On the Linear Programming Duals of Temporal Reasoning Problems

T. K. Satish Kumar
Information Sciences Institute
University of Southern California
tkskwork@gmail.com

Zhi Wang
Department of Computer Science
University of Southern California
zhiwang@usc.edu

Anoop Kumar   Craig Milo Rogers   Craig A. Knoblock
Information Sciences Institute
University of Southern California
{anoopk, rogers, knoblock}@isi.edu

Abstract

Temporal reasoning problems occur in many application domains of Artificial Intelligence; therefore, it is important for us to develop algorithms for solving them efficiently. While some problems like Simple Temporal Problems are known to be tractable, some other problems like Disjunctive Temporal Problems are known to be NP-hard. In this paper, we provide a Linear Programming (LP) *duality* perspective on temporal reasoning problems. In many cases, we show that their LP duals are the commonly-studied *flow* problems in Graph Theory. Using the general theory of LP duality, we develop novel algorithms for efficiently solving several temporal reasoning problems. We also show that other previously-known efficient methods in temporal reasoning also fit into this perspective of LP duality.

Introduction

Temporal reasoning problems occur in many application domains of Artificial Intelligence. For example, in medicine, temporal reasoning is used to aid doctors monitor the development of a disease and plan treatment accordingly (Augusto 2005); in job scheduling, it is used to schedule jobs on machines for makespan minimization (Ji, He, and Cheng 2007); in autonomous space exploration, it is used to schedule spacecraft operations (Knight et al. 2001); in transportation domains, it is used for throughput maximization (Feijer, Savla, and Frazzoli 2012); and in smart home and smart grid domains, temporal reasoning is used in appliance scheduling for energy cost minimization (Mohsenian-Rad and Leon-Garcia 2010). Given the extensive applications of temporal reasoning, it is imperative for us to develop efficient algorithms for solving temporal reasoning problems. In most cases, this necessitates the design of algorithms that have polynomial time complexities (preferably of low order).

Some problems in temporal reasoning, for example, Simple Temporal Problems (STPs) (Dechter, Meiri, and Pearl 1991), are known to be tractable, i.e, solvable in polynomial time. STPs are widely used for reasoning about difference constraints between the execution times of various events. An STP S is defined by a directed graph $\langle \mathcal{X}, \mathcal{E} \rangle$, where $\mathcal{X} = \{X_0, X_1, ...X_N\}$ is the set of nodes representing events, interchangeably used for their execution times, and $\mathcal{E}$ is the set of directed edges between them representing *simple temporal constraints*. X_0 is set to 0 to establish a frame of reference; and each directed edge $e_{ij} = \langle X_i, X_j \rangle \in \mathcal{E}$ is annotated with a pair of real numbers $[LB(e_{ij}), UB(e_{ij})]$, representing the simple temporal constraint $LB(e_{ij}) \leq X_j - X_i \leq UB(e_{ij})$. STPs can be solved in polynomial time using shortest path computations on their *distance graph* representations. In the distance graph representation, the constraint $X_j - X_i \leq \rho$ is represented as an edge from X_i to X_j annotated with ρ. The absence of negative cost cycles in the distance graph characterizes the consistency of the temporal constraints in STPs (Dechter, Meiri, and Pearl 1991). Since the distance graph can have negative cost edges, shortest paths are computed using the Bellman-Ford algorithm (Kleinberg and Tardos 2006). Improved algorithms for solving STPs have been developed by several authors (Xu and Choueiry 2003; Planken, de Weerdt, and van der Krogt 2008).

Some other problems in temporal reasoning, for example, Disjunctive Temporal Problems (DTPs) (Stergiou and Koubarakis 1998; Oddi and Cesta 2000), are known to be NP-hard. DTPs are more expressive than STPs since they allow disjunctions of the form $\bigvee_t \ (L_t \leq X_{j_t} - X_{i_t} \leq U_t)$ where $\forall\ t :\ X_{i_t}, X_{j_t} \in \mathcal{X}$. In planning and scheduling, DTPs arise in dispatchable execution (Shah and Williams 2008), threat resolution in partial order planning (Nguyen and Kambhampati 2001), and plan merging (de Weerdt 2003). Although restricted classes of DTPs can be solved in polynomial time (Kumar 2005; 2006), they generally require an exponential search space.

Several other temporal reasoning problems with or without resources, preferences, and/or controllability have been characterized for their complexities (Yorke-Smith, Venable, and Rossi 2003). For example, constructing resource envelopes in producer-consumer models can be done in polynomial time (Kumar 2003; Muscettola 2004), but job shop scheduling problems that reason about contention for resources are NP-hard (Smith and Cheng 1993). Similarly, temporal reasoning problems with convex preference functions (Morris et al. 2004) or piecewise constant preferences on individual variables (Kumar 2004) are tractable, whereas they are NP-hard for general functions.

In this paper, we provide a Linear Programming (LP) *duality* perspective on temporal reasoning problems. In many cases, we show that their LP duals are the commonly-studied

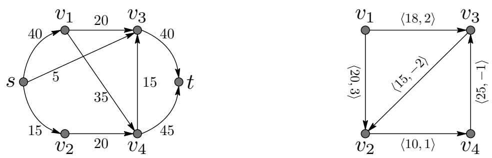

(a) (b)

Figure 1: (a) shows an instance of the maxflow problem where each edge is annotated with a nonnegative capacity; and (b) shows an instance of the MCCP where each edge is annotated with a nonnegative capacity and a cost of pushing a unit flow through it.

flow problems in Graph Theory. Their relationship to flow problems opens the possibility for solving them efficiently using state-of-the-art flow-based techniques. Moreover, the LP duality perspective can be used to unify the previous approaches for tractable temporal reasoning problems.

Background

In this section, we briefly describe background material that is relevant to this paper. We begin with the description of an LP problem and its dual (Chvátal 1983; Vazirani 2001, Chap. 12). An LP problem is the problem of maximizing (or minimizing) a linear objective function over a certain set of variables that are also constrained by a finite number of linear equations or linear inequalities. A *feasible* solution is one that satisfies all constraints; and an *optimal* solution is one that is feasible and has the maximum (minimum) value with respect to the objective function. Every LP problem is associated with a *dual*, referred to as its LP dual. The original LP problem is referred to as the *primal*; and it is well known that the dual of the dual is the primal. The variables in the dual are in one-to-one correspondence with the constraints in the primal, and vice versa.

If the primal is a maximization (minimization) problem, the dual is a minimization (maximization) problem. Several important theorems guide the relationship between a given primal and its dual. These include the *weak duality theorem*, the *strong duality theorem*, and the *complementary slackness theorem*. The weak duality theorem states that for a primal maximization (minimization) problem, any feasible solution has a value that is less than (greater than) or equal to the value of any feasible solution for its dual minimization (maximization) problem. The strong duality theorem states that the value of the optimal solution of the primal equals that of the dual. In addition, the complementary slackness theorem establishes that for a given optimal solution of a feasible and bounded primal, there exists an optimal solution of the dual such that if a primal constraint is not tightly satisfied, then the dual variable corresponding to the primal constraint is equal to 0; and for a given optimal solution of the dual, there exists an optimal solution of the primal such that for any dual variable that is not equal to 0, its corresponding primal constraint is tightly satisfied. Finally, if the value of the optimal solution of the primal is unbounded, then the dual is infeasible, and vice versa.

LP problems can be solved in polynomial time (Khachiyan 1980; Karmarkar 1984). Moreover, LP duality is a very powerful mathematical concept that is extensively used for analyzing combinatorial problems. LP formulations and their duals have been successfully used in the design of efficient polynomial-time exact/approximation algorithms for many combinatorial problems for which other techniques are unviable (Vazirani 2001).

The maxflow problem is a special case of an LP problem which can be formulated graph-theoretically. It is characterized by a directed graph $G = \langle V, E \rangle$, where $V = \{v_1, v_2, \ldots, v_n\}$ is the set of nodes, and E is the set of directed edges. There are two special nodes, $s \in V$ referred to as the source and $t \in V$ referred to as the sink. Each edge $e = \langle v_i, v_j \rangle \in E$ is annotated with a nonnegative capacity, $U(e)$. A *flow* is an assignment of a nonnegative number $f(e)$ for each edge, such that two kinds of constraints are satisfied: (capacity constraints) for each edge $e \in E$, $0 \leq f(e) \leq U(e)$; and (conservation constraints) for each node $v \in V \backslash \{s, t\}$, $\sum_{v_i : \langle v_i, v \rangle \in E} f(\langle v_i, v \rangle) = \sum_{v_i : \langle v, v_i \rangle \in E} f(\langle v, v_i \rangle)$. The goal is to maximize the total amount of flow emanating from s (or equivalently, sinking into t) (Kleinberg and Tardos 2006). Figure 1a shows an example of a maxflow problem where each edge is annotated with its capacity.

It is easy to see that the capacity constraints, the conservation constraints, and the objective function of a maxflow problem are all linear with respect to the flow variables $f(e)$'s. Therefore, it can be solved in polynomial time using a generic LP solver. However, because of its special structure, it can be solved very efficiently in strongly polynomial time (Orlin 2013). Most efficient procedures for the maxflow problem make use of the *residual graph* associated with a partial solution. The LP dual of a maxflow problem is the well-known *mincut* problem.

The *mincost circulation problem* (MCCP) is a generalization of the maxflow problem with two primary differences. First, there are no special nodes s and t, i.e., the conservation constraints are required to be true for all nodes. Second, each edge e is not only annotated with a nonnegative capacity $U(e)$ as before but is also annotated with a *cost* $c(e)$ that is free to be zero, positive, or negative. The goal is to minimize $\sum_{e \in E} c(e)f(e)$. The maxflow problem can be considered as a special case of the MCCP by adding a back edge from t to s of capacity $+\infty$ and cost $-\infty$. Like the maxflow problem, the MCCP also qualifies as a special case of an LP problem and can be solved in polynomial time using a generic LP solver and in strongly polynomial time using efficient algorithms that leverage its additional structure (Tardos 1985; Goldberg and Tarjan 1989; Orlin 1993). Figure 1b shows an example of an MCCP where each edge is annotated with a capacity as well as a cost.

The STP and its LP Dual

In this section, we describe the STP as an instance of the LP problem and examine its dual. We first present an alternative and interchangeable formulation of the STP for a simpler and more intuitive explanation of its dual. Specif-

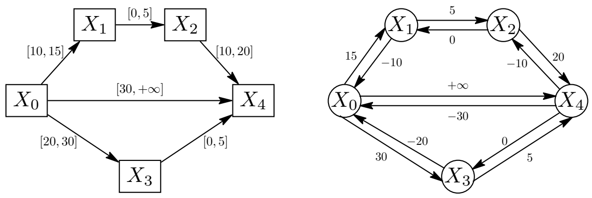

(a) (b)

Figure 2: (a) shows an STP in its original formulation. Each directed edge represents a simple temporal constraint specified using a lower and an upper bound. For example, X_1 and X_2 are connected by a directed edge that encodes the simple temporal constraint $0 \leq X_2 - X_1 \leq 5$. (b) shows the MCCP representation of the LP dual of the STP assuming that each edge is of infinite capacity and the annotation on it represents the cost of pushing a unit flow through it.

ically, we define an STP S as a directed graph $\langle \mathcal{X}, \mathcal{E}' \rangle$, where $\mathcal{X}$ remains the same as previously defined, but $\mathcal{E}'$, the set of directed edges, is defined as follows. For each edge $e_{ij} = \langle X_i, X_j \rangle \in \mathcal{E}$ annotated with $[LB(e_{ij}), UB(e_{ij})]$, we construct two edges in $\mathcal{E}'$, namely, e'_{ij} annotated with ρ_{ij} and e'_{ji} annotated with ρ_{ji}, where $\rho_{ij} = UB(e_{ij})$ and $\rho_{ji} = -LB(e_{ij})$. In this new formulation, any edge e'_{uv} with its annotation ρ_{uv} represents the simple temporal constraint $X_v - X_u \leq \rho_{uv}$.

The primal LP formulation of an STP is as follows:

$$\max \quad \sum_{i=1}^{N} 0 \cdot X_i \quad \text{s.t.} \tag{1}$$

$$\forall\, e'_{ij} \in \mathcal{E}' : X_j - X_i \leq \rho_{ij}. \tag{2}$$

We note that since the STP is a feasibility problem with a degenerate objective function, $X_0 = 0$ does not have to be explicitly encoded as it is merely a reference point. The LP dual of the above STP is as follows:

$$\min \quad \sum_{e'_{ij} \in \mathcal{E}'} \rho_{ij} \cdot f_{ij} \quad \text{s.t.} \tag{3}$$

$$\forall\, X_i \in \mathcal{X} : \quad \sum_{j : e'_{ji} \in \mathcal{E}'} f_{ji} - \sum_{j : e'_{ij} \in \mathcal{E}'} f_{ij} = 0 \tag{4}$$

$$\forall\, e'_{ij} \in \mathcal{E}' : \quad f_{ij} \geq 0. \tag{5}$$

Here, the f_{ij}'s are the dual variables corresponding to the primal constraints in Equation (2). They are nonnegative because the corresponding primal constraints are inequalities. Similarly, the linear constraints in Equation (4) are equalities because the corresponding primal variables are free variables, i.e., the X_i's can be zero, positive, or negative.

We note that the LP dual of the STP can be interpreted as an MCCP by treating the X_i's as the nodes and each dual variable f_{ij} as the flow variable on e'_{ij}. Indeed, f_{ij} is constrained to be nonnegative as required by the MCCP; and the linear equalities in Equation (4) correspond to the conservation constraints for all nodes. The ρ_{ij}'s can be interpreted as the costs on the edges and are allowed to be zero, positive, or negative. The capacity constraints, however, do not feature in the LP dual. This is equivalent to having an infinite capacity on each edge in the MCCP. Figure 2 shows an STP and the MCCP representation of its LP dual.

It is easy to see that an MCCP instance that does not have capacity constraints on its edges admits an optimal solution with unbounded value if and only if there is a negative cost cycle. In such a case, an increasing amount of flow can circulate through this cycle, driving down the value of a feasible flow indefinitely. By LP duality, the unbounded nature of the dual corresponds to the infeasibility of the primal. Therefore, the STP is feasible if and only if the MCCP interpretation of its LP dual does not have any negative cost cycles.

This characterization of the feasibility of an STP has been discovered in (Dechter, Meiri, and Pearl 1991) using the distance graph representation. Therefore, the distance graph representation of an STP is actually a representation of its LP dual.

STPs with Linear Objective Functions

In this section, we consider LP problems that have simple temporal constraints and a linear objective function. That is, they are of the following form:

$$\max \quad \sum_{i=1}^{N} w_i \cdot X_i \quad \text{s.t.} \tag{6}$$

$$\forall\, e'_{ij} \in \mathcal{E}' : X_j - X_i \leq \rho_{ij}. \tag{7}$$

We note that this LP problem has an optimal solution of unbounded value if the simple temporal constraints are feasible. This is because the reference point X_0 does not yet have a fixed value in this formulation while all the simple temporal constraints in Equation (7) are relative. Therefore, given any feasible solution, simply adding a fixed value to all variables preserves feasibility. However, doing so can increase the value of the objective function indefinitely and make it unbounded. To avoid this problem, we can either set X_0 to be equal to 0 using a constraint or we can formulate the following equivalent LP problem:

$$\max \quad \sum_{i=1}^{N} w_i \cdot (X_i - X_0) \quad \text{s.t.} \tag{8}$$

$$\forall\, e'_{ij} \in \mathcal{E}' : X_j - X_i \leq \rho_{ij}. \tag{9}$$

The LP dual of the above formulation is as follows:

$$\min \quad \sum_{e'_{ij} \in \mathcal{E}'} \rho_{ij} \cdot f_{ij} \quad \text{s.t.} \tag{10}$$

$$\forall X_i \in \mathcal{X} \backslash X_0 : \quad \sum_{j : e'_{ji} \in \mathcal{E}'} f_{ji} - \sum_{j : e'_{ij} \in \mathcal{E}'} f_{ij} = w_i \tag{11}$$

$$\forall\, e'_{ij} \in \mathcal{E}' : \quad f_{ij} \geq 0 \tag{12}$$

$$\sum_{j : e'_{j0} \in \mathcal{E}'} f_{j0} - \sum_{j : e'_{0j} \in \mathcal{E}'} f_{0j} = -\sum_{i=1}^{N} w_i. \tag{13}$$

As before, we can try to interpret the LP dual as an MCCP instance. Let $\delta(X_i)$ be the sum of the incoming flows to X_i minus the sum of the outgoing flows from X_i. We note that

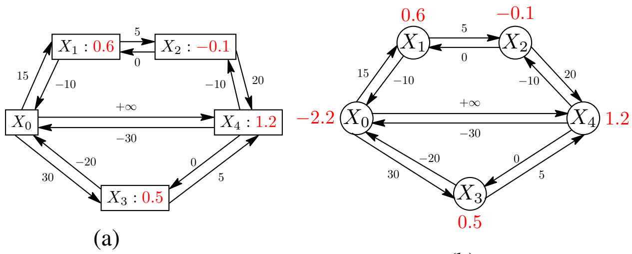

(a)

(b)

Figure 3: (a) shows the STP in Figure 2a with a linear objective function. Each node other than X_0 is annotated with a number w_i that corresponds to the term $w_i \cdot (X_i - X_0)$ in the objective function. (b) shows the MCCP representation of the LP dual of the STP where each edge is only annotated with its cost as in Figure 2b. Here, each node is annotated with $\delta(X_i)$.

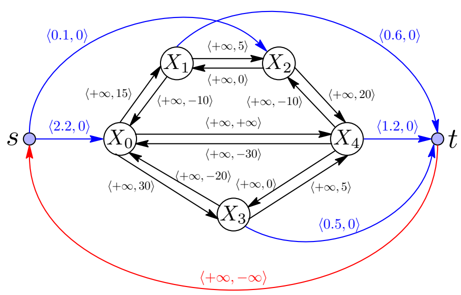

Figure 4: Shows a new MCCP without deficits and excesses for the example in Figure 3b. We add two special nodes, $s \in V$ referred to as the source, and $t \in V$ referred to as the sink. An edge of capacity $|\delta(X_i)|$ and cost zero is added to connect s with any node X_i that has a deficit $|\delta(X_i)|$; and an edge of capacity $|\delta(X_i)|$ and cost zero is added to connect any node X_i that has an excess $|\delta(X_i)|$ with t. Such edges are shown in blue. A back edge, shown in red, is added from t to s with capacity $+\infty$ and cost $-\infty$.

$\delta(X_i) = w_i$ for $X_i \in \mathcal{X} \backslash X_0$ and $\delta(X_0) = -\sum_{i=1}^{N} w_i$. If $\delta(X_i) < 0$, we say that there is a *deficit* $|\delta(X_i)| = -\delta(X_i)$ at X_i. If $\delta(X_i) > 0$, we say that there is an *excess* $|\delta(X_i)| = +\delta(X_i)$ at X_i. We note that the only places where the w_i's appear in the dual are in Equation (11) and Equation (13). Here, they compromise the conservation constraints by introducing a deficit or an excess at each node. Once again, the capacity constraints do not feature in the dual. Figure 3 shows an STP with a linear objective function and the MCCP representation of its LP dual with deficits and excesses at its individual nodes.

Figure 4 shows how these deficits and excesses at individual nodes can be reformulated as a new MCCP without deficits or excesses at any of the nodes. In the new MCCP, we first add two special nodes: $s \in V$ referred to as the *source*, and $t \in V$ referred to as the *sink*. We then add an edge of capacity $|\delta(X_i)|$ and cost zero to either connect s with any node X_i that has a deficit $|\delta(X_i)|$ or connect any node X_i that has an excess $|\delta(X_i)|$ with t. Finally, a back edge is added from t to s with capacity $+\infty$ and cost $-\infty$.

We now prove the equivalence of the new MCCP without deficits and excesses to the original MCCP with deficits and excesses. We start by arguing that any optimal circulation g^* for the new MCCP saturates all the edges emanating from s *and* all the edges leading into t. It is easy to see that g^* should saturate all the edges emanating from s *or* all the edges leading into t.[1] This is because the flow from s to t has to be maximized for the $-\infty$ cost incentive on the back edge from t to s. It now suffices to prove that $\sum_{X_i:\delta(X_i)<0} -\delta(X_i) = \sum_{X_i:\delta(X_i)>0} \delta(X_i)$. From Equation (11), we know that for any X_i other than X_0, $\delta(X_i) = w_i$. From Equation (13), we know that $\delta(X_0) = -\sum_{i=1}^{N} w_i$. Therefore, $\sum_{X_i:\delta(X_i)>0} \delta(X_i) - \sum_{X_i:\delta(X_i)<0} -\delta(X_i) = \sum_{X_i \in \mathcal{X} \backslash X_0 : w_i > 0} w_i - \sum_{X_i \in \mathcal{X} \backslash X_0 : w_i < 0} -w_i + \delta(X_0) = \sum_{X_i \in \mathcal{X} \backslash X_0} w_i - \sum_{i=1}^{N} w_i = 0$.

We then prove that if f^* is the optimal circulation for the original MCCP, it can be extended to an optimal circulation g^* for the new MCCP, and vice versa. Given f^*, let $g^*_{ij} = f^*_{ij}$ for all edges $e'_{ij} \in \mathcal{E}'$; let $g^*_{si} = |\delta(X_i)|$ for a newly added edge from s to X_i; and let $g^*_{it} = |\delta(X_i)|$ for a newly added edge from X_i to t. It is clear that g^* is a feasible circulation since deficits and excesses in the original MCCP are now compensated by the saturation of the new edges. g^* is also an optimal circulation for two reasons: (a) f^* is an optimal circulation through all edges that do not involve s or t; and (b) g^* saturates all the edges emanating from s *and* all the edges leading into t, and from the arguments in the previous paragraph, no other optimal circulation can push different amounts of flow on these edges. Conversely, given an optimal circulation g^* for the new MCCP, its projection onto the edges in $\mathcal{E}'$ is an optimal circulation f^* for the original MCCP. This can be easily proved by contradiction. Suppose there exists a circulation f' that is better than f^*. Then, f' can be extended to an optimal circulation g' for the new MCCP that is better than g^* (since all saturated edges from s and all saturated edges to t contribute equally to g' and g^*). This contradicts our assumption of an optimal g^*.

From the previous section, we know that the primal is infeasible if and only if the distance graph has a negative cost cycle. This is consistent with the fact that the same distance graph appears in the dual MCCP, in which we can circulate an infinite amount of flow and make the dual MCCP unbounded if and only if there is a negative cost cycle. We note that the edges added in the new MCCP cannot participate in an infinite circulation because of the bounded capacities on the edges emanating from s and the edges leading into t.

LP duality can also help us recognize an unbounded primal. We know that an unbounded primal is characterized by an infeasible dual. Consider a feasible circulation g for the new MCCP that saturates all the edges emanating from s and all the edges leading into t. Its projection onto the edges in $\mathcal{E}'$ is a feasible circulation for the original MCCP. Conversely, a

[1] For simplicity, we assume that there is a path from s to t. If this is not the case, trivial simple temporal constraints can be added to induce edges of high costs that make t reachable from s.

Algorithm 1: Shows a strongly-polynomial-time algorithm for solving STPs with linear objective functions using LP duality and flow-based techniques.

1 **Function** SOLVE-STP-WITH-LINEAR-OBJECTIVE
  Input: An LP problem that has simple temporal constraints and a linear objection function and is of the form in Equations (8) and (9);
  Output: An optimal solution to the LP problem;
2 • *Construct the LP dual and its MCCP interpretation:*
3 Construct the LP dual of the problem with f_{ij}'s as the dual variables;
4 Interpret the LP dual as an MCCP with deficits and excesses;
5 • *Reformulate the LP dual as a new MCCP without deficits and excesses:*
6 Add two special nodes, s referred to as the source, and t referred to as the sink;
7 Add a back edge from t to s with capacity $+\infty$ and cost $-\infty$;
8 For each node X_i that has a deficit $|\delta(X_i)|$, add an edge of capacity $|\delta(X_i)|$ and cost zero from s to X_i;
9 For each node X_i that has an excess $|\delta(X_i)|$, add an edge of capacity $|\delta(X_i)|$ and cost zero from X_i to t;
10 • *Solve the new MCCP:*
11 Use the strongly-polynomial-time algorithm in (Orlin 1993) to solve the new MCCP;
12 • *Tighten primal constraints using complementary slackness:*
13 For each dual variable $f_{ij} : e'_{ij} \in \mathcal{E}'$, if $f_{ij} \neq 0$, tighten the corresponding constraint in the primal such that $X_j - X_i = \rho_{ij}$;
14 • *Solve the new STP and return solution:*
15 Solve the new STP with tightened primal constraints using the Bellman-Ford algorithm;
16 **return** *any solution to this new STP*;

feasible circulation for the original MCCP is also a feasible circulation for the new MCCP when it is augmented with flow that saturates all the edges emanating from s and all the edges leading into t. Therefore, the primal is bounded if and only if the original MCCP is feasible, and if and only if it is possible to saturate all the edges emanating from s and all the edges leading into t in the new MCCP.

Algorithm 1 shows how to solve an STP with a linear objective function in strongly polynomial time when it is feasible and bounded. The algorithm is based on an interpretation of the LP dual of the problem as an MCCP with deficits and excesses (lines 2-4) and a subsequent reformulation of this MCCP to a new MCCP without deficits and excesses (lines 5-9). We showed that the new MCCP is equivalent to the LP dual of the STP with a linear objective function. Algorithm 1 solves the MCCP (lines 10-11) and makes use of its solution in accordance with the complementary slackness theorem to tighten those primal constraints whose corresponding dual variable does not equal 0 (lines 12-13).[2] The tightened primal is now a regular STP that can be solved easily using the Bellman-Ford algorithm (lines 14-15). The time complexity of Algorithm 1 is dominated by lines 11 and 15. Line 11 is of complexity $O(|\mathcal{E}| \log N(|\mathcal{E}| + N \log N))$, and line 15 is of complexity $O(N|\mathcal{E}|^2)$. Put together, the algorithm is of strongly polynomial time complexity.

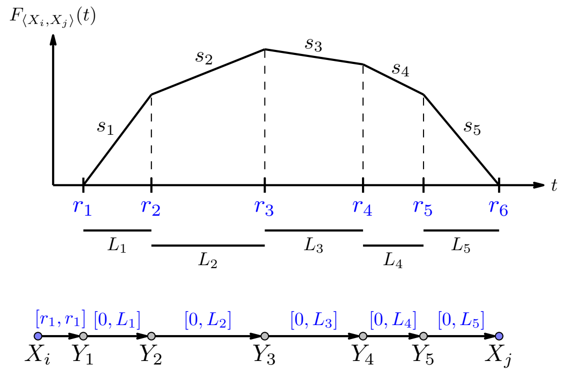

Figure 5: The top part of the figure shows a piecewise linear and convex $F_{\langle X_i,X_j \rangle}(t)$. Here, there are $K_{\langle X_i,X_j \rangle} = 6$ landmarks, denoted by $r_1, r_2, \ldots, r_6$. $L_1, L_2, \ldots, L_5$ are the lengths of the intervals between consecutive landmarks. $s_1, s_2, \ldots, s_5$ are the slopes of $F_{\langle X_i,X_j \rangle}(t)$ in these intervals. The bottom part of the figure shows a representation of the auxiliary variables and the linear constraints between them created for casting $F_{\langle X_i,X_j \rangle}(t)$ as the projection of an LP problem.

We note that the problem of minimizing the makespan of an STP and the problem of maximizing the throughput of an STP are both special cases of this LP problem. For makespan minimization, we can simply introduce a "finish" variable X_f such that $w_f = -1$ and for all $X_{i:i \neq f}$, $X_i - X_f \leq 0$ and $w_i = 0$. Such a formulation provides an incentive to schedule the finish event X_f as early as possible, thereby minimizing the makespan. For throughput maximization, we simply set $w_i = -1$ for all i. This is because throughput maximization is equivalent to the minimization of the average finishing time of all events. Although makespan minimization and throughput maximization for STPs were known to be tractable before, here, we demonstrate not only that they fall under the perspective of LP duality but also that Algorithm 1 solves a much more general problem for STPs with linear objective functions.

STPs with Piecewise Linear Convex Preferences

In this section, we examine an LP formulation of STPs with piecewise linear convex preferences. Such STPs with preferences (STPPs) have been studied in works such as (Yorke-

[2]The MCCP is feasible because the primal is bounded.

Smith, Venable, and Rossi 2003; Morris et al. 2004; Kumar 2007). We show that these STPPs are also amenable to more efficient algorithms based on LP duality. Consider such an STPP instance where each edge $e = \langle X_i, X_j \rangle$ is annotated with a function $F_e(t)$ that is piecewise linear and convex with respect to t. Figure 5 shows an example of such a function. The goal is to maximize the sum of the preferences $\sum_{e=\langle X_i,X_j\rangle\in\mathcal{E}} F_e(X_j - X_i)$.

Let the landmarks of $F_e(t)$ be $r^e_1, r^e_2, \ldots, r^e_{K_e}$ as indicated in Figure 5. Let L^e_m be the length of the interval between r^e_m and r^e_{m+1}. Let s^e_m be the slope of $F_e(t)$ in the interval between r^e_m and r^e_{m+1}. We assume that $X_j - X_i$ is restricted to be in the interval between r^e_1 and $r^e_{K_e}$, where $e = \langle X_i, X_j \rangle$.

We first observe that when $F_e(t)$ is piecewise linear and convex, the slopes of the line segments defining the contour of $F_e(t)$ are in decreasing order, e.g., $s_1 > s_2 \cdots > s_5$ in Figure 5. We can exploit this observation and reformulate $F_{\langle X_i,X_j\rangle}(X_j - X_i)$ as the projection of an LP problem—with auxiliary variables—onto the variables X_i and X_j as alluded to in (Kumar 2007). The value of $F_{\langle X_i,X_j\rangle}(X_j - X_i)$ is equal to the value of the following LP problem (where $e = \langle X_i, X_j \rangle$ is used for simplicity of notation):

$$\max \sum_{q=1}^{K_e-2} s_q \cdot (Y_{q+1} - Y_q) + s_{(K_e-1)} \cdot (X_j - Y_{(K_e-1)}) \quad \text{s.t.} \tag{14}$$

$$\forall 1 \le q \le (K_e - 2) : 0 \le Y_{q+1} - Y_q \le L_q \tag{15}$$

$$0 \le X_j - Y_{(K_e-1)} \le L_{(K_e-1)} \tag{16}$$

$$Y_1 - X_i = r^e_1. \tag{17}$$

We note that in the above LP problem, the variables are $Y_1, Y_2, \ldots, Y_{(K_e-1)}$. We now prove that the LP problem over these auxiliary variables represents $F_e(X_j - X_i)$. This is because the above LP problem has the same value as the following simplified LP problem:

$$\max \sum_{q=1}^{K_e-1} s_q \cdot \ell_q \quad \text{s.t.} \tag{18}$$

$$\forall\ 1 \le q \le (K_e - 1) : 0 \le \ell_q \le L_q \tag{19}$$

$$\ell_0 = r^e_1 \tag{20}$$

$$\sum_{q=0}^{K_e-1} \ell_q = X_j - X_i. \tag{21}$$

Here, ℓ_0 represents $(Y_1 - X_i)$, ℓ_q represents $(Y_{q+1} - Y_q)$ for $q : 1 \le q \le K_e - 2$, and $\ell_{(K_e-1)}$ represents $(X_j - Y_{(K_e-1)})$. It is easy to note that this LP problem is equivalent to the well-known Fractional Knapsack Problem that admits a greedy solution (Cormen et al. 2009). The polynomial-time greedy algorithm for this problem maximizes ℓ_1 until it hits its upper bound L_1, after which, it maximizes ℓ_2 until it hits its upper bound L_2, and so forth, until the running sum of ℓ's hits $(X_j - X_i)$. All remaining ℓ's are set to 0. Here, $(X_j - X_i - r^e_1)$ represents the capacity of the knapsack; L_q represents the available amount of commodity $q \ge 1$; and s_q represents the value of q.

We now show that the value of the optimal solution produced by this greedy algorithm matches $F_e(X_j - X_i)$. This is easy to prove by induction on the intervals between r^e_m and r^e_{m+1} for increasing values of m: In both the base case and the inductive step, since $s_1 > s_2 > \cdots > s_{(K_e-1)}$, in any interval between r^e_m and r^e_{m+1}, both $F_e(X_j - X_i)$ and the value of the solution produced by the greedy algorithm increase at the rate of s^e_m.

Having represented a piecewise linear convex preference function as the projection of an LP problem, we can now represent the entire STPP with piecewise linear convex preference functions as a composite LP problem. In other words, the STPP

$$\max \sum_{e=\langle X_i,X_j\rangle\in\mathcal{E}} F_e(X_j - X_i) \quad \text{s.t.} \quad X_0 = 0 \tag{22}$$

can now be written as the following LP problem between Equations (23) and (27) (where $e = \langle X^e_i, X^e_j \rangle$ is used for simplicity of notation):

$$\max \sum_{e\in\mathcal{E}} \sum_{q=1}^{K_e-2} s^e_q \cdot (Y^e_{q+1} - Y^e_q) + s^e_{(K_e-1)} \cdot (X^e_j - Y^e_{(K_e-1)}) \tag{23}$$

$$\forall e \in \mathcal{E}, 1 \le q \le (K_e - 2) : 0 \le Y^e_{q+1} - Y^e_q \le L^e_q \tag{24}$$

$$\forall e \in \mathcal{E} : 0 \le X^e_j - Y^e_{(K_e-1)} \le L^e_{(K_e-1)} \tag{25}$$

$$\forall e \in \mathcal{E} : Y^e_1 - X^e_i = r^e_1 \tag{26}$$

$$X_0 = 0. \tag{27}$$

We note that in this composite LP formulation of the STPP, the variables are the X's and the Y's. The constraints between them listed in Equations (24) - (27) are all simple temporal constraints. Moreover, the objective function is linear with respect to these variables. Therefore, this LP problem fits the case discussed in the previous section. Its dual can be interpreted as an MCCP as discussed before.

Algorithm 1 solves the above LP problem in strongly polynomial time, i.e., in $O(|\tilde{\mathcal{E}}| \log \tilde{N}(|\tilde{\mathcal{E}}| + \tilde{N} \log \tilde{N}) + \tilde{N}|\tilde{\mathcal{E}}|^2)$ time, as shown before. Here, $\tilde{N}$ represents the total number of variables in the above LP problem, given by $O(\sum_{e\in\mathcal{E}} K_e)$; and $|\tilde{\mathcal{E}}|$ represents the total number of simple temporal constraints in the above LP problem, given by $O(\max_{e\in\mathcal{E}} K_e \cdot |\mathcal{E}|)$. This time complexity is better than that of previous approaches which directly call a generic LP solver and do not have a strongly polynomial running time (Morris et al. 2004). In our approach, LP duality—instead of an LP solver—is used to design a low-order strongly-polynomial-time algorithm.

Conclusions and Future Work

In this paper, we provided an LP duality perspective on temporal reasoning problems. We showed that an STP can be formulated as an LP problem and that its LP dual can be interpreted as an MCCP instance. The MCCP representation can be used to characterize the feasibility of an STP, and this fits into the existing distance graph representation. We then provided an LP formulation for an STP with a linear objective function that generalizes makespan minimization and throughput maximization. We proved that the LP dual of the problem can be interpreted as an MCCP with deficits and excesses, which can then be reformulated as a new regular MCCP. We then designed a strongly-polynomial-time algorithm for solving such a problem using the general theory of LP duality and flow-based techniques. Furthermore, we studied STPPs with piecewise linear convex preference functions and reduced them to STPs with linear objective functions. This enabled us to provide a very efficient strongly-polynomial-time algorithm for an important class of STPPs as well.

There are many avenues for future work. One important avenue is to use LP duality-based strongly-polynomial-time algorithms for load scheduling problems in smart home and smart grid domains for energy cost minimization, which typically have piecewise linear cost functions that are not necessarily convex but are on individual variables. Another avenue is to study spatiotemporal reasoning problems—that are more expressive than regular temporal reasoning problems—from the perspective of LP duality.

Acknowledgment

This material is based upon work supported by the Air Force Research Laboratory (AFRL) and the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR0011-15-C-0138. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the official views or policies of the Department of Defense or the U.S. Government.

References

Augusto, J. C. 2005. Temporal reasoning for decision support in medicine. *Artificial Intelligence in Medicine* 33(1):1–24.

Chvátal, V. 1983. *Linear Programming*. W. H. Freeman and Co.

Cormen, T. H.; Leiserson, C. E.; Rivest, R. L.; and Stein, C. 2009. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition.

de Weerdt, M. 2003. *Plan Merging in Multi-Agent Systems*. Ph.D. Dissertation, Delft University of Technology.

Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. *Artificial Intelligence* 49(1-3):61–95.

Feijer, D.; Savla, K.; and Frazzoli, E. 2012. Strategic dynamic vehicle routing with spatio-temporal dependent demands. In *American Control Conference (ACC)*.

Goldberg, A. V., and Tarjan, R. E. 1989. Finding minimum-cost circulations by canceling negative cycles. *Journal of the ACM (JACM)* 36(4):873–886.

Ji, M.; He, Y.; and Cheng, T. E. 2007. Single-machine scheduling with periodic maintenance to minimize makespan. *Computers & Operations Research* 34(6):1764–1770.

Karmarkar, N. 1984. A new polynomial-time algorithm for linear programming. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing (STOC)*, 302–311.

Khachiyan, L. G. 1980. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics* 20(1):53–72.

Kleinberg, J., and Tardos, E. 2006. *Algorithm design*. Pearson Education.

Knight, S.; Rabideau, G.; Chien, S.; Engelhardt, B.; and Sherwood, R. 2001. Casper: Space exploration through continuous planning. *IEEE Intelligent Systems* 16(5):70–75.

Kumar, T. K. S. 2003. Incremental computation of resource-envelopes in producer-consumer models. In *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming (CP)*.

Kumar, T. K. S. 2004. A polynomial-time algorithm for simple temporal problems with piecewise constant domain preference functions. In *Proceedings of the 19th National Conference on Artificial intelligence (AAAI)*.

Kumar, T. K. S. 2005. On the tractability of restricted disjunctive temporal problems. In *Proceedings of the 15th International Conference on Automated Planning and Scheduling (ICAPS)*, 110–119.

Kumar, T. K. S. 2006. Tractable classes of metric temporal problems with domain rules. In *Proceedings of the 21st National Conference on Artificial intelligence (AAAI)*.

Kumar, T. K. S. 2007. Fast (incremental) algorithms for useful classes of simple temporal problems with preferences. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*.

Mohsenian-Rad, A.-H., and Leon-Garcia, A. 2010. Optimal residential load control with price prediction in real-time electricity pricing environments. *IEEE transactions on Smart Grid* 1(2).

Morris, P.; Morris, R.; Khatib, L.; Ramakrishnan, S.; and Bachmann, A. 2004. Strategies for global optimization of temporal preferences. In *Proceedings of the 10th International Conference on Principles and Practice of Constraint Programming (CP)*.

Muscettola, N. 2004. Incremental maximum flows for fast envelope computation. In *Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS)*.

Nguyen, X., and Kambhampati, S. 2001. Reviving partial order planning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*.

Oddi, A., and Cesta, A. 2000. Incremental forward checking for the disjunctive temporal problem. In *Proceedings of the 14th European Conference on Artificial Intelligence*.

Orlin, J. B. 1993. A faster strongly polynomial minimum cost flow algorithm. *Operations Research* 41(2):338–350.

Orlin, J. B. 2013. Max flows in $o(nm)$ time, or better. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, 765–774.

Planken, L.; de Weerdt, M.; and van der Krogt, R. 2008. P^3C: a new algorithm for the simple temporal problem. In

Proceedings of the 18th International Conference on Automated Planning and Scheduling (ICAPS), 256–263.

Shah, J. A., and Williams, B. C. 2008. Fast dynamic scheduling of disjunctive temporal constraint networks through incremental compilation. In *Proceedings of the 18th International Conference on Automated Planning and Scheduling (ICAPS)*.

Smith, S. F., and Cheng, C.-C. 1993. Slack-based heuristics for constraint satisfaction scheduling. In *Proceedings of the 11th National Conference on Artificial Intelligence (AAAI)*, 139–144.

Stergiou, K., and Koubarakis, M. 1998. Backtracking algorithms for disjunctions of temporal constraints. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI)*.

Tardos, É. 1985. A strongly polynomial minimum cost circulation algorithm. *Combinatorica* 5(3):247–255.

Vazirani, V. V. 2001. *Approximation Algorithms*. Springer.

Xu, L., and Choueiry, B. Y. 2003. A new efficient algorithm for solving the simple temporal problem. In *Proceedings of the 10th International Symposium on Temporal Representation and Reasoning and the Fourth International Conference on Temporal Logic*.

Yorke-Smith, N.; Venable, K. B.; and Rossi, F. 2003. Temporal reasoning with preferences and uncertainty. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*.