Bipolar Preference Problems: Framework, Properties and Solving Techniques

Stefano Bistarelli[1,2], Maria Silvia Pini[3], Francesca Rossi[3], and K. Brent Venable[3]

[1] Dipartimento di Scienze, Università "G. d'Annunzio", Pescara, Italy
bista@sci.unich.it
[2] Istituto di Informatica e Telematica, CNR, Pisa, Italy
Stefano.Bistarelli@iit.cnr.it
[3] Dipartimento di Matematica Pura ed Applicata, Università di Padova, Italy
{mpini,frossi,kvenable}@math.unipd.it

Abstract. Real-life problems present several kinds of preferences. We focus on problems with both positive and negative preferences, that we call *bipolar preference problems*. Although seemingly specular notions, these two kinds of preferences should be dealt with differently to obtain the desired natural behaviour. We technically address this by generalizing the soft constraint formalism, which is able to model problems with one kind of preferences. We show that soft constraints model only negative preferences, and we define a new mathematical structure which allows to handle positive preferences as well. We also address the issue of the compensation between positive and negative preferences, studying the properties of this operation. Finally, we extend the notion of arc consistency to bipolar problems, and we show how branch and bound (with or without constraint propagation) can be easily adapted to solve such problems.

1 Introduction

Many real-life problems contain statements which can be expressed as preferences. Our long-term goal is to define a framework where many kinds of preferences can be naturally modelled and efficiently dealt with. In this paper, we focus on problems which present positive and negative preferences, that we call *bipolar preference problems*.

Positive and negative preferences can be thought as two symmetric concepts, and thus one can think that they can be dealt with via the same operators. However, this would not model what one usually expects in real scenarios. In fact, usually combination of positive preferences should produce a higher (positive) preference, while combination of negative preferences should give us a lower (negative) preference.

When dealing with both kinds of preferences, it is natural to express also indifference, which means that we express neither a positive nor a negative preference over an object. Then, a desired behaviour of indifference is that, when combined with any preference (either positive or negative), it should not influence the overall preference.

Finally, besides combining preferences of the same type, we also want to be able to combine positive with negative preferences. We strongly believe that the most natural and intuitive way to do so is to allow for compensation. Comparing positive against negative aspects and compensating them w.r.t. their strength is one of the core features

F. Azevedo et al. (Eds.): CSCLP 2006, LNAI 4651, pp. 78–92, 2007.

of decision-making processes, and it is, undoubtedly, a tactic universally applied to solve many real life problems.

Positive and negative preferences might seem as just two different criteria to reason with, and thus techniques such as those usually adopted by multi-criteria optimization [13] could appear suitable for dealing with them. However, this interpretation would hide the fundamental nature of bipolar preferences, that is, positive preferences are naturally opposite of negative preferences. Moreover, in multi-criteria optimization it is often reasonable to use a Pareto-like approach, thus associating tuples of values to each solution, and comparing solutions according to tuple dominance. Instead, in bipolar problems, it would be very unnatural to force such an approach in all contexts, or to associate to a solution a preference which is neither a positive nor a negative one.

Soft constraints [5] are a useful formalism to model problems with quantitative preferences. However, they can only model negative preferences, since in this framework preference combination returns lower preferences. In this paper we adopt the soft constraint formalism based on semirings to model negative preferences. We then define a new algebraic structure to model positive preferences. To model bipolar problems, we link these two structures and we set the highest negative preference to coincide with the lowest positive preference to model indifference. We then define a combination operator between positive and negative preferences to model preference compensation, and we study its properties.

Non-associativity of preference compensation occurs in many contexts, thus we think it is too restrictive to focus just on associative environments. For example, non-associativity of compensation arises when either positive or negative preferences are aggregated with an idempotent operator (such as min or max), while compensation is instead non-idempotent (such as sum). Our framework allows for non-associativity, since we want to give complete freedom to choose the positive and negative algebraic structures. However, we also describe a technique that, given a negative structure, builds a corresponding positive structure and an associative compensation operator.

Finally, we consider the problem of finding optimal solutions of bipolar problems, by suggesting a possible adaptation of constraint propagation and branch and bound to the generalized scenario.

Summarizing, the main results are:

- a formal definition of an algebraic structure to model bipolar preferences;
- the study of the notion of compensation and of its properties (such as associativity);
- a technique to build a bipolar preference structure with an associative compensation operator;
- the adaptation of branch and bound to solve bipolar problems;
- the definition of bipolar propagation and its use within a branch and bound solver.

The paper is organized as follows. Section 2 recalls the main notions of semiring-based soft constraints. Section 3 describes how to model negative preferences using usual soft constraints, and how to model positive preferences. Section 4 shows how to model both positive and negative preferences, and Section 5 defines constraint problems with both positive and negative preferences. Section 6 shows that it is important to have a bipolar structure for expressing both positive and negative preferences, since expressing all the problems' requirements in a positive (or negative) form might lead to different optimal solutions. Section 7 shows that very often the compensation operator is not associative and it describes a technique to build a bipolar preference structure with an associative compensation operator. Section 8 shows how to adapt branch and bound to solve bipolar problems, how to define bipolar propagation and its use within a branch and bound solver. Finally, Section 9 describes the existing related work and gives some hints for future work.

Earlier versions of parts of this paper have appeared in [6].

2 Semiring-Based Soft Constraints

A soft constraint [5] is a classical constraint[1] 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. A c-semiring is a tuple $(A, +, \times, \mathbf{0}, \mathbf{1})$ s.t. A is a set 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. Given the relation $\leq_S$ over A s.t. $a \leq_S b$ iff $a + b = b$, $\leq_S$ 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.

Given a c-semiring $S = (A, +, \times, \mathbf{0}, \mathbf{1})$, 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 $def : D^{|con|} \rightarrow A$. 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. Given a subset of variables $I \subseteq V$, and a soft constraint $c = \langle def, con \rangle$, the projection of c over I, written $c \Downarrow_I$, is a new soft constraint $\langle def', con' \rangle$, where $con' = con \cap I$ and $def(t') = \sum_{\{t | t\downarrow_{con'} = t'\}} def(t)$. The scope, con', of the projection constraint contains the variables that con and I have in common, and thus $con' \subseteq con$. Moreover, the preference associated to each assignment to the variables in con', denoted with t', is the highest ($\sum$ is the additive operator of the c-semiring) among the preferences associated by def to any completion of t', t, to an assignment to con.

A soft constraint satisfaction problem (SCSP) is just a set of soft constraints over a set of variables. A classical CSP is just an SCSP where the chosen c-semiring is: $S_{CSP} = (\{false, true\}, \vee, \wedge, false, true)$. On the other hand, fuzzy CSPs can be modelled in the SCSP framework by choosing the c-semiring: $S_{FCSP} = ([0, 1], max, min, 0, 1)$. For weighted CSPs, the semiring is $S_{WCSP} = (R^+, min, +, +\infty, 0)$. Here preferences are interpreted as costs of which we want to minimize the sum. For probabilistic CSPs, the semiring is $S_{PCSP} = ([0, 1], max, \times, 0, 1)$. Here preferences are interpreted as probabilities and the aim is to maximize the joint probability.

Given an assignment to all the variables of an SCSP, we can compute its preference value by combining the preferences associated by each constraint to the subtuples of the assignments referring to the variables of the constraint. An optimal solution of an

SCSP is then a complete assignment t such that there is no other complete assignment t' with $pref(t) <_S pref(t')$.

3 Negative and Positive Preferences

The structure we use to model negative preferences is exactly a c-semiring, as defined in Section 2. In fact, in a c-semiring the element which acts as indifference is $\mathbf{1}$, since $\forall a \in A$, $a \times \mathbf{1} = a$. This element is the best in the ordering, which is consistent with the fact that indifference is the best preference when using only negative preferences. Moreover, in a c-semiring, combination goes down in the ordering, since $a \times b \leq a, b$. This can be naturally interpreted as the fact that combining negative preferences worsens the overall preference. From now on, we use $(N, +_n, \times_n, \bot_n, \top_n)$ as c-semiring to model negative preferences.

When dealing with positive preferences, we want two main properties to hold: combination should bring to better preferences, and indifference should be lower than all the other positive preferences. These properties can be found in the following structure.

Definition 1. *A positive preference structure is a tuple* $(P, +_p, \times_p, \bot_p, \top_p)$ *s.t.*

- P *is a set and* $\top_p, \bot_p \in P$*;*
- $+_p$*, the additive operator, is commutative, associative, idempotent, with* $\bot_p$ *as its unit element* ($\forall a \in P, a +_p \bot_p = a$) *and* $\top_p$ *as its absorbing element* ($\forall a \in P, a +_p \top_p = \top_p$)*;*
- $\times_p$*, the multiplicative operator, is associative, commutative and distributes over* $+_p$ ($a \times_p (b +_p c) = (a \times_p b) +_p (a \times_p c)$)*, with* $\bot_p$ *as its unit element and* $\top_p$ *as its absorbing element*[1].

Notice that the additive operator of this structure has the same properties as the corresponding one in c-semirings, and thus it induces a partial order over P in the usual way: $a \leq_p b$ iff $a +_p b = b$. This allows to prove that $+_p$ is monotone over $\leq_p$ and that it is the least upper bound in the lattice $(P, \leq_p)$. On the other hand, the multiplicative operator has different properties. More precisely, the best element in the ordering ($\top_p$) is now its absorbing element, while the worst element ($\bot_p$) is its unit element. This reflects the desired behavior of the combination of positive preferences.

Theorem 1. *Given a positive preference structure* $(P, +_p, \times_p, \bot_p, \top_p)$*, consider the relation* $\leq_p$ *over P. Then,* $\times_p$ *is monotone over* $\leq_p$ *(that is, for any* $a, b \in P$ *s.t.* $a \leq_p b$*, then* $a \times_p d \leq_p b \times_p d$*,* $\forall d \in P$*), and* $\forall a, b \in P$*,* $a \times_p b \geq_p a +_p b \geq_p a, b$*.*

Proof. Since $a \leq_p b$ iff $a +_p b = b$, then $b \times_p d = (a +_p b) \times_p d = (a \times_p d) +_p (b \times_p d)$. Thus $a \times_p d \leq_p b \times_p d$. Also, $a \times_p b$ = $a \times_p (b +_p \bot_p)$ = $(a \times_p b) +_p (a \times_p \bot_p)$ = $(a \times_p b) +_p a$. Thus $a \times_p b \geq_p a$ (the same for b). Finally: $a \times_p b \geq a, b$. Thus $a \times_p b \geq lub(a, b) = a +_p b$. Q.E.D.

In a positive preference structure, $\bot_p$ is the element modelling indifference. In fact, it is the worst one in the ordering and it is the unit element for the combination operator

[1] The absorbing nature of $\top_p$ can be derived from the other properties.

$\times_p$. These are exactly the desired properties for indifference w.r.t. positive preferences. The role of $\top_p$ is to model a very high preference, much higher than all the others. In fact, since it is the absorbing element of the combination operator, when we combine any positive preference a with $\top_p$, we get $\top_p$.

An example of a positive preference structure is $P_1 = (R^+, max, +, 0, +\infty)$, where preferences are positive reals. The smallest preference that can be assigned is 0. It represents the lack of any positive aspect and can thus be regarded as indifference. Preferences are aggregated taking the sum and are compared taking the max.

Another example is $P_2 = ([0,1], max, max, 0, 1)$. In this case preferences are reals between 0 and 1, as in the fuzzy semiring for negative preferences. However, the combination operator is max, which gives, as a resulting preference, the highest one among all those combined.

As an example of a partially ordered positive preference structure consider the Cartesian product of the two structures described above: $(R^+ \times [0,1], (max, max), (+, max), (0,0), (+\infty, 1))$. Positive preferences, here, are ordered pairs where the first element is a positive preference of type P_1 and the second one is a positive preference of type P_2. Consider for example the (incomparable) pairs $(8, 0.1)$ and $(3, 0.8)$. Applying the aggregation operator (+,max) gives the pair $(11, 0.8)$ which, as expected, is better than both pairs, since $max(8, 3, 11) = 11$ and $max(0.1, 0.8, 0.8) = 0.8$.

4 Bipolar Preference Structures

Once we are given a positive and a negative preference structure, a naive way to combine them would be performing the Cartesian product of the two structures. For example, if we have a positive structure P and a negative structure N, taking their Cartesian product would mean that, given a solution, it will be associated with a pair (p, n), where $p \in P$ is the overall positive preference and $n \in N$ is the overall negative preference. Such a pair is in general neither an element of P nor of N, so it is neither positive nor negative, unless one or both of p and n are the indifference element. Moreover, the ordering induced over these pairs is the well known Pareto ordering, which may induce a lot of incomparability among the solutions. These two features imply that compensation is not allowed at all. Instead, we believe that it should be allowed, if desired. We will therefore now describe a bipolar preference structure that allows for it.

Definition 2. *A bipolar preference structure is a tuple* $(N, P, +, \times, \bot, \Box, \top)$*, where*

- $(P, +_{|P}, \times_{|P}, \Box, \top)$ *is a positive preference structure;*
- $(N, +_{|N}, \times_{|N}, \bot, \Box)$ *is a c-semiring;*
- $+ : (N \cup P)^2 \longrightarrow (N \cup P)$ *is s.t.* $a_n + a_p = a_p$ *for any* $a_n \in N$ *and* $a_p \in P$*; this operator induces as partial ordering on* $N \cup P$*:* $\forall a, b \in P \cup N$*,* $a \leq b$ *iff* $a + b = b$*;*
- $\times : (N \cup P)^2 \longrightarrow (N \cup P)$ *is an operator (called the* compensation operator*) that, for all* $a, b, c \in N \cup P$*, satisfies the following properties:*
  - *commutativity:* $a \times b = b \times a$*;*
  - *monotonicity: if* $a \leq b$*, then* $a \times c \leq b \times c$*.*

In the following, we will write $+_n$ instead of $+_{|N}$ and $+_p$ instead of $+_{|P}$. Similarly for $\times_n$ and $\times_p$. Moreover, we will sometimes write $\times_{np}$ when operator $\times$ will be applied to a pair in $(N \times P)$.

Notice that bipolar structures generalize both negative and positive preference structures via a bipolar structure with a single positive/negative preference element.

Given the way the ordering is induced by $+$ on $N \cup P$, easily, we have $\bot \leq \Box \leq \top$. Thus, there is a unique maximum element (that is, $\top$), a unique minimum element (that is, $\bot$); the element $\Box$ is smaller than any positive preference and greater than any negative preference, and it is used to model indifference. The shape of a bipolar preference structure is shown in the following figure:

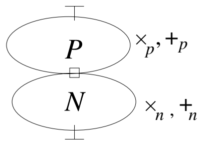

Notice that, although all positive preferences are strictly above negative preferences, our framework does not prevent from using the same scale, or partially overlapping scales, to represent positive and negative preferences.

A bipolar preference structure allows us to have different ways to model and reason about positive and negative preferences. In fact, we can have different lattices $(P, \leq_p)$ and $(N, \leq_n)$. This is common in real-life problems, where negative and positive statements are not necessarily expressed using the same granularity. For example, we could be satisfied with just two levels of negative preferences, while requiring several levels of positive preferences. Nevertheless, our framework allows to model cases in which the two structures are isomorphic, as well (see Section 7).

Notice that the combination of a positive and a negative preference is a preference which is higher than, or equal to, the negative one and lower than, or equal to, the positive one.

Theorem 2. *Given a bipolar preference structure* $(N, P, +, \times, \bot, \Box, \top)$*, for all* $p \in P$ *and* $n \in N$*,* $n \leq p \times n \leq p$*.*

Proof. For any $n \in N$ and $p \in P$, $\Box \leq p$ and $n \leq \Box$. By monotonicity of $\times$, we have: $n \times \Box \leq n \times p$ and $n \times p \leq \Box \times p$. Hence: $n = n \times \Box \leq n \times p \leq \Box \times p = p$. Q.E.D.

This means that the compensation of positive and negative preferences lies in one of the chains between the two combined preferences. Notice that all such chains pass through the indifference element $\Box$. Possible choices for combining strictly positive with strictly negative preferences are thus the average or the median operator, or also the minimum or the maximum.

Moreover, by monotonicity, we can show that if $\top \times \bot = \bot$, then the result of the compensation between any positive preference and the bottom element is the bottom element, and if $\top \times \bot = \top$, then the compensation between any negative preference and the top element is the top element.

Theorem 3. *Given a bipolar preference structure* $(N, P, +, \times, \bot, \Box, \top)$*, if* $\top \times \bot = \bot$*, then* $\forall p \in P$*,* $p \times \bot = \bot$*, while, if* $\top \times \bot = \top$*, then* $\forall n \in N$*,* $n \times \top = \top$*.*

Proof. Assume $\top \times \perp = \perp$. Since for all $p \in P$, $p \leq \top$, then, by monotonicity of $\times$, $p \times \perp \leq \top \times \perp = \perp$, hence $p \times \perp = \perp$.

Assume $\top \times \perp = \top$. Since for all $n \in N$, $\perp \leq n$, then, by monotonicity of $\times$, $\top = \top \times \perp \leq \top \times n$, hence $\top \times n = \top$. Q.E.D

In the following table we give three examples of bipolar preference structures, one for each row.

N,P	$+$	$\times$	$\perp, \Box, \top$
R^-, R^+	$+_p$=$+_n$=$+_{np}$= max	$\times_p$=$\times_n$=$\times_{np}$=sum	$-\infty$, 0, $+\infty$
$[-1, 0]$, $[0, 1]$	$+_p$=$+_n$=$+_{np}$=max	$\times_p$=max, $\times_n$=min, $\times_{np}$ =sum	-1, 0, 1
$[0, 1]$, $[1, +\infty]$	$+_p$=$+_n$=$+_{np}$=max	$\times_p$=$\times_n$=$\times_{np}$=prod	0, 1, $+\infty$

The structure described in the first row uses real numbers as positive and negative preferences. Compensation is obtained by summing the preferences, while the ordering is given by the max operator. In the second structure we have positive preferences between 0 and 1 and negative preferences between -1 and 0. Aggregation of positive preferences is max and of negative preferences is min, while compensation between positive and negative preferences is sum, and the order is given by max. In the third structure we use positive preferences between 1 and $+\infty$ and negative preferences between 0 and 1. Compensation is obtained by multiplying the preferences and ordering is again via max. Notice that if $\top \times \perp \in \{\top, \perp\}$, then compensation in the first and in the third structure is associative.

5 Bipolar Preference Problems

A bipolar constraint is just a constraint where each assignment of values to its variables is associated to one of the elements in a bipolar preference structure. A bipolar CSP (V, C) is then just a set of variables V and a set of bipolar constraints C over V. There could be many ways of defining the optimal solutions of a bipolar CSP. Here we propose one which avoids problems due to the possible non-associativity of the compensation operator, since compensation never involves more than two preference values.

Definition 3. *Given a bipolar preference structure* $(N, P, +, \times, \perp, \Box, \top)$*, a solution of a bipolar CSP* (V, C) *is a complete assignment to all variables in* V*, say* s*, and an associated preference which is computed as follows:* $pref(s) = (p_1 \times_p \ldots \times_p p_k) \times (n_1 \times_n \ldots \times_n n_l)$*, where* $p_i \in P$ *for* $i := 1, \ldots, k$ *and* $n_j \in N$ *for* $j := 1, \ldots, l$ *and where* $\exists\ \langle def_i, con_i\rangle \in C$ *s.t.* $p_i = def_i(s \downarrow_{con})$*, and* $\exists \langle def_j, con_j\rangle \in C$ *s.t.* $n_j = def_j(s \downarrow_{con})$*. A solution* s *is an optimal solution if there is no other solution* s' *with* $pref(s') > pref(s)$*.*

In this definition, the preference of a solution s is obtained by combining all the positive preferences associated to its projections over the constraints, by using $\times_p$, combining all the negative preferences associated to its projections over the constraints, by using $\times_n$, and then, combining the two preferences obtained so far (one positive and one negative) by using the operator $\times_{np}$.

Such a definition follows the same idea proposed in Chapter IV of [4] for evaluating the tendency of an act. Such an idea consists of summing up all the values of all the pleasures produced by the considered act on one side, and those of all the pains produced by it on the other, and then balancing these two resulting values in a value which can be on the side of pleasure or on the side of pain. If this value is on the side of pleasure, then the tendency of the act is good, otherwise the tendency is bad.

Consider the scenario in which we want to buy a car and we have preferences over some features. In terms of color, we like red, we are indifferent to white, and we hate black. Also, we like convertible cars a lot and we don't care much for big cars (e.g., SUVs). In terms of engines, we like diesel. However, we don't want a diesel convertible.

We represent positive preferences via positive integers, negative preferences via negative integers and we maximize the sum of all kinds of preferences. This can be modelled by a bipolar preference structure where $N = [-\infty, 0]$, $P = [0, +\infty]$, $+$ =max, $\times$=sum, $\perp = -\infty$, $\Box = 0, \top = +\infty$.

The following figure shows the structure (variables, domains, constraints, and preferences) of such a bipolar CSP, where preferences have been chosen to fit the informal specification above, and 0 is used to model indifference (also when tuples are not shown).

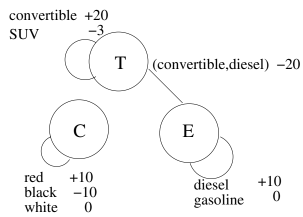

Consider solution s_1 =(red,convertible,diesel). We have $pref(s_1) = (def_1(\text{red}) \times def_2(\text{convertible}) \times def_3(\text{diesel})) \times def_4(\text{convertible, diesel}) = (10 + 20 + 10) + (-20) = 20$. We can see that the optimal solution is (red, convertible, gasoline) with global preference of 30.

Consider now a different bipolar preference structure, which differs from the previous one only for $\times_p$, which is now max. Now solution s_1 has preference $pref(s_1) = (def_1(\text{red}) \times def_2(\text{convertible}) \times def_3(\text{diesel})) \times def_4(\text{convertible, diesel}) = max(10, 20, 10) + (-20) = 0$. It is easy to see that now an optimal solution has preference 20. There are two of such solutions: one is the same as the optimal solution above, and the other one is (white,convertible, gasoline). The two cars have the same features except for the color. A white convertible is just as good as a red convertible because we decided to aggregate positive preference by taking the maximum elements rather than by summing them.

6 Positive Versus Negative Preferences

Positive and negative preferences look so similar that, even though we know they need different combination operators, we could wonder why we need two different structures to handle them. Why can't we just have one structure, for example the negative one, and transform each positive preference into a negative one? For example, if there are only two colors for cars, (i.e., red and blue), and we only like blue, instead of saying this using positive preferences (i.e., we like blue with a certain positive preference), we could phrase it using negative preferences (i.e., we don't like red with a certain negative preference). In other words, instead of associating a positive preference to blue and indifference to red, we could give a negative preference to red and indifference to blue.

In this section we will show that, by doing this, we could modify the solution ordering, thus representing a different optimization problem. Thus sometimes the two preference structures are needed to model the problems under consideration: using just one of them would not suffice.

It is easy to show that, by moving from a positive to a negative modelling of the same information, as we have done in the example above, all solutions get a lower preference value. In fact, in this transformation, a positive preference is replaced by indifference, or indifference is replaced by a negative preference. So, in any case, some preference is replaced by a lower one, and by monotonicity of the aggregation operators ($\times_n$, $\times_p$, and $\times$), the overall preference of the solutions is lower as well.

However, it is worth noting that this preference lowering might not preserve the ordering among solutions. That is, solutions that were ordered in a certain way before the modification, can be ordered in the opposite way after it. This is due to the fact that aggregation of positive and negative preferences may behave differently. The following example shows this phenomenon.

Consider the bipolar preference structure (R^-, R^+, max, $\times$, $-\infty$, 0, $+\infty$), where $\times$ is such that $\times_p$=$\times_{np}$= sum and $\times_n = min$. This means that we want to maximize the sum of positive preferences, maximize the minimal negative preference (thus negative preferences are handled as fuzzy constraints), that positive preferences are between 0 and $+\infty$, and negative preferences are between 0 and -∞. Compensation is via algebraic sum, thus values v and $-v$ are compensated completely (that is, the result of the compensation is 0), while the compensation of values v and $-v'$ is $v - v'$.

Consider now a bipolar CSP over this structure with four variables, say X, Y, Z, W, where each variable has a Boolean domain as follows: $D(X) = \{a, \bar{a}\}$, $D(Y) = \{b, \bar{b}\}$, $D(Z) = \{c, \bar{c}\}$, and $D(W) = \{d, \bar{d}\}$. Assume now that the preference of a is 2, of b is 1, of c is 2.4, and of d is 0.5, while the preference of the other elements is indifference (that is, 0 in this example). This means that we have expressed all our statements in a positive form.

Consider now two solutions s and s' as follows: $s = (a, b, \bar{c}, \bar{d})$ and $s' = (\bar{a}, \bar{b}, c, d)$. By computing the preference of s, we get $(2 + 1) + min(0, 0) = 3$, while for s' we get $min(0, 0) + (2.4 + 0.5) = 2.9$. Thus s is better than s'.

Assume now to express the same statements in negative terms assuming that if we like at level p an assignment t, then we dislike $\bar{t}$ at the same level p. Hence, the preference of $\bar{a}$ is -2, of $\bar{b}$ is -1, of $\bar{c}$ is -2.4, and of $\bar{d}$ is -0.5, while the preference of the other elements is 0. Now the preference of s is $(0 + 0) + min(-2.4, -0.5) = -2.4$, while the preference of s' is $min(-2, -1) + (0 + 0) = -2$. Thus s' is better than s.

7 Associativity of Preference Compensation

In general, the compensation operator $\times$ may be not associative. Here we list some sufficient conditions for the non-associativity of the $\times$ operator.

Theorem 4. *Given a bipolar preference structure* $(P, N, +, \times, \bot, \Box, \top)$*, operator* $\times$ *is not associative if at least one of the following two conditions is satisfied:*

- $\top \times \bot = c \in (N \cup P) - \{\top, \bot\}$*;*
- $\exists p \in P - \{\top, \Box\}$ *and* $n \in N - \{\bot, \Box\}$ *s.t.* $p \times n = \Box$ *and at least one of the following conditions holds:*
  - $\times_p$ *or* $\times_n$ *is idempotent;*
  - $\exists p' \in P - \{p, \top\}$ *s.t.* $p' \times n = \Box$ *or* $\exists n' \in N - \{n, \bot\}$ *s.t.* $p \times n' = \Box$*;*
  - $\top \times \bot = \bot$ *and* $\exists n' \in N - \{\bot\}$ *s.t.* $n \times n' = \bot$*;*
  - $\top \times \bot = \top$ *and* $\exists p' \in P - \{\top\}$ *s.t.* $p \times p' = \top$*;*
  - $\exists a, c \in N \cup P$ *s.t.* $a \times p = c$ *iff* $c \times n \neq a$ *(or* $\exists a, c \in N \cup P$ *s.t.* $a \times n = c$ *iff* $c \times p \neq a$*).*

Proof. If $c \in P - \{\top\}$, then $\top \times (\top \times \bot) = \top \times c = \top$, while $(\top \times \top) \times \bot = \top \times \bot = c$. If $c \in N - \{\bot\}$, then $\bot \times (\bot \times \top) = \bot \times c = \bot$, while $(\bot \times \bot) \times \top = \bot \times \top = c$.

Assume that $\exists p \in P - \{\top, \Box\}$ and $n \in N - \{\bot, \Box\}$ s.t. $p \times n = \Box$. If $\times_p$ is idempotent, then $p \times (p \times n) = p \times \Box = p$, while $(p \times p) \times n = p \times n = \Box$. Similarly if $\times_n$ is idempotent.

If $\exists p' \in P - \{p, \top\}$ s.t. $p' \times n = \Box$, then $(p \times n) \times p' = p'$, while $p \times (n \times p') = p$. Analogously, if $\exists n' \in N - \{n, \bot\}$ s.t. $p \times n' = \Box$.

If $\top \times \bot = \bot$, then, by Theorem 3, $p \times \bot = \bot$. If $\exists n' \in N - \{\bot\}$ s.t. $n \times n' = \bot$, then $(p \times n) \times n' = \Box \times n' = n'$, while $p \times (n \times n') = p \times \bot = \bot \neq n'$.

If $\top \times \bot = \top$, then, by Theorem 3, $n \times \top = \top$. If $\exists p' \in P - \{\top\}$ s.t. $p \times p' = \top$, then $(n \times p) \times p' = \Box \times p' = p'$, while $n \times (p \times p') = n \times \top = \top \neq p'$.

If $c \times n \neq a$, then $(a \times p) \times n = c \times n \neq a$, but $a \times (p \times n) = a \times \Box = a$. Analogously if $c \times p \neq a$. Q.E.D.

Notice that sufficient conditions refer to various aspects of a bipolar preference structure: properties of operators, shape of P and N orderings, the relation between $\times$ and the other operators. Since some of these conditions often occur in practice, it is not reasonable to require always associativity of $\times$.

It is however useful to be able to build bipolar preference structures where compensation is associative. It is obvious that, if we are free to choose any positive and any negative preference structure when building the bipolar framework, we will never be able to assure associativity of the compensation operator. Thus, to assure this, we must pose some restrictions on the way a bipolar preference structure is built.

We describe now a procedure to build positive preferences as inverses of negative preferences, that assures that the resulting bipolar preference structure has an associative compensation operator. To do that, $\times_n$ must be non-idempotent. The methodology is called *localization* and represents a standard systematic technique for adding multiplicative inverses to a (semi)ring [8].

Given a (semi)ring with carrier set N (representing, in our context, a negative preference structure), and a subset $S \subseteq N$, we can construct another structure with carrier set P (representing, for us, a positive preference structure), and a mapping from N to P which makes all elements in the image of S invertible in P. The localization of N by S is also denoted by $S^{-1}N$.

We can select any subset S of N. However, it is usual to select a subset S of N which is closed under $\times_n$, such that $\mathbf{1} \in S$ ($\mathbf{1}$ is the unit for $\times_n$, which represents indifference), and $\mathbf{0} \notin S$.

Given N and S, let us consider the quotient field of N w.r.t. S. This is denoted by $Quot(N, S)$, and will represent the carrier set of our bipolar structure. One can construct $Quot(N, S)$ by just taking the set of equivalence classes of pairs (n, d), where n and d are elements of N and S respectively, and the equivalence relation is: $(n, d) \equiv (m, b) \iff n \times_n b = m \times_n d$. We can think of the class of (n, d) as the fraction $\frac{n}{d}$.

The embedding of N in $Quot(N, S)$ is given by the mapping $f(n) = (n, \mathbf{1})$, thus the (semi)ring N is a subring of $S^{-1}N$ via the identification $f(a) = \frac{a}{\mathbf{1}}$.

The next step is to define the $+$ and $\times$ operator in $Quot(N, S)$, as function of the operators $+_n$ and $\times_n$ of N. We define $(n, d) + (m, b) = ((n \times_n b) +_n (m \times_n d), d \times_n b)$ and $(n, d) \times (m, b) = (m \times_n n, d \times_n b)$. By using the fraction representation we obtain the usual form where the addition and the multiplication of the formal fractions are defined according to the natural rules: $\frac{a}{s} + \frac{b}{t} = \frac{(a \times_n t) +_n (b \times_n s)}{s \times_n t}$ and $\frac{a}{s} \times \frac{b}{t} = \frac{a \times_n b}{s \times_n t}$.

It can be shown that the structure $(P, +_p, \times_p, \frac{\mathbf{1}}{\mathbf{1}}, \frac{\mathbf{1}}{0})$, where $P = \{\frac{\mathbf{1}}{a}$ s.t. $a \in (S \cup \{\mathbf{0}\})\}$, $+_p$ and $\times_p$ are the operators $+$ and $\times$ restricted over $\frac{\mathbf{1}}{S} \times \frac{\mathbf{1}}{S}$, $\frac{\mathbf{1}}{\mathbf{1}}$ is the bottom element in the induced order (notice that the element coincide with $\mathbf{1}$), and $\frac{\mathbf{1}}{\mathbf{0}}$ is the top element of the structure[2], is a positive preference structure. Moreover, $Quot(N, S) = P \cup N$, and it is the carrier of a bipolar preference structure $\langle P, N, +, \times, \mathbf{0}, \frac{\mathbf{1}}{\mathbf{1}}, \frac{\mathbf{1}}{\mathbf{0}} \rangle$ where $\times$ is an associative compensation operator by construction.

Notice that the first example of the table in Section 4, as well as the third example restricted to rational numbers, can be obtained via the localization procedure.

8 Solving Bipolar Preference Problems

Bipolar problems are NP-complete, since they generalize both classical and soft constraints, which are already known to be difficult problems [5]. In this section we will consider how to adapt some usual techniques for soft constraints to bipolar problems.

8.1 Branch and Bound

Preference problems based on c-semirings can be solved via a branch and bound technique (BB), possibly augmented via soft constraint propagation, which may lower the preferences and thus allow for the computation of better bounds [5].

In bipolar CSPs, we have both positive and negative preferences. However, if the compensation operator is associative, standard BB can be applied. Thus bipolar preferences can be handled without additional effort.

However, if compensation is not associative, the upper bound computation has to be slightly changed to avoid performing compensation before all the positive and the negative preferences have been collected. More precisely, each node of the search tree is associated to a positive and a negative preference, say p and n, which are obtained by aggregating all preferences of the same type obtained in the instantiated part of the problem. Next, all the best preferences (which may be positive or negative) in the

[2] This element is introduced ad hoc because $\mathbf{0}$ is not an unit and cannot be used to build its inverse.

uninstantiated part of the problem are considered. By aggregating those of the same type, we get a positive and a negative preference, say p' and n', which can be combined with the ones associated to the current node. This produces the upper bound $ub = (p \times_p p') \times (n \times_n n')$, where $p' = p_1 \times_p \ldots \times_p p_w$, $n' = n_1 \times_n \ldots \times_n n_s$, with $w + s = r$, where r is the number of uninstantiated variables/constraints. Notice that ub is computed via $r - 1$ aggregation steps and one compensation step.

On the other hand, when compensation is associative, we don't need to postpone compensation until all constraints have been considered. Thus, ub can be computed as $ub = a_1 \times \ldots \times a_{p+r}$, where $a_i \in N \cup P$ is the best preference found in a constraint of either the instantiated part of the problem (first p elements) or the uninstantiated part of the problem (last r elements). Thus ub can be computed via at most $p + r - 1$ steps among which there can be many compensation steps.

8.2 Bipolar Propagation

When looking for an optimal solution, BB can be helped by some form of partial or full constraint propagation. To see whether this can be done when solving bipolar problems as well, we must first understand what constraint propagation means in such problems. For sake of semplicity, we will focus here on arc-consistency.

Given any bipolar constraint, let us first define its negative version $neg(c)$, which is obtained by just replacing the positive preferences with indifference. Similarly, the positive version $pos(c)$ is obtained by replacing negative preferences with indifference.

A binary bipolar constraint c is then said negatively arc-consistent (NAC) iff $neg(c)$ is soft arc-consistent. If the binary constraint connects variables X and Y, let us call it c_{XY}, and let us call c_X the soft domain of X and c_Y the soft domain of Y. Then, being soft arc consistent means that $neg(c_X) = (neg(c_X) \times_n neg(c_Y) \times_n neg(c_{XY})) \Downarrow_X$ and $neg(c_Y) = (neg(c_X) \times_n neg(c_Y) \times_n neg(c_{XY})) \Downarrow_Y$. If this is not so, we can make a binary bipolar constraint NAC by modifying the soft domains of its two variables such that the two equations above hold. The modifications required can only decrease some preference values. Thus some negative preferences can become more negative than before. If operator $\times_n$ is idempotent, then such modifications generate a new constraint which is equivalent to the given one.

Let us now consider the positive version of a constraint. Let us also define an operation $\Uparrow_X$, which, taken any constraint c_S over variables S such that $X \in S$, computes a new constraint over X as follows: for every value a in the domain of X, its preference is computed by taking the greatest lower bound of all preferences given by c_S to tuples containing $X = a$. Then we say that a binary bipolar constraint is positively arc-consistent (PAC) iff $c_X = (pos(c_X) \times_p pos(c_Y) \times_p pos(c_{XY}))\Uparrow_X$ and $c_Y = (pos(c_X) \times_p pos(c_Y) \times_p pos(c_{XY}))\Uparrow_Y$. If this is not so, we can make a binary bipolar constraint PAC by modifying the soft domains of its two variables such that the two equations above hold. The modifications required can only involve the increase of some preference values. Thus some positive preferences can become more positive than before. If operator $\times_p$ is idempotent, such modifications generate a new constraint which is equivalent to the given one.

Finally, we say that a binary bipolar constraint is Bipolar Arc-Consistent (BAC) iff it is NAC and PAC. A bipolar constraint problem is BAC iff all its constraints are BAC.

If a bipolar constraint problem is not BAC, we can consider its negative and positive versions and achieve PAC and NAC on them. If both $\times_n$ and $\times_p$ are idempotent, this can be seen as the application of functions which are monotone, inflationary, and idempotent on a suitable partial order. Thus usual algorithms based on chaotic iterations [1] can be used, with the assurance of terminating and having a unique equivalent result which is independent of the order in which constraints are considered. However, this can generate two versions of the problem (of which one is NAC and the other one is PAC) which could be impossible to reconcile into a single bipolar problem.

The problem can be solved by achieving only partial forms of PAC and NAC in a bipolar problem. The basic idea is to consider the given bipolar problem, apply the NAC and PAC algorithms to its negative and positive versions, and then modify the preferences of the original problem only when the two new versions can be reconciled, that is, when at least one of the two new preferences is the indifference element. In fact, this means that, in one of the two consistency algorithms, no change has been made. If this holds, the other preference is used to modify the original one. This algorithm achieves a partial form of BAC, that we call p-BAC, and assures equivalence.

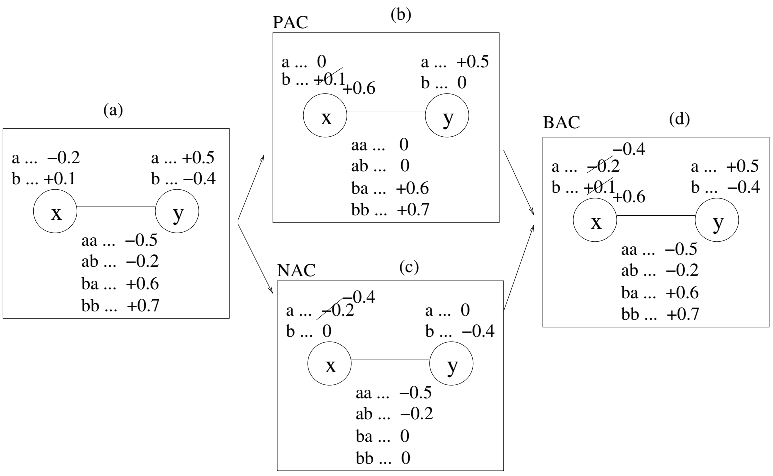

Fig. 1. How to make a bipolar constraint p-BAC

In Figure 1 it is shown how to make a bipolar constraint partially Bipolar Arc-Consistent. Part (a) shows a bipolar constraint, named c_X, over variable X, a bipolar constraint, named c_Y, over variable Y, and a bipolar constraint, named c_{XY}, linking X and Y. Preferences are modelled by the bipolar preference structure ($N = [-1, 0]$, $P = [0, 1]$, $+ = max$, $\times$, $\bot = -1$, $\Box = 0$, $\top = 1$), where $\times$ is such that $\times_p = max$, $\times_n = min$ and $\times_{np} = sum$. Since preferences are given independently in c_X, c_Y and c_{XY}, it is possible to give a low positive preference for a value of X (e.g., $X = b$) in c_X, a negative preference for a value of Y (e.g., $X = b$) in C_Y, but an high positive preference for the combination of such values in c_{XY}. In Part (b) we present the positive version of c_{XY}, that becomes PAC, by increasing the positive preference associated to

$X = b$ from $+0.1$ to $+0.6$. Part (c) presents the negative version of c_{XY}, that becomes NAC, by decreasing the negative preference associated to $X = a$ from -0.2 to -0.4. In Part (d) we show how to achieve p-BAC of c_{XY}. For obtaining p-BAC we must reconcile the modified preferences obtained in Part (b) and in Part (c) when it is possible. Since in this example it is always possible to reconcile such preferences, we obtain a bipolar constraint which is not only p-BAC, but also BAC.

In this approach we require idempotency of $\times_p$ and $\times_n$. However, we could apply arc-consistency also when such operators are not idempotent, by following the extended version of arc-consistency presented in [10,7,18].

Notice that our algorithm will possibly decrease some negative preferences and increase some positive preferences. Therefore, if we use constraint propagation to improve the bounds in a BB algorithm, it will actually sometimes produce worse bounds, due to the increase of the positive preferences. We will thus use only the propagation of negative preferences (that is, NAC) within a BB algorithm. Since the upper bound is just a combination of several preferences, and since preference combination is monotonic, lower preferences give a lower, and thus better, upper bound.

9 Related and Future Work

Bipolarity is an important topic in several fields, such as psychology [9,16,19,20] and multi-criteria decision making [15] and it has recently attracted interest in the AI community, especially in argumentation [17] and qualitative reasoning [2,3,11,12]. These works consider how two alternatives should be compared, given for each a set of positive arguments and a set of negative ones, but they don't analyze the question of combinatorial choice.

Two works in qualitative reasoning which are directly related to our approach are those described in [2,3]. In such papers a bipolar preference model based on a fuzzy-possibilistic approach is described where fuzzy preferences are considered and negative preferences are interpreted as violations of constraints. In this case precedence is given to negative preference optimization, and positive preferences are only used to distinguish among the optimals found in the first phase, thus not allowing for compensation.

Another related work is [14], which considers only totally ordered unipolar and bipolar preference scales, but not partially ordered bipolar scales like us. When the preferences are totally ordered, operators $\times_n$ and $\times_p$ described here correspond respectively to the $t-norm$ and $t-conorm$ used in [14]. Moreover, in [14] it is defined an operator, the $uninorm$, which can be seen as a restricted form of compensation and it is forced to always be associative.

We plan to develop a solver for bipolar CSPs, which should be flexible enough to accommodate for both associative and non-associative compensation operators. The outlined algorithms for BB, NAC, PAC, and p-BAC will also be implemented and tested over classes of bipolar problems.

We also intend to consider the presence of uncertainty in bipolar problems, possibly using possibility theory and to develop solving techniques for such scenarios. Another line of future research is the generalization of other preference formalisms, such as multicriteria methods and CP-nets, to deal with bipolar preferences and to study the relation between bipolarization and tradeoffs.

Acknowledgements

This work has been supported by Italian MIUR PRIN project “Constraints and Preferences” (n. 2005015491).

References

1. Apt, K.R.: Principles of Constraint Programming. Cambridge University Press, Cambridge (2003)
2. Benferhat, S., Dubois, D., Kaci, S., Prade, H.: Bipolar representation and fusion of preferences in the possibilistic logic framework. In: KR 2002, Morgan Kaufmann, San Francisco (2002)
3. Benferhat, S., Dubois, D., Kaci, S., Prade, H.: Bipolar possibility theory in prefrence modeling: representation, fusion and optimal solutions. Information Fusion, an Int. Journal on Multi-Sensor, Multi-Source Information Fusion (2006)
4. Bentham, J.: An Introduction to the Principles of Morals and Legislation. Oxford University Press, Oxford, 1789
5. Bistarelli, S., Montanari, U., Rossi, F.: Semiring-based constraint solving and optimization. Journal of the ACM 44(2), 201–236 (1997)
6. Bistarelli, S., Pini, M.S., Rossi, F., Venable, K.B.: Bipolar preference problems. In: ECAI-06 (poster), IOS Press, Amsterdam (2006)
7. Bistarelli, S., Gadducci, F.: Enhancing constraints manipulation in semiring-based formalisms. In: ECAI, pp. 63–67. IOS Press, Amsterdam (2006)
8. Bruns, W., Herzog, J.C.-M.: Rings. Cambridge University Press, Cambridge (1998)
9. Suci, G.J., Osgood, C.E., Tannenbaum, P.H.: The Measurement of Meaning. University of Illinois Press, Chicago (1957)
10. Cooper, M., Schiex, T.: Arc consistency for soft constraints. AI Journal 154(1-2), 199–227 (2004)
11. Dubois, D., Fargier, H.: On the qualitative comparison of sets of positive and negative affects. In: Godo, L. (ed.) ECSQARU 2005. LNCS (LNAI), vol. 3571, pp. 305–316. Springer, Heidelberg (2005)
12. Dubois, D., Fargier, H.: Qualitative decision making with bipolar information. In: KR’06, pp. 175–186 (2006)
13. Ehrgott, M., Gandibleux, X.: Multiple Criteria Optimization. State of the art annotated bibliographic surveys. Kluwer Academic Publishers, Dordrecht (2002)
14. Grabisch, M., de Baets, B., Fodor, J.: The quest for rings on bipolar scales. Int. Journ. of Uncertainty, Fuzziness and Knowledge-Based Systems (2003)
15. Grabisch, M., Labreuche, C.: Bi-capacities - parts i and ii. Fuzzy Sets and Systems 151, 211–260 (2005)
16. Gardner, W.L., Cacioppo, J.T., Berntson, G.G.: Beyound bipolar conceptualizations and measures: the case of attitudes and evaluative space. Personality and Social Psychology Review 1, 3–25 (1997)
17. Bonnenfon, J.F., Amgoud, L., Prade, H.: An argumentation-based approach to multiple criteria decision. In: Godo, L. (ed.) ECSQARU 2005. LNCS (LNAI), vol. 3571, pp. 10–17. Springer, Heidelberg (2005)
18. Larrosa, J., Schiex, T.: In the quest of the best form of local consistency for weighted csp. In: Proc. 18th IJCAI, pp. 239–244. Morgan Kaufman, San Francisco (2003)
19. Peters, E., Slovic, P., Finucane, M., MagGregor, D.G.: Rational actors or rational fools? Implications of the affect heuristic for behavioural economics. The Journal of Socio-Economics 31, 329–342 (2002)
20. Tversky, A., Kahneman, D.: Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty 5, 297–323 (1992)