NASA/TM–2020–220480

Dynamic Controllability of Partially Observable Temporal Plans

Arthur Bit-Monnot
University of Sassari, Sassari, Italy

Paul H. Morris
NASA Ames Research Center

May 2020

NASA STI Program. . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collection of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.
- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI Program, see the following:

- Access the NASA STI program home page at `http://www.sti.nasa.gov`
- E-mail your question to `help@sti.nasa.gov`
- Phone the NASA STI Information Desk at 757-864-9658
- Write to:
  NASA STI Information Desk
  Mail Stop 148
  NASA Langley Research Center
  Hampton, VA 23681-2199

NASA/TM–2020–220480

Dynamic Controllability of Partially Observable Temporal Plans

Arthur Bit-Monnot
University of Sassari, Sassari, Italy

Paul H. Morris
NASA Ames Research Center

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035-1000

May 2020

Acknowledgments

Thanks to Jeremy Frank for encouraging this study in connection with temporal reasoning methods for Planning and Scheduling. Also thanks to Dave Smith for facilitating the collaboration.

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an offical endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA STI Program / Mail Stop 148
NASA Langley Research Center
Hampton, VA 23681-2199
Fax: 757-864-6500

Dynamic Controllability of Partially Observable Temporal Plans

Arthur Bit-Monnot
University of Sassari
Sassari, Italy

Paul Morris
NASA Ames Research Center
Moffett Field, CA 94035, U.S.A.

NASA/TM-2020-220480

Abstract

The formalism of Simple Temporal Networks provides methods for evaluating the feasibility of temporal plans. The basic formalism deals with the consistency of quantitative temporal requirements on scheduled events. Over time, the formalism has been extended to handle exogenous events with varying degrees of observability. A major problem that has only been partially solved before now involves a combination of observable and unobservable events. In this paper, we present a sound and complete solution to this problem.

1 Introduction

Many applications (for example, the Remote Agent Experiment (Muscettola et al., 1998), as an early one) have drawn attention to the importance of quantitative reasoning about time in practical planning systems. In particular, a need has been felt for temporal representations to specify scheduling requirements that an agent needs to satisfy. In general, these requirements could involve exogenous events, as well as the agent's own actions, and these events might or might not be observable. A number of formalisms have been established in response to the need to model these kinds of problems, and algorithms have been developed to solve them. The formalism of Simple Temporal Networks and its extensions has been particularly useful in this regard. Nevertheless, there remain many unsolved problems and under-developed theory in this area. This paper provides a a greater understanding and solution for one such class of problems.

1.1 STNs and Extensions

A Simple Temporal Network (STN) (Dechter, Meiri, and Pearl, 1991) is a graph in which the edges are annotated with upper and lower numerical bounds. The nodes in the graph represent temporal events or *timepoints*, while the edges correspond to constraints on the durations between the events. Each STN is associated with a *distance graph* derived from the upper and lower bound constraints. An STN is consistent if and only if the distance graph does not contain a negative cycle. To avoid confusion with the distance graph, we will refer to edges in the STN as *links* while the term *edges* will be reserved for edges in the distance graph.

A Simple Temporal Network With Uncertainty (STNU), introduced by Vidal and Fargier (Vidal and Fargier, 1999; Vidal, 2000), is similar to an STN except the links are divided into two classes, *requirement links* and *contingent links*. Requirement links are temporal constraints that the agent must satisfy, like the links in an ordinary STN. Contingent links may be thought of as representing causal processes of uncertain duration, or periods from a reference time to exogenous events. Their finish timepoints, called here *contingent timepoints*, are controlled by Nature, subject to the limits imposed by the bounds on the contingent links. We will refer to the start timepoint of a contingent link as its *activation* timepoint. This may itself be a contingent timepoint if it is the finish point of some other contingent link. All other timepoints, called *executable timepoints*, are controlled by the agent, whose goal is to satisfy the bounds on the requirement links. Each contingent link is required to have finite positive upper and lower bounds. An STNU may be thought of as determining a family of STNs where the contingent links take on each of their possible durations; the individual STNs in the family are called *projections*. An STNU is said to be *Weakly Controllable* if every projection is consistent. Weak Controllability was shown to be in co-NP, and later proved to be co-NP Hard (Morris and Muscettola, 1999). However, this property does not support a generally useful execution strategy.

The uncontrolled timepoints in STNUs are generally assumed to be either all unobservable, or all observable when they occur, giving rise to different execution strategies. An STNU is *Strongly Controllable* if there is a single schedule that satisfies the requirements in *all* of the projections, and thus does not depend on observations. An STNU is said to be *Dynamically Controllable* (Vidal and Fargier, 1999; Morris, Muscettola, and Vidal, 2001; Hunsberger, 2009; Morris, 2014) if there is a strategy for scheduling each executable timepoint that depends only on observations that are available in the past *or present* at the time it is scheduled.[1] Whether an STNU is Dynamically Controllable or not can be determined by an algorithm that runs in cubic time (Morris, 2014). The algorithm tightens some constraints in a way that makes explicit limitations on the execution strategies due to the presence of contingent links

As, mentioned, an STN has an alternative representation as a *distance graph* (Dechter, Meiri, and Pearl, 1991). Similarly, there is a representation for an STNU called the *labeled distance graph* (Morris and Muscettola, 2005) that is exploited in the Dynamic Controllability algorithm.

A Partially Observable STNU (POSTNU) (Moffitt, 2007) is an STNU in which the contingent timepoints are further subdivided into observable and unobservable (hidden) timepoints. Thus, the controllability problem for a POSTNU may be regarded as a combination of Strong and Dynamic Controllability. In the general POSTNU problem, a contingent link may be activated by a hidden timepoint. In that case, if the endpoint is observable, the POSTNU semantics specifies that when it is observed, we learn only the *time* of the endpoint, not the *duration* of the link that was activated by the hidden

[1]The literature varies on whether observations can be reacted to instantaneously, or must be strictly in the past. The instantaneous variant was the original concept and has some technical advantages, including allowing an executable timepoint to be simultaneous with an observation.

timepoint. Of course we do learn (or can easily calculate) the time difference between the observed endpoint and any previous known time.

Moffitt's algorithm (Moffitt, 2007) for checking the controllability of a POSTNU is complete but not sound in that it might incorrectly label a POSTNU as controllable. Another algorithm, also relying on the compilation to STNUs is provided by (Bit-Monnot, Ghallab, and Ingrand, 2016) that is sound but only complete for a subclass of POSTNUs. A polynomial sound and complete algorithm for assessing the controllability of the general POSTNU problem has not previously been known. We present one here.

2 Formal Preliminaries

Formally, an STN may be described as a 4-tuple $\langle N, E, l, u\rangle$ where N is a set of nodes called *timepoints*, E is a set of edges called *links*, and l and u are functions mapping each edge into the lower and upper bounds of the interval of possible durations.

An STNU is a 5-tuple $\langle N, E, l, u, C\rangle$, where N, E, l, u are as in a STN, and C is a subset of the links: the *contingent* links, the others being called *requirement* links. Each contingent link e is required to satisfy $0 < l(e) \leq u(e) < \infty$.[2]

The durations of contingent links are assumed to vary independently; thus every combination that satisfies the upper and lower bounds of the contingent links gives rise to an STN called a *projection*. Contingent links are not allowed to share finish points. However, their finish points may be the start timepoints of other contingent links (thus potentially forming a branching tree structure). The finish timepoints of contingent links are called *contingent timepoints*. A contingent timepoint that does *not* start a new contingent link is called a *terminal* contingent timepoint.

A *schedule* is an assignment of times to all the timepoints. The *pre-history* of a specific time t with respect to a schedule T, denoted by $T\{\preceq t\}$, specifies the durations of all contingent links that have finished up to and including time t.

An *execution strategy* S is a mapping

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

where $\mathcal{P}$ is the set of projections and $\mathcal{T}$ is the set of schedules. An execution strategy S is *viable* if $S(p)$, henceforth written S_p, is consistent with p for each projection p. An execution strategy S is *dynamic* if for projections $p1$ and $p2$, and executable timepoint x where $S_{p1}(x) = t$, the strategy satisfies [3]

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

An STNU is *Dynamically Controllable* if there is a viable dynamic execution strategy.

A POSTNU is a 6-tuple $\langle N, E, l, u, C, O\rangle$ where N, E, l, u, C are as in a STNU. Here O is a subset of the contingent timepoints, called *observable* timepoints. For a

[2]We allow a contingent link with $l(e) = u(e)$ although for STNUs it essentially behaves the same as a requirement link with the same bounds. However, the behavior can be different in the POSTNU context.

[3]This incorporates the flaw correction of Hunsberger (2013) but also provides for an instantaneous reaction.

POSTNU, we will use the terminology *micro-projection* to describe each STN determined by the possible combinations of contingent links, which again are assumed to vary independently within their bounds. We will use the term *combo-link* to describe the result of composing a chain of contingent links. We may often refer to the original contingent links as *micro-links* to emphasize the distinction from combo-links. In general, the durations of separate combo-links may be correlated, unlike the case for micro-links. However, separate combo-links will still be independent if they do not share micro-links.

Recall that an STNU, hence also a POSTNU, does not allow two contingent links to have the same finish timepoint. Thus, every contingent timepoint is the finish timepoint of exactly one contingent link. Note that if we proceed from any contingent timepoint Y **backwards** through its chain of contingent links, we must eventually come to a non-hidden timepoint (i.e., a timepoint that is either observable or executable). We will refer to the unique first such timepoint X as the *closest non-hidden ancestor* of Y. Now suppose Y is observable. Consider the combo-link formed by composing the micro-links in the chain from X to Y. Since Y is observable and X is non-hidden the duration of the combo-link will be known when Y is observed. We will call the combo-links formed in this way from the observable timepoints Y and their closest non-hidden ancestors X the *macro-links* of the POSTNU.

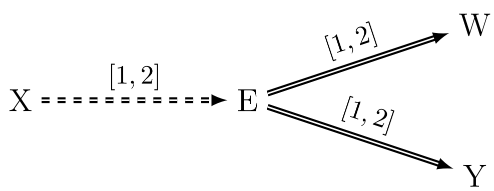

Figure 1: Example POSTNU where W and Y are observable contingent timepoints, E is a hidden contingent timepoint and X is an executable timepoint.

Consider the POSTNU of Figure 1 with two observables W and Y and a hidden timepoint E. The network has three micro-links XE, EW, EY and 8 micro projections where each A $\xRightarrow{[1,2]}$ B micro-link is replaced by either A $\xrightarrow{[1,1]}$ B or A $\xrightarrow{[2,2]}$ B. Note that X is the *closest non-hidden ancestor* of both W and Y. Observing the occurrence time of W and Y, respectively provide us with the duration of the *macro-links* XW and XY respectively. However, E being hidden means that the duration of XE cannot be directly observed. Finally, note that the durations of the macro-links XW and XY are correlated since they share the hidden contingent link XE. As a consequence, observing the duration of XW to be 4 (meaning XE = EY = 2), means that the possible values of XY are 3 and 4, depending on the value taken by EY.

Definition 1 *Two micro-projections* p_1 *and* p_2 *are* observationally equivalent *if the durations of all the macro-links are the same in* p_1 *and* p_2.

Clearly this relation is symmetric, reflexive, and transitive: thus, an equivalence relation. Its equivalence classes will be called *macro-projections*. Note that the macro-links have fixed durations in a macro-projection.

For example, in the POSTNU (where only E is hidden)

$$\mathrm{X} \xRightarrow{[1,10]} \mathrm{E} \xRightarrow{[1,10]} \mathrm{Y}$$

the set of micro-projections where XE and EY sum to 15, such as $6 + 9$, $10 + 5$, etc., constitute a macro-projection where XY = 15. Observe that the "extremal" macro-projections where XY = 2 and 20 each contain only one micro-projection.

For a POSTNU, an execution strategy S is *viable* if S_p is consistent with p for each micro-projection p. In order to define Dynamic Controllability for POSTNUs, we can now simply modify the definition of pre-history.

Definition 2 *.The* observational pre-history *of a specific time t with respect to a schedule T, denoted by $T\{\trianglelefteq t\}$, specifies the durations of all macro-links that have finished up to and including time t.*

A *dynamic strategy* S for a POSTNU is defined similarly as for an STNU except we require it to satisfy

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

instead, i.e., we replace $\preceq$ with $\trianglelefteq$ in the definition. Then, as for an STNU, a POSTNU is *Dynamically Controllable* if there is a viable dynamic execution strategy.

3 Analysis and General Approach

In a POSTNU, the contingent links (the micro-links) may be partitioned into separate groups, called *hidden groups*, whose elements are connected to each other by hidden timepoints. (Requirement links are ignored in determining the hidden groups.) For example, A $\Rightarrow$ E and E $\Rightarrow$ B will be in the same hidden group if E is a hidden timepoint.[4] If A is also hidden, the group will extend further in the backward direction. If B is hidden and starts one or more contingent links, the group will extend further in the forward direction and may branch. (Note however, that if A is *not* hidden, then A $\Rightarrow$ E_1 and A $\Rightarrow$ E_2 will *not* be in the same hidden group.) The boundaries of the group will be determined by timepoints that are terminal (i.e., not themselves the start of a contingent link), or non-hidden, or both.

Since the POSTNU definition does not permit two contingent links to have the same endpoint, the links in a hidden group G will form the edges of a tree rooted at some non-hidden timepoint A, which will be the closest non-hidden ancestor for all the timepoint nodes in the tree. The leaves of the tree will be either observable timepoints, or terminal hidden timepoints.

Each observable leaf will be called an *eye* of the hidden group. Notice that if a combo-link connects a hidden timepoint E to an eye, the observation of the eye provides information limiting the possible occurrence time of E to within the bounds of

[4]Formally, A $\Rightarrow$ E and E $\Rightarrow$ B are *hidden-related* in this case. The symmetric reflexive transitive closure of this relation is an equivalence relation, and the equivalence classes are the hudden groups.

the combo-link. Moreover, if another combo-link connects some other hidden timepoint E' to E, any limitation on E' may propagate to a limitation on E. This may be the case even if E is a terminal hidden timepoint.

Our approach for checking the Dynamic Controllability of a POSTNU essentially factors the problem into two parts: (1) solving a separate Strong Controllability problem for each of the macro-projections; and (2) feeding those solutions into an overall STNU-like Dynamic Controllability algorithm.

For part (1), the occurrences of hidden timepoints will be restricted to values consistent with the observations of the eyes. In that case, satisfying the requirement links is tantamount to solving special cases of the temporal decoupling problem (Hunsberger, 2002) where the observations are treated as additional constraints on Nature, which is viewed as the second agent. Given a fixed macro-projection P, the relative time of a hidden timepoint E within its hidden group G may vary over the micro-projections of P. However, the earliest and latest occurrences of E relative to G will be fixed in P, and determined by the observations. We will see that the requirement links on E can be effectively transferred to a new pair of timepoints related to the earliest and latest occurrences of E.

For part (2), the Strong Controllability solutions for the macro-projections can be expressed in terms of formulas with parameters that depend on the observations. These formulas can be interpreted as generalized versions of the labels used in the labeled distance graph of a standard STNU. We introduce corresponding generalizations of the STNU reduction rules. These are shown to be sound and complete for determining Dynamic Controllability.

We will prove the transformation steps and generalized reduction rules are sound by showing they leave the set of viable dynamic strategies unchanged. (This set may be empty if the POSTNU is not Dynamically Controllable.) Completeness will be a consequence of the fact that successful completion of the reduction process will make the projections dispatchable, and this implies Dynamic Controllability (Morris, 2014, 2016).

Many of the transformation steps involve replacing requirement constraints by equivalent ones. The following lemmas will be generally useful for showing the soundness of such steps, and we will apply them throughout without explicit mention.

Lemma 1 *Given a POSTNU, replacing a requirement constraint by an equivalent one leaves the set of viable dynamic strategies unchanged.*

Proof: The schedule chosen by a dynamic strategy depends only on the durations of the macro-links. Thus, replacing one requirement constraint by another will not change the schedule, and hence the dynamic nature of the strategy. In general, it might change the viability property, but this will not be the case if the replacement constraint is equivalent to the original. □

Lemma 2 *Given a POSTNU, adding a requirement constraint that is logically implied by the other constraints leaves the set of viable dynamic strategies unchanged.*

Proof: The proof is similar to that of the previous lemma. □

4 Transferring Requirement Constraints

The task of transferring requirement links away from hidden timepoints may be thought of as a type of multi-agent temporal decoupling problem (Hunsberger, 2002). Accordingly, we adopt a viewpoint where Nature is regarded as an agent scheduling the hidden groups subject to the observations, which are regarded as additional constraints on what Nature could have done. This leads to the following concept of a Nature STN.

Definition 3 *The* Nature STN *for a hidden group G with root timepoint X is an STN schema that contains an STN link for each micro-link in G, with the same bounds. In addition, for each eye Y of G, the network contains a rigid link from X to Y of length $\dot{Y}$, where $\dot{Y}$ is a variable representing the observed duration of the macro-link from X to Y.*

The links in the Nature STN that correspond to the micro-links will be called *concrete links*, while those corresponding to the observations will be called *observation links*.

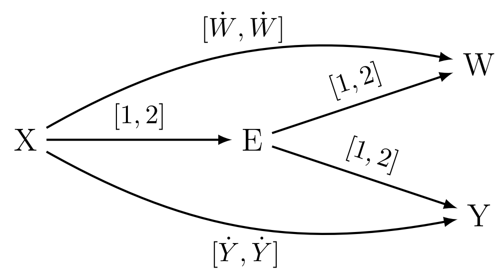

Figure 2: Nature STN schema of the hidden group depicted in Figure 1.

Note that the Nature STN schema instantiates to a specific STN for each macro-projection (where the observed durations have definite values). We assume the resulting STNs are all consistent; otherwise the model would be faulty, which is outside the scope of this paper. Thus, we can assume the existence of shortest paths that are well-defined and non-cyclic. We can also assume that the observed durations are non-negative since they result from the combination of contingent links that are non-negative by definition.

We wish to transfer requirement constraints from hidden timepoints to equivalent constraints on observable timepoints. The following lemma supports an intermediate step in that direction. Suppose E is a hidden timepoint in some hidden group G rooted at X.

Lemma 3 *Consider some fixed macro-projection P where $[e^-, e^+]$ are the inferred* **tightest** *bounds of the XE temporal distance in the Nature STN. Let $E_{lo} = X + e^-$ and $E_{hi} = X + e^+$. Suppose $[z^-, z^+]$ are the bounds on a requirement link $\mathcal{Z}$ that has E as the source timepoint and some other timepoint Z (with few restrictions[5]) as the target timepoint. Then $\mathcal{Z}$ can be replaced by the constraint $E_{hi} + z^- \leq Z \leq E_{lo} + z^+$.*

[5]We exclude from this analysis requirements between hidden timepoints in the same group. Their viability is independent of observation, and must be a logical consequence of the microlinks in the group.

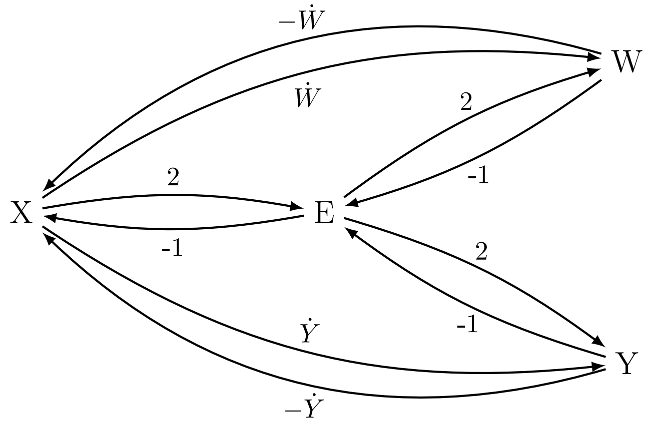

Figure 3: Distance graph of the Nature STN schema of Figure 2.

Proof: Observe that the micro-projections within the P macro-projection will contain all the possible values for E in the solutions of the Nature STN relative to P, i.e., all E in the range $[\mathrm{E_{lo}}, \mathrm{E_{hi}}]$. Note that:

$$(\forall E \in [\mathrm{E_{lo}}, \mathrm{E_{hi}}] : E + z^- \leq Z) \quad \equiv \quad \mathrm{E_{hi}} + z^- \leq Z$$
$$(\forall E \in [\mathrm{E_{lo}}, \mathrm{E_{hi}}] : Z \leq E + z^+) \quad \equiv \quad Z \leq \mathrm{E_{lo}} + z^+.$$

Thus, $\mathcal{Z}$ can be replaced by the equivalent $\mathrm{E_{hi}} + z^- \leq Z \leq \mathrm{E_{lo}} + z^+$. $\Box$

It is useful to think of $\mathrm{E_{lo}}$ and $\mathrm{E_{hi}}$ as *virtual* timepoints that have fixed offsets from X in each macro-projection (but the offsets may vary between macro-projections).

Motivated by the lemma, we are interested in lower/upper bounds on the distance from X to E in the Nature STN. This is not the same in general as the raw bounds on the combo-link from X to E because of the additional constraints imposed by the observations. We now develop symbolic formulas that express the bounds for each macro-projection in terms of the raw bounds and the observed durations.

Recall that upper bounds in an STN may be calculated as shortest path distances. Observe that there must exist a shortest path from X to E that contains at most one of the observation edges. (Otherwise the path would be cyclic since all the observation links have X as their start timepoint.) However, which observation, if any, is in the shortest path may depend on the relative values of the observation variables. Thus, we may write

$$\mathrm{SD}(X, E) = \min(\mathrm{CSD}(X, E), \mathrm{OSD}(X, E))$$

where SD is the shortest-path distance, CSD is the shortest-path distance over concrete edges only, and OSD is the shortest distance over paths that traverse a single observation edge, defined by

$$\mathrm{OSD}(X, E) = \min\{\dot{Y} + \mathrm{CSD}(Y, E) \mid Y \in \mathrm{eyes}(G)\}$$

where $\text{eyes}(G)$ denotes the set of eyes of G (and recall $\dot{Y} = Y - X$). We can summarize this as

$$\begin{aligned} \text{UB}(X, E) &= \min(\text{CUB}(X, E), \text{OUB}(X, E)) \\ \text{OUB}(X, E) &= \min\{\dot{Y} + \text{CUB}(Y, E) \mid Y \in \text{eyes}(G)\} \end{aligned}$$

where we have rewritten SD, OSD, and CSD using analogous upper bound notation. Looking at this, it is convenient to regard the root timepoint X as an additional special eye called the *root eye* such that $\dot{X} = 0$ always. (The other eyes will then be called *leaf* eyes to distinguish them.) Then we can rewrite the formula (where the notation $\text{Eyes}(G)$ includes the root eye) as

$$\text{UB}(X, E) = \min_{Y \in \text{Eyes}(G)} (\dot{Y} + \text{CUB}(Y, E)).$$

Analogously, for the shortest path distance from E to X, we have

$$\text{SD}(E, X) = \min_{Y \in \text{Eyes}(G)} (\text{CSD}(E, Y) - \dot{Y})$$

where $\dot{Y}$ is subtracted because we are now traversing the observation link in the opposite direction. Recalling that $\text{LB}(X, E) = -\text{SD}(E, X)$, we can rewrite this in terms of the lower bound as

$$\text{LB}(X, E) = \max_{Y \in \text{Eyes}(G)} (\text{CLB}(Y, E) + \dot{Y})$$

where LB, CLB, and OLB are lower bound notation analogous to that used for the upper bounds. Notice that min changes to max and the rigid link value $\dot{Y}$ ends up as an added term in the lower bound representation. Writing these together, we have

$$\begin{aligned} \text{UB}(X, E) &= \min_{Y \in \text{Eyes}(G)} (\dot{Y} + \text{CUB}(Y, E)) \\ \text{LB}(X, E) &= \max_{Y \in \text{Eyes}(G)} (\dot{Y} + \text{CLB}(Y, E)) \end{aligned}$$

as concise formulas.

Since $E_{\text{hi}} = X + \text{UB}(X, E)$ and $E_{\text{lo}} = X + \text{LB}(X, E)$, we could now apply Lemma 3 to transfer any requirements on E to requirements on E_{hi} and E_{lo}, whose fixed bounds with respect to X can be expressed in terms of formulas involving the observations. However, we will see that the formulas may involve negative offsets to the observations, which means they are not immediately observable. We postpone this issue to the next section, and close this section with an example.

Example 1 *Consider the following POSTNU where X and Z are executable timepoints, E is hidden, and Y is observable.*

$$\begin{tikzcd} X \arrow[r, dashed, "{[0,5]}"] & E \arrow[r, Rightarrow, "{[5,10]}"] \arrow[d, "{[0,10]}"] & Y \\ & Z & \end{tikzcd}$$

Even though E is not observable, the occurrence time of both X and Y provide indirect information on E. We are thus interested in reformulating the $\mathrm{E} \xrightarrow{[0,10]} \mathrm{Z}$ requirement link in terms of symbolic expressions involving the occurrence times of X and Y.

E is part of a hidden group composed of the micro-links XE and EY, for which the corresponding Nature STN distance graph is the following (note that Z is not part of the hidden group of E and hence not represented in the Nature STN).

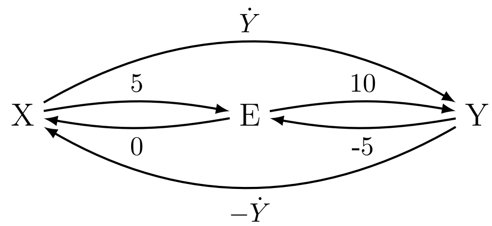

It is easy to see that the shortest distances XE and EX depend on $\dot{Y}$, the duration of the macro-link XY, resulting in the following shortest distance expressions:

$$SD(X,E) = \min\{5, \dot{Y} - 5\}$$
$$SD(E,X) = \min\{0, 10 - \dot{Y}\}$$

Interpreting these expression as the virtual timepoints E_{lo} and E_{hi}, we obtain the following definitions and corresponding bounds on Z.

$$E_{lo} = X + \max\{0, \dot{Y} - 10\} \qquad Z \leq E_{lo} + 10$$
$$E_{hi} = X + \min\{5, \dot{Y} - 5\} \qquad Z \geq E_{hi}$$

Let us give an intuitive interpretation of this result. We know that E occurs after X. Additionally, we know that E occurs at most 10 time units before Y. This is reflected in the definition of the virtual timepoint E_{lo} which should be interpreted as: it is not possible for E to occur before E_{lo}. The requirement of scheduling Z at most 10 time units after E_{lo} thus ensures that Z is scheduled at most 10 time units after E, which was our original constraint.

Similarly, we know that E occurs at most 5 time units after X and at least 5 time units before Y. This is reflected in the definition of the virtual timepoint E_{hi} which should be interpreted as: it is not possible for E to occur after E_{hi}. The requirement of scheduling Z after E_{hi} ensure that Z is scheduled after E which was our original constraint.

5 Compound Observables

Lemma 3 permits us to transfer requirements on hidden timepoints to virtual timepoints $\mathrm{E_{hi}}$ and $\mathrm{E_{lo}}$ that are fixed within each macro-projection, and depend only on the observation durations. One drawback, however, is that $\mathrm{E_{hi}}$ and $\mathrm{E_{lo}}$ may not be themselves observable.

For example, given an eye Y, the $\mathrm{CUB}(Y, E)$ value could be either negative or non-negative. (The E timepoint could be an ancestor or cousin of Y with respect to the tree structure of the hidden group.) If it is negative, then the value of $\dot{Y} + \mathrm{CUB}(Y, E)$ will not be known until later, when Y is observed. Consequently, we may not be able to tell whether

$$\mathsf{E}_{\mathsf{hi}} = X + \min_{Y \in \mathsf{Eyes}(G)} (\dot{Y} + \mathrm{CUB}(Y, E))$$

has occurred or not, *at the time* it occurs.

Note, however, that we can always choose some minimal fixed value $\delta_{\mathrm{hi}}(E) \geq 0$ such that $\mathrm{CUB}(Y, E) + \delta_{\mathrm{hi}}(E) \geq 0$ for every $Y \in \mathrm{Eyes}(G)$. The $\delta_{\mathrm{hi}}(E)$ value depends only on the weights of the concrete links and does not vary with the observations. In contrast to the E_{hi} value, $\mathsf{E}_{\mathsf{hi}} + \delta_{\mathrm{hi}}(E)$ will be observable. This suggests introducing a new virtual timepoint $\mathsf{E}_{\mathsf{hi+}} = \mathsf{E}_{\mathsf{hi}} + \delta_{\mathrm{hi}}(E)$. Similarly, we can introduce a virtual timepoint $\mathsf{E}_{\mathsf{lo+}} = \mathsf{E}_{\mathsf{lo}} + \delta_{\mathrm{lo}}(E)$ so that the offsets in the $\mathrm{LB}(X, E) + \delta_{\mathrm{lo}}(E)$ formula will be non-negative.

The following lemmas justify the transfer of requirement constraints from E_{hi} and E_{lo} to $\mathsf{E}_{\mathsf{hi+}}$ and $\mathsf{E}_{\mathsf{lo+}}$, respectively, and establish their observability. The lemmas are stated for completeness. We omit the proofs, which are immediate.

Lemma 4 *If $E_{hi+} = E_{hi} + \delta_{\mathrm{hi}}(E)$, then $E_{hi} + z^- \leq Z$ is equivalent to $E_{hi+} + (z^- - \delta_{\mathrm{hi}}(E)) \leq Z$. If $E_{lo+} = E_{lo} + \delta_{\mathrm{lo}}(E)$ then $Z \leq E_{lo} + z^+$ is equivalent to $Z \leq E_{lo+} + (z^+ - \delta_{\mathrm{lo}}(E))$.*

Lemma 5 *If an event $\dot{Y}$ is observable, then the event $\dot{Y} + p$, where $p \geq 0$ is some fixed value, is also observable. If every event in some set S is observable, then $\min_{y \in S}(y)$ and $\max_{y \in S}(y)$ are also observable.*

We can paraphrase $\min_{y \in S}(y)$ as "whichever is earliest of the events in S, and $\max_{y \in S}(y)$ as "whichever is latest of the events in S." Applying this interpretation to the

$$\mathsf{E}_{\mathsf{hi+}} - X = \min_{Y \in \mathsf{Eyes}(G)} (\dot{Y} + \mathrm{CUB}(Y, E) + \delta_{\mathrm{hi}}(E))$$
$$\mathsf{E}_{\mathsf{lo+}} - X = \max_{Y \in \mathsf{Eyes}(G)} (\dot{Y} + \mathrm{CLB}(Y, E) + \delta_{\mathrm{lo}}(E))$$

expressions, we can regard $\mathsf{E}_{\mathsf{hi+}}$ and $\mathsf{E}_{\mathsf{lo+}}$ as *compound observables*, derived from the eyes, with non-negative offsets calculated from the bounds of the contingent links. Repeated applications of the two lemmas have the effect of transferring the requirement constraints from hidden timepoints to equivalent constraints on compound observables.

Example 2 *Following on from Example 1, where we had*

$$E_{lo} = X + \max\{0, \dot{Y} - 10\} \qquad Z \leq E_{lo} + 10$$
$$E_{hi} = X + \min\{5, \dot{Y} - 5\} \qquad Z \geq E_{hi}$$

This could be interpreted as the following distance graph.

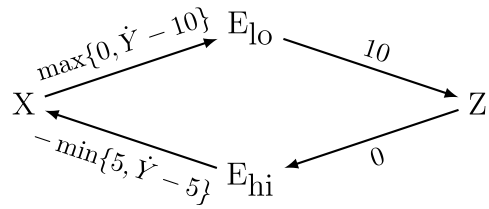

Consider the macro-projection where $\dot{Y} = 12$. Further assume that X is scheduled at the temporal origin ($X = 0$). In this macro-projection, E_{lo} would occur at time 2 while E_{hi} would occur at time 5. However this information is unknown until the duration $\dot{Y}$ is observed at time 12 (when the observable Y is observed).

We can now define the compound observables $E_{hi+} = E_{hi} + \delta_{\mathrm{hi}}(E)$ and $E_{lo+} = E_{lo} + \delta_{\mathrm{lo}}(E)$ by choosing the appropriate $\delta_{\mathrm{hi}}(E)$ and $\delta_{\mathrm{lo}}(E)$ terms such that all terms in the min/max expressions are non-negative. Choosing $\delta_{\mathrm{hi}}(E) = 5$ and $\delta_{\mathrm{lo}}(E) = 10$, we obtain the following expressions and corresponding distance graph.

$$E_{lo+} = X + \max\{10, \dot{Y}\} \qquad\qquad Z \leq E_{lo+}$$
$$E_{hi+} = X + \min\{10, \dot{Y}\} \qquad\qquad Z \geq E_{hi+} - 5$$

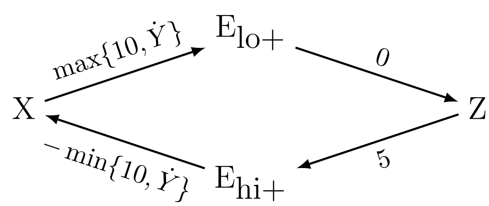

6 Generalized Distance Graph

6.1 Generalized labels

In this section, we interpret the max/min bound expressions for the compound observables as generalized labels analogous to those in the labeled distance graph of a conventional STNU. (See for example (Morris, 2014)). Pursuing this analogy, we will introduce Reduction Rules for the max/min expressions and prove their soundness.

Since the requirements have been transferred to observables, our further reasoning is in terms of the macro-projections. Thus, in the following, we may use the unqualified term “projection” to mean macro-projection, provided we remain mindful that their independence is limited by possibly shared micro-links.

Recall that a compound observable of the form $\mathrm{E}_{\mathrm{hi+}}$ provides a lower-bound

$$\min_{Y \in \mathrm{Eyes}(G)} \{ \dot{Y} + \mathrm{CUB}(Y, E) + \delta_{\mathrm{hi}}(E) \}$$

on the X to $\mathrm{E}_{\mathrm{hi+}}$ distance that could potentially combine with a $\mathrm{E}_{\mathrm{hi+}}$ to Z lower bound of some requirement. In our generalized distance graph, this will be represented as an edge with weight

$$\ell_{\mathrm{hi+}}(E) = -\min_{Y \in \mathrm{Eyes}(G)} \{ \dot{Y} + \mathrm{CUB}(Y, E) + \delta_{\mathrm{hi}}(E) \} \tag{1}$$

from $\mathrm{E}_{\mathrm{hi+}}$ to X. Notice that all the $\dot{Y} + \mathrm{CUB}(Y, E) + \delta_{\mathrm{hi}}(E)$ terms evaluate to non-negative numbers. Moreover, $\dot{Y}$ for the leaf eyes is always positive. Thus, the negated expression will evaluate to a negative number. The edge is analogous to an Upper-Case edge in an STNU distance graph, and for convenience we will sometimes call it an Upper-Case edge here.

Similarly, for a compound observable of the form $\mathrm{E}_{\mathrm{lo+}}$, there will be an edge with weight

$$\ell_{\mathrm{lo+}}(E) = \max_{Y \in \mathrm{Eyes}(G)} \{ \dot{Y} + \mathrm{CLB}(Y, E) + \delta_{\mathrm{lo}}(E) \} \tag{2}$$

from X to $\mathrm{E}_{\mathrm{lo+}}$. The $\dot{Y} + \mathrm{CLB}(Y, E) + \delta_{\mathrm{lo}}(E)$ terms all evaluate to non-negative numbers in each macro-projection. The edge is analogous to a Lower-Case edge in an STNU distance graph, and we will sometimes use that terminology here also.

When there are multiple leaf eyes, it is also possible for observations of one eye to provide information that restricts the possible time of occurrence of another eye *before* it has occurred, and this may be useful for an early determination that a requirement has been satisfied. This may happen if the macro-links for the two eyes share micro-links. The generalized labels for each eye Y must therefore reflect inferences from observations of the other eyes. As it turns out, the appropriate labels for Y are the same as if Y were also a hidden timepoint. The $\dot{Y}$ coming from the direct observation of Y is simply included as another term in the min/max formulas. Note, however, that the $\mathrm{CLB}(Y, Y)$ and $\mathrm{CUB}(Y, Y)$ terms will be 0.

Relationship with STNU It is instructive to regard an ordinary STNU as a special case of a POSTNU where there are no hidden timepoints. If we were to treat each contingent link X $\xRightarrow{[\mathrm{y}^-,\mathrm{y}^+]}$ Y as a degenerate "hidden group" with a single micro-link and a single leaf eye, the above analysis could still be applied, resulting in an upper-case label of the form $-\min(\dot{Y}, y^+)$ and a lower-case label of $\max(\dot{Y}, y^-)$. These may be compared to the conventional $Y{:}-y^+$ and $y : y^-$ labels of an STNU, and suggests a semantic interpretation of those labels. It is helpful to bear this comparison in mind when we consider reduction rules, in the next section.

6.2 Construction of the labeled distance graph

Given a POSTNU Π, we now give the complete procedure for the construction of the labeled distance graph $(V, Edges)$. The set of vertices V will be composed of the controllable and observable timepoints of Π as well as the two compound observables of each hidden timepoint of Π. The set of edges $Edges$ is obtained by *(i)* converting macro-links to upper and lower case edges, and *(ii)* transferring requirement links involving hidden timepoints to the corresponding compound observables.

The algorithm is detailed in Algorithm 1. Lines 8-14 add the compound observable timepoints and their edges to the network. Then lines 16-26 transfer the requirement constraints from the hidden timepoints to the new observables.

Algorithm 1 Construction of the labeled distance graph of a POSTNU Π

1: **procedure** DISTANCEGRAPH(Π)
2: $\quad V \leftarrow$ CONTROLLABLES(Π)
3: $\quad Edges \leftarrow \emptyset$
4: $\quad \{ G_1, \ldots, G_n \} \leftarrow$ ORDEREDHIDDENGROUPS(Π) ▷ Topologically ordered
5: $\quad$ **for all** G_i **do**
6: $\qquad X \leftarrow root(G_i)$ ▷ Split between lo/hi
7: $\qquad$ Construct nature STN of G_i
8: $\qquad$ **for all** non-root timepoint $Y \in G_i$ **do**
9: $\qquad\quad V \leftarrow V \cup \{ \mathrm{Y}_{\mathrm{lo+}}, \mathrm{Y}_{\mathrm{hi+}} \}$
10: $\qquad\quad Edges \leftarrow Edges \cup \{ \mathrm{X} \xrightarrow{\ell_{\mathrm{lo+}}(\mathrm{Y})} \mathrm{Y}_{\mathrm{lo+}}, \mathrm{Y}_{\mathrm{hi+}} \xrightarrow{\ell_{\mathrm{hi+}}(\mathrm{Y})} \mathrm{X} \}$
11: $\qquad$ **end for**
12: $\qquad$ **for all** $Y \in eyes(G_i)$ **do**
13: $\qquad\quad Edges \leftarrow Edges \cup \{ \mathrm{X} \xrightarrow{\ell_{\mathrm{lo+}}(\mathrm{Y})} \mathrm{Y}, \mathrm{Y} \xrightarrow{\ell_{\mathrm{hi+}}(\mathrm{Y})} \mathrm{X} \}$
14: $\qquad$ **end for**
15: $\quad$ **end for**
16: $\quad$ **for all** $\mathrm{A} \xrightarrow{\ell} \mathrm{B} \in$ REQUIREMENTS(Π) **do**
17: $\qquad$ **if** A is contingent **then**
18: $\qquad\quad A \leftarrow \mathrm{A}_{\mathrm{lo+}}$
19: $\qquad\quad \ell \leftarrow \ell - \delta_{\mathrm{lo}}(A)$
20: $\qquad$ **end if**
21: $\qquad$ **if** B is hidden **then**
22: $\qquad\quad B \leftarrow \mathrm{B}_{\mathrm{hi+}}$
23: $\qquad\quad \ell \leftarrow \ell + \delta_{\mathrm{hi}}(B)$
24: $\qquad$ **end if**
25: $\qquad Edges \leftarrow Edges \cup \{ \mathrm{A} \xrightarrow{\ell} \mathrm{B} \}$
26: $\quad$ **end for**
27: $\quad$ **return** $(V, Edges)$
28: **end procedure**

Complexity Observe that in a nature STN with m vertices, the number of edges e is bounded above by $4 \times (m - 1)$. In the worst case, computing the labels of the generalized distance graph require us to know the shortest distance between any two timepoints in a nature STN. This can be done in a single pass of Johnson's algorithm with a complexity of $O(m^2 \times log(m) + m \times e)$. Given that $e \in O(m)$ this further reduces to $O(m^2 \times log(m))$. This process must be repeated for each of the n nature STN.

Further observe that the nature STNs are built by partitioning the set of contingent edges and that there is exactly 1 contingent link per contingent timepoint. The total number of nodes across all nature STNs is thus bounded above by $2 \times |C|$ where $|C|$

is the number of contingent timepoints. The complexity of computing the generalized labeled in Algorithm 1 is thus $O(|C|^2 \times log|C|)$.

For completeness we must also incorporate the cost of transforming the requirement edges, each of which is treated exactly once in constant time. The overall complexity of Algorithm 1 is thus $O(|C|^2 \times log|C| + |E|)$, where $|E|$ is the number of requirement links in the original POSTNU. As we will see, this is superseded by the overall complexity of checking the Dynamic Controllability of the resulting network whose best known bound is cubic in the number of timepoints.

6.3 Generalized Label Properties

In this section, we note some properties of the generalized labels in the constructed labeled distance graph. These will be useful later when we introduce derived labels, and establish invariants that are preserved by the derivations. Here, we show the properties hold for the labels on the initial edges between root timepoints and observables. Recall the formulas

$$\mathrm{UB}(X, E) = \min_{Y \in \mathrm{Eyes}(G)} (\dot{Y} + \mathrm{CUB}(Y, E))$$
$$\mathrm{LB}(X, E) = \max_{Y \in \mathrm{Eyes}(G)} (\dot{Y} + \mathrm{CLB}(Y, E)).$$

We have the following lemmas that will be important for the derivations.

Lemma 6 *The following inequalities hold in general.*

$$\mathit{CUB}(X, E) \geq \mathit{CUB}(Y, E) \text{ for all } Y \in \mathit{eyes}(G)$$
$$\mathit{CLB}(X, E) \geq \mathit{CLB}(Y, E) \text{ for all } Y \in \mathit{eyes}(G)$$

Proof: Recall that $\mathrm{CLB}(P, Q) = -\mathrm{CUB}(Q, P)$ for any P and Q.

Consider any $Y \in \mathrm{eyes}(G)$. Then

$$\mathrm{CUB}(Y, E) \leq \mathrm{CUB}(Y, X) + \mathrm{CUB}(X, E) \leq \mathrm{CUB}(X, E)$$

by the triamgle inequality and $\mathrm{CUB}(Y, X) = -\mathrm{CLB}(X, Y) \leq 0$. Also, by the triangle inequality, we have

$$\mathrm{CUB}(E, X) \leq \mathrm{CUB}(E, Y) + \mathrm{CUB}(Y, X)$$

and so

$$\mathrm{CLB}(Y, E) + \mathrm{CLB}(X, Y) \leq \mathrm{CLB}(X, E).$$

Then

$$\mathrm{CLB}(Y, E) \leq \mathrm{CLB}(Y, E) + \mathrm{CLB}(X, Y) \leq \mathrm{CLB}(X, E).$$

□

The following concepts will be useful. With respect to the above formulas, we will say an eye $W \in \mathrm{Eyes}(G)$ is *UB active* for E if

$$\dot{W} + \mathrm{CUB}(W, E) = \min_{Y \in \mathrm{Eyes}(G)} (\dot{Y} + \mathrm{CUB}(Y, E)).$$

We will also say a subset $\mathcal{W}$ of Eyes(G) is *UB active* for E if it contains an eye that is UB active for E. Note that this implies

$$\min_{Y\in\mathcal{W}}(\dot{Y}+\text{CUB}(Y,E)) = \min_{Y\in\text{Eyes}(G)}(\dot{Y}+\text{CUB}(Y,E)).$$

Analogously, $W \in$ Eyes(G) is *LB active* for E if

$$\dot{W}+\text{CLB}(W,E) = \max_{Y\in\text{Eyes}(G)}(\dot{Y}+\text{CLB}(Y,E))$$

and we similarly extend this concept to subsets of Eyes(G).

The following lemma identifies special macro-projections P_{sup}, where the root eye is UB active, and P_{inf}, where the root eye is LB active.

Lemma 7 *There is a macro-projection P_{sup} where UB(X,E) = CUB(X,E), and a macro-projection P_{inf} where LB(X,E) = CLB(X,E) .*

Proof: The All-Max and All-Min projections, where the micro-links take on their maximum and minimum bounds, respectively, satisfy the conditions. □

Recall that macro-link observations have only limited independence from each other—their durations may be correlated if they involve shared micro-links. Nevertheless, the following lemma establishes some flexibility with respect to which observations dominate the max/min computations.

Lemma 8 *Suppose: $\mathcal{E}$ is a subset of Eyes(G) that includes the root eye; E is a hidden timepoint of G; and P is any macro-projection. Then there exists macro-projections P' and P'', in which $\dot{Y}$ is unchanged for $Y \in \mathcal{E}$, such that (a) $\mathcal{E}$ is UB active for E in P', and (b) $\mathcal{E}$ is LB active for E in P''.*

Proof: Consider the Nature STN instance Γ corresponding to P. Delete the rigid constraints corresponding to the W observations for W not in $\mathcal{E}$. This forms a new STN Γ'.

For part (a), by (Dechter, Meiri, and Pearl, 1991), there is a solution of Γ' where $\dot{W} = \text{UB}(X,W)$ for W not in $\mathcal{E}$. Consequently, we can add to Γ' rigid constraints of the form $\dot{W} = \text{UB}(X,W)$, for W not in $\mathcal{E}$, without creating an inconsistency, and we form P' accordingly.

Note that in Γ', for W not in $\mathcal{E}$,

$$\text{UB}(X,E) \le \text{UB}(X,W) + \text{UB}(W,E) \le \text{UB}(X,W) + \text{CUB}(W,E)$$

using the triangle inequality. It follows that, in P',

$$\min_{Y\in\mathcal{E}}(\dot{Y}+\text{CUB}(Y,E)) \le \dot{W} + \text{CUB}(W,E)$$

for W not in $\mathcal{E}$. Thus, $\mathcal{E}$ is UB active for E in P'.

For part (b), we analogously use the (Dechter, Meiri, and Pearl, 1991) solution where $\dot{W} = \text{LB}(X,W)$ to form P''. In Γ' we have

$$\text{LB}(X,E) \ge \text{LB}(X,W) + \text{LB}(W,E) \ge \text{LB}(X,W) + \text{CLB}(W,E)$$

so in P'' we have

$$\max_{Y \in \mathcal{E}} (\dot{Y} + \mathrm{CLB}(Y, E)) \geq \dot{W} + \mathrm{CLB}(W, E)$$

and thus $\mathcal{E}$ is LB active for E in P''. $\square$

6.3.1 Label Notation

For the rest of the paper, it is convenient to introduce a more compact notation for generalized labels, as follows. Suppose I is an index set for eyes(G). We can rewrite the initial labels as

$$\ell_{\mathrm{hi}+}(E) = -\min(v_0, \min_{i \in I}(\dot{Y}_i + v_i))$$

where $v_0 = \mathrm{CUB}(X, E) + \delta_{\mathrm{hi}}(E)$ and $v_i = \mathrm{CUB}(Y_i, E) + \delta_{\mathrm{hi}}(E)$. Similarly, we can rewrite

$$\ell_{\mathrm{lo}+}(E) = \max(u_0, \max_{j \in I}(\dot{Y}_j + u_j))$$

where $u_0 = \mathrm{CLB}(X, E) + \delta_{\mathrm{lo}}(E)$ and $u_i = \mathrm{CLB}(Y_i, E) + \delta_{\mathrm{lo}}(E)$.

This form of notation will be useful for derived labels as well as the initial labels originating from the Algorithm 1 construction (figure 1). It also makes it easier to treat the contribution from the root eye specially.

6.3.2 Observability Tightening for Initial Edges

The following are important results that exploit the sensitivity of dynamic strategies to observability. They emphasize the asymetry of dynamic controllability with respect to the directionality of time (in contrast to weak and strong controllability). Here we establish the results for the initial labels (before applying any reduction rules).

Suppose $q \geq 0$. Intuitively, a dynamic strategy cannot directly know a condition of the form $(\geq \dot{Y} - q)$ is satisfied until Y is observed. Consequently, $(\geq \dot{Y} - q)$ should be dynamically tightened to $(\geq \dot{Y})$. On the other hand, a condition of the form $(\leq \dot{Y} - q)$ is already false by the time Y is observed; thus, it should be regarded as not (directly) dynamically satisfiable.

The following lemmas make this precise. As context, we suppose Z is an executable timepoint and X is the root of some hidden group for which $\{Y_i\}$ is the indexed set of leaf eyes.

Lemma 9 *Suppose*

$$(-\ell_{\mathrm{hi}+}(E)) = \min(v_0, \min_j(\dot{Y}_j + v_j))$$

and a viable dynamic strategy satisfies $Z - X \geq (-\ell_{\mathrm{hi}+}(E)) - q$ for all projections, where $0 \leq q < v_0$, Then the strategy must also satisfy

$$Z - X \geq \min(v_0 - q, \min_j(\dot{Y}_j + \max(0, v_j - q)))$$

Proof: We are assuming $Z - X \geq (-\ell_{\mathrm{hi+}}(E)) - q$, that is, $Z - X \geq \min(v_0 - q, \max_i(\dot{Y}_i + v_i - q))$, for all projections.

To simplify notation, we set $q_i = v_i - q$. Then $q_0 \geq 0$, and $Z - X \geq \min(q_0, \min_i(\dot{Y}_i + q_i))$ for all projections. We need to show this implies

$$Z - X \geq \min(q_0, \min_i(\dot{Y}_i + q_i^+))$$

for all projections, where $q_i^+ = \max(0, q_i)$.

Our approach will be to show that if there is a projection that does not satisfy the condition, then there is a dynamically indistinguishable projection with a contradictory property. This establishes that the condition is satisfied for all projections.

Suppose P is a projection that does not satisfy the condition. Then

$$\min(q_0, \min_i(\dot{Y}_i + q_i)) \leq Z^P - X < \min(q_0, \min_i(\dot{Y}_i + q_i^+))$$

in P, where Z^P is the time assigned to Z in P by the viable dynamic strategy. It follows that $Z^P - X < \min(q_0, \min_{q_i \geq 0}(\dot{Y}_i + q_i))$ and $Z^P - X < \min_{q_i < 0}(\dot{Y}_i)$.

Let $\mathcal{E}$ be the subset of Eyes(G) defined by $\mathcal{E} = \{Y_i : Z^P - X < \dot{Y}_i + q_i \text{ in } P\}$. Thus, $\mathcal{E}$ includes the root eye $X = Y_0$ and is a superset of $\{Y_i : q_i \geq 0\}$ by our assumption. Since $Z^P - X < \min_{q_i < 0}(\dot{Y}_i)$, the Z^P value is assigned before the Y_i not in $\mathcal{E}$ have been observed.

By lemma 8 part (a), there is another projection P' where $\dot{Y}_i$ is unchanged for $Y_i \in \mathcal{E}$ and $\mathcal{E}$ is UB active. Let $M = \min_{Y_i \in \mathcal{E}}(\dot{Y}_i + q_i)$. (Note that the value of M is the same in P' and P.) Thus, $Z^P - X < M \leq \dot{Y}_i + q_i$ in P' for all i, including the Y_i not in $\mathcal{E}$ (since $\mathcal{E}$ is UB active). However,$\dot{Y}_i + q_i \leq Z^P - X$ in P for the Y_i not in $\mathcal{E}$. Thus, $\dot{Y}_i$ is increased in the transition from P to P' for $Y_i \notin \mathcal{E}$ (and is unchanged otherwise). Consequently, at time Z^P, the observables that are in the future in P are also in the furure in P', and those in the past are unchanged.

It follows that the pre-history at time Z^P is the same in P' as it is in P. Since the strategy is dynamic, therefore $Z^{P'} = Z^P$. Thus, $Z^{P'} - X < \dot{Y}_i + q_i$ in P' for all i, but that violates the constraint that $Z - X \geq \min(q_0, \min_i(\dot{Y}_i + q_i))$ for all projections. This contradiction establishes the result. $\Box$

Lemma 10 *Suppose*

$$\ell_{\mathrm{lo+}}(E)) = \max(u_0, \max_i(\dot{Y}_i + u_i))$$

and a viable dynamic strategy satisfies $Z - X \leq \ell_{\mathrm{lo+}}(E) - q$ for all projections, where $0 \leq q \leq u_0$. Then it must also satisfy

$$Z - X \leq \max(u_0 - q, \max_{u_i \geq q}(\dot{Y}_i + u_i - q))$$

(i.e., the terms where $(u_i - q)$ is negative may be dropped).

Proof: We are assuming $Z - X \leq \ell_{\mathrm{lo+}}(E) - q$, that is, $Z - X \leq \max(u_0 - q, \max_i(\dot{Y}_i + u_i - q))$, for all projections. The proof has similar steps to lemma 9 but there is an added complication, as we will see.

To simplify notation, we set $q_i = u_i - q$. Then $q_0 \geq 0$, and $Z - X \leq \max(q_0, \max_i(\dot{Y}_i + q_i))$ for all projections. We need to show this implies

$$Z - X \leq \max(q_0, \max_{q_i \geq 0}(\dot{Y}_i + q_i))$$

for all projections.

Our approach will be to show that if there is a projection that does not satisfy the condition, then there is a dynamically indistinguishable projection with a contradictory property. This establishes that the condition is satisfied for all projections.

Suppose P is a projection that does not satisfy the condition. Then

$$\max(q_0, \max_{q_i \geq 0}(\dot{Y}_i + q_i)) < Z^P - X \leq \max(q_0, \max_i(\dot{Y}_i + q_i))$$

in P, where Z^P is the time assigned to Z by the viable dynamic strategy.

Let $\mathcal{E}$ be the subset of Eyes(G) defined by $\mathcal{E} = \{Y_i : \dot{Y}_i + q_i < Z^P - X \text{ in } P\}$. Thus, $\mathcal{E}$ includes the root eye $X = Y_0$ and is a superset of $\{Y_i : q_i \geq 0\}$ by our assumption. Note also that $Z^P - X \leq \dot{Y}_i + q_i < \dot{Y}_i$ in P for any Y_i not in $\mathcal{E}$. Thus, Z^P is assigned before these Y_i have been observed. Without loss of generality, we may assume that $\mathcal{E}$ has a maximum size among projections P that do not satisfy the condition.

By lemma 8 part (b), there is another projection P' where $\dot{Y}_i$ is unchanged for $Y_i \in \mathcal{E}$ and $\mathcal{E}$ is LB active for E. Let $M = \max_{Y_i \in \mathcal{E}}(\dot{Y}_i + q_i)$. (Note that the value of M is the same in P' and P.) Thus, $\dot{Y}_i + q_i \leq M < Z^P - X$ in P' for all Y_i, including the Y_i not in $\mathcal{E}$. However, $Z^P - X \leq \dot{Y}_i + q_i$ in P for the Y_i not in $\mathcal{E}$. Thus, $\dot{Y}_i$ is reduced in the transition from P to P' for $Y_i \notin \mathcal{E}$ (and is unchanged otherwise).

Since the $\dot{Y}_i$ are reduced, rather than increased as in the proof of lemma 9, we require additional work to obtain a suitable projection indistinguishable from P.

Note that the projections of a POSTNU form a convex set. Thus, we can form a convex linear combination $P'' = \mu * P' + (1 - \mu) * P$ for any value μ such that $0 \leq \mu \leq 1$, where any contingent link that has duration d' in P' and duration d in P would have duration $d'' = \mu * d' + (1 - \mu) * d$ in P''. (Note that $d' = d$ implies $d'' = d' = d$.)

As μ transitions from 0 to 1, the intermediate projection P'' will transition between the properties of P and those of P'. Set $y_q = \min_{Y_i \notin \mathcal{E}}(\dot{Y}_i + q_i)$ and $y = \min_{Y_i \notin \mathcal{E}}(\dot{Y}_i)$. Note that $y_q < y$ (since the q_i are negative).. Then, as μ increases, both y_q and y move towards $Z^P - X$, and y_q eventually just passes it while y is still behind. Thus, we can choose μ and hence P'' such that: (1) $Z^P - X < \dot{Y}_i$ for all Y_i not in $\mathcal{E}$; and (2) $\dot{Y}_i + q_i < Z^P - X$ for some Y_i not in $\mathcal{E}$.

By (1) it follows that the pre-history at time Z^P is the same in P'' as it is in P. Since the strategy is dynamic, therefore $Z^{P''} = Z^P$. With respect to (2), note that $\dot{Y}_i + q_i < Z^P - X$ cannot be true for all Y_i not in $\mathcal{E}$, since that would violate the constraint that $Z - X \leq \max(u_0 - q, \max_i(\dot{Y}_i + u_i - q))$ for all projections. It follows that P'' would provide a counterexample to the lemma where $\mathcal{E}$ has a larger size than in P, which violates the maximum assumption. This contradiction establishes the result.

□

We may use the terms "max observability tightening" and "min observability tightening," respectively, for the types of observability tightening sanctioned by lemmas 10 and 9, respectively.

7 Generalized Reduction Rules

The reduction rules that we will introduce can be interpreted as generalizations of the "classical" STNU reduction rules. However, it is convenient to factor them in a slightly different way that makes more sense in the generalized context. In particular, the reduction rule applications will involve two separate types of transformation: (1) the addition of a logically implied constraint that leaves the entire set of viable strategies unchanged; and (2) an *observability tightening* step that may filter out some viable strategies, but not any dynamic viable strategies. The second step, which is mandatory when applicable, captures the removal of lower-case labels that is inherent in the classical Lower-Case and Cross-Case Reductions, and is also implicit in the classical Upper-Case Reduction. Its function is to replace derived constraints that are not observable by minimally strengthened constraints that are observable.

The overall purpose of STNU reduction rules is to try to make every projection dispatchable, since this implies Dynamic Controllability (Morris, Muscettola, and Vidal, 2001; Shah et al., 2007; Morris, 2014). In an STN, dispatchability can be achieved by applying "Plus/Minus" operations that compose non-negative edges with following negative edges. The aim is to ensure the path constraints are enforced by "vee-paths" where any non-negative edges follow any negative edges. If this process terminates without producing a negative cycle, then the resulting network is dispatchable (Morris, 2016). For an STNU, the negative edges include upper-case edges not subject to Label Removal, and the non-negative edges include lower-case edges. In effect, the reductions constitute the operations needed to create vee-paths, and their action is global across all of the projections.

The same approach is applicable to the generalized labels considered here. Thus, we have analogues of the Upper-Case, Lower-Case, and Cross-Case reductions, together with accompanying observability tightening steps. Actually, edges with ordinary numeric weights may be regarded as special cases of the generalized labels where the sets of leaf eyes are empty. Thus, we need only consider the generalized Cross-Case reduction.

The Cross-Case reduction involves a composition of a negative labeled edge (aka Upper Case) and a non-negative labeled edge (aka Lower Case), where the timepoints are either executable or observable, as follows:

$$\max(u_0, \max_i(\dot{Y}_i + u_i)) - \min(v_0, \min_j(\dot{W}_j + v_j))$$

where i and j range over suitable index sets for leaf eyes. The input edges to the reduction will be called the *parent* edges, and the resulting edge is the *daughter* edge. The u_i and v_i values for the initial edges are determined by the construction of the labeled distance graph. However, the derived edges will have derived values computed according the Cross-Case reduction rules. Nevertheless, certain properties of the labels will be maintained, as we see below.

7.1 Edges and Invariants

There are several invariants that hold in the original graph, and will be preserved by the reductions: These differ depending on whether the edge is Upper or Lower Case. Recall that the observables ($\dot{Y}_i$, etc.) are always positive.

7.1.1 Upper Case Edges

The label of an upper case edge A $\rightarrow$ B has the form

$$-\min(v_0, \min_j(\dot{W}_j + v_j))$$

with invariants:

1. The label can be expressed as

   $$-\min(v_0' - q, \min_j(\dot{W}_j + \max(0, v_j' - q)))$$

   where $-\min(v_0', \min_j(\dot{W}_j + v_j'))$ is one of the original upper case edges in the labeled distance graph as constructed by algorithm 1, and $v_0' > q \geq 0$.

2. $v_0 > 0$ and $\forall j \neq 0 : v_j \geq 0$. From previous item. Also implies $\dot{W}_j + v_j > 0$

3. The edge is negative. Implied by the above.

4. $v_0 \geq v_j$. Initially true (lemma 6) and then implied by the above.

5. Given B the target of this edge, then B is also the root observable timepoint for each $\dot{W}_j$ in the label. Initially true, and follows from details of the reductions.

7.1.2 Lower Case Edges

The label of a lower case edge A $\rightarrow$ B has the form

$$\max(u_0, \max_i(\dot{Y}_i + u_i))$$

with invariants:

1. The label can be expressed as

   $$\max(u_0' - q, \max_{u_i' \geq q}(\dot{Y}_i + u_i' - q))$$

   where $\max(u_0', \max_i(\dot{Y}_i + u_i'))$ is one of the original lower case edges as constructed by algorithm 1, and $u_0' \geq q \geq 0$.

2. $\forall i : u_i \geq 0$ (includes $i = 0$). From previous item. Also implies $\dot{Y}_i + u_i > 0$.

3. The edge is non-negative. Implied by the above invariants.

4. $u_0 \geq u_j$. Initially true (lemma 6) and then implied by the above.

5. Given A the source timepoint of this edge, A is also the root observable timepoint for each $\dot{Y}_i$ in the label. This will follow from the details of the reductions.

7.2 Observability Tightening for Derived Edges

The intuition behind observability tightening is as follows: if an execution constraint involves an event that cannot be observed when it is needed, then a dynamic strategy may instead satisfy a suitably tightened constraint that is observable. For example, suppose there is a deadline of $\max(A, B-5)$ for some timepoint Z. If B could happen at any time, but its time of occurrence is not known in advance, then it would be unsafe to wait until after A to execute Z, i.e., the deadline should be tightened to A. Similarly, a lower bound of $\min(A, B-5)$ might need to be tightened to $\min(A, B)$, depending on when $B-5$ can occur and be known to have occurred. These intuitions formed the basis for the reductions in the original dynamic controllability work (Morris, Muscettola, and Vidal, 2001).

Lemma 10 and lemma 9 established an observability tightening property for the initial generalized labels. We show here that the property continues to hold for the derived labels. The following notation will be useful. Given max and min labels $\ell_1 = \max(u_0, \max_i(\dot{Y}_i + u_i))$ and $\ell_2 = -\min(v_0, \min_j(\dot{W}_j + v_j))$ respectively, define

$$\text{MAXOT}(\ell_1) = \max(u_0, \max_{u_i \geq 0}(\dot{Y}_i + u_i))$$
$$\text{MINOT}(\ell_2) = -\min(v_0, \min_j(\dot{W}_j + \max(v_j, 0))).$$

The first Upper-Case Edge invariant then says the label can be expressed as MINOT$(\ell_1 + q)$ for a suitable q where ℓ_1 is the label of an initial max edge. Similarly the first Lower-Case Edge invariant says the label can be expressed as MAXOT$(\ell_2 - q)$ for suitable q and initial label ℓ_2,

Notice that

$$\text{MAXOT}[\,\text{MAXOT}(\ell - q1) - q2\,] = \text{MAXOT}(\ell - q1 - q2)$$
$$\text{MINOT}[\,\text{MINOT}(\ell + q1) + q2\,] = \text{MINOT}(\ell + q1 + q2)$$

for labels ℓ of the relevant max/min form. Thus, if the invariants hold, we can extend the applicability of lemma 10 and lemma 9 to derived labels.

It is convenient to refer to the u_0 and v_0 values as the *scalar* terms in the labels. Notice that observability tightening does not involve any modification of the scalar terms.. We may regard them as preserving a record of the q offset to the initial label.

We have an extension of lemma 7 to derived labels, assuming the invariants hold.

Lemma 11 *The derived label* $-\min(v_0, \min_j(\dot{W}_j + v_j))$ *reduces to* $-v_0$ *in the* P_{sup} *special projection. The derived label* $\max(u_0, \max_i(\dot{Y}_i + u_i))$ *reduces to* u_0 *in the* P_{inf} *special projection.*

Proof: First, note that the q modification applied uniformly to all terms does not disturb the special projections property, which holds for the initial labels. Second, if $v_0 = \min(v_0, \min_j(\dot{W}_j + v_j))$ then

$$v_0 \leq \min_j(\dot{W}_j + v_j) \leq \min_j(\dot{W}_j + \max(v_j, 0))$$

so the MINOT operation does not affect the property. Similarly, if $u_0 = \max(u_0, \max_i(\dot{Y}_i + u_i))$ then

$$u_0 \geq \max_i(\dot{Y}_i + u_i) \geq \max_{u_i \geq 0}(\dot{Y}_i + u_i)$$

so the MAXOT operation also preserves the property. □

7.3 Cross Case Reduction

In general, an expression of the form $\max_i(p_i) - \min_j(q_j)$ can be rewritten as $\max_{i,j}(p_i - q_j)$. Thus, in order for the result to be negative, *every one* of the $(p_i - q_j)$ terms must be negative; otherwise the result is non-negative. If the result is negative, it can be rewritten as $-min_{i,j}(q_j - p_i)$.

The cross case reduction requires combining two edges with max and min labels:

$$\mathrm{A} \xrightarrow{\max(u_0, \max_i(\dot{Y}_i + u_i))} \mathrm{B} \xrightarrow{-\min(v_0, \min_j(\dot{W}_j + v_j))} \mathrm{C} \tag{3}$$

into a single upper case or lower case edge A → C whose label should derive from:

$$\max(u_0, \max_i(\dot{Y}_i + u_i)) - \min(v_0, \min_j(\dot{W}_j + v_j)) \tag{4}$$

Same Root Our first task in composing the max and min labels is to deal with the case where the source timepoint of the edge with the max label coincides with the target of the edge with the min label (i.e. $A = C$ in Eq. (3)). We will call this the *same-root* case, since the root eye is the same for both labels. According to the invariants, this will also be the root timepoint for each of the $\dot{Y}_i$ and $\dot{W}_j$ observables.

In the same-root case, we are only really interested in determining whether the sign of the combination is negative or non-negative. If negative, then the POSTNU is not Dynamically Controllable. Otherwise the result derives a non-negative distance from the root timepoint to itself, which can be discarded as redundant.[6]

The combination can be conveniently written as

$$\max_{i \geq 0, j \geq 0}(\dot{Y}_i + (u_i - v_j) - \dot{W}_j) \tag{5}$$

where $\dot{Y}_0 = \dot{W}_0 = 0$. If any of the terms in the maximization is non-negative, then the combination is also. This sub-problem can be stated negatively as determining whether there is no macro-projection, which can be limited to the observables in the hidden group, such that:

$$\forall i, j : \dot{Y}_i - \dot{W}_j < v_j - u_i.$$

Given that $\dot{Y}$ is the shorthand for $Y - X$ where X is the root observable for Y and that all Y_i and W_j have the same root observable, the above inequalities can be rewritten as $Y_i - W_i < v_j - u_i$ which are essentially Simple Temporal constraints except that the inequalities are strict. The bounds on the micro-links may also be viewed as Simple Temporal constraints, and the problem reduces to determining whether this partially

[6] This is actually a generalization of the STNU criterion that the Cross-Case Reduction is not applicable when the lower and upper case edges involve the same label, since there $u_i = v_j = 0$.

strict STN has a solution. A standard STN algorithm such as Bellman-Ford (Cormen, Leiserson, and Rivest, 1990) can be used to look for either an explicit negative cycle, or an implied negative cycle where the bounds sum to zero and the cycle passes through one or more of the strict edges.[7] Note that if this derived STN is consistent, then the POSTNU is NOT Dynamically Controllable. If it is not consistent, then the same-root combination is discarded.

Different Root We now turn out attention to the non-same-root case (i.e. $A \neq C$ in Eq. (3)). This implies that the observables in the max label come from a different hidden group than those in the min label. Thus, there are no cross-correlations, durations of micro-links vary independently between the two groups, and we may speak of a *local projection* of a hidden group.

The cross-independence simplifies the analysis compared to the same-root case. In particular, observability tightening can be applied for an observable in one group provided only that it is applicable in *any* local projection of the other group.

The result of the

$$\max(u_0, \max_i(\dot{Y}_i + u_i)) - \min(v_0, \min_j(\dot{W}_j + v_j))$$

combination in the non-same-root case is as follows.

NEGATIVE RESULT (Upper Case): If $u_0 < v_0$, the result is

$$-\min(v_0 - u_0, \min_j(\dot{W}_j + \max(v_j - u_0, 0))).$$

Note the observation variables are all inherited from the negative parent edge. Since the target of the edge is also inherited from the negative parent, this preserves the invariant that it coincides with the root timepoint of the observables. Also note the result may be expressed as MINOT$(\ell + u_0)$ where ℓ is the negative parent label, preserving the representation invariant.

NON-NEGATIVE RESULT (Lower Case): If $u_0 \geq v_0$, the result is

$$\max(u_0 - v_0, \max_{u_i \geq v_0}(\dot{Y}_i + u_i - v_0))$$

Note the observation variables are now inherited from the non-negative parent edge. Since the source of the edge is also inherited from the non-negative parent, this preserves the invariant that it coincides with the root timepoint of the observables. Also note the result may be expressed as MAXOT$(\ell - v_0)$ where ℓ is the negative parent label, preserving the representation invariant.

The non-negative case may be viewed as a generalization of the STNU reduction cases that reduce away a lower-case edge. Here, the result is a lower-case edge that has fewer observation variables and/or smaller offsets.

[7] In practice, one could simply decrement the strict bounds by a tiny amount and use standard STN methods.

Theorem 1 *The Cross Case reduction rules are sound in both the negative and non-negative result cases.*

Proof:

Recall that we have $0 \leq u_i \leq u_0$ and $0 \leq v_j \leq v_0$ as invariants. In the following, we will use the symbol '$\approx$' to indicate a non-equivalent transformation step that nevertheless preserves the set of viable dynamic strategies because of observability tightening.

We are composing the max/min labels

$$\max(u_0, \max_{i>0}(\dot{Y}_i + u_i)) - \min(v_0, \min_{j>0}(\dot{W}_j + v_j))$$

corresponding to a X1 $\rightarrow$ Z $\rightarrow$ X2 combination.

First suppose $u_0 < v_0$. A viable dynamic strategy must satisfy the composed edge in the P_{inf} local projection for the Y_i group where the max label reduces to u_0 (lemma 11). Thus, it must satisfy

$$\begin{aligned} X2 - X1 &\leq u_0 - \min(v_0, \min_{j>0}(\dot{W}_j + v_j)) \\ &= -\min(v_0 - u_0, \min_{j>0}(\dot{W}_j + v_j - u_0)) \\ &\approx -\min(v_0 - u_0, \min_{j>0}(\dot{W}_j + \max(v_j - u_0, 0))). \end{aligned}$$

Next suppose $u_0 \geq v_0$. Then a viable dynamic strategy must satisfy the composed edge in the P_{sup} local projection for the W_j group where the min label reduces to v_0 (lemma 11). Thus, it must satisfy

$$\begin{aligned} X2 - X1 &\leq \max(u_0, \max_{i>0}(\dot{Y}_i + u_i)) - v_0 \\ &= \max(u_0 - v_0, \max_{i>0}(\dot{Y}_i + u_i - v_0)) \\ &\approx \max(u_0 - v_0, \max_{u_i \geq v_0}(\dot{Y}_i + u_i - v_0)) \end{aligned}$$

This establishes the soundness of the reductions in both the negative and non-negative result cases. $\square$

8 Completeness

Previous sections established the soundness of the generalized reduction rules. This section considers the issue of whether they can support a complete procedure for determining dynamic controllability for POSTNU networks. In this context, the hidden timepoints are irrelevant; we confine our attention to the non-hidden timepoints and their edges, including the labeled edges that have been added either by the preprocessing step, or by the reductions. The graph formed by these edges will be refered to as the *labeled distance graph*. Each macro-projection (hereafter called simply a projection) instantiates this labeled graph to the distance graph of an STN..

We have already seen, in the same-root case, that the reductions may lead to a negative self-loop. In the general case, they may also produce a negative cycle. If the reductions result in a negative cycle,[8] the POSTNU cannot be dynamically controllable. Since the reductions have been shown to preserve the set of viable dynamic strategies, derivation of a negative cycle implies that set is empty.

The remaining possibilities are that they terminate without producing a negative cycle, or do not terminate. We first deal with the non-termination possibility. This leads us to consider what it means to terminate. We will say the labeled distance graph is *quiescent* if there are no negative cycles, and either no reductions are applicable or any further reductions would not produce any new edges.

Lemma 12 *If the reductions do not produce a negative cycle, they will eventually result in quiescence.*

Proof: The new edges added by the reductions involve labels that either have fewer terms or smaller non-negative offsets. Furthermore, there is a lower bound on the size of the decrements.[9] Thus, the network must reach quiescence in a finite number of steps. □

Our main result in this section will be to show that reaching quiescence implies that the POSTNU is dynamically controllable. We will do this by proving the existence of a viable dynamic strategy. Recall that a strategy is a mapping from projections to schedules. The following lemmas will be helful for this purpose. We will say an STN distance-graph is *plus/minus closed* if composing a non-negative edge with a following negative edge does not yield a new or tighter edge.

Lemma 13 *When quiescence is reached for the POSTNU distance graph, all its projections are plus/minus closed.*

Proof: Consider a non-negative edge with a following negative edge in the distance graph of a projection. There will be a corresponding edge pair in the POSTNU distance graph. If the POSTNU distance graph is quiescent, then the cross-case reduction will have produced a labeled edge. That will correspond to an edge in the projection that is at least as tight as the composition of the original edges. (It may be tighter because of observability tightening.) □

Concepts from Morris (2014, 2016) that relate dynamic controllability of STNUs to dispatchability of their projections also play an important role for POSTNUs, using in particular the following definition and lemma.

Definition 4 *Given an STN distance graph, a vee-path is a path that consists of zero or more negative edges followed by zero or more non-negative edges.*

Notice that if a vee-path has negative total length, then it must begin with a negative edge.

Lemma 14 *If a consistent STN is plus/minus closed, then every path constraint is enforced by a vee-path. In that case the STN is said to be* dispatchable.

[8] To reduce verbosity, we will regard a negative self-loop as a special case of a negative cycle.

[9] Consider the smallest positive value of the minimal set that contains the magnitudes of all the original bounds and is closed under subtraction.

Proof: It is easy to see that repeated application of plus/minus compositions to a path will eventually result in a vee-path. □

Notice that if an inconsistent STN is plus/minus closed, then it has a cycle with all negative edges. (A similar observation was made in Nilsson, Kvarnström, and Doherty (2015) for STNUs.) Next we consider the concept of an earliest time schedule for a consistent STN. This has particular significance in the case of STNs that are plus/minus closed.

Definition 5 *Dechter, Meiri, and Pearl (1991) The* earliest-time schedule *S for an STN is obtained as follows. An initial node $\odot$ is added, and for each node X in the original STN we add an edge from X to $\odot$ of weight 0. Then $S(X)$ is defined as $-D(X)$ where $D(X)$ is the shortest-path distance from X to $\odot$.*

The node $\odot$ and its edges are merely used for convenience in defining S; they should not be considered part of the STN (although we may mention them in reasoning about the properties of S).

Note that since every node has a direct edge to $\odot$ of length 0, $S(X) \geq 0$ for all X. The following result has independent interest as well as being useful here.

Theorem 2 *For a dispatchable STN (i.e., a consistent STN that is plus/minus closed) and the earliest-time schedule S, every timepoint X either satisfies $S(X) = 0$ or there is a shortest path of all negative edges from X to some node Y such that $S(Y) = 0$.*

Proof: Suppose $S(X) > 0$. Then $D(X) < 0$ and there is a shortest path from X to $\odot$ that is negative. Since the STN is plus/minus closed, we can assume that the subpath up to the final edge is a vee-path. Note that the final edge to $\odot$ has a zero weight. Thus, the start timepoint Y of the final edge must satisfy $S(Y) = 0$. Without loss of generality, we may assume that Y is the node closest to X such that $S(Y) = 0$. Then $D(Z) < 0$ for every intermediate node Z on the path. Since the path from X to Y is a vee-path, it follows that every edge on that path has a negative weight. □

This remarkable result implies that for a dispatchable STN the non-negative edges are irrelevant as far as the earliest time schedule is concerned. It is only if one deviates from the earliest time schedule that the non-negative edges (and deadlines) matter.

Returning our attention to the POSTNU problem, we are now in a position to define our candidate strategy.

Definition 6 *Given a quiescent distance graph for a POSTNU, the* earliest time *execution strategy $\breve{S}$ (pronounced "S-breve") is defined for any projection p by $\breve{S}(p) = \breve{S}_p$ where $\breve{S}_p$ is the earliest time schedule for p.*

This leads us to our main theorem.

Theorem 3 *If the generalized reductions produce a quiescent distance graph, then the POSTNU is dynamically controllable.*

Proof: We will show that the earliest time execution strategy $\breve{S}$ is both viable and dynamic. Viability is immediate since the earliest time schedule for an STN is known to be a solution Dechter, Meiri, and Pearl (1991).

To establish that $\breve{S}$ is dynamic consider a projection p and an executable timepoint X. Suppose $\breve{S}_p(X) = t$. Let p' be any other projection such that $\breve{S}_p(\trianglelefteq t) = \breve{S}_{p'}(\trianglelefteq t)$. We must show that $\breve{S}_p(X) = \breve{S}_{p'}(X)$.

By theorem 2 the earliest-times schedules for p and p' (and thus the values of $\breve{S}_p(X)$ and $\breve{S}_{p'}(X)$) depend only on the negative edges in their distance graphs. Furthermore, the negative paths emanating from X necessarily lead to timepoints that are in the past (in any schedule). Since $\breve{S}_p(\trianglelefteq t) = \breve{S}_{p'}(\trianglelefteq t)$, any negative edges in those paths that arise from macro-links must have the same weights in p and p'. Also, any edges that arise from negative labeled edges have their weights determined by macro-links that have already finished by time t. Thus, the negative paths emanating from X in p are also there in p'.

The only possible confounding issue might be if an outgoing edge from X could be non-negative in p but negative in p'. Only the labeled edges can vary between projections, and this cannot be the case for a labeled edge: a label that has received a definite value by time t will continue to have that value thereafter.

We conclude that $\breve{S}_p(X) = \breve{S}_{p'}(X)$. $\square$

Corollary 3.1 *If a POSTNU is not dynamically controllable, then repeated application of the generalized reductions will lead to a negative cycle.*

Proof: Suppose the POSTNU is not dynamically controllable. By lemma 12, the reductions must terminate. If the reductions result in quiescence, the network would be dynamically controllable by the theorem, which is a contradiction. Thus, the reductions must lead to a negative cycle as the only other option. $\square$

It is of interest to consider a possible implementation design for the earliest time execution strategy.

8.0.1 Earliest Time Execution Design

We assume a separate thread is assigned to an executable timepoint X to manage its activation. The negative edges emanating from X constitute enabling conditions. The thread keeps a count of how conditions are still unsatisfied and executes X as soon as the count becomes zero. From the point of view of X, the target Y of a negative edge from X is something to be observed, even if Y is itself an executable timepoint. In general, therefore the edge will involve an enabling condition of the form

$$\min_{i \in I}(Y_i + u_i)$$

where the Y_i may be observable or executable timepoints (and where I may be a singleton).

The X thread will normally be sleeping. It will wake if any of the Y_i events occur. If that happens, the X thread will initiate a separate "alarm-clock" thread that will "ring" u_i units of time later. At any time there may be several such clocks that are active for a particular edge. The first one to ring will wake the X thread, at which point it will deactivate the other clocks for that edge, and will decrement the enablement count.

It is easy to see that the execution need only perform a bounded amount of active work for each $Y_i + u_i$ condition. Thus, the execution complexity is linear in the number of such conditions. For example, in a class of POSTNUs where the label sizes are bounded (such as STNUs), the complexity would be linear in the number of edges, and thus quadratic in the number of timepoints. This is similar to the complexity for STNUs of the latest time execution approach in Cairo and Rizzi (2017).

9 Closing Remarks

The focus of this paper has been theoretical with the goal of providing a sound and complete process for determining the Dynamic Controllability of partially observable STNUs. An obvious next step would be to adapt the STNU cubic algorithm of Morris (2014) with the goal of efficiently performing the reduction process. A more long-term objective would be to generalize the setting further, and replace Nature by a second agent, thus considering two-agent plans with limited communication. In that extended problem, the second agent might well have a more general STN than the Nature one, and this could affect many of the assumptions used in the current approach.

References

Bit-Monnot; Ghallab, M.; and Ingrand, F. 2016. Which contingent events to observe for the dynamic controllability of a plan. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2016)*, 3038–3044.

Cairo, M., and Rizzi, R. 2017. Dynamic controllability made simple. In *24th International Symposium on Temporal Representation and Reasoning (TIME 2017)*.

Cormen, T.; Leiserson, C.; and Rivest, R. 1990. *Introduction to Algorithms*. Cambridge, MA: MIT press.

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

Hunsberger, L. 2002. Algorithms for a temporal decoupling problem in multi-agent planning. In *Proc. of Eighteenth Nat. Conf. on Artificial Intelligence (AAAI-02)*.

Hunsberger, L. 2009. Fixing the semantics for dynamic controllability and providing a more practical characterization of dynamic execution strategies. In *Proc. of the IEEE 16th International Symposium on Temporal Representation and Reasoning (TIME-2009)*, 155–162.

Hunsberger, L. 2013. Magic loops in simple temporal networks with uncertainty - exploiting structure to speed up dynamic controllability checking. In *ICAART (2)*, 157–170.

Moffitt, M. D. 2007. On the partial observability of temporal uncertainty. In *Proc. of AAAI-07*.

Morris, P., and Muscettola, N. 1999. Managing temporal uncertainty through waypoint controllability. In *Proc. of Sixteenth Int. Joint Conf. on Artificial Intelligence (IJCAI-99)*.

Morris, P., and Muscettola, N. 2005. Dynamic controllability revisited. In *Proc. of AAAI-05*.

Morris, P.; Muscettola, N.; and Vidal, T. 2001. Dynamic control of plans with temporal uncertainty. In *Proc. of IJCAI-01*.

Morris, P. 2014. Dynamic controllability and dispatchability relationships. In *CPAIOR*.

Morris, P. 2016. The mathematics of dispatchability revisited. In *Proceedings of the 26th International Conference on Automated Planning and Scheduling (ICAPS-2016)*. AAAI.

Muscettola, N.; Nayak, P.; Pell, B.; and Williams, B. 1998. Remote agent: to boldly go where no AI system has gone before. *Artificial Intelligence* 103(1-2):5–48.

Nilsson, M.; Kvarnström, J.; and Doherty, P. 2015. Efficient Processing of Simple Temporal Networks with Uncertainty: Algorithms for Dynamic Controllability Verification. *Acta Informatica*.

Shah, J. A.; Stedl, J.; Williams, B. C.; and Robertson, P. 2007. A fast incremental algorithm for maintaining dispatchability of partially controllable plans. In Boddy, M. S.; Fox, M.; and Thibaux, S., eds., *ICAPS*, 296–303. AAAI.

Vidal, T., and Fargier, H. 1999. Handling contingency in temporal constraint networks: from consistency to controllabilities. *JETAI* 11:23–45.

Vidal, T. 2000. Controllability characterization and checking in contingent temporal constraint networks. In *Proc. of Seventh Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'2000)*.