Inference Rules for High-Order Consistency in Weighted CSP *

Carlos Ansótegui
DIEI, UdL
Lleida, Spain

María L. Bonet
LSI, UPC
Barcelona, Spain

Jordi Levy
IIIA, CSIC
Bellaterra, Spain

Felip Manyà
DIEI, UdL
Lleida, Spain

Abstract

Recently defined resolution calculi for Max-SAT and signed Max-SAT have provided a logical characterization of the solving techniques applied by Max-SAT and WCSP solvers. In this paper we first define a new resolution rule, called signed Max-SAT parallel resolution, and prove that it is sound and complete for signed Max-SAT. Second, we define a restriction and a generalization of the previous rule called, respectively, signed Max-SAT i-consistency resolution and signed Max-SAT (i, j)-consistency resolution. These rules have the following property: if a WCSP signed encoding is closed under signed Max-SAT i-consistency, then the WCSP is i-consistent, and if it is closed under signed Max-SAT (i, j)-consistency, then the WCSP is (i, j)-consistent. A new and practical insight derived from the definition of these new rules is that algorithms for enforcing high order consistency should incorporate an efficient and effective component for detecting minimal unsatisfiable cores. Finally, we describe an algorithm that applies directional soft consistency with the previous rules.

Introduction

The Weighted Constraint Satisfaction Problem (WCSP) is a well-suited framework for modelling real-life problems with soft constraints. WCSP is an optimization version of the CSP framework in which constraints are extended by associating *costs* to tuples. Solving a WCSP instance, which is NP-hard, consists of finding a complete assignment of minimal cost.

Exact solvers for WCSP typically implement either variable elimination algorithms (e.g. (Dechter 1999)) or branch and bound algorithms which enforce a certain degree of soft constraint propagation at each node of the search tree (e.g. (Larrosa & Schiex 2004; de Givry *et al.* 2005)). Given the relevance of both complete inference and incomplete inference in WCSP solvers, our aim in this paper is to define a new complete inference system for WCSP that could lead to improved variable elimination algorithms, and to define inference rules that capture soft local consistency properties stronger than the various forms of soft arc consistency defined in the literature.

Our work —which lies at the intersection of the communities of Multiple-Valued Logic, Satisfiability and Constraint Processing— is closely related to recent work on resolution inference rules for Max-SAT and WCSP (Larrosa & Heras 2005; Heras & Larrosa 2006; Bonet, Levy, & Manyà 2006; Ansótegui *et al.* 2007). On the one hand, the complete Max-SAT and signed Max-SAT inference rules have given rise to variable elimination algorithms for solving Max-SAT (Bonet, Levy, & Manyà 2006) and WCSP (Ansótegui *et al.* 2007) with an original notion of variable saturation. On the other hand, the soft arc consistency properties enforced by WCSP solvers (de Givry *et al.* 2005; Larrosa & Schiex 2004) have been expressed as derived inference rules in (Ansótegui *et al.* 2007), and instantiations of the Max-SAT resolution rule incorporated into Max-SAT solvers have given rise to important performance improvements (Larrosa & Heras 2005; Heras & Larrosa 2006). Our work is also closely related to the soft k-consistency properties and algorithms defined in (Cooper 2005). Actually, our results provide a logical framework for representing and analyzing the local consistency operations of (Cooper 2005).

The link between the logical machinery defined for Max-SAT and the graphical models used in WCSP is the many-valued clausal formalism known as signed CNF formulae, which provide a well-suited language for representing and solving WCSP. Signed CNF formula use a generalized notion of literal, called *signed literal*. A signed literal is an expression of the form $S{:}p$, where p is a propositional variable and S, its *sign*, is a subset of a domain N. The informal meaning of $S{:}p$ is "p takes one of the values in S". Signed CNF formulae have their origin in the community of automated theorem proving in many-valued logics, where they are used as a generic and flexible language for representing many-valued interpretations (Beckert, Hähnle, & Manyà 2000).

Signed CNF formulae exploit the structure of domains as in CSP/WCSP without losing the simplicity of clausal forms. As a result, we get a language more expressive than Boolean CNF formulae, and algorithms that extend, with a very low overhead, the techniques implemented in SAT/Max-SAT solvers in a natural way.

The contributions of the paper may be summarized as fol-

*This research has been partially supported by projects iDEAS (TIN2004-04343), Mulog (TIN2004-07933-C03-01/03) and IEA (TIN2006-15662-C02-02) funded by MEC.

lows: we first define a new resolution rule, called signed Max-SAT parallel resolution, and prove that it is sound and complete for signed Max-SAT. Second, we define a restriction and a generalization of the previous rule called, respectively, signed Max-SAT i-consistency resolution and signed Max-SAT (i, j)-consistency resolution. These rules have the following property: if a WCSP signed encoding is closed under signed Max-SAT i-consistency, then the WCSP is i-consistent, and if it is closed under signed Max-SAT (i, j)-consistency, then the WCSP is (i, j)-consistent. By the form of the signed Max-SAT parallel resolution rule and its variants we see that algorithms for enforcing high order consistency should incorporate the best component for detecting minimal unsatisfiable cores. Our last contribution is a description of an algorithm that applies directional soft consistency with the previous rules, and enforces directional i-consistency.

Preliminaries

Definition 1 *A* truth value set*, or* domain*, N is a non-empty finite set $\{i_1, i_2, \ldots, i_n\}$ where n denotes its cardinality. A* sign *is a subset $S \subseteq N$ of truth values. A* signed literal *is an expression of the form $S{:}p$, where S is a sign and p is a propositional variable. The* complement *of a signed literal l of the form $S{:}p$, denoted by $\overline{l}$, is $\overline{S}{:}p = (N \setminus S){:}p$. A* signed clause *is a disjunction of signed literals. The empty clause, denoted by $\square$, is a disjunction of zero literals. A* signed CNF formula *is a* multiset *of signed clauses. The empty multiset of clauses is denoted by $\emptyset$.*

Definition 2 *An* assignment *for a signed CNF formula is a mapping that assigns to every propositional variable an element of the truth value set. An assignment I* satisfies *a signed literal $S{:}p$ iff $I(p) \in S$,* satisfies *a signed clause C iff it satisfies at least one of the signed literals in C, and* satisfies *a signed CNF formula Γ iff it satisfies all clauses in Γ. A signed CNF formula is* satisfiable *iff it is satisfied by at least one assignment; otherwise it is* unsatisfiable*.*

Definition 3 *The* signed Max-SAT problem *for a signed CNF formula consists of finding an assignment that minimizes the number of falsified signed clauses.*

Definition 4 *A* minimal unsatisfiable core *of a signed CNF formula Γ is any unsatisfiable subset Γ' of Γ such that, if we remove any clause $C \in \Gamma'$, then $\Gamma' \setminus \{C\}$ is satisfiable.*

Definition 5 *A* constraint satisfaction problem (CSP) *instance is defined as a triple $\langle X, D, C\rangle$, where $X = \{x_1, \ldots, x_n\}$ is a set of variables, $D = \{d(x_1), \ldots, d(x_n)\}$ is a set of domains containing the values the variables may take, and $C = \{C_1, \ldots, C_m\}$ is a set of constraints. Each constraint $C_i = \langle S_i, R_i\rangle$ is defined as a relation R_i over a subset of variables $S_i = \{x_{i_1}, \ldots, x_{i_k}\}$, called the* constraint scope*. The relation R_i may be represented extensionally as a subset of the Cartesian product $d(x_{i_1}) \times \cdots \times d(x_{i_k})$.*

Definition 6 *An* assignment *v for a CSP instance $\langle X, D, C\rangle$ is a mapping that assigns to every variable $x_i \in X$ an element $v(x_i) \in d(x_i)$. An assignment v satisfies a constraint $\langle\{x_{i_1}, \ldots, x_{i_k}\}, R_i\rangle \in C$ iff $\langle v(x_{i_1}), \ldots, v(x_{i_k})\rangle \in R_i$.*

Definition 7 *A* Weighted CSP (WCSP) *instance is defined as a triple $\langle X, D, C\rangle$, where X and D are variables and domains as in CSP. A constraint C_i is now defined as a pair $\langle S_i, f_i\rangle$, where $S_i = \{x_{i_1}, \ldots, x_{i_k}\}$ is the constraint scope and $f_i : d(x_{i_1}) \times \cdots \times d(x_{i_k}) \to \mathbb{N}$ is a* cost function*. The cost of a constraint C_i induced by an assignment v in which the variables of $S_i = \{x_{i_1}, \ldots, x_{i_k}\}$ take values $b_{i_1}, \ldots, b_{i_k}$ is $f_i(b_{i_1}, \ldots, b_{i_k})$. An optimal solution to a WCSP instance is a complete assignment in which the sum of the costs of the constraints is minimal.*

Definition 8 *The Weighted Constraint Satisfaction Problem (WCSP) for a WCSP instance consists of finding an optimal solution for that instance.*

Definition 9 *The* signed encoding *of a WCSP instance $\langle X, D, C\rangle$ is the signed CNF formula over the domain $N = \bigcup_{x_i \in D} d(x_i)$ that contains for every possible tuple $\langle b_{i_1}, \ldots, b_{i_k}\rangle \in d(x_{i_1}) \times \cdots \times d(x_{i_k})$ of every constraint $\langle\{x_{i_1}, \ldots, x_{i_k}\}, f_i\rangle \in C$, $f_i(b_{i_1}, \ldots, b_{i_k})$ copies of the signed clause:*

$$\overline{\{b_{i_1}\}}{:}x_{i_1} \vee \cdots \vee \overline{\{b_{i_k}\}}{:}x_{i_k}.$$

An alternative encoding is to consider signed clauses with weights instead of allowing multiple copies of a clause. For the sake of clarity we use unweighted clauses. Nevertheless, any efficient implementation of the algorithms proposed should deal with weighted clauses. The extension of our theoretical results to weighted clauses is straightforward.

Proposition 10 *Solving a WCSP instance P is equivalent to solving the signed Max-SAT problem of its signed encoding; i.e., the optimal cost of P coincides with the minimal number of unsatisfied signed clauses of the signed encoding of P.*

PROOF: See (Ansótegui *et al.* 2007). ∎

Example 11 *Figure 1 shows a WCSP instance $\langle X, D, C\rangle$ and its signed encoding. The WCSP has the set of variables $X = \{x_1, x_2, x_3, x_4\}$ with domains $d(x_1) = d(x_2) = \{a, b, c\}$ and $d(x_3) = d(x_4) = \{a, b\}$. Unary costs are depicted inside small circles. Binary costs are depicted as labeled edges connecting the corresponding pair of values. The label of each edge is the corresponding cost. If two values are not connected, the binary cost between them is 0. The optimal cost is 1. The 8 signed clauses represent the initial WCSP instance.*

Complete Inference Rules for signed Max-SAT

We start by recalling two complete inference rules for solving the SAT problem of signed CNF formulae: the first —called signed binary resolution— is a straightforward generalization of the resolution rule, and the second —called signed parallel resolution— is a more efficient rule. Second, we recall a complete rule for solving signed Max-SAT, which is the natural extension to signed Max-SAT of the Boolean Max-SAT rule (Larrosa & Heras 2005; Bonet, Levy, & Manyà 2006). The rule for signed Max-SAT was defined in (Ansótegui *et al.* 2007), also showing that the existing soft arc consistency operations can be captured

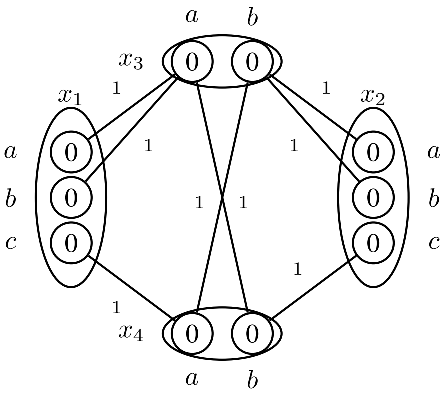

$$
\begin{array}{ll}
1: & \overline{\{a\}}:x_3 \vee \overline{\{a\}}:x_1\\
2: & \overline{\{a\}}:x_3 \vee \overline{\{b\}}:x_1\\
3: & \overline{\{a\}}:x_4 \vee \overline{\{c\}}:x_1\\
4: & \overline{\{b\}}:x_4 \vee \overline{\{a\}}:x_3\\
5: & \overline{\{a\}}:x_4 \vee \overline{\{b\}}:x_3\\
6: & \overline{\{b\}}:x_3 \vee \overline{\{a\}}:x_2\\
7: & \overline{\{b\}}:x_3 \vee \overline{\{b\}}:x_2\\
8: & \overline{\{b\}}:x_4 \vee \overline{\{c\}}:x_2
\end{array}
$$

Figure 1: A WCSP instance and its signed encoding.

by derived rules. Third, we define a new inference rule for signed Max-SAT which can be seen as the signed Max-SAT version of the signed parallel resolution rule, and prove its soundness and completeness. The completeness proof also provides a characterization of a family of complete rules for signed Max-SAT.

Signed binary resolution and signed parallel resolution are defined as follows (Hähnle 1994; 1996):

signed binary resolution

$$
\frac{\begin{array}{c}S{:}x \vee A\\ S'{:}x \vee B\end{array}}{S \cap S'{:}x \vee A \vee B}
\qquad\qquad
\frac{\emptyset{:}x \vee D}{D}
$$

signed parallel resolution

$$
\frac{\begin{array}{c}S_1{:}x \vee A_1\\ \cdots\\ S_k{:}x \vee A_k\end{array}}{A_1 \vee \cdots \vee A_k}
\qquad \text{whenever } \bigcap_{i=1}^{k} S_i = \emptyset
$$

In the above inference systems we assume w.l.o.g. that every variable in a clause appears only once collapsing different occurrences of a literal by computing the union of the supports.

It is possible to define the signed Max-SAT counterparts of the previous rules. The first one was defined and proved sound and complete in (Ansótegui *et al.* 2007). The second one is generalized in this paper.

Definition 12 *The* signed Max-SAT resolution *rule is defined as follows:*

$$
\begin{array}{l}
S{:}x \vee a_1 \vee \cdots \vee a_s\\
S'{:}x \vee b_1 \vee \cdots \vee b_t\\
\hline\hline
S \cap S'{:}x \vee a_1 \vee \cdots \vee a_s \vee b_1 \vee \cdots \vee b_t\\
S \cup S'{:}x \vee a_1 \vee \cdots \vee a_s \vee b_1 \vee \cdots \vee b_t\\
S{:}x \vee a_1 \vee \cdots \vee a_s \vee \overline{b_1}\\
S{:}x \vee a_1 \vee \cdots \vee a_s \vee b_1 \vee \overline{b_2}\\
\cdots\\
S{:}x \vee a_1 \vee \cdots \vee a_s \vee b_1 \vee \cdots \vee b_{t-1} \vee \overline{b_t}\\
S'{:}x \vee b_1 \vee \cdots \vee b_t \vee \overline{a_1}\\
S'{:}x \vee b_1 \vee \cdots \vee b_t \vee a_1 \vee \overline{a_2}\\
\cdots\\
S'{:}x \vee b_1 \vee \cdots \vee b_t \vee a_1 \vee \cdots \vee a_{s-1} \vee \overline{a_s}
\end{array}
$$

This inference rule is applied to multisets of clauses, and replaces *the premises of the rule by its conclusions.*

We say that the rule resolves *the variable x.*

The tautologies concluded by the rule like $N{:}x \vee A$ are removed from the resulting multiset. Also we substitute clauses like $S{:}x \vee S'{:}x \vee A$ by $(S \cup S'){:}x \vee A$, and clauses like $\emptyset{:}x \vee A$ by A.

For the sake of space, we can use the following more compact representation:

$$
\begin{array}{l}
S{:}x \vee A\\
S'{:}x \vee B\\
\hline\hline
S \cap S'{:}x \vee A \vee B\\
S \cup S'{:}x \vee A \vee B\\
S{:}x \vee A \vee \overline{B}\\
S'{:}x \vee \overline{A} \vee B
\end{array}
$$

Notice that $\overline{B}$, where B is a disjunction of signed literals, is not in CNF. Thus, an expression of the form $D \vee \overline{E}$, where D is a disjunction of signed literals and E is the disjunction of signed literals $e_1 \vee \cdots \vee e_t$, can be replaced by the following equivalent set of clauses:

$$
\begin{array}{l}
D \vee \overline{e_1}\\
D \vee e_1 \vee \overline{e_2}\\
\cdots\\
D \vee e_1 \vee \cdots \vee e_{t-1} \vee \overline{e_t}
\end{array}
$$

The second inference rule is new and is the one that corresponds to the signed parallel resolution rule in the signed Max-SAT framework:

Definition 13 *The* signed Max-SAT parallel resolution *rule is defined as follows:*

$$
\begin{array}{l}
S_1{:}x \vee D_1\\
\ldots\\
S_k{:}x \vee D_k\\
\hline\hline
D_1 \vee \ldots \vee D_k\\
\left\{\begin{array}{r}
(S_1 \cap \ldots \cap S_{t-1}) \cup S_t{:}x \vee D_1 \vee \ldots \vee D_t\\
(S_1 \cap \ldots \cap S_{t-1}){:}x \vee D_1 \vee \ldots \vee D_{t-1} \vee \overline{D_t}\\
S_t{:}x \vee \overline{D_1 \vee \ldots \vee D_{t-1}} \vee D_t
\end{array}\right\}_{t=2\ldots k}
\end{array}
$$

where D_i is a disjunction of signed literals, and $\{S_1{:}x, \ldots, S_k{:}x\}$ is a minimal unsatisfiable core. We call x the resolving variable.

We next prove that signed Max-SAT parallel resolution is a sound inference rule. In the context of Max-SAT rules, a rule is sound if the number of unsatisfied clauses in the premises coincides with the number of unsatisfied clauses in the conclusions for every truth assignment. Recall that applying a rule amounts to replacing the premises by the conclusions.

Theorem 14 *Signed Max-SAT parallel resolution is sound.*

PROOF: We show the soundness of the rule by proving that it is a derived rule of signed Max-SAT resolution. Given the initial k premises the multiset of conclusions can be obtained by exactly $k-1$ applications of signed Max-SAT resolution. The first application is on the first two premises. Then, for the l-th application, $2 \leq l \leq k-1$, one of the premises is the l+1-th premise, $S_{l+1}{:}x \vee D_{l+1}$, and the second premise is the first conclusion of the previous step, $S_1 \cap \ldots \cap S_l{:}x \vee D_1 \vee \ldots \vee D_l$. Note that here we use the fact that $\{S_1{:}x, \ldots, S_k{:}x\}$ is a minimal unsatisfiable core, since we assume that $S_1 \cap \cdots \cap S_l \neq \emptyset$. Then we replace these two premises by the following conclusions:

$$\begin{array}{l}
S_1 \cap \ldots \cap S_{l+1}{:}x \vee D_1 \vee \ldots \vee D_{l+1} \\
(S_1 \cap \ldots \cap S_l) \cup S_{l+1}{:}x \vee D_1 \vee \ldots \vee \overline{D_{l+1}} \\
(S_1 \cap \ldots \cap S_l){:}x \vee D_1 \vee \ldots \vee D_l \vee \overline{D_{l+1}} \\
S_{l+1}{:}x \vee \overline{D_1 \vee \ldots \vee D_l} \vee D_{l+1}
\end{array}$$

Notice that in the next step the second premise will be the first clause of the above conclusions. The first conclusion of the last step will be $S_1 \cap \ldots \cap S_k{:}x \vee D_1 \vee \ldots \vee D_k$. Since, $S_1{:}x, \ldots, S_k{:}x$ is a minimal unsatisfiable core, $S_1 \cap \ldots \cap S_k = \emptyset$. Therefore, the last application of the rule produces as a first conclusion $D_1 \vee \ldots \vee D_k$. ■

The completeness of the rule is a modification of the arguments of (Bonet, Levy, & Manyà 2006; Ansótegui *et al.* 2007). In order to sketch it we need some definitions and lemmas.

The first definition is the notion of saturation. This notion captures the idea that if a multiset of clauses $\mathcal{C}$ is saturated w.r.t. a variable x, then it does not make sense to keep applying the resolution rule with x as the resolving variable, either because the supports of x don't have an empty intersection (therefore the rule could not be applied) or the first clause of the conclusion is a tautology.

Definition 15 *A multiset of clauses $\mathcal{C}$ is said to be* saturated w.r.t. x *if for every subset of clauses $\{S_1{:}x \vee D_1, \ldots, S_m{:}x \vee D_m\} \subseteq \mathcal{C}$, it holds that either $S_1 \cap \cdots \cap S_m \neq \emptyset$ or there exist literals $S_{i_1}{:}y \in D_{i_1}, \ldots, S_{i_l}{:}y \in D_{i_l}, l \leq m$, such that $S_{i_1} \cup \cdots \cup S_{i_l} = N$. A multiset of clauses $\mathcal{C}'$ is a* saturation of $\mathcal{C}$ *w.r.t. x if $\mathcal{C}'$ is saturated w.r.t. x and $\mathcal{C} \vdash_x \mathcal{C}'$, i.e. $\mathcal{C}'$ can be obtained from $\mathcal{C}$ applying the inference rule resolving x finitely many times.*

Lemma 16 *Let $\mathcal{E}$ be a saturated multiset of clauses w.r.t. x. Let $\mathcal{E}'$ be the subset of clauses of $\mathcal{E}$ not containing x. Then, any assignment I satisfying $\mathcal{E}'$ (and not assigning x) can be extended to an assignment satisfying $\mathcal{E}$.*

PROOF: We have to extend I to satisfy the whole $\mathcal{E}$. In fact we only need to set the value of x. Let us partition the multiset $(\mathcal{E}-\mathcal{E}')$ (multiset of clauses that contain the variable x) into two multisets: $(\mathcal{E}-\mathcal{E}')_T$ the multiset already satisfied by I, and $(\mathcal{E}-\mathcal{E}')_F$ the multiset such that the partial assignment I doesn't satisfy any of the clauses.

Let $(\mathcal{E}-\mathcal{E}')_F = \{S_1{:}x \vee D_1, \ldots, S_k{:}x \vee D_k\}$. Since $\mathcal{E}$ is saturated, either $S_1 \cap \cdots \cap S_k \neq \emptyset$ or $D_1 \cup \cdots \cup D_k$ is a tautology. If $S_1 \cap \cdots \cap S_k \neq \emptyset$, we extend I with a value inside the intersection. If $S_1 \cap \cdots \cap S_k = \emptyset$, then $D_1 \cup \cdots \cup D_k$ is a tautology. Then, there exist $S_{i_1}{:}y \in D_{i_1}, \ldots, S_{i_l}{:}y \in D_{i_l}, l \leq m$, such that $S_{i_1} \cup \cdots \cup S_{i_l} = N$. Then, I satisfies one of these literals and the corresponding clause. So, some of the clauses are not in $(\mathcal{E}-\mathcal{E}')_F$. Contradiction. ■

Another ingredient of the proof is showing that the procedure of applying inference until saturation terminates. For that we need to define the notion of characteristic function of a multiset of clauses.

We assign a function $P : \mathbb{N} \rightarrow \{0,1\}$ to every signed literal, a function $P : \{0,1\}^n \rightarrow \{0,1\}$ to every clause, and a function $P : \{0,1\}^n \rightarrow \mathbb{N}$ to every multiset of clauses as follows.

Definition 17 *The* characteristic function *of a signed literal $L = \{i_1, \ldots, i_m\}{:}x$ is $P_L(x) = (1 - \{i_1\}{:}x) \cdot \ldots \cdot (1 - \{i_m\}{:}x)$.*

Definition 18 *For every clause $C = L_1 \vee \cdots \vee L_s$ we define its* characteristic function *as $P_C(\vec{x}) = P_{L_1}(x_1) \cdot \ldots \cdot P_{L_s}(x_s)$.*

For every multiset of clauses $\mathcal{C} = \{C_1, \ldots, C_m\}$, we define its characteristic function *as $P_{\mathcal{C}} = \Sigma_{i=1}^m P_{C_i}(\vec{x})$.*

Notice that, for every assignment $\mathcal{I}$, $P_C(\mathcal{I})$ is the number of clauses of $\mathcal{C}$ falsified by $\mathcal{I}$.

Also notice that the set of functions $\{0,1\}^n \rightarrow \mathbb{N}$, with the order relation: $f \leq g$ if for all x, $f(x) \leq g(x)$, defines a partial order between functions. The strict part of this relation, i.e. $f < g$ if for all x, $f(x) \leq g(x)$ and for some x, $f(x) < g(x)$, defines a well-founded order.

Lemma 19 *For every multiset of clauses $\mathcal{C}$ and every variable x, there exists a multiset $\mathcal{C}'$ such that $\mathcal{C}'$ is a saturation of $\mathcal{C}$ w.r.t. x.*

PROOF: (Sketch) By the soundness of the signed Max-SAT parallel resolution rule, every application of the rule replaces a multiset of clauses by another with the same characteristic function. But if we only look at the multisets containing the variable x, the characteristic function strictly decreases. This is because the first clause of the conclusion of the application of the rule does not contain the variable x and since it is not a tautology, its characteristic function is strictly greater than zero. Therefore when we apply the rule the characteristic function of the multiset of conclusions eliminating the first clause is strictly smaller than the characteristic function of the premises. ■

Now we are ready to state and sketch the proof of completeness.

Theorem 20 *Signed Max-SAT parallel resolution is complete; i.e., for any multiset of clauses $\mathcal{C}$, we have*

$$\mathcal{C} \vdash \underbrace{\Box, \ldots, \Box}_{m}, \mathcal{D}$$

where $\mathcal{D}$ is a satisfiable multiset of clauses, and m is the minimum number of unsatisfied clauses of $\mathcal{C}$.

PROOF: (Sketch) Let $x_1, \ldots, x_n$ be an ordering of the variables. We can saturate the set $\mathcal{C}$ with respect to the first variable x_1. By Lemma 19 this process terminates. The multiset that we obtain can be separated into two multisets. One contains the clauses where x_1 doesn't appear anymore, and we call it $\mathcal{C}_1$. The other contains the clauses where x_1 still appears, and we call $\mathcal{D}_1$.

Next we saturate the set $\mathcal{C}_1$ with respect to the second variable x_2. Again we separate the saturated multiset into two, $\mathcal{C}_2$ and $\mathcal{D}_2$, depending on whether x_2 appears or not. We continue this process until we saturate $\mathcal{C}_{n-1}$ with respect to x_n obtaining $\mathcal{C}_n$ and $\mathcal{D}_n$.

Notice that $\mathcal{C}$ is equivalent to $(\bigcup_{i=1}^{n} \mathcal{D}_i) \cup \mathcal{C}_n$. Clearly, $\mathcal{C}_n$ will be a multiset $\{\Box, \ldots, \Box\}$ of say m clauses. Then, m will be the minimum number of unsatisfied clauses of $\mathcal{C}$. On the other hand, $\mathcal{D}_1 \cup \ldots \cup \mathcal{D}_n$ will be a satisfiable multiset. This is proven by several applications of Lemma 16 extending an empty interpretation to one satisfying $\mathcal{D}_n$, and so on until $\mathcal{D}_1$. ■

In fact, our result is stronger than the statement of the previous theorem, because it characterizes a family of complete rules:

Corollary 21 *Any sound resolution rule for signed Max-SAT containing a non-tautological resolvent in which the resolving variable does not appear is complete.*

Local consistency via resolution

We are going to prove that the signed Max-SAT parallel resolution rule actually enforces i-consistency in WCSP. In order to do that, we first define (i,j)-*consistency* and i-*consistency* in WCSP.

Definition 22 *A WCSP is* (i,j)-consistent*, for $i \geq 0$ and $j \geq 1$, iff any consistent instantiation of i variables can be extended to a consistent instantiation of j additional variables.*

Definition 23 *A WCSP is* i-consistent *for, $i \geq 1$, iff it is $(i-1,1)$-consistent. A WCSP is* strong i-consistent*, for $i \geq 1$, iff it is k-consistent, for every $k, 1 \leq k \leq i$.*

For more details on soft k-consistency properties and algorithms see (Cooper 2005).

Now we will restrict the number of variables that appear in $D_1 \cup \cdots \cup D_k$ in the signed Max-SAT parallel resolution rule.

Definition 24 *The signed Max-SAT (i,1)-consistency resolution rule is the signed Max-SAT parallel resolution where exactly i variables appear in $D_1 \cup \cdots \cup D_k$ ($|var(D_1 \cup \cdots \cup D_k)| = i$). If $|var(D_1 \cup \cdots \cup D_k)| \leq i$ we call the resolution rule the* strong *signed Max-SAT (i,1)-consistency resolution rule.*

Lemma 25 *If a set of clauses is closed under the signed Max-SAT $(i-1,1)$-consistency resolution rule, then its corresponding WCSP is i-consistent.*

PROOF: Suppose that a set of clauses is closed by the signed Max-SAT $(i-1,1)$-consistency resolution rule, but its corresponding constraint network is not i-consistent. We have some tuple of $i-1$ variables $x_1, \ldots, x_{i-1}$ and $i-1$ consistent values $a_1, \ldots, a_{i-1}$ of their domains, and there exists also a variable x such that $a_1, \ldots, a_{i-1}$ can not be extended to x consistently. I.e. for any value b of the domain of x, the tuple of i values $a_1, \ldots, a_{i-1}, b$ for the variables $x_1, \ldots, x_{i-1}, x$ falsifies some constraint about a subset of such variables (where at least the variable x is present). Therefore, for any b, the tuple $a_1, \ldots, a_{i-1}, b$ is a nogood, and therefore there is a clause whose literals are a subset of $S_1{:}x_1 \vee \ldots \vee S_{i-1}{:}x_{i-1} \vee S{:}x$ where $\forall l\ 1 \leq l \leq i-1,\ a_l \notin S_l$ and $b \notin S$. Since the set of clauses is closed under the rule, and the intersection of the supports of x is empty, our set of clauses also contains a subclause of $S'_1{:}x_1 \vee \ldots \vee S'_{i-1}{:}x_{i-1}$ where $\forall l\ 1 \leq l \leq i-1,\ a_l \notin S'_l$. So the tuple $a_1, \ldots, a_{i-1}$ is a no good for $\langle x_1, \ldots, x_{i-1} \rangle$ and this contradicts the assumption. ■

Signed Max-SAT (i,j)-consistency resolution rule

Once we have introduced the signed Max-SAT (i,1)-consistency resolution rule, we are ready to introduce a more general rule, the signed Max-SAT (i,j)-consistency resolution rule. In this case, instead of just having one resolving variable we have exactly j resolving variables. For the sake of clarity one way of describing this rule is to *collapse* the j resolving variables, say $x_1, \ldots, x_j$, into one single variable x' whose domain is equal to the Cartesian product of the domains of the j original variables, $d(x') = d(x_1) \times \cdots \times d(x_j)$, where $d(x)$ is the domain of x. As a consequence, a disjunction of signed literals on the j variables, $S_1{:}x_1 \vee \cdots \vee S_j{:}x_j$ is replaced by the signed literal, $S'{:}x'$, where $S' = \overline{\overline{S_1} \times \cdots \times \overline{S_j}}$. Then, we can simply apply the (i,1)-signed Max-SAT resolution rule taking x' as the new resolving variable.

Definition 26 *The* signed Max-SAT (i,j)-consistency resolution *rule is defined as follows:*

$$
\begin{array}{l}
S_{1,1}{:}x_1 \vee \cdots \vee S_{1,j}{:}x_j \vee D_1 \\
\qquad\qquad\vdots \\
S_{k,1}{:}x_1 \vee \cdots \vee S_{k,j}{:}x_j \vee D_k \\
\hline\hline
D_1 \vee \ldots \vee D_k \\
\left\{ \begin{array}{r}
(S'_1 \cap \ldots \cap S'_{t-1}) \cup S'_t{:}x' \vee D_1 \vee \ldots \vee D_t \\
(S'_1 \cap \ldots \cap S'_{t-1}){:}x' \vee D_1 \vee \ldots \vee D_{t-1} \vee \overline{D_t} \\
S'_t{:}x' \vee \overline{D_1 \vee \ldots \vee D_{t-1}} \vee D_t
\end{array} \right\}_{t=2\ldots k}
\end{array}
$$

where $d(x') = d(x_1) \times \cdots \times d(x_j)$, $S'_r = \overline{\overline{S_{r,1}} \times \cdots \times \overline{S_{r,j}}}$, $k \geq r \geq 1$, $|var(D_1 \cup \cdots \cup D_k)| = i$, and $\{S'_1{:}x, \ldots, S'_k{:}x\}$ is a minimal unsatisfiable core.

Regarding the relation with the local consistencies in WCSP, we also say that if a set of clauses is closed under the signed Max-SAT (i,j)-consistency resolution rule, then its corresponding WCSP is (i,j)-consistent. Because of lack of space, the proof is omitted and will be provided in the long version of the paper.

$$
\begin{array}{rl}
\multicolumn{2}{l}{C_1} \\
\hline
4: & \overline{\{b\}} : x_4 \vee \overline{\{a\}} : x_3 \\
5: & \overline{\{a\}} : x_4 \vee \overline{\{b\}} : x_3 \\
6: & \overline{\{b\}} : x_3 \vee \overline{\{a\}} : x_2 \\
7: & \overline{\{b\}} : x_3 \vee \overline{\{b\}} : x_2 \\
8: & \overline{\{b\}} : x_4 \vee \overline{\{c\}} : x_2 \\
9: & \overline{\{a\}} : x_4 \vee \overline{\{a\}} : x_3 \\
\multicolumn{2}{l}{C_2} \\
\hline
4: & \overline{\{b\}} : x_4 \vee \overline{\{a\}} : x_3 \\
5: & \overline{\{a\}} : x_4 \vee \overline{\{b\}} : x_3 \\
9: & \overline{\{a\}} : x_4 \vee \overline{\{a\}} : x_3 \\
12: & \overline{\{b\}} : x_4 \vee \overline{\{b\}} : x_3 \\
\multicolumn{2}{l}{C_3} \\
\hline
15: & \overline{\{b\}} : x_4 \\
16: & \overline{\{a\}} : x_4 \\
\multicolumn{2}{l}{C_4} \\
\hline
17: & \square
\end{array}
\qquad
\begin{array}{rl}
\multicolumn{2}{l}{D_1} \\
\hline
10: & \{c\} : x_1 \vee \overline{\{a\}} : x_3 \vee \{a\} : x_4 \\
11: & \overline{\{c\}} : x_1 \vee \{a\} : x_3 \vee \overline{\{a\}} : x_4 \\
\multicolumn{2}{l}{D_2} \\
\hline
13: & \{c\} : x_2 \vee \overline{\{b\}} : x_3 \vee \{b\} : x_4 \\
14: & \overline{\{c\}} : x_2 \vee \{b\} : x_3 \vee \overline{\{b\}} : x_4 \\
\multicolumn{2}{l}{D_3} \\
\hline
 & \emptyset \\
\multicolumn{2}{l}{D_4} \\
\hline
 & \emptyset
\end{array}
$$

Figure 2: The application of the algorithm for local consistencies to the WCSP from Example 1.

Finally, notice that we get different instantiations of this rule by fixing the i and j parameters. In particular, if we set the i parameter to 0, we get an interesting instantiation where the rule has the empty clause as the first of its conclusions.

A directional local consistency algorithm

From the proof sketch of Theorem 20, we can extract an algorithm for applying complete and incomplete inference on WCSP. The only difference with the saturation procedure described in it is that we apply inference only to clauses of a bounded number of literals. This bound on the number of literals is the parameter i in the saturation function. Function $saturation(C_{s-1}, i, x_s)$ computes a saturation of C_{s-1} w.r.t. x_i applying the strong signed Max-SAT (i-1,1)-consistency resolution rule resolving x_s until it obtains a saturated set. Lemma 19 ensures that this process terminates, in particular that it does not cycle.

Function $partition(C, x_s)$ computes a partition of C, already saturated, into the subset of clauses containing x_s and the subset of clauses not containing x_s.

> **input:** A WCSP instance P, an index i
> $C_0 := signed\_encoding(P)$
> **for** $s := 1$ **to** n
> $\quad C := saturation(C_{s-1}, i, x_s)$
> $\quad \langle C_s, D_s \rangle := partition(C, x_s)$
> **endfor**
> **output:** $C_n \cup \bigcup_{s=1}^{n} D_s$

Given an initial WCSP instance P with n variables, the above algorithm returns an equivalent WCSP instance. The order on the saturation of the variables can be freely chosen, i.e. the sequence $x_1, \ldots x_n$ can be any enumeration of the variables. Notice that if i is the number of variables, the previous algorithm is complete.

Example 27 *Figure 2 shows the application of the algorithm to the WCSP from Example 1, with parameter* $i = 3$ *and the order on the variables given by the sequence* x_1, x_2, x_3, x_4. C_1 *and* D_1 *are obtained applying the signed Max-SAT (2,1)-consistency resolution rule on clauses* $1, 2, 3$ *with resolving variable* x_1. C_2 *and* D_2 *are obtained applying the signed Max-SAT (2,1)-consistency resolution rule on clauses* $6, 7, 8$ *with resolving variable* x_2. C_3 *is obtained applying the signed Max-SAT (1,1)-consistency on clauses* $4, 12$ *and* $5, 9$ *with the resolving variable* x_3. *Finally, the empty clause is obtained from clauses* $15, 16$ *with resolving variable* x_4.

Conclusions

We have defined a new resolution rule, called signed Max-SAT parallel resolution, and proved that it is sound and complete. Then, we have introduced a restriction and a generalization called signed Max-SAT i-consistency resolution and signed Max-SAT (i, j)-consistency resolution, respectively. If a WCSP signed encoding is closed under signed Max-SAT i-consistency, then the WCSP is i-consistent, and if it is closed under signed Max-SAT (i, j)-consistency, then the WCSP is (i, j)-consistent. A new and practical insight derived from the definition of these new rules is that algorithms for enforcing high order consistency should incorporate an efficient and effective component for detecting minimal unsatisfiable cores. Finally, we have described an algorithm that applies directional soft consistency with the previous rules, enforcing directional i-consistency.

References

Ansótegui, C.; Bonet, M.; Levy, J.; and Manyà, F. 2007. The logic behind weighted CSP. In *IJCAI-2007*, 32–37.

Beckert, B.; Hähnle, R.; and Manyà, F. 2000. The SAT problem of signed CNF formulas. In *Labelled Deduction*, volume 17 of *Applied Logic Series*. Kluwer. 61–82.

Bonet, M.; Levy, J.; and Manyà, F. 2006. A complete calculus for Max-SAT. In *SAT-2006*.

Cooper, M. C. 2005. High-order consistency in valued constraint satisfaction. *Constraints* 10:283–305.

de Givry, S.; Heras, F.; Larrosa, J.; and Zytnicki, M. 2005. Existencial arc consistency: Getting closer to full arc consistency in weighted CSPs. In *IJCAI-2005*, 84–89.

Dechter, R. 1999. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence* 113(1–2):41–85.

Hähnle, R. 1994. Efficient deduction in many-valued logics. In *ISMVL'94* , 240–249.

Hähnle, R. 1996. Exploiting data dependencies in many-valued logics. *J. of Applied Non-Classical Logics* 6:49–69.

Heras, F., and Larrosa, J. 2006. New inference rules for efficient Max-SAT solving. In *AAAI-2006*, 68–73.

Larrosa, J., and Heras, F. 2005. Resolution in Max-SAT and its relation to local consistency in weighted CSPs. In *IJCAI-2005*, 193–198.

Larrosa, J., and Schiex, T. 2004. Solving weighted CSP by maintaining arc-consistency. *Artificial Intelligence* 159(1–2):1–26.