MiniMaxSat: An efficient Weighted Max-SAT Solver ⋆

Federico Heras, Javier Larrosa, and Albert Oliveras
fheras@lsi.upc.edu, larrosa@lsi.upc.edu, oliveras@lsi.upc.edu

Universitat Politecnica de Catalunya,
Jordi Girona 1-3, 08034 Barcelona, Spain

Abstract. In this paper we introduce MINIMAXSAT, a new Max-SAT solver that is built on top of MINISAT+. It incorporates the best current SAT and Max-SAT techniques. It can handle hard clauses (clauses of mandatory satisfaction as in SAT), soft clauses (clauses whose falsification is penalized by a cost as in Max-SAT) as well as pseudo-boolean objective functions and constraints. Its main features are: learning and backjumping on hard clauses; resolution-based and substraction-based lower bounding; and lazy propagation with the two-watched literals scheme. Our empirical evaluation comparing a wide set of solving alternatives on a broad set of optimization benchmarks indicates that the performance of MINIMAXSAT is usually close to the best specialized alternative and, in some cases, even better.

1 Introduction

Max-SAT is the optimization version of SAT where the goal is to satisfy the maximum number of clauses. It is considered one of the fundamental combinatorial optimization problems and many important problems can be naturally expressed as Max-SAT. They include academic problems such as *max cut* or *max clique*, as well as real problems in domains like *routing*, *bioinformatics*, *scheduling*, *electronic markets*, *etc*...

There is a long tradition of theoretical work about the structural complexity [2] and approximability [3] of Max-SAT. Most of this work is restricted to the simplest case in which all clauses are equally important (*i.e.*, unweighted Max-SAT) and have a fixed size (mainly binary or ternary clauses). From a practical point of view, significant progress has been made in the last 3 years [4–9]. As a result, there is a handful of new solvers that can deal, for the first time, with medium-sized instances.

The main motivation of our work comes from the study of Max-SAT instances modelling real-world problems. We usually encounter three features:

- The satisfaction of all clauses does not have the same importance, so each clause needs to be associated with a weight that represents the cost of its violation. In the extreme case, which often happens in practice as observed in [10], there are clauses whose satisfaction is mandatory. They are usually modelled by associating a very high weight with them.
- Literals do not appear randomly along the clauses. On the contrary, it is easy to identify patterns, symmetries or other kinds of structures.

⋆ This paper includes and extends preliminary work from [1]

– In some problems there are mandatory clauses that reduce dramatically the number of feasible assignments, so the optimization part of the problem only plays a secondary role. However, in some other problems mandatory clauses are trivially satisfiable and the real difficulty lays on the optimization part.

When we look at current Max-SAT solvers, we find that none of them is robust over these three features. For instance, [8, 9] are restricted to formulas in which all clauses are equally important (i.e. unweighted Max-SAT), [4] is restricted to binary clauses, [6] seems to be efficient on very overconstrained problems (*i.e.*, only a small fraction of the clauses can be simultaneously satisfied), while [11] seems to be efficient on slightly overconstrained problems (*i.e.* almost all the clauses can be satisfied). The solver of [12], developed in parallel to the research described in this paper, can handle mandatory clauses and is the only one that incorporates some learning, so it seems to perform well on structured problems. However, all non-mandatory clauses must have the same weight. Finally, approaches based on translating a Max-SAT instance into a SAT instance and solve them with a SAT-solver seem to be effective in highly structured problems in which almost all clauses are mandatory [13, 14].

In this paper we introduce MINIMAXSAT, a new weighted Max-SAT solver that incorporates the current best SAT and Max-SAT techniques. It is build on top of MiniSAT+ [15], so it borrows its capability to deal with pseudo-boolean problems and all the MiniSAT [16] features processing mandatory clauses such as learning and backjumping. We have extended it allowing it to deal with weighted clauses, while preserving the two-watched literals lazy propagation method. The main original contribution of MINIMAXSAT is that it implements a novel and very efficient lower bounding technique. Specifically, it applies unit propagation in order to detect disjoint subsets of mutually inconsistent clauses like in [9]. Then it simplifies the problem like in [5, 17, 6] in order to increment the lower bound. However, while in [5, 17, 6] only the clauses that accomplish specific patterns are transformed, in MINIMAXSAT there is no need to define such patterns.

The structure of the paper is as follows: Section 2 provides preliminary definitions on SAT and Section 3 presents state-of-the-art solving techniques incorporated in a modern SAT-solver such as MINISAT. Then, Section 4 presents preliminary definitions on Max-SAT and Section 5 overviews MINIMAXSAT. After that, Sections 6 and 7 focus on its lower bounding and additional features, respectively. In Section 8 we present the benchmarks used in our empirical evaluation and we report the experimental results. Finally, Section 9 presents related work and Section 10 concludes and points out possible future work.

2 Preliminaries on SAT

In the sequel $X = \{x_1, x_2, \ldots, x_n\}$ is the set of boolean variables. A *literal* is either a variable x_i or its negation $\bar{x}_i$. The variable to which literal l refers is noted $var(l)$. Given a literal l, its negation $\bar{l}$ is $\bar{x}_i$ if l is x_i and is x_i if l is $\bar{x}_i$. A *clause* C is a disjunction of literals. The *size* of a clause, noted $|C|$, is the number of literals that it has. The set of variables that appear in C is noted $var(C)$. Sometimes we associate a subscript Greek letter to a clause (e.g. $(x_i \vee x_j)_\alpha$) in order to facilitate future references of such clause.

Algorithm 1: MiniSAT basic structure.

```
   Function MiniSAT() : boolean
1      InitQueue( ) ;
2      Loop
3          UP( ) ;
4          if Conflict then
5              AnalyzeConflict( ) ;
6              if Top Conflict then return false ;
               else
7                  LearnClause( ) ;
8                  Backjump( ) ;
9          else if all variables assigned then return true ;
10         else
11             l := SelectLiteral( ) ;
12             Enqueue(Q, l) ;
```

An *assignment* is a set of literals not containing a variable and its negation. Assignments of maximal size n are called *complete*, otherwise they are called *partial*. Given an assignment $\mathcal{A}$, a variable x is *unassigned* if neither x nor $\bar{x}$ belong to $\mathcal{A}$. Similarly, a literal l is *unassigned* if $var(l)$ is unassigned.

An assignment *satisfies* a literal iff it belongs to the assignment, it *satisfies* a clause iff it satisfies one or more of its literals and it *falsifies* a clause iff it contains the negation of all its literals. In the latter case we say that the clause is *conflicting* as it always happens with the empty clause, noted $\square$. A boolean formula $\mathcal{F}$ in *conjunctive normal form* (CNF) is a set of clauses representing their conjunction. A model of $\mathcal{F}$ is a complete assignment that satisfies all the clauses in $\mathcal{F}$.

If $\mathcal{F}$ has a model, we call it *satisfiable*, otherwise we say it is *unsatisfiable*. Moreover, if all complete assignments satisfy $\mathcal{F}$, we say that $\mathcal{F}$ is a *tautology*.

Clauses of size one are called *unit clauses* or simply *units*. When a formula contains a unit l, it can be simplified by removing all clauses containing l and removing $\bar{l}$ from all the clauses where it appears. The application of this rule until quiescence is called *unit propagation* (UP) and it is well recognized as a fundamental propagation technique in all current SAT solvers.

Another well-known rule is *resolution*, which, given a formula containing two clauses $(x \vee A), (\bar{x} \vee B)$ (called *clashing clauses*), allows one to add a new clause $(A \vee B)$ (called *the resolvent*).

3 Overview of MINISAT

In this section we overview MINISAT [16], a state-of-the-art SAT solver that incorporates the most effective current solving techniques. Its core algorithm is DPLL [18], a systematic depth-first search on the tree of assignments. An internal tree node is as-

Algorithm 2: Unit Propagation.

```
    Function UP(Q) : Conflict
        while (Q contains non-propagated literals) do
13          l := GetOldestLit(Q); MarkAsPropagated(l) ;
14          foreach clause C ∨ l̄ that becomes unit or falsified do
15              if C ∨ l̄ becomes a unit q then Enqueue(Q, q) ;
16              else if C ∨ l̄ becomes falsified then return Conflict ;
        return None ;
```

sociated to a partial assignment and its two successors are obtained by selecting an unassigned variable x and extending the current assignment with x and $\bar{x}$, respectively. At each visited node, new units are derived due to the application of unit propagation (UP). If that leads to a conflicting clause, MINISAT backtracks to a previous branching point, performing non-chronological backtracking and clause learning (as in [19]).

An algorithmic description of MINISAT appears in Algorithm 1. The algorithm uses a propagation queue Q which contains all units pending propagation and also contains a representation of the current assignment.

First, propagation queue Q is filled with the units contained in the original formula (line 1). The main loop starts in line 2 and at each iteration procedure UP is in charge of propagating all pending units (line 3). If a conflicting clause is found (line 4), the conflict is analyzed (line 5) and as a result a new clause is *learned* (i.e, inferred and recorded, line 7).

Then, MINISAT backtracks, using the propagation queue Q to undo the assignment until exactly one of the literals of the learned clause becomes unassigned (line 8). If one can further backtrack while still maintaining this condition, it is advantageous to do so (this is commonly referred to as *backjumping* or *non-chronological backtracking* [19]). If UP leads to no conflict, a new unassigned literal is selected to extend the current partial assignment. The new literal is added to Q (line 10) and a new iteration takes place.

The procedure stops when a complete assignment is found (line 9) or when a top level conflict is found (line 6). In the first case, the procedure returns $true$ which indicates that a model has been found, while in the second case it returns $false$ which means that no model exists for the input formula.

In the following we describe in more depth the different components of the search.

3.1 Unit Propagation

The aim of unit propagation is twofold: on the one hand, to find all clauses that have become units due to the current assignment, and on the other hand, to detect whether some clause has become conflicting. A more concrete procedure is given in Algorithm 2. While non-propagated literals exist in Q, it picks the oldest one l and marks it as propagated (line 13). Then all clauses containing $\bar{l}$ that may have become falsified or units are traversed (we will later describe how these clauses are detected). If one of

such clauses becomes a unit q, it is enqueued in Q to be propagated later (line 15). The procedure iterates until there are no more units to propagate or until a conflicting clause is found (line 16).

There are two types of literals in Q: *decision literals* are those that the algorithm has heuristically selected and assigned at a branching point (lines 11 and 12 in Algorithm 1); *consequence literals* are those that are added because they are logical consequences of a previous decision literal (line 15). MINISAT uses a non-standard queue to handle units pending propagation. Unlike classical queues, after fetching an element, it is not removed, but just marked as such. Consequently, Q is formed by two sets of elements: the already propagated literals and the literals pending propagation. The advantage of such strategy is that at any execution point, Q also contains the current assignment. Besides, the propagated literals of Q are divided into *decision levels*. Each decision level contains a decision literal and the set of its related consequences. Furthermore, a literal l is associated with the original clause that cause its propagation and it is noted as $l(\alpha)$; such a clause is usually referred to as the *reason* of l. Note that a decision literal l does not have a *reason* and will be represented as l^d.

Example 1 *Consider the formula* $\{(\bar{x}_1 \vee x_2)_\alpha, (\bar{x}_1 \vee x_3)_\beta, (\bar{x}_4 \vee \bar{x}_5)_\gamma\}$. *Before starting the execution, the propagation queue is empty* $Q = [\|]$. *We use the symbol* $\|$ *to separate propagated literals (on the left) from literals pending propagation (on the right). If literal* x_1 *is decided, it is added to Q. Before propagation the queue contains* $Q = [\|x_1^d]$. *UP will propagate* x_1 *and add two new consequences* x_2 *and* x_3. *The propagation queue is now* $Q = [x_1^d \| x_2(\alpha), x_3(\beta)]$ *and the current assignment is* $\{x_1, x_2, x_3\}$. *The propagation of* x_2 *and* x_3 *does not add new literals to Q, so it becomes* $Q = [x_1^d, x_2(\alpha), x_3(\beta) \|]$

If x_4 *is decided, UP will add a new consequence* $\bar{x}_5$. *After the propagation, we have* $Q = [x_1^d, x_2(\alpha), x_3(\beta), x_4^d, \bar{x}_5(\gamma) \|]$. *The current assignment is* $\{x_1, x_2, x_3, x_4, \bar{x}_5\}$. *Note that no more literals can be propagated and a complete assignment has been found. Note as well that Q contains two decision levels: the first one is formed by literals* x_1, x_2 *and* x_3 *while the second one is formed by literals* x_4 *and* $\bar{x}_5$.

Lazy Data Structures. As mentioned, the aim of UP is to detect all units and all conflicting clauses. Taking into account that this process typically takes as much as 80% of the total runtime of a SAT solver, it is important to design efficient data structures.

The first attempt was the use of *adjacency lists*. For each literal, one keeps the list of all clauses in which the literal appears. Then, upon the addition of a literal l to the assignment, only clauses containing $\bar{l}$ have to be traversed. The main drawback of further refinements to detect efficiently when a clause has become unit, such as keeping *counters* indicating the number of unassigned literals of a clause, is that they involved a considerable amount of work upon backtracking.

The method used by MINISAT is the two-watched literal scheme introduced in [20]. Its basic idea is that a clause cannot be unit or conflicting if (i) it has one satisfied literal or (ii) it has two unassigned literals.

The algorithm keeps two special literals for each clause, called the *watched literals*, initially two unassigned literals, and tries to maintain the invariant that always one satisfied literal or two unassigned literals are watched.

The invariant may be broken only if one of the two watched literals becomes falsified. In this case, the clause is traversed looking for another non-false literal to watch in order to restore the invariant. If one such literal cannot be found, the clause is declared to be true, unit or conflicting depending on the value of the other watched literal. Hence, when a literal l is added to the assignment, the clauses that may have become falsified or unit (line 14 in Algorithm 2) are only those clauses where $\bar{l}$ is watched.

The main advantage of such an approach is that no work on the clauses has to be done upon backtracking. However, the main drawback is that the only way to know how many literals are unassigned for a given clause is by traversing all its literals. Note that this information is used by other techniques such as the *Two-sided Jeroslow* branching heuristic (See Section 3.3).

Resolution Refutation Trees. When no branching point has been created and UP detects a conflict, an unsatisfiable subset of clauses $\mathcal{F}'$ can be determined using the information provided by Q. Since $\mathcal{F}'$ is unsatisfiable, the empty clause $\square$ can be derived from $\mathcal{F}'$ via resolution. Such resolution process is called a *refutation*. A refutation for an unsatisfiable clause set $\mathcal{F}'$ is a *resolution refutation tree* (or simply a *refutation tree*) if every clause is used exactly once during the resolution process.

A refutation tree Υ can be built from the propagation queue Q as follows: let C_0 be the conflicting clause. Traverse Q in a *LIFO* (*Last In First Out*) fashion until a clashing clause D_0 is found. Then resolution is applied between C_0 and D_0, obtaining resolvent C_1. Next, the traversal of Q continues until a clause D_1 that clashes with C_1 is found, giving resolvent C_2 and we iterate the process until the resolvent we obtain is the empty clause $\square$. The importance of refutation trees will become relevant in Section 6.

Example 2 *Consider the formula* $\mathcal{F} = \{(\bar{x}_1)_\alpha, (x_1 \vee x_4)_\beta, (x_1 \vee x_2)_\gamma, (x_1 \vee x_3)_\delta, (\bar{x}_2 \vee \bar{x}_3)_\varepsilon\}$. *If we apply unit propagation the unit clause* α *is enqueued producing* $Q = [\|\bar{x}_1(\alpha)]$. *Then* $\bar{x}_1$ *is propagated and* Q *becomes* $[\bar{x}_1(\alpha)\|x_4(\beta), x_2(\gamma), x_3(\delta)]$. *Then, literal* x_4 *is propagated and no new units are added to* Q.

After that, literal x_2 *is propagated and clause* ε *is found to be conflicting. Figure 1.a shows the state of* Q *after the propagation.*

Now we build the refutation tree. Starting from the tail of Q *the first clause clashing with the conflicting clause* ε *is* δ. *Resolution between* ε *and* δ *generates the resolvent* $x_1 \vee \bar{x}_2$. *The first clause clashing with it is* γ, *producing resolvent* x_1. *The next clause clashing with it is* α *and resolution generates* $\square$. *Figure 1.b shows the resulting resolution tree.*

3.2 Learning and Backjumping

Learning and backjumping are best illustrated with an example (see [19, 21] for a precise description):

Example 3 *Consider the formula* $\{(\bar{x}_1 \vee x_2)_\alpha, (\bar{x}_3 \vee x_4)_\beta, (\bar{x}_5 \vee \bar{x}_6)_\gamma, (\bar{x}_2 \vee \bar{x}_5 \vee x_6)_\delta\}$ *and the partial assignment* $\{x_1, x_2, x_3, x_4, x_5, \bar{x}_6\}$ *that leads to a conflict over clause* δ. *Suppose that the current propagation queue is* $Q = [x_1^d, x_2(\alpha), x_3^d, x_4(\beta), x_5^d, \bar{x}_6(\gamma)\|]$.

$$\mathcal{F} = \{(\bar{x}_1)_\alpha, (x_1 \vee x_4)_\beta, (x_1 \vee x_2)_\gamma, (x_1 \vee x_3)_\delta, (\bar{x}_2 \vee \bar{x}_3)_\varepsilon\}$$

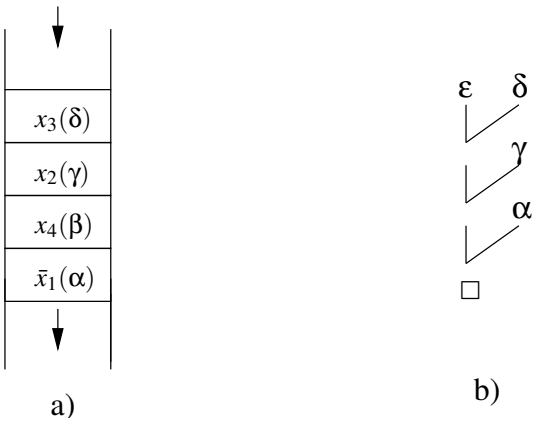

Fig. 1. Graphical representation of the propagation queue Q and a refutation tree Υ of example 2. On the top, the original formula $\mathcal{F}$. On the left, the propagation Q after step 1. On the right, the structure of the resolution tree computed in step 2.

In the example it is easy to see that decision x_1^d is incompatible with decision x_5^d. Such incompatibility can be represented with clause $(\bar{x}_1 \vee \bar{x}_5)$. Similarly, consequence x_2 is incompatible with decision x_5^d and it can be represented with the clause $(\bar{x}_2 \vee \bar{x}_5)$. Clause learning implements different techniques that are used to *discover* such implicit incompatibilities and adds them to the formula. Learned clauses can accelerate the subsequent search, since they can increase the potential of future UP executions. However, it has been observed that unrestricted clause learning can be impractical in some cases (recorded clauses consume memory and repeated recording may lead to its exhaustion). For this reason, current SAT-solvers incorporate different clause deletion policies in order to remove some of the learned clauses.

Learned clauses can also be used to backjump if their presence would have allowed a unit propagation at an earlier decision level. In this case, we say that the clause is *asserting* and backjumping can proceed by going back to that level and adding the unit propagated literal. Among the several automated ways of generating asserting clauses, MINISAT uses the so-called *First Unique Implication Point* (*1UIP*) [21].

3.3 Branching Heuristic

Branching occurs in the function `SelectLiteral` (Algorithm 1). When there are no more literals to propagate, this function chooses one variable from all the unassigned ones and assigns it a value. The importance of the branching heuristic is well known, since different branching heuristic may produce different-sized search trees.

Early branching heuristics includes the *Bohm's Heuristic*, the *Maximum Ocurrences on Minimum sized clauses* (MOM) and the *Two sided-Jeroslow Wang Heuristic*. Those heuristics try to choose the literal such that its assignment will generate the largest number of implications or that satisfy most clauses. All these heuristics are *state dependent*, that is, they use information about the state of the clauses given the current assignment. In most of them, such information is the number of unassigned literals for each clause. Hence, they were implemented jointly with data structures based on adjacency lists since they keep such information. For instance, the Two sided-Jeroslow Wang Heuristic computes for each literal l of F the following function:

$$J(l) = \sum_{\substack{C \in F \\ s.t.\ \ l \in C}} 2^{-|C|}$$

and selects the literal l that maximizes function $J(l)$.

As solvers become more efficient, updating metrics of state-dependent heuristics dominates the execution time. Hence MINISAT uses a slight modification of a state-independent heuristic first proposed in [20]. Such heuristic, called *Variable State Independent Decaying Sum* (*VSIDS*), selects the literal that appears more frequently over all clauses, but giving priority to recently learned clauses. The advantage of this heuristic is that metrics only have to be updated when clauses are learned or forgotten. Since this only occurs occasionally, its computation has very low overhead. The VSIDS heuristic suits perfectly with lazy data structures such as the two-watched literal scheme.

4 (Weighted) Max-SAT

A *weighted* clause is a pair (C,w), where C is a clause and w is the cost of its falsification, also called its *weight*. If a problem contains clauses that *must* be satisfied, we call such clauses *mandatory* or *hard* and associate with them a special weight $\top$. Non-mandatory clauses are also called *soft*. A *weighted formula* in *conjunctive normal form* (WCNF) is a set of weighted clauses. A *model* is a complete assignment that satisfies all mandatory clauses. The *cost of an assignment* is the sum of weights of the clauses that it falsifies. Given a WCNF formula $\mathcal{F}$, *Weighted* Max-SAT is the problem of finding a model of $\mathcal{F}$ of minimum cost. This cost will be called the *optimal cost of* $\mathcal{F}$. Note that if a formula contains only mandatory clauses, weighted Max-SAT is equivalent to classical SAT. If all the clauses have weight 1, we have the so-called (unweighted) Max-SAT problem. In the following, we will assume weighted Max-SAT.

We say that a weighted formula $\mathcal{F}'$ is a *relaxation* of weighted formula $\mathcal{F}$ (noted $\mathcal{F}' \sqsubseteq \mathcal{F}$) if the optimal cost of $\mathcal{F}'$ is less than or equal to the optimal cost in $\mathcal{F}$ (non-models are considered to have cost infinity). We say that two weighted formulas $\mathcal{F}'$ and $\mathcal{F}$ are *equivalent* (noted $\mathcal{F}' \equiv \mathcal{F}$) if $\mathcal{F}' \sqsubseteq \mathcal{F}$ and $\mathcal{F} \sqsubseteq \mathcal{F}'$.

Max-SAT simplification rules transforms a formula $\mathcal{F}$ into an equivalent, but presumably simpler formula $\mathcal{F}'$. All SAT simplification rules (e.g. unit propagation, tautology removal,...) can be directly applied to Max-SAT if restricted to mandatory clauses. However, several specific Max-SAT simplification rules exist [6]. For instance, if a formula contains clauses (C,u) and (C,v), they can be replaced by $(C,u+v)$. If it contains a clause $(C,0)$, it may be removed. If it contains a unit $(l,\top)$, it can be simplified by removing all (including soft) clauses containing l and removing $\bar{l}$ from all the clauses (including soft clauses) where it appears. The application of this rule until quiescence is the natural extension of *unit propagation* to Max-SAT.

The empty clause may appear in a weighted formula. If its weight is $\top$, it is clear that the formula does not have any model. If its weight is $w < \top$, the cost of any assignment will include that weight, so w is an obvious lower bound of the formula optimal cost. Weighted empty clauses and their interpretation in terms of lower bounds will become relevant in Section 6.

As shown by [6], the notion of resolution can be extended to weighted formulas as follows [1],

$$\{(x \vee A, u), (\bar{x} \vee B, w)\} \equiv \left\{ \begin{array}{l} (A \vee B, m), \\ (x \vee A, u-m), \\ (\bar{x} \vee B, w-m), \\ (x \vee A \vee \bar{B}, m), \\ (\bar{x} \vee \bar{A} \vee B, m) \end{array} \right\}$$

where A and B are arbitrary disjunctions of literals and $m = \min\{u,w\}$.

$(x \vee A, u)$ and $(\bar{x} \vee B, w)$ are called the *prior clashing clauses*, $(A \vee B, m)$ is called the *resolvent*, $(x \vee A, u-m)$ and $(\bar{x} \vee B, w-m)$ are called the *posterior clashing clauses*, and $(x \vee A \vee \bar{B}, m)$ and $(\bar{x} \vee \bar{A} \vee B, m)$ are called the *compensation clauses*. The effect of Max-SAT resolution, as in classical resolution, is to infer (namely, make explicit) a connection between A and B. However, there is an important difference between classical resolution and Max-SAT resolution. While the former yields the *addition* of a new clause, Max-RES is a transformation rule. Namely, it requires the *replacement* of the left-hand clauses by the right-hand clauses. The reason is that some cost of the prior clashing clauses must be substracted in order to *compensate* the new inferred information. Consequently, Max-RES is better understood as a *movement* of knowledge in the formula.

The resolution rule for Max-SAT preserves equivalence as defined in Section 2. The last two compensation clauses may lose the clausal form, so the following rule [6] may be needed to recover it:

$$CNF(A \vee \overline{l \vee B}, u) = \left\{ \begin{array}{rcl} A \vee \bar{l} & : & |B| = 0 \\ \{(A \vee \bar{l} \vee B, u)\} \cup CNF(A \vee \bar{B}, u) & : & |B| > 0 \end{array} \right.$$

Example 4 *If we apply weighted resolution to the following clauses* $\{(x_1 \vee x_2, 3), (\bar{x}_1 \vee \underline{x_2 \vee x_3}, 4)\}$ *we obtain* $\{(x_2 \vee x_2 \vee x_3, 3), (x_1 \vee x_2, 3-3), (\bar{x}_1 \vee x_2 \vee x_3, 4-3), (x_1 \vee x_2 \vee \overline{(x_2 \vee x_3)}, 3), (\bar{x}_1 \vee \bar{x}_2 \vee x_2 \vee x_3, 3)\}$. *The first clause can be simplified. The second clause can be omitted because it weight is zero. The fifth clause can be omitted because it*

[1] When A is the empty clause $\bar{A}$ represents a tautology.

Algorithm 3: MINIMAXSAT basic structure.

```
    Function Search() : integer
17      ub := LocalSearch() ;
18      InitQueue(Q) ;
19      Loop
20          Propagate() ;
21          if Hard Conflict then
                AnalyzeConflict() ;
                if Top Level Hard Conflict then return ub ;
                else
                    LearnClause() ;
                    Backjump() ;
22          else if Soft Conflict then
                ChronologicalBactrack() ;
                if End of Search then return ub ;
23          else if all variables assigned then
                ub := lb ;
24              if ub = 0 then return ub ;
25              ChronologicalBactrack() ;
                if End of Search then return ub ;
26          else
                l := SelectLiteral() ;
                Enqueue(Q, l) ;
```

is a tautology. We apply CNF rule to the fourth clause to obtain two new clauses $CNF(x_1 \vee x_2 \vee \overline{(x_2 \vee x_3)}, 3) = \{(x_1 \vee x_2 \vee \bar{x}_2 \vee x_3, 3), (x_1 \vee x_2 \vee \bar{x}_3, 3)\}$. *Note that the first new clause is a tautology. Therefore, we obtain the equivalent formula* $\{(x_2 \vee x_3, 3), (\bar{x}_1 \vee x_2 \vee x_3, 1), (x_1 \vee x_2 \vee \bar{x}_3, 3)\}$.

5 Overview of MINIMAXSAT

Given a WCNF formula (possibly containing hard and soft clauses), MINIMAXSAT returns the cost of the optimal model (or $\top$ if there is no model). Like MINISAT, it traverses depth-first the tree of assignments. At each search point, the algorithm tries to simplify the current formula and, ideally, detect a conflict, which would mean that the current partial assignment cannot be successfully extended. MINIMAXSAT distinguishes two types of conflicts: hard and soft. *Hard* conflicts indicate that there is no model extending the current partial assignment (namely, all the mandatory clauses cannot be simultaneously satisfied). Hard conflicts are detected taking only into account hard clauses and using the methods of MINISAT. When a hard conflict occurs, MINIMAXSAT learns a hard clause and backjumps as MINISAT would do. *Soft* conflicts in-

Algorithm 4: MiniMaxSAT propagation.

```
   Function MS-UP() : conflict
       while (Q contains non-propagated literals) do
27         l := GetOldestLit(Q); MarkAsPropagated(l) ;
28         lb := lb + V(l̄)) ;
29         if lb ≥ ub then return Soft Conflict ;
30         foreach Hard clause (C∨l̄, ⊤) that becomes unit or falsified do
31             if (C ∨ l̄, ⊤) becomes unit (q, ⊤) then Enqueue(Q, q) ;
32             else if (C ∨ l̄, ⊤) becomes falsified then return Hard Conflict ;
33         foreach Soft clause (C∨l̄, u) that becomes unit do
34             if (C∨l̄, u) becomes a unit (q, u) then V(q) := V(q) + u ;
       return None ;
   Function Propagate() : conflict
35     c := MS-UP( ) ;
36     if c = Hard or Soft Conflict then return c ;
37     improveLB( ) ;
38     if lb ≥ ub then return Soft Conflict ;
39     return None ;
```

dicate that the current partial assignment cannot be extended to an optimal assignment. In order to identify soft conflicts, the algorithm maintains two values during search:

- The cost of the best model found so far, which is an upper bound ub of the optimal solution.
- An underestimation of the best cost that can be achieved extending the current partial assignment into a model, which is a lower bound lb of the current subproblem.

A soft conflict is detected when $lb \geq ub$, because it means that the current assignment cannot lead to an optimal model. When a soft conflict is detected, the algorithm backtracks chronologically. Note that any soft clause (C, w) with $w \geq ub$ must be satisfied in an optimal assignment. Therefore, in the following we assume that such soft clauses are automatically transformed into hard clauses.

An algorithmic description of MINIMAXSAT is presented in *Algorithm 3*. Before starting the search, a good initial upper bound is obtained with a local search method (line 17) which may yield the identification of some new hard clauses. In our current implementation we use UBCSAT [22] with default parameters. The selected local search algorithm is *IROTS* (*Iterated Robust Tabu Search*). Next, the queue Q is initialized with all unit hard clauses in the resulting formula (line 18). The main loop starts in line 19 and each iteration is in charge of propagating all pending literals (line 20) and, if no conflict is detected, attempting the extension of the current partial assignment (line 26). Pending literals in Q are propagated in function `Propagate` (line 20), which may return a hard or soft conflict. If a hard conflict is encountered (line 21) the conflict is analyzed, a new hard clause is learned and backjumping is performed. This is done as introduced in Section 3. If a soft conflict is encountered (line 22) chronological backtracking is

performed. If no conflict is found (line 26), a literal is heuristically selected and added to Q for propagation in the next iteration. However, if the current assignment is complete (line 23), the upper bound is updated. Search stops if a zero-cost solution is found, since it cannot be further improved (line 24). Else, chronological backtracking is performed (line 25). Note that backjumping leads to termination if a top level hard conflict is found, while chronological backtracking leads to termination if the two values for the first assigned variable have been tried.

Algorithm 4 describes the propagation process (function `Propagate`). It uses an array $V(l)$ which accumulates the weight of all soft clauses that have become unit over l (namely, original clauses $(A \vee l, w)$ such that the current assignment falsifies A). First of all, it performs a Max-SAT-adapted form of unit propagation (`MS-UP`, line 35). `MS-UP` iterates over the non-propagated literals l in Q (line 27). Firstly, the cost of falsifying $\bar{l}$ (which is recorded in $V(\bar{l})$) is added to the lower bound (line 28). If the lower bound increment identifies a soft conflict, it is returned (line 29). Then, if a hard clause becomes a unit, the corresponding literal is added to Q for future propagation (line 31). Finally, if a soft clause becomes a unit clause (q,u) (line 33), its weight u is added to $V(q)$ (line 34). If during this process a hard conflict is detected, the function returns it (lines 32,36). Else, the algorithm attempts to detect a soft conflict with a call to procedure `improveLB` (line 37), and it returns the soft conflict if it is found (line 38). In the next section a detailed description of `improveLB` can be found. Finally, if no conflict is detected, the function returns *None* (line 39).

6 Lower Bounding in MiniMaxSat

In the following, we consider an arbitrary search state of MiniMaxSat before the call to `improveLB`. Such a search state is uniquely characterized by the current assignment. The current assignment determines the *current subformula* which is the original formula *conditioned* by the current assignment: If a clause contains a literal that is part of the current assignment, it is removed. Besides, all the literals whose negation appear in the current assignment are removed from the clauses where they appear.

The value of lb maintained by MiniMaxSat is precisely the aggregation of costs of all the clauses that have become empty due to the current assignment. Similarly, the value $V(l)$ is the aggregation of costs of all the clauses that have become unit over l due to the current assignment. Thus, the current subformula contains $(\square, lb)$ and $(l, V(l))$ for every l.

MiniMaxSat computes its lower bound by deriving new soft empty clauses $(\square, w)$ through a resolution process. Such clauses are added to the already existing clause $(\square, lb)$ producing an increment of the lower bound.

As a first step, `improveLB` replaces each occurrence of (l,u) and $(\bar{l},w)$ by $(l, u-m), (\bar{l}, w-m), (\square, m)$ (with $m = \min\{u,w\}$), which amounts to applying a restricted version of Max-SAT resolution known as Unit Neighborhood Resolution (UNR) [6].

It produces an immediate increment of the lower bound (*i.e.*, the weight of the empty clause, line 43) as it is illustrated in the following example,

Example 5 *Consider the current state is* $\{(\square,3),(x_1,1),(x_2,1),(\bar{x}_1,2),(\bar{x}_2,2),(x_1 \vee x_2,3)\}$. *UNR would resolve on clauses* $(x_1,1)$ *and* $(\bar{x}_1,2)$ *replacing them by* $(\bar{x}_1,1)$ *and* $(\square,1)$

Algorithm 5: Lower Bounding in MINIMAXSAT

```
   Function SUP() : conflict
40 |  InitQueue(Q) ;
   |  while (Q contains non-propagated literals) do
   |  |  l := GetFirstNonPropagatedLiteral(Q);  MarkAsPropagated(l) ;
41 |  |  foreach (Hard or Soft) Clause C ∨ l̄ that becomes unit or falsified do
   |  |  |  if C ∨ l̄ becomes a unit q then Enqueue(Q, q) ;
   |  |  |  else if C ∨ l̄ becomes falsified then return conflict ;
   |  return None ;
   Function improveLB() : lb
42 |  foreach (l, v), (l̄, w) ∈ F do
43 |  |  replace them by (l, v − m), (l̄, w − m), (□, m) with m := min(v, w) ;
44 |  while SUP() = conflict do
45 |  |  Υ := BuildTree() ;
46 |  |  m := minimum weight among clauses in Υ;
47 |  |  if Condition then ApplyResolution( Υ, m ) ;
48 |  |  else lb := lb + m; remove weight m from clauses in Υ;
```

(all other compensation clauses are removed because their weight is zero or they are tautologies). The two empty clauses can be grouped into $(\square, 3+1=4)$. *UNR would also resolve on clauses* $(x_2, 1)$ *and* $(\bar{x}_2, 2)$ *replacing them by* $(\bar{x}_2, 1)$ *and* $(\square, 1)$. *The two empty clauses can be grouped into* $(\square, 4+1=5)$. *So, the new equivalent formula is* $\{(\square, 5), (\bar{x}_1, 1), (\bar{x}_2, 1), (x_1 \vee x_2, 3)\}$ *with a higher lower bound of* 5.

As a second step `improveLB` executes a *simulation of unit propagation* (`SUP`, line 44) in which soft clauses are treated as if they were hard. First, `SUP` adds to Q all unit soft clauses (as if they were hard, line 40). Then, the new literals in Q are propagated. When new (hard or soft) clauses become unit, they are inserted in Q (line 41). If `SUP` yields a conflict, it means that there is a subset of (soft or hard) clauses that cannot be simultaneously satisfied. We showed in Section 3 that Q can be used to identify such subset and build a refutation tree Υ. `ImproveLB` computes such a tree (line 45). If we take into account again the weights of the clauses and apply Max-SAT resolution (Section 4) as dictated by Υ, one can see that it will produce a new clause $(\square, m)$, where m is the minimum weight among all the clauses in the tree (line 46). It means that the extension of the current partial assignment to the unassigned variables will have a cost of at least m.

Once the resolution tree Υ is computed, Max-SAT resolution can actually be done with the soft clauses and, as a result, the empty clause $(\square, m)$ can be derived.

It is important to remark that at each step in the Max-SAT resolution process we do not consider the minimum of the weight of the two clauses, but rather the minimum of all the clauses in the resolution tree. This is why m is passed as a parameter in line 47.

The result of the resolution process is the replacement of all the clauses in the leaves of Υ by $(\square, m)$ and the corresponding compensation clauses (function `ApplyResolution`

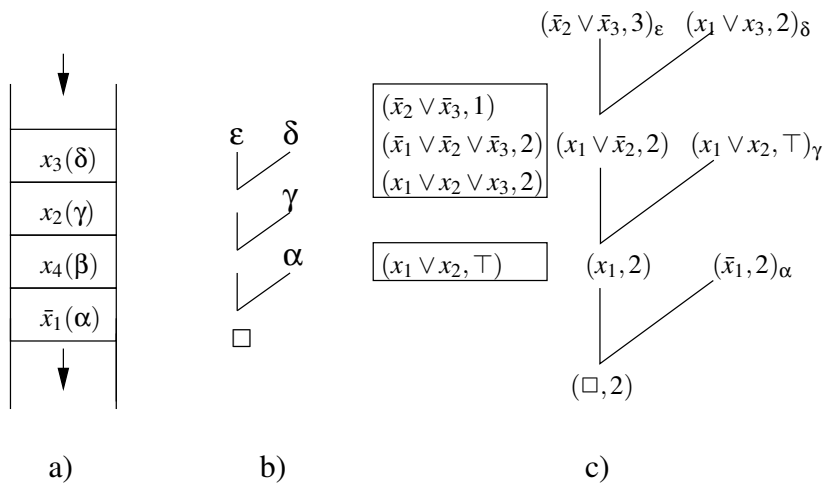

Fig. 2. Graphical representation of MINIMAXSAT lower bounding. On the top, the original current formula $\mathcal{F}$. On the left, the propagation Q after step 1. In the middle, the structure of the refutation tree computed by the simulation of UP in step 2. On the right, the effect of actually executing the Max-SAT resolution (step 3). The resulting formula $\mathcal{F}'$ appears bellow. If substraction-based lower bounding is performed, step 3 is replaced by a substraction of weights, producing formula $\mathcal{F}''$.

in line 47), thus obtaining an equivalent formula with a lower bound increment of m. We call this procedure *resolution-based* lower bounding.

Example 6 *Consider formula* $\mathcal{F} = \{(\bar{x}_1, 2)_\alpha, (x_1 \vee x_4, 1)_\beta, (x_1 \vee x_2, \top)_\gamma, (x_1 \vee x_3, 2)_\delta, (\bar{x}_2 \vee \bar{x}_3, 3)_\varepsilon\}$

Step 1. Apply SUP. *Initially, the unit clause* α *is enqueued producing* $Q = [\bar{x}_1(\alpha)\|]$. *Then* $\bar{x}_1$ *is propagated and* Q *becomes* $[\bar{x}_1(\alpha)\|x_4(\beta), x_2(\gamma), x_3(\delta)]$. *Literal* x_4 *is propagated and no new unit clauses are added. Literal* x_2 *is propagated producing* $Q = [\bar{x}_1(\alpha), x_4(\beta), x_2(\gamma)\|x_3(\delta)]$. *Since clause* ε *is conflicting,* SUP *stops. Figure 2.a shows the state of* Q *after the propagation.*

Step 2. Build the simulated refutation tree. *Starting from the tail of* Q *the first clause clashing with the conflicting clause* ε *is* δ. *Resolution between* ε *and* δ *generates the resolvent* $x_1 \vee \bar{x}_2$. *The first clause clashing with it is* γ, *producing resolvent* x_1. *The next*

clause clashing with it is α *and resolution generates* $\square$. *Figure 2.b shows the resulting resolution tree.*

Step 3. Apply Max-SAT resolution. *We apply Max-SAT resolution as indicated by the refutation tree computed in Step 2. Figure 2.c graphically shows the result of the process. Leaf clauses are the original (weighted) clauses involved in the resolution. Each internal node indicates a resolution step. The resolvents appear in the junction of the edges. Beside each resolvent, inside a box, there are the compensation clauses that must be added to the formula to preserve equivalence. Since clauses that are used in resolution must be removed, the resulting formula* $\mathcal{F}'$ *consists of the root of the tree (*$(\square, 2)$*) and all compensation clauses. That is, the resulting formula is* $\mathcal{F}' = \{(x_1 \vee x_4, 1), (x_1 \vee x_2, \top), (\bar{x}_2 \vee \bar{x}_3, 1), (\square, 2), (x_1 \vee x_2 \vee x_3, 2), (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3, 2)\}$. *The soundness of Max-SAT resolution guarantees that* $\mathcal{F} \equiv \mathcal{F}'$.

Remark 1 *All the transformations applied by the* resolution-based lower bounding *can be passed on to descendent nodes because the changes preserve equivalence. Nevertheless, transformations have to be restored when backtracking takes place.*

An alternative to problem transformation through resolution is to identify the lower bound increment m and then substract it from all the clauses that would have participated in the resolution tree. This procedure is similar to the lower bound computed in [8] and we call it *substraction-based* (line 48) lower bounding.

Example 7 *Consider formula* $\mathcal{F}$ *from the previous example. Steps 1 and 2 are identical. However, substraction-based lower bounding would replace Step 3 by Step 3' that substracts weight 2 from the clauses that appear in the refutation tree and then adds* $(\square, 2)$ *to the formula. The result is* $\mathcal{F}'' = \{(x_1 \vee x_4, 1), (x_1 \vee x_2, \top), (\bar{x}_2 \vee \bar{x}_3, 1), (\square, 2)\}$. *Note that* $\mathcal{F}'' \sqsubseteq \mathcal{F}$.

Remark 2 *All the substractions applied by the* substraction-based lower bounding *have to be restored before moving to a descendent node because they do not preserve equivalence.*

After the increment of the lower bound with either technique, procedure `SUP` can be executed again, which may yield new lower bound increments. The process is repeated until `SUP` does not detect any conflict.

When comparing the two previous approaches, we observe that resolution-based lower bounding has a larger overhead, because resolution steps need to be actually computed and their consequences must be added to the current formula and removed upon backtracking. However, the effort invested in the transformation may be well amortized because the increment obtained in the lower bound *becomes part of the current formula*, so it does not have to be *discovered* again and again by all the descendent nodes of the search. On the other hand, substraction-based lower bounding has a smaller overhead because resolution needs not to be actually computed. This also facilitates the context restoration upon backtracking.

MINIMAXSAT incorporates the two alternatives and chooses to apply one or the other heuristically (lines 47,48) depending on an specific *condition* (line 47). We observed that resolution-based lower bounding seems to be more effective if resolution is

only applied to low arity clauses. As a consequence, after the identification of the resolution tree, MINIMAXSAT applies resolution-based lower bounding only if the largest resolvent in the resolution tree has arity less than 4. Otherwise, it applies substraction-based lower bounding. See Section 8 for more details.

7 Additional features of MINIMAXSAT

7.1 Two-Watched literals

MINIMAXSAT extends the *two-watched literals scheme* to soft clauses. Recall that the main advantage of this technique, when applied to pure SAT problems, is that when backtracking takes place, no work has to be done on the clauses. Unfortunately, in the case of soft clauses some restoration needs to be done. When a soft clause becomes unit over literal l in function *MS-UP*, its weight is added to $V(l)$ and the clause is eliminated (or marked as eliminated) to avoid reusing it in the lower bounding procedure. These changes, as well as any addition to lb, have to be restored when backtracking is performed. However, note that during the executions of `SUP` (simulation of unit propagation) all clauses are considered as hard. In this case the two-watched literals works exactly as in a SAT solver with both hard and soft clauses. When an inconsistency is detected by `SUP` or it stops because there are no more literals to propagate, the initial state has to be recovered. In that situation restoring the initial state is completely overhead free.

7.2 Branching Heuristic

MINIMAXSAT incorporates two alternative branching heuristics. The first one is the VSIDS heuristic [20] disregarding soft clauses (that is, MINISAT'S default). This heuristic is likely to be good in structured problems in which learning and backjumping play a significant role, as well as in problems in which it is difficult to find models (namely, the satisfaction component of the problem is more difficult than the optimization component). Since this heuristic disregards soft clauses, it is likely to be ineffective in problems where it is easy to find models and the difficulty is to find the optimal one and proof its optimality. In the extreme case, where problems only contain soft clauses (every complete assignment is a model) the VSIDS heuristic is blind and therefore completely useless.

To overcome this limitation of VSIDS, MINIMAXSAT also incorporates the *Weighted Jeroslow* heuristic [17]. It is the extension of the SAT Jeroslow heuristic described in Section 3. Given a weighted formula F, for each literal l of F the following function is defined:

$$J(l) = \sum_{\substack{(C,w)\in F \\ s.t.\ \ l\in C}} 2^{-|C|} \cdot w$$

where mandatory clauses are assumed to have weight equal to the upper bound ub. The heuristic selects the literal with the highest value of $J(l)$. Its main disadvantage is that metrics need to be updated at each visited node. In combination with the two-watched literals this updating becomes expensive and does not seem to pay off in general. Thus,

in our current implementation of the heuristic, the $J(l)$ values are computed only at the root node and used throughout all the solving process. We found in our experiments that this heuristic is a good alternative in problems where the difficulty lies on the optimization part (e.g. problems with many models). MINIMAXSAT automatically changes from VSIDS to weighted Jeroslow if the problem does not contain any literal l such that there is some hard clause with l and some other hard clause with $\bar{l}$.

In both heuristics, if there is some literal l such that $V(l)+lb \geq ub$ at some node of the search tree, then $\bar{l}$ is the selected literal and l is never assigned.

7.3 Soft probing

Probing is a well-known SAT technique that allows the formulation of hypothetical scenarios [23]. The idea is to temporarily assume that l is a hard unit clause and then execute unit propagation. If UP yields a conflict, we know that any model extending the current assignment must contain $\bar{l}$. The process is iterated over all the literals until quiescence. Exhaustive experiments in the SAT context indicate that it is too expensive to probe during search [24, 23], so it is normally done as a pre-process in order to reduce the initial number of branching points.

We can easily extend this idea to Max-SAT. In that context, besides the *discovery* of unit hard clauses, it may be used to make explicit weighted unit clauses. We call it *soft probing*. As in SAT, the idea is to temporarily assume that l is a unit clause and then *simulate* unit propagation (*i.e.*, execute `SUP()`). Then, we build the resolution tree Υ from the propagation queue Q. If all the clauses in Υ are hard, we know that $\bar{l}$ must be added to the assignment. Else, we can reproduce Υ applying Max-SAT resolution with the weighted clauses and derive a unit clause $(\bar{l},m)$ where m is the minimum weight among the clauses in Υ. Having unit soft clauses upfront makes the future executions of `improveLB` much more effective in the subsequent search. Besides, if we derive both (l,u) and $(\bar{l},w)$, we can generate via unit neighborhood resolution (see Example 5) an initial non-trivial lower bound of $min\{u,w\}$. We tested soft probing during search and as a preprocessing in several benchmarks. We observed empirically that soft probing as a preprocessing was the best option as it is in SAT.

Example 8 *Consider formula* $\mathcal{F} = \{(x_1 \vee x_2, 1)_\alpha, (x_1 \vee x_3, 1)_\beta, (\bar{x}_2 \vee \bar{x}_3, 1)_\gamma\}$. *If we assume* $\bar{x}_1$ *by adding it to* Q *and then execute* `SUP` *a conflict is reached. We obtain* $Q = [\bar{x}_1^d, x_2(\alpha), x_3(\beta)]$ *and we detect that* γ *is a conflicting clause. The clauses involved in the refutation are* γ, β, *and* α. *Resolving clauses* γ *and* β *results in* $\{(x_1 \vee x_2, 1)_\alpha, (x_1 \vee \bar{x}_2, 1), (x_1 \vee x_2 \vee x_3, 1), (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3, 1)\}$. *The resolution of the previous resolvent and* α *produces the (equivalent) formula* $\mathcal{F}' = \{(x_1, 1), (x_1 \vee x_2 \vee x_3, 1), (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3, 1)\}$.

7.4 Pseudo-boolean functions

A *pseudo-boolean optimization problem* (*PBO*) [25, 15] has the form:

minimize $\sum_{j=1}^{n} c_j \cdot x_j$

subject to $\sum_{j=1}^{n} a_{ij} l_j \geq b_i, \quad i = 1 \ldots m$

where $x_j \in \{0,1\}$, l_j is either x_j or $1 - x_j$, and c_j, a_{ij} and b_i are non-negative integers. If MINIMAXSAT is provided with a PBO instance, it translates it into a Max-SAT formula as follows: each pseudo boolean constraint is translated into a set of *hard clauses*

using MINISAT+ [15] (the algorithm heuristically decides the most appropriate translation choosing among *adders*, *sorters* or *BDDs*). The objective function is translated into a set of *soft unit clauses*. Each summand $c_j \cdot x_j$ becomes a new soft unit clause $(\bar{x}_j, c_j)$. After the translation MINIMAXSAT is executed as usual.

8 Empirical Results

8.1 Benchmarks and Encodings

Having a good set of problems is fundamental to show the effectiveness of new solvers and tools. In the following, we present several problems and we explain how to encode them as Weighted Max-SAT.

Max-k-SAT A k-SAT CNF formula is a CNF formula in which all clauses have size k. We generated random unsatisfiable 2-SAT and 3-SAT formulas with the *Cnfgen* generator[2] and solved the corresponding MAX-SAT problem. In the benchmarks, we fixed the number of variables and varied the number of clauses.

Max-one Given a satisfiable CNF formula, *max-one* is the problem of finding a model with a maximum number of variables set to true. This problem, with a number of applications in circuit design, can be encoded as Max-SAT by considering the clauses in the original formula as mandatory and adding a weighted unary clause $(x_i, 1)$ for each variable in the formula. Note that solving this problem is much harder than solving the usual SAT problem, because the search cannot stop as soon as a model is found. The optimal model must be found and its optimality must be proved. We considered the max-one problem over two types of CNF formula: random 3-SAT instances of 120 variables (generated with *Cnfgen*), and structured satisfiable instances coming from the 2002 SAT Competition[3].

Max-cut Given a graph $G = (V, E)$, a *cut* is defined by a subset of vertices $U \subseteq V$. The size of a cut is the number of edges (v_i, v_j) such that $v_i \in U$ and $v_j \in V - U$. The *Max-cut* problem consists on finding a cut of maximum size. It can be encoded as Max-SAT associating one variable x_i to each graph vertex. Value *true* (respectively, *false*) indicates that vertex v_i belongs to U (respectively, to $V - U$). For each edge (v_i, v_j), there are two clauses $x_i \vee x_j, \bar{x}_i \vee \bar{x}_j$. Given a complete assignment, the number of violated clauses is $|E| - S$ where S is the size of the cut associated to the assignment. In our experiments we considered Max-Cut instances extracted from random graphs of 60 nodes with varying number of edges.

[2] A. van Gelder ftp://dimacs.rutgers.edu/pub/challenge/satisfiability/contributed/UCSC/instances
[3] http://www.satcompetition.org/2002/

Minimum Vertex Covering and Max-Clique Given a graph $G = (V,E)$, a *vertex covering* is a set $U \subseteq V$ such that for every edge (v_i, v_j) either $v_i \in U$ or $v_j \in U$. The size of a vertex covering is $|U|$. The *minimum vertex covering* problem consists in finding a covering of minimal size. It can be naturally formulated as (weighted) Max-SAT. We associate one variable x_i to each graph vertex v_i. Value *true* (respectively, *false*) indicates that vertex x_i belongs to U (respectively, to $V - U$). There is a binary weighted clause $(x_i \vee x_j, u)$ for each edge (v_i, v_j), where u is a number larger than $|V|$. It specifies that one or both of these two vertices have to be in the covering because there is an edge connecting them. There is a unary clause $(\bar{x}_i, 1)$ for each variable x_i, in order to specify that it is preferred not to add vertices to U. There is a simple way to transform minimum vertex coverings into max-cliques and vice-versa [26].

In our experiments, we considered maximum clique instances extracted from random graphs with 150 nodes and varying number of edges. We also considered the 66 Max-Clique instances from the DIMACS challenge[4].

Combinatorial Auctions A *combinatorial auction* is defined by a set of goods G and a set of bidders that bid for indivisible subsets of goods. Each bid i is defined by the subset of requested goods $G_i \subseteq G$ and the amount of money offered. The bid-taker, who wants to maximize its revenue, must decide which bids are to be accepted. Note that if two bids request the same good, they cannot be jointly accepted [27]. In its Max-SAT encoding, there is one variable x_i associated to each bid. There are unit clauses (x_i, u_i) indicating that if bid i is not accepted there is a loss of profit u_i. Besides, for each pair i, j of conflicting bids, there is a mandatory clause $(\bar{x}_i \vee \bar{x}_j, \top)$.

In our experiments, we used the CATS generator [28] that allows to generate random instances inspired from real-world scenarios. In particular, we generated instances from the *Regions*, *Paths* and *Scheduling* distributions. The number of goods was fixed to 60 and we increased the number of bids. By increasing the number of bids, instances become more constrained (namely, there are more conflicting pairs of bids) and harder to solve.

Miscellaneous We also considered the following sets of instances widely used in the literature:

- The unsatisfiable instances of the 2nd DIMACS Implementation Challenge [5] considered in [29, 8]: random 3-SAT instances (aim and dubois), pigeon hole problem (hole) and coloring problems (pret). Observe that all these instances are modelled as unweighted Max-SAT (i.e. all clauses have weight 1).
- *Max-CSP* random instances generated using the protocol specified in [30, 31]. We distinguish 4 different sets of problems: *Dense Loose* (*DL*), *Dense Tight* (*DT*), *Sparse Loose* (*SL*) and *Sparse Tight* (*ST*). Each set contain 10 instances with 3 values and 10 instances with 4 values per variable.

[4] ftp://dimacs.rutgers.edu/pub/challenge/graph/benchmarks/clique

[5] http://mat.gsia.cmu.edu/challenge.html

- *Planning* [32] and *graph coloring* [6] structured instances taken from a *Weighted Constraint Satisfaction Problem* (*WCSP*) repository [7].
- Problems taken from the 2006 pseudo-boolean evaluation [8]: *logic synthesis*,*misc* (*garden*), *routing*, *MPI* (*Minimum Prime Implicant)*, *MPS* (*miplib*). These instances are encoded to Max-SAT as specified in the previous section.

Note that Max-CSP, Planning and Graph Coloring instances are encoded into Max-SAT using the *direct encoding* [33].

8.2 Alternative solvers

We compare MINIMAXSAT with several optimizers from different communities. We restricted our comparison to freely available solvers. We considered the following ones:

- MAXSATZ [9, 34]. Unweighted Max-SAT solver. It was the best unweighted Max-SAT solver in the 2006 Max-SAT Evaluation.
- MAX-DPLL [17, 6]. Weighted Max-SAT solver. It is part of the TOOLBAR package. It was the best solver for weighted Max-SAT and the second best solver for unweighted Max-SAT in the 2006 Max-SAT Evaluation.
- TOOLBAR [35, 30, 29, 31]. It is a state-of-the-art Weighted CSP solver.
- PUEBLO 1.5 [25]. It is a pseudo-boolean solver. It ranked first on several categories of the 2005 Pseudo Boolean Evaluation.
- MINISAT+ [15]. It is a pseudo-boolean solver that translates the problems into SAT and solves them with MiniSAT. It ranked first on several categories of the 2005 Pseudo Boolean Evaluation.

Those instances taken from the pseudo-boolean evaluation were given in their original format to Pueblo and Minisat+. All other instances were translated from Max-SAT to PBO by partitioning the set of clauses into three sets: $\mathcal{H}$ contains the mandatory clauses $(C, \top)$, $\mathcal{W}$ contains the non-unary weighted clauses $(C, u < \top)$ and $\mathcal{U}$ contains the unary weighted clauses (l, u). For each hard clause $(C_j, \top) \in \mathcal{H}$ there is a pseudo boolean constraint $C'_j \geq 1$, where C'_j is obtained from C_j by replacing $\vee$ by $+$ and negated variables $\bar{x}$ by $1 - x$. For each non-unary weighted clause $(C_j, u_j) \in \mathcal{W}$ there is a pseudo boolean constraint $C'_j + r_j \geq 1$, where C'_j is computed as before, and r_j is a new variable that, when set to 1, trivially satisfies the constraint. Finally, the objective function to minimize is,

$$\sum_{(C_j, r_j) \in \mathcal{W}} u_j r_j + \sum_{(l_j, u_j) \in \mathcal{U}} u_j l_j$$

[6] http://mat.gsia.cmu.edu/COLORING02/benchmarks
[7] http://mulcyber.toulouse.inra.fr/plugins/scmcvs/cvsweb.php/benchs/?cvsroot=toolbar
[8] http://www.cril.univ-artois.fr/PB06/

8.3 Experimental Results

We divide the experiments in two parts. The purpose of the first part is to evaluate the impact of the different techniques of MINIMAXSAT and set the different parameters. Since some of the techniques can be effective in some benchmarks and useless or even counterproductive in some others [36], we aimed at finding a configuration such that MINIMAXSAT performs reasonably well on all the instances. The purpose of the second part is to compare MINIMAXSAT with alternative solvers. Since some of these solvers are specifically designed for some type of problems, we do not expect that MINIMAXSAT will outperform them. We rather want to show the robustness of MINIMAXSAT by showing that it is usually close in performance with the best alternative for each type of problems.

Results are presented in plots and tables. Regarding tables, the first column contains the name of the set of problems. The second column shows the number of instances. The rest of columns report the performance of the different solvers. Each cell contains the average CPU time that the solver required to solve all instances. If some solver could not solve all the instances of a set, a number inside brackets indicates the number of solved instances and the average *cpu* time only takes into account solved instances. If a cell contains a dash, it means that no instance could be solved within the time limit. Regarding plots, note that the legend goes in accordance with the performance of the solvers. The time limit was set to 900 seconds for each instance.

Our solver, written in C++, was implemented on top of MINISAT+ [15]. Executions were made on a 3.2 Ghz Xeon computer with Linux. In all the experiments with random instances, samples had 30 instances and plots report mean *cpu* time in seconds.

8.4 Setting the parameters of MiniMaxSat

In the following we evaluate in order the importance of the following techniques inside MINIMAXSAT: lower bounding, soft probing, branching heuristics and learning and backjumping.

Starting from a basic version that guides search with the Jeroslow branching heuristic and has the rest of techniques deactivated, we analyze them one by one. Each analysis studies one technique and incorporates all the previously analyzed ones with the corresponding tuned parameters. In the three first experiments we only consider little but challenging instances generated randomly in which lower bounding plays a fundamental role to solve them. Finally, we consider structured instances in which learning and backjumping is required to solve them.

Lower bounding In this experiment we analyze the impact of resolution-based lower bounding versus substraction-based lower bounding, as well as combined strategies. We considered the following combination of the two techniques: when SUP detects and inconsistency and the refutation tree is computed, we look at the resolvent with maximum size. If its size is less than or equal to a parameter K, then resolution-based lower bounding is applied, otherwise substraction-based lower bounding is applied.

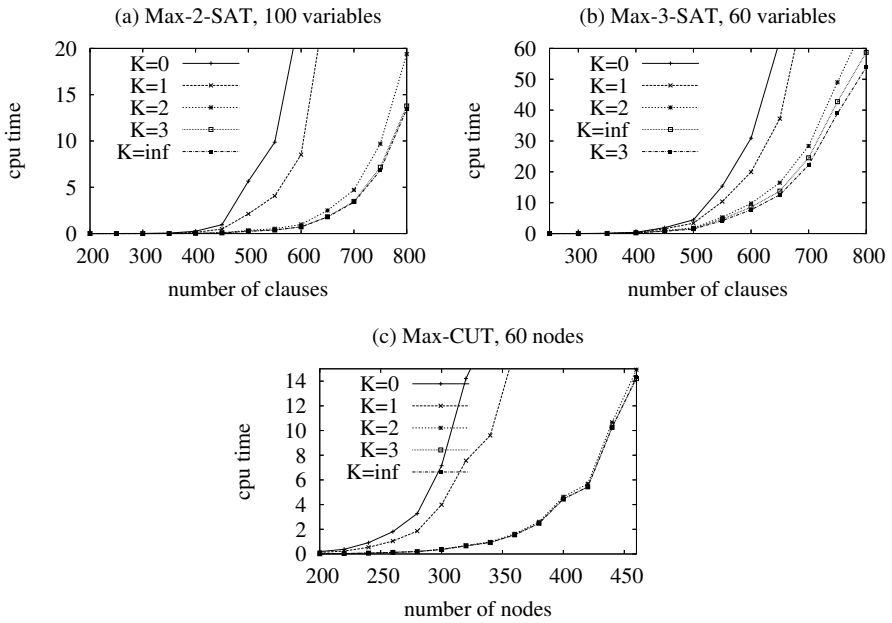

Fig. 3. Performance of MiniMaxSat with different mixed lower boundings ($K = 0, 1, 2, 3, \text{inf}$).

We tested $K = \{0, 1, 2, 3, 4, 5, \infty\}$. Note that $K = 0$ corresponds to pure substraction-based lower bounding (and therefore is similar to the approach of [8]), while $K = \infty$ corresponds to a pure resolution-based lower bounding.

The results are presented in figure 3. As can be seen, the pure substraction-based lower bounding $K = 0$ is always the worst option. Better results are obtained as K increases. However, the improvement stops (or nearly stops) when $K = 3$. When $K \geq 3$ no significant improvement is noticed. The plot omits the $K = 4$ and $K = 5$ case for clarity reasons. Since higher values of K may produce new clauses of higher size and this may cause overhead in some instances, we set $K = 3$ for the rest of the experiments.

Soft probing In our second experiment, we evaluate the impact of soft probing. In our preliminary experiments, we observed that soft probing was too time consuming, so we decided to limit soft probing as follows: Initially, we assign a *propagation level* of 0 to the variable to probe. Then, each new literal to propagate is assigned a *propagation level* $L+1$ if the literal that produces its propagation has level L. We limited probing to propagate literals with a maximum propagation level of M. We finally restricted $M \leq 2$ since it gives the best results. Note that a *propagation level* is not the same as a *decision level*.

We compare three alternatives: probing at each node of the search (S), probing as a pre-process before search (P) and no probing at all (N). The results, in Figure 4, indicates that probing during search is the worst option for Max-2-SAT and Max-3-SAT

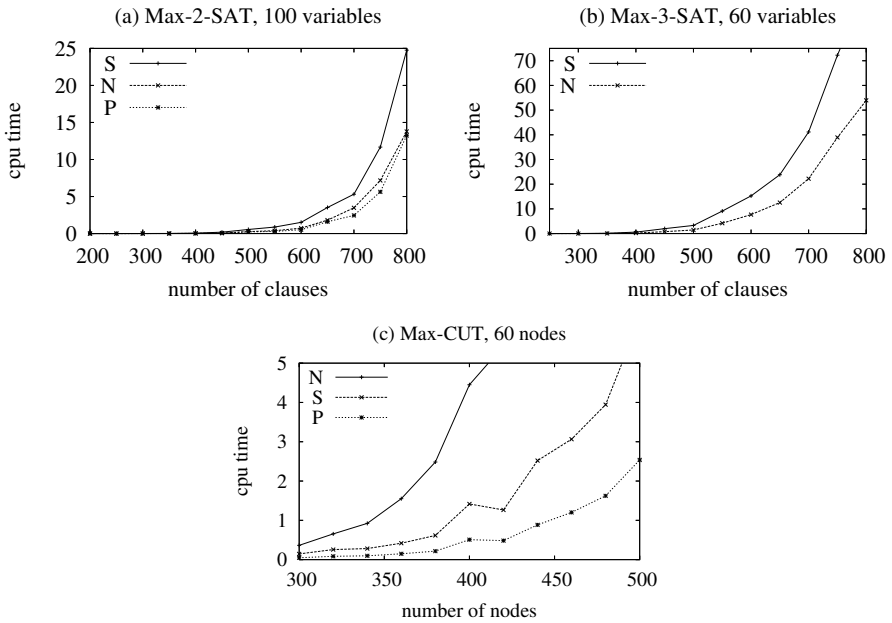

Fig. 4. Performance of MiniMaxSat without soft probing, with probing as preprocessing (P) and with probing during search (S).

while it produces some improvement in Max-CUT. Finally, probing as a preprocessing gives slightly improvement for Max-2-SAT and the best results for Max-CUT. Note that soft probing as a preprocessing on Max-3-SAT has no effect and is omitted from the plot (its results are similar to N). Given these results, we decided to include soft probing only as a preprocessing.

Jeroslow branching heuristic In the following experiment, we evaluate the importance of the weighted Jeroslow heuristic. Figure 5 shows the time difference between MINIMAXSAT with the Jeroslow heuristic as in the previous two experiments (Jeroslow) and without heuristic (None). The results indicates that guiding search with the Jeroslow heuristic gives important speed ups. Hence, we maintain the Jeroslow heuristic for MINIMAXSAT.

Learning, backjumping and VSIDS In the final experiment, we evaluate the importance of learning and backjumping. For these experiments we use structured instances, since it is well known that learning and backjumping are only useful in this type of problems. Besides, we also evaluate the importance of the VSIDS heuristic in combination with learning and backjumping. Recall that this heuristic was specially designed to work in cooperation with learning, so it is meaningless to analyze its effect by itself.

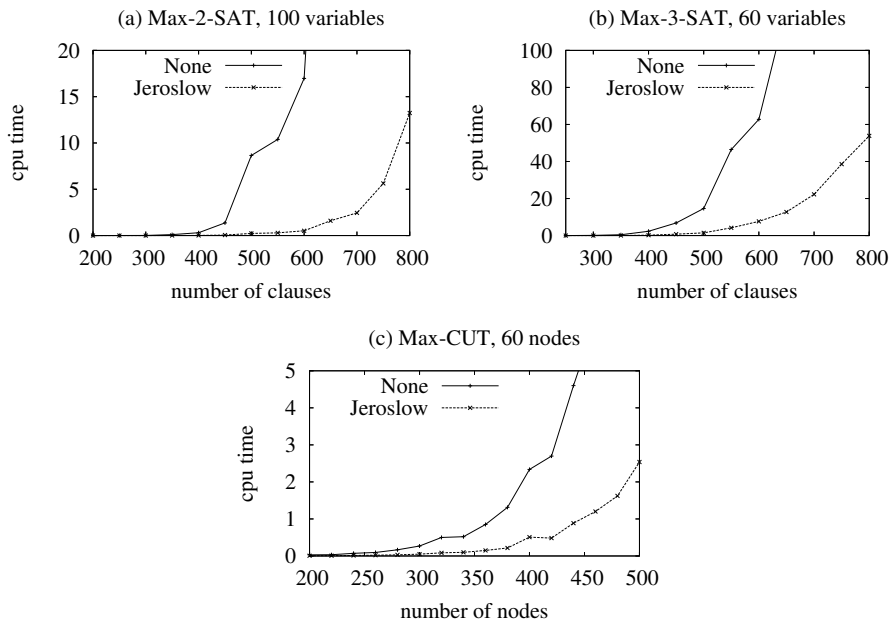

Fig. 5. Performance of MiniMaxSat without Heuristic (None) and with the Jeroslow heuristic computed in the root node of the search tree (Jeroslow).

Table 6 reports the results of this experiment. The third column reports results without learning and backjumping but with the lower bounding, probing and the Jeroslow heuristic (None). The fourth column reports results adding learning and backjumping to the previous version (Learning). The fifth column reports results adding learning, backjumping but changing the Jeroslow heuristic by the VSIDS heuristic (VSIDS). The results show that MINIMAXSAT without learning and backjumping (None) is clearly the worst option. Significant improvements are obtained when learning and backjumping (Learning) are added. Finally, adding the VSIDS heuristic (VSIDS) improve further the results specially on the routing instances. Based on those results, we incorporated learning and backjumping to MINIMAXSAT.

Regarding the branching heuristic, for problems in which literals appear in hard clauses with both polarities it applies the VSIDS heuristic, otherwise the Jeroslow heuristic is computed in the root of the search tree as stated in Section 7.

8.5 Comparison with other boolean optimizers

When reporting results, we will omit a solver if it cannot deal with the corresponding instances for technical reasons (e.g. it cannot deal with weighted clauses) or it performs extremely bad in comparison with the others.

Figure 7 contains plots with the results on different benchmarks. Plots *a* and *b* reports results on random unweighted Max-SAT instances. PUEBLO and MINISAT+ are

Problem	n. inst.	None	Learning	VSIDS
Max-One 3col	40	–	29.06	15.41
Max-One cnt	3	13.57(1)	119.53	6.58
Max-One dp	6	16.11(4)	40.03	28.63
Max-One ezfact32	10	654.94(2)	0.70	0.77
Routing S3	5	22.26(4)	1.02	0.10
Routing S4	10	–	410.61(2)	91.09(9)

Fig. 6. Structured instances.

orders of magnitude slower, so they are not included in the graphics. On Max-2-SAT (plot *a*), MINIMAXSAT lays between MAX-DPLL and MAXSATZ, which is the best option. On Max-3-SAT (plot *b*) MINIMAXSAT clearly outperforms MAX-DPLL and is very close to MAXSATZ, which is again the best. In both Max-2-SAT and Max-3-SAT MAXSATZ is no more than 3 times faster than MINIMAXSAT.

Plot *c* reports results on random Max-CUT instances. MINIMAXSAT performs slightly better than MAXSATZ, which is the second alternative.

Results regarding the unsatisfiable DIMACS instances are presented in Figure 8. Note that all these instances have optimum cost 1. Hence, as soon as MINIMAXSAT find a solution of cost 1, all the clauses are declared hard and learning and backjumping can be applied when hard conflicts arise. The results indicate that MAXSATZ and MAX-DPLL do not solve any instance on some sets (Pret150 and Aim200), while MINIMAXSAT solves all sets of instances with the best times in all of them, except for the hole instances in which MAXSATZ is slightly faster. If we encode these problems in the most advantageous way for PUEBLO and MINISAT+, that is, as decision problems rather than optimization problems they solve all the instances with similar times to MINIMAXSAT.

Plot *e* reports the results on Random Max-Clique instances. MINIMAXSAT is the best solver, up to an order of magnitude faster than MAX-DPLL, the second option. PUEBLO and MINISAT+ perform poorly again. Regarding the structured Dimacs instances, MINIMAXSAT is again the best option. It solves 36 instances within the time limit, while MAX-DPLL,MINISAT+ and PUEBLO solve 34, 22 and 18 respectively.

Plots *f*, *g* and *h* present the results on Combinatorial Auctions following different distributions. On the paths distribution, MINIMAXSAT is the best solver, twice faster than MAX-DPLL, which ranks second. On the regions distribution, MINIMAXSAT is the best solver while MAX-DPLL is the second best solver requiring double time. On the paths and regions distributions, PUEBLO and MINISAT+ perform very poorly. On the scheduling distribution, MINISAT+ is the best solver while MINIMAXSAT and MAX-DPLL are about one order of magnitude slower.

On random Max-One (plot *d*) MINIMAXSAT is the best solver by far. Almost all instances are solved instantly while PUEBLO and MAX-DPLL require up to 10 seconds in the most difficult instances. MINISAT+ performs very poorly. The results on structured Max-One instances are reported in Figure 9. MINISAT+ seems to be the fastest in general. MINIMAXSAT is close in performance to PUEBLO. Note, however, that in the *dp* instances, MINIMAXSAT is the system solving more instances.

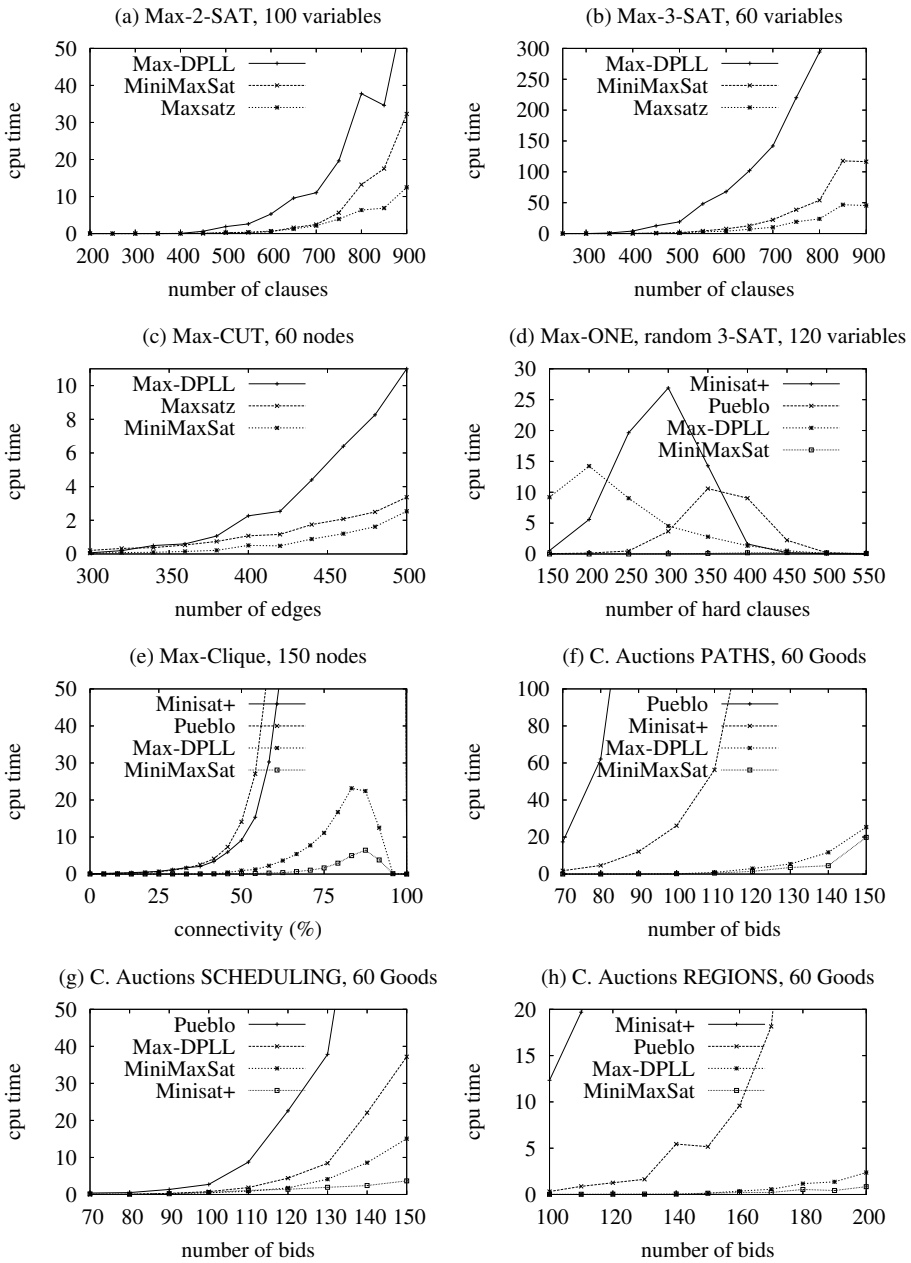

Fig. 7. Plots of different benchmarks. Note that the order in the legend goes in accordance with the performance of the solvers.

Problem	n. inst.	MINIMAXSAT	Maxsatz	Max-DPLL
Dubois	13	0.02	148.18(7)	174.33(6)
Pret60	4	0.07	10.06	22.00
Pret150	4	0.01	–	–
Hole	5	8.68	8.34	28.00
Aim50	8	0.00	0.01	0.00
Aim100	8	0.00	9.55	172.00
Aim200	8	0.00	–	–

Fig. 8. Unsatisfiable DIMACS instances.

On the planning instances (Fig. 10) PUEBLO is the best solver. MINIMAXSAT is the second best solver, TOOLBAR is the third and the last one is MINISAT+. This is not surprising since TOOLBAR does not perform learning over the hard constraints. Results regarding graph coloring instances are presented in Fig. 10. As can be observed, MINIMAXSAT is able to solve one more instance than TOOLBAR, while PUEBLO and MINISAT+ solve many less instances. On the Max-CSP problems (Fig. 10) TOOLBAR solves all the instances instantly while PUEBLO is the worst option unable to solve a lot of instances. MINIMAXSAT is clearly the second best solver and MINISAT+ is the third best performing solver. Note that both of them solve all the instances.

Results regarding the instances taken from the pseudo-boolean evaluation can be found in Figure 11. Note that this is the first time that a Max-SAT solver is tested on pseudo-boolean instances. Results indicate that no solver consistently outperforms the other and that MINIMAXSAT is fairly competitive with PUEBLO and MINISAT+.

From all these results we can conclude that MINIMAXSAT is a very robust Weighted Max-SAT solver. It is very competitive for pure optimization problems and for problems with lots of hard clauses and, sometimes, it is the best option.

As a final remark, note that MINIMAXSAT and almost all the previous benchmarks have been submitted to the *Second Max-SAT Evaluation 2007*, a co-located event of the *Tenth International Conference on Theory and Applications of Satisfiability Testing*. Hence, the interested reader will be able to find a more exhaustive comparison, including more instances and solvers, in the Second Max-SAT Evaluation 2007 web page[9].

9 Related Work

Some previous work has been done about incorporating SAT-techniques inside a Max-SAT solver. In [11] a lazy data structure to detect when clauses become unit is presented but it requires a static branching heuristic. In [37] different versions of a branch and bound are tested. One of these versions uses the two-watched literals, but it uses a very basic lower bounding. We can conclude that none of these previous approaches is as general as our extension of the two-watched literals. As far as we know, the rest of

[9] http://www.maxsat07.udl.es/

Problem	n. inst.	MINIMAXSAT	Pueblo	Minisat+
3col80	10	0.15	0.10	0.02
3col100	10	2.25	1.73	0.12
3col120	10	20.49	14.52	0.74
3col140	10	38.33	83.17	1.61
cnt	3	6.59	0.13	0.12
dp	6	28.81	1.19(3)	1.21(4)
ezfact32	10	0.77	0.34	0.33

Fig. 9. Structured Max-one instances.

Problem	n. inst.	Toolbar	MINIMAXSAT	Pueblo	Minisat+
Planning	71	4.02	3.81	0.16	7.40
Graph Coloring	22	49.29(16)	4.16(17)	68.50(11)	0.57(11)
Max-CSP DL	20	0.08	0.20	349.08(13)	8.60
Max-CSP DT	20	0.00	0.01	–	2.40
Max-CSP SL	20	0.01	0.03	123.67	0.48
Max-CSP ST	20	0.00	0.01	–	1.29

Fig. 10. Results for WCSP and Max-CSP instances.

Problem	n. inst.	MINIMAXSAT	Pueblo	Minisat+
misc	7	3.08(5)	8.51(5)	0.14(5)
Logic synthesis	17	82.55(2)	36.21(5)	253.93(5)
MPI	148	37.35(107)	32.04(101)	3.06(105)
MPS	16	22.65(5)	36.90(8)	8.50(8)
Routing	15	58.74(14)	5.96	13.09

Fig. 11. Results for pseudo-boolean instances.

Max-SAT solvers are based on *adjacency lists*. Therefore, they are presumably inefficient for unit propagation [38]. In [39] a Max-SAT branch and bound is enhanced with learning over hard constraints, but it is used in combination with simple lower bounding techniques. An improved version is presented in [12] with a more powerful lower bound, however it does not incorporate the two-watched literals scheme, backjumping, etc. To the best of our knowledge, no Max-SAT solver incorporates backjumping. Note that MINIMAXSAT restricts backjumping to the occurrence of hard conflicts. Related works on the integration of backjumping techniques into branch and bound include: [40] for *Weighted CSP*, [41] for pseudo-boolean optimization and [42] for *Sat Modulo Theories*.

Most Max-SAT solvers use variations of what we call substraction-based lower bounding. In most cases, they search for special patterns of mutually inconsistent subsets of clauses [4, 7, 11]. For efficiency reasons, these patterns are always restricted to small sets of small arity clauses (2 or 3 clauses or arity less than 3). MINIMAXSAT uses

a natural weighted extension of the approach proposed in [8]. It was the first one able to detect inconsistencies in arbitrarily large sets of arbitrarily large clauses.

The idea of what we call resolution-based lower bounding was inspired from the WCSP domain [35, 30, 29, 31] and it was first proposed in the Max-SAT context in [5] and further developed in [34, 17, 6]. In these works, only special patterns of fixed-size resolution trees were executed. The use of simulated unit propagation allows MINIMAXSAT to identify arbitrarily large resolution trees. In the following example, we present two inconsistent subsets of clauses that are detected by MINIMAXSAT and transformed into an equivalent formula while previous solvers cannot transform them since they are limited to specific patterns:

- $\{(x_1, w_1), (x_2, w_2), (x_3, w_3), (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3, w_4)\}$
- $\{(x_1, w_1), (\bar{x}_1 \vee x_2, w_2), (\bar{x}_1 \vee \bar{x}_2 \vee x_3, w_3), (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3 \vee x_4, w_4), (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3 \vee \bar{x}_4, w_5)\}$

In the first case, MINIMAXSAT replaces the clauses by $(\square, m)$ with $m = \min\{w_1, w_2, w_3, w_4\}$ and a set of compensation clauses. For the second case, MINIMAXSAT replaces it by $(\square, m)$ with $m = \min\{w_1, w_2, w_3, w_4, w_5\}$ and a set of compensation clauses. In both cases, the equivalence is preserved. However, other solvers in the literature detect those inconsistent subset of clauses but cannot transform the problem into an equivalent one [34] or simply cannot detect them [17].

Our probing method to derive weighted unit clauses is related to the $2 - RES$ and cycle rule of [17, 6] and to failed literals in [9]. Again, the use of simulated unit propagation allows MINIMAXSAT to identify arbitrarily large resolution trees.

10 Conclusions and Future Work

MINIMAXSAT is an efficient and very robust Max-SAT solver that can deal with hard and soft clauses as well as pseudo-boolean functions. It incorporates the best available techniques for each type of problems, so its performance is similar to the best specialized solver. Besides the development of MINIMAXSAT combining, for the first time, known techniques from different fields, the main original contribution of this paper is a novel lower bounding technique based on resolution. MINIMAXSAT lower bounding subsumes in a very clean and elegant way most of the approaches that have been proposed in the last years.

In this paper we use the information provided by the propagation queue (i) to determine a subset of inconsistent clauses and (ii) to determine a simple ordering in which resolution can be applied to increase the lower bound and generate an equivalent formula. However, this is not necessarily the best ordering to do so. It is easy to see that different orderings may generate resolvents and compensation clauses of different arities. If one selects the ordering that generates the smallest resolvents and compensation clauses the resulting formula may be presumably simpler. Future work concerns the study of such orderings, the development of VSIDS-like heuristics for soft clauses, backjumping techniques for soft conflicts.

References

1. Heras, F., Larrosa, J., Oliveras, A.: MiniMaxSat: A new weighted Max-SAT solver. In: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, SAT-07. (2007)
2. Papadimitriou, C.: Computational Complexity. Addison-Wesley, USA (1994)
3. Karloff, H.J., Zwick, U.: A 7/8-approximation algorithm for MAX3SAT? In: Proceedings of Annual IEEE Symposium on Foundations of Computer Science, FOCS-1997. (1997) 406–415
4. Shen, H., Zhang, H.: Study of lower bounds for Max-2-SAT. In: Proceedings of the National Conference on Artificial Intelligence, AAAI-04. (2004)
5. Larrosa, J., Heras, F.: Resolution in Max-SAT and its relation to local consistency for weighted CSPs. In: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-05. (2005)
6. Larrosa, J., Heras, F., de Givry, S.: A logical approach to efficient Max-SAT solving. In: Artificial Intelligence, To appear. (2007)
7. Xing, Z., Zhang, W.: Maxsolver: An efficient exact algorithm for (weighted) maximum satisfiability. Artificial Intelligence **164** (2005) 47–80
8. Li, C., Manyà, F., Planes, J.: Exploiting unit propagation to compute lower bounds in branch and bound Max-SAT solvers. In: Proceedings of the International Conference on Principles and Practice of Constraint Programming, CP-05. (2005) 403–414
9. Li, C.M., Manyà, F., Planes, J.: Detecting disjoint inconsistent subformulas for computing lower bounds for Max-SAT. In: Proceedings of the National Conference on Artificial Intelligence, AAAI-06. (2006)
10. Cha, B., Iwama, K., Kambayashi, Y., Miyazaki, S.: Local search algorithms for partial MAXSAT. In: Proceedings of the National Conference on Artificial Intelligence, AAAI-97. (1997) 263–268
11. Alsinet, T., Manyà, F., Planes, J.: Improved exact solvers for weighted Max-SAT. In: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, SAT-05. (2005) 371–377
12. Argelich, J., Manya, F.: Partial Max-SAT solvers with clause learning. In: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, SAT-07. (2007)
13. Fu, Z., Malik, S.: On solving the partial MAX-SAT problem. In: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, SAT-06. (2006) 252–265
14. Le Berre, D.: (The SAT4j project for Max-SAT)
15. Eén, N., Sörensson, N.: Translating pseudo-boolean constraints into SAT. Journal on Satisfiability, Boolean Modeling and Computation **2** (2006) 1–26
16. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, SAT-03. (2003) 502–518
17. Heras, F., Larrosa, J.: New inference rules for efficient Max-SAT solving. In: AAAI. (2006)
18. Davis, M., Logemann, G., Loveland, G.: A machine program for theorem proving. Communications of the ACM **5** (1962) 394–397
19. Silva, J.P.M., Sakallah, K.A.: Grasp - a new search algorithm for satisfiability. In: Proceedings of International Conference on Computer-Aided Design, ICCAD-1996. (1996) 220–227
20. Moskewicz, M.W., Madigan, C.F., Zhao, Y., Zhang, L., Malik, S.: Chaff: Engineering an efficient SAT solver. In: Proceedings of the Design Automation Conference, DAC-01. (2001) 530–535

21. Zhang, L., Madigan, C.F., Moskewicz, M.W., Malik, S.: Efficient conflict driven learning in boolean satisfiability solver. In: Proceedings of International Conference on Computer-Aided Design, ICCAD-2001. (2001) 279–285
22. Tompkins, D.A.D., Hoos, H.H.: UBCSAT: An implementation and experimentation environment for sls algorithms for SAT & MAX-SAT. In: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, SAT-04. (2004)
23. Lynce, I., Silva, J.P.M.: Probing-based preprocessing techniques for propositional satisfiability. In: Proceedings of International Conference on Tools with Artificial Intelligence, ICTAI-2003. (2003)
24. Le Berre, D.: Exploiting the real power of unit propagation lookahead. In: In LICS Workshop on Theory and Applications of Satisfiability Testing. (2001)
25. Sheini, H.M., Sakallah, K.A.: Pueblo: A hybrid pseudo-boolean SAT solver. Journal on Satisfiability, Boolean Modeling and Computation **2** (2006) 165–189
26. Fahle, T.: Simple and fast: Improving a branch-and-bound algorithm for maximum clique. In: Proceedings of European Symposium on Algorithms, ESA-2002. (2002) 485–498
27. Sandholm, T.: An algorithm for optimal winner determination in combinatorial auctions. In: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-99. (1999) 542–547
28. K. Leyton-Brown, M.P., Shoham, Y.: Towards a universal test suite for combinatorial auction algorithms. ACM Conference on Electronic Commerce (2000) 66–76
29. de Givry, S., Larrosa, J., Meseguer, P., Schiex, T.: Solving Max-SAT as weighted CSP. In: Proceedings of the International Conference on Principles and Practice of Constraint Programming, CP-03, Kinsale, Ireland, LNCS 2833. Springer Verlag (2003) 363–376
30. Larrosa, J., Schiex, T.: In the quest of the best form of local consistency for weighted CSP. In: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-03, Acapulco, Mexico (2003)
31. de Givry, S., Heras, F., Larrosa, J., Zytnicki, M.: Existential arc consistency: getting closer to full arc consistency in weighted CSPs. In: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-05, Edinburgh, U.K. (2005)
32. Cooper, M., Cussat-Blanc, S., de Roquemaurel, M., Régnier, P.: Soft arc consistency applied to optimal planning. In: Proceedings of the International Conference on Principles and Practice of Constraint Programming, CP-06. (2006) 680–684
33. Walsh, T.: SAT v CSP. In: Proceedings of the International Conference on Principles and Practice of Constraint Programming, CP-00. (2000) 441–456
34. Li, C., Manyà, F., Planes, J.: New inference rules for Max-SAT. In: Journal of Artificial Intelligence Research, To appear. (2007)
35. Larrosa, J.: Node and arc consistency in weighted CSP. In: Proceedings of the National Conference on Artificial Intelligence, AAAI-02. (2002) 48–53
36. Brglez, F., Li, X., Stallman, M.: The role of a skeptic agent in testing and benchmarking of sat algorithms. In: In Fifth Int. Symp. on the Theory and Applications of Satis ability Testing, Cincinnati, Ohio, USA (2002)
37. Argelich, J., Manyà, F.: Exact Max-SAT solvers for over-constrained problems. Journal of Heuristics **12** (2006) 375–392
38. Lynce, I., Silva, J.P.M.: Efficient data structures for backtrack search SAT solvers. Ann. Math. Artif. Intell. **43** (2005) 137–152
39. Argelich, J., Manyà, F.: Learning hard constraints in Max-SAT. In: Proceedings of the Annual ERCIM Workshop on Constraint Solving and Constraint Programming, CSCLP-06. (2006) 1–12
40. Zivan, R., Meisels, A.: Conflict directed backjumping for maxcsps. In: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-07. (2007)

41. Manquinho, V.M., Silva, J.P.M.: Satisfiability-based algorithms for boolean optimization. Ann. Math. Artif. Intell. **40** (2004) 353–372
42. Nieuwenhuis, R., Oliveras, A.: On SAT Modulo Theories and Optimization Problems. In: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, SAT-06. (2006) 156–169