

Temporal Planning with Preferences and Probabilities

Robert Morris[†], Paul Morris[†], Lina Khatib^{*†}

^{*}Kestrel Technology

[†]Computational Sciences Division

NASA Ames Research Center, MS 269-2

Moffett Field, CA 94035, USA

{morris, pmorris, lina}@email.arc.nasa.gov

Neil Yorke-Smith

Artificial Intelligence Center

SRI International

Menlo Park, CA 94025, USA

nysmith@ai.sri.com

Abstract

In an uncertain world, a rational planning agent must simultaneously reason with uncertainty about expected outcomes of actions and preferences for those outcomes. This work focuses on systematically exploring the interactions between preferences for the durations of events, and uncertainty, expressed as probability distributions about when certain events will occur. We expand previous work by introducing a means for representing events and durations that are not under the control of the planner, as well as quantitative beliefs about when those events are likely to occur. Two reasoning problems are introduced and methods for solving them proposed. First, given a desired overall preference level, compute the likelihood that a plan exists that meets or exceeds the specified degree of preference. Second, given an initial set of beliefs about durations of events, and preferences for times, infer a revised set of preferences that reflect those beliefs.

Introduction

Rational agents are capable of mentally exploring the interactions between what they believe and what they desire as outcomes of actions. More often than not, the value of the outcomes of actions cannot be described by a single attribute, but rather by attributes that combine to determine the overall value of the outcome (Keeney & Raifa 1993). Furthermore, the outcome of actions may not be known with certainty, as a result of the need to interact with the world.

Many practical planning or scheduling problems surround events that are not controlled by the planning agent. For example, Earth Science observation scheduling may involve assigning times for the remote sensing of an area of interest on the Earth before, during, or after a fire has occurred there. The start and end of the fire are not known with certainty at planning time, but Earth Science models might be available to estimate a set of times when fires are likely to occur. In addition, the scientific utility of an observation may vary based on when the observation is taken relative to the fire, resulting in preferences for temporal orderings and durations between planned events and uncontrollable events (Morris *et al.* 2004b). As automated planning matures as a software technology, new techniques inspired by decision theory are being integrated to address the fact that plans are executed in the world, with varying degrees of value to the planner based on their outcomes (Blythe 1999). A princi-

pled approach to scheduling problems as the above is essential for a decision-theoretic temporal planner that takes into account preferences when determining plan quality.

The goal in this paper is to devise systematic methods for exploring the interactions between temporal preferences and uncertainties. We introduce a framework that generalizes the *Simple Temporal Problem* (STP) formulation (Dechter, Meiri, & Pearl 1991), called the *Simple Temporal Problem with Preferences and Probabilities*, or STP³. One component of the generalization adds the capability to express preferences for times, following (Khatib *et al.* 2001). The other component allows for the designation of uncontrollable events and the associated probability space over times. We extend techniques previously used to solve temporal planning problems with preferences to identify solutions that are both globally preferred and highly probable.

Besides defining the STP³ framework, the contribution of this paper is to describe solutions to two practical reasoning problems arising from the interactions between probabilities and preferences. We extend techniques previously used to solve temporal problems with preferences to identify solutions that are both globally preferred and highly probable.

Decision-theoretic planning is surveyed by (Blythe 1999). Most approaches extend classical planning techniques or employ Markov Decision Processes (e.g. (Boutillier, Dean, & Hanks 1999)), in contrast to our constraint-based focus. Of work on temporal reasoning for planning, a characteristic example is (Hanks, Madigan, & Gavrin 1995), who, like us, consider exogenous events, but focus on eliciting probabilities and qualitative preferences from a human expert.

In the constraints literature, preferences are commonly represented using semiring-based formulations, the approach we adopt. An alternative formulation for qualitative preferences is CP-nets (Boutillier *et al.* 2004). Uncertainty has also been represented both qualitatively and quantitatively; probabilistic frameworks include that of (Fargier *et al.* 1995), which we adopt, and its extensions.

Generic constraint-based frameworks that account for both preferences and uncertainty include (Dubois, Fargier, & Prade 1996). Our work is distinguished by restricting attention to Simple Temporal constraints. Prior work in this line has considered STPs with preferences but no uncertainty (Khatib *et al.* 2001); and STPs with uncertainty constraints but no preferences (Morris, Muscettola, & Vidal

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE JUN 2005		2. REPORT TYPE		3. DATES COVERED 00-00-2005 to 00-00-2005	
4. TITLE AND SUBTITLE Temporal Planning with Preferences and Probabilities				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Artificial Intelligence Center,SRI International,333 Ravenswood Avenue,Menlo Park,CA,94025				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

2001; Tsamardinos, Pollack, & Ramakrishnan 2003). While (Rossi, Venable, & Yorke-Smith 2004) incorporate both aspects, that work considers only qualitative uncertainty, that is, with implied uniform distributions.

Example: Earth Science Campaign Observation Scheduling

An *Earth Science campaign* is a systematic set of activities undertaken to meet a particular science objective. Here, we present a hypothetical campaign based on a science objective to test an emissions model predicting the aerosols released by wildfires. Data on several variables must be gathered in order to accomplish the analysis, and several remote sensors, such as those on the Landsat satellite, provide data products at various spatial resolutions relevant to these variables. Preferred times for acquiring Landsat data for vegetation type for a region of interest in the northern hemisphere would be the prior June or July in the same year that the fire burned, when forested land can most easily be spectrally distinguished from grassland. For mapping aerosol concentration, images coincident to burning must be obtained; the Terra and/or Aqua satellites have relevant instruments. For the burned area, data should be acquired after (though not too long after) the fire is out, while for mapping vegetation moisture content, hyperspectral data from an EO-1 Hyperion instrument are relevant, and the most useful data would be that acquired just preceding the fire.

From this description, the inputs to a campaign planning problem potentially consist of the following characteristics:

- a set of temporal, spatial, and resource constraints on when and where images are to be taken;
- user preferences for when an observation should be taken;
- temporal ordering constraints between planned events and uncontrollable, exogenous events such as fires.

A reasonable goal, given these inputs, is to generate a concise representation of the set of solutions (assignments of times and sensing resources) that are maximally preferred and reflect a set of initial beliefs about when exogenous events are likely to occur. We formulate a framework capable of describing the problem and generating this output.

Simple Temporal Problems with Preferences and Probabilities

A *soft temporal constraint* depicts restrictions on the distance between arbitrary pairs of distinct events, and a user-specified preference for a subset of those distances. In Khatib et al. (Khatib *et al.* 2001), 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 I to a set A of admissible preference values.¹ When I is a single interval, a set of soft constraints defines a *Simple Temporal Problem with Preferences* (STPP), a generalization of a Simple Temporal Problem (Dechter, Meiri,

¹For the purposes of this paper, we assume the values in A are totally ordered, and that A contains designated values for minimum and maximum preference.




Figure 1: STP³ Representing the Fire Campaign Scenario

& Pearl 1991). An STPP can be depicted as a pair $\langle V, C \rangle$ where V is a set of variables representing events or other time-points, 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.

Similar to other recent approaches (Morris, Muscettola, & Vidal 2001; Tsamardinos, Pollack, & Ramakrishnan 2003; Rossi, Venable, & Yorke-Smith 2004), we extend the STPP framework to represent temporal uncertainty. First, we partition V into two groups: the *decision variables* V_d and the *parameters* V_u representing uncontrollable events. This partition induces a distinction between *decision constraints* (C_d) and *uncertainty constraints* (C_u): those constraints whose end-point is controllable (i.e. a decision variable), and those whose end-point is uncontrollable (i.e. a parameter). An uncertainty constraint depicts a duration between events as a continuous random variable. To ease the exposition, we assume that the uncertainty constraints are mutually independent²; this allows the constraints in C_u to be expressed in the form $\langle [a_{ij}, b_{ij}], p_{ij} \rangle$, where $p_{ij} : [a_{ij}, b_{ij}] \rightarrow [0, 1]$ is the probability density function over the designated interval. We call the framework $\langle V_d, V_u, C_d, C_u \rangle$, where C_d are soft constraints, a *Simple Temporal Problem with Preferences and Probabilities*, or STP³.

Example 1 *Earth Science Observation Problem.* Inputs: Variables in V_d standing for two controllable events consisting of taking an observation ($Obs1, Obs2$), and two uncontrollable events in V_u , the start and end of a fire (FS, FE) (for simplicity, observations are viewed as instantaneous), as shown in Figure 1. There is also an event TR representing the beginning of time. Soft constraints $f_1(t), f_2(t)$ in C_d are associated with the durations between $Obs1$ and FS , and between $Obs2$ and FE , respectively. For example, $f_1(t)$ may express that there is no value for taking $Obs1$ after the start of the fire (FS), and a preference for times that are as close to FS as possible. Similarly, $f_2(t)$ expresses a preference for $Obs2$ happening before FE as close as possible,

²For instance, imagine that the Earth Science planner maintains a Bayes network elsewhere to express the dependencies; each probability $p(t)$ is given implicitly by that network.

with a penalty if the observation is taken after the fire. Uncertainty constraints p_1, p_2 in C_u are associated with random variables representing the start time and the duration of the fire. These constraints are based on Earth Science models about fires in the area of interest. For example, p_1 may express a normal distribution over the range of times.

A *solution* to an STP³ is a set of assignments to $V = V_d \cup V_u$ that satisfies all the constraints in $C = C_d \cup C_u$. Given an STP³ P , let $Sol(P)$ be the set of all solutions to P . An arbitrary solution $s \in Sol(P)$ can be viewed as having two parts: s_d , the set of values assigned to V_d , and s_u , the set of values assigned to V_u .

Our goal is to develop efficient methods for generating a concise, graphical representation of subsets of $Sol(P)$ corresponding to highly likely, globally preferred solutions. This STP-based graphical representation is called a *flexible (temporal) plan*. Many planning systems use an STP-based representation of the temporal aspects of their plans (Smith, Frank, & Jónsson 2000).

Following previous efforts, methods for flexible temporal planning under uncertainty can be distinguished based on assumptions about the strategy to be applied in executing the flexible plan. A *static execution strategy* assumes no access to the values of s_u during plan execution; by contrast a *dynamic execution strategy* is applied as plan execution proceeds and the values of s_u are observed over time (Morris, Muscettola, & Vidal 2001; Rossi, Venable, & Yorke-Smith 2004). The results of this paper assume a static execution strategy; we defer discussions of planning for dynamic execution of STP³s to future work.

Component Solvers. The solution methods described below are based on different decompositions of an STP³ into component sub-problems for which efficient solution methods exist. As a final preliminary, we fix some terminology and briefly summarize these sub-problem solution methods. Given an STP³, the *underlying STPP* is the problem that results when a constraint $\{[a, b], p_{XY}\} \in C_u$ is replaced by the STP component constraint $[a, b]$. The *underlying Probabilistic STP* is the problem that results when each soft constraint $\{[a, b], f_{XY}\} \in C_d$ is replaced by the STP component constraint $[a, b]$. The *underlying STP* replaces all constraints in $C_d \cup C_u$ with their STP components.

Efficient solution methods for STPs are well-known (Dechter, Meiri, & Pearl 1991). A Simple Temporal Network (STN) is a graph of nodes representing the STP variables and edges labeled with the interval temporal constraints. Each STN is associated with a distance graph derived from the upper and lower bounds of the interval constraints. An STN is consistent iff the distance graph does not contain a negative cycle; this condition can be determined by applying a single-source shortest path algorithm such as Bellman-Ford. In addition to consistency, it is often useful to determine for an STN the *equivalent* STN (in terms of a set of solutions) in which all the intervals are as “tight” as possible. This *minimal network* can be determined by applying an All-Pairs Shortest Path (APSP) algorithm to the input network (Dechter, Meiri, & Pearl 1991).

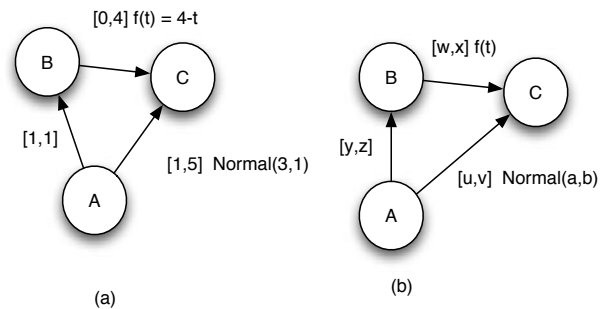


Figure 2: Illustrating the Interactions between Temporal Probabilities and Preferences

Previous efforts in solving STPPs have been based on identifying and applying criteria for “globally preferred solutions” such as “weakest link” (maximize the least preferred local preference), “pareto”, and “utilitarian” (Morris *et al.* 2004a). Developing efficient solvers has required local preference functions that are linear or semi-convex.³ One method for solving STPPs efficiently is called the *chop method*, first introduced in (Khatib *et al.* 2001). The chop method is a two-step search process of iteratively choosing a preference value α , “chopping” every preference function at that point, and then solving an underlying STP defined by the interval of temporal values whose preference values lie above the chop line, i.e. $\{x : f(x) \geq \alpha\}$; henceforth, we refer to this as the *chop interval*. The highest chop point that results in a solvable (i.e. consistent) STP produces a flexible plan whose solutions are exactly the optimal solutions of the original STPP (based on the criteria of weakest link). Binary search can be used to select candidate chop points, making the technique for solving the STPP tractable.

Assessing the Likelihood of Achieving Preferred Plans

This section and the next consider two practical reasoning problems involving the interactions of uncertainty and preferences about time, and demonstrate how under certain assumptions they can be solved efficiently using STP³s. The first problem addresses the question: *what are the chances of achieving a certain level of global preference, given my belief about the way the world will behave?* To illustrate, consider the simple STP³ in Figure 2(a). Here, $V_d = \{A, B\}$ and $V_u = \{C\}$, and there are two decision constraints, between B and C and between A and B . B is tightly constrained to occur exactly one time unit after A . The soft constraint BC prefers durations between B and C to be minimal (higher values more preferred); this is expressed by the preference function $f(t) = 4 - t$. The probability density function for AC is represented by specifying the named function (normal) with mean (3) and standard deviation (1).

³A function is *semi-convex* if drawing a horizontal line anywhere in the Cartesian plane of the graph of the function is such that the set of X such that $f(X)$ is not below the line forms an interval. Semi-convexity ensures that there is a single interval above any chop point, and hence that the resulting problem is an STP.




Figure 3: Why the Upper Bound May Not be Tight

Suppose an agent wants to infer the chances of there being a solution with an overall preference level of 2 or greater. We can answer this question by restricting assignments to BC with an f value of 2 or greater, and propagating the temporal constraints over the network. This means shrinking the BC interval to $[0, 2]$, which in turn shrinks AC to $[1, 3]$. Consequently, the answer to the posed question can be obtained by computing $P(1 \leq t \leq 3) = \int_1^3 p(t)dt$.

This technique can be generalized for arbitrary STP^3 s. Given an $STP^3 P$, to determine the probability of achieving a solution of global preference value γ or higher, we perform the following procedure:

1. Given an input STP^3 , chop each local preference function at the designated preference value γ . Form a new problem by replacing each associated interval with the resulting chop interval.
2. Determine the minimal network of the underlying STP of the new problem, using an APSP algorithm.
3. Compute the overall probability of the underlying probabilistic CSP. Assuming independence of the p_{ij} , the value to be computed is

$$\prod_{p_{ij}} P(a_{ij} \leq t \leq b_{ij}), \quad (1)$$

where for each uncertainty constraint, $[a_{ij}, b_{ij}]$ is the interval of the minimal network derived from step 2.

Provided step 3, which may be done using numerical integration, is of polynomial complexity, the whole method is polynomial. Steps 2 and 3 of this method resemble the method proposed in (Tsamardinou, Pollack, & Ramakrishnan 2003) for solving Probabilistic STPs. Unfortunately, it can be easily shown that the computed value provides only an upper bound on the probability that the solutions defined at that chop level or above will succeed. That this is not a tight upper bound can be demonstrated by a simple example, found in Figure 3. In this example, chopping the preference function at 10 and solving the underlying STP would not shrink the temporal bounds of the uncertainty links. Therefore, the probability of succeeding returned by this method would be 1, although in fact some of the probability mass is lost as a result of the chop.

Despite these limitations, an upper bound computation may be useful; if the bound is too low, the planner will be forced to “lower expectations” of the plan branch under consideration, i.e. its overall expected preference level.

A tighter bound would require examining the mass of the polytope defined by all the constraints (a similar observation was made in (Tsamardinou, Pollack, & Ramakrishnan 2003)). Applied to the previous example, we get $P((0 \leq AB \leq 10) \wedge (0 \leq BC \leq 10))$ from (1), but the true probability is $P((0 \leq AB \leq 10) \wedge (0 \leq BC \leq 10) \wedge (AB + BC \geq 10))$ or simply $P(AB + BC \geq 10)$, assuming the bounds. (Note that the AB and BC random variables are no longer independent under the condition $AB + BC \geq 10$.) We can reformulate this as $P(\bigvee_x (AB = x \wedge BC \geq 10 - x))$ and calculate it as

$$\int_0^{10} \left(\int_{10-x}^{10} p(y)dy \right) p(x)dx.$$

Inducing Preferences from Probabilities

In this section we consider a sort of dual problem to that posed in the previous section: *given current expectations about the world, how can preferences be systematically adjusted to fit with those expectations?* Intuitively, by answering this question, the planner can “factor out” the temporal uncertainty in the problem, resulting in a pure decision problem: because of the factoring, the solutions most preferred based on the induced preferences are also most likely.

This factoring process takes into account the initial preferences on decision constraints, combining them with the preferences induced from the uncertainty constraints. The core idea is to apply the concept of expected utility from decision analysis (Keeney & Raiffa 1993) to represent induced local preferences. Once the reasoning is complete, the “output” preferences on the decision constraints thus reflect both the preferences of the agent and its expectation about the uncertainty in the world.

The main result of this section will be to state a set of sufficient conditions for finding an efficient algorithm for this process for certain classes of STP^3 . The method consists of:

1. Given an input STP^3 , derive the minimal network of the underlying STP.
2. Apply a local consistency algorithm (discussed below) to the resulting STP^3 (i.e. with the tightened interval constraints) to compute the induced preferences.
3. Solve the underlying STPP of the resulting network using the chop solver to find the globally preferred solutions.

The set of solutions making up the flexible plan that results are the *expected globally preferred solutions*.

To examine the second step in more detail, we mimic the method of *triangular reduction* found in (Morris, Muscettola, & Vidal 2001), used to solve Simple Temporal Problems with Uncertainty (STPUs). We consider all STP^3 s as collections of *triangular subnetworks* of the form illustrated by Figure 2(b), where there is a single uncertainty constraint on AC with bounds $[u, v]$, and two decision constraints on AB and BC with bounds $[y, z]$ and $[w, x]$ respectively. As in the Earth Science example, A might be the beginning of

time, B might be the start of a planned observation, and C the onset of a fire. The goal is to compute the regression of p_{AC} over f_{BC} to find the induced soft decision constraint f_{AB} . (The case in which AB is also associated with a soft constraint can be handled as part of the general solution method discussed later.)

To handle the single triangle case, we need to consider three possible orderings between B and C . We assume that step 1 of the approach has been applied, so that the triangular network has been minimized. If B precedes C ($w \geq 0$), then the induced soft constraint is $\{[y, z], f_{AB}\}$, where

$$f_{AB}(t) = \int_u^v f(t' - t)p(t')dt'.$$

Although in general this function cannot be derived analytically, with certain restrictions placed on the shape of the preference function it may be possible to compute it directly. Alternatively, we can estimate it numerically (e.g. Monte Carlo integration), or even perform crude but fast estimation based on the expected value. If C precedes B ($x \leq 0$), then intuitively the planner does not require any knowledge about the expected time of C in order to deduce the preferred time to execute B dynamically (the soft constraint on AB in this case can be derived from that of BC). However, recall that we focus only on the situation of static execution, in which knowledge about C is not available at planning time. This means that the predictive models of the Precede case are relevant to planning the Follow case: the same technique can be followed. Finally, for static execution the same also applies if B and C are unordered ($w < 0, x > 0$).

To derive the induced constraints for general STP³ networks, we consider all triangles separately, propagating the effects of one operation to neighboring triangles, until the network is quiescent. Thus, the structure of the algorithm is similar to determining path-consistency in an STP network. Propagation requires combining local preference functions. The same combination operator as that used for determining local consistency for preference networks (Rossi *et al.* 2002) can be applied here for propagating soft constraints. After the network has reached quiescence, the planner can safely discard the probability density functions p_{XY} in C_u . Removing them results in the underlying STPP, which can be solved by the chop method (Khatib *et al.* 2001).

The following result summarizes these core ideas. It will be proved informally and illustrated by an example. Following terminology in (Morris, Muscettola, & Vidal 2001), an STP³ will be said to be *pseudo-controllable* if no interval in an uncertainty constraint is “squeezed” as the result of performing step 1 above (computing the minimal network). We refer to the STP³ that results from performing step 2 above as the *induced STP³*.

Theorem 1 Given an STP³ with the following properties:

1. The input preference functions p_{ij} are linear or semi-convex piecewise linear (intuitively, semi-convex piecewise linear means that there are no “V” shaped segments);
2. The STP³ is pseudo-controllable;
3. The probability distributions on the uncertainty constraints are normal;




Figure 4: Example of Induced Preferences

then, using the method described above, the set of expected globally preferred solutions to the initial STP³ can be computed in polynomial time.

The first condition of the theorem is needed to ensure that the induced STP³ has only functions that are semi-convex, which is required for the application of the chop solver method in step 3 (a polynomial-time procedure). Steps 2 and 3 are required to simplify the induced functions to linear functions involving expected values (see the example below). The conclusion of the proof consists of observing that the underlying procedures applied in the method (all-pairs shortest path, the local-consistency technique for deriving induced preferences, the chop solver, and numerical integration for determining the expected values) are all polynomial.

To illustrate step 2 of the method in the general case, consider the STP³ in Figure 4. This problem consists of two decision constraints on BC and BD with associated preference functions f, g defined, f clearly preferring larger durations between B and C , and g preferring smaller durations. Two uncertainty constraints on AC and AD consist of normal probability density functions p_1 and p_2 with means and standard deviations indicated in parentheses. The goal is to infer the induced preference function h on AB (the network is already minimal).

First, considering the triangle ABC , one induced function for h arises as follows:

$$\begin{aligned} h_1(t) &= \int_0^{10} f(t' - t)p_1(t')dt' \\ &= \int_0^{10} [t' - t]p_1(t')dt' \\ &= \int_0^{10} t'p_1(t')dt' - t \int_0^{10} p_1(t')dt'. \end{aligned}$$

Notice that because of the pseudo-controllability of the network (it being already minimal), the last equation reduces to $E(T_1) - t$, since then $\int_0^{10} p_1(t')dt' = 1$ and $\int_0^{10} t'p_1(t')dt' = E(T_1)$, where $E(T_1)$ is the expected value of the random variable T_1 associated with the duration. A similar derivation based on the triangle ABD then results in another induced function $h_2(t) = 10 - [E(T_2) - t]$. The

final induced function h becomes the combination of h_1 and h_2 : e.g. the intersection of the areas under the functions.

This approach can be generalized for regression over semi-convex piecewise linear preference functions. Let f_{BC} be the intersection of n linear segments $f_{BC}^1, \dots, f_{BC}^n$, where for each k , $[a_{BC}^k, b_{BC}^k]$ is the segment for which $f_{BC} = f_{BC}^k$. When regressing p_{AC} over f_{BC} to compute the induced preference function h_{AB} , we have:

$$\begin{aligned} h_{AB}(t) &= \int_{a^1}^{b^n} f_{BC}(t' - t)p_{AC}(t')dt' \\ &= \sum_{k=1, \dots, n} \int_{a^k}^{b^k} f_{BC}^k(t' - t)p_{AC}(t')dt', \end{aligned}$$

which simplifies to sums involving linear functions.

This example shows how with suitable restrictions on the shapes of the preference functions and on whether the all-pairs computation eliminates any of the probability mass, the computation of induced preferences can be made efficient.

Discussion and Future Work

We have examined temporal reasoning under the interactions of preferences and quantitative uncertainty in the context of constraint-based planning. In addition to the formulation of the STP³ framework, which augments the Simple Temporal Problem with both preferences and probabilities, the main contribution of this paper is to formulate two planning decision problems. Utilizing standard methods from decision theory, probability theory, and recent advances in constraint satisfaction, we have shown how flexible temporal plans can be generated that are most preferred based on what the planning agent believes about the expected times of events; and how the agent can update its preferences, given its beliefs.

Fundamentally, preferences and uncertainty are orthogonal aspects of the decision problem. Both planning decisions we have considered are approaches to combining the two aspects; which is most relevant depends on the aim of the planning agent and the questions being asked of it. The first decision, to evaluate the probability of a plan existing with at least a given preference, is useful to determine whether a plan branch can meet a minimum quality threshold. The second decision, to update preferences based on beliefs, is useful to factor the uncertainty into a single criterion for plan evaluation. Besides these two decision problems, the proposed framework can be applied to related problems; for instance, an agent might seek to determine the maximal preference level at which a solution exists with a given probability p . When $p = 1$ and the probabilities are uniform, this corresponds to certain forms of strong controllability addressed in (Rossi, Venable, & Yorke-Smith 2004).

Future theoretical efforts include characterizing more fully the computational complexity of STP³s, and refining the bound on the probability that a plan exists with given preference quality. In addition to implementing the methods described in this paper, our major next step is to extend the results here to address issues in planning under a dynamic

execution strategy. Of particular importance will be to examine the interactions between preferences and *wait constraints* that emerge when determining the controllability of flexible plans, as described in (Morris & Muscettola 2000).

Acknowledgments. We thank the reviewers for their comments. The work of the last author on this material is supported by the Defense Advanced Research Projects Agency, through the Department of the Interior, NBC, Acquisition Services Division, under Contract No. NBCHD030010.

References

- Blythe, J. 1999. Decision-theoretic planning. *AI Magazine* 20(2):37–54.
- Boutilier, C.; Brafman, R. I.; Domshlak, C.; Hoos, H. H.; and Poole, D. 2004. Preference-based constrained optimization with CP-Nets. *Computational Intelligence* 20:137–157.
- Boutilier, C.; Dean, T.; and Hanks, S. 1999. Decision-theoretic planning: Structural assumptions and computational leverage. *JAIR* 11:1–94.
- Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. *Artificial Intelligence* 49:61–95.
- Dubois, D.; Fargier, H.; and Prade, H. 1996. Possibility theory in constraint satisfaction problems: Handling priority, preference and uncertainty. *Applied Intelligence* 6(4):287–309.
- Fargier, H.; Lang, J.; Martin-Clouaire, R.; and Schiex, T. 1995. A constraint satisfaction framework for decision under uncertainty. In *Proc. of UAI 1995*, 167–174.
- Hanks, S.; Madigan, D.; and Gavrin, J. 1995. Probabilistic temporal reasoning with endogenous change. In *Proc. of UAI 1995*, 245–254.
- Keeney, R., and Raifa, H. 1993. *Decisions with Multiple Objectives*. Cambridge University Press.
- Khatib, L.; Morris, P.; Morris, R.; and Rossi, F. 2001. Temporal reasoning about preferences. In *Proc. of IJCAI-01*, 322–327.
- Morris, P., and Muscettola, N. 2000. Execution of temporal plans with uncertainty. In *Proc. of AAAI-2000*, 491–496.
- Morris, P.; Morris, R.; Khatib, L.; Ramakrishnan, S.; and Bachmann, A. 2004a. Strategies for global optimization of temporal preferences. In *Proc. of CP'04*, 408–422.
- Morris, R.; Dungan, J.; Khatib, L.; Bachman, A.; Gasch, J.; Hempel, P.; Williams, J.; Wood, T.; and Bruno, C. 2004b. Coordinated science campaign planning for earth observing missions. In *Proc. of ESTC'04*.
- Morris, P.; Muscettola, N.; and Vidal, T. 2001. Dynamic control of plans with temporal uncertainty. In *Proc. of IJCAI-01*, 494–502.
- Rossi, F.; Venable, K.; Khatib, L.; Morris, P.; and Morris, R. 2002. Two solvers for tractable temporal constraints with preferences. In *Proc. of Workshop on Preferences in AI and CP*.
- Rossi, F.; Venable, K.; and Yorke-Smith, N. 2004. Controllability of soft temporal constraint problems. In *Proc. of CP'04*, 588–603.
- Smith, D.; Frank, J.; and Jónsson, A. 2000. Bridging the gap between planning and scheduling. *Knowledge Eng. Review* 15(1).
- Tsamardinos, I.; Pollack, M. E.; and Ramakrishnan, S. 2003. Assessing the probability of legal execution of plans with temporal uncertainty. In *Proc. of ICAPS'03 Workshop on Planning Under Uncertainty and Incomplete Information*.