

Solving WCSP by Extraction of Minimal Unsatisfiable Cores

Christophe Lecoutre Nicolas Paris Olivier Roussel Sébastien Tabary
CRIL - CNRS, UMR 8188
Université Lille-Nord de France, Artois
rue de l'université, SP 16, F-62307 Lens, France
Email : {lecoutre,paris,roussel,tabary}@cril.fr

Abstract—Usual techniques to solve WCSP are based on cost transfer operations coupled with a branch and bound algorithm. In this paper, we focus on an approach integrating extraction and relaxation of Minimal Unsatisfiable Cores in order to solve this problem. We derive our approach in two ways: an incomplete, greedy, algorithm and a complete one.

I. INTRODUCTION

The Constraint Satisfaction Problem (CSP) consists in determining if a Constraint Network (CN) is satisfiable or not. The goal of this decision problem is to find a solution, i.e., a complete instantiation of the variables of the CN that satisfies each constraint. Sometimes, preferences among solutions need to be expressed. This is possible through the introduction of soft constraints. The framework called WCSP (Weighted CSP) allows us to handle such constraints: each soft constraint is defined by a cost function that associates a violation degree, called cost, with every possible instantiation of a subset of variables. The WCSP goal is to find a complete instantiation with a minimal combined cost of the soft constraints.

Most of the current methods to solve Weighted CNs (WCNs) are based on branch and bound tree search together with the use of soft local consistencies for estimating minimal costs of sub-problems during search. These local consistencies, e.g., AC* [1], [2], FDAC [3], EDAC [4], VAC [5], and OSAC [6], have been developed on the concept of equivalence-preserving transformations (cost transfer operations). Because they must combine costs, algorithms establishing soft local consistencies are often more complex than their CSP counterparts. In this paper, we propose an original approach for WCSP that consists in solving a WCN by iteratively generating and solving classical CNs. One advantage is to benefit from efficient CSP solvers developed in the community for more than two decades. Notice that reformulating WCNs into CNs has already been used quite successfully to enforce VAC [5], where (classical) arc consistency is established on a CN generated from a WCN in order to identify some relevant cost transfer operations.

In our approach, to solve a WCN, soft constraints are converted into hard constraints by retaining only the tuples

that have a specified cost. Actually, the principle is to enumerate a sequence of CNs according to an increasing cost order related to the WCN. When a CN is proved to be unsatisfiable, a Minimal Unsatisfiable Core (MUC) is extracted in order to identify the soft constraints whose reformulation into hard constraints needs to be relaxed, so as to finally obtain a solution. This approach comes in two versions: a complete algorithm and a greedy one (that is incomplete). Note that similar approaches have been successfully exploited for MaxSAT [7], [8], [9], [10].

This paper is organized as follows. After some usual definitions, we present the principle of our approach and the required data structures. Next, we completely describe a complete version of our approach, as well as a greedy one. Finally, after presenting some practical results, we conclude.

II. TECHNICAL BACKGROUND

A. Generalities

A *constraint network* (CN) P involves a finite set of n variables denoted by $vars(P)$ and a finite set of e constraints denoted by $cons(P)$. Each variable x has an associated domain, denoted by $dom(x)$, which contains the finite set of values that can be assigned to x ; the initial domain of x is denoted by $dom^{init}(x)$. Each (hard) constraint c involves an ordered set of variables, denoted $scp(c)$ and called *scope* of c , and is defined by a relation containing the set of tuples allowed for the variables in $scp(c)$. The arity of a constraint c is $|scp(c)|$. An *instantiation* I of a set $X = \{x_1, \dots, x_p\}$ of variables is a set $\{(x_1, a_1), \dots, (x_p, a_p)\}$ such that $\forall i \in 1..p, a_i \in dom^{init}(x_i)$. I is *valid* on P iff $\forall (x, a) \in I, a \in dom(x)$. A solution is a complete instantiation of $vars(P)$ (i.e., the assignment of a value to each variable) such that all the constraints are satisfied. P is satisfiable iff it admits at least one solution. For more information about constraint networks, see [11], [12], [13].

An unsatisfiable core of P is an unsatisfiable subset of constraints of P . A core C is a MUC (Minimal Unsatisfiable Core) iff each strict subset of constraints of C is satisfiable. Several methods to extract MUCs of constraint networks are presented in [14], [15], [16], [17]. Among the

different versions presented in [17], we have chosen for our implementation the *dichotomic* version, called *dcMUC*.

A *weighted constraint network* (WCN) W involves a finite set of n variables denoted by $vars(W)$ and a finite set of e soft constraints denoted by $cons(W)$, and has an associated value $k > 0$ which is either a natural integer or $+\infty$. Each soft constraint $w \in cons(W)$ has a scope $scp(w)$ and is defined as a cost function from $l(scp(w))$ to $\{0, \dots, k\}$, where $l(scp(w))$ (labelling) is the Cartesian product of the domains of the variables involved in w ; for each instantiation $I \in l(scp(w))$, the cost of I in w is denoted by $w(I)$.

When an instantiation is given the cost k it is said *forbidden*. Otherwise, it is permitted with the corresponding cost (0, being completely satisfactory). Costs are combined with the bounded addition \oplus defined as:

$$\forall a, b \in \{0, \dots, k\}, a \oplus b = \min(k, a + b)$$

The objective of Weighted Constraint Satisfaction Problem (WCSP) is, for a given WCN, to find a complete instantiation with a minimal cost. For more information on weighted constraint networks, see [18], [19].

Different variations of soft arc consistency for WCSP have been proposed during the last decade: AC* [1], [2], full directional arc consistency (FDAC) [3], existential arc consistency (EDAC) [4], virtual arc consistency (VAC) [5] and optimal soft arc consistency (OSAC) [6]. All the algorithms proposed to enforce these different levels of consistency use cost transfer operations (based on the concept of equivalence-preserving transformations) such as unary projection, projection and extension. In contrast, our approach is based on selecting tuples of a given cost which are called layers and defined below.

B. Layers and Fronts

In this paper, we focus on extensional soft constraints. These are soft table constraints where some tuples are explicitly represented with their costs in a table, and a default cost indicates the cost of all implicit tuples (i.e., those not present in the table); e.g., see [20]. We introduce now both the notions of layer and front that we use in our algorithm. Layers represent the different degrees of violation of a single constraint and fronts represent the currently admitted global degree of violation.

All tuples of a soft table constraint having the same cost can be grouped to form a subset called *layer*. The cost of any tuple from layer L is given by $cost(L)$. A particular layer corresponds to the default cost: this layer contains no tuple, but implicitly represents all tuples that do not explicitly appear elsewhere. The number of layers for a constraint w is given by $nbLayers(w)$. Within a constraint, we consider that layers are increasingly ordered according to their costs. Finally, for the sake of simplicity, we shall use indices, from 0 to $nbLayers(w) - 1$, to identify the different layers of a

constraint, and when the context is unambiguous, we shall not distinguish between indices and the layers they represent.

A *front* f is a function that maps each constraint of a WCN to one of its layers. A front represents a border between the layers that will be considered (allowed) at a given moment, and the layers that will be discarded (forbidden). We shall use an array notation for the fronts. So, $f[w]$ will represent the layer associated with the constraint w in the front f . The cost of a front f , $cost(f)$, is obtained by summing up all costs of the constraint layers corresponding to the front f :

$$cost(f) = \bigoplus_{w \in cons(W)} cost(f[w])$$

A front f' is the direct successor of a front f iff there exists a constraint w_i such that $f'[w_i] = f[w_i] + 1$ and $\forall j \neq i, f'[w_j] = f[w_j]$. We note $f \rightarrow f'$ iff f' is a direct successor of f and $f \rightarrow_* f'$ iff there exists a sequence of relations $f \rightarrow f_1, f_1 \rightarrow f_2, \dots, f_n \rightarrow f'$ (transitive closure).

Because the layers of any constraint are increasingly ordered following their costs, we deduce that:

$$f \rightarrow_* f' \Rightarrow cost(f) < cost(f')$$

We can observe that the set of all possible fronts and the direct successor relation form a lattice structure, where the lowest front (also noted \perp) associates with each constraint its first layer (of lowest cost) and the highest front (also noted \top) associates with each constraint its last layer (of highest cost).

III. GENERAL PRINCIPLE

Figure 1 illustrates the main idea of our approach for solving weighted constraint networks. First, from a WCN W , an initial constraint network P is built as follows: for each constraint of W , the tuples occurring in the layer of minimal cost are considered as allowed in P , while other tuples are considered as forbidden (function *toCN*). This process is similar to that proposed in [5]. Next, the CN P is solved using a constraint solver (through a call to the function *solveCN*). If P is satisfiable, then a solution is returned. Otherwise a MUC is extracted from P (through a call to the function *extractMUC*) and then exploited either in a greedy or complete version.

In the greedy version, some constraints of the MUC will be successively relaxed until satisfiability of the MUC is restored (through a call to the function *relax*). More precisely, a subset of layers of some constraints belonging to the MUC will be switched from the “forbidden” status to the “allowed” one. When constraints corresponding to the MUC are relaxed so as to become satisfiable, the process loops until the global satisfiability of P is reached. In this case, a solution is returned.

In the complete version, once a MUC is identified, we insert in a priority queue the direct successors of the current

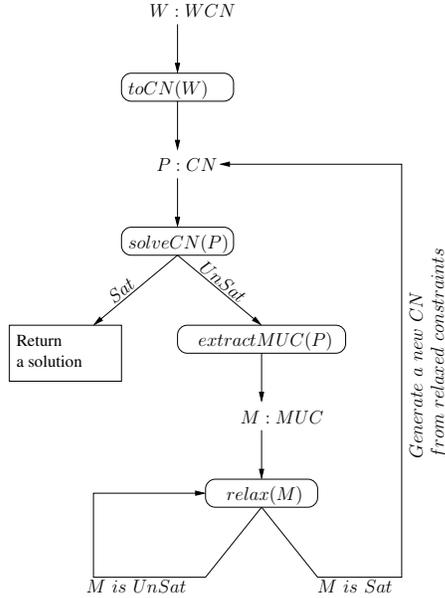


Figure 1. Principle of the iterative relaxation of MUCs.

front in order to break the MUC by relaxing at least one constraint. Fronts are extracted from the priority queue, translated to a CN, and successively solved until a satisfiable constraint network is obtained. In this case, an optimum solution is returned.

Interestingly, the principle described above can be generalized to take into account any kind of constraints: hard constraints from the CSP framework, “violable” constraints from the Weighted Max-CSP framework and soft constraints from the WCSP framework. When the function $toCN$ is called, we just need to indicate the translation process for each kind of constraints. For a hard constraint, no transformation is required (we have only one layer with a cost equal to 0). For a “violable” constraint, two layers are managed. The first layer, with a cost equal to 0, corresponds to the satisfaction of the “violable” constraint. The second one, with a cost equal to the weight of the “violable” constraint, corresponds to the violation of the constraint. In other words, our approach allows us to deal easily and naturally with these different frameworks. Due to lack of space, this is not further described in this paper.

IV. DATA STRUCTURES AND FUNCTIONS

Fronts are exploited by the function $toCN$ which builds a CN from a WCN. Two versions are considered. The first one, noted $toCN_{=}(W, f)$, builds a CN by selecting as allowed tuples for a soft constraint w those belonging to the layer $f[w]$; this layer is said allowed. The second one, noted $toCN_{\leq}(W, f)$, builds a CN by selecting as allowed tuples for a soft constraint w the tuples of the layers whose index is less than or equal to $f[w]$; these layers are said allowed.

In other words, for each hard constraint c built from a soft constraint w of W , the tuples allowed in c are those having a cost equal (or less than or equal) to the cost of the layer $f[w]$.

In practice, two representations of a hard constraint can be considered. On the one hand, when the layer corresponding to the default cost is not allowed, $toCN$ generates a hard positive constraint. The tuples of the hard constraint are allowed and correspond to the tuples of the allowed layers of the soft constraint. On the other hand, when the layer corresponding to the default cost is allowed, $toCN$ generates a hard negative constraint. The tuples of the hard constraint are forbidden and correspond to the tuples of the forbidden layers of the soft constraint in order to avoid to enumerate the implicit tuples having a cost equal to the default cost.

Figure 2 presents two constraints w_0 and w_1 with their layers. Constraint w_0 has a set of layers numbered from 0: layer 0 of minimal cost $cost_0$ contains the tuples τ_1 , τ_5 , and τ_7 , and the next layer (layer 1) of cost $cost_1$ contains the tuples τ_2 and τ_3 . In this example, $f[w_0]$ is equal to 1 and $f[w_1]$ is equal to 0. The CN returned by $toCN_{=}(W, f)$ contains a constraint c_0 (associated with w_0) with allowed tuples τ_2 and τ_3 and c_1 (associated with w_1) with allowed tuples τ_1 and τ_4 . The CN returned by $toCN_{\leq}(W, f)$ contains a constraint c_0 (associated with w_0) with allowed tuples τ_2 , τ_3 , τ_1 , τ_5 , and τ_7 and c_1 (associated with w_1) with allowed tuples τ_1 and τ_4 . Allowed layers of $toCN_{\leq}(W, f)$ are identified by a grey background in Figure 2.

We conclude this section by a short description of the other functions we need. The function $solveCN(P)$ solves a CN given in parameter and returns either a solution or \perp if the CN is unsatisfiable. The function $extractMUC(P)$ returns a MUC from an unsatisfiable CN P given in parameter. The function $cons(W, M)$ returns the set of soft

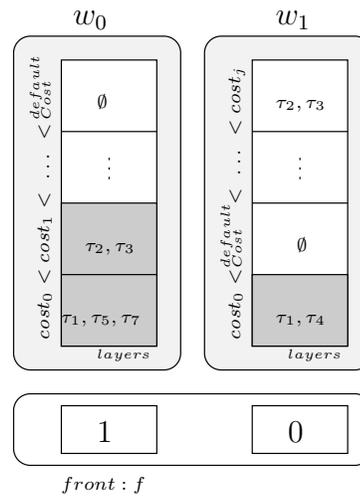


Figure 2. Description of data structures.

constraints of the WCN W which correspond to those of the MUC M (by construction any hard constraint maps to a soft constraint). The function $restrict(W, M)$ returns a WCN containing only the constraints of the WCN W corresponding to the constraints present in the MUC M .

V. ALGORITHMS

A. Introductory Complete Version

To facilitate the understanding of our approach, we present a first complete algorithm which does not exploit MUCs and thus is not efficient. Function `completeSearchNoMUC` starts with a first front which retains only layer 0 of each constraint (lines 1-3). Next, while there exists a front to explore, we extract from a priority queue the front with the minimal cost. The cost of a front is easily computed by summing the costs associated with the layers selected in each soft constraint. Then, we can build with the function $toCN_{=}$ a CN that we try to solve with $solveCN$. If a solution is found, it is optimal and the algorithm stops. Indeed, the algorithm enumerates fronts in increasing cost order (see Proposition 1, later) and guarantees that each previous front was unsatisfiable. If the tested CN was unsatisfiable, we enumerate the direct successors of the current front (lines 11-16) and insert them in the priority queue. Obviously, we have to check that the maximal layer for a given constraint has not been reached. Besides, fronts having a cost equal to k must be ignored.

Function `completeSearchNoMUC(W: WCN)`

```

1 foreach  $w \in cons(W)$  do
2    $f[w] \leftarrow 0$ 
3  $Q \leftarrow \{f\}$ 
4 while  $Q \neq \emptyset$  do
5    $f \leftarrow$  pick and delete least cost front in  $Q$ 
6    $P \leftarrow toCN_{=}(W, f)$ 
7    $sol \leftarrow solveCN(P)$ 
8   if  $sol \neq \perp$  then
9     return  $sol$ 
10  else
11    foreach  $w \in cons(W)$  do
12      if  $f[w] \neq nbLayers(w) - 1$  then
13         $f' \leftarrow f$ 
14         $f'[w] \leftarrow f'[w] + 1$ 
15        if  $cost(f') \neq k$  then
16           $Q \leftarrow Q \cup f'$ 

```

The priority queue Q used in this algorithm is a variant of the usual priority queues because it must ensure that fronts are extracted in increasing cost order but also that a given element is not inserted twice in the queue. For example, if we consider a network composed of two constraints each

having at least two layers, the first front will be the array $\langle 0, 0 \rangle$. During the first iteration, we shall insert in Q the neighbors $\langle 1, 0 \rangle$ and $\langle 0, 1 \rangle$. When these fronts are extracted from the queue, they both generate the same neighbor $\langle 1, 1 \rangle$ that must be inserted only once in Q to avoid redundant computations. This particular queue is obtained by slightly modifying a classical implementation of a priority queue with a binomial heap.

One can easily check that Algorithm `completeSearchNoMUC` enumerates all possible fronts when $solveCN$ never finds a solution. Indeed, if we ignore costs, this algorithm performs a breadth-first search enumerating fronts. Proposition 1 also guarantees that this algorithm does not loop endlessly and that fronts are enumerated in increasing cost order.

Proposition 1: The following properties are guaranteed by `completeSearchNoMUC`:

- A) fronts are enumerated in increasing cost order;
- B) once a front f is extracted from the priority queue, it can never be inserted again.

Proof: Let f_1 be a first front extracted from the queue and f_2 a front extracted from the queue after the extraction of f_1 .

A) Two cases are possible. Either (A.1) f_2 was already present in the queue when f_1 has been extracted or (A.2) f_2 has been inserted in the queue after the extraction of f_1 . In case (A.1), the priority queue ensures that $cost(f_1) \leq cost(f_2)$. In case (A.2), f_2 comes from a front f such that $f \rightarrow_* f_2$ and either $f = f_1$ or f was present in the queue when f_1 has been extracted. In both cases, $cost(f_1) \leq cost(f)$. As $f \rightarrow_* f_2 \Rightarrow cost(f) < cost(f_2)$, we conclude that $cost(f_1) < cost(f_2)$. Therefore, we obtain that in any case $cost(f_1) \leq cost(f_2)$.

B) Let us assume that $f_1 = f_2$. As the queue ensures that it never contains two identical fronts at a given time, we conclude that we are necessary in case (A.2). However, we proved that in this case, $cost(f_1) < cost(f_2)$, which contradicts $f_1 = f_2$. ■

It should be emphasized that in our approach, unary constraints are considered as normal constraints. Moreover, because only one layer of a constraint is considered at a given time, the CN generated are typically much smaller than the original WCN. One can expect that checking the satisfiability will be quite simple (hence fast) in most cases. One last observation is that this algorithm exploits a form of abstraction considering that, from a global point of view, there is no reason to distinguish tuples having the same cost.

To conclude this section, the drawback of this first version is the enumeration of all possible combinations of layers. In the worst-case, it enumerates a number of fronts equal to the product of the number of layers of each constraint, which corresponds to $\prod_{w \in cons(W)} nbLayers(w)$. Depending on the number of variables and constraints, this complexity may be much greater compared to a classical branch and

bound approach (i.e., when $\prod_{w \in \text{cons}(W)} \text{nbLayers}(w) \gg \prod_i |\text{dom}(x_i)|$). This is explained by the fact that, in our approach, the same instantiation can be explored several times in different tests of satisfiability. Nevertheless, even in this case, this approach can be effective if the optimum is low.

B. Exploiting MUCs

The complexity of the previous algorithm can be reduced considerably in practice by noticing that it is useless to move from a layer to the next one for a constraint which does not participate in the unsatisfiability of the CN. The idea is therefore to identify a MUC and to move to the next layer only for the constraints belonging to this MUC.

In the worst-case, the MUC extracted contains all the constraints of the problem and therefore the complexity in the worst-case remains unchanged. However in practice, on concrete instances, MUCs are often small. Hence, the generated neighborhood is much smaller and the practical complexity is significantly reduced.

1) *Complete Approach*: Function `completeSearch` (inspired by [8], [9], [10]) presents the algorithm modified in order to take profit of MUCs in a complete approach. The first steps of the algorithm are identical to those of `completeSearchNoMUC`. When the current CN is unsatisfiable, a MUC is identified, and we shall progressively allow the higher layers in the constraints of the MUC. More precisely, the algorithm enumerates all the direct successors of the current front in order to relax this MUC. For each constraint of the MUC (obtained by a call to $\text{cons}(W, M)$), the algorithm generates a successor f' from the current front which differs only by the incrementation of the allowed layer for this constraint. This new front is inserted in Q at a position depending on the value of $\text{cost}(f')$.

Proposition 2: Algorithm `completeSearch` is complete.

Proof: To simplify the proof, we shall use the term “front” to refer implicitly to the CN associated with a front f by $\text{toCN}_=(W, f)$. Let f_u be the front which is identified as unsatisfiable and M the MUC extracted from this front. The only difference between `completeSearch` and `completeSearchNoMUC` is that when a MUC is identified, we restrict the search by relaxing this MUC prior to any other operation.

If the WCN has no solution, restricting the search in this way obviously does not cause the loss of any solution. We can therefore only consider the case where the CN admits at least one solution S . Let f_S be the front corresponding to a solution S . The algorithm `completeSearchNoMUC` ensures that there always exists a front f in the priority queue such that $f \rightarrow_* f_S$. This is especially true for the first front inserted in the queue \perp and when this property is satisfied for a front extracted from the queue, it is satisfied for at least one of its direct successors (by definition of \rightarrow_*). Therefore, by recurrence, this property is always ensured.

Function `completeSearch(W: WCN)`

```

1 foreach  $w \in \text{cons}(W)$  do
2    $f[w] \leftarrow 0$ 
3  $Q \leftarrow \{f\}$ 
4 while  $Q \neq \emptyset$  do
5    $f \leftarrow$  pick and delete least cost front in  $Q$ 
6    $P \leftarrow \text{toCN}_=(W, f)$ 
7    $\text{sol} \leftarrow \text{solveCN}(P)$ 
8   if  $\text{sol} \neq \perp$  then
9     return  $\text{sol}$ 
10  else
11     $M \leftarrow \text{extractMUC}(P)$ 
12    foreach  $w \in \text{cons}(W, M)$  do
13      if  $f[w] \neq \text{nbLayers}(w) - 1$  then
14         $f' \leftarrow f$ 
15         $f'[w] \leftarrow f'[w] + 1$ 
16        if  $\text{cost}(f') \neq k$  then
17           $Q \leftarrow Q \cup \{f'\}$ 

```

If $f_u \not\rightarrow_* f_s$, the proposed restriction does not result in the loss of solution S . Therefore, we can further assume that $f_u \rightarrow_* f_s$. Hence, $\forall w \in \text{cons}(W), f_s[w] \geq f_u[w]$. Moreover, there necessarily exists a constraint $w_i \in \text{cons}(W, M)$ such that $f_s[w_i] > f_u[w_i]$, otherwise f_s will still be unsatisfiable. Let f' be the front defined by $f'[w_i] = f_u[w_i] + 1$ and $\forall w \in \text{cons}(W)$ such that $w \neq w_i, f'[w] = f_u[w]$. f' is a front which is generated by the algorithm and by construction $f' \rightarrow_* f_s$. ■

cost	tuples	cost	tuples	cost	tuples
100	$\{(x,c)\}$	100	default	100	$\{(y,c)\}$
10	$\{(x,b)\}$	5	$\{(x,c),(y,a)\}$	10	$\{(y,b)\}$
0	$\{(x,a)\}$	0	$\{(x,a),(y,b)\}$	0	$\{(y,a)\}$
	(a) w_x		(b) w_{xy}		(c) w_y

Figure 3. Example where it is required to enumerate all relaxations of a MUC to reach the optimum.

Figure 3 presents an example where one cannot simply relax a MUC in only one way, even if it was locally the optimal relaxation, without losing the optimal solution. In this example, a WCN is composed of three constraints w_x, w_{xy} and w_y . The first front selects the layers having a cost equal to 0, and therefore only retains $\{(x,a)\}, \{(x,a),(y,b)\}$ and $\{(y,a)\}$ as allowed tuples. Of course, in this case, the constraints w_{xy} and w_y constitute a MUC. There are two ways to relax this first MUC. The first relaxation switches to the next layer of w_{xy} (it is locally the best choice since it has a cost of 5) whereas the second one switches to the next layer of w_y . Considering the first case,

Function incompleteSearch(W : WCN)

```

1 foreach  $w \in cons(W)$  do
2    $f[w] \leftarrow 0$ 
3 repeat
4    $P \leftarrow toCN_{\leq}(W, f)$ 
5    $sol \leftarrow solveCN(P)$ 
6   if  $sol \neq \perp$  then
7     return  $sol$ 
8   else
9      $M \leftarrow extractMUC(P)$ 
10     $W' \leftarrow restrict(W, M)$ 
11     $f \leftarrow relax(W', f)$ 
12 until  $sol \neq \perp$ 

```

the new front obtained retains $\{(x, a)\}$, $\{(x, c), (y, a)\}$ and $\{(y, a)\}$ as allowed tuples. This time, the MUC contains the constraints w_x and w_{xy} . We can either relax w_{xy} to obtain a solution with a cost of 100, or relax w_x to obtain another MUC that can be relaxed in two different ways to obtain either a solution having a cost of 105, or a solution having a cost of 110.

The second relaxation of the first MUC consists in switching to the next layer of w_y . In this case, we directly obtain a solution having a cost of 10, which is the optimum. This clearly shows that all relaxations of a MUC are required.

2) *Greedy Approach*: In this part, we present an incomplete version of our approach based on the identification and relaxation of MUCs in order to solve WCNs.

From a WCN W , Function incompleteSearch returns a solution for a CN derived from W . First the structure *front* is initialized to 0 which corresponds to the first layer of each constraint of the WCN.

From a WCN W and a front f , we extract a CN and this network is then solved. If a solution is found, then it is returned by the function incompleteSearch (line 7). If the network has been proved unsatisfiable, it is necessary to relax some constraints of the WCN. To achieve this, the algorithm extracts a WCN W' from a computed MUC (lines 9 and 10). The front f is then updated by Function relax (line 11). This procedure will relax one (or several constraints) of W' in order to break the MUC associated with W' . This process loops until a solution is found for the constraint network associated with W .

Function $relax(W, f)$ updates the front f in order to allow new layers. The idea is to increment $f[w]$ for at least one constraint w , in such a way as to make the MUC satisfiable. In other words, for at least one of the constraints, we accept an increase of cost for this constraint. In the greedy approach used here, we keep doing this until the MUC is broken.

Algorithm relax uses a local priority queue Q_{loc} which is

Function relax(W : WCN, f : front): front

```

1  $Q_{loc} \leftarrow \{f\}$ 
2 while  $Q_{loc} \neq \emptyset$  do
3    $f \leftarrow$  pick and delete least cost front in  $Q_{loc}$ 
4    $P \leftarrow toCN_{\leq}(W, f)$ 
5   if  $solveCN(P) \neq \perp$  then
6     return  $f$ 
7   else
8      $M \leftarrow extractMUC(P)$ 
9     foreach  $w \in cons(W, M)$  do
10    if  $f[w] \neq nbLayers(w) - 1$  then
11     $f' \leftarrow f$ 
12     $f'[w] \leftarrow f'[w] + 1$ 
13    if  $cost(f') \neq k$  then
14     $Q_{loc} \leftarrow Q_{loc} \cup \{f'\}$ 

```

initialized with the front f . While Q_{loc} is not empty (line 2), we extract from the queue the front f having the lowest cost $cost(f)$. Function $toCN_{\leq}$ builds a constraint network from both the WCN and the front f previously selected (line 4). If this constraint network is satisfiable, f initially given in parameter is updated and returned by the algorithm relax. On the contrary, if the constraint network is unsatisfiable, it means that the relaxation is not sufficient and then the algorithm updates the queue of fronts by generating all the direct successors of f . For each constraint belonging to the WCN (obtained from the MUC), we generate a new front, different from the initial one, by incrementing the layer of one single constraint.

Note that in the different calls to the function relax, it is sufficient to work only on a subset of the constraints. Indeed, the constraints that can be relaxed are those belonging to the WCN W associated with the MUC ($restrict(W, M)$). In practice, we only work with the subset of the front which corresponds to constraints of the MUC, which saves space (for simplicity, this is not detailed in the algorithm).

In the incomplete version, we transform a WCN into a CN using function $toCN_{\leq}$ instead of $toCN_{=}$ which is used in the complete version. Indeed, in the complete version, all possible relaxations of a MUC are considered and the fronts are enumerated in increasing cost order. Hence, when we test the satisfiability of a CN associated with a front f , we know that all the CN associated with fronts f' such that $f' \rightarrow_* f$ have already been proved unsatisfiable. Hence, in the complete version, $toCN_{=}$ extracts a CN composed of the current layers of a front. In the incomplete version, all possible relaxations of a MUC are not tried. Therefore, we cannot ensure that, when we test the satisfiability of a CN associated with a front f , all the CNs associated with fronts f' such that $f' \rightarrow_* f$ have been tried and already proved

unsatisfiable. Function $toCN_{\leq}$ therefore extracts a CN from both the current and inferior layers of a front.

This approach doesn't guarantee the identification of an optimum solution. The reason is that all the relaxations of MUCs are not considered, contrary to the complete version (see Example 3). Indeed, we only identify the first relaxation which restores the satisfiability of a MUC.

VI. EXPERIMENTAL RESULTS

In order to show the practical interest of our approach, we have performed experiments using a computer with processors Intel(R) Core(TM) i7-2820QM CPU 2.30GHz. Our greedy approach, noted *GMR*, has been compared with two complete approaches with cost transfer algorithms enforcing EDAC, proposed by both our solver *AbsCon* and the solver *ToulBar2*. A time-out of 600 seconds was set per instance. The total CPU time necessary to solve each instance is given as well as the upper bound found

Instances		AbsCon		ToulBar2
		GMR	EDAC	EDAC
spot5-42	CPU	9.18	> 600	> 600
	UB	161,050	161,050	161,050
spot5-404	CPU	4.99	> 600	217
	UB	118	114	114
spot5-408	CPU	9.85	> 600	> 600
	UB	6,235	8,238	6,240
spot5-412	CPU	18.8	> 600	> 600
	UB	33,403	43,390	37,399
spot5-414	CPU	37.7	> 600	> 600
	UB	40,500	56,492	52,492
spot5-503	CPU	6.33	> 600	> 600
	UB	12,125	13,119	12,117
spot5-505	CPU	12	> 600	> 600
	UB	22,266	28,258	25,268
spot5-507	CPU	22.3	> 600	> 600
	UB	30,417	37,429	37,420
spot5-509	CPU	32.2	> 600	> 600
	UB	37,469	48,475	46,477
spot5-1401	CPU	76.5	> 600	> 600
	UB	483,109	513,097	516,095
spot5-1403	CPU	142.5	> 600	> 600
	UB	481,266	517,260	507,265
spot5-1407	CPU	552.6	> 600	> 600
	UB	492,614	517,623	507,633
spot5-1502	CPU	3.91	0.98	0.02
	UB	28,044	28,042	28,042
spot5-1504	CPU	67.6	> 600	> 600
	UB	175,311	204,314	198,318
spot5-1506	CPU	338	> 600	> 600
	UB	378,551	426,551	399,568

Table I
CPU TIME IN SECONDS AND BOUND FOUND BEFORE THE TIME-OUT FOR INSTANCES *spot5*.

(UB). If the execution of the algorithm is not yet finished before the time-out (CPU time greater than 600 seconds), the bound found at the end of the 600 seconds is given. Table I provides the results obtained for the serie *spot5* (except trivial instances) composed of constraints of arity less than or equal to 3. Table II provides the results obtained for the serie *celar*.

On instances of *spot5*, we observe that our approach provides better bounds than those obtained by the two complete approaches, except for the instances *spot5-404*, *spot5-503* and *spot5-1502*. Beyond the fact that our approach is not complete, this can be explained by the fact that these instances are relatively easy (involving less than 1,400 constraints) and cost transfer algorithms reach a better bound before the time-out. However, it is interesting to note that the bound found by our approach is not so different from these obtained by complete approaches, and this in a very short time. About other instances, considered as more difficult (in particular, more than 13,000 constraints for the instances

Instances		AbsCon		ToulBar2
		GMR	EDAC	EDAC
graph-05	CPU	16.6	> 600	0.62
	UB	221	4,645	221
scen-06	CPU	88.5	> 600	485.4
	UB	3,616	12,013	3,389
scen-06-16	CPU	60.6	> 600	237.7
	UB	4,149	11,286	3,277
scen-06-18	CPU	59.2	> 600	102.4
	UB	3,640	8,723	3,263
scen-06-20	CPU	68.5	> 600	67.9
	UB	3,402	8,594	3,163
scen-06-30	CPU	32.6	177.7	1.10
	UB	2,208	2,080	2,080
scen-07	CPU	209.9	> 600	> 600
	UB	426,423	31,230K	505,731
scen-07_10K	CPU	28.7	> 600	> 600
	UB	270K	17,000K	1,500K
celar6-sub2	CPU	23.3	> 600	4.74
	UB	2,927	2,746	2,746
celar6-sub3	CPU	26.6	> 600	15.7
	UB	3,271	3,279	3,079
celar6-sub4	CPU	41.5	> 600	28.7
	UB	3,704	5,178	3,230
celar7-sub2	CPU	60.8	> 600	17.8
	UB	283,955	252,436	173,252
celar7-sub3	CPU	48.7	> 600	93.4
	UB	414,161	1,342K	203,460
celar7-sub4	CPU	57	> 600	221
	UB	272,945	302,541	242,443

Table II
CPU TIME IN SECONDS AND BOUND FOUND BEFORE THE TIME-OUT FOR INSTANCES *celar*.

spot5-1403, *spot5-1407* and *spot5-1506*), we observe that our greedy approach provides a better bound than those of complete approaches in a time shorter than the time-out.

One can note that *GMR* is not as competitive as *ToulBar2* on some instances of *celar*, *scen* and *graph*. This is partly due to the fact that *ToulBar2* uses a constraint decomposition technique. Indeed, comparing *GMR* with a classical cost transfer algorithm (such as in our *AbsCon* solver), the approach described in this paper outperforms the EDAC version of *AbsCon*. However, on the *graph-05* instance, even if *GMR* is not so fast than *ToulBar2*, our approach is able to find the optimum upper bound.

The preliminary experiments we have conducted on the complete approach show that our current implementation is not yet competitive. It clearly appears that MUCs must be identified incrementally so that the computational effort for a given front benefits to the others. We are currently exploring the use of dynamic CSP algorithms to incrementally check the satisfiability of CNs and identify MUCs.

VII. CONCLUSION

In this paper, we have proposed an original approach for solving weighted constraint networks through successive resolutions of hard constraint networks. More precisely, CNs are obtained from WCNs by selecting tuples with a given cost. These CNs are enumerated in increasing cost order until a solution is found. To improve the complexity in practice of our approach, we identify a minimal unsatisfiable core for each unsatisfiable constraint network, in order to focus the relaxation on the sole constraints of the MUC. The approach is derived in both a complete and an incomplete, greedy algorithm. We have shown that our method based on MUCs extraction can be used in practice: the greedy algorithm obtains results which are comparable with other state of the art approaches. However, the complete algorithm is not yet competitive but we have identified several reasons. We are working on a new version using dynamic CSP algorithms in order to perform incremental computations. To conclude, we would like to emphasize that the proposed method allows a simple and natural integration of the CSP, (weighted) Max-CSP and WCSP frameworks. This can be done by generalizing in a straightforward way the construction of CNs generated during the resolution.

ACKNOWLEDGMENTS

This work has been supported by both CNRS and OSEO within the ISI project 'Pajero'.

REFERENCES

- [1] J. Larrosa, "Node and arc consistency in weighted CSP," in *Proceedings of AAAI'02*, 2002, pp. 48–53.
- [2] J. Larrosa and T. Schiex, "Solving weighted CSP by maintaining arc consistency," *Artificial Intelligence*, vol. 159, no. 1-2, pp. 1–26, 2004.
- [3] M. Cooper, "Reduction operations in fuzzy or valued constraint satisfaction," *Fuzzy Sets and Systems*, vol. 134, no. 3, pp. 311–342, 2003.
- [4] S. de Givry, F. Heras, M. Zytnicki, and J. Larrosa, "Existential arc consistency: Getting closer to full arc consistency in weighted CSPs," in *Proc. of IJCAI'05*, 2005, pp. 84–89.
- [5] M. Cooper, S. de Givry, M. Sanchez, T. Schiex, and M. Zytnicki, "Virtual arc consistency for weighted CSP," in *Proceedings of AAAI'08*, 2008, pp. 253–258.
- [6] M. Cooper, S. de Givry, M. Sanchez, T. Schiex, M. Zytnicki, and T. Werner, "Soft arc consistency revisited," *Artificial Intelligence*, vol. 174, no. 7-8, pp. 449–478, 2010.
- [7] A. Morgado, F. Heras, M. Liffiton, J. Planes, and J. Marques-Silva, "Iterative and core-guided MaxSAT solving: A survey and assessment," *Constraints*, pp. 1–57, 2013.
- [8] C. Ansotegui, M. Bonet, and J. Levy, "A new algorithm for weighted partial MaxSAT," in *Proc. of AAAI'10*, 2010, pp. 3–8.
- [9] Z. Fu and S. Malik, "On solving the partial MaxSAT problem," in *Proceedings of SAT'06*, 2006, pp. 252–265.
- [10] J. Marques-Silva and J. Planes, "Algorithms for maximum satisfiability using unsatisfiable cores," in *DATE'08*, 2008, pp. 408–413.
- [11] R. Dechter, *Constraint processing*. Morgan Kaufmann, 2003.
- [12] C. Lecoutre, *Constraint networks: techniques and algorithms*. ISTE/Wiley, 2009.
- [13] F. Rossi, P. van Beek, and T. Walsh, Eds., *Handbook of Constraint Programming*. Elsevier, 2006.
- [14] S. Wieringa, "Understanding, improving and parallelizing MUS finding using model rotation," in *Proc. of CP'12*, 2012, pp. 672–687.
- [15] J. de Siqueira and J. Puget, "Explanation-based generalisation of failures," in *Proceedings of ECAI'88*, 1988, pp. 339–344.
- [16] U. Junker, "QuickXplain: preferred explanations and relaxations for over-constrained problems," in *Proceedings of AAAI'04*, 2004, pp. 167–172.
- [17] F. Hemery, C. Lecoutre, L. Sais, and F. Boussemart, "Extracting MUCs from constraint networks," in *Proceedings of ECAI'06*, 2006, pp. 113–117.
- [18] S. Bistarelli, U. Montanari, F. Rossi, T. Schiex, G. Verfaillie, and H. Fargier, "Semiring-based CSPs and valued CSPs: Frameworks, properties, and comparison," *Constraints*, vol. 4, no. 3, pp. 199–240, 1999.
- [19] P. Meseguer, F. Rossi, and T. Schiex, "Soft constraints," in *Handbook of Constraint Programming*. Elsevier, 2006, ch. 9, pp. 281–328.
- [20] C. Lecoutre, N. Paris, O. Roussel, and S. Tabary, "Propagating soft table constraints," in *Proceedings of CP'12*, Quebec city, Canada, 2012, pp. 390–405.