

# 1 Extended Resolution Clause Learning via Dual 2 Implication Points

3 Sam Buss 

4 University of California, San Diego, La Jolla, CA, USA

5 Jonathan Chung 

6 Lorica Cybersecurity, Canada

7 Vijay Ganesh 

8 Georgia Institute of Technology, Atlanta, USA

9 Albert Oliveras 

10 Technical University of Catalonia, Barcelona

## 11 — Abstract —

12 We present a new extended resolution clause learning (ERCL) algorithm, implemented as part of a  
13 conflict-driven clause-learning (CDCL) SAT solver, wherein new variables are dynamically introduced  
14 as definitions for *Dual Implication Points* (DIPs) in the implication graph constructed by the solver  
15 at runtime. DIPs are generalizations of unique implication points and can be informally viewed as a  
16 pair of dominator nodes, from the decision variable at the highest decision level to the conflict node,  
17 in an implication graph. We perform extensive experimental evaluation to establish the efficacy of  
18 our ERCL method, implemented as part of the MapleLCM SAT solver and dubbed xMapleLCM,  
19 against several leading solvers including the baseline MapleLCM, as well as CDCL solvers such as  
20 Kissat 3.1.1, CryptoMiniSAT 5.11, and SBVA+CaDiCaL, the winner of SAT Competition 2023. We  
21 show that xMapleLCM outperforms these solvers on Tseitin and XORified formulas. We further  
22 compare xMapleLCM with GlucoseER, a system that implements extended resolution in a different  
23 way, and provide a detailed comparative analysis of their performance.

24 **2012 ACM Subject Classification** Theory of computation

25 **Keywords and phrases** SAT, Extended resolution, Conflict Analysis

26 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

27 **Funding** *Albert Oliveras*: Supported by grant PID2021-122830OB-C43, funded by MCIN/AEI/  
28 10.13039/501100011033 and by “ERDF: A way of making Europe” and by Barcelogic through  
29 research grant C-11423.

## 30 **1** Introduction

31 Over the last several years, Conflict-Driven Clause-Learning (CDCL) SAT solvers have had  
32 a dramatic impact on many fields including software engineering [23], security [39], and  
33 AI [22, 21]. As solvers continue to be adopted in increasing complex settings, the demand  
34 for greater efficiency and reasoning power by users continues unabated.

35 While developers continue to improve CDCL SAT solvers, it is simultaneously true that  
36 these solvers are provably no more powerful than the relatively weak general resolution (Res)  
37 proof system [1, 32], and therefore are fundamentally limited. Hence, solver developers have  
38 been actively researching novel algorithms that implement stronger proof systems that go  
39 beyond Res. Examples of such algorithms include satisfaction-driven clause-learning (SDCL)  
40 SAT solvers [17, 31], bounded variable addition (BVA) [26], symmetry breaking [35], and  
41 extended resolution (ER) solvers such as GlucoseER [3].

42 Continuing this trend of strong proof system implementations, we present a new extended  
43 resolution clause learning (ERCL) algorithm, incorporated into a CDCL SAT solver, where



© Sam Buss, Jonathan Chung, Vijay Ganesh and Albert Oliveras;  
licensed under Creative Commons License CC-BY 4.0

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 new variables are dynamically introduced as definitions for *dual implication points* (DIPs)  
 45 in the implication graph constructed by the solver at run time. The concept of a DIP is  
 46 best understood as a generalization of unique implication points (UIPs). Informally, a UIP  
 47 can be defined as a dominator node in an implication graph, corresponding to a variable at  
 48 the highest decision level (DL), that *dominates* all paths from the decision variable node at  
 49 the highest DL to the conflict node. By contrast, a DIP is a pair of dominator nodes in an  
 50 implication graph such that any path from the decision variable node at the highest DL to  
 51 the conflict node must pass through at least one node in the pair<sup>1</sup>.

52 Implementation of the ERCL algorithm requires several additional methods. First, we  
 53 need a method to identify DIPs, i.e., a technique that takes as input an implication graph  
 54 and outputs a DIP and does so in time linear in the size of the input. Second, we need a  
 55 technique replaces this DIP pair with a new variable and appropriately modifies the clause  
 56 learning algorithm to learn new clauses involving DIPs. Third, we need an ER framework,  
 57 built on top of a CDCL SAT solver, that enables new variable addition, ER clause addition  
 58 and deletion, etc. Finally, we need heuristics that specialize the above mentioned methods  
 59 in a variety of ways, such as clause learning and clause deletion policies that are based  
 60 on different kinds of DIPs. We implement all of these methods as part of the MapleLCM  
 61 solver [25], and refer to the resulting solver as xMapleLCM. In fact our proposed ERCL  
 62 method, and its implementation, is very general and easily extensible thus encouraging future  
 63 exploration and specialization efforts with a variety heuristics.

#### 64 Contributions.

- 65 **1. DIP:** We introduce the concept of dual implication points (DIP), a generalization of  
 66 UIPs in conflict graphs. We also came up with an algorithm that computes them in linear  
 67 time. However, due to space limitations and to the non-trivial nature of the procedure,  
 68 we discuss this in a separate paper [11].
- 69 **2. ERCL Algorithm:** We introduce a highly parameterizable DIP-based ERCL algorithm.  
 70 The existence of a multitude of different DIPs in a single conflict graph allows us to derive  
 71 a large variety of ERCL algorithms. This flexibility is crucial in adapting the procedure to  
 72 different scenarios, unlike previous methods that couple CDCL with extended resolution.
- 73 **3. xMapleLCM:** We present a highly extensible and general ER framework as part of  
 74 xMapleLCM, which allows developers to easily add their own new variable addition, ER  
 75 clause learning/deletion, and branching policies. Given that DIP-based ERCL is highly  
 76 flexible by nature, such a framework is mandatory to quickly prototype new procedures.
- 77 **4. Experiments:** We perform extensive empirical evaluation and ablation studies on 4  
 78 different classes of instances, namely, SAT Competition 2023 Main Track, random-k-XOR,  
 79 Tseitin, and interval matching, and compare xMapleLCM against leading solvers such as  
 80 Kissat 3.1.1 [8], Cryptominisat 5.11 [36], SBVA CaDiCaL [14], and the extended-resolution  
 81 solver GlucoseER [4]. Results show that on the last 3 sets of hard combinatorial formulas,  
 82 CDCL SAT solvers perform very poorly, whereas both xMapleLCM and GlucoseER excel  
 83 on them. Hence, we can now finally state that extended resolution can be added to CDCL  
 84 and improve the performance of these solvers, and that there are at least two completely  
 85 different ways to achieve this. Moreover, a simple heuristic dynamically allows us to  
 86 detect whether extended resolution is being helpful and turn back to standard CDCL  
 87 in order to get the best of both worlds. This technique enables xMapleLCM to perform  
 88 similarly to MapleLCM on the SAT Competition 2023 Main Track instances.

---

<sup>1</sup> While it is natural to generalize the concept of a DIP to K-Implication Points or *k*-IPs, we do not discuss them in this paper.

## 2 Related Work

The idea of using ER in SAT solving has been studied in various forms in the literature for nearly two decades. The closest approach to ours is GlucoseER [3], where extended variables are introduced dynamically during the CDCL search: whenever two consecutive learned lemmas are of the form  $\neg l_1 \vee C$  and  $\neg l_2 \vee C$ , with  $l_1$  and  $l_2$  being their UIPs, then an extended variable  $z \leftrightarrow l_1 \vee l_2$  is generated and any future lemma of the form  $l_1 \vee l_2 \vee D$  is replaced by  $z \vee D$ .

Our method differs significantly from GlucoseER in the way extended variables are identified: we choose DIPs, which can be seen as pairs of variables for which adding a definition would create a better first UIP (often written as 1UIP) in the conflict graph, whereas GlucoseER definitions are constructed using already existing UIPs. Also, unlike in GlucoseER, our approach does not always learn the standard 1UIP clause, but multiple clauses might be learned that take into account the newly introduced variable. On the other hand, there are similarities at the heart of both procedures: a certain restriction of ER is considered and introduced variables are used to shorten subsequently found clauses.

More recent uses of ER in SAT solvers are Bounded Variable Addition (BVA) [26] and its structured version SBVA [14], that due to clever strategies are able to identify new extended variables whose introduction can reduce the number of clauses. Essentially, the ultimate goal of BVA techniques is to reduce the size of the formula by reducing the number of clauses at the cost of adding a new variable. One big difference with our work is that this process is done only in a preprocessing step. Finally, one additional direction that has been researched is the development of a BDD-based solver to generate ER proofs [10].

Other approaches aiming at improving CDCL solvers by allowing them to use a more powerful proof system are related to the Propagation Redundancy notion [16, 18], either via preprocessing steps [34] or via the use of the SDCL algorithm [19, 31]. While these methods implement proof systems that are stronger than Res, many of them (without new variable addition) are known to be weaker than ER.

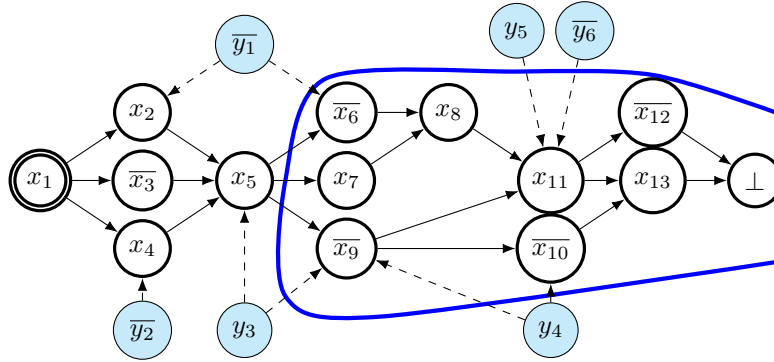
## 3 Preliminaries

We assume that the reader is familiar with the satisfiability (SAT) problem and the CDCL algorithm, and we refer her to the Handbook of Satisfiability for an excellent overview of these topics [27]. Below we focus on conflict analysis, which is the most relevant ingredient from the CDCL algorithm for this paper. We do so by means of the following example.

► **Example 1.** Consider the following clauses

- |  |   |  |
|--|---|--|
| (1) $y_1 \vee \neg x_1 \vee x_2$                             | (6) $\neg x_5 \vee x_7$                                     | (11) $\neg x_{11} \vee x_{12}$             |
| (2) $\neg x_1 \vee \neg x_3$                                 | (7) $x_6 \vee \neg x_7 \vee x_8$                            | (12) $x_{10} \vee \neg x_{11} \vee x_{13}$ |
| (3) $y_2 \vee \neg x_1 \vee x_4$                             | (8) $\neg y_3 \vee \neg y_4 \vee \neg x_5 \vee \neg x_9$    | (13) $x_{12} \vee \neg x_{13}$             |
| (4) $\neg y_3 \vee \neg x_2 \vee x_3 \vee \neg x_4 \vee x_5$ | (9) $\neg y_4 \vee x_9 \vee \neg x_{10}$                    |  |
| (5) $y_1 \vee \neg x_5 \vee \neg x_6$                        | (10) $\neg y_5 \vee y_6 \vee \neg x_8 \vee x_9 \vee x_{11}$ |  |

Assume that CDCL has constructed an assignment that contains, among others, literals  $\{\neg y_1, \neg y_2, y_3, y_4, y_5, \neg y_6\}$ . Since no propagation is possible, it now decides to add the decision literal  $x_1$ . Due to clause (1), we can unit propagate literal  $x_2$ , being (1) the reason of  $x_2$  and its antecedents  $\{\neg y_1, x_1\}$ . Similarly,  $\neg x_3$  is propagated due to reason (2), with antecedents  $\{x_1\}$ , and  $x_4$  due to reason (3) with antecedents  $\{\neg y_2, x_1\}$ . If we continue this process we eventually find that clause (13) is conflicting and we can construct the *conflict graph* in



■ **Figure 1** Conflict graph associated with Example 1. If 1UIP learning is applied, we generate lemma  $\neg y_4 \vee \neg y_5 \vee y_6 \vee x_5$ . White nodes belong to the current decision level, whereas blue ones are from previous decision levels.

127 Figure 1, where every literal in the current decision level has incoming edges corresponding  
 128 to its antecedents (except of course the decision literal). A special conflict node  $\perp$  with  
 129 incoming edges  $\{x_{12}, x_{13}\}$  represents conflicting clause (13).

130 The graph clearly shows that if we set  $x_1$  to true, together with the literals of previous  
 131 decision levels (the  $y$ 's), we obtain a conflict. However, the same happens with  $x_5$ , since  
 132 any path from  $x_1$  to the conflict necessarily goes through  $x_5$ . Literals with this property are  
 133 called *Unique Implication Points* (UIPs), of which we only have  $x_1$  and  $x_5$ . Since  $x_5$  is the  
 134 one closest to the conflict we call it *First Unique Implication Point* (1UIP) [40]. It is easy  
 135 to see that if we set  $x_5$  and the  $y$  literals that enter the cut delimited by the blue line, unit  
 136 propagation derives the same conflict. Hence, since they cannot be simultaneously true, we  
 137 can learn  $\neg y_3 \vee \neg y_4 \vee \neg y_5 \vee y_6 \vee x_5$ . The quality of a lemma can be assessed by its *Literal*  
 138 *Block Distance* (LBD) [4]: the number of different decision levels of the literals in the lemma.  
 139 The lower the LBD, the better the lemma. In our case, if we are at decision level 5,  $y_4$  and  
 140  $y_6$  belong to decision level 2, and  $y_5$  to decision level 4, the LBD of the lemma is 3. □

141 *Resolution.* Given two clauses  $l \vee C$  and  $\neg l \vee D$ , the *resolution* inference rule allows one to  
 142 derive the logical consequence  $C \vee D$ . It is well known that the lemma derived in Example 1  
 143 can be obtained via a series of resolution steps that start with the conflicting clause, and  
 144 resolve with reasons of literals in the reverse order in which they were added to the assignment.  
 145 In fact, it has been proved [33, 2] that resolution and CDCL (with restarts) are polynomially  
 146 equivalent, and hence classes of formulas that are hard for resolution; e.g., the pigeonhole  
 147 principle - PHP [15], or Tseitin formulas [38]) are also hard for CDCL SAT solvers.

148 *Extended Resolution (ER).* Given two literals  $l_1$  and  $l_2$ , the *extended resolution* [37] rule  
 149 allows us to introduce clauses representing the definition  $z \leftrightarrow l_1 \vee l_2$ . ER can be substantially  
 150 more powerful than resolution; for instance, it allows polynomial size proofs of PHP [13].  
 151 Incorporating ER to CDCL solvers could potentially enable them to solve such formulas  
 152 in polynomial time. However, we lack good methods to incorporate ER into CDCL proof  
 153 search. This is precisely the aim of this paper, namely, incorporating a restricted version of  
 154 ER into CDCL.

#### 155 4 Dual Implication Points

156 As discussed in the previous section, the unique implication points (UIPs) are crucial for  
 157 conflict clause learning in the CDCL algorithm. We now introduce a new concept of a *Dual*

Dual Implication Point (DIP)	Extension Variable	Post-DIP Learned Clause	Pre-DIP Learned Clause
$\overline{x_{12}}, x_{13}$	$z \leftrightarrow (\overline{x_{12}} \wedge x_{13})$	$\neg z$	$\neg x_5 \vee y_1 \vee \neg y_3 \vee \neg y_4 \vee \neg y_5 \vee y_6 \vee z$
$x_{11}, x_{13}$	$z \leftrightarrow (x_{11} \wedge x_{13})$	$\neg z$	$\neg x_5 \vee y_1 \vee \neg y_3 \vee \neg x_4 \vee \neg y_5 \vee y_6 \vee z$
$\overline{x_{10}}, x_{11}$	$z \leftrightarrow (\overline{x_{10}} \wedge x_{11})$	$\neg z$	$\neg x_5 \vee y_1 \vee \neg y_3 \vee \neg y_4 \vee \neg y_5 \vee y_6 \vee z$
$\overline{x_9}, x_{11}$	$z \leftrightarrow (\overline{x_9} \wedge x_{11})$	$\neg z \vee \neg y_4$	$\neg x_5 \vee y_1 \vee \neg y_3 \vee \neg y_4 \vee \neg y_5 \vee y_6 \vee z$
$x_8, \overline{x_9}$	$z \leftrightarrow (x_8 \wedge \overline{x_9})$	$\neg z \vee \neg y_4 \vee \neg y_5 \vee y_6$	$\neg x_5 \vee y_1 \vee \neg y_3 \vee \neg y_4 \vee z$

■ **Figure 2** The complete list of DIPs for the conflict graph of Figure 1. The DIP-learnable clauses involve the new extension variable  $z$  that can be introduced for that DIP. (The extension clauses defining  $z$  must also be learned; e.g., for the first line, the extension clauses express that  $z \leftrightarrow \overline{x_{12}} \wedge x_{13}$ .)


158 *Implication Point* (DIP) that gives a tool for analyzing the conflict graph. Its applications  
 159 include discovering new implied 2-clauses, introducing new variables by extension, and  
 160 learning clauses involving the extension variables. The idea behind a dual implication point  
 161 is that it consists of a pair of vertices (literals) in the conflict graph that disconnects or  
 162 “separates” the decision literal from the contradiction.<sup>2</sup> More precisely, a DIP is defined to be  
 163 a pair  $\{x, y\}$  of literals such that all paths in the conflict graph to the vertex  $\perp$  pass through  
 164 at least one of  $x$  and  $y$  and such that neither  $x$  nor  $y$  is a UIP. In contrast, UIP is a *single*  
 165 literal that separates the decision literal from the conflict.

166 We use the example in Figure 1 to illustrate the concept of DIPs and their potential  
 167 applications. Recall that  $x_5$  is the first UIP. In our applications, we are seeking DIPs between  
 168 the first UIP and the conflict node  $\perp$ . An obvious DIP is the pair  $\overline{x_{10}}$  and  $x_{11}$ , since it is  
 169 immediately clear that any path from  $x_5$  (or from  $x_1$ ) must pass through one of  $\overline{x_{10}}$  or  $x_{11}$ .  
 170 On the other hand, the pair  $\overline{x_{10}}$  and  $x_8$  is not a DIP since there are paths from  $x_1$  to  $\perp$  that  
 171 avoid these two literals; namely, any path that includes the edge from  $\overline{x_9}$  to  $x_{11}$ . There are  
 172 several other DIPs in Figure 1: a complete list is given in Figure 2.

173 Figure 2 also shows how a DIP pair can be used to introduce a variable  $z$  via extension, and  
 174 the associated clauses that can be learned. For example, in the third line, the new extension  
 175 variable  $z$  is introduced with the three clauses  $\neg z \vee \neg x_{10}$ ,  $\neg z \vee x_{11}$  and  $x_{10} \vee \neg x_{11} \vee \neg z$  which  
 176 express the condition  $z \leftrightarrow (\neg x_{10} \wedge x_{11})$ . From the conflict graph, this allows inferring the  
 177 clauses  $\neg z$  (the “post-DIP” learned clause) and  $\neg x_5 \vee y_1 \vee \neg y_3 \vee \neg y_4 \vee \neg y_5 \vee y_6 \vee z$  (the  
 178 “pre-DIP” learned clause). Since the post-DIP learned clause does not have any variables  
 179 from lower decision levels, we can also infer the 2-clause  $x_{10} \vee \neg x_{11}$  either instead of or in  
 180 addition to introducing  $z$  and the pre- and post-DIP clauses. Introducing 2-clauses in this  
 181 way might be helpful for CDCL solvers that do special processing of 2-clauses; for instance,  
 182 in the work of Bacchus et al. [5] or the recent work of Biere et al. [9] or Buss et al. [12].

183 In general, for any literals  $a, b$  that form a DIP in the above fashion, we may introduce  
 184 an extension variable  $z \leftrightarrow a \wedge b$ , and learn (1) a pre-DIP clause of the form  $\neg f \vee \neg C \vee z$  (i.e.  
 185  $f \wedge C \rightarrow z$ ), where  $f$  is the first UIP and  $C$  the set of literals from previous decision levels  
 186 that have an edge in the conflict graph to any literal appearing after the first UIP and no  
 187 later than the DIP pair; and (2) a post-DIP clause of the form  $\neg z \vee \neg D$  (i.e.  $z \wedge D \rightarrow \perp$ ),  
 188 where  $D$  contains those literals from previous decision levels with an edge to any literal  
 189 appearing strictly after the DIP pair.

<sup>2</sup> Perhaps “Dual Implication Pair” would be a better name than “Dual Implication Point”, since a DIP is a pair of literals. We use “Point” however to match the terminology of “Unique Implication Points”.



■ **Figure 3** Nodes  $a_4$ ,  $b_3$  and  $b_4$  are bypassed and thus not in any TVB. The path  $b_2, c, d, a_4$  is a crossing separator that prevents any of  $a_1$ ,  $a_2$  and  $a_3$  from being paired with  $b_5$  to form a TVB pair. All other pairs  $\{a'_i, b'_j\}$  are valid TVB pairs.

190 For example, the last line of the table in Figure 2 shows the case  $z \leftrightarrow (x_8 \wedge \overline{x_9})$ , where  
 191  $f$  is  $x_5$ , and  $C$  and  $D$  are  $\neg y_1 \wedge y_3 \wedge y_4$  and  $y_4 \wedge y_5 \wedge \neg y_6$ , respectively. In Section 5.2 we  
 192 discuss many possible ways that DIP extension variables and clauses may be introduced into  
 193 CDCL solvers. The rest of this section describes how to find DIPs.

#### 194 4.1 An Algorithm for Finding DIPs

195 A conventional CDCL algorithm maintains a trail of the literals set true, in the order they  
 196 were set, and this allows finding the UIP very quickly. Finding the DIPs is much more  
 197 complex than finding the UIP, and requires several traversals of the conflict graph between  
 198 the first-UIP and the conflict node. Nonetheless, it is possible to find all DIPs very quickly,  
 199 even in linear time.

200 We recast the DIP-finding problem in terms of a general directed graph  $G$ . Let  $G = (V, E)$   
 201 be a directed graph with two distinguished vertices  $s$  and  $t$ . We assume that  $s$  is 2-connected  
 202 to  $t$  in that there is a pair of vertex-disjoint paths  $\pi_1$  and  $\pi_2$  from  $s$  to  $t$  that share no vertices  
 203 apart from  $s$  and  $t$ . By the vertex-version of Menger’s theorem [28], either there are in fact  
 204 *three* vertex-disjoint paths from  $s$  to  $t$  or there is at least one pair of vertices  $\{a, b\}$  with  
 205 neither  $a$  nor  $b$  in  $\{s, t\}$  so that every path from  $s$  to  $t$  passes through at least one of  $a$  or  $b$ .  
 206 Such a pair  $\{a, b\}$ , if it exists, is called a *Two Vertex Bottleneck* (TVB).

207 Our goal is to find all possible TVBs efficiently and in linear time. An algorithm for this  
 208 is discussed in our companion paper [11]; for space reasons, we give only an abbreviated  
 209 discussion of the algorithm here. The first step is to find two vertex disjoint paths from  $s$   
 210 to  $t$ : this is done by greedily finding a path from  $s$  to  $t$  and then using an augmenting path  
 211 construction to find two vertex-disjoint paths from  $s$  to  $t$ . An example is shown in Figure 3,  
 212 where the two paths are  $a_0, \dots, a_\ell$  and  $b_0, \dots, b_k$  where  $a_0 = b_0 = s$  and  $a_\ell = b_k = t$ . These  
 213 paths are called  $\pi_a$  and  $\pi_b$ , respectively. Henceforth, a *path* is a directed path without any  
 214 repeated nodes. The *internal* vertices of a path  $\pi$  are the vertices on  $\pi$  other than the first  
 215 and last vertices. Two paths are said to be *vertex-disjoint* if they have no internal vertices in  
 216 common. A path  $\pi$  *avoids*  $\pi_a$  and  $\pi_b$  if it is vertex-disjoint from both paths.

217 Once the two vertex-disjoint paths are fixed, we can state the following definitions and  
 218 theorems:

219 ► **Definition 2.** [11] *A node  $a_i$  on  $\pi_a$  is bypassed if there are  $j < i < j'$  and a path  $\pi$  from*  
 220  *$a_j$  to  $a_{j'}$  such that  $\pi$  avoids  $\pi_a$  and  $\pi_b$ . A node  $b_i$  being bypassed is defined similarly.*

221 ► **Definition 3.** *Two nodes  $a_i$  and  $b_j$  have a crossing separator if there are nodes  $a_{i'}$  and  $b_{j'}$*   
 222 *joined by a path  $\pi$  that avoids both  $\pi_a$  and  $\pi_b$  such that either (a)  $i' < i$  and  $j' > j$  and  $\pi$  is*  
 223 *a path from  $a_{i'}$  to  $b_{j'}$ , or (b)  $i' > i$  and  $j' < j$  and  $\pi$  is a path from  $b_{j'}$  to  $a_{i'}$ .*

224 ► **Theorem 4.** [11] *For  $0 < i < \ell$  and  $0 < j < k$ , the two nodes  $a_i$  and  $b_j$  form a two-vertex*  
 225 *bottleneck (TVB) if and only if  $a_i$  and  $b_j$  do not have a crossing separator and neither  $a_i$*

226 *nor  $b_j$  is bypassed.*

227 Theorem 4 is proved in [11]. The theorem holds for both acyclic and cyclic directed  
228 graphs; however, for our applications to CDCL we are interested only in acyclic graphs since  
229 the conflict graph is always acyclic.

230 A consequence of Theorem 4 is that the set of all TVBs can be compactly represented  
231 in linear size, even though there can be quadratically many TVBs. Let  $a'_1, \dots, a'_{\ell'}$  be the  
232 subsequence of the internal nodes  $a_1, \dots, a_{\ell-1}$  of path  $\pi_a$  that, according to the conditions of  
233 Theorem 4, are in at least one TVB pair. Let  $b'_1, \dots, b'_{k'}$  be the corresponding subsequence  
234 of the internal nodes of  $\pi_b$ . Then, for each  $a'_i$  there are  $m \leq n$  such that  $a'_i$  forms a TVB  
235 pair with each  $b'_j$  with  $m \leq j \leq n$ . Dually, for each  $b'_i$  there are  $m \leq n$  such that  $b'_i$  forms a  
236 TVB pair with each  $a'_j$  with  $m \leq j \leq n$ . (See Figure 4.)

237 **► Example 5.** In the conflict graph of Figure 1, consider the portion of the graph between  
238 the first UIP  $x_5$  and the contradiction  $\perp$ . We can take path  $\pi_a$  to be  $x_5, x_7, x_8, x_{11}, \overline{x_{12}}, \perp$   
239 and path  $\pi_b$  to be  $x_5, \overline{x_9}, \overline{x_{10}}, x_{13}, \perp$ . The node  $x_7$  is bypassed by the path  $x_5, \overline{x_6}, x_8$ , so it  
240 cannot be part of a TVB. There are two crossing separator paths: the first is the edge from  
241  $\overline{x_9}$  to  $x_{11}$ ; the second is the edge from  $x_{11}$  to  $x_{13}$ . Therefore, the possible TVB pairs can be  
242 described in a table as:

Node on $\pi_a$	forms a TVB pair with	Node on $\pi_b$	forms a TVB pair with
$x_8$	$\overline{x_9}$	$\overline{x_9}$	$x_8, x_{11}$
$x_{11}$	$\overline{x_9}, \overline{x_{10}}, x_{13}$	$\overline{x_{10}}$	$x_{11}$
$\overline{x_{12}}$	$x_{13}$	$x_{13}$	$x_{11}, \overline{x_{12}}$

244 In this example, one node,  $x_7$ , was bypassed. It is also possible that non-bypassed nodes are  
245 eliminated just by the crossing separators. For example, if there were an additional edge  
246 from  $x_7$  to  $\overline{x_{10}}$ , then the crossing separator condition would imply that  $x_8$  and  $x_{11}$  are not  
247 part of any TVB pair. In this case,  $x_8$  and  $x_{11}$  would not be included among the  $a'_i$  nodes.


248 Theorem 4 is used in [11] to give an efficient, linear time algorithm for finding all TVBs.  
249 The algorithm has five phases. The first two phases find two vertex-disjoint paths from  $s$   
250 to  $t$ ; the next phase scans the graph from  $t$  to  $s$  to discover all relevant paths that avoid  
251  $\pi_a$  and  $\pi_b$ ; the fourth phase uses this to discard bypassed nodes and collect information on  
252 crossing separators; finally the fifth phase computes the compressed representation of all  
253 possible TVBs. Full details are given [11], which is available online in preprint form. In our  
254 experiments with the implication graphs constructed by the underlying CDCL solver, the  
255 time overhead in finding the TVBs is negligible.

## 256 **5 Extension Variables from Dual Implication Points**

257 This section discusses how DIPs can be used for introducing extension variables and implement  
258 the ERCL method for learning ER clauses in CDCL solver. We first present an example.

### 259 **5.1 An example with a grid Tseitin principle**

260 Given a graph where every vertex has a *charge*, a number which is 0 or 1, a Tseitin formula [37]  
261 is created by considering one variable per edge, and adding one constraint per vertex  $v$   
262 expressing that the sum of the variables of all edges incident to  $v$  modulo 2 is equal to its  
263 charge. The CNF version that we consider converts each *xor* constraint into clauses by simple



■ **Figure 4** An instance of the  $3 \times 3$ -grid Tseitin principle. The nodes are assigned a polarity in  $\{0, 1\}$ . The edges are labeled with variables.

264 enumeration of falsifying assignments. It is easy to see that the formula is unsatisfiable if and  
 265 only if the sum of all charges is even. Here we consider as a graph the  $3 \times 3$ -grid depicted in  
 266 Figure 4 and only one vertex has charge 1, so the clauses are unsatisfiable.

267 The first steps of the CDCL solver are as follows: First  $e_1$  is set true as a decision literal,  
 268 and  $\bar{e}_3$  is (unit) propagated. Second,  $e_2$  is set true as a decision literal, and  $\bar{e}_4$  and  $e_5$   
 269 are propagated. Third,  $e_6$  is set true as a decision literal, and  $e_8$  and  $e_{11}$  are propagated.  
 270 Fourth,  $e_7$  is set true as a decision literal, and  $\bar{e}_9$ ,  $\bar{e}_{10}$ , and  $e_{12}$  are propagated. This gives a  
 271 contradiction, since the clause  $e_{10} \vee \bar{e}_{12}$  (one of the two clauses from  $e_{10} \oplus e_{12} = 0$ ) is falsified.  
 272 Figure 5(a) shows the complete conflict graph at this point.

273 Examining the conflict graph at decision level 4, the first UIP is  $e_7$  and there are two DIPs  
 274 available,  $\{\bar{e}_9, \bar{e}_{10}\}$  and  $\{\bar{e}_{10}, e_{12}\}$ . Selecting the former DIP, we introduce a new variable  $x$   
 275 by extension as  $x \leftrightarrow \bar{e}_9 \wedge \bar{e}_{10}$ . We can learn the additional pre- and post-DIP clauses:  
 276  $e_4 \vee \bar{e}_5 \vee \bar{e}_6 \vee \bar{e}_7 \vee x$  and  $\bar{x} \vee \bar{e}_{11}$

277 We next backtrack to decision level 3, unsetting  $e_7$ ,  $e_9$ ,  $e_{10}$  and  $e_{12}$ . Unit propagation at  
 278 decision level 3 sets the new literal  $x$  and the first UIP  $\bar{e}_7$  false and then sets literals  $e_9$ ,  $e_{10}$ ,  
 279 and  $\bar{e}_{12}$  true. This yields a contradiction with the clause  $\bar{e}_{10} \vee e_{12}$ . In the conflict graph at  
 280 decision level 3, the first UIP is  $e_6$  and there are two DIPs available, namely  $\{\bar{e}_7, \bar{e}_{12}\}$  and  
 281  $\{e_{10}, \bar{e}_{12}\}$ . If we select the first one, then we introduce a new literal  $y$  defined by extension as  
 282  $y \leftrightarrow \bar{e}_7 \wedge \bar{e}_{12}$  and can in addition learn the clauses  $e_3 \vee e_4 \vee \bar{e}_5 \vee \bar{e}_6 \vee y$  and  $\bar{e}_5 \vee \bar{y}$ .

283 We do not carry this example further, but note that our experiments show that Tseitin  
 284 tautologies (not just on grid graphs) are examples where our experiments show the DIP  
 285 clause learning method is especially effective.

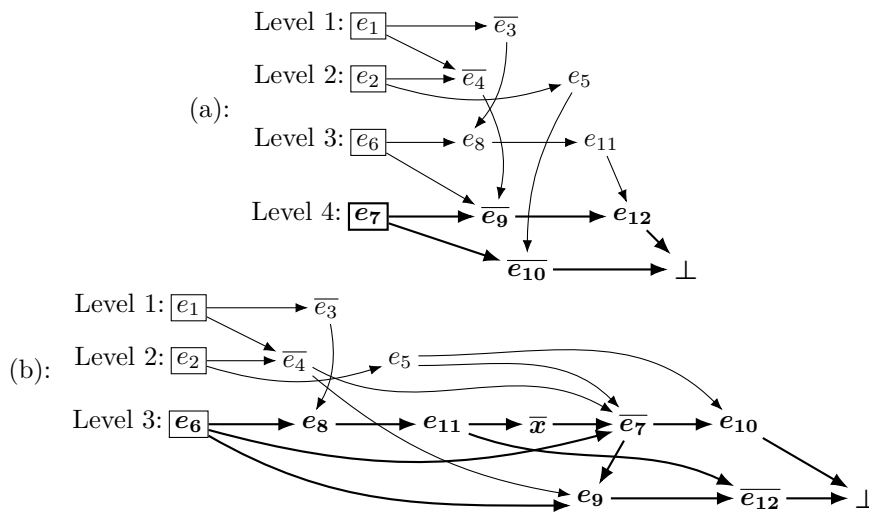
286 It is interesting to relate the Figure 5(b) to Theorem 4 on DIPs. In this example, there  
 287 is only one way to choose the two vertex-disjoint paths. Namely, to let  $\pi_a$  and  $\pi_b$  be the  
 288 paths  $e_6, e_8, e_{11}, \bar{x}, \bar{e}_7, e_{10}, \perp$  and  $e_6, e_9, \bar{e}_{12}, \perp$ . The edge from  $e_6$  to  $\bar{e}_7$  bypasses  $e_8, e_{11}$  and  $\bar{x}$ ;  
 289 and the edge from  $e_{11}$  to  $\bar{e}_{12}$  is a crossing separator. (The edge from  $\bar{e}_7$  to  $e_9$  is a “vacuous”  
 290 crossing separator that does not actually remove any possible DIP pairs.)

291 It should be evident from this example that many conflicts have DIPs; indeed our  
 292 experiments reported below show that approximately 2/3 of the conflicts have at least one  
 293 DIP and, very frequently there are quite a few choices for DIPs.

## 294 5.2 Extending CDCL with Dual Implication Points

295 The use of DIPs in conflict analysis opens a large spectrum of possibilities. This section  
 296 discusses some of them, with particular attention to the techniques that we have imple-  
 297 mented. Our present implementation, xMapleLCM, is a flexible framework that allows  
 298 one to implement extended-resolution based techniques in a simple way. It offers a set  
 299 of clearly-specified functions that facilitate determining which extension variables to add,





■ **Figure 5** The complete conflict graph from the example; (a) at the first conflict and (b) at the second conflict. Boxed literals are decision literals. Implications at the top decision level are in bold.

300 performing the corresponding addition, the possible posterior deletion or replacing definitions  
 301 in clauses. Additionally, a set of heuristic choices to control when and how these steps are  
 302 performed are provided, and replacing them by custom ones is a smooth task. We want to  
 303 remark that there are many more possible strategies for using DIPs than could be discussed.  
 304 We believe that in this paper we merely scratched the surface of heuristics for exploiting  
 305 DIPs, which is an indication of the potential of this approach.

306 **Choice of DIP.** As mentioned, there is possibly a quadratic number of DIPs. Even though  
 307 we could learn multiple DIPs at every conflict, with their corresponding lemmas, we decided  
 308 to choose only one. The first possibility we considered is to learn the DIP that is closest  
 309 to the conflict, as we do with UIPs. This may often create a short post-DIP conflict but a  
 310 long pre-DIP clause. Therefore, we considered the possibility of choosing a DIP that splits  
 311 the conflict graph into two balanced regions. Ideally, that would result in two equally long  
 312 pre and post-DIP clauses. Finally, we also implemented choosing a random DIP, to check  
 313 whether any of the other two schemes could outperform a random strategy. These heuristics  
 314 are referred to as **closest**, **middle**, and **random**, respectively.

315 **Filtering out bad-quality DIPs.** Learning a DIP whenever we find one would be too  
 316 prolific and overwhelm the solver. In order to determine whether the DIP chosen in the  
 317 previous step had to be discarded, the first possibility we explored was again inspired by  
 318 1-UIP learning, where learning glue clauses, i.e. having LBD equal to 2, is a desired situation.  
 319 The first filtering mechanism we implemented discarded all DIPs that did not have a glue  
 320 post-DIP clause. Another possibility we considered is to wait for a DIP to occur a certain  
 321 number of times before using it in DIP-based learning. In our implementation, we tried 20  
 322 and 5 occurrences, as the minimum threshold to introduce a DIP. A third possibility is to  
 323 use the activity-based heuristic of the literals in the DIP to assess its quality: DIPs whose  
 324 literals have high decision-heuristic scores should be prioritized. In our implementation,  
 325 we check whether the activity level of the current DIP is higher than the average activity  
 326 level of the 20 most recently encountered DIPs; if so, the current DIP is a candidate for  
 327 DIP-learning, otherwise it is discarded. We refer to these various techniques as **glue**, **occ5**  
 328 and **act**, respectively.

329 **Learning pre-DIP and post-DIP clauses.** In our implementation, we only considered  
 330 two variants: one that always learns both the pre- and the post-DIP clauses (**2-clause**) and  
 331 one that only learns the post-DIP clause (**1-clause**).

332 **Backjumping and asserting clauses.** Recall that when a new DIP extension variable  $z$  is  
 333 introduced, it is possible to learn a pre-DIP clause of the form  $\neg f \vee \neg C \vee z$  and a post-DIP  
 334 clause of the form  $\neg z \vee \neg D$ , where  $f$  is the first UIP and where  $C$  and  $D$  are conjunctions of  
 335 literals that were set true at lower levels. (Note that  $C$  and  $D$  may have literals in common.)  
 336 Letting  $\ell_C$  and  $\ell_D$  be the maximum of the levels at which literals in  $C$  and  $D$  (respectively)  
 337 were set, our implementation always backtracks to level  $\ell_D$ . This makes the post-DIP clause  
 338 asserting, so  $\neg z$  is set at decision level  $\ell_D$ . Furthermore, if  $\ell_C \leq \ell_D$  and the pre-DIP clause  
 339 is learned, then it is asserting and  $\neg f$  is set by unit propagation at decision level  $\ell_D$ , as it  
 340 should be.

341 It would make no sense to backtrack to level  $\ell_C$  when  $\ell_C < \ell_D$  since then neither  $\neg z$   
 342 nor  $\neg f$  would be propagated. However, another possible strategy would be to backtrack  
 343 to the maximum of the decision levels  $\ell_C$  and  $\ell_D$ . This would mean  $\neg z$  and  $\neg f$  are both  
 344 propagated. The disadvantage of this when  $\ell_C > \ell_D$  is that it would mean  $\neg z$  is asserted  
 345 by the post-DIP clause at level  $\ell_C$ , whereas it could have been propagated at the previous  
 346 decision level  $\ell_D$ . This breaks a usual invariant for CDCL solvers. This would only be  
 347 possible in a solver that permits chronological backtracking [30, 29]; xMapleLCM, however,  
 348 does not support this.

349 The previous reasoning needs some clarification for the case when the extension variable  
 350  $z$  with definition  $z \leftrightarrow l_1 \wedge l_2$  we want to introduce already exists. If  $z$  is undefined or defined  
 351 at the current decision level nothing changes. We know it cannot be true at some previous  
 352 level, because otherwise  $l_1$  and  $l_2$  would have been propagated at that level, and not in the  
 353 current one as all literals belonging to a DIP do. If it is false at some previous level, then  
 354 we have no guarantee that the pre or the post-DIP clause is asserting at any decision level.  
 355 Fortunately, that rarely happens in practice. However, we can always perform standard 1UIP  
 356 learning or try to apply DIP-based conflict analysis starting with the clause  $z \vee \neg l_1 \vee \neg l_2$   
 357 that we can guarantee is conflicting at the current decision level.

358 **Replacing literals by extended variables.** In xMapleLCM, every time a new lemma is  
 359 learned, we try to replace some of its literals by extended variables. An extended variable  
 360  $z \leftrightarrow l_1 \vee l_2$ , allows one to replace a lemma of the form  $l_1 \vee l_2 \vee C$  by  $z \vee C$ . This is done  
 361 by checking, for all pairs of literals in the lemma that appear in some extended variable  
 362 definition whether they are part of the same definition. Too long lemmas or lemmas that  
 363 have large LBD are discarded to mitigate the cost of this operation.

364 If one wants to introduce literal  $\neg z$  and obtain some reduction in formula size, the lemma  
 365 should be of the form  $\neg l_1 \vee C$  and there should be another clause of the form  $\neg l_2 \vee C$ . In this  
 366 case, they can be replaced by a single clause  $\neg z \vee C$ . However, this situation can be expensive  
 367 to detect since one has to traverse the whole database looking for a certain clause, and this  
 368 operation should be repeated for every literal in the lemma. For this reason, xMapleLCM  
 369 does not implement this.

370 **Deleting extended variables.** As it happens with lemma learning, where keeping too  
 371 many lemmas slows down unit propagation, managing too many extended variables might  
 372 also be counterproductive. Again following the analogy with learned lemmas, which are  
 373 useful at some point of the search but might become inactive after a while, it is natural to  
 374 think that extended variables follow the same behavior. All in all, it seems mandatory to  
 375 consider the deletion of extended variables.

376 Deletion of variables in xMapleLCM is scheduled to be performed every 1000 conflicts.

377 At that point, several strategies are possible: delete all variables, delete the ones with a  
 378 minimum decision-heuristic activity, or delete the worst  $k\%$  variables according to some  
 379 criterion (e.g. their decision-heuristic activity). In our implementation, we do the latter  
 380 with  $k = 50$ . Note that variables appearing in the right-hand side of an extended-variable  
 381 definition cannot be deleted. This is addressed by maintaining a counter for every variable  
 382 that corresponds to the number of definitions where it is involved.

## 383 **6 Experimental Evaluation**

384 We have implemented the DIP-based clause learning schemes described in Section 5.2 on top  
 385 of the xMapleLCM ER framework<sup>3</sup>. We started our experimental evaluation running a variant  
 386 of DIP-based learning on all benchmarks of the 2023 SAT Competition [6] and comparing  
 387 it to MapleLCM [25], the CDCL SAT solver on which it is based. Even though there did  
 388 not seem to be a systematic improvement on all benchmarks, a few families with important  
 389 speedups were identified. We start this section by analyzing the impact of DIP-based learning  
 390 on these families and then move to final considerations about the performance on the overall  
 391 2023 SAT competition benchmarks.

### 392 **6.1 Performance Analysis Methodology**

393 For each family, we start by describing the problem they encode. After that, we report on the  
 394 performance a variety of state-of-the-art solvers, each with some distinguished characteristic:  
 395 Kissat 3.1.1 [8] (an extremely efficient CDCL solver), Cryptominisat 5.11 [36] (support for  
 396 XOR reasoning), SBVA+CaDiCaL [14] (introduction of new variables via SBVA and winner  
 397 of the main track of the 2023 SAT Competition), GlucoseER [3] (extended-resolution based  
 398 CDCL solver) and a variant of xMapleLCM-DIP, which we refer to as the baseline, that we  
 399 describe next.

400 Finally, we evaluate the impact of the different techniques explained in Section 5.2. In  
 401 particular, we first consider as a **baseline** a version that (i) finds DIPs in the middle of  
 402 the conflict graph, (ii) only adds a DIP if it has occurred at least 20 times and (iii) always  
 403 learns the both pre- and post-DIP learning clauses. Different variants can be obtained by  
 404 changing only one of the three previous design decisions at a time. Regarding the type of  
 405 DIP used, we analyze the performance of **closest** and **random**, that is, the systems whose  
 406 only difference w.r.t. **baseline** is that the type of DIP is changed. Regarding the criterion  
 407 used to discard a DIP, we analyze the **glue**, **act** and **occ5** configurations, where the latter  
 408 discards any DIP that has not occurred at least 5 times. Finally, the **1-clause** variant only  
 409 learns the post-DIP clause. Our goal is to understand why certain configurations do not  
 410 work well on particular families.

#### 411 **6.1.1 Matching of Properly Intersecting Intervals**

412 We are given a sequence of numbers  $(a_1, \dots, a_n)$ , initially all set to zero, and a set of  
 413 operations, each consisting of assigning a certain number to a contiguous subsequence of  
 414 the  $a_i$ 's, defined by an interval. The goal is to perform each operation exactly once while  
 415 maximizing the number of pairs of consecutive numbers  $(a_i, a_{i+1})$  that are different at the  
 416 end of the process. Given some additional conditions, this can be formulated as maximum

---

<sup>3</sup> All sources used for this evaluation can be found in <https://github.com/chjon/xMapleSAT/tree/main>.

Baseline	1-Clause	Closest	Rand	Glue	Occ5	Act	
<b>PROPERLY INTERSECTING INTERVALS FORMULAS</b>							
SOLVED	6	5	<b>7</b>	<b>7</b>	0	0	0
BEST	1	1	2	<b>3</b>	0	0	0
Median (secs.)	2777	2566	2358	<b>2189</b>	3600	3600	3600
<b>TSEITIN FORMULAS</b>							
SOLVED	11	13	13	11	0	<b>15</b>	3
BEST	2	5	0	0	0	<b>8</b>	0
Median (secs.)	897	628	2707	2782	3600	<b>425</b>	3600
<b>(X)ORIFIED RANDOM <math>k</math>-xor FORMULAS</b>							
SOLVED	10	8	<b>11</b>	9	0	9	7
BEST	<b>5</b>	2	1	2	0	1	0
Median (secs.)	<b>430</b>	803	1080	820	3600	936	3112

■ **Table 1** Performance of DIP-learning variants on three sets of formulas. Median running times are in seconds and, for each row, the best performing system is in bold.

417 bipartite matching problem on a certain graph. By removing some edges from this graph, an  
 418 unsatisfiable problem is generated. A more detailed description can be found in [6].

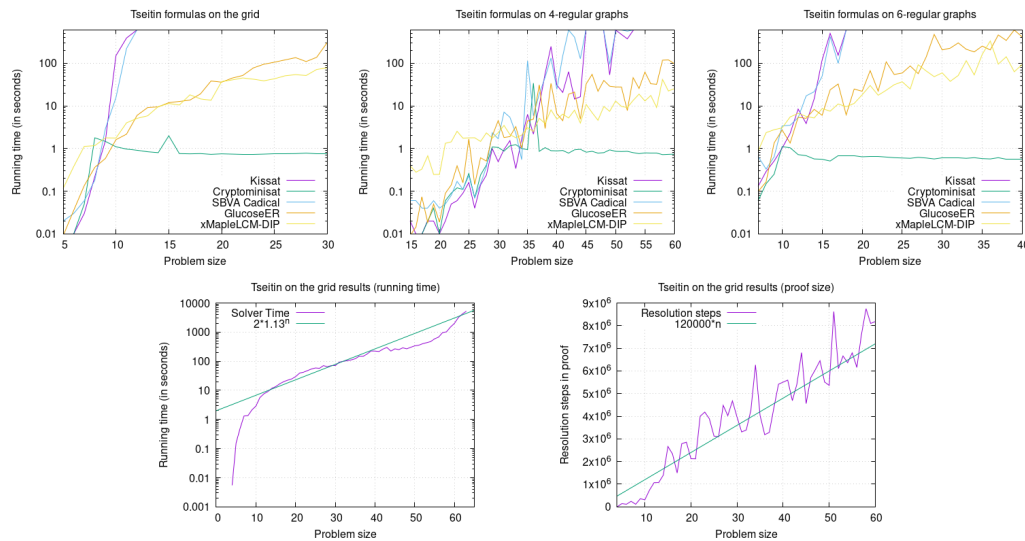
419 We downloaded the 23 instances submitted to the 2023 SAT Competition, using the  
 420 Global Benchmark Database [20], and observed that no system without extended-resolution  
 421 reasoning could solve any benchmark in a time limit of 1 hour. On the other hand, there are  
 422 14 benchmarks that GlucoseER could solve, of which our DIP-based variants could solve 7.  
 423 Detailed results for our variants can be seen in Table 1.

424 For this particular type of formulas, the **rand** variant performs the best, but detailed  
 425 results indicate that we cannot claim it is systematically faster. One clear conclusion we can  
 426 infer is that **glue**, **occ5** and **act** exhibit very poor performance. A more detailed analysis  
 427 on the runs of these systems reveal that, at least in these benchmarks, both **act** and **occ5**  
 428 apply DIP learning in many more conflicts than the other variants, whereas the distinctive  
 429 characteristic of **glue** is that the percentage of decisions on extended variables was extremely  
 430 low. As can be seen at the end of this section, the latter is a situation that tends to indicate  
 431 the DIP learning is not helping the solver in this context.

### 432 6.1.2 Tseitin Formulas

433 These formulas have already been described in Section 5.1. In this section, we consider  
 434 unsatisfiable formulas generated by CNFgen [24]. We started generating instances with grids  
 435 of size  $n \times n$  and ran all systems with a time limit of 300 seconds in order to analyze how  
 436 different solvers scale. The performance of all systems can be seen on the left plot in the  
 437 first row of Figure 6 (note the logarithmic scale on the  $y$  axis). A point  $(x, y)$  in a solver line  
 438 indicates that the solver took  $y$  seconds to solve the Tseitin formula on grid size  $x \times x$ .

439 These formulas are easy for Cryptominisat, since the application of Gaussian elimination  
 440 in a preprocessing step solves them. Regarding the rest of the systems, only our DIP-based  
 441 tool and GlucoseER have good scaling properties. We want to remark that this behavior  
 442 is not unique to Tseitin formulas on the grid. We generated Tseitin formulas over random  
 443 4-regular and 6-regular graphs and the conclusions are the same as can be seen in the other  
 444 two plots in the first row of Figure 6. These results start suggesting that, even though  
 445 GlucoseER and our DIP-based learning are quite different approaches to integrating ER into



■ **Figure 6** Performance analysis of different solvers on Tseitin formulas.

CDCL solvers, they seem to perform well on the same sets of benchmarks.


In order to understand whether the behavior of xMapleLCM-DIP was polynomial, we generated more challenging instances. Results can be seen in the left plot on the second row of Figure 6, with logarithmic scale on the  $y$  axis. There is little doubt that the runtime of our solver ends up being exponential w.r.t. the size of the problem. However, we went one step further and studied how large were the generated DRAT proofs. In the same row, but on the right, we show the number of resolution steps in the proof. Despite we have no theoretical support for that, it seems that although the solver takes exponential time in finding a proof, its size might be polynomial w.r.t. the problem size. That would indicate that our problem is having search heuristics that are not good enough to quickly find a short proof.

Finally, as we did in the previous section, we want to study the effect of the possible variants of DIP-based learning. For that purpose, we selected 5 challenging benchmarks of each type (grid, 4-regular, 6-regular) and evaluated the 7 DIP variants that we want to examine. Results can again be found in Table 1.

For Tseitin formulas, **occ5** performs much better than the rest. Its only difference is that the minimum number of occurrences for a DIP to be introduced is smaller; this suggests that, in these instances, being more aggressive in the introduction of DIP extension variables provides a competitive advantage. On the other hand, **glue** solves no instance. We realized that DIP-learning was hardly ever applied, that is, requiring glue clauses was too restrictive here. Regarding **act**, whose performance is similarly poor, the main noticeable difference in its behavior is exhibiting a very large number of conflicts where DIP learning is applied, between 20 and 50%, which is around ten times more than the variants that performed best.

### 6.1.3 (X)Orified Random $k$ -XOR Formulas

On last family where DIP-based systems perform very well are random  $k$ -xor formulas, where orification or xorification [7] has been applied, i.e., replacing variables by (x)ors of fresh variables. We used CNFgen to obtain 4 sets of formulas: xor constraints of length 3 applying orification with 2 and 3 variables, and also with xorification of 2 and 3 variables. All formulas we consider are unsatisfiable, the number of variables before (x)orification is equal to the number of clauses, and we increase such number to get progressively more difficult formulas.



■ **Figure 7** Performance analysis of different solvers on (x)orified random  $k$ -xor formulas.

475 We studied the scaling properties of all systems by running them on increasingly larger  
 476 benchmarks with a time limit of 300 seconds. Results can be seen in Figure 7, where the  
 477  $y$  axis has logarithmic scale. Maybe surprisingly, Cryptominisat can only benefit from its  
 478 preprocessing step in one of the families (*2-xorified*). Also, SBVA-CaDiCaL performs very well  
 479 in two of them (*2 and 3-orified*). Again, both GlucoseER and our DIP-based implementations  
 480 show the best average overall performance over all competitors.

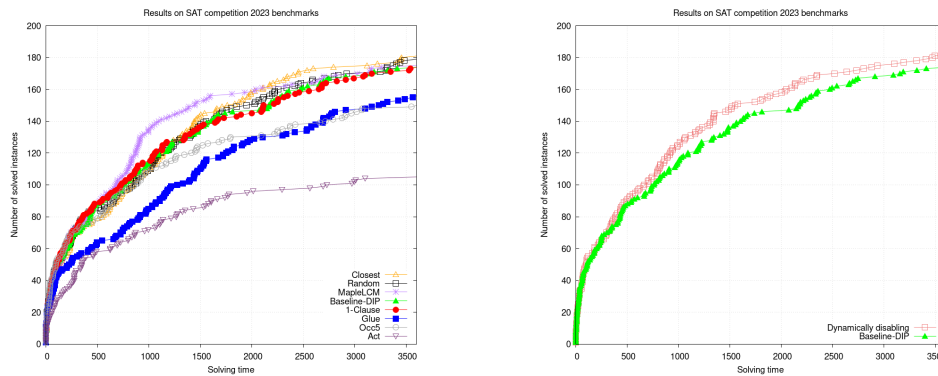
481 Finally, an analysis of the impact of the different DIP-learning variants can be found  
 482 in Table 1. Our baseline implementation seems to perform the best, but there are several  
 483 other configurations that perform very well. On the negative side, **glue** did not solve any  
 484 such formula within the time limit. The identifying characteristic of that version is that the  
 485 percentage of conflicts with DIP-learning is very low compared to the others (around 0.1%).

## 486 6.2 Performance Analysis on 2023 SAT Competition Benchmarks

487 Our experimental evaluation concludes with the lessons we have learned from executing our  
 488 DIP-variants<sup>4</sup> on all benchmarks from the 2023 SAT Competition. We first report on the  
 489 overhead caused by the DIP detection algorithm and the subsequent additional work to  
 490 retrieve the clauses to be learned. For our baseline DIP-based system, where two clauses are  
 491 learned and hence is the most computationally demanding method, in 4.5% of the benchmarks  
 492 the DIP-related work represented between 10 and 15% of the total runtime; in 11.5% of  
 493 the benchmarks between 5 and 10%; in 35.5% between 2.5 and in 48.5% less than 2.5%.  
 494 These data show that DIP computation does not significantly slow down the solver. Another  
 495 interesting information concerns the percentage of conflicts where there is at least one DIP,  
 496 which was on average 63%. This implies that we do need to have filtering mechanisms to  
 497 discard some of them. Otherwise, the search would be totally dominated by DIPs.

498 A cactus plot with the results of executing all our DIP variants on the 2023 SAT  
 499 Competition benchmarks with a time limit of 1 hours is in the left plot of Figure 8. Three

<sup>4</sup> We do not present a comparison with state-of-the-art solvers like Kissat or CaDiCaL since there are not many lessons to be learned. In a few words, our MapleLCM baseline CDCL solver is slower than those systems, and adding DIP-reasoning to it does not bridge this gap.



■ **Figure 8** Performance of DIP variants on 2023 SAT Competition benchmarks.

500 variants are much worse than the rest: **act**, **occ5** and **glue**, whereas the other are relatively  
 501 similar. These three configurations have a much larger percentage of conflicts where DIP-  
 502 based learning is applied than the rest, being **act** around 30% and the other two around 5%.  
 503 In the best performing variant this is around 1%, which seems to be a very low number and  
 504 one might wonder whether this is enough to have any effect at all. This can be answered  
 505 by analyzing the percentage of decisions on extended variables. For all benchmarks where  
 506 DIP-based systems outperform the baseline CDCL solver, this number is surprisingly large,  
 507 being usually over 10% of the decisions, whereas for most of the benchmarks is very low (in  
 508 70% of the benchmarks less than 1% of decisions are on extended variables). Moreover, this  
 509 does not only happen in problems with a very low number of initial variables, where the  
 510 introduction of a few extended variables could quickly dominate the decision heuristic.

511 This is remarkable since it shows that already existing decision heuristics like VSIDS  
 512 or LRB somehow infer whether the newly introduced variables improve the behavior of  
 513 the system. This is why we implemented, on top of our baseline DIP system, a procedure  
 514 that computes the percentage of decisions on extended variables. If after a given number of  
 515 conflicts it is still lower than 3%, it discontinues DIP learning and performs 1-UIP learning  
 516 from that moment onwards. This is a very preliminary step in the direction of trying to  
 517 automate the decision of whether to apply DIP reasoning or not, but the outcome, which is  
 518 found in the right plot of Figure 8, is very positive.

## 519 7 Conclusions and Future Work

520 We introduce a novel extended resolution clause learning (ERCL) algorithm, that when  
 521 implemented on top of a CDCL solver, turns out to be beneficial for a variety of problems,  
 522 in particular Tseitin, random  $k$ -xor and interval matching formulas. This is remarkable  
 523 since automating powerful proof systems is well-known to be a difficult task. Somewhat  
 524 surprisingly, the only previously existing attempt to incorporate ER into CDCL performs  
 525 similarly on the same instances. Considering the different nature of the two methods, this  
 526 clearly deserves further study. Further, we also introduce a new heuristic that allows our  
 527 ERCL solver xMapleSAT to perform similarly to the baseline CDCL solver MapleLCM,  
 528 thus being able to get the best of both worlds, i.e., the benefit of ERCL, without sacrificing  
 529 performance of the CDCL solver on, say, SAT Competition Main Track 2023 instances.  
 530 As future work, we plan to investigate a variety of machine learning based heuristics (e.g.,  
 531 branching) specialized for the ERCL method. Also, the use of  $k$ -IPs with  $k > 2$  is part of  
 532 our next steps. On the theoretical side, challenging questions like determining whether DIP-  
 533 or  $k$ -IP based ER simulates unrestricted ER are going to be central to our research efforts.

## 534 — References

- 535 1 Albert Atserias, Johannes Klaus Fichte, and Marc Thurley. Clause-learning algorithms with  
536 many restarts and bounded-width resolution. *Journal of Artificial Intelligence Research*,  
537 40:353–373, January 2011. Preliminary version in *SAT '09*.
- 538 2 Albert Atserias, Johannes Klaus Fichte, and Marc Thurley. Clause-learning algorithms  
539 with many restarts and bounded-width resolution. *J. Artif. Intell. Res.*, 40:353–373, 2011.  
540 doi:10.1613/jair.3152.
- 541 3 Gilles Audemard, George Katsirelos, and Laurent Simon. A restriction of extended resolution  
542 for clause learning sat solvers. In *AAAI*, 2010.
- 543 4 Gilles Audemard and Laurent Simon. Predicting learnt clauses quality in modern SAT solvers.  
544 In Craig Boutilier, editor, *IJCAI 2009, Proceedings of the 21st International Joint Conference  
545 on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, pages 399–404, 2009.  
546 URL: <http://ijcai.org/Proceedings/09/Papers/074.pdf>.
- 547 5 Fahiem Bacchus. Enhancing Davis Putnam with extended binary clause learning. In *Proc. 18th  
548 Annual Conference on Artificial Intelligence and 15th Conference on Innovative Applications  
549 of Artificial Intelligence*, pages 613–619. AAAI Press and MIT Press, 2002.
- 550 6 Tomas Balyo, Marijn Heule, Markus Iser, Matti Järvisalo, and Martin Suda, editors. *Pro-  
551 ceedings of SAT Competition 2023: Solver, Benchmark and Proof Checker Descriptions*.  
552 Department of Computer Science Series of Publications B. Department of Computer Science,  
553 University of Helsinki, Finland, 2023.
- 554 7 Eli Ben-Sasson. Size space tradeoffs for resolution. *SIAM Journal on Computing*, 38(6):2511–  
555 2525, 2009.
- 556 8 Armin Biere, Katalin Fazekas, Mathias Fleury, and Maximillian Heisinger. CaDiCaL, Kissat,  
557 Paracoba, Plingeling and Treengeling entering the SAT Competition 2020. In Tomas Balyo,  
558 Nils Froylyks, Marijn Heule, Markus Iser, Matti Järvisalo, and Martin Suda, editors, *Proc. of  
559 SAT Competition 2020 – Solver and Benchmark Descriptions*, volume B-2020-1 of *Department  
560 of Computer Science Report Series B*, pages 51–53. University of Helsinki, 2020.
- 561 9 Armin Biere, Nils Froylyks, and Wenzhi Wang. CadiBack: Extracting backbones with CaDiCal.  
562 In *Proc. 26th Intl. Conf. on Theory and Applications of Satisfiability Testing SAT*, pages  
563 3:1–12. Leibniz-Zentrum für Informatik, 2023.
- 564 10 Randal E. Bryant and Marijn J. H. Heule. Generating extended resolution proofs with a bdd-  
565 based SAT solver. *ACM Trans. Comput. Log.*, 24(4):31:1–31:28, 2023. doi:10.1145/3595295.
- 566 11 Sam Buss, Vijay Ganesh, and Albert Oliveras. A linear time algorithm for two-vertex  
567 bottlenecks. In preparation. Preliminary version available at [https://math.ucsd.edu/~sbuss/  
568 ResearchWeb/TwoVertBottlenecks](https://math.ucsd.edu/~sbuss/ResearchWeb/TwoVertBottlenecks), March 2024.
- 569 12 Sam Buss, Oliver Kullmann, and Virginia Vassilevska Williams. Dual depth first search for  
570 binary clause reasoning. Preliminary version, March 2024.
- 571 13 Stephen A. Cook. A short proof of the pigeon hole principle using extended resolution.  
572 *SIGACT News*, 8(4):28–32, 1976. doi:10.1145/1008335.1008338.
- 573 14 Andrew Haberlandt, Harrison Green, and Marijn J. H. Heule. Effective auxiliary variables via  
574 structured reencoding. In Meena Mahajan and Friedrich Slivovsky, editors, *26th International  
575 Conference on Theory and Applications of Satisfiability Testing, SAT 2023, July 4-8, 2023,  
576 Alghero, Italy*, volume 271 of *LIPICs*, pages 11:1–11:19. Schloss Dagstuhl - Leibniz-Zentrum  
577 für Informatik, 2023. URL: <https://doi.org/10.4230/LIPICs.SAT.2023.11>, doi:10.4230/  
578 LIPICs.SAT.2023.11.
- 579 15 Armin Haken. The intractability of resolution. *Theoretical computer science*, 39:297–308, 1985.
- 580 16 Marijn J. H. Heule, Benjamin Kiesl, and Armin Biere. Short proofs without new variables. In  
581 Leonardo de Moura, editor, *Automated Deduction - CADE 26 - 26th International Conference  
582 on Automated Deduction, Gothenburg, Sweden, August 6-11, 2017, Proceedings*, volume  
583 10395 of *Lecture Notes in Computer Science*, pages 130–147. Springer, 2017. doi:10.1007/  
584 978-3-319-63046-5\_9.



- 585 17 Marijn J. H. Heule, Benjamin Kiesl, and Armin Biere. Encoding redundancy for satisfaction-  
586 driven clause learning. In *Tools and Algorithms for the Construction and Analysis of Systems*  
587 - *25th International Conference (TACAS 2019), Part I*, Lecture Notes in Computer Science  
588 11427, pages 41–58. Springer Verlag, 2019.
- 589 18 Marijn J. H. Heule, Benjamin Kiesl, and Armin Biere. Strong extension-free proof systems. *J.*  
590 *Autom. Reason.*, 64(3):533–554, 2020. doi:10.1007/s10817-019-09516-0.
- 591 19 Marijn J. H. Heule, Benjamin Kiesl, Martina Seidl, and Armin Biere. Pruning through  
592 satisfaction. In Ofer Strichman and Rachel Tzoref-Brill, editors, *Hardware and Software:*  
593 *Verification and Testing - 13th International Haifa Verification Conference, HVC 2017, Haifa,*  
594 *Israel, November 13-15, 2017, Proceedings*, volume 10629 of *Lecture Notes in Computer Science*,  
595 pages 179–194. Springer, 2017. doi:10.1007/978-3-319-70389-3\_12.
- 596 20 Markus Iser and Carsten Sinz. A problem meta-data library for research in sat. In Daniel Le  
597 Berre and Matti J"arvisalo, editors, *Proceedings of SAT 2015 and 2018*,  
598 volume 59 of *EPiC Series in Computing*, pages 144–152. EasyChair, 2019. URL: <https://easychair.org/publications/paper/jQXv>, doi:10.29007/gdbb.
- 600 21 Henry A. Kautz. Deconstructing planning as satisfiability. In *Proceedings, The Twenty-First*  
601 *National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of*  
602 *Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*, pages 1524–  
603 1526. AAAI Press, 2006. URL: <http://www.aaai.org/Library/AAAI/2006/aaai06-241.php>.
- 604 22 Henry A. Kautz and Bart Selman. Planning as satisfiability. In Bernd Neumann, editor, *10th*  
605 *European Conference on Artificial Intelligence, ECAI 92, Vienna, Austria, August 3-7, 1992.*  
606 *Proceedings*, pages 359–363. John Wiley and Sons, 1992.
- 607 23 Daniel Kroening. Software verification. In Armin Biere, Marijn Heule, Hans van Maaren,  
608 and Toby Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers*  
609 *in Artificial Intelligence and Applications*, pages 791–818. IOS Press, 2021. doi:10.3233/  
610 FAIA201004.
- 611 24 Massimo Lauria, Jan Elffers, Jakob Nordström, and Marc Vinyals. Cnfgcn: A generator of  
612 crafted benchmarks. In Serge Gaspers and Toby Walsh, editors, *Theory and Applications of*  
613 *Satisfiability Testing - SAT 2017 - 20th International Conference, Melbourne, VIC, Australia,*  
614 *August 28 - September 1, 2017, Proceedings*, volume 10491 of *Lecture Notes in Computer*  
615 *Science*, pages 464–473. Springer, 2017. doi:10.1007/978-3-319-66263-3\_30.
- 616 25 Mao Luo, Chu-Min Li, Fan Xiao, Felip Manyà, and Zhipeng Lü. An effective learnt clause  
617 minimization approach for CDCL SAT solvers. In Carles Sierra, editor, *Proceedings of the*  
618 *Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne,*  
619 *Australia, August 19-25, 2017*, pages 703–711. ijcai.org, 2017. URL: [https://doi.org/10.](https://doi.org/10.24963/ijcai.2017/98)  
620 [24963/ijcai.2017/98](https://doi.org/10.24963/ijcai.2017/98), doi:10.24963/IJCAI.2017/98.
- 621 26 Norbert Manthey, Marijn Heule, and Armin Biere. Automated reencoding of boolean formulas.  
622 In Armin Biere, Amir Nahir, and Tanja E. J. Vos, editors, *Hardware and Software: Verification*  
623 *and Testing - 8th International Haifa Verification Conference, HVC 2012, Haifa, Israel,*  
624 *November 6-8, 2012. Revised Selected Papers*, volume 7857 of *Lecture Notes in Computer*  
625 *Science*, pages 102–117. Springer, 2012. doi:10.1007/978-3-642-39611-3\_14.
- 626 27 João Marques-Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers.  
627 In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of*  
628 *Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*,  
629 pages 133–182. IOS Press, 2021. doi:10.3233/FAIA200987.
- 630 28 Karl Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10:96–115, 1927.
- 631 29 Sibyle Möhle and Armin Biere. Backing backtracking. In *22nd Intl Conf. on Theory and*  
632 *Applications of Satisfiability Testing (SAT)*, Lecture Notes in Computer Science 11628, pages  
633 250–266. Springer, 2019.
- 634 30 Alexander Nadel and Vadim Ryvchin. Chronological backtracking. In *21st Intl Conf. on*  
635 *Theory and Applications of Satisfiability Testing (SAT)*, Lecture Notes in Computer Science  
636 10929, pages 111–121. Springer, 2018.

- 637 31 Albert Oliveras, Chunxiao Li, Darryl Wu, Jonathan Chung, and Vijay Ganesh. Learning  
638 shorter redundant clauses in SDCL using maxsat. In Meena Mahajan and Friedrich Slivovsky,  
639 editors, *26th International Conference on Theory and Applications of Satisfiability Testing,*  
640 *SAT 2023, July 4-8, 2023, Alghero, Italy*, volume 271 of *LIPICs*, pages 18:1–18:17. Schloss  
641 Dagstuhl - Leibniz-Zentrum für Informatik, 2023. URL: <https://doi.org/10.4230/LIPICs.SAT.2023.18>, doi:10.4230/LIPICs.SAT.2023.18.
- 643 32 Knot Pipatsrisawat and Adnan Darwiche. On modern clause-learning satisfiability solvers.  
644 *Journal of Automated Reasoning*, 44(3):277–301, 2010.
- 645 33 Knot Pipatsrisawat and Adnan Darwiche. On the power of clause-learning SAT solvers as  
646 resolution engines. *Artif. Intell.*, 175(2):512–525, 2011. doi:10.1016/j.artint.2010.10.002.
- 647 34 Joseph E. Reeves, Marijn J. H. Heule, and Randal E. Bryant. Preprocessing of propagation  
648 redundant clauses. In Jasmin Blanchette, Laura Kovács, and Dirk Pattinson, editors, *Automated*  
649 *Reasoning - 11th International Joint Conference, IJCAR 2022, Haifa, Israel, August 8-10, 2022,*  
650 *Proceedings*, volume 13385 of *Lecture Notes in Computer Science*, pages 106–124. Springer,  
651 2022. doi:10.1007/978-3-031-10769-6\_8.
- 652 35 Karem A. Sakallah. Symmetry and satisfiability. In Armin Biere, Marijn Heule, Hans  
653 van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume  
654 336 of *Frontiers in Artificial Intelligence and Applications*, pages 509–570. IOS Press, 2021.  
655 doi:10.3233/FAIA200996.
- 656 36 Mate Soos, Karsten Nohl, and Claude Castelluccia. Extending SAT solvers to cryptographic  
657 problems. In Oliver Kullmann, editor, *Theory and Applications of Satisfiability Testing -*  
658 *SAT 2009, 12th International Conference, SAT 2009, Swansea, UK, June 30 - July 3, 2009.*  
659 *Proceedings*, volume 5584 of *Lecture Notes in Computer Science*, pages 244–257. Springer,  
660 2009. doi:10.1007/978-3-642-02777-2\_24.
- 661 37 Grigori S Tseitin. On the complexity of derivation in propositional calculus. *Automation of*  
662 *reasoning: 2: Classical papers on computational logic 1967–1970*, pages 466–483, 1983.
- 663 38 Alasdair Urquhart. Hard examples for resolution. *J. ACM*, 34(1):209–219, 1987. doi:  
664 10.1145/7531.8928.
- 665 39 Yichen Xie and Alexander Aiken. Saturn: A SAT-Based Tool for Bug Detection. In *Proceedings*  
666 *of the 17th International Conference on Computer Aided Verification, CAV 2005*, pages 139–143,  
667 2005. doi:10.1007/11513988\_13.
- 668 40 Lintao Zhang, Conor F. Madigan, Matthew W. Moskewicz, and Sharad Malik. Efficient  
669 conflict driven learning in boolean satisfiability solver. In Rolf Ernst, editor, *Proceedings of*  
670 *the 2001 IEEE/ACM International Conference on Computer-Aided Design, ICCAD 2001,*  
671 *San Jose, CA, USA, November 4-8, 2001*, pages 279–285. IEEE Computer Society, 2001.  
672 doi:10.1109/ICCAD.2001.968634.