Encoding Partial Constraint Satisfaction in the Semiring-based Framework for Soft Constraints*

Stefano Bistarelli
Istituto di Informatica e Telematica,
CNR, Pisa, Italy
stefano.bistarelli@iit.cnr.it
Dipartimento di Scienze,
Universitá degli Studi "G. D'annunzio",
Pescara, Italy
bista@sci.unich.it

Eugene C. Freuder and Barry O'Sullivan
Cork Constraint Computation Centre
Department of Computer Science,
University College Cork, Ireland
{e.freuder|b.osullivan}@4c.ucc.ie

Abstract

The partial constraint satisfaction paradigm focuses on solving relaxations of problems that either do not admit solutions, or that are either impractical or impossible to solve completely. The semiring-based framework for soft constraints is a unifying model for a variety of extensions of the constraint satisfaction formalism. For example, the semiring-based framework can represent weighted, fuzzy, probabilistic and set-based constraint satisfaction problems. In this paper, we discuss how the semiring-based framework for soft constraints can be used to model partial constraint satisfaction problems. We show how the semiring framework can be used to capture a notion of distance between a solution and a problem based on the known distance metrics used in the partial constraint satisfaction literature. These solution-problem distance metrics can be seen as providing lower-bounds on the distance between a problem and its relaxation.

1. Introduction

Over the past 30 years constraint satisfaction has become an important paradigm in Artificial Intelligence [9]. Informally, a constraint satisfaction problem (CSP) is defined by a set of variables, each of which has a corresponding set of possible values called its domain, and the task is to find a value for each variable from its domain so that the set of constraints are satisfied. Constraint satisfaction is applicable to a wide variety of problems arising in scheduling, design, configuration, machine vision, temporal reasoning and planning [13].

Many extensions to the CSP paradigm have been proposed. In the context of this paper two of these extensions are particularly relevant: partial constraint satisfaction [6, 7] and the semiring-based framework for soft constraints [1, 3, 4].

Partial constraint satisfaction involves finding values for a subset of the variables that satisfy a subset of the constraints [7]. It can also be viewed as involving the weakening of a problem in order to accept more value combinations. Essentially, in partial constraint satisfaction we focus on relaxing a problem so that a satisfactory solution can be found. As we shall see in Section 2, three distance metrics were defined in the original paper on partial constraint satisfaction [7] that define various ways in which a problem can be relaxed[1]. The most popular of these distance metrics is known as Max-CSP: find an assignment of values to the variables that maximizes the number of constraints that are satisfied [8, 10, 14].

Some work already exists that relates the partial constraint satisfaction paradigm to soft constraints. Schiex *et al.* have defined a mapping for the Max-CSP distance metric to the valued constraint satisfaction framework for soft constraints [12]. We could easily combine this mapping with the known relationship between valued constraint satisfaction and the semiring-based framework, described in [2], to obtain a possible mapping from Max-CSPs to Soft-CSPs. However, this just gives us one form of partial constraint

* This work has received support from Enterprise Ireland under their Basic Research Grant Scheme (Grant Number SC/02/289) and their International Collaboration Programme (Grant Number IC/2003/88). Eugene C. Freuder is supported by Science Foundation Ireland (Grant Number 00/Pl.1/C075).

1 Other distance measures are possible of course, but in this paper we only consider those that were discussed explicitly in the original paper [7].

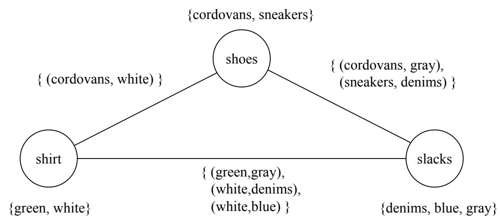

Figure 1. Example of an over-constrained problem.

satisfaction. In this paper we present a complete and uniform treatment of the three distance metrics underpinning partial constraint satisfaction, as presented in the literature.

2. Partial Constraint Satisfaction

To introduce partial constraint satisfaction, consider the problem presented in Figure 1, known as the Robot Dressing Problem [7]. The nodes in this graph represent the three variables – $shoes$, $shirt$ and $slacks$ – representing the items of clothing that must be chosen. Each node is also labeled with a set of values for the corresponding variable, i.e. the domain of each variable. The arcs that connect the variables are labeled with the legal combinations of values for each of the variables, i.e. the constraints between the variables. By analysis, we can see that this problem does not admit any solutions, i.e. it is over-constrained. Therefore, we need to find a "pragmatic" solution to this problem, which may involve tolerating at least one constraint being violated; or buying some more shoes, shirts or slacks; or accepting that certain clothing combinations are acceptable. This is the typical problem we face in partial constraint satisfaction. The various options we mentioned above for overcoming this problem are forms of constraint relaxation.

Taking a more formal view, Freuder and Wallace [7] define a partial constraint satisfaction problem as a triple: $\langle (P,U),(PS,\leq),(M,(N,S))\rangle$. Informally, the problem is defined in terms of an initial constraint satisfaction problem P, a set of "universes" U representing the possible values for each of the variables in P, a problem space PS containing P and a partial-order over problems, M is a distance metric on that space, and (N,S) are necessary and sufficient solution distances between P and an acceptable solution.

A *solution* to a partial constraint satisfaction problem can be defined as a problem P' from the problem space PS along with a solution to that problem where the distance of P' from P is less than N. A solution is sufficient if the distance is less than or equal to S. An *optimal solution* is one where the distance of P' from P is minimal over the problem space. Therefore, partial constraint satisfaction can be viewed as search through a problem space.

The partial-order defined over the problem space PS is defined in terms of the set of solutions to problems. Specifically, $P_1 \leq P_2$ iff $sols(P_1) \supseteq sols(P_2)$, where $sols(P)$ denotes the set of solutions to problem P. We can read $P_1 \leq P_2$ as "P_1 is obtained by weakening the constraints in P_2". When a problem is weakened, we mean that the constraints in the problem allow more consistent assignments and, as a consequence, the set of solutions may increase.

The manner in which a problem is weakened depends on the distance metric, M, that is used. A number of metrics have been proposed in the literature: solution subset distance, augmentation distance and Max-CSP distance. We shall discuss each briefly in turn below.

Solution Subset Distance – This distance metric is defined as the number of solutions not shared between problems P and P'. When $P' \leq P$ this metric reflects the number of solutions that have been added in order to relax problem P.

Augmentation Distance – This distance metric is slightly different to solution subset distance. In fact, this distance metric counts the number of constraint values that are not shared by problems P and P'. This represents the number of augmentations to the constraints in problem P that are required to reach its relaxation P'.

Max-CSP Distance – This is the most well-studied distance metric of the three. It involves finding a solution that violates the minimum number of constraints in the problem. The metric is normally defined as the number of constraints that are violated

The objective of relaxing a problem is to find more solutions. However, relaxing a problem does not always give rise to more solutions. It is quite possible that many relaxations of a problem do not admit solutions.

3. Soft Constraints

Several formalizations of the concept of *soft constraints* are currently available. In the following, we refer to the one based on c-semirings [1, 3, 4], which can be shown to generalize and express many of the others [2]. A soft constraint may be seen as a constraint where each instantiation of its variables has an associated value from a partially ordered set which can be interpreted as a set of preference values. Combining constraints will then have to take into account such additional values, and thus the formalism has also to provide suitable operations for combination ($\times$) and comparison ($+$) of tuples of values and constraints. This is why this formalization is based on the concept of c-semiring, which is just a set plus two operations.

Semirings. A semiring is a tuple $\langle A, +, \times, \mathbf{0}, \mathbf{1}\rangle$ such that: 1. A is a set and $\mathbf{0}, \mathbf{1} \in A$; 2. $+$ is commutative, associative and $\mathbf{0}$ is its unit element; 3. $\times$ is associative, distributes over $+$, $\mathbf{1}$ is its unit element and $\mathbf{0}$ is its absorbing element. A c-semiring is a semiring $\langle A, +, \times, \mathbf{0}, \mathbf{1}\rangle$ such that: $+$ is idempotent, $\mathbf{1}$ is its absorbing element and $\times$ is commutative. Let us consider the relation $\leq_S$ over A such that $a \leq_S b$ iff $a + b = b$. Then it is possible to prove that (see [3]): 1. $\leq_S$ is a partial order; 2. $+$ and $\times$ are monotone on $\leq_S$; 3. $\mathbf{0}$ is its minimum and $\mathbf{1}$ its maximum; 4. $\langle A, \leq_S\rangle$ is a complete lattice and, for all $a, b \in A$, $a + b = lub(a, b)$ (where lub is the *least upper bound*). Moreover, if $\times$ is idempotent, then: $+$ distributes over $\times$; $\langle A, \leq_S\rangle$ is a complete distributive lattice and $\times$ its glb (*greatest lower bound*). Informally, the relation $\leq_S$ gives us a way to compare semiring values and constraints. In fact, when we have $a \leq_S b$, we will say that *b is better than a*. In the following, when the semiring will be clear from the context, $a \leq_S b$ will be often indicated by $a \leq b$.

Constraint Problems. Given a semiring $S = \langle A, +, \times, \mathbf{0}, \mathbf{1}\rangle$ and an ordered set of variables V over a finite domain D, a *constraint* is a function which, given an assignment $\eta : V \to D$ of the variables, returns a value of the semiring. By using this notation we define $\mathcal{C} = \eta \to A$ as the set of all possible constraints that can be built starting from S, D and V.

Note that in this *functional* formulation, each constraint is a function (as defined in [4]) and not a pair (as defined in [3]). Such a function involves all the variables in V, but it depends on the assignment of only a finite subset of them. So, for instance, a binary constraint $c_{x,y}$ over variables x and y, is a function $c_{x,y} : V \to D \to A$, but it depends only on the assignment of variables $\{x, y\} \subseteq V$. We call this subset the *support* of the constraint. More formally, consider a constraint $c \in \mathcal{C}$. We define its support as $supp(c) = \{v \in V \mid \exists \eta, d_1, d_2 . c\eta[v := d_1] \neq c\eta[v := d_2]\}$, where

$$\eta[v := d]v' = \begin{cases} d & \text{if } v = v', \\ \eta v' & \text{otherwise.} \end{cases}$$

Note that $c\eta[v := d_1]$ means $c\eta'$ where η' is η modified with the assignment $v := d_1$ (that is the operator $[\ ]$ has precedence over application). Note also that $c\eta$ is the application of a constraint function $c : V \to D \to A$ to a function $\eta : D \to A$; what we obtain, is a semiring value $c\eta = a$.

A *soft constraint satisfaction problem* (SCSP) is a pair $\langle C, con\rangle$ where $con \subseteq V$ and C is a set of constraints: con is the set of variables of interest for the constraint set C, which however may concern also variables not in con. Note that a classical CSP is a SCSP where the chosen c-semiring is: $S_{CSP} = \langle\{false, true\}, \vee, \wedge, false, true\rangle$. Fuzzy CSPs [5, 11] can instead be modeled in the SCSP framework by choosing the c-semiring $S_{FCSP} = \langle[0, 1], max, min, 0, 1\rangle$. Many other "soft" CSPs (Probabilistic, weighted, ...) can be modeled by using a suitable semiring structure ($S_{prob} = \langle[0, 1], max, \times, 0, 1\rangle$, $S_{weight} = \langle\mathcal{R}, min, +, +\infty, 0\rangle$, ...).

Figure 2 shows the graph representation of a fuzzy CSP. Variables and constraints are represented respectively by nodes and by undirected (unary for c_1 and c_3 and binary for c_2) arcs, and semiring values are written to the right of the corresponding tuples. The variables of interest (that is the set con) are represented with a double circle. Here we assume that the domain D of the variables contains only elements a and b and c.

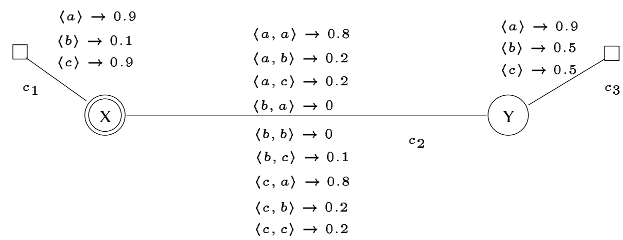

Figure 2. A fuzzy CSP.

Combining and projecting soft constraints. Given the set $\mathcal{C}$, the combination function $\otimes : \mathcal{C} \times \mathcal{C} \to \mathcal{C}$ is defined as $(c_1 \otimes c_2)\eta = c_1\eta \times_S c_2\eta$. Informally, combining two constraints means building a new constraint whose support involves all the variables of the original ones, and which associates with each tuple of domain values for such variables a semiring element which is obtained by multiplying the elements associated by the original constraints to the appropriate sub-tuples. It is easy to verify that $supp(c_1 \otimes c_2) \subseteq supp(c_1) \cup supp(c_2)$.

Given a constraint $c \in \mathcal{C}$ and a variable $v \in V$, the *projection* of c over $V - \{v\}$, written $c \Downarrow_{(V - \{v\})}$ is the constraint c' s.t. $c'\eta = \sum_{d \in D} c\eta[v := d]$.

Solutions. A *solution* of an SCSP $P = \langle C, con\rangle$ is the constraint $Sol(P) = (\bigotimes C) \Downarrow_{con}$. That is, we combine all constraints, and then project over the variables in con. In this way we get the constraint with support (not greater than) con which is "induced" by the entire SCSP. Note that when all the variables are of interest we do not need to perform any projection.

For example, the solution of the fuzzy CSP of Figure 2 associates a semiring element to every domain value of variable x. Such an element is obtained by first combining all the constraints together. For instance, for the tuple $\langle a, a\rangle$ (that is, $x = y = a$), we have to compute the minimum between 0.9 (which is the value assigned to $x = a$ in constraint c_1), 0.8 (which is the value assigned to $\langle x = a, y = a\rangle$ in c_2) and 0.9 (which is the value for $y = a$ in c_3). Hence, the resulting value for this tuple is 0.8. We can do the same work for tuple $\langle a, b\rangle \to 0.2$, $\langle a, c\rangle \to 0.2$, $\langle b, a\rangle \to 0$,

$\langle b, b\rangle \rightarrow 0$, $\langle b, c\rangle \rightarrow 0.1$, $\langle c, a\rangle \rightarrow 0.8$, $\langle c, b\rangle \rightarrow 0.2$ and $\langle c, c\rangle \rightarrow 0.2$. The obtained tuples are then projected over variable x, obtaining the solution $\langle a\rangle \rightarrow 0.8$, $\langle b\rangle \rightarrow 0.1$ and $\langle c\rangle \rightarrow 0.8$.

Solutions are constraints in themselves and can be ordered by extending the $\leq_S$ order. We say that a constraint c_1 is at least as constraining as constraint c_2 if $c_1 \sqsubseteq_S c_2$, where for any assignment η of variables then

$$c_1 \sqsubseteq_S c_2 \equiv c_1\eta \leq_S c_2\eta$$

Thus, if $c_1 \sqsubseteq_S c_2$ holds, then constraint c_1 may be thought of as a more restrictive replacement of constraint c_2, and c_2 as a relaxation of c_1.

4. From PCSPs to Semiring-based CSPs

One of the main difficulties we have to deal with when translating from the partial constraint satisfaction framework to the semiring framework is handling the different approaches to the reasoning about overconstraindness that are used in both paradigms. Below we shall compare the different approaches and describe how they can be reconciled.

Different measures of over-constrainedness – Freuder and Wallace define a notion of metric *between constraint problems* and a way to relax/weaken a problem. As discussed above, Freuder and Wallace in [7] note that each way can be seen as adding additional tuples to constraints.

In the semiring framework Bistarelli, Montanari and Rossi defined a $+$ operator in the semiring that can be used to compare instantiations of constraints (using $\leq_S$ or $\sqsubseteq_S$). In [1, 3] the $\sqsubseteq_S$ order has also been used to compare problems, by saying that $P_1 \sqsubseteq_S P_2$ if $\bigotimes C_1 \sqsubseteq \bigotimes C_2$ (when P_1 has constraints C_1 and P_2 has constraints C_2), where P_2 is a relaxation of P_1.

Representing PCSP distance in the SCSP semiring – The partial constraint satisfaction framework deals with crisp CSPs plus a notion of partial ordering amongst problems in terms of their solutions. If we relax problem P_1 we obtain a problem P_2 that is strictly better with respect to the distance defined in the PCSP. The key idea is to use the softness levels of semiring-based SCSPs to represent how much we relax a problem. Since the measure of "how much" we have relaxed a problem depends on the PCSP distance metric used, we have to use potentially different mappings from PCSP to SCSP and also different semirings to represent the different metrics of distance we wish to use.

Collecting Solutions added due to Relaxation – Recall that the PCSP framework deals with crisp CSPs. Therefore, the semiring level that is associated with each solution must not be regarded as a measure of the quality of a solution. Instead, the level associated with each complete assignment represents how far it is from the solutions of the initial crisp problem. This gives us a lower-bound on the distance between the original problem and its relaxation that has this particular complete assignment as a solution. Furthermore, based upon the distance between solution and problem, a distance metric between problems can be defined.

If we start from a crisp problem P that has some set of solutions, we would expect to find the same solutions in the corresponding SCSP problem P' with an associated semiring value $\mathbf{1}$, where $\mathbf{1}$ is the best element of the semiring. The semirings that we will use in the rest of the section will have 0 as the best level. Using this semiring, the higher the semiring level, the worse the solution will be with respect to the distance from the original problem. Solutions with semiring level worse than $\mathbf{1}$ (that is greater than 0) represent solutions that can be added to the problem only if it is relaxed. The level of the semiring for a given solution represents the amount of relaxation (distance) we have to perform in order to obtain this solution. If we want to collect all the solutions whose distance is within a given bound α, we just need to consider all solutions better than the bound α.

Below we define semiring-based mappings for the various metrics defined by Freuder and Wallace [7].

4.1. Max-CSP Distance

Freuder and Wallace defined the Max-CSP distance as the process of maximising the number of constraints that are satisfied by a solution [6, 7]. This notion can be easily cast in the semiring framework since this distance relation is already defined in terms of solutions (and not among problems). Therefore, maximizing the number of satisfied constraints is equivalent to minimizing the number of the violated ones. Following this idea we can map a crisp problem P with a set of constraints C to an SCSP problem P' with constraints C' over the semiring $S_{WCSP} = \langle \mathbb{N}^+, min, +, +\infty, 0\rangle$[2]. We can map each constraint c in P to a corresponding soft constraint c' in P' over the same variables such that $c'\eta = 0$ when the tuple η satisfies the constraint c, and $c'\eta = 1$ when the tuple η violates it.

Combining the constraints in the SCSP P' we obtain its solution. A solution η of the original problem P will lead to a solution of the SCSP P' with semiring level 0, while non-solutions of the original problem will lead to semiring levels higher (worse) than 0. If a solution η violates k constraints, the solution of the SCSP for η (that is $\bigotimes C'\eta$) will have semiring level k.

Theorem 1 *Consider a crisp CSP P and the corresponding SCSP P' over the semiring S_{WCSP} =*

2 This structure, with real instead of natural numbers, has already been proven to be a semiring in [1, 3].

$\langle \mathbb{N}^+, min, +, +\infty, 0\rangle$, *obtained using the mapping described above. Then,*

- *η is a solution of P iff η has semiring level 0 in P';*
- *η is violating k constraints in P iff η has semiring level k in P'.*

If we want to relax the initial problem P by a quantity α, that is we can afford to violate α constraints in solutions to the problem, this is captured in the SCSP P' by looking for all solutions with a semiring level that is better (lower) than α.

4.2. Solution Subset Distance

Freuder and Wallace defined the solution subset distance between two comparable problems P_1 and P_2, according to $\leq_S$, as the number of solutions of P_2 that are not solutions of P_1.

This distance metric compares two problems and describes how far the solutions of them are from each other. Before translating this distance into the SCSP framework we have to reformulate this distance between problems as a distance between solution and problem. As with the previous distance metric we want to associate with solutions that are already in the original, unrelaxed, problem the semiring level 0.

Since the solution subset distance is computed by looking only at the cardinality of the solution sets, adding a solution means an increase by just 1 of the cardinality of the set. Therefore, the distance between P and any solution η not in the original problem must be just 1.

Developing this idea we can map a crisp problem P with a set of constraints C to an SCSP P' with constraints C' over the semiring $S_{\{0,1\}} = \langle \{0, 1\}, min, max, 1, 0\rangle$[3]. We can then map each constraint c in P to a corresponding soft constraint c' in P' over the same variables such that $c'\eta = 0$ when the tuple η satisfies the constraint c, and $c'\eta = 1$ when the tuple η violates the constraint.

Combining the constraints in the SCSP P' we obtain its solution. A solution η of the original problem P will lead to a solution of the SCSP P' with semiring level 0, while non-solutions of the original problem will lead to semiring levels higher (worse) than 0. If an assignment η is not a solution of P, the solution of the SCSP for η (that is $\bigotimes C\eta$) will have semiring level 1.

Theorem 2 *Consider a crisp CSP P and the corresponding SCSP P' over the semiring $S_{\{0,1\}} = \langle \{0, 1\}, min, max, 1, 0\rangle$, obtained using the mapping described above. Then,*

- *η is a solution of P iff η has semiring level 0 in P';*
- *η is not a solution of P iff η has semiring level 1 in P'.*

Relaxing a problem using this distance metric means adding tuples to constraints such that new solutions will be added to the problem. Notice that using this distance we do not care about the number of constraint violations we have to tolerate in order to add a new assignment in the solution set. Regardless of how many constraints the new assignment η violates, we will have always $\bigotimes C'\eta = 1$. Therefore, we can only define two thresholds with this distance:

- *Threshold* 0: we do not want to add more solutions to the original problem;
- *Threshold* 1: we want to add additional solutions to the original problem.

4.3. Augmentation Distance

The augmentation distance metric was defined as the number of constraint augmentations required to relax the problem. To translate it to the semiring-based framework we need to change our point of view. Consider two problem P_1 and P_2 that are totally ordered with respect to each other. The augmentation distance between them is given by the sum of the "amount" each constraint in P_2 has been relaxed with respect to P_1. When dealing with a constraint represented extensionally, the amount of relaxation can be regarded as the number of tuples added. We will see at the end of this section what happens in the case where we have an intensional representation.

Suppose for the moment that constraints are represented extensionally. As usual, we want to associate with solutions to the original problem a distance (semiring level) of 0, while for assignments that are not solutions we wish to associate a semiring level that is worse (greater) than 0. If an assignment η has cost α, this means that we have to add α tuples to the constraints of the problem.

Therefore, we can map a crisp problem P with a set of constraints C to an SCSP P' with constraints C' over the semiring $S_{WCSP} = \langle \mathbb{N}^+, min, +, +\infty, 0\rangle$. Notice that the semiring we use is the same semiring used for the Max-CSP distance.

We can map each constraint c in P to a corresponding soft constraint c' in P' over the same variables such that $c'\eta = 0$ when the tuple η satisfies the constraint c, and $c'\eta = k$ when extending the constraint c in order to satisfy $c\eta$ will require adding an additional k tuples to c.

Note that when we deal with constraints represented in an extensional way, the value of k will be always 1, and we will obtain the same distance of Max-CSP. This follows from the fact that the number of constraints violated and the minimum number of tuple we must add are in fact the same.

3 The structure with the interval $[0, 1]$ instead of the set $\{0, 1\}$ has also been proven to be a semiring in [1, 3].

When dealing with an intensional representation of constraints the two distances are instead different. Consider the following example.

Example 1 *Consider a problem with the following intensional constraints:* $c_x : x > 5$ *and* $c_y : y = 3$*. The assignment* $\eta_1 = \{x := 6, y := 3\}$ *is a solution and in the corresponding soft CSP it must have value* 0*. Consider instead the assignments* $\eta_2 = \{x := 3, y := 3\}$ *and* $\eta_3 = \{x := 2, y := 3\}$*. For extensionally represented constraints both* η_2 *and* η_3 *relax constraint* c_x *by adding one more domain element (*$x := 3$ *and* $x := 2$*, respectively). However, since we are dealing with intensionally represented constraints it seems more realistic to assign a different relaxation cost to* η_2 *and* η_3*, because for* η_2 *we need to enlarge the constraint* $c_x : x > 5$ *to* $c'_x : x > 2$ *whilst for* η_3 *the minimum augmentation of the constraint has to be* $c'_x : x > 1$*.*

Therefore, we can map the crisp constraint problem P *into the soft* P'*, where* $c'_x(a) = 0$ *if* $a > 5$*, and* $c'_x(a) = 6 - a$ *if* $a \leq 5$*. In the above example the augmentation distance for* η_2 *is* 3 *while for* η_3 *is* 4*.*

Note that using the Max-CSP distance (see Section 4.1) both η_2 *and* η_3 *have the same semiring level* 1*, since in both cases only 1 constraint is violated.*

We can prove that the SCSP we obtain correctly represents this distance: a solution η of the original problem P will lead, in the SCSP P', to a semiring level 0; a non-solution of the original problem will lead instead to semiring levels bigger (worse) than 0. If a solution η will relax a constraint by an amount k, then the solution of the SCSP for η (that is $\bigotimes C'\eta$) will have a semiring level k.

Theorem 3 *Consider a crisp CSP* P *and the corresponding SCSP* P' *over the semiring* $S_{WCSP} = \langle \mathbb{N}^+, min, +, +\infty, 0 \rangle$*, obtained using the mapping described above. Then,*

- η *is a solution of* P *iff* η *has semiring level* 0 *in* P'*;*
- η *is violating the constraints for a global amount of* k *in* P *iff* η *has semiring level* k *in* P'*.*

If we want to relax the initial problem P by a quantity α (that is we can afford to reduce the tightness of the constraint by α) this is captured in the SCSP P' by looking for all solutions better (that is smaller) than α.

5. Conclusions

In this paper we have demonstrated how the semiring framework for soft constraints can be used to define the standard distance metrics used in partial constraint satisfaction. In particular, we have focused on capturing the notion of distance between a solution and a problem. These solution-problem distance measures can be seen as providing lower-bounds on the distance between a problem and its relaxation. The next step in our research agenda is to develop a general approach, within the semiring framework, to computing the distance between a problem and its relaxation based upon the approach we have presented here to computing the distances between a solution and a problem.

References

[1] S. Bistarelli. *Semirings for Soft Constraint Solving and Programming*, volume 2962 of *Lecture Notes in Computer Science*. Springer, 2004.

[2] S. Bistarelli, H. Fargier, U. Montanari, F. Rossi, T. Schiex, and G. Verfaillie. Semiring-based csps and valued csps: Frameworks, properties, and comparison. *Constraints*, 4(3):199–240, 1999.

[3] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based constraint solving and optimization. *Journal of ACM*, 44(2):201–236, 1997.

[4] S. Bistarelli, U. Montanari, and F. Rossi. Soft concurrent constraint programming. In *Proc. ESOP, Grenoble, France*, LNCS, pages 53–67. Springer-Verlag, 2002.

[5] J. Bowen, R. Lai, and D. Bahler. Lexical imprecision and fuzzy constraint networks. In *Proceedings of AAAI-92*, pages 616–621, July 1992.

[6] E. Freuder. Partial constraint satisfaction. In *Proceedings of IJCAI 89*, pages 278–283, 1989.

[7] E. Freuder and R. Wallace. Partial constraint satisfaction. *Artificial Intelligence*, 58(1–3):21–70, 1992.

[8] J. Larrosa, P. Meseguer, and T. Schiex. Maintaining reversable DAC for Max-CSP. *Artificial Intelligence*, 107:149–163, 1999.

[9] A. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8:99–118, 1977.

[10] J. Regin, T. Petit, C. Bessiere, and J.-F. Puget. New lower bounds of constraint violations for over-constrained problems. In T. Walsh, editor, *Proceedings of CP*, LNCS 2239, pages 332–345, 2001.

[11] T. Schiex. Possibilistic constraint satisfaction problems, or "how to handle soft constraints?". In *Proc. 8th Conf. of Uncertainty in AI*, pages 269–275, 1992.

[12] T. Schiex, H. Fargier, and G. Verfaille. Valued constraint satisfaction problems: hard and easy problems. In *Proceedings of the International joint Conference in AI*, 1995.

[13] M. Wallace. Practical applications of constraint programming. *Constraints*, 1(1–2):139–168, 1996.

[14] R. Wallace. Directed arc-consistency preprocessing as a strategy for maximal constraint satisfaction. In *Proceedings of ECAI*, pages 69–77, 1994.