Verifying Networks of Timed Processes

Parosh Aziz Abdulla Bengt Jonsson
Dept. of Computer Systems
P.O. Box 325
S-751 05 Uppsala, Sweden
{parosh,bengt}@docs.uu.se

Abstract

Over the last years there has been an increasing research effort directed towards the automatic verification of infinite state systems, such as timed automata, hybrid automata, data-independent systems, relational automata, Petri nets, and lossy channel systems. We present a method for deciding reachability properties of networks of timed processes. Such a network consists of an arbitrary set of identical timed automata, each with a single real-valued clock. Using a standard reduction from safety properties to reachability properties, we can use our algorithm to decide general safety properties of timed networks. To our knowledge, this is the first decidability result concerning verification of systems that are infinite-state in "two dimentions": they contain an arbitrary set of (identical) processes, and they use infinite data-structures, viz. real-valued clocks. We illustrate our method by showing how it can be used to automatically verify Fischer's protocol, a timer-based protocol for enforcing mutual exclusion among an arbitrary number of processes.

1 Introduction

The last decade has seen much progress with regard to automated verification of reactive programs. The most dramatic advances have been obtained for finite-state programs. However, methods and algorithms are now emerging for the automatic verification of infinite state programs. There are at least two ways in which a program can be infinite-state. A program can be infinite-state because it operates on data structures from a potentially infinite domain, e.g., integers, stacks, queues, etc. Nontrivial verification algorithms have been developed for several classes of such systems, notably timed automata [ACD90, AH89, Č92a], hybrid automata [Hen95], data-independent systems [JP93, Wol86], relational automata ([BBK77, Č92b, Č94]), Petri nets ([Jan90, JM95]), pushdown processes ([BS95]) and lossy channel systems [AJ96b, AK95]. A program can also be infinite-state because it is intended to run on a network with an arbitrary number of nodes, i.e., the the program is parameterized with respect to the topology of the network of nodes. In this case, one would like to verify correctness for any number of components and any interconnection topology. Verification algorithms have been developed for systems consisting of an unbounded number of similar or identical finite-state processes ([GS92]), and (using a manually supplied induction hypothesis) for more general classes of parameterized systems [CGJ95, KM89, WL89].

In this paper, we will present an algorithm for verifying safety properties of a class of programs, which we call *timed networks*. A timed network is a system consisting of an arbitrary set of processes, each of which is a finite-state system operating on a real-valued clock. Each process could roughly be considered as a timed automaton [ACD90] with a single clock. In addition, our model also allows a central finite-state process, called a *controller*. Timed networks embody both of the two reasons for being infinite-state: they use an infinite data structure (namely clocks which can assume values from the set of real numbers), and they are parameterized in allowing an arbitrary set of processes. To our knowledge, this is the first decidability result concerning verification of networks of infinite-state processes.

We present an algorithm for deciding reachability properties of timed networks. Using a standard reduction (described e.g., in [VW86, GW93]) from safety properties to reachability properties, we can use this algorithm to decide general safety properties of timed networks. To decide reachability, we adapt a standard symbolic verification algorithm which has been used e.g., in model-checking [QS82, CES86] and assertional verification [Sha93a]. A rough description of this method is that in order to check whether a state in some set F is reachable, we compute the set of all states from which a state in F is reachable. This computation is performed using a standard fixedpoint iteration, where for successively larger j we compute the set of states from which a state in F can be reached by a sequence of transitions of length less than or equal to j. More precisely, we obtain the $(j+1)$st approximation from the jth approximation by adding the *pre-image* of the jth approximation, i.e., the set of states from which a state in the jth approximation can be reached by some transition. If this procedure converges, one checks whether the result intersects the set of initial states of the model. The heart of our result is solving the following three problems:

- finding a suitable representation of infinite sets of states,
- finding a method for computing pre-images, and
- proving that the iteration always converges.

To represent sets of states, we use constraints which generalize the notion of regions used to verify properties of (non-parameterized) timed automata [ACD90]. A constraint represents conditions on a potentially unbounded number of processes and their clocks. In contrast to the situation for timed automata [ACD90], where for each program there are finitely many regions, there is in general an infinite number of constraints that can appear in the analysis of a given timed network. To handle this, we introduce an entailment ordering on contraints. The key step in our proof of decidability consists in proving that this relation is a well quasi-ordering, implying that the the above mentioned fixedpoint iteration converges.

Our results also demonstrate the strength and applicability of the general framework described in [AvJYK96, Fin90]. Using that framework, we can conclude the decidability of eventuality properties (of the form AFp in CTL). for timed networks, and the question of whether or not a timed network simulates or is simulated by a finite-state system. We will not further consider these questions in this paper.

Our model of timed networks is related to other formalisms for timed systems, notably time or timed Petri nets [MF76, GMMP91, BD91] and Timed CCS [Yi91]. Our decidability result can be translated to decidability results for variants of these formalisms. It is known that reachability is undecidable for time Petri nets. This is due to the inclusion of *urgency* in the Petri net model. Urgency means that a transition is forced to execute within a specified timeout period. In our model transitions can not be forced to occur; a timeout can only specify that a transition is executed within a specified time period *if* it is executed. Urgency allows the model to test for emptiness of a place, thus leading to undecidability. A similar difference holds in comparison with Timed CCS.

As an illustration of our method, we model Fischer's protocol [SBK92], and show how an automatic verification algorithm would go about verifying mutual exclusion. Several tools for verifying automata with a fixed number of clocks have been used to verify the protocol for an increasing number of processes (e.g., [ACHH92]). Kristoffersen et al. [KLL+97] describes an experiment where the number of processes is 50. In [LSW95], a constraint-based proof methodology is used to perform a manual verification of the protocol.

Outline The rest of the paper is structured as follows. In the next section, we present our model of timed networks. An overview of the reachability algorithm is presented in Section 3. In Section 4, we present our constraint system. In Section 5, we present a procedure for calculating the pre-image of a set of states which are represented by a constraint. In Section 6 we prove that the entailment ordering on the constraint system is a well quasi-ordering, which implies that our algorithm always terminates. An application of the reachability algorithm to the verification of Fischer's protocol is given in Section 7. In Section 8 we give some notes on the implementation of the reachability algorithm.

2 Timed Networks

In this section, we will define networks of timed processes. Intuitively, a network of timed processes consists of a *controller* and an arbitrarily large set of identical (timed) *processes*. The controller is a finite-state transition system. Each process has a finite-state control part, and an unbounded data structure, namely a real-valued clock[1]. The values of the clocks of the processes are incremented continuously at the same rate. In addition to letting time pass by incrementing the clocks, the network can change its configuration according to a finite number of *rules*. Each rule describes a set of transitions in which the controller and an arbitrary but fixed-size set of processes synchronize and simultaneously change their states. A rule may be conditioned on the control states of the participating controller and processes, and on conditions on the clock values of the participating processes. If the conditions for a rule are satisfied, the controller and each participating process may change its state and (optionally) reset its clock to 0.

We are interested in verifying correctness of a network regardless of its size. The actual object of study will therefore be a *family* of networks, where the number of processes is not given. A family merely defines the controller and process states together with a set of rules. The parameter (i.e., size) of the network will be introduced later, when we define configurations.

We use $\mathcal{N}$ and $\mathcal{R}^{\geq 0}$ to denote the sets of natural numbers and nonnegative reals respectively. For $n \in \mathcal{N}$, we use $\widehat{n}$ to denote the set $\{1, \ldots, n\}$. A *guarded command* is of the form $p(x) \longrightarrow op$, where $p(x)$ is a boolean combination of predicates of the form $k < x$, $k \leq x$, $k > x$, or $k \geq x$ for $k \in \mathcal{N}$, and $op \in \{reset, skip\}$.

Definition 2.1 A *family of timed networks* (timed network for short) is a triple $\langle C, Q, R \rangle$, where:

C is a finite set of *controller* states.

Q is a finite set of *process* states.

R is a finite set of *rules*. A rule r is of the form

$$\langle \langle c, c' \rangle, \langle q_1, stmt_1, q_1' \rangle, \ldots, \langle q_n, stmt_n, q_n' \rangle \rangle$$

where $c, c' \in C$, $q_i, q_i' \in Q$, and $stmt_i$ is a guarded command. □

Intuitively, the set C represents the set of states of the controller. The set Q represents the set of states of each of the identical processes. A rule r describes a set of transitions of the network. The rule is enabled if the state of the controller is c, and if there are n processes with states

[1]The controller could also be equipped with a timer, but this aspect is not central to our result, so we will omit it.

$q_1, \ldots, q_n$, respectively, whose clock values satisfy the corresponding guards. The rule is executed by simultaneously changing the state of the controller to c', changing the states of the n processes to $q_1', \ldots, q_n'$ respectively, and modifying values of the clocks according to the relevant guarded commands.

Definition 2.2 A *configuration* γ of a timed network $\langle C, Q, R \rangle$ is a quadruple of form $\langle I, c, \mathtt{q}, \mathtt{x} \rangle$, where I is a finite *index set*, $c \in C$, $\mathtt{q}: I \rightarrow Q$, and $\mathtt{x}: I \rightarrow \mathcal{R}^{\geq 0}$. □

Intuitively, I is the set of indices of processes in the network. The index set does not change when performing transitions. Each element in I will be used as an index to represent one particular process in the network. Thus, we can say that a timed network defines a family of networks parametrized by I[2]. The state of the controller is given by c, the states of the processes are given by the mapping $\mathtt{q}$ from indices to process states, and the clock values are given by the mapping $\mathtt{x}$ from indices to nonnegative real numbers.

A timed network changes its configuration by performing transitions. We will define a transition relation $\longrightarrow$ as the union of a *discrete* transition relation $\longrightarrow_D$, representing transitions caused by the rules, and a *timed* transition relation $\longrightarrow_T$ which represents the passage of time. The discrete relation $\longrightarrow_D$ will furthermore be the union of transition relations $\stackrel{r}{\longrightarrow}_D$ corresponding to each rule r, i.e., $\longrightarrow_D = \bigcup_{r \in R} \stackrel{r}{\longrightarrow}_D$.

Definition 2.3 Let $r = \langle \langle c, c' \rangle, \langle q_1, stmt_1, q_1' \rangle, \ldots, \langle q_n, stmt_n, q_n' \rangle \rangle$ be a rule where $stmt_i$ is of form $p_i(x) \longrightarrow op_i$ for $i = 1, \ldots, n$. Consider two configurations $\gamma = \langle I, c, \mathtt{q}, \mathtt{x} \rangle$ and $\gamma' = \langle I, c', \mathtt{q}', \mathtt{x}' \rangle$, with the same index sets, and where the controller states of γ and γ' are the same as the controller states in the rule r. We use $\gamma \stackrel{r}{\longrightarrow}_D \gamma'$ to denote that there is an injection $h: \widehat{n} \rightarrow I$ from indices of the rule r to indices of the network such that

1. $\mathtt{q}(h(i)) = q_i$, and $p_i(\mathtt{x}(h(i)))$ holds for each $i \in \widehat{n}$,
2. $\mathtt{q}'(h(i)) = q_i'$ for $i \in \widehat{n}$,
3. $\mathtt{q}'(j) = \mathtt{q}(j)$ for $j \in (I \setminus range(h))$,
4. $\mathtt{x}'(h(i)) = 0$ for $i \in \widehat{n}$ with $op_i = reset$,
5. $\mathtt{x}'(h(i)) = \mathtt{x}(h(i))$ for $i \in \widehat{n}$ with $op_i = skip$, and
6. $\mathtt{x}'(j) = \mathtt{x}(j)$ for $j \in (I \setminus range(h))$. □

[2] We can extend our model to include dynamic creation and destruction of processes, by allowing the set of indices in a configuration to change dynamically. Our decidability result holds also for such an extension. However, we will not consider that in the present paper.

The first condition asserts that r is enabled, i.e., that the process states $q_1, \ldots, q_n$ are matched by the corresponding process states in the configuration γ and that the corresponding guarded commands are enabled. The second condition means that in the transition from γ to γ', the states of processes that are matched (by h) with indices of r are changed according to r. The third condition asserts that the states of the other processes are unchanged. The fourth condition asserts that in the transition from γ to γ', the clock values of processes that are matched (by h) with indices of r are set to 0 if the corresponding guarded command contains *reset*, the fifth asserts that clocks are unchanged if the guarded command contains *skip*. The last condition asserts that clock values of unmatched processes are unchanged.

Let $\gamma = \langle I, c, \mathtt{q}, \mathtt{x}\rangle$ be configuration. For $\delta \in \mathcal{R}^{\geq 0}$, we use $\gamma^{+\delta}$ to denote the configuration $\langle I, c, \mathtt{q}, \mathtt{x}'\rangle$, where $x'(j) = x(j) + \delta$ for each $j \in I$.We say that γ performs a *timed transition* to a configuration γ', denoted $\gamma \longrightarrow_T \gamma'$, if there is a $\delta \in \mathcal{R}^{\geq 0}$ such that $\gamma' = \gamma^{+\delta}$. We use $\gamma \longrightarrow \gamma'$ to denote that either $\gamma \longrightarrow_D \gamma'$ or $\gamma \longrightarrow_T \gamma'$. We use $\stackrel{*}{\longrightarrow}$ to denote the reflexive transitive closure of $\longrightarrow$.

3 Overview of the Reachability Algorithm

In this section we define the reachability problem, and give an overview of our method for solving it. Given a timed network $\langle C, Q, R\rangle$ together with states $c_{init} \in C$ and $q_{init} \in Q$ which we call *the initial controller state* and *the initial process state* respectively, we define an *initial configuration* γ_{init} of the timed network $\langle C, Q, R\rangle$ as a configuration of the form $\langle I, c_{init}, \mathtt{q}_{init}, \mathtt{x}_{init}\rangle$ where $\mathtt{q}_{init}(j) = q_{init}$ and $\mathtt{x}_{init}(j) = 0$ for each $j \in I$. Thus, there is one initial configuration for each possible index set I. We say that a configuration γ is *reachable* if $\gamma_{init} \stackrel{*}{\longrightarrow} \gamma$, for some initial configuration γ_{init}. We say that a set Γ of configurations is *reachable* if there is a reachable $\gamma \in \Gamma$.

We will present an algorithm for deciding whether a set Γ of configurations of a timed network is reachable. Note that in general, Γ will contain configurations of networks with infinitely many different sizes, where the size of a configuration is given by its index set. This means that we ask whether there is *some* size of the network such that a configuration with this size (as given by its index set) is reachable. In a typical situation, we are interested in verifying that Γ is an unreachable set of "bad" configurations, irrespective of the size of the network. If we include in Γ the bad configurations of all possible network sizes, and if our analysis finds Γ to be unreachable, this means that we have verified that the configurations in Γ are unreachable for all possible sizes of the network. For instance, we can verify correctness of an n-process mutual exclusion algorithm for all values of n simultaneously.

In Section 4, we will define a class of constraints for representing sets of configurations. A constraint ϕ denotes a (possibly infinite) set $[\![\phi]\!]$ of configurations. A finite set $\Phi = \{\phi_1, \ldots, \phi_n\}$ of constraints denotes the union of the denotations of its elements, i.e., $[\![\Phi]\!] = \bigcup_{i=1}^{n} [\![\phi_i]\!]$. Formally, the reachability problem is defined as follows.

Instance: A timed network $\langle C, Q, R \rangle$, an initial controller state c_{init}, an initial process state q_{init} and a finite set Φ of constraints.

Question: Does $[\![\Phi]\!]$ contain a reachable configuration?

To check the reachability of Φ we perform a reachability analysis backwards. Let $pre(\phi)$ denote the set $\{\gamma \ : \ \exists \gamma' \in [\![\phi]\!] : \gamma \longrightarrow \gamma'\}$, and $pre(\Phi)$ denote the set $\{\gamma \ : \ \exists \gamma' \in [\![\Phi]\!] : \gamma \longrightarrow \gamma'\}$. Note that $pre(\Phi)$ is equivalent to $\cup_{\phi \in \Phi} pre(\phi)$. Starting from Φ we define the sequence $\Phi_0, \Phi_1, \Phi_2, \ldots$ of finite sets of constraints by $\Phi_0 = \Phi$ and $\Phi_{j+1} = \Phi_j \cup pre(\Phi_j)$. Intuitively, Φ_j denotes the set of configurations from which Φ is reachable by a sequence of at most j transitions. Note that the sequence is increasing i.e., that $[\![\Phi_0]\!] \subseteq [\![\Phi_1]\!] \subseteq [\![\Phi_2]\!] \subseteq \cdots$. In the next paragraph, we will prove that the iteration converges (using Theorem 6.4), i.e., that there is a k such that $[\![\Phi_k]\!] = [\![\Phi_{k+1}]\!]$, implying that $[\![\Phi_k]\!] = [\![\Phi_j]\!]$ for all $j \geq k$. It follows that Φ is reachable if and only if there is an initial configuration γ_{init} such that $\gamma_{init} \in [\![\Phi_k]\!]$, which is easily checked since Φ_k is a *finite* set of constraints.

To prove convergence, we introduce, in Definition 4.3, a quasi-order $\preceq$ on constraints by defining $\phi \preceq \phi'$ to denote that $[\![\phi']\!] \subseteq [\![\phi]\!]$. In Theorem 6.4, we will show that $\preceq$ is a well quasi-ordering on the set of constraints, i.e., that in any infinite sequence $\phi_0 \ \phi_1 \ \phi_2 \ \cdots$ of constraints, there are indices $i < j$ such that $\phi_i \preceq \phi_j$. This implies that any increasing sequence $[\![\Phi_0]\!] \subseteq [\![\Phi_1]\!] \subseteq [\![\Phi_2]\!] \subseteq \cdots$ of finite sets of constraints will converge, since otherwise we could extract an infinite sequence $\phi_0 \ \phi_1 \ \phi_2 \ \cdots$ of constraints (where ϕ_i is chosen such that $\phi_i \in \Phi_i$ but $[\![\phi_i]\!] \not\subseteq [\![\Phi_{i-1}]\!]$) for which there are no indices $i < j$ such that $\phi_i \preceq \phi_j$.

Summarizing, we have established the following theorem

Theorem 3.1 The reachability problem for families of timed networks is decidable

Proof. Follows from the preceding discussion, using Theorem 4.4 (decidability of $\preceq$), Theorem 5.8 (computability of *pre*) and Theorem 6.4 (well quasi-orderedness of $\preceq$). $\Box$

The following sections contain the above mentioned definitions and lemmas. In Section 4, we define the constraint system. In Section 5, we show that $pre(\phi)$ can be computed and represented by a finite set of constraints whenever ϕ is a contraint. Finally, in Section 6, we show that the relation $\preceq$ is a well quasi-ordering on the set of constraints.

4 A Constraint System for Timed Networks

In this section we introduce a constraint system for timed networks. Our constraint system generalizes the notion of regions, employed for the analysis of timed automata [ACD90]. We use a representation of constraints, which is similar to a representation of regions used by Godskesen [God94].

For a quasi-order $\sqsubseteq$[3] on some set, we use $a_1 \equiv a_2$ to denote that $a_1 \sqsubseteq a_2$ and $a_2 \sqsubseteq a_1$, and use $a_1 \sqsubset a_2$ to denote that $a_1 \sqsubseteq a_2$ and $a_2 \not\sqsubseteq a_1$. For a real number $x \in \mathcal{R}^{\geq 0}$, let $\lfloor x \rfloor$ denote its integer part, and let *fract*(x) denote its fractional part.

Definition 4.1 Let $\langle C, Q, R \rangle$ be a family of timed networks. Let *max* be the maximum constant occurring in the guarded commands in R. A *constraint* ϕ of $\langle C, Q, R \rangle$ is a tuple $\langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq \rangle$ where

- $c \in C$ is a controller state,
- m is a natural number, where $\widehat{m}$ intuitively denotes a set of indices of processes constrained by ϕ,
- $\mathbf{q} : \widehat{m} \mapsto Q$ is a mapping from indices to process states,
- $\mathbf{k} : \widehat{m} \mapsto \{0, \ldots, max\}$ maps each index to a natural number not greater than max,
- $\sqsubseteq$ is a quasi-order on the set $\widehat{m} \cup \{\bot, \top\}$ which satisfies
  - the elements $\bot$ and $\top$ are minimal and maximal elements of $\sqsubseteq$, respectively, with $\bot \sqsubset \top$ [4],
  - $j \equiv \bot$ or $j \equiv \top$ whenever $\mathbf{k}(j) = \max$, for $j \in \widehat{m}$, and
  - $\mathbf{k}(j) = \max$ whenever $j \equiv \top$, for $j \in \widehat{m}$. □

Intuitively, a constraint denotes a set of configurations of networks in the family. The constraint $\langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq \rangle$ represents the set of configurations with controller state c in which each index $j \in \widehat{m}$ represents a process which has control state $\mathbf{q}(j)$, for which $\mathbf{k}(j)$ is either *max* or the integer part of its clock, whichever is least, for which $j \equiv \bot$ iff the integer part of the clock is at most *max* and the fractional part of the clock is 0, and for which $j \equiv \top$ iff the clock value is more than *max*. Furthermore, the fractional parts of the clocks corresponding to indices j with $j \sqsubset \top$ are ordered exactly according to $\sqsubseteq$. This implies, among other things, that for clock values that are larger than *max*, a constraint gives no information about the difference between the actual clock value and *max*. The meaning of constraints is made formal in the following definition.

[3] A quasi-order is a reflexive and transitive relation.

[4] Note that $\bot, \top \notin \widehat{m}$.

Definition 4.2 Let $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq \rangle$ be a constraint and let $\gamma = \langle I, c, \mathtt{q}, \mathtt{x} \rangle$ be a configuration[5]. We define $\gamma \in [\![\phi]\!]$ to mean that there is an injection $h : \widehat{m} \mapsto I$ from the indices of ϕ to the indices of γ such that for all $j, j_1, j_2 \in \widehat{m}$

- $\mathtt{q}(h(j)) = \mathbf{q}(j)$,
- $\min(max, \lfloor \mathtt{x}(h(j)) \rfloor) = \mathbf{k}(j)$,
- $j \equiv \bot$ if and only if $\mathtt{x}(h(j)) \leq max$ and $fract(\mathtt{x}(h(j))) = 0$,
- $j \equiv \top$ if and only if $\mathtt{x}(h(j)) > max$, and
- if $j_1, j_2 \not\equiv \top$ then $fract(\mathtt{x}(h(j_1))) \leq fract(\mathtt{x}(h(j_2)))$ if and only if $j_1 \sqsubseteq j_2$. □

Note that a constraint ϕ defines conditions on states and clock values which should be satisfied by *some* set of processes (those represented by indices in $range(h)$) in the configuration γ in order for γ to be included in $[\![\phi]\!]$. The constraint puts no requirements on processes whose indices are outside $range(h)$.

Definition 4.3 Define the ordering $\preceq$ on constraints by $\phi \preceq \phi' \stackrel{\text{def}}{=} [\![\phi']\!] \subseteq [\![\phi]\!]$. □

Intuitively, $\phi \preceq \phi'$ means that ϕ' is "stronger" than ϕ, or that ϕ' "entails" ϕ. The following theorem shows how to compute $\preceq$.

Theorem 4.4 Let $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq \rangle$ and $\phi' = \langle c', m', \mathbf{q}', \mathbf{k}', \sqsubseteq' \rangle$ be constraints. We have $\phi \preceq \phi'$ if and only if there is an injection $g : \widehat{m} \mapsto \widehat{m'}$ such that

- $c = c'$,
- for all $j \in \widehat{m}$ we have
  - $\mathbf{q}'(g(j)) = \mathbf{q}(j)$,
  - $\mathbf{k}'(g(j)) = \mathbf{k}(j)$,
  - $g(j) \equiv' \bot$ iff $j \equiv \bot$,
  - $g(j) \equiv' \top$ iff $j \equiv \top$,
- if $j_1, j_2 \in \widehat{m}$ then $g(j_1) \sqsubseteq' g(j_2)$ if and only if $j_1 \sqsubseteq j_2$.

Proof. The proof can be found in the appendix. □

[5] Observe that the controller states are the same in ϕ and γ.

5 Computing *pre*

In this section we show, for a given constraint ϕ, how to compute $pre(\phi)$, defined as $\{\gamma \ : \ \exists\gamma' \in [\![\phi]\!] : \gamma \longrightarrow \gamma'\}$. Since the transition relation is the union of a discrete and a timed transition relation, we will compute $pre(\phi)$ as $pre_D(\phi) \cup pre_T(\phi)$, where $pre_D(\phi)$ is the set $\{\gamma \ : \ \exists\gamma' \in [\![\phi]\!] : \gamma \longrightarrow_D \gamma'\}$, and where $pre_T(\phi)$ is the set $\{\gamma \ : \ \exists\gamma' \in [\![\phi]\!] : \gamma \longrightarrow_T \gamma'\}$.

5.1 Computing pre_D

First a preliminary definition.

Definition 5.1 Let $p(x)$ be a guard and let $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq\rangle$ be a constraint. For $j \in \widehat{m}$, we define the relation $\langle\phi, j\rangle \models p(x)$, meaning that *$p$ is satisfied at index j in ϕ,* as follows:

- If $p(x)$ is of form $k \leq x$ for some $k \in \{0, \ldots, max\}$ then $\langle\phi, j\rangle \models p(x)$ iff $\mathbf{k}(j) \geq k$,
- If $p(x)$ is of form $k < x$ for some $k \in \{0, \ldots, max\}$ then $\langle\phi, j\rangle \models p(x)$ iff either $\mathbf{k}(j) > k$ or both $\mathbf{k}(j) = k$ and $\perp \sqsubset j$,
- If $p(x)$ is of form $k \geq x$ for some $k \in \{0, \ldots, max\}$ then $\langle\phi, j\rangle \models p(x)$ iff either $\mathbf{k}(j) < k$ or both $\mathbf{k}(j) = k$ and $j \equiv \perp$,
- If $p(x)$ is of form $k > x$ for some $k \in \{0, \ldots, max\}$ then $\langle\phi, j\rangle \models p(x)$ iff $\mathbf{k}(j) < k$.
- If $p(x)$ is of form $p_1(x) \wedge p_2(x)$ then $\langle\phi, j\rangle \models p(x)$ iff $\langle\phi, j\rangle \models p_1(x)$ and $\langle\phi, j\rangle \models p_2(x)$.
- Disjunction is treated analogously as conjunction. Negations are handled by pushing them inwards in the standard way before applying the above definitions. □

Lemma 5.2 Let $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq\rangle$ be a constraint and let $\gamma = \langle I, c, \mathtt{q}, \mathtt{x}\rangle$ be a configuration, such that $\gamma \in [\![\phi]\!]$. Let $h : \widehat{m} \mapsto I$ be an injection which satisfies the 5 conditions in Definition 4.2, i.e., h is an injection which shows why $\gamma \in [\![\phi]\!]$. Then

$$\langle\phi, j\rangle \models p(x) \qquad \text{iff} \qquad p(\mathtt{x}(h(j))) \text{ holds}$$

Proof. The proof can be found in the appendix. □

We will compute $pre_D(\phi')$ as $\cup_{r\in R} pre(r, \phi')$, where $pre(r, \phi')$ denotes the set $\{\gamma \ : \ \exists\gamma' \in [\![\phi']\!] : \gamma \stackrel{r}{\longrightarrow}_D \gamma'\}$ of configurations from which ϕ' is reachable through a single application of r[6].

[6] In order to be consistent with the notation in Section 2, we use the primed version of the constraint to refer to the constraint after a transition, and an unprimed version of the constraint to refer to the constraint before a transition.

Let $r = \langle\langle c, c'\rangle, \langle q_1, p_1(x) \longrightarrow op_1, q_1'\rangle, \ldots, \langle q_n, p_n(x) \longrightarrow op_n, q_n'\rangle\rangle$ and let $\phi' = \langle c', m', \mathbf{q}', \mathbf{k}', \sqsubseteq'\rangle$. We will compute a representation of $pre(r, \phi')$ as a finite set of constraints. Each constraint ϕ with $[\![\phi]\!] \subseteq pre(r, \phi')$ will be obtained from a particular way of matching indices of ϕ' with indices of r. Each such matching gives rise to a set of constraints ϕ with $[\![\phi]\!] \subseteq pre(r, \phi')$, namely those constraints that are consistent both with the conditions imposed by ϕ' according to Definition 4.2, and with the conditions imposed by r according to Definition 2.3.

A matching will be represented by a *partial* injection g' from $\widehat{m'}$ to $\widehat{n}$: each index $j \in domain(g')$ of ϕ' is matched with a unique index $g'(j)$ of r (note that $domain(g') \subseteq \widehat{m'}$). Indices in $(\widehat{m}' \setminus domain(g'))$ represent processes which are constrained by ϕ' but are not matched with any index of r. In the indices of ϕ, we must also include the $n - |range(g')|$ indices of r which are not matched with any index of ϕ'. Thus, let $m = m' + (n - |range(g')|)$ be the number of indices of ϕ. Define an *extension* g of g' to be a *surjective* partial injection $g : \widehat{m} \mapsto \widehat{n}$, with domain $domain(g) = domain(g') \cup (\widehat{m} \setminus \widehat{m'})$, such that $g(j) = g'(j)$ for each $j \in domain(g')$. [7] It follows from the definition of extension that $\widehat{m'} \setminus domain(g')$ is not in $domain(g)$ and that g in addition maps each $j \in (\widehat{m} \setminus \widehat{m'})$ to a unique $g(j) \in (\widehat{n} \setminus range(g'))$.

Lemma 5.3 If $\phi' = \langle c', m', \mathbf{q}', \mathbf{k}', \sqsubseteq'\rangle$ is a constraint and r is a rule, as above, then $pre(r, \phi')$ is the denotation of the set of constraints of form $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq\rangle$, for which there is an extension g of a partial injection g' from $\widehat{m'}$ to $\widehat{n}$, where $m = m' + (n - |range(g')|)$, which satisfies the following conditions[8].

1. $\mathbf{q}(j) = q_{g(j)}$ and $\langle \phi, j\rangle \models p_{g(j)}(x)$ for each $j \in domain(g)$,
2. $\mathbf{q}'(j) = q'_{g'(j)}$ for $j \in domain(g')$,
3. $\mathbf{q}'(j) = \mathbf{q}(j)$ for $j \in \widehat{m'} \setminus domain(g')$,
4. $\mathbf{k}'(j) = 0$ and $j \equiv' \perp$ if $j \in domain(g')$ and $op_{g'(j)} = reset$,
5. For all j such that either $j \in domain(g')$ and $op_{g'(j)} = skip$, or such that $j \in \widehat{m'} \setminus domain(g')$,
   we have $\mathbf{k}'(j) = \mathbf{k}(j)$, and $j \equiv' \perp$ iff $j \equiv \perp$, and $j \equiv' \top$ iff $j \equiv \top$.
6. For each j_1 and j_2 such that for $i = 1, 2$ either $j_i \in domain(g')$ and $op_{g'(j_i)} = skip$, or $j_i \in \widehat{m'} \setminus domain(g')$,
   we have $j_1 \sqsubseteq' j_2$ if and only if $j_1 \sqsubseteq j_2$. □

The above list of conditions captures the semantics of $\xrightarrow{r}_D$, given the correspondences between the indices of r, ϕ' and ϕ which are given by g and

[7] note that $\widehat{m} \setminus \widehat{m'} = \{m' + 1, m' + 2, \ldots, m\}$

[8] Note that we implicitly require the controller states c and c' of r to coincide with the controller states c and c' of ϕ and ϕ', respectively.

g'. Note the close correspondence between the conditions of the lemma and the conditions of transitions in Definition 2.3. The conditions on controller states are implicitly included by our notation, which requires that the controller states of ϕ and ϕ' be the controller states of r. Condition 1 state that r must be enabled in a configuration satisfying ϕ. Conditions 2 and 3 capture the conditions on states of the processes: after a transition, states of processes with indices in $domain(g')$ are constrained by 2; and processes with indices in $\widehat{m'} \backslash domain(g')$ are unaffected by the rule (condition 3). Condition 4 describes the effect of a *reset* statement: the clock value becomes 0 in ϕ'. Finally, conditions 5 and 6 assert that for indices that correspond to a *skip* statement, or for indices not matched by r (and hence unaffected by the transition), the clock values are unchanged by a transition.

Proof. The proof can be found in the appendix. □

5.2 Computing pre_T

First, we define a relation pre_t which we later use (Lemma 5.7) to compute pre_T.

Definition 5.4 For a constarint $\phi' = \langle c', m', \mathbf{q}', \mathbf{k}', \sqsubseteq' \rangle$ we define $pre_t(\phi')$ to be the denotation of the set of constraints of form $\phi = \langle c', m', \mathbf{q}', \mathbf{k}, \sqsubseteq \rangle$ satisfying either of the following two conditions.

1. for some $j \in \widehat{m}$ we have $j \equiv' \perp$, there is no $j \in \widehat{m}$ such that $j \equiv' \perp$ and $\mathbf{k}'(j) = 0$, and the following three conditions hold.
   - $\mathbf{k}(j) = \mathbf{k}'(j) - 1$ if $j \equiv' \perp$,
   - $\mathbf{k}(j) = \mathbf{k}'(j)$ if $j \not\equiv' \perp$,
   - $j_1 \sqsubseteq j_2$ if and only if either
     (a) $j_2 \equiv' \top$, or
     (b) $j_2 \equiv' \perp$ and $j_1 \not\equiv' \top$, or
     (c) $j_1 \sqsubseteq' j_2$ and $\perp \sqsubset' j_1, j_2 \sqsubset' \top$.
2. There is no $j \in \widehat{m}$ such that $j \equiv' \perp$ and the following four conditions hold:
   - $\mathbf{k} = \mathbf{k}'$,
   - whenever $\perp \sqsubset' j_1, j_2 \sqsubset' \top$ we have $j_1 \sqsubseteq j_2$ if and only if $j_1 \sqsubseteq' j_2$,
   - whenever $j \equiv' \top$ we have $j \equiv \perp$ or $j \equiv \top$,
   - $\sqsubset' \neq \sqsubset$. □

Intuitively, the first case captures the situation where there are indices with fractional parts of some clocks being 0. If no such clock has an integer part which is 0, time can move backwards by making these clocks (corresponding to the indices j with $j \equiv' \perp$) decrease their integer parts by one, and by making their fractional parts become the largest (wrp. to $\sqsubset$) among all clocks which are less than max. The second case captures the situation when no clocks have fractional parts equal to 0. In this situation, the smallest step backwards in time corresponds to preserving the relative order between the fractional parts of clocks which are less than max, and by doing either or both of

- making the fractional parts of all clocks that are minimal wrp. to $\sqsubset'$ become 0,
- making the fractional parts of some clocks with values larger than max become 0.

Lemma 5.5 There is no infinite sequence $\phi_0, \phi_1, \phi_2, \ldots$ of constraints, such that $\phi_{i+1} \in pre_t(\phi_i)$.

Proof. The proof can be found in the appendix. □

Definition 5.6 For a set of constraints Φ' and a natural number i, we define ${pre_t}^i(\Phi', i)$ inductively as follows.

- ${pre_t}^0(\Phi') = \Phi'$; and
- ${pre_t}^{i+1}(\Phi') = \{\phi \ : \ \exists \phi' \in {pre_t}^i(\Phi') : \phi \in pre_t(\phi')\}$.

We define ${pre_t}^*(\Phi') = \cup_{i \geq 0} {pre_t}^i(\Phi')$.

Sometimes we write ${pre_t}^*(\phi')$ instead of ${pre_t}^*(\{\phi'\})$

Lemma 5.7 If ϕ' is a constraint, then $pre_T(\phi')$ is the denotation of the set of constraints in the set ${pre_t}^*(\{\phi'\})$. In other words

$$pre_T(\phi') \ = \ [\![{pre_t}^*(\phi)]\!]$$

Proof. The proof can be found in the appendix. □

The computability of pre_T follows from Lemma 5.7 and Lemma 5.5.

5.3 Computing *pre*

By combining the rules for computing $pre_D(\phi)$ in Lemma 5.3 and the rules for computing $pre_T(\phi)$ in Lemma 5.7, we obtain the following theorem.

Theorem 5.8 If ϕ is a constraint, then we can compute a finite set Φ of constraints such that $[\![\Phi]\!] = pre(\phi)$.

6 The entailment ordering is a well quasi-ordering

In this section, we shall prove that the preorder $\preceq$ on constraints is a well quasi-ordering. We will first review some standard results from the literature concerning well quasi-orderings ([Hig52]), and then apply them to our constraint system.

Definition 6.1 Let A be a set. A quasi-order $\preceq$ on A is a binary relation over A which is reflexive and transitive. A quasi-order $\preceq$ is a *well quasi-ordering* (wqo) if in each infinite sequence $a_0\ a_1\ a_2\ a_3 \cdots$ of elements in A, there are indices $i < j$ such that $a_i \preceq a_j$. $\square$

We shall now restate two standard lemmas, which allow us to lift well quasi-orderings from elements to bags and to sequences. Let A^* denote the set of finite strings over A, and let Let A^B denote the set of finite bags over A. An element of A^* and of A^B can be represented as a mapping $w\colon \widehat{|w|} \mapsto A$ where $|w|$ is the size of the bag or the length of the sequence. Given a quasi-order $\preceq$ on a set A, define the quasi-order $\preceq^*$ on A^* by letting $w \preceq^* w'$ if and only if there is a monotone[9] injection $h\colon \widehat{|w|} \mapsto \widehat{|w'|}$ such that $w(j) \preceq w'(h(j))$ for $1 \leq j \leq |w|$. Define the quasi-order $\preceq^B$ on bags of A by $w \preceq^* w'$ if and only if there is a (not necessarily monotonic) injection $h\colon \widehat{|w|} \mapsto \widehat{|w'|}$ such that $w(j) \preceq w'(h(j))$ for $1 \leq j \leq |w|$.

Lemma 6.2 If $\preceq$ is a wqo on A, then $\preceq^*$ is a wqo on A^* and $\preceq^B$ is a wqo on A^B.

Proof. The proof can be found in [Hig52]. $\square$

Let ϕ be a constraint $\langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq\rangle$. For each $j \in \widehat{m} \cup \{\bot, \top\}$, define the *rank* of j in ϕ to be the number of equivalence classes of $\sqsubseteq$ which are less than or equal (wrp. to $\preceq$) to the equivalence class containing j. In other words, the rank of j is the maximum k such that there is a sequence $\bot \sqsubset j_1 \sqsubset \cdots \sqsubset j_k = j$. Note that the rank of $\top$ is equal to the number of equivalence classes of $\equiv$. Define the rank of ϕ as the rank of $\top$ in ϕ.

Let r be the rank of the constraint $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq\rangle$. For $i \in \widehat{r}$, define $\phi[i]$ to be the bag of pairs of the form $\langle q, k\rangle$ such that $u(j) = q$ and $\mathbf{k}(j) = k$ for some j with rank i in ϕ. Define the ordering $\preceq$ on these pairs to be the identity relation on pairs of the form $\langle q, k\rangle$. Since there are finitely many such pairs, $\preceq$ is trivially a well quasi-ordering.

Lemma 6.3 Let $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq\rangle$ and $\phi' = \langle c', m', \mathbf{q}', \mathbf{k}', \sqsubseteq'\rangle$ be constraints with ranks r and r'. We have $\phi \preceq \phi'$ if and only if $c = c'$, $\phi[0] \preceq^B \phi'[0]$, $\phi[r] \preceq^B \phi'[r']$, and there is a monotonic injection $h : \widehat{(r-1)} \mapsto \widehat{(r'-1)}$ such that $\phi[i] \preceq^B \phi'[h(i)]$ for all $i \in \widehat{(r-1)}$.

[9] meaning that $h(j_1) \leq h(j_2)$ if and only if $j_1 \leq j_2$

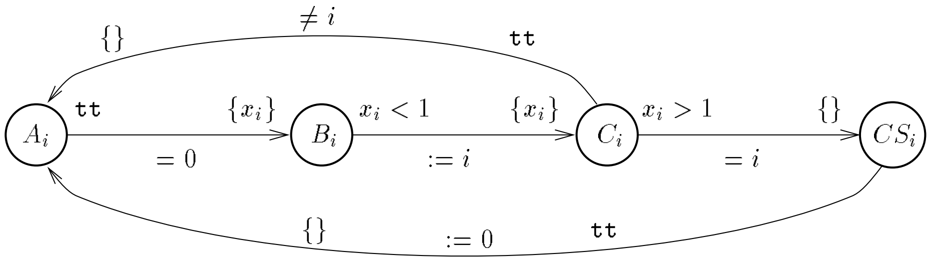

Figure 1: Fischer's Protocol for Mutual Exclusion.

Proof. The proof follows from the definitions. □

Theorem 6.4 The relation $\preceq$ on the set of constraints is a well quasi-ordering.

Proof. The proof follows from Lemma 6.3 and repeated application of Lemma 6.2. □

7 Example: Fischer's Protocol

As an illustration of our method, we model Fischer's protocol [SBK92], and show how an automatic verification algorithm would go about verifying mutual exclusion. Fischer's protocol has been used as a measure of the performance of tools for verification of timed automata. The example was suggested by Fred Schneider and has been verified manually (e.g., [AL92]) and using theorem provers (e.g., [Sha93b]). Several tools for verifying automata with a fixed number of clocks have been used to verify it for an increasing number of processes (e.g., [ACHH92]). Kristoffersen et al. [KLL$^+$97] describes an experiment where the number of processes is 50.

The purpose of the protocol is to guarantee mutual exclusion in a concurrent system consisting of an arbitrary number of processes, using clocks and a shared variable. Each process has a local clock, and runs a protocol before entering the critical section. Each process has a local control state, which in our model assumes values in the set $\{A, B, C, CS\}$ where A is the initial state and CS represents the critical section. The processes also read from and write to a shared variable whose value is either $\perp$ or the index of one of the processes. A description in a graphical pseudo-code (taken from [KLL$^+$97]) of the behavior of a process with index i is given in Figure 1.

Intuitively, the protocol behaves as follows: A process wishing to enter the critical section starts in state A. If the value of the shared variable is $\perp$, the process can proceed to state B and reset its local clock. From state B, the process can proceed to state C if the clock value is still less than 1. In other words, the clock implements a timeout which guarantees that the process either stays in state B at most one time unit, or gets stuck in B

forever. When moving from B to C, the process sets the value of the shared variable to its own index i and again resets its clock. From state C, the process can proceed to the critical section if the clock is strictly more than 1 and if the value of the shared variable is still i, the index of the process. Thus, in state C the clock enforces a delay which is longer than the length of the timeout in state B. Finally, when exiting the critical section, the process resets the shared variable to $\perp$. Processes that get stuck in state C can reenter the protocol by returning to state A. Since we do not intend to model liveness properties, such as e.g., absence of starvation, we do not impose requirements that force processes to change their state[10].

A rough argument for the correctness of the protocol goes as follows. The conditions on the shared variable ensure that a process cannot reach B if any other process is in C or CS. The timing conditions on the clocks ensure that a process cannot move from C to CS if some other process is still in B. Thus, if a set of processes start the mutual exclusion protocol and all arrive in C, then the process which was the last to enter C will read its own identity in the shared variable and enter the critical section.

$$
\begin{array}{ll}
\mathit{initiate}: & \langle\langle \mathit{udf}, \mathit{udf}\rangle, \langle A, x \geq 0 \longrightarrow \mathit{reset}, B\rangle\rangle \\
\\
\mathit{choose}_1: & \langle\langle \mathit{udf}, \mathit{df}\rangle, \langle B, x < 1 \longrightarrow \mathit{reset}, C^\dagger\rangle\rangle \\
\mathit{choose}_2: & \langle\langle \mathit{df}, \mathit{df}\rangle, \langle B, x < 1 \longrightarrow \mathit{reset}, C^\dagger\rangle, \langle q^\dagger, \mathit{skip}, q\rangle\rangle \\
\mathit{choose}_3: & \langle\langle \mathit{df}, \mathit{df}\rangle, \langle B^\dagger, x < 1 \longrightarrow \mathit{reset}, C^\dagger\rangle\rangle \\
\\
\mathit{enter}: & \langle\langle \mathit{df}, \mathit{df}\rangle, \langle C^\dagger, x > 1 \longrightarrow \mathit{skip}, CS^\dagger\rangle\rangle \\
\\
\mathit{fail}_1: & \langle\langle \mathit{udf}, \mathit{udf}\rangle, \langle C, \mathit{skip}, A\rangle\rangle \\
\mathit{fail}_2: & \langle\langle \mathit{df}, \mathit{df}\rangle, \langle C, \mathit{skip}, A\rangle\rangle \\
\\
\mathit{exit}_1: & \langle\langle \mathit{df}, \mathit{udf}\rangle, \langle CS^\dagger, \mathit{skip}, A\rangle\rangle \\
\mathit{exit}_2: & \langle\langle \mathit{df}, \mathit{udf}\rangle, \langle CS, \mathit{skip}, A\rangle, \langle q^\dagger, \mathit{skip}, q\rangle\rangle \\
\mathit{exit}_3: & \langle\langle \mathit{udf}, \mathit{udf}\rangle, \langle CS, \mathit{skip}, A\rangle\rangle
\end{array}
$$

Figure 2: Rules for Modeling Fischer's protocol

We will now make a model of the protocol in our timed networks formalism. The controller state is either *udf*, indicating that the value of the shared variable is undefined, or *df*, indicating that the value of the shared variable is defined. The set of process states is given by $\left\{A, B, C, CS, A^\dagger, B^\dagger, C^\dagger, CS^\dagger\right\}$. The states marked with † correspond to configurations where the value of the shared variable is equal to the index of that particular process.

[10]In fact, our formalism cannot express such requirements, although they can be added in terms of e.g., fairness constraints.

A straightforward translation of the description in Figure 1 yields the set of rules in Figure 2. We use q to denote an arbitrary process state. We use *skip* to denote the guarded command $0 \leq x \longrightarrow skip$.

Using the method described in this paper, it is possible to verify that there is never more than one process in its critical section. Below, we give a brief sketch of how a proof could be carried out using our method.

To verify mutual exclusion, let Φ_0 be the set of constraints with $m = 2$, which contain exactly two occurrences of CS or $CS^{\dagger}$. We refrain from writing out each individual constraint, since there are at least 18 constraints containing one occurrence of CS and one occurrence of $CS^{\dagger}$. Let us start a backwards reachability analysis from Φ_0. In each step of the analysis, we will apply *pre* to a set Φ of constraints by first running a rule r backwards and thereafter run time backwards, i.e., we will compute $pre_T(pre(r, \Phi))$. Applying pre_T will be done without explicitly mentioning it. We will present this procedure in the form of a table in Figure 3. The table contains sets of constraints, which are described by symbolic conditions on the constraints in the set. For instance, the symbolic condition

$$df, B, CS^{\dagger}, x_1 < 1$$

represents all constraints with $m = 2$, such that the controller state is df, such that $\mathbf{q}(1) = B$ (i.e.. the state of the first process is B), such that $\mathbf{q}(2) = CS^{\dagger}$ (i.e.. the state of the second process is $CS^{\dagger}$), and such that $\mathbf{k}$ and $\sqsubseteq$ are satisfied by clock states where the clock of the first process is less than one. (which means requiring that $\mathbf{k}(1) = 0$). For each set Φ of constraints, we enumerate the rules which can be used to generate pre-images, i.e., we enumerate the rules r such that $pre(r, \Phi)$ is non-empty. We add a matching between indices of the constraints and indices of the rule, and thereafter state which other sets of constraints are entailed by $pre_T(pre(r, \Phi))$.

8 Some Notes on Implementation

In this section, we present a refined version of the reachability algorithm of Section 3, which follows a paradigm of some of our earlier works [AJ96b, AvJYK96]. In comparison with the preceding discussion, the main difference is that instead of computing a sequence $\Phi_0\ \Phi_1\ \Phi_2\ \ldots$ of explicitly represented finite sets of constraints, we will maintain *one* single set of constraints. The basic step of our algorithm is to pick a constraint ϕ in the set, compute its pre-image $pre(\phi)$, and add those constraints in $pre(\phi)$ whose denotations can not easily be shown to be included in the denotation of the current set of constraints. It follows that this procedure will converge to a fixpoint, which denotes the same set of configurations as the denotation of the limit Φ_k in the above discussion. A pseudo-code description of this algorithm is shown in Figure 4. In the algorithm, we maintain two sets V and W of constraints.

Notation	Requirements on constraints	rule	matching	constraint entailed by *pre*
Φ_0	$*, CS^*, CS^*$	*enter*	$1 \mapsto 1$	Φ_1
		choose$_2$	$1 \mapsto 2,$	Φ_0
		exit$_2$	$1 \mapsto 2,$	Φ_0
Φ_1	$df, C^\dagger, CS^*$	*choose*$_2$	$1 \mapsto 1, 2 \mapsto 2$	Φ_2
		choose$_2$	$1 \mapsto 1,$	Φ_7
		choose$_1$	$1 \mapsto 1$	Φ_4
		choose$_3$	$1 \mapsto 1$	Φ_8
		enter	$2 \mapsto 1$	Φ_9
Φ_2	$df, B, CS^\dagger, x_1 < 1$	*enter*	$2 \mapsto 1$	Φ_3
		choose$_2$	$1 \mapsto 2$	Φ_8
Φ_3	$df, B, C^\dagger, x_1 < 1, x_2 > x_1$	*choose*$_1$		Φ_{11}
Φ_4	$udf, B, CS^*, x_1 < 1$	*initiate*	$1 \mapsto 1$	Φ_5
Φ_5	udf, A, CS^*	*exit*$_1$	$1 \mapsto 1$	Φ_0
		exit$_2$	$1 \mapsto 1$	Φ_0
		exit$_2$	$2 \mapsto 2$	Φ_5
		exit$_2$	$1 \mapsto 1, 2 \mapsto 2$	Φ_0
		exit$_2$	$1 \mapsto 2$	Φ_0
		exit$_3$	$1 \mapsto 1$	Φ_0
		fail$_1$	$1 \mapsto 1$	Φ_6
Φ_6	udf, C, CS^*	*enter*	$2 \mapsto 1$	Φ_{13}
		exit$_2$	$1 \mapsto 2$	Φ_0
Φ_7	$df, B, CS^*, q^\dagger, x_1 < 1$	*choose*$_1$	$3 \mapsto 1$	Φ_4
		choose$_2$	$3 \mapsto 1$	Φ_7
		choose$_2$	$1 \mapsto 2, 3 \mapsto 1$	Φ_7
		choose$_2$	$2 \mapsto 2$	Φ_2
		enter	$2 \mapsto 2$	Φ_3
Φ_8	$df, B^\dagger, CS^*, x_1 < 1$	*enter*	$2 \mapsto 1$	Φ_{11}
		choose$_1$		Φ_4
Φ_9	$df, C^\dagger, C^\dagger$	*choose*$_3$	$1 \mapsto 1$	Φ_{10}
		choose$_2$	$1 \mapsto 1$	Φ_{14}
		choose$_1$	$1 \mapsto 1$	Φ_{11}
Φ_{10}	$df, B^\dagger, C^\dagger, x_1 < 1$	*choose*$_3$	$2 \mapsto 1$	Φ_{15}
Φ_{11}	$udf, B, C^\dagger, x_1 < 1$	*initiate*	$1 \mapsto 1$	Φ_{12}
		exit$_2$	$1 \mapsto 2$	Φ_{10}
Φ_{12}	$udf, A, C^\dagger$	*fail*$_1$	$1 \mapsto 1$	Φ_{13}
		exit$_2$	$1 \mapsto 2$	Φ_1
		exit$_3$	$1 \mapsto 1$	Φ_1
Φ_{13}	$udf, C, C^\dagger$	*exit*$_2$	$1 \mapsto 1$	Φ_1
Φ_{14}	$df, B^\dagger, B^\dagger, x_1 < 1, x_2 < 1$	*choose*$_1$		Φ_{16}
Φ_{15}	$udf, B^\dagger, B^\dagger, B, x_1 < 1, x_2 < 1$	*initiate*		Φ_{17}
Φ_{16}	$udf, B^\dagger, B^\dagger, A, x_1 < 1, x_2 < 1$	*fail*$_1$		Φ_{18}
Φ_{17}	$udf, B^\dagger, B^\dagger, C, x_1 < 1, x_2 < 1$	*exit*$_2$	$3 \mapsto 2$	Φ_{10}

Figure 3: Table Representing a Verification of the protocol

```
Procedure Reachable(Φ)
var W , V : sets of constraints
begin
    W := Φ
    V := ∅
    while W ≠ ∅ do
        choose φ ∈ W
        if ⟦φ⟧ contains an initial configuration
        then return reachable
        else
            if ∃φ' ∈ V : φ' ⪯ φ
            then skip
            else
                V := V ∪ {φ}
                W := W ∪ pre(φ) \ {φ}
            fi
        fi
    od
    return unreachable
end
```

Figure 4: Algorithm for deciding reachability

The set V is the set of constraints whose predecessors have already been generated, and the set W is the set of constraints whose predecessors have not yet been generated. The algorithm repeatedly selects a constraint ϕ from W and adds $pre(\phi)$ to the set W of "unexplored" constraints. The constraint ϕ is moved from W to V. However, if $[\![\phi]\!]$ is included in $[\![\phi']\!]$ for some $\phi' \in V$, i.e., if $\phi' \preceq \phi$, then the addition of ϕ to V will not change $[\![V]\!]$. Therefore ϕ is simply discarded in this case. We also need not compute $pre(\phi)$ since this set is included in the already computed $pre(\phi')$ by monotonicity of pre.

Termination of this procedure follows from the fact that $\preceq$ is a well quasi-ordering. This implies that in any run of the algorithm, the test $\exists \phi' \in V : \phi' \preceq \phi$ can be false only a finite number of times, implying that eventually W will become empty.

From the small example in Section 7, it seems that an implementation which is based on our constraints will be rather inefficient, due to the explosion in the number of constraints, even for small numbers of max and the number of clocks in a constraint. In analogy with what is done for the analysis of timed automata, it appears more tractable to work with sets of constraints, which can be characterized by simple conditions on the states and clocks. Sets of constraints of the form that appear in Table 3 would probably be a suitable basic building block in an implementation. Such sets can be represented symbolically in a form similar to the representation used in Table 3. A similar representation is actually used in the verification of timed automata, where these blocks are often referred to as "zones".

In an implementation, we could use the algorithm of Figure 4, but with constraints replaced by sets of constraints of form that appear in Table 3. There is, however, one problem with this "more efficient" version of the reachability algorithm. The entailment ordering between sets of constraints is not a well quasi-ordering. This means that we cannot guarantee that the test $\exists\phi' \in V : \phi' \preceq \phi$ can be false only a finite number of times when ϕ represents a set of constraints. One way to guarantee termination is to replace this test by $[\![\phi]\!] \subseteq [\![V]\!]$. This test is, however, much more expensive to compute. On the other hand, our analysis of Fischer's protocol suggests that the simpler version of the test may suffice in many practical situations.

References

[ACD90] R. Alur, C. Courcoubetis, and D. Dill. Model-checking for real-time systems. In *Proc. 5^{th} IEEE Int. Symp. on Logic in Computer Science*, pages 414–425, Philadelphia, 1990.

[ACHH92] R. Alur, C. Courcoubetis, T. Henzinger, and P.-H. Ho. Hybrid automata: An algorithmic approach to the specification and verificationof hybrid systems. In Grossman, Nerode, Ravn, and Rischel, editors, *Hybrid Systems*, number 736 in Lecture Notes in Computer Science, pages 209–229, 1992.

[AH89] R. Alur and T. Henzinger. A really temporal logic. In *Proc. 30^{th} Annual Symp. Foundations of Computer Science*, pages 164–169, 1989.

[AJ96a] Parosh Aziz Abdulla and Bengt Jonsson. Undecidable verification problems for programs with unreliable channels. *Information and Computation*, 130(1):71–90, 1996.

[AJ96b] Parosh Aziz Abdulla and Bengt Jonsson. Verifying programs with unreliable channels. *Information and Computation*, 127(2):91–101, 1996.

[AK95] Parosh Aziz Abdulla and Mats Kindahl. Decidability of simulation and bisimulation between lossy channel systems and finite state systems. In Lee and Smolka, editors, *Proc. CONCUR '95, 6^{th} Int. Conf. on Concurrency Theory*, volume 962 of *Lecture Notes in Computer Science*, pages 333 – 347. Springer Verlag, 1995.

[AL92] M. Abadi and L. Lamport. An old-fashioned recipe for real time. In de Bakker, Huizing, de Roever, and Rozenberg, editors, *Real-Time: Theory in Practice*, volume 600 of *Lecture Notes in Computer Science*, 1992.

[AvJYK96] Parosh Aziz Abdulla, Karlis Čerāns, Bengt Jonsson, and Tsay Yih-Kuen. General decidability theorems for infinite-state systems. In *Proc.* 11^{th} *IEEE Int. Symp. on Logic in Computer Science*, pages 313–321, 1996.

[BBK77] J. M. Barzdin, J. J. Bicevskis, and A. A. Kalninsh. Automatic construction of complete sample systems for program testing. In *IFIP Congress, 1977*, 1977.

[BD91] B. Berthomieu and M. Diaz. Modeling and verification of time dependent systems using time petri nets. *IEEE Trans. on Software Engineering*, 17(3):259–273, 1991.

[BS95] O. Burkart and B. Steffen. Composition, decomposition, and model checking of pushdown processes. *Nordic Journal of Computing*, 2(2):89–125, 1995.

[CES86] E.M. Clarke, E.A. Emerson, and A.P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specification. *ACM Trans. on Programming Languages and Systems*, 8(2):244–263, April 1986.

[CGJ95] E. M. Clarke, O. Grumberg, and S. Jha. Verifying parameterized networks using abstraction and regular languages. In Lee and Smolka, editors, *Proc. CONCUR '95,* 6^{th} *Int. Conf. on Concurrency Theory*, volume 962 of *Lecture Notes in Computer Science*, pages 395–407. Springer Verlag, 1995.

[Č92a] K. Čerāns. Decidability of bisimulation equivalence for parallel timer processes. In *Proc. Workshop on Computer Aided Verification*, volume 663 of *Lecture Notes in Computer Science*, pages 302–315, 1992.

[Č92b] K. Čerāns. Feasibility of finite and infinite paths in data dependent programs. In *LFCS'92*, volume 620 of *Lecture Notes in Computer Science*, pages 69–80, 1992.

[Č94] K. Čerāns. Deciding properties of integral relational automata. In Abiteboul and Shamir, editors, *Proc. ICALP '94*, volume 820 of *Lecture Notes in Computer Science*, pages 35–46. Springer Verlag, 1994.

[Fin90] A. Finkel. Reduction and covering of infinite reachability trees. *Information and Computation*, (89):144–179, 1990.

[GMMP91] C. Ghezzi, D. Mandrioli, S. Morasca, and M. Pezzè. A unified high-level petri net formalism for time-critical systems. *IEEE Trans. on Software Engineering*, 17(2):160–172, 1991.

[God94] J.C. Godskesen. *Timed Modal Specifications*. PhD thesis, Aalborg University, 1994.

[GS92] S. M. German and A. P. Sistla. Reasoning about systems with many processes. *Journal of the ACM*, 39(3):675–735, 1992.

[GW93] P. Godefroid and P. Wolper. Using partial orders for the efficient verification of deadlock freedom and safety properties. *Formal Methods in System Design*, 2(2):149–164, 1993.

[Hen95] T.A. Henzinger. Hybrid automata with finite bisimulations. In *Proc. ICALP '95*, 1995.

[Hig52] G. Higman. Ordering by divisibility in abstract algebras. *Proc. London Math. Soc.*, 2:326–336, 1952.

[Jan90] P. Jančar. Decidability of a temporal logic problem for petri nets. *Theoretical Computer Science*, 74:71–93, 1990.

[JM95] P. Jančar and F. Moller. Checking regular properties of petri nets. In *Proc. CONCUR '95, 6^{th} Int. Conf. on Concurrency Theory*, pages 348–362, 1995.

[JP93] B. Jonsson and J. Parrow. Deciding bisimulation equivalences for a class of non-finite-state programs. *Information and Computation*, 107(2):272–302, Dec. 1993.

[KLL$^+$97] K.J. Kristoffersen, F. Larroussinie, K. G. Larsen, P. Pettersson, and W. Yi. A compositional proof of a real-time mutual exclusion protocol. In *TAPSOFT '97 7th International Joint Conference on the Theory and Practice of Software Development*, Lecture Notes in Computer Science, Lille, France, April 1997. Springer Verlag.

[KM89] R.P. Kurshan and K. McMillan. A structural induction theorem for processes. In *Proc. 8^{th} ACM Symp. on Principles of Distributed Computing, Canada*, pages 239–247, Edmonton, Alberta, 1989.

[LSW95] K.G. Larsen, B. Steffen, and C. Weise. Fischer's protocol revisited: a simples proof using modal constraints. In *4th DIMACS Workshop on Verification and Control of Hybrid Systems*, New Brunswick, New Jersey, Oct. 1995.

[MF76] P. Merlin and D.J. Farber. Redoverability of communication protocols - implications of a theoretical study. *IEEE Trans. on Computers*, COM-24:1036–1043, Sept. 1976.

[QS82] J.P. Queille and J. Sifakis. A temporaal logic to deal with fairness in transition systems. In *Proc. 23^{rd} Annual Symp. Foundations of Computer Science*, pages 217–225, 1982.

[SBK92] F. B. Schneider, Bloom B, and Marzullo K. Putting time into proof outlines. In de Bakker, Huizing, de Roever, and Rozenberg, editors, *Real-Time: Theory in Practice*, volume 600 of *Lecture Notes in Computer Science*, 1992.

[Sha93a] A.U. Shankar. An introduction to assertional reasoning for concurrent systems. *Computing Surveys*, 25(3):225–262, Sept. 1993.

[Sha93b] N. Shankar. Verification of real-time systems using PVS. In Courcoubetis, editor, *Proc. 5^{th} Int. Conf. on Computer Aided Verification*, volume 697 of *Lecture Notes in Computer Science*, pages 280–291, 1993.

[VW86] M. Y. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *Proc. 1^{st} IEEE Int. Symp. on Logic in Computer Science*, pages 332–344, June 1986.

[WL89] P. Wolper and V. Lovinfosse. Verifying properties of large sets of processes with network invariants (extended abstract). In Sifakis, editor, *Proc. Workshop on Computer Aided Verification*, number 407 in Lecture Notes in Computer Science, pages 68–80, 1989.

[Wol86] Pierre Wolper. Expressing interesting properties of programs in propositional temporal logic (extended abstract). In *Proc. 13^{th} ACM Symp. on Principles of Programming Languages*, pages 184–193, Jan. 1986.

[Yi91] Wang Yi. CCS + Time = an interleaving model for real time systems. In Leach Albert, Monien, and Rodriguez Artalejo, editors, *Proc. ICALP '91*, volume 510 of *Lecture Notes in Computer Science*. Springer Verlag, 1991.

A Appendix - Proof of Some Lemmas

A.1 Proof of Theorem 4.4

[If:] Assume that the conditions of the theorem hold and that $\langle I, c, \mathtt{q}, \mathtt{x} \rangle \in [\![\phi']\!]$. It follows that there is a mapping $h' : \widehat{m'} \mapsto I$ which satisfies the conditions in Definition 4.2. Define $g : \widehat{m} \mapsto I$ by $g \stackrel{\text{def}}{=} h' \circ h$. It is easy to check that g satisfies the conditions of Definition 4.2, implying that $\langle I, c, \mathtt{q}, \mathtt{x} \rangle \in [\![\phi]\!]$.

[Only if:] Assume that $\phi \preceq \phi'$. It is trivial to see that $[\![\phi]\!] \neq \emptyset$ for any constraint ϕ. Let $\gamma \in [\![\phi']\!]$, implying that $\gamma \in [\![\phi]\!]$ and that there is a mapping $h' : \widehat{m'} \mapsto I$ which satisfies the conditions in Definition 4.2. Since $\gamma \in [\![\phi]\!]$, there is also a mapping $h : \widehat{m} \mapsto I$ which satisfies the conditions in Definition 4.2. First observe that $\phi \preceq \phi'$ trivially implies $m \leq m'$. Define the mapping $g : \widehat{m} \mapsto \widehat{m'}$ by $g \stackrel{\text{def}}{=} (h')^{-1} \circ h$. It is easy to see that g is an injection which satisfies the conditions in the theorem. $\square$

A.2 Proof of Lemma 5.2

The proof is structured according to the structure of $p(x)$.

- If $p(x)$ is of form $k \leq x$ for some $k \in \{0, \ldots, max\}$ then $\langle \phi, j \rangle \models p(x)$ is equivalent to $\mathbf{k}(j) \geq k$, which by Definition 4.2 is equivalent to $\min(max, \lfloor \mathtt{x}(h(j)) \rfloor) \geq k$, which, since $k \leq max$, is equivalent to $\mathtt{x}(h(j)) \geq k$, which is the same as $p(\mathtt{x}(h(j)))$.
- The other cases are analogous.

$\square$

A.3 Proof of Lemma 5.3

We state an auxiliary lemma, which will be used in the proof of Lemma 5.8.

Lemma A.1 Let $\gamma = \langle I, c, \mathtt{q}, \mathtt{x} \rangle$ be a configuration, let $m \in \mathcal{N}$ with $m \leq |I|$ be a natural number which is at most the size of I, and let $h : \widehat{m} \mapsto I$ be an injection. Then there is a unique constraint $\phi_{\gamma,h}$ which satisfies the 5 conditions in Definition 4.2.

Proof. It can be seen that the conditions in Definition 4.2 uniquely define the components of $\phi_{\gamma,h}$ under the assumptions in the lemma. $\square$

Proof of the Lemma [If:] We show that $pre(r, \phi')$ is a subset of the denotations of constraints satisfying the conditions in the lemma. Thus, assume that $\gamma \in pre(r, \phi')$, i.e., that there is a $\gamma' \in [\![\phi']\!]$ with $\gamma \xrightarrow{r}_D \gamma'$. We must prove that $\gamma \in [\![\phi]\!]$ for some constraint ϕ which satisfies the conditions of the lemma. Let $\gamma = \langle I, c, \mathtt{q}, \mathtt{x}\rangle$ and $\gamma' = \langle I, c', \mathtt{q}', \mathtt{x}'\rangle$, where I is the set of indices in γ and γ'. By the definition of $\xrightarrow{r}_D$, there is an injection $f\colon \widehat{n} \mapsto I$ satisfying the conditions of Definition 2.3. Since $\gamma' \in [\![\phi']\!]$, there is an injection $h' : \widehat{m'} \mapsto I$ satisfying the conditions in Definition 4.2. Define the partial injection $g' : \widehat{m'} \mapsto \widehat{n}$ by $g' \stackrel{\mathrm{def}}{=} (f)^{-1} \circ h'$. The domain of g' is the set of indices $j \in \widehat{m'}$ such that $h'(j) \in range(f)$. Let $m = m' + (n - range(h)|)$ and let $g : \widehat{m} \mapsto \widehat{n}$ be an (arbitrary) extension of g', i.e., g is a surjective partial injection. Extend h' to an injection $h : \widehat{m} \mapsto I$ by letting $h = h'$ on $\widehat{m'}$ and $h = f \circ g$ on $domain(g)$. Note that h is well-defined since $f \circ g = f \circ ((f)^{-1} \circ h') = h'$ on $domain(g) \cap \widehat{m'}$. We will take the sought constraint ϕ to be $\phi_{\gamma,h}$, which by Lemma A.1 satisfies the conditions in Definition 4.2, with h as the injection. In order to help the reader, Figure 5 shows the injections between the different index sets, and their relationships. The figure contains the (partial) injections between the index set I of the involved configurations, the index sets of the the constraints ϕ and ϕ', and the index set of the rule r.

$$\begin{tikzcd}
\phi \arrow[r, "g"] \arrow[dr, "h"'] & r \arrow[d, "f"] & \phi' \arrow[l, "g'"'] \arrow[dl, "h'"] \\
& I &
\end{tikzcd}$$

Figure 5: Injections and partial injections between index sets.

We must check that $\phi_{\gamma,h}$ and g satisfies the conditions of the lemma. Let us treat the conditions one by one.

1. Let $j \in domain(g)$. We have that $\mathbf{q}(j)$ equals $\mathtt{q}(h(j))$ by the definition of $\phi_{\gamma,h}$, which equals $\mathtt{q}(f(g(j)))$ by the definition of h, which by condition 1 in Definition 2.3 equals $q_{g(j)}$ (noting that $g(j) \in \widehat{n}$). We have by condition 1 of Definition 2.3 that $p_{g(j)}(\mathtt{x}(f(g(j))))$ holds, which by Lemma 5.2 implies $\langle \phi, j\rangle \models p_{g(j)}(x)$.

2. Let $j \in domain(g')$. We have that $\mathbf{q}'(j)$ equals $\mathtt{q}'(h'(j))$ since $\gamma' \in [\![\phi']\!]$, which equals $q'(f(g'(j)))$ by the definition of h', which by condition 2 in Definition 2.3 equals $q'_{g'(j)}$.

3. Let $j \in \widehat{m'} \setminus domain(g')$, i.e., $h'(j) \notin range(f)$. By condition 2 in Definition 2.3 we have $\mathtt{q}'(h'(j)) = \mathtt{q}(h'(j))$, which by $\gamma' \in [\![\phi']\!]$ and

Definition 4.2 implies $\mathbf{q}'(j) = \mathtt{q}(h'(j))$, which from the definition of $\phi_{\gamma,h}$ (noting that $h = h'$ on $\widehat{m'} \setminus domain(g')$) implies $\mathbf{q}'(j) = \mathbf{q}(j)$.

4. Let $j \in domain(g')$ and $op_{g'(j)} = reset$. We have by condition 4 of Definition 2.3 that $\mathtt{x}'(f(g'(j))) = 0$ holds, which since $h' = f \circ g'$ implies $\mathtt{x}'(h'(j)) = 0$, which by $\gamma' \in [\![\phi']\!]$ and Definition 4.2 implies $\mathbf{k}'(j) = 0$ and $j \equiv' \perp$.

5 and 6. Let $j \in domain(g')$ and $op_{g'(j)} = skip$. By condition 5 of Definition 2.3 we have that $\mathtt{x}'(f(g'(j))) = \mathtt{x}(f(g'(j)))$, which since $h' = f \circ g'$ implies $\mathtt{x}'(h'(j)) = \mathtt{x}(h'(j))$. If $j \in \widehat{m'} \setminus domain(g')$, i.e., $h'(j) \notin range(f)$, then by condition 6 in Definition 2.3 we have $\mathtt{x}'(h'(j)) = \mathtt{x}(h'(j))$. In both cases, the conclusion $\mathtt{x}'(h'(j)) = \mathtt{x}(h'(j))$ implies conditions 5. and 6, using the definition of $\phi_{\gamma,h}$, using $\gamma' \in [\![\phi']\!]$, and using Definition 4.2.

[If:] Assume that there are ϕ, g', and g which satisfy the conditions of the lemma with $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq \rangle$. We must prove that $[\![\phi]\!] \subseteq pre(r, \phi')$. Assume that $\gamma \in [\![\phi]\!]$. Let ϕ, ϕ', γ, and γ' be as before. By $\gamma \in [\![\phi]\!]$ there is an injection $h : \widehat{m} \mapsto I$ which satisfies the conditions in Definition 4.2. Define the injection $f : \widehat{n} \mapsto I$ by $f \stackrel{\text{def}}{=} h \circ g^{-1}$. If $i \in \widehat{n}$ then $i = g(j)$ for some $j \in domain(g)$. By condition 1 of the lemma, we have $\mathbf{q}(j) = q_{g(j)}$ and $\langle \phi, j \rangle \models p_{g(j)}(x)$. By $\gamma \in [\![\phi]\!]$ and Lemma 5.2 we have $\mathbf{q}(j) = \mathtt{q}(h(j))$ and $p_{g(j)}(\mathtt{x}(h(j)))$, implying $\mathtt{q}(f(i)) = q_i$ and $p_i(\mathtt{x}(h(j)))$. Thus, the first condition in Definition 2.3 is satisfied, implying that there is a unique γ' satisfying the conditions for $\gamma \xrightarrow{r}_D \gamma'$.

Let h' be the restriction of h to $\widehat{m'}$. We will now check that $\gamma' \in [\![\phi']\!]$ by proving that the conditions in Definition 4.2 are satisfied, using h' as the injection. The proof is structured into cases, depending on the index j.

- If $j \in domain(g')$ then by condition 2 we have $\mathbf{q}'(j) = q'_{g'(j)}$, and by condition 2 in Definition 2.3 we have $\mathtt{q}'(f(g'(j))) = q'_{g'(j)}$. We conclude that $\mathtt{q}'(h'(j))) = \mathtt{q}'(f(g'(j))) = q'_{g'(j)} = \mathbf{q}'(j)$.
  If $j \in \widehat{m'} \setminus domain(g')$, then by condition 3 we have $\mathbf{q}'(j) = \mathbf{q}(j)$. By $\gamma \in [\![\phi]\!]$ and condition 1 in Definition 4.2 we have $\mathtt{q}(h(j))) = \mathbf{q}(j)$, and by condition 3 in Definition 2.3 we have $\mathtt{q}(h(j))) = \mathtt{q}'(h(j)))$ since $h(j) \notin range(f)$. We conclude $\mathtt{q}'(h(j))) = \mathbf{q}'(j)$,
  This concludes the proof of the first condition in Definition 4.2.
- Let $j \in domain(g')$ and $op_{g'(j)} = reset$. Then by condition 4 we have $\mathbf{k}'(j) = 0$ and $j \equiv' \perp$. By condition 4 of Definition 2.3 we have $\mathtt{x}'(h'(j)) = \mathtt{x}'(f(g'(j))) = 0$, implying that conditions 2, 3, and 4 in Definition 4.2 are satisfied.
- Let j be such that either $j \in domain(g')$ and $op_{g'(j)} = skip$, or such that $j \in \widehat{m'} \setminus domain(g')$. By condition 5 and 6 of Definition 2.3 we have $\mathtt{x}'(h'(j)) = \mathtt{x}'(f(g'(j))) = \mathtt{x}(f(g'(j))) = \mathtt{x}(h(j))$. Since $\gamma \in [\![\phi]\!]$,

the fact that γ satisfies conditions 2 - 4 of Definition 2.3 implies that γ' also does. The fact that γ satisfies conditions 5 for indices j_1 and j_2 implies that γ' also does in the case that for $i = 1, 2$ we have either $j_i \in domain(g')$ and $op_{g'(j_i)} = skip$ or $j_i \in \widehat{m'} \setminus domain(g')$. If $j_i \in domain(g')$ and $op_{g'(j_i)} = reset$ this follows by noting that $\mathtt{x}'(h'(j)) = 0$ was shown in the previous case.

□

A.4 Proof of Lemma 5.5

For a constraint $\phi = \langle c, m, \mathbf{q}, \mathbf{k}, \sqsubseteq \rangle$, we say that ϕ is *of type 0* if there is a $j \in \widehat{m}$ such that $j \equiv \bot$, otherwise we say that ϕ is *of type 1*.

Suppose that we have an infinite sequence $\phi_0, \phi_1, \phi_2, \ldots$ of constraints, such that $\phi_{i+1} \in pre_t(\phi_i)$, for $i \geq 0$.

Let $\phi_i = \langle c_i, m_i, \mathbf{q}_i, \mathbf{k}_i, \sqsubseteq_i \rangle$, and let $K_i = \sum_{j \in \widehat{m_i}} \mathbf{k}_i(j)$. From Definition 5.4, it follows that for each $i \geq 0$, one the following holds.

- ϕ_i is of type 0, ϕ_{i+1} is of type 1, and $K_i > K_{i+1}$.
- ϕ_i is of type 1, ϕ_{i+1} is of type 0, and $K_i = K_{i+1}$.

This implies that $K_i > K_{i+2}$ for each $i \geq 0$, which is a contradiction, since each K_i is a natural number. □

A.5 Proof of Lemma 5.7

First, we give some preliminary definitions and auxiliary lemmas.

Let $\gamma = \langle I, c, \mathtt{q}, \mathtt{x} \rangle$ be a configuration. By $\mathtt{x}_{\min}(\gamma)$ we mean $\min\{\mathtt{x}(j) \; : \; j \in I\}$. For $\delta \in \mathcal{R}^{\geq 0}$, with $\delta \leq \mathtt{x}_{\min}(\gamma)$, we use $\gamma^{-\delta}$ to denote the configuration $\langle I, c, \mathtt{q}, \mathtt{x}' \rangle$, where $x'(j) = x(j) - \delta$ for each $j \in I$.

Lemma A.2 For a configuration γ' and a constraint ϕ', with $\gamma' \in [\![\phi']\!]$ and $\mathtt{x}_{\min}(\gamma') > 0$, there is a $\delta' \in \mathcal{R}^{\geq 0}$ and a constraint ϕ such that $\phi \in pre_t(\phi')$, and one of the following three conditions holds

1. $\gamma^{-\delta} \in [\![\phi']\!]$, for all $0 \leq \delta \leq \mathtt{x}_{\min}(\gamma)$;
2. $\gamma^{-\delta} \in [\![\phi]\!]$, for all $\delta : 0 < \delta \leq \delta'$; or
3. $\gamma^{-\delta} \in [\![\phi']\!]$, for all $\delta : 0 \leq \delta < \delta'$, and $\gamma^{-\delta'} \in [\![\phi]\!]$.

Proof. Let $\gamma' = \langle I', c', \mathtt{q}', \mathtt{x}' \rangle$, and $\phi = \langle c', m', \mathbf{q}', \mathbf{k}', \sqsubseteq' \rangle$. Let $h : \widehat{m} \mapsto I$ be an injection satisfying the conditions stated in Definition 4.2. There are three cases:

- If there is a $j \in \widehat{m}$ where $\mathit{fract}(x(h(j))) = 0$. Define $\delta_1 = \min\{\mathit{fract}(x(h(j))) : j \in \widehat{m} \text{ and } \mathit{fract}(x(h(j))) > 0\}$. We take δ' such that $0 < \delta' < \delta_1$ if δ_1 is well-defined, and $0 < \delta' < 1$ otherwise[11]. We define ϕ to be the (unique) constraint $\langle c', m', \mathbf{q}', \mathbf{k}, \sqsubseteq \rangle$ satisfying condition 1 in Definition 5.4. It follows that condition 2 in the lemma is satisfied.

- If there is no $j \in \widehat{m}$ where $fract(x(h(j))) = 0$, and there is at least on $j \in \widehat{m}$ such that $x(h(j)) < \max + 1$. We take $\delta' = \min\{\mathit{fract}(x(h(j))) : j \in \widehat{m} \text{ and } x(h(j)) < \max + 1\}$. If $\delta' > \mathtt{x}_{\min}(\gamma')$ then condition 1 is satisfied, otherwise we define γ to be the (unique) constraint $\langle I', c', \mathtt{q}', \mathtt{x} \rangle$ where

  - whenever $\perp \sqsubset' j_1, j_2 \sqsubset' \top$, we have $j_1 \sqsubseteq j_2$ if and only if $j_1 \sqsubseteq' j_2$,
  - $j \equiv \perp$ if and only if $\mathit{fract}(x(h(j))) = \delta'$ and $x(h(j)) < \max + 1$.
  - $j \equiv \top$ if and only if $x(h(j)) > \max + \delta'$.

  From Definition 5.4 it follows that condition 3 of the lemma is satisfied.

- If $x(h(j)) \geq \max + 1$ for each $j \in \widehat{m}$. Define $\delta' = \min\{x(h(j)) - \max : j \in \widehat{m}\}$. If $\delta' > \mathtt{x}_{\min}(\gamma)$ then condition 1 is satisfied, otherwise we define γ to be the (unique) constraint $\langle I', c', \mathtt{q}', \mathtt{x} \rangle$ where

  - $j \equiv \perp$ if and only if $x(h(j)) = \max + \delta'$,
  - $j \equiv \top$ if and only if $x(h(j)) > \max + \delta'$.

  From Definition 5.4 it follows that condition 3 of the lemma is satisfied.

□

Corollary A.3 For configurations γ and γ', and a constraint ϕ', if $\gamma' \in [\![\phi']\!]$ and $\gamma \longrightarrow_T \gamma'$, then there exists a constraint ϕ such that $\gamma \in [\![\phi]\!]$ and $\phi \in \mathit{pre}_t{}^*(\phi')$.

Lemma A.4 For a configuration γ, and constraints ϕ and ϕ', if $\phi \in \mathit{pre}_t(\phi')$ and $\gamma \in [\![\phi]\!]$, then there is a configuration γ' such that $\gamma' \in [\![\phi']\!]$ and $\gamma \longrightarrow_T \gamma'$.

Proof. Let $\gamma = \langle I, c, \mathtt{q}, \mathtt{x} \rangle$. We define $\gamma' = \gamma^{+\delta}$, where δ is defined according to one of the following two cases.

- If ϕ and ϕ' satisfy the conditions of case 1 in Definition 5.4, then we define $\delta < 1 - \max\{\mathit{fract}(x(j)) : j \in I \text{ and } x(j) \leq \max\}$.

- If ϕ and ϕ' satisfy the conditions of case 2 in Definition 5.4. We define $\delta_1 = \max\{\mathit{fract}(x(j)) : j \in I \text{ and } x(j) \leq \max\}$. We take $\delta = 1 - \delta_1$ if δ_1 is well-defined, and take δ arbitrarily otherwise[12].

[11] δ' will be undefined if there is no $j \in \widehat{m}$ such that $fract(x(h(j))) > 0$. We can also take $\delta_1 = \min\{\mathit{fract}(x(h(j))) : j \in \widehat{m} \text{ and } \mathit{fract}(x(h(j))) > 0 \text{ and } x(h(j)) < \max + 1\}$.

[12] δ' will be undefined if there is no $j \in I$ with $x(j) \leq \max$.

□

Corollary A.5 For a configuration γ, and constraints ϕ and ϕ', if $\phi \in {pre_t}^*(\phi')$ and $\gamma \in [\![\phi]\!]$, then there is a configuration γ' such that $\gamma' \in [\![\phi']\!]$ and $\gamma \longrightarrow_T \gamma'$.

Proof of Lemma 5.7 The proof follows form Corollary A.3 and Corollary A.5. □