



**HAL**  
open science

# Enhancing Reasoning with the Extension Rule in CDCL SAT Solvers

Rodrigue Konan Tchinda, Clémentin Tayou Djamegni

► **To cite this version:**

Rodrigue Konan Tchinda, Clémentin Tayou Djamegni. Enhancing Reasoning with the Extension Rule in CDCL SAT Solvers. 2021. hal-02554519v4


**HAL Id: hal-02554519**

**<https://hal.science/hal-02554519v4>**

Preprint submitted on 25 Aug 2021 (v4), last revised 5 Nov 2021 (v5)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



# Enhancing Reasoning with the Extension Rule in CDCL SAT Solvers

Rodrigue Konan Tchinda and Clémentin Tayou Djamegni

Department of Mathematics and Computer Science  
University of Dschang  
Dschang  
Cameroon  
{rodriguekonanktr, dtayou}@gmail.com



**ABSTRACT.** The extension rule first introduced by G. TSEITIN is a simple but powerful rule that, when added to *resolution*, leads to an exponentially stronger proof system known as *extended resolution*(ER). Despite the outstanding theoretical results obtained with ER, its exploitation in practice to improve SAT solvers' efficiency still poses some challenging issues. There have been several attempts in the literature aiming at integrating the extension rule within CDCL SAT solvers but the results are in general not as promising as in theory. An important remark that can be made on these attempts is that most of them focus on reducing the sizes of the proofs using the extended variables introduced in the solver. We adopt in this work a different view. We see extended variables as a means to enhance reasoning in solvers and therefore to give them the ability of reasoning on various semantic aspects of variables. Experiments carried out on the 2018 SAT competition's benchmarks show the use of the extension rule in CDCL SAT solvers to be practically beneficial for both satisfiable and unsatisfiable instances.

**RÉSUMÉ.** La règle d'extension introduite pour la première fois par G. TSEITIN est une règle simple mais puissante qui, ajoutée à la *résolution*, conduit à un système de preuves plus puissant appelé *résolution étendue* (ER). Malgré les résultats théoriques remarquables obtenus avec ER, son exploitation pratique pour améliorer l'efficacité des solveurs SAT pose encore quelques problèmes. Plusieurs tentatives visant à intégrer la règle d'extension aux solveurs CDCL SAT existent dans la littérature, mais les résultats ne sont en général pas aussi prometteurs qu'en théorie. Une remarque importante à faire sur ces tentatives est qu'elles se concentrent pour la plupart sur la réduction de la taille des preuves à l'aide des variables étendues introduites dans le solveur. Nous adoptons dans ce travail un point de vue différent. Nous considérons les variables étendues comme un moyen d'améliorer le raisonnement dans les solveurs et donc de leur donner la capacité de raisonner sur différents aspects sémantiques des variables. Les expérimentations réalisées sur les instances tirées de la compétition SAT 2018 montrent que l'utilisation de la règle d'extension dans les solveurs CDCL est bénéfique aussi bien pour les instances satisfiables que celles insatisfiables.

**KEYWORDS :** SAT, CDCL, extension rule

**MOTS-CLÉS :** SAT, CDCL, règle d'extension



---

## 1. Introduction

The Boolean satisfiability problem (SAT) consists in deciding whether a given propositional logic formula — generally expressed in conjunctive normal form or CNF — admits a model or not. There have been tremendous advances in its resolution during the last two decades and nowadays, SAT solvers are used in industry to solve several challenging problems. The key of this great success lies in a very subtle combination of several features within the so-called CDCL (Conflict-Driven Clause Learning) [12, 20, 21, 6] SAT solvers. The latter include conflicts analysis with clause learning, efficient unit propagation through watched literals, dynamic branching/polarity heuristics and sporadic restarts. SAT has also attracted theoreticians since it was the first problem proved to be NP-complete [4]. Hence, the existence or non-existence of an efficient algorithm for SAT will definitely give the answer to the question  $P =? NP$  which is one of the seven millennium prize problems stated by the *Clay Mathematics Institute* for which an award of 1 Million USD is given to anyone solving one of them.

Despite their current great efficiency, there are still some instances that are out of the reach of current CDCL SAT solvers. The increasing quest of efficiency is achieved by equipping solvers with new techniques and heuristics but the latter are limited from a theoretical perspective. Indeed, CDCL SAT solvers can be formalized as proof systems and it has been shown that the resulting proof system is p-equivalent to general resolution [3, 16] which is known to have an exponential lower bound [8]. This means that CDCL SAT solvers cannot do better than what can be done with general resolution and in particular, that exponential lower bounds known for resolution hold for CDCL SAT solvers as well. To overcome this limitation, a promising research direction is to equip solvers with proof systems that are stronger than resolution. One such proof system is *extended resolution* (ER) which makes use of the *extension rule*.

Several work aiming at integrating the extension rule within CDCL SAT solvers exist in the literature [9, 1]. Most of them focus on the use of the extension rule as a means to reduce the size of the proofs produced by solvers using extended variables introduced in the latter. Seen like that, it might appear that the extension rule is only beneficial for unsatisfiable formulas. However, extended variables introduced in the solver can be seen as a means to increase the level of abstraction and hence enhance reasoning in the solver so that it helps improve the resolution of both satisfiable and unsatisfiable formulas.

We are interested in this work in designing a new integration scheme of the extension rule within CDCL SAT solvers in order to enhance their reasoning capabilities by the use of the extension rule.

The main contributions of this paper are the following: (1) we design a new integration scheme of the extension rule within CDCL SAT solvers called Extended CDCL (ECDCL in short) aiming at enhancing solvers' reasoning. (2) We prove that the substitution of extended literals performed on an asserting clause does not alter its asserting characteristic nor its asserting level. (3) We implemented ECDCL on top of a state-of-the-art SAT solver and conducted an empirical evaluation.

The rest of this paper is organized as follows: Section 2 presents the necessary background for understanding the contribution. In Section 3 we review some related work. Our contribution is given in section 4 and empirically evaluated in Section 5. We finally conclude our work in Section 6 while outlining some future research directions.

---

## 2. Background

A Boolean variable is one that domain is  $\{true, false\}$ . A literal is either a Boolean variable  $x$  or its negation  $\neg x$ . A clause  $c$  is a finite disjunction of literals ( $c = l_1 \vee \dots \vee l_k$ ) and a CNF formula  $\mathcal{F}$  is a finite conjunction of clauses ( $\mathcal{F} = c_1 \wedge \dots \wedge c_n$ ). A CNF formula can be also seen as a set of clauses where each clause is thought of as a set of literals. In this way, we can use set operators on CNF formulas and clauses. An interpretation  $\mathcal{I}$  of a CNF formula  $\mathcal{F}$  is a function that maps each variable in  $\mathcal{F}$  to a truth value in  $\{true, false\}$ . A clause is said to be satisfied under an interpretation  $\mathcal{I}$  if at least one of its literals is satisfied under  $\mathcal{I}$ . A CNF formula is said to be satisfied under an interpretation  $\mathcal{I}$  if all its clauses are satisfied under  $\mathcal{I}$ . A CNF formula  $\mathcal{F}$  is satisfiable if there can be found an interpretation under which it is satisfied; otherwise it is unsatisfiable. Given a CNF formula  $\mathcal{F}$  and a literal  $l$ , we write  $\mathcal{F}_{|l} = \{c \mid c \in \mathcal{F}, \{l, \neg l\} \cap c = \emptyset\} \cup \{c \setminus \{\neg l\} \mid c \in \mathcal{F}, \neg l \in c \text{ and } l \notin c\}$ .  $\mathcal{F}_{|l}$  denotes the simplified formula obtained from  $\mathcal{F}$  by removing all clauses  $c \in \mathcal{F}$  such that  $l \in c$  and  $\neg l$  from clauses containing it. This simplification can be extended to a set of literals  $\{l_1, \dots, l_k\}$ ; thereby  $\mathcal{F}_{|\{l_1, \dots, l_k\}}$  is the formula obtained from  $\mathcal{F}$  by successively applying the previous simplification rule on  $l_1, l_2, \dots$  and  $l_k$  i.e.  $\mathcal{F}_{|\{l_1, \dots, l_k\}} = (\dots(\mathcal{F}_{|l_1})_{|l_2} \dots)_{|l_k}$ . *Unit propagation* is the application of the rule  $\mathcal{F}_{|x}$  for each unit clause  $\{x\} \in \mathcal{F}$  until a clause in  $\mathcal{F}$  is falsified or  $\mathcal{F}$  does not contain unit clauses anymore.

The Boolean satisfiability problem (SAT) consists in deciding whether a given CNF formula is satisfiable or not. The latter definition, which considers only formulas in the CNF representation, is not a restriction since every propositional logic formula can be efficiently translated into an equisatisfiable CNF formula [19].

The most widespread algorithm today for solving SAT is known as CDCL (*Conflict-Driven Clause Learning*) [12, 20, 21, 6]. The principle of CDCL can be summarized as follows: the algorithm performs a sequence of unit propagations until a fixed point is reached (i.e. no further unit propagation can be made) or a conflict is found (i.e. a clause is falsified). If no conflict was found, the algorithm proceeds by making a decision and subsequently increases the decision level. Each assigned variable is associated to a *decision level* and a literal  $l$  is said to be of *level*  $k$  if it is the  $k$ th decision literal or is deduced by unit propagation after setting the  $k$ th decision literal. If a conflict is found, then procedure *analyze* is invoked to examine it in order to produce an *asserting clause* (i.e. a clause that is falsified under the interpretation being constructed and that contains only one literal of the conflicting decision level) as well as the decision level at which the solver must backtrack in order to continue the search. Afterward, the solver learns the asserting clause and backtracks accordingly. From time to time, the algorithm performs restarts which consist in backtracking at *decision level zero* and begin a new search while keeping some information of the previous round (such as learned clauses, variable activities, etc.) which might help speed up the new search.

We formally characterize the state of a CDCL SAT solver by the tuple  $(\mathcal{F}, \Delta, \delta)$  where  $\mathcal{F}$  is the formula being solved,  $\Delta$  the learned clause database and  $\delta$  the partial interpretation being constructed by the solver. We denote the level of a literal  $x$  relatively to a solver state  $S$  by  $level(x)$ . Given an asserting clause  $c$  w.r.t. a state  $S$ , its *asserting literal* is the literal with the highest decision level and its *asserting level* is the second highest decision level of literals in  $c$ . In this paper, the state should be clear from the context when not explicitly specified.

A *propositional proof system* is a polynomial time algorithm  $V$ , such that for every propositional formulas  $\mathcal{F}$ ,  $\mathcal{F}$  is unsatisfiable iff there exists a string  $P$  (a proof of unsatisfiability or refutation of  $\mathcal{F}$ ) such that  $V$  accepts the input  $(\mathcal{F}, P)$ . In the rest of this paper, we omit the word *propositional* and refer to *propositional proof system* simply as *proof system*. Given two proof systems  $V_1$  and  $V_2$ ,  $V_1$  p-simulates  $V_2$  iff there exists a polynomial-time computable function  $f$  such that  $V_2$  accepts  $(\mathcal{F}, P)$  iff  $V_1$  accepts  $(\mathcal{F}, f(P))$ . Two proof systems are p-equivalent if they p-simulate each other. The size of the proof from a given formula is defined as the number of inference steps in the proof. A well-known proof system is *resolution* (also referred to as *general resolution*) which makes use of the resolution rule [17] as inference rule. The strength of *resolution* can be further increased by adding the *extension rule*.

The *extension rule* first introduced by TSEITIN [19] allows the use of literals as abbreviation for longer formulas. Concretely, let  $\mathcal{F}$  be a CNF formula,  $\{x, l_1, l_2\}$  be a set of literals such that neither  $x$  nor  $\neg x$  appears in  $\mathcal{F}$ . The extension rule allows to introduce definitions of the form  $x \leftrightarrow l_1 \vee l_2$  by adding the clauses  $\neg x \vee l_1 \vee l_2$ ;  $x \vee \neg l_1$  and  $x \vee \neg l_2$  to  $\mathcal{F}$ . This rule when added to *resolution*, turns it to an exponentially stronger proof system known as *extended resolution* (ER). A typical example of formulas that are hard for resolution are pigeonhole formulas which do not admit any short (i.e. polynomial size) resolution proof [8]. However, short proofs of pigeonhole formulas exist when using the extension rule [5]. The challenge when using the extension rule is to determine which variables to choose for extension so as to produce short proofs. Even with the right variable choices, the resolution steps that should be performed to achieve this goal still constitute an important issue.

---

### 3. Related Work

There have been several work attempting to integrate the extension rule within CDCL SAT solvers. AUDEMARD *et al.* [1] argued that significant advances in SAT solving must come from implementation of stronger proof systems since exponential lower bounds are known for resolution [8, 13]. They used a restriction of ER called *Local Extended Resolution* (LER) by introducing the extension  $z \leftrightarrow l_1 \vee l_2$  if there exist previously derived clauses in the form  $\neg l_1 \vee \alpha$  and  $\neg l_2 \vee \beta$  where  $\alpha$  and  $\beta$  are disjunction of literals such that  $l \in \alpha \Rightarrow \neg l \notin \beta$ . A clear limitation of this approach is that it uses clauses of particular form that might seldom appear in the set of derived clauses. In addition, looking for such clauses can be difficult and costly. For the latter reasons, the authors in their implementation restricted this lookup to a small window of recent clauses, only looking for those of the form  $\neg l_1 \vee \alpha$  and  $\neg l_2 \vee \alpha$ . Some implementation-level optimizations of LER are proposed in [11].

HUANG [9] proposed Extended Clause Learning (ECL), a general scheme which is a modification of the CDCL algorithm where decisions to use the extension rule might be made (guided by a heuristic) when the number of assigned literals is greater than 2. Besides ECL, they proposed a concrete heuristic where the extension rule was used after learning clauses  $\gamma$  of size greater than 2. Concretely, if the decision to make an extension is taken, then  $\gamma$  is split into  $\alpha \vee \beta$  such that  $|\alpha| \geq 2$  and  $|\beta| > 0$  and the solver learns the clauses  $x \vee \beta$ ,  $x \leftrightarrow \alpha$  where  $x$  is a fresh variable. A restart is performed after each extension introduced in the solver. The drawback of this is that it alters the restart strategy of the solver. Hence, if the heuristic used to decide the time to make extensions is not well designed, it might compromise the completeness of the CDCL SAT solver. For instance,

if we decide to make an extension after each conflict, the solver will never reach more than one conflict and the search will hardly progress in this situation.

In [10], JABBOUR *et al.* proposed a method that mimics the principle behind extended resolution by detecting hidden Boolean functions introduced in the CNF during the encoding phase [15, 7] and by using them to shorten learned clauses through substitution. This approach however does not use fresh variables at all and substitution is restricted to only literals which are the input arguments of a detected Boolean function.

---

## 4. Extended CDCL

### 4.1. Motivation

The extension rule in combination with resolution has been theoretically shown to be useful for shortening the proof size of unsatisfiable formulas. The work mentioned in the literature try to reproduce this result in practice but the outcome of most of them turns to be limited as it seldom matches the expectations. When looking at the extension rule as a means to reduce the size of the proof, it might seem that it will be useful only for unsatisfiable formulas. We adopt here a different view. Extensions are introduced within a solver in order to increase its reasoning capabilities with the ultimate goal of enhancing solving times. When solving a CNF formula, current CDCL SAT solvers proceed by assuming a selected variable to be *true* or *false* and by evaluating the consequences of this assumption on the formula being solved. Proceeding this way limits reasoning to a single semantic aspect of variables, notably their truth values. We want the solver to be able to carry out reasoning on other semantic aspects of formulas' variables. That is, we want the solver in addition, to make other types of assumptions such as assuming that two or more variables are equivalent, simultaneously *true* or *false*, one variable implies the other etc. To achieve this without modifying the way solvers proceed, we are going to use extensions to encapsulate these semantic aspects. Hence, the algorithm of the solver will not change since it will still continue to carry out reasoning as usual; that is, assigning *true/false* values to variables. The difference however will be that when this reasoning is performed on an extended variable, it will denote other semantic aspects. For instance, suppose the extension  $x \leftrightarrow l_1 \Leftrightarrow l_2$  has been made in the solver. When the solver picks the extended variable  $x$  and assigns it value *true*, this means that it is assuming  $l_1$  and  $l_2$  to be equivalent. Hence, after setting  $x$  to *true*, anytime in the search where one of the variables in  $\{l_1, l_2\}$  will be given a value, then the other will also be given the same value via unit propagation. In this way, we could expect an improvement of reasoning in the solver.

### 4.2. Description of the integration scheme

In order to integrate the extension rule in the CDCL framework, we should answer the following questions: which variables should be chosen for extension and when should we perform these extensions?

We propose to use extensions at restarts and to choose for each extension, two variables from two different decision levels (in our implementation, we chose the most active decision variables i.e. the ones with the highest VSIDS scores [14]). At this level, extensions can be performed using any binary connective; for instance  $x \leftrightarrow l_1 \vee l_2$ ,  $x \leftrightarrow l_1 \wedge l_2$ ,

$x \leftrightarrow l_1 \Rightarrow l_2$ ,  $x \leftrightarrow l_1 \Leftrightarrow l_2$  etc. The solver can then make assumptions on the extended variables resulting in an implicit meaning on variables on which extension is performed (for instance, they are/are not equivalent, one implies/does not imply the other, they are both *false/true*, etc.) and evaluates the consequences on the other variables of the formula being solved. By choosing variables of different decision levels instead of any two variables, we aim to give the solver the ability to carry out reasoning on variables that apparently seem not to be dependent and avoid some useless extensions such as extensions where the value of one literal is already known or extensions where the extended variable is immediately forced to a given value. We further provide a substitution mechanism that can be cheaply performed in order to favor the use of extended variables within the solver. This substitution consists in replacing any pair  $\{l_1, l_2\}$  of literals in a clause  $c$  by the literal  $l$  provided that the extension  $l \leftrightarrow l_1 \vee l_2$  has been made in the solver. This substitution can be seen as an application of hyper-resolution [2] involving the original clause and some binary clauses encoding the extensions. It has as effect the shortening of the clause as well as the increase of its propagation power. In fact, as illustrated in [1], if we consider the clause  $c = l_1 \vee l_2 \vee \alpha$  and the extension  $l \leftrightarrow l_1 \vee l_2$ , then the clause  $c' = l \vee \alpha$  obtained from  $c$  by replacing  $l_1 \vee l_2$  with  $l$  unlike the clause  $c$  itself will become unit once all literals in  $\alpha$  are set to *false*. The substitution mechanism is performed on each asserting clause derived after conflict analysis and Proposition 4.1 ensures that after substitution, the resulting clause will remain asserting and the asserting level will be kept.

**Proposition 4.1.** *Let  $S = (\mathcal{F}, \Delta, \delta)$  be the state of a solver and  $c = a \vee \alpha$  an asserting clause w.r.t.  $S$  where  $a$  is the asserting literal. Let  $x \leftrightarrow l_1 \vee l_2$  be an extension such that  $\{l_1, l_2\} \subseteq \alpha$ . Then, the clause  $c' = a \vee x \vee \beta$  where  $\beta = \alpha \setminus \{l_1, l_2\}$  is asserting w.r.t.  $S$ . Furthermore, the asserting levels of  $c$  and  $c'$  are identical.*

The proof of Proposition 4.1 uses the following lemma:

**Lemma 4.1.** *Let  $x \leftrightarrow l_1 \vee l_2$  be an extension introduced in a solver. If literals  $x, l_1$  and  $l_2$  are all assigned, with  $x$  set to *false*, then the decision level of  $x$  is the maximum decision level of  $l_1$  and  $l_2$ .*

*Proof.* Since  $x = \text{false}$  and  $x \leftrightarrow l_1 \vee l_2$ , then  $l_1$  and  $l_2$  are set to *false* as well. If  $x$  is first set to *false* by the solver or set to *false* after setting either of  $l_1$  or  $l_2$  to *false* (no matter the decision level), then the values of  $l_1$  and/or  $l_2$  will immediately be deduced by unit propagation through the clauses  $x \vee \neg l_1$  and  $x \vee \neg l_2$ . In this case, the decision levels of  $x$  and  $l_1$  or  $x$  and  $l_2$  will be identical. If  $l_1$  and  $l_2$  are first set to *false*, then  $x = \text{false}$  will be inferred by unit propagation via the clause  $\neg x \vee l_1 \vee l_2$ . This occurs as soon as the last literal of  $\{l_1, l_2\}$  is assigned. In either of the previous cases, the level of  $x$  is the same as that of the most recently assigned literal of  $\{l_1, l_2\}$ , that is  $level(x)$  is the maximum of  $level(l_1)$  and  $level(l_2)$ .  $\square$

*Proof of Prop. 4.1.*  $c = a \vee \alpha$  is an asserting clause, hence all its literals are *false*. Since  $x \leftrightarrow l_1 \vee l_2$  is an extension introduced in the solver and  $\{l_1, l_2\} \subseteq c$ , then  $x$  is *false* as well. By Lemma 4.1,  $level(x) = \max(level(l_1), level(l_2))$ . In addition,  $level(l_1) < level(a)$  and  $level(l_2) < level(a)$ , hence,  $level(x) < level(a)$  which means that  $a$  still has the highest decision level in  $c' = a \vee x \vee \beta$  where  $\beta = \alpha \setminus \{l_1, l_2\}$ .  $c'$  therefore remains asserting. Furthermore, the second highest decision level of literals in  $c'$  remains the same as in  $c$  since  $\max\{level(y), y \in \alpha\} = \max\{level(y), y \in (\beta \cup \{x\})\}$ .  $\square$

---

**Algorithm 1: Extended CDCL**

---

**Input:** A CNF formula  $\mathcal{F}$ **Result:** SAT or UNSAT

```
1 begin
2    $dl \leftarrow 0; \Delta \leftarrow \emptyset;$ 
3   while true do
4      $conf \leftarrow \text{unitPropagation}(\mathcal{F} \cup \Delta);$ 
5     if  $conf \neq \text{null}$  then
6       if  $dl = 0$  then return UNSAT;
7        $(c, btLevel) \leftarrow \text{analyze}(conf);$  /* return an asserting clause and
8         the backtracking level */
9        $c \leftarrow \text{substituteExtendedLits}(c);$ 
10       $\Delta \leftarrow \Delta \cup \{c\};$  /* learn clause c */
11      backtrack to  $btLevel;$ 
12    else
13      if all variables are assigned then return SAT;
14      if time to restart then
15        if extension needed and not reached max number of extensions then
16          let  $\delta$  be the current interpretation;
17          choose a fresh variable  $x$  and a not yet extended pair  $\{l_1, l_2\}$  from  $\delta$ 
18          such that  $level(l_1) > 0, level(l_2) > 0, level(l_1) \neq level(l_2);$ 
19           $\mathcal{F} \leftarrow \mathcal{F} \cup \text{clauses}(x \leftrightarrow l_1 \circ l_2);$  /* encode the extension as
20          clauses and add to the formula */
21        restart();
22      pick an unassigned variable and assign it a value;
23       $dl \leftarrow dl + 1;$ 
```

---

---

**Algorithm 2: substituteExtendedLits**

---

**Input:** An asserting clause  $c$ **Result:** A clause

```
1 begin
2    $a \leftarrow$  the asserting literal of  $c;$ 
3    $c' \leftarrow \{a\};$ 
4    $c \leftarrow c \setminus \{a\};$ 
5   while  $\exists \{l_1, l_2\} \subseteq c$  such that the extension  $l \leftrightarrow l_1 \vee l_2$  is present in the solver do
6      $c \leftarrow c \setminus \{l_1, l_2\};$ 
7      $c' \leftarrow c' \cup \{l\};$ 
8    $c' \leftarrow c' \cup c;$ 
9   return  $c';$ 
```

---



The resulting scheme that we call *Extended CDCL* (ECDL in short) is described in Algorithm 1. All in this algorithm are as in CDCL except that we introduce extensions at restarts (lines 14–17) and substitution of literals with extended literals for asserting clauses derived from conflicts (line 8). At line 17, the extension operator  $\circ$  can be any binary connective and in this paper we take  $\circ \in \{\vee, \wedge, \Rightarrow, \Leftrightarrow\}$ . Note that an extension is not systematically added at each restart but only when the solver finds it necessary (through a heuristic) and when the maximum number of extensions (introduced as a parameter of the algorithm) is not yet reached. Hence, by *extension needed* we mean the moment the solver decides to make an extension typically through a heuristic. Function *substituteExtendedLits* which carries out substitution is described in Algorithm 2. In order to perform substitution, all extensions are converted so that they use the connective  $\vee$ . Hence, the extensions  $l \leftrightarrow l_1 \wedge l_2$  and  $l \leftrightarrow l_1 \Rightarrow l_2$  are respectively converted to  $\neg l \leftrightarrow \neg l_1 \vee \neg l_2$  and  $l \leftrightarrow \neg l_1 \vee l_2$ . As far as the extension  $l \leftrightarrow l_1 \Leftrightarrow l_2$  is concerned, we use two auxiliary variables  $\{l', l''\}$  for its conversion which lead to the extensions  $\neg l \leftrightarrow \neg l' \vee \neg l''$ ,  $l' \leftrightarrow \neg l_1 \vee l_2$  and  $l'' \leftrightarrow l_1 \vee \neg l_2$ . It is worth mentioning that substitution here is not performed on unary or binary clauses since it requires that the clause contains at least two literals in addition to the asserting literal. Notice that function *substituteExtendedLits* will return the clause as is if called with an unary or a binary clause or even when no extension has yet been introduced in the solver. Another thing to point out in Algorithm 1 is that extensions are made with only two literals. This does not mean that reasoning in solvers is limited to at most two literals at a time. In fact, in Algorithm 1, extensions can be made using other extended literals as well. Hence, extensions with more literals such as  $l \leftrightarrow l_1 \circ l_2 \circ \dots \circ l_n$  can be represented by a sequence of two-literal extensions. The drawback here is that several fresh variables need to be introduced in the solver for that.

Unlike Extended Clause Learning (ECL) which performs a restart after each extension made after conflicts analysis, ECDCL does not alter the restart policy of the CDCL solver and hence eliminates the completeness issue related to an uncontrolled restart strategy. Furthermore, since we limit the number of extensions that can be used, ECDCL remains complete.

---

## 5. Experimental Results

In order to evaluate ECDCL, we conducted a first stage of experiments on the 400 application benchmarks<sup>1</sup> drawn from the 2018 SAT Competition<sup>2</sup>. For these experiments, we implemented ECDCL on top of the state-of-the-art CDCL SAT solver *Glucose-3.0*<sup>3</sup>. *Glucose-3.0* was chosen because it is one of the most used today as a base for many CDCL SAT solvers. We distinguished several different versions obtained by varying the type of extension and the maximum number of extended variables allowed in the solver: We designated by *Glucose-3.0\_exT\_◦\_maxExtVars\_K* the version of our solver where the extension operator is  $\circ$  with  $\circ \in \{\vee, \wedge, \Rightarrow, \Leftrightarrow\}$  and the maximum number of extended variables allowed set to  $K \in \{100, 500, 1000\}$ . In our implementations, we used a very straightforward heuristic to decide when the extensions are made: it consists of making an extension at every restart performed when the current interpretation contains at least two decision levels, starting from the beginning of the search until the maximum number of

---

1. <http://sat2018.forsyte.tuwien.ac.at/benchmarks/Main.zip>

2. <http://www.satcompetition.org/>

3. <https://www.labri.fr/perso/lSimon/downloads/software/glucose-3.0.tgz>

allowed extensions is reached. This heuristic deserves a further study but we leave it for future investigations. When performing an extension  $l \leftrightarrow l_1 \circ l_2$ , we set the VSIDS score of  $l$  to the sum of the scores of  $l_1$  and  $l_2$  so that when making a decision in the solver, the extended variable  $l$  is prioritized over  $l_1$  and  $l_2$ . It is worth mentioning that in our modified solvers, all the other parameters of *Glucose-3.0* were left unchanged. The base solver, referred to as *Glucose-3.0\_default* was used with its default configuration. Preprocessing was enabled for all versions of *Glucose-3.0* used in the experiments including the default one.

All our experiments were carried out on the StarExec<sup>4</sup> [18] cluster infrastructure running Red Hat Enterprise Linux Server version 7.2 (Maipo). Each node of this infrastructure has 128 GB of memory and two Intel processors with 4 cores (2.4 GHz) each. For each solver, we set a time limit of 1800 seconds and a memory limit of 24GB for the resolution of each benchmark.

The results obtained at the end of these experiments are summarized in Table 1. The table indicates for each extension type and maximum number of extended variables allowed, the number of satisfiable instances solved (#S), the number of unsatisfiable instances (#U), the total number of instances solved (#T) as well as the PAR-2 score. The PAR-2 score is defined as the *sum of all runtimes for solved instances + 2 × timeout for unsolved instances*. The line *default* in the table represents the performance of *Glucose-3.0\_default* which are repeated for every limit on the number of extended variables to ease the comparisons. We see in this table that all our solvers outperformed the original solver on the total number of instances solved (up to 14 additional instances solved for our best performing solver). The second remark is that our solvers are very efficient on satisfiable instances (our best performing solver on satisfiable instances solved 12 more satisfiable instances than the original solver) meaning that the extension rule greatly helped improve the resolution of satisfiable instances.

	Number of extended vars ≤ 100				Number of extended vars ≤ 500				Number of extended vars ≤ 1000			
	#S	#U	#T	PAR-2	#S	#U	#T	PAR-2	#S	#U	#T	PAR-2
exT_∨	<b>83</b>	<b>71</b>	<b>154</b>	<b>2436.36</b>	<b>82</b>	69	<b>151</b>	<b>2454.71</b>	76	<b>69</b>	145	2499.94
exT_∧	73	<b>71</b>	144	2515.47	78	<b>70</b>	148	2472.11	75	68	143	2518.84
exT_⇒	80	70	150	2455.51	79	<b>70</b>	149	2484.04	75	67	142	2521.33
exT_⇔	74	67	141	2531.56	75	68	143	2505.41	<b>84</b>	68	<b>152</b>	<b>2456.92</b>
default	72	68	140	2526.22	72	68	140	2526.22	72	68	140	2526.22

**Table 1.** Performance of *Glucose-3.0* and *Glucose-3.0\_exT\_◦\_maxExtVars\_K* on the 2018 SAT Competition’s benchmarks when substitution is carried out on all learned clauses

This table also shows that the number of solved instances generally decreases as the limit on the number of extended variables increases when substitution is carried out on all learned clauses. There is however an exception for the extension operator  $\Leftrightarrow$  where the performance increased with the limit on the number of extended variables.

When considering another performance metric, notably the average PAR-2 score<sup>5</sup> currently used for ranking solvers at SAT competitions, we also notice that all our solvers outperformed the original solver except for *Glucose-3.0\_exT\_⇔\_maxExtVars\_100*.

Fig. 1 shows the cactus plots of all the solvers on the 400 application benchmarks of the 2018 SAT competition. It clearly appears on these plots that the use of the exten-

4. <https://www.starexec.org/>

5. Solvers with the lowest scores are the best performing




Figure 1. Cactus plots, for the extension type, 1 =  $\vee$ , 2 =  $\wedge$ , 3 =  $\Rightarrow$ , 4 =  $\Leftrightarrow$




Figure 2. Scatter plot *Glucose-3.0\_exT\_1\_maxExtVars\_100* vs *Glucose-3.0\_default*

sion rule was beneficial for our solvers since they still outperformed the original solver *Glucose-3.0\_default* for various solving time limits in term of number of solved instances. Fig. 2 shows the scatter plot comparing the original solver *Glucose-3.0\_default* and our best performing integration namely *Glucose-3.0\_exT\_1\_maxExtVars\_100*. It appears on this plot that the run time of many instances has been significantly improved: this is no-

ticeable through the presence of many points in the upper triangle that are far from the diagonal. However there are still a number of instances where the original solver is faster.

We noticed during our experiments that the substitution operations were very time consuming despite our optimizations. To address this issue, we decided to restrict the substitutions to the learned clauses that were most likely to remain in solvers for a longer time. Thus, we restricted the substitutions to clauses whose LBD values at the time of their learning were less than or equal to an empirically chosen threshold set to 20. With this new optimization, we conducted a second phase of experiments, this time with an extended set of benchmarks including those of the SAT competitions 2018<sup>6</sup> and 2020<sup>7</sup>. Moreover, we integrated ECDCL with this optimization in one of the current top performing SAT solvers, *MapleLCMDistChronoBT*<sup>8</sup>, in order to evaluate the impact on its performance. The versions of the solvers used in these experiments were named in the same way as before, i.e. *MapleLCMDistChronoBT\_exT\_◦\_maxExtVars\_K* is the version of *MapleLCMDistChronoBT* integrating ECDCL where the extension operator is ◦ with ◦ ∈ {∨, ∧, ⇒, ⇔} and the maximum number of extended variables allowed set to  $K ∈ \{100, 500, 1000\}$ . The default version of *MapleLCMDistChronoBT* as far as it is concerned was named *MapleLCMDistChronoBT\_default*. For all these versions of *MapleLCMDistChronoBT* including the default one used in our experiments, we set the --chrono option to -1 and enabled preprocessing. The experiments were conducted under the same conditions as described in the first phase and the results obtained are reported in Table 2 and Table 3. Each of these tables has two parts respectively for the benchmarks of the SAT Competition 2018 (SC2018) and for the SAT Competition 2020 (SC2020).

		Number of extended vars ≤ 100				Number of extended vars ≤ 500				Number of extended vars ≤ 1000			
		#S	#U	#T	PAR-2	#S	#U	#T	PAR-2	#S	#U	#T	PAR-2
SC2018	exT_∨	<b>78</b>	<b>72</b>	<b>150</b>	<b>2465.88</b>	80	<b>74</b>	<b>154</b>	<b>2426.68</b>	80	<b>74</b>	<b>154</b>	<b>2427.34</b>
	exT_∧	72	71	143	2494.51	<b>83</b>	71	154	2433.98	<b>81</b>	70	151	2448.30
	exT_⇒	75	69	144	2493.24	79	73	152	2455.69	80	69	149	2466.28
	exT_⇔	71	71	142	2498.26	81	71	152	2443.35	73	72	145	2490.11
	default	72	68	140	2526.22	72	68	140	2526.22	72	68	140	2526.22
SC2020	exT_∨	<b>50</b>	61	111	2778.94	<b>54</b>	64	<b>118</b>	<b>2732.46</b>	50	62	112	2765.89
	exT_∧	49	67	<b>116</b>	<b>2749.18</b>	50	60	110	2788.62	<b>51</b>	62	113	2757.20
	exT_⇒	45	<b>68</b>	113	2761.84	52	<b>66</b>	<b>118</b>	2747.27	43	63	106	2821.51
	exT_⇔	49	67	<b>116</b>	2750.60	47	64	111	2773.05	58	63	<b>121</b>	<b>2706.09</b>
	default	47	65	112	2773.01	47	65	112	2773.01	47	<b>65</b>	112	2773.01

**Table 2.** Performance of Glucose-3.0 and Glucose-3.0\_exT\_◦\_maxExtVars\_K on the 2018 and 2020 SAT Competitions’ benchmarks when substitutions are carried out only on learned clauses having LBD lower than or equal to 20

We can see in Table 2 that for the SAT Competition 2018 benchmarks, the PAR-2 scores as well as the number of solved unsatisfiable instances have been significantly improved compared to those of the first phase experiments in Table 1. This clearly indicates that it is more favorable to perform substitutions on clauses with low LBD scores. In Tables 2 and 3 it appears that the most notable improvements still reside at the level of satisfiable instances both for those of the 2018 and the 2020 SAT Competitions. We see for instance in Table 2 that the solver *Glucose-3.0\_exT\_∧\_maxExtVars\_500* (resp. *Glucose-3.0\_exT\_∨\_maxExtVars\_500*) solved 11 (resp. 7) satisfiable instances

6. <http://sat2018.forsyte.tuwien.ac.at/benchmarks/Main.zip>

7. <https://satcompetition.github.io/2020/downloads/sc2020-main.uri>

8. <http://sat2018.forsyte.tuwien.ac.at/solvers/>

		Number of extended vars $\leq 100$				Number of extended vars $\leq 500$				Number of extended vars $\leq 1000$			
		#S	#U	#T	PAR-2	#S	#U	#T	PAR-2	#S	#U	#T	PAR-2
SC2018	exT_ $\vee$	114	<b>86</b>	200	2025.28	111	84	195	2045.07	113	84	197	2042.22
	exT_ $\wedge$	108	<b>86</b>	194	2066.72	<b>119</b>	<b>88</b>	<b>207</b>	<b>1953.83</b>	<b>116</b>	87	<b>203</b>	1986.98
	exT_ $\Rightarrow$	<b>118</b>	<b>86</b>	<b>204</b>	<b>1983.04</b>	111	<b>88</b>	199	2038.26	110	<b>88</b>	198	2034.79
	exT_ $\Leftrightarrow$	<b>118</b>	84	202	1993.29	117	85	202	1987.66	112	86	198	2031.29
	default	115	<b>86</b>	201	1986.70	115	86	201	1986.70	115	86	201	<b>1986.70</b>
SC2020	exT_ $\vee$	<b>63</b>	74	<b>137</b>	<b>2585.92</b>	57	69	126	2656.81	60	70	130	2636.93
	exT_ $\wedge$	60	73	133	2616.19	54	70	124	2670.61	58	70	128	2648.66
	exT_ $\Rightarrow$	62	71	133	2608.55	61	72	133	2601.37	59	72	131	2617.42
	exT_ $\Leftrightarrow$	59	71	130	2621.29	<b>62</b>	72	<b>134</b>	2610.89	<b>62</b>	72	<b>134</b>	2599.70
	default	59	<b>75</b>	134	2592.51	59	<b>75</b>	<b>134</b>	<b>2592.51</b>	59	<b>75</b>	<b>134</b>	<b>2592.51</b>

**Table 3.** Performance of *MapleLCMDistChronoBT* and *MapleLCMDistChronoBT\_exT\_ $\circ$ \_maxExtVars\_K* on the 2018 and 2020 SAT Competitions’ benchmarks when substitutions are carried out only on learned clauses having LBD lower than or equal to 20

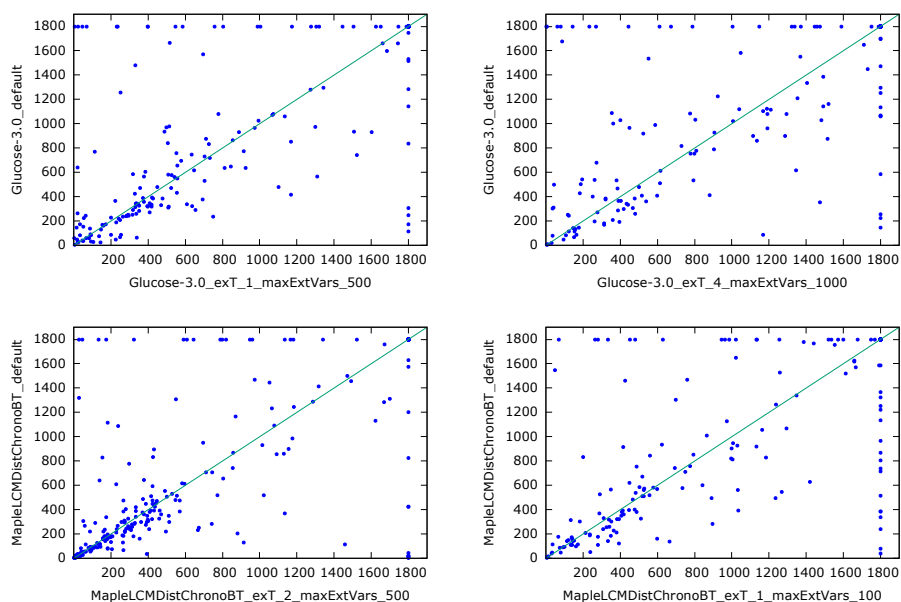
of the SC2018 (resp. SC2020) more than the base solver *Glucose-3.0\_default* and in Table 3 that *MapleLCMDistChronoBT\_exT\_ $\wedge$ \_maxExtVars\_500* (resp. *MapleLCMDistChronoBT\_exT\_ $\vee$ \_maxExtVars\_100*) solved 4 (resp. 4) more satisfiable instances of the SC2018 (resp. SC2020) than the base solver *MapleLCMDistChronoBT\_default*. Even if we can see clear improvements on the unsatisfiable instances of the SAT Competition 2018 in Table 2, we still notice a decrease in performance at the level of those of the SAT Competition 2020. This same observation is made on Table 3 for the solvers *MapleLCMDistChronoBT\_exT\_ $\circ$ \_maxExtVars\_K*. From Table 3, we can see that the gain in performance for the solvers *MapleLCMDistChronoBT\_exT\_ $\circ$ \_maxExtVars\_K* is not as important as that for *Glucose-3.0\_exT\_ $\circ$ \_maxExtVars\_K* solvers and even sometimes we observe a decrease in performance. This is justifiable because the base solver *MapleLCMDistChronoBT* is a highly optimized and therefore more difficult to improve.

The scatter plots in Fig 3 compare the performance of some of our best integrations with the base solvers. We see on these plots in general that many instances were solved with run times almost similar to those of the base solvers. This is noticeable through the presence of several points near the diagonals. However, with the presence of many points above the diagonals, we note that the solving times of many instances have been improved. It is also important to mention that for all these plots, there are many points at the top and right borders indicating the instances solved by one of the solvers and not by the other.

## 6. Conclusion and Future Work

We presented in this paper a new integration scheme of the extension rule within CDCL called *extended CDCL* (ECDCL) which, unlike the state-of-the-art integrations, uses extensions to enhance reasoning in solvers. ECDCL also allows to substitute literals in asserting clauses with extended literals while preserving their asserting nature as well as the asserting levels. We showed experimentally that the extension rules helps improve the resolution of satisfiable instances and in some cases unsatisfiable instances as well.

This work opens doors to many other investigations. For instance, an interesting research direction is to determine if our algorithm seen as a proof system is theoretically strictly more powerful than general resolution. Additionally, it might be interesting to see whether the proof system implemented in our algorithm p-simulates ER. Some extensions



**Figure 3.** Scatter plots comparing some of our best integrations with the base solvers. At the left (resp. right) we have the comparison using the SC2018 (resp. SC2020) benchmarks

introduced in the solver might already exist in the original formula in the form of Boolean function [10]. So it would be interesting to detect and use them instead of making new extensions with fresh variables which unnecessarily increase the formula size.

---

## Acknowledgments

The authors wish to thank the anonymous reviewers for their insightful comments and valuable suggestions on a previous version of this paper.

---

## References

- [1] Gilles Audemard, George Katsirelos, and Laurent Simon. A restriction of extended resolution for clause learning sat solvers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24, pages 15–20, 2010.
- [2] Fahiem Bacchus. Enhancing davis putnam with extended binary clause reasoning. *AAAI/IAAI*, 2002:613–619, 2002.
- [3] Paul Beame, Henry Kautz, and Ashish Sabharwal. Towards understanding and harnessing the potential of clause learning. *Journal of Artificial Intelligence Research*, 22:319–351, 2004.
- [4] Stephen A Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158. ACM, 1971.

- [5] Stephen A. Cook. A short proof of the pigeon hole principle using extended resolution. *SIGACT News*, 8(4):28–32, October 1976.
- [6] Carla P Gomes, Bart Selman, Henry Kautz, et al. Boosting combinatorial search through randomization. *AAAI/IAAI*, 98:431–437, 1998.
- [7] Éric Grégoire, Richard Ostrowski, Bertrand Mazure, and Lakhdar Saïs. Automatic extraction of functional dependencies. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 122–132. Springer, 2004.
- [8] Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297 – 308, 1985. Third Conference on Foundations of Software Technology and Theoretical Computer Science.
- [9] Jinbo Huang. Extended clause learning. *Artificial Intelligence*, 174(15):1277–1284, October 2010.
- [10] Saïd Jabbour, Jerry Lonlac, and Lakhdar Saïs. Extending resolution by dynamic substitution of boolean functions. In *24th International Conference on Tools with Artificial Intelligence*, volume 1, pages 1029–1034. IEEE, 2012.
- [11] Norbert Manthey. Extended resolution in modern sat solving. In *Joint Automated Reasoning Workshop and Deduktionstreffen*, pages 26–27, 2014.
- [12] JP Marques-Silva and KA Sakallah. Grasp—a new search algorithm for satisfiability. *iccad*. pages 220–227. IEEE Computer Society Press, 1996.
- [13] Mladen Mikša and Jakob Nordström. Long proofs of (seemingly) simple formulas. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 121–137. Springer, 2014.
- [14] Matthew W Moskewicz, Conor F Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. Chaff: Engineering an efficient sat solver. In *Proceedings of the 38th annual Design Automation Conference*, pages 530–535. ACM, 2001.
- [15] Richard Ostrowski, Éric Grégoire, Bertrand Mazure, and Lakhdar Saïs. Recovering and exploiting structural knowledge from cnf formulas. In Pascal Van Hentenryck, editor, *Principles and Practice of Constraint Programming*, pages 185–199, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [16] Knot Pipatsrisawat and Adnan Darwiche. On the power of clause-learning sat solvers as resolution engines. *Artificial Intelligence*, 175(2):512–525, 2011.
- [17] John Alan Robinson et al. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965.
- [18] Aaron Stump, Geoff Sutcliffe, and Cesare Tinelli. Starexec: a cross-community infrastructure for logic solving. In *International Joint Conference on Automated Reasoning*, pages 367–373. Springer, 2014.
- [19] G Tseitin. On the complexity of derivation in propositional calculus. *Studies in Constrained Mathematics and Mathematical Logic*, 1968.
- [20] Hantao Zhang. Sato: An efficient prepositional prover. In *Automated Deduction-CADE-14*, pages 272–275. Springer, 1997.
- [21] Lintao Zhang, Conor F Madigan, Matthew H Moskewicz, and Sharad Malik. Efficient conflict driven learning in a boolean satisfiability solver. In *Proceedings of the International Conference on Computer-Aided Design*, pages 279–285. IEEE Press, 2001.