

Generalizing Temporal Controllability

Michael D. Moffitt and Martha E. Pollack

Department of Electrical Engineering and Computer Science

University of Michigan

Ann Arbor, MI 48109, USA

{mmoffitt, pollackm}@eecs.umich.edu

Abstract

In this paper, we focus on extending the expressive power of constraint-based temporal reasoning formalisms. We begin with the well-known Simple Temporal Problem with Uncertainty, and incorporate three extensions: *prior observability*, in which the values of uncontrollable events become known prior to their actual occurrence; *partial shrinkage*, in which an observation event triggers the reduction of a contingent temporal interval; and a generalization of partial shrinkage to requirement links, making it possible to express certain types of uncertainty that may arise even when the time points in a problem are themselves fully controllable. We describe levels of controllability in the resulting formalism, the Generalized STPU, and relate this formalism to related developments in disjunctive temporal reasoning. Throughout, we motivate our approach with simple, real-world examples that illustrate the limitations of existing formalisms and the flexibility of our proposed extensions.

1 Introduction

Several recent studies have addressed the problem of uncertainty in constraint-based temporal reasoning [Vidal and Ghallab, 1996; Vidal and Fargier, 1999; Morris *et al.*, 2001]. In this line of research, traditional temporal networks [Dechter *et al.*, 1991] are extended to include two classes of time points: *controllable* ones, whose time of occurrence can be set by the execution agent, and *uncontrollable* ones that are instead determined by an external force, often called “Nature”. As a result, the problem of consistency is transformed into one of *controllability*. Recent efforts have made significant progress toward increasing both the expressivity [Venable and Yorke-Smith, 2005] and efficiency [Morris and Muscettola, 2005] of constraint-based temporal reasoning with uncertainty.

However, significant limitations in expressive power do exist in the formalisms developed to date. In particular, they have all assumed that temporal uncertainty is resolved in a specific way: the execution agent discovers the duration of a contingent link only at the exact moment of its conclusion (or equivalently, it discovers the time of a uncontrollable time

point only at the time of its occurrence). This is not an accurate reflection of how such knowledge is obtained in many practical situations: one often learns the time of an initially uncertain event before its actual execution, and can use that knowledge to make informed decisions about the remainder of the plan. In addition, one may acquire only partial information about the time of an upcoming event. Finally, and perhaps surprisingly, certain varieties of uncertainty may exist even for problems that should be modeled entirely with controllable time points, something that again cannot be captured by previous formalisms.

In this paper, we generalize the problem of controllability in constraint-based temporal reasoning. We begin with the Simple Temporal Problem with Uncertainty, an increasingly popular framework that models uncertain events and contingent constraints. We then consider three progressive extensions to this formalism. The first introduces the notion of *prior observability*, in which the values of uncontrollable events become known prior to their actual occurrence. The second captures situations in which an observation event triggers the reduction of a contingent temporal interval, a form of *partial observability* that has not yet been captured in previous work. The third and final extension generalizes this notion to the requirement links, making it possible to express certain additional types of uncertainty. Finally, we draw relationships between our extended framework and related developments in disjunctive temporal reasoning, and suggest promising avenues of continued research. Throughout the text, we make use of practical, real-world examples that illustrate the limitations of existing formalisms and the flexibility of our proposed extensions.

2 Background

2.1 Simple Temporal Problems

The *Simple Temporal Problem* [Dechter *et al.*, 1991] (STP) is defined by a pair $\langle X, E \rangle$, where each X_i in X designates a time point, and each E_{ij} in E is a constraint of the form

$$l_{ij} \leq x_i - x_j \leq u_{ij}$$

with $x_i, x_j \in X$ and $l_{ij}, u_{ij} \in \mathfrak{R}$. An STP is said to be *consistent* if there exists a *solution* $S : X \rightarrow \mathfrak{R}$ that satisfies all constraints in E .

Each STP has an associated graph-based encoding, which contains a node for each time point, and a directed edge for

each inequality having weight u_{ij} or $-l_{ij}$ (depending on its direction). An STP is consistent if and only if its corresponding graph contains no negative cycles; this can be determined in polynomial time by computing its All Pairs Shortest Path matrix and examining the entries along the main diagonal.

2.2 Simple Temporal Problems with Uncertainty

The STP models situations in which the agent in charge of plan execution is in complete control of all time points. The *Simple Temporal Problem with Uncertainty* (STPU) is an extension of the STP that relaxes this assumption. Specifically, the STPU is defined as a tuple $\langle X_C, X_U, E, C \rangle$, where:

- X_C and X_U are sets of *controllable* and *uncontrollable* time points, respectively. Their union, $X_C \cup X_U$, forms an entire set X of time points.
- E is a set of *requirement links*, where each E_{ij} is of the form $l_{ij} \leq x_i - x_j \leq u_{ij}$.
- C is a set of *contingent links*, where each C_{ij} is of the form $l_{ij} \leq x_i - x_j \leq u_{ij}$ and $x_i \in X_U$.¹

The contingent links in the STPU can be regarded as representing causal processes whose durations are uncertain, and thus their endpoints (the uncontrollable time points) are determined by some external force. The remaining time points are in the control of the agent, who is charged with the task of assigning them in such a way as to satisfy the requirement links.

It is often convenient to refer to a *projection* p of the STPU [Vidal and Ghallab, 1996], which is simply an STP obtained by replacing the interval of each contingent link $[l, u]$ with a particular fixed bound $[b, b]$ where $l \leq b \leq u$. A *schedule* T is defined as a mapping

$$T : X \rightarrow \mathfrak{R}$$

where $T(x)$ is the *time* of time point x . A schedule is deemed *consistent* if it satisfies all links. The *prehistory* of a time point x with respect to a schedule T , denoted $T\{\prec x\}$, specifies the durations of all contingent links that finish prior to x ; it can be determined directly from any schedule.

Finally, we define an *execution strategy* S as a mapping:

$$S : \mathcal{P} \rightarrow \mathcal{T}$$

where \mathcal{P} is the set of all projections, and \mathcal{T} is the set of all schedules. An execution strategy S is *viable* if $S(p)$, henceforth written S_p , is consistent with p for each projection p .

2.3 Controllability of the STPU

With the addition of uncontrollable events and contingent links in the STPU, the previously defined notion of consistency for the STP is no longer sufficient. Instead, we explore various flavors of *controllability*, which have all been defined previously in [Vidal, 2000]. To illustrate these concepts, we put forth the following example (which will later be extended):

Example 1: ACME Appliances is delivering a new refrigerator to your apartment at some (unknown) time T_D between

¹As in prior work, we assume $0 < l_{ij} < u_{ij} < \infty$ for each contingent link, and that contingent links do not share endpoints.

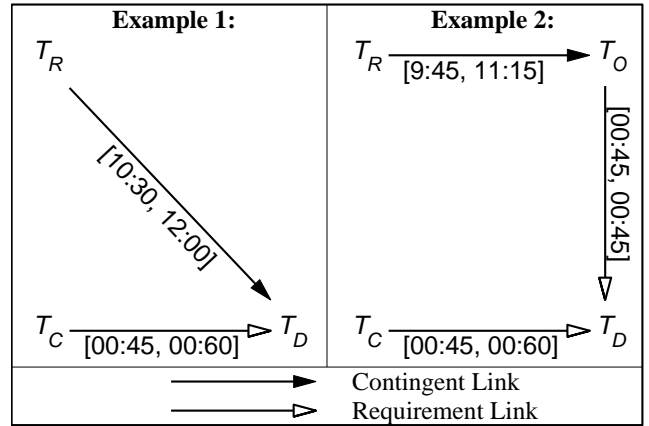


Figure 1: Networks corresponding to Examples 1 and 2.

10:30 and 12:00. At time T_C , you begin to clean out the items in your current fridge; to guarantee that all items can be removed without any spoiling, this must happen no later than 45 minutes, and no earlier than 60 minutes, before they arrive. \square

The network corresponding to this example is depicted graphically in Figure 1 (left). The time point T_R is a temporal *reference point* representing midnight, and is used to express constraints with respect to wall-clock time.

Weak Controllability

An STPU is said to be *Weakly Controllable* if there is a viable execution strategy; in other words, for every possible projection, there exists a consistent solution. Our example is indeed weakly controllable; if we happen to know what Nature will assign to T_D , we can obtain a consistent solution by setting T_C to $T_D - x$ for any x between 45 and 60 minutes.

Strong Controllability

An STPU is said to be *Strongly Controllable* if there is a viable execution strategy S such that

$$S_{p1}(x) = S_{p2}(x)$$

for each controllable time point x and all projections $p1$ and $p2$, i.e., there exists a single consistent solution that satisfies every possible projection. Our example is clearly not strongly controllable; for instance, there is no solution for the case when $T_D = 10:30$ that will also work for the case when $T_D = 11:30$.

Dynamic Controllability

The most interesting and useful type of controllability is that of dynamic controllability, a concept that exploits the temporal nature of plan execution. Specifically, an STPU is said to be *Dynamically Controllable* if there is a viable execution strategy S such that:

$$S_{p1}\{\prec x\} = S_{p2}\{\prec x\} \Rightarrow S_{p1}(x) = S_{p2}(x)$$

for each controllable time point x and all projections $p1$ and $p2$. In other words, there must exist a strategy that depends on the outcomes of only those uncontrollable events that have occurred in the past (and not on those which have yet to occur). Our current example is not dynamically controllable,

since we need to set the time of T_C before we learn that T_D has occurred.

Of the three types of controllability, dynamic controllability has been the most extensively studied. Recently, it was shown to be computable in polynomial time [Morris and Muscettola, 2005].

3 Adding Observation Events

The concept of dynamic controllability assumes that the values of uncontrollable events become known at the same time as their occurrence; that is, we learn when they happen, only when they actually happen. However, as we see in the following example, the temporal uncertainty of an event is sometimes resolved earlier than its actual execution.

Example 2: Consider Example 1; however, now ACME Appliances agrees to phone you exactly 45 minutes before they arrive. \square

In this new scenario, we discover the time of the uncontrollable event *before* that particular time arrives. While obtaining this information earlier has no impact on the weak or strong controllability of the problem, it serves to enable dynamic controllability (e.g., we can begin clean out the fridge immediately after receiving the phone call).

Fortunately, prior observability of this form can be modeled with the existing STPU framework. We achieve this by introducing an uncontrollable *observation event* – for instance, T_O – and making it, rather than T_D , the endpoint of a contingent link. A requirement link can then be introduced between the new observation action and the original uncertain event. Therefore, the act of “splicing” contingent links in this fashion makes it possible for the agent to reason about the foreknowledge of events. The STPU for the current example is shown graphically in network of Figure 1 (right).

Despite the fact that our current notion of prior observability is a mere modeling issue, this concept will prove to be essential in extensions to follow.

4 Partial Shrinkage of Contingent Links

In the STPU, the occurrence of an uncontrollable event can be thought of as establishing *complete shrinkage* of a contingent link; that is, a constraint $x_i - x_j \in [l, u]$ is squeezed to a single point $x_i - x_j \in [b, b]$. However, many real-world problems demand a generalization of this mechanism, in which an intermediate reduction is achieved. We refer to this as *partial shrinkage*.

Example 3: Again consider Example 1; however, ACME Appliances agrees to phone you very early in the morning to inform you of within which of the three half-hour blocks they intend to arrive (e.g., they may say “sometime between 11:00 and 11:30”). \square

Unlike Example 2, here the time of the phone call is only a portion of the information obtained. In addition, there is knowledge conveyed during the observation event that triggers the reduction of a contingent link: the original interval $[l, u]$ is reduced to a subinterval $[l', u']$, where $l' \neq u'$. (For example, the interval $[10:30, 12:00]$ may become $[11:00, 11:30]$). This reduction becomes known simultaneously with the execution of T_O (the phone call). Unfortu-

nately, the reduction is not sufficient to achieve dynamic controllability. However, if ACME instead provided a 15-minute time window on their arrival, then the reduction would supply enough information to guarantee dynamic controllability (despite the fact that the exact delivery time would still not be known precisely).

This scenario clearly addresses a different kind of knowledge acquisition than the previous examples. Partial observability is a very general phenomenon that commonly arises in a broad range of practical instances that involve complex sequences of decision making. For instance, consider a Mars rover domain where the texture of the upcoming terrain is unknown. As the rover moves forward, it may execute a sensing action to establish whether the ground will be rough or smooth. While this information might not allow it to calculate the exact travel time, it could enable the estimation of a more accurate window on the time of arrival. The need to model constraint “shrinkage” of this type motivates the development of an extension to the STPU formalism.

4.1 The STPU⁺

We define the STPU⁺ as a tuple $\langle X_C, X_U, E, C, O \rangle$, where:

- X_C and X_U are as in the STPU (and X is their union).
- E is a set of *requirement links*, as in the STPU.
- C is a set of *generalized contingent links*, where each C_{ij} is of the form:

$$x_i - x_j \in \{d_{ij1} : [l_{ij1}, u_{ij1}], \dots, d_{ijn_{ij}} : [l_{ijn_{ij}}, u_{ijn_{ij}}]\}$$

and $x_i \in X_U$. Each sub-interval d_{ijk} is referred to as a *reduction*.²

- $O : X_O \rightarrow C$ is a mapping from observation events to generalized contingent links, where $X_O \subseteq X$. Each event in X_O specifies (at the time of its occurrence) the reduction of its corresponding constraint in C .

Turning to Example 3, let us assume that the phone call will occur sometime between 7:00 and 7:30. We can then encode Example 3 as an STPU⁺ by creating a generalized contingent link C_{deliv} containing three possible reductions:

$$T_D - T_R \in \{[10:30, 11:00], [11:00, 11:30], [11:30, 12:00]\}$$

and establishing that this link will be resolved at the time of the phone call:

$$O(T_O) = C_{deliv}$$

Regular STPU-style contingent links can be modeled as well by encoding them as generalized contingent links with singleton intervals. An example of this is the contingent link on the time of the phone call:

$$T_O - T_R \in \{[7:00, 7:30]\}$$

In Figure 2 (left) we provide a graphical interpretation of this example. The generalized contingent link is labeled with each

²As will be discussed in Section 7, these generalized links resemble – but should not be confused with – the disjunctive constraints of a TCSP, DTP, or DTPU. Note also that the sub-intervals in a generalized link may overlap, although in our examples we use non-overlapping intervals without loss of generality.

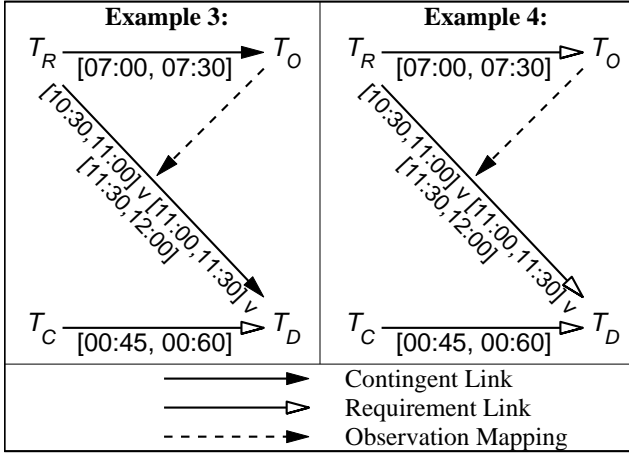


Figure 2: Networks corresponding to Examples 3 and 4.

of the possible reductions, and a dotted arrow is added to signify that the reduction occurs simultaneously with T_O .

One might wonder why we enumerate the set of all allowed partial shrinkages, rather than leave the potential reductions completely unspecified. Although the latter alternative might seem to be more general, we will see in the next section how it precludes the possibility of generalized notions of controllability.

5 Partial Shrinkage of Requirement Links

The generalization in the STPU⁺ is limited to the contingent link, an interval whose endpoint is an uncontrollable event. As we demonstrate in the following example, this expressive power can be further extended to express uncertainty over requirement links that constrain pairs of purely controllable time points.

Example 4: Consider another variation of Example 3. This time, ACME tells *you* to call them sometime between 7:00 and 7:30. When you do, they’ll tell you which half-hour block works best for them, but then allow *you* to specify the exact time in that block for delivery.³ □

Note that in this final variation (depicted in Figure 2 right), all events involved – including the phone call (T_O) and the delivery time (T_D) – are both *controllable*. Yet there is still uncertainty that can only be resolved dynamically, i.e., during the execution of the plan. In particular, you do not know what constraint on delivery time must be satisfied prior to the phone call, even though the exact time is ultimately up to you. Further, the nature of this uncertainty interacts directly with the level of controllability that can be achieved. For instance, our current example is dynamically controllable even with half-hour intervals, as we can decide T_D and thus T_C once the delivery interval becomes known. However, if instead our phone call had to be placed at 10:20, we would

³If this example seems contrived, imagine a situation in which a friend says that he wants to get together this weekend, but isn’t sure yet whether he has to work on Saturday or Sunday. Once he finds out, he calls you, and lets you select the time to meet on his day off.

have a problem if ACME selected the interval [10:30, 11:00]: even though we can choose to schedule the delivery as late as 11:00, we now only have 40 minutes to clean out the fridge (whereas we need 45).

We can now see the advantage of explicitly specifying the set of allowed reductions. In the current example, because we know up front that the possible resolutions of the $T_R \rightarrow T_D$ link are each a half-hour in length, we are able to determine that the problem is dynamically controllable. If, instead, any arbitrary reduction were possible, we could never conclude this: indeed, in the worst case, the “reduction” might not result in any shrinkage at all.⁴

6 The Generalized STPU

We define the *Generalized STPU* as a tuple $\langle X_C, X_U, E, C, O \rangle$, where:

- X_C and X_U are as in the STPU (and X is their union).
- E is a set of *generalized requirement links*, where each E_{ij} is of the form:
$$x_i - x_j \in \{d_{ij1} : [l_{ij1}, u_{ij1}], \dots, d_{ijn_{ij}} : [l_{ijn_{ij}}, u_{ijn_{ij}}]\}$$
- C is a set of *generalized contingent links*, where each C_{ij} is of the form:
$$x_i - x_j \in \{d_{ij1} : [l_{ij1}, u_{ij1}], \dots, d_{ijn_{ij}} : [l_{ijn_{ij}}, u_{ijn_{ij}}]\}$$
and $x_i \in X_U$.
- $O : X_O \rightarrow E \cup C$ is a mapping from observation events to generalized links, where $X_O \subseteq X$.

In the Generalized STPU, we augment both requirement links and contingent links in an identical fashion, expanding them from single intervals to sets of intervals (the reductions), and associating the moment at which a single reduction is selected with observation events. The definitions of *projection* and *schedule* can remain unchanged from the STPU. In addition to these, we introduce a new notion, an *observation trace* v , which is a record of the generalized link reductions dynamically selected (by Nature):

$$\langle L_1, d_{1k_1} \rangle, \dots, \langle L_n, d_{nk_n} \rangle$$

where $L = E \cup C$ is the set of all generalized links. We also need to extend the notion of a *prehistory*: $T \prec x$ (the prehistory of x with respect to schedule T) now specifies both the durations of all contingent links that finish prior to x , and the selected reduction of all generalized links whose associated observation event occurs before x .

Finally, we define a *generalized execution strategy* S to be a mapping:

$$S : \mathcal{P} \times \mathcal{V} \rightarrow \mathcal{T}$$

where \mathcal{P} and \mathcal{T} are as before, and \mathcal{V} is the set of all possible observation traces. A generalized execution strategy S is *viable* if $S(p, v)$, henceforth written $S_{p,v}$, is consistent with p and v for each mutually consistent pair of a projection p and an observation trace v .

⁴One could consider extensions to our formalism that allow for more flexible specifications of possible reductions (for instance, involving reductions that comprise sets of intervals).

6.1 Generalizing Controllability

With these new definitions in hand, we can now specify the three levels of controllability for the Generalized STPU. A Generalized STPU can be said to be *Weakly Controllable* if there exists a viable generalized execution strategy; this is equivalent to saying that every pair $\langle p, v \rangle$ of a projection p and an observation trace v is consistent.

A Generalized STPU is *Strongly Controllable* if there is a viable generalized execution strategy S such that

$$S_{p1,v1}(x) = S_{p2,v2}(x)$$

for each executable time point x , all projections $p1$ and $p2$, and all observation traces $v1$ and $v2$. As before, this means that a single schedule is guaranteed to satisfy all constraints, regardless of the choices (assignments to uncontrollable events *or* partial shrinkages of generalized links) that Nature makes.⁵

Finally, a Generalized STPU is said to be *Dynamically Controllable* if there is a viable execution strategy S such that:

$$S_{p1,v1}\{\prec x\} = S_{p2,v2}\{\prec x\} \Rightarrow S_{p1,v1}(x) = S_{p2,v2}(x)$$

for each executable time point x , all projections $p1$ and $p2$, and all observation traces $v1$ and $v2$. In other words, there exists a strategy that depends on the outcomes of only those uncontrollable events and partial shrinkages that have occurred in the past (and not on those which have yet to occur).

7 Disjunctions, Generalized Links, and Controllability

As mentioned earlier, there is an apparent similarity between the reductions of a generalized (requirement or contingent) link, and the disjuncts in a constraint of a Binary Temporal CSP [Dechter *et al.*, 1991] or Disjunctive Temporal Problem (DTP) [Stergiou and Koubarakis, 1998]. However, the semantics of the two types of links are in fact quite different from one another.

Recall that the DTP is a generalization of the STP, in which each constraint C_i may consist of a disjunction of simple temporal constraints:

$$l_{i1} \leq x_{i1} - y_{i1} \leq u_{i1} \vee \dots \vee l_{in_i} \leq x_{in_i} - y_{in_i} \leq u_{in_i}$$

A solution to a DTP is an assignment to its time points that satisfies every constraint; of course, since the constraints are disjunctive, what is required is that *some* disjunct from each constraint be satisfied.

In contrast, consider a Generalized STPU; for the moment, to facilitate comparison with the DTP (which has only controllable time points), assume that it is also comprised exclusively of controllables. Because Nature may select *any* reduction (or disjunct) for each constraint, it will be controllable only if there is a set of solutions, at least one of which satisfies each combination of selected reductions. In other

⁵It is worth noting that strong controllability can be easily disproven in cases where there exists a generalized requirement link containing a non-overlapping pair of reductions.

	Agent Picks a Duration	Nature Picks a Duration
Agent Picks a Disjunct	Disjunctive Requirement DTP Link	Disjunctive Contingent DTPU Link
Nature Picks a Disjunct	Generalized Requirement STPU Link	Generalized Contingent STPU Link

Table 1: Classification of “Disjunctive” Links

words, we need to be able to satisfy *each* of the disjuncts in the generalized requirement link.

Further properties of the solution set will define the level of controllability: the Generalized STPU will be strongly controllable if the set of solutions is a singleton (so that the same assignment to controllables “works” regardless of the decisions made by Nature); it will be dynamically controllable if the solutions for sets of selections that are identical prior to time point x make an identical assignment to x ; and it will otherwise be weakly controllable.

These distinctions are summarized in Table 1. Interestingly, the difference in what is needed to satisfy standard disjunctive links (e.g., of the DTP) versus generalized links closely parallels the difference in what is needed to satisfy requirement versus contingent links of either flavor. Given a problem with only requirement links, what is needed for consistency is the existence of a solution that “selects” a particular legal duration for each constraint. However, once contingent links are added, controllability depends upon there being solutions that satisfy *each* possible combination of durations that the contingent links might take.

Note that just as the Generalized STPU with controllable time points serves as a counterpart to the TCSP or DTP, the full Generalized STPU (which allows contingent time points) parallels the recently proposed DTPU [Venable and Yorke-Smith, 2005]. This analysis reveals that there are additional scenarios that can be modeled using this general approach: for example, by introducing observation events and partial shrinkage into DTPU’s, we can model cases in which both Nature and the agent can “pick” reduced intervals from contingent links.

8 Other Related Work

The addition of observation events is similar to the approach taken in Conditional Temporal Problems [Tsamardinos *et al.*, 2003]. However, in that line of research, the observation actions serve to remove events from the execution plan, rather than to specify the resolution of constraints. Furthermore, pure temporal uncertainty is not modeled in the CTP, as all time points are required to be controllable.

Also, there are some relationships between our Generalized STPU and the Uncertain CSP [Yorke-Smith and Gervet, 2003], a formalism that augments a standard CSP $\langle X, D, C \rangle$

with a set of uncertain coefficients Λ and a set of corresponding uncertainty sets \mathcal{U} . Since the presence of disjunctions in Temporal CSPs often results in a meta-CSP reformulation (where the constraints become meta-level variables whose domains are their associated disjuncts), one can view Nature's selection of a disjunct as the assignment to a finite-domain coefficient from its respective uncertainty set. However, the UCSP cannot model the underlying semantics of these values, which in our case are temporal differences, nor can it model scenarios where values are assigned online in a dynamic environment.

9 Conclusion and Future Work

In this paper, we have generalized the problem of controllability in constraint-based temporal reasoning. Beginning with the Simple Temporal Problem with Uncertainty, we have progressively extended the formalism to take into account prior observability and partial shrinkage of both contingent and requirement links. These extensions make it possible to capture more complex sequences of dynamic knowledge acquisition, and even to express some types of uncertainty that arise when the temporal events are themselves fully controllable.

This line of research opens the door to several avenues of continued progress. First of all, our focus so far has been entirely on expressive power, and significant work remains to be done to develop algorithms for effectively computing controllability (especially dynamic controllability) in the Generalized STPU. Second, our current extension requires a strong coupling between a generalized link and a single observation event that selects a specific reduction from within a generalized link; in some cases, it may be useful to allow reductions to be spread out over multiple observation events. Also, we believe that the formal relationship between generalized links in our Generalized STPU and the disjunctive constraints in the DTPU could potentially be exploited in a hybrid formalism.

Acknowledgements

We thank the anonymous reviewers for their helpful and constructive feedback. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. NBCHD030010 and the Air Force Office of Scientific Research under Contract No. FA9550-04-1-0043. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the view of DARPA, the Department of Interior-National Business Center, or the United States Air Force.

References

- [Dechter *et al.*, 1991] Rina Dechter, Itay Meiri, and Judea Pearl. Temporal constraint networks. *Artificial Intelligence*, 49(1-3):61–95, 1991.
- [Morris and Muscettola, 2005] Paul H. Morris and Nicola Muscettola. Temporal dynamic controllability revisited. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-2005)*, pages 1193–1198, 2005.
- [Morris *et al.*, 2001] Paul H. Morris, Nicola Muscettola, and Thierry Vidal. Dynamic control of plans with temporal uncertainty. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-2001)*, pages 494–502, 2001.
- [Stergiou and Koubarakis, 1998] Kostas Stergiou and Manolis Koubarakis. Backtracking algorithms for disjunctions of temporal constraints. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98)*, pages 248–253, 1998.
- [Tsamardinos *et al.*, 2003] Ioannis Tsamardinos, Thierry Vidal, and Martha E. Pollack. CTP: A new constraint-based formalism for conditional, temporal planning. *Constraints*, 8(4):365–388, 2003.
- [Venable and Yorke-Smith, 2005] Kristen Brent Venable and Neil Yorke-Smith. Disjunctive temporal planning with uncertainty. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-2005)*, pages 1721–1722, 2005.
- [Vidal and Fargier, 1999] Thierry Vidal and Helene Fargier. Handling contingency in temporal constraint networks: From consistency to controllabilities. *Journal of Experimental and Theoretical Artificial Intelligence*, 11(1):23–45, 1999.
- [Vidal and Ghallab, 1996] Thierry Vidal and Malik Ghallab. Dealing with uncertain durations in temporal constraint networks dedicated to planning. In *12th European Conference on Artificial Intelligence (ECAI-1996)*, pages 48–54, 1996.
- [Vidal, 2000] Thierry Vidal. Controllability characterization and checking in contingent temporal constraint networks. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR 2000)*, pages 559–570, 2000.
- [Yorke-Smith and Gervet, 2003] Neil Yorke-Smith and Carmen Gervet. Certainty closure: A framework for reliable constraint reasoning with uncertainty. In *Proceedings 9th International Conference on Principles and Practice of Constraint Programming (CP 2003)*, pages 769–783, 2003.