Optimizing Temporal Preferences

Lina Khatib[1,2] **Paul Morris**[2] **Robert Morris**[2]
1. Kestrel Technology
2. Computational Sciences Division
NASA Ames Research Center, MS 269-2
Moffett Field, CA 94035

Abstract

This paper focuses on problems where the objective is to maximize a set of local preferences for times at which events occur. Previous work by the authors and others has resulted in a formalization of a subclass of these problems into a generalization of Temporal Constraint Satisfaction Problems, using a semi-ring as the underlying structure for reasoning about preferences. A tractable strategy for combining and comparing preferences was proposed, wherein computing global preferences consists of taking the minimum of the component preference values. This strategy, which we here dub the "weakest link optimization" (WLO) strategy, is a coarse method for comparing solutions, and can easily be demonstrated to have drawbacks. To compensate for these limitations, we make WLO more robust by combining it with a process that involves iteratively committing to temporal values in the set of optimal solutions, and concomitantly fixing the preference value for the projection of the solution to the local preference. We prove the value of this strategy by showing that the resulting solutions are also Pareto optimal.

Introduction

The notion of *softness* has been applied in the planning or scheduling literature to describe either a constraint or planning goal, indicating that either can be satisfied to matters of degree. Sometimes softness is naturally associated with the constraints used in defining the problem. 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 that the instrument can be on over a given length of time, 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 must be off for a certain duration. One way to express the soft constraint is to say that the duration of time the instrument is on should be as close as possible to that specified in the constraint. Reasoning about soft constraints for planning or scheduling is for the purpose of finding a solution that satisfies the constraint 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 well-known television game show. WLO was chosen for reasons of computational efficiency. Specifically, it 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). 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 an approach to making WLO more robust by combining it with an iterative strategy for solving soft constraint reasoning problems. The process involves repeatedly restricting temporal values for the weakest links, resetting their preference values, and again applying the WLO procedure to the reduced problem that results from these changes. The intuition is simple: by re-starting WLO iteratively on the reduced problem, WLO might be able to improve the preference values of the remaining temporal variables, thus compensating for the myopia of WLO. We motivate this technique with an example from a Mars Rover planning domain.

The remainder of this paper is organized as follows. In Section 2, we summarize the soft constraint framework introduced previously; this review serves to motivate the current work. We then illustrate 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 section 3, which formalizes this strategy and proves that any solutions generated by an application of this strategy is in the set of Pareto optimal solutions for the original problem.

Reasoning about preferences with soft constraints

In (Khatib *et al.* 2001), a *soft temporal constraint* is defined as a pair $\langle I, f \rangle$, where I is a set of intervals $\{[a, b], a \le b\}$ and f is a function from $\bigcup I$ to a set A of preference values. To compare and combine values from this set, A is organized 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.

Soft temporal constraints give rise to a class of constrained optimization problems called Temporal Constraint Satisfaction Problems with Preferences (TCSPPs). A TCSPP can simply be viewed as a generalization of a classical TCSP with soft constraints. In TCSPs (Dechter *et al.* 1991), a temporal constraint depicts restrictions either on the start times of events (in which case the constraints are unary), or on the distance between pairs of distinct events (in which case they are binary). For example, a unary constraint over a variable X representing an event, restricts the domain of X, representing its possible times of occurrence; then the interval constraint is shorthand for $(a_1 \le X \le b_1) \vee \ldots \vee (a_n \le X \le b_n)$. A binary constraint over X and Y, restricts the values of the distance $Y - X$, in which case the constraint can be expressed as $(a_1 \le Y - X \le b_1) \vee \ldots \vee (a_n \le Y - X \le b_n)$. 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 TCSP in which each constraint is defined by a single (convex) interval is called a Simple Temporal Problem (STP).

As with classical TCSPs, the interval component of a soft temporal constraint depicts restrictions either on the start times of events, or on the distance between pairs of distinct events. The class of TCSPPs in which each constraint consists of a single interval is called *Simple Temporal Problems with Preferences* (STPPs). A *solution* to a TCSPP 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 $\le_S$ over A to compare preference values of arbitrary assignments; $a \le_S b$ can be read *b is more preferred than a*. Classical Temporal CSPs can be seen as a special case of TCSPP, with "soft" constraints that assign the "best" (1) preference value to each element in the domain, and the "worst" (0) value to everything else.

The optimal solutions of a TCSPP are those solutions which have the best preference value in terms of the ordering $\le_S$. *Weakest Link Optimization* (WLO) is formalized via the semiring $S_{WLO} = \langle A, max, min, \mathbf{0}, \mathbf{1} \rangle$. Thus, where $a, b \in A$, $a + b = max(a, b)$ and $a \times b = min(a, b)$, and $\mathbf{1}$ ($\mathbf{0}$) is the best (worst) preference value. Given a solution t in a TCSPP 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" in a WLO solution t is any minimum f 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 ($\le_S$ 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 of functions 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). STTPs 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 (same) 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

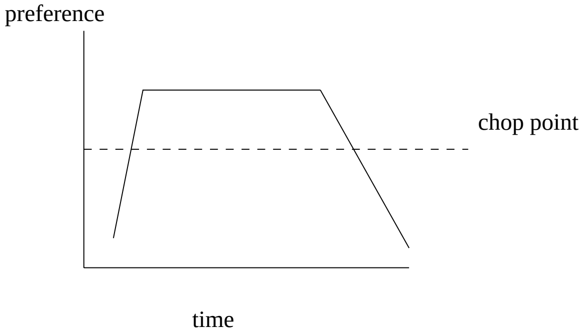

Figure 1: "Chopping" a semi-convex function

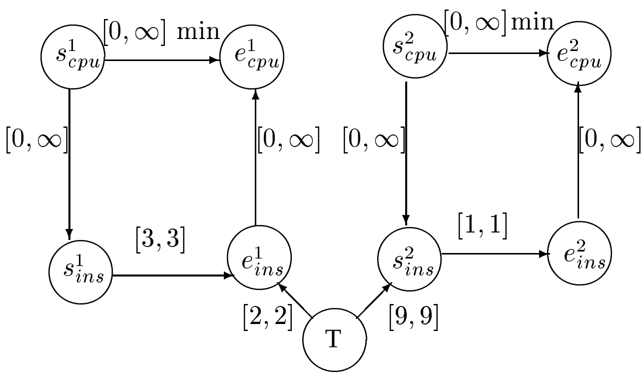

Figure 2: The STPP for the Rover Science Planning Problem where T is any timepoint

search can be used to select candidate chop points, making the technique for solving the STTP tractable (see [Khatib, *et al.*, 2001] for more details on all this). The second step, solving the induced STP, can be performed by transforming the graph associated with this STP into a distance graph, then solving two single-source shortest path problems on the distance graph. 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. The complexity of this phase is $O(en)$ (using the Bellman-Ford algorithm (Cormen *et al.* 1990)).

The problem with WLO

Formalized in this way, WLO offers 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 s^1_{ins}, e^1_{ins} and s^2_{ins}, e^2_{ins} 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 on the CPU duration constraints is the negated identity function, $f(t) = -t$; 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 solutions 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, $f_1, \ldots, f_m$ as criteria requiring simultaneous optimization, and let $s = [t_1, \ldots, t_n]$ and $s' = [t'_1, \ldots t'_m]$ 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.[1]

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 smaller values. (This strategy would also ensure Pareto optimality, since every minimum sum solution is Pareto optimal.) This policy might be dubbed "utilitarian." The main drawback to this alternative is a loss of tractability. 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 for ensuring that a local consistency approach is applicable to the soft constraint reasoning problem.

Of course, it would still be 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.

[1]This phenomenon is often referred to in the literature as the "drowning effect."

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 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. (This is reminiscent of the Weakest Link game show, where each round eliminates a weakest link.)

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 is one in which the optimal value v for the preference function associated with the constraint is the same as the optimal value for the global solution using WLO, and furthermore, v is not the “best” preference value (i.e., $v < \mathbf{1}$, where $\mathbf{1}$ 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. $\mathbf{1}$) 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 STTP, P. Where C is a set of soft constraints, the STTP (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 $\langle [a_{opt}, b_{opt}], f_{best} \rangle$, 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.

To illustrate the main theoretical result of this paper, Figure 4 shows the relationship between the WLO-optimal solutions and the Pareto-optimal solutions.

We now proceed to prove the main result, which is that the subset of solutions of the input STTP returned by WLO+ is contained in the intersection of WLO-optimal and Pareto-optimal solutions. Formally, given an

Inputs: an STPP $P = (V, C)$
Output:
A 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)

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

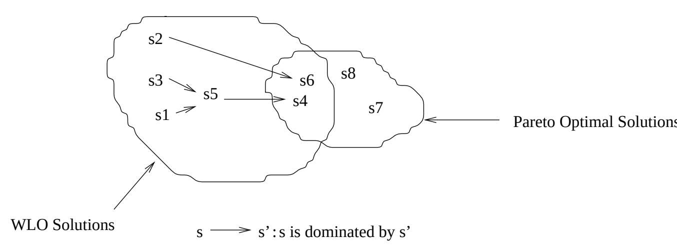

Figure 4: Relationships between Solution Spaces for STTPs that are WLO or Pareto Optimal

STTP P, let $Sol_{WLO}(P), Sol_{PAR}(P)$ 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 $Sol_{WLO+}(P) \subseteq Sol_{WLO}(P) \cap Sol_{PAR}(P)$. *Moreover, if P has any solution, then $Sol_{WLO+}(P)$ is nonempty.*

Proof:

First note that after a weakest link is processed in steps (4) to (6), it will never again be a candidate for a weakest link (since its preference is set to f_{best}). Thus, the algorithm will terminate when all the soft constraints in C_P have preferences that have f_{best} value over their entire domain.

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 initial preference value for s on c cannot be f_{best}. It follows that at some point during the course of the algorithm, c will become a weakest link. Since s survives until then and s' dominates s, it follows that s' also survives. At that time, s will be excluded by step (6) since it is not WLO optimal, contradicting the assumption that $s \in Sol_{WLO+}(P)$. 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 so-

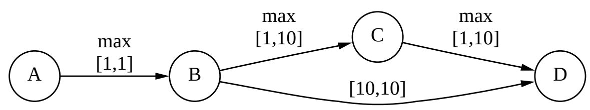

Figure 5: Single WLO+ Solution.

lutions 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. $\square$

Corollary 1.1 *If P has any solution, the set $Sol_{WLO}(P) \cap Sol_{PAR}(P)$ is nonempty.*

Although the algorithm determines a nonempty set of solutions that are both WLO optimal and Pareto optimal, the set may 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. 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.)

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. The only other restriction still required is that of the semi-convexity of the preference functions. This restriction is needed because the "chopping" method assumes that the domain above the chop point defines a STP, which implies that the preference function is semi-convex. A possible extension to this work would be to examine ways to relax this restriction in order to solve more general constrained optimization problems.

Another interesting observation is that in each iteration, the restriction to the WLO optimal solutions means that the weakest links are restricted to intervals where the original preference function has a constant (optimal) value. It follows that the surviving solutions in $SOL_{WLO+}(P)$ all have the same (original) preference value for each soft constraint. Let the *signature* of a solution be the vector of values of its projection on each constraint. Then the observation may be restated as saying that the set of solutions in the STP returned by the WLO+ strategy all have the same signature.

A consequence of this observation is that the WLO+ strategy produces solutions that are what will be termed *utilitarian optimal with respect to the set of WLO+ solutions*.[2] In the current context, a utilitarian strategy is one that seeks to maximize the sum of the local preferences, and a *utilitarian optimal* solution to a problem is one that is optimal with respect to this strategy. To be utilitarian optimal with respect to the set of WLO+ solutions means that, among only the WLO+ solutions, the given solution is utilitarian optimal. However, the WLO+ solutions are not necessarily utilitarian optimal with respect to all solutions or even the WLO solutions. For example, if in figure 5, the preference function is $f(t) = t^2$, a utilitarian optimal WLO solution would be given by BC $= 1$ and CD $= 9$, which is better than the WLO+ solution in the utilitarian sense.

Thus, there are utilitarian optimal solutions that are missed by WLO+. These could include solutions that are not even WLO-optimal but do better overall in meeting local preferences. Nonetheless, the results of this paper show that under certain restrictions it is possible to construct efficient soft constraint-based systems to perform meaningful optimization on preferences.

[2] A similar statement can be made regarding *lexicographic* optimality (Bistarelli *et al.* 1999).

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.

Acknowledgments

The authors wish to thank Francesca Rossi and Jeremy Frank for helpful suggestions and comments. We are also grateful for referee comments that helped identify an error in the original paper.

References

S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based Constraint Solving and Optimization. *Journal of the ACM*, 44(2):201–236, March 1997.

S. Bistarelli, H. Fargier, U. Montanari, F. Rossi, T. Schiex, and G. Verfaillie. Semiring-based CSPs and valued CSPs: Frameworks, properties and comparison. *Constraints*, 4:199–240, 1999.

T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT press, Cambridge, MA, 1990.

R. Dechter, I. Meiri, and J. Pearl. Temporal constraint networks. *Artificial Intelligence*, 49, 1991:61–95, 1991.

L. Khatib, P. Morris, R. Morris and F. Rossi. Temporal Reasoning about Preferences. In *Proceedings of IJCAI-01*. Morgan Kaufman, 1995.