

Multiagent Simple Temporal Problem: The Arc-Consistency Approach

Shufeng Kong,¹ Jae Hee Lee,¹ Sanjiang Li^{1,2}

¹Centre for Quantum Software and Information, FEIT, University of Technology Sydney, Australia

²UTS-AMSS Joint Research Laboratory, AMSS, Chinese Academy of Sciences, China
shufeng.kong@student.uts.edu.au, {jaehee.lee, sanjiang.li}@uts.edu.au

Abstract

The Simple Temporal Problem (STP) is a fundamental temporal reasoning problem and has recently been extended to the Multiagent Simple Temporal Problem (MaSTP). In this paper we present a novel approach that is based on enforcing arc-consistency (AC) on the input (multiagent) simple temporal network. We show that the AC-based approach is sufficient for solving both the STP and MaSTP and provide efficient algorithms for them. As our AC-based approach does not impose new constraints between agents, it does not violate the privacy of the agents and is superior to the state-of-the-art approach to MaSTP. Empirical evaluations on diverse benchmark datasets also show that our AC-based algorithms for STP and MaSTP are significantly more efficient than existing approaches.

1 Introduction

The Simple Temporal Problem (STP) (Dechter, Meiri, and Pearl 1991) is arguably the most well-known quantitative temporal representation framework in AI. The STP considers time points as the variables and represents temporal information by a set of unary or binary constraints, each specifying an interval on the real line. Since its introduction in 1991, the STP has become an essential sub-problem in planning or scheduling problem (Barták, Morris, and Venable 2014).

While the STP is initially introduced for a single scheduling agent and is solved by centralized algorithms, many real-world applications involve multiple agents who interact with each other to find a solution like the following example:

Example 1. When Alice is looking for a position at company X , she might need to arrange an interview appointment with X . Suppose that her colleague Bob is also applying for the position and Alice and Bob are both applying for another position at another company Y . To represent and solve such an interview scheduling problem, we need a multiagent framework (see Figure 1 for an illustration).

Recently, the extension of STP to multiagent STP (MaSTP) has been provided in (Boerkoel and Durfee 2013), which presents a formal definition of the MaSTP as well as

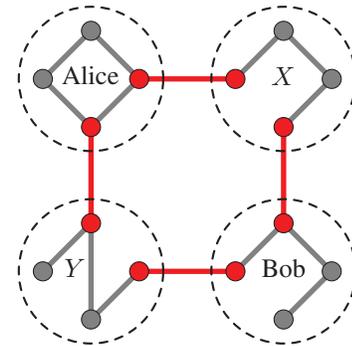


Figure 1: An illustration of Example 1. Alice, Bob, company X , company Y are four agents, each owning a local simple temporal network. The circles represent variables and edges constraints. Red edges represent constraints that are shared by two different agents.

a distributed algorithm, called $D\Delta$ PPC, for computing the complete joint solution space.

However, as $D\Delta$ PPC is based on the P^3C algorithm (Planken, de Weerd, and van der Krogt 2008), which triangulates the input constraint graph, it has the drawback of creating new constraints between agents that are possibly not directly connected. In Figure 1, $D\Delta$ PPC triangulates the inner cycle by adding at least one new constraint either between X and Y or between Alice and Bob. Neither of these new constraints are desirable, as they introduce constraints between two previously not directly connected agents and thus present a threat to the privacy of the relevant agents.

As the recent technological advancements have allowed for solving larger problems that are highly interwoven and dependent on each other, efficiency and privacy have become critical requirements. To address this challenge, we propose a new approach to solve the MaSTP, which is based on *arc-consistency*.

A constraint R between two variables x, y is called arc-consistent (AC), if for every value d_x from the domain of x there is a value d_y in the domain of y such that $(d_x, d_y) \in R$. While AC is an important tool for solving finite (multi-agent) constraint satisfaction problems (CSPs) (Montanari 1974; Baudot and Deville 1997; Nguyen and Deville 1998;

Hamadi 2002) at first glance it is not clear how it can be applied to solving CSPs with real domains such as the STP, because either the existing AC algorithms are fine-grained and work with each single element of a domain to enforce AC, which is impossible for real domains, or they are coarse-grained, but cannot guarantee their termination, as real domains can be infinitely refined when constraints are propagated.¹

Our contributions in this paper are as follows.

- We provide the first AC-based approach for solving STP and analyze its computational complexity (Section 3).
- We provide the first AC-based approach for solving multi-agent STP, which preserves the privacy of the agents, and analyze its computational complexity (Section 4).
- We experimentally show that both our centralized and distributed algorithms outperform their existing counterparts for solving STP (Section 5).

The next section gives a formal introduction to STP.

2 The Simple Temporal Problem

This section briefly introduces the STP. Details can be found in (Dechter, Meiri, and Pearl 1991).

The *simple temporal problem* (STP) is a constraint satisfaction problem where each constraint is a set of linear inequalities of the form

$$a_{vw} \leq w - v \leq b_{vw}, \quad (1)$$

where a_{vw}, b_{vw} are constants and v, w are variables defined on a continuous domain representing time points. The constraint in (1) is abbreviated as $I_{vw} = [a_{vw}, b_{vw}]$. As (1) is equivalent to $-b_{vw} \leq v - w \leq -a_{vw}$, we also obtain $I_{vw} = I_{vw}^{-1} = [-b_{vw}, -a_{vw}]$. The *domain* of each variable v is an interval $I_v = [a_v, b_v]$, where I_v could be a singleton or empty. Assume that o is a special *auxiliary* variable that represents the fixed zero temporal point. Then the domain I_v can also be regarded as a constraint from o to v and $I_v = [a_v, b_v] = [a_{ov}, b_{ov}] = I_{ov}$.

Algebraic operations on STP constraints are defined as follows. The *intersection* of two STP constraints defined on variables v, w yields a new constraint over v, w that represents the conjunction of the constraints. It is defined as $I_{vw} \cap I'_{vw} := [\max\{a_{vw}, a'_{vw}\}, \min\{b_{vw}, b'_{vw}\}]$.

The *composition* of an STP constraint I_{vu} over variables v, u and another STP constraint I_{uw} over u, w yields a new STP constraint over v, w that is inferred from the other two constraints and is defined as $I_{vu} \otimes I_{uw} := [a_{vu} + a_{uw}, b_{vu} + b_{uw}]$. Here we require that $[a, b] \otimes \emptyset = \emptyset$ for any $a \leq b$.

Remark 1. For STP constraints, the composition and intersection are associative and, as noted in (Dechter, Meiri, and Pearl 1991), composition distributes over *non-empty* intersection for intervals, i.e., $I \otimes (J \cap K) = (I \otimes J) \cap (I \otimes K)$ for any three intervals I, J, K such that $J \cap K \neq \emptyset$.

¹Dechter, Meiri, and Pearl (1991) for example, suggest discretizing the domains to overcome this issue, in which case the total number of constraint propagations would depend on the sizes of the domains. The performance of our AC algorithm for (multi-agent) STP does not depend on the sizes of the domains.

Definition 1. An instance of STP is called a *simple temporal network* (STN) and is a tuple $\langle V, D, C \rangle$, where V is a finite set of variables, $D = \{I_v \mid v \in V\}$ is a set of intervals, and C is a set of STP constraints defined on V .

We assume that all variables in V appear in C and at *most* one constraint exists between any pair of variables v and w . Moreover, if $I_{vw} = [a, b]$ is the constraint in C from v to w , we always assume that the constraint $I_{wv} = I_{vw}^{-1} = [-b, -a]$ is also in C . As previously mentioned, the domain I_v of each variable v can be regarded as either a unary constraint, or a binary constraint $I_{ov} = I_v$, where o is a fixed variable representing the zero time point.

An STN naturally induces a graph in the following sense.

Definition 2. The *constraint graph* $G_{\mathcal{N}} = (V, E)$ of an STN $\mathcal{N} = \langle V, D, C \rangle$ is an undirected graph, where the set E of edges consists of constrained unordered pairs of variables in C , i.e.,

$$E = \{\{v, w\} \mid v, w \in V, v \neq w, I_{vw} \in C\}.$$

Let $G_{\mathcal{N}} = (V, E)$ be the constraint graph of an STN \mathcal{N} . We can use a labelled *directed* graph to illustrate \mathcal{N} , where for any undirected edge $\{v, w\} \in E$ there is exactly one directed edge (v, w) that is labelled with the corresponding interval $[a_{vw}, b_{vw}]$.

A *path* π from v to w in $G_{\mathcal{N}}$ is a sequence of variables u_0, u_1, \dots, u_k such that $v = u_0, w = u_k$, and $\{u_s, u_{s+1}\}$ is an edge in E for each $s = 0, \dots, k-1$ (k is called the length of π). We write $\otimes \pi$ for the composition of all these $I_{u_s, u_{s+1}}$, i.e.,

$$\otimes \pi = I_{u_0, u_1} \otimes I_{u_1, u_2} \otimes \dots \otimes I_{u_{k-1}, u_k} \quad (2)$$

If $v = w$, then we call π a *cycle* at v . For a cycle π , let $[a, b] = \otimes \pi$. We call π a *negative cycle* if $b < 0$.

Definition 3. A *solution* of an STN $\mathcal{N} = \langle V, D, C \rangle$ is an assignment, that assigns to each variable $v \in V$ a time point from $I_v \in D$ such that all constraints in C are satisfied. \mathcal{N} is said to be *consistent* if \mathcal{N} has a solution. Two STNs are said to be *equivalent* if they have the same solution set.

Definition 4 (Minimality). Let $\mathcal{N} = \langle V, D, C \rangle$ be a consistent STN and let v and w be variables in V . A constraint I_{vw} from v to w is said to be *minimal* if every assignment that assigns time points from domains I_v and I_w to v and w , respectively, and satisfies I_{vw} can be extended to a solution of \mathcal{N} . A domain I_v of $v \in V$ is said to be minimal if every assignment of a time point from I_v to v can be extended to a solution of \mathcal{N} . We say \mathcal{N} is minimal if every constraint in C as well as every domain in D is minimal (note that, since we regard domains as constraints between the zero time point o and variables, we also require the domains to be minimal).

3 Solving the STP with Arc-Consistency

In this section we show that enforcing arc-consistency is sufficient to solve the STP.

Definition 5. Let $\mathcal{N} = \langle V, D, C \rangle$ be an STN. Suppose v and w are two variables in V , I_v and I_w are, respectively, their domains, and I_{vw} is a constraint in C from v to w .

We say that I_{vw} is *arc-consistent* (AC) (relative to I_v and I_w) if for any $t_v \in I_v$ there exists some $t_w \in I_w$ such that $t_w - t_v \in I_{vw}$, i.e., $a_{vw} \leq t_w - t_v \leq b_{vw}$. We say that \mathcal{N} is AC if both I_{vw} and I_w are AC for every constraint $I_{vw} \in C$.

An STN $\mathcal{N}' = \langle V, D', C \rangle$ with $D' = \{I'_v \mid v \in V\}$ is called the *AC-closure* of \mathcal{N} , if \mathcal{N}' is the largest arc-consistent STN which is equivalent to \mathcal{N} , in the sense that for every other arc-consistent STN $\mathcal{N}'' = \langle V, D'', C \rangle$ with $D'' = \{I''_v \mid v \in V\}$, we have that $I''_v \subseteq I'_v$ for all $v \in V$.

Lemma 1. *Let $\mathcal{N} = \langle V, D, C \rangle$ be an STN and $v, w \in V$ two variables that are constrained by I_{vw} in C . Then I_{vw} is arc-consistent relative to I_v and I_w , iff $I_v \subseteq I_w \otimes I_{vw}$.*

Proof. It suffices to show that

$$I_w \otimes I_{vw} = \{x \in \mathbb{R} \mid \exists y \in I_w \text{ s.t. } y - x \in I_{vw}\} \quad (3)$$

Let $I_v = [a, b]$, $I_w = [c, d]$ and $I_{vw} = [e, f]$. Then

$$\begin{aligned} & \{x \in \mathbb{R} \mid \exists y \in I_w \text{ s.t. } y - x \in I_{vw}\} \\ &= \{x \in \mathbb{R} \mid \exists c \leq y \leq d \text{ s.t. } e \leq y - x \leq f\} \\ &= \{x \in \mathbb{R} \mid \exists c \leq y \leq d \text{ s.t. } y - f \leq x \leq y - e\} \\ &= \{x \in \mathbb{R} \mid c - f \leq x \leq d - e\} \\ &= [c, f] \otimes [-f, -e] = I_w \otimes I_{vw}, \end{aligned}$$

which proves Eq. (3). \square

Lemma 2. *Let $\mathcal{N} = \langle V, D, C \rangle$ be an arc-consistent STN and $v, w \in V$ two variables that are constrained by I_{vw} in C . Then $I_v \subseteq I_w \otimes I_{vw}$.*

Proof. This follows directly from Lemma 1 and that I_{vw} is AC relative to I_v and I_w . \square

The following result directly follows from Lemma 2.

Corollary 3. *Let $\mathcal{N} = \langle V, D, C \rangle$ be an arc-consistent STN. Let π be a path in \mathcal{N} from w to v . Then $I_v \subseteq I_w \otimes \bigotimes \pi$.*

Lemma 4. *Let $\mathcal{N} = \langle V, D, C \rangle$ be an arc-consistent STN and v, w variables in V . If \mathcal{N} is consistent, then $I_v \subseteq I_w \otimes I_{vw}^m$, where I_{vw}^m is the minimal constraint from w to v .*

Proof. Since \mathcal{N} is consistent, I_{vw}^m is nonempty. Recall that I_{vw}^m is the intersection of the compositions along all paths in \mathcal{N} from w to v (cf. (Dechter, Meiri, and Pearl 1991, §3)) and composition distributes over non-empty intersection for intervals. The result follows directly from Corollary 3. \square

Lemma 5 ((Shostak 1981)). *Suppose $\mathcal{N} = \langle V, D, C \rangle$ is an STN. Then \mathcal{N} is inconsistent if and only if there exists a negative cycle.*

Lemma 6. *Given a consistent STN $\mathcal{N} = \langle V, D, C \rangle$ with $n = |V|$, for any path π of length $\geq n$ there is a path π' of length $< n$ such that $\bigotimes \pi' \subseteq \bigotimes \pi$.*

Proof. Since the length of π is $\geq n$, π must have a cycle at a variable v . As the cycle is not negative, removing the cycle and leaving only v in the path results in a path π' with $\bigotimes \pi' \subseteq \bigotimes \pi$. Repeating this procedure until there is no cycle gives the desired result. \square

Lemma 7. *Let $\mathcal{N} = \langle V, D, C \rangle$ be an STN and \mathcal{N}' its AC-closure. Then \mathcal{N} is consistent iff \mathcal{N}' has no empty domain.*

Proof. We prove \mathcal{N} is inconsistent iff \mathcal{N}' has an empty domain. As \mathcal{N} and \mathcal{N}' are equivalent, if \mathcal{N}' has an empty domain, then \mathcal{N} is inconsistent.

Now suppose \mathcal{N} is inconsistent. Then by Lemma 5, there exists a negative cycle π in \mathcal{N} at some w such that $\bigotimes \pi = [l, h]$ with $h < 0$. Now let v be a variable in \mathcal{N} with $I_{vw} = [e, f]$ and let $I'_v = [a, b]$, $I'_w = [c, d]$ be the domains of v and w in \mathcal{N}' , respectively. Choose $k \in \mathbb{N}$ sufficiently large, such that $kh < b - d - f$. Then, by Lemma 4 we have

$$\begin{aligned} I'_v &\subseteq I'_w \otimes \left(\bigotimes \pi^k \otimes I_{vw} \right) \\ &= [c, d] \otimes ([kl, kh] \otimes [e, f]) \\ &= [c + kl + e, d + kh + f], \end{aligned} \quad (4)$$

where π^k is the concatenation of k copies of path π . Because $kh < b - d - f$, (4) is possible only if I'_v is empty. \square

Theorem 8. *Let $\mathcal{N} = \langle V, D, C \rangle$ be a consistent STN and \mathcal{N}' its AC-closure. Then all domains in \mathcal{N}' are minimal.*

Proof. If the constraint graph $G_{\mathcal{N}}$ is connected, i.e., for any two variables v, w , there is a path in $G_{\mathcal{N}}$ that connects v to w , then we may replace the constraint from v to w with the nonempty minimal constraint I_{vw}^m (or add I_{vw}^m , if there was no constraint between v and w). We write the refined STN as \mathcal{N}^* . For any two variables v, w , by Lemma 4, I_v is contained in $I_w \otimes I_{vw}^m$ and I_w is contained in $I_v \otimes I_{vw}^m$. This shows that \mathcal{N}^* is the same as the minimal STN of \mathcal{N} , and thus, establishes the minimality of each I_v .

In case the constraint graph is disconnected, we consider the restriction of \mathcal{N} to its connected components instead. The same result applies. \square

Two special solutions can be constructed if \mathcal{N} is arc-consistent and has no empty domain.

Proposition 9. *Let $\mathcal{N} = \langle V, D, C \rangle$ be an arc-consistent STN with $D = \{I_v \mid v \in V\}$ and $I_v = [a_v, b_v]$ for each v . If no I_v is empty, then the assignments $A = \{a_v \mid v \in V\}$ and $B = \{b_v \mid v \in V\}$ are two solutions of \mathcal{N} .*

Proof. Let $\mathcal{N}' = \langle V, D', C \rangle$ be the minimal STN of \mathcal{N} . By Theorem 8, we have $D' = D$ and \mathcal{N}' is equivalent to \mathcal{N} . The above claim follows as the assignments $A = \{a_v \mid v \in V\}$ and $B = \{b_v \mid v \in V\}$ are two solutions of the minimal STN \mathcal{N}' (cf. (Dechter 2003, Corollary 3.2)). \square

Theorem 10. *Enforcing AC is sufficient to solve STP.*

Proof. Let \mathcal{N} be an STN and \mathcal{N}' its AC-closure. If \mathcal{N}' has an empty domain, then \mathcal{N} has no solution by Lemma 7. If \mathcal{N}' does not have an empty domain, then we can use Proposition 9 to find a solution. \square

Remark 2. (i) As solving an STN is equivalent to solving a system of linear inequalities, the solution set of an STN is a convex polyhedron. Thus any convex combination of the two solutions A and B is again a solution of the STN. (ii) Enforcing AC can in essence find all solutions of an STN: Suppose

Algorithm 1: ACSTP

Input : An STN $\mathcal{N} = \langle V, D, C \rangle$ and its constraint graph $G = (V, E)$, where $|V| = n$.

Output: An equivalent network that is AC, or “inconsistent”.

```
1  $Q \leftarrow \emptyset$ 
2 for  $k \leftarrow 1$  to  $n$  do
3   foreach  $v \in V$  do
4      $I'_v \leftarrow I_v$ 
5     foreach  $w \in V$  s.t.  $\{v, w\} \in E$  do
6        $I_v \leftarrow I_v \cap I_w \otimes I_{vw}$ 
7     if  $I_v = \emptyset$  then return “inconsistent”
8     if  $I'_v = I_v$  then  $Q \leftarrow Q \cup \{v\}$ 
9     else  $Q \leftarrow Q \setminus \{v\}$ 
10  if  $\#Q = n$  then return  $\mathcal{N}$ 
11 return “inconsistent”
```

\mathcal{N} is arc-consistent and has no empty domain. We pick an arbitrary variable v that has not been instantiated yet, then assign any value from D_v to v , and enforce AC on the resulting network. We repeat this process until all variables are instantiated. (iii) Proposition 9 can also be obtained by first showing that STP constraints are both max/min-closed, and then using the result in (Jeavons and Cooper 1995, Thm 4.2), which states that the AC-closure of a constraint network over max/min-closed constraints have the maximal and the minimal values of the domains as two solutions. As a consequence of this, Theorem 8 can also be obtained, because the solution set of an STN is convex (cf. Remark 2 (i)).

A Centralized AC Algorithm for the STP

In this section we propose an AC algorithm, called ACSTP, to solve STNs. The algorithm is presented as Algorithm 1.

Theorem 11. *Given an input STN \mathcal{N} , Algorithm 1 returns “inconsistent” if \mathcal{N} is inconsistent. Otherwise, it returns the AC-closure of \mathcal{N} .*

Proof. We first note that intersection and composition of constraints do not change the solution set of the input STN \mathcal{N} . This has two implications: First, if a domain I_v becomes empty during the process of the algorithm, then the solution set of \mathcal{N} is empty and \mathcal{N} is inconsistent. Second, if the algorithm terminates and its output \mathcal{N}' is AC, then \mathcal{N}' is the AC-closure of \mathcal{N} . Consequently, it suffices to show that if the algorithm terminates and returns \mathcal{N}' , then \mathcal{N}' is AC.

We first consider the case, where the algorithm returns \mathcal{N}' in line 10 at the k th iteration of the for-loop (lines 2–10) for some $1 \leq k \leq n$. We show that \mathcal{N}' is AC. Let I_v^k be the domain of v obtained after the k th iteration of the for-loop. Due to lines 6 and 8, we have for all $\{v, w\} \in E$ that $I_v^k \subseteq I_v^{k-1} \cap (I_w^{k-1} \otimes I_{vw})$ and $I_w^{k-1} = I_w^k$. Thus we have for $\{v, w\} \in E$ that $I_v^k \subseteq I_w^k \otimes I_{vw}$, which is by Lemma 1 equivalent to saying that I_{vw} is AC w.r.t. domains I_v^k and I_w^k . Hence, the output \mathcal{N}' is AC.

Now suppose that the algorithm exited in line 11 returning “inconsistent”. Thus, at the n th iteration of the for-loop we have $\#Q < n$ in line 10. We prove that \mathcal{N} is inconsistent by contradiction. Assume that \mathcal{N} is consistent. For any $v \in V$ and any $k \geq 1$, we write Π_v^k for the set of paths from o (the auxiliary variable denoting the zero time point) to v with length $\leq k$ in the constraint graph of \mathcal{N} . We claim

$$I_v^{k-1} \subseteq \bigcap_{\pi \in \Pi_v^k} \pi \quad (5)$$

for any $k \geq 1$. Then, with I_v^m being the minimal domain of v , we have

$$I_v^m \subseteq I_v^{n-1} \subseteq \bigcap_{\pi \in \Pi_v^n} \pi = I_v^m,$$

because I_v^m is the intersection of the compositions along all paths in \mathcal{N} from o to v (cf. (Dechter, Meiri, and Pearl 1991, §3)), where it suffices to only build compositions along paths of length $\leq n$ by Lemma 6. Thus $I_v^n = I_v^{n-1} = I_v^m$ for all $v \in V$, which is a contradiction to our assumption that at the n th iteration of the for-loop we have $\#Q < n$ in line 10.

We now prove (5) by using induction on k . First, for $k = 1$, since Π_v^1 contains only one path of length 1 (i.e., the edge $\{o, v\}$), we have $I_v^0 = I_v = \bigcap_{\pi \in \Pi_v^1} \pi$. Now suppose (5) is true for $k - 2$ for all $w \in V$. Then by line 6 and our induction hypothesis we have

$$\begin{aligned} I_v^{k-1} &\subseteq I_v^{k-2} \cap \left(\bigcap_w I_w^{k-2} \otimes I_{vw} \right) \\ &\subseteq I_v^{k-2} \cap \left(\bigcap_w \left(\bigcap_{\pi \in \Pi_w^{k-1}} \pi \right) \otimes I_{vw} \right) \\ &\subseteq \left(\bigcap_{\pi \in \Pi_v^{k-1}} \pi \right) \cap \left(\bigcap_{(\pi \in \Pi_v^k) \wedge (|\pi| \geq 2)} \pi \right) \\ &= \bigcap_{\pi \in \Pi_v^k} \pi, \end{aligned}$$

which proves (5). \square

Theorem 12. *Algorithm 1 runs in time $O(en)$, where e is the number of edges of the constraint graph of the input STN and n is the number of variables.*

Proof. There are at most n iterations of the for-loop and each iteration involves $O(e)$ operations. \square

Remark 3. Algorithm 1 can also be understood as computing the shortest path from a source vertex o to every other vertex v and the shortest path from every other vertex v to the source vertex o . This can be realized in time $O(en)$ by using a shortest path tree algorithm with negative cycle detection (cf. (Tarjan 1983, Section 7.2) and (Korte and Vygen 2012, Section 7.1)).

4 Solving the MaSTP with Arc-Consistency

In this section we extend ACSTP to a distributed algorithm DisACSTP to solve multiagent simple temporal networks (MaSTNs).

Definition 6. (Boerkoel and Durfee 2013) A *multiagent simple temporal network (MaSTN)* is a tuple $\mathcal{M} = \langle \mathcal{P}, C^X \rangle$, where

- $\mathcal{P} = \{\mathcal{N}_i \mid i = 1, \dots, p\}$ is a set of *local* STNs, where each $\mathcal{N}_i = \langle V_i, D_i, C_i \rangle$ is an STN belonging to agent i and we require that $V_i \cap V_j = \emptyset$ for any two different agents $i, j = 1, \dots, p$.
- C^X is a set of *external* constraints, where each constraint is over two variables belonging to two different agents.

Constraint graphs for MaSTNs can be defined analogously as that for STNs, where we use E^X for the set of edges corresponding to constraints in C^X . See Figure 1 for an illustration. In Figure 1, the edges in E^X are represented as red lines.

Definition 7. Suppose $\mathcal{M} = \langle \mathcal{P}, C^X \rangle$ is an MaSTN. Let $I_{vw} \in C^X$ with $v \in V_i, w \in V_j$ be an external constraint. We say that I_{vw} is an external constraint of agent i , and write C_i^X for the set of external constraints of agent i . We call v and w a *shared* and an *external* variable of agent i , respectively. We write V_i^X for the set of external variables of agent i . In Figure 1, the vertices for shared variables are represented as red circles.

DisACSTP is presented in Algorithm 2. In DisACSTP each agent i gets as input its portion \mathcal{N}_i of the input MaSTN \mathcal{M} and the set C_i^X of its external constraints, and runs its own algorithm. Similar to ACSTP, DisACSTP updates the domains of \mathcal{N}_i at each iteration of the for-loop and maintains a queue Q_i to record the information about the unchanged domains. When a domain becomes empty during the updates, then the agent can terminate the algorithm and conclude that the input MaSTN \mathcal{M} is inconsistent. There are however aspects in DisACSTP that are different from ACSTP, which stem from the fact that in MaSTP an agent cannot have the global knowledge of the states of other agents' processes without sharing certain information with other agents. These aspects are the following:

1. The total number n of the variables in the input MaSTN is initially not known to individual agents. This, however, can easily be determined using an echo algorithm (Chang 1982). We can therefore regard n as given as an input to DisACSTP.
2. As the agents may run their processes at different paces, at each iteration of the for-loop (lines 2–34), they synchronize the domains of their external variables (lines 3–4). Otherwise, some agents might use stale external domains and make wrong conclusions.
3. When a domain becomes empty while running DisACSTP, an agent broadcasts (lines 9–11) this information to other agents so that they can terminate their algorithms as soon as possible.

Algorithm 2: DisACSTP

Input : \mathcal{N}_i : agent i 's portion of MaSTN \mathcal{M} ;
 V_i^X : the set of agent i 's external variables;
 C_i^X : the set of agent i 's external constraints;
 $parent(i)$: the parent of agent i w.r.t. $T(\mathcal{M})$;
 $children(i)$: the children of agent i w.r.t. $T(\mathcal{M})$;
 $T(\mathcal{M})$;
 n : the number of variables of \mathcal{M} .

Output: Agent i 's portion of the AC-closure of \mathcal{M} or “inconsistent”.

```

1  $Q_i \leftarrow \emptyset$ 
2 for  $k \leftarrow 1$  to  $n$  do
3   Send the domains of the shared variables to the
   neighbors.
4   Receive the domains of the external variables
   from the neighbors.
5   foreach  $v \in V_i$  do
6      $I'_v \leftarrow I_v$ 
7     foreach  $w \in V_i \cup V_i^X$  s.t.  $\{v, w\} \in E_i \cup E_i^X$ 
8       do
9          $I_v \leftarrow I_v \cap I_w \otimes I_{vw}$ 
10        if  $I_v = \emptyset$  then
11          Broadcast “inconsistent”.
12          return “inconsistent”
13        if  $I'_v = I_v$  then  $Q_i \leftarrow Q_i \cup \{v\}$ 
14        else  $Q_i \leftarrow Q_i \setminus \{v\}$ 
15  if  $\#Q_i = \#v$  then
16    if  $root(i)$  then
17      Send inquiry (“Are all  $Q_i$  full?”,  $k$ ) to
18       $children(i)$ 
19    while true do
20       $m \leftarrow \text{RECEIVEMESSAGE}()$ 
21      if  $m$  is domains of external variables
22      from a neighbor then
23        break
24      if  $m$  is inquiry (“Are all  $Q_i$  full?”,  $k$ ) then
25        if  $leaf(i)$  then
26          Send feedback (“yes”,  $k$ ) to
27           $parent(i)$ 
28        else Send  $m$  to  $children(i)$ 
29      if  $m$  is feedback (“yes”,  $k$ ) then
30        if all feedbacks received from
31         $children(i)$  then
32          if  $root(i)$  then
33            Broadcast “arc-consistent”
34            return  $\mathcal{N}_i$ 
35          else Send  $m$  to  $parent(i)$ 
36      if  $m$  is “arc-consistent” then
37        return  $\mathcal{N}_i$ 
38      if  $m$  is “inconsistent” then
39        return “inconsistent”
40  return “inconsistent”

```

4. If the queue Q_i of an agent i is full (i.e., it contains all of the agent’s variables in V_i) after an iteration of the for-loop, then the agent shares this information with all other agents in \mathcal{M} so as to jointly determine whether the queues of all agents are full and the network is arc-consistent (lines 15–16 and 21–30).
5. If the queue Q of an agent is *not* full after an iteration of the for-loop, then the agent broadcasts this information to all other agents, so that they can move to the next iteration of the for-loop as soon as possible.

All the preceding aspects are subject to communication of certain information between agents. DisACSTP coordinates this communication while (i) preserving the privacy of each agent and (ii) reducing the duration of any idle state of an individual agent. Concretely:

- Each agent shares information only with the agents who are connected through an external constraint. We call them the *neighbors* of the agent. This neighborhood-relationship among the agents induces a graph that we call henceforth the *agent graph*.
 - Each agent shares with its neighbors *only* the domains of its shared variables. No other information is shared (such as its network structure, constraints, private variables and their domains) and only the neighbors w.r.t. the agent graph can share the information. This property is a critical advantage over D Δ PPC (Boerkoel and Durfee 2013), as D Δ PPC often creates new external constraints during the process and reveal more private information of the agents than necessary.
 - Each agent uses a *broadcasting* mechanism to share global properties of the input MaSTN, i.e., an agent first sends a message (e.g., “inconsistent”) to its neighbors, then the neighbors forward the message to their neighbors and so on, until all agents receive the message. To reduce the number of messages, duplicates are ignored by the agents.
- An agent i broadcasts the following messages: “arc-consistent”, “inconsistent” and “ Q_i is not full”, where the last message is indirectly broadcasted by agent i skipping lines 14–34 and moving to the next iteration of the for-loop and then sending its shared domains to its neighbors. This initiates a chain reaction among the idle neighbors of agent i who have not moved to the next iteration yet, as they quit the idle states (lines 19–20) and move to the next iteration of the for-loop and then send also their shared domains to their idle neighbors (lines 3–4).
- There is a dedicated agent who checks at each iteration of its for-loop (given its queue is full) whether the queues of all other agents are full at the same iteration. This dedicated agent is determined by building a minimal spanning tree (e.g., by using an echo algorithm (Chang 1982)) $T(\mathcal{M})$ of the agent graph. The agent who is the root (henceforth the *root agent*) of this tree becomes then the dedicated agent.

The root agent sends an inquiry to its children to check whether the queues of all its descendants are full

(lines 15–16). The inquiry is then successively forwarded by the descendants whose queues are full. We have to distinguish here between two cases:

- (1) If all descendants’ queues are *full*, then the inquiry reaches all the leaf agents and returns back as feedbacks (lines 22–23) until the root agent receives all the feedbacks (lines 25–30) and broadcasts “arc-consistency”.
- (2) If a descendant’s queue is *not full*, then the descendant moves on to the next iteration of the for-loop and initiates a chain reaction among other agents by sending the domains of its shared variables to its neighbors (cf. the second paragraph of the third bullet point).

Due to the properties so far considered, DisACSTP is guaranteed to simulate the behavior of ACSTP while allowing concurrent domain update operations.

Theorem 13. *Let $\mathcal{M} = \langle \mathcal{P}, C^X \rangle$ be an MaSTN. Let \mathcal{N}_{\max} be a network with $e_{\max} = \max\{e_i + e_i^X \mid 1 \leq i \leq p\}$, where e_i and e_i^X are the number of edges of the constraint graph of \mathcal{N}_i and the number of external constraints of agent i , respectively. Then Algorithm 2 enforces AC on \mathcal{M} in time $O(e_{\max}n)$.*

5 Evaluation

In this section we experimentally compare our algorithms against the state-of-the-art algorithms for solving STNs. For centralized algorithms, we compare our ACSTP algorithm against Planken, de Weerd, and van der Krogt’s P³C algorithm (2008); for distributed algorithms, we compare our DisACSTP algorithm against Boerkoel and Durfee’s D Δ PPC algorithm (2013). All experiments for distributed algorithms used an asynchronous simulator in which agents are simulated by processes which communicate only through message passing and default communication latency is assumed to be zero. Our experiments were implemented in Python 3.6 and carried out on a computer with an Intel Core i5 processor with a 2.9 GHz frequency per CPU, 8 GB memory ².

As measures for comparing performances we use the number of constraint checks and the number of non-concurrent constraint checks (NCCCs) performed by the centralized algorithms and the distributed algorithms, respectively. Given an STN $\mathcal{N} = \langle V, D, C \rangle$, a constraint check is performed when we compute relation $r \leftarrow I_{vw} \cap (I_{vu} \otimes I_{uw})$ and check if $r = I_{vw}$ or $r \not\subseteq I_{vw}$.

ACSTP vs. P³C

Datasets We selected instances from the benchmark datasets of STNs used in (Planken, de Weerd, and van der Krogt 2012) for evaluations. We considered the scale-free graphs (**Scale-free-1**) with 1000 vertices and density parameter varying from 2 to 50. We also considered the scale-free graphs (**Scale-free-2**) with varying vertex count. The scale-free density parameter for this selection is 5. Beside these artificially constructed graphs, we also considered

²The source code for our evaluation can be found in <https://github.com/sharingcodes/MaSTN>

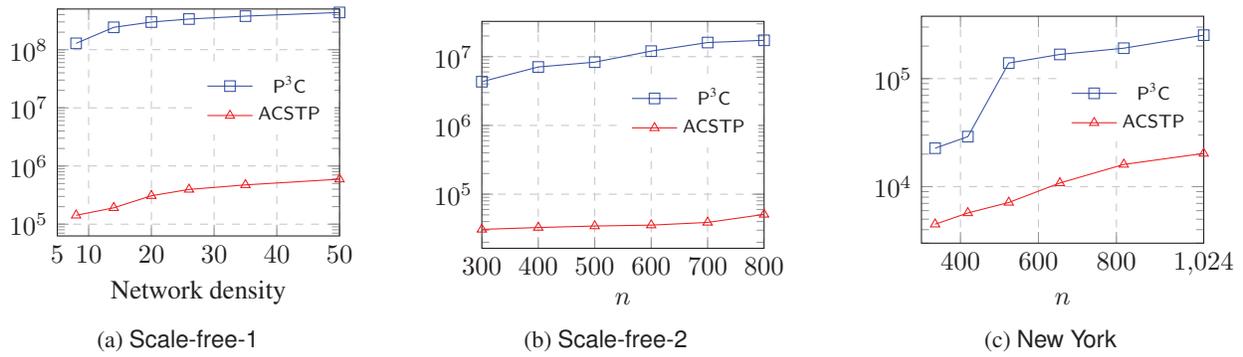


Figure 2: Evaluation of ACSTP and P³C. The y -axes (on the log scale) represent the number constraint checks.

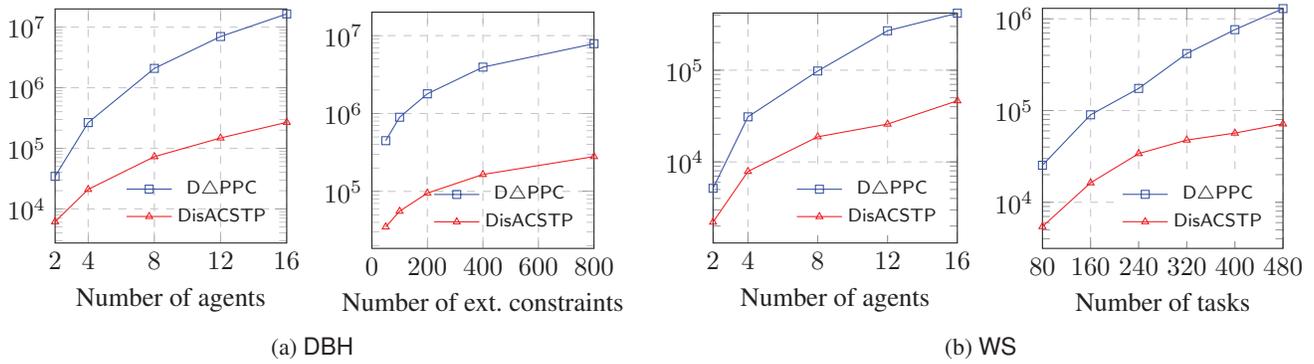


Figure 3: Evaluation of DisACSTP and D Δ PPC. The y -axis (on the log scale) represent the number of NCCCs.

graphs that are based on the road network of New York City (New York). This dataset contains 170 graphs on 108–3906 vertices, 113–6422 edges.

Results The results are presented in Figure 2, where base-10 log scales are used for the y -axes. For the scale-free graphs we observe that ACSTP is 100–1000 times faster than P³C. The dataset New York only contains very sparse networks (each network’s density is less than 1%), thus both algorithms could easily solve these networks. However, we still observe that ACSTP is about 5–12 times faster than P³C.

DisACSTP vs. D Δ PPC

Datasets We selected instances from the benchmark datasets of MaSTNs used in (Boerkoel and Durfee 2013) for evaluations. The first dataset BDH was randomly generated using the multiagent adaptation of Hunsberger’s (2002) random STN generator. Each MaSTN has N agents each with start time points and end time points for 10 activities, which are subject to various local constraints. In addition, each MaSTN has X external constraints. We evaluated the algorithms by varying the number of agents ($N \in \{2, 4, 8, 12, 16\}$, $X = 50 \times (N - 1)$) and the total number of external constraints ($N = 16$, $X \in \{100, 200, 400, 800\}$).

The second dataset WS is derived from a multiagent factory scheduling domain (Wilcox and Shah 2012), where N agents are working together to complete T tasks in a man-

ufacturing environment. We evaluated algorithms by varying the number of agents ($N \in \{2, 4, 8, 12, 16\}$, $T = 20 \times N$) and the total number of tasks ($N = 16$, $T \in \{80, 160, 240, 320, 400, 480\}$).

Results The results are presented in Figure 3, where base-10 log scales are again used for the y -axes. For the DBH random networks (Figure 3a) we observe that DisACSTP is 5–30 times faster than D Δ PPC. For the WS scheduling networks (Figure 3b) DisACSTP is 2–10 times faster than D Δ PPC. For both datasets we observe that, with increasing x -values, the y -values (i.e., NCCCs) for DisACSTP grow slower than those for D Δ PPC.

6 Conclusion

In this paper we presented a novel AC-based approach for solving the STP and the MaSTP. We have shown that arc-consistency is sufficient for solving an STN. Considering that STNs are defined over infinite domains, this result is rather surprising. Our empirical evaluations showed that the AC-based algorithms are significantly more efficient than their PC-based counterparts. This is mainly due to the fact that PC-based algorithms add many redundant constraints in the process of triangulation. More importantly, since our AC-based approach does not impose new constraints between agents that are previously not directly connected, it respects as much privacy of these agents as possible. We should note here that even though our distributed algo-

rithm DisACSTP showed remarkable performance, it can be further fine-tuned by using different termination detection mechanisms (cf. (Mattern 1987) and (Raynal 2013, Ch. 14)).

It would be interesting to see how the result in this paper can be used for solving the general disjunctive temporal problems (Stergiou and Koubarakis 2000). Potential extensions of our paper also include adapting our AC algorithms to incremental algorithms for the STP (Planken, de Weerd, and Yorke-Smith 2010), dynamic situations (Morris, Muscettola, and Vidal 2001) and uncertainty (Venable and Yorke-Smith 2005).

Acknowledgments

We thank the anonymous reviewers, who pointed out the connections to max/min-closed constraints and the shortest-path problem. The work of SL was partially supported by NSFC (No. 11671244), and the work of JL was partially supported by the Alexander von Humboldt Foundation.

References

- Barták, R.; Morris, R. A.; and Venable, K. B. 2014. An Introduction to Constraint-Based Temporal Reasoning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 8(1):1–121.
- Baudot, B., and Deville, Y. 1997. Analysis of distributed arc-consistency algorithms. Technical Report 97-07, Université catholique de Louvain.
- Boerkoel, J. C., and Durfee, E. H. 2013. Distributed reasoning for multi-agent simple temporal problems. *Journal of Artificial Intelligence Research* 47:95–156.
- Chang, E. J. H. 1982. Echo Algorithms: Depth Parallel Operations on General Graphs. *IEEE Transactions on Software Engineering* SE-8(4):391–401.
- Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. *Artificial Intelligence* 49(1-3):61–95.
- Dechter, R. 2003. *Constraint processing*. San Francisco: Morgan Kaufmann Publishers.
- Hamadi, Y. 2002. Optimal distributed arc-consistency. *Constraints* 7(3-4):367–385.
- Hunsberger, L. 2002. Algorithms for a Temporal Decoupling Problem in Multi-agent Planning. In *Eighteenth National Conference on Artificial Intelligence*, 468–475. Menlo Park, CA, USA: AAAI Press.
- Jeavons, P. G., and Cooper, M. C. 1995. Tractable constraints on ordered domains. *Artificial Intelligence* 79(2):327–339.
- Korte, B., and Vygen, J. 2012. *Combinatorial Optimization*, volume 21 of *Algorithms and Combinatorics*. Springer Berlin Heidelberg.
- Mattern, F. 1987. Algorithms for distributed termination detection. *Distributed computing* 2(3):161–175.
- Montanari, U. 1974. Networks of constraints: Fundamental properties and applications to picture processing. *Information Sciences* 7:95–132.
- Morris, P.; Muscettola, N.; and Vidal, T. 2001. Dynamic Control of Plans with Temporal Uncertainty. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI’01*, 494–499. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Nguyen, T., and Deville, Y. 1998. A distributed arc-consistency algorithm. *Science of Computer Programming* 30(1-2):227–250.
- Planken, L.; de Weerd, M.; and van der Krogt, R. 2008. P³C: A New Algorithm for the Simple Temporal Problem. In *Proceedings of the Eighteenth International Conference on International Conference on Automated Planning and Scheduling, ICAPS’08*, 256–263. Sydney, Australia: AAAI Press.
- Planken, L. R.; de Weerd, M. M.; and van der Krogt, R. P. 2012. Computing all-pairs shortest paths by leveraging low treewidth. *Journal of Artificial Intelligence Research* 43:353–388.
- Planken, L.; de Weerd, M.; and Yorke-Smith, N. 2010. Incrementally Solving STNs by Enforcing Partial Path Consistency. In *Proceedings of the Twentieth International Conference on International Conference on Automated Planning and Scheduling, ICAPS’10*, 129–136. Toronto, Ontario, Canada: AAAI Press.
- Raynal, M. 2013. *Distributed Algorithms for Message-Passing Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Shostak, R. 1981. Deciding linear inequalities by computing loop residues. *J. ACM* 28(4):769–779.
- Stergiou, K., and Koubarakis, M. 2000. Backtracking algorithms for disjunctions of temporal constraints. *Artificial Intelligence* 120(1):81–117.
- Tarjan, R. E. 1983. *Data Structures and Network Algorithms*. Number 44 in CBMS-NSF Regional Conference Series in Applied Mathematics. Philadelphia, Pa: Society for Industrial and Applied Mathematics.
- Venable, K. B., and Yorke-Smith, N. 2005. Disjunctive Temporal Planning with Uncertainty. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence, IJCAI’05*, 1721–1722. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Wilcox, R., and Shah, J. 2012. Optimization of Multi-Agent Workflow for Human-Robot Collaboration in Assembly Manufacturing. In *Infotech@Aerospace 2012*. American Institute of Aeronautics and Astronautics.