

Revisiting Classical Dynamic Controllability

A Tighter Complexity Analysis

Mikael Nilsson, Jonas Kvarnström, and Patrick Doherty

Department of Computer and Information Science,
Linköping University, SE-58183 Linköping, Sweden
{mikni, jonkv, patdo}@ida.liu.se
<http://www.ida.liu.se>

Abstract. Simple Temporal Networks with Uncertainty (STNUs) allow the representation of temporal problems where some durations are uncontrollable (determined by nature), as is often the case for actions in planning. It is essential to verify that such networks are dynamically controllable (DC) – executable regardless of the outcomes of uncontrollable durations – and to convert them to an executable form. We use insights from incremental DC verification algorithms to re-analyze the original, classical, verification algorithm. This algorithm is the entry level algorithm for DC verification, based on a less complex and more intuitive theory than subsequent algorithms. We show that with a small modification the algorithm is transformed from pseudo-polynomial to $O(n^4)$ which makes it still useful. We also discuss a change reducing the amount of work performed by the algorithm.

1 BACKGROUND

Time and concurrency are increasingly considered essential in planning and multi-agent environments, but temporal representations vary widely in expressivity. For example, Simple Temporal Problems (STPs, [1]) allow us to efficiently determine whether a set of *timepoints* (events) can be assigned real-valued *times* in a way consistent with a set of *constraints* bounding temporal distances between timepoints. The start and end of an action can be represented as timepoints, but its possible durations can only be represented as an STP constraint if the execution mechanism can *choose* durations arbitrarily within the given bounds. Usually, exact durations are instead chosen by nature and agents must generate plans that work regardless of the eventual outcomes.

STPs with Uncertainty (STPUs, [2]) capture this aspect by introducing *contingent* timepoints corresponding to the end of actions. Associated with these are *contingent* temporal constraints that correspond to possible action durations to be decided by nature. One must then find a way to assign times to ordinary *controlled* timepoints (determine when to start actions) so that for *every* possible outcome for the contingent constraints (action durations), all ordinary *requirement* constraints (corresponding to STP constraints) are satisfied.

If an STPU allows us to schedule controlled timepoints (actions to be started) incrementally given that we receive information when a contingent timepoint occurs (an action ends), it is *dynamically controllable* (DC) and can be efficiently executed by a

dispatching algorithm [3]. Conversely, guaranteeing that constraints are satisfied when executing a non-DC plan is impossible; it would require information about future duration outcomes.

We will not describe dispatch algorithms for STPs and STPUs here. They can be found in [3, 4]. The idea behind them is to keep track of which events that are allowed to be executed, since their predecessors have executed, and to execute these in a way that satisfies all constraints towards predecessors. When we in the future mention that a network is *dispatchable* we mean that it is in a form which can directly be executed by one of the existing dispatch algorithms.

Although several algorithms for verifying the dynamic controllability of STPUs have been published [4–7] we will focus our attention on the first which is often referred to as classical or MMV [4]. The algorithm is easily implemented, it captures the intuition behind STPUs and has a direct correctness proof. It is also the entry level algorithm for verification. We will show that its run-time is not as thought before, pseudo-polynomial, but $O(n^4)$ through a small modification – the algorithm merely needs to stop earlier. This result shows that the algorithm is quite fast and still useful.

The intuition behind the analysis is that not all of MMV’s derivations and tightenings are necessary: Only a certain *core* of derivations actually matters for verifying dynamic controllability, and when the STPU is DC, this core is free of cyclic derivations. This can be exploited through a small change to MMV. Stopping at the right time also preserves another aspect of MMV: the result is dispatchable.

Outline. After providing some fundamental definitions (Section 2), we describe the MMV algorithm (Section 3). We also present the FastIDC algorithm, which will provide intuitions for our analysis of MMV (Section 4). We compare the derivations made by the two algorithms (Section 5) and analyze the length of FastIDC derivation chains (Section 6), resulting in the new algorithm GlobalDC (Section 7) which runs in $O(n^4)$. GlobalDC is in fact identical to a slightly modified MMV algorithm (Section 8).

2 TEMPORAL PROBLEMS

We start with defining some fundamental concepts.

Definition 1. A *simple temporal problem* (STP, [1]) consists of a number of real variables x_1, \dots, x_n and constraints $T_{ij} = [a_{ij}, b_{ij}], i \neq j$ limiting the temporal distance $a_{ij} \leq x_j - x_i \leq b_{ij}$ between the variables.

We will work with STPs in graph form, with timepoints represented as nodes and constraints as labeled edges. They are then referred to as Simple Temporal Networks (STNs). We will also make use of the fact that any STN can be represented as an equivalent *distance graph* [1]. Each constraint $[u, v]$ on an edge AB in an STN is represented as two *corresponding* edges in its distance graph: AB with weight v and BA with weight $-u$. Computing the all-pairs-shortest-path (APSP) distances in the distance graph yields a *minimal representation* containing the tightest distance bounds that are implicit in the original problem [1]. This directly corresponds to the tightest interval constraints $[u', v']$ implicit in the original STN.




Fig. 1. Example STNU.

If the distance graph has a negative cycle, then no assignment of timepoints to variables satisfies the STN: It is *inconsistent*. Otherwise it is consistent and can be *executed*: Its events can be assigned time-points so that all constraints are satisfied.

Definition 2. A *simple temporal problem with uncertainty (STPU)* [2] consists of a number of real variables x_1, \dots, x_n , divided into two disjoint sets of controlled timepoints R and contingent timepoints C . An STPU also contains a number of requirement constraints $R_{ij} = [a_{ij}, b_{ij}]$ limiting the distance $a_{ij} \leq x_j - x_i \leq b_{ij}$, and a number of contingent constraints $C_{ij} = [c_{ij}, d_{ij}]$ limiting the distance $c_{ij} \leq x_j - x_i \leq d_{ij}$. For the constraints C_{ij} we require that $x_j \in C$, $0 < c_{ij} < d_{ij} < \infty$.

STPUs in graph form are called STNs with Uncertainty (STNUs). An example is shown in Figure 1. In this example a man wants to cook for his wife. He does not want her to wait too long after she returns home, nor does he want the food to wait too long. These two requirements are captured by a single requirement constraint, whereas the uncontrollable durations of shopping, driving home and cooking are captured by the contingent constraints. The question is whether the requirements can be guaranteed regardless of the outcomes of the uncontrollable durations.

In addition to the types of constraints already existing in an STNU, some algorithms can also generate *wait constraints* that make certain implicit requirements explicit for use in further computations.

Definition 3. Given a contingent constraint between A and B and a requirement constraint from A to C , the $\langle B, t \rangle$ annotation on the constraint AC indicates that execution of the timepoint C is not allowed to take place until after *either* B has occurred *or* t units of time have elapsed since A occurred. This constraint is called a **wait constraint** [4], or **wait**, between A and C .

As there are events whose occurrence we cannot fully control, consistency is not sufficient for an STNU to be executable. However, suppose that for a given STNU there exists a **dynamic execution strategy** [4] that can assign timepoints to controllable events during execution, given that at each time, it is known which contingent events have already occurred. The STNU is then **dynamically controllable** [4] (**DC**) and can be executed. In Figure 1 a dynamic execution strategy is to start cooking 10 time units after receiving a call that the wife starts driving home. This guarantees that cooking is done within the required time, since she will arrive at home 35 to 40 time units after starting to drive and the dinner will be ready 35 to 40 time units after she started driving.

Algorithm 1: The MMV Algorithm

```

Boolean procedure determineDC()
repeat
  if not pseudo-controllable then
    | return false
  else
    | forall the triangles ABC do
    |   tighten ABC using the tightenings in Figure 2
    | end
until no tightenings were found
return true

```

3 THE MMV ALGORITHM

Algorithm 1 shows a reformulated and clarified [5] version of the classical “MMV” algorithm [4]. Note that these versions share the same worst case complexity.

The algorithm builds on the concept of *pseudo-controllability* [4], a necessary but not sufficient requirement for dynamic controllability. To test for pseudo-controllability the STNU is first converted to an STN by converting all contingent constraints into requirement constraints. The STN then has to be put in its minimal representation (see Section 2). If the STN is inconsistent, the corresponding STNU cannot be consistently executed and is not DC. If the STN is consistent but a constraint corresponding to a contingent constraint in the STNU became tighter in the minimal representation, the contingent constraint is *squeezed* [4]. Nature can then place the uncontrollable outcome of the contingent constraint outside what is allowed by the STN representation, causing execution to fail. Therefore the STNU is not DC. Conversely, if the minimal representation is consistent and does not squeeze any corresponding contingent constraint, the STNU is *pseudo-controllable*, but may still fail to be DC.

MMV additionally uses STNU-specific *tightening rules*, also called *derivation rules*, which make constraints that were previously implicit in the STNU explicit (Figure 2). Each tightening rule can be applied to a “triangle” of nodes if the constraints and requirements of the rule are matched. The result of applying a tightening is a new or tightened constraint, shown as bold edges in the leftmost part of the triangle.

Algorithm 1 is centered around a loop where it first verifies pseudo-controllability and transfers all tighter constraints found by the associated APSP calculation into the STNU, then applies all possible tightenings. If an STNU is not DC, the tightenings will eventually produce sufficient explicit constraints for the pseudo-controllability test to detect this [4].

The complexity of MMV is said to be $O(Un^3)$ where U is a measure of the size of the domain (the number of constraints and the size of constraint bounds). This comes from a cost of $O(n^3)$ per iteration and the fact that each iteration must tighten at least one constraint leading in the extreme to a negative cycle. Since the complexity bound depends on the size of constraint bounds, it is pseudo-polynomial.

If MMV labels an STNU as DC, the processed STNU can be executed by the dispatcher in [4].

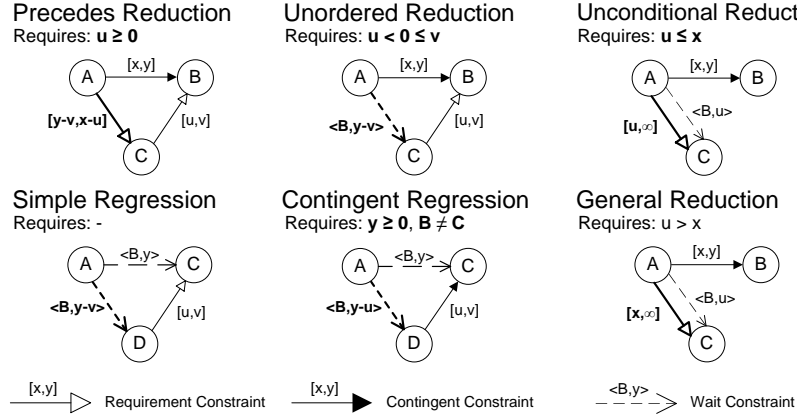


Fig. 2. Tightenings (derivations) of the MMV algorithm.

4 THE FASTIDC ALGORITHM

The property of dynamic controllability is “monotonic” in the sense that if an STNU is not DC, it can never be made DC by further adding or tightening constraints. Therefore, the *non-incremental* verification performed by MMV is equivalent to starting with an empty STNU (which is trivially DC) and *incrementally* adding one edge at a time, verifying at each step that the STNU remains DC.

We will exploit this fact to compare MMV to the incremental FastIDC algorithm [8, 9], which will allow us to draw certain conclusions about MMV. First, though, we will present and explain FastIDC itself, specifically its tightening / edge-addition aspect (since loosening or removing edges will not be required here). As the original version of this algorithm was incorrect in certain cases, we use the corrected version shown in algorithm 2 as our starting point [10]. A proof that this version is correct can be found in [11].

FastIDC has three main differences compared to the MMV algorithm.

1: Representation. FastIDC does not work in the standard STNU representation but uses an *extended distance graph* [12], analogous to the distance graphs sometimes used for STNs. Requirement edges and contingent edges are then translated into pairs of edges of the corresponding type in a manner similar to what was previously described for STNs.

Definition 4. An *extended distance graph (EDG)* is a directed multi-graph with weighted edges of 5 kinds: positive requirement, negative requirement, positive contingent, negative contingent and conditional.

The *conditional* edges mentioned above, first used by [12], are used to represent the *waits* that can be derived by MMV. The direction of a conditional edge is intentionally opposite to that of the wait it encodes. This makes the conditional edge more similar to a negative requirement edge in the same direction, the difference being the condition.

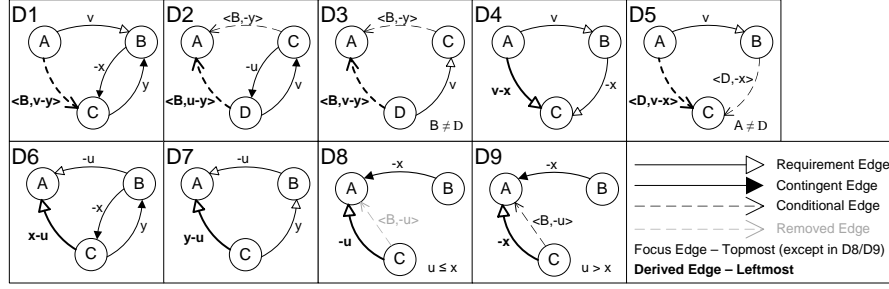


Fig. 3. FastIDC derivation rules D1-D9.

Definition 5. A *conditional edge* CA annotated $\langle B, -w \rangle$ encodes a conditional constraint: C must execute after B or at least w time units after A , whichever comes first. The node B is called the *conditioning node* of the constraint/edge.

2: Derivation rules. Partly due to the new representation, FastIDC uses different derivation rules. These are shown in EDG form in Figure 3, where we have numbered two rules (D8–D9) that were unnumbered in the original publication, but shown to be needed [11].

3: Traversal order. FastIDC uses a significantly different graph traversal order. MMV traverses a graph iteratively, and in each iteration, it considers *all* “triangles” in a graph in arbitrary order. FastIDC, in contrast, uses the concept of *focus edges*. A focus edge is an edge that was tightened and may lead to other constraints being tightened. FastIDC only applies derivation rules to focus edges. If this leads to new tightened edges it will recursively continue to apply the derivation rules until quiescence. Intuitively, this guarantees that all possible consequences of any tightening are covered by the algorithm.

FastIDC Details. Being incremental, FastIDC assumes that at some point a dynamically controllable STNU was already constructed (for example, the empty STNU is trivially DC). Now one or more requirement edges e_1, \dots, e_n have been added or tightened, together with zero or more contingent edges and zero or more new nodes, resulting in the graph G . FastIDC should then determine whether G is DC.

The algorithm works in the EDG of the STNU. First it adds the newly modified or added requirement edges to a queue, Q (a contingent edge must be added before any other constraint is added to its target node and is then handled implicitly through requirement edges). The queue is sorted in order of decreasing distance to the *temporal reference* (TR), a node always executed before all other nodes at time zero. Therefore nodes close to the “end” of the STNU will be dequeued before nodes closer to the “start”. This will to some extent prevent duplication of effort by the algorithm, but is not essential for correctness or for understanding the derivation process.

In each iteration an edge e_i is dequeued from Q .

A positive loop (an edge of positive weight from a node to itself) represents a trivially satisfied constraint that can be skipped. A negative loop entails that a node must be executed before itself, which violates DC and is reported.

Algorithm 2: FastIDC – sound version

```

function FAST-IDC ( $G, e_1, \dots, e_n$ )
   $Q \leftarrow$  sort  $e_1, \dots, e_n$  by distance to temporal reference
    (order important for efficiency, not correctness)
  for each modified edge  $e_i$  in ordered  $Q$  do
    if IS-POS-LOOP ( $e_i$ ) then SKIP  $e_i$ 
    if IS-NEG-LOOP ( $e_i$ ) then return false
    for each rule (Figure 3) applicable with  $e_i$  as focus do
      if edge  $z_i$  in  $G$  is modified or created then
        Update  $CCGraph$ 
        if Negative cycle created in  $CCGraph$  then return false
        if  $G$  is squeezed then return false
        if not FAST-IDC ( $G, z_i$ ) then
          return false
      end
    end
  end
return true

```

If e_i is not a loop, FastIDC determines whether one or more of the derivation rules in Figure 3 can be applied with e_i as focus. The topmost edge in the figure is the focus in all rules except D8 and D9, where the focus is the conditional edge $\langle B, -u \rangle$. Note that rule D8 is special: The derived requirement edge represents a stronger constraint than the conditional focus edge, so the conditional edge is removed.

For example, consider rule D1. This rule will be matched if e_i is a positive requirement edge, there is a negative contingent edge from its target B to some other node C , and there is a positive contingent edge from C to B . Then a new constraint (the bold edge) can be derived. This constraint is only added to the EDG if it is strictly tighter than any existing constraint between the same nodes.

More intuitively, D1 represents the situation where an action is started at C and ends at B , with an uncontrollable duration in the interval $[x, y]$. The focus edge AB represents the fact that B , the end of the action, must not occur more than v time units after A . This can be represented more explicitly with a conditional constraint AC labeled $\langle B, v - y \rangle$: If B has occurred (the action has ended), it is safe to execute A . If at most $v - y$ time units remain until C (equivalently, at least $y - v$ time units have passed after C), no more than v time units can remain until B occurs, so it is also safe to execute A .

Whenever a new edge is created, the corrected FastIDC tests whether a cycle containing only negative edges is generated. The test is performed by keeping the nodes in an incrementally updated topological order relative to negative edges. The unlabeled graph which is used for keeping the topological order is called the $CCGraph$. It contains the same nodes as the EDG and has an edge between two nodes iff there is a negative edge between them in the EDG. See [10] for further information.

After this a check is done to see if the new edge *squeezes* a contingent constraint. Suppose FastIDC derives a requirement edge BA of weight w , for example $w = -12$, representing the fact that B must occur at least 12 time units after A . Suppose there is

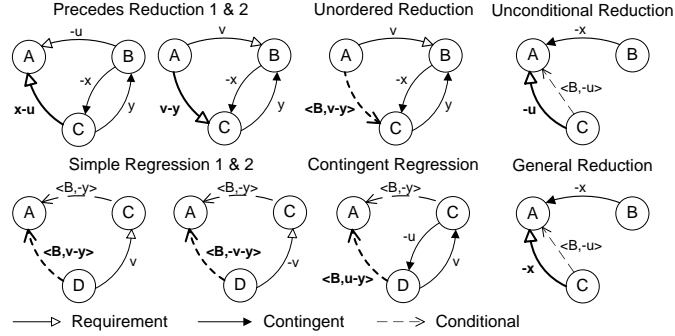


Fig. 4. Classical derivations in EDG format.

also a contingent edge BA of weight $w' > w$, for example $w' = -10$, representing the fact that an action started at A and ending at B may in fact take as little as 10 time units to execute. Then there are situations where nature may violate the requirement edge constraint, and the STNU is not DC.

If the tests are passed and the edge is tighter than any existing edges in the same position, FastIDC is called recursively to take care of any derivations caused by this new edge. Although perhaps not easy to see at a first glance, all derivations lead to new edges that are closer to the temporal reference. Derivations therefore have a direction and will eventually stop. When no more derivations can be done the algorithm returns true to testify that the STNU is DC. If FastIDC returns true after processing an EDG this EDG can be directly executed by the dispatcher in [4].

5 COMPARING FASTIDC / MMV

To compare the derivation rules used by MMV to those of FastIDC, we first need a translation into EDG format. This is shown in Figure 4 where as before the bold edges are derived. *Precedes reduction* is split in two since it adds two edges. *Simple regression* is also split in two, one version regressing over a positive edge and one regressing over a negative edge. All variables used as weights are considered positive, i.e., $-u$ is a negative number (with unconditional reduction as an exception). The additional requirements from Figure 2 still apply but are omitted for clarity. Most are encoded by the edge types – for instance in unordered reduction, only a positive requirement edge can match the rule, making the $v > 0$ requirement implicit. We now see the following similarities:

- Precedes Reduction 1 (PR1) is identical to D6.
- Unordered reduction is equivalent to D1. However without the extra requirement ($u \geq 0$) used by MMV to distinguish between applying PR2 and unordered reduction, FastIDC will always apply unordered reduction, even when MMV instead would apply PR2. It can be shown that if the situation calls for an application of




Fig. 5. Simple regression when the edge is negative.

PR2, FastIDC derives the same edge as MMV through conversion of the conditional edge resulting from D1 into a requirement edge (via unconditional reduction, D8). If the application of PR2 directly leads to non-DC detection, FastIDC also detects this directly. So PR and unordered reduction are handled by D1, D6 and D8 together.

- Simple regression 1 is equivalent to D3 and D5. The only difference between D3 and D5 is which edge is regarded as focus.
- Contingent regression is identical to D2.
- Unconditional Reduction is identical to D8.
- General Reduction is identical to D9.

Thus, the only significant differences are:

- FastIDC derivations has no counterpart to Simple Regression 2.
- D4 and D7 have no counterpart rules in MMV. These derive shortest path distances towards earlier nodes in the STNU. This derivation is present and handled by the APSP calculation in MMV.

We see that MMV does everything that the FastIDC derivations do, and also applies SR2 and a complete APSP calculation.

It can in fact be seen that SR2 is not needed, not even by MMV. Figure 5 shows the situation where a conditional edge CA is regressed over an incoming negative requirement edge DC . Adding a constraint DA to “bridge” two consecutive negative edges is always redundant both for execution and for DC verification. From an execution perspective this is easily seen since C is always executed before D which ensures that the chain of constraints is respected without the addition of DA . From a verification perspective this can be seen since the derived constraint is in fact weaker than the two original constraints. If B is executed before the wait expires the DA constraint “forgets” about the $-v$ part of the constraint which must still be fulfilled. If the wait expires both paths require D to be at least $v + y$ time units after A and the constraint is redundant.

6 FOCUS PROPAGATION

If we apply rules D1–D9 in Figure 3, every derived edge has a uniquely defined “parent”: The focus edge of the derivation rule. Unless this edge was already present in the original graph, it (recursively) also has a parent. This leads to the following definition.

Definition 6. *Edges that are derived through Figure 3 derivations are part of a **derived chain**, where the parent of each edge is the focus edge used to derive it.*

Table 1. The derived edges compared to the focus edges.

Rule	Effect
D1	The target of the derived edge is an earlier node.
D2,D6	The source of the derived edge is an earlier node.
D3,D7	The source of the derived edge is an earlier or unordered node.
D4,D5	The target of the derived edge is an earlier node.
D8,D9	The derived edge connects the same nodes.

We observe the following:

- A contingent constraint orders the nodes it constrains. In EDG form we see this by the fact that the target of a negative contingent edge is always executed before its source.
- Either D8 or D9 is applicable to any conditional edge. Thus there will always be an order between its nodes set by the negative requirement edge from D8/D9: The target node of a conditional constraint is always executed before its source.

This leads directly to the facts in Table 1. Here, node n_1 is considered *earlier* than n_2 if n_1 must be executed before n_2 in every dynamic execution strategy and for all duration outcomes. Similarly, node n_1 is considered *unordered* relative to n_2 if their order can differ depending on strategy or outcome.

We now consider the structure of derived chains in DC STNUs. The focus will be on the direction and weight of each derived edge, ignoring whether edges are negative, positive, requirement or conditional (but still keeping track of contingent edges).

Lemma 1. *Suppose all rules in Figure 3 are applied to the EDG of a dynamically controllable STNU until no more rules are applicable. Then, all derived chains are **acyclic**: No derivation rule has generated an edge having the same source and target as an ancestor of its parent edge along the current chain.*

Proof. Note that by the definition of acyclicity we allow “cycles” of length 1. These can only be created by applications of D8–D9 in a DC STNU.

For D1–D7, each derived edge shares one node with its parent focus edge, but has another source or target. We can then track how the source and target of the focus edge changes through the chain.

Table 1 shows that only derivation rules D1, D4 or D5 result in a different *target* for the derived edge compared to the focus edge. The new target has always “moved” along a negative edge, so it must be executed earlier than the target of the focus edge. Since the STNU is DC, its associated STN cannot have negative cycles. Thus, if the target changes along a chain, it cannot “cycle back” to a previously visited target.

Rules D2, D3, D6 and D7 result in a different *source* for the derived edge. This source may be earlier *or* later than the source of the focus edge, so these rules can be applied in a sequence where the source of the focus edge “leaves” a node n and eventually “returns”. Suppose that this happens and the target n' has not changed. This must occur through applications of rules D2, D3, and/or D6–D9. No such derivation step decreases the weight of the focus edge. Therefore, when the source returns to n , the new edge to be derived between n and n' cannot be tighter than the one that already exists. No new edge is actually derived. Thus, if the source changes along a chain, it cannot “cycle back” to a previously visited source. \square

This fact together with the previous lemma limits the length of a derived chain to $2n^2$ since we have at most n^2 distinct ordered source/target pairs and can at most have one application of D8/D9 in-between source/target movements. The use of chains to reach an upper bound on iterations is inspired by [5] where an upper bound of $O(n^5)$ is reached for MM.

Note that FastIDC derivations together with local consistency checks and global cycle detection is sufficient to guarantee that all implicit constraints represented by a chain of negative edges are respected, or non-DC is reported. There is no need to add these implicit constraints but the next proof will make use of the fact that they exist.

Some derivations carried out by FastIDC can be proven not to affect the DC verification process, and hence we would like to avoid doing these. These can both be derivations of weaker constraints and constraints that are implicitly checked even if they are not explicitly present in the EDG. In order to single out the needed derivations we define *critical chains*.

Definition 7. A *critical chain* is a derived chain in which all derivations are needed to correctly classify the STNU. If any derivation in the chain was missing, a non-DC STNU might be misclassified as DC.

Given a focus edge, one or more derivations may be applicable. Those that would extend the current critical chain into a non-critical one can be skipped without affecting classification. We therefore identify some criteria that are satisfied in all critical chains.

Theorem 1. Given a DC STNU:

1. A D1 derivation for a specific contingent constraint C can only be part of a critical chain once.
2. At most one derivation of type D2 and D6 involving a specific contingent constraint C can be part of a critical chain.

Proof (Proof sketch:). Part 1 is shown as in the proof of lemma 1: The target cannot come back for another D1 application to the same contingent node.

We use Figure 6 to illustrate the situation when D2 or D6 is applied over the contingent ab constraint. The rightmost part of this figure is an arbitrary triangle abc where one of the rules is applicable, while the leftmost part is motivated by the proof below.

In the following we do not care if the edges are conditional or requirement: Only the weights of the derived edges are important. We follow a critical chain and see how




Fig. 6. Situation where D2 or D6 is applied.

the source and target change as we continually derive new edges. Applying D2 or D6 gives a new edge ac where the source changes from b to a . We now investigate how derivations can move the source back to b and show that all derivations using the edge which resulted from moving the source back to b are redundant. We already know that the source can only move back to b if the target moves from c . Otherwise there would be a cycle contradicting lemma 1. So there must be a list, $\langle c, \dots, y \rangle$ of one or more nodes that the target moves along. Since the source moves only over positive edges (using the weight of the negative in case of contingent) there must be another list $\langle a, \dots, x \rangle$ that the source moves over before reaching b again. The final edge derived before reaching b is xy , whose edge will be a sum of negative weights along $\langle c, \dots, y \rangle$ where negative requirement edges and positive contingent edges contribute, and positive weights along $\langle a, \dots, x \rangle$ where positive requirement edges and negative contingent edges contribute. For the source to return to b , the weight of xy must be negative and there must be a positive edge bx . Then we can apply a rule deriving the edge by . We can determine that this edge is redundant by applying derivations to it. If by is positive it is redundant since there is a tighter implicit constraint along the strictly negative bcy path, as discussed before the theorem. If by is negative we apply derivation to move the source towards x . In this way we continue to apply derivations until we get a positive edge zy or the source reaches x . If this happens the derived edge must have a larger value than the already present xy edge, and be redundant, or we have derived a cycle contradicting lemma 1. This can also be seen by observing that derivations start with the weight of xy , which can only increase along the derivation chain.

If we instead get a positive edge zy along the derivations we can show that there is a tighter constraint implicit here. We know $z \neq x$. When first deriving xy there was a negative edge from z to some node t in the $\langle c, \dots, y \rangle$ list. If $t = y$ we arrive with a larger weighted edge (positive) ty this time and it is redundant. If $t \neq y$ there is an implicit tighter negative constraint zy . So again the zt edge is redundant.

So by is already explicitly or implicitly covered and hence redundant for DC-verification. Therefore it is not part of a critical chain. \square

This entails that along a critical chain each contingent constraint can only be part of at most two derivations: One using D1 and one using D2 or D6.

Algorithm 3: The GlobalDC Algorithm

```

function GLOBAL-DC ( $G - STNU$ )
   $Interesting \leftarrow \{\text{All edges of } G\}$ 
  repeat
    for each edge  $e$  in  $G$  do
       $Interesting \leftarrow Interesting \setminus \{e\}$ 
      for each rule (Figure 3) applicable with  $e$  as focus do
        Derive new edges  $z_i$ 
        for each added edge  $z_i$  do
           $Interesting \leftarrow Interesting \cup \{z_i\}$ 
          if not locally consistent then return false
          if negative cycle created then return false
        end
      end
    end
  until  $Interesting$  is empty
  return true

```

7 GlobalDC

We will apply the theorem above to the new algorithm GlobalDC (Algorithm 3). Given a full STNU this algorithm applies the derivation rules of Figure 3 globally, i.e., with all edges as focus in all possible *triangles* (giving an iteration $O(n^3)$ run-time). It does this until there are no more changes detected over a global iteration. The structure of GlobalDC is hence directly inspired by the Bellman-Ford algorithm [13]. Non-DC STNUs are detected in the same way as FastIDC, by checking locally that there is no squeeze of contingent constraints and globally that there is no negative cycle.

This full DC algorithm can be compared with how an incremental algorithm (FastIDC) could be used to verify full DC, i.e., by adding edges from the full graph one at a time and doing derivations until done. Note that the order in which the derivation rules are applied to edges does not affect the correctness of FastIDC, only its run-time.

Given a DC STNU, GlobalDC will use the same derivation rules as FastIDC and therefore cannot generate tighter constraints. Since the same mechanism is used for detecting non-DC STNUs, both FastIDC and GlobalDC will indicate that the STNU is DC.

Given a non-DC STNU, there exists a sequence of derivations that will let FastIDC decide this. Since GlobalDC performs all possible derivations in each iteration, it will do all derivations that FastIDC does in the same sequence. Again, the same mechanism is used for detecting non-DC STNUs, and both FastIDC and GlobalDC will indicate that the STNU is non-DC.

The key to analyzing the complexity of GlobalDC is the realization that we can stop deriving new constraints as soon as we have derived all critical chains: These are the only derivations that are required for detecting whether the STNU is DC or not.

The target of derived edges must eventually move. It can move at most n times, since it always moves to a node guaranteed to execute earlier. In-between two such




Fig. 7. Example graph in quiescence.




Fig. 8. Derivations resulting from adding the $i \rightarrow e$ edge.

moves the source can move between at most n nodes. Between each move of the source there can be one application of D8/D9, resulting in a chain of length $2n$ between each of the target moves. Together this bounds the longest critical chain by $2n^2$.

An example will illustrate how we can shrink the length of critical chains. Figure 7 shows a graph where no more derivations can be made. In Figure 8 a negative edge ie is added to the graph and GlobalDC is used to update the graph with this increment.

Figure 9 shows the critical chain of edge ac at this point. Here we see as mentioned before that the source of the derived edge can move many times in sequence without the target moving in-between. In the example chain this is shown by the sequential D7 derivations. For requirement edges in general such a sequence may also include D4 derivations. Conditional edges can also induce sequences of moving sources through derivation rules D3 and D5.

All these derivations leading to sequential movement of the source require it to pass over requirement edges. If we had access to the shortest paths along requirement edges all these movements could in fact be derived in one global iteration. The source would be moved to all destinations at once and would not be replaced later since it had already followed a shortest path making the derived edge as tight as possible. Of course derivation rules may change the shortest paths, but if we added an APSP calculation to every global iteration we would compress the critical chains so that there would be no repeated application of sources moving along requirement constraints.

Figure 10 shows how several applications of D7 and two of D3 are compressed by the availability of shortest path edges.



Fig. 9. The critical chain of edge ac , derived in Figure 8.



Fig. 10. Critical chain compressed using shortest paths.

GlobalDC with the addition of APSP calculations in each iteration is still sound and complete since the APSP calculations only make more implicit constraints explicit. The run-time complexity is also preserved since each iteration was already $O(n^3)$ (applying rules to all focus edges). We now give an upper bound of the critical chain length:

Theorem 2. *The length of the longest critical chain in GlobalDC with APSP is $\leq 7n$.*

Proof. To be able to prove this we need the results of theorem 1. We will refer to derivations that can only occur once along a critical chain, i.e. D1, D2 and D6, as *limited derivations*.

What is the longest sequence in a critical chain consisting only of requirement edges such that it does not use any limited derivations? The only non-limited derivation rules that result in a requirement edge are D4, D7 and D8/D9. The last two require a conditional edge as focus, and can therefore only be at the start of such a sequence. We know that due to APSP there can only be one of D4/D7 in a row. Therefore the longest requirement-only sequence not using limited derivations starts with D8/D9 which is followed by D4/D7 for a total length of 2.

The longest sequence consisting of only conditional edges not using limited derivations must start with D5. It can then be continued only by D3. As we have access to shortest paths there can be at most one D3 in any sequence of only conditional edges.

In summary the longest sequences of the same type, requirement or conditional, not using limited derivations, are of length 2.

It is not possible to interleave the length-2 sequences of conditional edges with requirement edges more than once without changing the conditioning node of the conditional edges. To see this suppose we have a requirement edge which derives a conditional edge conditioned on B . This means that the edge is pointing towards A being the start of the contingent duration ending in B . If derivations now takes this edge into a requirement edge this edge must point towards A as well since the only way of going from conditional to requirement is via D8/D9 which preserves the target. If the target of the requirement edge later were to move (such targets only move forwards) it would become impossible to later invoke D5 for going back to conditional, because D5 requires the requirement edge to point towards a node that is after A . So in order for derivations to come back to a conditional edge again by D5 the target must stay at A . But then D5 cannot be applicable, for the same reason: It must point towards a node after A . So it is not possible to interleave these sequences.

This gives us the longest possible sequence without using limited derivations. It starts with a requirement sequence followed by a conditional sequence again followed by a requirement sequence. Such a sequence can have a length of at most 6. An issue here is that if a conditional edge conditioned on for instance B is part of the chain a D1 derivation involving B cannot also occur in the chain since this contingent constraint has already been passed. This means that it does not matter which of derivation D1 or D5 is used to introduce a conditioning node into the chain. The limitation applies to them both.

In conclusion this lets us construct an upper bound on the number of derivations in a critical chain. We have sequences of length 6 and these are interleaved with the n derivations of type D2 and D6 for a total of at most $7n$ derivations. \square

Therefore all critical chains will have been generated after at most $7n$ iterations of GlobalDC. If we can iterate $7n$ times without detecting that an STNU is non-DC, it must be DC. With a limitation of $7n$ iterations, GlobalDC verifies DC in $O(n^4)$.

8 A REVISED MMV ALGORITHM

We have described a new algorithm called GlobalDC and seen that it is $O(n^4)$. Compared to MMV, the following similarities and differences exist.

First, GlobalDC and MMV both interleave the application of derivation rules with the calculation of APSP distances and the detection of local inconsistencies and negative cycles. In MMV some of this is hidden in the pseudo-controllability test, but the actual conditions being tested are equivalent.

Second, GlobalDC works in an EDG whereas MMV works in an STNU extended with wait constraints. These structures represent the same underlying constraints and the difference is not essential.

Third, GlobalDC lacks SR2, which is half of the original Simple Regression (SR) rule. Making this change in MMV will greatly speed it up in practice. Since it runs in an APSP graph it is reasonable to expect, on average, half of the nodes to be after a derived wait. This change will then cut the needed regression in MMV to half of that of the original version.

Fourth, GlobalDC stops after $7n$ iterations. Given the similarities above and the fact that the theorem about critical chain lengths directly carries over, MMV can also stop after $7n$ iterations without affecting correctness. The modified MMV can then decide DC in $O(n^4)$ time. We formulate this as a theorem.

Theorem 3. *The classical MMV algorithm for deciding dynamic controllability of an STNU can, with the small modifications shown in Algorithm 4, decide dynamic controllability in time $O(n^4)$.*

9 RELATED AND FUTURE WORK

Recently several papers [14, 15] have examined the use of Timed Game Automata (TGA) for both verification and execution of STNUs. These solutions work on a smaller

Algorithm 4: The revised MMV Algorithm

```

function Revised-MMV( $G - STNU$ )
   $Interesting \leftarrow \{\text{All edges of } G\}$ 
   $iterations \leftarrow 0$ 
  repeat
    if not pseudo-controllable( $G$ ) then
      | return false
    Compare edges and add all edges which were changed since last iteration to
     $Interesting$ 
    for each edge  $e$  in  $Interesting$  do
      |  $Interesting \leftarrow Interesting \setminus \{e\}$ 
      | for each triangle  $ABC$  containing  $e$  do
      | | tighten  $ABC$  according to Figure 4 except SR2
      | end
    end
     $iterations \leftarrow iterations + 1$ 
  until  $Interesting$  is empty or  $iterations = 7n$ 
  return true

```

scale and do not exploit the inherent structure of STNUs as distance graphs. Therefore they are more useful in networks that are small in size but involve choice and resources which cannot be handled by pure STNU algorithms.

At the time most of the research presented here was conducted the fastest algorithm for verifying dynamic controllability of an STNU was $O(n^4)$ [6]. Very recent activities however have converged on an algorithm which performs this in $O(n^3)$ [7, 16].

Since there are now so many algorithms available for verifying DC it is important to find good benchmarks that can be used both to identify weaknesses but also to establish run-times of the algorithms in relevant use cases. This constitutes a large study that need to be carried out by the community in the near future. Also, execution of the produced networks need to be investigated further in the spirit of [17, 18].

10 CONCLUSION

We have proven that with a small modification the classical “MMV” dynamic controllability algorithm, which in its original form is pseudo-polynomial, finishes in $O(n^4)$ time. The modified algorithm is still a viable option for determining whether an STNU is dynamically controllable. Compared to other algorithms, it offers a simpler and more intuitive theory. It is also an entry level algorithm which many familiarize with before implementing more advanced algorithms. As such it is an excellent choice for regression testing of the more complicated algorithms.

In this paper we also showed that there is no reason for MMV to regress over negative edges, a result that can be used to improve performance further.

ACKNOWLEDGMENTS This work is partially supported by the Swedish Research Council (VR) Linnaeus Center CADICS, the ELLIIT network organization for Information and Com-

munication Technology, the Swedish Foundation for Strategic Research (CUAS Project), the EU FP7 project SHERPA (grant agreement 600958), and Vinnova NFFP6 Project 2013-01206.

References

- [1] R. Dechter, I. Meiri, and J. Pearl, “Temporal Constraint Networks,” *Artificial Intelligence*, vol. 49, no. 1-3, pp. 61–95, 1991.
- [2] T. Vidal and M. Ghallab, “Dealing with Uncertain Durations in Temporal Constraint Networks Dedicated to Planning,” in *Proc. ECAI*, 1996.
- [3] N. Muscettola, P. Morris, and I. Tsamardinos, “Reformulating temporal plans for efficient execution,” in *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 1998.
- [4] P. Morris, N. Muscettola, and T. Vidal, “Dynamic Control of Plans with Temporal Uncertainty,” in *Proc. IJCAI*, 2001.
- [5] P. Morris and N. Muscettola, “Temporal dynamic controllability revisited,” in *Proc. AAAI*, 2005.
- [6] P. Morris, “A Structural Characterization of Temporal Dynamic Controllability,” in *Proc. CP*, 2006.
- [7] —, “Dynamic Controllability and Dispatchability Relationships,” in *Proc. CPAIOR*, 2014.
- [8] J. Stedl and B. Williams, “A Fast Incremental Dynamic Controllability Algorithm,” in *Proc. ICAPS Workshop on Plan Execution*, 2005.
- [9] J. A. Shah, J. Stedl, B. C. Williams, and P. Robertson, “A Fast Incremental Algorithm for Maintaining Dispatchability of Partially Controllable Plans,” in *Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS)*, 2007. [Online]. Available: <http://dblp.uni-trier.de/db/conf/aips/icaps2007.html/#ShahSWR07>
- [10] M. Nilsson, J. Kvarnström, and P. Doherty, “Incremental Dynamic Controllability Revisited,” in *Proc. ICAPS*, 2013.
- [11] M. Nilsson, J. Kvarnström, and P. Doherty, “Classical Dynamic Controllability Revisited: A Tighter Bound on the Classical Algorithm,” in *Proc. ICAART*, 2014.
- [12] J. L. Stedl, “Managing Temporal Uncertainty Under Limited Communication: A Formal Model of Tight and Loose Team Coordination,” Master’s thesis, Massachusetts Institute of Technology, 2004.
- [13] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson, *Introduction to Algorithms*. McGraw-Hill Higher Education, 2001.
- [14] A. Cimatti, L. Hunsberger, A. Micheli, and M. Roveri, “Using Timed Game Automata to Synthesize Execution Strategies for Simple Temporal Networks with Uncertainty,” in *Proc. AAAI*, 2014.
- [15] A. Cesta, A. Finzi, S. Fratini, A. Orlandini, and E. Tronci, “Analyzing Flexible Timeline-based Plans,” in *Proc. ECAI*, 2010.
- [16] M. Nilsson, J. Kvarnström, and P. Doherty, “EfficientIDC: A Faster Incremental Dynamic Controllability Algorithm,” in *Proc. ICAPS*, 2014.
- [17] L. Hunsberger, “A Fast Incremental Algorithm for Managing the Execution of Dynamically Controllable Temporal Networks,” in *Proc. TIME*, 2010.
- [18] —, “A Faster Execution Algorithm for Dynamically Controllable STNUs,” in *Proc. TIME*. IEEE, 2013, pp. 26–33.