

Limits of CDCL Learning via Merge Resolution

Marc Vinyals Chunxiao Li Noah Fleming Antonina Kolokolova
 Vijay Ganesh

18th April 2023

Abstract

In their seminal work, Atserias et al. and independently Pipatsrisawat and Darwiche in 2009 showed that CDCL solvers can simulate resolution proofs with polynomial overhead. However, previous work does not address the tightness of the simulation, i.e., the question of how large this overhead needs to be. In this paper, we address this question by focusing on an important property of proofs generated by CDCL solvers that employ *standard learning schemes*, namely that the derivation of a learned clause has at least one inference where a literal appears in both premises (aka, a merge literal). Specifically, we show that proofs of this kind can simulate resolution proofs with at most a linear overhead, but there also exist formulas where such overhead is necessary or, more precisely, that there exist formulas with resolution proofs of linear length that require quadratic CDCL proofs.

1 Introduction

Over the last two decades, CDCL SAT solvers have had a dramatic impact on many areas of software engineering [CGP⁺08], security [DVT07, XA05], and AI [BF97]. This is due to their ability to solve very large real-world formulas that contain upwards of millions of variables and clauses [MLM21]. Both theorists and practitioners have expended considerable effort in understanding the CDCL algorithm and the reasons for its unreasonable effectiveness in the context of practical applications. While considerable progress has been made, many questions remain unanswered.

Perhaps the most successful set of tools for understanding the CDCL algorithm come from proof complexity, and a highly influential result is the one that shows that idealized models of CDCL can polynomially simulate the resolution proof system, proved independently by Atserias, Fichte, and Thurley [AFT11], and Pipatsrisawat and Darwiche [PD11], building on initial results by Beame et al. [BKS04] and Hertel et al. [HBPV08]. (See also a recent alternative proof by Beyersdorff and Böhm [BB21].) Such simulation results are very useful because they reassure us that whenever a formula has a short resolution proof then CDCL with the right choice of heuristics can reproduce it.

Recent models make assumptions that are closer to real solvers, but pay for that with a polynomial overhead in the simulation. A series of papers have focused on understanding which of the assumptions are needed for these simulations to hold, often using and/or introducing refinements of resolution along the way. For instance, the question of whether restarts are needed, while still open, has been investigated at length, and the pool resolution [Van05] and RTL [BHJ08] proof systems were devised to capture proofs produced by CDCL solvers that do not restart. The importance of decision heuristics has also been explored recently, with results showing that neither static [MPR20] nor VSIDS-like [Vin20] ordering of variables are enough to simulate resolution in full generality (unless VSIDS scores are periodically erased [LFV⁺20]). In the case of static ordering, the (semi-)ordered resolution proof system [MPR20] was used to reason about such variants of CDCL solvers.

But even if we stay within the idealized model, it is not clear how efficient CDCL is in simulating resolution. The analysis of Pipatsrisawat and Darwiche gives an $O(n^4)$ overhead—that is, if a formula

over n variables has a resolution refutation of length L , then a CDCL proof with no more than $O(n^4L)$ steps exists. Beyersdorff and Böhm [BB21] improved the overhead to $O(n^3)$, but we do not know what the optimal is. Furthermore, to the best of our knowledge, prior to our paper, we did not even know if the overhead can be avoided altogether.

1.1 Learning Schemes in CDCL and Connection with Merges

A common feature of CDCL solvers is the use of 1-empowering learning schemes [PD08, AFT11]: that is, they only learn clauses which enable unit propagations that were not possible before. An example of 1-empowering learning scheme is the popular 1UIP learning scheme [MS99]. To model this behavior we build upon a connection between 1-empowerment, and merges [And68], i.e., resolution steps involving clauses with shared literals.

Nearly every CDCL solver nowadays uses the First Unique Implication Point (1UIP) learning scheme, where conflict analysis starts with a clause falsified by the current state of the solver and sequentially resolves it with clauses responsible for unit propagations leading to the conflict, until the clause becomes *asserting*, i.e., unit immediately upon backjumping.

Descriptions of early implementations of CDCL solvers [MS99, MMZ⁺01] already remark on the importance of learning an asserting clause, since that nudges the solver towards another part of the search space, and consequently early alternative learning schemes explored learning many kinds of asserting clauses. First observe that conflict analysis can be extended to produce other asserting clauses that appear after the 1UIP during conflict analysis such as intermediate UIPs and the last UIP [BS97]. The early solver GRASP can even learn multiple UIP clauses from a single conflict. While there is empirical evidence that it is often best to stop conflict analysis at the 1UIP [ZMMM01], recent work has identified conditions where it is advantageous to continue past it [FB20] (see also the discussion of learning schemes therein).

Ryan [Rya04, §2.5] also observed empirically that clause quality is negatively correlated with the length of the conflict analysis derivation and considered the opposite approach, that is, learning clauses that appear before the 1UIP during conflict analysis in addition to the 1UIP. This approach is claimed to be useful for some empirical benchmarks but, like any scheme that learns multiple clauses, slows down Boolean constraint propagation (BCP) in comparison to a scheme that learns just the 1UIP.

Later works provide a more theoretically oriented approach to understanding the strength of 1UIP and to learning clauses that appear before the 1UIP [DHN07, PD08]. In particular, and highly relevant for our discussion, Pipatsrisawat and Darwiche identified 1-empowerment as a fundamental property of asserting clauses. Furthermore they identified a connection between 1-empowering clauses and merges, and used the simplicity of checking for merges as an approximation for 1-empowerment.

An orthogonal approach is to extend the 1UIP derivation by resolving it with clauses other than those that would usually be used during conflict analysis [ABH⁺08]. A prominent example is clause minimization [SB09], where literals are eliminated from the 1UIP clause by resolving it with the appropriate input clauses, independently of their role in the conflict, so the resultant clause that is actually learned is a shorter and therefore stronger version of the 1UIP.

Furthermore, a relation between merges and unit-resolution completeness has also been observed in the context of knowledge compilation [dV94]. Finally, the amount of merges directly inferable from a formula (i.e., in a single resolution step) has been proposed, under the name of mergeability, as a measure to help explain the hardness of a formula based on both controlled experiments as well as analysis of real-world instances [ZMW⁺18].

To summarize, merges are relevant in the context of CDCL learning schemes for the following reason: all practical CDCL learning schemes either produce a 1-empowering clause or extend one, and since 1-empowering clauses always contain a merge in its derivation, we have that all practical

learning schemes produce a clause that contains a merge in its derivation, which is exactly the property imposed by the proof systems we introduce below.

1.2 Our Contributions

As mentioned earlier, we build upon a connection between 1-empowerment and merges [PD08, AFT11], and introduce a proof system RMA (for “resolution with merge ancestors”) which includes CDCL with an arbitrary 1-empowering learning scheme. The “merge ancestors” in the name of this system comes from the fact that for any 1-empowering clause, at least one step in its resolution derivation must resolve two clauses that share a common literal: a *merge* step in the sense of [And68]. Clause minimization procedures, as long as they are applied on top of 1-empowering clauses, are also modelled by RMA.

We prove that, on the one hand, RMA is able to simulate resolution only with a linear overhead. On the other hand, we show a quadratic separation between resolution and RMA, that is there exist formulas with resolution proofs of linear length that require RMA proofs of quadratic length. That is, we show that CDCL may be polynomially worse than resolution because of the properties of a standard learning scheme, but that the blow-up due to these properties is not more than linear.

We also consider weaker proof systems, all of which contain 1UIP (and do so with finer granularity), but not necessarily other asserting learning schemes. A technical point of interest is that we work with proof systems that are provably not closed under restrictions, which is unusual in proof complexity. This fact forces our proof to exploit syntactic properties of the proof system, as opposed to relying on more convenient semantic properties.

2 Preliminaries

A literal is either a variable $x^1 = x$ or its negation $x^0 = \bar{x}$. A clause is a disjunction of literals, and a CNF formula is a conjunction of clauses. The support of a clause or $\text{vars}(C)$ is the set of variables it contains. A resolution derivation from a formula F is a sequence of clauses $\eta = C_1, \dots, C_L$ such that C_i is either an axiom in F or it is the conclusion of applying the resolution rule

$$\text{Res}(A \vee x, B \vee \bar{x}) = A \vee B$$

on two premises C_j, C_k with $j, k < i$. The variable x that appears with opposite signs in the premises of a resolution inference is called the pivot. If furthermore there is a *literal* common to A and B the resolvent is called a merge. If instead of being the result of a syntactic inference we allow C_i to be any clause semantically implied by C_j and C_k , even if C_j and C_k might not be resolvable, then we say η is a semantic resolution derivation. A derivation is a refutation if its last clause is the empty clause \perp . We denote $\eta[a, b] = \{C_i \in \eta \mid i \in [a, b]\}$.

We assume that every clause in a derivation is annotated with the premises it is obtained from, which allows us to treat the proof as a DAG where vertices are clauses and edges point from premises to conclusions. When this DAG is a tree we call a derivation tree-like, and when it is a centipede (i.e., a maximally unbalanced tree) we call it input.

A derivation is unit if in every inference at least one of the premises is a unit clause consisting of a single literal. Since neither input nor unit resolution are complete proof systems, we write $F \vdash_i C$ (respectively $F \vdash_1 C$) to indicate that there exists an input (resp. unit) resolution derivation of C from F .

A clause C syntactically depends on an axiom A with respect to a derivation η if there is a path from A to C in the DAG representation of η . This does not imply that A is required to derive C , since a different derivation might not use A .

A restriction to variables is a mapping $\rho : X \rightarrow X \cup \{0, 1\}$, successively extended to literals, clauses, formulas, and refutations, simplifying where needed. We write $\rho(x) = *$ as a shorthand for $\rho(x) = x$. It is well-known that if η is a resolution derivation from F and ρ is a restriction, then $\eta \upharpoonright_\rho$ is a semantic resolution derivation from $F \upharpoonright_\rho$.

It is convenient to leave satisfied clauses in place in a derivation that is the result of applying a restriction to another derivation so that we can use the same indices to refer to both derivations. To do that we use the symbol \perp and treat it as a clause that always evaluates to true, is not supported on any set, does not depend on any clause, and cannot be syntactically resolved with any clause.

A semantic derivation can be turned into a syntactic derivation by ignoring unnecessary clauses. Formally, if η is a semantic resolution derivation, we define its syntactic equivalent $s(\eta)$ as the following syntactic resolution derivation. Let $C \in \eta$ and let A and B be the parents of C . If $s(A) \models C$ we set $s(C) = s(A)$, analogously with $s(B)$. Otherwise we set $s(C) = \text{Res}(s(A), s(B))$. It is not hard to see that for each $C_i \in \eta$, $s(C_i) \models C_i$.

2.1 CDCL

We need to define a few concepts from CDCL proofs. An in-depth treatment can be found in the Handbook of Satisfiability [BN21]. Fix a CNF F , also known as clause database. A trail τ is a sequence of tuples $(x_{j_i} = b, C_i)$ where C_i is either a clause in F or the special symbol d representing a decision. We denote by $\alpha_{<i}$ the assignment $\{x_{j_i} = b \mid i' < i\}$, and we denote by $\text{dl}(i) = \text{dl}(i-1) + \llbracket C_i = d \rrbracket$ the decision level at position i , that is the number of decisions up to i . We mark the position of the last decision in a trail by i^* .

A trail is valid if for every position i that is not a decision we have that $C_i \upharpoonright_{\alpha_{<i}} = x_{j_i}^b$ and for every decision i we have that for every clause $C \in F$ such that $C \upharpoonright_{\alpha_{<i}} = x^b$, the literal x^b appears in the trail before i . In particular, for every position $i' < i$ with $\text{dl}(i') < \text{dl}(i)$ we have $C_i \upharpoonright_{\alpha_{<i'}} \neq x_{j_i}^b$.

A clause C is asserting if it is unit at the last decision in the trail, that is $C \upharpoonright_{\alpha_{<i^*}} = x^b$. It is 1-empowering if C is implied by F and can lead to new unit propagations after being added to F , that is if there exists a literal $\ell \in C$ such that for some $A \in \{\perp, \ell\}$, it holds that $F \wedge C \setminus \ell \not\models_1 A$. If a clause is not 1-empowering then we say it is absorbed by F .

Given a clause $D_{|\tau|}$ falsified by a trail τ , the conflict derivation is an input derivation $D_{|\tau|}, \dots, D_{k+1}, D_k$ where $D_{i-1} = \text{Res}(D_i, C_i)$ if $x_{j_i} \in D_i$, and $D_{i-1} = D_i$ otherwise. The first (i.e., with the largest index) asserting clause in the derivation is called the 1UIP. Note that D_{i^*} is always asserting (because D_i is falsified by $\alpha_{<i}$ for $i^* \leq i \leq |\tau|$ and D_{i^*} is not falsified by $\alpha_{<i^*}$), therefore we can assume that the 1UIP always has index at least i^* .

We call a sequence of input derivations *input-structured* if the last clause of each derivation can be used as an axiom in successive derivations. The last clause of each but the last derivation is called a lemma. A CDCL derivation is an input-structured sequence of conflict derivations, where learned clauses are lemmas. This definition is similar to that of Resolution Trees with Input Lemmas [BHJ08], with the difference that the sequence only needs to be ordered, without imposing any further tree-structure on the global proof.

The following Lemmas highlight the practical relevance of merges by relating them to 1UIP, asserting, and 1-empowering clauses.

Lemma 2.1 ([PD08, Proposition 2]). *If a clause is asserting, then it is 1-empowering.*¹

Lemma 2.2 ([AFT11, Lemma 8]). *If $A \vee x$ and $B \vee \bar{x}$ are absorbed but $A \vee B$ is 1-empowering, then $A \vee B$ is a merge. In particular, if a clause is 1-empowering, then it contains a merge in its derivation.*

Lemma 2.3. *The 1UIP clause is a merge.*

¹The original result does not prove 1-consistency, but the proof is analogous.

Proof. Let $D_j = \text{Res}(C_{j+1}, D_{j+1})$ be the 1UIP. On the one hand, since every clause in the trail contains at least two literals at the same decision level it appears in, C_{j+1} contains two literals at the last decision level. On the other hand, any clause that is not in the trail also contains two literals at the last decision level, and in particular $D_{|\tau|}$. Since $|D_{i+1} \setminus D_i| \leq 1$ and D_{j+1} is not asserting, it also contains two literals at the last decision level.

We accounted for 4 literals at the last decision level present in the premises of D_j , of which 2 are not present in the conclusion because they are the pivots. In order for D_j to contain only one literal at the last decision level, the remaining two literals must be equal. \square

3 Proof Systems

We define our proof systems in terms of the *input-structured* framework. Every resolution proof can be thought of as being input-structured if we consider it as a sequence of unit-length input resolutions and every clause as a lemma; it is when we impose restrictions on which clauses are permitted as lemmas that we obtain different proof systems. The diagram in Figure 1 can help keeping track of the proof systems.




Figure 1: Relations between proof systems. A solid arrow $A \rightarrow B$ indicates that A simulates B with no overhead. A dashed arrow $A \dashrightarrow B$ indicates that A simulates B with no overhead, but B requires linear overhead to simulate A . Statements proving separations are referenced.

Andrews’ definition of merge resolution [And68] considers tree-like proofs with the additional restriction that in every inference at least one premise is an axiom or a merge. He also observes that such derivations can be made input-structured.

Observation 3.1 ([And68]). *A tree-like merge resolution derivation can be decomposed into an input-structured sequence where all the lemmas are merges.*

This observation is key when working with such derivations, as is apparent in Sections 4 and A, to the point that we use as an alternative way to define merge resolution.

Andrews’ main result is that the merge restriction does not affect tree-like resolution.

Lemma 3.2 ([And68, Lemma 5]). *If there is a tree-like resolution derivation of C of length L where at most the root is a merge, then there is an input resolution derivation of some $C' \subseteq C$ of length at most L .*

Theorem 3.3 ([And68, Theorem 1]). *If there is a tree-like resolution derivation of C of length L , then there is a tree-like merge resolution derivation of some $C' \subseteq C$ of length at most L .*

If we lift the tree-like restriction from the input-structured view of merge resolution proofs we obtain a proof system between tree- and DAG-like resolution where clauses can be reused (i.e., have outdegree larger than 1) if and only if they are merges or, in other words, lemmas in the input-structured decomposition. We call this proof system Resolution with Merge Lemmas and refer to it with the acronym RML.

Definition 3.4. A RML derivation is an input-structured sequence of unit resolution derivations where all lemmas are merges.

CDCL refutations produced by solvers that use the 1UIP learning scheme are in RML form, as a consequence of Lemma 2.3. We can also generalize RML to allow reusing clauses that contain a merge anywhere in their derivation. We call this proof system Resolution with Merge Ancestors, or RMA for short.

Definition 3.5. A RMA derivation is an input-structured sequence of unit resolution derivations where all derivations but the last contain a merge.

Note that by Lemma 3.2 it does not matter if we require the sequence of derivations of an RMA derivation to be input derivations or if we allow general trees. In fact, our lower bound results hold for a more general proof system where we only ask that every clause with outdegree larger than 1 has an ancestor that is a merge. Such proof system does not have a simple input structure, but can rather be thought of as a sequence of tree-like resolution derivations whose roots are merges, followed by a standard resolution derivation using the roots of the previous derivations as axioms.

To make the connection back to CDCL, we can define a proof system called Resolution with Empowering Lemmas that captures CDCL refutations produced by solvers that use any asserting learning scheme or 1-empowering learning scheme.

Definition 3.6. Let C_1, \dots, C_{L-1} be the lemmas of an input-structured sequence of unit derivations. The sequence is a Resolution with Empowering Lemmas (REL) derivation of a formula F if C_i is 1-empowering with respect to $F \cup \{C_j : j < i\}$ for all $i \in [1, L - 1]$.

It follows from Lemmas 2.1 and 2.2 that such refutations are in RMA form.

Observation 3.7. *A REL derivation is a RMA derivation.*

It might seem more natural to work with the REL proof system rather than its merge-based counterparts, since REL is defined exactly through the 1-empowering property. However, while the merge property is easy to check because it is local to the derivation at hand, we can only determine if a clause is 1-empowering by looking at the full history of the derivation, in particular what the previous lemmas are. This makes REL too cumbersome to analyse. Furthermore, CDCL refutations produced apply a clause minimization scheme on top of an asserting clause might not be in REL form, but they are still in RMA form.

A further property of input derivations produced by a CDCL solver is that once a variable is resolved, it does not appear later in the derivation.

Definition 3.8. A resolution derivation η is strongly regular if for every resolution step i , the pivot variable x_i is not part of the support of any clause $C_i \in \eta[i, L]$. A sequence of derivations is locally regular if every derivation in the sequence is strongly regular. A LRML derivation (resp. LRMA) is a locally regular RML derivation (resp. RMA).

Finally we can consider derivations that have empowering, merge lemmas and are locally regular. These still include 1UIP proofs.

Definition 3.9. A LREML derivation is a derivation that is both LRML and REL.

It follows from the simulation of resolution by CDCL [PD11, AFT11] that all (DAG-like) proof systems we defined polynomially simulate standard resolution. In Section 4 we make this simulation more precise and prove that the simulation overhead can be made linear, and in Section 5 that the simulation is optimal because there exist formulas that have resolution refutations of linear length but require RMA refutations of quadratic length.

4 Simulation

As an auxiliary tool to simulate resolution in RML we define the input-resolution closure of a set D , denoted $\text{Cl}_i(D) = \{C \mid \exists C' \subseteq C, D \vdash_i C'\}$, as the set of clauses derivable from D via input resolution plus weakening. It is well-known that, since input resolution derivations can be assumed to be strongly regular without loss of generality, we can also assume them to be at most linear in the number of variables.

Observation 4.1. *If D is a CNF formula over n variables and $C \in \text{Cl}_i(D)$ then there is a strongly regular input resolution derivation of some $C' \subseteq C$ from D of length at most n .*

Combining Theorem 3.3 with the idea that in order to simulate a resolution derivation we do not need to generate each clause, but only do enough work so that in the following steps we can pretend that we had derived it [PD11, AFT11], we can prove that merge resolution simulates resolution with at most a multiplicative linear overhead in the number of variables.

Theorem 4.2. *If F is a CNF formula over n variables that has a resolution refutation of length L then it has a RML refutation of length $O(nL)$.*

Proof. Let $\pi = (C_1, \dots, C_L)$ be a resolution refutation. We construct a sequence of sets D_0, \dots, D_L with the following properties.

1. $D_t \setminus F$ is the set of lemmas in a RML derivation of length at most $(2n + 1)t$.
2. $\pi[1, t] \subseteq \text{Cl}_i(D_t)$.

This is enough to prove the theorem: since $\perp \in D_t$ we can obtain \perp from D_t in length n , so the total length of the refutation is $(2n + 1)L + n$.

We build the sets by induction, starting with $D_0 = F$. Assume we have built D_t and let $C = C_{t+1} = \text{Res}(A, B)$ with $A, B \in \pi[1, t]$. If $C \in \text{Cl}_i(D_t)$ we set $D_{t+1} = D_t$ and we are done. Otherwise, by induction we have $A, B \in \text{Cl}_i(D_t)$, therefore by Observation 4.1 there are input resolution derivations of $A' \subseteq A$ and $B' \subseteq B$ of length at most n . Since neither $A' \models C$ nor $B' \models C$, A' and B' can be resolved and therefore there is a tree-like derivation η of $C' \subseteq C$ from D_t of length at most $2n + 1$. By Theorem 3.3 there is a tree-like merge resolution derivation η' of $C'' \subseteq C$ from D_t of length at most $2n + 1$. By Observation 3.1 the derivation η' can be decomposed into a sequence of input derivations of total length at most $2n + 1$. Let E be the lemmas in that sequence and set $D_{t+1} = D_t \cup E$. We have that $C \in \text{Cl}_i(F \cup E) \subseteq \text{Cl}_i(D_{t+1})$, and that we can obtain E from D_t in at most $2n + 1$ steps. Thus D_{t+1} has all the required properties. \square

We can be a bit more precise with the description of the simulation if we look at the structure of η before applying Theorem 3.3. Let A_M and B_M be the last merges in the input derivation of A' and B' respectively, and let $E = \{A_M, B_M\}$.

Now consider the fragment of the input derivation of A' from A_M to A' , analogously with B' . We have a tree-like derivation of C' where at most the root is a merge, therefore we can apply Lemma 3.2 directly instead of Theorem 3.3 and obtain an input resolution derivation of $C'' \subseteq C$ from $E \cup F$.

If we also make sure that the input derivations of A' and B' are strongly regular, we have that LRML can also simulate resolution with the same $O(n)$ overhead as RML.

An analogous result can be obtained for LREML from the following lemma.

Lemma 4.3 ([PD11]). *If F absorbs $A \vee x$ and $B \vee \bar{x}$, then $F \vdash_i C' \subseteq A \vee B$.*

Corollary 4.4. *If F is a CNF formula over n variables that has a resolution refutation of length L then it has a LREML refutation of length $O(nL)$.*

Proof. The proof follows the general structure of Theorem 4.2, except that we use a sequence of steps D_t^j in order to construct D_t . Our induction hypothesis is that D_t^j can be derived from D_t in p inference steps in LREML, and that A' and B' can be derived from D_t^j in q steps, with $p + q \leq 2n$.

The base case $D_t^0 = D_t$ is trivial.

For the inductive case, assume that the input derivations leading to A' and B' are strongly regular without loss of generality. By Lemma 4.3 either A' or B' is 1-empowering, say A' . Let C be the first 1-empowering clause in the derivation of A' . By Lemma 2.2 C is a merge, therefore we can take $D_t^{j+1} = D_t^j \cup \{C\}$. \square

5 Separation

We prove the following separation between standard resolution and RMA.

Theorem 5.1. *There exists a family of formulas F_n over $O(n \log n)$ variables and $O(n \log n)$ clauses that have resolution refutations of length $O(n \log n)$ but every RMA refutation requires length $\Omega(n^2 \log n)$.*

5.1 Formula

Let ℓ, m, n be positive integers. We have variables x_i for $i \in [m\ell - 1]$ and $w_{j,k}$ for $j \in [\ell]$ and $k \in [n]$. For convenience we define $x_0 = 1$ and $x_{m\ell} = 0$, which are not variables. Let $X = \{x_i \mid i \in [m\ell - 1]\}$, $W_j = \{w_{j,k} \mid k \in [n]\}$ and $W = \bigcup_{j \in [\ell]} W_j$. For each $j \in [\ell]$ we build the following gadget:

$$w_{j,k} = w_{j,k+1} \quad \text{for } k \in [n-1] \quad (1)$$

Each equality is expanded into the two clauses $B_{j,k,1} = w_{j,k} \vee \overline{w_{j,k+1}}$ and $B_{j,k,0} = \overline{w_{j,k}} \vee w_{j,k+1}$, and we collectively call them $\mathcal{W} = \{B_{j,k,b} \mid j \in [\ell], k \in [n-1], b \in \{0, 1\}\}$. Observe that the j -th gadget implies $w_{j,1} = w_{j,n}$. Additionally we build the following gadget:

$$(w_{1,1} = w_{1,n}) \rightarrow x_1 \quad (2)$$

$$(w_{i,1} = w_{i,n}) \rightarrow (x_{i-1} \rightarrow x_i) \quad \text{for } i \in [2, m\ell - 1] \quad (3)$$

$$(w_{\ell,1} = w_{\ell,n}) \rightarrow \overline{x_{m\ell-1}} \quad (4)$$

where $\hat{i} \in [\ell]$ denotes the canonical form of $i \pmod{\ell}$. Each constraint is expanded into the two clauses $A_{i,1} = w_{\hat{i},1} \vee w_{\hat{i},n} \vee \overline{x_{i-1}} \vee x_i$ and $A_{i,0} = \overline{w_{\hat{i},1}} \vee \overline{w_{\hat{i},n}} \vee \overline{x_{i-1}} \vee x_i$, and we collectively call them $\mathcal{X} = \{A_{i,b} \mid i \in [m\ell], b \in \{0, 1\}\}$. The resulting formula is called $F_{\ell,m,n}$.

5.2 Upper Bound

It is not hard to see that there is a resolution refutation of $F_{\ell,m,n}$ of length $O(\ell \cdot (m+n))$. Indeed, we first derive the two clauses representing $w_{j,1} = w_{j,n}$ for each $j \in [\ell]$, which requires $O(n\ell)$ steps:

$$\frac{\frac{w_{j,1} \vee \overline{w_{j,2}} \quad w_{j,2} \vee \overline{w_{j,3}}}{w_{j,1} \vee \overline{w_{j,3}}}}{\vdots} \frac{\frac{w_{j,1} \vee \overline{w_{j,n-1}} \quad w_{j,n-1} \vee \overline{w_{j,n}}}{w_{j,1} \vee \overline{w_{j,n}}}}{\quad} \quad (5)$$

Then we resolve each of the \mathcal{X} axioms with one of these clauses, appropriately chosen so that we obtain pairs of clauses of the form $w_i^b \vee \overline{x_{i-1}} \vee x_i$ for $i \in [m\ell]$, and resolve each pair to obtain the chain of implications $x_1, \dots, x_i \rightarrow x_{i+1}, \dots, \overline{x_{n\ell-1}}$ in $O(m\ell)$ steps.

$$\frac{\frac{w_{i,1} \vee \overline{w_{i,n}} \quad w_{i,1} \vee w_{i,n} \vee \overline{x_{i-1}} \vee x_i}{w_{i,1} \vee \overline{x_{i-1}} \vee x_i} \quad \frac{\overline{w_{i,1}} \vee w_{i,n} \quad \overline{w_{i,1}} \vee \overline{w_{i,n}} \vee \overline{x_{i-1}} \vee x_i}{\overline{w_{i,1}} \vee \overline{x_{i-1}} \vee x_i}}{\overline{x_{i-1}} \vee x_i} \quad (6)$$

Since we have derived a chain of implications $x_1, x_1 \rightarrow x_2, \dots, x_{m\ell-1} \rightarrow x_{m\ell-1}, \overline{x_{m\ell-1}}$ we can complete the refutation in $O(m\ell)$ more steps. Let us record our discussion.

Lemma 5.2. $F_{\ell,m,n}$ has a resolution refutation of length $O(\ell \cdot (m+n))$.

Before we prove the lower bound let us discuss informally what are the natural ways to refute this formula in RML, so that we understand which behaviours we need to rule out.

If we try to reproduce the previous resolution refutation, since we cannot reuse the clauses representing $w_{j,1} = w_{j,n}$ because they are not merges, we have to rederive them each time we need them, which means that it takes $O(mn\ell)$ steps to derive the chain of implications $x_1, \dots, x_i \rightarrow x_{i+1}, \dots, \overline{x_{n\ell-1}}$. We call this approach *refutation 1*. This refutation has merges (over $w_{i,1}, x_{i-1}$, and x_i) when we produce $w_{i,1}^b \vee \overline{x_{i-1}} \vee x_i$, and (over x_{i-1} and x_i) when we produce $\overline{x_{i-1}} \vee x_i$, but since we never reuse these clauses the refutation is in fact tree-like.

An alternative approach, which we call *refutation 2*, is to start working with the \mathcal{X} axioms instead. In this proof we clump together all of the repeated constraints of the form $w_{j,1} \neq w_{j,n}$ for every $j \in [\ell]$, and then resolve them out in one go. In other words, we first derive the sequence of constraints

$$D_i = \left(\bigvee_{i \in [\min(i,\ell)]} w_{i,1} \neq w_{i,n} \right) \vee x_i \quad \text{for } i \in [m\ell], \quad (7)$$

where D_i can be obtained from D_{i-1} and the pair of \mathcal{X} axioms $A_{i,b}$, then resolve away the inequalities from $D_{m\ell} = \bigvee_{j \in [\ell]} w_{j,1} \neq w_{j,n}$ using the \mathcal{W} axioms. However, representing any of the constraints D_i for $i \geq \ell$ requires 2^ℓ clauses, which is significantly larger than $m\ell$ and even superpolynomial for large enough ℓ , so this refutation is not efficient either. Note that this refutation has merges (over W variables) each time that we derive D_i with $i \geq \ell$.

A third and somewhat contrived way to build a refutation is to derive the pair of clauses representing $w_{j,1} = w_{j,n}$ using a derivation whose last step is a merge, so that they can be reused. Each of these clauses can be derived individually in $O(mn\ell)$ steps, for a total of $O(mn\ell^2)$ steps, by slightly adapting *refutation 1*, substituting each derivation of $x_i \rightarrow x_{i+1}$ by a derivation of $w_{j,1} \vee \overline{w_{j,n}} \vee \overline{x_i} \vee x_{i+1}$ whenever $i \equiv j \pmod{\ell}$ so that at the end we obtain $w_{j,1} \vee \overline{w_{j,n}}$ instead of the

empty clause. Such a substitution clause can be obtained, e.g., by resolving $w_{j,1} \vee w_{j,2} \vee \overline{x_i} \vee x_{i+1}$ with $\overline{w_{j,2}} \vee \overline{w_{j,n}} \vee \overline{x_i} \vee x_{i+1}$ as follows

$$\begin{array}{c}
\frac{w_{j,2} \vee \overline{w_{j,3}} \quad w_{j,3} \vee \overline{w_{j,4}}}{w_{j,2} \vee \overline{w_{j,4}}} \\
\vdots \\
\frac{w_{j,2} \vee \overline{w_{j,n-1}} \quad w_{j,n-1} \vee \overline{w_{j,n}}}{w_{j,2} \vee \overline{w_{j,n}}} \\
\frac{\frac{w_{j,2} \vee \overline{w_{j,n}} \quad w_{i,1} \vee w_{i,n} \vee \overline{x_{i-1}} \vee x_i}{w_{i,1} \vee w_{i,2} \vee \overline{x_{i-1}} \vee x_i} \quad \frac{w_{i,1} \vee \overline{w_{i,2}} \quad \overline{w_{i,1}} \vee \overline{w_{i,n}} \vee \overline{x_{i-1}} \vee x_i}{\overline{w_{i,2}} \vee \overline{w_{i,n}} \vee \overline{x_{i-1}} \vee x_i}}{w_{i,1} \vee \overline{w_{i,n}} \vee \overline{x_{i-1}} \vee x_i}
\end{array} \tag{8}$$

After deriving $w_{j,1} = w_{j,n}$ as merges we follow the next steps of refutation 1 and complete the refutation in $O(m\ell)$ steps. We call this *refutation 3*.

Observe that the minimum length of deriving the clauses representing $w_{j,1} = w_{j,n}$ is only $O(n)$, even in RML, so if we only used the information that refutation 3 contains these clauses we would only be able to bound its length by $\Omega(\ell \cdot (m + n))$. Therefore when we compute the hardness of deriving a clause we need to take into account not only its semantics but how it was obtained syntactically.

5.3 Lower Bound

Before we begin proving our lower bound in earnest we make two useful observations.

Lemma 5.3. *Let η be a resolution derivation that only depends on the \mathcal{W} axioms. Then η does not contain any merges, and all clauses are supported on W .*

Proof. We prove by induction that every clause in η is of the form $w_{j,k} \vee \overline{w_{j,k'}}$ with $k \neq k'$. This is true for the axioms. By induction hypothesis, a generic resolution step over $w_{j,k}$ is of the form

$$\frac{w_{j,k} \vee \overline{w_{j,k'}} \quad \overline{w_{j,k}} \vee w_{j,k''}}{w_{j,k''} \vee \overline{w_{j,k'}}} \tag{9}$$

and in particular is not a merge. □

Lemma 5.4. *Let η be a resolution derivation of a clause C supported on W variables that uses an \mathcal{X} axiom. Then η uses at least one $A_{i,b}$ axiom for each $i \in [m\ell]$.*

Proof. We prove the contrapositive and assume that there is an axiom $A_{i,b}$ that is used, and either both $A_{i+1,0}$ and $A_{i+1,1}$ are not used, or both $A_{i-1,0}$ and $A_{i-1,1}$ are not. In the first case the literal x_i appears in every clause in the path from $A_{i,b}$ to C , contradicting that C is supported on W variables. Analogously with literal $\overline{x_{i-1}}$ in the second case. □

Our first step towards proving the lower bound is to rule out that refutations like refutation 2 can be small, and to do that we show that wide clauses allow for very little progress. This is a common theme in proof complexity, and the standard tool is to apply a random restriction to a short refutation in order to obtain a narrow refutation. However, RMA is not closed under restrictions, as we prove later in Corollary 5.12, and because of this we need to argue separately about which merges are preserved.

Let us define the class of restrictions that we use and which need to respect the structure of the formula. A restriction is an autarky [MS85] with respect to a set of clauses D if it satisfies every clause that it touches; in other words for every clause $C \in D$ either $C \upharpoonright_\rho = 1$ or $C \upharpoonright_\rho = C$. A restriction is k -respecting if it is an autarky with respect to \mathcal{W} axioms, we have $F_{\ell,m,n} \upharpoonright_\rho \cong F_{k,m,n}$ up to variable

renaming, and every X variable is mapped to an X variable. Our definition of a narrow clause is also tailored to the formula at hand, and counts the number of different W -blocks that a clause C mentions. Formally $\mu(C) = |\{j \in [\ell] \mid \exists x_{j,k} \in \text{vars}(C)\}|$.

Lemma 5.5. *Let π be a resolution refutation of $F_{\ell,m,n}$ of length $L = o((4/3)^{\ell/8})$. There exists an $\ell/4$ -respecting restriction ρ such that every clause in $\pi \upharpoonright_\rho$ has $\mu(C) \leq \ell/8$.*

Proof. We use the probabilistic method. Consider the following distribution \mathcal{J} over $\{0, 1, *\}^\ell$: each coordinate is chosen independently with $\Pr[J_i = 0] = \Pr[J_i = 1] = 1/4$, $\Pr[J_i = *] = 1/2$. Given a random variable $J \sim \mathcal{J}$ sampled according to this distribution, we derive a random restriction ρ as follows: $\rho(w_{j,i}) = J_j$, $\rho(x_i) = *$ if $J_i = *$, and $\rho(x_i) = \rho(x_{i-1})$ otherwise (where $\rho(x_0) = 1$).

Observe that $F_{\ell,m,n} \upharpoonright_\rho \cong F_{|J^{-1}(*)|,m,n}$ up to variable renaming, and by a Chernoff bound we have $\Pr[|J^{-1}(*)| < \ell/4] \leq e^{-\ell/16}$.

We also have, for every clause $C \in \pi$ with $\mu(C) > \ell/8$, that

$$\Pr[C \upharpoonright_\rho \neq 1] \leq (3/4)^{\mu(C)} \leq (3/4)^{\ell/8} . \quad (10)$$

Therefore by a union bound the probability that $|J^{-1}(*)| < \ell/4$ or that any clause has $\mu(C \upharpoonright_\rho) > \ell/8$ is bounded away from 1 and we conclude that there exists a restriction ρ that satisfies the conclusion of the lemma. \square

Note that $s(\pi \upharpoonright_\rho)$ is a resolution refutation of $F_{n,\ell} \upharpoonright_\rho$, but not necessarily a RMA refutation, therefore we lose control over which clauses may be reused². Nevertheless, we can identify a fragment of $s(\pi \upharpoonright_\rho)$ where we still have enough information.

Lemma 5.6. *There exists an integer t such that $\psi = s(\pi[1, t] \upharpoonright_\rho)$ is a resolution derivation of a clause supported on W variables that depends on an \mathcal{X} axiom and where no clause supported on W variables is reused.*

Proof. Let $C_t \in \pi$ be the first clause that depends on an \mathcal{X} axiom and such that $D_t = s(C_t \upharpoonright_\rho)$ is supported on W , which exists because \perp is one such clause.

By definition of t , we have that every ancestor $D_k \in \psi$ of D_t that is supported on W variables corresponds to a clause C_k in π that only depends on \mathcal{W} axioms, hence by Lemma 5.3 C_k is not a merge. By definition of RMA C_k is not reused, and by construction of $s(\cdot)$ neither is D_k .

It remains to prove that D_t depends on an \mathcal{X} axiom. Since C_t depends on an \mathcal{X} axiom, at least one of its predecessors C_p and C_q also does, say C_p . By definition of t , $D_p = s(C_p \upharpoonright_\rho)$ is not supported on W , and hence by Lemma 5.3 either D_p depends on an \mathcal{X} axiom or $D_p = 1$. Analogously, if C_q also depends on an \mathcal{X} axiom then so does $D_q = s(C_q \upharpoonright_\rho)$ (or it is 1) and we are done. Otherwise C_q is of the form $w_{j,k} \vee \overline{w_{j,k'}}$ and is either satisfied by ρ or left untouched. In both cases we have that $D_q \not\equiv C_t \upharpoonright_\rho$ (trivially in the first case and because D_q contains the pivot while C_t does not in the second), hence D_t depends on D_p . \square

Note that C_t may be semantically implied by the \mathcal{W} axioms, and have a short derivation as in refutation 3, therefore we are forced to use syntactic arguments to argue that deriving C_t using an \mathcal{X} axiom takes many resolution steps.

The next step is to break ψ into m (possibly intersecting) parts, each corresponding roughly to the part of ψ that uses \mathcal{X} axioms with variables in an interval of length ℓ (by Lemma 5.4 we can assume that ψ contains axioms from every interval). To do this we use the following family of restrictions defined for $i \in [n]$:

$$\sigma_i(x_{i'}) = \begin{cases} 1 & \text{if } i' \leq i\ell \\ * & \text{if } i\ell < i' \leq (i+1)\ell \\ 0 & \text{if } (i+1)\ell < i' \end{cases} \quad \sigma_i(w_{i',j}) = * \quad (11)$$

²Recall that $s(\pi)$ is the syntactic equivalent of π .

Let $X_i = X \cap \sigma_i^{-1}(\ast)$ and note that $F_{\ell,m,n} \upharpoonright_{\sigma_i} \cong F_{\ell,1,n}$.

Clauses in ψ with many X variables could be tricky to classify, but intuitively it should be enough to look at the smallest positive literal and the largest negative literal, since these are the hardest to eliminate. Therefore we define $r(C)$ to be the following operation on a clause: literals over W variables are left untouched, all positive X literals but the smallest are removed, and all negative X literals but the largest are removed. Formally,

$$r\left(\bigvee_{i \in A} x_i \vee \bigvee_{i \in B} \overline{x_i} \vee \bigvee_{(i,j) \in C} w_{i,j}^{b_{i,j}}\right) = x_{\min A} \vee \overline{x_{\max B}} \vee \bigvee_{(i,j) \in C} w_{i,j}^{b_{i,j}} \quad (12)$$

where $x_{\min A}$ (resp. $\overline{x_{\max B}}$) is omitted if A (resp. B) is empty.

We need the following property of $r(C)$.

Lemma 5.7. *If $C \upharpoonright_{\sigma_i} \neq 1$ and $\text{vars}(r(C)) \cap X_i = \emptyset$ then $C \upharpoonright_{\sigma_i}$ is supported over W variables.*

Proof. The hypothesis that $\text{vars}(r(C)) \cap X_i = \emptyset$ implies that the smallest positive X literal in C is either not larger than $i\ell$ or larger than $(i+1)\ell$, but the hypothesis that $C \upharpoonright_{\sigma_i} \neq 1$ rules out the first case. Therefore all positive X literals are falsified by σ_i . Analogously the largest negative X literal is not larger than $i\ell$ and all negative X literals are also falsified. \square

We define each part ψ_i to consist of all clauses $C \in \sigma$ such that C is

1. an \mathcal{X} axiom not satisfied by σ_i ; or
2. the conclusion of an inference with pivot in X_i ; or
3. the conclusion of an inference with pivot in W that depends on an \mathcal{X} axiom if $r(C)$ contains a variable in X_i ; or
4. the conclusion of an inference with pivot in W that does not depend on \mathcal{X} axioms if the *only* immediate successor of C is in ψ_i .

This is the point in the proof where we use crucially that the original derivation is in RMA form: because clauses that do not depend on \mathcal{X} axioms are not merges, they have only one successor and the definition is well-formed.

Ideally we would like to argue that parts ψ_i are pairwise disjoint. This is not quite true, but nevertheless they do not overlap too much.

Lemma 5.8. *Let ψ and $\{\psi_i \mid i \in [\ell]\}$ be as discussed above. Then $2|\psi| \geq \sum_i |\psi_i|$.*

Proof. Axioms may appear in at most two different ψ_i , and clauses obtained after resolving with an X pivot in only one. The only other clauses that depend on an \mathcal{X} axiom and may appear in different ψ_i are obtained after resolving with a W pivot, but since $r(C)$ only contains two X variables, such clause only may appear in two different ψ_i . Finally, clauses that do not depend on an \mathcal{X} axiom appear in the same ψ_i as one clause of the previous types, and therefore at most two different parts. \square

To conclude the proof we need to argue that each ψ_i is large. The intuitive reason is that ψ_i must use one \mathcal{X} axiom for each $j \in [(i\ell, (i+1)\ell]$, which introduces a pair of W variables from each W_j block, but since no clause contains more than $\ell/8$ such variables, we need to use enough \mathcal{W} axioms to remove the aforementioned W variables. Formally the claim follows from these two lemmas.

Lemma 5.9. *For each $i \in [\ell]$ there exists an integer t_i such that $s(\psi_i[1, t_i] \upharpoonright_{\sigma_i})$ is a resolution derivation of a clause supported on W variables that depends on an \mathcal{X} axiom.*

Proof. Let C_{t_i} be the first clause in ψ_i that depends on an \mathcal{X} axiom and such that $C_{t_i} \upharpoonright_{\sigma_i}$ is supported on W variables. We prove that t_i is well-defined, that $\psi_i[1, t_i] \upharpoonright_{\sigma_i}$ is a valid semantic resolution derivation, and that $D_{t_i} = s(\psi_i \upharpoonright_{\sigma_i})$ depends on an \mathcal{X} axiom.

Our induction hypothesis is that for $k \leq t_i$ (or any k if t_i does not exist), if the clause $C_k \in \psi$ depends on an \mathcal{X} axiom and is not satisfied by σ_i , then there exists a clause $C_{k'} \in \psi_i$ with $k' \leq k$ that implies C_k modulo σ_i , that is $C_{k'} \upharpoonright_{\sigma_i} \models C_k \upharpoonright_{\sigma_i}$, and depends on an \mathcal{X} axiom (over ψ).

If the induction hypothesis holds then t_i is well-defined: since C_t is not satisfied by σ_i and depends on an \mathcal{X} axiom there exists a clause $C_{t'} \in \psi_i$ that depends on an \mathcal{X} axiom and such that $C_{t'} \upharpoonright_{\sigma_i} \models C_t \upharpoonright_{\sigma_i} = C_t$, which is supported on W variables.

The base case is when C_k is a non-satisfied \mathcal{X} axiom, where we can take $C_{k'} = C_k$. For the inductive case let C_p and C_q be the premises of C_k in ψ . If exactly one of the premises, say C_p , is non-satisfied and, furthermore, depends on an \mathcal{X} axiom, then by the induction hypothesis we can take $C_{k'} = C_p$. Otherwise we need to consider a few subcases. If the pivot is an X variable then both premises depend on an \mathcal{X} axiom (by Lemma 5.3), hence neither premise is satisfied. It follows that the pivot is unassigned by σ_i , and therefore we can take $C_{k'} = C_k$.

If the pivot is a W variable then, because σ_i only assigns X variables, neither premise is satisfied. We have two subcases: if exactly one premise depends on an \mathcal{X} axiom, say C_p , then $C_{p'}$ is present in ψ_i , and by construction of ψ_i the other premise C_q is present in ψ_i if and only if the conclusion C_k is. If both premises depend on an \mathcal{X} axiom then both $C_{p'}$ and $C_{q'}$ are present in ψ_i .

Therefore in the two latter subcases it is enough to prove that $C_k \in \psi_i$, since then we can take $C_{k'} = C_k$ and we have that $C_k \upharpoonright_{\sigma_i}$ follows from a valid semantic resolution step. Indeed by Lemma 5.7 $C_k \upharpoonright_{\sigma_i}$ is a clause supported on W variables, which by definition of C_{t_i} implies that $k = t_i$. However, since the pivot is a W variable, $C_{p'} \upharpoonright_{\sigma_i}$ is also supported on W variables and, together with the fact that $C_{p'}$ depends on an \mathcal{X} axiom, this contradicts that C_{t_i} is the first such clause.

This finishes the first induction argument and proves that $\psi[1, t_i] \upharpoonright_{\sigma_i}$ is a valid semantic derivation; it remains to prove that D_{t_i} depends on an X axiom over $s(\psi_i \upharpoonright_{\sigma_i})$. We prove by a second induction argument that for every clause $D_k \in s(\psi_i[1, t_i] \upharpoonright_{\sigma_i})$, if C_k depends on an \mathcal{X} axiom then so does D_k . The base case, when D_k is an axiom, holds.

For the inductive case fix $C_k, E_k = C_k \upharpoonright_{\sigma_i}$, and $D_k = s(E_k)$, and let $E_p = C_p \upharpoonright_{\sigma_i}$ and $E_q = C_q \upharpoonright_{\sigma_i}$ be the premises of E_k in $\psi_i \upharpoonright_{\sigma_i}$. When both C_p and C_q depend on an X axiom, then by hypothesis so do D_p and D_q and we are done. We only need to argue the case when one premise C_p depends on an X axiom and the other premise C_q does not. In that case, because σ_i only affects X variables, all the axioms used in the derivation of C_q are left untouched by σ_i , therefore we have that $s(\sigma_i(C_q)) = C_q$, which contains the pivot used to derive C_k and therefore does not imply $s(\sigma_i(C_k))$. By construction of $s(\cdot)$, $s(\sigma_i(C_k))$ depends on $s(\sigma_i(C_p))$. \square

Lemma 5.10. *Let η be a resolution derivation from $F_{\ell,1,n}$ of a clause C supported on W variables that depends on an \mathcal{X} axiom. Then $|\eta| \geq (n-2)(\ell - \mu(C))/2$.*

Proof. By Lemma 5.4 we can assume that η uses at least one $A_{j,b}$ axiom for each $j \in [\ell]$.

Let $J = \{j \in [\ell] \mid \exists w_{j,k} \in \text{vars}(C)\}$ be the set of W blocks mentioned by C . We show that for each $j \in \bar{J} = [\ell] \setminus J$ at least $(n-2)/2$ axioms over variables in W_j appear in η , which makes for at least $(n-2)|\bar{J}|/2 = (n-2)(\ell - \mu(C))/2$ axioms.

Fix $j \in \bar{J}$ and assume for the sake of contradiction that less than $(n-2)/2$ axioms over variables in W_j appear in η . Then there exists $k \in [2, n-1]$ such that variable $w_{j,k}$ does not appear in η . Rename variables as follows: $w_{j,k'} \mapsto y_{k'}$ for $k' < k$, and $w_{j,k'} \mapsto \overline{y_{k'-n}}$ for $k' > k$. Then we can prove by induction, analogously to the proof of Lemma 5.3, that every clause derived from axiom $A_{j,b}$ is of the form $y_{k'} \vee \overline{y_{k'-n}} \vee D$ where D are literals supported outside W_j . Since that includes C , it contradicts our assumption that $j \notin J$. \square

To conclude the proof of Theorem 5.1 we simply need to put the pieces together.

Proof of Theorem 5.1. We take as the formula family $F_{\ell=48\log n, n, n}$, for which a resolution refutation of length $O(n \log n)$ exists by Lemma 5.2.

To prove a lower bound we and assume that a RMA refutation π of length $L \leq n^3 = 2^{16\ell} = o((4/3)^{8\ell})$ exists; otherwise the lower bound trivially holds. We apply the restriction given by Lemma 5.5 to π and we use Lemma 5.6 to obtain a resolution derivation ψ of a clause supported on W variables that uses an \mathcal{X} axiom. We then break ψ into m parts ψ_i , each of size at least $n\ell/16$ as follows from Lemmas 5.9 and 5.10. Finally by Lemma 5.8 we have $|\pi| \geq |\psi| \geq mn\ell/32 = \Omega(n^2 \log n)$. \square

5.4 Structural Consequences

Theorem 5.1 immediately gives us two structural properties of RML and RMA. One is that proof length may decrease when introducing a weakening rule.

Corollary 5.11. *There exists a family of formulas over $O(n \log n)$ variables and $O(n \log n)$ clauses that have RML with weakening refutations of length $O(n \log n)$ but every RMA refutation requires length $\Omega(n^2 \log n)$.*

Proof. Consider the formula $F_n \wedge \bar{z}$, where F_n is the formula given by Theorem 5.1 and z is a new variable. If we weaken every clause $C \in F_n$ to $C \vee z$ then we can derive $F \vee z \vdash z$ in $O(n \log n)$ RML steps because each inference is a merge. However, if we cannot do weakening, then \bar{z} cannot be resolved with any clause in F_n and the lower bound of Theorem 5.1 applies. \square

The second property is that RML and RMA are not *natural* proof systems in the sense of [BKS04] because proof length may increase after a restriction.

Corollary 5.12. *There exists a restriction ρ and a family of formulas over $O(n \log n)$ variables and $O(n \log n)$ clauses that have RML refutations of length $O(n \log n)$ but every RMA refutation of $F_n \upharpoonright_\rho$ requires length $\Omega(n^2 \log n)$.*

Proof. Consider the formula $G_n = (F_n \vee z) \wedge \bar{z}$, where F_n is the formula given by Theorem 5.1, $F \vee z = \{C \vee z \mid C \in F\}$, and z is a new variable. As in the proof of Corollary 5.11 there is a RML derivation of z of length $O(n \log n)$ steps, while $G_n \upharpoonright_\rho = F_n$. \square

6 Further Separations

We can separate the different flavours of merge resolution that we introduced using a few variations of $F_{\ell, m, n}$ where we add a constant number of redundant clauses for each $i \in [\ell]$. We consider these different clauses part of \mathcal{W} .

Upper bounds all follow the same pattern. We first show on a case-by-case basis how to obtain $w_1 \bar{w}_n$ and $\bar{w}_1 w_n$ as lemmas, and then proceed as in Section 5.2.

Towards proving lower bounds we are going to generalize the lower bound part of the proof of Theorem 5.1 to apply to these variations as well. Fortunately we only require a few local modifications.

First, we need to prove an equivalent of Lemma 5.3, which we do on a case-by-case basis.

Second, we need to show that k -respecting restrictions can be extended to the new variables. For each block J_i , since the new clauses are semantically subsumed by $w_{i,1} = w_{i,2}$, there exists a way to map the new variables into $w_{i,1}$ and $w_{i,2}$ so that the result of the restriction is the same as if we had started with clauses $\bar{w}_{i,1} \vee w_{i,2}$ and $w_{i,1} \vee \bar{w}_{i,2}$, which are already part of \mathcal{W}_i . That is, the formula that we work with after Lemma 5.6 is a copy of an unaltered $F_{\ell', m', n'}$ formula.

The only part of the lower bound that depends on the specific subsystem of Resolution is Lemma 5.6; afterwards all the information we use is that no clause supported on W variables is reused. Furthermore, the only property of the subsystem that we use in the proof of Lemma 5.6 is that Lemma 5.3 applies. Therefore, the modifications we just outlined are sufficient for the lower bound to go through.

6.1 Separation between RMA and LRMA

Proposition 6.1. *There exists a family of formulas over $O(n \log n)$ variables and $O(n \log n)$ clauses that have RMA refutations of length $O(n \log n)$ but every LRMA refutation requires length $\Omega(n^2 \log n)$.*

The separating formula is $F_{m,n,\ell}^{(1)}$, where we add to $F_{m,n,\ell}$ clauses

$$\overline{w_{i,1}} \vee w_{i,2} \vee \overline{z_i}, \quad (\text{C1})$$

$$\overline{w_{i,2}} \vee z_i, \quad (\text{C2})$$

$$w_{i,1} \vee \overline{w_{i,2}} \vee \overline{y_i}, \quad (\text{C3})$$

$$w_{i,2} \vee y_i, \quad (\text{C4})$$

for each $i \in [\ell]$. The new variables can be assigned as $z_i = w_{i,1}$ and $y_i = \overline{w_{i,1}}$ to obtain the original formula back.

The upper bound follows from the following lemma.

Lemma 6.2. *Clauses $w_{i,1} \vee \overline{w_{i,n}}$ and $\overline{w_{i,1}} \vee w_{i,n}$ can be derived as lemmas from $F_{m,n,\ell}^{(1)}$ in length $O(n)$ in RMA.*

Proof. We resolve clause $\overline{w_{i,1}} \vee w_{i,2}$ first with (C2) and then (C1) in order to obtain $\overline{w_{i,1}} \vee w_{i,2}$ as a merge, then derive $\overline{w_{i,1}} \vee w_{i,n}$, having a merge as its ancestor, so it can be remembered. Analogously starting from $w_{i,1} \vee \overline{w_{i,2}}$, (C3), and (C4) we can obtain $w_{i,1} \vee \overline{w_{i,n}}$ as a lemma. \square

The following observation is useful for the lower bound.

Lemma 6.3. *Let C and D be clauses with two pairs of opposite literals. Then C and D cannot appear in the same locally regular input derivation.*

Proof. Let $C = x \vee y \vee C'$ and $D = \overline{x} \vee \overline{y} \vee D'$. Assume wlog that C is the first clause out of C and D to appear in the derivation. If x or y are used as pivots before D , then the locally regular condition prevents using D as an axiom. Otherwise $x \vee y$ appears in the derivation since the time C is used, which also prevents using D . \square

The equivalent of Lemma 5.3 is the following.

Lemma 6.4. *Let η be a LRMA derivation that only depends on \mathcal{W} axioms. Then no clause in η can be reused.*

Proof. We can only obtain a merge using one of (C1) or (C3), assume wlog (C1) is the first of these to be used in the derivation. By Lemma 6.3 neither (C2) nor (C3) appear in the derivation. We can show by induction that we can only obtain clauses of the form $\overline{w_{i,j}} \vee \overline{w_{i,j'}} \vee \overline{z_i}$ or $y_i \vee \overline{w_{i,j}} \vee \overline{z_i}$, never as a merge. \square

6.2 Separation between RML/LRMA and LRML

Proposition 6.5. *There exists a family of formulas over $O(n \log n)$ variables and $O(n \log n)$ clauses that have RML and LRMA and refutations of length $O(n \log n)$ but every LRML refutation requires length $\Omega(n^2 \log n)$.*

The separating formula is $F_{m,n,\ell}^{(2)}$, where we add to $F_{m,n,\ell}$ clauses

$$z_i \vee \overline{w_{i,1}} \vee w_{i,2}, \quad (\text{C1})$$

$$\overline{z_i} \vee \overline{w_{i,1}} \vee w_{i,2}, \quad (\text{C2})$$

$$y_i \vee w_{i,1} \vee \overline{w_{i,2}}, \quad (\text{C3})$$

$$\overline{y_i} \vee w_{i,1} \vee \overline{w_{i,2}}, \quad (\text{C4})$$

for each $i \in [\ell]$. The new variables can be assigned as $z_i = 1$ and $y_i = 1$ to obtain the original formula back.

The upper bounds follow respectively from the following lemmas.

Lemma 6.6. *Clauses $w_{i,1} \vee \overline{w_{i,n}}$ and $\overline{w_{i,1}} \vee w_{i,n}$ can be derived as lemmas from $F_{m,n,\ell}^{(2)}$ in length $O(n)$ in RML.*

Proof. We first resolve clauses $\overline{w_{n-1}} \vee w_{i,n}, \overline{w_{i,n-2}} \vee w_{i,n-1}, \dots, \overline{w_{i,2}} \vee w_{i,3}$, (C1) to obtain $z_i \vee \overline{w_{i,1}} \vee w_{i,n}$. We continue the input derivation resolving with (C2) to obtain $\overline{w_{i,1}} \vee w_{i,2} \vee w_{i,n}$. We then resolve with $\overline{w_{i,2}} \vee w_{i,3}, \overline{w_{i,3}} \vee w_{i,4}, \dots, \overline{w_{i,n-1}} \vee w_{i,n}$ to obtain $\overline{w_{i,1}} \vee w_{i,n}$ as a merge over $w_{i,n}$. Analogously we can obtain $w_{i,1} \vee \overline{w_{i,n}}$. \square

Lemma 6.7. *Clauses $w_{i,1} \vee \overline{w_{i,n}}$ and $\overline{w_{i,1}} \vee w_{i,n}$ can be derived as lemmas from $F_{m,n,\ell}^{(2)}$ in length $O(n)$ in LRMA.*

Proof. We resolve clauses (C1) and (C2) to obtain $\overline{w_{i,1}} \vee w_{i,2}$, which is a merge, then derive $\overline{w_{i,1}} \vee w_{i,n}$, having a merge as its ancestor, so it can be used as a lemma. Analogously starting from (C3) and (C4) we can obtain $w_{i,1} \vee \overline{w_{i,n}}$ as a lemma. \square

The equivalent of Lemma 5.3 is the following.

Lemma 6.8. *Let η be a LRML derivation that only depends on \mathcal{W} axioms. Then no clause in η can be reused.*

The proof idea is that the only merge we can obtain involves the z_i or the y_i variable. If we just resolve the two clauses over such a variable we obtain a clause we already had, so this is useless. Otherwise we are resolving one of w_2 away, which would be reintroduced at the time of resolving z_i away, and that is not allowed by the SR condition.

Proof. We can only obtain a merge by using one of the new clauses (C1)–(C4). If we resolve either pair of clauses over y_i or over z_i then we obtain a clause that was already present in the formula, and therefore we may preprocess such derivation away.

Otherwise consider the first step in the derivation where one of the new clauses is used as a premise, assume wlog it is (C1). That step is with a clause of the form $\overline{w_{i,2}} \vee w_{i,j}$, and we obtain a clause of the form $z_i \vee \overline{w_{i,1}} \vee w_{i,j}$, which is not a merge. That clause can be possibly resolved over $w_{i,j}$ ($j > 2$) to obtain other clauses of the same form, neither of which is a merge, but it cannot be resolved over y_i, z_i , or $w_{i,1}$ because that step would reintroduce variable $w_{i,2}$. \square

6.3 Separation between LRML and REL

Proposition 6.9. *There exists a family of formulas over $O(n \log n)$ variables and $O(n \log n)$ clauses that have LRML refutations of length $O(n \log n)$ but every REL refutation requires length $\Omega(n^2 \log n)$.*

The separating formula is $F_{m,n,\ell}^{(3)}$, where we add to $F_{m,n,\ell}$ clauses

$$\overline{w_{i,1}} \vee \overline{w_{i,2}} \vee w_{i,3}, \tag{C1}$$

$$w_{i,1} \vee w_{i,2} \vee \overline{w_{i,3}} \tag{C2}$$

for each $i \in [\ell]$. If we assign $w_{i,2} = w_{i,1}$ we obtain a copy of $F_{m,n-1,\ell}$ which, even if technically it is not the same formula we started with, is enough for our purposes.

The upper bound follows from the following lemma.

Lemma 6.10. *Clauses $w_{i,1} \vee \overline{w_{i,n}}$ and $\overline{w_{i,1}} \vee w_{i,n}$ can be derived as lemmas from $F_{m,n,\ell}^{(3)}$ in length $O(n)$ in LRML.*

Proof. We resolve (C1) with $\overline{w_{i,3}} \vee w_{i,4}, \dots, \overline{w_{i,n-1}} \vee w_{i,n}$ to obtain $\overline{w_{i,1}} \vee \overline{w_{i,2}} \vee w_{i,n}$, then with $\overline{w_{i,1}} \vee w_{i,2}$ to obtain $\overline{w_{i,1}} \vee w_{i,n}$ as a merge. Analogously starting from (C2) we can obtain $w_{i,1} \vee \overline{w_{i,n}}$ as a lemma. \square

The equivalent of Lemma 5.3 is the following.

Lemma 6.11. *Let η be a REL derivation that only depends on \mathcal{W} axioms. Then no clause in η can be reused.*

Proof. Observe that every derivable clause has width at least 2. Let C be any derivable clause and ℓ any literal in C . We have that $\alpha = \overline{C} \setminus \ell$ is not empty. However, assigning any variable $w_{i,j}$ immediately propagates all variables, hence ℓ is not empowering. \square

7 Concluding Remarks

In this paper, we address the question of the tightness of simulation of resolution proofs by CDCL solvers. Specifically, we show that RMA, among other flavours of DAG-like merge resolution, simulates standard resolution with at most a linear multiplicative overhead. However, contrary to what we see in the tree-like case, this overhead is necessary. While the proof systems we introduce help us explain one source of overhead in the simulation of resolution by CDCL, it is not clear if they capture it exactly. In other words, an interesting future direction would be to explore whether it is possible for CDCL to simulate some flavour of merge resolution with less overhead than what is required to simulate standard resolution.

Acknowledgements

The authors are grateful to Yuval Filmus and a long list of participants in the program *Satisfiability: Theory, Practice, and Beyond* at the Simons Institute for the Theory of Computing for numerous discussions. This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing.

A Tree-like Merge Resolution

For completeness we informally sketch the proofs of Lemma 3.2 and Theorem 3.3, which can be found in full detail in [And68].

Lemma A.1 (Lemma 3.2, restated). *If there is a tree-like resolution derivation of C of length L where at most the root is a merge, then there is an input resolution derivation of some $C' \subseteq C$ of length at most L .*

Proof (sketch). We prove by induction on $|\eta|$ that for every axiom $E \in \eta$ there exists an input derivation of C' that uses a subset of the axioms of η where E is the topmost axiom. As intermediate objects we allow clauses in this derivation to contain opposite literals; these are cleaned up later.

Let $C = \text{Res}(A \vee x, B \vee \overline{x})$, and let η_1 and η_2 be the derivations used to infer $A \vee x$ and $B \vee x$ respectively. Assume wlog that $E \in \eta_1$. Since η_2 does not contain any merges there exists a unique path from $B \vee \overline{x}$ to an axiom $D \vee \overline{x} \in \eta_2$, where all clauses contain \overline{x} . Note that other clauses in η_2 might still contain x or \overline{x} . We replace $D \vee \overline{x}$ by D in η_2 (and consequently remove all the occurrences of \overline{x} in the aforementioned path) and we obtain a valid derivation η_3 of B . We apply the induction hypothesis to η_1 and η_3 to obtain two unit derivations η_4 and η_5 of $A' \vee x \subseteq A \vee x$ and $B' \subseteq B$ whose topmost leaves are E and D . We replace D by $D \vee A'$ in η_5 and obtain a unit derivation η_6 of $B' \vee A' \subseteq B' \vee A'$. We stitch together η_4 and η_6 by observing that $\text{Res}(A' \vee x, D \vee \overline{x}) = A' \vee D$,




Figure 2: Proof of Lemma A.1




Figure 3: Proof of Theorem A.2

which is the only axiom in η_6 not present in the original axioms, and obtain a unit derivation η_7 of $B \vee A'' = C' \subseteq C$ that only uses original axioms.

Finally, and outside the inductive argument, we get rid of clauses that contain opposite literals by replacing any such clause by 1 to obtain a semantic derivation η_8 . Its syntactic counterpart $s(\eta_8)$ satisfies the conclusion of the lemma. \square

Theorem A.2 (Theorem 3.3, restated). *If there is a tree-like resolution derivation of C of length L , then there is a merge resolution derivation of some $C' \subseteq C$ of length at most L .*

Proof (sketch). The proof is by induction on the number of merges. The base case when there are no merges follows by Lemma A.1. Otherwise let ψ be a subtree where exactly the root C is a merge. Let ψ' be the input resolution derivation of C' given by Lemma A.1, let D be the last merge in ψ' , and let ω and ω' be the fragments of ψ' from D to C' and up to D respectively. We replace ψ by ω in η to obtain a refutation η' that uses D as an axiom (note that in replacing C by C' we may have to prune away parts of η). Because η' has one less merge we can apply the induction hypothesis and obtain a merge resolution derivation ψ'' . Finally we replace the axiom D by the derivation ω' . \square

References

- [ABH⁺08] Gilles Audemard, Lucas Bordeaux, Youssef Hamadi, Saïd Jabbour, and Lakhdar Sais. A generalized framework for conflict analysis. In Hans Kleine Büning and Xishun Zhao, editors, *Theory and Applications of Satisfiability Testing - SAT 2008, 11th International*

- Conference, SAT 2008, Guangzhou, China, May 12-15, 2008. Proceedings, volume 4996 of *Lecture Notes in Computer Science*, pages 21–27. Springer, 2008. doi:[10.1007/978-3-540-79719-7_3](https://doi.org/10.1007/978-3-540-79719-7_3).
- [AFT11] Albert Atserias, Johannes Klaus Fichte, and Marc Thurley. Clause-learning algorithms with many restarts and bounded-width resolution. *Journal of Artificial Intelligence Research*, 40:353–373, January 2011. Preliminary version in SAT ’09.
- [And68] Peter B. Andrews. Resolution with merging. *J. ACM*, 15(3):367–381, 1968.
- [BB21] Olaf Beyersdorff and Benjamin Böhm. Understanding the relative strength of QBF CDCL solvers and QBF resolution. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPICs*, pages 12:1–12:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:[10.4230/LIPICs.ITCS.2021.12](https://doi.org/10.4230/LIPICs.ITCS.2021.12).
- [BF97] Avrim L Blum and Merrick L Furst. Fast planning through planning graph analysis. *Artificial intelligence*, 90(1-2):281–300, 1997.
- [BHJ08] Samuel R. Buss, Jan Hoffmann, and Jan Johannsen. Resolution trees with lemmas: Resolution refinements that characterize DLL-algorithms with clause learning. *Logical Methods in Computer Science*, 4(4:13), December 2008.
- [BKS04] Paul Beame, Henry Kautz, and Ashish Sabharwal. Towards understanding and harnessing the potential of clause learning. *Journal of Artificial Intelligence Research*, 22:319–351, December 2004. Preliminary version in *IJCAI ’03*.
- [BN21] Sam Buss and Jakob Nordström. Proof complexity and SAT solving. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, chapter 7, pages 233–350. IOS Press, 2nd edition, 2021. doi:[10.3233/FAIA200990](https://doi.org/10.3233/FAIA200990).
- [BS97] Roberto J. Bayardo Jr. and Robert Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI ’97)*, pages 203–208, July 1997.
- [CGP⁺08] Cristian Cadar, Vijay Ganesh, Peter M Pawlowski, David L Dill, and Dawson R Engler. EXE: Automatically Generating Inputs of Death. *ACM Transactions on Information and System Security (TISSEC)*, 12(2):1–38, 2008.
- [DHN07] Nachum Dershowitz, Ziyad Hanna, and Alexander Nadel. Towards a better understanding of the functionality of a conflict-driven SAT solver. In João Marques-Silva and Karem A. Sakallah, editors, *Theory and Applications of Satisfiability Testing - SAT 2007, 10th International Conference, Lisbon, Portugal, May 28-31, 2007, Proceedings*, volume 4501 of *Lecture Notes in Computer Science*, pages 287–293. Springer, 2007. doi:[10.1007/978-3-540-72788-0_27](https://doi.org/10.1007/978-3-540-72788-0_27).
- [dV94] Alvaro del Val. Tractable databases: How to make propositional unit resolution complete through compilation. In Jon Doyle, Erik Sandewall, and Pietro Torasso, editors, *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning (KR’94)*. Bonn, Germany, May 24-27, 1994, pages 551–561. Morgan Kaufmann, 1994.

- [DVT07] Julian Dolby, Mandana Vaziri, and Frank Tip. Finding Bugs Efficiently With a SAT Solver. In *Proceedings of the 6th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 195–204, 2007. doi:[10.1145/1287624.1287653](https://doi.org/10.1145/1287624.1287653).
- [FB20] Nick Feng and Fahiem Bacchus. Clause size reduction with all-uir learning. In Luca Pulina and Martina Seidl, editors, *Theory and Applications of Satisfiability Testing - SAT 2020 - 23rd International Conference, Alghero, Italy, July 3-10, 2020, Proceedings*, volume 12178 of *Lecture Notes in Computer Science*, pages 28–45. Springer, 2020. doi:[10.1007/978-3-030-51825-7_3](https://doi.org/10.1007/978-3-030-51825-7_3).
- [HBPV08] Philipp Hertel, Fahiem Bacchus, Toniann Pitassi, and Allen Van Gelder. Clause learning can effectively P-simulate general propositional resolution. In *Proceedings of the 23rd National Conference on Artificial Intelligence (AAAI '08)*, pages 283–290, July 2008.
- [LFV⁺20] Chunxiao Li, Noah Fleming, Marc Vinyals, Toniann Pitassi, and Vijay Ganesh. Towards a complexity-theoretic understanding of restarts in sat solvers. In *Theory and Applications of Satisfiability Testing–SAT 2020: 23rd International Conference, Alghero, Italy, July 3–10, 2020, Proceedings 23*, pages 233–249. Springer, 2020.
- [MLM21] João Marques-Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pages 133–182. IOS Press, 2021. doi:[10.3233/FAIA200987](https://doi.org/10.3233/FAIA200987).
- [MMZ⁺01] 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 '01)*, pages 530–535, June 2001.
- [MPR20] Nathan Mull, Shuo Pang, and Alexander A. Razborov. On CDCL-based proof systems with the ordered decision strategy. In *Proceedings of the 23rd International Conference on Theory and Applications of Satisfiability Testing (SAT '20)*, volume 12178 of *Lecture Notes in Computer Science*, pages 149–165. Springer, July 2020.
- [MS85] Burkhard Monien and Ewald Speckenmeyer. Solving satisfiability in less than 2^n steps. *Discret. Appl. Math.*, 10(3):287–295, 1985. doi:[10.1016/0166-218X\(85\)90050-2](https://doi.org/10.1016/0166-218X(85)90050-2).
- [MS99] João P. Marques-Silva and Karem A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 48(5):506–521, May 1999. Preliminary version in ICCAD '96.
- [PD08] Knot Pipatsrisawat and Adnan Darwiche. A new clause learning scheme for efficient unsatisfiability proofs. In Dieter Fox and Carla P. Gomes, editors, *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13-17, 2008*, pages 1481–1484. AAAI Press, 2008. URL: <http://www.aaai.org/Library/AAAI/2008/aaai08-243.php>.
- [PD11] Knot Pipatsrisawat and Adnan Darwiche. On the power of clause-learning SAT solvers as resolution engines. *Artificial Intelligence*, 175(2):512–525, February 2011. Preliminary version in CP '09.
- [Rya04] Lawrence Ryan. Efficient algorithms for clause-learning SAT solvers. Master's thesis, Simon Fraser University, 2004.

- [SB09] Niklas Sörensson and Armin Biere. Minimizing learned clauses. In *Proceedings of the 12th International Conference on Theory and Applications of Satisfiability Testing (SAT '09)*, volume 5584 of *Lecture Notes in Computer Science*, pages 237–243. Springer, July 2009.
- [Van05] Allen Van Gelder. Pool resolution and its relation to regular resolution and DPLL with clause learning. In *Proceedings of the 12th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR '05)*, volume 3835 of *Lecture Notes in Computer Science*, pages 580–594. Springer, 2005.
- [Vin20] Marc Vinyals. Hard examples for common variable decision heuristics. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI '20)*, pages 1652–1659, February 2020.
- [XA05] Yichen Xie and Alexander Aiken. Saturn: A SAT-Based Tool for Bug Detection. In *Proceedings of the 17th International Conference on Computer Aided Verification, CAV 2005*, pages 139–143, 2005. doi:10.1007/11513988\13.
- [ZMMM01] Lintao Zhang, Conor F. Madigan, Matthew W. Moskewicz, and Sharad Malik. Efficient conflict driven learning in Boolean satisfiability solver. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD '01)*, pages 279–285, November 2001.
- [ZMW⁺18] Edward Zulkoski, Ruben Martins, Christoph M. Wintersteiger, Jia Hui Liang, Krzysztof Czarnecki, and Vijay Ganesh. The effect of structural measures and merges on SAT solver performance. In John N. Hooker, editor, *Principles and Practice of Constraint Programming - 24th International Conference, CP 2018, Lille, France, August 27-31, 2018, Proceedings*, volume 11008 of *Lecture Notes in Computer Science*, pages 436–452. Springer, 2018. doi:10.1007/978-3-319-98334-9\29.