

# 1 Certified Branch-and-Bound MaxSAT Solving

2 **Dieter Vandesande**

3 Vrije Universiteit Brussel, Belgium

4 **Jordi Coll**

5 Universitat de Girona, Spain

6 **Chu-Min Li**

7 Université de Picardie Jules Verne, France

8 **Bart Bogaerts**

9 Vrije Universiteit Brussel, Belgium

## 10 — Abstract —

11 Over the past few decades, there has been a remarkable improvement in the performance of  
12 combinatorial solvers, which has led to their practical usage in real-world applications. In some of  
13 these applications, the correctness of the solver's result is of utmost importance. Unfortunately, the  
14 reality is different: the algorithmic complexity of modern solvers enables bugs to sneak into the  
15 source code. For Satisfiability Checking (SAT), this problem was mitigated by letting the solver write  
16 down a formal proof of correctness of the obtained solution, which is also known as proof logging.  
17 However, for more expressive fields such as MaxSAT, which is the optimization variant of SAT, proof  
18 logging had not yet had its breakthrough until recently, when the proof system VeriPB was put  
19 forward as a good candidate to serve as a general-purpose proof system for MaxSAT solvers. In this  
20 paper, we show how VeriPB can be used as a proof system to let the Branch-and-Bound MaxSAT  
21 solver MaxCDCL produce proofs of optimality for its solutions. This is the first state-of-the-art  
22 MaxSAT solver that implements the Branch-and-Bound approach, opposed to the SAT-based solvers  
23 with proof logging in earlier work. We also show how to use VeriPB to add proof logging for an  
24 encoding of the model improving constraint into CNF based on Multi-valued Decision Diagrams  
25 (MDD).

26 **2012 ACM Subject Classification** Theory of computation → Logic

27 **Keywords and phrases** Combinatorial optimization, maximum satisfiability, branch-and-bound,  
28 certifying algorithms, proof logging

29 **Digital Object Identifier** 10.4230/LIPIcs.Soft.2024.2024.

30 **Funding** This work was partially supported by Fonds Wetenschappelijk Onderzoek – Vlaanderen  
31 (project G070521N). This work is also supported by grant PID2021-122274OB-I00 funded by grant  
32 MICIU/AEI/10.13039/501100011033 and by ERDF, EU.

## 33 **1 Introduction**

34 With ever increasingly efficient solvers being developed in various areas concerned with  
35 combinatorial search and optimization, we have now effectively arrived in a situation where  
36 NP hard problems are routinely tackled in practice. This maturation has resulted in solvers  
37 being deployed in various practical use cases, including safety-critical applications or making  
38 life-affecting decisions (e.g., verifying software that drives our transportation infrastructure  
39 [16], checking correctness of plans for the operation of the reaction control system of the  
40 space shuttle [32], or matching donors and recipients for kidney transplants [29]). For this  
41 reason, it is of utmost importance that the results produced by such solvers are guaranteed  
42 to be correct. Unfortunately, this is not always the case; in fact, there have been numerous  
43 reports of solvers outputting infeasible solutions, or falsely claiming optimality or the absence  
44 of solutions [2, 8, 9, 13, 18, 26].



© Dieter Vandesande, Jordi Coll, Chu-Min Li, Bart Bogaerts;  
licensed under Creative Commons License CC-BY 4.0

First International Workshop on Discrete Optimization with Soft Constraints.

Editors: Simon de Givry and Javier Larrosa; Article No. ; pp. :1–:15

Leibniz International Proceedings in Informatics



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

45 This problem calls for a systematic solution. The most evident solution would be to use  
 46 *formal verification*; that is, using a proof assistant [3, 31, 34, 15] to prove that a solver is  
 47 correct. However, in practice, using formal verification means sacrificing performance [17, 22],  
 48 which is precisely what has led to the success of combinatorial optimization.

49 Instead, what we believe to be the way forward is to use *certifying algorithms* [30], an  
 50 idea that is also known as *proof logging* in the setting of combinatorial optimization. With  
 51 proof logging, the solver at hand does not only produce an answer (e.g., an optimal solution  
 52 for optimization problems), but also a *proof of correctness* of this answer that can easily (in  
 53 terms of the size of the proof) be verified by an independent tool (known as the *proof checker*).  
 54 Next to guaranteeing correctness, proof logging is also very useful as a software development  
 55 methodology: it facilitates advanced *testing* and *debugging* of solver code. Moreover, the  
 56 produced proofs can be seen as an auditable record of how and why a certain conclusion was  
 57 reached.

58 Proof logging has been pioneered in the field of Boolean satisfiability (SAT), where  
 59 numerous proof formats and proof checkers (including formally verified checkers) have seen  
 60 the light of day [4, 21, 23, 37, 14], with a breakthrough moment when it was used to resolve  
 61 the Pythagorean triple problem, resulting in the “largest math proof ever” [24]. Moreover,  
 62 for years, proof logging has been mandatory in (the main tracks) of the yearly solving  
 63 competition, reflecting the fact that the community effectively demands that all SAT solvers  
 64 be certifying.

65 In this paper, we are concerned with *maximum satisfiability* (MaxSAT), the optimization  
 66 variant of SAT. In MaxSAT, proof logging is less well-spread. Only recently was the VERIPB  
 67 proof system proposed as a general-purpose proof logging methodology for MaxSAT [36].  
 68 VERIPB has been successfully applied for adding proof logging to MaxSAT solvers employing  
 69 *solution-improving search* [36, 35] (where a SAT oracle is repeatedly queried each time  
 70 searching for a solution that improves upon the previously best found solution) as well as for  
 71 *core-guided search* (where a SAT oracle is repeatedly queried under ever-relaxing optimistic  
 72 assumptions). Next to these two search paradigms, in MaxSAT also *implicit hitting set* and  
 73 *branch-and-bound* are used in state-of-the-art MaxSAT solvers.

74 In this paper, we are concerned with bringing VERIPB-based certification to branch-and-  
 75 bound search [28], thereby covering another search paradigm for MaxSAT solving. Modern  
 76 branch-and-bound solvers combine conflict-driven clause learning (CDCL) [33] with a clever  
 77 bounding function that determines whether the current search branch is “hopeless”, meaning  
 78 that there are no assignments that improve upon the best solution found so far. Since it  
 79 is well-known how to certify CDCL search with VERIPB, the main challenge is to certify  
 80 the conclusions of the bounding function. This bounding function makes use of look-ahead  
 81 search to generate (conditional) unsatisfiable cores (literals that cannot be true together).  
 82 These cores are then combined to get an estimate of the best possible objective value that is  
 83 still achievable. Especially in the unweighted case, this combination of cores relies on very  
 84 subtle arguments, which we clearly spell out and formalize in the VERIPB proof system.

85 In order to bring this certification to practice, there are more hurdles to overcome:  
 86 state-of-the-art branch-and-bound solvers employ several other clever tricks to speed up the  
 87 solving process, such as pre-processing methods as well as integrating ideas from other search  
 88 paradigms. One particular technique that proved to be challenging is the use of *multi-valued*  
 89 *decision diagrams* (MDDs) [6] in order to create a CNF encoding of a solution-improving  
 90 constraint (which is enabled or disabled heuristically depending on the size of the instance at  
 91 hand). The MDD-based encoding used by MaxCDCL generalizes the encoding of Binary  
 92 Decision Diagrams (BDDs) (by allowing splits on sets of variables instead of a single variable)

93 proposed in [1]. From the perspective of proof logging, the main challenge here is to prove  
 94 that these constraints are indeed equivalent. As an example, consider the constraints

$$95 \quad 12x_1 + 5x_2 + 4x_3 \geq 6$$

96 and

$$97 \quad 12x_1 + 5x_2 + 4x_3 \geq 9.$$

98 Taking into account that the variables take values in  $\{0, 1\}$ , it is not hard to see that these  
 99 constraints are equivalent: all combinations of truth assignments that lead to the left-hand  
 100 side taking a value at least six, must assign it a value of at least nine. In other words: the  
 101 left-hand side cannot take values 6, 7, or 8. In general, checking whether such a linear  
 102 expression can take a specific value (e.g., 7) is well-known to be NP hard, but BDD (and  
 103 MDD) construction algorithms will often detect this efficiently (in the worst-case this can  
 104 take exponential time, which is why such MDD-based encodings will only be enabled for  
 105 constraints with few variables and small coefficients). The main question of interest for us is  
 106 how to convince a proof checker of the fact that these two constraints are equivalent, without  
 107 doing a substantial amount of additional work. We achieve this using an algorithm that  
 108 proves this property for all nodes in an MDD in a linear pass over its representation.

109 We implemented some of our algorithms in MAXCDCL, which is the only branch-and-  
 110 bound solver that participated in last year's MaxSAT evaluation [25], where it proved to  
 111 be one of the top performing solvers (it won the weighted track and ended second on the  
 112 unweighted track). We intend to experimentally evaluate this work, but currently it is still  
 113 work in progress: while the theory has been worked out and a partial implementation is  
 114 available, we have no concrete results to report yet.

115 The rest of this paper is structured as follows. In Section 2, we recall some preliminaries  
 116 about MaxSAT solving and VERIPB-based proof logging. Section 3 is devoted to presenting  
 117 the core of the MAXCDCL algorithm, focusing on look-ahead-based bounding, as well as  
 118 explaining how to get a certifying version of this. In Section 4, we explain how BDDs and  
 119 MDDs are used to encode PB constraints. Section 5 concludes the paper. Formal details  
 120 and proofs are often omitted, but can be found in the supplementary material.

## 121 **2 Preliminaries**

122 We first recall some concepts from pseudo-Boolean optimization and MaxSAT solving.  
 123 Afterwards, we introduce the VERIPB proof system. For a full exposition, we refer the reader  
 124 to [11, 27, 7].

125 In this paper, all variables are assumed to be *Boolean*; meaning they take a value in  
 126  $\{0, 1\}$ . A *literal*  $\ell$  is a Boolean variable  $x$  or its negation  $\bar{x}$ . A *pseudo-Boolean constraint*  $C$  is  
 127 a 0–1 integer linear inequality  $\sum_i w_i \ell_i \geq A$ . Without loss of generality, we will often assume  
 128 our constraints to be *normalized*, meaning that the  $\ell_i$  are different literals and all coefficients  
 129  $w_i$  and the degree  $A$  are non-negative. A *formula* is a conjunction of PB constraints. A  
 130 *clause* is a special case of a PB constraint having all  $w_i$  and  $A$  equal to one. A *cardinality*  
 131 *constraint* is a PB constraint where all  $w_i$  are one. If  $L$  is a set of literals, we will sometimes  
 132 simply write  $L$  as a shorthand for  $\sum_{\ell \in L} \ell$  and thus write constraints such as  $L + 3K \geq 42$ ,  
 133 meaning  $\sum_{\ell \in L} \ell + \sum_{\ell \in K} 3 \cdot \ell \geq 42$ .

134 A (*partial*) *assignment*  $\alpha$  is a (*partial*) function from the set of variables to  $\{0, 1\}$ ; it is  
 135 extended to literals in the obvious way. We write  $C|_\alpha$  for the constraint obtained from  $C$  by  
 136 substituting all assigned variables  $x$  by their assigned value  $\alpha(x)$ . A constraint  $C$  is *satisfied*

## XX:4 Certified Branch-and-Bound MaxSAT Solving

137 under  $\alpha$  if  $\sum_{\alpha(\ell_i)=1} w_i \geq A$ . A formula  $F$  is satisfied under  $\alpha$  if all of its constraints are. We  
138 say that  $F$  *implies*  $C$  (and write  $F \models C$ ) if all assignments that satisfy  $F$  also satisfy  $C$ .

139 A *pseudo-Boolean optimization* instance consists of a formula  $F$  and a linear term  
140  $\mathcal{O} = \sum_i v_i b_i$  (called the *objective*) to be minimized, where  $v_i$  are integers and  $b_i$  are different  
141 literals. If  $\ell$  is a literal, we write  $w_{\mathcal{O}}(\ell)$  for the weight of  $\ell$  in  $\mathcal{O}$ , i.e.,  $w_{\mathcal{O}}(b_i) = v_i$  for all  $i$   
142 and  $w_{\mathcal{O}}(\ell) = 0$  for all other literals. In this paper, we view MaxSAT as the special case of  
143 pseudo-Boolean optimization where  $F$  is a conjunction of clauses.

144 For a pseudo-Boolean optimization instance  $(F, \mathcal{O})$ , the VERIPB proof system maintains  
145 a *proof configuration*  $(\mathcal{C}, \mathcal{D})$  where  $\mathcal{C}$  (standing for *core*) and  $\mathcal{D}$  (standing for *derived*) are  
146 sets of constraints (initialized as  $F$  and  $\emptyset$ , respectively). Constraints can be moved from  $\mathcal{D}$   
147 to  $\mathcal{C}$  but not vice versa. We are allowed to update the configuration using the cutting planes  
148 proof system [12]:

149 **Literal Axioms:** For any literal, we can add  $\ell_i \geq 0$  to  $\mathcal{D}$ .

150 **Linear Combination:** Given two PB constraints  $C_1$  and  $C_2$  in  $\mathcal{C} \cup \mathcal{D}$ , we can add a positive  
151 linear combination of  $C_1$  and  $C_2$  to  $\mathcal{D}$ .

152 **Division:** Given the normalized PB constraint  $\sum_i w_i \ell_i \geq A$  in  $\mathcal{C} \cup \mathcal{D}$  and a positive integer  
153  $c$ , we can add the constraint  $\sum_i \lceil w_i/c \rceil \ell_i \geq \lceil A/c \rceil$  to  $\mathcal{D}$ .

154 VERIPB has some additional rules, which are briefly discussed below. We refer to [7, 20]  
155 for more details on these rules. There is a rule for dealing with optimization statements:

156 **Objective Improvement:** Given an assignment  $\alpha$  that satisfies  $\mathcal{C}$ , we can add the constraint  
157  $\mathcal{O} < \mathcal{O}|_{\alpha}$  to  $\mathcal{C}$ .

158 This rule states that once a solution is found, we search for strictly better solutions. It is  
159 also possible to rewrite the objective:

160 **Objective Reformulation** Given a linear term  $\mathcal{O}'$  in  $\mathcal{O}$ ,  $\mathcal{O}$  can be rewritten by replacing  $\mathcal{O}'$   
161 with  $\mathcal{O}'_{\text{new}}$  if we have shown that  $\mathcal{O}' \geq \mathcal{O}'_{\text{new}}$  and  $\mathcal{O}' \leq \mathcal{O}'_{\text{new}}$ .

162 VERIPB allows deriving non-implied constraints with a generalization of the RAT rule [26]  
163 (which is common in proof systems for SAT). This rule makes use of a *substitution*, which  
164 maps every variable to 0, 1, or a literal. Applying a substitution  $\omega$  on a constraint  $C$  results  
165 in the constraint  $C|_{\omega}$  obtained from  $C$  by replacing each  $x$  by  $\omega(x)$ .

166 **Redundance-based strengthening:** If  $\mathcal{C} \wedge \mathcal{D} \wedge \neg C \models \mathcal{O} \leq \mathcal{O}|_{\omega} \wedge (\mathcal{C} \wedge C)|_{\omega}$ , we add  $C$  to  $\mathcal{C}$  if  
167 the strengthening-to-core mode is enabled and to  $\mathcal{D}$  if it is disabled.

168 Intuitively, this rule can be used to show that  $\omega$ , when applied to assignments instead of  
169 formulas, maps any solution of  $\mathcal{C}$  that does not satisfy  $C$  to a solution of  $\mathcal{C}$  that does satisfy  
170  $C$  and that has an objective value that is at least as good. The most important use case of  
171 the redundance-rule is **reification**: for any PB constraint  $C$  and any fresh variable  $v$ , not  
172 used before, two applications of the redundance rule allow us to derive PB constraints that  
173 express  $v \Rightarrow C$  and  $v \Leftarrow C$ . Finally, VERIPB has rules for deleting constraints in a way that  
174 guarantees that no better-than-optimal value can be found:

175 **Deletion** A constraint in  $\mathcal{D}$  can be deleted at any time. For the deletion of a constraint  
176  $C \in \mathcal{C}$  it has to be shown that  $C$  can be derived by redundance-based strengthening from  
177  $\mathcal{C} \setminus \{C\}$ .

### 3 Certification for MaxCDCL

179 In this section, we present the MaxCDCL algorithm. We start with the core version of the  
180 algorithm and afterwards discuss various extensions to it. With each algorithm/extension,  
181 we immediately discuss what needs to be done to integrate proof logging support for it.

### 3.1 The Core of MaxCDCL

MAXCDCL combines branch-and-bound search and conflict-driven clause learning (CDCL). It maintains the objective value  $v^*$  of the best found solution so far, which is initialized as  $+\infty$ . It performs standard CDCL search (branch). However, when a partial assignment is visited where it is clear that the value of the objective will be at least  $v^*$ , the search is interrupted and a new clause forcing the solver to backtrack is learned (bound).

VERIPB-based certification for CDCL-based SAT solvers is easy to obtain since VERIPB extends the DRAT proof system. Moreover, this has been done several times before [36, 19, 7], making the certification of the *branch* part straightforward. The *bound* part, which lies at the heart of MAXCDCL, on the other hand, has never been certified before. It consists of a clever way of detecting that it is indeed “clear” that all extensions of the current assignment that satisfy  $F$  have an objective value that is too high. To this end, in the general weighted case, given a current assignment  $\alpha$ , it uses a lookahead mechanism that is designed to generate a set  $\mathcal{K}$  of *weighted local cores*: tuples  $(w, K)$  where  $w \in \mathbb{N}$  is the weight of the core, and  $K$  is a core relative to  $\alpha$ : a set of negations of objective literals such that  $F \wedge \alpha \wedge K \models \perp$ , where  $\perp$  denotes the trivially false constraint  $0 \geq 1$ . In other words a local core guarantees that, given the assignments in  $\alpha$ , the underlying objective literals cannot all be false (or in other words: at least one of the underlying objective literals incurs cost). The set  $\mathcal{K}$  should moreover satisfy the properties that

1. For each objective literal  $\ell$ ,

$$\sum_{(w,K) \in \mathcal{K} \wedge \bar{\ell} \in K} w \leq w_{\mathcal{O}}(\ell),$$

i.e., the total weight of all cores containing a literal does not exceed the weight of the underlying objective literal in the objective.

2. The total weight of  $\mathcal{K}$  exceeds the current upper bound:

$$\sum_{(w,K) \in \mathcal{K}} w \geq v^*.$$

When these conditions are satisfied, since in each assignment that extends  $\alpha$  at least one literal in each  $K$  is false, in each such assignment

$$v^* \leq \sum_{(w,K) \in \mathcal{K}} w \leq \sum_{(w,K) \in \mathcal{K}} w \cdot \left( \sum_{\bar{\ell} \in K} \ell \right) = \sum_{\ell} \left( \sum_{(w,K) \in \mathcal{K} \wedge \bar{\ell} \in K} w \right) \cdot \ell \leq \mathcal{O} \quad (1)$$

and we indeed see that the value of the objective in such an assignment can never (strictly) improve upon the currently best-found value. This situation is referred to as a *soft conflict*.

At this point, the solver can learn the clause  $\neg\alpha$ , stating that the current assignment is hopeless, but in practice, from the way the local cores are generated, we get more precise information. The way this is done is as follows. During the lookahead phase, the solver maintains a temporary objective  $\mathcal{O}^t$  that is initialized to be equal to  $\mathcal{O}$ . The set of cores is initialized with trivially falsified unit cores of the form  $(w(\ell), \{\bar{\ell}\})$ , for all objective literals  $\ell$  such that  $\ell \in \alpha$ . These represent the cost incurred by the current assignment  $\alpha$ . Starting from the current assignment  $\alpha$ , in the lookahead phase, all unassigned objective literals  $\ell_1, \ell_2, \dots$  are one at a time falsified (set to their non-cost-incurring polarity) followed by application of unit propagation. As soon as

- either an objective literal  $\ell'$  is propagated to be cost-incurring,

## XX:6 Certified Branch-and-Bound MaxSAT Solving

222 ■ or a conflict is found

223 we can learn a new clause that tells us that under  $\alpha$ , the literals  $\ell_i$  that were falsified (and  
224 in the first case also the literal  $\ell'$ ) cannot be false together. At this point, we can apply  
225 standard conflict analysis techniques to minimize this learned clause  $C_K$  corresponding to  
226 the local core  $K$  we will learn.  $C_K$  consists of (1) a set of negations of literals in  $\alpha$  (denoted  
227  $reason(K)$ ) and (2) a set of objective literals, where  $K$  is the set of the negations of these  
228 literals. For the initial cores of the form  $(w(\ell), \{\bar{\ell}\})$ , we define  $reason(K) = \{\bar{\ell}\}$ . The  
229 literals in  $reason(K)$  are called the reason for  $K$  since they tell us precisely which part  
230 of the current assignment was relevant for deriving the local core  $K$ . The weight  $w_K$  of  
231 cores  $K$  is chosen to be the  $\min_{\bar{\ell} \in K} w_{\mathcal{O}^t}(\ell)$  and  $\mathcal{O}^t$  is rewritten to  $\mathcal{O}^t - w_K \sum_{\bar{\ell} \in K} \ell$ . The  
232 lookahead then starts again from scratch, but this time only setting variables still in  $\mathcal{O}^t$  to  
233 false (thus excluding objective literals whose weights have been completely “consumed” by  
234 the previous core). This procedure is guaranteed to indeed produce a set of local cores that  
235 satisfy condition 1 above. If condition 2 is not satisfied (i.e., the total weight of the produced  
236 cores does not exceed the upper bound), then MAXCDCL continues search as if nothing  
237 happened. Otherwise, it learns a clause  $C_{\mathcal{K}} \subseteq \neg\alpha$ , which is

$$238 \quad \bigcup_{(w,K) \in \mathcal{K}} reason(K).$$

240 Indeed, the literals in  $\neg C_{\mathcal{K}}$  are the literals in the current assignment that prevent the  
241 solver from finding a (strictly) improving solution. After applying standard conflict analysis  
242 techniques on the clause  $C_{\mathcal{K}}$  for further simplification, it is learned by the solver, the solver  
243 backtracks and continues its search.

### 244 Hardening

245 If at some point, we end up in a state where for some literal  $\ell$  currently still unassigned,

$$246 \quad w_{\mathcal{O}^t}(\ell) + \sum_{(w,K) \in \mathcal{K}} w \geq v^*,$$

247 then a clause can be learned that propagates  $\ell$  to be false. The reasoning for this is that in  
248 the assignment  $\alpha \cup \ell$ , we would have found a soft conflict. From the proof logging perspective,  
249 hardening does not require special treatment, so in what follows we focus on soft conflicts.

### 250 Proof-Logging for Lookahead-Based Bounding with Soft Conflicts

251 The VERIPB proof of lookahead-based bounding we produce is heavily inspired by equation  
252 (1). Essentially, what happens is that we will reproduce this derivation inside the VERIPB  
253 proof, and as such we will derive that  $\alpha \Rightarrow \mathcal{O} \geq v^*$ . Combining this with the model-improving  
254 constraint then allows to derive a clause that excludes the current assignment  $\alpha$ . All further  
255 minimizations that happen to this clause are done using standard conflict analysis (i.e., using  
256 resolution and hence, can be de done using explicit cutting planes derivations).

### 257 3.2 Literal Unlocking for Cardinality Reasoning

258 In the unweighted case, during the lookahead-based bounding a smarter mechanism is  
259 employed. In this case, the objective to minimize is of the form

$$260 \quad \sum_{\ell \in objLits} \ell.$$



261 Given a previous upper bound  $v^*$ , in this case the lookahead procedure described above would  
 262 yield a set of  $v^*$  disjoint cores. Instead of searching for a set of disjoint cores, MAXCDCL  
 263 would in this situation search for a set of disjoint *cardinality constraints*: a set  $\mathcal{L}$  of tuples  
 264  $(b, L)$  such that:

- 265 1. Each  $L$  is a set of objective literals
  - 266 2. For each  $(b, L)$  in  $\mathcal{L}$ , it holds that  $F \wedge \alpha \models L \geq b$  (we remind the reader that we write  
 267  $L \geq b$  for  $\sum_{\ell \in L} \ell \geq b$ ).
  - 268 3. For each pair  $(b, L)$  and  $(b', L')$  in  $\mathcal{L}$ ,  $L \cap L' = \emptyset$ .
  - 269 4. The total weight exceeds the current upper bound:  $\sum_{(b, L) \in \mathcal{L}} b \geq v^*$ .
- 270 If these four conditions are satisfied, we can see that indeed

$$271 \quad F \wedge \alpha \models \sum_{\ell \in \text{objLits}} \ell \geq \sum_{(b, L) \in \mathcal{L}, \ell \in L} \ell \geq \sum_{(b, L) \in \mathcal{L}} b \geq v^*.$$

272 In other words, each assignment more precise than  $\alpha$  will have an objective value that does  
 273 not improve upon the best found so far and we can learn a clause that excludes the current  
 274 assignment  $\alpha$ .

275 The construction of such a set of disjoint cardinality constraints is achieved building on  
 276 the following result.

277 ► **Definition 1.** Let  $(b, L)$  be a cardinality constraint,  $U$  a subset of  $L$ , and  $L'$  a set of objective  
 278 literals disjoint from  $L$ . We say that  $L'$  unlocks  $(b, L)$  on  $U$  if  $|U| \geq b$  and  $F \wedge \alpha \wedge \bigwedge_{\ell \in L'} \bar{\ell} \models \ell'$   
 279 for each  $\ell' \in U$ .

280 ► **Proposition 2.** Assume  $\langle (b_i, L_i) \rangle_{i \in [1, k]}$  is a sequence of cardinality constraints such that  
 281  $F \wedge \alpha \models L_i \geq b_i$  for each  $i$ . Furthermore let  $L$  be a set of objective literals disjoint from all  
 282 the  $L_i$  and let  $U_i \subseteq L_i$  be such that  $L \cup \bigcup_{j < i} L_j \setminus U_j$  unlocks  $U_i$  for each  $i$ . If additionally

$$283 \quad F \wedge \alpha \models L + \sum_i (L_i \setminus U_i) \geq 1,$$

284 then

$$285 \quad F \wedge \alpha \models L + \sum_i L_i \geq \sum_i b_i + 1.$$

286 This proposition tells us that instead of searching for a new core that is disjoint from all  
 287 the previously found cores (or more general, cardinality constraints), we can search for one  
 288 that overlaps with some of the previous constraints, provided that all these constraints get  
 289 unlocked. In practice this leads to a procedure whereby we falsify one by one all as-of-yet  
 290 unassigned objective literals. As soon as unit propagation allows us to conclude that one of  
 291 the previous cardinality constraints is unlocked, we are again allowed to assign the remaining  
 292 literals from that cardinality constraint. If at some point, this yields a conflict, we can cut to  
 293 the core of the conflict using standard conflict analysis techniques and learn a new cardinality  
 294 constraint that increases the total bound by one. Let us illustrate this on a small example.

295 ► **Example 3.** Assume  $x_1, \dots, x_{10}$  are objective literals (to be minimized). During the  
 296 lookahead phase, the following happens:

- 297 ■ Assigning  $x_1 = x_2 = x_3 = x_4 = 0$  results in a conflict. After analyzing this conflict the  
 298 cardinality constraint (clause, in this case)

$$299 \quad x_1 + x_2 + x_4 \geq 1 \tag{2}$$

300 is learned.

## XX:8 Certified Branch-and-Bound MaxSAT Solving

301 ■ Assigning  $x_3 = x_5 = 0$  propagates  $x_6 = 1$ , meaning that the cardinality constraint

$$302 \quad x_3 + x_5 + x_6 \geq 1 \quad (3)$$

303 is learned.

304 ■ Assigning  $x_7 = x_8 = 0$  and running unit propagation yields  $x_1 = 1$ . At this point, the  
305 core (2) is unlocked. Further assigning  $x_2 = x_4 = 0$  unit propagates to a conflict. As  
306 such we can replace core (2) by

$$307 \quad x_1 + x_2 + x_4 + x_7 + x_8 \geq 2. \quad (4)$$

308 In combination with (3), this tells us that the value of the objective is at least three.

309 ■ Assigning  $x_9 = 0$  unit propagates that  $x_1 = x_3 = 1$ . This is not enough to unlock (4), but  
310 (3) is unlocked at this point. Assigning  $x_5 = x_6 = 0$  propagates that  $x_7 = 1$ , meaning  
311 that also (4) gets unlocked. Further assigning  $x_2 = 0$  results in a conflict, meaning that  
312 we can replace (3) and (4) by

$$313 \quad \sum_{1 \leq i \leq 9} x_i \geq 4.$$

### 314 Proof-Logging for Literal Unlocking

315 To prove correctness of these reasoning steps we will, as before make use of the fact that  
316 whenever unit propagation derives a new literal or contradiction, a clause can be learned  
317 that expresses precisely this propagation. The final clause to be learned can then be derived  
318 following the next proposition. In this proposition, we make abstraction of the current  
319 assignment  $\alpha$ , but the proposition directly generalizes to that case by appending  $\sum_{\ell \in \alpha} M\bar{\ell}$   
320 to the left-hand side of every constraint (for a large enough constant  $M$ ). The effect of this  
321 is that the constraint is made conditional on  $\alpha$ .

322 ► **Proposition 4.** *Let  $L_i|_{1 \leq i \leq k}$  and  $L$  be pairwise disjoint sets of objective literals and  $b_i|_{1 \leq i \leq k}$   
323 natural numbers. Assume  $U_i \subseteq L_i$  with  $|U_i| = b_i$  for each  $i$  and write  $R_i$  for  $L_i \setminus U_i$ . From  
324 the constraints*

$$325 \quad L_i \geq b_i \quad \text{for each } i \quad (5)$$

$$326 \quad L + \sum_{j < i} R_j + \ell \geq 1 \quad \text{for each } i \text{ and each } \ell \in U_i \quad (6)$$

$$327 \quad L + \sum_j R_j \geq 1 \quad (7)$$

328 there is a cutting planes derivation that derives

$$329 \quad L + \sum_{j > i} U_j + \sum_j R_j \geq 1 + \sum_{j > i} b_j \quad (8)$$

330 for each  $i \in \{0, \dots, k\}$ . In particular, taking  $i = 0$ , there is a cutting planes derivation of

$$331 \quad L + \sum_j L_j \geq 1 + \sum_j b_j. \quad (9)$$

332 Please note that (5) represents the previously derived cardinality constraints, (6) are the  
333 constraints that guarantee that  $L \cup \bigcup_{j < i} R_j$  unlocks  $(b_i, L_i)$  on  $U_i$  and (7) is the constraint  
334 representing the final contradiction that is derived. The constraint (9), in case the sum of  
335 all the  $b_i$  is large enough, allows us indeed to conclude that there are no assignments that  
336 improve upon the best value found so far. While Proposition 4 only specifies the existence of  
337 a cutting planes derivation, the proof is constructive and provides all the details necessary to  
338 automate this construction.



## 4 Certified BDD-Based Encodings of PB Constraints

A Binary Decision Diagram (BDD) is a (node- and edge-)labeled graph with

- Two leaves, labeled true (**t**) and false (**f**), respectively
- Each internal node is labeled with a variable and has two outgoing edges, labeled true (**t**) and false (**f**), respectively.

Each node in a BDD represents a Boolean function. If  $\eta$  is a node labeled  $x$  with true child  $\eta_{\mathbf{t}}$  and false child  $\eta_{\mathbf{f}}$ , it maps any (total) interpretation  $I$  to  $n(I)$ , which is defined as  $\eta_{\mathbf{t}}(I)$  if  $x \in I$  and as  $\eta_{\mathbf{f}}(I)$  if  $\bar{x} \in I$ . The true and false leaf nodes represent a tautology and contradiction respectively.

A BDD is *ordered* if there is a total order of the variables such that each path through the BDD respects this order and it is *reduced* if two conditions hold: (1) no node has two identical children and (2) no two nodes have the same label, **t**-child and **f**-child. For a fixed variable ordering, each Boolean function has a unique ordered and reduced representation as a BDD, i.e., ordered and reduced BDDs form a canonical representation of Boolean functions. From now on, we will fix the standard variable ordering (ordering the  $x_i$  based on their index), but the ideas presented here work for an arbitrary ordering. When we write  $\text{bdd}(x, \eta_{\mathbf{t}}, \eta_{\mathbf{f}})$  we mean a node in the BDD labeled  $x$  with true child  $\eta_{\mathbf{t}}$  and false child  $\eta_{\mathbf{f}}$ . If the BDD is reduced, this node is unique if it exists.

In MaxCDCL, BDDs are used, heuristically, to encode the solution-improving constraint  $\mathcal{O} \leq v^*$  in SAT. This is done in two phases: first a BDD is constructed, then a set of clauses is generated from this BDD.

### Phase 1: Construction of a BDD for $\mathcal{O} \leq v^*$

For simplicity of the presentation, we will assume here without loss of generality that  $\mathcal{O}$  is of the form  $\sum_{i=1}^n a_i x_i$  (with all the  $a_i$  positive and all the  $x_i$  variables). The standard way to create a reduced and ordered BDD for  $\mathcal{O} \leq v^*$  is to create BDDs  $\eta_{\mathbf{t}}$  for  $\sum_{i>1}^n a_i x_i \leq v^* - a_1$  and  $\eta_{\mathbf{f}}$  for  $\sum_{i>1}^n a_i x_i \leq v^*$ , while making sure the diagram remains reduced, and combine them into the node  $\text{bdd}(x_1, \eta_{\mathbf{t}}, \eta_{\mathbf{f}})$ . However, this approach will always require an exponential number of calls; to avoid this in many cases (but not in the worst-case), a dynamic programming approach was developed. All the PB constraints for which we will create a BDD node are of the form

$$\sum_{i \geq k}^n a_i x_i \leq A. \quad (10)$$

The key observation is that in case  $\sum_{i \geq k}^n a_i x_i$  cannot take the value  $A$  (because of the  $x_i$  being Boolean variables), then this is in fact equivalent to  $\sum_{i \geq k}^n a_i x_i \leq A - 1$ . In fact, for each such pseudo-Boolean constraint there is a (possibly unbounded) interval  $[l, u]$  such that for all  $b \in [l, u]$ , (10) is equivalent to  $\sum_{i \geq k}^n a_i x_i \leq b$ , and moreover, this interval can be computed “bottom-up” from the BDD structure. We will denote this as

$$\sum_{i \geq k}^n a_i x_i \leq [l, u]. \quad (11)$$

The dynamic programming approach now consists in keeping track of this interval for each translated PB constraint and reusing already created BDDs whenever possible. In other words, memoization is used for two things: (1) for looking up whenever a new BDD needs to

## XX:10 Certified Branch-and-Bound MaxSAT Solving

382 be created for a formula of the form (10) whether that combination of  $k$  and  $A$  is already  
383 captured by a previously created BDD and (2) for looking up whether a node with two  
384 specified children already exists. To illustrate this last point, note that the formulas

$$385 \quad 2x + 7y + 7z + 7u \leq [9, 13]$$

386 and

$$387 \quad 7y + 7z + 7u \leq [7, 13]$$

388 are equivalent and hence should have the same representation in a reduced BDD-based  
389 representation.

### 390 Phase 2: A CNF Encoding From the BDD

391 Given a BDD that represents a PB constraint, constructed as discussed above, we can get a  
392 CNF encoding as follows:

393 ■ For each internal node  $\eta$  in the BDD, a new variable  $v_\eta$  is created; intuitively this variable  
394 is true only when the Boolean function is true. In practice for the two leaf nodes no  
395 variable is created but their truth value is filled in directly. However, in the proofs below  
396 we will, to avoid case splitting pretend that a variable exists for each node.

397 ■ For each internal node  $\eta = \text{bdd}(x, \eta_t, \eta_f)$ , the clauses

$$398 \quad \bar{v}_{\eta_t} \wedge x \Rightarrow \bar{v}_\eta \tag{12}$$

399 and

$$400 \quad \bar{v}_{\eta_f} \wedge \bar{x} \Rightarrow \bar{v}_\eta \tag{13}$$

401 are added (encoding precisely the two cases in which the Boolean function in this node  
402 can be violated). If all the coefficients are positive, the second clause can be further  
403 simplified to  $\bar{v}_{\eta_f} \Rightarrow \bar{v}_\eta$ .

404 ■ Finally, for the topf node  $v_\top$  representing  $\mathcal{O} \leq v^*$ , the unit clause  $v_\top$  is added.

### 405 Certifying the BDD-Based CNF Encoding

406 Our strategy for certifying this works as follows.

407 **First** we introduce the new variables using their pseudo-Boolean definitions and immedi-  
408 ately prove that the lower and upper bound in the interval yield the same Boolean function.  
409 This can be done by structural induction on the BDD. Formally, for each node  $\eta$ , representing  
410 the PB constraints (11), we will show that we can derive

$$411 \quad v_\eta \Rightarrow \sum_{i \geq k}^n a_i x_i \leq l, \quad \text{which is} \quad \left( \sum_{i \geq k}^n a_i - l \right) \cdot \bar{v}_\eta + \sum_{i \geq k}^n -a_i x_i \geq -l \tag{14}$$

412 and

$$413 \quad v_\eta \Leftarrow \sum_{i \geq k}^n a_i x_i \leq u, \quad \text{which is} \quad (u + 1) \cdot v_\eta + \sum_{i \geq k}^n a_i x_i \geq u + 1. \tag{15}$$

414 Showing this is the hard part of the derivation.

415 **Secondly**, we derive the desired clauses from the BDD. This can be done using a  
416 straightforward cutting planes derivation.

417 **Certifying Reduced BDDs**

418 Since we are working with *reduced* BDDs, it can happen that different constraints are  
 419 represented by the same node because they are the same underlying Boolean function. In this  
 420 case, what we need to show is that these Boolean functions are indeed equivalent. Specifically,  
 421 the situation we can arrive at is that we will find a node that represents both the following  
 422 constraints

$$423 \quad \sum_{i=\beta}^n a_i x_i \leq [l_1, u] \qquad \sum_{i=\alpha}^n a_i x_i \leq [l_2, u]$$

425 with  $\alpha < \beta$  and with  $l_2 = l_1 + \sum_{i=\alpha}^{\beta-1} a_i$ . What we show then is that there is a cutting planes  
 426 derivation that from the constraints

$$427 \quad v \Rightarrow \sum_{i=\beta}^n a_i x_i \leq l_1 \qquad (16)$$

$$428 \quad v \Leftarrow \sum_{i=\beta}^n a_i x_i \leq u \qquad (17)$$

429 we can derive the constraints

$$431 \quad v \Rightarrow \sum_{i=\alpha}^n a_i x_i \leq l_2 \qquad (18)$$

$$432 \quad v \Leftarrow \sum_{i=\alpha}^n a_i x_i \leq u \qquad (19)$$

433 In other words, we obtain two instantiations of the constraints (14) and (15) for a single  
 434 variable  $v$ , but for the different constraints. Since the rest of our proof logging procedures  
 435 only rely on having derived these two constraints, this is sufficient.  
 436

437 **Using Multi-Valued Decision Diagrams (MDD)**

438 MAXCDCL not only makes use of BDDs for encoding the solution-improving constraint,  
 439 but also of MDDs. The idea is as follows. In some cases, MAXCDCL can infer implicit  
 440 at-most-one constraints. These are constraints of the form  $\sum_{i \in I} x_i \leq 1$  where the  $x_i$  are  
 441 literals in the objective  $\mathcal{O}$ . The detection of such constraints is common in MaxSAT solvers,  
 442 and certification for it has been described in [2], so we will not repeat this here. Now  
 443 assume that a set of disjoint at-most-one constraints has been found. In an MDD, instead of  
 444 branching on *single variable* in each node, we will branch on a set of variables for which an  
 445 at-most-one constraint has previous been derived. This means a node does not have two, but  
 446  $|I| + 1$  children: one for each variable in the set and one for the case where none of them  
 447 is true. Otherwise, the construction and ideas remain the same. As far as *certification* is  
 448 concerned: essentially all proofs continue to hold; the main difference is that there where  
 449 case splitting is used, we will now split into  $|I| + 1$  cases instead of two (splitting on whether  
 450 the variable decided on in that node is true or false) and then use the at most one constraint  
 451 to derive that the conclusion must hold in any case.

452 **5 Conclusions and Future Work**

453 In this paper, we have for the first time presented certification for branch-and-bound MaxSAT  
 454 solving. This work fits in an ongoing effort to show that VERIPB-based proof logging is

455 feasible for different MaxSAT solving paradigms. Together with previous work, three out  
 456 of the four major paradigms are now covered, the major omission being implicit hitting-  
 457 set search. While we see no major theoretic obstacles, an important hurdle preventing  
 458 proof logging for this last paradigm is the fact that state-of-the art solvers make use of  
 459 commercial closed source MIP solvers for computing their hitting sets. Unless these solvers  
 460 are equipped with proof logging capabilities, or if they are replaced by an open-source  
 461 alternative, certification will remain out of reach.

462 Our quest to extend MAXCDCL with proof logging capabilities also resulted in an  
 463 investigation of BDD-based CNF encodings of pseudo-Boolean constraints. This relates to  
 464 the work [10], who has used BDDs for DRAT-based proof logging. In contrast, the focus of  
 465 our work is not on developing general proof logging methods for arbitrary BDD operations,  
 466 but specialized procedures for the constructions occurring in the constructions of BDDs (and  
 467 more generally MDDs) for encoding PB constraints. For this pseudo-Boolean proof logging  
 468 turned out to be very practical as it enabled us to reason about intervals of bounds.

469 In the near future, we plan to finish our implementation and experimentally evaluate the  
 470 performance of our proof logging procedures.

## 471 ——— References ———

- 472 **1** Ignasi Abío, Robert Nieuwenhuis, Albert Oliveras, Enric Rodríguez-Carbonell, and Valentin  
 473 Mayer-Eichberger. A new look at BDDs for pseudo-Boolean constraints. *J. Artif. Intell. Res.*,  
 474 45:443–480, 2012. URL: <https://doi.org/10.1613/jair.3653>, doi:10.1613/JAIR.3653.
- 475 **2** Jeremias Berg, Bart Bogaerts, Jakob Nordström, Andy Oertel, and Dieter Vandesande.  
 476 Certified core-guided MaxSAT solving. In Brigitte Pientka and Cesare Tinelli, editors,  
 477 *Automated Deduction - CADE 29 - 29th International Conference on Automated Deduction,*  
 478 *Rome, Italy, July 1-4, 2023, Proceedings*, volume 14132 of *Lecture Notes in Computer Science*,  
 479 pages 1–22. Springer, 2023. doi:10.1007/978-3-031-38499-8\_1.
- 480 **3** Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development -*  
 481 *Coq'Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science.  
 482 An EATCS Series. Springer, 2004. doi:10.1007/978-3-662-07964-5.
- 483 **4** Armin Biere. Tracecheck. <http://fmv.jku.at/tracecheck/>, 2006.
- 484 **5** Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of*  
 485 *Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*.  
 486 IOS Press, 2021. doi:10.3233/FAIA336.
- 487 **6** Miquel Bofill, Jordi Coll, Josep Suy, and Mateu Villaret. An mdd-based SAT en-  
 488 coding for pseudo-Boolean constraints with at-most-one relations. *Artif. Intell. Rev.*,  
 489 53(7):5157–5188, 2020. URL: <https://doi.org/10.1007/s10462-020-09817-6>, doi:10.  
 490 1007/S10462-020-09817-6.
- 491 **7** Bart Bogaerts, Stephan Gocht, Ciaran McCreesh, and Jakob Nordström. Certified dominance  
 492 and symmetry breaking for combinatorial optimisation. *J. Artif. Intell. Res.*, 77:1539–1589,  
 493 2023. doi:10.1613/jair.1.14296.
- 494 **8** Robert Brummayer and Armin Biere. Fuzzing and delta-debugging SMT solvers. In Ofer Strich-  
 495 man Bruno Dutertre, editor, *Proceedings of the 7th International Workshop on Satisfiability*  
 496 *Modulo Theories*, SMT '09, page 1–5, New York, NY, USA, 2009. Association for Computing  
 497 Machinery. doi:10.1145/1670412.1670413.
- 498 **9** Robert Brummayer, Florian Lonsing, and Armin Biere. Automated testing and debugging of  
 499 SAT and QBF solvers. In Ofer Strichman and Stefan Szeider, editors, *Theory and Applications*  
 500 *of Satisfiability Testing - SAT 2010, 13th International Conference, SAT 2010, Edinburgh,*  
 501 *UK, July 11-14, 2010. Proceedings*, volume 6175 of *Lecture Notes in Computer Science*, pages  
 502 44–57. Springer, 2010. doi:10.1007/978-3-642-14186-7\_6.

- 503 10 Randal E. Bryant. Tbuddy: A proof-generating BDD package. In Alberto Griggio and Neha  
504 Rungta, editors, *22nd Formal Methods in Computer-Aided Design, FMCAD 2022, Trento,*  
505 *Italy, October 17-21, 2022*, pages 49–58. IEEE, 2022. URL: [https://doi.org/10.34727/](https://doi.org/10.34727/2022/isbn.978-3-85448-053-2_10)  
506 [2022/ISBN.978-3-85448-053-2\\_10](https://doi.org/10.34727/2022/ISBN.978-3-85448-053-2_10), doi:10.34727/2022/ISBN.978-3-85448-053-2\_10.
- 507 11 Sam Buss and Jakob Nordström. Proof complexity and SAT solving. In Biere et al. [5], pages  
508 233–350. doi:10.3233/FAIA200990.
- 509 12 William J. Cook, Collette R. Coullard, and György Turán. On the complexity of cutting-plane  
510 proofs. *Discret. Appl. Math.*, 18(1):25–38, 1987. doi:10.1016/0166-218X(87)90039-4.
- 511 13 William J. Cook, Thorsten Koch, Daniel E. Steffy, and Kati Wolter. A hybrid branch-and-  
512 bound approach for exact rational mixed-integer programming. *Math. Program. Comput.*,  
513 5(3):305–344, 2013. doi:10.1007/s12532-013-0055-6.
- 514 14 Luís Cruz-Filipe, Marijn J. H. Heule, Warren A. Hunt Jr., Matt Kaufmann, and Peter Schneider-  
515 Kamp. Efficient certified RAT verification. In Leonardo de Moura, editor, *Automated Deduction*  
516 *- CADE 26 - 26th International Conference on Automated Deduction, Gothenburg, Sweden,*  
517 *August 6-11, 2017, Proceedings*, volume 10395 of *Lecture Notes in Computer Science*, pages  
518 220–236. Springer, 2017. doi:10.1007/978-3-319-63046-5\_14.
- 519 15 Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming  
520 language. In André Platzer and Geoff Sutcliffe, editors, *Automated Deduction - CADE 28*  
521 *- 28th International Conference on Automated Deduction, Virtual Event, July 12-15, 2021,*  
522 *Proceedings*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer,  
523 2021. doi:10.1007/978-3-030-79876-5\_37.
- 524 16 Nelson Guimarães Ferreira and Paulo Sérgio Muniz Silva. Automatic verification of safety  
525 rules for a subway control software. In Alexandre Mota and Arnaldo V. Moura, editors,  
526 *Proceedings of the Seventh Brazilian Symposium on Formal Methods, SBMF 2004, Recife,*  
527 *Pernambuco, Brazil, November 29 - December 1, 2004*, volume 130 of *Electronic Notes in*  
528 *Theoretical Computer Science*, pages 323–343. Elsevier, 2004. URL: [https://doi.org/10.](https://doi.org/10.1016/j.entcs.2005.03.017)  
529 [1016/j.entcs.2005.03.017](https://doi.org/10.1016/j.entcs.2005.03.017), doi:10.1016/J.ENTCS.2005.03.017.
- 530 17 Mathias Fleury. *Formalization of logical calculi in Isabelle/HOL*. PhD thesis, Saar-  
531 land University, Saarbrücken, Germany, 2020. URL: [https://tel.archives-ouvertes.fr/](https://tel.archives-ouvertes.fr/tel-02963301)  
532 [tel-02963301](https://tel.archives-ouvertes.fr/tel-02963301).
- 533 18 Xavier Gillard, Pierre Schaus, and Yves Deville. Solvercheck: Declarative testing of con-  
534 straints. In Thomas Schiex and Simon de Givry, editors, *Principles and Practice of Constraint*  
535 *Programming - 25th International Conference, CP 2019, Stamford, CT, USA, September 30*  
536 *- October 4, 2019, Proceedings*, volume 11802 of *Lecture Notes in Computer Science*, pages  
537 565–582. Springer, 2019. doi:10.1007/978-3-030-30048-7\_33.
- 538 19 Stephan Gocht, Ruben Martins, Jakob Nordström, and Andy Oertel. Certified CNF translations  
539 for pseudo-Boolean solving. In Kuldeep S. Meel and Ofer Strichman, editors, *25th International*  
540 *Conference on Theory and Applications of Satisfiability Testing, SAT 2022, August 2-5, 2022,*  
541 *Haifa, Israel*, volume 236 of *LIPICs*, pages 16:1–16:25. Schloss Dagstuhl - Leibniz-Zentrum für  
542 Informatik, 2022. doi:10.4230/LIPICs.SAT.2022.16.
- 543 20 Stephan Gocht and Jakob Nordström. Certifying parity reasoning efficiently using pseudo-  
544 Boolean proofs. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021,*  
545 *Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The*  
546 *Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual*  
547 *Event, February 2-9, 2021*, pages 3768–3777. AAAI Press, 2021. URL: [https://ojs.aaai.](https://ojs.aaai.org/index.php/AAAI/article/view/16494)  
548 [org/index.php/AAAI/article/view/16494](https://ojs.aaai.org/index.php/AAAI/article/view/16494).
- 549 21 Evguenii I. Goldberg and Yakov Novikov. Verification of proofs of unsatisfiability for CNF  
550 formulas. In *2003 Design, Automation and Test in Europe Conference and Exposition*  
551 *(DATE 2003), 3-7 March 2003, Munich, Germany*, pages 10886–10891. IEEE Computer  
552 Society, 2003. URL: <https://doi.ieeecomputersociety.org/10.1109/DATE.2003.10008>,  
553 doi:10.1109/DATE.2003.10008.

- 554 22 Marijn Heule, Matti Järvisalo, Martin Suda, Markus Iser, and Tomáš Balyo. The 2022  
555 international SAT competition. <https://satcompetition.github.io/2022/>, 2022.
- 556 23 Marijn Heule, Warren A. Hunt Jr., and Nathan Wetzler. Trimming while checking clausal  
557 proofs. In *Formal Methods in Computer-Aided Design, FMCAD 2013, Portland, OR, USA,*  
558 *October 20-23, 2013*, pages 181–188. IEEE, 2013. URL: [https://ieeexplore.ieee.org/  
559 document/6679408/](https://ieeexplore.ieee.org/document/6679408/).
- 560 24 Marijn J. H. Heule, Oliver Kullmann, and Victor W. Marek. Solving and verifying the Boolean  
561 Pythagorean triples problem via cube-and-conquer. In Nadia Creignou and Daniel Le Berre,  
562 editors, *Theory and Applications of Satisfiability Testing - SAT 2016 - 19th International*  
563 *Conference, Bordeaux, France, July 5-8, 2016, Proceedings*, volume 9710 of *Lecture Notes in*  
564 *Computer Science*, pages 228–245. Springer, 2016. doi:10.1007/978-3-319-40970-2\_15.
- 565 25 Matti Järvisalo, Jeremias Berg, Ruben Martins, and Andreas Niskanen. MaxSAT evaluation  
566 2023. <https://maxsat-evaluations.github.io/2023/>, 2023.
- 567 26 Matti Järvisalo, Marijn Heule, and Armin Biere. Inprocessing rules. In Bernhard Gramlich,  
568 Dale Miller, and Uli Sattler, editors, *Automated Reasoning - 6th International Joint Conference,*  
569 *IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings*, volume 7364 of *Lecture Notes*  
570 *in Computer Science*, pages 355–370. Springer, 2012. doi:10.1007/978-3-642-31365-3\_28.
- 571 27 Chu Min Li and Felip Manyà. MaxSAT, hard and soft constraints. In Biere et al. [5], pages  
572 903–927. doi:10.3233/FAIA201007.
- 573 28 Chu-Min Li, Zhenxing Xu, Jordi Coll, Felip Manyà, Djamal Habet, and Kun He. Boosting  
574 branch-and-bound MaxSAT solvers with clause learning. *AI Commun.*, 35(2):131–151, 2022.  
575 doi:10.3233/AIC-210178.
- 576 29 David F. Manlove and Gregg O’Malley. Paired and altruistic kidney donation in the UK:  
577 algorithms and experimentation. *ACM J. Exp. Algorithmics*, 19(1), 2014. doi:10.1145/  
578 2670129.
- 579 30 Ross M. McConnell, Kurt Mehlhorn, Stefan Näher, and Pascal Schweitzer. Certifying al-  
580 gorithms. *Comput. Sci. Rev.*, 5(2):119–161, 2011. doi:10.1016/j.cosrev.2010.09.009.
- 581 31 Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL - A Proof Assistant*  
582 *for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer, 2002.  
583 doi:10.1007/3-540-45949-9.
- 584 32 Monica L. Nogueira, Marcello Balduccini, Michael Gelfond, Richard Watson, and Matthew  
585 Barry. An a-Prolog decision support system for the space shuttle. In I. V. Ramakrishnan,  
586 editor, *Practical Aspects of Declarative Languages, Third International Symposium, PADL*  
587 *2001, Las Vegas, Nevada, USA, March 11-12, 2001, Proceedings*, volume 1990 of *Lecture Notes*  
588 *in Computer Science*, pages 169–183. Springer, 2001. doi:10.1007/3-540-45241-9\_12.
- 589 33 João P. Marques Silva and Karem A. Sakallah. GRASP: A search algorithm for propositional  
590 satisfiability. *IEEE Trans. Computers*, 48(5):506–521, 1999. doi:10.1109/12.769433.
- 591 34 Konrad Slind and Michael Norrish. A brief overview of HOL4. In Otmane Aït Mohamed,  
592 César A. Muñoz, and Sofiène Tahar, editors, *Theorem Proving in Higher Order Logics, 21st*  
593 *International Conference, TPHOLs 2008, Montreal, Canada, August 18-21, 2008. Proceedings*,  
594 volume 5170 of *Lecture Notes in Computer Science*, pages 28–32. Springer, 2008. doi:  
595 10.1007/978-3-540-71067-7\_6.
- 596 35 Dieter Vandesande. Towards certified MaxSAT solving: Certified MaxSAT solving with  
597 SAT oracles and encodings of pseudo-Boolean constraints. Master’s thesis, Vrije Uni-  
598 versiteit Brussel (VUB), 2023. URL: [https://researchportal.vub.be/nl/studentTheses/  
599 towards-certified-maxsat-solving](https://researchportal.vub.be/nl/studentTheses/towards-certified-maxsat-solving).
- 600 36 Dieter Vandesande, Wolf De Wulf, and Bart Bogaerts. QMaxSATpb: A certified MaxSAT  
601 solver. In Georg Gottlob, Daniela Incezan, and Marco Maratea, editors, *Logic Programming*  
602 *and Nonmonotonic Reasoning - 16th International Conference, LPNMR 2022, Genova, Italy,*  
603 *September 5-9, 2022, Proceedings*, volume 13416 of *Lecture Notes in Computer Science*, pages  
604 429–442. Springer, 2022. doi:10.1007/978-3-031-15707-3\_33.

- 605 37 Nathan Wetzler, Marijn Heule, and Warren A. Hunt Jr. DRAT-trim: Efficient checking and  
606 trimming using expressive clausal proofs. In Carsten Sinz and Uwe Egly, editors, *Theory*  
607 *and Applications of Satisfiability Testing - SAT 2014 - 17th International Conference, Held*  
608 *as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014.*  
609 *Proceedings*, volume 8561 of *Lecture Notes in Computer Science*, pages 422–429. Springer,  
610 2014. doi:10.1007/978-3-319-09284-3\_31.