

# Reducing the Size of Resolution Proofs in Linear Time <sup>\*</sup>

Omer Bar-Ilan<sup>1</sup> Oded Fuhrmann<sup>2</sup> Shlomo Hoory<sup>3</sup> Ohad Shacham<sup>4</sup>  
Ofer Strichman<sup>5</sup>

<sup>1</sup> Rafael, Israel [omerb@rafael.co.il](mailto:omerb@rafael.co.il)

<sup>2</sup> Google, California, USA [frmoded@gmail.com](mailto:frmoded@gmail.com)

<sup>3</sup> Syntact Design, Israel [hoorys@gmail.com](mailto:hoorys@gmail.com)

<sup>4</sup> IBM Haifa Research Laboratory, Israel [ohads@il.ibm.com](mailto:ohads@il.ibm.com)

<sup>5</sup> Information Systems Engineering, IE, Technion, Haifa, Israel.  
[offers@ie.technion.ac.il](mailto:offers@ie.technion.ac.il)

**Abstract.** DPLL-based SAT solvers progress by implicitly applying binary resolution. The resolution proofs that they generate are used, after the SAT solver's run has terminated, for various purposes. Most notable uses in formal verification are: extracting an *unsatisfiable core*, extracting an *interpolant*, and detecting clauses that can be reused in an *incremental satisfiability* setting (the latter uses the proof only implicitly, during the run of the SAT solver). Making the resolution proof smaller can benefit all of these goals: it can lead to smaller cores, smaller interpolants, and smaller clauses that are propagated to the next SAT instance in an incremental setting. We suggest two methods that are linear in the size of the proof for doing so. Our first technique, called RECYCLE-UNITS, uses each learned constant (unit clause) ( $x$ ) for simplifying resolution steps in which  $x$  was the pivot, prior to when it was learned. Our second technique, called RECYCLE-PIVOTS, simplifies proofs in which there are several nodes in the resolution graph, one of which dominates the others, that correspond to the same pivot. Our experiments with industrial instances show that these simplifications reduce the core by  $\approx 5\%$  and the proof by  $\approx 13\%$ . It reduces the core less than competing methods such as RUN-TILL-FIX, but whereas our algorithms are linear in the size of the proof, the latter and other competing techniques are all exponential as they are based on SAT runs. If we consider the size of the proof (the resolution graph) as being polynomial in the number of variables (it is not necessarily the case in general), this gives our method an exponential time reduction comparing to existing tools for small core extraction. Our experiments show that this result is evident in practice more so for the second method: rarely it takes more than a few seconds, even when competing tools time out, and hence it can be used as a cheap proof post-processing procedure.

---

<sup>\*</sup> This article extends [3] with proofs and a more elaborated discussion of related work. The first three authors were in IBM when this research was conducted.

## 1 Introduction

DPLL-based SAT solving is the most used back-end engine for model checking and satisfiability modulo theories (SMT). While SAT solvers exist from at least the 1960's, and learning through derivation of conflict clauses exists from at least the 1980's (while recognized as implicitly derived by resolution), only in 2003 the question of how to produce a resolution proof from a run of a DPLL solver was addressed in practice [26] and implemented. Most modern solvers nowadays are capable of producing such proofs. Further, there are now decision heuristics that are based on an understanding of the DPLL process as a resolution-based proof engine rather than as a search engine. Examples are Ryan's thesis [20] and his SAT solver SIEGE, which bias conflict-clause generation towards those that will more likely lead to other resolutions, and the work by Gershman et al. [9] on a model for explaining and designing decision heuristics based on an understanding of the SAT solving process as a resolution engine.

The resolution proofs that modern SAT solvers generate are used in a broad range of applications in formal verification and elsewhere. Prominent examples are:

- The resolution proof can be used for extracting an *unsatisfiable core*, which can then be used, e.g., in a proof-based abstraction-refinement procedure, as shown by Amla and McMillan [1, 2]. The unsatisfiable core was also used in the past for detecting the reasons for unsatisfiability in an underapproximation/refinement process for model-checking [11]. In the context of satisfiability modulo theories (SMT): the core is used for finding small *explanations* (see, for example, the recent work by Cimatti et al. [5], who suggest to invoke a tool for minimizing resolution proofs for this purpose) and in theory-specific decision procedures, such as the abstraction-based procedures for Presburger and bitvector formulas proposed by Kroening et al. and Bryant et al., respectively [13, 4].
- The resolution proof can be used for extracting an *interpolant* as part of a complete model-checking technique, as suggested by McMillan in [16].
- The proof can be used for detecting clauses that can be reused in an *incremental satisfiability* setting [21, 24], such as the one used in Bounded Model-Checking. The analysis of the proof is done implicitly, during the run of the SAT solver (in contrast to the first two uses), and is required in order to check whether all the clauses that were used for resolving a particular conflict clause are still expected to be present in the new SAT instance. If yes, this conflict clause can be reused.

Making the proof smaller by removing some of the nodes and removing literals from other nodes can benefit all of these goals. In all of these applications, however, reducing the size of the core/interpolant/reused-clauses has an influence on the overall run time which is somewhat unpredictable and only vaguely known. Further, the proof reduction component is called many times during the overall solving process. As a result, best overall results are most likely to be achieved with a limited investment in such reductions.

The script RUN-TILL-FIX by Zhang and Malik [26] extracts a core and attempts to minimize it by simply running the SAT solver on it repeatedly until a fix-point is reached. Achieving fast reductions with this tool is not always possible, as it requires a normal SAT run. The goal of achieving fast reductions was first addressed in the work by Gershman et al. [8], based on analyzing the resolution graph and trying to restructure it with the aid of a SAT solver. The tasks the SAT solver is asked to solve by their script TRIM-TILL-FIX are closely related, and hence incremental satisfiability makes this process relatively fast. There are also many published techniques for finding *minimal* cores, i.e., an unsatisfiable subset of clauses from which no clause can be removed without making it satisfiable (this is also known in the literature by the name MUS, for Minimal Unsatisfiable Subformula). Some works in this direction from the last four years are [15, 18, 12, 17, 6, 10], all of which are worst-case exponential. The complexity of the decision problem corresponding to finding the minimal unsatisfiable core is  $D^P$ -complete<sup>1</sup> [19]. A minimal core (in contrast to the *minimum* core) is not unique and depends on the starting point, like all the methods we are aware of (including the current one). Therefore whether the core found is minimal or just ‘small’, as indicated in [8], has little significance in practice.

RUN-TILL-FIX is linear in the size of the proof (extracting the core) and exponential in the number of variables (rerunning the SAT solver). Our method is only linear (or quadratic, depending on which of the two suggested methods is used) in the size of the proof. The proof itself can be exponential in the number of variables, and hence in the worst case our procedure is still exponential in the number of variables, like RUN-TILL-FIX. But the latter still has the SAT phase which our method doesn’t. The significance of the exponential gap is that our techniques typically reduce the core significantly less than the exponential methods, but do so much faster. They also reduce the size of the proof itself whereas a procedure like RUN-TILL-FIX, as our experiments show, typically increases it. We therefore consider them as useful tools in the context of short time-outs and where the size of the proof matters.

The first method we suggest is called RECYCLE-UNITS. The idea is to remove edges from the resolution graph by using information that is inferred by the SAT solver only *after* the resolutions at the nodes adjacent to these branches were made originally. For example, if  $(x)$  is a unit clause that was learned by the SAT solver, it can be used for simplifying resolution inferences that used  $x$  as the pivot prior to learning this clause. If not carefully applied, however, this may lead to circular reasoning.

The second method is called RECYCLE-PIVOTS. It is based on the following observation. For simplicity assume that the resolution graph is a tree. Let  $n_1$  and  $n_2$  be two nodes on the same path in the proof tree such that  $n_1$  is closer to the sink node. Further, assume that the pivot variable associated with both nodes is the same. Our convention is that proofs progress from top to bottom, from the

---

<sup>1</sup>  $D^P$  is the class containing all languages that can be considered as the difference between two languages in NP, or equivalently, the intersection of a language in NP with a language in co-NP.

premises (also called the *axioms*) of the proof to its consequent. We also follow the convention by which the right parent of each node contains the negative phase of the pivot variable, and the left parent contains the positive phase of this variable. Assume that  $n_2$  is on the right branch of  $n_1$ . In this case, we will show, the left incoming branch of  $n_2$  can be pruned, and the proof rewritten without it, in a way that the resulting proof is a legal resolution proof with a smaller core. Since in practice resolution graphs are DAGs and not just trees, the procedure is somewhat more complicated as we later describe.

The two techniques tighten the proof, which means that the resulting proof uses a subset of the core used by the original proof, and there is an injective mapping from the target to the source nodes, such that each target node is either equal to or subsumes the source node to which it is mapped. As a theoretical curiosity, we note that the resolution proofs generated by our techniques cannot necessarily be generated by any modern SAT solver and hence by any core reduction technique that is based on rerunning such a solver. More details about this issue can be found in the appendix.

The rest of the paper is structured as follows: Section 2 summarizes some preliminaries necessary for the description of the technique (we assume, however, that the reader is familiar with SAT basics). Section 3 and 4 describe the two techniques by giving pseudo code and intuitive explanation of their correctness. Section 5 is dedicated to a formal proof of their correctness. We conclude in Sect. 6 with a description of the experiments we conducted and their results.

## 2 Preliminaries

A literal is a Boolean variable or its negation. A pair of literals corresponding to a variable and its negation are called *complementary*. A clause is a (possibly empty) disjunction of literals, and a CNF formula is a conjunction of clauses. For convenience we occasionally refer to a clause as a set of literals.

### 2.1 Inference by Resolution

The process in which DPLL SAT solvers infer new conflict clauses can be interpreted as applying the *binary resolution* inference rule:

$$\frac{(l \vee l_1 \vee \dots \vee l_n) \quad (\neg l \vee l'_1 \vee \dots \vee l'_{n'})}{(l_1 \vee \dots \vee l_n \vee l'_1 \vee \dots \vee l'_{n'})} \text{ (RESOLUTION)}. \quad (1)$$

The variable  $l$  is called the *pivot variable* (also called ‘resolution variable’ in the literature). The resulting clause may contain multiple occurrences of the same literal, so we assume that such redundancy is removed.

Let *Res* be a function that receives two clauses with complementary literals as input, and returns the consequent of the resolution rule applied to these two clauses, as output. More formally:

$$\text{Given clauses } C_1 = (l \vee l_1 \vee \dots \vee l_n) \text{ and } C_2 = (\neg l \vee l'_1 \vee l'_2 \vee \dots \vee l'_{n'}),$$

$$\text{Res}(C_1, C_2) = (l_1 \vee \dots \vee l_n \vee l'_1 \vee \dots \vee l'_n).$$

Resolution is known to be a sound and complete proof system for CNF formulas. Specifically, a CNF formula  $\varphi$  is unsatisfiable if and only if there exists a resolution proof of the empty clause using  $\varphi$ 's clauses as premises. We use the formulation of [1] for resolution proofs:

**Definition 1 (Resolution proof of unsatisfiability).** *A resolution proof of unsatisfiability of a set of clauses  $C$  is a directed acyclic graph (DAG)  $(V, E)$ , where the nodes represent clauses, and each node  $c \in V$  is either a root and  $c \in C$ , or has exactly two parents  $c_1, c_2$  such that  $c$  is their resolvent. The graph has a single sink, which represents the empty clause.<sup>2</sup>*

Since resolution proofs are graphs, we use the terms *proof* and *resolution graph* interchangeably.

Example of a resolution graph can be seen in Fig. 1[a]. Every resolution proof can be represented by a resolution graph. If a resolution graph has a single sink, it is called the *consequent* of the proof. The root nodes are the premises of the proof. For a given node, the root nodes that can reach it on the resolution graph are called its *core*. Specifically, if the consequent is an empty clause then its core is called the *unsatisfiable core*.

The resolution described so far is known by the name *general resolution*. Two well-known restrictions of general resolutions that we will mention later on are *tree resolution*, which means that the resolution graph is a tree rather than a DAG, and Tseitin's *regular resolution* [22], which means that along each path no variable is used twice as a pivot. Both of these restrictions may entail a penalty of an exponent in the size of the proof. In other words, there are classes of formulas that can be proven with general resolution in a polynomial number of steps, but only with an exponential number of steps in a tree resolution proof or regular resolution proof.

How do resolution proofs relate to proofs of SAT solvers? Modern DPLL SAT solvers generate *conflict clauses* [25] during their run, which are implicitly inferred from other clauses by a chain of (general) resolutions. Hence, the roots of the resolution proof that is given by solvers are labeled with original clauses, and the internal nodes are labeled with conflict clauses. We refer the reader to [1, 26] for more details.

We are going to use a variant of the resolution graph in which a single parent is allowed, if this parent is associated with the same clause as the child. The resolution proof corresponding to this graph is derived by first eliminating all nodes with a single parent (removing such a node  $n$  and connecting  $n$ 's single parent to  $n$ 's children), and continuing as before. This extension is convenient for simplifying the algorithm and later on the proofs, but is not essential.

<sup>2</sup> Note that we use the convention by which the edges are in the direction of the proof, i.e., from premises to consequents. For practical purposes it is common to build this graph with edges (pointers) pointing in the other direction, because this facilitates a search for the core.

### 3 Recycling Learned Unit Clauses

Some of the conflict clauses learned during the run of a SAT solver are unit clauses, e.g.,  $(x)$ . In such a case we say that the SAT solver inferred the *constant value* of the variable constrained by this clause ( $x = true$  in this case). These constants can be used to rewrite those parts of the proof that were inferred prior to learning these constants, an action that reduces the overall size of the proof and its core. Recall that the SAT solver can only apply resolution to clauses in its clause database, and that this database is continuously updated with new conflict clauses. Further, observe that the solver can only use resolution variables which at that point in time are unassigned. If at a later stage of the computation the same resolution variable is proved to have a constant value then the proof can be regenerated taking this information into account.

Algorithm 1 presents RECYCLE-UNITS, which is the first of a two-step algorithm for recycling unit clauses. The second step is RECONSTRUCT-PROOF, which is presented in Algorithm 2. The two algorithms use the following notation.  $P$  is a resolution proof of the empty clause. For a given node  $n$  in  $P$ ,  $n.C$  is the clause represented by  $n$ ,  $n.L$  and  $n.R$  are the left and right parents of  $n$  respectively, and  $n.piv$  is the pivot variable used to resolve  $n.C$  from  $n.L.C$  and  $n.R.C$ . Recall that we use the convention by which the left parent includes the positive phase of the pivot, and the right parent includes its negative phase. If  $l$  is a literal, we denote by  $var(l)$  its corresponding variable. When  $n.C$  is a unit clause, we sometimes refer to it as a literal rather than a clause, when the meaning is clear from the context. For example  $var(n.C)$  is the variable corresponding to the literal in the unit clause  $n.C$ .

**Algorithm 1:** RECYCLE-UNITS (PROOF  $P$ )

```

1: Let  $U$  be the set of nodes representing constants proved in  $P$ ;
2: for each  $u \in U$  do
3:   Mark the ancestors of  $u$ ;
4:   for each unmarked  $n \in P$  do
5:     if  $n.piv = u.C$  then
6:        $n.L := u$ ;
7:     else if  $n.piv = \neg u.C$  then
8:        $n.R := u$ ;

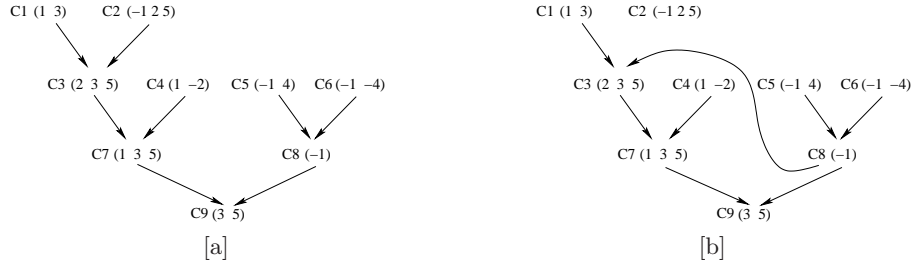
```

RECYCLE-UNITS iterates over all constants that were proved in  $P$ . Let  $u$  be a node representing such a constant, and assume for now that this constant is a positive literal (i.e.,  $u.C$  is a positive literal). First, in line 3, RECYCLE-UNITS marks the ancestors of  $u$ , i.e., the nodes that can reach  $U$  in  $P$ . Then, it searches unmarked nodes for those that represent a resolution step using  $var(u.C)$  as pivot. Let  $n$  be such a node. According to our convention  $u.C \in n.L.C$  and  $\neg u.C \in n.R.C$ . In line 6 RECYCLE-UNITS changes the left parent of  $n$  from  $n.L$

to  $u$ , i.e., it disconnects the edge  $(n.L, n)$  and connects instead  $(u, n)$ . If  $u.C$  is a negative literal, then the edge is shifted from  $(n.R, n)$  to  $(u, n)$  in line 8.

At this stage the proof is no longer a legal resolution proof, and hence a reconstruction phase begins by calling RECONSTRUCT-PROOF, which appears in Alg. 2.

*Example 1.* Consider the partial resolution graph that is depicted in Fig. 1a. The unit clause C8 is proven only after the resolution of C3, which uses the variable ‘1’ (the unit of C8) as its pivot. RECYCLE-UNITS begins by marking clauses that were used for proving C8 – clauses C5 and C6 in this case. It then identifies C3 as an unmarked node that uses ‘1’ as pivot and rewires the proof – in this case disconnects  $(C2, C3)$  and adds instead the edge  $(C8, C3)$ , as can be seen in Fig. 1b. It is left to reconstruct the proof so it becomes a legitimate resolution proof once again.



**Fig. 1.** [a] Part of a resolution proof [b] After RECYCLE-UNITS and before RECONSTRUCT-PROOF.

□

*Implementation of RECYCLE-UNITS* The most time-consuming component of Algorithm 1 is Step 3. Recall that the purpose of this step is to prevent cycles after we connect a unit clause to another node. Rather than traversing the graph backwards each time (which would make this method quadratic in the size of the graph), our implementation maintains at each unit node  $u$  pointers to every descendant unit  $u'$  such that there is no other unit between them. Denote by  $G_U$  the resulting graph of units, i.e., the nodes of  $G_U$  are the units in the resolve graph and the edges are defined by the list of pointers that we maintain with each such node. In Step 3, when considering connecting a unit clause  $u$  to another node  $c$ , our tool temporarily makes this connection and checks if it can reach itself on  $G_U$ . If the answer is yes, we undo this temporary connection and continue. Otherwise we keep the connection and update the list of pointers in the units that are parents of  $u$  so they now also point to  $u$ . This method makes the solution quadratic in the number of units. Note that we can always bound the number

of units that we consider if it becomes not cost-effective and that in practice the number of units is small comparing to the size of the proof.

Let us now shift the focus to RECONSTRUCT-PROOF. The resolution graph given to RECONSTRUCT-PROOF as input has a subset of the nodes of the original resolution graph. This is because only edges are shifted up to this point, and these shifts may disconnect parts of the graph (e.g., in Example 1, node C2 has been disconnected). Note that only the graph connected to the sink node is sent to RECONSTRUCT-PROOF.

RECONSTRUCT-PROOF, appearing in Alg. 2, is a recursive procedure that, given a proof node  $n$  – initially the sink node of the “broken” proof – reconstructs a legal resolution proof of  $n$  (we will define formally the nature of this broken proof in Sect. 5). We can think of  $n$  as a pointer (in the algorithm below it appears as sent by reference, C++ style), which in the end of the procedure will point to the sink of a legal resolution proof of that node. The base of the recursion are the root nodes. When  $n$  is not a root, there are several cases. If the pivot  $n.piv$  is still present in its parents then  $n.C$  is a resolution between them (see lines 9 and 10). If it is only present in one of them, say the positive phase is present in  $n.L.C$ , then the other node ( $n.R$ ) replaces  $n$ . This is because we know that  $n.R.C$  subsumes  $n.C$ . If it is contained in neither of its parents, then both  $n.L.C$  and  $n.R.C$  subsume  $n.C$  and hence either one of them can replace  $n.C$ .

*Example 2.* The graphs in Fig. 2 show the steps of reconstruction for the resolution graph in Fig. 1b.

□

## 4 Recycling Pivots

The second technique we present is linear in the size of the proof and is also based on two steps, as in the previous case: in the first step we remove edges from the proof, and in the second step we reconstruct the proof using the same algorithm RECONSTRUCT-PROOF. The algorithm is based on an observation of Urquhart [23] that along each path from root to sink, there is no need for resolving on the same variable more than once. If there is such a situation, then the redundant resolution steps can be avoided and consequently some of the branches of the proof can be pruned away. The correctness of this algorithm will be proven in Sect. 5. We note that this does not result in a regular resolution proof as defined in Sect. 2, because we apply it only to some parts of the graph as described below. If the resolution graph we start with happens to be a tree, then the result of applying this step is indeed a regular resolution proof. Our algorithm can be seen as a realization of an idea that was originally described in the proof of Lemma 5.1 in [23], stating that “a tree resolution refutation of minimal size is regular”.

For simplicity of the presentation, assume for now that the proof is a tree rather than a DAG. Consider a node  $n$  with a pivot  $n.piv$ . We propagate  $\neg n.piv$

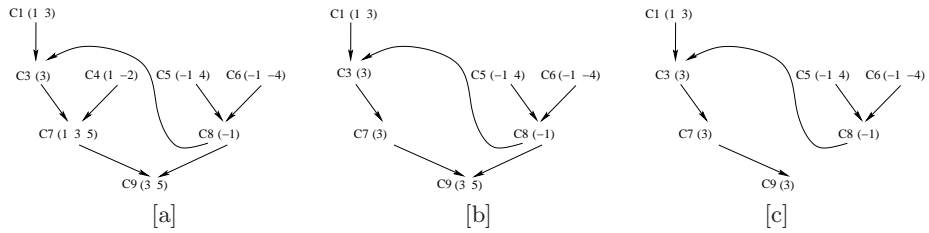


**Algorithm 2:** RECONSTRUCT-PROOF (NODE  $n$ )

```

1: if  $n$  visited then return
2: mark  $n$  as visited;
3: if  $n$  is a root then return
4: if  $n$  has a single parent  $s \in \{L, R\}$  then
5:   RECONSTRUCT-PROOF ( $n.s$ );
6:    $n := n.s$ ;
7:   return ;
7: RECONSTRUCT-PROOF ( $n.L$ );
8: RECONSTRUCT-PROOF ( $n.R$ );
9: if  $n.piv \in n.L.C$  and  $\neg n.piv \in n.R.C$  then
10:   $n.C := Res(n.L.C, n.R.C)$ ;
11: else if  $n.piv \in n.L.C$  and  $\neg n.piv \notin n.R.C$  then
12:   $n.C := n.R.C$ ;
13:   $n.L := nil$ 
14: else if  $n.piv \notin n.L.C$  and  $\neg n.piv \in n.R.C$  then
15:   $n.C := n.L.C$ ;
16:   $n.R := nil$ ;
17: else
18:   $side :=$  one of  $\{L, R\}$ ;  $otherside :=$  other side;  $\triangleright$  Choose heuristically
19:   $n.C := n.side.C$ ;
20:   $n.otherside := nil$ ;

```



**Fig. 2.** The recursive steps, from left to right, of reconstructing the resolution graph in Fig. 1b. Recall that single parent nodes are consistent with our definition of resolution graphs.

up the right branch (recall that according to our convention  $n.R.C$  contains  $\neg n.piv$ ) in a set called *Removable-Literals*, or  $RL$  for short. Similarly, we propagate  $n.piv$  up the left branch. Consider the right branch: if along this branch there is another node  $n'$  such that  $n'.piv = n.piv$  then we can replace  $n'$  with  $n'.R$ . This means that the branch starting at  $n'.L$  is pruned. The correctness of this operation is tied to the second step, which is the proof reconstruction in RECONSTRUCT-PROOF. Note that  $n'.R.C$  is contained in  $n'.C$  other than  $\neg n.piv$ . RECONSTRUCT-PROOF will effectively propagate  $\neg n.piv$  down the branch until inevitably reaching  $n$  (a result of our assumption that this is a tree). At node  $n$ ,  $\neg n.piv$  will disappear again due to the resolution on  $n.piv$  at  $n$ . As a result  $n.C$  will subsume its original version.

In practice the input proof can be a DAG. This may cause a situation in which our node  $n'$  has paths to the sink not through  $n$ . This, in turn, may cause a situation that RECONSTRUCT-PROOF propagates  $\neg n.piv$  all the way down to the sink, which contradicts our goal of producing a proof with an equal or stronger consequent. Another possible problem is that  $n'$  has paths to the sink node through both incoming edges of  $n$ , which nullifies our suggested technique. There are two possible solutions to this problem. One, which is the solution taken in Alg. 3 (see line 4) and also in our implementation, is to propagate  $RL$  up (towards the roots) only as long as it is a tree, i.e., we stop at a node that has multiple children. A more complicated solution is to check whether all paths from  $n'$  to the sink go through the edge  $(n.R, n)$ . This can be done by computing *dominance* relation in the graph, for example with the Lengauer-Tarjan algorithm[14] (which runs in  $O(|E|\log|V|)$  time). In order to achieve the maximum reduction of the proof size one must recalculate the dominance relation after each edge removal. This might be computationally expensive, so recomputing this relation only a bounded number of times may be a better choice. We leave this investigation for future work.

Given a resolution proof  $P$ , Alg. 3 is called with its sink node  $n$  and  $RL = \emptyset$ , and modifies  $P$ .

*Example 3.* Assume Fig. 3a represents the proof pointed to by  $n$ , the input for RECYCLE-PIVOTS. RECYCLE-PIVOTS propagates up the removable literals (denoted by  $RL$  in the drawing) and, owing to the fact that ‘2’ is the pivot of C3 and that ‘2’ is in  $RL$ , removes the edge  $(C1,C3)$ . The proof after calling RECONSTRUCT-PROOF is depicted in Fig. 3b.

□

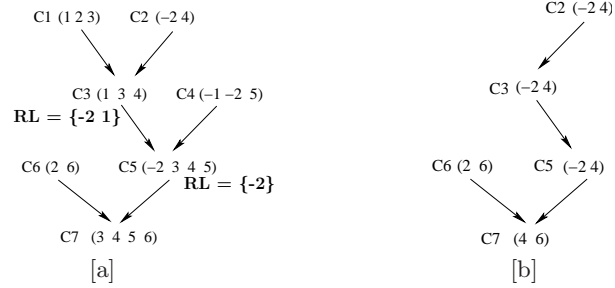
Let us conclude this section by mentioning that this technique was used recently by some of the authors in [7] in a different context: extending a resolution proof of correctness over bounded paths (e.g., Bounded Model Checking formulas) into an inductive proof over unbounded paths. Specifically, it reorders segments of resolution trees that have a “comb”-like topology, i.e., fully ordered sequences of resolution. This technique only works if the input “comb” is regular, thus the first step is applying the algorithm described in this section.

**Algorithm 3:** RECYCLE-PIVOTS ( $n, RL$ )

```

1: if  $n$  visited then return
2: Mark  $n$  as visited
3: if root then return
4: if  $n$  has more than one child then  $RL := \{\}$ 
5: if  $n.piv \notin RL$  and  $\neg n.piv \notin RL$  then
6:   RECYCLE-PIVOTS( $n.L, RL \cup \{n.piv\}$ )
7:   RECYCLE-PIVOTS( $n.R, RL \cup \{\neg n.piv\}$ )
8: else if  $\neg n.piv \in RL$  then
9:    $n.L := \text{nil};$ 
10:  RECYCLE-PIVOTS( $n.R, RL$ );
11: else
12:   $n.R := \text{nil};$ 
13:  RECYCLE-PIVOTS( $n.L, RL$ );

```



**Fig. 3.** For the input proof presented in drawing [a], RECYCLE-PIVOTS erases the edge (C1,C3) and then calls RECONSTRUCT-PROOF, which results in the graph in drawing [b].  $RL$  denotes the Removable-Literals sets.

## 5 Proofs

Our goal is to prove the following two theorems.

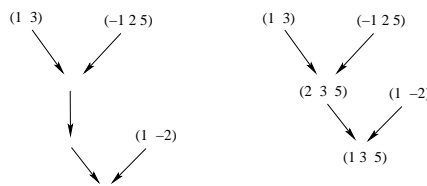
**Theorem 1.** *Let  $P$  be a resolution unsatisfiability proof on the premises  $A$ . Let  $P'$  be the result obtained by running RECYCLE-UNITS and then RECONSTRUCT-PROOF on  $P$ . Then  $P'$  is a valid resolution unsatisfiability proof on the premises  $A'$ , where  $A' \subseteq A$ .*

**Theorem 2.** *Let  $P$  be a resolution unsatisfiability proof on the premises  $A$ . Let  $P'$  be the result obtained by running RECYCLE-PIVOTS and then RECONSTRUCT-PROOF on  $P$ . Then,  $P'$  is a valid resolution unsatisfiability proof on the premises  $A'$ , where  $A' \subseteq A$ .*

The proof of correctness relies on a notion of *e-resolution*, which we will soon define. As we will prove, the graphs produced by RECYCLE-UNITS and RECYCLE-PIVOTS are e-resolution graphs, and RECONSTRUCT-PROOF transforms them back to resolution graphs.

It is convenient to represent resolution proofs as DAGs in which only the roots are labeled with clauses. The clauses labeling the other nodes, including the consequents, can be inferred from the topology of the graph and the roots, and hence are not considered as part of the representation. It is important, in order to understand the proofs that follow, to understand that the consequence of this convention is that *changing the topology of a proof (as done by Alg.1 and Alg.3) implicitly changes its internal labels*. We do not recalculate the labels explicitly after Alg.1 and 3 because this will be done anyway by RECONSTRUCT-PROOF.

As an example, the right drawing of Fig. 4 represents the resolution graph corresponding to the graph on its left, which is based on our convention. Recall that single parents are allowed. The convention in that case is that the child inherits the label from the parent.



**Fig. 4.** The two graphs represent equivalent resolution proofs. The convention we use in this section corresponds to the left figure.

The notion of e-resolution is defined based on this convention, with the difference that each internal node  $n$  with two parents is labeled with a variable

$n.piv$ ,<sup>3</sup> which in our case is the pivot used to infer  $n.C$  at that node in the original proof. For nodes with two parents, instead of Eq. (1), e-resolution uses a more relaxed inference rule, of which (1) is a special case. Using the convention that  $n.piv$  may occur only in  $n.L.C$ , and  $\neg n.piv$  may occur only in  $n.R.C$ , the e-resolution inference rule is:

$$\begin{aligned} n.C &= \text{e-resolution}_{n.piv}(n.L.C, n.R.C) \\ &= (n.L.C \setminus \{n.piv\}) \cup (n.R.C \setminus \{\neg n.piv\}). \end{aligned} \quad (2)$$

For example, for a node  $n$  labeled with a pivot  $x$ , and parents  $n.L.C = (y_1 \vee y_2)$  and  $n.R.C = (\neg x \vee y_3)$  we have  $n.C = (y_1 \vee y_2 \vee y_3)$ .

Every connected DAG with not more than two parents per node that its roots are labeled with clauses and its internal nodes are labeled with a literal, is an e-resolution graph. Hence clearly Alg. 1 and 3 implicitly generate such graphs, because, starting from a resolution graph, they only replace a parent (Alg. 1) or remove edges (Alg. 3). The pivot variables in the original resolution graph label the remaining internal nodes. It is left to prove that the consequent of these graphs is the empty clause, and that RECONSTRUCT-PROOF generates a resolution proof of the empty clause from them.

### 5.1 The RECONSTRUCT-PROOF Algorithm

In the following we denote by  $\leq$  the subsumption relation, e.g.,  $C_1 \leq C_2$  means that  $C_1$  has a subset of the literals of  $C_2$ . By  $A(P)$  we denote the set of premises (roots) of a proof  $P$ .

**Lemma 1.** *Let  $P$  be the e-resolution proof pointed to by the node  $n$ , and let  $P'$  be the proof pointed to by the same node after applying RECONSTRUCT-PROOF( $n$ ). Let  $C$  and  $C'$  denote the clause  $n.C$  in  $P$  and  $P'$ , respectively. Then  $P'$  is a resolution proof of  $C'$ , where  $C' \leq C$  and  $A(P') \subseteq A(P)$ .*

*Proof.* The claim follows by induction on the height of the proof (the length of the longest path from a root to the sink). The base case is when  $n$  points to a root, a case in which the algorithm does not change anything because of line 3, which implies that the lemma holds trivially. For the step, we consider several cases. If  $n$  has a single parent, the algorithm is applied recursively to this parent, which means that it returns a proof that satisfies the lemma by the induction hypothesis. Otherwise,  $n$  has two parents. The proof continues by a straight-forward analysis of the four possible cases for a node  $n$  as described by the table in Fig. 5. In the first case RECONSTRUCT-PROOF applies simple resolution to clauses that satisfy the lemma by the induction hypothesis, and hence clearly satisfies the lemma. The other three cases simply copy either  $n.L.C$  or  $n.R.C$ , which, by the induction hypothesis, point to a legitimate resolution

<sup>3</sup> This is in contrast to a standard resolution graph, in which the pivots can be inferred from the topology and premise.

proof. It is left to show that  $C' \leq C$ . Consider the second case, for example. Let  $C = \text{e-resolution}_{\text{piv}}(C_l, C_r) = (C_l \setminus \{\text{piv}\}) \cup (C_r \setminus \{\neg\text{piv}\})$  for some clauses  $C_l, C_r$ . In that case  $C' = C'_l$ , where  $C'_l \leq C_l$  by the induction hypothesis and furthermore  $\text{piv} \notin C'_l$  by the definition of case 2. Hence  $C' = C'_l \leq C_l \setminus \{\text{piv}\} \leq C$ . The other two cases are proven in a similar way. Finally,  $A(P') \subseteq A(P)$  is implied by the fact that RECONSTRUCT-PROOF only removes nodes.

$n.C$	$n.\text{piv} \in n.L.C$	$\neg n.\text{piv} \in n.R.C$
$\text{Res}(n.L.C, n.R.C)$	yes	yes
$n.L.C$	no	yes
$n.R.C$	yes	no
$n.R.C$ or $n.L.C$	no	no

**Fig. 5.** The four cases discussed in the proof of Lemma 1.

## 5.2 The RECYCLE-UNITS Algorithm

RECYCLE-UNITS is a special case of a more general procedure, which we call **subsumption**. It generalizes RECYCLE-UNITS in the sense that it does not only recycle unit clauses, rather any learnt clause that subsumes other clauses in the proof. Our proof will refer to subsumption, and the correctness of RECYCLE-UNITS is then implied.

Consider two nodes  $p, m$  in the proof  $P$  such that  $p.C \leq m.s.C$  for some side  $s \in \{L, R\}$ . Then, as long as no cycle is produced, we can set  $m.s = p$  so that  $m$  uses the stronger  $p$  instead of its original parent  $m.s$ . In fact, the next definition and lemma show that one can perform more than one such subsumption operation in parallel, and the definition of subsumption can be slightly strengthened by considering the pivot variable. Indeed, RECYCLE-UNITS as listed in Algorithm 1 performs such substitutions with an arbitrary order.

**Definition 2 (e-subsumption).** *Let  $P$  be an e-resolution proof. We say that the node  $p$  e-subsumes the  $k$  nodes  $m_1, \dots, m_k$  with sides  $s_1, \dots, s_k$  (i.e.,  $s_i \in \{L, R\}$ ), if the following conditions hold for every  $i$ : The node  $m_i$  has two parents;  $m_i$  is not an ancestor of  $p$ ; and*

$$p.C \leq \begin{cases} m_i.L.C \cup \{m_i.\text{piv}\} & \text{if } s_i = L \\ m_i.R.C \cup \{\neg m_i.\text{piv}\} & \text{if } s_i = R. \end{cases}$$

**Lemma 2 (Parallel e-subsumption).** *Let  $P$  be an e-resolution unsatisfiability proof, and let  $p, m_1, \dots, m_k, s_1, \dots, s_k$  be as above. Then the proof  $P'$  obtained by setting  $m_i.s_i = p$  for all  $i$  is an e-resolution unsatisfiability proof with  $A(P') \subseteq A(P)$ .*

*Proof.* Let  $P, P'$  be as above. We first argue that  $P'$  contains no cycles. Indeed, assume such a cycle exists. Obviously the cycle must contain some new edge. However, since all the new edges emanate from  $p$ , we can assume that the cycle contains exactly one new edge, say the edge from  $p$  to  $m_1$ . However, the rest of the cycle is a path from  $m_1$  to  $p$ , which is a contradiction to the assumption that  $m_1$  is not an ancestor of  $p$ .

Second, we argue that  $P'$  is an e-resolution unsatisfiability proof. The proof is by induction on the height of  $n$  with the induction claim:  $n.C' \leq n.C$ . Here we let  $n.C, n.C'$  denote the consequent clauses at the node  $n$  when applying the e-resolution inference rule for the proof  $P, P'$  respectively. As the claim trivially holds for root nodes and for nodes with in-degree one, we can restrict our attention to some node  $n$  with two parents. By induction  $n.L.C' \leq n.L.C \cup \{n.piv\}$  and  $n.R.C' \leq n.R.C \cup \{\neg n.piv\}$ .<sup>4</sup> Therefore, by (2) we conclude that  $n.C' \leq n.C$  as claimed.

(Proof of Theorem 1)

*Proof.* The claim follows by applying Lemma 2 once for every iteration starting in line 2 of the RECYCLE-UNITS algorithm, and subsequently applying Lemma 1 once.

The only point that needs some extra elaboration is that the assumptions of Lemma 2 hold each time it is applied. Indeed, the algorithm RECYCLE-UNITS only considers the case of subsumption by a unit clause  $u$ . Furthermore, the algorithm checks for the condition  $u = n.piv$  if  $s_i = L$  and  $u = \neg n.piv$  if  $s_i = R$ , which is strictly stronger than required by the Lemma.

### 5.3 The RECYCLE-PIVOTS Algorithm

The main component we need for proving Theorem 2 is the following Lemma.

**Lemma 3.** *Given a set of literals  $RL$  and a node  $n$  in a resolution proof  $P$ , let  $n.C$  and  $n.C'$  be the clause at  $n$  before and after applying RECYCLE-PIVOTS( $n, RL$ ). Then  $n.C' \leq n.C \cup RL$ .*

Note that specifically when  $n.C$  is the empty clause and  $RL$  is the empty set – the parameters with which we call RECYCLE-PIVOTS – it is implied that  $n.C'$  is the empty clause.

*Proof.* The proof is by induction on the height of  $P$ , i.e., length of a longest path from root to sink. For the base case,  $n$  is a root, and  $n.C' = n.C$  so the lemma holds trivially. Given some non-root node  $n$  we consider three cases:

1. lines 6-7, when ( $n.piv \notin RL$  and  $\neg n.piv \notin RL$ ) or  $n$  has more than one child,
2. lines 9-10, when  $\neg n.piv \in RL$  and  $n$  has one child, and

<sup>4</sup> Note that the stronger claim  $n.L.C' \leq n.L.C$  may not be correct since  $n.piv$  may have been added to  $n.L.C$  if  $n.L$  is one of the  $k$  nodes modified in the transition from  $P$  to  $P'$ .

3. line 12-13, when  $n.piv \in RL$  and  $n$  has one child.

In the first case by the induction hypothesis for the call  $\text{RECYCLE-PIVOTS}(n.L, RL \cup n.piv)$  we have  $n.L.C' \leq n.L.C \cup (RL \cup n.piv)$ . Similarly, for the call  $\text{RECYCLE-PIVOTS}(n.L, RL \cup \neg n.piv)$  we have  $n.R.C' \leq n.R.C \cup (RL \cup \neg n.piv)$ .

$$\begin{aligned} n.C' &= \text{e-resolution}_{piv}(n.L.C', n.R.C') \\ &\leq \text{e-resolution}_{piv}(n.L.C \cup (RL \cup n.piv), n.R.C \cup (RL \cup \neg n.piv)) \\ &= ((n.L.C \cup RL) \setminus n.piv) \cup ((n.R.C \cup RL) \setminus \neg n.piv) . \end{aligned}$$

Since  $n.C$  is the resolution of  $n.L.C$  with  $n.R.C$  by  $n.piv$ , this is equal to

$$n.C \cup RL .$$

Hence  $n.C' \leq n.C \cup RL$  as required.

In the second case by the induction hypothesis  $n.R.C' \leq n.R.C \cup RL$  so:

$$n.C' = n.R.C' \leq n.R.C \cup RL \leq (n.C \cup \{\neg n.piv\}) \cup RL = n.C \cup RL .$$

The last equality is owing to the fact that in the second case  $\neg n.piv \in RL$ .

The proof of the third case is similar.  $\square$

(proof of Theorem 2)

*Proof.* Given some unsatisfiability proof  $P$  with sink  $n$ , the initial call to  $\text{RECYCLE-PIVOTS}$  is  $n$  and  $RL = \emptyset$ . By Lemma 3 we have  $n.C' \leq n.C \cup RL$ . Since both  $RL$  and  $n.C$  are empty sets it follows that  $n.C'$  is empty, i.e. the resulting proof  $P'$  is an unsatisfiability e-resolution proof. The fact that  $A(P') \leq A(P)$  follows from the fact that the algorithm can only remove a single parent from any given node, and hence no new roots are introduced.  $\square$

## 6 Experimental results and conclusions

We ran our reductions on the IBM benchmark suite, which comprises 63 different designs. On each design we ran bounded model checking with a bound  $k$  initially set to 10 and then incremented by 5, up to  $k = 100$  or a bug was found. This gave us 630 unsatisfiable instances. From those we chose only the instances that take 10 seconds or more for  $\text{RUN-TILL-FIX}$ .<sup>5</sup> This left us with 67 proofs. We set the timeout for each run to be 1800 seconds. We used a 64-bit machine with 8 GB memory, and 2x2.4Ghz Opteron dual core.

The results appear in Fig. 6. What can be concluded from them is that the linear reductions we proposed are on one hand fast (especially  $\text{RECYCLE-PIVOTS}$ ),

<sup>5</sup> This creates a bias against run-till-fix, but recall that we are not competing against run-till-fix – we only check whether our methods can be helpful when run-till-fix fails with a short time-out.



Reduction	Time	Roots				Nodes			
		Before	After	Per sec	Ratio	Before	After	per sec	Ratio
RUN-TILL-FIX	8095	1002924	533941	57.9	0.53	11830898	17677419	-722.2	1.49
units	1002.5	1002924	997674	5.2	0.99	11830898	11513195	316.9	0.97
pivots	32.5	1002924	953585	1518.6	0.95	11830898	10464394	42059.2	0.88
units + pivots	1235.8	1002924	949279	43.4	0.95	11830898	10247401	1281.3	0.87

**Fig. 6.** Reduction in proof roots and nodes. Run time is cumulative for 63 unsatisfiable runs. The ‘Per sec’ columns indicate the number of removed roots / nodes per second. The ‘Ratio’ columns indicate the ratio between the number of roots (or nodes) before and after the reduction.

but their effectiveness in reducing roots is small: only  $\approx 5\%$  of the roots are removed, comparing to 47% with RUN-TILL-FIX. When it comes to the size of the proof itself, it turns out that RUN-TILL-FIX increases the number of proof nodes substantially (by 49%) whereas our reductions decrease their size by 13%. The size of the proof can be relevant when computing interpolants, and indeed as future work it is left to check its effectiveness with an interpolation-based model-checker.

To conclude, we showed two techniques for fast preprocessing (perhaps it should be called postprocessing) of resolution proofs. They cannot replace RUN-TILL-FIX if the main goal is the reduction in core regardless of the time it takes, but they can complement it or even replace it in a realm of short time outs. As indicated above, it is more valuable in scenarios in which decreasing the proof size rather than its core is what matters.

## References

1. N. Amla and K. McMillan. Automatic abstraction without counterexamples. In H. Garavel and J. Hatcliff, editors, *TACAS’03*, volume 2619 of *Lect. Notes in Comp. Sci.*, 2003.
2. N. Amla and K. L. McMillan. A hybrid of counterexample-based and proof-based abstraction. In *Formal Methods in Computer-Aided Design, 5th International Conference, FMCAD 2004*, pages 260–274, 2004.
3. O. Bar-Ilan, O. Fuhrmann, S. Hoory, O. Shacham, and O. Strichman. Linear-time reductions of resolution proofs. In H. Chockler and A. Hu, editors, *Haifa Verification Conference (HVC’08)*, volume 5394 of *Lect. Notes in Comp. Sci.*, pages 114 – 128, 2008.
4. R. E. Bryant, D. Kroening, J. Ouaknine, S. A. Seshia, O. Strichman, and B. Brady. Deciding bit-vector arithmetic with abstraction. In O. Grumberg and M. Huth, editors, *13th Intl. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’07)*, volume 4424 of *Lect. Notes in Comp. Sci.*, pages 358–372, March 2007.
5. A. Cimatti, A. Griggio, and R. Sebastiani. A simple and flexible way of computing small unsatisfiable cores in sat modulo theories. In *SAT*, pages 334–339, 2007.

6. N. Dershowitz, Z. Hanna, and A. Nadel. A scalable algorithm for minimal unsatisfiable core extraction. In *SAT*, volume 4121 of *Lecture Notes in Computer Science*, pages 36–41. Springer, 2006.
7. O. Fuhrman and S. Hoory. On extending bounded proofs to inductive proofs. In *Proc. 21<sup>st</sup> Intl. Conference on Computer Aided Verification (CAV'09)*, 2009. (to appear).
8. R. Gershman, M. Koifman, and O. Strichman. Deriving small unsatisfiable cores with dominators. In *Proc. 18<sup>th</sup> Intl. Conference on Computer Aided Verification (CAV'06)*, number 4144 in *Lect. Notes in Comp. Sci.*, pages 109–122, 2006.
9. R. Gershman and O. Strichman. Haifasat: A new robust SAT solver. In S. Ur, E. Bin, and Y. Wolfsthal, editors, *First International Haifa Verification Conference*, volume 3875 of *Lect. Notes in Comp. Sci.*, pages 76 – 89. Springer-Verlag, 2005.
10. É. Grégoire, B. Mazure, and C. Piette. Local-search extraction of muses. *Constraints*, 12(3):325–344, 2007.
11. O. Grumberg, F. Lerda, O. Strichman, and M. Theobald. Proof-guided underapproximation-widening for multi-process systems. In *POPL '05: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 122–131. ACM Press, 2005.
12. J. Huang. Mup: A minimal unsatisfiability prover. In *Proc. of the 10<sup>th</sup> Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 432–437, 2005.
13. D. Kroening, J. Ouaknine, S. Seshia, and O. Strichman. Abstraction-based satisfiability solving of Presburger arithmetic. In R. Alur and D. Peled, editors, *Proc. 16<sup>th</sup> Intl. Conference on Computer Aided Verification (CAV'04)*, number 3114 in *LNCS*, pages 308–320, Boston, MA, July 2004. Springer-Verlag.
14. T. Lengauer and R. E. Tarjan. A fast algorithm for finding dominators in a flow-graph. *ACM Trans. Program. Lang. Syst.*, 1(1):121–141, 1979.
15. I. Lynce and J. Marques-Silva. On computing minimum unsatisfiable cores. In *Proceedings of the International Symposium on Theory and Applications of Satisfiability Testing*, pages 305–310, 2004.
16. K. McMillan. Interpolation and sat-based model checking. In J. Warren A. Hunt and F. Somenzi, editors, *cav03*, *Lect. Notes in Comp. Sci.*, Jul 2003.
17. M. N. Mneimneh, I. Lynce, Z. S. Andraus, J. P. M. Silva, and K. A. Sakallah. A branch-and-bound algorithm for extracting smallest minimal unsatisfiable formulas. In *SAT*, volume 3569 of *Lecture Notes in Computer Science*, pages 467–474. Springer, 2005.
18. Y. Oh, M. N. Mneimneh, Z. S. Andraus, K. A. Sakallah, and I. L. Markov. Amuse: a minimally-unsatisfiable subformula extractor. In *DAC '04*, pages 518–523, 2004.
19. C. H. Papadimitriou and D. Wolfe. The complexity of facets resolved. *J. Comput. Syst. Sci.*, 37(1):2–13, 1988.
20. L. Ryan. Efficient algorithms for clause-learning SAT solvers. Master's thesis, Simon Fraser University, 2004.
21. O. Shtrichman. Pruning techniques for the SAT-based bounded model checking problem. In *proc. of the 11th Conference on Correct Hardware Design and Verification Methods (CHARME'01)*, Edinburgh, Sept. 2001.
22. G. Tseitin. On the complexity of proofs in poropositional logics. In J. Siekmann and G. Wrightson, editors, *Automation of Reasoning: Classical Papers in Computational Logic 1967-1970*, volume 2. Springer-Verlag, 1983. Originally published 1970.
23. A. Urquhart. The complexity of propositional proofs. *Bulletin of Symbolic Logic*, 1(4):425–467, 1995.

24. J. Whittimore, J. Kim, , and K. Sakallah. Satire: a new incremental satisfiability engine. In *In IEEE/ACM Design Automation Conference (DAC)*, 2001.
25. L. Zhang, C. Madigan, M. Moskewicz, and S. Malik. Efficient conflict driven learning in a Boolean satisfiability solver. In *ICCAD*, 2001.
26. L. Zhang and S. Malik. Extracting small unsatisfiable cores from unsatisfiable boolean formulas. In *In Sixth International Conference on Theory and Applications of Satisfiability Testing (SAT2003)*, S. Margherita Ligure, 2003.

## A SAT solvers cannot produce all resolution proofs

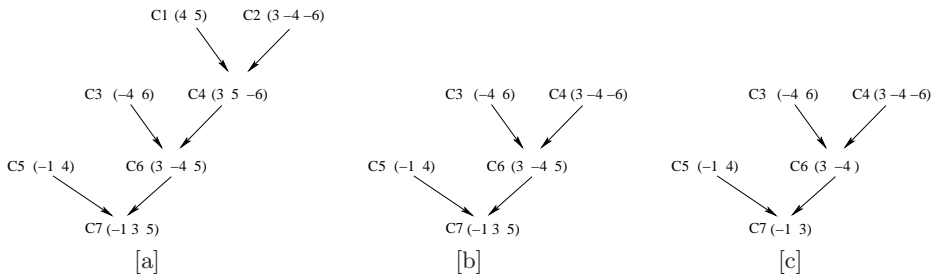
Modern DPLL solvers that restrict learning to asserting clauses (virtually all competitive solvers that we are aware of) are restricted in the resolution proofs that they can generate. In asserting clauses there is exactly one literal from the decision level in which the conflict occurred. Recall that learned clauses are the internal nodes in the resolution.

On the other hand, some pairs of literals can be forced by the constraints to be assigned in the same level. Consider, for example, the clause set

$$(-1 \ 2) \ (1 \ -2) \ (-2 \ 3) \ (2 \ -3)$$

forcing 1, 2 and 3 to be equal. Resolving the first and third clauses yields  $C = (-1 \ 3)$ . But since these variables imply each other, they must have the same decision level, which means that a clause such as  $C$  above cannot be an asserting clause.

The proof reconstruction can result in a proof of a clause such as  $C$ , however. Assume we also have the clauses  $(-1 \ 4)$ ,  $(-4 \ 6)$ ,  $(4 \ 5)$  and  $(3 \ -4 \ -6)$ , and that the resolution graph includes the subgraph depicted in Fig. 7a. The other two graphs show the process of applying RECYCLEPIVOTS, ending with a proof of  $C$ . RECYCLEPIVOTS detects that the pivot variable 4 is used in both  $C7$  and  $C4$ , removes the edge  $(C1, C4)$ , copies  $C2$  to  $C4$ , and continues from there as shown in drawings b and c in the same figure.



**Fig. 7.** Recycling pivots can generate proofs that cannot be generated by a SAT solver.

The restriction to asserting clauses does not change the worst-case complexity of the resulting resolution proofs, however, at least not owing to a scenario as demonstrated above. The reason is that for each clause that cannot be asserting according to the pattern above, there is an asserting clause that subsumes it that is still implied by the formula. In the example above rather than learning  $C$ , an asserting clause can be a unit  $(-1)$  or  $(3)$ . If  $C$  is implied, then so is each of these two clauses.