail π for $C_{S,S}^{J,R}$ up to the variable l from the second highest decision from the last learnt clause before the first restart call. Now, observe that because the decision and variable selection orders are static, once $C_{S,S}^{J,R}$ restarts, it will force it to repeat the same decisions and unit propagations that brought it to the trail π . Suppose that this is not the case and consider the first literal on which the trails differ. This difference could not be caused by a unit propagation as the solver has not learned any new clauses since the restart. Thus, it must have been caused by a decision. However, because the clause databases are the same, this would contradict the static variable and value order. Therefore, this restart can be ignored, and we obtain a run of $C_{S,S}^{J,R}$ with $d - 1$ restarts without increasing the run-time. The proof follows by induction. Once all restarts have been removed, the result is a valid run of $C_{S,S}^J$. \square

Note that in the proof of Theorem 17, not only we argue that $C_{S,S}^J$ is p -equivalent to $C_{S,S}^{J,R}$, we also show that the two configurations produce the same run. The crucial observation is that given any state of $C_{S,S}^{J,R}$, we can produce a run of $C_{S,S}^J$ which ends in the same state. In other words, our proof not only suggests that $C_{S,S}^{J,R}$ is equivalent to $C_{S,S}^J$ from a proof theoretic point of view, it also implies that the two configurations are equivalent for satisfiable formulas.

6.2 Equivalence Between DPLL Solvers with ND Variable Selection on UNSAT Formulas

We show that when considered as a proof system, a DPLL solver with non-deterministic dynamic variable selection, arbitrary value selection and no restarts ($D_{ND,*}^T$) is p -equivalent to DPLL solver with non-deterministic dynamic variable and value selection and restarts ($D_{ND,ND}^{T,R}$), and hence, transitively p -equivalent to tree-like resolution—the restriction of general resolution where each consequent can be an antecedent in only one later inference.

Theorem 18. $D_{ND,*}^T \sim_p D_{ND,ND}^T$.

Proof. To show that $D_{ND,ND}^T$ p -simulates $D_{ND,*}^T$, we argue that every proof of $D_{ND,ND}^T$ can be converted to a proof of same size in $D_{ND,*}^T$. Let F be an unsatisfiable formula. Recall that a run of $D_{ND,ND}^T$ on F begins with non-deterministically picking some variable x to branch on, and a truth value to assign to x . W.l.o.g. suppose that the solver assigns x to 1. Thus, the solver will first refute $F[x = 1]$ before backtracking and refuting $F[x = 0]$.

To simulate a run of $D_{ND,ND}^T$ with $D_{ND,*}^T$, since variable selection is non-deterministic, $D_{ND,*}^T$ also chooses the variable x as the first variable to branch on. If the value selection returns $x = \alpha$ for $\alpha \in \{0, 1\}$, then the solver focus on the restricted formula $F[x = \alpha]$ first. Because there is no clause learning, whether $F[x = 1]$ or $F[x = 0]$ is searched first does not affect the size of the search space for the other. The proof follows by recursively calling $D_{ND,*}^T$ on $F[x = 1]$ and $F[x = 0]$. The converse direction follows since every run of $D_{ND,*}^T$ is a run of $D_{ND,ND}^T$. \square

Corollary 19. $D_{ND,*}^T \sim_p D_{ND,ND}^{T,R}$.

Proof. This follows from the fact that $D_{ND,ND}^{T,R} \sim_p D_{ND,ND}^T$. Indeed, with non-deterministic branching and without clause learning, restarts cannot help. If ever $D_{ND,ND}^{T,R}$ queries a variable $x = \alpha$ for $\alpha \in \{0, 1\}$ and then later restarts to assign it to $1 - \alpha$, then $D_{ND,ND}^T$ ignores the part of the computation when $x = \alpha$ and instead immediately non-deterministically chooses $x = 1 - \alpha$. \square

It is interesting to note that while the above result establishes a p -equivalence between DPLL solver models $D_{ND,*}^T$ and $D_{ND,ND}^{T,R}$, the following corollary implies that DPLL solvers with non-deterministic variable and randomized value selection are exponentially separable for satisfiable instances.

6.3 Separation Result for Drunk DPLL Solvers

We show that DPLL solvers with non-deterministic variable selection, randomized value selection and no restarts ($D_{ND,RD}^T$) is exponentially weaker than the same configuration with restarts ($D_{ND,RD}^{T,R}$).

Corollary 20. $D_{ND,RD}^T$ runs exponentially slower on the class of satisfiable formulas $Ladder_n(G, f)$ than $D_{ND,RD}^{T,R}$, with high probability.

The separation follows from the fact that our proof of the upper bound from Theorem 3 does not use the fact the solver has access to clause learning, which means the solver $D_{ND,RD}^{T,R}$ can also find a satisfying assignment for $Ladder_n(G, f)$ in time $O(n^2)$, except with exponentially small probability. On the other hand, the lower bound from Theorem 3 immediately implies an exponential lower bound for $D_{ND,RD}^T$, since $D_{ND,RD}^T$ is strictly weaker than $C_{ND,RD}^T$.

6.4 Separation Result for CDCL Solvers with WDLS

Finally, we state an observation of Robert Robere [19] on restarts in the context of the Weak Decision Learning Scheme (WDLS).

Definition 21 (WDLS). Upon deriving a conflict, a CDCL solver with WDLS learns a conflict clause which is the disjunction of the negation of the decision variables on the current assignment trail.

Theorem 22. $C_{ND,ND}^J + \text{WDLS}$ is exponentially weaker than $C_{ND,ND}^{J,R} + \text{WDLS}$.

Proof. The solver model $C_{ND,ND}^J$ with WDLS is only as powerful as $D_{ND,ND}^T$, since each learnt clause will only be used once for propagation after the solver backtracks immediately after learning the conflict clause, and remains satisfied for the rest of the solver run. This is exactly how $D_{ND,ND}^T$ behaves under the same circumstances. On the other hand, WDLS is an asserting learning scheme [17], and hence satisfies the conditions of the main theorem in [18], proving that CDCL with any asserting learning scheme and restarts p-simulates general resolution. Thus, we immediately have $C_{ND,ND}^{J,R}$ with WDLS is exponentially more powerful than the same solver but with no restarts (for unsatisfiable instances). \square

7 Related Work

Previous Work on Theoretical Understanding of Restarts: Buss et al. [8] and Van Gelder [21] proposed two proof systems, namely regWRTI and pool resolution respectively, with the aim of explaining the power of restarts in CDCL solvers. Buss et al. proved that regWRTI is able to capture exactly the power of CDCL solvers with *non-greedy BCP* and without restarts and Van Gelder proved that pool resolution can simulate certain configurations of DPLL solvers with clause learning. As both pool resolution and regWRTI are strictly more powerful than regular resolution, a natural question is whether formulas that exponentially separate regular and general resolution can be used to prove lower bounds for pool resolution and regWRTI, thus transitively proving lower bounds for CDCL solvers without restarts. However, since Bonet et al. [7] and Buss and Kołodziejczyk [9] proved that all such candidates have short proofs in pool resolution and regWRTI, the question of whether CDCL solvers without restarts are as powerful as general resolution still remains open.

Previous Work on Empirical Understanding of Restarts: The first paper to discuss restarts in the context of DPLL SAT solvers was by Gomes and Selman [10]. They proposed an explanation for the power of restarts popularly referred to as “heavy-tailed explanation of restarts”. Their explanation relies on the observation that the runtime of randomized DPLL SAT solvers on satisfiable instances, when invoked with different random seeds, exhibits a heavy-tailed distribution. This means that the probability of the solver exhibiting a long runtime on a given input and random seed is non-negligible. However, because of the heavy-tailed distribution of solver runtimes, it is likely that the solver may run quickly on the given input for a different random seed. This observation was the motivation for the original proposal of the restart heuristic in DPLL SAT solvers by Gomes and Selman [10].

Unfortunately, the heavy-tailed explanation of the power of restarts does not lift to the context of CDCL SAT solvers. The key reason is that, unlike DPLL solvers, CDCL solvers save solver state (e.g., learnt clauses and variable activities) across restart boundaries. Additionally, the efficacy of restarts has been

observed for both deterministic and randomized CDCL solvers, while the heavy-tailed explanation inherently relies on randomness. Hence, newer explanations have been proposed and validated empirically on SAT competition benchmarks. Chief among them is the idea that “restarts compact the assignment trail during its run and hence produce clauses with lower literal block distance (LBD), a key metric of quality of learnt clauses” [14].

Comparison of Our Separation Results with Heavy-Tailed Explanation of Restarts: A cursory glance at some of our separation results might lead one to believe that they are a complexity-theoretical analogue of the heavy-tailed explanation of the power of restarts, since our separation results are over randomized solver models. We argue this is not the case. First, notice that our main results are for drunk CDCL solvers that save solver state (e.g., learnt clauses) across restart boundaries, unlike the randomized DPLL solvers studied by Gomes et al. [10]. Second, we make no assumptions about independence (or lack thereof) of branching decisions across restart boundaries. In point of fact, the variable selection in the CDCL model we use is non-deterministic. Only the value selection is randomized. More precisely, we have arrived at a separation result without relying on the assumptions made by the heavy-tailed distribution explanation, and interestingly we are able to prove that the “solver does get stuck in a bad part of the search space by making bad value selections”. Note that in our model the solver is free to go back to “similar parts of the search space across restart boundaries”. In fact, in our proof for CDCL with restarts, the solver chooses the same variable order across restart boundaries.

8 Conclusions

In this paper, we prove a series of results that establish the power of restarts (or lack thereof) for several models of CDCL and DPLL solvers. We first showed that CDCL solvers with backtracking, non-deterministic dynamic variable selection, randomized dynamic value selection, and restarts are exponentially faster than the same model without restarts for a class of satisfiable instances. Second, we showed CDCL solvers with VSIDS variable selection and phase saving without restarts are exponentially weaker than the same solver with restarts, for a family of unsatisfiable formulas. Finally, we proved four additional smaller separation and equivalence results for various configurations of DPLL and CDCL solvers.

By contrast to previous attempts at a “theoretical understanding the power of restarts” that typically assumed that variable and value selection heuristics in solvers are non-deterministic, we chose to study randomized or real-world models of solvers (e.g., VSIDS branching with phase saving value selection) that enabled us to more effectively isolate the power of restarts. This leads us to believe that the efficacy of restarts becomes apparent only when the solver models considered have weak heuristics (e.g., randomized or real-world deterministic) as opposed to models that assume that all solver heuristics are non-deterministic.

References

1. Alekhnovich, M., Hirsch, E.A., Itsykson, D.: Exponential lower bounds for the running time of DPLL algorithms on satisfiable formulas. *J. Autom. Reasoning* **35**(1–3), 51–72 (2005)
2. Atserias, A., Fichte, J.K., Thurley, M.: Clause-learning algorithms with many restarts and bounded-width resolution. *J. Artif. Intell. Res.* **40**, 353–373 (2011)
3. Beame, P., Kautz, H.A., Sabharwal, A.: Towards understanding and harnessing the potential of clause learning. *J. Artif. Intell. Res.* **22**, 319–351 (2004)
4. Beame, P., Pitassi, T.: Propositional proof complexity: past, present, and future. In: Paun, G., Rozenberg, G., Salomaa, A. (eds.) *Current Trends in Theoretical Computer Science, Entering the 21th Century*, pp. 42–70. World Scientific, Singapore (2001)
5. Ben-Sasson, E., Wigderson, A.: Short proofs are narrow—resolution made simple. *J. ACM* **48**(2), 149–169 (2001)
6. Biere, A., Heule, M., van Maaren, H.: *Handbook of Satisfiability*, vol. 185. IOS press, Amsterdam (2009)
7. Bonet, M.L., Buss, S., Johannsen, J.: Improved separations of regular resolution from clause learning proof systems. *J. Artif. Intell. Res.* **49**, 669–703 (2014)
8. Buss, S.R., Hoffmann, J., Johannsen, J.: Resolution trees with lemmas: resolution refinements that characterize DLL algorithms with clauselearning. *Log. Methods Comput. Sci.* **4**(4) (2008)
9. Buss, S.R., Kołodziejczyk, L.: Small stone in pool. *Log. Methods Comput. Sci.* **10**(2), 16:1–16:22 (2014)
10. Gomes, C.P., Selman, B., Crato, N., Kautz, H.: Heavy-tailed phenomena in satisfiability and constraint satisfaction problems. *J. Autom. Reasoning* **24**(1–2), 67–100 (2000). <https://doi.org/10.1023/A:1006314320276>
11. Hertel, P., Bacchus, F., Pitassi, T., Gelder, A.V.: Clause learning can effectively P-simulate general propositional resolution. In: Fox, D., Gomes, C.P. (eds.) *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI 2008)*, pp. 283–290. AAAI Press (2008)
12. Krajčček, J.: *Proof Complexity, Encyclopedia of Mathematics and Its Applications*, vol. 170. Cambridge University Press, Cambridge (2019)
13. Li, C., Fleming, N., Vinyals, M., Pitassi, T., Ganesh, V.: Towards a complexity-theoretic understanding of restarts in SAT solvers (2020)
14. Liang, J.H., Oh, C., Mathew, M., Thomas, C., Li, C., Ganesh, V.: Machine learning-based restart policy for CDCL SAT solvers. In: Beyersdorff, O., Wintersteiger, C.M. (eds.) *SAT 2018. LNCS*, vol. 10929, pp. 94–110. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-94144-8_6
15. Marques-Silva, J.P., Sakallah, K.A.: GRASP: a search algorithm for propositional satisfiability. *IEEE Trans. Comput.* **48**(5), 506–521 (1999)
16. Moskewicz, M.W., Madigan, C.F., Zhao, Y., Zhang, L., Malik, S.: Chaff: engineering an efficient SAT solver. In: *Proceedings of the 38th annual Design Automation Conference*, pp. 530–535. ACM (2001)
17. Pipatsrisawat, K., Darwiche, A.: A new clause learning scheme for efficient unsatisfiability proofs. In: *AAAI*, pp. 1481–1484 (2008)
18. Pipatsrisawat, K., Darwiche, A.: On the power of clause-learning SAT solvers as resolution engines. *Artif. Intell.* **175**(2), 512–525 (2011)
19. Robere, R.: Personal communication (2018)
20. Urquhart, A.: Hard examples for resolution. *J. ACM* **34**(1), 209–219 (1987)

21. Gelder, A.: Pool resolution and its relation to regular resolution and DPLL with clause learning. In: Sutcliffe, G., Voronkov, A. (eds.) LPAR 2005. LNCS (LNAI), vol. 3835, pp. 580–594. Springer, Heidelberg (2005). https://doi.org/10.1007/11591191_40
22. Vinyals, M.: Hard examples for common variable decision heuristics. In: Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI 2020), February 2020
23. Williams, R., Gomes, C., Selman, B.: On the connections between backdoors, restarts, and heavy-tailedness in combinatorial search. *Structure* **23**, 4 (2003)
24. Williams, R., Gomes, C.P., Selman, B.: Backdoors to typical case complexity. In: Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI 2003), pp. 1173–1178 (2003)