

Quantifying Degrees of Controllability in Temporal Networks with Uncertainty

Shyan Akmal* and Savana Ammons* and Hemeng Li* and James C. Boerkoel Jr.

Human Experience and Agent Teamwork Lab (heatlab.org)

Harvey Mudd College

Claremont, California 91711

{sakmal, sammons, mli, boerkoel}@hmc.edu

Abstract

Controllability for Simple Temporal Networks with Uncertainty (STNUs) has thus far been limited to three levels: strong, dynamic, and weak. Because of this, there is currently no systematic way for an agent to assess just how far from being controllable an uncontrollable STNU is. We use a new geometric interpretation of STNUs to introduce the degrees of strong and dynamic controllability – continuous metrics that measure how far a network is from being controllable. We utilize these metrics to approximate the probabilities that an STNU can be dispatched successfully offline and online respectively. We introduce new methods for predicting the degrees of strong and dynamic controllability for uncontrollable networks. In addition, we show empirically that both metrics are good predictors of the actual dispatch success rate.

Introduction

When tasked with making a schedule, a planner would ideally have the capability to pick exact times for every event that could occur. In practice however, autonomous agents rarely have control over all events affecting their plans.

For example, imagine a chemist named Dr. V running a small experiment. She first combines 20 mL of chemical W and chemical X in a beaker. The exact amount of time it takes these chemicals to react is uncertain: all Dr. V knows is that it will take between twenty and thirty-one minutes for the reaction to finish (unfortunately, not much is known about the reaction rates of chemicals W and X). Within ten minutes of this reaction completing, she must add 20 mL of chemical Y to the mixture, which catalyzes a new reaction taking thirty to thirty-five minutes. Finally, Dr. V must collect a product, precipitate Z, from the solution within ten minutes of the reaction completing.

To run the experiment successfully, Dr. V needs to schedule several different events. She is able to select times for some events, such as the addition of chemical Y, but other events, such as the completion times of the reactions, are not under her control. There are multiple different strategies Dr. V can use to deal with this uncertainty in timing and ensure the experiment’s success. The effectiveness of different types of scheduling strategies is related to the scheduling

problem’s *controllability* – how easily the experiment can be successfully executed given its inherent uncertainty. Currently, all forms of controllability discussed in the literature are discrete – they tell Dr. V whether her experiment is controllable or not, but do not tell her “how” controllable or uncontrollable the experiment is.

In this paper, we extend the notion of controllability by introducing two new continuous metrics on temporal networks: the degrees of strong and dynamic controllability. These metrics are motivated by a geometric interpretation of STNUs as pairs of polytopes and characterize how far a network is from being controllable given different types of execution strategies. These metrics naturally relate to the probability that a network can be successfully dispatched. In that context, we define optimization problems for determining the probability of successful dispatch given online/offline plans, and offer approximate solutions to these problems.

Background

Simple Temporal Networks

A *Simple Temporal Network* (STN) is a tuple $S = \langle T, C \rangle$, where T is the set of temporal events t_i , and C is the set of binary constraints on T . Each element in C is of the form $t_j - t_i \leq c_{ij}$, for some $c_{ij} \in \mathbb{R}$ (Dechter, Meiri, and Pearl 1991). By convention, the event $t_0 \in T$ in an STN always represents a fixed reference point assigned time zero. Because of this, when we say that an STN has n temporal events, we really mean that it has n time-points in addition to this reference event. A solution to the STN is an assignment of values to the timepoints t_i that satisfies all constraints.

Networks with Uncertainty

A *Simple Temporal Network with Uncertainty* (STNU) is an STN that explicitly models uncertainty. In an STNU, the set of events T is partitioned into a set of controllable events T^c and a set of uncontrollable events T^u . An agent is allowed to schedule specific times for the events in T^c , but the times for events in T^u are determined by “Nature,” a force external to the agent. For every $t_j \in T^u$, there exists a unique $t_i \in T^c$ forming a *contingent constraint* of the form $t_j - t_i \in [\ell_j, u_j]$. We use C_c to denote the set of these contingent constraints, and C_r to denote the remaining *requirement* constraints between events in T . Then an STNU is defined as a quadruple

*Primary authors listed alphabetically but contributed equally. Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: An STNU representation of Dr. V's experiment.

$\langle T^c, T^u, C_r, C_c \rangle$ satisfying the aforementioned properties.

In an STNU, a *decision* is an assignment of time values to each of the controllable events. A *realization* is a selection of values for contingent edges (relative start times of the uncontrollable events) assigned by Nature. The set of all possible realizations forms the *realization space* Ω for the STNU. If a decision does not violate any of the constraints between pairs of controllable events, we say the decision is *admissible*. A decision together with a realization determines a schedule σ for the network. If σ satisfies all constraints in the STNU, we call σ a *valid* schedule.

For example, Dr. V's experiment from the introduction can be modeled as an STNU. There are five events of interest: the beginning of the first reaction, the conclusion of this initial reaction, the addition of chemical Y, the completion of the subsequent reaction, and the extraction of the precipitate. If we measure time relative to the experiment's start, we can refer to these events as t_i , where i ranges from 0 to 4 respectively. Events t_0, t_2 and t_4 are under Dr. V's control, while events t_1 and t_3 are uncontrollable. This STNU representation of the experiment is depicted in Figure 1. Events in the network are drawn as nodes. Requirement and contingent edges are drawn as straight and curvy arrows respectively. The edges are labeled by intervals indicating the lower and upper bounds of the STNU's constraints.

Controllability

An STNU is controllable when an agent has a reasonable way of working around the uncertainty in the network to schedule events. Prior research has focused primarily on detecting three types of controllability: strong, dynamic, and weak. We focus on strong and dynamic controllability only, since they are important for dispatch. We omit weak controllability from our discussion because it is less relevant to STNU execution (Vidal and Fargier 1999). Henceforth in this paper, we use the term "uncontrollable STNU" to refer to *both* STNUs that are not strongly controllable and those that are not dynamically controllable.

An STNU is *strongly controllable* if there is a single fixed decision the agent can make that guarantees success regardless of how Nature behaves. More formally, a network is strongly controllable if there exists a decision δ such that for all $\omega \in \Omega$, the schedule determined by δ and ω is valid (Vidal and Fargier 1999). We call such a decision a *strong decision*. Using the example from the introduction, one can check that Dr. V's experiment is not strongly controllable

because there is no fixed decision she could make to guarantee the experiment's success before it takes place.

An STNU is *dynamically controllable* if an agent can always make decisions in real-time to guarantee success. This means that during execution, if the agent picks times for controllable events while observing the values taken by all *past* events, then there exists a strategy it can follow to find a valid schedule. A formal definition is given by Hunsberger (2009). For example, Dr. V's experiment is dynamically controllable, because she could wait until each reaction finishes and then proceed to the next step immediately. This corresponds to scheduling $t_2 = t_1$ and $t_4 = t_3$, and involves "reacting" dynamically to previously observed timepoints. This strategy guarantees the experiment's success.

Visualizing Controllability

An STN with n temporal events determines a region in \mathbb{R}^n as follows. Each axis corresponds to the value of a particular event, and every constraint in the STN is a linear inequality, which consequently defines a half-space. The solution space of the STN is the intersection of these half-spaces, and therefore a convex polytope. If the STN has a finite *makespan* – time during which all its temporal events must occur – then this polytope is bounded. Points within the polytope correspond to valid schedules of the STN. Previous research has leveraged this geometric perspective to determine ways of assessing a network's flexibility (Huang et al. 2018) and detect weak controllability (Cimatti, Micheli, and Roveri 2015). Motivated by this, we characterize STNUs as geometric objects to inform our ideas on controllability.

Geometry of STNUs

An STNU with n controllable events and m uncontrollable events can be represented geometrically by a pair of regions \mathcal{P} and \mathcal{P}' in \mathbb{R}^{n+m} , where $\mathcal{P}' \subseteq \mathcal{P}$. The larger region \mathcal{P} is the set of all schedules (including those outside the solution space) that can arise from an admissible decision. These are the situations the agent could find itself in. The smaller region \mathcal{P}' is the set of all valid schedules: the situations the agent desires to end up in.




Figure 2: The polytopes associated with network S .

Comparing the relative locations of \mathcal{P} and \mathcal{P}' (in particular, the amount of \mathcal{P} that “sticks out” of \mathcal{P}') gives a sense of how difficult it is to control for the uncertainty present in the STNU. This geometric perspective motivates our consideration of *volume* as a way to measure uncertainty in scheduling problems. Since \mathcal{P}' is the shape we get if we treat the STNU as if it were just an STN, by viewing each contingent edge as if it were a requirement edge, \mathcal{P}' is a convex polytope. Similarly, \mathcal{P} is the region formed by the inequalities if we ignore all requirement constraints involving an uncontrollable event, so it is also a convex polytope. As an example, consider a network \mathcal{S} with one uncontrollable event t_3 and two controllable events t_1 and t_2 . This network has requirement constraints $t_1 - t_0 \in [0, 8]$, $t_2 - t_0 \in [0, 12]$, $t_3 - t_0 \in [0, 16]$, and $t_1 - t_2 \in (-\infty, 8]$, as well as one contingent constraint $t_3 - t_2 \in [0, 6]$. Using the method we just described, \mathcal{S} determines two polytopes, shown in Figure 2. The entire polytope is \mathcal{P} , while the sub-region in the lower left is \mathcal{P}' .

Geometry of Uncertainty & Controllability

To supplement this geometric view, we associate polytopes independently to the set of uncontrollable and controllable events in the network. The regions for uncontrollables and controllables provide information on the uncertainty and controllability of the network respectively. Each uncontrollable timepoint depends on a unique contingent edge, and the possible values each contingent edge can take on form the realization space Ω defined earlier. If we define \mathcal{U} to be the set of contingent intervals in the STNU, it follows that Ω is just the Cartesian product of the intervals in \mathcal{U} . Geometrically then, Ω is an m -dimensional hyperrectangle.

The set of all strong decisions in an STNU can be determined in polynomial time (Vidal and Fargier 1999). This region is determined by tightening each constraint in the STNU by considering the “worst-case” values for the contingent edges. For example, if τ is a controllable event and $t' - t \in [l_t, u_t]$ is a contingent edge, then the original constraint $\tau - t' \leq c$, involving the uncontrollable t' , would be replaced with $\tau - t \leq c + l_t$, a constraint between two controllables. The area determined by these tightened inequalities is the *strongly controllable region* of the network. This region is the set of *strong* decisions, and is a polytope in \mathbb{R}^n since it is described completely by inequalities on the controllable events. The original STNU is strongly controllable exactly when the strongly controllable region is nonempty.

Beyond Discrete Categories of Controllability

Currently, controllability is categorical—either a network is strongly/dynamically controllable or it is not. This limited categorization does not always yield sufficient information for applications. Ideally, we would like to know the likelihood of success using offline dispatch (in strong controllability) or online dispatch (in dynamic controllability) methods for a given temporal network. In practice, the problem of finding dispatch techniques that maximize probability of success is infeasible, so we focus on finding good approximations. We first investigate this idea in the context of strong controllability. Our “degree of controllability” characterization provides a new way to understand and evaluate existing

dispatch and approximate controllability approaches, particularly in the context of STNUs. Moreover, our geometric framing offers new insights that allow us to improve previous approaches for working with STNUs and provides a more rigorous definition of what it means to be “maximally strongly controllable.”

Degree of Strong Controllability

In a strongly controllable STNU, execution is simple. The agent can guarantee success simply by finding and committing to a single strong decision. This makes strong controllability a highly desirable property for STNUs. When a network is not strongly controllable, the traditional approach has been to check for dynamic controllability and, if this check is successful, seek out a dynamic execution strategy for dispatch. Any such strategy necessarily keeps track of when uncontrollable events occur, and is thus more complicated than pre-committing to one particular decision. Applications that require offline planning or low-overhead are not amenable to these more involved strategies. In these cases, it is useful to understand *how* controllable a network is, because it may be worthwhile to pre-commit to a single decision if the agent knows this will yield successful execution *most*, even if not all, of the time.

For instance, we observed earlier that Dr. V could ensure her experiment runs successfully by scheduling t_2 and t_4 as soon as the times for t_1 and t_3 are realized. This plan requires Dr. V to stay in the lab and monitor all reactions, so she can act as soon as they terminate. Suppose however, that Dr. V is busy the morning of the experiment. She would prefer to pick fixed times beforehand to schedule t_2 and t_4 , so she could drop by the lab at only those specific times, and attend to other tasks during the rest of the experiment. As noted earlier, the experiment is not strongly controllable, so there is no such fixed strategy Dr. V could employ to guarantee success. However, if Dr. V sets $t_2 = 30$ and $t_4 = 65$, it turns out that the experiment has greater than a 90% chance of success (this admissible decision guarantees success if the first reaction does not take over thirty minutes). This probability might be high enough that, even though success is not guaranteed, Dr. V will employ this simpler offline strategy.

Definition. Let S be an STNU with realization space Ω . For a given subset Ω' of Ω , we let S' denote the network with the same required constraints as S , but realization space Ω' . Then the **Degree of Strong Controllability (DSC)** for S is a continuous, 0-dimensional metric (i.e., value in $[0, 1]$) defined as the maximum possible value of

$$\frac{\text{Vol}(\Omega')}{\text{Vol}(\Omega)}$$

taken over all axis-parallel hyperrectangles $\Omega' \subseteq \Omega$ such that S' is strongly controllable. Here, $\text{Vol}(\cdot)$ measures the volume of a region.

If we interpret contingent constraints as representing uniform uncertainty, the DSC is intuitively the probability that S “ends up” being controllable. In this way, the degree measures the maximum proportion of the realization space that an execution strategy could account for. Networks with

lower degrees are “less controllable”, while networks with degrees closer to 1 have a fixed decision that is likely to yield successful execution. For example, a strongly controllable STNU has DSC equal to 1. Because Ω' is a hyperrectangle, it is the Cartesian product of m intervals, each of which is necessarily a subinterval of one of S 's contingent intervals (otherwise Ω' would not be a subset of Ω). We denote the collection of these intervals by \mathcal{U}' , and observe that this collection uniquely determines Ω' (and vice-versa).

The DSC of an STNU is a lower bound on the maximum probability of obtaining a valid schedule for the given network by committing to an admissible decision before any events have executed. This is because there always exists a decision that can be paired with every realization in Ω' (which is a subset of Ω) to create a valid schedule. So, by comparing Ω' 's volume to Ω 's volume, we are calculating the *proportion* of time that a random realization yields a valid schedule when paired with an admissible decision.

Estimating DSC

To compute the DSC, it suffices to find an Ω' of maximum volume such that S' is strongly controllable. The strong controllability of S' can be verified by checking whether a set of linear constraints are satisfied (Vidal and Fargier 1999). However, solving this general problem is difficult, since maximizing a nonlinear polynomial over linear constraints is NP-hard (Motzkin and Straus 1965). Thus we propose our approximation method for finding the optimal Ω' .

Degree of Strong Controllability Linear Program (DSC-LP): Our DSC-LP is adapted from the maximum subinterval problem discussed by Wilson et al. (2014). The goal of DSC-LP is to find a set of contingent subintervals \mathcal{U}' that approximately maximize the volume of Ω' and a fixed decision that ensures success in the associated STNU S' . Given an STNU, $S = \langle T^c, T^u, C_r, C_c \rangle$, a set of maximal subintervals for S can be obtained by solving the following LP:

$$\begin{aligned}
& \text{minimize: } \sum_{t \in T^u} \frac{\epsilon_t^- + \epsilon_t^+}{u_t - l_t} \\
& \text{subject to: } t^- \leq t^+ && \forall t \in T^u \quad (1) \\
& t^- = t^+ && \forall t \in T^c \quad (2) \\
& t^+ - t'^- \leq c && \forall (t - t' \leq c) \in C_r \quad (3) \\
& t^- - t'^- = l_t + \epsilon_t^- && \forall (t - t' \geq l_t) \in C_r \quad (4) \\
& t^+ - t'^+ = u_t - \epsilon_t^+ && \forall (t - t' \leq u_t) \in C_r \quad (5) \\
& \epsilon_t^-, \epsilon_t^+ \geq 0 && \forall t \in T^u \quad (6) \\
& t_0^- = t_0^+ = 0 && \quad (7)
\end{aligned}$$

Like Wilson et al., we introduce two variables t^- and t^+ for each event t , representing the lower bound and upper bound of the time interval in which event t occurs and constrain these variables so that they respect all requirement constraints (lines 1-3,7). Our LP differs from (Wilson et al. 2014) in three key ways. First, it allows us to modify the upper and lower bounds of contingent intervals by different amounts (lines 4-6). We do this by introducing two variables

ϵ_t^- and ϵ_t^+ for every contingent constraint $t - t' \in [l_t, u_t]$, which represent the amount the lower and upper bounds of the original contingent interval are tightened by to guarantee strong controllability. The resulting contingent subinterval becomes $t - t' \in [l_t + \epsilon_t^-, u_t - \epsilon_t^+]$. Second, our LP sets $t^- = t^+$ for all $t \in T^c$. By doing this, we are guaranteed to have a specific schedule returned by the LP. The third and final way our LP differs from the original is that it uses a new objective function, which sums the amount each contingent interval is decreased by, normalized by its original length. Let $l_t = u_t - l_t$ denote the uncontrollable timepoint t 's contingent interval length, and let $\epsilon_t = \epsilon_t^+ + \epsilon_t^-$ be the amount that interval is shrunk by. Then, the objective function is equal to

$$\sum_{t \in T^u} \frac{\epsilon_t}{l_t}.$$

This choice is motivated by considering the amount of volume lost

$$f(\epsilon) = \prod_{t \in T^u} l_t - \prod_{t \in T^u} (l_t - \epsilon_t)$$

in going from Ω to Ω' as a function of the ϵ_t variables, and setting the objective function equal to a constant multiple of the gradient ∇f evaluated at $\epsilon = \vec{0}$. Thus, minimizing the objective function corresponds to minimizing the first order approximation of the amount of volume lost by shrinking the contingent intervals.

Since Ω and Ω' are the Cartesian products of elements in \mathcal{U} and \mathcal{U}' in this case, the volumes of these regions are just the product of lengths of the intervals in \mathcal{U} and \mathcal{U}' respectively. Our LP yields a choice of \mathcal{U}' , so the solution to our LP provides an approximate value for DSC. Moreover, the values taken on by t^+ variables, for $t \in T^c$, in the optimal solution of DSC-LP determine a decision for the network that is a strong decision for the S' . We call a decision returned by our LP in this way a **strong-LP decision** for the network. If a network has high DSC, the strong-LP decision is likely a good strategy for this network.

There are other natural choices of LPs one could use to approximate DSC. For example, we could minimize the total amount of uncertainty removed (as Wilson et al. (2014) does), minimize the maximum amount of uncertainty removed from a single contingent interval, or maximize the minimum length over all contingent subintervals. An empirical comparison of these methods, not reproduced here, shows that the DSC-LP provides better approximations than these other approaches. In particular, the use of volume as probability, motivated by the relationship between Ω and \mathcal{P} and \mathcal{P}' , seems more effective than looser heuristics and metrics like flexibility for estimating robustness.

Related Approaches

The DSC is related to previous work done on dispatch rates of *Probabilistic STNs (PSTNs)* (Tsamardinos 2002; Brooks et al. 2015). PSTNs augment the STNU framework by providing probability distributions over the uncertain duration of contingent edges, which allows evaluating approaches in terms of *risk*, the probability a dispatch strategy

might fail to successfully execute, or *robustness*, the probability of execution success. Indeed, there are many existing approximate approaches for minimizing risk, all of which appeal to the idea of reducing the problem to a strongly controllable STNU that maximizes likelihood of successful dispatch (Fang, Yu, and Williams 2014; Santana et al. 2016; Lund et al. 2017). Thus, our DSC can be considered as the ‘robustness’ of a special case of PSTNs, where all contingent constraints have a uniform distribution with a fixed interval. However, our work is the first to characterize and empirically compare against (see the Empirical Evaluation section) a theoretically optimal solution to this problem.

Degree of Dynamic Controllability

Although strong controllability is valuable because it allows an agent to rely on a static schedule, dynamic controllability applies to a wider range of STNUs (i.e., all strongly controllable networks are dynamically controllable, but not vice-versa). In cases where an STNU is not dynamically controllable, an agent may still wish to attempt to dispatch the schedule. For, even when success is not guaranteed, if the probability of successful execution is high enough, dispatching on an uncontrollable network may be worthwhile. To assess how safe it is to attempt dispatch, the agent will need some measure of *how* dynamically controllable the network is. This motivates the following definition.

Definition. Let S be an STNU with realization space Ω . For a given subset Ω' of Ω , we let S' denote the network with the same required constraints as S , but realization space Ω' . Then the **Degree of Dynamic Controllability (DDC)** for S is defined to be the maximum possible value of

$$\frac{\text{Vol}(\Omega')}{\text{Vol}(\Omega)}$$

taken over *all* $\Omega' \subseteq \Omega$ such that S' is dynamically controllable and where $\text{Vol}(\cdot)$ measures the volume of a region.

The key distinction between the DDC and the DSC is that Ω' is allowed to be any measurable subset of Ω . In particular, Ω' is not constrained to be an axis-parallel hyperrectangle. This means that the network S' is not necessarily an STNU (which always has a hyperrectangular realization space) in the traditional sense, and instead is more accurately described as an STN together with contingent edges and a realization space. The values taken on by the contingent edges of such a network are determined by randomly sampling from the space Ω' (so that in this more general setting, the values taken by the contingent edges are not necessarily independent). This notion can be formalized appropriately. However, because it would lead us too far astray from the main goals and results of this paper, we avoid describing a complete definition of these “generalized STNUs.”

We first show how we can lower bound the DDC by solving a particular type of temporal relaxation problem. This is done using Algorithm 1, which optimally relaxes a network’s constraints until it is dynamically controllable. From here, we compute Ω' ’s volume simply by multiplying the weights of the relaxed contingent intervals, as we did when computing the DSC. This is equivalent to approximating the

volume of the optimal Ω' using an inscribed hyperrectangle. In general this method underestimates the true maximum volume, so in the second subsection we improve this predicted probability via a normal approximation technique, and then conclude by discussing related approaches.

Optimal Dynamic Controllability Relaxations

Dynamic controllability of STNUs can be checked in polynomial time (Morris and Muscettola 2005). The fastest current algorithms for determining DC search for a *conflict*: a series of constraints that, when taken together, cannot simultaneously be satisfied by any dynamic execution strategy. Previous work characterized conflicts as a special type of negative cycle occurring in the labeled distance graph of an STNU (Morris 2006).

Bhargava, Vaquero, and Williams (2017) built off this work and introduced a cubic algorithm to extract conflicts from networks. They identified conflicts with the aim of solving *temporal relaxation problems*. A relaxation problem involves taking an uncontrollable network and determining how to relax the network’s constraints in a minimal fashion to produce a new, controllable network.

In our context, where relaxing requirement constraints is not permitted, an STNU conflict consists of a set of contingent edges, together with a *resolution constant* κ that indicates the total amount the contingent edges must be shrunk by to resolve the conflict (the latter can be obtained from an identified conflict). We now present a solution (presented as Algorithm 1) to our version of the relaxation problem, where only contingents can be relaxed (by shrinking their intervals). If we have an STNU with only one single conflict, our solution is optimal in the sense that it shrinks intervals to minimize the amount of volume of Ω lost while still resolving the conflict.

Suppose the conflict consists of intervals I_1, I_2, \dots, I_p with respective lengths $\ell_1 \leq \ell_2 \leq \dots \leq \ell_p$. We want to shrink each interval I_j to an interval of length ℓ'_j such that the sum of the ℓ'_j s is less than or equal to the quantity

$$\ell' = \left(\sum_{j=1}^p \ell_j \right) - \kappa$$

since this is exactly what it means for the conflict to be resolved. Thus, resolving the conflict while maximizing the volume of the reduced realization space is equivalent to finding values $\ell'_j \in [0, \ell_j]$ with $\sum_{j=1}^p \ell'_j = \ell'$ that maximize the product $\prod_{j=1}^p \ell'_j$. This optimization problem has an analytic solution, which we now describe. Let q be the smallest index such that

$$\ell' \leq \ell_1 + \ell_2 + \dots + \ell_{q-1} + (p - q + 1)\ell_q. \quad (8)$$

Then taking $\ell'_j = \ell_j$ for $j < q$ and

$$\ell'_j = \frac{\ell' - \ell_1 - \dots - \ell_{q-1}}{p - q + 1}$$

for $j \geq q$ solves the maximization problem. In this solution, the largest contingent intervals are all shrunk to the same

size, while smaller intervals are left alone. Intuitively, the solution reduces the realization space to make it look more like a hypercube.

Algorithm 1: Optimal Relax (ODC-RELAX) Strategy

Input : $conflicts$, a sorted list of contingent intervals in the detected STNU conflict
Input : κ , resolution constant
Var : p , length of $conflicts$
Var : ℓ , a list in which $\ell[j]$ is the length of the contingent interval, $conflicts[j]$

Initialization :

- 1 $relaxations \leftarrow \{\}$;
- 2 $S \leftarrow sum(\ell[1..p]) - \kappa$;
- 3 $q \leftarrow \text{None}$;

OPTIMALRELAX

- 4 **for** $j \leftarrow 1$ **to** p **do**
- 5 $s \leftarrow sum(\ell[1..j]) + (p - j + 1) * \ell[j]$;
- 6 **if** $s \geq S$ **then**
- 7 $q = j$;
- 8 **break**;

- 9 $A \leftarrow (S - sum(\ell[1..q])) / (p - q + 1)$;
- 10 **for** $j \leftarrow 1$ **to** p **do**
- 11 $edge \leftarrow conflicts[j]$;
- 12 $relaxations[edge] \leftarrow j < q? \ell[j] : A$;
- 13 **return** $relaxations$;

We prove optimality by induction on q . When $q = 1$, all of the ℓ_j are less than or equal to ℓ'/p by definition of q . Now, the AM-GM inequality asserts that for a list of nonnegative reals with a fixed sum, the product is maximized when all the numbers are equal. So, setting $\ell'_j = \ell'/p$ satisfies the constraints and maximizes the product in this scenario.

Now take some fixed $q > 1$, and suppose the result is correct for all integers less than q . We claim there exists an optimal choice of the ℓ'_j with $\ell'_1 = \ell_1$. Suppose to the contrary this were not true. Then, take an optimal selection of values $\ell'_1, \ell'_2, \dots, \ell'_p$ (this exists because the domain of the ℓ_j s is compact). If $\ell'_1 < \ell_1$ in this assignment, then (8) implies that there must exist an index k with $\ell'_k \geq \ell_1$. It follows that there exists a small positive ϵ for which replacing ℓ'_1 with $\ell'_1 + \epsilon$ and ℓ'_k with $\ell'_k - \epsilon$ preserves the sum of the ℓ'_j while increasing their product. This contradicts the optimality of our initial choice, so $\ell'_1 = \ell_1$ as claimed. Now the problem the problem reduces to maximizing $\prod_{j=2}^p \ell'_j$ subject to the constraint that $\sum_{j=2}^p \ell'_j = \ell' - \ell_1$. Finally, using the inductive hypothesis proves that our claimed solution is optimal.

Optimality is guaranteed only for networks with a single conflict. In the case of multiple conflicts our algorithm is unlikely to yield a globally optimal solution. However, our “greedy” approach of minimizing volume lost at each step remains a useful heuristic for resolving conflicts while minimizing probability mass lost through relaxation.

Estimating DDC

Although solving relaxation problems is useful, we are also interested in accurately computing the probability of successful dispatch with a dynamic execution strategy. Previous research has attempted to measure this likelihood (called “robustness” in the literature) by relaxing to strongly/dynamically controllable networks. As noted in (Cui et al. 2015) however, the volume of the realization space in a relaxed network does not accurately estimate the relevant probability, but instead bounds it below. The relaxation approach underestimates the probability too severely to compute DDC.




Figure 3: The realization space for S' . The inner box, the realization space of the relaxed network, underestimates the volume of realizations that lead to valid schedules.

For example, consider an STNU S' with two uncontrollables t_1 and t_3 , and one controllable event t_2 . This network has requirement constraints $t_1 - t_2 \leq 0$ and $t_3 - t_0 \leq 3$, along with contingent constraints $t_1 - t_0 \in [0, 2]$ and $t_3 - t_2 \in [0, 2]$. Network S' is not dynamically controllable, but by inspection we see that the optimal execution strategy schedules $t_2 = t_1$. In this case, if $a_1 = t_1 - t_0$ and $a_2 = t_3 - t_0$ denote the realized lengths of the contingent edges, then as depicted in Figure 3 our schedule succeeds exactly when $a_1 + a_2 \leq 3$, which happens 87.5% of the time. In contrast, if we wanted to relax S' to get a dynamically controllable network while maximizing volume, both intervals for the contingent edges would shrink from $[0, 2]$ to $[0, 3/2]$. This shrunk realization space takes up 56.25% of the original volume, which is significantly smaller than the true probability of succeeding during dispatch. As a side note, this effect tends to be even more prominent in higher dimensions – geometrically, this is because inscribed hyperrectangles are bad at approximating volumes of simplices.

Motivated by the above example, we leverage the notion of conflicts to directly approximate the probability of successful dispatch using the **normal approximation to DDC**. Suppose we have an STNU with a single conflict preventing it from being dynamically controllable. Carry over the notation from the previous section and additionally let a_j denote the difference between the realized value during dispatch and the lower bound of the contingent interval I_j for

$j = 1, 2, \dots, p$. Then if the inequality

$$\sum_{i=1}^p a_i \leq \ell'$$

holds, the conflict is “avoided” in the sense that the observed realization could have come from a controllable STNU. This idea, that dynamic execution strategies succeed if conflicts are avoided during dispatch, is the crucial insight that allows us to estimate the DDC. We may view each a_j as an independent random variable drawn from the uniform distribution on $[0, \ell_j]$. Then the Central Limit Theorem implies that the sum of the a_j (the left hand side of the above inequality) can be approximated as coming from a normal distribution $\mathcal{N}(\mu, \sigma^2)$, where $\mu = (\ell_1 + \dots + \ell_p) / 2$ and $\sigma^2 = (\ell_1^2 + \dots + \ell_p^2) / 12$ are the sums of the means and variances of the a_j s respectively. Thus the probability the above inequality holds can be computed using the standard normal CDF Φ . For uncontrollable STNUs with a single conflict, the normal approximation accurately estimates the probability of successful dispatch. In the case of multiple conflicts we use the normal approximation to estimate the probabilities of avoiding each conflict separately, and then multiply the obtained probabilities. Taking this product implicitly treats each conflict as if it were independent of the others (no sharing contingent edges among conflicts). Because conflicts are not generally independent, this technique should behave worse in networks with multiple conflicts. However, as detailed in our empirical evaluation, the normal approximation is still useful in these cases.

Related Approaches

The ODC-RELAX algorithm builds on top of work done by Cui et al. (2015). Those authors use a mixed integer LP to solve a broad range of relaxation problems related to dynamic controllability. Our approach tackles a specific variant of their general formulation, but in this special case obtains an analytic solution that is provably, locally optimal and can be computed in $O(p \log p)$ time for an STN with p conflicts. Moreover, ODC-RELAX can naturally be used as a subroutine in the conflict-directed search methods in (Yu, Fang, and Williams 2014) and the RELAXSEARCH method of Bhargava, Vaquero, and Williams (2017) to augment existing algorithms.

Although work done by (Yu, Fang, and Williams 2015) also evaluates the probability of successful dispatch in a dynamic controllability setting, to the best of our knowledge, our normal approximation for DDC is the first known approach for estimating the rate of online dispatch without appealing to relaxations and nonlinear solvers. Because our definition of DDC does not constrain the shape of Ω' , it effectively captures the notion of robustness geometrically. Note that in our framework, we assume agents in networks are mandated to satisfy all requirement constraints, so that violating any one constraint is just as bad as violating all constraints simultaneously. This is consistent with the methodology of Yu, Fang, and Williams (2015), but contrasts with previous work (Cui et al. 2015; Yu, Fang, and Williams 2014; Bhargava, Vaquero, and Williams 2017) that

allows both requirement *and* contingent constraints to be relaxed, and research (Rossi, Venable, and Yorke-Smith 2003) that maximizes the number of constraints satisfied under preference schemes.

Empirical Evaluation

To evaluate our approaches¹ for assessing DSC and DDC, we ran experiments on STNUs derived from the PSTNs provided in the publicly available *ROVERS* and *CAR-SHARING* datasets (Santana et al. 2016). Each PSTN in the *ROVERS* dataset was converted to an STNU by replacing contingent edges drawn from $\mathcal{N}(\mu, \sigma^2)$ with intervals $[\mu - 2\sigma, \mu + 2\sigma]$. Uniformly distributed contingent edges were preserved. STNUs derived from the *ROVERS* dataset were used to evaluate DSC. The PSTNs from *CAR-SHARING* were converted to STNUs by preserving the network structure, but varying the lengths of edges to keep the derived STNUs consistent yet not dynamically controllable.

Success Rate of DSC-LP Decision

We examine how well our DSC metric actually measures success rate for the strong-LP decision. For each STNU, we solved the DSC-LP for that network. We took the decision returned by the program and used it as a dispatch strategy. After sampling uniformly from the contingent edges, we picked random realizations 50,000 times for each STNU. Then, we recorded the success rate as the proportion of the time the strong-LP decision, together with a random realization, led to a valid schedule. This will determine if the strong-LP decision produces better-than-average success rates (compared to other decisions) when used during execution. In Figure 4a this success rate is plotted against the approximate DSC value returned by the DSC-LP. The graph displays a clear linear trend ($r = 0.999$), showing that the LP-predicted DSC of an STNU closely matches the rate of successfully executing the network using the strong-LP decision. This also validates the notion that the strong-decision produces better than average execution results when paired with a realization. Next, we discuss how our LP-predicted DSC compares with the true degree of strong controllability.

Optimality of LP-DSC

We now check how accurate our DSC-LP approximation is by comparing it to the true optimization problem of maximizing volume. Selecting subintervals to maximize volume is a nonlinear optimization problem. We solved these problems using Baron optimization software accessed through the NEOS server (Dolan 2001; Tawarmalani and Sahinidis 2005) to compute the true DSC for all problem instances. We compared this value against the approximate degree found using the LP. 4b demonstrates the correlation between our LP-predicted DSC approximation and the solution to the optimization problem. It demonstrates that at low degrees of strong controllability our LP may underestimate the degree, but at values greater than 0.5 the LP approximation is extremely accurate ($r = 0.996$). This result is consistent

¹All code and problem instances available for download from <https://github.com/HEATlab/Prob-in-Ctrl>




Figure 4: Evaluation of approximate DSC and DDC metrics.

with our expectation that the approximation method works best for STNUs with high DSC values. Although DSC-LP is slightly less accurate at lower DSC values, when our approximation method returns a low value for the DSC we expect the true DSC to be low as well. In these cases where the DSC is quite low, committing to a fixed decision is never a good strategy to employ, so the lack of precision in these cases is perhaps not so concerning. We are more interested in networks where we can pre-commit to a decision that yields successful execution most of the time, and 4b shows that the LP does give accurate approximations for those STNUs. We note that DSC-LP seems to perform better on networks that are dynamically controllable than those that are not. It is not clear why this pattern holds, but it might be an artifact of DSC-LP being more suited to STNUs with the structure of networks from *ROVERS* rather than *CAR-SHARING*.

Success Rate of Approximate DDC

Finally, we examine how well our estimate of DDC measures the actual success rate of STNU dispatch. For each STNU, we computed the estimate of DDC using the normal approximation method. To measure the dispatch rate empirically, we used the early-first dispatch strategy discussed in (Nilsson, Kvarnström, and Doherty 2014) on all uncontrollable STNUs derived from the datasets. This strategy is guaranteed to succeed on dynamically controllable networks, and thus was a natural choice for testing a network’s DDC. In each trial, we simulated online dispatch 50,000 times (with realizations chosen randomly). Success rate was recorded as the proportion of the time we obtained a valid schedule. In 4c, success rate is plotted against against the approximate DDC. The graph displays a clear linear trend ($r = 0.952$), indicating that the normal approximation to DDC accurately tracks the actual success rate. The approach also worked well for STNUs that had multiple separate conflicts, marked as darker points in the graph.

For two outliers in the bottom right corner of the plot, the normal approximation seems to drastically overestimate the success rate. In those two instances, after we employed a different dispatch strategy which waited longer before scheduling certain events, the success rates for the two STNUs were within 2% of the normal approximation. So even in these cases, the predicted DDC matches the *maximum* success rate

achievable with dynamic strategies. In the few other cases where the normal approximation method overestimated the empirical success rate, it is possible that the variant of the early first strategy we used was not the optimal choice.

Conclusion

In this paper, we applied a geometric view of STNUs and expanded the notion of controllability to a continuous measure. These new degrees of strong and dynamic controllability assess how far an STNU is from being controllable. We used this definition in the cases of strong and dynamic controllability to produce the Degree of Strong Controllability (DSC) and the Degree of Dynamic Controllability (DDC) metrics, which provide information on the maximum probability of success using certain types of offline and online strategies respectively. In doing so, we found an efficient LP for approximating DSC, presented a locally optimal solution to a variant of the relaxation problem discussed by Bhargava, Vaquero, and Williams (2017), and provided a normal approximation method for estimating DDC. These contributions present a unified, geometric way of tracking the robustness of networks in the context of controllability.

In the future, it would be interesting to define a continuous metric that interpolates between strong, dynamic, and weak controllability. Although these metrics extend the definition of controllability in useful ways, they currently draw no direct connection between strong and dynamic controllability. That is, even though strong controllability implies dynamic controllability, the DSC value does not provide any information about whether a network is dynamically controllable. This could yield an explicit way of classifying states intermediate between dynamic and strong controllability and offer a more refined way of comparing different networks. In addition, the constructions presented in this paper can likely be extended to apply to PSTNs. Currently, our work with the DDC relies on the assumption that all contingent links have uniform probability distributions. However, our normal approximation method is still applicable in cases where uncertainty is drawn from non-uniform distributions, and it would be interesting to modify the approach to work in more general settings. Furthermore, while our LP returns a decision that tends to yield a successful execution for networks with a high DSC, we have no way of generating such an on-

line dispatch strategy for networks with a high DDC. Future work in this subject could focus on developing an algorithm that returns a good decision for a *dynamic* execution strategy. Doing so would greatly improve the planner's ability to ensure successful dispatch. Lastly, our relaxation algorithm and normal approximation for the DDC currently do not handle STNUs with multiple conflicts well, as they overlooks the possibility of conflicts sharing contingent edges. Further work could attempt to figure out how to contend with these dependencies between conflicts in a way that more accurately assesses the probability of successful dispatch.

Acknowledgments

Funding for this work was graciously provided by the National Science Foundation under grant IIS-1651822 and also the Rose Hills Foundation. Thanks to the anonymous reviewers, HMC faculty, staff, and fellow HEATlab members for their support and constructive feedback. We thank Nikhil Bhargava for providing assistance with the implementation of algorithms presented in (Bhargava, Vaquero, and Williams 2017). Finally, we thank Jordan Abrahams, Susan Martonosi, and Mohamed Omar for offering their expertise in temporal networks, optimization, and convex geometry respectively.

References

- Bhargava, N.; Vaquero, T.; and Williams, B. 2017. Faster conflict generation for dynamic controllability. In *Proc. of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*, 4280–4286.
- Brooks, J.; Reed, E.; Gruver, A.; and Boerkoel, J. C. 2015. Robustness in probabilistic temporal planning. In *Proc. of the 29th National Conference on Artificial Intelligence (AAAI-15)*, 3239–3246.
- Cimatti, A.; Micheli, A.; and Roveri, M. 2015. An SMT-based approach to weak controllability for disjunctive temporal problems with uncertainty. *Artificial Intelligence* 224:1–27.
- Cui, J.; Yu, P.; Fang, C.; Haslum, P.; and Williams, B. C. 2015. Optimising bounds in simple temporal networks with uncertainty under dynamic controllability constraints. In *Proc. of the 25th International Conference on Automated Planning and Scheduling (ICAPS-15)*, 52–60.
- Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. In *Artificial Intelligence* 49, 61–95.
- Dolan, E. D. 2001. NEOS server 4.0 administrative guide. *CoRR* cs.DC/0107034.
- Fang, C.; Yu, P.; and Williams, B. C. 2014. Chance-constrained probabilistic simple temporal problems. In *Proc. of the 28th National Conference on Artificial Intelligence (AAAI-16)*, 2264–2270.
- Huang, A.; Lloyd, L.; Omar, M.; and Boerkoel, J. 2018. New perspectives on flexibility in simple temporal planning. In *Proc. of the 28th International Conference on Automated Planning and Scheduling (ICAPS-18)*, 123–131.
- Hunsberger, L. 2009. Fixing the semantics for dynamic controllability and providing a more practical characterization of dynamic execution strategies. In *Proc. of the 16th International Symposium on Temporal Representation and Reasoning (TIME-09)*, 155–162.
- Lund, K.; Dietrich, S.; Chow, S.; and Boerkoel, J. 2017. Robust execution of probabilistic temporal plans. In *Proc. of the 31st National Conference on Artificial Intelligence (AAAI-17)*, 3597–3604.
- Morris, P., and Muscettola, N. 2005. Temporal dynamic controllability revisited. In *Proc. of the 20th National Conference on Artificial Intelligence (AAAI-05)*, 1193–1198.
- Morris, P. 2006. A structural characterization of temporal dynamic controllability. In *Proc. of Principles and Practice of Constraint Programming (CP-06)*, 375–389.
- Motzkin, T. S., and Straus, E. G. 1965. Maxima for graphs and a new proof of a theorem of turán. *Journal canadien de mathématiques* 17(0):533–540.
- Nilsson, M.; Kvarnström, J.; and Doherty, P. 2014. Classical dynamic controllability revisited - a tighter bound on the classical algorithm. In *Proc. of 6th International Conference on Agents and Artificial Intelligence (ICAART-14), 6-8 March 2014, Angers, France*, 130–141.
- Rossi, F.; Venable, B.; and Yorke-Smith, N. 2003. Simple temporal problems with preferences and uncertainty. In *Proc. CP 2003 Workshop on Online Constraint Solving: Handling Change and Uncertainty*.
- Santana, P.; Vaquero, T.; Toledo, C.; Wang, A.; Fang, C.; and Williams, B. 2016. Paris: a polynomial-time, risk-sensitive scheduling algorithm for probabilistic simple temporal networks with uncertainty. In *Proc. of the 26th International Conference on Automated Planning and Scheduling (ICAPS-16)*, 267–275.
- Tawarmalani, M., and Sahinidis, N. V. 2005. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming* 103:225–249.
- Tsamardinos, I. 2002. A probabilistic approach to robust execution of temporal plans with uncertainty. In *Methods and Applications of Artificial Intelligence*, 97–108.
- Vidal, T., and Fargier, H. 1999. Handling contingency in temporal constraint networks: from consistency to controllabilities. In *Journal of Experimental and Theoretical Artificial Intelligence*, 0–32.
- Wilson, M.; Klos, T.; Witteveen, C.; and Huisman, B. 2014. Flexibility and decoupling in simple temporal networks. In *Artificial Intelligence* 214, 26–44.
- Yu, P.; Fang, C.; and Williams, B. 2014. Resolving uncontrollable conditional temporal problems using continuous relaxations. In *Proc. of 24th International Conference on Automated Planning and Scheduling (ICAPS-14)*, 341–349.
- Yu, P.; Fang, C.; and Williams, B. 2015. Resolving over-constrained probabilistic temporal problems through chance constraint relaxation. In *Proc. of the 29th AAAI Conference on Artificial Intelligence (AAAI-15)*, 3425–3431.