Strategies for Global Optimization of Temporal Preferences

Paul Morris[3] Robert Morris[3] Lina Khatib[1,3]
Sailesh Ramakrishnan[2,3] Andrew Bachmann[2,3]

1. Kestrel Technology
2. QSS Group Inc.
3. Computational Sciences Division

NASA Ames Research Center
Moffett Field, CA 94035
pmorris@email.arc.nasa.gov

Abstract. A temporal reasoning problem can often be naturally characterized as a collection of constraints with associated local preferences for times that make up the admissible values for those constraints. Globally preferred solutions to such problems emerge as a result of well-defined operations that compose and order temporal assignments. The overall objective of this work is a characterization of different notions of global temporal preference within a temporal constraint reasoning framework, and the identification of tractable sub-classes of temporal reasoning problems incorporating these notions. This paper extends previous results by refining the class of useful notions of global temporal preference that are associated with problems that admit of tractable solution techniques. This paper also resolves the hitherto unanswered question of whether the solutions that are globally preferred from a *utilitarian* criterion for global preference can be found tractably. A technique is described for identifying and representing the entire set of utilitarian-optimal solutions to a temporal problem with preferences.

1 Introduction

Many temporal reasoning problems can be naturally characterized as collections of constraints with associated local preferences for times that make up the admissible values for those constraints. For example, one class of vehicle routing problems [14] consists of constraints on requested service pick-up or delivery that allow flexibility in temporal assignments around a specified fixed time; solutions with assignments that deviate from this time are considered feasible, but may incur a penalty. Similarly, dynamic scheduling problems [12], whose constraints may change over time, thus potentially requiring solution revision, often induce preferences for revised solutions that deviate minimally from the original schedule.

To effectively solve such problems, it is necessary to be able to order the space of assignments to times based on some notion of global preference, and to have a mechanism to guide the search for solutions that are globally preferred. Such a framework arises as a simple generalization of the Simple Temporal Problem (STP) [5], in which temporal constraints are associated with a local preference function that maps admissible times into values; the result is called *Simple Temporal Problem with Preferences (STPP)* [9]. Globally optimal solutions to STPPs emerge as a result of well-defined operations that compose and order partial solutions.

Different concepts of composition and comparison result in different characterizations of global optimality. Past work has introduced three notions of global preference: Weakest Link (maximize the least preferred time), Pareto, and utilitarian. Much of the work to date has been motivated by the overall goal of finding tractable solutions to temporal optimization problems with realistic global preference criteria. In particular, NASA is motivated to create systems that will automatically find optimally preferred solutions to problems in the rover planning domain [3], where the goal is to devise plans for investigating a number of scientifically promising science targets.

In addition to reviewing the STPP framework (section 2), this paper extends previous results motivated by the overall goal of identifying useful notions of global preference that correspond to problems that can be solved tractably. First, we introduce a new category of global optimality called *stratified egalitarian* optimality, and prove that it precisely characterizes the subset of Pareto optimal solutions returned by a tractable technique called WLO+ introduced previously (section 3). Second, we provide an affirmative answer to the question of whether the utilitarian optimal solutions to temporal preference problems can be also found tractably within this framework. A technique is described for identifying and representing the whole set of utilitarian-optimal solutions to a temporal reasoning problem with preferences (section 4). This paper closes with a summary of experiments (section 5) and a discussion of future work.

2 Simple Temporal Problems with Preferences

A *temporal constraint* determines a restriction on the distance between an arbitrary pair of distinct events. In [9], a *soft temporal constraint* between events i and j is defined as a pair $\langle I, f_{ij} \rangle$, where I is a set of intervals $\{[a,b], a \leq b\}$ and f_{ij} is a *local preference function* from $\bigcup I$ to a set A of admissible preference values. For the purposes of this paper, we assume the values in A are totally ordered, and that A contains designated values for minimum and maximum preferences.

When I is a single interval, a set of soft constraints defines a *Simple Temporal Problem with Preferences* (STPP), a generalization of Simple Temporal Problems [5]. An STPP can be depicted as a pair (V, C) where V is a set of variables standing for temporal events or timepoints, and $C = \{\langle [a_{ij}, b_{ij}], f_{ij} \rangle\}$ is a set of soft constraints defined over V. An STPP, like an STP, can be organized as a network of variables representing events, and links labeled with constraint information. A *solution* to an STPP is a complete assignment to all the variables that satisfies the temporal constraints.

A soft temporal constraint $\langle [a_{ij}, b_{ij}], f_{ij} \rangle$ results from defining a preference function f_{ij} over an interval $[a_{ij}, b_{ij}]$. Clearly, removing the preference functions from the set of constraints making up an STPP *P* results in an STP; we call this the STP *underlying P*.

We define a *preference vector* of all the local preference values associated with a set $F = \{f_{ij}\}$ of local preference functions and a solution S. Formally, let $f_{ij}(S)$ refer to

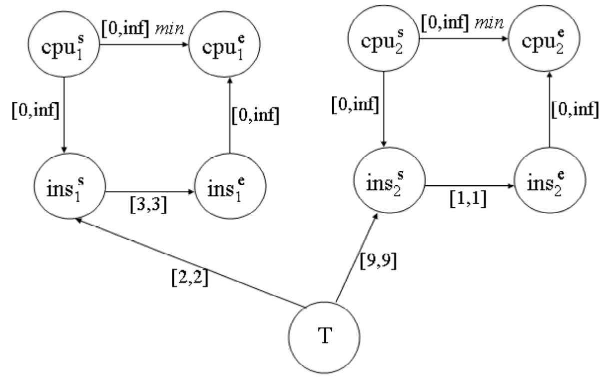

Fig. 1. The STPP for a Rover Science Planning Problem (T is any timepoint)

the preference value assigned by f_{ij} to the temporal value that S assigns to the distance between events i and j, and let

$$\begin{array}{rl} U_{F(S)} = \langle f_{12}(S), & f_{13}(S), \ldots\ f_{1n}(S), \\ & f_{23}(S), \ldots\ f_{2n}(S), \\ & \quad\quad\ddots\ \vdots \\ & \quad\quad\quad f_{n-1,n}(S)\rangle \end{array}$$

be the *preference vector* associated with F and S. In what follows the context will permit us to write U_S instead of $U_{F(S)}$ without ambiguity, and U_S^k will refer to the k^{th} preference value of U_S.

For an example of an STPP, consider a simple Mars rover planning problem, illustrated in Figure 1. 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 $\text{ins}_1^s, \text{ins}_1^e$ and $\text{ins}_2^s, \text{ins}_2^e$ respectively). The event T depicts a reference time point (sometimes referred to as "the beginning of time") that allows for constraints to be specified on the start times for events. 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 possible durations 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., $min_{ij}(t) = -t$; thus higher preference values, i.e. shorter durations, are preferred.

A solution to an STPP has a *global preference value*, obtained by combining the local preference values using operations for composition and comparison. Optimal solutions to an STPP are those solutions which have the best preference value in terms of the ordering induced by the selected comparison operator. Solving STPPs for globally preferred assignments has been shown to be tractable, under certain assumptions about

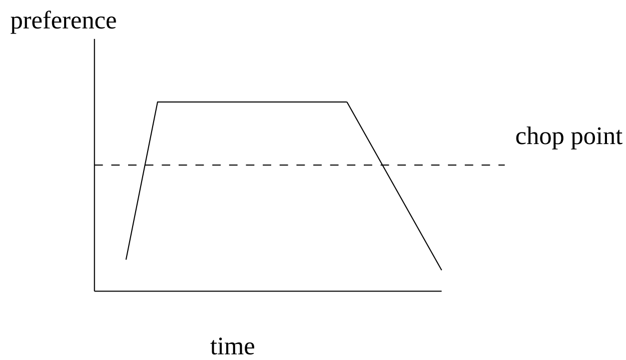

Fig. 2. "Chopping" a semi-convex function.

the "shape" of the local preference functions and about the operations used to compose and compare solutions.

For example, first consider a class of local preference functions that includes any function such 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. This class of *semi-convex* functions includes linear, convex, and also some step functions.

Second, consider an STPP solver based on the notion of *Weakest Link Optimization* (WLO). This framework consists of an operator for composing preference values in A based on the minimal value of the component values. This framework induces an evaluation of solutions based on a single, "weakest link" value. Given preference vectors U_S and $U_{S'}$ corresponding to distinct solutions S and S', we will say that S is *Weakest-Link-Optimal (WLO) -preferred* to S', or S' is *WLO-dominated* by S, if $min(U_{S'}) < min(U_S)$, where $min(U)$ returns the minimum value of vector U. *WLO-optimal solutions* are those to which no other solutions are WLO-preferred.

STPPs with semi-convex preference functions for WLO-optimal solutions can be solved tractably by a process called the *chop method*. This method is based on the act of "chopping" a preference function (Figure 2). Semi-convexity implies that the set of times for which the preference function returns a value above a selected chop point forms a convex interval; call this interval the *chop-induced constraint*. For a set of preference functions in an STPP, chopping all of them at the same preference value induces a Simple Temporal Problem, namely, of finding a set of assignments that satisfies all the chop-induced constraints. A binary search will return the largest preference value v_{opt} for which a solution to the induced STP exists; it can been shown that the solutions at v_{opt} are *WLO-optimal*.

Because the chop method returns the solution to an STP, its output is a *flexible temporal plan*, i.e., a set of solutions that have the same WLO-optimal value. Plan flexibility is often considered important in ensuring robustness in an execution environment that is uncertain [11]. Nonetheless, the WLO criterion for globally preferred solutions has the disadvantage of being "myopic", in the sense that it bases its evaluation on a single value. This feature can be shown to limit its usefulness in solving real temporal planning problems. The rover example in Figure 1 can be used to illustrate this myopia. 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 solutions that have the WLO-optimal value includes solutions in which the CPU duration for the second sensing event varies from 1 to 3 time units (again, since WLO bases its evaluation solely on the least preferred value). 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, can be clearly viewed as a limitation.

Less myopic global preference criteria can be defined. For example, we can say that S' *Pareto-dominates* S if for each j, $U_S^j \leq U_{S'}^j$ and for some k, $U_S^k < U_{S'}^k$. The *Pareto optimal set* of solutions is the set of non-Pareto-dominated solutions. Similarly, we can say that S' *utilitarian-dominates* S if $\sum_j U_S^j < \sum_j U_{S'}^j$, and the *utilitarian optimal set* of solutions is the set of non-utilitarian-dominated solutions.

In a previous result [10], it was shown that a restricted form of Pareto-optimality can be achieved by an iterative application of the chop method. The intuition is that if a constraint solver could "ignore" the links that contribute the weakest link values (i.e. the values that determined the global solution evaluation), then it could eventually recognize solutions that dominate others in the Pareto sense. The links to be ignored are called *weakest link constraints*: formally, they comprise all links in which the optimal value for the preference function associated with the constraint is the same as the WLO value for the global solution. Formalizing the process of "ignoring" weakest link values is a two-step process of committing the flexible solution to consist of the interval of optimal temporal values, and reinforcing this commitment 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 WLO-optimally preferred; and
- replacing the preference function by one that assigns the highest (most preferred) value 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 resulting algorithm, called WLO+ returns, in polynomial time, a Simple Temporal Problem (STP) whose solutions are a nonempty subset of the WLO-optimal, Pareto-optimal solutions to an STPP. The algorithm WLO+ from [10] is reproduced in Figure 3 for completeness. Where C is a set of soft constraints, the STPP (V, C_P) is solved (step 3) using the chop approach. In step 5, we depict the soft constraint that results from the two-step process described above as $\langle [a_{opt}, b_{opt}], f_{best} \rangle$, where $[a_{opt}, b_{opt}]$ is the interval of temporal values that are optimally preferred, and f_{best} is the preference function that returns the most preferred preference value for any input value. 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.

WLO+, applied to the rover example in Figure 1, finds a Pareto optimal solution in two iterations of the while loop. In the first iteration, the weakest link is that between the start and end of the first CPU event. WLO+ deletes this link and replaces it with one with the interval $[3, 3]$ and the local preference function f_{best}. This new STPP is

Inputs: an STPP $P = (V, C)$
Output:
An STP (V, C_P) whose solutions are Pareto optimal for P.
(1) $C_P = C$
(2) while there are weakest link soft constraints in C_P do
(3) Solve (V, C_P)
(4) Delete all weakest link soft constraints from C_P
(5) For each deleted constraint $\langle [a, b], f \rangle$,
(6) add $\langle [a_{opt}, b_{opt}], f_{best} \rangle$ to C_P
(7) Return (V, C_P)

Fig. 3. STPP solver WLO+ returns a solution in the Pareto optimal set of solutions

then solved on the second iteration, whereby the WLO-optimal solution with the CPU duration of 1 is generated. The solution to this STPP is a Pareto-optimal solution to the original problem.

WLO+ was a positive result in the search for tractable methods for finding globally preferred solutions based on less myopic criteria for global preference than WLO-optimality. We now proceed to refine and expand these results in two ways: first by offering a more concise characterization of the class of solution returned by WLO+, and secondly, by showing how restricted classes of STPP with a utilitarian criterion for global preference can be solved tractably.

3 WLO+ and Stratified Egalitarianism

As noted in the previous section, the set of solutions returned by running WLO+ on an STPP is a subset of the set of Pareto Optimal Solutions for that problem. In this section, we present a concise description of this set. By doing so, it is revealed that WLO+ is based on a useful concept of global preference.

We introduce a concept of global preference called Stratified Egalitarianism (SE). Consider again two preference vectors U_S and $U_{S'}$ associated with solutions S and S'. We will say S' SE-dominates S at preference level (or stratum) x if:

- $U^i_S < x$ implies $U^i_{S'} \geq U^i_S$.
- There exists an i such that $U^i_S < x$ and $U^i_{S'} > U^i_S$.
- $U^i_S \geq x$ implies $U^i_{S'} \geq x$.

We say that S' SE-dominates S (without further qualification) if there is any level x such that S' SE-dominates S at x. It is not hard to see that the SE-dominance relation is antisymmetric and transitive[1], thus inducing a partial ordering of solutions. A solution S' will be said to be SE-optimal if it is not SE-dominated. Note that if a solution S' Pareto-dominates S, then S' SE-dominates S at the "highest" level of the $U_{S'}$ vector. Thus, SE-optimality implies Pareto optimality. Furthermore, if S' dominates S in the

[1] The proof makes use of the requirement that the preference values be totally ordered.

WLO ordering, then S' SE-dominates S at the "lowest" level of the $U_{S'}$ vector. Thus, SE-optimality also implies WLO optimality.

Using an economic metaphor to ground intuition, x represents a sort of *poverty line*, and a "policy" S' has a better overall quality than S if some members below the poverty line in S are improved in S', even if some of those above the poverty line in S are made worse off in S' (as long as they do not drop below the poverty line). This metaphor suggests that SE-optimality could be a reasonable criterion for specifying globally preferred solutions.

Economists have considered some notions of egalitarian optimality, but have rejected them as being "non-rational" because they do not imply Pareto optimality. Note that SE-optimality does, however, meet this rationality criterion, but we are unaware of any consideration of the SE preference ordering in the economics literature.

We now prove that the WLO+ algorithm finds exactly the SE-optimal solutions.

Theorem 1. *The set of solutions returned by WLO+ is precisely the set of SE-optimal solutions.*

Proof. Consider a solution S not returned by WLO+, i.e., one that is eliminated at some iteration of the WLO+ algorithm; let the optimal value (i.e., value of the weakest link) of the set of solutions be v at that iteration. Let S' be any survivor at that iteration. There must be some link i such that $U_S^i < v$ (otherwise S wouldn't be eliminated). But $U_{S'}^i \geq v$ since S' survives. Thus, $U_{S'}^i > U_S^i$. Note also that $U_{S'}^j \geq v$ for all links j. Thus, for any value k such that $U_S^k \leq v$, we have $U_{S'}^k \geq U_S^k$. It follows that S is dominated at stratum v.

Conversely, suppose S is dominated at some stratum v but, for the sake of contradiction, suppose S is not excluded from the set of solutions returned by WLO+. From the assumption that S is dominated at stratum v, there exists an S' and i such that $v > U_S^i$ and $U_{S'}^i > U_S^i$, and for any j, $U_{S'}^j \leq v$ implies $U_{S'}^j \geq U_S^j$. During the execution of the WLO+ algorithm, an increasing sequence V of preference values $v_1, v_2, \ldots, v_N = 1$ (where 1 is the "best" preference value) is created, representing the WLO optimal values at each iteration. Clearly, $U_S^i < 1$ (where 1 is the "best" preference value), so one of the Vs must exceed U_S^i. Suppose v_K is the smallest element in V such that $v_K > U_S^i$. Note that S would be removed at this iteration, as a result of its being not WLO optimal, unless the preference function for link i had been reset at an iteration $J < K$. But that function would get reset only if i was a weakest link at J. Then $v_J \leq U_S^i$ since $J < K$, and v_K is the smallest V such that $v_K > U_S^i$. Note however, that for all links j, either $U_{S'}^j \geq v > v_J$ or $U_{S'}^j \geq U_S^i$. Thus, S' would have survived to this iteration if S had. However, $U_{S'}^i > U_S^i \geq v_J$, which contradicts the fact that i is a weakest link. $\square$

3.1 SE versus Leximin

Another global optimality criterion discussed in [8]) is leximin optimality. The leximin ordering compares the minimum value of two preference vectors, then the second lowest and so on, until it finds a difference; the ordering of the first such mismatch determines the leximin ordering of the vectors. It is not difficult to show that SE-dominance implies leximin-dominance, but the converse is not true in general. For example, the preference vector $\langle 5, 1\rangle$ dominates $\langle 1, 3\rangle$ in the leximin ordering but not in the SE ordering. (The SE ordering cares about individuals but leximin does not.)

Nevertheless, it is possible to prove a partial converse as follows.

Theorem 2. *If $\boldsymbol{x} > \boldsymbol{y}$ in leximin-order, then $\boldsymbol{z} > \boldsymbol{y}$ in SE-order, where $\boldsymbol{x}$, $\boldsymbol{y}$, and $\boldsymbol{z}$ are preference vectors, and $\boldsymbol{z} = (\boldsymbol{x} + \boldsymbol{y})/2$ is the average of $\boldsymbol{x}$ and $\boldsymbol{y}$.*

Proof. Note that the coordinate ordering is arbitrary, so we can assume without loss of generality that $\boldsymbol{y}$ is sorted in increasing order. (This simplifies the notation below.)

Suppose $\boldsymbol{x} > \boldsymbol{y}$ in leximin-order. Let $k = \min\{j \mid x_j \neq y_j\}$. Note that for $j \geq k$, we have $x_j \geq y_k$ since otherwise $\boldsymbol{y} > \boldsymbol{x}$ in leximin-order instead of vice versa. Also $y_j \geq y_k$ for $j \geq k$, since $\boldsymbol{y}$ is sorted. Thus, for $j \geq k$, we have $z_j \geq y_k$. It is also easy to see that if x_j and y_j are unequal, which can only happen for $j \geq k$, then one of them must exceed y_k, so $z_j > y_k$ in this case. It follows that either $x_j = y_j$ or $z_j > y_k$. In particular, $z_k > y_k$.

Now let $v = \min\{z_j \mid z_j > y_k\}$. (Note that v is well-defined since $j = k$ satisfies the condition.) We claim that $\boldsymbol{z} >_{\text{SE}} \boldsymbol{y}$ at level v. To see this, note from the previous paragraph that for all j either $x_j = y_j$, in which case $z_j = y_j$, or $z_j > y_k$. In the latter case, $z_j \geq v$ by the definition of v. Since also $z_k > y_k$ and $v > y_k$, this establishes the result.

□

It is well-known that the solutions to an STPP form a convex set, so if S and S' are solutions, then $(S + S')/2$ is also a solution. Furthermore, if the preference functions are convex, then $U^j_{(S+S')/2} \geq (U^j_S + U^j_{S'})/2$ for all j. It follows that if a solution S is leximin-dominated by a solution S', then it is SE-dominated by $(S + S')/2$, which is also a solution. Thus, in the setting of an STPP with convex preference functions, leximin-optimality coincides with SE-optimality. (However, our previous remarks show they are not equivalent in a more general setting.)

4 Utilitarian Optimality

Perhaps the most natural criterion for global optimality is *utilitarian*, where the global value of a solution is the sum of the local values. In this section, we consider applying a utilitarian global optimality criterion to the temporal preference problem. We show that determining the set of all utilitarian optimal solutions as an STP is tractable in the case where all the preference functions are convex and piecewise linear. Piecewise linear preference functions characterize soft constraints in many real scheduling problems; for example, in vehicle routing (where the best solutions are close to desired start times) and in dynamic rescheduling (where the goal is to find solutions that minimally perturb the original schedule).

We first consider the problem of finding a *single* utilitarian optimal solution. (Some constructions related to this have previously appeared in the literature [1, 12]. Our main contribution in this respect will be what follows, where the whole set of solutions is determined as an STP.)

Consider an STPP with preferences $F = \{f_{ij}\}$, and assume that the goal is to find a utilitarian optimal solution S, i.e. where $\sum_{ij} f_{ij}(S)$ is optimal. Suppose each f_{ij} is

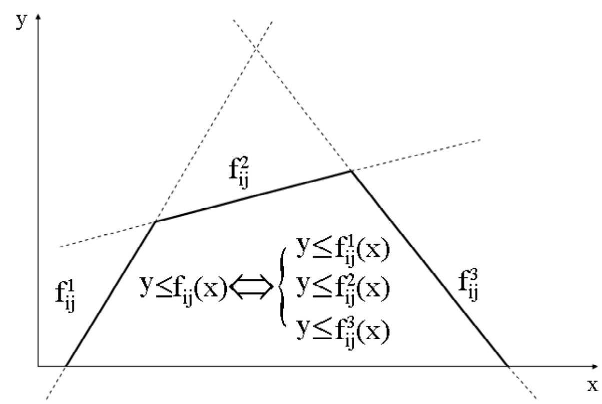

Fig. 4. Convex Piecewise Linear Function

convex and piecewise linear. Thus, there is a sequence of intersecting line segments that make up f_{ij}. We will denote the individual linear functions corresponding to the segments by $f_{ij}^1, f_{ij}^2, \ldots, f_{ij}^{m_{ij}}$, as illustrated in Figure 4.

In this case, we will show that the utilitarian optimization problem can be reduced to a Linear Programming Problem (LPP), which is known to be solvable in polynomial time by Karmarkar's Algorithm [4]. This result generalizes the observation in [9] that STPPs with linear preference functions can be mapped into LPPs.

Since the f's are convex, notice that $y \leq f_{ij}(x)$ if and only if $y \leq f_{ij}^1(x) \wedge y \leq f_{ij}^2(x) \wedge \ldots y \leq f_{ij}^{m_{ij}}(x)$. (See Figure 4.) For the LPP, we introduce an auxiliary variable Z_{ij} for each f_{ij}, together with m_{ij} additional linear constraints of the form

$$Z_{ij} \leq f_{ij}^k(S).$$

We also introduce a set of variables $X = \{X_1, X_2, \ldots X_n\}$ for the nodes in the STP. Note that X_i and X_j, respectively, correspond to the start and end points of the edge associated with f_{ij}. An interval $[p_{ij}, q_{ij}]$ denotes the domain of f_{ij}.

The complete LPP can now be formulated as follows. The indices are assumed to range over their available values, which should be clear from the above discussion. Note that ij in $\{f_{ij}\}$ and $\{Z_{ij}\}$ range over the edges associated with preferences. This could be a small subset of the entire edges in real applications. Finally, we introduce a variable S_{ij} for each temporal distance assignment in a solution.

- Variables: $\{X_i\}$, $\{S_{ij}\}$, and $\{Z_{ij}\}$.
- Constraints (conjunctive over all values of the indices):
  1. $S_{ij} = X_j - X_i$
  2. $p_{ij} \leq S_{ij} \leq q_{ij}$
  3. $Z_{ij} \leq f_{ij}^k(S)$
- Objective Function: $\sum_{ij} Z_{ij}$

Theorem 3. *The solution to the LPP as formulated above provides a utilitarian optimal solution to the STPP.*

Proof. Consider the candidate STPP solution S obtained from the values of the $\{X_i\}$ variables in an optimal solution of the LPP. Clearly, the constraints (items 1 and 2) guarantee that S satisfies the STP underlying the STPP. It only remains to show that it is optimal in the utilitarian ordering for the STPP. From the constraints in item 3, we see that $Z_{ij} \leq f_{ij}^k(S)$ for each linear component k and hence $Z_{ij} \leq f_{ij}(S)$. We claim that $Z_{ij} = f_{ij}(S)$. To see this, note that the Z_{ij} variables can be varied independently without affecting the constraints in items 1 and 2. If $Z_{ij} < f_{ij}(S)$, then the objective function can be increased, without violating any constraints, by increasing Z_{ij} to $f_{ij}(S)$, which contradicts the assumption that the solution is already optimal. Thus, $Z_{ij} = f_{ij}(S)$ for each ij, and so $\sum_{ij} Z_{ij} = \sum_{ij} f_{ij}(S)$

Suppose now there was a better solution S' for the STPP in terms of the utilitarian ordering. Then $\sum_{ij} f_{ij}(S') > \sum_{ij} f_{ij}(S) = \sum_{ij} Z_{ij}$. Observe that we can now formulate a better solution to the LPP based on S' (where we set $Z'_{ij} = f_{ij}(S')$), which is a contradiction. Thus, the S obtained from the LPP is also optimal for the STPP. $\square$

The previous result shows that, for an STPP with preference functions that are convex and piecewise linear, a single solution can be obtained by mapping the problem into an LPP. An interesting question presents itself: is there a compact representation for the *entire* set of the utilitarian optimal solutions to the STPP? In particular, can the set be represented as the solutions to an STP, as can the corresponding set for SE-optimality?

This question is answered in the affirmative by the theorem that follows. As it turns out, whereas solving a primal LPP problem gives a single solution, solving instead a *dual* LPP problem [13] provides the entire set of solutions. Specifically, the dual solution is used to find additional temporal constraints on the STP that underlies the STPP so that its solutions are all and only the optimal solutions to the STPP.

Theorem 4. *Suppose an STPP* P *has preference functions that are convex and piecewise linear. Then the set of all utilitarian optimal solutions can be represented as the solutions to an STP that is formed by adding constraints to the STP underlying* P.

Proof. We map the STPP into an LPP in the same way as before. Note that we can apply certain results from linear programming theory [13, page 101]: the set of solutions to an LPP coincides with one of the faces of the (hyper) polyhedron that defines the feasible region.[2] Note also that the faces can be obtained by changing some of the inequalities to equalities. In a well-known result, the indices of the constraints that change to equalities can be obtained by solving the *dual* of the original LPP.

There are two kinds of inequalities in the LPP that are not already equalities: edge bounds and preference value bounds. In the former case, changing an edge bound to an equality instead of an inequality can be accomplished by adding a simple temporal constraint. In the latter case, an inequality of the form $Z_{ij} \leq f_{ij}^k(S)$ is changed to the equality $Z_{ij} = f_{ij}^k(S)$. This change can be accomplished by restricting the solution to be within the bounds of the f_{ij}^k "piece" of the piecewise linear preference function, which can also be performed through adding a simple temporal constraint. Figure 5 demonstrates this process. A piecewise-linear function with three pieces is displayed. One of the pieces, f, has become part of an equality preference constraint $Z = f(X)$ as

[2] In this context, the term *face* includes vertices, edges, and higher-dimensional bounding surfaces of the polyhedron, as well as the whole polyhedron.

the result of solving the dual LPP. The consequence of this update is to add the temporal bound $[a, b]$ to the STP underlying the original STPP. This bound limits the duration of the edge to be that of the piece of the original preference function that has become an equality. We make the following claims.

1. No temporal value outside the interval $[a, b]$ can be part of an optimal solution; and
2. Every solution of the restricted STP is an optimal solution of the original STPP.

The first claim is obvious from the figure: if there are solutions which contain temporal values outside the bound, they must receive a preference value less than the linear function of the selected piece (by the convexity of the preference function); hence they are not optimal, since this piece is satisfied with equality in all the optimal solutions.

To see that the second claim is true, consider any solution S of the restricted STP. We can extend this to a feasible solution of the (primal) LPP by setting $Z_{ij} = f_{ij}(S)$ for each i and j. Note that $f_{ij}(S) = f_{ij}^k(S)$ for each preference edge that has been restricted to a k piece as discussed above, so $Z_{ij} = f_{ij}^k(S)$ will be satisfied in these cases. Thus, the extended solution is in the optimal face of the LPP, and hence S is optimal for the STPP.

Thus, from the information provided by the dual solution to the LPP, a new STP is formed whose solutions are all and only the utilitarian-optimal solutions of the original STPP.

□

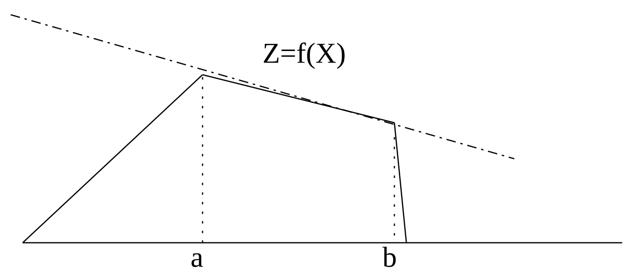

Fig. 5. Squeezing a Temporal Bound by Adding a Preference Equality

This theorem suggests an approach to transforming an STPP into an STP all of whose solutions are utilitarian-optimal for the original problem. First, formulate an LPP as described in the theorem. Second, solve the dual of the LPP to identify the LPP constraints that change from being inequalities to being equalities. Third, map these changes to the STP that underlies the original STPP, again as described above.

As an example, consider the STPP shown in figure 6 (left) with nodes A, B, and C and edges $x1$, $x2$, and $x3$ for which the domains are $[0, 10]$. The preference function f on both $x1$ and $x2$ is given by

$$f(x) = \begin{cases} x \text{ for } 0 \leq x \leq 6 \\ 6 \text{ for } 6 \leq x \leq 10 \end{cases}$$

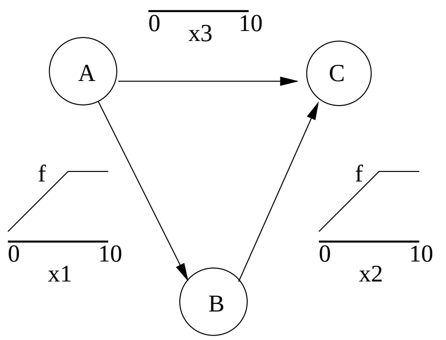

	Primal Variables					
	x1	x2	x3	z1	z2	
Dual Variables	Primal Constraints					
y1	x1					$\leq$ 10
y2		x2				$\leq$ 10
y3			x3			$\leq$ 10
y4	-x1			+z1		$\leq$ 0
y5				z1		$\leq$ 6
y6		-x2			+z2	$\leq$ 0
y7					z2	$\leq$ 6
y8	x1	+x2	-x3			$\leq$ 0
y9	-x1	-x2	+x3			$\leq$ 0

Fig. 6. STPP to LPP example. The primal objective is $\max : (z1 + z2)$ while the dual objective is $\min : (10y1 + 10y2 + 10y3 + 6y5 + 6y7)$.

Note that this is convex and piecewise-linear with 2 pieces. There is no preference (or constant preference) for $x3$.

We can obtain a single optimal solution to this STPP by solving an LPP:

$$\text{maximize } cX : AX \leq b$$

where $X = (x1, x2, x3, z1, z2)$ and $z1$ and $z2$ are the preferences for $x1$ and $x2$ respectively. This is shown in more detail on the right of figure 6.[3] Each optimal solution has a value of 10 for the objective. However, our goal is to find the additional constraints needed to define the STP that characterizes *all* the optimal solutions. For this, we need only find a single solution to the dual problem. The dual solution has $\{y3, y4, y6, y8\}$ as the variables with non-zero values. As a consequence of the Duality Theorem, we conclude that the inequalities corresponding to $y3$, $y4$, $y6$, and $y8$, shown in figure 6, are satisfied as equalities. From ($x3 = 10$) we conclude that $x3$ satisfies [10,10]. From ($-x1 + z1 = 0$) we conclude that $x1$ is restricted to [0,6] (a single piece of the preference function), and similarly for $x2$ using ($-x2 + z2 = 0$). By computing the minimal STP, we can further restrict $x1$ and $x2$ to $[4, 6]$.

5 Experimental Results

Experiments were conducted comparing the performance of WLO+ with a Simplex LP solver. As noted above, WLO+ generates flexible plans in polynomial (worse-case cubic) time that are SE-optimal. The simplex algorithm applied to temporal planning generates fixed utilitarian-optimal plans and is known to perform well in practice but takes exponential time in the worst case. In these experiments, we were interested in comparing both the run-time performance of the two approaches, as well as the quality of WLO+ solutions with respect to an utilitarian evaluation, on a variety of randomly

[3] For expository reasons, this is simplified from the earlier formulation. In particular, we have eliminated the node variables and consolidated the graph constraints into cycle constraints.

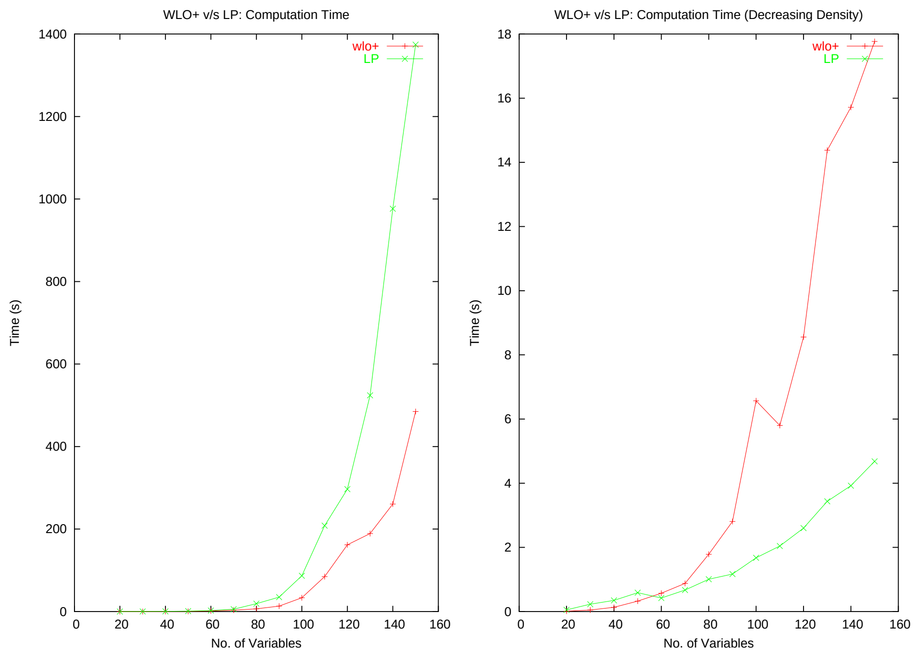

Fig. 7. Simplex vs. WLO+: Time to Solution

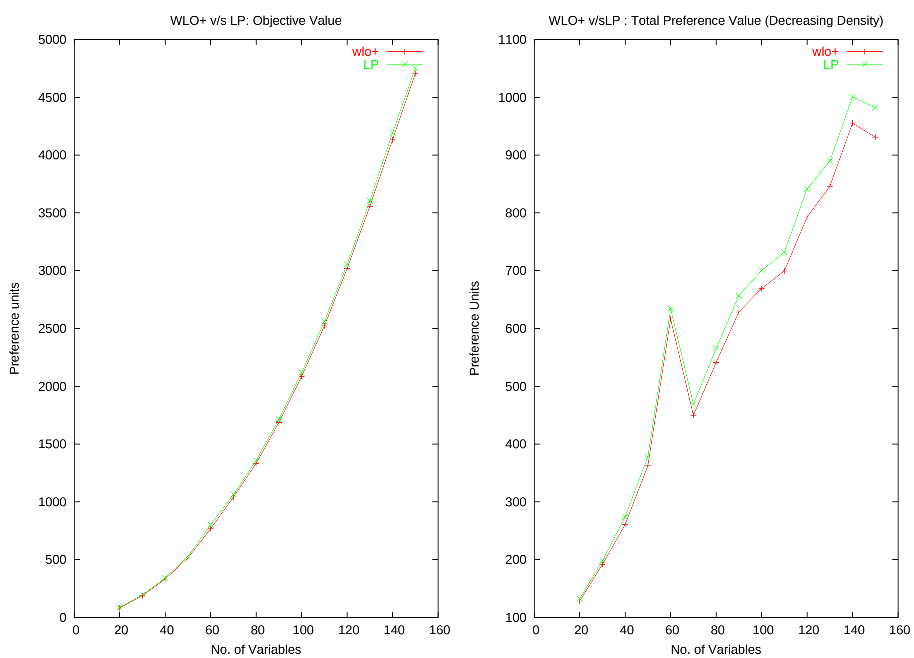

Fig. 8. Simplex vs. WLO+: Solution Quality

generated problems. The results summarized in this section are intended to be preliminary in nature.

A random problem generator was constructed to generate an STPP to be solved by WLO+. A convertor routine is applied to the problem to construct the equivalent LP, in the manner discussed above. The random problem is generated from a seed consisting of a grounded solution. All the programs were compiled optimized on a dual processor 3.06 GHz Linux box. All times reported below include the time to solve the problem but exclude the times to convert the inputs. The LP solver utilized was the *lp_solve* free MIP solver.

The solvers were tested on problems of varying constraint density. (The density here is determined as the ratio of the number of constraints compared to the number that a complete graph would have.) In one set of experiments, the densities were fixed at 10, 50, or 80%. Five problem instances for each density were generated, and the results were averaged to form a single data point. In another set of experiments, the density varied with the problem size (in order to obtain sparse graphs), using the formula $1600/N$, where N is the number of STP nodes. (This keeps the ratio of constraints to nodes constant.) Problem sizes varied between 20 and 150 variables.

The results are shown in Figures 7 and 8, where the graphs on the left of the page show results where densities do not vary with respect to problem size, and those on the right show results when density varies with size. The top two graphs compare solution times for the two solvers, and the bottom two graphs compare solution quality. WLO+ was shown to be faster than the LP solver, on average, and this improvement seemed to increase with problem size. However, LP tended to out-perform WLO+ on sparse problems. This result is somewhat surprising, given the fact that WLO+ uses the Bellman-Ford shortest-path algorithm to solve the underlying STP. Bellman-Ford is designed to perform well in sparse graphs. Further analysis is required to interpret this result.

With respect to solution quality, the SE-optimal solutions generated by WLO+ were, on average, within 90% of the utilitarian-optimal value. These results suggest that WLO+ offers a feasible alternative to LP-based solution techniques, sacrificing tolerable amounts of solution quality for an increase in speed.

6 Discussion and Conclusion

The work reported here contributes to the overall goal of increasing the adeptness of automated systems for planning and scheduling. The objectives of this work overlap with those of a number of diverse research efforts. First, this work offers an alternative approach for reasoning about preferences to approaches based on multi-objective decision theory [2]. Specifically, the characterizations of optimization problems and their properties resemble those found in [7]. The work in this paper also contributes to, and builds upon, the on-going effort to extend CSP algorithms and representations to solve optimization problems or problems where knowledge is uncertain (for example, [6]). Finally, the focus on solving problems involving piecewise-linear constraints has similarities to other efforts more grounded in Operations Research (for example, [1]).

References

1. F. Ajili and H. El Sakkout. A probe-based algorithm for piecewise linear optimization in scheduling. *Annals of Operations Research*, 118:25–48, 2003.
2. F. Bacchus and A. Grove. Graphical models for preference and utility. In *Proceedings of the Eleventh Conference of Uncertainty in Artificial Intelligence*, pages 3–10. Morgan Kaufmann, 1995.
3. J. Bresina, K. Golden, D. Smith, and R. Washington. Increased flexibility and robustness of mars rovers. In *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation for Space, ESTEC*, Noordwijk, Netherlands, 1999.
4. T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT press, 1990.
5. R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*, 49:61–95, 1991.
6. D. Dubois, H. Fargier, and H. Prade. Possibility theory in constraint satisfaction problems: Handling priority, preference and uncertainty. *Applied Intelligence*, 6(4):287–309, 1996.
7. M. Ehrgott. Discrete decision problems, multiple criteria optimization classes and lexicographic max-ordering. In *Trends in Multicriteria Decision Making, volume 465 of Lecture Notes in Economics and Mathematical Systems*, pages 31–44. Springer, 1998.
8. Ulrich Junker. Preference-based search and multi-criteria optimization. In *Proceedings of AAAI-02*. AAAI Press, 2003.
9. L. Khatib, P. Morris, R. Morris, and F. Rossi. Temporal reasoning about preferences. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, Seattle (WA, USA), 2001. Morgan Kaufmann, San Francisco, CA.
10. L. Khatib, P. Morris, R. Morris, and B. Venable. Tractable pareto optimization of temporal preferences. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, Acupulco, Mexico, 2003. Morgan Kaufmann, San Francisco, CA.
11. N. Muscettola, P. Morris, and I. Tsamardinos. Reformulating temporal plans for efficient execution. In *Proceedings of Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, 1998.
12. H. El Sakkout, M. Wallace, and T. Richards. Minimal perturbance in dynamic scheduling. In *Proceedings of the 13th European Conference on Artificial Intelligence (ECAI-98)*, pages 504–508, 1998.
13. A. Schrijver. *Theory of linear and integer programming*. John Wiley and Sons, NY, 1986.
14. Paolo Toth and Daniele Vigo. *The Vehicle Routing Problem*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, PA, 2000.