Time-based Dynamic Controllability of Disjunctive Temporal Networks with Uncertainty: A Tree Search Approach with Graph Neural Network Guidance

Kevin Osanlou[1,2,3,4], Jeremy Frank[1], J. Benton[1], Andrei Bursuc[5], Christophe Guettier[2], Eric Jacopin[6] and Tristan Cazenave[3]

[1] NASA Ames Research Center   [2] Safran Electronics & Defense   [3]LAMSADE, Paris-Dauphine
[4] Universities Space Research Association   [5]valeo.ai   [6]CREC Saint-Cyr Coetquidan
{kevin.osanlou, jeremy.d.frank, j.benton}@nasa.gov
{kevin.osanlou, christophe.guettier}@safrangroup.com
andrei.bursuc@valeo.com
eric.jacopin@st-cyr.terre-net.defense.gouv.fr
tristan.cazenave@lamsade.dauphine.fr

Abstract

Scheduling in the presence of uncertainty is an area of interest in artificial intelligence due to the large number of applications. We study the problem of dynamic controllability (DC) of disjunctive temporal networks with uncertainty (DTNU), which seeks a strategy to satisfy all constraints in response to uncontrollable action durations. We introduce a more restricted, stronger form of controllability than DC for DTNUs, time-based dynamic controllability (TDC), and present a tree search approach to determine whether or not a DTNU is TDC. Moreover, we leverage the learning capability of a message passing neural network (MPNN) as a heuristic for tree search guidance. Finally, we conduct experiments for which the tree search shows superior results to state-of-the-art timed-game automata (TGA) based approaches. We observe that using an MPNN for tree search guidance leads to a significant increase in solving performance and scalability to harder DTNU problems.

1 Introduction

Temporal Networks (TN) are a common formalism to represent and reason about temporal constraints over a set of time points (e.g. start/end of activities in a scheduling problem). The Simple Temporal Networks with Uncertainty (STNUs) (Tsamardinos 2002) (Vidal and Fargier 1999) explicitly incorporate qualitative uncertainty into temporal networks. Considerable work has resulted in algorithms to determine whether or not all timepoints can be scheduled, either upfront or reactively, in order to account for uncertainty (e.g. (Morris and Muscettola 2005), (Morris 2014)). In particular, an STNU is *dynamically controllable* (DC) if there is a reactive strategy in which controllable timepoints can be executed either at a specific time, or after observing the occurrence of an uncontrollable timepoint. Cimatti et al. (Cimatti, Micheli, and Roveri 2016) investigate the problem of DC for Disjunctive Temporal Networks with Uncertainty (DTNUs), which generalize STNUs. Figure 1a shows an example of two DTNUs γ and γ' on the left side; a_i are controllable timepoints, u_j are uncontrollable timepoints. Timepoints are variables which can take on any value in $\mathbb{R}$. Constraints be-
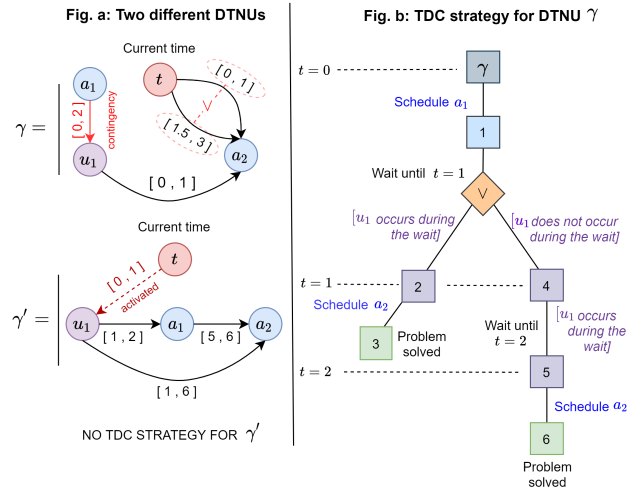
Figure 1: **Two example DTNUs γ and γ'.** In both examples, timepoints a_1 and a_2 are controllable; u_1 is uncontrollable. Black arrows and their intervals (valued in $\mathbb{R}$) represent time constraints between timepoints; the light red arrow and its interval contingency links. The dashed dark red arrow in γ' implies u_1 has already been activated and will occur in the specified interval. A TDC strategy is displayed for γ: the root node of the strategy is γ while other nodes are sub-DTNUs except the $\vee$ node which lists transitional possibilities. DTNU γ', on the other hand, is an example of a DTNU which is DC but not TDC.

tween timepoints characterize both a minimum and maximum time distance separating them, likewise valued in $\mathbb{R}$. The key difference between STNUs and DTNUs lies in the *disjunctions* that yields more choice points for consistent scheduling, especially reactively.

The complexity of DC checking for DTNUs is $PSPACE$-complete (Bhargava and Williams 2019), making this a very challenging problem. The difficulty in proving or disproving DC arises from the need to check all possible combinations of disjuncts in order to handle all possible occurrence outcomes of the uncontrollable timepoints. The best previous approaches for this problem use timed-game automata (TGAs) and Satisfiability Modulo Theories (SMTs), described in (Cimatti, Micheli, and Roveri 2016).

Recently, applications such as image classification have benefited from learning techniques such as Convolutional Neural Networks (CNNs) (Krizhevsky, Sutskever, and Hinton 2012). A new emerging trend of neural networks, graph-based neural networks (GNNs), have been proposed as an extension of CNNs to graph-structured data; recent variants based on spectral graph theory include (Defferrard, Bresson, and Vandergheynst 2016), (Li et al. 2016), (Kipf and Welling 2017). These GNNs take advantage of relational properties between nodes for classification, but do not take into account potential edge weights. In newer approaches, Message Passing Neural Networks (MPNNs) with architectures such as in (Battaglia et al. 2016), (Gilmer et al. 2017) and (Kipf et al. 2018) use embeddings comprising edge weights within each computational layer. We focus our interest on these architecture types as DTNUs can be formalized as graphs with edge distances representing time constraints.

In this work, we study DC checking of DTNUs as a search problem, express states as graphs, and use MPNNs to learn heuristics based on previously solved DTNUs to guide search. The key contributions of our approach are the following. **(1)** We introduce a time-based form of dynamic controllability (TDC) and a tree search approach to identify TDC strategies. We informally show that TDC implies DC, but the opposite is not generally true. **(2)** We define a relevant way of using an MPNN architecture for handling DTNU scheduling problems and use it as heuristic for guidance in the tree search. Moreover, we define a self-supervised training scheme to train the MPNN based on solving randomly generated DTNUs with short timeouts to limit search duration. **(3)** We introduce constraint propagation rules which enable us to enforce time domain restrictions for variables in order to ensure soundness of strategies found. We carry out experiments which show the tree search algorithm offers improved scalability over the best previous DC-solving approach evaluated in (Cimatti, Micheli, and Roveri 2016), PYDC-SMT. Moreover, we show the tree search is improved upon significantly by the learned MPNN heuristic on harder DTNUs.

2 Time-based Dynamic Controllability

A DC *strategy* for a DTNU either executes controllable timepoints at a specific time, or reacts to the occurrence of an uncontrollable timepoint. We present our TDC formalism here. A TDC strategy executes controllable timepoints at specific times under the assumption that some uncontrollable timepoints may occur or not in a given time interval. Each interval in a TDC strategy can have an arbitrary duration. Thus, controllable timepoints are usually executed at the end of the same interval, regardless of if and when uncontrollable timepoints occur inside the interval. Within a given interval, TDC also leaves open the choice to execute a controllable timepoint at the same time as the occurrence of an uncontrollable timepoint, which we call reactive execution.

Nonetheless, TDC is less flexible than a DC strategy which can wait for an uncontrollable timepoint to occur before making a new decision. It does not allow, for example, *delayed* reactive execution of the controllable timepoint.

TDC is a subset of DC, and a stronger form on controllability: TDC implies DC. As described below, representing reactive DC strategies in TDC may require the tree to become arbitrarily large, so DC does not imply TDC. DTNU γ' in figure 1a shows an example of an STNU which is not TDC but DC. In this example, uncontrollable timepoint u_1 is activated, *i.e.* the controllable timepoint associated to u_1 in the contingency links has been executed. Moreover, it is known that u_1 occurs between t and $t + 1$, where t is the current time. The interval $[t, t + 1]$ is referred to as the *activation time interval* for u_1. Controllable timepoint a_1 must be executed at least 1 time unit after u_1, and controllable timepoint a_2 at least 5 time units after a_1. However, controllable timepoint a_2 cannot be executed later than 6 time units after u_1. A valid DC strategy waits for u_1 to occur, then schedules a_1 exactly 1 time unit later, and a_2 5 time units after a_1. However, for any TDC strategy, there is no wait duration small enough while waiting for u_1 to happen that does not violate these constraints. There will always be some strictly positive lapse of time between the moment u_1 occurs and the end of the wait. The exact execution time of u_1 during the wait is unknown: a TDC strategy therefore assumes u_1 happened at the end of the wait when trying to schedule a_1 at the earliest. Therefore, the earliest time a_1 can be scheduled in a TDC strategy is 1 time unit after the end of the wait, which is too late.

3 Tree Search Preliminaries

We introduce here the tree search algorithm. Intuitively, the approach discretizes uncontrollable durations, *i.e.* durations when one or several uncontrollable timepoints can occur, into multiple reduced intervals. These reduced intervals are then used to account for possible outcomes of uncontrollable timepoints and adapt the scheduling strategy accordingly.

The root of the search tree is the DTNU, and other tree nodes are either sub-DTNUs of the DTNU or logical nodes (*OR, AND*) which represent the presence or lack of control over transition into children nodes. At a given DTNU tree node, decisions such as executing a controllable timepoint or waiting for a period of time develop children DTNU nodes for which these decisions are propagated to constraints. The TDC controllability of a *leaf* DTNU, *i.e.* a sub-DTNU for which all controllable timepoints have been executed and uncontrollable timepoints are assumed to have occurred in specific intervals, indicates whether or not this sub-DTNU has been solved at the end of the scheduling process. We also refer to the TDC controllability of a DTNU node in the search tree as its *truth attribute*. Lastly, the search logically combines TDC controllability of children DTNUs to determine TDC controllability for parent nodes. We give a simple example of a TDC strategy for a DTNU γ in figure 1.

Let $\Gamma = \{A, U, C, L\}$ be a DTNU. A is the list of controllable timepoints, U the list of uncontrollable timepoints, C the list of constraints and L the list of contingency links. The root node of the search tree built by the algorithm is Γ. There are four different types of nodes in the tree. Furthermore, each node possesses a *truth* attribute, as explained in §4.4, which is initialized to *unknown* and can be set to either *true* or *false* . The different types of tree nodes are listed

below and shown in figure 2.

DTNU nodes. Any DTNU node other than the original problem Γ corresponds to a sub-problem of Γ at a given point in time t, for which some controllable timepoints may have already been scheduled in upper branches of the tree, some amount of time may have passed and some uncontrollable timepoints may have occurred. A DTNU node is made of the same timepoints A and U, constraints C and contingency links L as the original DTNU Γ. It also carries a schedule memory S of what exact time, or during what time interval, scheduled timepoints were executed during previous decisions in the tree. Lastly, the node also keeps track of the activation time intervals of activated uncontrollable timepoints B. The schedule memory S is used to create an updated list of constraints C' resulting from the propagation of the execution time or execution time interval of timepoints in constraints C as described in §4.5. A non-terminal DTNU node, *i.e.* a DTNU node for which all timepoints have not been scheduled, has exactly one child node: a *d-OR* node.

_OR_ nodes. When a choice can be made at time t, this transition control is represented by an *OR* node. We distinguish two types of such nodes, *d-OR* and *w-OR* . For *d-OR* nodes, the first type of choice available is which controllable timepoint a_i to execute. This leads to a DTNU node. The other type of choice is to wait a period of time (§4.2) which leads to a *WAIT* node. *w-OR* nodes can be used for *reactive wait strategies*, *i.e.* to stipulate that some controllable timepoints will be scheduled reactively during waits (§4.1). The parent of a *w-OR* node is therefore a *WAIT* node and its children are *AND* nodes, described below.

WAIT nodes. These nodes are used after a decision to wait a certain period of time Δ_t. The parent of a *WAIT* node is a *d-OR* node. A *WAIT* node has exactly one child: a *w-OR* node, which has the purpose of exploring different reactive wait strategies. The uncertainty management related to uncontrollable timepoints is handled by *AND* nodes.

AND nodes. This type of node represents a lack of transition control over children nodes. It is used after a wait decision is taken and a reactive wait strategy is decided, represented consecutively by a *WAIT* and *w-OR* node. Each child node of the *AND* node is a DTNU node at time $t+\Delta_t$, where t is the time before the wait and Δ_t the wait duration. Furthermore, each child node represents an outcome of how uncontrollable timepoints may unfold, and is built from the set of *activated* uncontrollable timepoints (uncontrollable timepoints that have been started by the execution of their controllable timepoint) whose occurrence time interval overlaps the wait. If there are l activated uncontrollable timepoints, then there are at most 2^l *AND* node children, representing each element of the power set of activated uncontrollable timepoints (§4.2).

Figure 2 illustrates how a sub-problem of Γ, referred to as $DTNU_{O,P,t}$, is developed. Here, $O \subset A$ is the set of controllable timepoints that have already been executed, $P \subset U$ the set of uncontrollable timepoints which have occurred, and t the time. This root node transitions into a *d-OR* node. The *d-OR* node in turn is developed into several
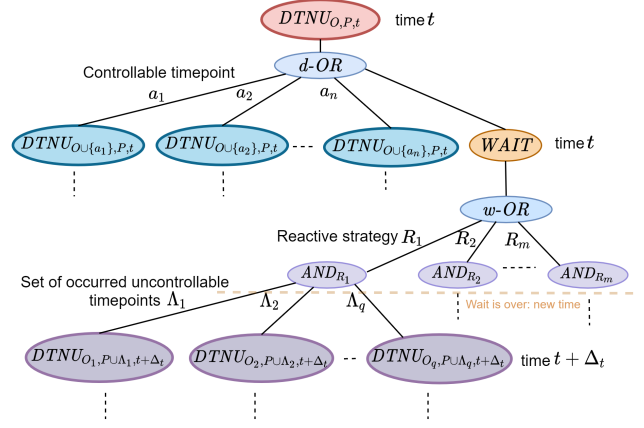

Figure 2: **Basic structure of the search tree describing how a DTNU node $DTNU_{O,P,t}$ is developed.** $DTNU_{O,P,t}$ (placed at the root of the tree) refers to a DTNU where O is the set of controllable timepoints that have already been executed, P the set of uncontrollable timepoints that have occurred, and t the time. Each branch a_i refers to a controllable timepoint a_i, R_i to a reactive strategy during the wait, and Λ_i to a combination of uncontrollable timepoints which can occur during the wait.

children nodes $DTNU_{O \cup \{a_i\},P,t}$ and a *WAIT* node. Each node $DTNU_{O \cup \{a_i\},P,t}$ corresponds to a sub-problem which is obtained from the execution of controllable timepoint a_i at time t. The *WAIT* node refers to the process of waiting a given period of time, Δ_t in the figure, before making the next decision. The *WAIT* node leads directly to a *w-OR* node which lists different wait strategies R_i. If there are l activated uncontrollable timepoints, there are 2^l subsets of uncontrollable timepoints Λ_i that could occur. Each AND_{R_j} node has one sub-problem DTNU for each Λ_i.

Each sub-problem $DTNU_{O_i,P \cup \Lambda_i,t+\Delta_t}$ of the node AND_{R_j} is a DTNU at time $t + \Delta_t$ for which all uncontrollable timepoints in Λ_i are assumed to have happened during the wait period, *i.e.* in the time interval $[t, t + \Delta_t]$. Additionally, some controllable timepoints may have been reactively executed during the wait and may now be included in the set of scheduled controllable timepoints O_i. Otherwise, $O_i = O$.

Two types of leaf nodes exist in the tree. The first type is a node $DTNU_{A,U,t}$ for which all controllable timepoints $a_i \in A$ have been scheduled and all uncontrollable timepoints $u_i \in U$ have occurred. The second type is a node $DTNU_{A \setminus A',U,t}$ for which all uncontrollable timepoints $u_i \in U$ have occurred, but some controllable timepoints $a_i \in A'$ have not been executed. The constraint satisfiability test of the former type of leaf node is straightforward: all execution times of all timepoints are propagated to constraints in the same fashion as in §4.5. The leaf node's truth attribute is set to *true* if all constraints are satisfied, *false* otherwise. For the latter type, we propagate the execution times of all uncontrollable timepoints as well as all scheduled controllable timepoints in the same way, and obtain an updated set of constraint C'. This leaf node, $DTNU_{A \setminus A',U,t}$, is therefore characterized as $\{A', \emptyset, C', \emptyset\}$ and is a DTN. We add the constraints $a_i' \geq t, \forall a_i' \in A'$

and use a mixed integer linear programming solver (Cplex 2009) to solve the DTN. If a solution is found, the execution time values for each $a'_i \in A'$ are stored and the leaf node's truth value is set to *true*. Otherwise, it is set to *false*. After a truth value is assigned to the leaf node, the truth propagation function defined in §4.4 is called to logically infer truth value properties for parent nodes.

Lastly, the search algorithm explores the tree in a depth-first manner. At each *d-OR* , *w-OR* and *AND* node, children nodes are visited in the order they are created. Once a child node is selected, its entire subtree will be processed by the algorithm before the other children are explored. Some simplifications made in the exploration are detailed in §11.6 in the appendix.

4   Tree Search Characteristics

4.1   Reactive scheduling during waits

Some situations may arise when instant scheduling of a controllable timepoint is necessary as soon as an uncontrollable timepoint occurs to satisfy a constraint. We designate as a *conjunct* a constraint relationship of the form $v_i - v_j \in [x, y]$ or $v_i \in [x, y]$, where v_i, v_j are timepoints and $x, y, \in \mathbb{R}$. We refer to a constraint where several conjuncts are linked by $\vee$ operators as a *disjunct*. If at any given DTNU node in the tree there is an activated uncontrollable timepoint u with the potential to occur during the next wait and there is at least one unscheduled controllable timepoint a such that a conjunct of the form $u - a \in [0, y], y \geq 0$ is present in the constraints, a reactive wait strategy is considered that schedules a as soon as u occurs. Let $\Phi = \{\phi_1, \phi_2, ..., \phi_s\} \subset A$ be the complete set of unscheduled controllable timepoints for which there are conjunct clauses $u - \phi_i \in [0, y]$. We denote as $R_1, R_2, ..., R_m$ all possible combinations of elements taken from Φ, including the empty set. As depicted in Figure 2, we account for potential reactive wait strategies by using a *w-OR* node. The child node AND_{R_i} of the *w-OR* node resulting from the combination R_i has a reactive wait strategy for which all controllable timepoints in R_i will be immediately executed at the moment u occurs during the wait, if it does. If u doesn't occur, no controllable timepoint is reactively scheduled during the wait.

4.2   Wait action

When a wait decision of duration Δ_t is taken at time t for a given DTNU node, two different categories of uncontrollable timepoints are considered to account for all transitional possibilities:

- $Z = \{\zeta_1, \zeta_2, ..., \zeta_l\}$ is a set of timepoints that could either happen during the wait, or afterwards, i.e. the end of the activation time interval for each ζ_i is greater than $t + \Delta_t$.

- $H = \{\eta_1, \eta_2, ..., \eta_m\}$ is a set of timepoints that are certain to happen during the wait, i.e. the end of the activation time interval for each η_i is less than or equal to $t + \Delta_t$.

There are $q = 2^l$ number of different possible combinations (empty set included) $V_1, V_2, ..., V_q$ for elements taken from Z. For each combination V_i, the set $\Lambda_i = H \cup V_i$ is

created. The union $\bigcup\limits_{i=1}^{q} \Lambda_i$ refers to all possible combinations of uncontrollable timepoints which can occur by $t + \Delta_t$. In figure 2, for each *AND* node, the combination Λ_i leads to a DTNU sub-problem $DTNU_{O_i, P \cup \Lambda_i, t + \Delta_t}$ for which the uncontrollable timepoints in Λ_i are considered to have occurred between t and $t + \Delta_t$ in the schedule memory S. In addition, any potential controllable timepoint ϕ planned to be instantly scheduled in a reactive wait strategy R_i in response to an uncontrollable timepoint u in Λ_i will also be considered to have been scheduled between t and $t + \Delta_t$ in S. The only exception is when checking constraint satisfiability for the conjunct $u - \phi \in [0, y]$ which required the reactive scheduling, for which we assume ϕ executed at the same time as u, thus the conjunct is considered satisfied automatically.

4.3   Wait Eligibility and Period

The way time is discretized is fundamental and holds direct implications on the search space explored and the capability of the algorithm to find TDC strategies. Longer waits make the search space smaller, but carry the risk of missing key moments where a decision is needed. On the other hand, smaller waits can make the search space too large to explore. We explain when the wait action is eligible, and how the wait duration is computed.

Eligibility   At least one of the following criteria has to be met for a *WAIT* node to be added as child of a *d-OR* node:

- There is at least one activated uncontrollable timepoint for the parent DTNU node.

- There is at least one conjunct of the form $v \in [x, y]$, where v is a timepoint, in the constraints of the parent DTNU node.

These criteria ensure that the search tree will not develop branches below *WAIT* nodes when waiting is not relevant, *i.e.* when a controllable timepoint necessarily needs to be scheduled. It also prevents the tree search from getting stuck in infinite *WAIT* loop cycles.

Wait Period   We define the wait duration Δ_t at a given *d-OR* node eligible for a wait dynamically by examining the updated constraint list C' of the parent DTNU and the activation time intervals B of its activated uncontrollable timepoints. Let t be the current time for this DTNU node. The wait duration is defined by comparing t to elements in C' and B to look for a minimum positive value defined by the following three rules:

First rule   For each activated time interval $u \in [x, y]$ in B, we select $x - t$ or $y - t$, whichever is smaller and positive, and we keep the smallest value δ_1 found over all activated time intervals.

Second rule   For each conjunct $v \in [x, y]$ in C', where v is a timepoint, we select $x - t$ or $y - t$, whichever is smaller and positive, and we keep the smallest value δ_2 found over all conjuncts.

Third rule This rule is used to determine timepoints which need to be scheduled ahead of time by chaining constraints together. Intuitively, when a conjunct $v \in [x, y]$ is in C', it means v has to be executed when $t \in [x, y]$ to satisfy this conjunct. However, v could be linked to other timepoints by constraints which require them to happen before v. These timepoints could in turn be linked to yet other timepoints in the same way, and so on. The purpose of the third rule is to chain backwards to identify potential timepoints which start this chain and potential time intervals in which they need to be executed. The following mechanism is used: for each conjunct $v \in [x, y]$ in C' found in 2), we apply a recursive backward chain function to both (v, x) and (v, y). We detail here how it is applied to (v, x), the process being the same for (v, y). Conjuncts of the form $v - v' \in [x', y'], x' \geq 0$ in C' are searched for. For each conjunct found, we add to a list two elements, $(v', x - x')$ and $(v', x - y')$. We also select $x - x' - t$ or $x - y' - t$, whichever is smaller and positive, as potential minimum candidate. The backward chain function is called recursively on each element of the list, proceeding the same way. We keep the smallest candidate δ_3. Figure 8 in the appendix illustrates an application of this process.

We set $\Delta_t = \min(\delta_1, \delta_2, \delta_3)$ as the wait duration. This duration is stored inside the *WAIT* node.

4.4 Truth Value Propagation

In this section, we describe how truth attributes of nodes are related to each other. The truth attribute of a tree node represents its TDC controllability, and the relationships shared between nodes make it possible to define sound strategies. When a leaf node is assigned a truth attribute β, the tree search is momentarily stopped and a propagator function `PropagateTruth()` is called. This function recursively propagates β onto upper parent nodes. A parent node ω is selected recursively and we distinguish the following cases:

- The parent ω is a DTNU node or a *WAIT* node: ω is assigned β.

- The parent ω is a *d-OR* or *w-OR* node: If $\beta = true$, then ω is assigned *true*. If $\beta = false$ and all children nodes of ω have *false* attributes, ω is assigned *false*. Otherwise, the propagation stops.

- The parent ω is an *AND* node: If $\beta = false$, then ω is assigned *false*. If $\beta = true$ and all children nodes of ω have *true* attributes, ω is assigned *true*. Otherwise, the propagation stops.

After the propagation algorithm finishes, the tree search algorithm resumes where it was temporarily stopped. If a truth attribute has reached the root node of the tree, the tree search algorithm will be swiftly ended due to the branch cuts implemented in §11.6 in the appendix. A *true* attribute reaching the root node of the tree means a TDC strategy has been found. A *false* attribute means none could be found. The pseudocode for the `PropagateTruth()` function is given in Algorithm 1 in the appendix.

4.5 Constraint Propagation

Decisions taken in the tree define when controllable timepoints are executed and also bear consequences on the execution time of uncontrollable timepoints. We explain here how these decisions are propagated into constraints, as well as the concept of '*tight bound*'. Let C' be the list of updated constraints for a DTNU node ψ for which the parent node is ω. We distinguish two cases. Either ω is a *d-OR* node and ψ results from the execution of a controllable timepoint a_i, or ω is an *AND* node and ψ results from a wait of Δ_t time units. In the first case, let t be the execution time of a_i. The updated list C' is built from the constraints of the parent DTNU of ψ in the tree. If a conjunct contains a_i and is of the form $a_i \in [x, y]$, this conjunct is replaced with *true* if $t \in [x, y]$, *false* otherwise. If the conjunct is of the form $v_j - a_i \in [x, y]$, we replace the conjunct with $v_j \in [t + x, t + y]$. The other possibility is that ψ results from a wait of Δ_t time at time t, with a reactive wait strategy R_j. In this case, the new time is $t + \Delta_t$ for ψ. As a result of the wait, some uncontrollable timepoints $u_i \in \Lambda_i$ may occur, and some controllable timepoints $a_i \in R_j$ may be executed reactively during the wait. Let $v_i \in \Lambda_i \cup R_j$ be these timepoints occurring during the wait. The execution time of these timepoints is considered to be in $[t, t + \Delta_t]$. For uncontrollable timepoints $u_i' \in \Lambda_i' \subset \Lambda_i$ for which the activation time ends at $t + \Delta_{t_i}' < t + \Delta_t$, and potential controllable timepoints a_i' instantly reacting to these uncontrollable timepoints, the execution time is further reduced and considered to be in $[t, t + \Delta_{t_i}']$. We define a concept of *tight bound* to update constraints which restricts time intervals in order to account for all possible values v_i can take between t and $t + \Delta_t$. For all conjuncts $v_j - v_i \in [x, y]$, we replace the conjunct with $v_j \in [t + \Delta_t + x, t + y]$. Intuitively, this means that since v_i can happen at the latest at $t + \Delta_t$, v_j can not be allowed to happen before $t + \Delta_t + x$. Likewise, since v_i can happen at the earliest at t, v_j can not be allowed to happen after $t + y$. Finally, if $t + \Delta_t + x > t + y$, the conjunct is replaced with *false*. Also, the process can be applied recursively in the event that v_j is also a timepoint that occurred during the wait, in which case the conjunct would be replaced by *true* or *false*. In any case, any conjunct obtained of the form $a_j \in [x', y']$ is replaced with *false* if $t + \Delta_t > y'$. Finally, if all conjuncts inside a disjunct are set to *false* by this process, the constraint is violated and the DTNU is no longer satisfiable.

5 Learning-based Heuristic

We present our learning model and explain how it provides tree search heuristic guidance. Our learning architecture originates from (Gilmer et al. 2017). It uses message passing rules allowing neural networks to process graph-structured inputs where both vertices and edges possess features. Authors of this architecture carry out node classification experiments in quantum chemistry and achieve state-of-the-art results on a molecular property prediction benchmark. Here, we first define a way of converting DTNUs into graph data. Then, we process the graph data with our MPNN and explain how the output is used to guide the tree search.

Let $\Gamma = \{A, U, C, L\}$ be a DTNU. We start by explaining how we turn Γ into a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. First, we convert all time values from absolute to relative with the assumption the current time for Γ is $t = 0$. We search all converted time intervals $[x_i, y_i]$ in C and L for the highest interval bound

value d_{max}, *i.e.* the farthest point in time. We then proceed to normalize every time value in C and L by dividing them by d_{max}. As a result, every time value becomes a real number between 0 and 1. Next, we convert each controllable timepoint $a \in A$ and uncontrollable timepoint $u \in U$ into graph nodes with corresponding *controllable* or *uncontrollable* node features. The time constraints in C and contingency links in L are expressed as edges between nodes with 10 different edge distance classes $(0 : [0, 0.1), 1 : [0.1, 0.2), ..., 9 : [0.9, 1])$. We also use additional edge features to account for edge types (constraint, disjunction, contingency link, direction sign for lower and upper bounds). Moreover, intermediary nodes are used with a distinct node feature in order to map possible disjunctions in constraints and contingency links. We also add a *WAIT* node with a distinct node feature which implicitly designates the act of waiting a period of time. The graph conversion of DTNU γ is characterized by three elements: the matrix of all node features X_v, the adjacency matrix of the graph X_e and the matrix of all edge features X_w.

Let f be the mathematical function for our MPNN and θ its set of parameters. Our function f stacks 5 graph convolutional layers from (Gilmer et al. 2017) coupled with the $\mathrm{ReLU}(\cdot) = \max(0, \cdot)$ piece-wise activation function (Glorot, Bordes, and Bengio 2011). The `sigmoid` function $\sigma(\cdot) = \frac{1}{1+\exp(-\cdot)}$ is then used to obtain a list of probabilities π over all nodes in $\mathcal{G}$: $f_\theta(X_v, X_e, X_w) = \pi$. The probability of each node v in π corresponds to the likelihood of transitioning into a TDC DTNU from the original DTNU Γ by taking the action corresponding to v. If v represents a controllable timepoint a in Γ, its corresponding probability in π is the likelihood of the sub-DTNU resulting from the execution of a being TDC. If v represents a *WAIT* decision, its probability refers to the likelihood of the *WAIT* node having a *true* attribute, *i.e.* the likelihood of all children DTNUs resulting from the wait being TDC (with the wait duration rules set in §4.3). We call these two types of nodes *active* nodes. Otherwise, if v is another type of node, its probability is not relevant to the problem and ignored. Our MPNN is trained on DTNUs generated and solved in §6 only on active nodes by minimizing the cross-entropy loss:

$$\frac{1}{m} \sum_{i=1}^{m} \sum_{j=1}^{q} -Y_{ij} \log(f_\theta(X_i)_j) - (1 - Y_{ij}) \log(1 - f_\theta(X_i)_j)$$

Here $X_i = (X_{i_v}, X_{i_e}, X_{i_w})$ is DTNU number i among a training set of m examples, Y_{ij} is the TDC controllability (1 or 0) of active node j for DTNU number i.

Lastly, the MPNN heuristic is used in the following way in the tree search. Once a *d-OR* node is reached, the parent DTNU node is converted into a graph and the MPNN f is called upon the corresponding graph elements X_v, X_e, X_w. Active nodes in output probabilities π are then ordered by highest values first, and the tree search visits the corresponding children tree nodes in the suggested order, preferring children with higher likelihood of being TDC first.

6 Randomized Simulations for Heuristic Training

We leverage a learning-based heuristic to guide the tree search and increase its effectiveness. A key component in learning-based methods is the annotated training data. We generate such data in automatic manner by using a DTNU generator to create random DTNU problems and solving them with a modified version of the tree search. We store results and use them for training the MPNN. We detail now our data generation strategy.

For training purposes, we create DTNUs which have a number of controllable timepoints ranging from 10 to 20 and a number of uncontrollable timepoints ranging from 1 to 3. The DTNUs are generated in the following way. For interval bounds of constraint conjuncts or contingency links, we always randomly generate real numbers within $[0, 100]$. We restrict the number of conjuncts inside a disjunct to 5 at most. A random number $n_1 \in [10, 20]$ of controllable timepoints and a random number $n_2 \in [1, 3]$ of uncontrollable timepoints are selected. Each uncontrollable timepoint is randomly linked to a different controllable timepoint with a contingency link. Next, we iterate over the list of timepoints, and for each timepoint v_i not appearing in constraints or contingency links, we add a disjunct for which at least one conjunct constrains v_i. The type of conjunct is selected randomly from either a *distance* conjunct $v_i - v_j \in [x, y]$ or a *bounded* conjunct $v_i \in [x, y]$. On the other hand, if v_i was already present in the constraints or contingency links, we add a disjunct constraining v_i with only a 20% probability. In order to solve these random DTNUs, we modify the tree search as follows. For a DTNU Γ, the first *d-OR* child node is developed as well as its children $\psi_1, \psi_2, ..., \psi_n \in \Psi$. The modified tree search will explore each ψ_i multiple times (ν times at most), each time with a timeout of τ seconds. Here we set $\nu = 25$ and $\tau = 3$. For each exploration of ψ_i, children nodes of any *d-OR* node encountered in the corresponding subtree are explored randomly each time. If ψ_i is proved to be either TDC or non-TDC during an exploration, the next explorations of the same child ψ_i are called off and the truth attribute β_i of ψ_i is updated accordingly. The active node number k, corresponding to the decision leading to ψ_i from DTNU Γ's *d-OR* node, is updated with the same value, *i.e.* $Y_k = \beta_i$. However, if every exploration times out, ψ_i is assumed non-TDC and Y_k is set to *false*. Once each ψ_i has been explored, the pair $\langle \mathcal{G}(\Gamma), (Y_1, Y_2, ..., Y_n) \rangle$ is stored in the training set, where $\mathcal{G}(\Gamma)$ is the application of the graph conversion of Γ described in §5.

The assumption of non-TDC controllability for children nodes for which all explorations time out is good enough. The output of the MPNN is a probability for each child node of the *d-OR* node of the input DTNU. These children nodes are visited in the suggested order when the heuristic is active, until one is found to be TDC, and no child is ever discarded. The trained MPNN will tend to give higher probabilities for children nodes for which explorations often found a TDC strategy before timeout, and lower probabilities for ones where explorations often ended up with a timeout.

7 Strategy Execution

A strategy found by the tree search for a DTNU Γ is sound and guarantees constraint satisfiability if executed in the following manner. Let $\mathcal{Q}$ be the system interacting with the environment, executing controllable timepoints and observing how uncontrollable timepoints unfold. At each DTNU node in the tree, $\mathcal{Q}$ will move on to the child *d-OR* node. The child node ψ_i of the *d-OR* node which was found by the strategy to have a *true* attribute is selected. If ψ_i is a DTNU node, $\mathcal{Q}$ executes the corresponding controllable timepoint a_i and moves on to ψ_i. On the other hand, if ψ_i is a *WAIT* node, $\mathcal{Q}$ moves on to ψ_i, reads the wait duration time Δ_t stored in ψ_i and moves on to the child *w-OR* node. The child node AND_{R_j} of the *w-OR* node which has a *true* attribute is selected, and $\mathcal{Q}$ will wait Δ_t time units with the reactive wait strategy R_j. After the wait is over, $\mathcal{Q}$ observes the list of all uncontrollable timepoints Λ_i which occurred, deduces which DTNU child node of the AND_{R_j} node it transitioned into, and moves on to that node.

By following these guidelines, the final tree node $\mathcal{Q}$ transitions into is necessarily a leaf node with a *true* attribute, *i.e.* a node for which all constraints are satisfied. This is due to the fact that for *d-OR* and *w-OR* nodes $\mathcal{Q}$ visits, $\mathcal{Q}$ chooses to transition into a child node with a *true* attribute. For *AND* nodes $\mathcal{Q}$ visits, all children DTNU nodes have a *true* attribute, so $\mathcal{Q}$ transitions into a child node with a *true* attribute regardless of how uncontrollable timepoints unfold.

8 Related Works

The use of learning-based heuristics has recently become increasingly popular for planning, combinatorial and network modeling problems. Recent works applied to network modeling and routing problems include (Rusek et al. 2019), (Chen et al. 2018), (Xu et al. 2018), (Kool and Welling 2018). Recently, GNNs have become a popular extension of CNNs. Essentially, their ability to represent problems with a graph structure and the resulting node permutation invariance makes them convenient for some applications. We refer the reader to (Wu et al. 2019) for a complete survey on GNNs. In combinatorial optimization, GNNs can benefit both approximate and exact solvers. In (Li, Chen, and Koltun 2018), authors combine tree search, GNNs and a local search algorithm to achieve state-of-the-art results for approximate solving of NP-hard problems such as the maximum independent set problem. On the other hand, (Gasse et al. 2019) use a GNN for branch and bound variable selection for exact solving of NP-hard problems and achieve superior results to previous learning approaches. In path-planning problems with NP-hard constraints, (Osanlou et al. 2019) use a GNN to predict an upper bound for a branch and bound solver and outperform an A*-based planner coupled with a problem-suited handcrafted heuristic. Lastly, (Ma et al. 2018) call a GNN for the selection of a planner inside a portfolio for STRIPS planning problems and outperform the leading learning-based approach which was based on a CNN (Sievers et al. 2019). In most works, GNNs seem to offer generalization to bigger problems than they are trained on. Results from our experiments are in line with this obser-
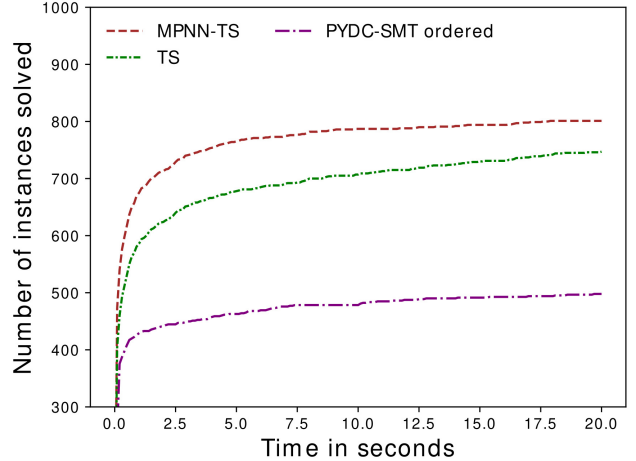

Figure 3: **Experiments on (Cimatti, Micheli, and Roveri 2016)'s benchmark from which DTNs and STNs have been removed.** The X-axis represents the allocated time in seconds and the Y-axis the number of instances in the benchmark each solver can solve within the corresponding allocated time. Timeout is set to 20 seconds per instance.

vation.

9 Experiments

We carry out experiments to evaluate the efficiency of the proposed tree search approach and the effect of the MPNN's guidance. We also compare these methods to a DC solver from (Cimatti, Micheli, and Roveri 2016). TDC is a subset of DC and a more restrictive form of controllability: non-TDC controllability does not imply non-DC controllability. In that sense, a TDC solver can be expected to offer better performance that a DC counterpart in exchange for potentially being unable to find a strategy when a DC algorithm would. In this section, we refer to the tree search algorithm as TS, the tree search algorithm guided by the trained MPNN up to the 15^{th} (respectively X^{th}) *d-OR* node depth-wise in the tree as MPNN-TS (respectively MPNN-TS-X) and the most efficient DC solver from (Cimatti, Micheli, and Roveri 2016) as PYDC-SMT ordered.

First, we use the benchmark in the experiments of (Cimatti, Micheli, and Roveri 2016) from which we remove DTNs and STNs. We compare TS, MPNN-TS and PYDC-SMT on the resulting benchmark which is comprised of 290 DTNUs and 1042 STNUs. Here, Limiting the maximum depth use of the MPNN to 15 offers a good trade off between guidance gain and cost of calling the heuristic. Results are given in Figure 3. We observe that TS solves roughly 50% more problem instances than PYDC-SMT within the allocated time (20 seconds). In addition, TS solves 56% of all instances while the remaining ones time out. Among solved instances, a strategy is found for 89% and the remaining 11% are proved non-TDC. On the other hand, PYDC-SMT solves 37% of all instances. A strategy is found for 85% of PYDC-SMT's solved instances while the remaining 15% are proved non-DC. Finally, out of all instances PYDC-SMT
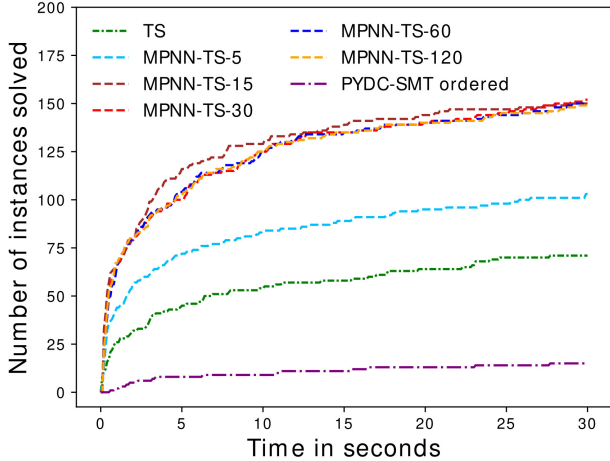
Figure 4: **Experiments on benchmark B_1.** Axes are the same as in figure 3. Timeout is set to 30 seconds per instance.

Figure 5: **Experiments on benchmark B_3.** Axes are the same as in figure 3. Timeout is set to 180 seconds.

solves, TS solves 97% accurately with the same conclusion, *i.e.* TDC when DC and non-TDC when non-DC. The use of the heuristic leads to an additional $+6\%$ problems solved within the allocated time. We argue this small increase is essentially due to the fact that most problems solved in the benchmark are small-sized problems with few timepoints which are solved quickly.

For further evaluation of the heuristic, we create new benchmarks using the DTNU generator described in §6 with varying number of timepoints. These benchmarks contain fewer DTNU instances which are quick to solve and more harder instances. Each benchmark contains 500 randomly generated DTNUs which have 1 to 3 uncontrollable timepoints. Moreover, each DTNU has 10 to 20 controllable timepoints in the first benchmark B_1, 20 to 25 in the second benchmark B_2 and 25 to 30 in the last benchmark B_3. Experiments on B_1, B_2 and B_3 are respectively shown in figure 4, 6c (in the appendix) and 5. We note that for all three benchmarks, no solver ever proves non-TDC or non-DC controllability before timing out due to the larger size of these problems.

In these benchmarks, PYDC-SMT does not perform well on B_1 and cannot solve any instance on B_2 and B_3. TS does not perform well on B_2 and only solves 2 instances on B_3. However, we see a significantly higher gain from the use of the MPNN for TS, varying with the maximum depth use. At best depth use, the gain is $+91\%$ instances solved for B_1, $+980\%$ instances solved for B_2 and $+1150\%$ instances solved for B_3. The more timepoints instances have, the more worthwhile heuristic guidance appears to be. Indeed, the optimal maximum depth use of the MPNN in the tree increases with the problem size: 15 for B_1, 60 for B_2 and 120 for B_3. We argue this is due to the fact that more timepoints results in a wider search tree overall, including in deeper sections where heuristic use was not necessarily worth its cost for smaller problems. Furthermore, the MPNN is trained on randomly generated DTNUs which have 10 to 20 controllable timepoints. The promising gains shown by experiments on
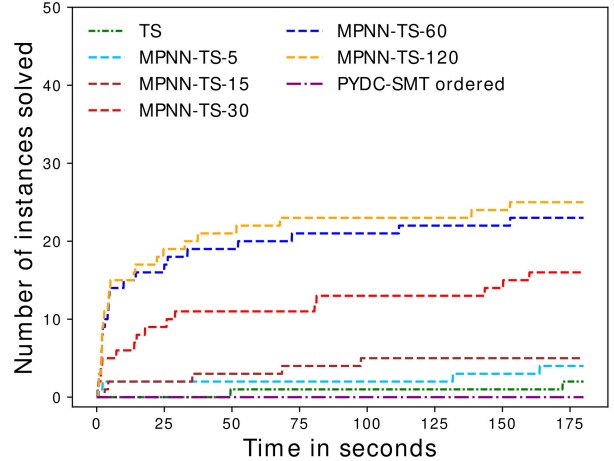
B_2 and B_3 suggest generalization of the MPNN to bigger problems than it is trained on.

The tree search approach presented in this work presents a good trade off between search completeness and effectiveness: almost all examples solved by PYDC-SMT from (Cimatti, Micheli, and Roveri 2016)'s benchmark are solved with the same conclusion, and many more which could not be solved are. Moreover, the TDC approach scales up better to problems with more timepoints, and the tree structure allows the use of learning-based heuristics. Although these heuristics are not key to solving problems of big scales, our experiments suggest they can still provide a high increase in efficiency.

10   Conclusion

We introduced a new type of controllability, time-based dynamic controllability (TDC), and a tree search approach for solving disjunctive temporal networks with uncertainty (DTNU) in TDC. Strategies are built by discretizing time and exploring different decisions which can be taken at different key points, as well as anticipating how uncontrollable timepoints can unfold. We defined constraint propagation rules which ensure soundness of strategies found. We showed that the tree search approach is able to solve DTNUs in TDC more efficiently than the state-of-the-art dynamic controllability (DC) solver, PYDC-SMT, with almost always the same conclusion. Lastly, we created MPNN-TS, a solver which combines the tree search with a heuristic function based on message passing neural networks (MPNN) for guidance. The MPNN is trained with a self-supervised strategy based on a variant of the tree search algorithm. The use of the MPNN allows significant improvement of the tree search on harder DTNU problems, notably on DTNUs of bigger size than those used for training the MPNN.

References

Battaglia, P.; Pascanu, R.; Lai, M.; Rezende, D. J.; et al. 2016. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, 4502–4510.

Bhargava, N., and Williams, B. C. 2019. Complexity bounds for the controllability of temporal networks with conditions, disjunctions, and uncertainty. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 6353 – 6357.

Chen, X.; Guo, J.; Zhu, Z.; Proietti, R.; Castro, A.; and Yoo, S. 2018. Deep-rmsa: A deep-reinforcement-learning routing, modulation and spectrum assignment agent for elastic optical networks. In *2018 Optical Fiber Communications Conference and Exposition (OFC)*, 1–3. IEEE.

Cimatti, A.; Micheli, A.; and Roveri, M. 2016. Dynamic controllability of disjunctive temporal networks: Validation and synthesis of executable strategies. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Cplex, I. I. 2009. V12. 1: User's manual for cplex. *International Business Machines Corporation* 46(53):157.

Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, 3844–3852.

Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research* 12(Jul):2121–2159.

Gasse, M.; Chételat, D.; Ferroni, N.; Charlin, L.; and Lodi, A. 2019. Exact combinatorial optimization with graph convolutional neural networks. In *Advances in Neural Information Processing Systems*, 15554–15566.

Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1263–1272. JMLR. org.

Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 315–323.

Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.

Kipf, T.; Fetaya, E.; Wang, K.-C.; Welling, M.; and Zemel, R. 2018. Neural relational inference for interacting systems. *arXiv preprint arXiv:1802.04687*.

Kool, W., and Welling, M. 2018. Attention solves your tsp. *arXiv preprint arXiv:1803.08475*.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.

Li, Y.; Tarlow, D.; Brockschmidt, M.; and Zemel, R. S. 2016. Gated graph sequence neural networks. In *International Conference on Learning Representations*.

Li, Z.; Chen, Q.; and Koltun, V. 2018. Combinatorial optimization with graph convolutional networks and guided tree search. In *Advances in Neural Information Processing Systems*, 536—-545.

Ma, T.; Ferber, P.; Huo, S.; Chen, J.; and Katz, M. 2018. Adaptive planner scheduling with graph neural networks. *CoRR* abs/1811.00210.

Morris, P., and Muscettola, N. 2005. Temporal dynamic controllability revisited. In *Proceedings of the 22^{nd} National Conference on Artificial Intelligence*.

Morris, P. 2014. Dynamic controllability and dispatchability relationships. In *Proceedings of the IInternational Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, 464 – 479.

Osanlou, K.; Bursuc, A.; Guettier, C.; Cazenave, T.; and Jacopin, E. 2019. Optimal solving of constrained path-planning problems with graph convolutional networks and optimized tree search. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3519–3525. IEEE.

Rusek, K.; Suárez-Varela, J.; Mestres, A.; Barlet-Ros, P.; and Cabellos-Aparicio, A. 2019. Unveiling the potential of graph neural networks for network modeling and optimization in sdn. In *Proceedings of the 2019 ACM Symposium on SDN Research*, 140–151.

Sievers, S.; Katz, M.; Sohrabi, S.; Samulowitz, H.; and Ferber, P. 2019. Deep learning for cost-optimal planning: Task-dependent planner selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7715–7723.

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. *Journal of Experimental and Theoretical Artificial Intelligence* 11(1):23 – 45.

Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Yu, P. S. 2019. A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*.

Xu, Z.; Tang, J.; Meng, J.; Zhang, W.; Wang, Y.; Liu, C. H.; and Yang, D. 2018. Experience-driven networking: A deep reinforcement learning based approach. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, 1871–1879. IEEE.

11 Appendix

11.1 Plots
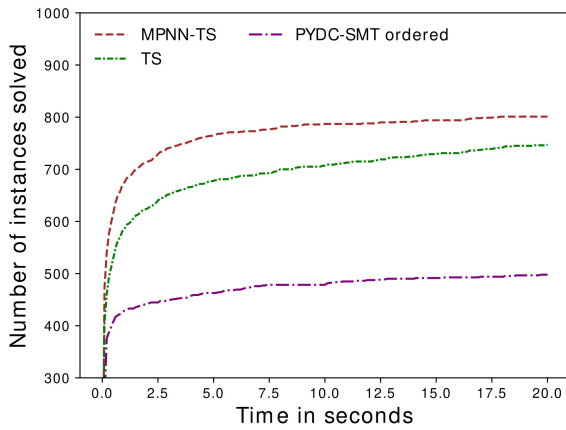
(a) Experiments on (Cimatti, Micheli, and Roveri 2016)'s benchmark from which the DTNs and STNs have been removed. The X-axis represents the allocated time in seconds and the Y-axis the total number of instances that each solver can solve within the corresponding allocated time. Timeout is set to 20 seconds per instance.
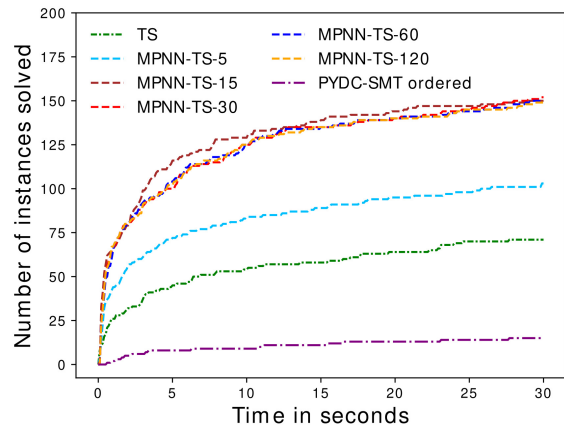
(b) Experiments on benchmark B_1. Axes are the same as in figure 6a. Timeout is set to 30 seconds per instance.
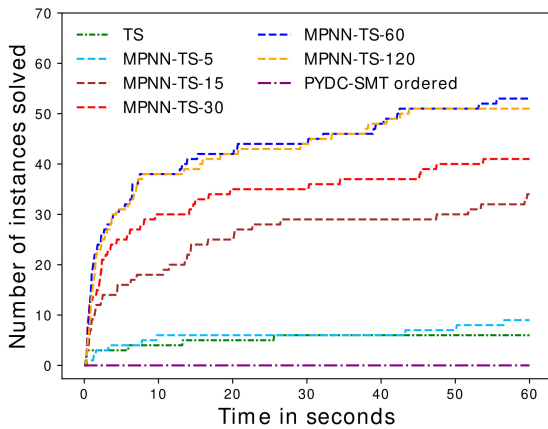
(c) Experiments on benchmark B_2. Axes are the same as in figure 6a. Timeout is set to 60 seconds per instance.
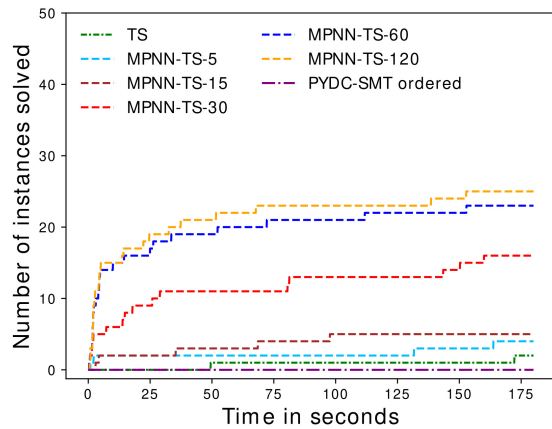
(d) Experiments on benchmark B_3. Axes are the same as in figure 6a. Timeout is set to 180 seconds per instance.

Figure 6: **Summary of experiments on benchmarks**

11.2 Simplified Example

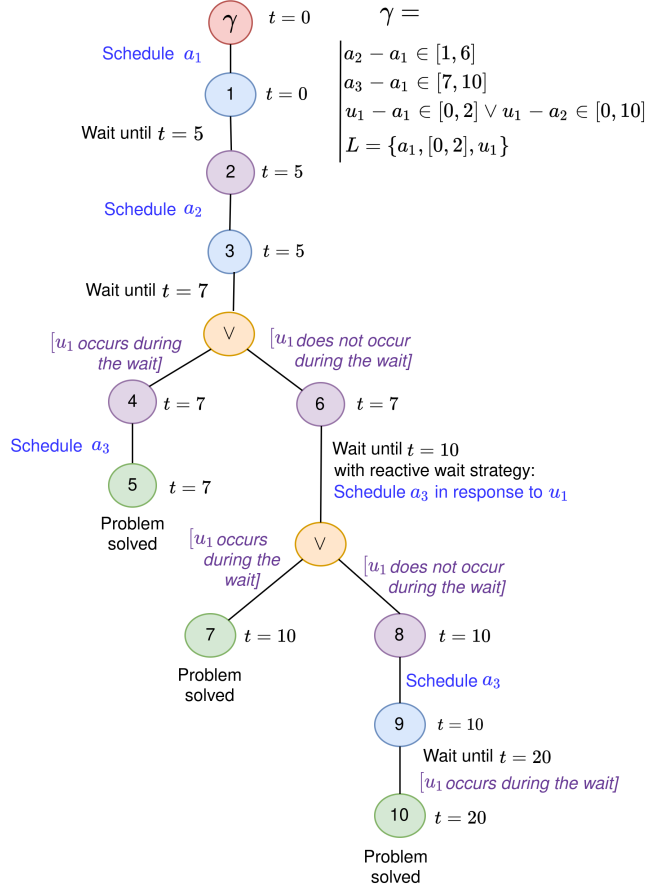Figure 7 is a simplified example of a TDC strategy of the example DTNU from (Cimatti, Micheli, and Roveri 2016).

Figure 7: **Simplified TDC strategy of a DTNU Γ.** For space reasons, we only give a summarized copy of the strategy found. Branches leading to unsolved cases are excluded, and we do not include *d-OR* , *w-OR* , *AND* and *WAIT* nodes. The node γ is the original DTNU. Other nodes are sub-DTNUs, except the $\vee$ node which aims to list transitional possibilities, and should be interpreted in the figure as an *AND* node.

11.3 Truth Value Propagation Algorithm

We present in this section Algorithm 1. This algorithm is called to propagate a truth value in the tree. The propagation is done in an ascending way: truth values are inferred from the leaves of the tree towards the root.

Algorithm 1 Truth Value Propagation

```
 1: function PROPAGATETRUTH(TREENODE ψ)
 2:     ω ← parent(ψ)                              ▷ 1*
 3:     if ω = null then
 4:         return
 5:     if isDTNU(ω) or isWAIT(ω) then             ▷ 2*
 6:         ω.truth ← ψ.truth
 7:         propagateTruth(ω)
 8:     else if isOR(ω) then                       ▷ 3*
 9:         if ψ.truth = True then
10:             ω.truth ← True
11:             propagateTruth(ω)
12:         else
13:             if ∀σi, σi.truth = False then      ▷ 4*
14:                 ω.truth ← False
15:                 propagateTruth(ω)
16:     else if isAND(ω) then                      ▷ 5*
17:         if ψ.truth = False then
18:             ω.truth ← False
19:             propagateTruth(ω)
20:         else
21:             if ∀σi, σi.truth = True then       ▷ 4*
22:                 ω.truth ← True
23:                 propagateTruth(ω)
```

---

[1]* parent(x): Returns the parent node of x, $null$ if none.

[2]* isDTNU(x): Returns *True* if x is a DTNU node, *False* otherwise; isWait(x): Returns *True* if x is a *WAIT* node, *False* otherwise.

[3]* isOR(x): Returns *True* if x is an *d-OR* or *w-OR* node, *False* otherwise.

[4]* σ_i: Child number i of ω. For a *d-OR* or *w-OR* node, in the case where ψ is *false* but not all other children of ω are *false* the propagation stops. Likewise, for an *AND* node and in the case where ψ is *true* but not all other children of ω are *true* , the propagation stops.

[5]* isAnd(x): Returns *True* if x is an *AND* node, *False* otherwise.

11.4 Tree Search Algorithm

We give the simplified pseudocode for the tree search in Algorithm 2

Algorithm 2 Tree Search

1: **function** EXPLORE(ψ)
2:     **if** $parent(\psi).truth \neq unknown$ **then**
3:        **return**
4:     **if** isDTNU(ψ) **then**
5:        updateConstraints(ψ) $\triangleright$ [6]∗
6:        **if** IsLeaf(ψ) **then** $\triangleright$ [7]∗
7:           propagateTruth(ψ)
8:           **return**
9:        Create *d-OR* child ψ'
10:       explore(ψ')
11:     **if** isOR(ψ) **then**
12:        Create list of all children Ψ' $\triangleright$ [8]∗
13:        **for** $\psi' \in \Psi'$ **do**
14:           explore(ψ')
15:     **if** isAND(ψ) **then**
16:        Create list of all children Ψ' $\triangleright$ [9]∗
17:        **for** $\psi' \in \Psi'$ **do**
18:           explore(ψ')
19:     **if** isWAIT(ψ) **then**
20:        create *w-OR* child ψ'
21:        explore (ψ')
22: **function** MAIN(DTNU γ)
23:     explore(γ)
24:     **if** $\gamma.truth = True$ **then**
25:        **return** $True$
26:     **else**
27:        **return** $False$

---

[6]∗ updateConstraints(x): Updates the constraints of DTNU node x.

[7]∗ isLeaf(x): Sets the truth value of x to *true* and returns *true* if all constraints are satisfied. Sets the truth value to *false* and returns *true* if a constraint is violated. If no truth value can be inferred at this stage with the updated constraints, a second check is run to determine if all uncontrollable timepoints have occurred. In this case, the corresponding DTN is solved, the truth value of x is updated accordingly, and the function returns *true* . Otherwise, if no logical outcome can be inferred for the current state of the constraints and there remains at least one uncontrollable timepoint, this function returns *false* .

[8]∗ If this is a *d-OR* node, the list Ψ' contains all the children DTNU nodes resulting from either the decision of scheduling a controllable timepoint, or the *WAIT* node resulting from a wait if available. If this is a *w-OR* node, Ψ' contains all AND_{R_j} nodes, each of which possess a reactive wait strategy R_j

[9]∗ Here, the list Ψ' contains all DTNUs resulting from all possible combinations $\Lambda_1, \Lambda_2, ..., \Lambda_q$ of uncontrollable timepoints which have the potential to occur during the current wait.

11.5 Wait Period

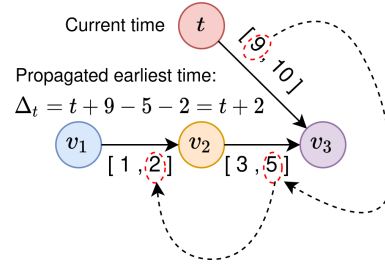The following figure gives an example of the third rule used to compute a wait duration.

Figure 8: **Application of the 3^{rd} rule to determine a wait duration.** Variables v_1, v_2 and v_3 are timepoints. Here, v_2 is constrained to execute in the time interval $[1, 2]$ after v_1, v_3 in $[3, 5]$ after v_2 and v_3 in $[t + 9, t + 10]$. It is suggested not to wait longer than 2 units of time at t: an execution of v_1 at $t + 2$, followed by an execution of v_2 at $t + 4$ opens a window of opportunity for v_3 to execute at $t + 9$.

11.6 Optimization Rules

The following rules are added to make branch cuts when possible.

Constraint Check. When a DTNU node is explored and the updated list of constraints C' is built according to §4.5, if a disjunct is found to be *false* , C' will no longer be satisfiable. All the subtree which can be developed from the DTNU will only have leaf nodes for which this is the case as well. Therefore, the search algorithm will not develop this subtree.

Symmetrical subtrees. Some situations can lead to the development of the exact same subtrees. A trivial example, for a given DTNU node at a time t, is the order in which a given combination of controllable timepoints $a_1, a_2, ..., a_k$ is taken before taking a wait decision. Regardless of what order these timepoints are explored in the tree before moving to a *WAIT* node, they will be considered executed at time t. Therefore, when taking a wait decision, it is checked that all preceding controllable timepoints executed before the previous wait are a combination of timepoints that has not been not tested yet.

Truth Checks. Before exploring a new node for which the truth attribute is set to *unknown*, the truth attribute of the parent node is also checked. The node is only developed if the parent node's truth attribute is set to *unknown*. In this manner, when children of a tree node are being explored depth-first and the exploration of a child node leads to the assignment of a truth value to the tree node, the remaining unexplored children can be left unexplored.

11.7 Implementation details

Our MPNN architecture is made of 5 graph convolutional layers from (Gilmer et al. 2017) with 32 neurons in each layer. For each layer, we use a two-layer MLP (multi layer perceptron) with 128 neurons in the hidden layer to compute

$$\gamma = \left| \begin{array}{l} a_2 - u_1 \in [0,1] \\ a_2 \in [0,1] \vee a_2 \in [1.5,3] \\ L = \{a_1, [0,2], u_1\} \end{array} \right.$$

$$\gamma' = \left| \begin{array}{l} a_2 - u_1 \in [0,0.33] \\ a_2 \in [0,0.33] \vee a_2 \in [0.5,1] \\ L = \{a_1, [0,0.66], u_1\} \end{array} \right.$$
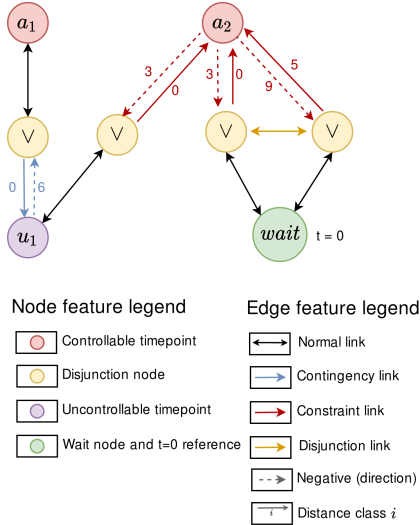

Figure 9: **Conversion of a DTNU γ into a graph.** γ' is the normalized DTNU. Edge distances are expressed as distance classes. To distinguish between lower and upper bounds in intervals, we introduce an additional *negative directional sign* feature.

node and edge embeddings. In addition, we use batch normalization after each graph layer and apply the $\mathrm{ReLU}(\cdot) = \max(0, \cdot)$ activation function. The input of our MPNN is the graph conversion of a DTNU. Figure 9 illustrates an example of graph conversion. We use 10 different edge distance classes: $0 : [0, 0.1), 1 : [0.1, 0.2), ..., 9 : [0.9, 1]$. Training is done with the *adagrad* optimizer (Duchi, Hazan, and Singer 2011) and an initial learning rate 10^{-4} on a dataset comprised of 30.000 instances generated as described in §6. We split the data into a training set comprised of 25.000 instances and a cross-validation set comprised of 5.000 instances. We add a dropout regularization layer with a *keep rate* 0.9 before the output layer to reduce overfitting.