Soft Constraint Problems with Incompleteness

Mirco Gelain, Maria Silvia Pini, Francesca Rossi, and Kristen Brent Venable

Dipartimento di Matematica Pura ed Applicata, Università di Padova, Italy
E-mail: {mgelain,mpini,frossi,kvenable}@math.unipd.it

Abstract. We consider soft constraint problems where some of the preferences may be unspecified. This models, for example, situations with several agents providing the data, or with possible privacy issues. In this context, we study how to find an optimal solution without having to wait for all the preferences. In particular, we define an algorithm to find a solution which is necessarily optimal, that is, optimal no matter what the missing data will be, with the aim to ask the user to reveal as few preferences as possible. Experimental results show that in many cases a necessarily optimal solution can be found by eliciting a small number of preferences.

1 Introduction

Traditionally, tasks such as scheduling, planning, and resource allocation have been tackled using several techniques, among which constraint reasoning is one of the winning ones: the task is represented by a set of variables, their domains, and a set of constraints, and a solution of the problem is an assignment to all the variables in their domains such that all constraints are satisfied. Preferences or objective functions have been used to extend such scenario and allow for the modelling of constraint optimization, rather than satisfaction, problems. However, what is common to all these approaches is that the data (variables, domains, constraints) are completely known before the solving process starts.

On the contrary, the increasing use of web services and in general of multi-agent applications demands for the formalization and handling of data that is only partially known when the solving process works, and that can be added later, for example via elicitation. In many web applications, data may come from different sources, which may provide their piece of information at different times. Also, in multi-agent settings, data provided by some agents may be voluntarily hidden due to privacy reasons, and only released if needed to find a solution to the problem.

Recently, some lines of work have addressed these issues by allowing for open settings in CSPs: both open CSPs [7, 9] and interactive CSPs [13] work with domains that can be partially specified, and in dynamic CSPs [6] variables, domains, and constraints may change over time. It has been shown that these approaches are closely related. In fact, interactive CSPs can be seen as a special case of both dynamic and open CSPs [12].

Here we consider the same issues but we focus on constraint optimization problems rather than CSPs, thus looking for an optimal solution rather than any solution. In

particular, we consider problems where constraints are replaced by soft constraints, in which each assignment to the variables of the constraint has an associated preference coming from a preference set [1]. In this setting, for the purpose of this paper we assume that variables, domains, and constraint topology are given at the beginning, while the preferences can be partially specified and possibly added during the solving process.

Constraint optimization has also been considered in the context of open CSPs in a cost-minimization setting [8]. However, the incompleteness considered is on which values belong to domains, and not on the preference of tuples.

There are several application domains where such setting is useful. One regards the fact that quantitative preferences, and needed in soft constraints, may be difficult and tedius to provide for a user. Another one concerns multi-agent settings, where agents agree on the structures of the problem by they may provide their preferences on different parts of the problem at different times. Finally, some preferences can be initially hidden because of privacy reasons.

Formally, we take the soft constraint formalism when preferences are totally ordered and we allow for some preferences to be left unspecified. Although some of the preferences can be missing, it could still be feasible to find an optimal solution. If not, then we ask the user and we start again from the new problem with some added preferences.

More precisely, we consider two notions of optimal solution: *possibly optimal* solutions are assignments to all the variables that are optimal in *at least one way* in which currently unspecified preferences can be revealed, while *necessarily optimal* solutions are assignments to all the variables that are optimal in *all ways* in which currently unspecified preferences can be revealed. This notation comes from multi-agent preference aggregation [14], where, in the context of voting theory, some preferences are missing but still one would like to declare a winner.

Given an incomplete soft constraint problem (ISCSP), its set of possibly optimal solutions is never empty, while the set of necessarily optimal solutions can be empty. Of course what we would like to find is a necessarily optimal solution, to be on the safe side: such solutions are optimal regardless of how the missing preferences would be specified. However, since such a set may be empty, in this case there are two choices: either to be satisfied with a possibly optimal solution, or to ask users to provide some of the missing preferences and try to find, if any, a necessarily optimal solution of the new ISCSP. In this paper we follow this second approach, and we repeat the process until the current ISCSP has at least one necessarily optimal solution.

To achieve this, we employ an approach based on branch and bound which first checks whether the given problem has a necessarily optimal solution (by just solving the completion of the problem where all unspecified preferences are replaced by the worst preference). If not, then finds the most promising among the possibly optimal solutions (where the promise is measured by its preference level), and asks the user to reveal the missing preferences related to such a solution. This second step is then repeated until the current problem has a necessarily optimal solution.

We implemented our algorithm and we tested it against classes of randomly generated binary fuzzy ISCSPs, where, beside the usual parameters (number of variables, domain size, density, and tightness) we added a new parameters measuring the percentage of unspecified preferences in each constraint and domain. The experimental results

show that in many cases a necessarily optimal solution can be found by eliciting a small amount of preferences. In particular, in no case the percentage of elicited preferences exceeded 35% of the total number of missing preferences.

2 Soft constraints

A soft constraint [1] is just a classical constraint [5] where each instantiation of its variables has an associated value from a (totally or partially ordered) set. This set has two operations, which makes it similar to a semiring, and is called a c-semiring. More precisely, a c-semiring is a tuple $\langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$ such that: A is a set, called the carrier of the c-semiring, and $\mathbf{0}, \mathbf{1} \in A$; $+$ is commutative, associative, idempotent, $\mathbf{0}$ is its unit element, and $\mathbf{1}$ is its absorbing element; $\times$ is associative, commutative, distributes over $+$, $\mathbf{1}$ is its unit element and $\mathbf{0}$ is its absorbing element.

Consider the relation $\leq_S$ over A such that $a \leq_S b$ iff $a + b = b$. Then: $\leq_S$ is a partial order; $+$ and $\times$ are monotone on $\leq_S$; $\mathbf{0}$ is its minimum and $\mathbf{1}$ its maximum; $\langle A, \leq_S \rangle$ is a lattice and, for all $a, b \in A$, $a + b = lub(a, b)$. Moreover, if $\times$ is idempotent, then $\langle A, \leq_S \rangle$ is a distributive lattice and $\times$ is its glb.

Informally, the relation $\leq_S$ gives us a way to compare (some of the) tuples of values and constraints. In fact, when we have $a \leq_S b$, we will say that b *is better than* a. Thus, $\mathbf{0}$ is the worst value and $\mathbf{1}$ is the best one.

Given a c-semiring $S = \langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$, a finite set D (the domain of the variables), and an ordered set of variables V, a constraint is a pair $\langle def, con \rangle$ where $con \subseteq V$ and $def : D^{|con|} \to A$. Therefore, a constraint specifies a set of variables (the ones in con), and assigns to each tuple of values of D of these variables an element of the semiring set A. A soft constraint satisfaction problem (SCSP) is just a set of soft constraints over a set of variables.

Many known classes of satisfaction or oprimization problem can be cast in this formalism. A classical CSP is just an SCSP where the chosen c-semiring is: $S_{CSP} = \langle \{false, true\}, \vee, \wedge, false, true \rangle$. On the other hand, fuzzy CSPs [15, 11] can be modeled in the SCSP framework by choosing the c-semiring: $S_{FCSP} = \langle [0, 1], max, min, 0, 1 \rangle$. For weighted CSPs, the semiring is $S_{WCSP} = \langle \Re^+, min, +, +\infty, 0 \rangle$. Here preferences are interpreted as costs from 0 to $+\infty$, which are combined with the sum and compared with min. Thus the optimization criterion is to minimize the sum of costs. For probabilistic CSPs [10], the semiring is $S_{PCSP} = \langle [0, 1], max, \times, 0, 1 \rangle$. Here preferences are interpreted as probabilities ranging from 0 to 1, which are combined using the product and compared using max. Thus the aim is to maximize the joint probability.

Given an assignment t to all the variables of an SCSP, we can compute its preference value $pref(t)$ by combining the preferences associated by each constraint to the subtuples of the assignments referring to the variables of the constraint. More precisely, $pref(P, s) = \Pi_{\langle idef, con \rangle \in C} def(s_{\downarrow con})$, where Π refers to the $\times$ operation of the semiring and $s_{\downarrow con}$ is the projection of tuple s on the variables in con.

For example, in fuzzy CSPs, the preference of a complete assignment is the minimum preference given by the constraints. In weighted constraints, it is instead the sum of the costs given by the constraints.

An optimal solution of an SCSP is then a complete assignment t such that there is no other complete assignment t'' with $pref(t) <_S pref(t'')$. The set of optimal solutions of an SCSP P will be written as $Opt(P)$.

3 Incomplete Soft Constraint Problems (ISCSPs)

Informally, an incomplete SCSP, written ISCSP, is an SCSP where the preferences of some tuples in the constraints, and/or of some values in the domains, are not specified. In detail, given a set of variables V with finite domain D, and c-semiring $S = \langle A, +, \times, 0, 1 \rangle$ with a totally ordered carrier, we extend the SCSP framework to incompleteness by the following definitions.

Definition 1 (incomplete soft constraint). *Given a set of variables V with finite domain D, and a c-semiring $\langle A, +, \times, 0, 1 \rangle$, an incomplete soft constraint is a pair $\langle idef, con \rangle$ where $con \subseteq V$ is the scope of the constraint and $idef : D^{|con|} \longrightarrow A \cup \{?\}$ is the preference function of the constraint. All tuples mapped into $?$ by $idef$ are called incomplete tuples.*

In an incomplete soft constraint, the preference function can either specify the preference value of a tuple by assigning a specific element from the carrier of the c-semiring, or leave such preference unspecified. Formally, in the latter case the associated value is $?$. A soft constraint is a special case of an incomplete soft constraint where all the tuples have a specified preference.

Example 1. Assume a travel agency is planning Alice and Bob's honeymoon, having only some information on their preferences. Alice and Bob live in the US, the candidate destinations are the Maldive islands and the Caribbean, and they can decide to go by ship or by plane. Going to Maldives by ship would take very long, thus they have a high preference to go by plane and a low preference to go by ship. Also, they have been told that a cruise in the Caribbean is very nice. Thus they have a high preference to go there by ship. However, they don't have any preference on going there by plane. We can model this scenario by using the fuzzy c-semiring $\langle [0, 1], max, min, 0, 1 \rangle$, two variables T (standing for $Transport$) and D (standing for $Destination$) with domains $D(T) = \{p, sh\}$ (p stands for $plane$ and sh for $ship$) and $D(D) = \{m, c\}$ (m stands for $Maldives$, c for $Caribbean$), and an incomplete soft constraint $\langle idef, con \rangle$ with $con = \{T, D\}$ and with preference function as shown in Figure 1. The only incomplete tuple is (p, c).

Definition 2 (incomplete soft constraint problem (ISCSP)). *An incomplete soft constraint problem is a pair $\langle C, V, D \rangle$ where C is a set of incomplete soft constraints over the variables in V with domain D. Given an ISCSP P, we will denote with $IT(P)$ the set of all incomplete tuples in P.*

Definition 3 (completion). *Given an ISCSP P, a completion of P is an SCSP P' obtained from P by associating to each incomplete tuple in every constraint an element of the carrier of the c-semiring. A completion is partial if some preference remains unspecified. We will denote with $C(P)$ the set of all possible completions of P and with $PC(P)$ the set of all its partial completions.*

Example 2. Consider again Example 1. Assume that for the considered season the Maldives are slightly preferrable to the Caribbean. Also, Alice and Bob don't mind travelling by plane and have never travelled by ship. Thus they have a reasonably high
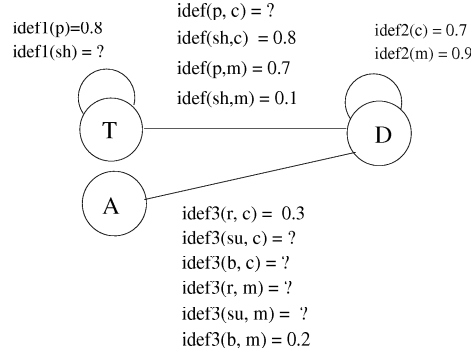
Fig. 1. An ISCSP.

preference to plane as a way of transport, while they don't give any preference to ship. Moreover, as far as accommodations, which can be in a standard room, or a suite, or a bungalow, assume that a suite in the Maldives is too expensive for the young couple while a standard room in the Caribbean is not special enough for a honeymoon. To model this new information we use a variable A (standing for *Accommodation*) with domain $D(A) = \{r, su, b\}$ (r stands for *room*, su for *suite* and b for *bungalow*), and three constraints: two unary incomplete soft constraints, $\langle idef1, \{T\} \rangle$, $\langle idef2, \{D\} \rangle$ and a binary incomplete soft constraint $\langle idef3, \{A, D\} \rangle$. Their definition is shown in Figure 1. The set of incomplete tuples of this problem is $IT(P) = \{(sh), (p, c), (su, c), (su, m), (r, m), (b, c)\}$.

Definition 4 (preference of an assignment). *Given an ISCSP $P = \langle C, V, D \rangle$ and an assignment s to all its variables we denote with $pref(P, s)$ the preference of s in P. In detail, $pref(P, s) = \Pi_{<idef,con> \in C | idef(s_{\downarrow con}) \neq ?} idef(s_{\downarrow con})$.*

The preference of an assignment s in an incomplete problem is thus obtained by combining the known preferences associated to the projections of the assignment, that is, of the appropriated subtuples in the constraints. The projections which have unspecified preferences are simply ignored. Given an assignment s to all the varibales, the set of its projections with unspecified preference is denoted by $it(s)$.

Example 3. Consider the two assignments $s_1 = (p, m, b)$ and $s_2 = (p, m, su)$, we have that $pref(P, s_1)) = min(0.8, 0.7, 0.9, 0.2) = 0.2$, while $pref(P, s_2) = min(0.8, 0.7, 0.9) = 0.7$. However, while the preference of s_1 is fixed, since none of its projections is incomplete, the preference of s_2 may become lower that 0.7 depending on the preference of the incomplete tuple (su, m).

As shown by the example, the presence of incompleteness generates a partition of the set of assignments into two sets: those which have a certain preference which is independent of how incompleteness is resolved, and those whose preference is only an upperbound, in the sense that it can be lowered in some completions.

Given an ISCSP P, we will denote the first set of assignments as $Fixed(P)$ and the second with $Unfixed(P)$. In Example 3, $Fixed(P) = \{s_1\}$, while all other assignments belong to $Unfixed(P)$.

In SCSPs we have that an assignment is an optimal solution if its global preference is undominated. This notion can be generalized to the incomplete setting. In particular, when some preferences are unknown, we will speak of necessarily and possibly optimal solutions, that is, assignments which are undominated in all (resp., some) completions.

Definition 5 (necessarily and possibly optimal solution). *Given an ISCSP $P = \langle C, V, D \rangle$, an assignment $s \in D^{|V|}$ is a necessarily (resp, possibly) optimal solution iff $\forall Q \in C(P)$ (resp., $\exists Q \in C(P)$ such that) $\forall s' \in D^{|V|}$, $pref(Q, s') \not> pref(Q, s)$.*

Given an ISCSP P, we will denote with $NOS(P)$ (resp., $POS(P)$) the set of necessarily (resp., possibly) optimal solutions of P. Notice that, while $POS(P)$ is never empty, in general $NOS(P)$ may be empty. In particular, $NOS(P)$ is empty whenever the available preferences do not allow to determine the relation between an assignment and all the others.

Example 4. In the ISCSP P of Figure 1, we can easily see that $NOS(P) = \emptyset$ since, given any assignment, it is possible to construct a completion of P in which it is not an optimal solution. On the other hand, $POS(P)$ contains all assignments not including tuple (sh, m). In fact, such a tuple has preference 0.1 and it drowns the preference of any assignment containing it below the preference of fixed solution $s_1 = (p, m, b)$ (i.e., below 0.2). Thus, in all completions, s_1 dominates any such assignment. Instead, for any assignment s not including (sh, m), we can complete P in order to make s optimal, for example by setting the preferences of all the incomplete tuples of s to $pref(s, P)$ and the preferences of all other incomplete tuples to 0.

4 Characterizing POS(P) and NOS(P)

In this section we investigate how to characterize the set of necessarily and possibly optimal solutions of an ISCSP given the preferences of the optimal solutions of two of the completions of P. In particular, given an ISCSP P defined on c-semiring $\langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$, we consider:

- the SCSP $P_0 \in C(P)$, called the **0**-completion of P, obtained from P by associating preference $\mathbf{0}$ to each tuple of $IT(P)$.
- the SCSP $P_1 \in C(P)$, called the **1**-completion of P, obtained from P by associating preference $\mathbf{1}$ to each tuple of $IT(P)$.

Let us indicate respectively with $pref_0$ and $pref_1$ the preference of an optimal solution of P_0 and P_1. Due to the monotonicity of $\times$, and since $\mathbf{0} \leq \mathbf{1}$, we have that $pref_0 \leq pref_1$.

In the following theorem we will show that, if $pref_0 > \mathbf{0}$, there is a necessarily optimal solution of P iff $pref_0 = pref_1$, and in this case $NOS(P)$ coincides with the set of optimal solutions of P_0.

Theorem 1. *Given an ISCSP P and the two completions $P_0, P_1 \in C(P)$ as defined above, if $pref_0 > \mathbf{0}$ we have that:*

- *$NOS(P) \neq \emptyset$ iff $pref_1 = pref_0$;*

– *if* $NOS(P) \neq \emptyset$ *then* $NOS(P) = Opt(P_0)$.

Proof. Since we know that $pref_0 \leq pref_1$, if $pref_0 \neq pref_1$ then $pref_1 > pref_0$. We prove that, if $pref_1 > pref_0$, then $NOS(P) = \emptyset$. Let us consider any assignment s of P. Due to the monotonicity of $\times$, for all $P' \in C(P)$, we have $pref(P', s) \leq pref(P_1, s) \leq pref_1$.

- If $pref(P_1, s) < pref_1$, then s is not in $NOS(P)$ since P_1 is a completion of P where s is not optimal.
- If instead $pref(P_1, s) = pref_1$, then, since $pref_1 > pref_0$, we have $s \in Unfixed(P)$. Thus we can consider completion P_1' obtained from P_1 by associating preference $\mathbf{0}$ to the incomplete tuples of s. In P_1' the preference of s is $\mathbf{0}$ and the preference of an optimal solution of P_1' is, due to the monotonicity of $\times$, at least that of s in P_0, that is $pref_0 > \mathbf{0}$. Thus $s \notin NOS(P)$.

Next we consider when $pref_0 = pref_1$. Clearly $NOS(P) \subseteq Opt(P_0)$, since any assignment which is not optimal in P_0 is not in $NOS(P)$. We will show that $NOS(P) \neq \emptyset$ by showing that any $s \in Opt(P_0)$ is in $NOS(P)$. Let us assume, on the contrary, that there is $s \in Opt(P_0)$ such that $s \notin NOS(P)$. Thus there is a completion P' of P with an assignment s' with $pref(P', s') > pref(P', s)$. By construction of P_0, any assignment $s \in Opt(P_0)$ must be in $Fixed(P)$. In fact, if it had some incomplete tuple, its preference in P_0 would be $\mathbf{0}$, since $\mathbf{0}$ is the absorbing element of $\times$. Since $s \in Fixed(P)$, $pref(P', s) = pref(P_0, s) = pref_0$. By construction of P_1 and monotonicity of $\times$, we have $pref(P_1, s') \geq pref(P', s')$. Thus the contradiction $pref_1 \geq pref(P_1, s') \geq pref(P', s') > pref(P', s) = pref_0$. This allows us to conclude that $s \in NOS(P) = Opt(P_0)$. $\square$

In the theorem above we have assumed that $pref_0 > \mathbf{0}$. The case in which $pref_0 = \mathbf{0}$ needs to be treated separately. We consider it in the following theorem.

Theorem 2. *Given ISCSP* $P = \langle C, V, D \rangle$ *and the two completions* $P_0, P_1 \in C(P)$ *as defined above, assume* $pref_0 = \mathbf{0}$. *Then:*

- *if* $pref_1 = \mathbf{0}$, $NOS(P) = D^{|V|}$;
- *if* $pref_1 > \mathbf{0}$, $NOS(P) = \{s \in Opt(P_1) | \forall s' \in D^{|V|}$ *with* $pref(P_1, s') > \mathbf{0}$ *we have* $it(s) \subseteq it(s')\}$.

The formal proof is omitted for lack of space. However, we give the informal intuition. In words, the theorem above says that, if $pref_0 = \mathbf{0}$ and $pref_1 > \mathbf{0}$, then an assignment is a necessarily optimal solution only if it is optimal in P_1 and if the set of its incomplete tuples is contained in the set of incomplete tuples of all other assignments in $Unfixed(P)$. Intuitively, if some assignment s' has an incomplete tuple which is not part of another assignment s, then we can make s' dominate s in a completion by setting all the incomplete tuples of s' to $\mathbf{1}$ and all the remaining incomplete tuples of s to $\mathbf{0}$. In such a completion s is not optimal. Thus s is not a necessarily optimal solution.

However, if the tuples of s are a subset of the incomplete tuples of all other assignments then it is not possible to lower s without lowering all other tuples even further. This means that s is a necessarily optimal solution.

We now turn our attention to possible optimal solutions. Given a c-semiring $\langle A, +, \times, \mathbf{0}, \mathbf{1} \rangle$, it has been shown in [2] that either the $\times$ operator is idempotent or it is strictly monotonic. In the following two theorems we show that such a distinction on the c-semiring plays a key role in the characterization of $POS(P)$ where P is an ISCSP.

In particular, if $\times$ is idempotent, then the possible optimal solutions are the assignments with preference in P between $pref_0$ and $pref_1$. If, instead, $\times$ is strictly monotonic, then the possibliy optimal solutions have preference in P between $pref_0$ and $pref_1$ and dominate all the assignments which have as set of incomplete tuples a subset of their incomplete tuples.

Theorem 3. *Given an ISCSP P defined on a c-semiring with idempotent $\times$ and the two completions $P_0, P_1 \in C(P)$ as defined above, if $pref_0 > \mathbf{0}$ we have that: $POS(P) = \{s \in D^{|V|} | pref_0 \leq pref(P, s) \leq pref_1\}$.*

Proof. First we show that any s such that $pref_0 \leq pref(P, s) \leq pref_1$ is in $POS(P)$. Let us consider the completion of P, P', obtained by associating preference $pref(P, s)$ to all the incomplete tuples of s and $\mathbf{0}$ to all other incomplete tuples of P. For any other assignment s' we can show that it never dominates s:

- $s' \in Fixed(P)$ and thus $pref(P', s') = pref(P_0, s') \leq pref_0 \leq pref(P, s)$;
- $s' \in Unfixed(P)$ and
  - $it(s') \not\subseteq it(s)$, then $pref(P', s') = \mathbf{0}$ since in P' the incomplete tuples in $it(s')$ which are not in $it(s)$ have been associated to preference $\mathbf{0}$;
  - $it(s') \subseteq it(s)$. By construction of P' and since $\times$ is idempotent and associative we have that: $pref(P', s) = (pref(P, s) \times (\Pi_{|it(s)|} pref(P, s))) = pref(P, s)$ and $pref(P', s') = (pref(P, s') \times (\Pi_{|it(s')|} pref(P, s))) = pref(P, s') \times pref(P, s)$. Since $\times$ is intensive, $pref(P', s') = (pref(P, s') \times pref(P, s)) \leq pref(P, s) = pref(P', s)$.

Thus in P' no assignment dominates s. This means that $s \in POS(P)$.

We will now show that if s $in POS(P)$, $pref_0 \leq pref(P, s) \leq pref_1$. By construction of P_1 and due to monotonicity of $\times$ we have that there is no assignment s with $pref(P, s) > pref_1$. Thus, to conclude the proof of the theorem, we only need to show that any assignment s such that $pref(P, s) < pref_0$ is not in $POS(P)$. First let us notice that $\forall P' \in C(P)$ and for any assignment s we have that $pref(P', s) \leq pref(P, s)$ due to the intensive property of $\times$. Let us now consider any assignment s' such that $pref(P_0, s') = pref_0$. By construction of P_0, any assignment $s' \in Opt(P_0)$ must be in $Fixed(P)$, otherwise if it had some incomplete tuple its preference in P_0 would be $\mathbf{0}$, since $\mathbf{0}$ is the absorbing element of $\times$. Since $s' \in Fixed(P)$, $pref(P', s') = pref(P_0, s') = pref_0, \forall P' \in C(P)$. Thus, $\forall P' \in C(P), pref(P', s) \leq pref(P, s) < pref_0 = pref(P', s')$. This shows that if $pref(P, s) < pref_0$ then $s \notin POS(P)$. $\square$

Theorem 4. *Given an ISCSP P defined on a c-semiring with a strictly monotonic $\times$ and the two completions $P_0, P_1 \in C(P)$ as defined above, if $pref_0 > \mathbf{0}$ we have that: an assignment $s \in POS(P)$ iff $pref_0 \leq pref(P, s) \leq pref_1$ and $pref(P, s) = max\{ pref(P, s') | it(s') \subseteq it(s)\}$.*

Proof. Let us first show that if assignment s is such that $pref_0 \le pref(P,s) \le pref_1$ and $pref(P,s) = max\{pref(P,s')|it(s') \subseteq it(s)\}$ it is in $POS(P)$. We must show there is a completion of P where s is undominated. Let us consider completion P' obtained associating preference **1** to all the tuples in $it(s)$ and **0** to all the tuples in $IT(P)\backslash it(s)$. First we notice that $pref(P',s) = pref(P,s)$, since **1** is the unit element of $\times$. Let us consider any other assignment s'. Then we have one of the following:

- $it(s') = \emptyset$, which means that $s' \in Fixed(P)$ and thus $pref(P',s') = pref(P_0,s')$ $\le pref_0 \le pref(P,s) = pref(P',s)$;
- $it(s') \not\subseteq it(s)$, which means that there is at least one incomplete tuple of $it(s')$ which is associated to **0**. Since **0** is the absorbing element of $\times$, $pref(P',s') = \mathbf{0}$ and thus $pref(P',s') < pref_0 \le pref(P',s)$;
- $it(s') \subseteq it(s)$, in this case $pref(P',s') = pref(P,s')$ since all tuples in $it(s')$ are associated to **1** in P'. However since $pref(P,s) = max\{pref(P,s')|it(s') \subseteq it(s)\}$, $pref(P',s') \le pref(P',s)$.

We can thus conclude that s is not dominated by any assignment in P'. Hence $s \in POS(P)$.

Let us now prove the other direction by contradiction. If $pref(P,s) < pref_0$ then we can conclude as in the previous proof. We must prove that if $pref_0 \le pref(P,s) \le pref_1$ and $pref(P,s) < max\{pref(P,s')|it(s') \subseteq it(s)\}$ then s is not in $POS(P)$. In any completion P' of P we have that $pref(P',s) = pref(P,s)\times it\text{-}pref(P',s)$ and $pref(P',s') = pref(P,s')\times it\text{-}pref(P',s')$ where $it\text{-}pref(P',s)$ (resp. $it\text{-}pref(P',s'$)) is the combination of the preferences associated to the incomplete tuples in $it(s)$ (resp. $it(s')$). Since $it(s') \subseteq it(s)$, for any completion P' we have that $it\text{-}pref(P',s) \le it\text{-}pref(P',s')$. Moreover, let s'' be such that $pref(P,s'') = max\{pref(P,s')|it(s') \subseteq it(s)\}$. Then we have that for any completion P', $pref(P',s'') > pref(P',s)$ since $pref(P,s'') > pref(P,s)$ and $it\text{-}pref(P',s'') \ge it\text{-}pref(P',s)$ and $\times$ is strictly monotonic. Thus, if $pref_0 \le pref(P,s) \le pref_1$ and $pref(P,s) < max\{pref(P,s')| it(s') \subseteq it(s)\}$, then s is not in $POS(P)$. $\square$

In contrast to $NOS(P)$, when $pref_0 = \mathbf{0}$ we can immediately conclude that $POS(P) = D^{|V|}$, independently of the nature of $\times$, since all assignments are optimal in P_0.

Corollary 1. *Given an ISCSP $P = \langle C, V, D \rangle$, if $pref_0 = \mathbf{0}$, then $POS(P) = D^{|V|}$.*

The results given in this section can be summarized as follows:

- when $pref_0 = \mathbf{0}$
    - not enough information to compute $NOS(P)$ (by Theorem 2);
    - $POS(P) = D^{|V|}$ (by Corollary 1);
- when $pref_0 = pref_1 = \mathbf{0}$
    - $NOS(P) = D^{|V|}$ (by Theorem 2);
    - $POS(P) = D^{|V|}$ (by Corollary 1) ;
- when $\mathbf{0} = pref_0 < pref_1$
    - $NOS(P) = \{s \in Opt(P_1)|\forall s' \in D^{|V|}$ with $pref(P_1,s') > \mathbf{0}$ we have $it(s) \subseteq it(s')\}$ (by Theorem 2);

- $POS(P) = D^{|V|}$ (by Corollary 1);
  - when $0 < pref_0 = pref_1$
    - $NOS(P) = Opt(P_0)$ (by Theorem 1);
    - if $\times$ is idempotent: $POS(P) = \{s \in D^{|V|} | pref_0 \leq pref(P,s) \leq pref_1\}$ (by Theorem 3);
    - if $\times$ is strictly monotonic: $POS(P) = \{s \in D^{|V|} | pref_0 \leq pref(P,s) \leq pref_1, pref(P,s) = max\{ pref(P,s') | it(s') \subseteq it(s)\}\}$ (by Theorem 4);
  - when $0 < pref_0 < pref_1$
    - $NOS(P) = \emptyset$ (by Theorem 1);
    - $POS(P)$ as for the case when $0 < pref_0 = pref_1$.

5 A solver for ISCSPs

We want to find a necessarily optimal solution of the given problem, if it exists. In most cases, however, the available information will only allow to determine the set of possibly optimal solutions. In such cases, preference elicitation is needed to discriminate among such assignments in order to determine a necessarily optimal one of the new problem with the elicited preferences. In this section we describe an algorithm, called *Find-NOS*, to achieve this task.

Algorithm 1: Find-NOS

Input: an ISCSP P
Output: an ISCSP Q, an assignment s, a preference p
$P_0 \leftarrow P[?/0]$
$s_0, pref_0 \leftarrow BB(P_0, -)$
$s_1 \leftarrow s_0$
$pref_1 \leftarrow pref_0$
$s_{max} \leftarrow s_0$
$pref_{max} \leftarrow pref_0$
repeat
    $P_1 \leftarrow P[?/1]$
    if $pref_1 > pref_{max}$ **then**
        $s_{max} \leftarrow s_1$
        $pref_{max} \leftarrow pref_1$
    $s_1, pref_1 \leftarrow BB(P_1, pref_{max})$
    if $s_1 \neq nil$ **then**
        $S \leftarrow it(s_1)$
        $P \leftarrow \text{Elicit}(P,S)$
        $pref_1 \leftarrow pref(P,s_1)$
until $s_1 \neq nil$;
return $P, s_{max}, pref_{max}$

Algorithm *Find-NOS* takes in input an ISCSP P over a totally ordered c-semiring and returns an ISCSP Q which is a partial completion of P, and an assignment $s \in$

$NOS(Q)$ together with its preference p. Given an ISCSP P, *Find-NOS* first checks if $NOS(P)$ is not empty, and, if so, it returns P, $s \in NOS(P)$, and its preference. If instead $NOS(P) = \emptyset$, it starts eliciting the preferences of some incomplete tuples.

In detail, *Find-NOS* first computes the **0**-completion of P, written as $P[?/0]$, called P_0, and applies Branch and Bound (BB) to it. This allows to find an optimal solution of P_0, say s_0, and its preference $pref_0$.

In our notation, the application of the BB procedure has two parameters: the problem to which it is applied, and the starting bound. When BB is applied without a starting bound, we will write $BB(P, -)$. When the BB has finished, it returns a solution and its preference. If no solution is found, we assume that the returned items are both nil.

Variables s_1 and $pref_1$ (resp., s_{max} and $pref_{max}$) represent the optimal solution and the corresponding preference of the **1**-completion of the current problem (written $P[?/1]$) (resp., the best solution and the corresponding preference found so far). At the beginning, such variables are initialized to s_0 and $pref_0$.

The main loop of the algorithm, achieved through the **repeat** command, computes the 1-completion, denoted by P_1, of the current problem. In the first iteration the condition of the first **if** is not satisfied since $pref_1 = pref_{max} = pref_0$. The execution thus proceeds by applying BB to P_1 with bound $pref_{max} = pref_0 \geq \mathbf{0}$. This allows us to find an optimal solution of P_1 and its corresponding preference, assigned to s_1 and $pref_1$. If BB fails to find a solution, s_1 is nil. Thus the second **if** is not executed and the algorithm exits the loop and returns P, $s_{max} = s_0$, and $pref_{max} = pref_0$.

If instead BB applied to P_1 with bound $pref_{max}$ does not fail, then we have that $pref_0 < pref_1$. Now the algorithm elicits the preference of some incomplete tuples, via procedure *Elicit*. This procedure takes an ISCSP and a set of tuples of variable assignments, and asks the user to provide the preference for such tuples, returning the updated ISCSP. The algorithm calls procedure *Elicit* over the current problem P and the set of incomplete tuples of s_1 in P. After elicitation, the new preference of s_1 is computed and assigned to $pref_1$.

Since $s_1 \neq nil$, a new iteration begins, and BB is applied with initial bound given by the best preference between $pref_1$ and $pref_{max}$. Moreover, if $pref_1 > pref_{max}$, then s_{max} and $pref_{max}$ are updated to always contain the best solution and its preference. Iteration continues until the elicited preferences are enough to make BB fail to find a solution with a better preference w.r.t. the previous application of BB. At that point, the algorithm returns the current problem and the best solution found so far, together with its preference.

Theorem 5. *Given an ISCSP P in input, algorithm* Find-NOS *always terminates and returns an ISCSP Q such that $Q \in PC(P)$, an assignment $s \in NOS(Q)$, and its preference in Q.*

Proof. At each iteration, either $pref_{max}$ increases or, if it stays the same, a new solution will be found since after elicitation the preference of s_1 has not increased. Thus, either $pref_{max}$ is so high that BB doesn't find any solution, or all the optimal solutions have been considered. In both cases the algorithm exits the loop.

At the end of its execution, the algorithm returns the current partial completion of given problem and a solution s_{max} with the best preference seen so far $pref_{max}$. The

repeat command is exited when $s_1 = nil$, that is, when $BB(P[?/1], pref_{max})$ fails. In this situation, $pref_{max}$ is the preference of an optimal solution of the **0**-completion of the current problem P. Since BB fails on $P[?/1]$ with such a bound, by monotonicity of the $\times$ operator, $pref_{max}$ is also the preference of an optimal solution of $P[?/1]$. By Theorems 1 and 2, we can conclude that $NOS(P)$ is not empty. If $pref_{max} = \mathbf{0}$, then $NOS(P)$ contains all the assignments and thus also s_0. The algorithm correctly returns the same ISCSP given in input, assignment s_0 and its preference $pref_0 = \mathbf{0}$. If instead $\mathbf{0} < pref_{max}$, again the algorithm is correct, since by Theorem 1 we know that $NOS(P) = Opt(P[?/0])$, and since $s_{max} \in Opt(P[?/0])$. $\square$

Notice also that the algorithm performs preference elicitation only on solutions which are possibly optimal in the current partial completion of the given problem (and thus also in the given problem). In fact, by Theorems 3 and 4, any optimal solution of the 1-completion of the current partial completion Q is a possibly optimal solution of Q. Thus no useless work is done to elicit preferences related to solutions which cannot be necessarily optimal for any partial completion of the given problem. This also means that our algorithm works independently of the properties of the $\times$ operator.

6 Experimental setting and results

We have implemented Algorithm *Find-NOS* in Java and we have tested it on randomly generated ISCSPs with binary constraints and based on the Fuzzy c-semiring. To generate such problems, we use the following parameters:

- n: number of variables;
- m: cardinality of the domain of each variable;
- d: density of the constraints, that is, the percentage of binary constraints present in the problem w.r.t. the total number of possible binary constraints that can be defined on n variables;
- t: tightness, that is, the percentage of tuples with preference 0 in each constraint, w.r.t. the total number of tuples (m^2 since we have only binary constraints), and in each domain;
- i: incompleteness, that is, the percentage of incomplete tuples (formally, tuples with preference ?) in each constraint and in each domain.

For example, if the generator is given in input $n = 10$, $m = 5$, $d = 50$, $t = 10$, and $i = 30$, it will generate a binary ISCSP with 10 variables, each with 5 elements in the domain, 22 constraints on a total of $45 = n(n-1)/2$, 2 tuples with preference 0 and 7 incomplete tuples over a total of 25 in each constraint, and 1 missing preference in each domain.

We have generated classes of ISCSPs by varying one parameter at a time, and fixing the other ones. The varying parameters are the number of variables, the density, and the incompleteness. When the number of variables varies (from $n = 5$ to $n = 20$, with step 3), we set $m = 5$, $d = 50$, $t = 10$, and $i = 30$. When we vary the density (from $d = 10$ to $d = 80$ with step 5), we set $n = 10$, $m = 5$, $t = 10$, and $i = 30$. Finally, when we vary the incompleteness (from $i = 10$ to $i = 80$ with step 5), we set $n = 10$, $m = 5$, $d = 50$, and $t = 10$.

In all the experiments, we have measured the number of tuples elicited by Algorithm *Find-NOS*. We also show the percentage of elicited tuples over the total number of incomplete tuples of the problem in input. For each fixed value of all the parameters, we show the average of the results obtained for 50 different problem instances, each given in input to *Find-NOS* 10 times. This setting is necessary since we have two kinds of randomness: the usual one in the generation phase and a specific one when eliciting preferences.


Fig. 2. Number and percentage of elicited preferences, as a function of the number of variables. Fixed parameters: $m = 5$, $d = 50$, $t = 10$, $i = 30$.

Figure 2 shows the absolute number and the percentage of elicited preferences when the number of variables varies. As expected, when the number of variables increases, the absolute number of elicited preferences increases as well, since there is a growth of the total number of incomplete tuples. However, if we consider the percentage of elicited tuples, we see that it is not affected by the increase in the number of variables. In particular, the percentage of elicited preferences remains stable around 22%, meaning that, regardless of the number of variables, the agent is asked to reveal only 22 preferences over 100 incomplete tuples. A necessarily optimal solution can be thus found leaving 88% of the missing preferences unrevealed.

Similar results are obtained when density varies (see Figure 3). We can see that the absolute number of elicited preferences grows when density increases. The maximum number of elicited preferences reached is however lower that the maximum reached when varying the variables (see Figure 2(a)). The reason for this is that the largest problems considered when varying the number of variables have more incomplete tuples than the largest obtained when varying the density. In fact, a problem with $n = 20$, given the fixed parameters, has around 685 incomplete tuples, 165 of which (about 22%) are elicited. On the other hand, a problem with $d = 80$, given the fixed parameters, has around 262 incomplete tuples, 55 (about 22%) of which are elicited. This is coherent with the fact that the results on the percentage of elicited preferences when varying the density and the number of variables are very similar.

The last set of experiments vary the percentage of incompleteness (see Figure 4). As for density and number of variables, the absolute number of elicited preferences grows when the percentage of incompleteness increases. The maximum number of elicited

Fig. 3. Number and percentage of elicited preferences, as a function of the density. Fixed parameters: $n = 10$, $m = 5$, $t = 10$, $i = 30$.

Fig. 4. Number and percentage of elicited preferences, as a function of the incompleteness. Fixed parameters: $n = 10$, $m = 5$, $t = 10$, $i = 30$.

preferences reached is close to that reached when varying the variables. However, the number of incomplete tuples of the problems with $i = 70$ is around 460 and thus smaller than that of problems with $n = 20$. Thus the percentage of elicited preferences is larger in problems with $i = 70$. This is confirmed by the corresponding result for the percentage of elicited preferences, which is shown to be around 35%. Additionally, the percentage of elicited preferences follows a slightly increasing trend as the percentage of incompleteness in the problem grows. However, it maintains itself below 35%, which means that in the worst case, where 70% of the tuples are incomplete, we are able to find a necessary optimal solution leaving 46% of the total number of tuples unspecified.

These experimental results show that it is indeed possible to find a necessarily optimal solution while forcing the user to reveal only a small percentage of the missing preferences. This is very promising both in terms of elicitation-related costs and also when concerned with privacy issues.

7 Conclusions and future work

We consider problems modelled via soft constraints with totally ordered and possibly unspecified preferences, and propose to solve them via an approach based on systematic search. Experimental results show that a small amount of preferences has to be revelead before being able to find an optimal solution.

Future work should consider partially ordered preferences and also other ways to express preferences, such as qualitative ones a la CP nets [4, 3], as well as other kinds of missing data, such as those considered in dynamic, interactive, and open CSPs. Moreover, other solving approaches can be considered, such as those based on local search rather than systematic search.

References

1. S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based constraint solving and optimization. *Journal of the ACM*, 44(2):201–236, mar 1997.
2. S. Bistarelli, U. Montanari, F. Rossi, T. Schiex, G. Verfaillie, and H. Fargier. Semiring-based csps and valued csps: Frameworks, properties, and comparison. *Constraints*, 4(3), 1999.
3. C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.
4. C. Boutilier, R. I. Brafman, H. H. Hoos, and D. Poole. Reasoning with conditional ceteris paribus preference statements. In *UAI*, pages 71–80, 1999.
5. R. Dechter. *Constraint processing*. Morgan Kaufmann, 2003.
6. R. Dechter and A. Dechter. Belief maintenance in dynamic constraint networks. In *AAAI*, pages 37–42, 1988.
7. B. Faltings and S. Macho-Gonzalez. Open constraint satisfaction. In *CP*, volume 2470 of *Lecture Notes in Computer Science*, pages 356–370. Springer, 2002.
8. B. Faltings and S. Macho-Gonzalez. Open constraint optimization. In *CP*, volume 2833 of *Lecture Notes in Computer Science*, pages 303–317. Springer, 2003.
9. B. Faltings and S. Macho-Gonzalez. Open constraint programming. *Artif. Intell.*, 161(1-2):181–208, 2005.
10. H. Fargier and J. Lang. Uncertainty in constraint satisfaction problems: a probalistic approach. In M. Clarke, R. Kruse, and S. Moral, editors, *Symbolic and Quantitative Approaches to Reasoning and Uncertainty, European Conference, ECSQARU'93, Granada, Spain, November 8-10, Proceedings*, volume 747 of *Lecture Notes in Computer Science*, pages 97–104. Springer, 1993.
11. H. Fargier, T. Schiex, and G. Verfaille. Valued Constraint Satisfaction Problems: Hard and Easy Problems. In *IJCAI-95*, pages 631–637. Morgan Kaufmann, 1995.
12. S. Macho González, C. Ansótegui, and P. Meseguer. On the relation among open, interactive and dynamic csp. In *The Fifth Workshop on Modelling and Solving Problems with Constraints (IJCAI'05)*, 2005.
13. E. Lamma, P. Mello, M. Milano, R. Cucchiara, M. Gavanelli, and M. Piccardi. Constraint propagation and value acquisition: Why we should do it interactively. In *IJCAI*, pages 468–477, 1999.
14. J. Lang, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Winner determination in sequential majority voting. In *IJCAI*, pages 1372–1377, 2007.
15. Z. Ruttkay. Fuzzy constraint satisfaction. In *Proceedings 1st IEEE Conference on Evolutionary Computing*, pages 542–547, Orlando, 1994.