

# *Anytime answer set optimization via unsatisfiable core shrinking*

MARIO ALVIANO and CARMINE DODARO

*Department of Mathematics and Computer Science, University of Calabria, Italy*  
(e-mail: {alviano,dodaro}@mat.unical.it)

*submitted 6 May 2016; revised 8 July 2016; accepted 22 August 2016*

---

## Abstract

Unsatisfiable core analysis can boost the computation of optimum stable models for logic programs with weak constraints. However, current solvers employing unsatisfiable core analysis either run to completion, or provide no suboptimal stable models but the one resulting from the preliminary disjoint cores analysis. This drawback is circumvented here by introducing a progression based shrinking of the analyzed unsatisfiable cores. In fact, suboptimal stable models are possibly found while shrinking unsatisfiable cores, hence resulting into an anytime algorithm. Moreover, as confirmed empirically, unsatisfiable core analysis also benefits from the shrinking process in terms of solved instances.

**KEYWORDS:** answer set programming, weak constraints, unsatisfiable cores

---

## 1 Introduction

Answer set programming (ASP) is a declarative formalism for knowledge representation and reasoning based on stable model semantics (Gelfond and Lifschitz 1991; Niemelä 1999; Marek *et al.* 2008; Lifschitz 2008; Eiter *et al.* 2009; Brewka *et al.* 2011), for which efficient implementations such as CLASP (Gebser *et al.* 2015a), CMODELS (Lierler and Maratea 2004; Giunchiglia *et al.* 2006), DLV (Alviano *et al.* 2010), and WASP (Alviano *et al.* 2013; Dodaro *et al.* 2011) are available. ASP programs are associated with classical models satisfying a stability condition: only necessary information is included in a model of the input program under the assumptions provided by the model itself for the *unknown knowledge* in the program, where unknown knowledge is encoded by means of *default negation*. Moreover, the language includes several constructs to ease the representation of real world knowledge, among them aggregates (Simons *et al.* 2002; Liu *et al.* 2010; Bartholomew *et al.* 2011; Faber *et al.* 2011; Ferraris 2011; Alviano and Faber 2013; Gelfond and Zhang 2014; Alviano and Leone 2015; Alviano *et al.* 2015d; Alviano *et al.* 2016).

Reasoning in presence of unknown knowledge is quite common for rational agents acting in the real world. It is also common that real world agents cannot meet all their desiderata, and therefore ASP programs may come with *weak constraints* for

representing numerical preferences over jointly incompatible conditions (Buccafurri *et al.* 2000; Simons *et al.* 2002; Gebser *et al.* 2011a; Gebser *et al.* 2011b). Stable models are therefore associated with a cost given by the sum of the weights of the unsatisfied weak constraints, so that stable models of minimum cost are *preferred*. It is important here to stress the meaning of the word preferred: any stable model describes a plausible scenario for the knowledge represented in the input program, even if it may be only an admissible solution of non optimum cost. In fact, many rational agents would still accept suboptimal solutions, possibly with an estimate on the maximum distance to the optimum cost. This flexibility is also justified by the intrinsic complexity of the problem: the computation of an optimum stable model requires in general at least linearly many calls to a  $\Sigma_2^P$  oracle (Buccafurri *et al.*), and it is therefore practically unfeasible for the hardest instances.

According to the above observations, a good algorithm for answer set optimization should produce better and better stable models during the computation of an optimum stable model. Algorithms having this property are called *anytime* in the literature (Alviano *et al.* 2014; Bliem *et al.* 2016), and among them there is *linear search sat-unsat*, the first algorithm introduced for answer set optimization (Buccafurri *et al.* 2000). In short, any stable model is searched so to initialize an overestimate of the optimum cost, and new stable models of improved cost are iteratively searched until an unsatisfiability arises, proving optimality of the latest found stable model. The algorithm is simple and anytime, but in practice quite inefficient in computing stable models of reasonable cost. Moreover, the algorithm does not provide any underestimate, or error estimation. (Section 3.1 provides additional details within this respect.)

Answer set optimization is often boosted by algorithms based on *unsatisfiable core* analysis, such as OLL (Andres *et al.* 2012) and ONE (introduced by Alviano *et al.* 2015b for MaxSAT, and adapted to ASP in Section 3.2). These algorithms start by searching a stable model satisfying all weak constraints, which would be therefore an optimum stable model. On the other hand, if there is no stable model of this kind, a subset of the weak constraints that cannot be jointly satisfied is identified. Such a set is called *unsatisfiable core*, and essentially evidences that any optimum stable model must sacrifice at least one of the desiderata expressed by the weak constraints. Hence, unsatisfiable cores provide underestimates of the cost of optimum stable models. Moreover, the program can be modified by replacing the weak constraints in the core with new weak constraints that essentially express a preference for stable models satisfying all but one of the original weak constraints, and anyhow the largest number of them, so that the process can be reiterated.

The main drawback of these algorithms is that they run to completion, and therefore provide no intermediate suboptimal stable models for the hardest instances that cannot be solved in reasonable time. In ASP, current alternatives are based on *disjoint cores analysis*, and on *stratification*. Disjoint cores analysis amounts to identifying unsatisfiable cores with empty intersection, and terminates with a (suboptimal) stable model. After that, the usual core based algorithm is run to completion, providing no further intermediate stable models. With stratification, instead, only weak constraints of the greatest weight are first considered, so that

the underestimate may be better improved by the identified unsatisfiable cores. After all cores have been analyzed, a (suboptimal) stable model will be found, and weak constraints with the second greatest weight will be also considered. The process is repeated, hence producing a (suboptimal) stable model at each iteration, possibly improving the current overestimate. However, no intermediate stable model is produced before completing each iteration, which means that few overestimate improvements occur when few different weights occur in the input. Actually, when all weak constraints have the same weight, the algorithm runs to completion providing no intermediate stable models.

The main questions addressed in this paper are therefore the following: How to produce better and better stable models while running core based algorithms for answer set optimization? Should any overhead be paid for obtaining such an anytime behavior? The answer we propose for the first question is the following: Unsatisfiable cores are often non-minimal, and their sizes can be significantly reduced by a few additional oracle calls, where each call may either return a smaller core, or a stable model possibly improving the current overestimate. Within this respect, we implemented two strategies, referred to as *linear* and *reiterated progression based shrinking* (Section 3.3). Concerning the second question, the overhead introduced by the additional oracle calls is often mitigated by the performance gain obtained thanks to the smaller unsatisfiable cores that the algorithm has to analyze. Indeed, we provide empirical evidence that often the running time of our core based algorithm sensibly decreases when core shrinking is performed (Section 4). The advantage of introducing our strategy for core shrinking is also confirmed by a comparison with CLASP (Gebser *et al.* 2015a): even if our solver, WASP (Alviano *et al.* 2015a), is in general slower than CLASP at completing stable model searches, its performance is sufficiently improved by core shrinking that the two solvers are almost on par in terms of solved instances, with the crucial difference that WASP provides both overestimates and underestimates during the computation, while ones or the others are produced by CLASP only after running to completion.

## 2 Background

Let  $\mathcal{A}$  be a set of (propositional) *atoms* comprising  $\perp$ . A *literal* is an atom  $p$  possibly preceded by zero or more occurrences of the *default negation* symbol  $\sim$ . A *sum* aggregate is of the form:

$$\text{SUM}[w_1 : \ell_1, \dots, w_n : \ell_n] \odot k \quad (1)$$

where  $n, k, w_1, \dots, w_n$  are nonnegative integers,  $\ell_1, \dots, \ell_n$  are literals, and  $\odot \in \{<, \leq, \geq, >, =, \neq\}$ . If  $w_1 = \dots = w_n = 1$ , (1) is also called *count* aggregate, and denoted as  $\text{COUNT}[\ell_1, \dots, \ell_n] \odot k$ . A *rule*  $r$  is an implication  $H(r) \leftarrow B(r)$ , where  $H(r)$  is a disjunction of atoms, and  $B(r)$  is a conjunction of literals and aggregates.  $H(r)$  and  $B(r)$  are called *head* and *body* of  $r$ , and abusing of notation also denote the sets of their elements. If  $H(r) \subseteq \{\perp\}$ , then  $r$  is called *integrity constraint*. A *program*  $\Pi$  is a set of rules. Let  $At(\Pi)$  denote the set of atoms occurring in  $\Pi$ .

An *interpretation*  $I$  is a set of atoms not containing  $\perp$ . Relation  $\models$  is inductively defined as follows: for  $p \in \mathcal{A}$ ,  $I \models p$  if  $p \in I$ ;  $I \models \sim \ell$  if  $I \not\models \ell$ ; for  $A$  of the form (1),  $I \models A$  if  $\sum_{i \in [1..n], I \models \ell_i} w_i \odot k$ ; for a rule  $r$ ,  $I \models B(r)$  if  $I \models \ell$  for all  $\ell \in B(r)$ , and  $I \models r$  if  $I \cap H(r) \neq \emptyset$  whenever  $I \models B(r)$ ; for a program  $\Pi$ ,  $I \models \Pi$  if  $I \models r$  for all  $r \in \Pi$ .  $I$  is a *model* of a literal, aggregate, rule, or program  $\pi$  if  $I \models \pi$ . Note that  $I \not\models \perp$ , and  $I \models \sim \perp$ , for any interpretation  $I$ ; let  $\top$  be a shortcut for  $\sim \perp$ .

The *reduct*  $\Pi^I$  of a program  $\Pi$  with respect to an interpretation  $I$  is obtained from  $\Pi$  as follows: (i) any rule  $r$  such that  $I \not\models B(r)$  is removed; (ii) any negated literal  $\ell$  such that  $I \not\models \ell$  is replaced by  $\perp$ ; (iii) any negated literal  $\ell$  such that  $I \models \ell$  is replaced by  $\top$ . An interpretation  $I$  is a *stable model* of a program  $\Pi$  if  $I \models \Pi$ , and there is no  $J \subset I$  such that  $J \models \Pi^I$ . Let  $SM(\Pi)$  denote the set of stable models of  $\Pi$ . A program  $\Pi$  is *coherent* if  $SM(\Pi) \neq \emptyset$ ; otherwise,  $\Pi$  is *incoherent*.

### Example 1

Let  $\Pi_1$  be the program comprising the following five rules:

$$a \vee c \leftarrow \sim b, \sim d \quad a \leftarrow \sim b, c \quad c \leftarrow a, b \quad b \leftarrow a, c \quad d \leftarrow \sim \sim d$$

Its stable models are  $I_1 = \{a\}$  and  $I_2 = \{d\}$ , and the associated program reducts are the following:  $\Pi_1^{I_1} = \{a \vee c \leftarrow\}$  and  $\Pi_1^{I_2} = \{d \leftarrow\}$ , where  $\top$  in rule bodies is omitted for simplicity. ■

A *weak constraint*  $r$  is of the form  $w @ l \leftarrow B(r)$ , where  $B(r)$  is a conjunction of literals and aggregates, and  $w, l$  are positive integers denoted respectively *weight*( $r$ ) and *level*( $r$ ). For a multiset  $\mathcal{W}$  of weak constraints, and  $l \geq 1$ , let  $\mathcal{W}^l$  denote the multiset of weak constraints in  $\mathcal{W}$  whose level is  $l$ . The *cost* of an interpretation  $I$  for  $\mathcal{W}^l$  is  $\mathcal{W}^l(I) := \sum_{r \in \mathcal{W}^l, I \models B(r)} \text{weight}(r)$ . For any pair  $I, J$  of interpretations,  $J <_{\mathcal{W}} I$  if there is  $l \geq 1$  such that both  $\mathcal{W}^l(J) < \mathcal{W}^l(I)$ , and  $\mathcal{W}^{l+i}(J) \leq \mathcal{W}^{l+i}(I)$  for all  $i \geq 1$ .  $I$  is an *optimum stable model* of a program  $\Pi$  with respect to a multiset  $\mathcal{W}$  of weak constraints if  $I \in SM(\Pi)$ , and there is no  $J \in SM(\Pi)$  such that  $J <_{\mathcal{W}} I$ . Let  $OSM(\Pi, \mathcal{W})$  denote the set of optimum stable models of  $\Pi$  with respect to  $\mathcal{W}$ .

### Example 2

Continuing with Example 1, let  $\mathcal{W}_1$  be  $\{1 @ 2 \leftarrow d, 2 @ 1 \leftarrow a, 2 @ 1 \leftarrow b, 1 @ 1 \leftarrow c\}$ . Hence,  $OSM(\Pi_1, \mathcal{W}_1)$  only contains  $I_1$ , with  $\mathcal{W}_1^2(I_1) = 0$  and  $\mathcal{W}_1^1(I_1) = 2$ . Indeed,  $I_2$  is such that  $\mathcal{W}_1^2(I_2) = 1$  and  $\mathcal{W}_1^1(I_2) = 0$ , and therefore it is discarded because  $I_1 <_{\mathcal{W}_1} I_2$ . ■

## 3 Optimum stable model search

The computational problem analyzed in this paper is referred to as *optimum stable model search*: Given a (coherent) program  $\Pi$  and a multiset of weak constraints  $\mathcal{W}$ , compute an optimum stable model  $I^* \in OSM(\Pi, \mathcal{W})$ . Currently available algorithms for this problem are either inefficient, as linear search sat-unsat (Section 3.1), or not anytime, as algorithms based on unsatisfiable core analysis; ONE is among the core-based algorithms originally introduced for MaxSAT (Alviano *et al.* 2015c), and is adapted to ASP in Section 3.2. A concrete strategy to obtain anytime variants of

---

**Algorithm 1:** Linear search sat-unsat

---

**Input** : A coherent program  $\Pi$ , and a nonempty multiset of weak constraints  $\mathcal{W}$ .

**Output**: An optimum stable model  $I^* \in OSM(\Pi, \mathcal{W})$ .

```

1  $V := At(\Pi)$ ; // visible atoms
2 while  $\mathcal{W} \neq \emptyset$  do
3    $l := \max\{i \in \mathbb{N}^+ \mid \mathcal{W}^i \neq \emptyset\}$ ;  $\Pi := \Pi \cup relax(\mathcal{W}^l)$ ;
    $ub := 1 + \sum_{r \in \mathcal{W}^l} weight(r)$ ;
4   if  $I^*$  is defined then  $ub := \mathcal{W}^l(I^*)$ ;
5    $(res, I) := solve(\Pi \cup \{\perp \leftarrow \text{SUM}[(weight(r) : \sim s_r) \mid s_r \in soft(\mathcal{W}^l)] \geq ub\})$ ;
6   if  $res$  is COHERENT then  $I^* := I \cap V$ ;  $ub := \mathcal{W}^l(I)$ ; goto 5;
7   ;  $\Pi := \Pi \cup \{\perp \leftarrow \text{SUM}[(weight(r) : \sim s_r) \mid s_r \in soft(\mathcal{W}^l)] \neq ub\}$ ;
    $\mathcal{W} := \mathcal{W} \setminus \mathcal{W}^l$ ;
8 return  $I^*$ ;

```

---

core-based algorithms is then presented in Section 3.3; it consists in a progression (or linear) search to shrink unsatisfiable cores that may also discover better and better stable models.

**3.1 Inefficiencies in linear search sat-unsat**

An immediate algorithm for addressing this problem is known as *linear search sat-unsat*: a first stable model is searched, and new stable models of improved cost are iteratively computed until an incoherence arises, proving optimality of the last computed model. To ease the definition of such an algorithm, weak constraints are transformed into *relaxed* integrity constraints by introducing fresh literals called *soft literals*: whenever the body of a weak constraint is true, the associated soft literal is inferred false so to satisfy the relaxed integrity constraint. Formally, for a weak constraint  $r$ , its relaxed form is  $\perp \leftarrow B(r) \wedge s_r$ , where  $s_r$  is the soft literal associated with  $r$ , and defined by a *choice rule*  $s_r \leftarrow \sim \sim s_r$ . For a multiset of weak constraints  $\mathcal{W}$ , let  $relax(\mathcal{W})$  be the set of relaxed integrity constraints and choice rules obtained from the weak constraints in  $\mathcal{W}$ , and  $soft(\mathcal{W})$  be the set of associated soft literals.

*Example 3*

Continuing with Example 2,  $relax(\mathcal{W}_1^2)$  is  $\{\perp \leftarrow d, s_1\} \cup \{s_1 \leftarrow \sim \sim s_1\}$ , where  $s_1$  is the soft literal in  $soft(\mathcal{W}_1^2)$ , and  $relax(\mathcal{W}_1^1)$  is  $\{\perp \leftarrow a, s_2; \perp \leftarrow b, s_3; \perp \leftarrow c, s_4\} \cup \{s_i \leftarrow \sim \sim s_i \mid i \in [2..4]\}$ , where  $s_2, s_3, s_4$  are the soft literals in  $soft(\mathcal{W}_1^1)$ . ■

Linear search sat-unsat is reported as Algorithm 1. Levels are processed one at a time, from the greatest to the smallest, by relaxing weak constraints (line 3) and iteratively searching for stable models improving the current upper bound (lines 5–6). When the upper bound cannot be further improved, an integrity constraint is added to the program in order to discard all stable models of cost different from the upper bound (line 7), so that the next level can be processed correctly.

*Example 4*

Continuing with the previous examples, Algorithm 1 with input  $\Pi_1$  and  $\mathcal{W}_1$  would relax weak constraints in  $\mathcal{W}_1^2$ , and set  $ub$  to 2. After that, a stable model of the following program:

$$\begin{array}{l} a \vee c \leftarrow \sim b, \sim d \quad a \leftarrow \sim b, c \quad c \leftarrow a, b \quad b \leftarrow a, c \quad d \leftarrow \sim \sim d \\ \perp \leftarrow d, s_1 \quad s_1 \leftarrow \sim \sim s_1 \quad \perp \leftarrow \text{SUM}[1 : \sim s_1] \geq 2 \end{array}$$

is searched, and  $I_2 = \{d\}$  may be found. In this case,  $ub$  is decreased to 1, and a stable model satisfying the new constraint  $\perp \leftarrow \text{SUM}[1 : \sim s_1] \geq 1$  is searched. Eventually,  $I'_1 = \{a, s_1\}$  is found, and  $ub$  is set to 0. The next stable model search trivially fails because of  $\perp \leftarrow \text{SUM}[1 : \sim s_1] \geq 0$ , and therefore the cost for level 2 is fixed to 0 by adding  $\perp \leftarrow \text{SUM}[1 : \sim s_1] \neq 0$ . Weak constraints of level 1 are relaxed, and  $ub$  is set to  $\mathcal{W}_1^1(I_1) = 2$ . A stable model of the following program:

$$\begin{array}{l} a \vee c \leftarrow \sim b, \sim d \quad a \leftarrow \sim b, c \quad c \leftarrow a, b \quad b \leftarrow a, c \quad d \leftarrow \sim \sim d \\ \perp \leftarrow d, s_1 \quad s_i \leftarrow \sim \sim s_i \quad (\forall i \in [1..4]) \quad \perp \leftarrow \text{SUM}[1 : \sim s_1] \neq 0 \\ \perp \leftarrow a, s_2 \quad \perp \leftarrow b, s_3 \quad \perp \leftarrow c, s_4 \quad \perp \leftarrow \text{SUM}[2 : \sim s_2, 2 : \sim s_3, 1 : \sim s_4] \geq 2 \end{array}$$

is searched, but none is found. The algorithm terminates returning  $I'_1 \cap \{a, b, c, d\} = \{a\} = I_1$ . ■

The main drawback of linear search sat-unsat is that in practice it is quite inefficient. In fact, the sum aggregate introduced in order to improve the current upper bound (line 4 of Algorithm 1) allows the solver to discard several stable models, but often does not provide sufficiently strong evidences to help the refutation process. It is because of this fact that linear search sat-unsat is often unable to prove optimality of a stable model, even if it is provided externally. In fact, an integrity constraint of the form  $\perp \leftarrow \text{SUM}[(w_i : \ell_i) \mid 1 \leq i \leq n] \geq k$  is satisfied by exponentially many interpretations. Even when all weights are 1, the number of satisfying interpretations is in the order of  $\binom{n}{k-1}$ . It is a prohibitive number even for small values of  $k$  when  $n$  is a large integer, as it usually happens for Algorithm 1, where  $n$  is essentially the number of weak constraints.

### 3.2 Almost silent unsatisfiable core analysis

The idea underlying algorithms based on unsatisfiable core analysis is the following: a set of soft literals that cannot be jointly satisfied is identified, and the reason of unsatisfiability is removed by adjusting the program and the soft literals; the process is repeated until a stable model is found, which is also guaranteed to be optimum. To ease the definition of such an algorithm, the notion of unsatisfiable core is given in terms of *assumptions*: modern ASP solvers accept as input a set  $S$  of atoms, called assumptions, in addition to the usual logic program  $\Pi$ , and return a stable model  $I$  of  $\Pi$  such that  $S \subseteq I$ , if it exists; otherwise, they return a set  $C \subseteq S$  such that  $\Pi \cup \{\perp \leftarrow \sim p \mid p \in C\}$  is incoherent, which is called *unsatisfiable core*.

---

**Algorithm 2:** Unsatisfiable core analysis with ONE

---

**Input** : A coherent program  $\Pi$ , and a nonempty multiset of weak constraints  $\mathcal{W}$ .

**Output**: An optimum stable model  $I^* \in OSM(\Pi, \mathcal{W})$ .

```

1  $V := At(\Pi)$ ;   for  $p \in \mathcal{A}$  do  $w(p) := 0$ ;
2 while  $\mathcal{W} \neq \emptyset$  do
3    $l := \max\{i \in \mathbb{N}^+ \mid \mathcal{W}^i \neq \emptyset\}$ ;  $\Pi := \Pi \cup relax(\mathcal{W}^l)$ ;  $lb := 0$ ;  $ub := \infty$ ;
    $stratum := \infty$ ;
4   for  $s_r \in soft(\mathcal{W}^l)$  do  $w(s_r) := weight(r)$ ;           // soft literals
   initialization
5   if  $I^*$  is defined then  $ub := \mathcal{W}^l(I^*)$ ; HARDENING( $\Pi, w, lb, ub$ );
6    $stratum := \max_{p \in \mathcal{A}, w(p) < stratum} w(p)$ ;           // next stratum
7    $(res, I, C) := solve(\Pi, \{p \in \mathcal{A} \mid w(p) \geq stratum\})$ ;
8   if  $res$  is INCOHERENT then
16     Let  $C$  be  $\{p_0, \dots, p_n\}$  (for some  $n \geq 0$ ), and  $s_1, \dots, s_n$  be fresh atoms;
17      $lb := lb + stratum$ ;
18     for  $i \in [0..n]$  do  $w(p_i) := w(p_i) - stratum$ ;           // remainders
19     for  $i \in [1..n]$  do  $w(s_i) := stratum$ ;           // new soft literals
20      $\Pi := \Pi \cup \{s_i \leftarrow \sim s_i \mid i \in [1..n]\} \cup \{\perp \leftarrow s_i, \sim s_{i+1} \mid i \in [1..n-1]\}$ 
        $\cup \{\perp \leftarrow COUNT[p_0, \dots, p_n, \sim s_1, \dots, \sim s_n] < n\}$ ;           // relax core
21     HARDENING( $\Pi, w, lb, ub$ ); goto 7;
22   if  $\mathcal{W}^l(I) < ub$  then  $I^* := I \cap V$ ;  $ub := \mathcal{W}^l(I)$ ; HARDENING( $\Pi, w, lb, ub$ );
23   if  $\exists p \in \mathcal{A}$  such that  $1 \leq w(p) < stratum$  then goto 6;
24    $\mathcal{W} := \mathcal{W} \setminus \mathcal{W}^l$ ;
25 return  $I^*$ ;

```

---



---

**Procedure** HARDENING( $\Pi, w, lb, ub$ )

---

1 **for**  $p \in \mathcal{A}$  **such that**  $lb + w(p) > ub$  **do**  $\Pi := \Pi \cup \{\perp \leftarrow \sim p\}$ ;  $w(s_r) := 0$ ;

---

*Example 5*

Consider program  $\Pi_1$  from Example 1 with the addition of  $relax(\mathcal{W}_1^2)$  and  $relax(\mathcal{W}_1^1)$ :

$$\begin{aligned}
 & a \vee c \leftarrow \sim b, \sim d \quad a \leftarrow \sim b, c \quad c \leftarrow a, b \quad b \leftarrow a, c \quad d \leftarrow \sim \sim d \\
 & \perp \leftarrow d, s_1 \quad \perp \leftarrow a, s_2 \quad \perp \leftarrow b, s_3 \quad \perp \leftarrow c, s_4 \quad s_i \leftarrow \sim \sim s_i \quad (\forall i \in [1..4])
 \end{aligned}$$

If  $\{s_1, \dots, s_4\}$  is the set of assumptions, the unsatisfiable cores are  $\{s_1, s_2\}$ , and its supersets. ■

The algorithm presented in this paper is ONE, reported as Algorithm 2 (lines 9–15 will be *injected* later to shrink unsatisfiable cores). Every atom  $p$  is associated with a weight  $w(p)$ , initially set to zero (line 1), and levels are processed from the greatest to the smallest by relaxing constraints (line 3). Note that soft literals are associated with nonzero weights (line 4), and are processed per stratum, i.e., the greatest weights are processed first (line 6). A stable model containing all soft literals in the current stratum is then searched (line 7). If an unsatisfiable core  $\{p_0, \dots, p_n\}$  is returned

(lines 8 and 16), since at least one of  $p_0, \dots, p_n$  must be false in any optimum stable model, the lower bound is increased by the minimum weight among  $w(p_0), \dots, w(p_n)$  (or equivalently by *stratum*; line 17). Such a quantity is removed from  $p_0, \dots, p_n$  (line 18), and assigned to  $n$  new soft literals  $s_1, \dots, s_n$  (line 19). The new soft literals, and the following constraint (line 20):

$$\perp \leftarrow \text{COUNT}[p_0, \dots, p_n, \sim s_1, \dots, \sim s_n] < n$$

enforce the next call to function *solve* to search for a stable model satisfying at least  $n$  literals among  $p_0, \dots, p_n$ . Moreover, note that symmetry breakers of the form  $\perp \leftarrow s_i, \sim s_{i+1}$  are also added to  $\Pi$ , so that  $s_i$  is true if and only if at least  $n - i + 1$  literals among  $p_0, \dots, p_n$  are true.

The current stratum is then processed again (line 21), until a stable model is found. In this case, the upper bound is possibly improved (line 22), and the stratum is extended to soft literals of smaller weight (line 23), if there are. Otherwise, the processed stratum covers all soft literals, and weak constraints of the current level are removed so that the next level can be considered (line 24). Note that the algorithm also includes the HARDENING procedure, which essentially enforces truth of soft literals that are guaranteed to belong to all optimum stable models.

#### Example 6

Initially, the weak constraint in  $\mathcal{W}_1^2$  is relaxed, and the soft literal  $s_1$  is associated with weight  $w(s_1) = 1$ . A stable model for the assumption  $\{s_1\}$  and the following program:

$$\begin{aligned} a \vee c \leftarrow \sim b, \sim d \quad a \leftarrow \sim b, c \quad c \leftarrow a, b \quad b \leftarrow a, c \quad d \leftarrow \sim \sim d \\ \perp \leftarrow d, s_1 \quad s_1 \leftarrow \sim \sim s_1 \end{aligned}$$

is searched, and  $I_1' = \{a, s_1\}$  is found. Hence,  $ub$  is set to 0, and  $s_1$  is hardened by adding  $\perp \leftarrow \sim s_1$ . Weak constraints in  $\mathcal{W}_1^1$  are relaxed,  $w$  is such that  $w(s_2) = 2$ ,  $w(s_3) = 2$  and  $w(s_4) = 1$ , and  $ub$  is set to  $\mathcal{W}_1^1(I_1) = 2$ . A stable model for the assumptions  $\{s_2, s_3\}$  and the following program:

$$\begin{aligned} a \vee c \leftarrow \sim b, \sim d \quad a \leftarrow \sim b, c \quad c \leftarrow a, b \quad b \leftarrow a, c \quad d \leftarrow \sim \sim d \quad \perp \leftarrow \sim s_1 \\ \perp \leftarrow d, s_1 \quad \perp \leftarrow a, s_2 \quad \perp \leftarrow b, s_3 \quad \perp \leftarrow c, s_4 \quad s_i \leftarrow \sim \sim s_i \quad (\forall i \in [1..4]) \end{aligned}$$

is searched, and an unsatisfiable core, say  $\{s_2, s_3\}$ , is returned. The lower bound  $lb$  is set to 2, the program is extended with the following rules:

$$s_5 \leftarrow \sim \sim s_5 \quad \perp \leftarrow \text{COUNT}[s_2, s_3, \sim s_5] < 1$$

and function  $w$  is now such that  $w(s_2) = 0$ ,  $w(s_3) = 0$ ,  $w(s_4) = 1$  and  $w(s_5) = 2$ . Since  $lb = ub$ , all soft literals are hardened. The algorithm terminates returning  $I_1' \cap \{a, b, c, d\} = \{a\} = I_1$ . ■

The algorithm described in this section is almost silent, as it essentially runs to completion without printing any suboptimal stable models but those computed when changing stratum or level. This fact is actually limiting the use of core-based algorithms in ASP, as common instances usually have one or few levels, and also weights are often uniform. In particular, when all weak constraints have the same level and the same weight, the algorithm is completely silent.



---

**Algorithm 3:** Unsatisfiable core shrinking with reiterated progression

---

```

9  $m := -1; pr := 1;$ 
10 Let  $C$  be  $\{p_0, \dots, p_n\}$  (for some  $n \geq 0$ );
11  $(res, I, C') := solve\_with\_budget(\Pi, \{p_i \mid i \in [0..m + pr]\});$ 
12 if  $res$  is INCOHERENT then  $C := C'$ ; // smaller core found
13 if  $res$  is COHERENT and  $\mathcal{W}^l(I) < ub$  then  $I^* := I \cap V; ub := \mathcal{W}^l(I)$ 
14 if  $m + 2 \cdot pr \geq |C| - 1$  then  $m := m + pr; pr := 1/2;$  // reiterate
    progression
15 if  $m + 2 \cdot pr < |C| - 1$  then  $pr := 2 \cdot pr; \text{goto } 10;$  // increase progression

```

---

*Example 7*

Consider now a multiset  $\mathcal{W}_2$  of weak constraints obtained from  $\mathcal{W}_1$  by setting all weights and levels to 1. Algorithm 2 starts with assumptions  $\{s_1, \dots, s_4\}$  and the following program:

$$\begin{aligned}
 & a \vee c \leftarrow \sim b, \sim d \quad a \leftarrow \sim b, c \quad c \leftarrow a, b \quad b \leftarrow a, c \quad d \leftarrow \sim \sim d \\
 & \perp \leftarrow d, s_1 \quad \perp \leftarrow a, s_2 \quad \perp \leftarrow b, s_3 \quad \perp \leftarrow c, s_4 \quad s_i \leftarrow \sim \sim s_i \quad (\forall i \in [1..4])
 \end{aligned}$$

Hence, an unsatisfiable core, say  $\{s_1, \dots, s_4\}$ , is returned, and the following rules are added:

$$\begin{aligned}
 & s_i \leftarrow \sim \sim s_i \quad (\forall i \in [5..7]) \quad \perp \leftarrow s_5, \sim s_6 \quad \perp \leftarrow s_6, \sim s_7 \\
 & \perp \leftarrow \text{COUNT}[s_1, s_2, s_3, s_4, \sim s_5, \sim s_6, \sim s_7] < 3
 \end{aligned}$$

where  $s_5, s_6, s_7$  are the new soft literals. The assumptions are now  $\{s_5, s_6, s_7\}$ , and either  $I''_1 = \{a, s_1, s_3, s_4, s_5, s_6, s_7\}$  or  $I''_2 = \{d, s_2, s_3, s_4, s_5, s_6, s_7\}$  is found. The algorithm terminates because  $lb = ub = \mathcal{W}^l_2(I''_1) = \mathcal{W}^l_2(I''_2) = 1$ . Note that no stable model was found before running to completion, and in general Algorithm 2 has to analyze several unsatisfiable cores to terminate. ■

### 3.3 Unsatisfiable core shrinking and upper bounds

Unsatisfiable cores returned by function *solve* are not subset minimal in general. The non-minimality of the unsatisfiable core is justified both theoretically and practically: linearly many coherence checks are required in general to verify the minimality of an unsatisfiable core, hence giving a  $\Delta^P_3$ -complete problem; on the other hand, extracting an unsatisfiable core after a stable model search failure is quite easy and usually implemented by identifying the assumptions involved in the refutation. The non minimality of the analyzed unsatisfiable cores may affect negatively the performance of subsequent calls to function *solve* due to aggregation over large sets. However, it also gives an opportunity to improve Algorithm 2: the size of unsatisfiable cores can be reduced by performing a few stable model searches within a given budget on the running time. In more detail, Algorithm 3 is injected between lines 8 and 16 of Algorithm 2. It implements a progression search in the unsatisfiable core  $\{p_0, \dots, p_n\}$ : the size of the assumptions passed to function *solve\_with\_budget* is doubled at each call (line 15), and the progression is reiterated when all assumptions are covered (line 14). If function *solve\_with\_budget* terminates

within the given budget, it either returns a smaller unsatisfiable core (line 12), or a stable model that possibly improves the current upper bound (line 13).

#### Example 8

Consider  $\mathcal{W}_2$  from Example 7, and the unsatisfiable core  $\{s_1, \dots, s_4\}$  returned after the first call to function *solve*. The shrinking process searches a stable model with assumption  $\{s_1\}$ , and  $\{a, s_1\} \cup X$  (for some  $X \subseteq \{s_3, s_4\}$ ) may be found within the allotted budget. In any case, a stable model satisfying the assumptions  $\{s_1, s_2\}$  is searched, and the unsatisfiable core  $\{s_1, s_2\}$  may be returned if the budget is sufficient. Otherwise, the progression is reiterated, and one more soft literal is added to the assumptions. Hence,  $\{s_1, s_2, s_3\}$  may be returned as an unsatisfiable core, or the budget may be insufficient and the original unsatisfiable core will be processed. ■

As an alternative, the shrinking procedure reported in Algorithm 3 can be modified as follows: variable *pr* is not doubled in line 15, but instead it is incremented by one, i.e.,  $pr := pr + 1$ . The resulting procedure is called *linear based shrinking*. For unsatisfiable cores of size 4 or smaller, as those considered in Example 8, the two shrinking procedures coincide, while in general linear based shrinking performs more stable model searches.

#### Theorem 1

Let  $C$  be an unsatisfiable core. Function *solve\_with\_budget* is invoked  $O((\log |C|)^2)$  times by progression based shrinking, and  $O(|C|)$  times by linear based shrinking.

#### Proof

The worst case occurs when *solve\_with\_budget* never returns *COHERENT*, either because the tested set of assumptions is not an unsatisfiable core, or because the allotted budget is insufficient. This is the case, for example, if  $C$  is already minimal. For progression based shrinking, line 14 of Algorithm 3 is first executed after  $k := \lceil \log |C| \rceil$  executions of lines 11 and 15. After that, and in the worst case, the process is repeated on half of the literals in  $C$ , for  $k - 1$  executions of lines 11 and 15. Hence, for a total of  $k(k + 1)/2$  executions of function *solve\_with\_budget*. For linear based shrinking, instead, line 14 of Algorithm 3 is first executed after  $|C| - 1$  executions of lines 11 and 15. After that, the process terminates because  $m \geq |C| - 1$ . □

## 4 Implementation and experiments

Algorithm ONE (Alviano *et al.* 2015c) has been implemented in WASP (Alviano *et al.* 2015a), an ASP solver also supporting, among other algorithms, linear search sat-unsat. The implementation of ONE optionally includes the two shrinking procedures described in Section 3.3, so that both underestimates and overestimates can be produced by WASP in any case, weighted or unweighted. Currently, the time budget of function *solve\_with\_budget* is fixed to 10 seconds, but the architecture of WASP can easily accommodate alternative options, such as a budget proportional to the time required to find the unsatisfiable core to be shrank.

WASP also implements *disjoint cores analysis*, which is essentially a preliminary step where only soft literals associated with weak constraints in the input are passed as assumptions to function *solve*, while new soft literals introduced by the analysis of detected cores are temporarily ignored. Disjoint cores analysis terminates with the detection of a stable model, and after that algorithm ONE is run as usual not distinguishing between initial and new soft literals.

In order to assess empirically the impact of these shrinking procedures to the performance of WASP, instances (with weak constraints) from the ASP Competition 2015 (Gebser *et al.* 2015b) were considered. Moreover, WASP was also compared with CLASP (Gebser *et al.* 2015a), which implements linear search sat-unsat (strategy *bb, 1*) and OLL (strategy *usc; usc, 1* with disjoint cores analysis; Andres *et al.* 2012), a core based algorithm that inspired the definition of ONE. The experiments were run on an Intel Xeon 2.4 GHz with 16 GB of memory, and time and memory were limited to 20 minutes and 15 GB, respectively.

An overview of the obtained results is given in Table 1, where the number of solved instances is reported. As a first comment, the fact that CLASP is in general faster than WASP to complete stable model searches is confirmed by comparing the performance of the two solvers running linear search sat-unsat (LINSU in the table; CLASP solves 29 instances more than WASP) or the core based algorithms (difference of 12 instances, reduced to 9 if disjoint cores are computed). This gap is completely filled by adding the shrinking procedures if disjoint cores are computed, and significantly reduced otherwise: 6 instances with linear based shrinking (LSHR in the table), and 3 instances with reiterated progression (PSHR in the table). It is also important to observe that the performance of WASP is boosted by disjoint cores only for instances of Video Streaming, where there are many strata and few unsatisfiable cores, which is the worst case for stratification.

Table 1 also shows the number of wins of each tested strategy, where a strategy wins on a tested instance if it terminates in the allotted time, or it finds the smallest upper bound and none of the tested strategies terminated. Within this respect, the linear search sat-unsat algorithm implemented by CLASP has a significant performance improvement, in particular for instances of Steiner Tree and Video Streaming. However, even according to this measure, a better performance is obtained by CLASP using OLL and disjoint cores analysis, which is also reached by WASP thanks to the core shrinking procedures introduced in this paper.

Concerning the average execution time of the tested algorithms, a cactus plot is reported in Fig. 1 (left; WASP using LINSU is not reported because it solves less than 50 instances). The graph highlights that core based algorithms are faster than linear search sat-unsat in more testcases. Moreover, and more important, the addition of core shrinking does not add overhead to WASP. This aspect is even more clear in the scatter plots on the right, where the impact of PSHR is shown instance-by-instance. It can be observed that the majority of points are below the diagonal, meaning that often core shrinking provides a performance gain, and only in a few cases it introduces overhead. The main reason for this performance improvement is that shrinking a core often implies that subsequently found unsatisfiable cores are smaller: The cumulative number of literals in the analyzed cores is reduced by at

Table 1. Solved instances, and wins in parentheses (180 testcases; 20 for each problem).

Problem	CLASP			WASP				WASP+DISJ		
	LINSU	OLL	OLL+DISJ	LINSU	ONE	LSHR	PSHR	ONE	LSHR	PSHR
ADF	20 (20)	19 (19)	19 (19)	18 (18)	14 (14)	17 (17)	18 (18)	14 (14)	18 (18)	18 (18)
ConnectedStillLife	7 (10)	11 (11)	13 (13)	0 (0)	13 (13)	10 (10)	10 (12)	13 (13)	9 (11)	10 (10)
CrossingMin	7 (7)	19 (19)	19 (19)	2 (2)	19 (19)	19 (19)	19 (19)	19 (19)	19 (20)	19 (19)
MaxClique	0 (4)	16 (16)	16 (16)	0 (0)	10 (10)	14 (14)	16 (16)	11 (11)	15 (15)	16 (16)
MaxSAT	7 (7)	15 (15)	15 (15)	5 (5)	19 (19)	20 (20)	20 (20)	18 (18)	20 (20)	20 (20)
SteinerTree	3 (16)	1 (1)	1 (1)	1 (1)	1 (5)	1 (3)	1 (3)	0 (1)	0 (0)	0 (1)
SystemSynthesis	0 (1)	4 (11)	4 (11)	0 (4)	0 (2)	0 (2)	0 (2)	0 (2)	0 (2)	0 (2)
ValvesLocation	16 (17)	9 (9)	11 (11)	11 (11)	15 (17)	16 (19)	16 (19)	14 (14)	16 (16)	16 (16)
VideoStreaming	14 (19)	9 (9)	9 (9)	8 (8)	0 (0)	0 (0)	0 (0)	9 (10)	9 (9)	9 (9)
<b>Total</b>	74 (101)	103 (110)	107 (114)	45 (49)	91 (99)	97 (104)	100 (109)	98 (102)	106 (111)	108 (111)

Table 2. Number of solved instances within a given error estimation (140 testcases).

$\epsilon(ub, lb)$		WASP			WASP+DISJ			CLASP + best $lb$ by WASP	
		ONE	LSHR	PSHR	ONE	LSHR	PSHR	LINSU	OLL+DISJ
$\leq$	0.00% (0)	77	80	82	84	88	90	77	84
$\leq$	6.25% (1/16)	77	87	88	86	95	94	90	87
$\leq$	12.50% (1/8)	77	91	93	86	101	97	96	88
$\leq$	25.00% (1/4)	77	97	95	92	105	103	99	99
$\leq$	50.00% (1/2)	78	97	97	102	105	105	99	101
$\leq$	100.00% (1)	78	97	97	104	105	105	107	105

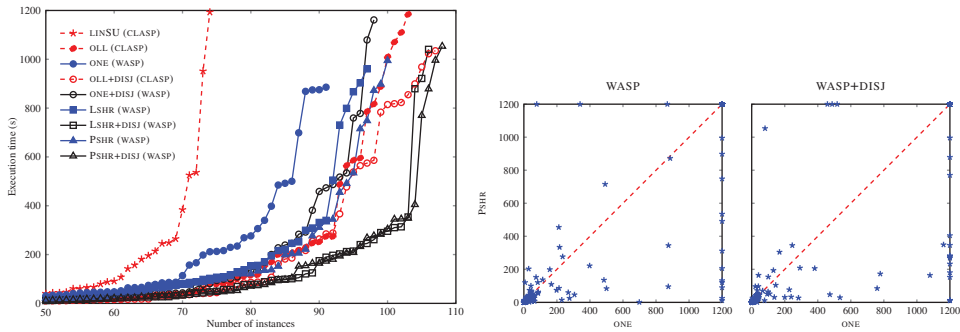


Fig. 1. Solved instances within a bound on the running time (cactus plot on the left), and instance-by-instance comparison of ONE with and without PSHR (scatter plots on the right).

least 68% when shrinking is performed (excluding Steiner Tree, System Synthesis and Video Streaming, for which WASP found few unsatisfiable cores). The budget is reached at least once in each problem, and often no more than 2 times, with a peak of 20–25 times on average for instances of Max Clique and Still Life.

Another advantage of unsatisfiable core shrinking is that better and better stable models are possibly discovered while computing an optimum stable model. In order to measure the impact of our strategies within this respect, let us define the *estimate error*  $\epsilon$  of the last stable model produced by Algorithm 2 as follows:

$$\epsilon(ub, lb) := \begin{cases} \frac{ub-lb}{lb} & \text{if } ub \neq \infty \text{ and } lb \neq 0; \\ \infty & \text{if } ub = \infty, \text{ or both } ub \neq 0 \text{ and } lb = 0; \\ 0 & \text{if } ub = lb = 0. \end{cases}$$

Hence, the cost associated with the stable model returned by Algorithm 2 is at most  $\epsilon(ub, lb)$  times greater than the cost of an optimum stable model. Such a measure is applicable to testcases comprising weak constraints of the same level, which is the case for all problems in our benchmark but ADF and System Synthesis.

Table 2 reports the number of instances for which WASP produced a stable model within a given error estimate. In particular, the first row shows the number of instances for which an optimum stable model was computed (error estimate is 0). The last row, instead, shows the number of instances solved with error estimate bounded by 1, and smaller values for the error estimate are considered in the

intermediate rows. It is interesting to observe that without shrinking and disjoint cores analysis, the number of solved instances only increases by 1, confirming the fact that stratification alone is not sufficient to produce suboptimal stable models of reasonable cost. Better results are instead obtained by adding core shrinking, which gives an increase of up to 17 instances, many of which with an error estimate bounded by 6.25%. Another interesting observation concerns disjoint cores analysis. The stable model produced after the analysis of all disjoint cores is already sufficient to obtain an error estimate bounded by 100% for many tested instances. However, many of these stable models have an error estimate greater than 25%. Also in this case, adding core shrinking leads to better results.

For the sake of completeness, also CLASP is included in Table 2. However, since CLASP does not print any lower bound, the best value for *lb* produced by WASP is combined with the upper bounds given by CLASP running LINSU and OLL. If an error estimate of 100% is acceptable, then the number of stable models produced by CLASP is aligned with WASP, or even better. However, when the error estimate must be less or equal than 50%, the combination of disjoint cores analysis and core shrinking implemented by WASP leads to better results in this benchmark.

## 5 Related work

Weak constraints are the analogous of soft clauses in MaxSAT (Cha *et al.* 1997), which was also considered as target language by some ASP solver (Liu *et al.* 2012). The main difference is that weak constraints are associated with levels and weights, while soft clauses can be only associated with weights. However, levels can be simulated by properly modifying weights as follows:

1. Let  $s := 1 + \sum_{r \in \mathcal{W}^1} \text{weight}(r)$ , and let  $l$  be the smallest integer such that  $l \geq 2$  and  $\mathcal{W}^l \neq \emptyset$ .
2. For all weak constraints  $r$  in  $\mathcal{W}^l$ , set  $\text{weight}(r) := \text{weight}(r) \cdot s$  and  $\text{level}(r) := 1$ .
3. Repeat steps 1–2 while there are weak constraints of level different from 1.

Actually, the above process was reverted in the MaxSAT literature in order to take advantage from the “identification of levels” (Argelich *et al.* 2009). In fact, taking into account levels is crucial for some algorithms, as for example linear search sat-unsat (i.e., Algorithm 1).

### Example 9

Consider the following simple program  $\Pi_2$ , for any  $m \geq 1$  and any  $n \geq 1$ :

$$p_{ij} \leftarrow \sim \sim p_{ij} \quad 1 @ i \leftarrow \sim \sim p_{ij} \quad \forall i \in [1..m], \forall j \in [1..n]$$

In the worst case, Algorithm 1 has to find  $n$  stable models before understanding that the optimum cost of level  $m$  is 0. Similarly for the other levels, for a total of  $m \cdot n$  stable models in the worst case. If levels are ignored, then in the worst case the upper bound is improved by one after each stable model search, and therefore exponentially many stable models are computed. ■

The first algorithms based on unsatisfiable core analysis were introduced in MaxSAT (Fu and Malik 2006; Marques-Silva and Manquinho 2008; Marques-Silva and Planes 2008; Manquinho *et al.* 2009; Ansótegui *et al.* 2009), and subsequently ported to ASP in the experimental solver UNCLASP (Andres *et al.* 2012), where also OLL was presented. The main drawback of OLL is that it may add new constraints if soft literals introduced by the analysis of previous cores belong to subsequently detected unsatisfiable cores, and these new constraints only minimal differ from those already added to the input program. This drawback was later circumvented in the MaxSAT solver MSCG using a smart encoding based on sorting networks (Morgado *et al.* 2014).

Algorithm OLL was also implemented in the ASP solver WASP (Alviano *et al.* 2015b), and compared with other algorithms from MaxSAT literature such as PMRES (Narodytska and Bacchus 2014). These two algorithms, OLL and PMRES, are also the origin of two other algorithms for MaxSAT, namely ONE and  $\kappa$  (Alviano *et al.* 2015c). In fact, ONE is a simplification of OLL, and adds exactly one constraint for each analyzed core; whether new soft literals will be part of other unsatisfiable cores is irrelevant for ONE, as all the information required to complete the computation is already encoded in the added constraint.

Concerning  $\kappa$ , instead, it is a generalization of ONE based on the observation that aggregates involving a huge number of literals are handled inefficiently in the refutation process. This fact was first observed by Narodytska and Bacchus, who proposed PMRES: the inefficiency is circumvented because PMRES introduces constraints of size 3. The more general idea underlying  $\kappa$  is to bound the size of the new constraints to a given constant  $k$ . Hence,  $\kappa$  combines ONE and PMRES, which are obtained for special values of  $k$  (respectively, infinity and 3).

Finally, it is important to observe that current core based algorithms in ASP cannot continuously improve both lower and upper bound. Alternatives for obtaining such a behavior are based on combinations of different algorithms, either sequentially or in parallel (Alviano *et al.* 2015b; Gebser *et al.* 2015a). The strategy suggested in this paper is instead to improve core based algorithms by shrinking unsatisfiable cores. Within this respect, the proposed shrinking procedures are also original, and the closest approaches in MaxSAT iteratively remove one literal from the unsatisfiable core, either obtaining a smaller unsatisfiable core, or a necessary literal in the processed unsatisfiable core (Nadel 2010; Nadel *et al.* 2014).

## 6 Conclusion

The combination of ASP programs and weak constraints is important to ease the modeling of optimization problems. However, the computation of optimum stable models is often very hard, and suboptimal stable models may be the only affordable solutions in some cases. Despite that fact, the most efficient algorithms developed so far, which are based on unsatisfiable core analysis, produce few (or even no) stable models while running to completion.

A concrete strategy to improve current ASP solvers is presented in this paper: better and better stable models can be produced if unsatisfiable cores are shrunk.

The overhead due to the shrinking process is limited by introducing a budget on the running time, and eventually a performance gain is obtained thanks to the reduced size of the analyzed unsatisfiable cores. On the instances of the Sixth ASP Competition, our implementation is often able to provide (suboptimal) stable models with a guarantee of distance to the optimum cost of around 10%.

### Acknowledgement

This work was partially supported by the National Group for Scientific Computation (GNCS-INDAM), by the Italian Ministry of University and Research under PON project “Ba2Know (Business Analytics to Know) Service Innovation - LAB”, No. PON03PE\_00001\_1, and by the Italian Ministry of Economic Development under project “PIUCultura (Paradigmi Innovativi per l’Utilizzo della Cultura)” n. F/020016/01–02/X27.

### References

- ALVIANO, M., DODARO, C., FABER, W., LEONE, N. AND RICCA, F. 2013. WASP: A native ASP solver based on constraint learning. In P. Cabalar and T. C. Son (Eds.), *Logic Programming and Nonmonotonic Reasoning, 12th International Conference, LPNMR 2013, Corunna, Spain, September 15-19, 2013. Proceedings*, Volume 8148 of *Lecture Notes in Computer Science*, Springer, 54–66.
- ALVIANO, M., DODARO, C., LEONE, N. AND RICCA, F. 2015a. Advances in WASP. In F. Calimeri, G. Ianni and M. Truszczynski (Eds.), *Proceedings of Logic Programming and Nonmonotonic Reasoning - 13th International Conference, LPNMR 2015*, Volume 9345 of *Lecture Notes in Computer Science*, Springer, 40–54.
- ALVIANO, M., DODARO, C., MARQUES-SILVA, J. AND RICCA, F. 2015b. Optimum stable model search: algorithms and implementation. *Journal of Logic and Computation*.
- ALVIANO, M., DODARO, C. AND RICCA, F. 2014. Anytime Computation of Cautious Consequences in Answer Set Programming. *Theory and Practice of Logic Programming* 14, 4-5, 755–770.
- ALVIANO, M., DODARO, C. AND RICCA, F. 2015c. A MaxSAT Algorithm Using Cardinality Constraints of Bounded Size. In Q. Yang and M. Wooldridge (Eds.), *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015*, AAAI Press, 2677–2683.
- ALVIANO, M. AND FABER, W. 2013. The complexity boundary of answer set programming with generalized atoms under the FLP semantics. In P. Cabalar and T. C. Son (Eds.), *Logic Programming and Nonmonotonic Reasoning, 12th International Conference, LPNMR 2013, Corunna, Spain, September 15-19, 2013. Proceedings*, Volume 8148 of *Lecture Notes in Computer Science*, Springer, 67–72.
- ALVIANO, M., FABER, W. AND GEBSER, M. 2015d. Rewriting recursive aggregates in answer set programming: back to monotonicity. *Theory and Practice of Logic Programming* 15, 4-5, 559–573.
- ALVIANO, M., FABER, W., LEONE, N., PERRI, S., PFEIFER, G. AND TERRACINA, G. 2010. The disjunctive datalog system DLV. In O. de Moor, G. Gottlob, T. Furche and A. J. Sellers (Eds.), *Datalog Reloaded - First International Workshop, Datalog 2010, Oxford, UK, March 16-19, 2010. Revised Selected Papers*, Volume 6702 of *Lecture Notes in Computer Science*, Springer, 282–301.



- ALVIANO, M., FABER, W. AND STRASS, H. 2016. Boolean functions with ordered domains in answer set programming. In D. Schuurmans and M. P. Wellman (Eds.), *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, AAAI Press, 879–885.
- ALVIANO, M. AND LEONE, N. 2015. Complexity and compilation of gz-aggregates in answer set programming. *Theory and Practice of Logic Programming* 15, 4-5, 574–587.
- ANDRES, B., KAUFMANN, B., MATHEIS, O. AND SCHAUB, T. 2012. Unsatisfiability-based optimization in clasp. In A. Dovier and V. S. Costa (Eds.), *Technical Communications of the 28th International Conference on Logic Programming, ICLP 2012*, Volume 17 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 211–221.
- ANSÓTEGUL, C., BONET, M. L. AND LEVY, J. 2009. Solving (weighted) partial maxsat through satisfiability testing. In O. Kullmann (Ed.), *Proceedings of Theory and Applications of Satisfiability Testing, SAT 2009*, Volume 5584 of *Lecture Notes in Computer Science*, Springer, 427–440.
- ARGELICH, J., LYNCE, I. AND SILVA, J. P. M. 2009. On Solving Boolean Multilevel Optimization Problems. In C. Boutilier (Ed.), *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009*, 393–398.
- BARTHOLOMEW, M., LEE, J. AND MENG, Y. 2011. First-order semantics of aggregates in answer set programming via modified circumscription. In *Logical Formalizations of Commonsense Reasoning, Papers from the 2011 AAAI Spring Symposium, Technical Report SS-11-06, Stanford, California, USA, March 21-23, 2011*. AAAI.
- BLIEM, B., KAUFMANN, B., SCHAUB, T. AND WOLTRAN, S. 2016. ASP for Anytime Dynamic Programming on Tree Decompositions. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*. AAAI Press.
- BREWKA, G., EITER, T. AND TRUSZCZYNSKI, M. 2011. Answer set programming at a glance. *Commun. ACM* 54, 12, 92–103.
- BUCCAFURRI, F., LEONE, N. AND RULLO, P. 2000. Enhancing Disjunctive Datalog by Constraints. *IEEE Trans. Knowl. Data Eng.* 12, 5, 845–860.
- CHA, B., IWAMA, K., KAMBAYASHI, Y. AND MIYAZAKI, S. 1997. Local search algorithms for partial MAXSAT. In *Fourteenth National Conference on Artificial Intelligence and Ninth Innovative Applications of Artificial Intelligence Conference*, AAAI Press / The MIT Press, 263–268.
- DODARO, C., ALVIANO, M., FABER, W., LEONE, N., RICCA, F. AND SIRIANNI, M. 2011. The birth of a WASP: preliminary report on a new ASP solver. In F. Fioravanti (Ed.), *Proceedings of the 26th Italian Conference on Computational Logic, Pescara, Italy, August 31 - September 2, 2011*, Volume 810 of *CEUR Workshop Proceedings*, CEUR-WS.org, 99–113. .
- EITER, T., IANNI, G. AND KRENNWALLNER, T. 2009. Answer Set Programming: A Primer. In S. Tessaris, E. Franconi, T. Eiter, C. Gutierrez, S. Handschuh, M. Rousset and R. A. Schmidt (Eds.), *Reasoning Web. Semantic Technologies for Information Systems, 5th International Summer School 2009, Tutorial Lectures*, Volume 5689 of *Lecture Notes in Computer Science*, Springer, 40–110.
- FABER, W., PFEIFER, G. AND LEONE, N. 2011. Semantics and complexity of recursive aggregates in answer set programming. *Artif. Intell.* 175, 1, 278–298.
- FERRARIS, P. 2011. Logic programs with propositional connectives and aggregates. *ACM Trans. Comput. Log.* 12, 4, 25.
- FU, Z. AND MALIK, S. 2006. On Solving the Partial MAX-SAT Problem. In A. Biere and C. P. Gomes (Eds.), *Proceedings of Theory and Applications of Satisfiability Testing, SAT 2006*, Volume 4121 of *Lecture Notes in Computer Science*, Springer, 252–265.

- GEBSER, M., KAMINSKI, R., KAUFMANN, B., ROMERO, J., AND SCHAUB, T. 2015a. Progress in clasp Series 3. In F. Calimeri, G. Ianni and M. Truszczynski (Eds.), *LPNMR 2015*, Volume 9345 of *Lecture Notes in Computer Science*, Springer, 368–383.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2011a. Multi-Criteria Optimization in Answer Set Programming. In J. P. Gallagher and M. Gelfond (Eds.), *Technical Communications of the 27th International Conference on Logic Programming, ICLP 2011*, Volume 11 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 1–10.
- GEBSER, M., KAMINSKI, R. AND SCHAUB, T. 2011b. Complex optimization in answer set programming. *Theory and Practice of Logic Programming* 11, 4-5, 821–839.
- GEBSER, M., MARATEA, M. AND RICCA, F. 2015b. The Design of the Sixth Answer Set Programming Competition. In F. Calimeri, G. Ianni and M. Truszczynski (Eds.), *Proceedings of Logic Programming and Nonmonotonic Reasoning - 13th International Conference, LPNMR 2015*, Volume 9345 of *Lecture Notes in Computer Science*, Springer, 531–544.
- GELFOND, M. AND LIFSCHITZ, V. 1991. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Comput.* 9, 3/4, 365–386.
- GELFOND, M. AND ZHANG, Y. 2014. Vicious circle principle and logic programs with aggregates. *Theory and Practice of Logic Programming* 14, 4-5, 587–601.
- GIUNCHIGLIA, E., LIERLER, Y. AND MARATEA, M. 2006. Answer set programming based on propositional satisfiability. *J. Autom. Reasoning* 36, 4, 345–377.
- LIERLER, Y. AND MARATEA, M. 2004. Cmodels-2: Sat-based answer set solver enhanced to non-tight programs. In V. Lifschitz and I. Niemelä (Eds.), *Logic Programming and Nonmonotonic Reasoning, 7th International Conference, LPNMR 2004, Fort Lauderdale, FL, USA, January 6-8, 2004, Proceedings*, Volume 2923 of *Lecture Notes in Computer Science*, Springer, 346–350.
- LIFSCHITZ, V. 2008. What Is Answer Set Programming? In D. Fox and C. P. Gomes (Eds.), *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008*, AAAI Press, 1594–1597.
- LIU, G., JANHUNEN, T. AND NIEMELÄ, I. 2012. Answer Set Programming via Mixed Integer Programming. In G. Brewka, T. Eiter and S. A. McIlraith (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012*. AAAI Press.
- LIU, L., PONTELLI, E., SON, T. C. AND TRUSZCZYNSKI, M. 2010. Logic programs with abstract constraint atoms: The role of computations. *Artif. Intell.* 174, 3-4, 295–315.
- MANQUINHO, V. M., SILVA, J. P. M. AND PLANES, J. 2009. Algorithms for weighted boolean optimization. In O. Kullmann (Ed.), *Proceedings of Theory and Applications of Satisfiability Testing, SAT 2009*, Volume 5584 of *Lecture Notes in Computer Science*, Springer, 495–508.
- MAREK, V. W., NIEMELÄ, I. AND TRUSZCZYNSKI, M. 2008. Logic programs with monotone abstract constraint atoms. *Theory and Practice of Logic Programming* 8, 2, 167–199.
- MARQUES-SILVA, J. AND MANQUINHO, V. M. 2008. Towards more effective unsatisfiability-based maximum satisfiability algorithms. In H. K. Büning and X. Zhao (Eds.), *Proceedings of Theory and Applications of Satisfiability Testing, SAT 2008*, Volume 4996 of *Lecture Notes in Computer Science*, Springer, 225–230.
- MARQUES-SILVA, J. AND PLANES, J. 2008. Algorithms for Maximum Satisfiability using Unsatisfiable Cores. In *Design, Automation and Test in Europe, DATE 2008*, Munich, Germany, IEEE, 408–413.
- MORGADO, A., DODARO, C. AND MARQUES-SILVA, J. 2014. Core-Guided MaxSAT with Soft Cardinality Constraints. In *Proceedings of Principles and Practice of Constraint Programming - 20th International Conference, CP 2014*, Lyon, France, Springer, 564–573.

- NADEL, A. 2010. Boosting minimal unsatisfiable core extraction. In R. Bloem and N. Sharygina (Eds.), *Proceedings of 10th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2010*, IEEE, 221–229.
- NADEL, A., RYVCHIN, V. AND STRICHMAN, O. 2014. Accelerated Deletion-based Extraction of Minimal Unsatisfiable Cores. *JSAT* 9, 27–51.
- NARODYTSKA, N. AND BACCHUS, F. 2014. Maximum Satisfiability Using Core-Guided MaxSAT Resolution. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, Québec City, Canada, AAAI Press, 2717–2723.
- NIEMELÄ, I. 1999. Logic Programs with Stable Model Semantics as a Constraint Programming Paradigm. *Ann. Math. Artif. Intell.* 25, 3-4, 241–273.
- SIMONS, P., NIEMELÄ, I. AND SOININEN, T. 2002. Extending and implementing the stable model semantics. *Artif. Intell.* 138, 1-2, 181–234.