

# Tractable Pareto Optimization of Temporal Preferences

Lina Khatib<sup>1,2</sup> Paul Morris<sup>2</sup> Robert Morr  
1. Kestrel Technology  
2. Computational Sciences Division  
NASA Ames Research Center, MS 269-2  
Moffett Field, CA 94035

Kristen Brent Venable  
Department of Pure and Applied Mathematics  
University of Padova  
via G. Belzoni, 7  
35131 Padova, Italy

## Abstract

This paper focuses on temporal constraint problems where the objective is to optimize a set of local preferences for when events occur. In previous work, a subclass of these problems has been formalized as a generalization of Temporal CSPs, and a tractable strategy for optimization has been proposed, where global optimality is defined as maximizing the minimum of the component preference values. This criterion for optimality, which we call "Weakest Link Optimization" (WLO), is known to have limited practical usefulness because solutions are compared only on the basis of their worst value; thus, there is no requirement to improve the other values. To address this limitation, we introduce a new algorithm that re-applies WLO iteratively in a way that leads to improvement of all the values. We show the value of this strategy by proving that, with suitable preference functions, the resulting solutions are Pareto Optimal.

## 1 Introduction

The notion of *softness* has been applied to either a constraint or planning goal, indicating that either can be satisfied to matters of degree. It is not hard to find applicable real world problems for such a notion. For example, in an earth orbiting spacecraft, sensitive instruments like imagers have *duty cycles*, which impose restrictions on the amount of use of the instrument. A duty cycle is typically a complex function based on both the expected lifetime of the instrument, as well as short term concerns such as the amount of heat it can be exposed to while turned on. Duty cycles impose constraints on the duration of the periods for which the instrument can be on, but it is natural to view this duration as flexible. For example, this restriction might be waived to capture an important event such as an active volcano. Thus, the flexibility of the duty cycle "softens" the constraint that the instrument cannot be on beyond a certain duration. Reasoning about soft constraints for planning or scheduling is for the purpose of finding a solution that satisfies the constraints to the highest degree possible.

For temporal reasoning problems, a simple method for evaluating the global temporal preference of a solution to

a Temporal CSP involving local temporal preferences was introduced in [Khatib *et al.*, 2001], based on maximizing the minimally preferred local preference for a time value. Because the locally minimally preferred assignment can be viewed as a sort of "weakest link" with respect to the global solution, we dub this method "weakest link optimization" (WLO), in the spirit of the television game show. WLO can be formalized using a generalization of Simple Temporal Problems (STPs), called STPs with Preferences (STPPs), that preserves the capability to tractably solve for solutions (with suitable preference functions associated with the temporal constraints). Unfortunately, as often occurs, this efficiency has a price. Specifically, WLO offers an insufficiently fine-grained method for comparing solutions, for it is based on a single value, viz., the "weakest link." It is consequently easy to conceive of examples where WLO would accept intuitively inferior solutions because of this myopic focus. Although it is possible to consider more robust alternatives to a WLO strategy for evaluating solutions, it is not clear whether any of these methods would preserve the computational benefits of WLO. This impasse is the starting point of the work described in this paper.

We propose here to make WLO more robust by combining it with an iterative strategy for solving STPPs. The process involves repeatedly restricting temporal values for the weakest links, resetting their preference values, and applying the WLO procedure to the reduced problem that results from these changes. The intuition is a simple one, and we motivate this technique with an example from a Mars Rover planning domain. In Section 2, we summarize the soft constraint problem solver based on WLO introduced previously. We then illustrate in section 3 the deficiencies of WLO on a simple example, which also reveals the intuition underlying the proposed strategy for overcoming this deficiency. The main contribution of this paper is discussed in sections 4 and 5, which formalize this strategy and prove that any solution generated by an application of this strategy is in the set of Pareto optimal solutions for the original problem.

## 2 Reasoning about preferences with soft constraints

This section reviews the material first presented in [Khatib *et al.*, 2001]. There, a class of constrained optimization prob-

lems, called Temporal Constraint Satisfaction Problems with Preferences (TCSPPs), was first defined. A TCSP is a generalization of classical TCSPs which allows for a representation of soft constraints. In classical TCSPs [Dechter *et al.*, 1991], a unary constraint over a variable  $X$  representing an event restricts the domain of  $X$ , representing its possible times of occurrence; the constraint is then shorthand for  $(a_1 \leq X \leq b_1) \vee \dots \vee (a_n \leq X \leq b_n)$ . Similarly, a binary constraint over  $X$  and  $Y$  restricts the values of the distance  $Y - X$ , in which case the constraint can be expressed as  $(a_1 \leq Y - X \leq b_1) \vee \dots \vee (a_n \leq Y - X \leq b_n)$ . A uniform, binary representation of all the constraints results from introducing a variable  $X_0$  for the *beginning of time*, and recasting unary constraints as binary constraints involving the distance  $X - X_0$ .

A *soft temporal constraint* is a pair  $\langle I, f \rangle$ , where  $I$  is a set of intervals  $\{[a, b], a \leq b\}$  of temporal values, and  $f$  is a function from  $(U^I$  to a set  $A$  of values. Intuitively,  $f$  expresses local preferences for temporal values based on the value it assigns from  $A$ . For example, the soft constraint represented by  $\{ \{[1, 3], [8, 12]\}, \min \}$  can be interpreted to mean that the temporal assignments must be selected from either of the intervals in the set, and the function  $\min$  assigns a greater preference to smaller values. The cardinality of the set  $A$ , i.e., the number of distinct preference values, reflects the ability to discriminate among degrees of preference for temporal assignments. The class of TCSPPs in which each soft constraint consists of a single interval is called *Simple Temporal Problems with Preferences* (STPPs).

Local preferences combine to form global preferences for complete assignments. To formalize these operations,  $A$  can be structured in the form of a *c-semiring* [Bistarelli *et al.*, 1997]. A *semiring* is a tuple  $\langle A, +, \times, 0, 1 \rangle$  such that

- $A$  is a set and  $0, 1 \in A$ ;
- $+$ , the additive operation, is commutative, associative and 0 is its identity element ( $a+0 = a$ );
- $\times$ , the multiplicative operation, is associative, distributes over  $+$ , 1 is its identity element and 0 is its absorbing element ( $a \times 0 = 0$ ).

A *c-semiring* is a semiring in which  $+$  is idempotent (i.e.,  $a + a = a, a \in A$ ), 1 is its absorbing element, and  $\times$  is commutative. The semi-ring representation of operations on preference values is used as part of the proof of tractability for restricted sub-classes of TCSP, which occurs below.

A *solution* to a TCSP is a complete assignment to all the variables that satisfies the temporal constraints. An arbitrary assignment of values to variables has a *global preference value*, obtained by combining the local preference values using the semiring operations. A *c-semiring* induces a partial order relation  $\leq_S$  over  $A$  to compare preference values of arbitrary assignments;  $a \leq_S b$  can be read *b is more preferred than a*. Classical Temporal CSPs can be seen as a special case of TCSP, with "soft" constraints that assign the "best" preference value to each element in the domain, and the "worst" value to everything else. The optimal solutions of a TCSP are those solutions which have the best preference value in terms of the ordering  $\leq_S$ .

*Weakest Link Optimization* (WLO) is formalized via the semiring  $S_{WLO} = \langle A, \max, \min, 0, 1 \rangle$ , that is, where for  $a, b \in A$ ,  $a + b = \max(a, b)$  and  $a \times b = \min(a, b)$ , and 1 (0) is the best (worst) preference value. Given a solution  $t$  in a TCSP with semiring  $S_{WLO}$ , let  $T_{ij} = \langle I_{i,j}, f_{i,j} \rangle$  be a soft constraint over variables  $X_i, X_j$  and  $(v_i, v_j)$  be the projection of  $t$  over the values assigned to variables  $X_i$  and  $X_j$  (abbreviated as  $(v_i, v_j) = t_{\downarrow X_i, X_j}$ ). The corresponding preference value given by  $f_{ij}$  is  $f_{ij}(v_j - v_i)$ , where  $v_j - v_i \in I_{i,j}$ . The global preference value of  $t$ ,  $val(t)$ , is defined as  $val(t) = \min\{f_{ij}(v_j - v_i) \mid (v_i, v_j) = t_{\downarrow X_i, X_j}\}$ . Thus, a "weakest link value" for a solution  $t$  is any minimum  $f_{ij}(v_j - v_i)$  that determines  $val(t)$ , and the *WLO-optimal solutions* to a problem are the ones that have a maximum weakest link value.

As with classical (binary) CSPs, TCSPPs can be arranged to form a network of nodes representing variables, and edges labeled with constraint information. Given a network of soft constraints, under certain restrictions on the properties of the semiring, it can be shown that local consistency techniques can be applied in polynomial time to find an equivalent minimal network in which the constraints are as explicit as possible. The restrictions that suffice for this result apply to

1. the "shape" of the preference functions used in the soft constraints;
2. the multiplicative operator  $\times$  (it should be idempotent); and
3. the ordering of the preference values ( $<_S$  must be a total ordering).

The class of restricted preference functions that suffice to guarantee that local consistency can be meaningfully applied to soft constraint networks is called *semi-convex*. This class includes linear, convex, and also some step functions. All of these functions have the property that if one draws a horizontal line anywhere in the Cartesian plane of the graph of the function, the set of  $X$  such that  $f(X)$  is not below the line forms an interval. Semi-convexity is preserved under the operations performed by local consistency (intersection and composition). STPPs with semiring  $S_{WLO}$  can easily be seen to satisfy these restrictions.

The same restrictions that allow local consistency to be applied are sufficient to prove that STPPs can be solved tractably. Finding an optimal solution of the given STPP with semi-convex preference functions reduces to a two-step search process consisting of iteratively choosing a preference value, "chopping" every preference function at that point, then solving a STP defined by considering the interval of temporal values whose preference values lies above the chop line (semi-convexity ensures that there is a single interval above the chop point, hence that the problem is indeed an STP). Figure 1 illustrates the chopping process. It has been shown that the "highest" chop point that results in a solvable STP in fact produces an STP whose solutions are exactly the optimal solutions of the original STPP. Binary search can be used to select candidate chop points, making the technique for solving the STPP tractable. The second step, solving the induced STP, can be performed by transforming the graph associated




Figure 1: "Chopping" a semi-convex function

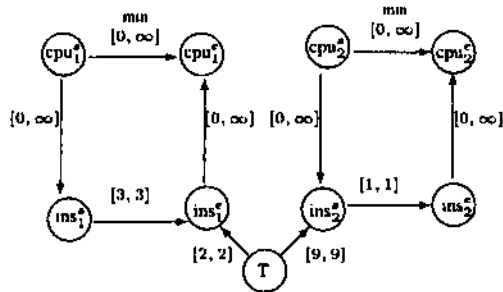


Figure 2: The STPP for the Rover Science Planning Problem where T is any timepoint

with this STP into a distance graph, then solving two single-source shortest path problems on the distance graph. (The solutions to these provide upper and lower time bounds for each event. If the problem has a solution, then for each event it is possible to arbitrarily pick a time within its time bounds, and find corresponding times for the other events such that the set of times for all the events satisfy the interval constraints.) For  $N$  nodes and  $E$  edges, the complexity of this phase is  $O(EN)$  (using the Bellman-Ford algorithm [Cormen *et al*, 1990]).

### 3 The problem with WLO

Formalized in this way, WLO offers what amounts to a coarse method for comparing solutions, one based on the minimal preference value over all the projections of the solutions to local preference functions. Consequently, the advice given to a temporal solver by WLO may be insufficient to find solutions that are intuitively more globally preferable. For example, consider the following simple Mars rover planning problem, illustrated in Figure 2. The rover has a sensing instrument and a CPU. There are two sensing events, of durations 3 time units and 1 time unit (indicated in the figure by the pairs of nodes labeled  $ins_1^e, ins_1^s$  and  $ins_2^e, ins_2^s$  respectively). There is a hard temporal constraint that the CPU be on while the instrument is on, as well as a soft constraint that the CPU should be on as little as possible, to conserve power. This constraint is expressed in the STPP as a function from temporal values indicating the duration that the CPU is on, to preference values. For simplicity, we assume that the preference function  $min$  on the CPU duration constraints is the negated identity function; i.e.,  $min(t) = -t$ ; thus higher preference values, i.e. shorter durations, are preferred. Because the CPU must be on at least as long as the sensing events, any globally preferred solution using WLO has preference value -3. The set of solu-

tions that have the optimal value includes solutions in which the CPU duration for the second sensing event varies from 1 to 3 time units. The fact that WLO is unable to discriminate between the global values of these solutions, despite the fact that the one with 1 time unit is obviously preferable to the others, is a clear limitation of WLO.

One way of formalizing this drawback of WLO is to observe that a WLO policy is not *Pareto Optimal*. To see this, we reformulate the set of preference functions of a STPP,  $f_1, \dots, f_m$  as criteria requiring simultaneous optimization, and let  $s = [t_1, \dots, t_n]$  and  $s' = [t'_1, \dots, t'_m]$  be two solutions to a given STPP.  $s'$  dominates  $s$  if for each  $j$ ,  $f_j(t_j) \leq f_j(t'_j)$  and for some  $k$ ,  $f_k(t_k) < f_k(t'_k)$ . In a Pareto optimization problem, the *Pareto optimal set* of solutions is the set of non-dominated solutions. Similarly, let the *WLO-optimal set* be the set of optimal solutions that result from applying the chopping technique for solving STPPs described above. Clearly, applying WLO to an STPP does not guarantee that the set of WLO-optimal solutions is a Pareto optimal set. In the rover planning problem, for example, suppose we consider only solutions where the CPU duration for the first sensing event is 3. Then the solution in which the CPU duration for the second sensing event is 1 time unit dominates the solution in which it is 2 time units, but both are WLO-optimal, since they have the same weakest link value.<sup>1</sup>

Assuming that Pareto-optimality is a desirable objective in optimization, a reasonable response to this deficiency is to replace WLO with an alternative strategy for evaluating solution tuples. A natural, and more robust alternative evaluates solutions by summing the preference values, and ordering them based on preferences towards larger values. (This strategy would also ensure Pareto optimality, since every maximum sum solution is Pareto optimal.) This policy might be dubbed "utilitarian." The main drawback to this alternative is that the ability to solve STPPs tractably is no longer apparent. The reason is that the formalization of utilitarianism as a semiring forces the multiplicative operator (in this case, *sum*), not to be idempotent (i.e.,  $a + a \neq a$ ), a condition required in the proof that a local consistency approach is applicable to the soft constraint reasoning problem.

Of course, it is still possible to apply a utilitarian framework for optimizing preferences, using either local search or a complete search strategy such as branch and bound. Rather than pursuing this direction of resolving the problems with WLO, we select another approach, based on an algorithm that interleaves flexible assignment with propagation using WLO.

### 4 An algorithm for Pareto Optimization

The proposed solution is based on the intuition that if a constraint solver using WLO could iteratively "ignore" the weakest link values (i.e. the values that contributed to the global solution evaluation) then it could eventually recognize solutions that dominate others in the Pareto sense. For example, in the Rover Planning problem illustrated earlier, if the weakest link value of the global solution could be "ignored," the WLO solver could recognize that a solution with the CPU on

<sup>1</sup>This phenomenon is often referred to in the literature as the "drowning effect."

Inputs: STPP  $P = (V, C)$

Output:

STP  $(V, C_P)$  whose solutions are Pareto optimal for  $P$ .

(1)  $C_P = C$

(2) Do

(3) Solve  $(V, C_P)$  using WLO

(4) Delete all weakest link soft constraints from  $C_P$

(5) For each deleted constraint  $\{(a, b), f\}$ ,

(6) add  $\{[a_{opt}, b_{opt}], f_{best}\}$  to  $C_P$

(7) until (3)-(6) leave  $(V, C_P)$  unchanged

(8) Return  $(V, C_P)$

Figure 3: STPP solver WLO+ returns a solution in the Pareto optimal set of solutions

for 1 time unit during the second instrument event is to be preferred to one where the CPU is on for 2 or 3 time units.

We formalize this intuition by a procedure wherein the original STPP is transformed by iteratively selecting what we shall refer to as a *weakest link constraint*, changing the constraint in such a way that it can effectively be "ignored," and solving the transformed problem. A weakest link (soft) constraint for a WLO set of solutions is one in which the preference value of its duration in all the WLO solutions is the same as the chop level  $v$  of the optimal STP using WLO. For example, after applying WLO to the problem in Figure 2, the CPU duration constraint associated with the first sensing event will be a weakest link, since it now has a fixed preference value of -3. However, the CPU constraint for the second event will not be a weakest link since its preference value can still vary from -3 to -1.

We also define a weakest link constraint to be *open* if  $v$  is not the "best" preference value (i.e.,  $v < \mathbf{1}$ , where 1 is the designated "best" value among the values in the semi-ring).

Formalizing the process of "ignoring" weakest link values is a two-step process of restricting the weakest links to their intervals of optimal temporal values, while eliminating their WLO restraining influence by resetting their preferences to a single, "best" value. Formally, the process consists of:

- Squeezing the temporal domain to include all and only those values which are optimally preferred; and
- Replacing the preference function by one that assigns the most preferred value (i.e. 1) to each element in the new domain.

The first step ensures that only the best temporal values are part of any solution, and the second step allows WLO to be re-applied to eliminate Pareto-dominated solutions from the remaining solution space.

The algorithm WLO+ (Figure 3) returns a Simple Temporal Problem (STP) whose solutions are contained in the WLO-optimal, Pareto-optimal solutions to the original STPP,  $P$ . Where  $C$  is a set of soft constraints, the STPP  $(V, C_P)$  is solved (step 3) using the chopping approach described earlier. In step 5, we denote the soft constraint that results from the two-step process described above as  $\{[a_{opt}, b_{opt}], f_{best}\}$ , where  $[a_{opt}, b_{opt}]$  is the interval of temporal values that are




Figure 4: Relationships between Solution Spaces for STPPs that are WLO or Pareto Optimal

optimally preferred, and  $f_{best}$  is the preference function such that  $f_{best}(v) = \mathbf{1}$  for any input value  $v$ . Notice that the run time of WLO+ is  $O(|C|)$  times the time it takes to execute  $Solve(V, C_P)$ , which is a polynomial.

We now proceed to prove the main result, in two steps. In this section we assume the existence of weakest links at every iteration of the WLO+ algorithm, and show that the subset of solutions of the input STPP returned by WLO+ is contained in the intersection of WLO-optimal and Pareto-optimal solutions. In the next section we show that, given additional restrictions on the shape of the preference functions, such weakest links can be shown to always exist.

Given an STPP  $P$ , let  $Sol_{WLO}(P)$  (resp.  $Sol_{PAR}(P)$ ) be the set of WLO-optimal (respectively, Pareto-Optimal) solutions of  $P$ , and let  $Sol_{WLO+}(P)$  be the set of solutions to  $P$  returned by WLO+. Then the result can be stated as follows.

**Theorem 1** *If a weakest link constraint is found at each stage of WLO+, then  $Sol_{WLO+}(P) \subseteq Sol_{WLO}(P) \cap Sol_{PAR}(P)$ . Moreover, if  $P$  has any solution, then  $Sol_{WLO+}(P)$  is nonempty.*

Proof:

First note that after an open weakest link is processed in steps (4) to (6), it will never again be an open weakest link (since its preference is reset to  $f_{best}$ ). Since the theorem assumes a weakest link constraint is found at each stage of WLO+, the algorithm will terminate when the weakest link constraint is not open, i.e., when all the soft constraints in  $C_P$  have WLO preferences that equal the best (1) value.

Now assume  $s \in Sol_{WLO+}(P)$ . Since the first iteration reduces the set of solutions of  $(V, C_P)$  to  $Sol_{WLO}(P)$ , and each subsequent iteration either leaves the set unchanged or reduces it further, it follows that  $s \in Sol_{WLO}(P)$ . Now suppose  $s \notin Sol_{PAR}(P)$ . Then  $s$  must be dominated by a Pareto optimal solution  $s'$ . Let  $c$  be a soft constraint in  $C$  for which  $s'$  is superior to  $s$ . Thus, the preference value of the duration assigned by  $s$  to  $c$  cannot be 1. It follows that at some point during the course of the algorithm,  $c$  must become an open weakest link. Since  $s$  is in  $Sol_{WLO+}(P)$ , it survives until then, and so it must provide a value for  $c$  that is equal to the chop level. However, since  $s'$  dominates  $s$ ,  $s'$  must also survive until then. But this contradicts the assumption that  $c$  is a weakest link constraint, since  $s'$  has a value greater than the WLO chop level. Hence,  $s$  is in  $Sol_{PAR}(P)$ , and so in  $Sol_{WLO}(P) \cap Sol_{PAR}(P)$ .

Next suppose the original STPP  $P$  has at least one solution. To see that  $Sol_{WLO+}(P)$  is nonempty, observe that the modifications in steps (4) to (6), while stripping out solutions




Figure 5: A unique WLO+ Solution.

that are not WLO optimal with respect to  $(V, C_P)$ , do retain all the WLO optimal solutions. Clearly, if there is any solution, there is a WLO optimal one. Thus, if the  $(V, C_P)$  in any iteration has a solution, the  $(V, C_P)$  in the next iteration will also have a solution. Since we are assuming the first  $(V, C_P)$  ( $— (V, C)$ ) has a solution, it follows by induction that  $Sol_{WLO+}(P)$  is nonempty.  $\square$

The theorem shows that it is possible to maintain the tractability of WLO-based optimization while overcoming some of the restrictions it imposes. In particular, it is possible to improve the quality of the flexible solutions generated within an STPP framework from being WLO optimal to being Pareto optimal.

Although the algorithm determines a nonempty set of solutions that are both WLO optimal and Pareto optimal, the set might not include all such solutions. Consider the example in figure 5. Assume the preference function for all soft constraints is given by  $f(t) = t$ , i.e., longer durations are preferred (signified by the *max* label on the edges). The WLO+ algorithm will retain a single solution where BC and CD are both 5. However, the solution where BC = 2 and CD = 8, which is excluded, is also both Pareto optimal and WLO optimal. (Note that AB, with a fixed value of 1, is the weakest link.)

Many optimization schemes seek what is known as *utilitarian* optimality, where the objective is to maximize the sum of the local preferences. However, the WLO+ solutions are not necessarily utilitarian optimal with respect to all solutions or even the WLO solutions. For example, in figure 5, if the preference function is  $f(t) = t^2$ , a utilitarian optimal WLO solution would be given by BC = 1 and CD = 9, but WLO+ will still return the solution where BC and CD are both 5.

We can summarize the position taken in this paper by saying that utilitarian strategies, while attractive in many ways, are apparently intractable. The WLO+ approach provides some of the same benefit at lower cost. For example, non-competing constraints are fully optimized by WLO+. For competing constraints, WLO+ tends to divide the preferences as equally as possible. In some applications, this might be more desirable than a utilitarian allocation.

## 5 Existence of Weakest Links

In this section we show that under suitable conditions, a weakest link constraint always exists. This involves a stronger requirement than for WLO: the preference functions must be convex, not merely semi-convex. This would include linear functions, cycloids, and upward-pointing parabolas, for example, but not Gaussian curves, or step functions. (Later on, we will see this requirement can be relaxed somewhat so that Gaussians can be permitted.)

Before proceeding, we note that while a solution  $s$  of an

STP  $P$  is defined in terms of an assignment to each variable, it also determines a value  $s(e)$  for each edge  $e$ , given by  $s(e) = s(Y) - s(X)$  where  $X$  and  $Y$  are the start and end variables of  $e$ , respectively. We will use this notation in what follows.

Now consider any consistent STP  $P$ . The *minimal network* [Dechter *et al.*, 1991] corresponding to  $P$  is another STP  $P'$ . The constraints between any two points  $X$  and  $Y$  in  $P'$  are formed by intersecting the constraints induced by all possible paths between  $X$  and  $Y$  in  $P$ .

In the following, a preference function  $f$  is said to be *convex* if  $\{ \langle x, y \rangle \mid y \leq f(x) \}$  is a convex set. The claim of the existence of weakest links can be stated as follows:

**Theorem 2** *Let  $P$  be an STPP with continuous domains and convex preference functions. Then there will be at least one weakest link constraint for the WLO optimal set of solutions.*

*Proof:*

Consider the (minimal) STP  $P_{opt}$  that corresponds to the optimal chopping level for  $P$  (as described in the WLO algorithm). Suppose there is no weakest link constraint. Then for each edge constraint  $c$  there is a solution  $s_c$  to  $P_{opt}$  such that  $f(s_c(c)) > opt$ , where  $f$  is the preference function for the edge.<sup>2</sup>

Let  $\bar{s}$  be the average of all the  $s_c$  solutions, i.e.

$$\bar{s}(X) = 1/|E| \sum_{c \in E} s_c(X) \text{ for all } X \in V$$

where  $E$  is the set of edges and  $V$  is the set of temporal variables. By the linearity of the constraints, it is easy to show that  $\bar{s}$  is also a solution to  $P_{opt}$ . For example, if  $s_c(Y) - s_c(X) \geq a$  for all  $e$ , then  $\bar{s}(Y) - \bar{s}(X) = 1/|E| \sum_{c \in E} (s_c(Y) - s_c(X)) \geq (1/|E|)|E|a = a$ .

Since  $\bar{s}$  is a solution, we must have  $f(\bar{s}(e)) \geq opt$  for all edges  $e$ . Notice, however, that there must be some edge  $e$  such that  $f(\bar{s}(e)) = opt$ , otherwise  $\bar{s}$  would be a solution with value greater than the optimal value. It follows that  $\bar{s}(e) \neq s_c(e)$ , since we already know that  $f(s_c(e)) > opt$ .

Thus, either  $\bar{s}(e) < s_c(e)$  or  $\bar{s}(e) > s_c(e)$ . We consider only the case where  $\bar{s}(e) < s_c(e)$ . (The proof is similar in the other case.) Note also that  $\min_{c \in E} s_c(e) < \bar{s}(e)$ , since the minimum of a set of numbers can only be equal to the average if all the numbers are equal. It follows that  $\min_{c \in E} s_c(e) < \bar{s}(e) < s_c(e)$ . However,  $f(\min_{c \in E} s_c(e)) \geq opt$ ,  $f(\bar{s}(e)) = opt$ , and  $f(s_c(e)) > opt$ . This violates the convexity of  $f$ , which establishes the theorem.  $\square$

The operations in the WLO+ algorithm preserve the convexity property of the preference functions. Each stage of WLO+ repeats a WLO calculation. Thus, theorem 2 implies

**Corollary 2.1** *Suppose  $P$  is an STPP with continuous domains and convex preference functions. Then a weakest link is found at each iteration of WLO+.*

An example of why the existence result does not apply more generally to semi-convex functions is found in figure 6. The

<sup>2</sup>To avoid excessive subscripting, we suppress the implied  $e$  subscript on  $f$  here and in what follows. In all cases, the applicable preference function will be clear from the context.




Figure 6: An STPP with no weakest link

STPP in the figure contains two semi-convex step-like functions with optimal preference values associated with durations  $t$  and  $t'$ . Assume the STPP is minimal, and that the assignment  $e = t, e' = t'$  is inconsistent. Then the highest possible chop point is  $p$ , and no weakest link exists, i.e., for neither  $e$  nor  $e'$  is it the case that, for every solution  $s$ ,  $p$  is the value returned by the preference function associated with that constraint for the duration assigned by  $s$ .

## 6 Discussion and Related Work

An examination of the proof of theorem 2 shows that the weakest link constraint exists under a somewhat less restrictive condition than convexity: it is enough, assuming semi-convexity, to require that plateaus (subintervals of non-zero length where the preference function is constant) can only occur at the global maximum of the preference function. This means, for example, that the theorem is applicable in principle to any semi-convex smooth function such as a Gaussian curve.

However, in the practical setting of a computer program where numbers are computed to a finite precision and continuous curves are approximated, some adjustments may need to be made. Note that a representation as a discretized step function does not satisfy the no-plateau condition. An alternative is to treat a discretized function as corresponding to a piecewise linear function where the linear segments join successive points on the discretized graph. Even there, the long tails of a Gaussian curve may get approximated by horizontal segments. However, generally we can trim the domain of the curve to eliminate the flat tails without excluding all the solutions. In that case, the discretized Gaussian is acceptable. (Note that figure 6 could be simulated by an example involving extreme Gaussians where the tails are essential for the solution.)

Note that preferences such as longest or shortest durations, or closest to a fixed time, which appear to be the most useful in practice, can be easily modeled within this framework.

WLO+ has been implemented and tested on randomly generated problems, where each semi-convex preference function is a quadratic  $ax^2 + bx + c$ , with randomly selected parameters and  $a \leq 0$ . We compared the best solution found after applying WLO+ with the quality of the earliest solution found using the chop solver, using the utilitarian measure of quality (i.e., summing preference values). An average improvement of between 6 and 10% was observed, depending on constraint density (more improvement on lower density problems). Future research will focus on the application of WLO+ to the rover science planning domain.

The results described here are clearly relevant to any effort whose objective is representing and reasoning about preferences and utility. A detailed survey of this vast literature is clearly beyond our scope; here we provide pointers to work that exhibits significant overlap. First, the idea of extending CSPs to solve multi-criteria optimization problems is proposed in [Torrens and Faltings, 2002]; this work also uses Pareto-optimality as a criterion for ordering solutions. Second, the idea of applying the notion of degrees of satisfaction to solving temporal reasoning problems has been applied previously [Dubois and Prade, 1989]. Third, a number of graphical-based representations of local preferences have appeared; in [Bacchus and Grove, 1995], for example, an approach is taken based on drawing connections between preferences and probabilities, as expressed in a Bayesian network. Finally, for a survey of AI-based approaches to preferences and utility, with an emphasis on qualitative approaches, the reader is referred to [Doyle and Thomason, 1999].

## 7 Summary

This paper has presented a reformulation of problems in the optimization of temporal preferences using a generalization of Temporal CSPs. The practical context from which this effort arose is temporal decision-making in planning, where associated with domains representing temporal distances between events is a function expressing preferences for some temporal values over others. The work here extends previous work by overcoming limitations in the approach that arose when considerations of efficiency in reasoning with preferences resulted in coarseness in the evaluation procedure for global temporal assignments.

## References

- [Bistarelli et al., 1997] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based Constraint Solving and Optimization. *Journal of the ACM*, 44(2): 201-236, March 1997.
- [Bacchus and Grove, 1995] F. Bacchus and A. Grove. Graphical Models of Preferences and Utility. *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, 3-10, Morgan Kaufmann, 1999.
- [Cormen et al., 1990] T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT press, Cambridge, MA, 1990.
- [Dechter et al., 1991] R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*, 49, 61-95, 1991.
- [Doyle and Thomason, 1999] J. Doyle and R. Thomason. Background to Qualitative Decision Theory. *AI Magazine*, 55-68, Summer, 1999.
- [Dubois and Prade, 1989] D. Dubois and H. Prade. Processing Fuzzy Temporal Knowledge. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(4), 1989.
- [Khatib et al., 2001] L. Khatib, P. Morris, R. Morris and F. Rossi. Temporal Reasoning about Preferences. In *Proceedings of IJCAI-01*. Morgan Kaufman, 1995.
- [Torrens and Faltings, 2002] M. Torrens and B. Faltings. Using Soft CSPs for Approximating Pareto-Optimal Solution Sets. In *AAAI Workshop Proceedings Preferences in AI and CP: Symbolic Approaches*, AAAI Press, 2002.