

Solving satisfiability problems with preferences

Emanuele Di Rosa · Enrico Giunchiglia ·
Marco Maratea

Published online: 3 July 2010
© Springer Science+Business Media, LLC 2010

Abstract Propositional satisfiability (SAT) is a success story in Computer Science and Artificial Intelligence: SAT solvers are currently used to solve problems in many different application domains, including planning and formal verification. The main reason for this success is that modern SAT solvers can successfully deal with problems having millions of variables. All these solvers are based on the Davis–Logemann–Loveland procedure (DLL). In its original version, DLL is a decision procedure, but it can be very easily modified in order to return one or all assignments satisfying the input set of clauses, assuming at least one exists. However, in many cases it is not enough to compute assignments satisfying all the input clauses: Indeed, the returned assignments have also to be “optimal” in some sense, e.g., they have to satisfy as many other constraints—expressed as preferences—as possible. In this paper we start with qualitative preferences on literals, defined as a partially ordered set (poset) of literals. Such a poset induces a poset on total assignments and leads to the definition of optimal model for a formula ψ as a minimal element of the poset on the models of ψ . We show (i) how DLL can be extended in order to return *one* or *all* optimal models of ψ (once converted in clauses and assuming ψ is satisfiable), and (ii) how the same procedures can be used to compute optimal models wrt a qualitative preference on formulas and/or wrt a quantitative preference on literals or formulas. We implemented our ideas and we tested the resulting system on a variety of very challenging structured benchmarks. The results indicate that our

This work extends the results presented in [19, 24].

E. Di Rosa (✉) · E. Giunchiglia · M. Maratea
DIST, University of Genoa, Viale Causa, 13, 16145 Genoa, Italy
e-mail: emanuele@dist.unige.it

E. Giunchiglia
e-mail: enrico@dist.unige.it

M. Maratea
e-mail: marco@dist.unige.it

implementation has comparable performances with other state-of-the-art systems, tailored for the specific problems we consider.

Keywords Satisfiability · Preferences

1 Introduction

Propositional satisfiability (SAT) is a success story in Computer Science and Artificial Intelligence: SAT solvers are currently used to solve problems in many different application domains, including planning [30], formal verification [6], and many others, such as RNA folding, hand-writing recognition, graph isomorphism and sudoku problems. The main reason for this success is that modern SAT solvers can successfully deal with problems having millions of variables [32, 33].¹ All these solvers are based on the Davis–Logemann–Loveland procedure (DLL) [17]. The original version of DLL is a decision procedure—given a finite set of clauses φ , DLL returns whether φ is satisfiable or not—but DLL can be easily modified in order to return one or all assignments satisfying φ , assuming at least one exists. However, in many cases, it is not enough to compute one or more satisfying assignments: Indeed, the returned assignments have also to be “optimal” in some sense, e.g., they have to satisfy as many other constraints—expressed as preferences—as possible. For example, in standard MIN-ONE (resp. MIN-ONE_⊆), given a satisfiable instance, the goal is to find a satisfying assignment in which the set of variables assigned to true is of minimal cardinality (resp. subset-minimal). In standard MAX-SAT (resp. MAX-SAT_⊆), the goal is to find an assignment satisfying as many clauses as possible, i.e., such that the set of satisfied clauses is of maximal cardinality (resp. subset-maximal). In the partial version of MIN-ONE/MIN-ONE_⊆ (resp. MAX-SAT/MAX-SAT_⊆) the optimization is performed on a subset of the variables (resp. clauses) of the instance.

In this paper we start considering the simple model in which preferences are expressed as a partially ordered set (poset) of literals as in, e.g., [50, 54]. Such a poset induces a poset on total assignments and leads to the definition of optimal model for a formula ψ as a minimal element of the poset on the models of ψ . Given a qualitative preference on literals and a finite set of clauses φ , we show how DLL can be easily modified in order to return an optimal model of φ , assuming φ is satisfiable. The simple idea for computing *one* optimal model, is to force DLL branching heuristic in order to follow the partial order on the literals. The idea of computing “optimal” (according to some given definition) models by modifying the heuristic in order to follow the expressed preferences on literals has been already proposed in [12] for SAT and in [8] for acyclic CP-nets [9]. There are however some important differences in the underlying formalism used in [8, 12] for expressing preferences—and thus on the procedures based on these formalisms—wrt ours:²

1. In the language: Both [12] and [8] allow for expressing preferences on literals, but in these approaches it is not possible to rank the preferences according to

¹<http://www.satcompetition.org/>

²In the case of [9], we consider the simple case in which variables are Boolean and preferences are not conditional.

- a partial order. For instance, these approaches allow for directly expressing a preference in which literals l_1 and l_2 are assigned to true, but do not allow for expressing that having l_1 assigned to true is preferred to having l_2 assigned to true. Further, in [8, 12] the set of preferences has to be consistent, while we do not make this assumption.
2. In the semantics: Even considering the case in which preferences are expressed as a consistent set S of literals, the order on models induced by S in [8, 12] is different from ours. For example, given a language with two variables x_1 and x_2 , and assuming that our only preference is to have x_1 assigned to true, given an assignment μ_1 (resp. μ_2) assigning both x_1 and x_2 to true (resp. false), μ_1 is preferred to μ_2 according to our semantics (see Section 2), while this is not the case for the semantics in [8, 12] (see Definition 7 in [12] and Section 2.3 in [8]).

We then extend our procedure in order to find more, and possibly all, optimal models. As in [12], the idea is to add to the input formula a constraint imposing that the new models have not to follow μ in the partial order: Thus, assuming we have already generated a non empty set of optimal models and we are interested in more, differently from the procedures in [8, 47] for CP-nets, our algorithm for generating a new optimal model μ never requires a dominance test to see if there exists another model which is preferred to μ . Finally, we show how the same procedures can be extended to compute optimal models wrt a qualitative preference on formulas and a quantitative preference on literals or formulas. Indeed, this is a trivial consequence of the fact that all these concepts (qualitative/quantitative preference on literals/formulas and also their mixing) can be reduced to the basic framework of qualitative preference on literals.

We implemented our ideas in MINISAT [21], and we called *nOPTSAT* the resulting system.³ In order to comparatively test our system, we focused our experimental analysis on MAX-SAT/MAX-SAT_⊆ and MIN-ONE/MIN-ONE_⊆ problems. Indeed, this is a very challenging—if not the most challenging—test bench for our implementation:

1. for MIN-ONE/MAX-SAT, a wide variety of recently developed, customized systems are available, e.g., those in the last Pseudo-Boolean (PB) and Max-SAT Evaluations;⁴ and
2. in these problems, the number of preferences is very high (equal to the number of clauses in standard MAX-SAT/MAX-SAT_⊆ and to the number of variables in standard MIN-ONE/MIN-ONE_⊆): As we have already shown in [25], in the context of planning as satisfiability, the more preferences we have, the more the performances of our system are negatively affected.

Despite the above, our analysis shows that *nOPTSAT* has comparable performances with respect to other state-of-the-art systems on MIN-ONE/MAX-SAT problems. In the case of MIN-ONE_⊆/MAX-SAT_⊆ problems, we consider the only other implementation available for standard MAX-SAT_⊆, and here again we show that our system compares

³Available at <http://www.star.dist.unige.it/~emanuele/nOPTSAT/>.

⁴See <http://www.cril.univ-artois.fr/PB09/> and <http://www.maxsat.udl.cat/09/>, respectively.

well. However, we remark that our results goes far beyond the $\text{MIN-ONE}/\text{MIN-ONE}_{\subseteq}$ and $\text{MAX-SAT}/\text{MAX-SAT}_{\subseteq}$ cases, which

1. are the simplest cases of preferences on literals and on formulas respectively, since they have an empty partial order on preferences; and
2. are the most difficult problems for our system, given the very high number of preferences they have.

Indeed, we allow for preferences to be partially ordered and in many applications the number of preferences is relatively low, as, e.g., in planning with soft goals [23, 25].

Summing up, given a set of clauses φ , the main contributions of the paper are:

1. We show how it is possible to easily extend DLL in order to compute one optimal model of φ wrt a qualitative preference on literals: We allow for inconsistent set of preferences and for a partial order on the preferred literals.
2. We extend the procedure in order to compute and return more than one, and possibly all, optimal models of φ wrt a given qualitative preference: Our procedure does not require any dominance test.
3. We show how qualitative preferences on formulas and quantitative preferences on literals or formulas can be reduced to the basic framework of qualitative preferences on literals: This allows to use our procedures also in these extended settings, and also for solving problems with mixed qualitative and quantitative preferences.
4. We implemented these ideas on top of MINISAT and we comparatively tested our system on a variety of structured $\text{MIN-ONE}/\text{MIN-ONE}_{\subseteq}$ and $\text{MAX-SAT}/\text{MAX-SAT}_{\subseteq}$ problems against various state-of-the-art systems, tailored for such problems: Despite the generality of our procedure, the results indicate that our system has comparable performances wrt the others.

The paper is structured as follows. In Section 2, we present the formalism we are using for expressing qualitative preferences on literals, and we show how these preferences induce a preference on the set of total assignments and thus also on the models of any given formula. In Section 3, we first present OPT-DLL , i.e., DLL modified in order to compute an optimal model of a finite set of clauses wrt a qualitative preference on literals, and then NOPT-DLL , i.e., OPT-DLL extended in order to compute all optimal models. In Section 4, we show how it is possible to reduce quantitative/qualitative preference on formulas to qualitative preference on literals. In the same section, we give examples showing how it is possible to represent and solve problems in which qualitative and quantitative preferences are mixed. Section 5 is devoted to the implementation details and the comparative experimental analysis of the ideas presented. The paper ends in Section 6 with the conclusions.

2 Satisfiability and qualitative preferences

Consider a finite set P of variables, called *signature*. A *literal* is a variable x or its negation $\neg x$. A *formula* or *constraint* is either a variable or a combination of formulas using the n -ary connectives \wedge and \vee for conjunction and disjunction, respectively ($n \geq 0$); the n -ary connective \equiv for equivalence ($n \geq 2$); and the unary connective \neg for negation.

For example, given the five variables $AtWork_0, AtWork_1, Car_0, Bus_0, Bike_0$, in which the subscript models the time step, the following formula

$$AtWork_1 \equiv \neg AtWork_0 \equiv (Car_0 \vee Bus_0 \vee Bike_0), \tag{1}$$

models the fact that if we perform the action of taking the car (Car_0) or the bus (Bus_0) or the bike ($Bike_0$) then we move from home ($\neg AtWork_0$) to work ($AtWork_1$), or viceversa, while the formula

$$\neg AtWork_0 \wedge AtWork_1 \tag{2}$$

models the fact that at time 0 we are at home (or not at work) and at time 1 we are at work (variables have the obvious meaning). The two equations do not rule out the possibility of executing two or even three actions at the same time (e.g., to have a model in which both Car_0 and Bus_0 are true). Indeed, these models can be easily eliminated by adding additional constraints.

In the following, given a literal l, \bar{l} denotes $\neg l$ if l is a variable, and the variable in l otherwise. If S is a set of literals,

$$\bar{S} = \{\bar{l} : l \in S\}.$$

An *assignment* is a consistent set of literals.

Consider an assignment μ .

If $l \in \mu$, we say that both l and \bar{l} are *assigned* by μ , to true and false, respectively. μ is *total* if each literal is assigned by μ . If μ is total, we say that μ *satisfies*

- a variable x if $x \in \mu$;
- a formula ψ if μ satisfies ψ according to the truth tables of the propositional connectives;
- a set of formulas if μ satisfies each formula in the set.

A *model* of a set of formulas is an assignment satisfying the set.

For instance, (1) has 16 models, while considering (1) and (2) the models become 7. In the following, we represent a total assignment as the set of variables assigned to true, and we write $\mu \models \psi$ to indicate that μ is a model of ψ . For instance, $\{Car_0, AtWork_1\}$ represents the total assignment in which the only variables assigned to true are Car_0 and $AtWork_1$.

As we already said in the introduction, in many applications not all the models of a set of formulas are considered to be equally good. For example, in (1) we may prefer the models in which as few actions as possible are performed, i.e., we may prefer the models having a maximal intersection with $\{\neg Car_0, \neg Bus_0, \neg Bike_0\}$: This intuitively corresponds to considering the two models in which we do not move from/to home as optimal. Further, in case it is not possible to have a model with $\neg Car_0, \neg Bus_0, \neg Bike_0$ (as with (1) and (2)), we may want to express the additional information that we have better to give up on $\neg Car_0$ than the other two, i.e., that of the three preferences, $\neg Car_0$ is the least important: With such additional information, we expect $\{Car_0, AtWork_1\}$ to be the only optimal model for (1) and (2).

Such additional information about which model is preferred to the other ones can be expressed via qualitative preferences on literals. A *qualitative preference on literals*

is a (strict) partially ordered set of literals.⁵ We recall that a (strict) partially ordered set (or poset) is a pair $\langle S, < \rangle$ whose first element is a set and $<$ is a (strict) partial order on S , i.e., a binary relation satisfying the following two properties:

1. *Irreflexivity*: For each $a \in S, a \not< a$.
2. *Transitivity*: For each $a, b, c \in S, a < b$ and $b < c$ implies $a < c$.⁶

Given a poset $\langle S, < \rangle$,

- if for each two distinct $a, b \in S, a < b$ or $b < a$ then $<$ is said to be a *total order*;
- an element $a \in S$ is said to be *minimal (in $\langle S, < \rangle$)* if there is no $b \in S$ with $b < a$:
It is clear that if S is finite then the poset has at least one minimal element; and
- if $S' \subseteq S$ then $\langle S', <' \rangle$ is also a poset, where $<'$ is $<$ restricted to the literals in S' .

It is common to represent a poset $\langle S, < \rangle$ as the direct acyclic graph (DAG) whose vertexes are the elements in S , and with an arc from a to b if and only if $a < b$ and there is no c with $a < c < b$.

In a qualitative preference on literals $\langle S, < \rangle$, S is the *set of preferences* and intuitively represents the set of literals that we would like to have satisfied, while $<$ models the relative importance of the preferences. Notice that the set S can be inconsistent: If S contains both a variable x and its negation $\neg x$ —assuming compatibility with the constraints and that neither $x < \neg x$ nor $\neg x < x$ —we expect to have an optimal model with x true and another optimal model with x false. For example,

1. the qualitative preference

$$\langle \{\neg Car_0, \neg Bus_0, \neg Bike_0, Bike_0\}, \emptyset \rangle$$

models our preference to get (assuming compatibility with the underlying constraints) two optimal models μ_1 and μ_2 , satisfying $\{\neg Car_0, \neg Bus_0, \neg Bike_0\}$ and $\{\neg Car_0, \neg Bus_0, Bike_0\}$, respectively;

2. while

$$\langle \{\neg Car_0, \neg Bus_0, \neg Bike_0\}, \{\neg Bus_0 < \neg Car_0, \neg Bike_0 < \neg Car_0\} \rangle \tag{3}$$

models our preference for not moving from/to home, and in which, of the three preferences, $\neg Car_0$ is the least important.

Consider a qualitative preference on literals $\langle S, < \rangle$. The partial order on S can be extended to the set of total assignments as follows [54]: Given two total assignments μ and μ' , μ is preferred to μ' or μ dominates μ' ($\mu < \mu'$) if and only if

1. μ satisfies at least one preference which is not satisfied by μ' , i.e., there exists a literal $l \in S$ with $l \in \mu$ and $\bar{l} \in \mu'$; and
2. the preferences satisfied by μ' and not by μ are less preferred to those satisfied by μ and not by μ' , i.e., for each literal $l \in S \cap (\mu' \setminus \mu)$, there exists a literal $l' \in S \cap (\mu \setminus \mu')$ such that $l' < l$.

⁵The given definition of qualitative preference on literals generalizes the ones given in [24].

⁶If $<$ is irreflexive and transitive then it is also antisymmetric, i.e., if $a < b$ then $b \not< a$.

From the definition, it is clear that, for any two total assignments μ and μ' :

1. If $S \cap \mu = S \cap \mu'$ then $\mu \not< \mu'$: In particular, if the set S of preferences is empty, every model is optimal.
2. If $S \cap \mu' \subset S \cap \mu$ then $\mu < \mu'$: Every optimal model has a maximal intersection with S . In the case $<$ is empty, every model with a maximal intersection with S is optimal.

$\langle S, < \rangle$ induces a partial order on the set of total assignments, as stated by the following theorem, similar to Theorem 7 in [55].

Theorem 1 *Let $\langle S, < \rangle$ be a qualitative preference on literals. The relation $<$ extended to the set of total assignments is a partial order.*

Proof For a literal l and a total assignment μ , define

$$\text{dom}(l, \mu) = \{l' : l' \in S \cap \mu, l' < l\}.$$

We have to show that the relation $<$ on the set of total assignments is irreflexive and transitive. Let μ, μ' and μ'' be three total assignments.

1. *Irreflexivity*: Clearly $\mu \not< \mu$, because of the first condition in the definition of $<$.
2. *Transitivity*: $\mu < \mu'$ and $\mu' < \mu''$ implies $\mu < \mu''$. We have to prove that for each literal $l' \in S \cap (\mu'' \setminus \mu)$, there exists a literal $l \in S \cap (\mu \setminus \mu')$ such that $l < l'$. Considering the set

$$S' = S \cap (\mu'' \setminus \mu),$$

it is enough to show that for each literal l' which is minimal in $\langle S', < \rangle$, there exists a literal $l \in S \cap (\mu \setminus \mu')$ such that $l < l'$.

Let l' be a minimal element in $\langle S', < \rangle$.

There are two cases:

- (a) $l' \in \mu'$ and $\text{dom}(l', \mu'') \subseteq \mu'$: Since $l' \in \mu'$ and $\mu < \mu'$, let l be a literal in $S \cap \mu$ such that $l < l'$, and $l \notin \mu'$. Because of the second initial assumption, $l \notin \mu''$ and thus the thesis.
- (b) $l' \notin \mu'$ or $\text{dom}(l', \mu'') \not\subseteq \mu'$: The set $(\{l'\} \cup \text{dom}(l', \mu'')) \setminus \mu'$ is not empty: Let l'' be a minimal element of this set according to $<$. Notice that either $l'' = l'$ or $l'' < l'$. In both cases, $l'' \in \mu''$, $l'' \notin \mu'$, and

$$\text{dom}(l'', \mu'') \subseteq \mu'. \tag{4}$$

Since $\mu' < \mu''$, there exists a literal $l''' \in \mu' \setminus \mu''$ with $l''' < l''$ and thus $l''' < l'$. If $l''' \in \mu$, l''' is the literal l we are looking for and thesis follows. If $l''' \notin \mu$, since $\mu < \mu'$ there exists a literal $l \in S \cap (\mu \setminus \mu')$ such that $l < l'''$ and thus $l < l'$. This literal is not in μ' and thus, because of (4), it is not in μ'' as well. □

Since the set M of models of an arbitrary set of constraints is a subset of the set of total assignments, then also $\langle M, < \rangle$ is a partially ordered set. For example, given the constraints (1) and (2), and the qualitative preference (3), the partial order on models is represented by the DAG in Fig. 1.

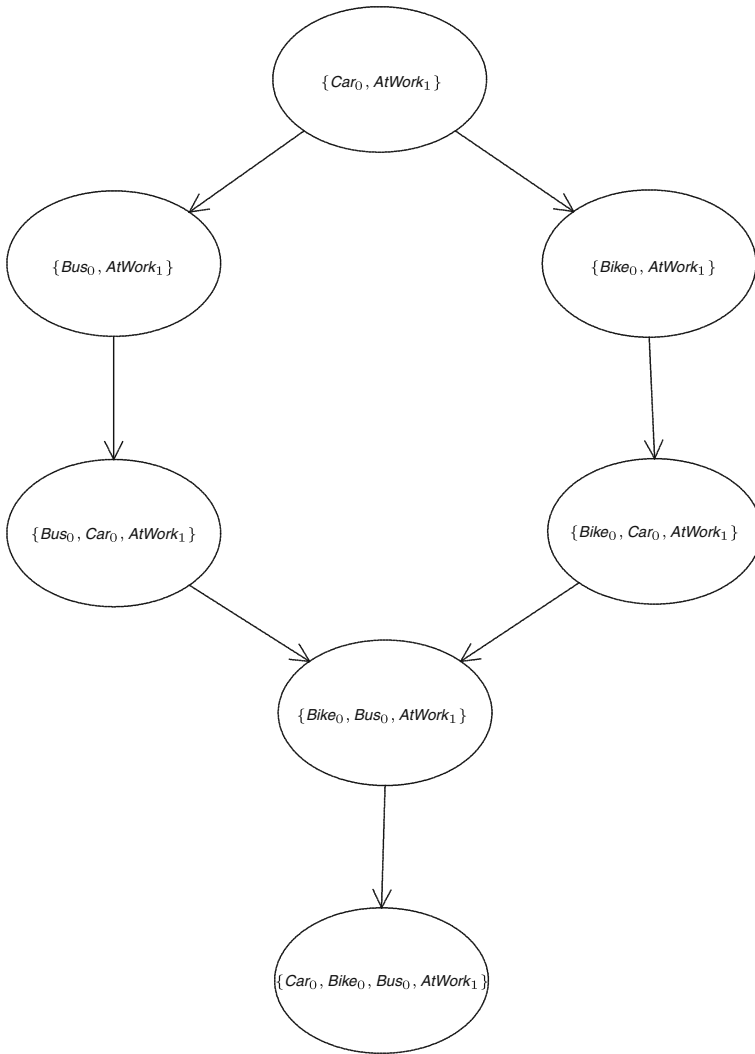


Fig. 1 DAG representation of the partial order on the models of (1) and (2), wrt (3)

However, there can be partial orders on models (or on total assignments) which are not extensions of partial order on literals. For example, assuming we have only two variables x_0, x_1 , the total order on the set of total assignments $\{\bar{x}_1, \bar{x}_0\} < \{x_1, \bar{x}_0\} < \{x_1, x_0\} < \{\bar{x}_1, x_0\}$ can not be obtained as the result of the extension of a partial order on a subset of the literals in $\{x_0, \bar{x}_0, x_1, \bar{x}_1\}$. In Section 4 we extend the formalism in order to express qualitative preferences on formulas. With preferences on formulas, given that a total assignment μ can be represented as the conjunction C_μ of the literals in μ , it is possible to capture any partial order on the set of total assignments: Trivially, given any two total assignments μ_1 and μ_2 for which we want $\mu_1 < \mu_2$, imposing the preference on formulas $C_{\mu_1} < C_{\mu_2}$ leads to the desired order on μ_1 and μ_2 .

Given that $\mu < \mu'$ encodes the fact that μ is preferred to μ' , it is natural to define a model as *optimal* if it is a minimal element of the partially ordered set of models. A finite set of constraints and a qualitative preference defines an *optimization problem*, consisting in determining an optimal model of the set of constraints. As it is clear from Fig. 1, the assignment $\{Car_0, AtWork_1\}$ is the only solution for the optimization problem consisting of the constraints (1), (2), and the qualitative preference (3).

3 Solving optimization problems with DLL

Given a finite set of constraints and a qualitative preference on literals, we show how it is possible to compute an optimal model by a simple modification of the Davis–Logemann–Loveland procedure (DLL) [17], and then how to extend DLL in order to compute all optimal models.

DLL is at the basis of current state-of-the-art procedures for checking the satisfiability of a finite set of constraints, which extend it in many ways, including conflict driven backjumping and clause learning, see, e.g., [41, 45]. However, DLL does not directly handle arbitrary formulas, but finite sets of clauses (the clauses to be interpreted in conjunction among them), where a *clause* is a finite set of literals (the literals to be interpreted in disjunction among them). This is not a limitation because of well known and efficient clause form transformation procedures (see, e.g., [27, 46, 53]). Thus, from here on—with the exception of the text in the formal statements—we assume we are able to use DLL and our procedures also on formulas, implicitly assuming that formulas are converted into a set of clauses beforehand.

3.1 Computing an optimal model with DLL

Our procedure, that we call OPT-DLL, is a simple modification of DLL in order to take into account the given qualitative preference on literals $\langle S, < \rangle$.

The pseudo-code of OPT-DLL is represented in Fig. 2, where:

- φ and $\langle S, < \rangle$ are global variables storing the input set of clauses and the qualitative preference on literals, respectively;
- μ is the current assignment, initially empty;

Fig. 2 The algorithm of OPT-DLL for computing one optimal solution

```

 $\varphi :=$  (a finite set of clauses in the signature  $P$ );
 $\langle S, < \rangle :=$  (a qualitative preference on literals);

function OPT-DLL()      return OPT-DLL-REC( $\emptyset$ )

function OPT-DLL-REC( $\mu$ )
1  if ( $\emptyset \in \varphi_\mu$ ) return FALSE;
2  if ( $\mu$  is total) return  $\mu$ ;
3  if ( $\{l\} \in \varphi_\mu$ ) return OPT-DLL-REC( $\mu \cup \{l\}$ );
4   $l :=$  ChooseLiteral( $\mu$ );
5   $v :=$  OPT-DLL-REC( $\mu \cup \{l\}$ );
6  if ( $v \neq$  FALSE) return  $v$ ;
7  return OPT-DLL-REC( $\mu \cup \{\bar{l}\}$ ).
    
```

- φ_μ is the result of simplifying the input set of clauses wrt the assignment μ , i.e., φ_μ is the set of clauses obtained from φ by (i) deleting the clauses $C \in \varphi$ such that $C \cap \mu \neq \emptyset$, and (ii) substituting the other clauses $C \in \varphi$ with $C \setminus \bar{\mu}$;
- *ChooseLiteral*(μ) returns a literal l unassigned by μ and such that
 - either $l \in S$ and each literal l' with $l' \prec l$ is assigned by μ ;
 - or, if each literal in S is assigned by μ , l is an arbitrary literal in φ_μ , selected by any given heuristic.

It is easy to see that if there are no preferences (i.e., if the set S is empty) OPT-DLL is the standard DLL. On the other hand, if the set of preferences is not empty, the search tree is explored in a way to ensure that the returned model (assuming the input formula is satisfiable) is optimal. For instance, assuming we have the qualitative preference

$$\{\{\neg Bus_0, \neg Bike_0\}, \{\neg Bus_0 \prec \neg Bike_0\}\}$$

(modeling the fact that we prefer (i) to not take the bus; (ii) to not take the bike; and (iii) to not go by bus more than to not go by bike) OPT-DLL looks for a model extending

1. $\{\neg Bus_0, \neg Bike_0\}$; if no such model exists, OPT-DLL looks for a model extending
2. $\{\neg Bus_0, Bike_0\}$; if no such model exists, OPT-DLL looks for a model extending
3. $\{Bus_0, \neg Bike_0\}$; if no such model exists, OPT-DLL looks for a model extending
4. $\{Bus_0, Bike_0\}$; if no such model exists, OPT-DLL returns false.

In words, OPT-DLL first looks for a model where both actions of going by bus and bike are false; then one in which we use the bike and not the bus; then one in which we use the bus and not the bike; and only finally for models where we use both the bus and the bike. If ψ is the conjunction of (1) and (2), the model of ψ returned with such exploration is $\{Car_0, AtWork_1\}$, while it is $\{Bike_0, AtWork_1\}$ if we also consider the constraint $\neg Car_0$.

OPT-DLL returns an optimal assignment if the input formula is satisfiable, and FALSE otherwise, as stated by the following theorem, easy consequence of Theorem 4, which is stated and proved in the next subsection.

Theorem 2 *Let $\varphi, \langle S, \prec \rangle$ and OPT-DLL as in Fig. 2. OPT-DLL() returns an optimal model of φ wrt $\langle S, \prec \rangle$ if φ is satisfiable, and FALSE otherwise.*

Consider a satisfiable and finite set of clauses φ .

Depending on the literal returned by *ChooseLiteral*(μ), different optimal models are computed and returned by OPT-DLL. For instance, given the qualitative preference

$$\{\{\neg Car_0, \neg Bus_0, \neg Bike_0\}, \emptyset\}, \tag{5}$$

each of the three optimal models of (1) and (2) wrt (5) can be returned by OPT-DLL, which one depending on the specific implementation of *ChooseLiteral*: If the first two literals assigned at line 4 in OPT-DLL are in

1. $\{\neg Car_0, \neg Bus_0\}$, then the returned optimal model is $\{Bike_0, AtWork_1\}$;
2. $\{\neg Car_0, \neg Bike_0\}$, then the returned optimal model is $\{Bus_0, AtWork_1\}$;
3. $\{\neg Bus_0, \neg Bike_0\}$, then the returned optimal model is $\{Car_0, AtWork_1\}$.

In this example, each optimal model is a possible output of OPT-DLL. There are however optimization problems in which some of the optimal models cannot be returned by OPT-DLL. Consider, for instance, an optimization problem with qualitative preference

$$\langle \{x_0, x_1, x_2, x_3\}, \{x_0 < x_1, x_2 < x_3\} \rangle$$

and constraints imposing that if one variable in $\{x_0, x_2\}$ is true then the other three variables have to be false,⁷ e.g.,

$$(\bar{x}_0 \vee \bar{x}_1) \wedge (\bar{x}_0 \vee \bar{x}_2) \wedge (\bar{x}_0 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee \bar{x}_1) \wedge (\bar{x}_2 \vee \bar{x}_3).$$

The above formula has three models, i.e., $\{x_0\}$, $\{x_2\}$ and $\{x_1, x_3\}$, and all of them are optimal. However, OPT-DLL can return only two of them, namely either $\{x_0\}$ or $\{x_2\}$, which one depending on the implementation of *ChooseLiteral*.

The fact that there can be some optimal model that cannot be returned by OPT-DLL is not a limitation if the goal is to compute *one* optimal model. If the goal is to compute more than one or *all* optimal models, we will see in the next subsection how to generalize OPT-DLL.

3.2 Computing all optimal models with DLL

Consider a formula ψ and a qualitative preference $\langle S, < \rangle$.

The problem of computing all optimal models of ψ wrt $\langle S, < \rangle$ can be solved by

1. determining and printing an optimal model μ of ψ by imposing an ordering on the splitting heuristic, as in the previous subsection;
2. adding to the input formula a new formula which prunes the assignments which are dominated by μ ; and
3. returning FALSE in order to continue the search for other optimal models.

The idea to compute all optimal models by adding constraints pruning the models dominated by the already computed optimal models, has been already proposed in [12]. Other works which exploit further techniques to eliminate previously computed solutions in SAT include, e.g., [29, 43, 49] in the context of symbolic model checking [42]. However, the framework used in [12] assumes a consistent set of preferences, does not allow for ordering the preferences and ultimately has a different semantics.

For the above procedure, given an assignment μ , we have to define a formula whose models are the assignments dominated by μ (wrt $\langle S, < \rangle$). Such a formula is

$$\left(\bigvee_{l \in S \cap \mu} \bar{l} \right) \wedge \left(\bigwedge_{l \in S \cap \bar{\mu}} \left(\bar{l} \vee \bigvee_{l' \in S \cap \mu, l' < l} \bar{l}' \right) \right). \tag{6}$$

Theorem 3 *Let $\langle S, < \rangle$ be a qualitative preference on literals. A total assignment μ dominates a total assignment μ' wrt $\langle S, < \rangle$ if and only if μ' satisfies (6).*

⁷For instance $\{x_0, x_1, x_2, x_3\}$ can model the fact that we like to have wine (x_0), beer (x_1), fish (x_2) and pizza (x_3); we like wine more than beer, and fish more than pizza. Due to budget limitations, if we buy fish or wine we cannot afford anything else.

Proof According to the definition, given two total assignments μ and μ' , $\mu < \mu'$ iff

1. there exists a literal $l \in S$ with $l \in \mu$ and $\bar{l} \in \mu'$; and
2. for each literal $l \in S \cap (\mu' \setminus \mu)$, there exists a literal $l' \in S \cap (\mu \setminus \mu')$ such that $l' < l$.

The first condition corresponds to the formula

$$\bigvee_{l \in S \cap \mu} \bar{l}.$$

The second condition corresponds to the formula

$$\bigwedge_{l \in S \cap \bar{\mu}} \left(\bar{l} \vee \bigvee_{l' \in S \cap \mu, l' < l} \bar{l}' \right).$$

The conjunction of the above two formulas is (6). □

As examples of the application of Theorem 3, consider a total assignment μ :

1. If $S \cap \mu = \emptyset$ then the formula (6) is equivalent to the empty disjunction, i.e., FALSE: Indeed, if μ does not satisfy any preference, no assignment is dominated by μ ;
2. If $S \subseteq \mu$ then the formula (6) is equivalent to $\bigvee_{l \in S} \bar{l}$: Each assignment which does not satisfy all the preferences is dominated by μ ;
3. If $< = \emptyset$ then, the formula (6) is equivalent to $\bigvee_{l \in S \cap \mu} \bar{l} \wedge \bigwedge_{l \in S \cap \bar{\mu}} \bar{l}$: Each assignment satisfying a strict subset of the set of preferences satisfied by μ , is dominated by μ .

Notice that if μ_1 dominates μ_2 and ψ_1 (resp. ψ_2) is the formula (6) computed for μ_1 (resp. μ_2), then ψ_2 entails ψ_1 , i.e., the models of ψ_2 are a subset of the models ψ_1 : This is a simple consequence of the fact that if $\mu_1 < \mu_2$ then μ_1 dominates a superset of the total assignments dominated by μ_2 .

Thanks to Theorem 3, it is possible to generalize OPT-DLL in Fig. 2 in order to return all the optimal models of a finite set of clauses φ . The resulting procedure is represented in Fig. 3. In the figure, let P be the signature of φ , $Reason(\mu)$

Fig. 3 The algorithm of nOPT-DLL for computing all optimal solutions

$\varphi :=$ (a finite set of clauses in the signature P);
 $\langle S, < \rangle :=$ (a qualitative preference on literals);

function nOPT-DLL() **return** nOPT-DLL(\emptyset)

function nOPT-DLL(μ)
 1 **if** ($\emptyset \in \varphi_\mu$) **return** FALSE;
 2 **if** (μ is total)
 3 Print($\mu \cap (P \cup \bar{P})$);
 4 $\varphi := \varphi \cup Reason(\mu)$;
 5 **return** FALSE;
 6 **if** ($\{l\} \in \varphi_\mu$) **return** nOPT-DLL($\mu \cup \{l\}$);
 7 $l := ChooseLiteral(\mu)$;
 8 $v := nOPT-DLL(\mu \cup \{l\})$;
 9 **if** ($v \neq FALSE$) **return** v ;
 10 **return** nOPT-DLL($\mu \cup \{\bar{l}\}$).

corresponds to the negation of (6), i.e., $Reason(\mu)$ is a finite set of clauses—possibly in a signature P' extending P —such that

1. for each total assignment μ satisfying the negation of (6), there exists one assignment μ' in P' extending μ and satisfying $Reason(\mu)$;
2. for each total assignment μ' in P' satisfying $Reason(\mu)$, the restriction of μ' to P satisfies the negation of (6).

Such a set of clauses can be computed starting from the negation of (6) using the already mentioned clause form transformations [27, 46, 53].

As an example, consider the optimization problem discussed at the end of previous subsection, having

$$\{x_0, x_1, x_2, x_3, \{x_0 < x_1, x_2 < x_3\}\}$$

as qualitative preference, while the constraint

$$(\bar{x}_0 \vee \bar{x}_1) \wedge (\bar{x}_0 \vee \bar{x}_2) \wedge (\bar{x}_0 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee \bar{x}_1) \wedge (\bar{x}_2 \vee \bar{x}_3),$$

imposes that the only models are $\{x_0\}$, $\{x_2\}$ and $\{x_1, x_3\}$. *nOPT-DLL*

1. starts determining and printing the optimal model $\{x_0\}$ or $\{x_2\}$, which one depending on whether $ChooseLiteral(\emptyset)$ returns x_0 or x_2 . Assuming $ChooseLiteral(\emptyset) = x_0$, *nOPT-DLL*

- (a) computes and prints the first optimal model $\mu_0 = \{x_0\}$,
- (b) computes (6) for μ_0 , i.e.,

$$\bar{x}_0 \wedge (\bar{x}_1 \vee \bar{x}_0) \wedge \bar{x}_2 \wedge \bar{x}_3,$$

- (c) adds to the input set of clauses a set of clauses corresponding to the negation of the previous formula, e.g.,

$$x_0 \vee x_2 \vee x_3;$$

2. backtracks setting \bar{x}_0 , and continues the search looking for models extending the partial assignment $\{\bar{x}_0\}$. Assuming $ChooseLiteral(\{\bar{x}_0\}) = x_2$, *nOPT-DLL*

- (a) determines and prints the second optimal model $\mu_1 = \{x_2\}$,
- (b) computes (6) for μ_1 , i.e.,

$$\bar{x}_2 \wedge (\bar{x}_3 \vee \bar{x}_2) \wedge \bar{x}_0 \wedge \bar{x}_1,$$

- (c) adds to the input set of clauses a set of clauses corresponding to the negation of the previous formula, e.g.,

$$x_2 \vee x_0 \vee x_1;$$

3. backtracks setting \bar{x}_2 , and continues the search looking for models extending the partial assignment $\{\bar{x}_0, \bar{x}_2\}$. *nOPT-DLL*

- (a) determines and prints the third (and last) optimal model $\mu_2 = \{x_1, x_3\}$,
- (b) computes (6) for μ_2 , i.e.,

$$(\bar{x}_1 \vee \bar{x}_3) \wedge \bar{x}_0 \wedge \bar{x}_2,$$

- (c) adds to the input set of clauses a set of clauses corresponding to the negation of the previous formula, e.g.,

$$(x_1 \vee x_0 \vee x_2) \wedge (x_3 \vee x_0 \vee x_2); \tag{7}$$

4. backtracks setting

- (a) either \bar{x}_1 and thus the current assignment is $\{\bar{x}_0, \bar{x}_2, \bar{x}_1\}$
- (b) or \bar{x}_3 and thus the current assignment is $\{\bar{x}_0, \bar{x}_2, \bar{x}_3\}$.

In both cases the formula (7) is contradicted, and given that the search tree has been entirely explored, *nOPT-DLL* terminates returning FALSE.

nOPT-DLL prints all and only the optimal models of φ wrt $\langle S, < \rangle$ as stated by the following theorem.

Theorem 4 Consider Fig. 3, and let $\varphi, \langle S, < \rangle$ and *nOPT-DLL* as in the figure. *nOPT-DLL*() prints all and only the optimal models of φ wrt $\langle S, < \rangle$.

Proof Assume φ is satisfiable, otherwise *nOPT-DLL* prints nothing and the thesis trivially holds.

Let μ_1, \dots, μ_n ($n \geq 1$) be the models printed by *nOPT-DLL*, listed according to the order in which they are printed (thus μ_1 is the first model printed by *nOPT-DLL*). The first observation is that if a model μ of the input set of clauses is not printed by *nOPT-DLL* then μ is not optimal. Indeed, if μ is not printed, then μ falsifies one of the clauses added to φ when a model μ' is determined and printed, and this implies that $\mu' < \mu$. Thus, the set $T = \{\mu_1, \dots, \mu_n\}$ is a superset of the set of optimal models.

It remains to be showed that each model μ in T is optimal, i.e., that for each other model $\mu' \in T, \mu' \not< \mu$.

For any i, j with $1 \leq i < j \leq n, \mu_j \not< \mu_i$. To show this,

1. let $l_1; \dots; l_k; \overline{l_{k+1}}; \overline{l_{k+2}}; \dots; l_{|P|}$ be the sequence of literals in μ_i listed according to the order in which they are assigned by *nOPT-DLL*, and
2. let $l_1; \dots; l_k; \overline{l_{k+1}}; \overline{l'_{k+2}}; \dots; l'_{|P|}$ be the sequence of literals in μ_j listed according to the order in which they are assigned by *nOPT-DLL*,

($k \geq 0$). Thus, l_{k+1} is the first literal which is assigned differently by μ_i and μ_j , i.e., l_{k+1} has been assigned at line 8 and $\overline{l_{k+1}}$ has been assigned at line 10 after $l_1; \dots; l_k$ by *nOPT-DLL*. There are two cases:

1. $l_{k+1} \in S$ and thus by definition of *ChooseLiteral* all the literals $l < l_{k+1}$ are assigned by $\{l_1, \dots, l_k\}$ and thus in the same way by μ_i and μ_j . Hence $\mu_j \not< \mu_i$.
2. $l_{k+1} \notin S$ and thus by definition of *ChooseLiteral* all the literals in S are assigned by $\{l_1, \dots, l_k\}$ and thus in the same way by μ_i and μ_j . Hence $\mu_j \not< \mu_i$.

Thus, each model μ_i is optimal:

1. For each model μ_k with $k < i$ (i.e., printed before μ_i), $\mu_k \not< \mu_i$ because μ_i satisfies the formula (6) added to φ when μ_k has been computed.
2. For each model μ_k with $k > i$ (i.e., printed after μ_i), we have shown that $\mu_k \not< \mu_i$.

□

4 Quantitative and qualitative preferences on formulas and their mixing

In this section, we first introduce quantitative preferences on literals and then we generalize the concept of preferences on literals to preferences on formulas, showing how all these notions (both separately or mixed) can be reduced to the basic framework of qualitative preference on literals.

4.1 Quantitative preferences on literals

Given a set of preferences S and a formula ψ , if it is not possible to satisfy both S and ψ , a standard approach to model the relative importance of the preferences in S , is to define a function $c : S \mapsto \mathcal{N}^+$: Intuitively, $c(l)$ is the reward for satisfying $l \in S$. A pair $\langle S, c \rangle$ is a *quantitative preference* and a model μ of ψ is *optimal* if it maximizes the *objective function* defined as⁸

$$\sum_{l \in S \cap \mu} c(l). \tag{8}$$

In the literature, such kind of problem is also known as Binate Covering Problem [16], recently generalized in [39].

Considering (1) and (2), if we have the preferences

1. $\{\neg Bike_0, \neg Bus_0, \neg Car_0\}$, assuming the reward function c is constant, then the optimal models are $\{Bike_0, AtWork_1\}, \{Bus_0, AtWork_1\}, \{Car_0, AtWork_1\}$.
2. $\{\neg Bike_0, \neg Bus_0, \neg Car_0\}$, if we assume $c(\neg Bike_0) = 2$ while $c(\neg Bus_0) = 1$ and $c(\neg Car_0) = 1$, then the optimal models are $\{Bus_0, AtWork_1\}$ and $\{Car_0, AtWork_1\}$.

Consider a quantitative preference $\langle S', c \rangle$ and a satisfiable set of clauses φ' .

The problem of finding an (resp. all) optimal model (resp. models) of φ' wrt $\langle S', c \rangle$ can be solved again using OPT-DLL (resp. NOPT-DLL) as core engine. The basic idea is to encode the value of the objective function (8) as a sequence of bits b_{n-1}, \dots, b_0 and then consider the qualitative preference $\langle \{b_{n-1}, \dots, b_0\}, \{b_i < b_j : 0 \leq j < i < n\} \rangle$. In more details, let $adder(S', c)$ be a set of clauses such that:

1. If $n = \lceil \log_2(\sum_{l \in S'} c(l) + 1) \rceil$, $adder(S', c)$ contains n new variables b_{n-1}, \dots, b_0 ; and
2. A total assignment μ satisfies φ' iff there exists a unique total assignment μ' to the variables in φ' and in $adder(S', c)$ such that
 - (a) μ' extends μ and satisfies both φ' and $adder(S', c)$, and
 - (b) $\sum_{l \in S' \cap \mu} c(l) = \sum_{i=0}^{n-1} \mu'(b_i) \times 2^i$, where $\mu'(b_i)$ is 1 if $b_i \in \mu'$, and is 0 otherwise.

⁸Assuming we want $c(l) < 0$ for some $l \in S$, we can replace l with \bar{l} in S and define $c(\bar{l}) = -c(l)$: The set of optimal models does not change. Given $\langle S, c \rangle$ and assuming we are interested in minimizing the objective function (8), we can consider the quantitative preference $\langle \bar{S}, c' \rangle$ with $c'(l) = c(\bar{l})$, and then look for a model maximizing $\sum_{l \in \bar{S} \cap \mu} c'(l)$.

If the above conditions are satisfied, we say that $adder(S', c)$ is a *Boolean encoding of $\langle S', c \rangle$ with output b_{n-1}, \dots, b_0* . $adder(S', c)$ can be realized in polynomial time in many ways, see, e.g., [56].

In the above hypotheses, if

1. φ is the set of clauses in φ' or in $adder(S', c)$, and
2. $\langle S, < \rangle$ is the qualitative preference $\langle \{b_{n-1}, \dots, b_0\}, \{b_i < b_j : 0 \leq j < i < n\} \rangle$

then OPT-DLL (resp. *noPT-DLL*) returns one optimal model (resp. prints all the optimal models) of φ' wrt $\langle S', c \rangle$. The following theorem formally states this result for *noPT-DLL*.

Theorem 5 *Let φ' be a set of clauses and let $\langle S', c \rangle$ be a quantitative preference on literals. Let $adder(S', c)$ be a Boolean encoding of $\langle S', c \rangle$ with output b_{n-1}, \dots, b_0 . If*

1. φ is the set of clauses in φ' or in $adder(S', c)$, and
2. $\langle S, < \rangle$ is the qualitative preference $\langle \{b_{n-1}, \dots, b_0\}, \{b_i < b_j : 0 \leq j < i < n\} \rangle$
3. M is the set of models of φ printed by *noPT-DLL* in Fig. 3,

then the assignments in M , restricted to the signature of φ' , are all the optimal models of φ' wrt $\langle S', c \rangle$.

Proof The qualitative preference $\langle \{b_{n-1}, \dots, b_0\}, \{b_i < b_j : 0 \leq j < i < n\} \rangle$ induces a partial order on the models of $\varphi' \cup adder(S', c)$ according to which $\mu < \mu'$ if and only if

$$\sum_{i=0}^{n-1} \mu(b_i) \times 2^i > \sum_{i=0}^{n-1} \mu'(b_i) \times 2^i,$$

i.e., if and only if

$$\sum_{l \in S' \cap \mu} c(l) > \sum_{l \in S' \cap \mu'} c(l).$$

□

As an application of the above theorem, given a quantitative preference with preferences $S' = \{\neg Bike_0, \neg Bus_0, \neg Car_0\}$ and reward function c always returning 1, then $adder(S', c)$ has two bits b_1, b_0 as output; and the models of $adder(S', c)$ satisfy

$$\begin{aligned} b_0 &\equiv (\neg Bike_0 \equiv \neg Bus_0 \equiv \neg Car_0), \\ b_1 &\equiv ((\neg Bike_0 \wedge \neg Bus_0) \vee (\neg Bike_0 \wedge \neg Car_0) \vee (\neg Bus_0 \wedge \neg Car_0)). \end{aligned} \tag{9}$$

The optimal models of (1), (2) and (9), given the qualitative preference $\langle \{b_1, b_0\}, b_1 < b_0 \rangle$ are $\{Bike_0, AtWork_1, b_1\}$, $\{Bus_0, AtWork_1, b_1\}$, $\{Car_0, AtWork_1, b_1\}$, i.e., the models of (1), (2) whose objective value is 2.

4.2 Qualitative and quantitative preferences on formulas

So far, a preference is a literal, and we have seen how it is possible to use *DLL* to find optimal models wrt both qualitative and quantitative preferences on literals. We now show that the hypothesis that preferences are literals can be waved, i.e.,

that it is possible to generalize the previous concepts and results from literals to arbitrary formulas. The basic idea is to introduce definitions [53] or “names” [46] for the formulas at hand (see also [18, 48]).

First, we define a *qualitative preference on formulas* to be a pair $\langle S, \prec \rangle$ where S is a finite set of formulas and \prec a (strict) partial order on S . The set S of preferences does not need to be consistent. Then, as in Section 2, the partial order on S induces a partial order on the sets of total assignments according to which, if μ and μ' are two total assignments $\mu \prec \mu'$ if and only if

1. there exists a formula $\psi \in S$ satisfied by μ and not by μ' ; and
2. for each formula $\psi' \in S$ satisfied by μ' and not by μ , there exists a formula $\psi \in S$ satisfied by μ and not by μ' such that $\psi \prec \psi'$.

It is easy to see that if the formulas in S are literals, then the above definition coincides with the one given in Section 2. It is also straightforward to generalize the result of Theorem 1 saying that the if $\langle S, \prec \rangle$ is a qualitative preference on formulas, the relation \prec extended to the set of total assignments is a partial order.

For example,

$$\begin{aligned} & \{ \{ AtWork_0 \vee \neg AtWork_1 \vee \neg Bus_0, AtWork_0 \vee \neg AtWork_1 \vee \neg Car_0, \\ & \neg AtWork_0 \vee AtWork_1 \vee \neg Bus_0, \neg AtWork_0 \vee AtWork_1 \vee \neg Car_0 \}, \\ & \{ AtWork_0 \vee \neg AtWork_1 \vee \neg Bus_0 \prec AtWork_0 \vee \neg AtWork_1 \vee \neg Car_0, \\ & \neg AtWork_0 \vee AtWork_1 \vee \neg Car_0 \prec \neg AtWork_0 \vee AtWork_1 \vee \neg Bus_0 \} \end{aligned}$$

models the preference in which

1. we prefer to use neither the car nor the bus for moving from/to home, but
2. we prefer to use the car more than the bus for moving to work, and
3. we prefer to use the bus more than the car for moving to home.

A model μ of a formula ψ is *optimal wrt a qualitative preference on formulas* $\langle S, \prec \rangle$ if μ is a minimal element of the partial order on the models of ψ .

Consider a formula ψ and a qualitative preference on formulas $\langle S, \prec \rangle$.

Instead of ψ and $\langle S, \prec \rangle$ we can consider

1. the qualitative preference on literals $\langle L_S, \prec_S \rangle$, where
 - L_S has a newly introduced variable x_α for each formula $\alpha \in S$, and
 - $x_\alpha \prec_S x_\beta$ if and only if $\alpha \prec \beta$; and
2. the formula

$$\psi \wedge \bigwedge_{\alpha \in S} (x_\alpha \equiv \alpha). \tag{10}$$

Then, if

$$\mu_S = \mu \cup \{x_\alpha : \alpha \in S, \mu \models \alpha\} \cup \{\neg x_\alpha : \alpha \in S, \mu \not\models \alpha\}$$

it is straightforward to see that a model μ of ψ is optimal wrt the qualitative preference on formulas $\langle S, \prec \rangle$ iff μ_S is an optimal model of (10) wrt the qualitative preference on literals $\langle L_S, \prec_S \rangle$. It is also easy to see that (10) can be simplified to

$$\psi \wedge \bigwedge_{\alpha \in S} (\neg x_\alpha \vee \alpha) \tag{11}$$

and we obtain again the desired correspondence between the models of ψ and (11).

Introducing definitions [53] or “names” [46] for the formulas in the preferences allows us also to reduce quantitative preferences on formulas (defined in the obvious way) to quantitative preferences on literals. Further, it allows us to use OPT-DLL and NOPT-DLL as engines for computing optimal models of ψ given a qualitative/quantitative preference on formulas. Similar modeling approaches have been presented in, e.g., [1, 3].

Notice that in our approach, quantitative preferences are reduced to qualitative ones. An advantage of this reduction is that it makes also possible to mix the two, e.g., we can ask (we assume b_{n-1}, \dots, b_0 to be the output bits of $adder(S', c)$):

1. Which among the optimal models according to a qualitative preference $\langle S, \prec \rangle$, are optimal according to a quantitative preference $\langle S', c \rangle$: Such assignments correspond to the optimal models of $\psi \wedge adder(S', c)$ wrt the qualitative preference

$$\langle S \cup \{b_{n-1}, \dots, b_0\}, \prec \cup \{b_i \prec b_j : 0 \leq j < i < n\} \cup \{\alpha \prec b_i : \alpha \in S, 0 \leq i < n\} \rangle.$$

The above preference forces OPT-DLL to consider first $\langle S, \prec \rangle$ and then $\langle S', c \rangle$.

2. or which among the optimal models according to a quantitative preference $\langle S', c \rangle$, are optimal according to a qualitative preference $\langle S, \prec \rangle$: Such assignments correspond to the optimal models of $\psi \wedge adder(S', c)$ wrt the qualitative preference

$$\langle S \cup \{b_{n-1}, \dots, b_0\}, \prec \cup \{b_i \prec b_j : 0 \leq j < i < n\} \cup \{b_i \prec \alpha : \alpha \in S, 0 \leq i < n\} \rangle.$$

The above preference forces OPT-DLL to consider first $\langle S', c \rangle$ and then $\langle S, \prec \rangle$.

For example, when buying a computer, assuming that we have a qualitative preference $\langle S, \prec \rangle$ on components, we may want to know

1. which are the least expensive computers among the ones which are optimal according to $\langle S, \prec \rangle$,
2. or, alternatively, what is the optimal computer according to $\langle S, \prec \rangle$ among the least expensive ones.

5 Implementation and experimental results

In order to test the viability of our approach we implemented our ideas in MINISAT [21], the 2005 version, winner of the SAT 2005 competition on the industrial benchmarks category (together with the SAT/CNF minimizer SATELITE [20]). Such choice is motivated by our interest in solving, in particular, large structured problems coming from applications in general, planning and formal verification in particular. It has to be noted that MINISAT is one of the most famous and efficient implementation of what is nowadays called a Conflict-Driven Clause Learning (CDCL) SAT solver [44]. However, this choice is driven by the interest on solving structured instances; other SAT solvers, e.g., SATZ, does not have learning incorporated, and may perform (much) better on different (e.g., random or crafted) benchmarks. Indeed, the algorithm in Fig. 2 is focused on DLL, which is the basic part of almost all available SAT solvers. Issues related to the use of SAT solvers with non-chronological backtracking and learning to solve optimization problems have been first studied in [37], and then in more recent publications by the same authors.

For comparative benchmarking, we focused on $\text{MIN-ONE}/\text{MIN-ONE}_{\subseteq}$ and $\text{MAX-SAT}/\text{MAX-SAT}_{\subseteq}$ problems, precisely defined as below:

1. Given a set S of variables, in $\text{MIN-ONE}_{\subseteq}^S(\varphi)$ (resp. $\text{MIN-ONE}^S(\varphi)$) the goal is to find one or all the optimal models of φ wrt the qualitative (resp. quantitative) preference on literals $\langle \bar{S}, \emptyset \rangle$ (resp. $\langle \bar{S}, c \rangle$, where c is a constant function). When all variables are considered, i.e., when $S = P$, we say that we are considering a *standard* $\text{MIN-ONE}/\text{MIN-ONE}_{\subseteq}$ problem, and a *partial* version of the problem otherwise.
2. Given a set S of clauses, in $\text{MAX-SAT}_{\subseteq}^S(\varphi)$ (resp. $\text{MAX-SAT}^S(\varphi)$) the goal is to find one or all the optimal models of φ wrt the qualitative (resp. quantitative) preference on formulas $\langle S, \emptyset \rangle$ (resp. $\langle S, c \rangle$, where c is a constant function). When there are no hard constraints, i.e., when $\varphi = \emptyset$, we say that we are considering a *standard* $\text{MAX-SAT}/\text{MAX-SAT}_{\subseteq}$ problem, and a *partial* version of the problem otherwise.

It is thus clear that these two categories of problems are the simplest cases of preferences on literals ($\text{MIN-ONE}/\text{MIN-ONE}_{\subseteq}$) and on formulas ($\text{MAX-SAT}/\text{MAX-SAT}_{\subseteq}$), having an empty partial order. Further,

1. for MIN-ONE and MAX-SAT problems, highly tuned tools are available, designed for international competitions; and
2. the number of preferences is very high, equal to the number of clauses in $\text{MAX-SAT}/\text{MAX-SAT}_{\subseteq}$ and to the number of variables in $\text{MIN-ONE}/\text{MIN-ONE}_{\subseteq}$.

Given the above, the goal of the experimental analysis,

1. on $\text{MIN-ONE}_{\subseteq}/\text{MAX-SAT}_{\subseteq}$ problems, is to show that our approach for computing optimal models wrt qualitative preferences is viable also when the number of preferences is very high; and
2. on $\text{MIN-ONE}/\text{MAX-SAT}$ problems, is to show that our reduction from quantitative to qualitative preferences is viable also when the number of preferences is very high.

Problems with a high number of preferences are particularly interesting because the more preferences we have, the more the performances of the underlying SAT solver are negatively affected, see, e.g., [25]. Further, the availability of very efficient tools for $\text{MIN-ONE}/\text{MAX-SAT}$ gives a good reference point to evaluate the results, also for $\text{MIN-ONE}_{\subseteq}/\text{MAX-SAT}_{\subseteq}$: Indeed, an optimal solution for MIN-ONE (resp. MAX-SAT) is also optimal for $\text{MIN-ONE}_{\subseteq}$ (resp. $\text{MAX-SAT}_{\subseteq}$). However, we want to remark that the applicability of our ideas go far beyond, allowing for solving problems with any partial order on preferences, for which problems no implemented system is available for comparative analysis.

Starting from MINISAT , the modifications needed in order to solve $\text{MIN-ONE}_{\subseteq}$ problems have been minor: In practice, we have modified the vsIDS -like heuristic of MINISAT in order to first branch on the literals in the set of preferences. Analogously for $\text{MAX-SAT}_{\subseteq}$, once preferences on formulas are reduced to preferences on literals. $\text{MIN-ONE}/\text{MAX-SAT}$ problems also required the implementation of a function *adder* as specified in Section 4.1. As we already said, there are various ways to implement such a function. We used the method described in [56], which takes linear time in the size

of the input. We also experimented with the method described in [4] for encoding the objective function. This last encoding has some interesting computational properties, (e.g., it is efficiently coupled with the unit-propagation technique implemented in all DLL-based SAT solvers), but the resulting encoding is quadratic in the size of the input, and for this reason—in our experience—it is usually very effective on small instances, but not practical for instances of medium and big size, like the majority of the ones we used.

Beside the modification in the heuristic, we had also to modify MINISAT internal pre-processor in order to disable the assignment of pure literals.⁹

Concerning the other solvers for MAX-SAT/MIN-ONE, we initially considered both dedicated solvers for MAX-SAT problems—like BF [7], MAXSOLVER [57], TOOLBAR [18, 31] ver. 3.0, MAXSATZ version submitted to the 2007 Evaluation [34], MINIMAXSAT (abbreviated with MMSAT in the Tables) ver. 1.0 [26], MSU1.2 [40]—and generic PB solvers—like OPBDP ver. 1.1.1 [5], PBS ver. 2.1 and ver. 4 [2], MINISAT+ (abbreviated with MSAT+ in the Tables) based on MINISAT ver. 1.13 [22], GLPPB ver. 0.2 (by the same authors of PUEBLO [51]),¹⁰ BSOLO ver. 3.0.17 [38]. All these systems are among the state-of-the-art solvers for MAX-SAT or PB problems (see the results of the last evaluations). For the standard MAX-SAT_⊆ problems, we also considered CAMUS [35]. CAMUS is a system for computing all Minimal Unsatisfiable Subsets (MUSes) of a given formula which, as first step, computes all the MAX-SAT_⊆ optimal solutions: Of course, in the tables, we consider only the time CAMUS takes for generating the MAX-SAT_⊆ solutions.

Each solver has been run using its default settings. All the experiments have been run on a Linux box equipped with a Pentium IV 3.2 GHz processor and 1GB of RAM. In the Tables, “TIME” indicates that the solver does not solve the instance within 1,800 s; “MEM” indicates that the solver requires more than the allocated 800 MB; “SF” indicates that the solver exits abnormally; “–” indicates that the solver returns an incorrect result. Moreover, best performing system(s) on each benchmark/domain are emphasized, in bold: For a domain, we count number of instances solved, with ties broken by mean CPU time (as customary in Max-SAT evaluations).

Considering the dedicated solvers for MAX-SAT, we discarded BF, MAXSOLVER and TOOLBAR after an initial analysis because they seem to be tailored for relatively small typically randomly generated problems, and are thus not suited to deal with most of the problems we consider. Concerning the PB solvers, we do not show the results for OPBDP, PBS ver 2.1 and ver. 4, and also GLPPB because they are almost always slower than the other systems on the instances that we consider.

The next two subsections show the results for MIN-ONE/MIN-ONE_⊆ and for MAX-SAT/MAX-SAT_⊆ problems, respectively.

⁹A literal l is *pure* in a set of clauses φ if \bar{l} does not belong to any clause in φ . If l is pure in φ , φ is satisfiable if and only if $\varphi_{(l)}$ is. However, in an optimization problem, it may be the case that there exist optimal models with \bar{l} assigned to true, and thus the necessity to disable the assignment of pure literals in MINISAT internal pre-processor.

¹⁰<http://www.eecs.umich.edu/~hsheini/pueblo/>

5.1 MIN-ONE and MIN-ONE_⊆

The results for MIN-ONE/MIN-ONE_⊆ problems are reported in Table 1, on a variety of well-known, publicly available satisfiable benchmarks. In particular, (1–5) are Formal Verification instances ((1–2) from the Beijing’96 competition, (3–5) by Ofer Shtrichman); (6–14) are planning problems from SATPLAN; (15–20) are Data Encryption Standard instances; (21–26) are quasi group instances. The domains contain, in general, more benchmarks than the one we show: The instances showed representative of the behavior of the systems, and, in general, smaller (resp. bigger) instances are easy (resp. hard) to solve for most of the systems.

The first column of the table contains the number of the instance, followed by its name in the second column. The third column is the optimal value of the instance. Then, columns 4–7 show the results of each system on the MIN-ONE version of the problem, while the last column shows *nOPTSAT* results on the MIN-ONE_⊆ version of the problem.

The columns for all the systems but *nOPTSAT*, show the CPU time (in seconds) that the system takes to compute an optimal solution. For *nOPTSAT*, we show both the time that *nOPTSAT* takes to compute the first solution (subcolumns T_1), and (in subcolumn *#Sols*) the number of optimal solutions computed by *nOPTSAT* in 3,600 s, together

Table 1 Results on MIN-ONE (columns 3–7) and MIN-ONE_⊆ (column 8) problems

Instance	#C	MSAT+	BSOLO	MMSAT	<i>nOPTSAT</i>		<i>nOPTSAT</i>			
					T_1	<i>#Sols</i>	#C _⊆	T_1	<i>#Sols</i>	
1	bcomp5	39	0.4	4.98	0.22	1.74	3,360 A	40	0	21,600 A
2	bmax6	61	8.42	1,401.72	1.41	430.33	24 T	62	0	183,072 T
3	ibm2	940	19.73	19.96	2.04	TIME	0 T	966	0.02	33,132 M
4	ibm3	6,356	TIME	TIME	TIME	25.37	16 A	6,371	0.39	5,679 M
5	gal8		SF	MEM	TIME	TIME	0 T	9,372	0.96	1,540 M
6	3blocks	56	0.29	0.5	3.85	0.56	1 A	60	0.02	174 A
7	4blocksb	66	0.24	0.65	4.87	1.3	4 A	66	0.05	4 A
8	4blocks	108	50.94	353.32	1,086.05	TIME	1 A	110	0.16	55,097 A
9	large.c	265	0.96	2.84	35.24	1.18	1 A	265	0.16	6 A
10	large.d	431	7.71	51.05	506.6	42.73	4 A	432	0.81	106 A
11	log.a	135	1.39	TIME	3.07	1.88	11,305 M	135	0.02	108,151 M
12	log.b	138	8.99	TIME	3.03	6.1	16,728 M	138	0.02	108,675 M
13	rock.a	65	0.2	1.18	0.4	1.58	13,100 A	65	0.01	13,100 A
14	rock.b	69	0.27	0.55	0.41	1.92	902 A	69	0.01	902 A
15	r2b3.1	141	0.2	0.07	3.36	0.12	2 A	141	0.02	2 A
16	r2b3.2	138	0.08	0.08	2	0.11	1 A	138	0.02	1 A
17	r3b1.1	119	1.3	5.93	22.96	28.77	1 A	119	0.14	1 A
18	r3b1.2	126	0.82	5.62	25.04	0.93	1 A	126	0.13	1 A
19	r3b2.1	217	0.46	1.32	47.03	0.72	1 A	217	0.08	1 A
20	r3b2.2	206	0.53	1.39	42.48	0.58	1 A	206	0.06	1 A
21	qg1-8	64	31.06	414.21	131.28	78.73	16 A	64	0.84	16 A
22	qg2-7	49	0.27	1.19	2.29	0.28	14 A	49	0.13	14 A
23	qg2-8	64	21.83	200.82	111.11	43.9	2 A	64	1.03	2 A
24	qg3-8	64	0.1	0.41	1.29	0.32	18 A	64	0.02	18 A
25	qg4-9	81	19.36	77.19	98.31	50.58	194 A	81	0.02	194 A
26	qg5-11	121	0.43	1.12	4	0.34	5 A	121	0.1	5 A

Bold entries highlight the winner solver for each problem (specified in rows)

with the indication about whether *nOPTSAT* computed all the solutions before the time out (indicated with “A”), or whether *nOPTSAT* exceeded the time out (indicated with “T”), or whether *nOPTSAT* exceeded the available memory (indicated with “M”).¹¹ Looking at the results, we see that *nOPTSAT* compares well wrt the other solvers when computing one solution: *MINISAT+/BSOLO/MMSAT/nOPTSAT* are not able to find an optimal solution in 2/4/2/3 cases respectively, and *nOPTSAT* is the only system able to solve the “ibm3” instance. From column *#Sols* it also emerges that our solver is able to compute all the optimal solutions for many problems (21 out of 26), exceeding the available memory in two cases and timing out in three cases. Considering the five cases in which *nOPTSAT* is not able to compute all the optimal solutions with the given resources, we have to stress that in general a user is interested to a relatively small number of solutions (say 20), possibly generated one after the other upon request. Thus, a more reasonable reading of the results in column “*#Sols*” is that our system is able to compute all or at least 20 optimal solutions in all cases but two.

Concerning the results for *MIN-ONE_⊆* problems, these are reported in the last column only for *nOPTSAT* being the only system able to directly deal with these problems. The first two subcolumns show the optimal value returned and the time taken by *nOPTSAT* when computing the first optimal solution; the last subcolumn *#Sols* has the same meaning as before. Looking at the results, the first observation is that almost all the problems are solved in less than 1s, and that also the *gal8* problem is solved. However, comparing *nOPTSAT* results on *MIN-ONE* and *MIN-ONE_⊆* problems, two other observations are in order:

1. Considering the performances in solving *MIN-ONE* and *MIN-ONE_⊆* problems, we see that the latter are solved in much less time. This could have been expected given that handling *MIN-ONE* problems requires the encoding of adders counting the number of variables set to true, and many of the examples have more than a thousand variables (the “*gal8*” instance has > 58,000 variables: The resulting adder has > 270,000 variables and > 1,300,000 clauses).
2. Considering the optimal value of the first solution found (columns *#C* and *#C_⊆*) we see that for most instances *#C* = *#C_⊆*, and, even when not equal, the two values are very close but for rows 3, 4, 6. This points out that *nOPTSAT* for *MIN-ONE_⊆* can be used as a good approximation algorithm for *MIN-ONE* problems.
3. Considering the number of optimal solutions (columns *#Sols*), *nOPTSAT* is able to determine many more optimal solutions in the case of *MIN-ONE_⊆* problems. Interestingly, in six cases *nOPTSAT* is not able to find all solutions in *MIN-ONE_⊆* problems (compared to the five problems for *MIN-ONE*), and in five out of the six cases, *nOPTSAT* exceeds the available memory: This can be easily explained by the fact that whenever a solution is found, a set of clauses is added to the input formula in order to rule out the solutions which are dominated by the solutions found. However, as we already said, we do not expect that users want to generate, e.g., 183,072 optimal models of problem “*bmax6*”, but rather a much smaller number.

¹¹Notice that we have two different time outs: 1,800 s for computing one solution, for all the systems. For *nOPTSAT*, we raised the time out to 3,600 s when trying to compute all solutions. This explains why for the problem “4blocks” (#8 in the table) for *nOPTSAT* subcolumn *T₁* we have the value “TIME” indicating that the system did not compute any solution in 1,800 s, and “1 A” in column “*#Sols*” indicating that the system computed the only optimal solution in less than 3,600 s.

Table 2 Results on partial MIN-ONE (columns 3–7) and partial MIN-ONE_⊆ (column 8) problems

Instance	#C	MSAT+	BSOLO	MMSAT	nOPTSAT		nOPTSAT			
					T ₁	#Sols	#C _⊆	T ₁	#Sols	
27	air15	58	16.37	5.35	22.75	9.53	1,474 M	58	0.3	764 M
28	block5-2	16	1.58	31.79	36.96	30.55	27 A	16	0.14	27 A
29	dep7	23	78.92	TIME	TIME	187.84	1,458 M	23	120.6	1,743 M
30	driv10	20	8.39	415.89	902.16	17.19	512 A	20	42.77	512 A
31	free3	21	14.69	169.24	803.63	7.68	2 A	21	2.71	2 A
32	log6-9	24	2.72	65.05	8.87	23.22	14,171 M	28	0.01	7,863 M
33	mprime2	9	TIME	327.73	TIME	49.45	786 M		MEM	0 M
34	mprime5	11	TIME	TIME	TIME	129.08	663 M	11	75.97	917 M
35	myst2	9	160.63	236.32	TIME	28.71	1,130 M	9	1,340.94	571 M
36	opt11	216	TIME	MEM	TIME	159.93	533 M	216	972.95	85 M
37	path5	30	25.48	1,194.79	50.46	112.57	5,046 M	30	0.06	3,196 M
38	phil29	330	4.38	10.63	339.26	1.48	38,196 T	330	0.84	20,961 M
39	pipe6	8	16.15	838.54	940.27	238.65	9,715 M	8	1.76	6,320 M
40	pipet6	8	232.83	TIME	TIME	121.09	5,152 M	8	29.7	3,274 M
41	psr29	21	18.03	234.09	TIME	35.12	1,999 M		TIME	0 T
42	psr31	19	20.11	172.78	TIME	24.68	632 M	19	5.86	750 M
43	psr47	27	25.23	187.86	922.92	83.59	2,197 M	27	2.56	3,286 M
44	sat3	13	16.92	559.94	195.75	22.29	192 A	13	3.3	192 A
45	stor7	14	64.41	1,305.37	436.97	59.06	7,900 M	14	0.19	6,994 M
46	truck2	17	39.31	TIME	703.9	35.55	6,510 M	17	1.2	2,063 M
47	zeno8	15	123.3	1,175.87	TIME	156.13	1,523 M	15	6.88	2,459 M

Bold entries highlight the winner solver for each problem (specified in rows)

We also considered partial MIN-ONE/MIN-ONE_⊆ planning problems generated with SATPLAN 2004, release of 10 Feb. 2006.¹² SATPLAN works as follows: Given a planning problem Π and a makespan *n* (initially set to 0), it

1. generates a corresponding SAT formula Π_{*n*}, and checks Π_{*n*} for satisfiability;
2. if Π_{*n*} is satisfiable then SATPLAN stops and a plan with optimal makespan is returned;
3. otherwise, *n* is increased and the process is repeated.

In our experiments, we selected various planning problems from previous International Planning Competitions (IPCs); we considered the first satisfiable instance generated by SATPLAN, and we fixed the set *S* of preferences to be the set of action variables: With such preferences, we are looking for a plan with as few action variables as possible set to true. In Table 2 we show the results, where columns have the same meaning as in Table 1.

On these partial MIN-ONE planning problems, nOPTSAT performs very well: MIN-ISAT+/BSOLO/MMSAT/nOPTSAT are not able to find an optimal solution in 3/5/9/0 cases respectively, and nOPTSAT is the only system able to solve the “mprime5” and “opt11” instances. From column #Sols it emerges that nOPTSAT is able to compute all the optimal solutions only in a few cases (four out of 21), exceeding the available memory in 16 cases and timing out in one case.

¹²<http://www.cs.rochester.edu/u/kautz/satplan/index.htm>

Considering also *nOPTSAT* results on partial $\text{MIN-ONE}_{\subseteq}$ (last column), it is no longer true that *nOPTSAT* is faster when solving a partial $\text{MIN-ONE}_{\subseteq}$ problem than the corresponding partial MIN-ONE problem. It is also interesting to notice that, given the available resources in time and memory, for many problems *nOPTSAT* is able to find more optimal solutions in the MIN-ONE case than in the $\text{MIN-ONE}_{\subseteq}$ one (e.g., for “air15”). We believe that this is mostly due to the specific structure of the problems in which

1. in the MIN-ONE case, *nOPTSAT* is able to quickly determine the optimal value V for the solution, i.e., *nOPTSAT* is able to quickly determine the non existence of a plan with less than V actions, and then *nOPTSAT* is free to select the V actions,
2. in the $\text{MIN-ONE}_{\subseteq}$ case, *nOPTSAT* has to pay the price of always setting the variables in the preferences to FALSE , i.e., *nOPTSAT* has always to decide to not execute an action, even though there’s a strong evidence (given by the score of the heuristic) that the action has to belong to the plan being built.

5.2 MAX-SAT and $\text{MAX-SAT}_{\subseteq}$

For $\text{MAX-SAT}/\text{MAX-SAT}_{\subseteq}$ problems we considered non random benchmarks from the Max-SAT Evaluations 2007 and 2008. The results are shown in Table 3 and in Table 4 for the standard and partial case respectively, organized as in the report of such evaluations: Each row corresponds to a domain of benchmarks indicated in the first column,¹³ the second column contains the number of instances,¹⁴ and the remaining columns report, for each system, the number of solved instances in parenthesis, and the average time taken to solve them: So, for example, in the “spinglass” domain there are five instances, and MAXSATZ solves three of them, taking 33.19 s on average. For *nOPTSAT*, we show the data as for the other systems in subcolumn T_1 , and the number of times *nOPTSAT*

1. is able to compute all solutions (first number in subcolumn (A/T/M)),
2. exceed the available time (second number in subcolumn (A/T/M)),
3. exceed the available memory (third number in subcolumn (A/T/M)).

For example, in the MAX-SAT case of the “spinglass” domain, *nOPTSAT*

1. is able to compute one optimal solution in one case taking (on average) 7.52 s, and
2. when computing all the optimal solutions, in one case terminates normally (meaning that *nOPTSAT* is able to compute all the optimal solutions); in four cases times out, in zero cases exceeds the available memory.

¹³The majority of domains are (partial) Max-SAT problems originally designed in PB format and submitted to PB evaluations, and then expressed as Max-SAT problems. The instances we have given to PB solvers are the ones that come from a translation first in the input format of *nOPTSAT*, and then in the PB format: This is because we did want to evaluate the systems on the very same formulation. Results for PB solvers on the original PB instances may thus be different.

¹⁴The industrial domain contains 112 instances: The 94 mentioned in Table 3 do not contain the biggest instances. Nonetheless, the considered instances are very hard for the solvers considered and already big, i.e., in the order of tens of MB: Thus, the remaining instances are likely to reach the memory limit.

Table 3 Results on standard MAX-SAT (columns 3–8) and standard MAX-SAT_⊆ (columns 9–10) problems from the MAX-SAT Evaluation 2008

Domain	#I	MSAT+	BSOLO	MAXSATZ	MMSAT	MSU1.2	#OPTSAT		CAMUS		#OPTSAT	
							T ₁	A/T/M	T ₁	A/T/M	T ₁	A/T/M
Industrial	94	40.96(1)	219.27(1)	95.18(3)	7.505(2)	60.38(85)	262.72(7)	(7/87/0)	98.93 (54)	(5/89/0)	492.39 (58)	(2/90/2)
Spinglass	5	0.86(1)	76.57(1)	33.19(3)	1.09(3)	0.0(0)	7.52(1)	(1/4/0)	109 (1)	(0/5/0)	0.00 (5)	(0/5/0)
dimacs_mod	62	247.54(7)	0.01(2)	59.27(52)	194.52(52)	0.03(3)	67.44(4)	(2/60/0)	0.006 (2)	(2/60/0)	0.01 (62)	(1/61/0)

Bold entries highlight the winner solver for each problem (specified in rows)

Table 4 Results on partial MAX-SAT (columns 3–7) and partial MAX-SAT_≡ (column 8) problems from the MAX-SAT Evaluation 2008 (first five domains) and from the MAX-SAT Evaluation 2007 (last four domains)

domain	#I	MSAT+	BSOLO	MMSAT	MSU1.2	#OPTSAT		A/T/M	
						T ₁	A/T/M	T ₁	A/T/M
bcp-fir	59	48.82(22)	44.89(10)	142.05(14)	166.47(48)	217.43(44)	2.44(59)	(15/12/32)	(37/49)
bcp-mtg	215	0.2(215)	48.33(190)	28.3(208)	19.15(174)	0.2(215)	8.91(215)	(81/6/128)	(57/100/58)
bcp-syn	74	34.27(32)	169.22(21)	61.63(29)	60.84(32)	16.46(26)	0.04(74)	(16/50/8)	(4/26/44)
pbo-mqc/nenc	128	124.68(120)	248.03(95)	307.38(67)	54.72(47)	156.24(80)	564.64(48)	(69/52/7)	(9/92/27)
pbo-mqc/nlogenc	128	39.19(119)	268.95(111)	279.79(104)	109.29(70)	218.03(96)	52.33(110)	(77/43/8)	(14/25/89)
PSEUDO/primes	148	11.52(104)	22.23(94)	62.08(107)	128.7(24)	33.11(104)	4.60(130)	(90/47/11)	(86/17/45)
PSEUDO/routing	15	43.74(15)	373.73(8)	109.49(14)	0.85(15)	58.65(15)	11.22(15)	(0/0/15)	(0/0/15)
MAXONE/struct	60	2.02(58)	40.96(60)	22.5(60)	0.30(1)	546.23(50)	1.45(60)	(56/4/0)	(56/0/4)
MAXCLIQ/struct	62	154.39(22)	248.26(14)	61.97(36)	1.56(4)	57.05(19)	0.11(62)	(17/45/0)	(10/4/48)

Bold entries highlight the winner solver for each problem (specified in rows)

MAXSATZ (resp. CAMUS) can only handle standard MAX-SAT (resp. MAX-SAT_⊆) problems and thus it does not appear in Table 4. Extended versions of MAXSATZ, namely W-MAXSATZ and INCWMAXSATZ, can handle partial Max-SAT problems: However they have not been added to Table 4 because, while being very effective on random benchmarks, other solvers we evaluated perform better on non random benchmarks, as witnessed by the results of Max-SAT Evaluations 2007 and 2008.

Looking at the result for MAX-SAT, the first observation is that there is not a system clearly outperforming the others: Different systems perform better than the others in different domains, and this is in line with the results of the last evaluations. Indeed, there is a great variety in the techniques used by the solvers and thus we can expect such a variety in the performances, depending on the specific features of the domain at hand. Considering *nOPTSAT* on the MAX-SAT problems, its performances are comparable to those of the other systems: If we rank systems on each domain according first to the number of problems solved and then to the average CPU time reported, we see that *nOPTSAT* is among the top three systems in eight out of 12 domains, being a top system in the *bcp-mtg* domain.

On the other hand, considering the MAX-SAT_⊆ problems, *nOPTSAT* is able to solve almost always all the problems with the notable exception of the *pbo-mqc/nenc* domain, in which *nOPTSAT* is able to solve more problems in the quantitative than in the qualitative setting. The positive results in the qualitative setting wrt the ones in the quantitative setting echo the ones presented in the previous subsection for MIN-ONE_⊆/MIN-ONE. However, differently from the MIN-ONE_⊆/MIN-ONE case, the optimal result (called $\#C_{\subseteq}$ in the previous tables) of the first solution computed by *nOPTSAT* in the MAX-SAT_⊆ case differs, sometimes significantly, from the optimal result in the MAX-SAT_⊆ case (i.e., $\#C$ and $\#C_{\subseteq}$ in the previous tables): In many cases $\#C_{\subseteq}$ is about $0.5 \times \#C$, and only for a few problems $\#C_{\subseteq} = \#C$. Compared to CAMUS, *nOPTSAT* has better performances when the goal is to find one optimal solution: CAMUS on the *industrial/spinglass/dimacs_mod* domains, is able to find an optimal solution in 54/1/2 cases respectively, compared to the 58/5/62 of *nOPTSAT*. However, for the same domains, CAMUS is able to find all optimal solutions in 5/0/2 cases, while *nOPTSAT* corresponding results are 2/0/1.

Summing up the results on MAX-SAT in this subsection, we have seen that *nOPTSAT* performances are comparable to those of the other systems, which have been specifically designed for solving MAX-SAT problems. On standard MAX-SAT_⊆ problems, *nOPTSAT* performs better than CAMUS when computing one solution, but worse when computing all solutions. By further analyzing the results of *nOPTSAT* on MAX-SAT, we have noticed a relation between the number of unsatisfied clauses in the optimal solution, and the performance of *nOPTSAT*: Performance seems to be better when there are “few” unsatisfied clauses in the solution.

However, we have to remind that the benchmarks we have considered in this but also in previous subsection are characterized by a very high number of preferences. Indeed, in many applications, including planning with soft goals [23], we can expect problems with a few (in the order of tens) preferences, and with such a low number of preferences there is hardly any difference between the performances of *nOPTSAT* and those of MINISAT when run on the original SAT instance (see [25]).

Finally, our system is not tuned in any way: For example, there is no pre-computation of an upper or lower bound for cutting initial portions of the search tree. Herewith we mention how we could use a (pre)computed lower bound for a

MIN-ONE problem in our algorithm. Consider a formula ψ ; assume the output of the *adder* formula are the variables b_{n-1}, \dots, b_0 ; and let $lb > 0$ the pre-computed lower bound. It is relatively easy to add to ψ a set of clauses which inhibit the configurations of bits corresponding to a value o s.t. $o \leq lb$. For example, if $lb = 5$ and $n = 4$, the added clauses are those corresponding to the formula:

$$b_3 \vee (b_2 \wedge b_1)$$

e.g.,

$$(b_3 \vee b_2) \quad (b_3 \vee b_1).$$

With such clauses, the initial branches of *NOPTSAT* search tree corresponding to values $o \leq lb = 5$ are not explored.

6 Conclusions

In this paper we showed that DLL can be used to solve satisfiability problems in the presence of qualitative preferences on literals/formulas by simply imposing an ordering on the literals to be used while branching. The computation of all optimal models requires adding a formula pruning the generation of dominated models. We also showed how it is possible to reduce quantitative preferences to qualitative ones by a Boolean encoding of the objective function. We implemented our ideas in *MINISAT* and we showed that the resulting system compares well with other state-of-the-art systems even on *MIN-ONE*/*MIN-ONE_C*/*MAX-SAT*/*MAX-SAT_C* problems.

There is a huge literature on qualitative and quantitative preferences, see, e.g., the proceedings of the last “Multi-disciplinary Workshop on Advances in Preference Handling” [13], or the 2007 AAI tutorial on “Representing, Eliciting, and Reasoning with Preferences” by Ronen Brafman and Carmel Domshlak.¹⁵ In this paper, we have deeply analyzed the relation between our proposal and other works that deal with computing “optimal” models in SAT and CSP [8, 12], and relate to papers that use our same concept of optimality [54, 55].

Considering our approach for reducing quantitative to qualitative preferences in general, and for solving *MIN-ONE*/*MAX-SAT* problems in particular, the idea of translating the objective function (via an *adder* function) into a set of clauses, to be added to the input formula, has already been introduced and discussed in, e.g., [5, 14, 36], implemented in *MINISAT+* [22] and applied to, e.g., planning [11] and telecommunication feature subscription [15]. However, wrt *MINISAT+*, in our approach we do not run the SAT solver multiple times (one for each different value of the objective function we want to test) till the optimal value is found. Instead, we run the solver *once*, and the modification of the heuristic guarantees that the first model our solver finds is also optimal. Moreover, again in comparison to *MINISAT+*, which is the system whose behavior is “closer” to our, there are several other differences, i.e., (i) *MINISAT+* can not reuse the learned clauses from previous calls, while for us this is inherited within our approach; (ii) we can compute “all”

¹⁵ Available at <http://iew3.technion.ac.il/~dcarmel/tutorial/>.

optimal solutions; and (iii) we can deal with qualitative preferences. Also, MINISAT+ computes an initial bound by calling a SAT solver.

Finally, we acknowledge that our approach for handling qualitative and quantitative preferences has some drawbacks:

1. The underlying SAT solver can not be used as black-box but modifications in the branching heuristics have to be implemented.
2. In SAT, it is known that imposing an ordering on the branching heuristics can lead to significant degradations in performances (see, e.g. [10, 28]).
3. In the quantitative case, we need an *adder* formula encoding the value of the objective function, which can be of significant size.

Despite the above, we have seen that our system compares well with other state-of-the-art systems even on MIN-ONE/MAX-SAT problems. Further, in practice, significant degradations in the performances of the underlying SAT solver shows up only when the number of preferences is very high. Finally, the modifications to the SAT solver are minimal and limited to the branching heuristics: Because of this, when a new, more efficient SAT solver will be available, it will be relatively easy to modify it thus gaining further efficiency.

Acknowledgements The authors would like to thank Carmel Domshlak and Derek Long for useful discussions related to the topic of the paper, and the anonymous reviewers for their useful suggestions.. This work is partially supported by MIUR.

References

1. Aloul, F. A., Ramani, A., Markov, I. L., & Sakallah, K. A. (2002). Generic ILP versus specialized 0-1 ILP: An update. In *Proc. ICCAD* (pp. 450–457).
2. Aloul, F. A., Ramani, A., Markov, I. L., & Sakallah, K. A. (2002). PBS: A backtrack search Pseudo-Boolean solver. In *Proc. SAT*.
3. Amgoud, L., Cayrol, C., & Le Berre, D. (1996). Comparing arguments using preference ordering for argument-based reasoning. In *Proc. ICTAI* (pp. 400–403).
4. Bailleux, O., & Bouffkhad, Y. (2003). Efficient CNF encoding of Boolean cardinality constraints. In *Proc. CP* (pp. 108–122).
5. Barth, P. (1995). *A Davis-Putnam enumeration algorithm for linear Pseudo-Boolean optimization*. Technical report, Max Plank Institute for Computer Science, MPI-I-95-2-2003.
6. Biere, A., Cimatti, A., Clarke, E., & Zhu, Y. (1999). Symbolic model checking without BDDs. In *Proc. TACAS*.
7. Borchers, B., & Furman, J. (1998). A two-phase exact algorithm for Max-SAT and weighted Max-SAT problems. *Journal of Combinatorial Optimization*, 2(4), 299–306.
8. Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., & Poole, D. (2004). Preference-based constrained optimization with CP-nets. *Computational Intelligence*, 20(2), 137–157.
9. Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., & Poole, D. (2004). CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21, 135–191.
10. Buresh-Oppenheim, J., & Pitassi, T. (2003). The complexity of resolution refinements. In *Proc. LICS* (pp. 138–147).
11. Büttner, M., & Rintanen, J. (2005). Satisfiability planning with constraints on the number of actions. In *Proc. ICAPS* (pp. 292–299).
12. Castell, T., Cayrol, C., Cayrol, M., & Le Berre, D. (1996). Using the Davis and Putnam procedure for an efficient computation of preferred models. In *Proc. ECAI* (pp. 350–354).
13. Chomicki, J., Conitzer, V., Junker, U., & Perny, P. (Eds.) (2008). *4th Multidisciplinary Workshop on Advances in Preference Handling (MPREF'08)*. Menlo Park: AAAI.

14. Codish, M., Lagoon, V., & Stuckey, P. J. (2008). Logic programming with satisfiability. *Theory and Practice of Logic Programming*, 8(1), 121–128.
15. Codish, M., Lagoon, V., & Stuckey, P. J. (2008). Telecommunications feature subscription as a partial order constraint problem. In *Proc. ICLP 2008* (pp. 749–753).
16. Coudert, O. (1996). On solving covering problems. In *Proc. DAC* (pp. 197–202).
17. Davis, M., Logemann, G., & Loveland, D. W. (1962). A machine program for theorem proving. *Communication of ACM*, 5(7), 394–397.
18. De Givry, S., Larrosa, J., Meseguer, P., & Schieueux, T. (2003). Solving Max-SAT as weighted CSP. In *Proc. CP* (pp. 363–376).
19. Di Rosa, E., Giunchiglia, E., & Maratea, M. (2008). Computing all optimal solutions in satisfiability problems with preferences. In *Proc. CP* (pp. 603–607).
20. Eén, N., & Biere, A. (2005). Effective preprocessing in SAT through variable and clause elimination. In *Proc. SAT* (pp. 61–75).
21. Eén, N., & Sörensson, N. (2003). An extensible SAT-solver. In *Proc. SAT* (pp. 502–518).
22. Eén, N., & Sörensson, N. (2006). Translating Pseudo-Boolean constraints into SAT. *Journal on Satisfiability, Boolean Modeling and Computation*, 2, 1–26.
23. Gerevini, A., Haslum, P., Long, D., Saetti, A., & Dimopoulos, Y. (2009). Deterministic planning in the 5th IPC: PDDL3 and experimental evaluation of the planners. *Artificial Intelligence*, 173(5–6), 619–668.
24. Giunchiglia, E., & Maratea, M. (2006). Solving optimization problems with DLL. In *Proc. ECAI* (pp. 377–381).
25. Giunchiglia, E., & Maratea, M. (2007). Planning as satisfiability with preferences. In *Proc. AAAI* (pp. 987–992).
26. Heras, F., Larrosa, J., & Oliveras, A. (2008). MiniMaxSat: A new weighted Max-SAT solver. *Journal of Artificial Intelligence Research*, 31, 1–32.
27. Jackson, P., & Sheridan, D. (2004). Clause form conversions for Boolean circuits. In *Proc. SAT* (pp. 183–198).
28. Järvisalo, M., Junttila, T., & Niemelä, I. (2005). Unrestricted vs restricted cut in a tableau method for Boolean circuits. *Annals of Mathematics and Artificial Intelligence*, 44(4), 373–399.
29. Jin, H., & Somenzi, F. (2005). Prime clauses for fast enumeration of satisfying assignments to boolean circuits. In *Proc. DAC* (pp. 750–753).
30. Kautz, H., & Selman, B. (1992). Planning as satisfiability. In *Proc. ECAI* (pp. 359–363).
31. Larrosa, J., & Schieueux, T. (2003). In the quest of the best form of local consistency for weighted CSP. In *Proc. IJCAI 2003* (pp. 239–244).
32. Le Berre, D., & Simon, L. (2004). Fifty-five solvers in Vancouver: The SAT 2004 competition. In *Proc. SAT (selected papers)* (pp. 321–344).
33. Le Berre, D., & Simon, L. (2006). Preface to the special volume on the SAT 2005 competitions and evaluation. *Journal on Satisfiability, Boolean Modeling and Computation*, 2, 1–14.
34. Li, C. M., Manyá, F., & Planes, J. (2007). New inference rules for Max-SAT. *Journal of Artificial Intelligence Research*, 30, 321–359.
35. Liffiton, M. H., & Sakallah, K. A., (2008). Algorithms for computing minimal unsatisfiable subsets of constraints. *Journal of Automated Reasoning*, 40(1), 1–33.
36. Manquinho, V. M., Flores, P. F., Marques Silva, J. P., & Oliveira, A. L., (1997). Prime implicant computation using satisfiability algorithms. In *Proc. ICTAI* (pp. 232–239).
37. Manquinho, V. M., & Marques-Silva, J. P. (2000). On solving boolean optimization with satisfiability-based algorithms. In *Proc. AMAI*.
38. Manquinho, V. M., & Marques-Silva, J. P. (2006). On using cutting planes in Pseudo-boolean optimization. *Journal on Satisfiability, Boolean Modeling and Computation*, 2, 209–219.
39. Manquinho, V. M., Marques Silva, J. P. & Planes, J. (2009). Algorithms for weighted boolean optimization. In *Proc. SAT 2009* (pp. 495–508).
40. Marques-Silva, J., & Planes, J. (2008). Algorithms for maximum satisfiability using unsatisfiable cores. In *Proc. DATE* (pp. 408–413).
41. Marques-Silva, J. P., & Sakallah, K. A. (1996). GRASP—a new search algorithm for satisfiability. In *Proc. ICCAD* (pp. 220–227).
42. McMillan, K. L. (1993). *Symbolic model checking: An approach to the state explosion problem*. Dordrecht: Kluwer Academic.
43. McMillan, K. L. (2002). Applying SAT methods in unbounded symbolic model checking. In *Proc. CAV* (pp. 250–264).
44. Mitchell, D. G. (2005). A SAT solver primer. *Bulletin of the European Association for Theoretical Computer Science*, 85, 112–132.

45. Moskewicz, M. W., Madigan, C. F., Zhao, Y., Zhang, L., & Malik, S. (2001). Chaff: Engineering an efficient SAT Solver. In *Proc. DAC* (pp. 530–535).
46. Plaisted, D. A., & Greenbaum, S. (1986). A structure-preserving clause form translation. *Journal of Symbolic Computation* 2, 293–304.
47. Prestwich, S. D., Rossi, F., Venable, K. B., & Walsh, T. (2005). Constraint-based preferential optimization. In *Proc. AAAI* (pp. 461–466).
48. Ramirez, M., & Geffner, H. (2007). Structural relaxations by variable renaming and their compilation for solving MinCostSAT. In *Proc. CP* (pp. 605–619).
49. Ravi, K., & Somenzi, F. (2004). Minimal assignments for bounded model checking. In *Proc. TACAS* (pp. 31–45).
50. Sakama, C., & Inoue, K. (2000). Prioritized logic programming and its application to common-sense reasoning. *Artificial Intelligence*, 123(1–2), 185–222.
51. Sheini, H. M., & Sakallah, K. A. (2005). Pueblo: A modern Pseudo-Boolean Sat solver. In *Proc. DATE* (pp. 684–685).
52. Siekmann, J., & Wrightson, G. (Eds.) (1983). *Automation of reasoning: Classical papers in computational logic* (Vol. 1–2, pp. 1967–1970). New York: Springer.
53. Tseitin, G. S. (1970). On the complexity of proofs in propositional logics. *Seminars in Mathematics*, 8, 466–483 (Reprinted in [52]).
54. Van Nieuwenborgh, D., Heymans, S., & Vermeir, D. (2004). On programs with linearly ordered multiple preferences. In *Proc. ICLP* (pp. 180–194).
55. Van Nieuwenborgh, D., & Vermeir, D. (2006). Preferred answer sets for ordered logic programs. *Theory and Practice of Logic Programming*, 6(1–2), 107–167.
56. Warners, J. P. (1998). A linear-time transformation of linear inequalities into conjunctive normal form. *Information Processing Letters*, 68(2), 63–69.
57. Xing, Z., & Zhang, W. (2005). MaxSolver: An efficient exact algorithm for (weighted) maximum satisfiability. *Artificial Intelligence*, 164(1–2), 47–80.