

The Theory of Timed I/O Automata

Dilsun K. Kaynar and Nancy Lynch*
MIT Computer Science and Artificial Intelligence Laboratory

Roberto Segala
Dipartimento di Informatica, Università di Verona

Frits Vaandrager†
Nijmegen Institute for Computing and Information Sciences, University of Nijmegen

November 23, 2004

Abstract

This paper presents the *Timed Input/Output Automaton (TIOA)* modeling framework, a basic mathematical framework to support description and analysis of timed systems. An important feature of this model is its support for decomposing timed system descriptions. In particular, the framework includes a notion of *external behavior* for a timed I/O automaton, which captures its discrete interactions with its environment. The framework also defines what it means for one TIOA to *implement* another, based on an inclusion relationship between their external behavior sets, and defines notions of *simulations*, which provide sufficient conditions for demonstrating implementation relationships. The framework includes a *composition* operation for TIOAs, which respects external behavior, and a notion of *receptiveness*, which implies that a TIOA does not block the passage of time.

*Corresponding author's email address: dilsun@theory.lcs.mit.edu. Manuscript available at: <http://theory.lcs.mit.edu/tds/reflist.html>. Research supported by DARPA/AFOSR MURI Contract F49620-02-1-0325, DARPA SEC contract F33615-01-C-1850, NSF ITR contract CCR-0121277, and Air Force Aerospace Research-OSR Contract F49620-00-1-0097.

†Supported by EU IST project IST-2001-35304: Advanced Methods for Timed Systems (AMETIST) and PROGRESS project TES4999: Verification of Hard and Softly Timed Systems (HaaST).

Contents

1	Introduction	5
1.1	Overview	5
1.2	Related work	8
1.3	Paper Organization	9
2	Mathematical Preliminaries	9
2.1	Functions and Relations	9
2.2	Sequences	10
2.3	Partial Orders	11
2.4	A Basic Graph Lemma	12
2.5	Untimed Automata	12
3	Describing Timed System Behavior	12
3.1	Time	13
3.2	Static and Dynamic Types	13
3.3	Trajectories	14
3.3.1	Basic Definitions	14
3.3.2	Prefix Ordering	15
3.3.3	Concatenation	15
3.4	Hybrid Sequences	16
3.4.1	Basic Definitions	17
3.4.2	Prefix Ordering	18
3.4.3	Concatenation	18
3.4.4	Restriction	19
4	Timed Automata	20
4.1	Definition of Timed Automata	20
4.2	Executions and Traces	31
4.3	Special Kinds of Timed Automata	35
4.3.1	Basic constraints	35

4.3.2	Alur-Dill Automata	37
4.4	Implementation Relationships	39
4.5	Simulation Relations	41
4.5.1	Forward Simulations	42
4.5.2	Refinements	46
4.5.3	Backward Simulations	47
4.5.4	History Relations	50
4.5.5	Prophecy Relations	53
5	Operations on Timed Automata	55
5.1	Composition	55
5.1.1	Definitions and Basic Results	55
5.1.2	Substitutivity Results	61
5.2	Hiding	62
5.3	Extending Timed Automata with Bounds	65
5.4	Untiming	71
5.4.1	State Congruence	73
5.4.2	Definition of the Untiming Operation	73
5.4.3	Basic Results	74
5.4.4	An Equivalence Relation for Alur-Dill Automata	77
6	Properties for Timed Automata	81
6.1	Definitions and Basic Results	81
6.1.1	Safety and Liveness Properties	81
6.1.2	Machine-closure	83
6.1.3	Special kinds of properties	86
6.2	Implementation Relationships	89
6.3	Simulation Relations	89
6.4	Composition	96
6.4.1	Definitions and Basic Results	96
6.4.2	Substitutivity Results	96

7	Timed I/O Automata	97
7.1	Definition of Timed I/O Automata	97
7.2	Executions and Traces	98
7.3	Special Kinds of Timed I/O Automata	98
7.3.1	Feasible and I/O Feasible TIOAs	98
7.3.2	Progressive TIOAs	99
7.3.3	Receptive Timed I/O Automata	100
7.4	Implementation Relationships	102
7.5	Simulation Relations	102
8	Operations on Timed I/O Automata	102
8.1	Composition	102
8.1.1	Definitions and Basic Results	103
8.1.2	Substitutivity Results	104
8.1.3	Composition of Special Kinds of TIOAs	114
8.2	Hiding	115
9	Properties for Timed I/O Automata	116
9.1	Definitions and Basic Results	116
9.2	Composition	117
9.3	Receptiveness for Properties	119
10	Conclusions	121
A	Notational Conventions	123

1 Introduction

1.1 Overview

Timed computing systems are systems in which desirable correctness or performance properties of the system depend on the timing of events, not just on the order of their occurrence. A typical timed system consists of computer components, which operate in discrete steps, and timing-related components such as physical or logical clocks, whose behavior involve continuous transformation over time. Timed systems are employed in a wide range of domains including communications, embedded systems, real-time operating systems, and automated control. Many applications involving timed systems have strong safety, reliability and predictability requirements, which makes it important to have methods for systematic design of systems and rigorous analysis of timing-dependent behavior.

In this paper, we introduce a basic mathematical framework – the *Timed Input/Output Automaton* modeling framework – to support description and analysis of timed systems. A *Timed I/O Automaton (TIOA)* is a kind of nondeterministic, possibly infinite-state, state machine. The state of a TIOA is described by a valuation of state variables that are internal to the automaton. The state of a TIOA can change in two ways: instantaneously by the occurrence of a *discrete transition*, which is labeled by a discrete action, or according a *trajectory*, which is a function that describes the evolution of the state variables over intervals of time. Trajectories may be continuous or discontinuous functions.

The TIOA framework supports decomposition of system description and analysis. A key to this decomposition is the rigorously-defined notion of *external behavior* for timed I/O automata. The external behavior of each TIOA is defined by a simple mathematical object called a *trace*—essentially, a sequence of actions interspersed with time-passage steps. *Abstraction* and *parallel composition* are other important notions for decomposition of system description and analysis.

For abstraction, the framework includes notions of *implementation* and *simulation*, which can be used to view timed systems at multiple levels of abstraction, starting from a high-level version that describes required properties, and ending with a low-level version that describes a detailed design or implementation. In particular, the TIOA framework defines what it means for one TIOA, \mathcal{A} , to *implement* another TIOA, \mathcal{B} , namely, any trace that can be exhibited by \mathcal{A} is also allowed by \mathcal{B} . In this case, \mathcal{A} might be more deterministic than \mathcal{B} , in terms of either discrete transitions or trajectories. For instance, \mathcal{B} might be allowed to perform an output action at an arbitrary time before noon, whereas \mathcal{A} produces the same output sometime between 10 and 11AM. The notion of a *simulation relation* from \mathcal{A} to \mathcal{B} provides a sufficient condition for demonstrating that \mathcal{A} implements \mathcal{B} . A simulation relation is defined to satisfy three conditions, one relating start states, one relating discrete transitions, and one relating trajectories of \mathcal{A} and \mathcal{B} .

For parallel composition, the framework provides a *composition operation*, by which TIOAs modeling individual timed system components can be combined to produce a model

for a larger timed system. The model for the composed system can describe interactions among the components, which involves joint participation in discrete transitions. Composition requires certain “compatibility” conditions, namely, that each output action be controlled by at most one automaton, and that internal actions of one automaton cannot be shared by any other automaton. The composition operation respects traces, for example, if \mathcal{A}_1 implements \mathcal{A}_2 then the composition of \mathcal{A}_1 and \mathcal{B} implements the composition of \mathcal{A}_2 and \mathcal{B} . Composition also satisfies *projection* and *pasting* results, which are fundamental for compositional design and verification of systems: a trace of a composition of TIOAs “projects” to give traces of the individual TIOAs, and traces of components are “pastable” to give behaviors of the composition.

A formal modeling framework needs to support the statement and verification of both *safety and liveness properties* if it is to be of general practical use. A safety property specifies the absence of certain undesirable events, while a liveness property specifies that certain desirable events eventually occur. The TIOA modeling framework defines the notions of safety and liveness properties for a TIOA, and what it means for a pair of safety and liveness properties to be *machine-closed*. Machine-closure refers to the condition that a liveness property does not impose safety constraints beyond those already imposed by the safety property, and is usually considered to be a reasonable condition to satisfy in defining safety and liveness properties for a system.

The proof of many interesting liveness properties for concurrent systems requires some assumption about each activity in the system getting “enough” chances to make progress. *Fairness properties* are special kinds of liveness properties that express this informal idea. The TIOA framework includes notions of weak and strong fairness, and results that state under which conditions the fair traces of a TIOA can be shown to be included in the fair traces of another.

An interesting complication that arises in the timed setting is the possibility that a state machine could exhibit the so called *Zeno behavior*, by allowing time to approach a finite point in time without quite reaching it, or by scheduling infinitely many discrete actions to happen in a finite amount of time. The TIOA framework includes a notion of *receptiveness*, which is used to classify automata that do not contribute to producing Zeno behavior, and which is preserved by composition. Receptiveness of a TIOA, \mathcal{A} , in the TIOA framework is defined in terms of the existence of a strategy, which is defined as a subautomaton of \mathcal{A} that chooses some of the evolutions from each state of \mathcal{A} . This simple notion of a strategy is used also in the statement of results that identify the conditions under which the outcome of a system’s interactions with its environment satisfies a liveness property.

The TIOA modeling framework presented in this paper has evolved from the recently introduced *Hybrid Input/Output Automaton (HIOA)* modeling framework for hybrid systems [22] by Lynch, Segala and Vaandrager. Our approach is based on the assumption that a timed system can be viewed as a special kind of a hybrid system where the continuous transformation is limited to internal system components that determine the timing

of events. Therefore, we define a TIOA as a restricted HIOA where the only essential difference between an HIOA and a TIOA is that an HIOA may have *external variables* to model the continuous information flowing into and out of the system, in addition to state variables. A major consequence of this definition is that the communication between TIOAs is restricted to shared-action communication only. The TIOA model does not impose any further restrictions on the expressive power of the HIOA model.

We have undertaken the project of developing this new modeling framework even though there are several timed automaton models that extend the basic I/O automaton model [29, 36, 27, 25], because we have observed that the new HIOA modeling framework of Lynch, Segala and Vaandrager offered a way of improving and simplifying previous work on timed I/O automaton models [36, 27, 25]. For example, the use of trajectories as first-class objects to represent the external behavior of a timed automaton, the definition of a strategy as an automaton rather than a two-player game, and the variable structure on states are all new features that were motivated by what we learned in developing the HIOA framework and that gave rise to more elegant definitions and simpler proofs for timed automata.

We intend the TIOA model to serve as a general semantic framework in which previous results for timed I/O automata [27, 29, 36, 25] and other related models [6, 28, 32, 11] can be re-cast in a style that is upwardly compatible with the new HIOA model. Limiting the communication to discrete interactions is an apt choice since the previous timed I/O automaton models also adopt this type of communication. On the other hand, by avoiding any further restrictions on the general hybrid model, we obtain an expressive model suitable for specifying complex timing behavior. For example, our model does not require variables to be either discrete or to evolve at the same rate as real-time as in some other models [6, 32]. Consequently, algorithms such as clock synchronization algorithms that use local clocks evolving at different and varying rates can be formalized naturally in our framework.

The fact that HIOAs subsume TIOAs as a special class does not eliminate the need for having a separate modeling framework for timed systems. First, having no external variables in the TIOA model gives rise to considerable simplifications in the theory. For example, proving that the composition of two timed automata is a well-defined automaton becomes simpler in the absence of external variables; no extra compatibility conditions as in the general HIOA framework are needed to obtain the desirable composition theorems for TIOAs.

Second, we believe that focusing on the TIOA model presented in this paper is compatible with our longer-term goal of developing a unified I/O automaton model that can address timing-dependent, probabilistic and general hybrid behavior in a common framework. We are planning to start out with a probabilistic model with discrete interactions only, and then extend the model to handle timing-dependent behavior, and only at later stages consider continuous interactions. It would be harder to integrate probabilistic mechanisms into the full hybrid model than it would be to integrate them into the TIOA model

presented here.

1.2 Related work

One of the widely-used formal frameworks for timed systems is that of Alur-Dill timed automata [6, 4]. An Alur-Dill automaton is a finite directed multigraph augmented with a finite set of clock variables. The semantics of such a timed automaton are defined as a state transition system in which each state consists of a location and a clock valuation. Clocks are assumed to change at the same time as real-time. The aim of facilitating automated verification based on reachability analysis seems to be the main motivation for the restrictions on the expressive power of the model. The timed automaton model presented in this paper is more expressive than the model of Alur-Dill automata. In our model, there are no finiteness assumptions and no restrictions imposed on the dynamic type of variables. We give a semantics for Alur-Dill automata by using a restricted class of our timed automata. Alur-Dill timed automata have been extensively studied with a formal language theoretic-view. Our focus, on the other hand, has been to develop a general formal framework with a well-defined notion of external behavior, parallel composition and abstraction that supports reasoning with simulation relations.

Uppaal [32, 21] is a widely-used modeling and verification tool for timed systems. It supports the description of systems as a network of Alur-Dill timed automata and enhances that model with CCS-style communication [30] along with other notions such as committed and urgent locations. Uppaal also supports communication via shared variables. Uppaal has a sophisticated model-checker that explores the whole state space of the modeled system to verify timing properties. Therefore, finiteness assumptions are built into the model to make such verification possible and the operations on clocks are restricted. For example, it is not possible to add the current value of a clock to a message as a timestamp when it is placed in a buffer. One of our plans for the near future is to work on a formal semantics for Uppaal based on some variation of our restricted hybrid I/O automaton model. There are several small mismatches due to the style of communication and notions such as committed locations but we intend to investigate to what extent we can use the communication mechanisms of our automata to model these formally. We could, for example, allow a non-empty set of external variables with restricted dynamic types and seek restrictions on the use of shared variables in Uppaal which would allow us to view these variables as external variables in the HIOA sense.

A slight generalization of Alur-Dill timed automata are the linear hybrid automata of [5]. In this model, apart from clocks that progress with rate 1, one can also use continuous variables whose derivatives are contained in some arbitrary interval. A well-known model checking tool for linear hybrid automata is HyTech [17]. The input language of HyTech can easily be translated into our TIOA model.

The timed I/O automaton modeling framework presented in this paper can be used to express models that use lower and upper time bounds on tasks or actions [29, 28].

Our framework includes an operation for adding time bounds on a subset of the actions of a timed automaton. As a result of this operation, lower bounds are transformed to appropriate preconditions for transitions and upper bounds are transformed to stopping conditions for trajectories.

An interesting timed automaton model called “Clock GTA ” has been introduced in [11]. The model was used for describing algorithms that behave in accordance with their timing constraints in certain intervals but may exhibit timing failures for some other intervals. The possibility of expressing such an ability turns out to be crucial for performance and fault-tolerance analysis for practical algorithms [11, 26]. We are interested in finding a systematic way of describing such behavior with our new timed I/O automaton model.

1.3 Paper Organization

The rest of this paper is organized as follows. Section 2 contains mathematical preliminaries. Section 3 defines notions that are useful for describing the behavior of timed systems, most importantly, trajectories and timed sequences. Section 4 defines *timed automata (TAs)*, which contain all of the structure of TIOAs except for the classification of external actions as inputs or outputs. It also defines external behavior for TAs and implementation and simulation relationships between TAs. Section 5 presents composition and hiding operations for TAs, along with operations for untiming and adding bounds that relate TIOAs to other timed automaton models. Section 6 presents definitions and results on the classification of properties of TAs as safety and liveness properties. Section 7 defines *timed I/O automata (TIOAs)* by adding an input/output classification to TAs, and extends the theory of TAs to TIOAs. It also defines special kinds of TIOAs such as progressive and receptive TIOAs. Section 8 presents compositionality results for TIOAs in general, and for the special classes of progressive and receptive TIOAs. Section 9 presents a theory for properties for TIOAs focusing on receptiveness for properties. Examples are included throughout.

2 Mathematical Preliminaries

In this section, we give basic mathematical definitions and notation that will be used as a foundation for our definitions of timed automata and timed I/O automata. These definitions involve functions, sequences, partial orders, and untimed automata.

2.1 Functions and Relations

If f is a function, then we denote the domain and range of f by $dom(f)$ and $range(f)$, respectively. If also S is a set, then we write $f \upharpoonright S$ for the restriction of f to S , that is, the function g with $dom(g) = dom(f) \cap S$ such that $g(c) = f(c)$ for each $c \in dom(g)$.

We say that two functions f and g are *compatible* if $f \upharpoonright \text{dom}(g) = g \upharpoonright \text{dom}(f)$. If f and g are compatible functions then we write $f \cup g$ for the unique function h with $\text{dom}(h) = \text{dom}(f) \cup \text{dom}(g)$ satisfying the condition: for each $c \in \text{dom}(h)$, if $c \in \text{dom}(f)$ then $h(c) = f(c)$ and if $c \in \text{dom}(g)$ then $h(c) = g(c)$. More generally, if F is a set of pairwise compatible functions then we write $\bigcup F$ for the unique function h with $\text{dom}(h) = \bigcup \{\text{dom}(f) \mid f \in F\}$ satisfying the condition: for each $f \in F$ and $c \in \text{dom}(f)$, $h(c) = f(c)$.

If f is a function whose range is a set of functions and S is a set, then we write $f \downarrow S$ for the function g with $\text{dom}(g) = \text{dom}(f)$ such that $g(c) = f(c) \upharpoonright S$ for each $c \in \text{dom}(g)$. The restriction operation \downarrow is extended to sets of functions by pointwise extension. Also, if f is a function whose range is a set of functions, all of which have a particular element d in their domain, then we write $f \downarrow d$ for the function g with $\text{dom}(g) = \text{dom}(f)$ such that $g(c) = f(c)(d)$ for each $c \in \text{dom}(g)$.

We say that two functions f and g whose ranges are sets of functions are *pointwise compatible* if for each $c \in \text{dom}(f) \cap \text{dom}(g)$, $f(c)$ and $g(c)$ are compatible. If f and g have the same domain and are pointwise compatible, then we denote by $f \dot{\cup} g$ the function h with $\text{dom}(h) = \text{dom}(f)$ such that $h(c) = f(c) \cup g(c)$ for each c .

A relation over sets X and Y is defined to be any subset of $X \times Y$. If R is a relation, then we denote the domain and range of R by $\text{dom}(R)$ and $\text{range}(R)$, respectively. A relation over X and Y is *total* over X if $\text{dom}(R) = X$. We say that a relation R over X and Y is *image-finite* if for each $x \in X$, $R(x)$ is finite.

2.2 Sequences

Let S be any set. A *sequence* over S is a function from a downward-closed subset of $\mathbb{Z}^>$ to S . Thus, the domain of a sequence is either the set of all positive integers, or is of the form $\{1, \dots, k\}$ for some k . In the first case we say that the sequence is infinite, and in the second case finite. We use $|\sigma|$ to denote the cardinality of $\text{dom}(\sigma)$, number of elements in σ . The sets of finite and infinite sequences over S are denoted by S^* and S^ω , respectively. Concatenation of a finite sequence with a finite or infinite sequence is denoted by juxtaposition. We use λ to denote the empty sequence, that is, the sequence with the empty domain. The sequence containing one element $c \in S$ is abbreviated as c . We say that a sequence σ is a *prefix* of a sequence ρ , denoted by $\sigma \leq \rho$, if $\sigma = \rho \upharpoonright \text{dom}(\sigma)$. Thus, $\sigma \leq \rho$ if either $\sigma = \rho$, or σ is finite and $\rho = \sigma\sigma'$ for some sequence σ' . If σ is a nonempty sequence then $\text{head}(\sigma)$ denotes the first element of σ and $\text{tail}(\sigma)$ denotes σ with its first element removed. Moreover, if σ is finite, then $\text{last}(\sigma)$ denotes the last element of σ and $\text{init}(\sigma)$ denotes σ with its last element removed. Let σ and σ' be sequences over S . Then σ' is a subsequence of σ provided that there exists a monotone increasing function $f : \text{dom}(\sigma') \rightarrow \text{dom}(\sigma)$ such that $\sigma'(i) = \sigma(f(i))$ for all $i \in \text{dom}(\sigma')$. If $1 \leq j_1 \leq j_2 \leq |\sigma|$, then we define $\sigma(j_1 \dots j_2)$ to be the subsequence of σ obtained by extracting the elements in positions j_1, \dots, j_2 ; that is, σ' is the subsequence obtained from function f of length $j_2 - j_1 + 1$, where $f(i) = i + j_1 - 1$ for all i .

2.3 Partial Orders

We recall some basic definitions and results regarding partial orders, and in particular, complete partial orders (cpo) from [15, 16]. A *partial order* is a set S together with a binary relation \sqsubseteq that is reflexive, antisymmetric, and transitive. In the sequel, we usually denote posets by the set S without explicit mention to the binary relation \sqsubseteq .

A subset $P \subseteq S$ is *bounded (above)* if there is a $c \in S$ such that $d \sqsubseteq c$ for each $d \in P$; in this case, c is an *upper bound* for P . A *least upper bound (lub)* for a subset $P \subseteq S$ is an upper bound c for P such that $c \leq d$ for every upper bound d for P . If P has a lub, then it is necessarily unique, and we denote it by $\sqcup P$. A subset $P \subseteq S$ is *directed* if every finite subset Q of P has an upper bound in P . A poset S is *complete*, and hence is a *complete partial order (cpo)* if every directed subset P of S has a lub in S .

We say that $P' \subseteq S$ *dominates* $P \subseteq S$, denoted by $P \sqsubseteq P'$, if for every $c \in P$ there is some $c' \in P'$ such that $c \sqsubseteq c'$. We use the following two simple lemmas, adapted from [16] [Lemmas 3.1.1 and 3.1.2].

Lemma 2.1 *If P, P' are directed subsets of a cpo S and $P \sqsubseteq P'$ then $\sqcup P \sqsubseteq \sqcup P'$.*

Lemma 2.2 *Let $P = \{c_{ij} \mid i \in I, j \in J\}$ be a doubly indexed subset of a cpo S . Let P_i denote the set $\{c_{ij} \mid j \in J\}$ for each $i \in I$. Suppose*

1. P is directed,
2. each P_i is directed with lub c_i , and
3. the set $\{c_i \mid i \in I\}$ is directed.

Then $\sqcup P = \sqcup \{c_i \mid i \in I\}$.

A finite or infinite sequence of elements, $c_0 c_1 c_2 \dots$, of a partially ordered set (S, \sqsubseteq) is called a *chain* if $c_i \sqsubseteq c_{i+1}$ for each non-final index i . We define the *limit* of the chain, $\lim_{i \rightarrow \infty} c_i$, to be the lub of the set $\{c_0, c_1, c_2, \dots\}$ if S contains such a bound; otherwise, the limit is undefined. Since a chain is a special case of a directed set, each chain of a cpo has a limit.

A function $f : S \rightarrow S'$ between posets S and S' is *monotone* if $f(c) \sqsubseteq f(d)$ whenever $c \sqsubseteq d$. If f is monotone and P is a directed set, then the set $f(P) = \{f(c) \mid c \in P\}$ is directed as well. If f is monotone and $f(\sqcup P) = \sqcup f(P)$ for every directed P , then f is said to be *continuous*.

An element c of a cpo S is *compact* if, for every directed set P such that $c \sqsubseteq \sqcup P$, there is some $d \in P$ such that $c \sqsubseteq d$. We define $K(S)$ to be the set of compact elements of S . A cpo S is *algebraic* if every $c \in S$ is the lub of the set $\{d \in K(S) \mid d \sqsubseteq c\}$. A simple example of an algebraic cpo is the set of finite or infinite sequences over some given domain, equipped with the prefix ordering. Here the compact elements are the finite sequences.

2.4 A Basic Graph Lemma

Lemma 2.3 *Let G be an infinite directed graph that satisfies the following properties.*

1. G has finitely many roots.
2. Each node of G has finite outdegree.
3. Each node of G is reachable from some root of G .

Then, there is an infinite path in G starting from some root.

Proof: The proof is an extension of König's Lemma [20]. ■

2.5 Untimed Automata

An untimed automaton (UA) \mathcal{A} is defined as a tuple $(Q, \Theta, E, H, \mathcal{D})$ which consists of:

- A set Q of *states*.
- A non-empty set $\Theta \subseteq Q$ of *start states*.
- A set E of *external actions* and a set H of *internal actions*, disjoint from each other. We write $A \triangleq E \cup H$.
- A set $\mathcal{D} \subseteq Q \times A \times Q$ of *discrete transitions*.

An *execution fragment* of an untimed automaton \mathcal{A} is either a finite sequence $s_0 a_1 s_1 a_2 \cdots a_n s_n$ or an infinite sequence $s_0 a_1 s_1 a_2 \cdots$, of alternating states and actions of \mathcal{A} such that (s_k, a_{k+1}, s_{k+1}) is in \mathcal{D} for every non-final index k where $k \geq 0$. An *execution fragment* beginning with a start state is called an *execution*. If σ is an execution fragment of \mathcal{A} , then the trace of σ is defined as the subsequence of σ consisting of all the external actions.

If σ is a finite execution fragment of an automaton \mathcal{A} and σ' is any execution fragment of \mathcal{A} that begins with the last state of σ , then we write $\sigma \frown \sigma'$ to represent the sequence obtained by concatenating σ and σ' , eliminating the duplicate occurrence of the last state of σ . It is easy to see that, $\sigma \frown \sigma'$ is also an execution fragment of \mathcal{A} .

3 Describing Timed System Behavior

In this section, we give basic definitions that are useful for describing discrete and continuous changes to the system's state. The key notions are *static* and *dynamic types* for variables, *trajectories*, and *hybrid sequences*. Most of the material in this section comes from the paper on the HIOA modeling framework [22]. The reader is referred to [22] for the proofs that are not included here.

3.1 Time

Throughout this paper, we fix a *time axis* \mathbb{T} , which is a subgroup of $(\mathbb{R}, +)$, the real numbers with addition. We assume that every infinite, monotone, bounded sequence of elements of \mathbb{T} has a limit in \mathbb{T} . The reader may find it convenient to think of \mathbb{T} as the set \mathbb{R} of real numbers, but the set \mathbb{Z} of integers and the singleton set $\{0\}$ are also examples of allowed time axes. We define $\mathbb{T}^{\geq 0} \triangleq \{t \in \mathbb{T} \mid t \geq 0\}$.

An *interval* J is a nonempty, convex subset of \mathbb{T} . We denote intervals as usual: $[t_1, t_2] = \{t \in \mathbb{T} \mid t_1 \leq t \leq t_2\}$, $[t_1, t_2) = \{t \in \mathbb{T} \mid t_1 \leq t < t_2\}$ etc. An interval J is *left-closed* (*right-closed*) if it has a minimum (resp., maximum) element, and *left-open* (*right-open*) otherwise. It is *closed* if it is both left-closed and right-closed. We write $\min(J)$ and $\max(J)$ for the minimum and maximum elements, respectively, of an interval J (if they exist), and $\inf(J)$ and $\sup(J)$ for the infimum and supremum, respectively, of J in $\mathbb{R} \cup \{-\infty, \infty\}$. For $K \subseteq \mathbb{T}$ and $t \in \mathbb{T}$, we define $K + t \triangleq \{t' + t \mid t' \in K\}$. Similarly, for a function f with domain K , we define $f + t$ to be the function with domain $K + t$ satisfying, for each $t' \in K + t$, $(f + t)(t') = f(t' - t)$.

In some definitions and theorems in the paper where we use \mathbb{R} as the time domain we assume that the relation \leq on \mathbb{R} extends to a relation on $\mathbb{R} \cup \{\infty\}$ such that $\infty \leq \infty$ and for all $t \in \mathbb{R}$, $t < \infty$.

3.2 Static and Dynamic Types

We assume a universal set V of *variables*. A variable represents a location within the state of a system. For each variable v , we assume both a *static type*, which gives the set of values it may take on, and a *dynamic type*, which gives the set of trajectories it may follow. Formally, for each variable v we assume the following:

- $type(v)$, the *static type* of v . This is a nonempty set of values.
- $dtype(v)$, the *dynamic type* of v . This is a set of functions from left-closed intervals of \mathbb{T} to $type(v)$ that satisfies the following properties:
 1. (*Closure under time shift*)
For each $f \in dtype(v)$ and $t \in \mathbb{T}$, $f + t \in dtype(v)$.
 2. (*Closure under subinterval*)
For each $f \in dtype(v)$ and each left-closed interval $J \subseteq dom(f)$, $f \upharpoonright J \in dtype(v)$.
 3. (*Closure under pasting*)
Let $f_0 f_1 f_2, \dots$ be a sequence of functions in $dtype(v)$ such that, for each index i such that f_i is not the final function in the sequence, $dom(f_i)$ is right-closed and $\max(dom(f_i)) = \min(dom(f_{i+1}))$. Then the function f defined by $f(t) \triangleq f_i(t)$, where i is the smallest index such that $t \in dom(f_i)$, is in $dtype(v)$.

Example 3.1 (Discrete variables) Let v be any variable and let *Constant* be the set of constant functions from a left-closed interval of \mathbb{T} to $\text{type}(v)$. Then *Constant* is closed under time shift and subinterval. If the dynamic type of v is obtained by closing *Constant* under the pasting operation, then v is called a *discrete* variable. This is essentially the same as the definition of a discrete variable in [28]. ■

Example 3.2 (Analog variables) Assume that $\mathbb{T} = \mathbb{R}$. Let v be any variable whose static type is an interval of \mathbb{R} and *Continuous* be the set of continuous functions from a left-closed interval of \mathbb{T} to $\text{type}(v)$. Then *Continuous* is closed under time shift and subinterval. If the dynamic type of v is obtained by closing *Continuous* under the pasting operation, then v is called an *analog* variable. ■

Example 3.3 (Standard real-valued function classes) If we take $\mathbb{T} = \mathbb{R}$ and $\text{type}(v) = \mathbb{R}$, then other examples of dynamic types can be obtained by taking the pasting closure of standard function classes from real analysis, the set of differentiable functions, the set of functions that are differentiable k times (for any k), the set of smooth functions, the set of integrable functions, the set of L^p functions (for any p), the set of measurable locally essentially bounded functions [37], or the set of all functions. ■

Standard function classes are closed under time shift and subinterval, but not under pasting. A natural way of defining a dynamic type is as the pasting closure of a class of functions that is closed under time shift and subinterval. In such a case, it follows that the new class is closed under all three operations.

3.3 Trajectories

In this subsection, we define the notion of a *trajectory*, define operations on trajectories, and prove simple properties of trajectories and their operations. A trajectory is used to model the evolution of a collection of variables over an interval of time.

3.3.1 Basic Definitions

Let V be a set of variables, that is, a subset of \mathbb{V} . A *valuation* \mathbf{v} for V is a function that associates with each variable $v \in V$ a value in $\text{type}(v)$. We write $\text{val}(V)$ for the set of valuations for V . Let J be a left-closed interval of \mathbb{T} with left endpoint equal to 0. Then a *J-trajectory* for V is a function $\tau : J \rightarrow \text{val}(V)$, such that for each $v \in V$, $\tau \downarrow v \in \text{dtype}(v)$. A *trajectory* for V is a J -trajectory for V , for any J . We write $\text{trajs}(V)$ for the set of all trajectories for V .

A trajectory for V with domain $[0, 0]$ is called a *point* trajectory for V . If \mathbf{v} is a valuation for V then $\wp(\mathbf{v})$ denotes the point trajectory for V that maps 0 to \mathbf{v} . We say

that a J -trajectory is *finite* if J is a finite interval, *closed* if J is a (finite) closed interval, *open* if J is a right-open interval, and *full* if $J = \mathbb{T}^{\geq 0}$. If T is a set of trajectories, then $\text{finite}(T)$, $\text{closed}(T)$, $\text{open}(T)$, and $\text{full}(T)$ denote the subsets of T consisting of all the finite, closed, open, and full trajectories in T , respectively.

If τ is a trajectory then $\tau.\text{ltime}$, the *limit time* of τ , is the supremum of $\text{dom}(\tau)$. We define $\tau.\text{fval}$, the *first valuation* of τ , to be $\tau(0)$, and if τ is closed, we define $\tau.\text{lval}$, the *last valuation* of τ , to be $\tau(\tau.\text{ltime})$. For τ a trajectory and $t \in \mathbb{T}^{\geq 0}$, we define

$$\begin{aligned}\tau \sqsubseteq t &\triangleq \tau \upharpoonright [0, t], \\ \tau \triangleleft t &\triangleq \tau \upharpoonright [0, t), \\ \tau \sqsupseteq t &\triangleq (\tau \upharpoonright [t, \infty)) - t.\end{aligned}$$

Note that, since dynamic types are closed under time shift and subintervals, the result of applying the above operations is always a trajectory, except when the result is a function with an empty domain. By convention, we also write $\tau \sqsubseteq \infty \triangleq \tau$ and $\tau \triangleleft \infty \triangleq \tau$.

3.3.2 Prefix Ordering

Trajectory τ is a *prefix* of trajectory v , denoted by $\tau \leq v$, if τ can be obtained by restricting v to a subset of its domain. Formally, if τ and v are trajectories for V , then $\tau \leq v$ iff $\tau = v \upharpoonright \text{dom}(\tau)$. Alternatively, $\tau \leq v$ iff there exists a $t \in \mathbb{T}^{\geq 0} \cup \{\infty\}$ such that $\tau = v \sqsubseteq t$ or $\tau = v \triangleleft t$. If $\tau \leq v$ then clearly $\text{dom}(\tau) \subseteq \text{dom}(v)$. If T is a set of trajectories for V , then $\text{pref}(T)$ denotes the *prefix closure* of T , defined by:

$$\text{pref}(T) \triangleq \{\tau \in \text{trajs}(V) \mid \exists v \in T : \tau \leq v\}.$$

We say that T is *prefix closed* if $T = \text{pref}(T)$.

The following lemma gives a simple domain-theoretic characterization of the set of trajectories over a given set V of variables:

Lemma 3.4 *Let V be a set of variables. The set $\text{trajs}(V)$ of trajectories for V , together with the prefix ordering \leq , is an algebraic cpo. Its compact elements are the closed trajectories.*

3.3.3 Concatenation

The concatenation of two trajectories is obtained by taking the union of the first trajectory and the function obtained by shifting the domain of the second trajectory until the start time agrees with the limit time of the first trajectory; the last valuation of the first trajectory, which may not be the same as the first valuation of the second trajectory, is

the one that appears in the concatenation. Formally, suppose τ and τ' are trajectories for V , with τ closed. Then the *concatenation* $\tau \frown \tau'$ is the function given by

$$\tau \frown \tau' \triangleq \tau \cup (\tau' \upharpoonright (0, \infty) + \tau.ltime).$$

Because dynamic types are closed under time shift and pasting, it follows that $\tau \frown \tau'$ is a trajectory for V . Observe that $\tau \frown \tau'$ is finite (resp., closed, full) if and only if τ' is finite (resp., closed, full). Observe also that concatenation is associative.

The following lemma, which is easy to prove, shows the close connection between concatenation and the prefix ordering.

Lemma 3.5 *Let τ and v be trajectories for V with τ closed. Then*

$$\tau \leq v \iff \exists \tau' : v = \tau \frown \tau'.$$

Note that if $\tau \leq v$, then the trajectory τ' such that $v = \tau \frown \tau'$ is unique except that it has an arbitrary value for $\tau'.fval$. Note also that the “ \Leftarrow ” implication in Lemma 3.5 would not hold if the first valuation of the second argument, rather than the last valuation of the first argument, were used in the concatenation.

We extend the definition of concatenation to any (finite or countably infinite) number of arguments. Let $\tau_0 \tau_1 \tau_2 \dots$ be a (finite or infinite) sequence of trajectories such that τ_i is closed for each nonfinal index i . Define trajectories $\tau'_0, \tau'_1, \tau'_2, \dots$ inductively by

$$\begin{aligned} \tau'_0 &\triangleq \tau_0, \\ \tau'_{i+1} &\triangleq \tau'_i \frown \tau_{i+1} \text{ for nonfinal } i. \end{aligned}$$

Lemma 3.5 implies that for each nonfinal i , $\tau'_i \leq \tau'_{i+1}$. We define the *concatenation* $\tau_0 \frown \tau_1 \frown \tau_2 \dots$ to be the limit of the chain $\tau'_0 \tau'_1 \tau'_2 \dots$; existence of this limit follows from Lemma 3.4.

3.4 Hybrid Sequences

In this subsection, we introduce the notion of a *hybrid sequence*, which is used to model a combination of changes that occur instantaneously and changes that occur over intervals of time. Our definition is parameterized by a set A of *actions*, which are used to model instantaneous changes and instantaneous synchronizations with the environment, and a set V of *variables*, which are used to model changes over intervals of time. We also define some special kinds of hybrid sequences and some operations on hybrid sequences, and give basic properties.

3.4.1 Basic Definitions

Fix a set A of actions and a set V of variables. An (A, V) -sequence is a finite or infinite alternating sequence $\alpha = \tau_0 a_1 \tau_1 a_2 \tau_2 \dots$, where

1. each τ_i is a trajectory in $\text{trajs}(V)$,
2. each a_i is an action in A ,
3. if α is a finite sequence then it ends with a trajectory, and
4. if τ_i is not the last trajectory in α then $\text{dom}(\tau_i)$ is closed.

A *hybrid sequence* is an (A, V) -sequence for some A and V .

Since the trajectories in a hybrid sequence can be point trajectories our notion of hybrid sequence allows a sequence of discrete actions to occur at the same real time, with corresponding changes of variable values. An alternative approach is described in [34], where state changes at a single real time are modeled using a notion of “superdense time”. Specifically, hybrid behavior is modeled in [34] using functions from an extended time domain, which includes countably many elements for each real time, to states.

If α is a hybrid sequence, with notation as above, then we define the *limit time* of α , $\alpha.\text{ltime}$, to be $\sum_i \tau_i.\text{ltime}$. A hybrid sequence α is defined to be:

- *time-bounded* if $\alpha.\text{ltime}$ is finite.
- *admissible* if $\alpha.\text{ltime} = \infty$.
- *closed* if α is a finite sequence and the domain of its final trajectory is a closed interval.
- *Zeno* if α is neither closed nor admissible, that is, if α is time-bounded and is either an infinite sequence, or else a finite sequence ending with a trajectory whose domain is right-open.
- *non-Zeno* if α is not Zeno.

For any hybrid sequence α , we define the *first valuation* of α , $\alpha.\text{fval}$, to be $\text{head}(\alpha).\text{fval}$. Also, if α is closed, we define the *last valuation* of α , $\alpha.\text{lval}$, to be $\text{last}(\alpha).\text{lval}$, that is, the last valuation in the final trajectory of α .

If α is a hybrid sequence of the form $\tau_0 a_1 \tau_1 a_2 \tau_2 \dots$, we use $\text{actions}(\alpha)$ to denote the sequence $a_1 a_2 a_3 \dots$, which is obtained by discarding the trajectories in α .

If α is a closed (A, V) -sequence, where $V = \emptyset$ and $\beta \in \text{trajs}(\emptyset)$, we call $\alpha \frown \beta$ a *time-extension* of α .

3.4.2 Prefix Ordering

We say that (A, V) -sequence $\alpha = \tau_0 a_1 \tau_1 \dots$ is a *prefix* of (A, V) -sequence $\beta = v_0 b_1 v_1 \dots$, denoted by $\alpha \leq \beta$, provided that (at least) one of the following holds:

1. $\alpha = \beta$.
2. α is a finite sequence ending in some τ_k ; $\tau_i = v_i$ and $a_{i+1} = b_{i+1}$ for every i , $0 \leq i < k$; and $\tau_k \leq v_k$.

Like the set of trajectories over V , the set of (A, V) -sequences is an algebraic cpo:

Lemma 3.6 *Let V be a set of variables and A a set of actions. The set of (A, V) -sequences, together with the prefix ordering \leq , is an algebraic cpo. Its compact elements are the closed (A, V) -sequences.*

3.4.3 Concatenation

Suppose α and α' are (A, V) -sequences with α closed. Then the *concatenation* $\alpha \frown \alpha'$ is the (A, V) -sequence given by

$$\alpha \frown \alpha' \triangleq \text{init}(\alpha) (\text{last}(\alpha) \frown \text{head}(\alpha')) \text{tail}(\alpha').$$

(Here, *init*, *last*, *head* and *tail* are ordinary sequence operations.)

Lemma 3.7 *Let α and β be (A, V) -sequences with α closed. Then*

$$\alpha \leq \beta \Leftrightarrow \exists \alpha' : \beta = \alpha \frown \alpha'.$$

Note that if $\alpha \leq \beta$, then the (A, V) -sequence α' such that $\beta = \alpha \frown \alpha'$ is unique except that it has an arbitrary value in $\text{val}(V)$ for $\alpha'.\text{fval}$.

As we did for trajectories, we extend the concatenation definition for (A, V) -sequences to any finite or infinite number of arguments. Let $\alpha_0 \alpha_1 \dots$ be a finite or infinite sequence of (A, V) -sequences such that α_i is closed for each nonfinal index i . Define (A, V) -sequences $\alpha'_0, \alpha'_1, \dots$ inductively by

$$\begin{aligned} \alpha'_0 &\triangleq \alpha_0, \\ \alpha'_{i+1} &\triangleq \alpha'_i \frown \alpha_{i+1} \text{ for nonfinal } i. \end{aligned}$$

Lemma 3.7 implies that for each nonfinal i , $\alpha'_i \leq \alpha'_{i+1}$. We define the *concatenation* $\alpha_0 \frown \alpha_1 \dots$ to be the limit of the chain $\alpha'_0 \alpha'_1 \dots$; existence of this limit is ensured by Lemma 3.6.

3.4.4 Restriction

Let A and A' be sets of actions and let V and V' be sets of variables. The (A', V') -restriction of an (A, V) -sequence α , denoted by $\alpha \upharpoonright (A', V')$, is obtained by first projecting all trajectories of α on the variables in V' , then removing the actions not in A' , and finally concatenating all adjacent trajectories. Formally, we define the (A', V') -restriction first for closed (A, V) -sequences and then extend the definition to arbitrary (A, V) -sequences using a limit construction. The definition for closed (A, V) -sequences is by induction on the length of those sequences:

$$\begin{aligned} \tau \upharpoonright (A', V') &= \tau \downarrow V' \text{ if } \tau \text{ is a single trajectory,} \\ \alpha a \tau \upharpoonright (A', V') &= \begin{cases} (\alpha \upharpoonright (A', V')) a (\tau \downarrow V') & \text{if } a \in A', \\ (\alpha \upharpoonright (A', V')) \frown (\tau \downarrow V') & \text{otherwise.} \end{cases} \end{aligned}$$

It is easy to see that the restriction operator is monotone on the set of closed (A, V) -sequences. Hence, if we apply this operation to a directed set, the result is again a directed set. Together with Lemma 3.6, this allows us to extend the definition of restriction to arbitrary (A, V) -sequences by:

$$\alpha \upharpoonright (A', V') = \sqcup \{ \beta \upharpoonright (A', V') \mid \beta \text{ is a closed prefix of } \alpha \}.$$

Lemma 3.8 (A', V') -restriction is a continuous operation.

Lemma 3.9 $(\alpha_0 \frown \alpha_1 \frown \dots) \upharpoonright (A, V) = \alpha_0 \upharpoonright (A, V) \frown \alpha_1 \upharpoonright (A, V) \frown \dots$

Lemma 3.10 $(\alpha \upharpoonright (A, V)) \upharpoonright (A', V') = \alpha \upharpoonright (A \cap A', V \cap V')$.

Lemma 3.11 Let α be a hybrid sequence A a set of actions and V a set of variables.

1. α is time-bounded if and only if $\alpha \upharpoonright (A, V)$ is time-bounded.
2. α is admissible if and only if $\alpha \upharpoonright (A, V)$ is admissible.
3. If α is closed then $\alpha \upharpoonright (A, V)$ is closed.
4. If α is non-Zeno then $\alpha \upharpoonright (A, V)$ is non-Zeno.

Example 3.12 (A Zeno execution with a closed (A, V) -restriction) In order to understand why we have an implication in only one direction in items 3 and 4, consider the Zeno sequence α of the form $\wp(\mathbf{v}) a \wp(\mathbf{v}) a \wp(\mathbf{v}) \dots$. Let A be a set such that $a \notin A$ and let V consist of the variables in $\text{dom}(\mathbf{v})$. Obviously, $\alpha \upharpoonright (A, V)$, which is $\wp(\mathbf{v})$, is closed, and hence also non-Zeno. This shows that the fact that $\alpha \upharpoonright (A, V)$ is closed (resp., non-Zeno) does not imply that α is closed (resp., non-Zeno). ■

4 Timed Automata

In this section, as a preliminary step toward defining timed I/O automata, we define a slightly more general *timed automaton* model. In timed automata, actions are classified as external or internal, but external actions are not further classified as input or output; the input/output distinction is added in Section 7. We define how timed automata execute and define implementation and simulation relations between timed automata.

4.1 Definition of Timed Automata

A timed automaton is a state machine whose states are divided into *variables*, and that has a set of discrete *actions*, some of which may be internal and some external. The state of a timed automaton may change in two ways: by *discrete transitions*, which change the state atomically, and by *trajectories*, which describe the evolution of the state over intervals of time. The discrete transitions are labeled with actions; this will allow us to synchronize the transitions of different timed automata when we compose them in parallel. The evolution described by a trajectory may be described by continuous or discontinuous functions.

Formally, a *timed automaton (TA)* $\mathcal{A} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$ consists of:

- A set X of *internal variables*.
- A set $Q \subseteq \text{val}(X)$ of *states*.
- A nonempty set $\Theta \subseteq Q$ of *start states*.
- A set E of *external actions* and a set H of *internal actions*, disjoint from each other. We write $A \triangleq E \cup H$.
- A set $\mathcal{D} \subseteq Q \times A \times Q$ of *discrete transitions*.
We use $\mathbf{x} \xrightarrow{a}_{\mathcal{A}} \mathbf{x}'$ as shorthand for $(\mathbf{x}, a, \mathbf{x}') \in \mathcal{D}$. Here and elsewhere, we sometimes drop the subscript and write $\mathbf{x} \xrightarrow{a} \mathbf{x}'$, when we think \mathcal{A} should be clear from the context. We say that a is *enabled* in \mathbf{x} if $\mathbf{x} \xrightarrow{a} \mathbf{x}'$ for some \mathbf{x}' . We say that a set C of actions is *enabled* in a state \mathbf{x} if some action in C is enabled in \mathbf{x} .
- A set \mathcal{T} of trajectories for X such that $\tau(t) \in Q$ for every $\tau \in \mathcal{T}$ and $t \in \text{dom}(\tau)$. Given a trajectory $\tau \in \mathcal{T}$ we denote $\tau.\text{fval}$ by $\tau.\text{fstate}$ and, if τ is closed, we denote $\tau.\text{lval}$ by $\tau.\text{lstate}$. When $\tau.\text{fstate} = \mathbf{x}$ and $\tau.\text{lstate} = \mathbf{x}'$, we sometimes write $\mathbf{x} \xrightarrow{\tau}_{\mathcal{A}} \mathbf{x}'$. We require that the following axioms hold:

T0 (*Existence of point trajectories*)
If $\mathbf{x} \in Q$ then $\wp(\mathbf{x}) \in \mathcal{T}$.

T1 (*Prefix closure*)

For every $\tau \in \mathcal{T}$ and every $\tau' \leq \tau$, $\tau' \in \mathcal{T}$.

T2 (*Suffix closure*)

For every $\tau \in \mathcal{T}$ and every $t \in \text{dom}(\tau)$, $\tau \geq t \in \mathcal{T}$.

T3 (*Concatenation closure*)

Let $\tau_0 \tau_1 \tau_2 \dots$ be a sequence of trajectories in \mathcal{T} such that, for each nonfinal index i , τ_i is closed and $\tau_i.\text{lstate} = \tau_{i+1}.\text{fstate}$. Then $\tau_0 \frown \tau_1 \frown \tau_2 \dots \in \mathcal{T}$.

Thus, a timed automaton is essentially a hybrid automaton in the sense of [22] in which W , the set of external variables, is empty. (The only difference is the addition of the axiom **T0**, which does not affect any of the results of [22].) This definition differs from previous definitions of timed automata [25, 36] in two major respects. First, the states are structured using variables, which have dynamic types with specific closure properties. The variable structure is convenient for writing specifications and the dynamic types are useful in analyzing continuous evolution of the state. Second, the set of trajectories is defined as an explicit component of an automaton. In the previous definitions, time-passage was represented by special time-passage actions and trajectories were defined implicitly, as auxiliary functions describing the effects of time-passage actions on states.

Notation: We often denote the components of a TA \mathcal{A} by $X_{\mathcal{A}}$, $Q_{\mathcal{A}}$, $\Theta_{\mathcal{A}}$, $E_{\mathcal{A}}$, etc., and the components of a TA \mathcal{A}_i by X_i , Q_i , Θ_i , E_i , etc. We sometimes omit these subscripts, where no confusion seems likely. In examples we typically specify sets of trajectories using differential and algebraic equations and inclusions. Below we explain a few notational conventions that help us in doing this. Suppose the time domain \mathbb{T} is \mathbb{R} , τ is a (fixed) trajectory over some set of variables V , and $v \in V$. With some abuse of notation, we use the variable name v to denote the function $\tau \downarrow v$ in $\text{dom}(\tau) \rightarrow \text{type}(v)$, which gives the value of v at all times during trajectory τ . Similarly, we view any expression e containing variables from V as a function with domain $\text{dom}(\tau)$. Suppose that v is a variable and e is a real-valued expression containing variables from V . Using these conventions we can say, for example, that τ satisfies the algebraic equation

$$v = e$$

which means that, for every $t \in \text{dom}(\tau)$, $v(t) = e(t)$, that is, the constraint on the variables expressed by the equation $v = e$ holds for each state on trajectory τ . Now suppose also that e , when viewed as a function, is integrable. Then we say that τ satisfies

$$d(v) = e$$

if, for every $t \in \text{dom}(\tau)$, $v(t) = v(0) + \int_0^t e(t') dt'$. Equivalently, for every $t_1, t_2 \in \text{dom}(\tau)$ such that $t_1 \leq t_2$, $v(t_2) = v(t_1) + \int_{t_1}^{t_2} e(t') dt'$. Note that this interpretation of the differential

equation makes sense even at points where v is not differentiable. A similar interpretation of differential equations is used by Polderman and Willems [35], who call functions defined in this way “weak solutions”.

We generalize this notation to handle inequalities as well as equalities. Suppose that v is a variable and e is a real-valued expression containing variables from V . The inequality

$$e \leq v$$

means that, for every $t \in \text{dom}(\tau)$, $e(t) \leq v(t)$. That is, the constraint expressed by the inequality $e \leq v$ holds for each state of trajectory τ . Similarly, the inequality

$$v \leq e$$

means that, for every $t \in \text{dom}(\tau)$, $v(t) \leq e(t)$. Now suppose that e is integrable when viewed as a function. Then we say that τ satisfies

$$e \leq d(v)$$

if, for every $t_1, t_2 \in \text{dom}(\tau)$ such that $t_1 \leq t_2$, $v(t_1) + \int_{t_1}^{t_2} e(t') dt' \leq v(t_2)$, and τ satisfies

$$d(v) \leq e$$

if, for every $t_1, t_2 \in \text{dom}(\tau)$ such that $t_1 \leq t_2$, $v(t_2) \leq v(t_1) + \int_{t_1}^{t_2} e(t') dt'$.

Conventions for automata specifications: In all the examples of this paper we assume that $\mathbb{T} = \mathbb{R}$. The static type of a variable v is always written explicitly. Discrete and analog variables are designated using the keywords **discrete** and **analog** respectively. The definition of what it means for a variable to be discrete or analog is given in Examples 3.1 and 3.2. Although timed automata may contain variables that are neither discrete nor analog, none of our examples use such variables.

The transitions are specified in precondition-effect style. A **precondition** clause specifies the enabling condition for an action. The **effect** clause contains a list of statements that specify the effect of performing that action on the state. All the statements in an effect clause are assumed to be executed sequentially in a single indivisible step. The absence of a specified precondition for an action means that the action is always enabled and the absence of a specified effect means that performing the action does not change the state.

The trajectories are specified by using a variation of the language presented in [31]. A **satisfies** clause contains a list of predicates that must be satisfied by all the trajectories. This clause is followed by a **stops when** clause. If the predicate in this clause becomes

Automaton $TimedChannel(b, M)$ **where** $b \in \mathbb{R}^+$

Variables X : **discrete** $queue$, a finite sequence of elements of $M \times \mathbb{R}$ **initially** empty
 analog $now \in \mathbb{R}$ **initially** 0

States Q : $val(X)$

Actions A : **external** $send(m), receive(m)$ **where** $m \in M$

Transitions \mathcal{D} : **external** $send(m)$
 effect
 add $(m, now + b)$ to $queue$

 external $receive(m)$
 precondition
 $\exists u. (m, u)$ is first element of $queue$
 effect
 remove first element of $queue$

Trajectories \mathcal{T} : **satisfies**
 constant($queue$)
 $\mathbf{d}(now) = 1$
 stops when
 $\exists(m, u) \in queue. (now = u)$

Figure 1: Time-bounded channel

true at a point t in time, then t must be the limit time of the trajectory. When there is no stopping condition for trajectories we omit the **stops when** clause. We write $\mathbf{d}(v) = e$ for $d(v) = e$, $\mathbf{d}(v) \leq e$ for $d(v) \leq e$ and $e \leq \mathbf{d}(v)$ for $e \leq d(v)$. If the value of a variable is constant throughout a trajectory then we write **constant**(v). If the evolution of a variable follows a continuous function throughout a trajectory then we write **continuous**(v).

Example 4.1 (Time-bounded channel) The automaton in Figure 2 is the specification of a reliable FIFO channel that delivers its messages within a certain time bound, represented by the automaton parameter b of type \mathbb{R}^+ . The other automaton parameter M is an arbitrary type parameter that represents the type of messages communicated by the channel.

The discrete variable $queue$ is used to hold pairs consisting of a message that has been sent and its delivery deadline. The analog variable now is used to describe real time.

Every $send(m)$ transition adds to the queue a new pair whose first component is m and whose second component is the deadline $now + b$. A $receive(m)$ transition can occur only when m is the first message in the queue and it results in the removal of the first message from the queue.

```

automaton TimedChannel(b: Real, M)
  signature
    external send(m), receive(m) where m ∈ M
  states
    queue: Queue[M] := {}
    now: Real := 0
    initially b > 0
  transitions
    external send(m)
      eff
        queue := append((m, now+b), queue)
    external receive(m)
      pre
        \exists u (m, u) = head(queue)
      eff
        queue := tail(queue)
  trajectories
    stop when \exists (m, u) ∈ queue (now = u)
    evolve
      d(now)=1

```

Figure 2: Time-bounded channel

The trajectory specification shows that the discrete variable *queue* is kept constant by trajectories and that the variable *now* increases with rate 1, that is, at the same rate as real time. The stopping condition implies that, within a trajectory, time cannot pass beyond the point where *now* becomes equal to the delivery deadline of some message in the queue. ■

Example 4.2 (Periodic sending process) The automaton in Figure 3 is the specification of a process that sends messages periodically, every u time units, where u is an automaton parameter of type $\mathbb{R}^{\geq 0}$. The type parameter M represents the type of the messages sent by the process.

The analog variable *clock* is a timer whose value records the amount of time that has elapsed since it was last reset to 0. A $send(m)$ transition can occur only when $clock = u$, and it causes *clock* to be reset. The trajectory specification says that *clock* increases at the same rate as real time and time cannot pass beyond the point where $clock = u$. ■

Example 4.3 (Periodic sending process with failures) The specification of the $PeriodicSend(u, M)$ process from Example 4.2 does not model failures. We now consider

Automaton $PeriodicSend(u, M)$ where $u \in \mathbb{R}^{\geq 0}$

Variables X : **analog** $clock \in \mathbb{R}$ **initially** 0

States Q : $val(X)$

Actions A : **external** $send(m)$ **where** $m \in M$

Transitions \mathcal{D} : **external** $send(m)$
 precondition
 $clock = u$
 effect
 $clock := 0$

Trajectories \mathcal{T} : **satisfies**
 $d(clock) = 1$
 stops when
 $clock = u$

Figure 3: Periodic sending process

a variant of $PeriodicSend(u, M)$ where the process may fail and stop doing any discrete actions. The specification of this new automaton is given in Figure 4.

The discrete variable *failed* in automaton $PeriodicSend2$ is a boolean flag that records whether the process is failed. It is initialized to `false` and is set to `true` when a *fail* action occurs. The trajectory specification of $PeriodicSend2$ shows that time can advance without any bound when the process is failed.

■

Example 4.4 (Timeout process) The automaton $Timeout(u, M)$ in Figure 5 is the specification of a process that awaits the receipt of a message from another process. If u time units elapse without such a message arriving, $Timeout(u, M)$ performs a *timeout* action, thereby “suspecting” the other process. When a message arrives it “unsuspects” the other process. $Timeout(u, M)$ may suspect and unsuspect repeatedly.

The discrete variable *suspected* is a flag that shows whether $Timeout(u, M)$ suspects that the other process is failed. The variable *clock* is a timer that records the amount of time that has elapsed since the receipt of the last message.

A *receive(m)* transition can occur at any time; this causes the variable *clock* to be reset and the flag *suspected* to be set to `false`. If *clock* reaches u before the arrival of a message then the *timeout* action becomes enabled. The process sets *suspected* to `true` as a result of a *timeout*.

The discrete variable *suspected* remains constant throughout each trajectory. The trajectory specification also shows that *clock* increases at the same rate as real time and,

Automaton *PeriodicSend2*(u, M) **where** $u \in \mathbb{R}^+$

Variables X : **discrete** $failed \in Bool$ **initially** $false$
 analog $clock \in \mathbb{R}$ **initially** 0

States Q : $val(X)$

Actions A : **external** $send(m)$ **where** $m \in M$
 external $fail$

Transitions \mathcal{D} : **external** $send(m)$
 precondition
 $\neg failed$
 $clock = u$
 effect
 $clock := 0$

external $fail$
 effect
 $failed := true$

Trajectories \mathcal{T} : **satisfies**
 constant($failed$)
 $d(clock) = 1$
 stops when
 $\neg failed$ and $clock = u$

Figure 4: Periodic sending process with failures

if $suspected = false$, then time cannot go beyond the point where $clock = u$. Note that if $suspected = true$, there is no restriction on the amount of time that can elapse. ■

Example 4.5 (Fischer’s mutual exclusion algorithm) The automaton presented in Figures 6 and 7 is the specification of a shared memory mutual exclusion algorithm which uses a single shared variable that can be read and written by all the participants. The automaton parameters u_{set} and l_{check} represent upper and lower time bounds for the set_i and $check_i$ actions respectively. We assume that $u_{set} < l_{check}$. The parameter I represents the set of indices of processes that participate in the algorithm and is required to be finite.

The shared variable x can be assigned any value in I or the special value \perp . If a process is in the critical region, then the variable x contains the index of that process. If all users are in the remainder region, then the variable x contains the value \perp . The array variable pc records the program counters of all processes. The array variable $lastset$ keeps track of the deadlines by which the processes’ set actions must occur. Similarly, the array variable $firstcheck$ keeps track of the earliest time the processes’ $check$ actions may occur.

Automaton $Timeout(u, M)$ **where** $u \in \mathbb{R}^+$

Variables X : **discrete** $suspected \in Bool$ **initially** $false$
 analog $clock \in \mathbb{R}$ **initially** 0

States Q : $val(X)$

Actions A : **external** $receive(m)$ **where** $m \in M$
 external $timeout$

Transitions \mathcal{D} : **external** $receive(m)$
 effect
 $clock := 0$
 $suspected := false$

external $timeout$
 precondition
 $\neg suspected$
 $clock = u$
 effect
 $suspected := true$

Trajectories \mathcal{T} : **satisfies**
 constant($suspected$)
 $\mathbf{d}(clock) = 1$
 stops when
 $clock = u$ **and** $\neg suspected$

Figure 5: Timeout

The analog variable now models real time.

The transition definitions for external actions try_i , $test_i$, $crit_i$, $exit_i$ are straightforward. When a process performs one of these actions, its program counter is updated to record the region entered by the process. The most interesting transition definitions are $test_i$, set_i and $check_i$ since they are the ones that involve timing constraints of the algorithm. When a process i performs a $test$ action and observes x to be \perp , it sets $lastset[i]$ to $now + u_{set}$. This sets the deadline for the performance of the set_i action. Note that this deadline is enforced through the stopping condition in the trajectory specification. The transition set_i sets $firstcheck[i]$ to $now + l_{check}$. The value of $firstcheck[i]$ determines the earliest time $check_i$ may occur. The $check_i$ action is enabled only when the current time has at least this value.

The trajectory specification says that the values of discrete variables are kept constant by trajectories. The stopping condition implies that if the value of now reaches the value of $lastset[i]$ for some process i at some point in time, then that point must be the limit time of the trajectory.

■

Type $PcValue$ = enumeration of $rem, test, set, check, leavetry, crit, leaveexit$

Automaton $FischerME(u_{set}, l_{check}, I)$ where $u_{set} \in \mathbb{R}^{\geq 0}, l_{check} \in \mathbb{R}^{\geq 0}, u_{set} < l_{check}$

Variables X : **discrete** $x \in I \cup \{\perp\}$ **initially** \perp
 discrete pc , an array of elements of $PcValue$ indexed by I
 initially $\forall i \in I. pc[i] = rem$
 discrete $lastset$, an array of elements of $\mathbb{R} \cup \{\infty\}$ indexed by I
 initially $\forall i \in I. lastset[i] = \infty$
 discrete $firstcheck$, an array of elements of type \mathbb{R}
 initially $\forall i \in I. firstcheck[i] = 0$
 analog $now \in \mathbb{R}$ **initially** 0

States Q : $val(X)$

Actions A : **external** $try_i, crit_i, exit_i, rem_i$
 internal $test_i, set_i, check_i, reset_i$ **where** $i \in I$

Figure 6: Fischer’s mutual exclusion algorithm: Variables, states, and actions

Example 4.6 (Clock synchronization) The automaton in Figure 8 is the specification of a single process in a clock synchronization algorithm. Each process has a physical clock and generates a logical clock. The goal of the algorithm is to achieve “agreement” and “validity” among the logical clock values. Agreement means that the logical clocks are close to one another. Validity means that the logical clocks are within the range of the physical clocks.

The algorithm is based on the exchange of physical clock values between different processes in the system. The parameter u determines the frequency of sending messages. Processes in the system are indexed by the elements of a finite set I . $ClockSync(u, \rho)_i$ has a physical clock $physclock$, which may drift from the real time with a drift rate bounded by ρ . It uses the variable $maxother$ to keep track of the largest physical clock value of the other processes in the system. The variable $nextsend$ records when it is supposed to send its physical clock to the other processes. The logical clock, $logclock$, is defined to be the maximum of $maxother$ and $physclock$. Formally $logclock$ is a *derived variable*, which is a function whose value is defined in terms of the state variables.

A $send(m)_i$ transition is enabled when $m = physclock$ and $nextsend = physclock$. It causes the value of $nextsend$ to be updated so that the next send can occur when $physclock$ has advanced by u time units. The transition definition for $receive(m)_{j,i}$ specifies the effect of receiving a message from another process j in the system. Upon the receipt of a message m from j , i sets $maxother$ to the maximum of m and the current value of $maxother$, thereby updating its knowledge of the largest physical clock value of other processes in the system.

The trajectory specification is slightly different from that in the previous examples. In this example, the analog variable $physclock$ does not change at the same rate as real time

<p>Transitions \mathcal{D} :</p> <p>external try_i precondition $pc[i] = rem$ effect $pc[i] := test$</p> <p>internal $test_i$ precondition $pc[i] = test$ effect if $x = \perp$ then $pc[i] := set$ $lastset[i] := now + u_{set}$</p> <p>internal set_i precondition $pc[i] = set$ effect $x := i$ $pc[i] := check$ $lastset[i] := \infty$ $firstcheck[i] := now + l_{check}$</p> <p>internal $check_i$ precondition $pc[i] = check$ $now \geq firstcheck[i]$ effect if $x = i$ then $pc[i] := leavetry$ else $pc[i] := test$</p>	<p>external $crit_i$ precondition $pc[i] = leavetry$ effect $pc[i] := crit$</p> <p>external $exit_i$ precondition $pc[i] = crit$ effect $pc[i] := reset$</p> <p>internal $reset_i$ precondition $pc[i] = reset$ effect $x := \perp$ $pc[i] := leaveexit$</p> <p>external rem_i precondition $pc[i] = leaveexit$ effect $pc[i] := rem$</p>
<p>Trajectories \mathcal{T} :</p> <p>satisfies constant(x) constant(pc) constant($lastset$) constant($firstcheck$) $\mathbf{d}(now) = 1$ stops when $\exists i \in I. now = lastset[i]$</p>	

Figure 7: Fischer's mutual exclusion algorithm: Transitions and trajectories

Automaton $ClockSync(u, \rho)_i$ **where** $u \in \mathbb{R}^+$, $0 \leq \rho < 1$, $i \in I$

Variables X : **analog** $physclock \in \mathbb{R}$ **initially** 0
 discrete $nextsend \in \mathbb{R}$ **initially** 0
 discrete $maxother \in \mathbb{R}$ **initially** 0

Derived variables: $logclock = \max(maxother, physclock)$

States Q : $val(X)$

Actions A : **external** $send(m)_i, receive(m)_{j,i}$ **where** $m \in \mathbb{R}$, $j \in I$, $j \neq i$

Transitions \mathcal{D} : **external** $send(m)_i$
 precondition
 $m = physclock$
 $physclock = nextsend$
 effect
 $nextsend := nextsend + u$

 external $receive(m)_{j,i}$
 effect
 $maxother := \max(maxother, m)$

Trajectories \mathcal{T} : **satisfies**
 constant($nextsend$)
 constant($maxother$)
 continuous($physclock$)
 $1 - \rho \leq \mathbf{d}(physclock) \leq 1 + \rho$
 stops when
 $physclock = nextsend$

Figure 8: Clock synchronization

but it drifts with a rate that is bounded by ρ . The periodic sending of physical clocks to other processes is enforced through the stopping condition in the trajectory specification. Time is not allowed to pass beyond the point where $\text{physclock} = \text{nextsend}$. ■

4.2 Executions and Traces

We now define execution fragments, executions, trace fragments, and traces, which are used to describe automaton behavior. An *execution fragment* of a timed automaton \mathcal{A} is an (A, V) -sequence $\alpha = \tau_0 a_1 \tau_1 a_2 \tau_2 \dots$, where (1) each τ_i is a trajectory in \mathcal{T} , and (2) if τ_i is not the last trajectory in α then $\tau_i.\text{lstate} \xrightarrow{a_{i+1}} \tau_{i+1}.\text{fstate}$. An execution fragment records what happens during a particular run of a system, including all the instantaneous, discrete state changes and all the changes to the state that occur while time advances. We write $\text{frags}_{\mathcal{A}}$ for the set of all execution fragments of \mathcal{A} .

If α is an execution fragment, with notation as above, then we define the *first state* of α , $\alpha.\text{fstate}$, to be $\alpha.\text{fval}$. An *execution fragment* of a timed automaton \mathcal{A} from a state \mathbf{x} of \mathcal{A} is an execution fragment of \mathcal{A} whose first state is \mathbf{x} . We write $\text{frags}_{\mathcal{A}}(\mathbf{x})$ for the set of execution fragments of \mathcal{A} from \mathbf{x} . An execution fragment α is defined to be an *execution* if $\alpha.\text{fstate}$ is a start state, that is, $\alpha.\text{fstate} \in \Theta$. We write $\text{execs}_{\mathcal{A}}$ for the set of all executions of \mathcal{A} . If α is a closed (A, V) -sequence then we define the *last state* of α , $\alpha.\text{lstate}$, to be $\alpha.\text{lval}$.

If α is an execution fragment, then β is a *suffix* of α provided that there exists α' such that $\alpha' \frown \beta = \alpha$ and $\alpha'.\text{lstate} = \beta.\text{fstate}$.

A state of \mathcal{A} is *reachable* if it is the last state of some closed execution of \mathcal{A} . A property that is true for all reachable states of an automaton is called an *invariant assertion*, or *invariant*, for short.

Lemma 4.7 *Let $\alpha_0 \alpha_1 \dots$ be a finite or infinite sequence of execution fragments of \mathcal{A} such that, for each nonfinal index i , α_i is closed and $\alpha_i.\text{lstate} = \alpha_{i+1}.\text{fstate}$. Then $\alpha_0 \frown \alpha_1 \frown \dots$ is an execution fragment of \mathcal{A} .*

Proof: Follows easily from the definitions, using axiom **T3**. ■

Lemma 4.8 *Let α and β be execution fragments of \mathcal{A} with α closed. Then*

$$\alpha \leq \beta \iff \exists \alpha' \in \text{frags}_{\mathcal{A}} : \beta = \alpha \frown \alpha'.$$

Proof: Implication “ \Leftarrow ” follows directly from the corresponding implication in Lemma 3.7. Implication “ \Rightarrow ” follows from the definitions and **T2**. ■

The external behavior of a timed automaton is captured by the set of “traces” of its execution fragments, which record external actions and the trajectories that describe the intervening passage of time. A trace consists of alternating external actions and trajectories over the empty set of variables, \emptyset ; the only interesting information contained in these trajectories is the amount of time that elapses.

Formally, if α is an execution fragment, then the *trace* of α , denoted by $trace(\alpha)$, is the (E, \emptyset) -restriction of α , $\alpha \upharpoonright (E, \emptyset)$. A *trace fragment* of a timed automaton \mathcal{A} from a state \mathbf{x} of \mathcal{A} is the trace of an execution fragment of \mathcal{A} whose first state is \mathbf{x} . We write $tracefrags_{\mathcal{A}}(\mathbf{x})$ for the set of trace fragments of \mathcal{A} from \mathbf{x} . Also, we define a *trace* of \mathcal{A} to be a trace fragment from a start state, that is, the trace of an execution of \mathcal{A} , and write $traces_{\mathcal{A}}$ for the set of traces of \mathcal{A} .

In the earlier timed automaton models [25, 36], execution fragments were defined in a similar style to the one presented here, that is, as an alternating sequence of trajectories and actions. However, the traces were not derived from execution fragments by a simple restriction to external actions and the empty set of variables. Rather, a trace was defined as a sequence consisting of actions paired with their time of occurrence together with a limit time. The new definition increases uniformity; the definitions, results and proof techniques for hybrid sequences apply to both execution fragments and traces.

We now revisit some of the automata presented earlier in this section and give sample executions and traces for these automata.

Example 4.9 (Periodic sending process) Consider the automaton $PeriodicSend(u, M)$ from Example 4.2 where u is instantiated to the real number 3 and the message type parameter M is instantiated to the set $\{m_1, m_2 \dots\}$. The following sequence is an execution of the automaton:

$$\alpha = \tau_0 \text{ send}(m_1) \tau_1 \text{ send}(m_2) \tau_2 \text{ send}(m_3) \tau_3 \dots$$

where $\tau_i : [0, 3] \rightarrow val(\{clock\})$ are defined such that $\tau_i(t)(clock) = t$ for all $t \in [0, 3]$.

The functions τ_i are defined for closed intervals of length 3, starting at time 0. They describe the evolution of the variable *clock*, which is 0 at the start of each τ_i and increases with rate 1 for 3 time units. The discrete *send* events occur periodically, every 3 time units and reset the *clock* variable to 0.

The trace of the above execution fragment, $trace(\alpha)$, is the sequence

$$\tau'_0 \text{ send}(m_1) \tau'_1 \text{ send}(m_2) \tau'_2 \text{ send}(m_3) \tau'_3 \dots$$

where $\tau_i' : [0, 3] \rightarrow \text{val}(\emptyset)$.

Since the range of each function τ_i' contains only the function with the empty domain, $\text{trace}(\alpha)$ does not contain any information about what happens to the value of clock as time progresses. Since the domains of each τ_i and τ_i' are identical, α and $\text{trace}(\alpha')$ express the same information about the amount of time that elapses between discrete steps. ■

Example 4.10 (Timeout process) We now present an execution of the automaton $\text{Timeout}(u, M)$ from Example 4.4 where the the maximum waiting time u for a message is 5 and the message alphabet M is the set $\{m_1, m_2\}$. The following finite sequence is an execution of $\text{Timeout}(u, M)$:

$$\alpha = \tau_0 \text{ receive}(m_1) \tau_1 \text{ timeout } \tau_2 \text{ receive}(m_2) \tau_3 \text{ timeout } \tau_4$$

where $\text{Val} = \text{val}(\{\text{suspected}, \text{clock}\})$ and the functions $\tau_0, \tau_1, \tau_2, \tau_3, \tau_4$ are defined as follows:

$$\tau_0 : [0, 2] \rightarrow \text{Val} \text{ where } \tau_0(t)(\text{suspected}) = \text{false} \text{ and } \tau_0(t)(\text{clock}) = t \text{ for all } t \in [0, 2].$$

$$\tau_1 : [0, 5] \rightarrow \text{Val} \text{ where } \tau_1(t)(\text{suspected}) = \text{false} \text{ and } \tau_1(t)(\text{clock}) = t \text{ for all } t \in [0, 5].$$

$$\tau_2 : [0, 1] \rightarrow \text{Val} \text{ where } \tau_2(t)(\text{suspected}) = \text{true} \text{ and } \tau_2(t)(\text{clock}) = 5 + t \text{ for all } t \in [0, 1].$$

$$\tau_3 : [0, 5] \rightarrow \text{Val} \text{ where } \tau_3(t)(\text{suspected}) = \text{false} \text{ and } \tau_3(t)(\text{clock}) = t \text{ for all } t \in [0, 5].$$

$$\tau_4 : [0, \infty) \rightarrow \text{Val} \text{ where } \tau_4(t)(\text{suspected}) = \text{true} \text{ and } \tau_4(t)(\text{clock}) = 5 + t \text{ for all } t \in [0, \infty).$$

In this sample execution, the first awaited message arrives at time 2. Since no other message arrives within the next 5 time units, the process performs a timeout. A new message arrives 1 time unit after the timeout and the variable clock is reset to 0. Since no new message arrives in the next 5 time units the process performs another timeout. The time elapses forever after this timeout since no further message arrives.

This example illustrates that the automaton $\text{Timeout}(u, M)$ can perform multiple timeout transitions. Another point to note is that the sample execution consists of a finite (A, V) -sequence ending with a trajectory, as opposed to an infinite sequence as in Example 4.9 . The final trajectory here is a trajectory whose domain is right open and the execution is admissible and non-Zeno. Replacing τ_4 with a function on a closed interval would yield a non-Zeno execution that is not admissible.

The trace of the execution α can be obtained by letting the range of τ_i be the set consisting of the function with the empty domain, as we did in the previous example. That is, by hiding the values of the internal variables clock and suspected during trajectories. ■

Example 4.11 (Time-bounded channel) Consider the time-bounded channel automaton from Example 4.1. It is easy to observe that time cannot pass beyond any delivery

deadline recorded in the message queue and that each deadline in the queue is less than or equal to the sum of the current time and the bound b . This property can be stated as an invariant assertion as follows.

Invariant 1 : *In any reachable state \mathbf{x} of automaton $TimedChannel(b, M)$, for all (m, u) in $\mathbf{x}(queue)$, $\mathbf{x}(now) \leq u \leq \mathbf{x}(now) + b$.*

Such an invariant can be proved by induction. Recall that reachable states are the final states of closed executions. Axioms **T1** and **T2** allow us to view any closed execution as a concatenation of closed execution fragments, $\alpha_0 \frown \alpha_1 \frown \dots \frown \alpha_k$, where every α_i is either a closed trajectory or a discrete action surrounded by point trajectories, and where $\alpha_i.lstate = \alpha_{i+1}.fstate$ for $0 \leq i \leq k - 1$. The invariant can then be proved using induction on the length k of the sequence of execution fragments α_i . ■

Example 4.12 (Fischer’s mutual exclusion) The main safety property that needs to be satisfied by the automaton $FischerME$ from Example 4.5 is mutual exclusion. This safety property can be expressed as an invariant assertion:

Invariant 1 : *In any reachable state \mathbf{x} of $FischerME(u_{set}, l_{check}, I)$, there do not exist $i \in I$ and $j \in I$ such that $\mathbf{x}(pc)[i] = crit$ and $\mathbf{x}(pc)[j] = crit$.*

Even though the invariant does not refer to time, its proof depends on the timing constraints of the automaton. For example, the following auxiliary invariant can be used in proving Invariant 4.12:

Invariant 2 : *In any reachable state \mathbf{x} of $FischerME(u_{set}, l_{check}, I)$, if $pc[i] = check$, $x = i$, and $pc[j] = set$, then $firstcheck[i] > lastset[j]$.*

This invariant states that if the program counter of process i has the value *check*, the program counter of process j has the value *set*, and the variable x has the value i , then i will allow enough time for j to set x to j , before performing the check. If this timing constraint were not satisfied, it would be possible for i to check that $x = i$ before j sets x to j . Both of the processes would then observe x to contain their own index and enter the critical region. ■

Lemma 4.13 *If α is an execution of \mathcal{A} then*

1. α is time-bounded if and only if $trace(\alpha)$ is time-bounded.
2. α is admissible if and only if $trace(\alpha)$ is admissible.
3. If α is closed then $trace(\alpha)$ is closed.
4. If α is non-Zeno then $trace(\alpha)$ is non-Zeno.

Proof: It follows directly from the restriction of (A,V)-sequences. ■

Lemma 4.14 *If β is a trace of \mathcal{A} then*

1. *If β is closed then there exists an execution α of \mathcal{A} such that $\text{trace}(\alpha) = \beta$ and α is closed.*
2. *If β is non-Zeno then there exists an execution α of \mathcal{A} such that $\text{trace}(\alpha) = \beta$ and α is non-Zeno.*

Proof: For the first part of the theorem, let $\beta = \text{trace}(\alpha)$ be a closed trace of \mathcal{A} . By definition of a trace, we know that $\beta.\text{ltime} = \alpha.\text{ltime}$. We also know that α is either closed or has a suffix which is an infinite sequence of alternating point trajectories and actions. Now, let α' be the least closed prefix of α such that $\alpha'.\text{ltime} = \beta.\text{ltime}$. Clearly, α' is a closed execution of \mathcal{A} .

For the second part of the theorem, observe that a non-Zeno trace is either closed or admissible. Let $\beta = \text{trace}(\alpha)$. For the case where β is closed, we have already shown how we can find a closed execution. For the case where $\beta = \text{trace}(\alpha)$ is admissible, we know that $\alpha.\text{ltime} = \infty$. Hence, α is admissible, as needed. ■

Example 4.15 (Constructing a closed execution from a closed trace) Consider the Zeno hybrid sequence $\alpha = \wp(\mathbf{v}) a \wp(\mathbf{v}) a \wp(\mathbf{v}) \dots$ given in Example 3.12. Suppose that α is an execution of \mathcal{A} and that a is an internal action of \mathcal{A} . Then, $\text{trace}(\alpha) = \wp(\mathbf{v}')$ where $\wp(\mathbf{v}')$ is a trajectory over the empty set of variables. However, the fact that $\text{trace}(\alpha)$ is closed does not imply that α is closed. Thus, we see why we have a one way implication in item 3 of Lemma 4.13. On the other hand, we can construct a closed execution of \mathcal{A} with trace $\wp(\mathbf{v}')$ as explained in the proof of Lemma 4.14. The execution consisting of the point trajectory $\wp(\mathbf{v}')$ is a closed execution of \mathcal{A} with trace $\wp(\mathbf{v}')$. ■

4.3 Special Kinds of Timed Automata

This section describes several restricted forms of timed automata. In Section 4.3.1 we give definitions that are needed for theorems later in the paper. In Section 4.3.2 we formulate the timed automata of Alur and Dill [4, 6] as a special case of our timed automata.

4.3.1 Basic constraints

Timed Automata with Finite Internal Nondeterminism: We are sometimes interested in bounding the amount of internal nondeterminism in a timed automaton. Thus, we say that a timed automaton \mathcal{A} has *finite internal nondeterminism (FIN)* provided that:

1. The set Θ of start states is finite, and
2. For every state \mathbf{x} of \mathcal{A} and every trace fragment β of \mathcal{A} from \mathbf{x} , the set $\{\alpha.lstate \mid \alpha \in frags_{\mathcal{A}}(\mathbf{x}) \wedge trace(\alpha) = \beta\}$ is finite.

Example 4.16 (Automata with FIN) The automata $TimedChannel(u, M)$, $PeriodicSend(u, M)$, $PeriodicSend2(u, M)$ and $Timeout(u, M)$ given in Section 4.1 all have FIN. The first property of the definition of FIN is satisfied since each of these automata has a unique start state. The second property follows from the fact that in each automaton, for every state \mathbf{x} and every trace fragment β from \mathbf{x} , there is a unique execution fragment α such that $trace(\alpha) = \beta$. ■

Example 4.17 (Automata without FIN) We now show that $FischerME(u_{set}, l_{check}, I)$ and $ClockSync(a, \rho)_i$ do not have FIN. For each automaton, we specify a trace, describe the set of all executions that have the specified trace, and argue that the second property in the definition of FIN fails for the chosen trace.

Let \mathbf{x} be the start state of $FischerME(u_{set}, l_{check}, I)$ and $\beta = \tau_0 try_1 \tau_1$ be a trace of the same automaton where the domains of the functions τ_0 and τ_1 are, respectively, the single point interval $[0, 0]$ and the interval $[0, u]$, and the range of both functions is the set consisting of the function with the empty domain. For any execution α , $trace(\alpha) = \beta$, if and only if $\alpha.ltime = u$, try_1 occurs at time 0, and all the actions in α that occur after try_1 are internal actions. There are infinitely many different times that the internal actions may occur, and infinitely many values $lastcheck$ and $firstcheck$ could have, by the time u . Therefore, the set $\{\alpha.lstate \mid \alpha \in frags_{\mathcal{A}}(\mathbf{x}) \wedge trace(\alpha) = \tau_0 try_1 \tau_1\}$ is not finite and $FischerME(u_{set}, l_{check}, I)$ does not have FIN.

Now, let \mathbf{x} be the start state of $ClockSync(a, \rho)_i$ where $\mathbf{x}(physclock) = \mathbf{x}(nextsend) = \mathbf{x}(maxother) = 0$ and $\beta = \tau_0 send(0) \tau_1$ be a trace of $ClockSync(a, \rho)_i$ where the domains of functions τ_0 and τ_1 are, respectively, the interval $[0, 0]$ and the interval $[0, u]$, and the range of both functions is the set consisting of the function with the empty domain. For any α in which $send(0)$ occurs at time 0 and is followed by a trajectory τ such that $\tau.ltime = u$, we have $trace(\alpha) = \beta$. For any such α , $\alpha.lstate(physclock)$ can be any value in the interval $[u(1 - \rho), u(1 + \rho)]$. Therefore, the set $\{\alpha.lstate \mid \alpha \in frags_{\mathcal{A}}(\mathbf{x}) \wedge trace(\alpha) = \tau_0 send(0) \tau_1\}$ is not finite and $ClockSync(a, \rho)_i$ does not have FIN. ■

The following lemma states that if a timed automaton has FIN, then its set of traces is limit-closed.

Lemma 4.18 *Suppose that timed automaton \mathcal{A} has FIN and $\mathbf{x} \in Q$. Suppose that $\beta_1 \beta_2 \dots$ is a chain of trace fragments of \mathcal{A} from \mathbf{x} . Then the hybrid sequence $\lim_i \beta_i$ is a trace fragment of \mathcal{A} from \mathbf{x} .*

Proof: This is analogous to the proof of Lemma 4.3 of [25]. Suppose that \mathcal{A} is a timed automaton that has FIN, \mathbf{x} is a state of \mathcal{A} , and $\beta_1 \beta_2 \dots$ is a chain of trace fragments of \mathcal{A} from \mathbf{x} . We define a relation *after* between trace fragments from \mathbf{x} and states of \mathcal{A} : $after = \{(\beta, \mathbf{y}) \mid \exists \alpha \in frags_{\mathcal{A}}(\mathbf{x}). trace(\alpha) = \beta \wedge \alpha.lstate = \mathbf{y}\}$.

We construct a directed graph G whose nodes are pairs $(\beta_i, \mathbf{y}) \in after$ where β_i is an element of the given chain. In G , there is an edge from (β_i, \mathbf{y}) to $(\beta_{i+1}, \mathbf{y}')$ exactly if $\beta_{i+1} = \beta_i \frown \gamma$ such that $\gamma = trace(\alpha)$ for some $\alpha \in frags_{\mathcal{A}}(\mathbf{y})$, and $\alpha.lstate = \mathbf{y}'$. By the definition of property FIN, there are finitely many roots of G . By the definition of FIN and the construction of G , each node of G has finite outdegree.

We claim that each node (β_i, \mathbf{y}) of G is reachable from some root (β_1, \mathbf{z}) for some \mathbf{z} . By definition of the node set, there exists $\alpha \in frags_{\mathcal{A}}(\mathbf{x})$ such that $trace(\alpha) = \beta_i$ and $\alpha.lstate = \mathbf{y}$. Choose $\alpha' \in frags_{\mathcal{A}}(\mathbf{x})$ to be a prefix of α such that $trace(\alpha') = \beta_1$ and let $\mathbf{z} = \alpha'.lstate$. By definition of the edge set of G , (β_i, \mathbf{y}) is reachable from (β_1, \mathbf{z}) .

Hence, G satisfies the hypotheses of Lemma 2.3, which implies that there is an infinite execution fragment starting from \mathbf{x} whose trace is $\lim_i \beta_i$. Lemma 2.3 is an extension of Konig’s lemma. ■

There are two references to automata with FIN later in the paper. The first one is in Theorem 4.20, which lists some sufficient conditions for establishing an implementation relationship between two automata. The second reference appears in the discussion about the kinds of automata that satisfy the assumptions of Theorem 8.7.

Feasible Timed Automata: A timed automaton \mathcal{A} is *feasible* provided that, for every state \mathbf{x} of \mathcal{A} , there exists an admissible execution fragment of \mathcal{A} from \mathbf{x} .

Feasibility is a basic requirement that any “reasonable” timed automaton should satisfy. Theorems 4.20, 6.11 and 7.2 establish some results about feasible automata.

Timing-Independent Timed Automata: A timed automaton \mathcal{A} is said to be *timing-independent* provided that all its state variables are discrete variables, and its set of trajectories is exactly the set of constant-valued functions over left-closed time intervals with left endpoint 0.

We refer to timing-independent automata later in Example 6.5 and in our discussion about Corollary 8.8.

4.3.2 Alur-Dill Automata

The timed automaton framework of Alur and Dill [4, 6] is widely used in the formal modeling and verification of timed systems. An Alur-Dill timed automaton is a finite directed multigraph augmented with a finite set of *clock* variables. The nodes and edges

of this multigraph are called *locations* and *switches*, respectively. Locations are generally used to represent different modes of operation of the automaton, whereas the clocks are used in expressing timing constraints. Each switch has an associated *clock constraint*, which is a predicate on clock valuations that constrains when the switch may be taken. The semantics of such a timed automaton are defined as a state transition system in which each state consists of a location and a clock valuation. A transition between states occurs as a result of a switch or time passage.

Alur and Dill restrict the form of clock constraints in order to make the reachability problem (the problem of determining whether some target location is reachable) decidable: a clock constraint can be either a simple constraint comparing a clock variable to a rational constant, or a conjunction of simple constraints.

In this section, we define a version of the Alur-Dill timed automaton model as a special case of our TA model. Our formulation relaxes the restrictions on the form of clock constraints.

We assume that $\mathbb{T} = \mathbb{R}$ and define an Alur-Dill (AD) timed automaton as a TA $\mathcal{A} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$ that satisfies the following conditions:

1. X is partitioned into two sets X_d and X_c where X_d is a set of discrete variables and X_c is a set of analog variables. We call the variables in X_c *clock variables*.
2. If $\mathbf{x} \in \Theta$, then for every $x \in X_c$, $\mathbf{x}(x) = 0$.
3. If $(\mathbf{x}, a, \mathbf{x}') \in \mathcal{D}$, then for every $x \in X_c$, either $\mathbf{x}'(x) = 0$ or $\mathbf{x}'(x) = \mathbf{x}(x)$.
4. Each trajectory $\tau \in \mathcal{T}$ satisfies the following conditions:
 - (a) For every $x \in X_d$, x is constant in τ .
 - (b) For every $x \in X_c$, $d(x) = 1$.

Thus, in an AD timed automaton, the set of internal variables consists of discrete variables, which together represent the locations, and analog variables, which correspond to the clocks. In the initial states, all the clocks have value 0. A discrete transition either resets a clock or leaves it unchanged. The evolution of variables during a time interval is described by trajectories. In an AD automaton, the discrete variables are constant throughout a trajectory and clocks increase at the same rate as real time.

Example 4.19 (An AD automaton) We revisit a timed automaton example from [4]. We first present the timed automaton using the original graphical notation of Alur and Dill, as in [4], and then redefine it as an AD timed automaton, using the notational conventions we have been using in our other examples.

In the following multigraph, each switch is annotated with a symbol from a specified alphabet of *labels*, a constraint involving clock variables, and a statement that shows which

clocks are reset to 0 as a result of a location switch. Note that some switches have no reset statements, meaning that the switch has no effect on the clock variables.

The multigraph has four locations, s_0, s_1, s_2 , and s_3 , and two clocks, x and y . A location switch, represented by an arrow annotated with a label a, b, c , or d , can be performed only when the constraint on the same arrow is satisfied. For example, the automaton can change its location from s_3 to s_1 , following the switch labeled with a , when the clock variable y has a value smaller than 1. The clock variable y is reset as an effect of this location switch.




Figure 9 includes the expression of this multigraph as an AD automaton using our notational conventions. In the automaton AD, the discrete variable loc keeps track of the current location in the multigraph and the analog variables x and y represent the clocks. The actions of AD correspond to the labels in the original multigraph. Preconditions in transition definitions are used to express clock constraints associated with switches. Effects clauses in transition definitions are used to describe location changes and resetting of clocks. The trajectory specification describes the effect of time passage on the location and the clocks. ■

It is easy to check that the automaton AD, given in Figure 9, is an AD automaton. It satisfies the four conditions required to be classified as an AD automaton: (1) the set of internal variables X can be partitioned into two sets X_d and X_c where $X_d = \{loc\}$ and $X_c = \{x, y\}$. (2) The clock variables x and y are initially 0. (3) The transition definitions either reset a clock or leave it unchanged. (4) The discrete variable loc is constant throughout trajectories while x and y increase at rate 1.

4.4 Implementation Relationships

Timed automata \mathcal{A}_1 and \mathcal{A}_2 are *comparable* if they have the same external interface, that is, if $E_1 = E_2$. If \mathcal{A}_1 and \mathcal{A}_2 are comparable then we say that \mathcal{A}_1 *implements* \mathcal{A}_2 ,

Automaton AD

Variables X : **discrete** $loc \in \{s_0, s_1, s_2, s_3\}$ **initially** s_0
 analog $x \in \mathbb{R}$ **initially** 0
 analog $y \in \mathbb{R}$ **initially** 0

States Q : $val(X)$

Actions A : **external** a, b, c, d

Transitions \mathcal{D} : **external** a
 precondition
 $(loc = s_0 \text{ and } x > 0) \text{ or } (loc = s_3 \text{ and } y < 1)$
 effect
 $loc := s_1$
 $y := 0$

external b
 precondition
 $loc = s_1 \text{ and } y = 1$
 effect
 $loc := s_2$

external c
 precondition
 $(loc = s_1 \text{ and } x < 1) \text{ or } (loc = s_2 \text{ and } x < 1)$
 effect
 $loc := s_3$

external d
 precondition
 $loc = s_3 \text{ and } x > 1$

Trajectories \mathcal{T} : **satisfies**
 constant (loc)
 $\mathbf{d}(x) = 1$
 $\mathbf{d}(y) = 1$

Figure 9: An AD automaton

denoted by $\mathcal{A}_1 \leq \mathcal{A}_2$, if the traces of \mathcal{A}_1 are included among those of \mathcal{A}_2 , that is, if $\text{traces}_{\mathcal{A}_1} \subseteq \text{traces}_{\mathcal{A}_2}$.¹

Other preorders between timed automata could also be used as implementation relationships, for example, if \mathcal{A}_1 and \mathcal{A}_2 are comparable timed automata, we could consider:

- Every closed trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 .
- Every admissible trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 .
- Every non-Zeno trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 .

Theorem 4.20 *Let \mathcal{A}_1 and \mathcal{A}_2 be comparable TAs.*

1. *If every closed trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 and \mathcal{A}_2 has FIN, then $\mathcal{A}_1 \leq \mathcal{A}_2$.*
2. *If every admissible trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 and \mathcal{A}_1 is feasible, then every closed trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 .*
3. *If every admissible trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 , \mathcal{A}_1 is feasible, and \mathcal{A}_2 has FIN, then $\mathcal{A}_1 \leq \mathcal{A}_2$.*

Proof: Part 1 follows from Lemma 4.18.

For Part 2, consider a closed trace β of \mathcal{A}_1 . By feasibility of \mathcal{A}_1 , we may extend β to an admissible trace β' of \mathcal{A}_1 . Then by assumption, β' is also a trace of \mathcal{A}_2 . By prefix closure of the set of traces, β is a trace of \mathcal{A}_2 .

Part 3 follows from Parts 1 and 2. ■

4.5 Simulation Relations

In this section, we define simulation relations between timed automata. Simulation relations may be used to show that one TA implements another, in the sense of inclusion of sets of traces. We define two types of simulation relations: forward and backward simulations.

Forward simulations are more commonly used than backward simulations because they are easier to think about and are general enough to cover most interesting situations that arise in practice. Backward simulations are sometimes necessary, in particular, when non-deterministic choices are resolved earlier in the specification than in the implementation. In proving implementation relations, we prefer to use forward simulation relations whenever they exist, since backward simulations are harder to think about.

¹In [25, 14, 23, 24], definitions of the set of traces of an automaton and of one automaton implementing another are based on closed and admissible executions only. The results we obtain in this paper using the newer, more inclusive definition imply corresponding results for the earlier definition. For example, we have the following property: If $\mathcal{A}_1 \leq \mathcal{A}_2$ then the set of traces that arise from closed or admissible executions of \mathcal{A}_1 is a subset of the set of traces that arise from closed or admissible executions of \mathcal{A}_2 . This follows from Lemmas 4.13 and 4.14.

4.5.1 Forward Simulations

Let \mathcal{A} and \mathcal{B} be comparable TAs. A *forward simulation* from \mathcal{A} to \mathcal{B} is a relation $R \subseteq Q_{\mathcal{A}} \times Q_{\mathcal{B}}$ satisfying the following conditions, for all states $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ of \mathcal{A} and \mathcal{B} , respectively:

1. If $\mathbf{x}_{\mathcal{A}} \in \Theta_{\mathcal{A}}$ then there exists a state $\mathbf{x}_{\mathcal{B}} \in \Theta_{\mathcal{B}}$ such that $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$.
2. If $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$ and α is an execution fragment of \mathcal{A} consisting of one action surrounded by two point trajectories, with $\alpha.fstate = \mathbf{x}_{\mathcal{A}}$, then \mathcal{B} has a closed execution fragment β with $\beta.fstate = \mathbf{x}_{\mathcal{B}}$, $trace(\beta) = trace(\alpha)$, and $\alpha.lstate R \beta.lstate$.
3. If $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$ and α is an execution fragment of \mathcal{A} consisting of a single closed trajectory, with $\alpha.fstate = \mathbf{x}_{\mathcal{A}}$, then \mathcal{B} has a closed execution fragment β with $\beta.fstate = \mathbf{x}_{\mathcal{B}}$, $trace(\beta) = trace(\alpha)$, and $\alpha.lstate R \beta.lstate$.

Forward simulation relations induce a preorder between timed automata.

Theorem 4.21 *Let \mathcal{A}, \mathcal{B} and \mathcal{C} be comparable TAs. If R_1 is a forward simulation from \mathcal{A} to \mathcal{B} and R_2 is a forward simulation from \mathcal{B} to \mathcal{C} , then $R_2 \circ R_1$ is a forward simulation from \mathcal{A} to \mathcal{C} .*

The definition of a forward simulation from \mathcal{A} to \mathcal{B} yields a correspondence for open trajectories of \mathcal{A} :

Lemma 4.22 *Let \mathcal{A} and \mathcal{B} be comparable TAs and let R be a forward simulation from \mathcal{A} to \mathcal{B} . Let $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ be states of \mathcal{A} and \mathcal{B} , respectively, such that $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$. Let α be an execution fragment of \mathcal{A} from state $\mathbf{x}_{\mathcal{A}}$ consisting of a single open trajectory. Then \mathcal{B} has an execution fragment β with $\beta.fstate = \mathbf{x}_{\mathcal{B}}$ and $trace(\beta) = trace(\alpha)$.*

Proof: Let τ be the single open trajectory in α . Using axioms **T1** and **T2**, we construct an infinite sequence $\tau_0 \tau_1 \dots$ of closed trajectories of \mathcal{A} such that $\tau = \tau_0 \frown \tau_1 \frown \dots$. Then, working recursively, we construct a sequence $\beta_0 \beta_1 \dots$ of closed execution fragments of \mathcal{B} such that $\beta_0.fstate = \mathbf{x}_{\mathcal{B}}$ and, for each i , $\tau_i.lstate R \beta_i.lstate$, $\beta_i.lstate = \beta_{i+1}.fstate$, and $trace(\tau_i) = trace(\beta_i)$. This construction uses induction on i , using Property 3 of the definition of a forward simulation in the induction step. Now let $\beta = \beta_0 \frown \beta_1 \frown \dots$. By Lemma 4.7, β is an execution fragment of \mathcal{B} . Clearly, $\beta.fstate = \mathbf{x}_{\mathcal{B}}$. By Lemma 3.9 applied to both α and β , $trace(\beta) = trace(\alpha)$. Thus β has the required properties. \blacksquare

Theorem 4.23 *Let \mathcal{A} and \mathcal{B} be comparable TAs and let R be a forward simulation from \mathcal{A} to \mathcal{B} . Let $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ be states of \mathcal{A} and \mathcal{B} , respectively, such that $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$. Then $tracefrags_{\mathcal{A}}(\mathbf{x}_{\mathcal{A}}) \subseteq tracefrags_{\mathcal{B}}(\mathbf{x}_{\mathcal{B}})$.*

Proof: Suppose that δ is the trace of an execution fragment of \mathcal{A} that starts from $\mathbf{x}_{\mathcal{A}}$; we prove that δ is also a trace of an execution fragment of \mathcal{B} that starts from $\mathbf{x}_{\mathcal{B}}$. Let $\alpha = \tau_0 a_1 \tau_1 a_2 \tau_2 \dots$ be an execution fragment of \mathcal{A} such that $\alpha.fstate = \mathbf{x}_{\mathcal{A}}$ and $\delta = trace(\alpha)$. We consider cases:

1. α is an infinite sequence.

Using axioms **T1** and **T2**, we can write α as an infinite concatenation $\alpha_0 \frown \alpha_1 \frown \alpha_2 \dots$, in which the execution fragments α_i with i even consist of a trajectory only, and the execution fragments α_i with i odd consist of a single discrete step surrounded by two point trajectories.

We define inductively a sequence $\beta_0 \beta_1 \dots$ of closed execution fragments of \mathcal{B} , such that $\beta_0.fstate = \mathbf{x}_{\mathcal{B}}$ and, for all i , $\beta_i.lstate = \beta_{i+1}.fstate$, $\alpha_i.lstate R \beta_i.lstate$, and $trace(\beta_i) = trace(\alpha_i)$. We use Property 3 of the definition of a simulation for the construction of the β_i 's with i even, and Property 2 for the construction of the β_i 's with i odd. Let $\beta = \beta_0 \frown \beta_1 \frown \beta_2 \dots$. By Lemma 4.7, β is an execution fragment of \mathcal{B} . Clearly, $\beta.fstate = \mathbf{x}_{\mathcal{B}}$. By Lemma 3.9, $trace(\beta) = trace(\alpha)$. Thus β has the required properties.

2. α is a finite sequence ending with a closed trajectory.

Similar to the first case.

3. α is a finite sequence ending with an open trajectory.

Similar to the first case, using Lemma 4.22. ■

Corollary 4.24 *Let \mathcal{A} and \mathcal{B} be comparable TAs and let R be a forward simulation from \mathcal{A} to \mathcal{B} . Then $traces_{\mathcal{A}} \subseteq traces_{\mathcal{B}}$.*

Proof: Suppose $\beta \in traces_{\mathcal{A}}$. Then $\beta \in tracefrags_{\mathcal{A}}(\mathbf{x}_{\mathcal{A}})$ for some start state $\mathbf{x}_{\mathcal{A}}$ of \mathcal{A} . Property 1 of the definition of simulation implies the existence of a start state $\mathbf{x}_{\mathcal{B}}$ of \mathcal{B} such that $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$. Then Theorem 4.23 implies that $\beta \in tracefrags_{\mathcal{B}}(\mathbf{x}_{\mathcal{B}})$. Since $\mathbf{x}_{\mathcal{B}}$ is a start state of \mathcal{B} , this implies that $\beta \in traces_{\mathcal{B}}$, as needed. ■

Example 4.25 (Time-bounded channels) Consider two instances of the specification in Figure 2, $TimedChannel(b_1, M)$ and $TimedChannel(b_2, M)$ where $b_1 \leq b_2$. We define a forward simulation R from $TimedChannel(b_1, M)$ to $TimedChannel(b_2, M)$ below. If \mathbf{x} is a state of $TimedChannel(b_1, M)$ and \mathbf{y} is a state of $TimedChannel(b_2, M)$, then $\mathbf{x} R \mathbf{y}$ provided that the following conditions are satisfied:

1. $\mathbf{x}(now) = \mathbf{y}(now)$.
2. $|\mathbf{x}(queue)| = |\mathbf{y}(queue)|$.

3. $\forall i. 1 \leq i \leq |\mathbf{x}(queue)|$, if $\mathbf{x}(queue)(i) = (m, u_1)$ then $\mathbf{y}(queue)(i) = (m, u_2)$ and $u_1 \leq u_2$.

We can prove that R is a forward simulation from the automaton $TimedChannel(b_1, M)$ to the automaton $TimedChannel(b_2, M)$ by showing that R satisfies each of the three properties in the definition of a forward simulation relation. In each automaton there is a unique initial state that maps the variable *now* to 0 and *queue* to the empty sequence. It is obvious that the initial states, which are identical, are related by R and so the first property is satisfied.

For the rest of the proof, we let \mathbf{x} and \mathbf{y} be, respectively, states of $TimedChannel(b_1, M)$ and $TimedChannel(b_2, M)$ such that $\mathbf{x} R \mathbf{y}$. In order to show that the second property is satisfied, we need to consider two cases, one for each discrete action that may be performed by $TimedChannel(b_1, M)$.

If $TimedChannel(b_1, M)$ performs a *send*(m) action, and the state changes from \mathbf{x} to \mathbf{x}' then we need to find an execution fragment β of $TimedChannel(b_2, M)$ from \mathbf{y} ending in \mathbf{y}' , such that $\mathbf{x}' R \mathbf{y}'$ and $trace(\beta)$ is the same as the trace of $\wp(\mathbf{x}) \text{ send}(m) \wp(\mathbf{y})$. The execution fragment $\beta = \wp(\mathbf{y}) \text{ send}(m) \wp(\mathbf{y}')$ satisfies the required conditions. This follows from the hypothesis that $\mathbf{x} R \mathbf{y}$ and the definition of R , using the fact that the effect of a *send*(m) action of $TimedChannel(b_1, M)$, $TimedChannel(b_2, M)$ are, respectively, adding the entry $(m, now + b_1)$ to $\mathbf{x}(queue)$, and $(m, now + b_2)$ to $\mathbf{y}(queue)$ where $b_1 \leq b_2$.

If $TimedChannel(b_1, M)$ performs a *receive*(m) action, and the state changes from \mathbf{x} to \mathbf{x}' then we need to show that *receive*(m) is also enabled in \mathbf{y} and that there is an execution fragment with the required properties that ends in a state \mathbf{y}' such that $\mathbf{x}' R \mathbf{y}'$. In order to show that *receive*(m) is enabled in \mathbf{y} , we use the hypothesis that $\mathbf{x} R \mathbf{y}$, which implies that the first element of $\mathbf{y}(queue)$ is of the form (m, u) for some u . The execution fragment $\wp(\mathbf{y}) \text{ receive}(m) \wp(\mathbf{y}')$ of $TimedChannel(b_1, M)$ can be shown to satisfy the required conditions.

For the third property, we consider a closed trajectory τ of $TimedChannel(b_1, M)$ with $\tau.fstate = \mathbf{x}$ and show that there exists a closed execution fragment β of the automaton $TimedChannel(b_2, M)$ with $\beta.fstate = \mathbf{y}$, $trace(\beta) = trace(\tau)$, and $\tau.lstate = \beta.lstate$. It is easy to check that the trajectory τ' of $TimedChannel(b_2, M)$ with $\tau'.fstate = \mathbf{y}$ and $\tau'.ltime = \tau.ltime$ satisfies the required conditions. \blacksquare

Example 4.26 (Time-bounded channel that keeps all messages) In this example we define a variant of $TimedChannel(b, M)$ from Example 4.1 called $TimedChannel2(b, M)$. The main difference between $TimedChannel(b, M)$ and $TimedChannel2(b, M)$ is that the message queue in $TimedChannel2(b, M)$ is implemented using a finite sequence of (message, delivery deadline) pairs *queue* and a pointer *ptr* that points to the next element that is to be delivered. Hence, the internal variables of $TimedChannel2(b, M)$ consist of *queue*, *now* and *ptr*. The variable *ptr* initially has value 1, which indicates that it

Automaton $SendVal(u, \rho)_i$ where $u \in \mathbb{R}^+$, $0 \leq \rho < 1$, $i \in I$

Variables X : **discrete** $counter \in \mathbb{R}$ **initially** 0
 analog $now \in \mathbb{R}$ **initially** 0

States Q : $val(X)$

Actions A : **external** $send(m)_i, receive(m)_{j,i}$ where $m \in \mathbb{R}$, $j \in I$, $j \neq i$

Transitions \mathcal{D} : **external** $send(m)_i$
 precondition
 $m = counter \times u$
 $counter \times u / (1 + \rho) \leq now$
 effect
 $counter := counter + 1$

 external $receive(m)_{j,i}$

Trajectories \mathcal{T} : **satisfies**
 constant($counter$)
 $d(now) = 1$
 stops when
 $now = counter \times u / (1 - \rho)$

Figure 10: Clock synchronization

is pointing to the first element in the sequence. A $send(m)$ action causes messages and deadlines to be added to the sequence as in $TimedChannel(b, M)$. A $receive(m)$ causes ptr to be incremented to make it point to the next element in the sequence instead of removing the first element. The automaton $TimedChannel(b, M)$ can be viewed as an optimized implementation of $TimedChannel2(b, M)$.

We define below a forward simulation R from $TimedChannel(b, M)$ to $TimedChannel2(b, M)$. If \mathbf{x} is a state of $TimedChannel(b, M)$ and \mathbf{y} is a state of $TimedChannel2(b, M)$, then $\mathbf{x} R \mathbf{y}$ provided that the following conditions are satisfied:

1. $\mathbf{x}(now) = \mathbf{y}(now)$.
2. $\mathbf{x}(queue) = \mathbf{y}(queue)(\mathbf{y}(ptr) \dots |\mathbf{y}(queue)|)$.

■

Example 4.27 (Clock synchronization) In this example, we define a forward simulation from $ClockSync(u, \rho)_i$ of Figure 8 to an automaton that sends multiples of u . The specification of this automaton, which is called $SendVal(u, \rho)$, is given in Figure 10. We

assume that the subscripts representing process indices in both automata are drawn from the same finite set I .

The variable *counter* keeps track of which multiple of u is to be sent next, and variable *now* contains the current time. The automaton parameter ρ is used in the precondition of the *send* and the stopping condition of the trajectory definition, to enforce bounds on the times of occurrence of *send*.

We now define a forward simulation R from the automaton $ClockSync(u, \rho)_i$ to the automaton $SendVal(u, \rho)$ where u and ρ are actual parameters. If \mathbf{x} is a state of the automaton $ClockSync(u, \rho)_i$ and \mathbf{y} is a state of $SendVal(u, \rho)$, then $\mathbf{x} R \mathbf{y}$ provided that the following conditions are satisfied:

1. $\mathbf{y}(now)(1 - \rho) \leq \mathbf{x}(physclock) \leq \mathbf{y}(now)(1 + \rho)$.
2. $\mathbf{y}(counter) = \mathbf{x}(nextsend)/u$.

■

4.5.2 Refinements

Let \mathcal{A} and \mathcal{B} be comparable TAs. A *refinement* from \mathcal{A} to \mathcal{B} is a function $F \subseteq Q_{\mathcal{A}} \times Q_{\mathcal{B}}$, satisfying the following conditions, for all states $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ of \mathcal{A} and \mathcal{B} , respectively:

1. If $\mathbf{x}_{\mathcal{A}} \in \Theta_{\mathcal{A}}$ then $F(\mathbf{x}_{\mathcal{A}}) \in \Theta_{\mathcal{B}}$.
2. If α is an execution fragment of \mathcal{A} consisting of one action surrounded by two point trajectories, with $\alpha.fstate = \mathbf{x}_{\mathcal{A}}$, then \mathcal{B} has a closed execution fragment β with $\beta.fstate = F(\mathbf{x}_{\mathcal{A}})$, $trace(\beta) = trace(\alpha)$, and $\beta.lstate = F(\alpha.lstate)$.
3. If α is an execution fragment of \mathcal{A} consisting of a single closed trajectory, with $\alpha.fstate = \mathbf{x}_{\mathcal{A}}$, then \mathcal{B} has a closed execution fragment β with $\beta.fstate = F(\mathbf{x}_{\mathcal{A}})$, $trace(\beta) = trace(\alpha)$, and $\beta.lstate = F(\alpha.lstate)$.

Theorem 4.28 *Let \mathcal{A} and \mathcal{B} be two TAs and suppose $R \subseteq Q_{\mathcal{A}} \times Q_{\mathcal{B}}$. Then R is a refinement from \mathcal{A} to \mathcal{B} if and only if R is a forward simulation from \mathcal{A} to \mathcal{B} and R is a function.*

Theorem 4.29 *Let \mathcal{A}, \mathcal{B} and \mathcal{C} be comparable TAs. If R_1 is a refinement from \mathcal{A} to \mathcal{B} and R_2 is a refinement from \mathcal{B} to \mathcal{C} , then $R_2 \circ R_1$ is a refinement from \mathcal{A} to \mathcal{C} .*

An isomorphism from \mathcal{A} to \mathcal{B} is a refinement F from \mathcal{A} to \mathcal{B} such that F^{-1} is a refinement from \mathcal{B} to \mathcal{A} . We say that two automata \mathcal{A} and \mathcal{B} are *isomorphic*, if there exists an isomorphism from \mathcal{A} to \mathcal{B} (or, equivalently from \mathcal{B} to \mathcal{A}).

4.5.3 Backward Simulations

Let \mathcal{A} and \mathcal{B} be comparable TAs. A *backward simulation* from \mathcal{A} to \mathcal{B} is a total relation $R \subseteq Q_{\mathcal{A}} \times Q_{\mathcal{B}}$ satisfying the following conditions, for all states $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ of \mathcal{A} and \mathcal{B} , respectively:

1. If $\mathbf{x}_{\mathcal{A}} \in \Theta_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$ then $\mathbf{x}_{\mathcal{B}} \in \Theta_{\mathcal{B}}$.
2. If $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$ and α is an execution fragment of \mathcal{A} with $\alpha.lstate = \mathbf{x}_{\mathcal{A}}$, consisting of one discrete action surrounded by two point trajectories, then \mathcal{B} has a closed execution fragment β with $\beta.lstate = \mathbf{x}_{\mathcal{B}}$, $trace(\beta) = trace(\alpha)$, and $\alpha.fstate R \beta.fstate$.
3. If $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$ and α is an execution fragment of \mathcal{A} with $\alpha.lstate = \mathbf{x}_{\mathcal{A}}$, consisting of one trajectory, then \mathcal{B} has a closed execution fragment β with $\beta.lstate = \mathbf{x}_{\mathcal{B}}$, $trace(\beta) = trace(\alpha)$, and $\alpha.fstate R \beta.fstate$.

Backward simulations induce a preorder between timed automata.

Theorem 4.30 *Let \mathcal{A}, \mathcal{B} and \mathcal{C} be comparable TAs. If R_1 is a backward simulation from \mathcal{A} to \mathcal{B} and R_2 is a backward simulation \mathcal{B} to \mathcal{C} , then $R_2 \circ R_1$ is a backward simulation from \mathcal{A} to \mathcal{C} .*

Theorem 4.31 *Let \mathcal{A} and \mathcal{B} be comparable TAs and let R be a backward simulation from \mathcal{A} to \mathcal{B} . Let $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ be states of \mathcal{A} and \mathcal{B} , respectively, such that $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$. Let β be the trace of a closed execution fragment of \mathcal{A} from $\mathbf{y}_{\mathcal{A}}$ with last state $\mathbf{x}_{\mathcal{A}}$. Then there exists $\mathbf{y}_{\mathcal{B}}$ such that β is also the trace of a closed execution fragment of \mathcal{B} from $\mathbf{y}_{\mathcal{B}}$ with last state $\mathbf{x}_{\mathcal{B}}$ and $\mathbf{y}_{\mathcal{A}} R \mathbf{y}_{\mathcal{B}}$.*

Proof: Fix some R , $\mathbf{x}_{\mathcal{A}}$, $\mathbf{x}_{\mathcal{B}}$ and β satisfying the conditions in the statement of the theorem. Let $\alpha \in frags_{\mathcal{A}}(\mathbf{y}_{\mathcal{A}})$ for some state $\mathbf{y}_{\mathcal{A}}$ of \mathcal{A} with $trace(\alpha) = \beta$. By using the axioms **T1** and **T2**, we can write α as the concatenation of a sequence of closed execution fragments, $\alpha = \alpha_0 \frown \alpha_1 \frown \dots \frown \alpha_n$, where each α_i is either a closed trajectory or an action surrounded by two point trajectories, and $\alpha_i.lstate = \alpha_{i+1}.fstate$ for $0 \leq i \leq n$.

By using the definition of a backward simulation, working backwards from α_n , we can construct an execution fragment $\alpha' = \alpha'_0 \frown \alpha'_1 \frown \dots \frown \alpha'_n$ from a state $\mathbf{y}_{\mathcal{B}}$ of \mathcal{B} such that (a) $\alpha'.lstate = \mathbf{x}_{\mathcal{B}}$, (b) for all i , $0 \leq i \leq n$, $\alpha_i.fstate R \alpha'_i.fstate$ and $trace(\alpha'_i) = trace(\alpha_i)$, (c) for all i , $0 \leq i \leq n-1$, $\alpha'_i.lstate = \alpha'_{i+1}.fstate$. Using Lemma 4.7, we can see that α' is an execution fragment of \mathcal{B} . By Lemma 3.9, $trace(\alpha) = trace(\alpha')$ as needed. ■

Corollary 4.32 *Let \mathcal{A} and \mathcal{B} be comparable TAs and let R be a backward simulation from \mathcal{A} to \mathcal{B} . Then every closed trace of \mathcal{A} is a trace of \mathcal{B} .*

Proof: Suppose R is a backward simulation from \mathcal{A} to \mathcal{B} and β is a closed trace of \mathcal{A} . Then $\beta = \text{trace}(\alpha)$ for some closed execution α of \mathcal{A} . Let $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{y}_{\mathcal{A}}$ be the first and last states of α respectively. By the totality of relation R , there exists some state $\mathbf{y}_{\mathcal{B}}$ of \mathcal{B} such that $\mathbf{y}_{\mathcal{A}} R \mathbf{y}_{\mathcal{B}}$. By Theorem 4.31, there exists $\mathbf{x}_{\mathcal{B}}$ of \mathcal{B} such that β is the trace of a closed execution fragment of \mathcal{B} from $\mathbf{x}_{\mathcal{B}}$ with last state $\mathbf{y}_{\mathcal{B}}$ and $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$. Property 1 of the definition of a backward simulation relation implies that $\mathbf{x}_{\mathcal{B}}$ is a start state of \mathcal{B} . It follows that $\beta \in \text{traces}_{\mathcal{B}}$, as needed. ■

Theorem 4.33 *Let \mathcal{A} and \mathcal{B} be comparable TAs and let R be an image-finite backward simulation from \mathcal{A} to \mathcal{B} . Then $\text{traces}_{\mathcal{A}} \subseteq \text{traces}_{\mathcal{B}}$.*

Proof: Let $\beta \in \text{traces}_{\mathcal{A}}$. If β is closed then Corollary 4.32 implies that β is a trace of \mathcal{B} . From now on we assume β is not closed.

Let $\alpha \in \text{execs}_{\mathcal{A}}$ with $\text{trace}(\alpha) = \beta$. Note that any such α is either an infinite sequence $\tau_0 a_1 \tau_1 \dots$ or a finite sequence $\tau_0 a_1 \tau_1 \dots \tau_n$ where the final trajectory τ_n is right open. In either case, using the axioms **T1** and **T2**, we can construct an infinite sequence $\alpha_0 \alpha_1 \dots$ of closed execution fragments such that $\alpha = \alpha_0 \hat{\ } \alpha_1 \hat{\ } \dots$ where α_0 is a point trajectory, each α_i is either a closed trajectory or an action surrounded by two point trajectories, and $\alpha_i.\text{lstate} = \alpha_{i+1}.\text{fstate}$ for each i , $0 \leq i$.

We construct a directed graph G whose nodes are pairs (\mathbf{x}, i) consisting of a state of \mathcal{B} and an index such that $(\alpha_i.\text{lstate}, \mathbf{x}) \in R$. In G , there is an edge from (\mathbf{x}, i) to (\mathbf{x}', j) exactly if $j = i + 1$ and there is an $\alpha' \in \text{frags}_{\mathcal{B}}(\mathbf{x})$ with $\text{trace}(\alpha') = \text{trace}(\alpha_{i+1})$ such that $\alpha'.\text{lstate} = \mathbf{x}'$. Since R is image-finite there are finitely many roots of G . By image-finiteness of R and the definition of the edge set, each node has finite outdegree. By using the definition of a backward simulation and the edge set of G , we can show that each node (\mathbf{x}, i) is reachable from some root node $(\mathbf{z}, 0)$ for some start state \mathbf{z} of \mathcal{B} .

The directed graph G satisfies the hypotheses of Lemma 2.3, which implies that there is an infinite path in G starting from a root. An edge from a node (\mathbf{x}, i) to $(\mathbf{x}', i + 1)$ along this infinite path corresponds to a closed execution fragment γ_{i+1} of \mathcal{B} for i , $0 \leq i$ such that $\gamma_{i+1}.\text{fstate} = \mathbf{x}$, $\gamma_{i+1}.\text{lstate} = \mathbf{x}'$ and $\text{trace}(\gamma_{i+1}) = \text{trace}(\alpha_{i+1})$. By Lemma 4.7, $\gamma = \gamma_1 \hat{\ } \gamma_2 \hat{\ } \dots$ is an execution of \mathcal{B} and by Lemma 3.9, $\text{trace}(\gamma) = \text{trace}(\gamma_1) \hat{\ } \text{trace}(\gamma_2) \hat{\ } \dots$. Since $\text{trace}(\gamma_{i+1}) = \text{trace}(\alpha_{i+1})$ for all i , $0 \leq i$, and α_0 is a point trajectory, by Lemma 3.9, we get $\text{trace}(\gamma) = \text{trace}(\alpha) = \beta$. ■


Example 4.34 (A backward simulation relation) This example illustrates the difference between forward and backward simulations. We consider two automata \mathcal{A} and

\mathcal{B} and show that a forward simulation from \mathcal{A} to \mathcal{B} does not exist while we exhibit a backward simulation from \mathcal{A} to \mathcal{B} .

Let \mathcal{A} and \mathcal{B} be two comparable automata specified below. The trajectories consist of a set of point trajectories. This implies that the automaton does not allow time to pass — everything happens at time 0.

- $V_{\mathcal{A}} = \{stateA\}$ and $V_{\mathcal{B}} = \{stateB\}$ where:
 $stateA$ is a discrete variable with $type(stateA) = \{x_{\mathcal{A}}, y_{\mathcal{A}}, q_{\mathcal{A}}, s_{\mathcal{A}}\}$, and
 $stateB$ is a discrete variable with $type(stateB) = \{x_{\mathcal{B}}, y_{\mathcal{B}}, y'_{\mathcal{B}}, q_{\mathcal{B}}, s_{\mathcal{B}}\}$.
- $Q_{\mathcal{A}} = val(V_{\mathcal{A}})$ and $Q_{\mathcal{B}} = val(V_{\mathcal{B}})$. We write $\mathbf{x}_{\mathcal{A}}$ for the valuation that maps $stateA$ to $x_{\mathcal{A}}$, $\mathbf{y}_{\mathcal{A}}$ for the valuation that maps $stateA$ to $y_{\mathcal{A}}$, etc. Similarly, we write $\mathbf{x}_{\mathcal{B}}$ for the valuation that maps $stateB$ to $x_{\mathcal{B}}$, $\mathbf{y}_{\mathcal{B}}$ for the valuation that maps $stateB$ to $y_{\mathcal{B}}$, etc.
- $\Theta_{\mathcal{A}} = \{\mathbf{x}_{\mathcal{A}}\}$ and $\Theta_{\mathcal{B}} = \{\mathbf{x}_{\mathcal{B}}\}$.
- $E_{\mathcal{A}} = E_{\mathcal{B}} = \{a, b, c\}$ and $H_{\mathcal{A}} = H_{\mathcal{B}} = \emptyset$.
- $\mathcal{D}_{\mathcal{A}} = \{(\mathbf{x}_{\mathcal{A}}, a, \mathbf{y}_{\mathcal{A}}), (\mathbf{y}_{\mathcal{A}}, b, \mathbf{q}_{\mathcal{A}}), (\mathbf{y}_{\mathcal{A}}, c, \mathbf{s}_{\mathcal{A}})\}$, and
 $\mathcal{D}_{\mathcal{B}} = \{(\mathbf{x}_{\mathcal{B}}, a, \mathbf{y}_{\mathcal{B}}), (\mathbf{x}_{\mathcal{B}}, a, \mathbf{y}'_{\mathcal{B}}), (\mathbf{y}_{\mathcal{B}}, b, \mathbf{q}_{\mathcal{B}}), (\mathbf{y}'_{\mathcal{B}}, c, \mathbf{s}_{\mathcal{B}})\}$.
- $\mathcal{T}_{\mathcal{A}} = \{\wp(\mathbf{v}) \mid \mathbf{v} \in Q_{\mathcal{A}}\}$, and $\mathcal{T}_{\mathcal{B}} = \{\wp(\mathbf{v}) \mid \mathbf{v} \in Q_{\mathcal{B}}\}$

The following are representations of automata \mathcal{A} and \mathcal{B} as directed multigraphs. The nodes in the graph represent states and the edges represent discrete transitions where a label on an edge stands for the action involved in the transition.



An obvious candidate for a forward simulation from \mathcal{A} to \mathcal{B} is the relation $R = \{(\mathbf{x}_{\mathcal{A}}, \mathbf{x}_{\mathcal{B}}), (\mathbf{y}_{\mathcal{A}}, \mathbf{y}_{\mathcal{B}}), (\mathbf{y}_{\mathcal{A}}, \mathbf{y}'_{\mathcal{B}}), (\mathbf{q}_{\mathcal{A}}, \mathbf{q}_{\mathcal{B}}), (\mathbf{s}_{\mathcal{A}}, \mathbf{s}_{\mathcal{B}})\}$. However, observe that even though $\mathbf{y}_{\mathcal{A}}$ and $\mathbf{y}_{\mathcal{B}}$ are related by R , the execution fragment $\wp(\mathbf{y}_{\mathcal{A}}) c \wp(\mathbf{s}_{\mathcal{A}})$ of \mathcal{A} cannot be matched by any execution fragment of \mathcal{B} starting with state $\mathbf{y}_{\mathcal{B}}$. Similarly, even though $\mathbf{y}_{\mathcal{A}}$ and $\mathbf{y}'_{\mathcal{B}}$ are related by R , the execution fragment $\wp(\mathbf{y}_{\mathcal{A}}) b \wp(\mathbf{q}_{\mathcal{A}})$ of \mathcal{A} cannot be matched by any execution fragment of \mathcal{B} starting with $\mathbf{y}'_{\mathcal{B}}$. Therefore, R is not a forward simulation. In fact, there is no forward simulation relation from \mathcal{A} to \mathcal{B} : there are finitely many possibilities for forward simulations from \mathcal{A} to \mathcal{B} and we see that none of them is a forward simulation by examining all the possibilities. The main reason for this is that

while \mathcal{A} makes the nondeterministic choice between performing b or c after performing a , \mathcal{B} makes its choice earlier at the same time it performs a .

There is, however, a backward simulation from \mathcal{A} to \mathcal{B} : the relation R defined above is a backward simulation. \blacksquare

4.5.4 History Relations

A relation $R \subseteq Q_{\mathcal{A}} \times Q_{\mathcal{B}}$ is a *history relation* from \mathcal{A} to \mathcal{B} if R is a forward simulation from \mathcal{A} to \mathcal{B} and R^{-1} is a refinement from \mathcal{B} to \mathcal{A} . History relations induce a preorder between timed automata.

An automaton \mathcal{B} is obtained from an automaton \mathcal{A} by *adding history variables* if there exists a set of variables V such that

1. $V_{\mathcal{B}} = V_{\mathcal{A}} \cup V$ and $V_{\mathcal{A}} \cap V = \emptyset$,
2. $Q_{\mathcal{B}} \subseteq \text{val}(V_{\mathcal{B}})$ such that $Q_{\mathcal{B}} \upharpoonright V_{\mathcal{A}} \subseteq Q_{\mathcal{A}}$, and
3. The relation $\{(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in Q_{\mathcal{B}} \text{ and } \mathbf{y} \upharpoonright V_{\mathcal{A}} = \mathbf{x}\}$ is a history relation from \mathcal{A} to \mathcal{B} .

The method of adding history variables is typically used to make it possible to establish an implementation relationship using a refinement. If a refinement does not exist from a low-level automaton to a higher-level one, it can often be made to exist by adding history variables to the low-level automaton.

Example 4.35 (Adding history variables to obtain a refinement) We cannot show that $\textit{TimedChannel}(b, M)$ is an implementation of $\textit{TimedChannel2}(b, M)$ from Example 4.26 by using a refinement. This is because we have no way of specifying what the subsequence before the pointer should be in $\textit{TimedChannel2}(b, M)$ when relating the states of the two automata. This example shows how we can add history variables to $\textit{TimedChannel}(b, M)$ (actually, we add just one variable) to obtain a new automaton that is related to $\textit{TimedChannel2}(b, M)$ by a refinement.

Let \textit{log} be a discrete variable whose static type is the same as the static type of *queue* in $\textit{TimedChannel}(b, M)$ and let the initial value of \textit{log} be the empty sequence. We define a new automaton $\textit{TimedChannelH}(b, M)$ whose set of variables consists of the variables of $\textit{TimedChannel}(b, M)$ and the variable \textit{log} . The rest of the definition of $\textit{TimedChannelH}(b, M)$ is the same as $\textit{TimedChannel}(b, M)$ except for the transition definition for *receive*(m). A *receive*(m) event in $\textit{TimedChannelH}(b, M)$ not only removes the first message from the message queue but also appends this message to the sequence contained in \textit{log} .

Let V_1, V_2 be the set of variables and Q_1, Q_2 be the set of states of $\textit{TimedChannel}(b, M)$ and $\textit{TimedChannelH}(b, M)$ respectively. It is easy to verify that the relation $\{(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in$

Q_2 and $\mathbf{y} \upharpoonright V_1 = \mathbf{x}$ is a history relation from $TimedChannel(b, M)$ to $TimedChannelH(b, M)$. This means that $TimedChannelH(b, M)$ is obtained from $TimedChannel(b, M)$ by adding a history variable.

We now define a refinement F from $TimedChannelH(b, M)$ to $TimedChannel2(b, M)$ as follows. In our definition we assume the following conventions. Concatenation on the left corresponds to putting an element on the front of a queue. Recall also that we use juxtaposition for concatenation of sequences. If \mathbf{x} is a state of $TimedChannelH(b, M)$ and \mathbf{y} is a state of $TimedChannel2(b, M)$, then $F(\mathbf{x}) = \mathbf{y}$ where:

1. $\mathbf{y}(now) = \mathbf{x}(now)$.
2. $\mathbf{y}(queue) = \mathbf{x}(log) \mathbf{x}(queue)$ such that $|\mathbf{x}(log)| = \mathbf{y}(ptr) - 1$.

■

Whenever an automaton \mathcal{B} is obtained from \mathcal{A} by adding history variables, then there exists a history relation from \mathcal{A} to \mathcal{B} by definition. Theorem 4.36 states that the converse also holds, if isomorphic automata are considered.

Theorem 4.36 *Let \mathcal{A} and \mathcal{B} be two comparable TAs such that $V_{\mathcal{A}}$ and $V_{\mathcal{B}}$ are disjoint. Suppose that there is a history relation from \mathcal{A} to \mathcal{B} . Then, there exists an automaton \mathcal{C} that is isomorphic to \mathcal{B} and is obtained from \mathcal{A} by adding history variables.*

Proof: Let R be a history relation from \mathcal{A} to \mathcal{B} . Define automaton \mathcal{C} as follows:

- $V_{\mathcal{C}} = V_{\mathcal{A}} \cup V_{\mathcal{B}}$.
- $Q_{\mathcal{C}} = \{\mathbf{x} \in val(V_{\mathcal{C}}) \mid (\mathbf{x} \upharpoonright V_{\mathcal{A}}, \mathbf{x} \upharpoonright V_{\mathcal{B}}) \in R\}$.
- $\Theta_{\mathcal{C}} = \{\mathbf{x} \in Q_{\mathcal{C}} \mid \mathbf{x} \upharpoonright V_{\mathcal{B}} \in \Theta_{\mathcal{B}}\}$.
- $E_{\mathcal{C}} = E_{\mathcal{B}}$ and $H_{\mathcal{C}} = H_{\mathcal{B}}$.
- $\mathbf{x} \xrightarrow{a}_{\mathcal{C}} \mathbf{y}$ if and only if $\mathbf{x} \upharpoonright V_{\mathcal{B}} \xrightarrow{a}_{\mathcal{B}} \mathbf{y} \upharpoonright V_{\mathcal{B}}$.
- $\mathbf{x} \xrightarrow{\tau}_{\mathcal{C}} \mathbf{y}$ if and only if $\mathbf{x} \upharpoonright V_{\mathcal{B}} \xrightarrow{\tau_1}_{\mathcal{B}} \mathbf{y} \upharpoonright V_{\mathcal{B}}$ where $\tau_1 = \tau \downarrow V_{\mathcal{B}}$.

Let $F : Q_{\mathcal{C}} \rightarrow Q_{\mathcal{B}}$ be defined such that $F(\mathbf{x}) = \mathbf{x} \upharpoonright V_{\mathcal{B}}$ for all $\mathbf{x} \in Q_{\mathcal{C}}$. The function F is an isomorphism from \mathcal{C} to \mathcal{B} : It is easy to check that F is a refinement from \mathcal{C} to \mathcal{B} . We can also easily verify that F^{-1} is a refinement from \mathcal{B} to \mathcal{C} , by definition of \mathcal{C} and the fact that R^{-1} is a function from the states of \mathcal{B} to the states of \mathcal{A} .

Now, we verify that \mathcal{C} is obtained from \mathcal{A} by adding history variables. Let $V_{\mathcal{B}}$ be the variable set V required in the definition of a history variable and let $R' = \{(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in Q_{\mathcal{C}} \wedge \mathbf{y} \upharpoonright V_{\mathcal{A}} = \mathbf{x}\}$. We need to show that R' is a history relation from \mathcal{A} to \mathcal{C} .

1. R' is a forward simulation from \mathcal{A} to \mathcal{C} .
By definitions of the relations F , R' and the automaton \mathcal{C} , $R' = F^{-1} \circ R$. Since F^{-1} is a refinement from \mathcal{B} to \mathcal{C} , by Theorem 4.28, we know that it is a forward simulation from \mathcal{B} to \mathcal{C} . Since R is a forward simulation from \mathcal{A} to \mathcal{B} , by Theorem 4.21 we have R' is a forward simulation from \mathcal{A} to \mathcal{C} , as needed.
2. R'^{-1} is a refinement from \mathcal{C} to \mathcal{A} .
By definitions of the relations F , R' and the automaton \mathcal{C} , $R'^{-1} = R^{-1} \circ F$. Since F is a refinement from \mathcal{C} to \mathcal{B} and R^{-1} is a refinement from \mathcal{B} to \mathcal{A} , by Theorem 4.29, we have R'^{-1} is a refinement from \mathcal{C} to \mathcal{A} , as needed.

■

The following theorem shows that forward simulations are essentially the same as history relations combined with refinements.

Theorem 4.37 *Let \mathcal{A} and \mathcal{B} be two comparable TAs such that $V_{\mathcal{A}}$ and $V_{\mathcal{B}}$ are disjoint. There is a forward simulation from \mathcal{A} to \mathcal{B} if and only if there exists a TA \mathcal{C} such that there is a history relation from \mathcal{A} to \mathcal{C} and a refinement from \mathcal{C} to \mathcal{B} .*

Proof: To prove the implication \Rightarrow , suppose R is a forward simulation from \mathcal{A} to \mathcal{B} . Let \mathcal{C} be an automaton defined as follows:

- $V_{\mathcal{C}} = V_{\mathcal{A}} \cup V_{\mathcal{B}}$.
- $Q_{\mathcal{C}} = \{\mathbf{x} \in \text{val}(V_{\mathcal{C}}) \mid (\mathbf{x} \upharpoonright V_{\mathcal{A}}, \mathbf{x} \upharpoonright V_{\mathcal{B}}) \in R\}$.
- $\Theta_{\mathcal{C}} = \{\mathbf{x} \in Q_{\mathcal{C}} \mid \mathbf{x} \upharpoonright V_{\mathcal{A}} \in \Theta_{\mathcal{A}} \wedge \mathbf{x} \upharpoonright V_{\mathcal{B}} \in \Theta_{\mathcal{B}}\}$.
- $E_{\mathcal{C}} = E_{\mathcal{A}}$ and $H_{\mathcal{C}} = H_{\mathcal{A}}$.
- $\mathbf{x} \xrightarrow{a}_{\mathcal{C}} \mathbf{y}$ if and only if both of the following conditions hold:
 1. $\mathbf{x} \upharpoonright V_{\mathcal{A}} \xrightarrow{a}_{\mathcal{A}} \mathbf{y} \upharpoonright V_{\mathcal{A}}$.
 2. There exists an execution fragment β of \mathcal{B} such that $\beta.\text{fstate} = \mathbf{x} \upharpoonright V_{\mathcal{B}}$, $\beta.\text{lstate} = \mathbf{y} \upharpoonright V_{\mathcal{B}}$, and $\text{trace}(\beta) = \text{trace}(\wp(\mathbf{x}) \ a \ \wp(\mathbf{y}))$.
- $\mathbf{x} \xrightarrow{\tau}_{\mathcal{C}} \mathbf{y}$ if and only if both of the following conditions hold:
 1. $\tau_1 = \tau \downarrow V_{\mathcal{A}} \in \mathcal{T}_{\mathcal{A}}$ and $\mathbf{x} \upharpoonright V_{\mathcal{A}} \xrightarrow{\tau}_{\mathcal{A}} \mathbf{y} \upharpoonright V_{\mathcal{A}}$.
 2. $\tau_2 = \tau \downarrow V_{\mathcal{B}} \in \mathcal{T}_{\mathcal{B}}$ and $\mathbf{x} \upharpoonright V_{\mathcal{B}} \xrightarrow{a}_{\mathcal{B}} \mathbf{y} \upharpoonright V_{\mathcal{B}}$.

Let π_A and π_B be the functions that restrict states of C to, respectively, V_A and V_B . It follows from the definitions that π_A^{-1} is a history relation from \mathcal{A} to \mathcal{C} and π_B is a refinement from \mathcal{C} to \mathcal{B} .

For the implication \Leftarrow , suppose that there is a history relation from \mathcal{A} to \mathcal{C} and that there is a refinement from \mathcal{C} to \mathcal{B} . Then, by definition of a history relation, we know that there is a forward simulation from \mathcal{A} to \mathcal{C} . We also know that there is a forward simulation from \mathcal{C} to \mathcal{B} by Theorem 4.28. It follows that there is a forward simulation from \mathcal{A} to \mathcal{B} , as needed. ■

Example 4.38 (Theorem 4.37 applied to time-bounded channels) In Example 4.26, we demonstrated a forward simulation from the automaton $TimedChannel(b, M)$ to the automaton $TimedChannel2(b, M)$. Theorem 4.37 implies that there exists an automaton \mathcal{A} such that there is a history relation from $TimedChannel(b, M)$ to \mathcal{A} and a refinement from \mathcal{A} to $TimedChannel2(b, M)$. The automaton $TimedChannelH(b, M)$ from Example 4.35 is a witness for \mathcal{A} . ■

4.5.5 Prophecy Relations

A relation $R \subseteq Q_A \times Q_B$ is a *prophecy relation* from \mathcal{A} to \mathcal{B} if R is a backward simulation from \mathcal{A} to \mathcal{B} and R^{-1} is a refinement from \mathcal{B} to \mathcal{A} . Prophecy relations induce a preorder between timed automata.



An automaton \mathcal{B} is obtained from an automaton \mathcal{A} by *adding prophecy variables* if there exists a set of variables V such that

1. $V_B = V_A \cup V$ and $V_A \cap V = \emptyset$,
2. $Q_B \subseteq val(V_B)$ such that $Q_B \upharpoonright V_A \subseteq Q_A$, and
3. The relation $\{(\mathbf{x}, \mathbf{y}) \mid \mathbf{y} \in Q_B \text{ and } \mathbf{y} \upharpoonright V_A = \mathbf{x}\}$ is a prophecy relation from \mathcal{A} to \mathcal{B} .

Example 4.39 (Adding prophecy variables to obtain a refinement) In this example we consider adding a prophecy variable to the automaton \mathcal{A} from Example 4.34. Let \mathcal{C} be an automaton defined as follows:

- $V_C = V_A \cup \{v\}$ where v is a discrete variable with $type(v) = \{b, c\}$.

- $Q_C = \{\mathbf{x}_C, \mathbf{x}'_C, \mathbf{y}_C, \mathbf{y}'_C, \mathbf{q}_C, \mathbf{s}_C\}$ such that
 - $\mathbf{x}_C \upharpoonright V_A = \mathbf{x}_A$ and $\mathbf{x}_C \upharpoonright \{v\} = b$
 - $\mathbf{x}'_C \upharpoonright V_A = \mathbf{x}_A$ and $\mathbf{x}'_C \upharpoonright \{v\} = c$
 - $\mathbf{y}_C \upharpoonright V_A = \mathbf{y}_A$ and $\mathbf{y}_C \upharpoonright \{v\} = b$
 - $\mathbf{y}'_C \upharpoonright V_A = \mathbf{y}_A$ and $\mathbf{y}'_C \upharpoonright \{v\} = c$
 - $\mathbf{q}_C \upharpoonright V_A = \mathbf{q}_A$ and $\mathbf{q}_C \upharpoonright \{v\} = b$
 - $\mathbf{s}_C \upharpoonright V_A = \mathbf{s}_A$ and $\mathbf{s}_C \upharpoonright \{v\} = c$
- $\Theta_C = \{\mathbf{x}_C, \mathbf{x}'_C\}$.
- $E_C = \{a, b, c\}$.
- $\mathcal{D}_C = \{(\mathbf{x}_C, a, \mathbf{y}_C), (\mathbf{x}'_C, a, \mathbf{y}'_C), (\mathbf{y}_C, b, \mathbf{q}_C), (\mathbf{y}'_C, c, \mathbf{s}_C)\}$.
- $\mathcal{T}_C = \{\emptyset(\mathbf{v}) \mid \mathbf{v} \in Q_C\}$.



The relation $R = \{(\mathbf{x}_A, \mathbf{x}_C), (\mathbf{x}_A, \mathbf{x}'_C), (\mathbf{y}_A, \mathbf{y}_C), (\mathbf{y}_A, \mathbf{y}'_C), (\mathbf{q}_A, \mathbf{q}_C), (\mathbf{s}_A, \mathbf{s}_C)\}$ is a backward simulation from \mathcal{A} to \mathcal{C} and R^{-1} is a refinement. Therefore, \mathcal{C} is obtained by adding a prophecy variable to \mathcal{A} . Note that there is no refinement from \mathcal{A} to \mathcal{B} defined in Example 4.34. However, the relation $F = \{(\mathbf{x}_C, \mathbf{x}_B), (\mathbf{x}'_C, \mathbf{x}_B), (\mathbf{y}_C, \mathbf{y}_B), (\mathbf{y}'_C, \mathbf{y}'_B), (\mathbf{q}_C, \mathbf{q}_B), (\mathbf{s}_C, \mathbf{s}_B)\}$ is a refinement from \mathcal{C} to \mathcal{B} . ■

Theorem 4.40 *Let \mathcal{A} and \mathcal{B} be two comparable TAs such that V_A and V_B are disjoint. Suppose that there is a prophecy relation from \mathcal{A} to \mathcal{B} . Then, there exists an automaton \mathcal{C} that is isomorphic to \mathcal{B} and is obtained from \mathcal{A} by adding prophecy variables.*

Proof: The proof is analogous to the proof of Theorem 4.36. We assume a backward simulation relation R instead of a forward simulation relation. We construct the automaton \mathcal{C} as in Theorem 4.36 and verify that it is obtained from \mathcal{A} by adding a prophecy variable. ■

Theorem 4.41 *Let \mathcal{A} and \mathcal{B} be two comparable TAs such that $V_{\mathcal{A}}$ and $V_{\mathcal{B}}$ are disjoint. There is a backward simulation from \mathcal{A} to \mathcal{B} if and only if there exists a TA \mathcal{C} such that there is a prophecy relation from \mathcal{A} to \mathcal{C} and a refinement from \mathcal{C} to \mathcal{B} .*

Proof: The proof is analogous to the proof of Theorem 4.37. We assume a backward simulation relation R instead of a forward simulation. The construction of the automaton \mathcal{C} and the reasoning that follows are similar. ■

Example 4.42 (Theorem 4.41 applied to Examples 4.34 and 4.39) In Example 4.34, we demonstrated a backward simulation from \mathcal{A} to \mathcal{B} . Theorem 4.41 implies that there exists an automaton \mathcal{C} such that there is a prophecy relation from \mathcal{A} to \mathcal{C} and a refinement from \mathcal{C} to \mathcal{B} . The automaton \mathcal{C} defined in Example 4.39 constitutes a witness for \mathcal{C} . ■

5 Operations on Timed Automata

In this section, we introduce four kinds of operations on timed automata: parallel composition, hiding, adding lower and upper bounds for tasks, and untiming.

5.1 Composition

5.1.1 Definitions and Basic Results

The composition operation for timed automata allows an automaton representing a complex system to be constructed by composing automata representing individual system components. Our composition operation identifies external actions with the same name in different component automata. When any component automaton performs a discrete step involving an action a , so do all component automata that have a as an external action. The composition operator for timed automata is simpler than it is for general hybrid automata since all the variables in a timed automaton are internal.²

Formally, we say that timed automata \mathcal{A}_1 and \mathcal{A}_2 are *compatible* if $H_1 \cap A_2 = H_2 \cap A_1 = \emptyset$ and $X_1 \cap X_2 = \emptyset$. If \mathcal{A}_1 and \mathcal{A}_2 are compatible then their *composition* $\mathcal{A}_1 \parallel \mathcal{A}_2$ is defined to be the structure $\mathcal{A} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$ where

- $X = X_1 \cup X_2$.
- $Q = \{\mathbf{x} \in \text{val}(X) \mid \mathbf{x} \upharpoonright X_i \in Q_i, i \in \{1, 2\}\}$.

²The composition operation for general hybrid automata requires external variables to be identified as well as external actions. When any component automaton follows a particular trajectory for an external variable v , then so do all component automata of which v is an external variable.

- $\Theta = \{\mathbf{x} \in Q \mid \mathbf{x} \upharpoonright X_i \in \Theta_i, i \in \{1, 2\}\}$.
- $E = E_1 \cup E_2$ and $H = H_1 \cup H_2$.
- For each $\mathbf{x}, \mathbf{x}' \in Q$ and each $a \in A$, $\mathbf{x} \xrightarrow{a}_{\mathcal{A}} \mathbf{x}'$ iff for $i \in \{1, 2\}$, either (1) $a \in A_i$ and $\mathbf{x} \upharpoonright X_i \xrightarrow{a}_i \mathbf{x}' \upharpoonright X_i$, or (2) $a \notin A_i$ and $\mathbf{x} \upharpoonright X_i = \mathbf{x}' \upharpoonright X_i$.
- $\mathcal{T} \subseteq \text{trajs}(X)$ is given by $\tau \in \mathcal{T} \Leftrightarrow \tau \downarrow X_i \in \mathcal{T}_i, i \in \{1, 2\}$.

Theorem 5.1 *If \mathcal{A}_1 and \mathcal{A}_2 are timed automata then $\mathcal{A}_1 \parallel \mathcal{A}_2$ is a timed automaton.*

Lemma 5.2 *Let $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$ and let α be an execution fragment of \mathcal{A} . Then $\alpha \upharpoonright (A_1, X_1)$ and $\alpha \upharpoonright (A_2, X_2)$ are execution fragments of \mathcal{A}_1 and \mathcal{A}_2 , respectively. Furthermore,*

1. α is time-bounded iff both $\alpha \upharpoonright (A_1, X_1)$ and $\alpha \upharpoonright (A_2, X_2)$ are time-bounded.
2. α is admissible iff both $\alpha \upharpoonright (A_1, X_1)$ and $\alpha \upharpoonright (A_2, X_2)$ are admissible.
3. α is closed iff both $\alpha \upharpoonright (A_1, X_1)$ and $\alpha \upharpoonright (A_2, X_2)$ are closed.
4. α is non-Zeno iff both $\alpha \upharpoonright (A_1, X_1)$ and $\alpha \upharpoonright (A_2, X_2)$ are non-Zeno.
5. α is an execution iff both $\alpha \upharpoonright (A_1, X_1)$ and $\alpha \upharpoonright (A_2, X_2)$ are executions.

Lemma 5.3 *Let $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$, and let α be an execution fragment of \mathcal{A} . Then, for $i = 1, 2$, $\text{trace}(\alpha) \upharpoonright (E_i, \emptyset) = \text{trace}(\alpha \upharpoonright (A_i, X_i))$.*

The following theorem is a fundamental theorem that relates the set of traces of a composed automaton to the sets of traces of its components. Set inclusion in one direction expresses the idea that a trace of a composition “projects” to yield traces of the components. Set inclusion in the other direction expresses the idea that traces of components can be “pasted” to yield a trace of the composition.

Theorem 5.4 *Let $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$. Then $\text{traces}_{\mathcal{A}}$ is exactly the set of (E, \emptyset) -sequences whose restrictions to \mathcal{A}_1 and \mathcal{A}_2 are traces of \mathcal{A}_1 and \mathcal{A}_2 , respectively.*

That is, $\text{traces}_{\mathcal{A}} = \{\beta \mid \beta \text{ is an } (E, \emptyset)\text{-sequence and } \beta \upharpoonright (E_i, \emptyset) \in \text{traces}_{\mathcal{A}_i}, i \in \{1, 2\}\}$.

Notation: The compatibility conditions for composition require the set of internal variables of each automaton to be disjoint from the set of internal variables of all the other automata in the composition. We use a general scheme to disambiguate the internal variables of components in order to avoid possible name clashes that can violate the compatibility conditions. If \mathcal{A} is the name of an automaton and v is an internal variable of \mathcal{A} , then we refer to this variable as $\mathcal{A}.v$ in the composite automaton.

Example 5.5 (Periodic sending process with timeouts) Let \mathcal{C} be the composition of three automata from Examples 4.1, 4.2 and 4.4:

$$C = \text{PeriodicSend}(u_1, M) \parallel \text{TimedChannel}(b, M) \parallel \text{Timeout}(u_2, M)$$

where $M = \{m_1, \dots, m_n\}$ and $b + u_1 < u_2$. The following sequence is a trace of C .

$$\alpha = \tau_0 \text{ send}(m_1) \tau_1 \text{ receive}(m_1) \tau_2 \text{ send}(m_2) \tau_3 \text{ receive}(m_2) \tau_4 \dots$$

where e is the set consisting of the function with the empty domain and

$$\tau_0 : [0, u_1] \rightarrow e \quad \tau_1 : [0, b] \rightarrow e \quad \tau_2 : [0, u_1 - b] \rightarrow e \quad \tau_3 : [0, b] \rightarrow e \quad \tau_4 : [0, u_1 - b] \rightarrow e$$

The following invariant states that C never performs a *timeout* action.

Invariant 1 : *In any reachable state \mathbf{x} of C , $\mathbf{x}(\text{Timeout.suspected}) = \text{false}$.*

In order to prove this invariant we can use an auxiliary invariant such as the one below, which establishes the fact that every message is delivered before the variable *now*, which keeps track of real-time, reaches the point at which a *timeout* action occurs.

Invariant 2 :

1. *if $\mathbf{x}(\text{TimedChannel.queue})$ is not empty then*
 $\mathbf{x}(\text{TimedChannel.queue})(1) < \mathbf{x}(\text{TimedChannel.now}) + u_2 - \mathbf{x}(\text{Timeout.clock})$.
2. *if $\mathbf{x}(\text{TimedChannel.queue})$ is empty then*
 $u_1 - \mathbf{x}(\text{PeriodicSend.clock}) + b < u_2 - \mathbf{x}(\text{Timeout.clock})$.

■

Example 5.6 (Periodic sending process with failures and timeouts) In this example, we consider a composite automaton defined exactly like the one in Example 5.5 except that the automaton $\text{PeriodicSend}(u_1, M)$ is replaced with $\text{PeriodicSend2}(u_1, M)$. Let $C = \text{PeriodicSend2}(u_1, M) \parallel \text{TimedChannel}(b, M) \parallel \text{Timeout}(u_2, M)$. The following sequence is a trace of C .

$$\tau_0 \text{ send}(m_1) \tau_1 \text{ receive}(m_1) \tau_2 \text{ fail} \tau_3 \text{ timeout} \tau_4$$

where e is the set consisting of the function with the empty domain and

$$\tau_0 : [0, u_1] \rightarrow e \quad \tau_1 : [0, b] \rightarrow e \quad \tau_2 : [0, b] \rightarrow e \quad \tau_3 : [0, u_2 - b] \rightarrow e \quad \tau_4 : [0, \infty) \rightarrow e$$

According to this sample trace, the first message sent by the periodic sending process is received exactly b time units after it is sent. The periodic sending process fails $2b$ time units after sending its first message. The timeout process performs a *timeout* since no second message arrives within the next u_2 time units after the receipt of the first message.

The following invariant states that a *timeout* performed by C can be used to conclude that the sender process has failed.

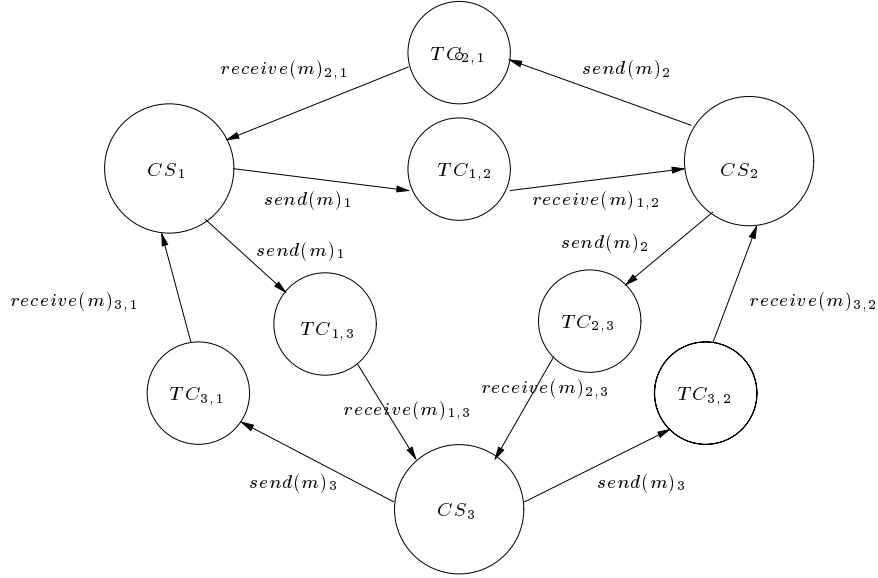
Invariant 1 : *Let $C = \text{PeriodicSend2}(u_1, M) \parallel \text{TimedChannel}(b, M) \parallel \text{Timeout}(u_2, M)$ and assume that $b + u_1 < u_2$. In any reachable state \mathbf{x} of C , if $\mathbf{x}(\text{Timeout.suspected}) = \text{true}$ then $\mathbf{x}(\text{PeriodicSend2.failed}) = \text{true}$.*

The automaton C is guaranteed to perform a *timeout* to signal the failure of a process, within a specified amount of time after the occurrence of a fail event. The following is a formal statement of this property.

Let α be an execution of C and let t be the point in time at which a *fail* event occurs in α . Then α includes a *timeout* event that occurs in the interval $(t + b, t + b + u_2]$. ■

Example 5.7 (Clock synchronization) In this example we consider the composition of three clock synchronization automata with six time-bounded channel automata. A graphical representation of the composite automaton is given below. The abbreviation CS_i represents the automaton $\text{ClockSync}(u, \rho)_i$. The abbreviation $TC_{i,j}$ represents the timed channel that communicates messages from $\text{ClockSync}(u, \rho)_i$ to $\text{ClockSync}(u, \rho)_j$. We assume that the time-bounded channel automata used in this composition are defined as in Example 4.1 where *receive* and *send* actions in each instance are renamed such that they can be shared with clock synchronization automata. Let C be

$$\text{ClockSync}(u, \rho)_1 \parallel \text{ClockSync}(u, \rho)_2 \parallel \text{ClockSync}(u, \rho)_3 \parallel \\ \text{TimedChannel}(b, M)_1 \parallel \dots \parallel \text{TimedChannel}(b, M)_6 \text{ where } M = \mathbb{R}^+.$$



A physical clock diverges from real time at the largest rate when it evolves with rate $1 + \rho$ or $1 - \rho$. For example, if a physical clock evolves with rate $1 + \rho$, then at time t , its value is $t(1 + \rho)$. Hence, the largest possible difference between a physical clock and the real time is $t\rho$. This property is stated by the invariant below.

Invariant 1 : *In any reachable state \mathbf{x} of C , at any time $t \in \mathbb{T}$, for any $i \in \{1, 2, 3\}$, $|\mathbf{x}(\text{ClockSync}(u, \rho)_i.\text{physclock}) - t| \leq t\rho$.*

Two physical clocks in C diverge at the largest rate when one evolves with rate $1 + \rho$ and the other with $1 - \rho$. It follows from Invariant 1 that, at any time t the largest possible difference between the physical clock values for two processes is $2t\rho$. This property is formalized by the following invariant.

Invariant 2 : *In any reachable state \mathbf{x} of C , at any time $t \in \mathbb{T}$, for any $i, j \in \{1, 2, 3\}$, $|\mathbf{x}(\text{ClockSync}(u, \rho)_i.\text{physclock}) - \mathbf{x}(\text{ClockSync}(u, \rho)_j.\text{physclock})| \leq 2t\rho$ where $i, j \in \{1, 2, 3\}$.*

The following invariant states that in any reachable state there exists a process j such that the logical clock of each other process in the system is smaller than or equal to the physical clock of j . This follows from the definition of a logical clock and the fact that physical clocks always increase.

Invariant 3 : *In any reachable state \mathbf{x} of C , there exists $j \in \{1, 2, 3\}$ such that for all $i \in \{1, 2, 3\}$, $\mathbf{x}(\text{ClockSync}(u, \rho)_i.\text{logclock}) \leq \mathbf{x}(\text{ClockSync}(u, \rho)_j.\text{physclock})$.*

The following invariant states that in any reachable state there exists a process j such that the logical clock of each other process in the system is larger than or equal to the physical clock of j . This follows from the definition of a logical clock.

Invariant 4 : *In any reachable state \mathbf{x} of C , there exists $j \in \{1, 2, 3\}$ such that for all $i \in \{1, 2, 3\}$, $\mathbf{x}(\text{ClockSync}(u, \rho)_i.\text{logclock}) \geq \mathbf{x}(\text{ClockSync}(u, \rho)_j.\text{physclock})$.*

Invariants 3 and 4 together are called validity properties. They express the condition that all the logical clocks remain in an envelope bounded by the maximum and minimum physical clock values in the system.

The following invariant formalizes the property that all the logical clocks at a given time lie within the envelope formed by the largest and the smallest physical clock values in the system. It follows from Invariants 1, 3 and 4 that any point in this envelope can diverge from real time t by at most $t\rho$ time units.

Invariant 5 : *In any reachable state \mathbf{x} of C , at any time $t \in \mathbb{T}$, for any $i \in \{1, 2, 3\}$ $|\mathbf{x}(\text{ClockSync}(u, \rho)_i.\text{logclock}) - t| \leq t\rho$.*

Finally, we state a property about the agreement of logical clocks in C .

Invariant 6 : *In any reachable state \mathbf{x} of C , for $i, j \in \{1, 2, 3\}$, $|\mathbf{x}(\text{ClockSync}(u, \rho)_i.\text{logclock}) - \mathbf{x}(\text{ClockSync}(u, \rho)_j.\text{logclock})| \leq u + b(1 + \rho)$.*

To see why Invariant 6 holds, fix j to be a process with the largest physical clock in \mathbf{x} , and fix i to be any other process. Let v_j, v_i be the logical clock values of j and i respectively in state \mathbf{x} . Note that v_j is also the physical clock value of j in \mathbf{x} . By Invariant 3, we know that $v_i \leq v_j$. To show Invariant 6, it suffices to show that $v_j - v_i \leq u + b(1 + \rho)$.

Let α be a finite execution that leads to state \mathbf{x} . There are two cases to consider.

1. Some message sent by j arrives at i in α .

Consider the last such message and let v_1 be the value that it contains. Let v_2 be the newly adjusted logical clock value of i immediately after the message arrives. We know that $v_i \geq v_2 \geq v_1$.

If j sends a later message to i in α , then it sends the next later message when its physical clock has value $v_1 + u$. By assumption, this message does not arrive at i . Therefore, the real time that elapses after sending it must be at most b . It follows that the physical clock increase of j since sending this message is at most $b(1 + \rho)$ and so $v_j \leq v_1 + u + b(1 + \rho)$. On the other hand, if j does not send a later message to i in α , then $v_j \leq v_1 + u$. In either case, we have $v_j \leq v_1 + u + b(1 + \rho)$. Since $v_i \geq v_1$, we have $v_j - v_i \leq u + b(1 + \rho)$, as needed for Invariant 6.

2. No message sent by j arrives at i in α .

Since the first send occurs at time 0 and b is the largest possible communication delay, the fact that i has not received the first message sent by j at time 0 implies that $t \leq b$. Since both clocks start at 0, we have $v_j \leq b(1 + \rho)$ and $v_i \geq 0$. Therefore, $v_j - v_i \leq u + b(1 + \rho)$, which suffices for Invariant 6. ■

5.1.2 Substitutivity Results

Theorem 5.4, which relates the set of traces of a composed automaton to the set of traces of component automata, is fundamental for compositional reasoning. We now introduce another important class of results, substitutivity results, that are useful for decomposing verification of composite automata. These results are best understood by viewing one of the components of a composition as the system and the other as the environment with which the system interacts.

The following result states that if a TA \mathcal{A}_1 can be shown to implement another one \mathcal{A}_2 , with no assumptions about their environments, then \mathcal{A}_1 can be shown to implement \mathcal{A}_2 in a given environment \mathcal{B} .

Theorem 5.8 *Suppose \mathcal{A}_1 , \mathcal{A}_2 and \mathcal{B} are TAs, \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with \mathcal{B} . If $\mathcal{A}_1 \leq \mathcal{A}_2$ then $\mathcal{A}_1 \parallel \mathcal{B} \leq \mathcal{A}_2 \parallel \mathcal{B}$.*

Corollary 5.9 *Suppose \mathcal{A}_1 , \mathcal{A}_2 , \mathcal{B}_1 , and \mathcal{B}_2 are TAs, \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, \mathcal{B}_1 and \mathcal{B}_2 have the same external actions, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with each of \mathcal{B}_1 and \mathcal{B}_2 . If $\mathcal{A}_1 \leq \mathcal{A}_2$ and $\mathcal{B}_1 \leq \mathcal{B}_2$ then $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.*

We can strengthen Corollary 5.9 slightly by the following corollary: if \mathcal{A}_1 implements \mathcal{A}_2 in an environment \mathcal{B}_2 , then \mathcal{A}_1 composed with an environment that is more restrictive than \mathcal{B}_2 (whose set of external behaviors is smaller than that of \mathcal{B}_2), implements \mathcal{A}_2 composed with \mathcal{B}_2 .

Corollary 5.10 *Suppose \mathcal{A}_1 , \mathcal{A}_2 , \mathcal{B}_1 , and \mathcal{B}_2 are TAs, \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, \mathcal{B}_1 and \mathcal{B}_2 have the same external actions, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with each of \mathcal{B}_1 and \mathcal{B}_2 . If $\mathcal{A}_1 \parallel \mathcal{B}_2 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ and $\mathcal{B}_1 \leq \mathcal{B}_2$ then $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.*

Proof: Let $\beta \in \text{traces}_{\mathcal{A}_1 \parallel \mathcal{B}_1}$. By Theorem 5.4, $\beta \upharpoonright (E_{\mathcal{A}_1}, \emptyset) \in \text{traces}_{\mathcal{A}_1}$ and $\beta \upharpoonright (E_{\mathcal{B}_1}, \emptyset) \in \text{traces}_{\mathcal{B}_1}$. Since $\mathcal{B}_1 \leq \mathcal{B}_2$, $\beta \upharpoonright (E_{\mathcal{B}_1}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. Since \mathcal{B}_1 and \mathcal{B}_2 have the same external actions, it follows that $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. We have $\beta \upharpoonright (E_{\mathcal{A}_1}, \emptyset) \in \text{traces}_{\mathcal{A}_1}$ and $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. By Theorem 5.4, $\beta \in \text{traces}_{\mathcal{A}_1 \parallel \mathcal{B}_2}$. Since $\mathcal{A}_1 \parallel \mathcal{B}_2 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ by assumption, $\beta \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_2}$, as needed. ■

For other preorders, we also get substitutivity results, for example:

Theorem 5.11 *Suppose \mathcal{A}_1 , \mathcal{A}_2 and \mathcal{B} are TAs, \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with \mathcal{B} .*

1. If every closed trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 then every closed trace of $\mathcal{A}_1 \parallel \mathcal{B}$ is a trace of $\mathcal{A}_2 \parallel \mathcal{B}$.
2. If every admissible trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 then every admissible trace of $\mathcal{A}_1 \parallel \mathcal{B}$ is a trace of $\mathcal{A}_2 \parallel \mathcal{B}$.
3. If every non-Zeno trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 then every non-Zeno trace of $\mathcal{A}_1 \parallel \mathcal{B}$ is a trace of $\mathcal{A}_2 \parallel \mathcal{B}$.

Example 5.12 (A counterexample for a desirable substitutivity theorem) Suppose \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, \mathcal{B}_1 and \mathcal{B}_2 have the same external actions, and that each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with each of \mathcal{B}_1 and \mathcal{B}_2 . If we view \mathcal{A}_2 and \mathcal{B}_2 as specifications and want to prove that $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$, it would be useful to have a theorem that says if $\mathcal{A}_1 \parallel \mathcal{B}_2 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ and $\mathcal{A}_2 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ then $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$. That is, if \mathcal{A}_1 implements \mathcal{A}_2 in the context of \mathcal{B}_2 and \mathcal{B}_1 implements \mathcal{B}_2 in the context of \mathcal{A}_2 , we would like to conclude that $\mathcal{A}_1 \parallel \mathcal{B}_1$ implements $\mathcal{A}_2 \parallel \mathcal{B}_2$. We show by means of a counterexample that it is impossible to prove such a theorem.

Consider the definitions of automata $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1, \mathcal{B}_2$ in Figures 11 and 12. All automata have the same set of actions, consisting of the external actions a and b . \mathcal{A}_1 can perform an arbitrary number of bs , and can perform an a provided that the count of as and the count of bs are equal. \mathcal{A}_1 allows the count of as to increase to one more than the count of bs .

\mathcal{B}_1 can perform an arbitrary number of as , and can perform a b provided that the count of as is one more than the count of bs . \mathcal{B}_1 allows the count of bs to reach the count of as .

\mathcal{A}_2 has an infinite number of start states, each giving a different finite bound on the number of a actions it can perform. Similarly, \mathcal{B}_2 has an infinite number of start states, each giving a different finite bound on the number of b actions it can perform.

Clearly, $\mathcal{A}_1 \parallel \mathcal{B}_2 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$, and $\mathcal{A}_2 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$. On the other hand, $\mathcal{A}_1 \parallel \mathcal{B}_1$ can perform an infinite sequence of alternating as and bs , which is not allowed by the specification $\mathcal{A}_2 \parallel \mathcal{B}_2$. This implies that $\mathcal{A}_1 \parallel \mathcal{B}_1$ does not implement $\mathcal{A}_2 \parallel \mathcal{B}_2$. ■

In Section 8, we revisit the substitutivity issue and prove Theorem 8.8, a variant of the desirable theorem considered in the above example, by assuming certain conditions on the environments \mathcal{A}_2 and \mathcal{B}_2 .

5.2 Hiding

We define one hiding operation for timed automata, which hides external actions: if $E \subseteq E_{\mathcal{A}}$, then $\text{ActHide}(E, \mathcal{A})$ is the TA \mathcal{B} that is equal to \mathcal{A} except that $E_{\mathcal{B}} = E_{\mathcal{A}} - E$ and $H_{\mathcal{B}} = H_{\mathcal{A}} \cup E$.

Automaton \mathcal{A}_1

Variables X : **discrete** $counta \in \mathbb{Z}$ **initially** 0
 discrete $countb \in \mathbb{Z}$ **initially** 0

States Q : $val(X)$

Actions A : **external** a, b

Transitions \mathcal{D} : **external** a
 precondition
 $countb = counta$
 effect
 $counta := counta + 1$

 external b
 effect
 $countb := countb + 1$

Trajectories \mathcal{T} : $\{\varphi(\mathbf{x}) \mid \mathbf{x} \in Q\}$

Automaton \mathcal{B}_1

Variables X : **discrete** $counta \in \mathbb{Z}$ **initially** 0
 discrete $countb \in \mathbb{Z}$ **initially** 0

States Q : $val(X)$

Actions A : **external** a, b

Transitions \mathcal{D} : **external** b
 precondition
 $counta = countb + 1$
 effect
 $countb := countb + 1$

 external a
 effect
 $counta := counta + 1$

Trajectories \mathcal{T} : $\{\varphi(\mathbf{x}) \mid \mathbf{x} \in Q\}$

Figure 11: Automata \mathcal{A}_1 and \mathcal{B}_1

Automaton \mathcal{A}_2

Variables X : **discrete** $maxcount \in \mathbb{Z}^{\geq 0}$ **initially** arbitrary
 discrete $counta \in \mathbb{Z}^{\geq 0}$ **initially** 0

States Q : $val(X)$

Actions A : **external** a, b

Transitions \mathcal{D} : **external** a
 precondition
 $counta < maxcount$
 effect
 $counta := counta + 1$

 external b

Trajectories \mathcal{T} : $\{\wp(\mathbf{x}) \mid \mathbf{x} \in Q\}$

Automaton \mathcal{B}_2

Variables X : **discrete** $maxcount \in \mathbb{Z}^{\geq 0}$ **initially** arbitrary
 discrete $countb \in \mathbb{Z}^{\geq 0}$ **initially** 0

States Q : $val(X)$

Actions A : **external** a, b

Transitions \mathcal{D} : **external** b
 $countb < maxcount$
 effect
 $countb := countb + 1$

 external a

Trajectories \mathcal{T} : $\{\wp(\mathbf{x}) \mid \mathbf{x} \in Q\}$

Figure 12: Automata \mathcal{A}_2 and \mathcal{B}_2

Lemma 5.13 *If $E \subseteq E_{\mathcal{A}}$ then $\text{ActHide}(E, \mathcal{A})$ is a TA.*

Lemma 5.14 *If \mathcal{A} is a TA and $E \subseteq E_{\mathcal{A}}$ then $\text{traces}_{\text{ActHide}(E, \mathcal{A})} = \{\beta \upharpoonright (E_{\mathcal{A}} - E, \emptyset) \mid \beta \in \text{traces}_{\mathcal{A}}\}$.*

The following theorem states that the hiding operation respects the implementation relation.

Theorem 5.15 *Suppose \mathcal{A} and \mathcal{B} are TAs with $\mathcal{A} \leq \mathcal{B}$, and suppose $E \subseteq E_{\mathcal{A}}$. Then $\text{ActHide}(E, \mathcal{A}) \leq \text{ActHide}(E, \mathcal{B})$.*

5.3 Extending Timed Automata with Bounds

In this section, we define a new class of automata, “TA with bounds” where the basic definition of a timed automaton is extended with the notion of a task and a pair of bounds (a lower and an upper bound) for each task. We then define an operation that transforms a given TA with bounds to another TA. This operation supports specifying a system by thinking in terms of tasks and bounds as in the timed automata of Merritt, Modugno, and Tuttle [29] and the phase transition systems of Maler, Manna and Pnueli [28].

In defining the operation for extending timed automata with bounds, we restrict attention to a class of automata where the enabling and disabling of actions during trajectories follow certain rules. Specifically, our operation is defined on automata in which each action is enabled or disabled throughout an entire trajectory, or becomes enabled once during a trajectory and remains so until the end of that trajectory. The given restrictions ensure that the result of applying the operation to a TA is another TA and that the resulting TA satisfies the restrictions.

Let \mathcal{A} be a TA, C a set of actions of \mathcal{A} , and \mathcal{T} the set of trajectories of \mathcal{A} . We say that \mathcal{T} is *well-formed* with respect to C if each $\tau \in \mathcal{T}$ satisfies one of the following conditions:

1. For all $t \in \text{dom}(\tau)$, C is enabled in $\tau(t)$.
2. For all $t \in \text{dom}(\tau)$, C is disabled in $\tau(t)$.
3. There exists $t \in \text{dom}(\tau)$ such that for all $t' \in [0, t)$, C is disabled in $\tau(t')$ and for all $t' \in \text{dom}(\tau) - [0, t)$, C is enabled in $\tau(t')$.

A TA with bounds, $\mathcal{A} = (\mathcal{B}, C, l, u)$ consists of:

- A timed automaton $\mathcal{B} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$.
- A set $C \subseteq E \cup H$ of actions called a *task*; we assume that \mathcal{T} is well-formed with respect to C .

- A lower time bound l and an upper time bound u for C . We require that the following axioms are satisfied for l and u :

B1 $l \in \mathbb{R}^{\geq 0}$ and $u \in \mathbb{R}^{\geq 0} \cup \{\infty\}$.

B2 $l \leq u$.

Lower and upper bounds are used to specify how much time is allowed to pass between the enabling and the performance of an action. If l is the lower bound for a task C , then an action in C must remain enabled at least for l time units before being performed. If u is the upper bound for a task C , then an action in C can remain enabled at most u time units without being performed: it must either be performed or become disabled within u time units.

We now define an operation *Extend*, which transforms a TA \mathcal{A} with bounds to another TA \mathcal{A}' that incorporates the new bounds, in addition to the timing constraints already present in \mathcal{A} . Let $\mathcal{A} = (\mathcal{B}, C, l, u)$ be a TA with bounds where $\mathcal{B} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$. Then *Extend*(\mathcal{A}) is the TA $\mathcal{A}' = (X', Q', \Theta', E', H', \mathcal{D}', \mathcal{T}')$ such that the components of \mathcal{A}' consist of:

- $X' = X \cup \{now, first, last\}$ where:
 1. $now, first,$ and $last$ are new variables that do not appear in X .
 2. now is an analog variable such that $type(now) = \mathbb{R}$.
 3. $first$ and $last$ are discrete variables where $type(first) = \mathbb{R}$ and $type(last) = \mathbb{R} \cup \{\infty\}$.
- $Q' = \{\mathbf{x} \in val(X') \mid \mathbf{x} \upharpoonright X \in Q\}$.
- Θ' consists of all the states $\mathbf{x} \in Q'$ that satisfy the following conditions:
 1. $\mathbf{x} \upharpoonright X \in \Theta$.
 2. $\mathbf{x}(now) = 0$.
 3. $\mathbf{x}(first) = \begin{cases} l & \text{if } C \text{ is enabled in } \mathbf{x} \upharpoonright X, \\ 0 & \text{otherwise.} \end{cases}$
 $\mathbf{x}(last) = \begin{cases} u & \text{if } C \text{ is enabled in } \mathbf{x} \upharpoonright X, \\ \infty & \text{otherwise.} \end{cases}$
- $E' = E$ and $H' = H$. We write $A' \triangleq E' \cup H'$.
- If $a \in (E \cup H)$ then $(\mathbf{x}, a, \mathbf{x}') \in \mathcal{D}'$ exactly if all of the following conditions hold:
 1. $(\mathbf{x} \upharpoonright X) \xrightarrow{a}_{\mathcal{A}} (\mathbf{x}' \upharpoonright X)$.
 2. $\mathbf{x}'(now) = \mathbf{x}(now)$.

3. (a) If $a \in C$, then $\mathbf{x}(first) \leq \mathbf{x}(now)$.
 - (b) If C is enabled both in $\mathbf{x} \upharpoonright X$ and $\mathbf{x}' \upharpoonright X$ and $a \notin C$, then $\mathbf{x}(first) = \mathbf{x}'(first)$ and $\mathbf{x}(last) = \mathbf{x}'(last)$.
 - (c) If C is enabled in $\mathbf{x}' \upharpoonright X$ and either C is not enabled in $\mathbf{x} \upharpoonright X$ or $a \in C$, then $\mathbf{x}'(first) = \mathbf{x}(now) + l$ and $\mathbf{x}'(last) = \mathbf{x}(now) + u$.
 - (d) If C is not enabled in $\mathbf{x}' \upharpoonright X$, then $\mathbf{x}'(first) = 0$ and $\mathbf{x}'(last) = \infty$.
- \mathcal{T}' is a set that consists of all $\tau \in \text{trajs}(X')$ that satisfy the following conditions:
 1. $(\tau \downarrow X) \in \mathcal{T}$.
 2. $d(now) = 1$.
 3. (a) If for all $t \in \text{dom}(\tau)$, C is enabled in $\tau \downarrow X(t)$ then $first$ and $last$ are constant throughout τ .
 - (b) If for all $t \in \text{dom}(\tau)$, C is disabled in $\tau \downarrow X(t)$ then $first$ and $last$ are constant throughout τ .
 - (c) If for all $t' \in [0, t)$, C is disabled in $\tau(t')$ and for all $t' \in \text{dom}(\tau) - [0, t)$, C is enabled in $\tau(t')$ then
 - i. $first$ and $last$ are constant in $[0, t)$.
 - ii. $\tau(t)(first) = \tau(t)(now) + l$ and $\tau(t)(last) = \tau(t)(now) + u$.
 - iii. $first$ and $last$ are constant in $\text{dom}(\tau) - [0, t)$.
 - (d) $now \leq last$.

The transformation is based on the idea of augmenting the state of the original automaton with a variable to represent current time (now) and the earliest time ($first$) and the latest time ($last$) a task can be performed. All these variables represent time in absolute terms. Item 3(a) in the definition of \mathcal{D}' expresses the new lower bound constraint and Item 3(d) in the definition of \mathcal{T}' the new upper bound constraint.

Let \mathcal{A} be a TA with bounds (\mathcal{B}, C, l, u) . In a start state \mathbf{x} of $\text{Extend}(\mathcal{A})$, the variables $first$ and $last$ are initialized to l and u respectively, if C is enabled in \mathbf{x} . If C is not enabled in \mathbf{x} , then $first$ is set to 0 and $last$ is set to ∞ . Items 3(c) in the definition of \mathcal{D}' and 3(c) in the definition of \mathcal{T}' show how the variables $first$ and $last$ are updated. When C becomes newly enabled by a discrete transition or when a C action leads to a state in which C is enabled, $first$ is set to $now + l$ and $last$ is set to $now + u$. The variables $first$ and $last$ are updated similarly when C becomes newly enabled in the course of a trajectory.

Theorem 5.16 *Suppose that $\mathcal{A} = (\mathcal{B}, C, l, u)$ is a TA with bounds. Then $\text{Extend}(\mathcal{A})$ is a TA with a set of trajectories that is well-formed with respect to C .*

Proof: The proof follows from the definitions of TA and the operation Extend . Step 3(a) in the definition of \mathcal{D}' adds a new lower bound constraint, which makes enabling

start at some particular time. Step 3(b) in the definition of \mathcal{T}' , adds a new upper bound constraint, which stops trajectories at a particular time and which does not add any enabling or disabling to trajectories. ■

In the rest of this section, we sometimes speak of variables, states and traces of a TA with bounds. If $\mathcal{A} = (\mathcal{B}, C, l, u)$ is a TA with bounds, variables, states and traces of \mathcal{A} refer to, respectively, the states and the traces of the underlying automaton \mathcal{B} .

Theorem 5.17 *Suppose $\mathcal{A} = (\mathcal{B}, C, l, u)$ is a TA with bounds. Then $\text{traces}_{\text{Extend}(\mathcal{A})} \subseteq \text{traces}_{\mathcal{A}}$.*

Proof: Let $F : Q' \rightarrow Q$ be defined as follows: $F(\mathbf{x}) = \mathbf{x} \upharpoonright X$ where X is the set of internal variables of \mathcal{A} . It is easy to check that F is a refinement from $\text{Extend}(\mathcal{A})$ to \mathcal{A} . By Theorem 4.28 and Corollary 4.24, we conclude that $\text{traces}_{\text{Extend}(\mathcal{A})} \subseteq \text{traces}_{\mathcal{A}}$. ■

Lemma 5.18 *Suppose that \mathcal{A} is a TA with bounds. For any reachable state \mathbf{x} of $\text{Extend}(\mathcal{A})$, if C is enabled in $\mathbf{x} \upharpoonright X$ in \mathcal{A} , then $\mathbf{x}(\text{last}) \leq \mathbf{x}(\text{now}) + u$.*

Proof: Consider a closed execution α of $\text{Extend}(\mathcal{A})$. Using the axioms **T1** and **T2** for trajectories, we can write α as a concatenation of closed execution fragments $\alpha_0 \frown \alpha_1 \frown \dots \frown \alpha_k$ where α_0 is a point trajectory, and each α_i for $i \geq 1$ is either a trajectory or a discrete action surrounded by two point trajectories such that for all $0 \leq i \leq k-1$, $\alpha_i.\text{lstate} = \alpha_{i+1}.\text{fstate}$. We prove the invariant by induction on the length k of the sequence of execution fragments.

For the base case, suppose that C is enabled in $\alpha_0.\text{fstate} \upharpoonright X$. Since α is an execution, we know that $\alpha_0.\text{fstate}$ is a start state of $\text{Extend}(\mathcal{A})$. By definition of $\text{Extend}(\mathcal{A})$, $\alpha_0.\text{fstate}(\text{last}) = u$. Since $\alpha_0.\text{fstate}(\text{now}) = 0$, $\alpha_0.\text{fstate}(\text{last}) \leq \alpha_0.\text{fstate}(\text{now}) + u$, as required.

For the inductive step, we assume that the property is true for the sequence $\alpha_0 \frown \alpha_1 \frown \dots \frown \alpha_k$ and show that it is true in the sequence α_{k+1} in $\alpha_0 \frown \alpha_1 \frown \dots \frown \alpha_k \frown \alpha_{k+1}$. There are two cases to consider depending on whether α_{k+1} is a discrete action surrounded by two point trajectories or a trajectory.

1. α_{k+1} is an action a surrounded by two point trajectories. Suppose that C is enabled in $\alpha_{k+1}.\text{lstate}$. There are two subcases to consider:

- (a) C is enabled in $\alpha_k.\text{lstate} \upharpoonright X$ and $a \notin C$.

Then, $\alpha_{k+1}.\text{fstate}(\text{last}) = \alpha_k.\text{fstate}(\text{last})$ and $\alpha_{k+1}.\text{fstate}(\text{now}) = \alpha_k.\text{fstate}(\text{now})$. By inductive hypothesis, $\alpha_k.\text{lstate}(\text{last}) \leq \alpha_k.\text{lstate}(\text{now}) + u$. Therefore, $\alpha_{k+1}.\text{lstate}(\text{last}) \leq \alpha_{k+1}.\text{lstate}(\text{now}) + u$, as needed.

(b) C is disabled in $\alpha_k.lstate \upharpoonright X$ or $a \in C$.

Then, by definition of $\text{Extend}(\mathcal{A})$, $\alpha_{k+1}.lstate(last) = \alpha_{k+1}.lstate(now) + u$, which suffices.

2. α_{k+1} is a trajectory.

Suppose that C is enabled in $\alpha_{k+1}.lstate \upharpoonright X$ in \mathcal{A} . There are two subcases to consider:

(a) C is enabled in $\alpha_{k+1}.fstate \upharpoonright X$ in \mathcal{A} .

By inductive hypothesis $\alpha_{k+1}.fstate(last) \leq \alpha_{k+1}.fstate(now) + u$. By the well-formedness assumption, we know that C must be enabled throughout α_{k+1} and by definition of $\text{Extend}(\mathcal{A})$ $last$ is constant throughout α_{k+1} . Since the value of now increases, it is easy to see that $\alpha_{k+1}.lstate(last) \leq \alpha_{k+1}.lstate(now) + u$.

(b) C is disabled in $\alpha_{k+1}.fstate \upharpoonright X$ in \mathcal{A} .

Then, since it is enabled in $\alpha_{k+1}.lstate \upharpoonright X$ by the well-formedness assumption, it becomes enabled at some point t in the domain of α_{k+1} and remains enabled thereafter. Therefore, $\alpha_{k+1}(t)(last) = \alpha_{k+1}(t)(now) + u$, by definition of $\text{Extend}(\mathcal{A})$. Since $last$ remains constant after it is set and the value of now increases, $\alpha_{k+1}.lstate(last) \leq \alpha_{k+1}.lstate(now) + u$ holds. ■

The theorem below shows that the executions of an automaton obtained by applying the transformation Extend to a TA with bounds respect the time bounds specified by the lower bound l and the upper bound u .

Theorem 5.19 *Let $\mathcal{A} = (\mathcal{B}, C, l, u)$ be a TA with bounds. Then,*

1. *There does not exist a closed execution fragment α of $\text{Extend}(\mathcal{A})$ from a reachable state, where $\alpha.ltime > u$, C is enabled in \mathcal{A} in all the states of $\alpha \upharpoonright (A, X)$ and no action in C occurs in α .*
2. *There does not exist a closed execution fragment α of $\text{Extend}(\mathcal{A})$ from a reachable state, where $\alpha.ltime < l$, such that C is not enabled in \mathcal{A} in the first state of $\alpha \upharpoonright (A, X)$ and an action in C occurs in α .*

Proof:

1. Suppose, for the sake of contradiction, that there exists a closed execution fragment $\alpha = \tau_0 a_1 \tau_1 a_2 \dots \tau_n$ of $\text{Extend}(\mathcal{A})$ from a reachable state, where $\alpha.ltime > u$, C is enabled in \mathcal{A} in all the states of $\alpha \upharpoonright (A, X)$ and none of the a_i in α is in C . By definition of trajectories for $\text{Extend}(\mathcal{A})$ it must be the case that $\alpha.lstate(now) \leq \alpha.lstate(last)$.

Since C is enabled in \mathcal{A} in all states in α , by Lemma 5.18 we have $\alpha.fstate(last) \leq \alpha.fstate(now) + u$. By definition of $\text{Extend}(\mathcal{A})$, $last$ remains constant throughout α ; therefore, $\alpha.lstate(last) = \alpha.fstate(last)$. Since $\alpha.fstate(last) \leq \alpha.fstate(now) + u$, it follows that $\alpha.lstate(last) \leq \alpha.fstate(now) + u$. By definition of α , we have $\alpha.lstate(now) = \alpha.fstate(now) + \alpha.ltime$. It follows that $\alpha.fstate(now) + \alpha.ltime \leq \alpha.fstate(now) + u$. This implies $\alpha.ltime \leq u$. But this gives us the needed contradiction since $\alpha.ltime > u$.

2. We assume that α is a closed execution fragment of $\text{Extend}(\mathcal{A})$ from a reachable state where $\alpha.ltime < l$, such that C is not enabled in \mathcal{A} in the first state of α and an action in C occurs in α . Let $(\mathbf{x}, a, \mathbf{x}')$ be the first discrete transition of $\text{Extend}(\mathcal{A})$ in α such that $a \in C$. We show that the condition $\mathbf{x}(first) \leq \mathbf{x}(now)$, which has to hold for the discrete transition to occur, cannot be true, hence arrive at a contradiction. By Theorem 5.16, the set of trajectories of $\text{Extend}(\mathcal{A})$ is well-formed with respect to C . Therefore, C can become enabled by either a discrete transition or during a trajectory, and remains enabled until the occurrence of $(\mathbf{x}, a, \mathbf{x}')$.

- (a) C becomes enabled by a discrete transition and remains enabled in \mathcal{A} until the occurrence of $(\mathbf{x}, a, \mathbf{x}')$.

Let $(\mathbf{y}, b, \mathbf{y}')$ be the discrete transition of \mathcal{A} that enables C . By item 3(c) in the definition of \mathcal{D}' we know that $first$ is set to $\mathbf{y}(now) + l$ when C becomes enabled. By item 3(b) in the definition of \mathcal{D}' and 3(a) in the definition of \mathcal{T}' , we know that it remains constant so that $\mathbf{x}(first) = \mathbf{y}(now) + l$. Since $(\mathbf{x}, a, \mathbf{x}')$ is a discrete transition of $\text{Extend}(\mathcal{A})$, it must be the case that $\mathbf{x}(first) \leq \mathbf{x}(now)$. Since $\mathbf{x}(now) \leq \mathbf{y}(now) + \alpha.ltime$ and $\mathbf{x}(first) = \mathbf{y}(now) + l$ it follows that $\mathbf{y}(now) + l \leq \mathbf{y}(now) + \alpha.ltime$. But we know by assumption that $\alpha.ltime < l$ which gives the needed contradiction.

- (b) C becomes enabled at some point in the course of a trajectory τ and remains enabled in \mathcal{A} until the occurrence of $(\mathbf{x}, a, \mathbf{x}')$.

Let \mathbf{y} be a state in the range of τ where C becomes enabled. By item 3(c) in the definition of \mathcal{T}' we know that $first$ is set to $\mathbf{y}(now) + l$ when C becomes enabled and it remains constant in τ so that $\mathbf{x}(first) = \mathbf{y}(now) + l$. By item 3(b) in the definition of \mathcal{D}' and 3(a) in the definition of \mathcal{T}' , we know that $first$ remains constant until the occurrence of $(\mathbf{x}, a, \mathbf{x}')$. Since $(\mathbf{x}, a, \mathbf{x}')$ is a discrete transition of $\text{Extend}(\mathcal{A})$, it must be the case that $\mathbf{x}(first) \leq \mathbf{x}(now)$. Since $\mathbf{x}(now) \leq \mathbf{y}(now) + \alpha.ltime$ and $\mathbf{x}(first) = \mathbf{y}(now) + l$ it follows that $\mathbf{y}(now) + l \leq \mathbf{y}(now) + \alpha.ltime$. But we know by assumption that $\alpha.ltime < l$ which gives the needed contradiction.

■

Example 5.20 (Fischer’s mutual exclusion algorithm specified using tasks and bounds)

In Example 4.5 we presented the specification of Fischer’s mutual exclusion algorithm as a TA. This example illustrates an alternative way of specifying the same algorithm by using a TA with bounds.

Recall that, formally, we define a TA with bounds as a TA augmented with a single task along with lower and upper bounds for that task. The automaton in Figure 13 is, however, augmented with a set of tasks and bounds. This is for notational convenience and the automaton in Figure 13 should be viewed as the automaton representing the cumulative result of adding in successive steps two tasks for each $i \in I$. We assume that *Extend* is applied once for each task. That is, we start with the timing-independent version of *FischerME*, apply *Extend* to the automaton augmented with the task $\{set_i\}$ to add the lower bound 0 and the upper bound u_{set} , then apply *Extend* to the resulting automaton augmented with $\{check_i\}$ to add the lower bound l_{check} and the upper bound ∞ . Such two successive applications are allowed since the result of the first application of *Extend* satisfies the the well-formedness conditions for the set of trajectories.

The result of these successive applications yields an automaton similar to the one in Example 4.5. The only difference is that the mechanical application of the transformation would reset the value of $firstcheck[i]$ to 0 as an effect of $check_i$ while we do not reset $firstcheck[i]$ explicitly in 4.5, when it becomes disabled. This is because we make use of the facts that the value of $firstcheck[i]$ is used only in determining whether $check_i$ is enabled and that $check_i$ becomes enabled only in the poststate of set_i which also sets the value of $firstcheck[i]$. Note that this discrepancy does not give rise to any difference in the behaviors of the two automata. ■

5.4 Untiming

We define an “untiming” operation that transforms a timed automaton to an untimed automaton of the kind defined in Section 2.5. The idea behind this operation is to reduce the state space of a timed automaton by identifying those states that are equivalent in the sense that they give rise to similar discrete behavior. The executions of the untimed automaton obtained as a result of applying the untiming operation to a TA, \mathcal{A} , preserve the order of discrete actions of \mathcal{A} but forget the possible time passage between them. This operation has its roots in a similar operation defined in [6, 4] but we do not deal with the finiteness of the resulting state space and ease of reachability analysis, as those papers do. Instead, we aim to understand the main ideas of the untiming operation of [6, 4] using our more general framework.

The untiming operation uses the notion of congruence defined below to determine equivalence classes of states.

Type $PcValue = \text{enumeration } rem, test, set, check, leavetry, crit, leaveexit$

Automaton $FischerME2(u_{set}, l_{check}, I)$ **where** $u_{set} \in \mathbb{R}^{\geq 0}, l_{check} \in \mathbb{R}^{\geq 0}, u_{set} \leq l_{check}$

Variables X : **discrete** pc , an array of elements of $PcValue$ indexed by I
initially $\forall i \in I. pc[i] = rem$
discrete $x \in I \cup \{\perp\}$ **initially** $x = \perp$

States Q : $val(X)$

Actions A : **external** $try_i, crit_i, exit_i, rem_i$
internal $test_i, set_i, check_i, reset_i$ **where** $i \in I$

Transitions \mathcal{D} :

<p>external try_i precondition $pc[i] = rem$ effect $pc[i] := test$</p>	<p>external $crit_i$ precondition $pc[i] = leavetry$ effect $pc[i] := crit$</p>
<p>internal $test_i$ precondition $pc[i] = test$ effect if $x = \perp$ then $pc[i] := set$</p>	<p>external $exit_i$ precondition $pc[i] = crit$ effect $pc[i] := reset$</p>
<p>internal set_i precondition $pc[i] = set$ effect $x := i$ $pc[i] := check$</p>	<p>internal $reset_i$ precondition $pc[i] = reset$ effect $x := \perp$ $pc[i] := leaveexit$</p>
<p>internal $check_i$ precondition $pc[i] = check$ effect if $x = i$ then $pc[i] := leavetry$ else $pc[i] := test$</p>	<p>external rem_i precondition $pc[i] = leaveexit$ effect $pc[i] := rem$</p>

Trajectories \mathcal{T} : $\{\tau \in \text{trajs}(X) \mid pc \text{ and } x \text{ constant in } \tau\}$

Tasks C : $\forall i \in I. \{set_i\}, \{check_i\}$

Bounds B : $\forall i \in I. \text{lower}(\{set_i\}) = 0, \text{upper}(\{set_i\}) = u_{set}$
 $\forall i \in I. \text{lower}(\{check_i\}) = l_{check}, \text{upper}(\{check_i\}) = \infty$

Figure 13: FischerME with bounds

5.4.1 State Congruence

Let $\mathcal{A} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$ be a TA. An equivalence relation $R \subseteq Q \times Q$ is a *congruence* for \mathcal{A} if, for all actions $a \in (E \cup H)$ and trajectories $\tau \in \mathcal{T}$ the following hold:

1. If $\mathbf{x} R \mathbf{y}$ and $\mathbf{x} \in \Theta$ then $\mathbf{y} \in \Theta$.
2. If $\mathbf{x} R \mathbf{y}$ and $\mathbf{x} \xrightarrow{a} \mathbf{x}'$ then there exists a state \mathbf{y}' such that $\mathbf{y} \xrightarrow{a} \mathbf{y}'$ and $\mathbf{x}' R \mathbf{y}'$.
3. If $\mathbf{x} R \mathbf{y}$, and $\mathbf{x} \xrightarrow{\tau} \mathbf{x}'$ then there exists a state \mathbf{y}' and a trajectory τ' such that $\mathbf{y} \xrightarrow{\tau'} \mathbf{y}'$ and $\mathbf{x}' R \mathbf{y}'$.

The relation R partitions Q into equivalence classes. In the rest of this section, we use $[\mathbf{x}]$ to denote the equivalence class of $\mathbf{x} \in Q$, that is $[\mathbf{x}] = \{\mathbf{y} \mid \mathbf{x} R \mathbf{y}\}$.

5.4.2 Definition of the Untiming Operation

Given a TA $\mathcal{A} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$ and a congruence $R \subseteq Q \times Q$ for \mathcal{A} , the untiming operation yields an untimed automaton $\text{Untime}(\mathcal{A}, R) = (Q', \Theta', E', H', \mathcal{D}')$ where

- $Q' = \{[\mathbf{x}] \mid \mathbf{x} \in Q\}$.
- $\Theta' = \{[\mathbf{x}] \mid \mathbf{x} \in \Theta\}$.
- $E' = E$.
- $H' = H \cup \{\pi\}$ where π is a special action representing time passage.
- $\mathcal{D}' \subseteq Q' \times A' \times Q'$ where $A' \triangleq E' \cup H'$ such that
 1. $s \xrightarrow{a} s' \in \mathcal{D}'$ if and only if there exists $(\mathbf{x}, a, \mathbf{x}') \in \mathcal{D}$ where $[\mathbf{x}] = s$ and $[\mathbf{x}'] = s'$.
 2. $s \xrightarrow{\pi} s' \in \mathcal{D}'$ if and only if there exists $\tau \in \mathcal{T}$ where τ is closed, $[\tau.fstate] = s$ and $[\tau.lstate] = s'$.

Example 5.21 ($\text{Untime}(AD, R)$) In this example we define a congruence for the automaton AD from Example 4.19 and give the result of applying the untiming operation to AD by using this congruence. Let I be the set of intervals $\{(0, 1), (1, \infty)\}$. Let R be an equivalence relation defined as follows. $\mathbf{x} R \mathbf{y}$ if the following conditions hold:

1. $\mathbf{x} \upharpoonright X_d = \mathbf{y} \upharpoonright X_d$.
2. For every $x \in X_c$, either $\mathbf{x}(x), \mathbf{y}(x) \in J$ for some $J \in I$ or $\mathbf{x}(x) = \mathbf{y}(x) = i$ for some integer i .

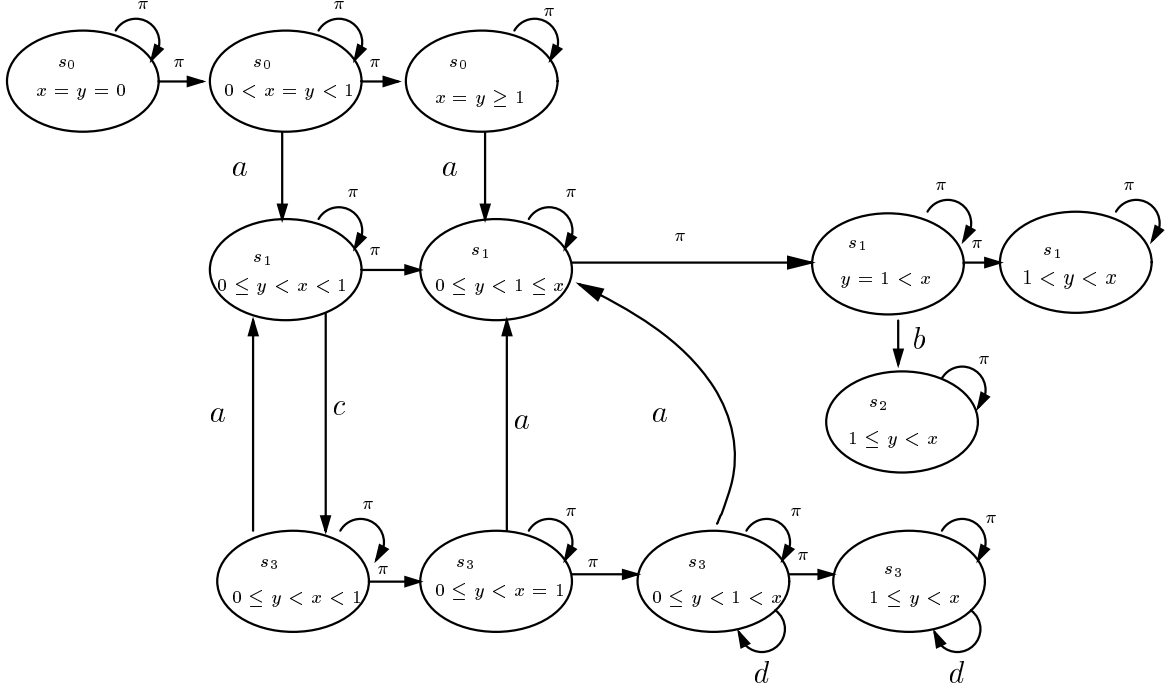


Figure 14: $\text{Untime}(AD, R)$

3. For every $z, w \in X_c$, $\mathbf{x}(z) > \mathbf{x}(w)$ if and only if $\mathbf{y}(z) > \mathbf{y}(w)$.

R is a congruence for the automaton AD from Example 4.19. Figure 14 contains a graphical representation of $\text{Untime}(AD, R)$. Each node in the graph represents a state of $\text{Untime}(AD, R)$, that is, an equivalence class of states of AD with respect to R . The annotations within the nodes are used to define the equivalence class. For example, a node that is annotated with s_0 and $x = y = 0$ denotes the set of states $\{\mathbf{x} \in Q_{AD} \mid \mathbf{x}(\text{loc}) = s_0, \mathbf{x}(x) = 0, \text{ and } \mathbf{x}(y) = 0\}$.

■

5.4.3 Basic Results

In this section we present some results that establish a correspondence between the executions of a TA and those of the corresponding untimed automaton.

The lemma below states that the trace of discrete events in an execution fragment of a timed automaton is also exhibited by some execution fragment of the corresponding untimed automaton.

Lemma 5.22 *Suppose \mathcal{A} is a TA and R is a congruence for \mathcal{A} . If α is an execution fragment of \mathcal{A} , then $\text{Untime}(\mathcal{A}, R)$ has an execution fragment α' such that $\alpha'.fstate = [\alpha.fstate]$ and $\text{trace}(\alpha') = \text{actions}(\text{trace}(\alpha))$.*

Proof: We consider the following cases:

1. α is an infinite sequence.

Using axioms **T1** and **T2** we can write α as an infinite concatenation $\alpha_0 \hat{\ } \alpha_1 \hat{\ } \dots$, in which each execution fragment α_i is either a trajectory with $\alpha_i.ltime > 0$ or a single discrete action surrounded by two point trajectories, and for every $i \geq 0$, $\alpha_i.lstate = \alpha_{i+1}.fstate$.

We define a sequence $\alpha'_0 \alpha'_1 \dots$ of execution fragments of $\text{Untime}(\mathcal{A}, R)$ such that

- (a) If α_i is a trajectory, then $\alpha'_i = (s, \pi, s')$ where $s = [\alpha_i.fstate]$ and $s' = [\alpha_i.lstate]$ (recall that we use $[\mathbf{x}]$ to denote the equivalence class of \mathbf{x} with respect to R).
- (b) If α_i is a single discrete action a surrounded by two point trajectories, then $\alpha'_i = (s, a, s')$ where $s = [\alpha_i.fstate]$, $s' = [\alpha_i.lstate]$.

It is immediate from the definition of $\text{Untime}(\mathcal{A}, R)$ in Section 5.4.2 that each α'_i constructed above is an execution fragment of $\text{Untime}(\mathcal{A}, R)$ and that $\alpha'.fstate = [\alpha.fstate]$. By definitions of concatenation and execution fragments for untimed automata from Section 2.5 we have that $\alpha'_0 \hat{\ } \alpha'_1 \hat{\ } \dots$ is an execution fragment of $\text{Untime}(\mathcal{A}, R)$. By definitions of the operators *trace* for untimed automata from Section 2.5, and for timed automata from Section 4, and *discrete* from Section 3 we have $\text{trace}(\alpha') = \text{actions}(\text{trace}(\alpha))$, as needed.

2. α is a finite sequence ending with a closed trajectory.
Similar to the first case.

3. α is a finite sequence ending with an open trajectory.
The sequence α' can be constructed similarly to the first case except for the last trajectory τ_n in α . Taking α'_n to be the empty sequence gives the required result.

■

Corollary 5.23 *Suppose \mathcal{A} is a TA and R is a congruence for \mathcal{A} . If α is an execution of \mathcal{A} , then $\text{Untime}(\mathcal{A}, R)$ has an execution α' such that $\text{trace}(\alpha') = \text{actions}(\text{trace}(\alpha))$.*

Proof: Let α be an execution of \mathcal{A} . We know by Lemma 5.22 that $\text{Untime}(\mathcal{A}, R)$ has an execution α' such that $\text{trace}(\alpha') = \text{actions}(\text{trace}(\alpha))$ and $\alpha'.fstate = [\alpha.fstate]$. Since α is an execution of \mathcal{A} , $\alpha.fstate \in Q_{\mathcal{A}}$. Then by the definition in Section 5.4.2, $\alpha'.fstate \in \Theta'$ and therefore α' is an execution of $\text{Untime}(\mathcal{A}, R)$, as needed. ■

The following lemma states that, for every execution fragment α of $\text{Untime}(\mathcal{A}, R)$ and for every state \mathbf{x} that is in the equivalence class represented by the first state of α , it is possible to derive an execution fragment of \mathcal{A} from \mathbf{x} that exhibits the same discrete trace as $\text{Untime}(\mathcal{A}, R)$.

Lemma 5.24 *Suppose \mathcal{A} is a TA and R is a congruence for \mathcal{A} . If α is an execution fragment of $\text{Untime}(\mathcal{A}, R)$ and \mathbf{x} is a state of \mathcal{A} such that $[\mathbf{x}] = \alpha.fstate$, then \mathcal{A} has an execution fragment α' from \mathbf{x} such that $\text{trace}(\alpha) = \text{actions}(\text{trace}(\alpha'))$.*

Proof:

1. α is an infinite sequence of the form $s_0 a_1 s_1 a_2 s_2 \dots$

The sequence α can be written as the concatenation $\alpha_0 \frown \alpha_1 \frown \alpha_2 \dots$ of execution fragments (s_i, a_{i+1}, s_{i+1}) for $i \geq 0$. We define α' inductively as the concatenation $\alpha'_0 \frown \alpha'_1 \frown \alpha'_2 \dots$ where $[\alpha'_0.fstate] = \alpha.fstate$ and for every $i \geq 0$, $\alpha'_i.lstate = \alpha'_{i+1}.fstate$ and $[\alpha'_i.lstate] = s_i$ as follows:

- (a) $\alpha'_0 = \wp(\mathbf{x})$. By axiom **T0**, α'_0 is an execution fragment of \mathcal{A} . Since $\alpha'_0.fstate = \mathbf{x}$ by construction of α'_0 and $[\mathbf{x}] = \alpha.fstate$ by definition of \mathbf{x} , we have $[\alpha'_0.fstate] = \alpha.fstate$. Since $\alpha'_0.lstate = \mathbf{x}$ by construction of α'_0 and $[\mathbf{x}] = \alpha.fstate$ by definition of \mathbf{x} and $\alpha.fstate = s_0$ by the assumed structure of α we have $[\alpha'_0.lstate] = s_0$.
- (b) For $i \geq 1$, if α_{i-1} is (s_{i-1}, a_i, s_i) where $a_i \in (A' \setminus \{\pi\})$, then define α'_i to be $\wp(\alpha'_{i-1}.lstate) a_i \wp(\mathbf{y})$ where $(\alpha'_{i-1}.lstate, a_i, \mathbf{y}) \in \mathcal{D}$ and $[\mathbf{y}] = s_i$. We need to show that \mathcal{A} has such an execution fragment α'_i . For $i \geq 1$, consider $\alpha_{i-1} = (s_{i-1}, a_i, s_i)$. By definition of $\text{Untime}(\mathcal{A}, R)$ in Section 5.4.2, there must be some $(\mathbf{z}, a_i, \mathbf{z}') \in \mathcal{D}$ such that $[\mathbf{z}] = s_{i-1}$ and $[\mathbf{z}'] = s_i$. By inductive hypothesis $[\alpha'_{i-1}.lstate] = s_{i-1}$. Since $[\alpha'_{i-1}.lstate] = s_{i-1} = [\mathbf{z}]$ we know by the definition of state congruence in Section 5.4.1 that there exists \mathbf{y} such that $(\alpha'_{i-1}.lstate, a_i, \mathbf{y}) \in \mathcal{D}$ and $[\mathbf{y}] = [\mathbf{z}'] = s_i$. Therefore, $\alpha'_i = \wp(\alpha'_{i-1}.lstate) a_i \wp(\mathbf{y})$ is an execution fragment of \mathcal{A} where $\alpha'_i.fstate = \alpha'_{i-1}.lstate$ and $[\alpha'_i.lstate] = s_i$.
- (c) For $i \geq 1$, if α_{i-1} is (s_{i-1}, a_i, s_i) where a_i is the π action, then define α'_i to be τ where $\tau \in \mathcal{T}$, $\tau.fstate = \alpha'_{i-1}.lstate$ and $[\tau.lstate] = s_i$. We need to show that \mathcal{A} has such an execution fragment α'_i . For $i \geq 1$, consider $\alpha_{i-1} = (s_{i-1}, a_i, s_i)$. By definition of $\text{Untime}(\mathcal{A}, R)$ in Section 5.4.2, there must be some trajectory τ' such that τ' is closed, $[\tau'.fstate] = s_{i-1}$ and $[\tau'.lstate] = s_i$. By inductive hypothesis $[\alpha'_{i-1}.lstate] = s_{i-1}$. Since $[\alpha'_{i-1}.lstate] = s_{i-1} = [\tau'.fstate]$ we know by the definition of state congruence in Section 5.4.1 that there exists τ where $\tau.fstate = \alpha'_{i-1}.lstate$ and $[\tau.lstate] = s_i = [\tau'.lstate]$. Therefore, $\alpha'_i = \tau$ is an execution fragment of \mathcal{A} where $\alpha'_i.fstate = \alpha'_{i-1}.lstate$ and $[\alpha'_i.lstate] = s_i$.

By construction of α' , we have $\alpha.fstate = [\alpha'.fstate]$. Since $\alpha'_i.lstate = \alpha'_{i+1}.fstate$ for all $i \geq 0$, we know by Lemma 4.7 that $\alpha' = \alpha'_0 \frown \alpha'_1 \frown \alpha'_2 \dots$ is an execution fragment of \mathcal{A} . It is easy to check that $trace(\alpha) = actions(trace(\alpha'))$.

2. α is a finite sequence of the form $s_0 a_1 s_1 a_2 s_2 \dots s_n$.

The proof is similar to the previous case. ■

Corollary 5.25 *Suppose \mathcal{A} is a TA and R is a congruence for \mathcal{A} . If α is an execution of $Untime(\mathcal{A}, R)$, and \mathbf{x} is a state of \mathcal{A} such that $[\mathbf{x}] = \alpha.fstate$, then \mathcal{A} has an execution α' from \mathbf{x} such that $trace(\alpha) = actions(trace(\alpha'))$.*

Proof: Let α be an execution of $Untime(\mathcal{A}, R)$ and \mathbf{x} be a state of \mathcal{A} such that $[\mathbf{x}] = \alpha.fstate$. By Lemma 5.24, we know that \mathcal{A} has an execution fragment α' from \mathbf{x} such that $trace(\alpha) = actions(trace(\alpha'))$. Since α is an execution, $\alpha.fstate \in \Theta'$. By the definition of $Untime(\mathcal{A}, R)$ in Section 5.4.2, we know that $\mathbf{x} \in \Theta$, and therefore α' is an execution of \mathcal{A} , as needed. ■

5.4.4 An Equivalence Relation for Alur-Dill Automata

In [6, 4] Alur and Dill present a region construction technique that allows an infinite state space to be reduced to a finite state space by using an equivalence relation on states. Our untiming operation is based on a similar idea. It aims to reduce the state space by identifying those states that exhibit “equivalent” behavior. Our operation, however, does not use a fixed equivalence relation. Rather, it is parameterized by equivalence relations that meet our congruence criteria.

In this section we formulate the equivalence relation of Alur and Dill presented in [6] in our framework and show that it is a congruence for an AD automaton under a certain set of assumptions. Recall that our definition of AD automata (see Section 4.3.2) does not impose any restrictions on the form of clock constraints. Adopting such a general definition and seeking a minimal set of assumptions required for the proof allows us to identify which restrictions were incorporated into the model of Alur and Dill mainly to ensure that the resulting region automata have a finite state space.

Let $\mathcal{A} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$ be an AD timed automaton where X is partitioned into two sets: X_d of discrete variables and X_c of clock variables. Let I be the set of intervals and P be the set of points in the time domain $\mathbb{T} = \mathbb{R}$ defined as follows:

$$I = \{(t_1, t_1 + 1) \mid t_1 \in \mathbb{N}\}.$$

$$P = \mathbb{N}.$$

Now, we define an equivalence relation \sim over Q . In our definition we use the notation $fr(\mathbf{v})$ for the fractional part of a value \mathbf{v} . Two states $\mathbf{x}, \mathbf{y} \in Q$ are related, written $\mathbf{x} \sim \mathbf{y}$, if the following conditions hold:

1. $\mathbf{x} \upharpoonright X_d = \mathbf{y} \upharpoonright X_d$.
2. For every $x \in X_c$, either $\{\mathbf{x}(x), \mathbf{y}(x)\} \subseteq J$ for some $J \in I$ or $\mathbf{x}(x) = \mathbf{y}(x) = i$ for some $i \in P$.
3. For every $z, w \in X_c$, $fr(\mathbf{x}(z)) > fr(\mathbf{x}(w))$ if and only if $fr(\mathbf{y}(z)) > fr(\mathbf{y}(w))$.

The first property in the definition of \sim requires that a discrete variable have the same value in two related states. The second property involves clock variables. If a clock variable has a value that falls between a pair of consecutive integers, then its value must be between the same integers in a related state. Likewise, if a clock variable has an integer value, it must have the same value in a related state. The third property states that the ordering of the fractional parts of different clock variables must be the same across related states.

The following theorem states that the relation \sim defined above is a congruence for an AD automaton \mathcal{A} if the same discrete actions can be performed from two equivalent states with the same effect.

Theorem 5.26 *Assume for an AD automaton \mathcal{A} that whenever $\mathbf{x} \sim \mathbf{y}$ for two states $\mathbf{x}, \mathbf{y} \in Q$, and $\mathbf{x} \xrightarrow{a} \mathbf{x}' \in \mathcal{D}$, then there exists $\mathbf{y} \xrightarrow{a} \mathbf{y}' \in \mathcal{D}$ such that*

- $\mathbf{x}' \upharpoonright X_d = \mathbf{y}' \upharpoonright X_d$.
- For every $x \in X_c$, $\mathbf{x}'(x) = 0$ if and only if $\mathbf{y}'(x) = 0$.

Then relation \sim is a congruence for \mathcal{A} .

Proof: We establish the three properties of congruence defined in Section 5.4.1 for the relation \sim .

1. Suppose $\mathbf{x} \sim \mathbf{y}$ and $\mathbf{x} \in \Theta$. By definition of AD automata from Section 4.3.2, if $\mathbf{x} \in \Theta$ then for all $x \in X_c$, $\mathbf{x}(x) = 0$. Since $\mathbf{x} \sim \mathbf{y}$, for all $x \in C$, we have $\mathbf{y}(x) = 0$, and $\mathbf{x} \upharpoonright X_d = \mathbf{y} \upharpoonright X_d$. It follows that $\mathbf{x} = \mathbf{y}$, and therefore $\mathbf{y} \in \Theta$ as needed.
2. Suppose $\mathbf{x} \sim \mathbf{y}$ and $\mathbf{x} \xrightarrow{a} \mathbf{x}'$ where a is a discrete action. By assumption there exists \mathbf{y}' such that $\mathbf{y} \xrightarrow{a} \mathbf{y}'$. It remains to show that $\mathbf{x}' \sim \mathbf{y}'$. We do this by establishing the three properties in the definition of \sim .

- (a) The first property is immediate from the assumptions.
- (b) For the second property, we are required to show that for all $x \in X_c$, either $\mathbf{x}'(x)$ and $\mathbf{y}'(x)$ are in the same interval or have the same integer value. We fix x and consider two cases:
- i. $\mathbf{x}'(x) = 0$.
By assumption $\mathbf{x}'(x) = 0$ if and only if $\mathbf{y}'(x) = 0$. Clearly, $\mathbf{x}'(x)$ and $\mathbf{y}'(x)$ have the same integer value 0.
 - ii. $\mathbf{x}'(x) \neq 0$.
By definition of AD automata from Section 4.3.2, $\mathbf{x}'(x) = \mathbf{x}(x)$. Since $\mathbf{x}'(x) \neq 0$ if and only if $\mathbf{y}'(x) \neq 0$ by assumption, we have $\mathbf{y}'(x) \neq 0$, and by definition of AD automata we have $\mathbf{y}'(x) = \mathbf{y}(x)$. Since $\mathbf{x} \sim \mathbf{y}$ by hypothesis, $\mathbf{y}(x)$ and $\mathbf{x}(x)$ are in the same interval. Since $\mathbf{y}'(x) = \mathbf{y}(x)$ and $\mathbf{x}(x) = \mathbf{x}'(x)$, $\mathbf{x}'(x)$ and $\mathbf{y}'(x)$ are in the same interval, as needed.
- (c) For the third property, we are required to show that for any $z, w \in C$, the ordering between the fractional parts of z and w in \mathbf{x}' is preserved in \mathbf{y}' . For a fixed z and a fixed w consider the following cases:
- i. Neither z nor w is reset by action a .
Then, $\mathbf{x}'(z) = \mathbf{x}(z)$ and $\mathbf{x}'(w) = \mathbf{x}(w)$. Since $\mathbf{x} \sim \mathbf{y}$, we know that $fr(\mathbf{x}(z)) > fr(\mathbf{x}(w))$ if and only if $fr(\mathbf{y}(z)) > fr(\mathbf{y}(w))$. It follows that $fr(\mathbf{x}'(z)) > fr(\mathbf{x}'(w))$ if and only if $fr(\mathbf{y}'(z)) > fr(\mathbf{y}'(w))$, as needed.
 - ii. Both z and w are reset by action a .
By assumption we have $\mathbf{x}'(z) = 0$ if and only if $\mathbf{y}'(z) = 0$ and $\mathbf{x}'(w) = 0$ if and only if $\mathbf{y}'(w) = 0$. Since $fr(\mathbf{x}'(z)) = fr(\mathbf{x}'(w)) = fr(\mathbf{y}'(z)) = fr(\mathbf{y}'(w)) = 0$, it is obvious that the ordering between the fractional parts of the clocks in \mathbf{x}' is preserved in \mathbf{y}' .
 - iii. One of the clocks is reset by action a .
Without loss of generality, let the clock that is reset be z . That is, $\mathbf{x}'(z) = 0$ and $\mathbf{x}'(w) = \mathbf{x}(w)$. Then, either $fr(\mathbf{x}'(w)) = 0$ or $fr(\mathbf{x}'(w)) \neq 0$. First, suppose $fr(\mathbf{x}'(w)) = 0$. Then, $fr(\mathbf{x}'(z)) = fr(\mathbf{x}'(w))$. Since $fr(\mathbf{x}'(w)) = 0$, $\mathbf{x}'(w) = v$ for an integer v . By case (b), we have $\mathbf{y}'(w) = v$ and hence $fr(\mathbf{y}'(w)) = 0$. It follows that $fr(\mathbf{y}'(z)) = fr(\mathbf{y}'(w))$. Now, suppose that $fr(\mathbf{x}'(w)) \neq 0$. Then $fr(\mathbf{x}'(z)) < fr(\mathbf{x}'(w))$. By assumption which says for all $x \in X_c$, $\mathbf{x}'(x) = 0$ if and only if $\mathbf{y}'(x) = 0$, we have $\mathbf{y}'(z) = 0$. Since $fr(\mathbf{x}'(w)) \neq 0$, by the same assumption we get $\mathbf{y}'(w) \neq 0$. It follows that $fr(\mathbf{y}'(z)) < fr(\mathbf{y}'(w))$. Hence, we have shown that the ordering between the fractional parts of the clocks in \mathbf{x}' is preserved in \mathbf{y}' .
3. Suppose $\mathbf{x} \sim \mathbf{y}$ and $\mathbf{x} \xrightarrow{\tau} \mathbf{x}'$ where τ is a trajectory. We need to show that we can find trajectory τ' such that $\mathbf{x}' \sim \mathbf{y}'$ where $\mathbf{y}'(x) = \mathbf{y}(x) + \tau'.ltime$ for all $x \in X_c$. We do this by establishing the three properties in the definition of \sim .
- (a) The first property is immediate from the assumption.

- (b) For the second property, we are required to show that for all $x \in X_c$, either $\mathbf{x}'(x)$ and $\mathbf{y}'(x)$ are in the same interval or have the same integer value. We consider the following cases:
- i. Zero time passage ($\tau.ltime = 0$).
Clearly, τ' with $\tau'.ltime = 0$ results in $\mathbf{y}' = \mathbf{y}$. Since $\mathbf{x} \sim \mathbf{y}$ by hypothesis, we have $\mathbf{x}' \sim \mathbf{y}'$, as needed.
 - ii. $\tau.ltime > 0$ and τ does not make any clock reach an integer boundary.
 - A. Some clocks remain in the same interval.
Let *Cross* be the set of clocks that crossed to a new interval and let *NotCross* be the set of clocks that did not cross to a new interval. We need to make sure that τ' that we choose makes all elements of *Cross* cross to a new interval in \mathbf{y}' and all elements of *NotCross* remain in the same interval, while preserving the ordering of fractional parts of clock values across two equivalent states. Consider the set $\{t - \mathbf{y}(z) \mid z \in \text{Cross}, \mathbf{x}'(z) \in (t, t + 1)\}$ and define m to be the maximum element of this set if it is non-empty and to be 0 if it is empty. Now, consider the set $\{(t + 1) - \mathbf{y}(w) \mid w \in \text{NotCross}, \mathbf{x}(w), \mathbf{x}'(w) \in (t, t + 1)\}$ and define n to be minimum element of this set. It is easy to check that for any τ' such that $m < \tau'.ltime < n$, property 2 holds for \mathbf{x}' and \mathbf{y}' .
 - B. All clocks cross to a new interval.
Let $m, n \in \mathbb{T}$ be respectively, the maximum and minimum elements of the set $\{t - \mathbf{y}(x) \mid \mathbf{x}'(x) \in (t, t + 1)\}$. Taking τ' such that $m < \tau'.ltime < n + 1$ gives the required result.
 - iii. $\tau.ltime > 0$ and τ makes some clocks reach an integer boundary.
Let *Reach* be the set of clocks that reached an integer boundary. Observe that for any two elements z and w of *Reach* it must be the case that $fr(\mathbf{x}(z)) = fr(\mathbf{x}(w))$. Now, take some $x \in \text{Reach}$ and let $m = (t - \mathbf{y}(x))$ where $t = \mathbf{x}'(x)$. Any τ' such that $\tau'.ltime = m$ gives us the required result. It is clear that such a τ' makes all the clocks in *Reach* reach an integer boundary. For any $z \in \text{Reach}$ and any clock w that has not reached an integer boundary in \mathbf{x}' , it must be the case that $fr(\mathbf{x}(z)) > fr(\mathbf{x}(w))$. By hypothesis and the third property of \sim , we also know that $fr(\mathbf{y}(z)) > fr(\mathbf{y}(w))$. It follows that w does not reach an integer boundary in \mathbf{y}' , as required. In the case where w is a clock that has crossed an integer boundary in \mathbf{x}' , we observe that $fr(\mathbf{x}(z)) < fr(\mathbf{x}(w))$ holds and conclude that the τ' we have chosen makes w cross the same integer boundary in \mathbf{y}' .
- (c) For the third property, we need to show that the τ' we defined for each case above, ensures that the ordering between the fractional parts of the clocks in \mathbf{x}' is preserved in \mathbf{y}' .

By property 2, which we have established for \mathbf{x}' and \mathbf{y}' , we know that, for any $x \in X_c$ if τ leads to $\mathbf{x}'(x) \in J$ then τ' has the same effect on \mathbf{y} such that

$\mathbf{y}'(x) \in J$. Similarly, if τ makes a clock cross reach an integer boundary in the evolution from \mathbf{x} to \mathbf{x}' , that is $\mathbf{x}'(x) = t$ then τ' yields $\mathbf{y}'(x) = t$. Since $\mathbf{x} \sim \mathbf{y}$, by property 3, we also know that the ordering between the fractional parts of clocks in \mathbf{x} and \mathbf{y} are the same. We know that in τ' all the clocks increase by the same amount. It follows that the ordering between the fractional parts of clocks is the same in \mathbf{x}' and \mathbf{y}' are the same. ■

6 Properties for Timed Automata

In this section, we define what we mean by a property for a timed automaton, describe some types of properties that are typically specified and proved for systems, and state some results about composition of automata with properties.

6.1 Definitions and Basic Results

A *property* P for a timed automaton \mathcal{A} is defined to be any subset of the execution fragments of \mathcal{A} . We write $execs_{(\mathcal{A}, P)}$ for the set of executions of \mathcal{A} in P , $traces_{(\mathcal{A}, P)}$ for the set of traces of executions of \mathcal{A} in P , and $tracefrags_{(\mathcal{A}, P)}$ for the set of traces of execution fragments of \mathcal{A} in P .

6.1.1 Safety and Liveness Properties

[[Nancy: We should ask Frits and Roberto to consider/approve the changed discussion of safety and liveness properties, and other significant changes we are making near the end of the paper.]]

A property P for a TA \mathcal{A} is said to be a *safety* property if it is closed under prefix and limits of execution fragments. In other words, if an execution fragment satisfies a safety property P , then so do all its prefixes, and if all the executions in a “chain” of successive extensions satisfy P , then so does the “limit” of the chain. Safety properties represent requirements that should be maintained by the system throughout its execution.

We say that an automaton \mathcal{A} satisfies a safety property S if every execution of \mathcal{A} is in S . Typically, the satisfaction of a safety property by an automaton is proved by induction. One shows that the property holds in any trivial execution fragment consisting of a point trajectory and that it is preserved by discrete steps and trajectories of the automaton.

A property P for \mathcal{A} is defined to be a *liveness* property provided that for any closed execution fragment α of \mathcal{A} , there exists an execution fragment β such that $\alpha \frown \beta \in P$. In

other words, no matter how \mathcal{A} behaves for a finite period of time, it is still possible for it to continue in some way and satisfy P .

We say that an automaton \mathcal{A} satisfies a liveness property L if every “maximal” execution of \mathcal{A} (an execution α such that there exists no execution of which α is a proper prefix) is in L . Typically, the proof of the satisfaction of a liveness property by an automaton involves the use of proof rules of a temporal logic, or progress functions from states to a well-founded set that measure the distance from the desired goal.

These definitions of safety and liveness are analogous to those considered for untimed systems in [3, 8, 10], and have also been adopted in the few models for timed systems that have addressed the classification of properties as safety and liveness properties [36, 1]. In order to support the definitions for our model we establish the following results, stated formally in Theorems 6.1 and 6.4: (1) The classes of safety and liveness properties are disjoint, (2) Every property can be expressed as the intersection of a safety and a liveness property.

The following theorem states that no property of a timed automaton can be both a safety and a liveness property, except for the special case where the property consists of all the execution fragments of the automaton.

Theorem 6.1 *Let \mathcal{A} be a TA. If P is both a safety property and a liveness property for \mathcal{A} , then $P = frags_{\mathcal{A}}$.*

Proof: Suppose that P is both a safety and a liveness property for \mathcal{A} and let α be any execution fragment of \mathcal{A} . We show $\alpha \in P$. Now consider the following cases:

1. α is a closed execution fragment.
Then, since P is a liveness property, there exists β such that $\alpha \frown \beta \in P$. Since P is also a safety property and is prefix-closed by definition, it must be that $\alpha \in P$.
2. α is an infinite sequence or a finite sequence ending with a right-open trajectory.
Then, α can be expressed as the limit of a chain of closed execution fragments $\alpha_0 \alpha_1 \alpha_2 \dots$. In case (1) we have established that for all $i \geq 0$, $\alpha_i \in P$. Since P is a safety property, the limit of this chain, which is α , must be in P .

Cases (1) and (2) together imply that $P = frags_{\mathcal{A}}$. ■

Let \mathcal{A} be a TA and P be a property for \mathcal{A} . We define $safe(P)$ to be the prefix- and limit-closure of the property P .

Lemma 6.2 *Let \mathcal{A} be a TA. For any property P for \mathcal{A} , $safe(P)$ is a safety property for \mathcal{A} .*

Proof: Immediate from the definitions of $\text{safe}(P)$ and of a safety property. ■

Lemma 6.3 *Let \mathcal{A} be a TA and P be a property for \mathcal{A} . If α is a closed execution fragment and $\alpha \in \text{safe}(P)$ then α is a prefix of some element in P .*

The following theorem states that any property for an automaton can be expressed as the intersection of a safety and a liveness property for that automaton.

Theorem 6.4 *Let \mathcal{A} be a TA. If P is a property for \mathcal{A} , then there exists a safety property S and a liveness property L for \mathcal{A} such that $P = S \cap L$.*

Proof: Let $S = \text{safe}(P)$. By Lemma 6.2, we know that S is a safety property for \mathcal{A} . Let $L = P \cup \{\alpha \mid \alpha \in \text{frags}_{\mathcal{A}}, \alpha \text{ is closed and no execution fragment of the form } \alpha \frown \beta \text{ is in } P\}$. We now show that L is a liveness property. Let α be a closed execution fragment of \mathcal{A} . If there exists some execution fragment β of \mathcal{A} such that $\alpha \frown \beta \in P$, then $\alpha \frown \beta \in L$ because $P \subseteq L$. On the other hand, if there is no execution fragment β such that $\alpha \frown \beta \in P$, then α is explicitly defined to be in L . Hence, we have shown that any closed execution fragment of \mathcal{A} has an extension in L as needed.

In order to conclude $P = S \cap L$, we need to show that $P \subseteq S \cap L$ and that $S \cap L \subseteq P$. $P \subseteq S \cap L$ is immediate from the definitions of S and L . We now show that $S \cap L \subseteq P$. Let α be an execution fragment in $S \cap L$ and suppose for the sake of contradiction that $\alpha \notin P$. Since $\alpha \in L - P$, by definition of L , α is closed and there exists no execution fragment β such that $\alpha \frown \beta \in P$. Since $\alpha \in S$ and α is closed, by Lemma 6.3, α must be a prefix of an execution fragment in P . This gives the needed contradiction. ■

6.1.2 Machine-closure

Consider a safety property S and a liveness property L for an automaton \mathcal{A} . It is in general desirable that L does not itself impose safety constraints, beyond those already imposed by S . To achieve this, L should be defined so that every closed execution in S can be extended to some execution that is in both S and L . The notion of machine-closure is used to formalize this condition. The pair of properties (S, L) is defined to be *machine-closed* provided that, for every closed execution fragment $\alpha \in S$, there exists β such that $\alpha \frown \beta \in S \cap L$.

Example 6.5 (A non-machine-closed pair of properties) Consider the timing-independent TA \mathcal{A} , given in Figure 15, whose set of state variables consists of a single discrete variable *countb*, and whose set of trajectories is exactly the set of constant-valued functions over left-closed time intervals with left endpoint 0. The automaton \mathcal{A} can per-

Automaton \mathcal{A}

Variables X : **discrete** $countb \in \mathbb{Z}$ **initially** 0

States Q : $val(X)$

Actions A : **external** a, b

Transitions \mathcal{D} : **external** a
 precondition
 $countb = 0$

external b
 effect
 $countb := countb + 1$

Trajectories \mathcal{T} : **satisfies**
 constant($countb$)

Figure 15: Machine closure

form b any time and it can perform a provided that it has not performed b . Now, consider the liveness property L for \mathcal{A} that consists of all the executions with infinitely many discrete actions and the safety property S for \mathcal{A} that consists of all the executions containing at most one b event. Then, since b disables all future a s, the intersection of L and S contains all the executions of \mathcal{A} with infinitely many a events and no b events.

Now, consider a closed execution α in S whose last action is b . This implies that α has no extension that contains an a , since by assumption the occurrence of b disables a . The only way of extending α to an execution $\alpha \frown \alpha'$ that contains infinitely many discrete actions is to perform infinitely many b s, but this would yield an execution $\alpha \frown \alpha'$ in $L - S$. Hence, the pair (S, L) is not machine-closed. ■

The above example illustrates that if a pair of safety and liveness properties for an automaton is not machine-closed, then the automaton may exhibit an anomaly. Namely, after some prefixes, the automaton may not be able to meet its liveness requirement without violating its safety requirement. This phenomenon has been observed in several studies on the classification of properties for untimed systems, including those by Dederichs and Weber [10], and Abadi and Lamport [1]. These studies suggest that the problem lies in defining the intended safety and liveness properties independently from one another. If the above-mentioned anomaly is to be avoided, a pair of safety and liveness properties need to be defined so that the pair is machine-closed.

The following theorem states that a pair of a safety and a liveness property for an automaton is machine-closed if the liveness property is defined as a subset of the safety property.

Theorem 6.6 *Let \mathcal{A} be a TA, S be a safety property and L be a liveness property for \mathcal{A} such that $L \subseteq S$. Then the pair (S, L) is machine-closed.*

Proof: Let α be a closed execution fragment in S . Since L is a liveness property for \mathcal{A} , there exists β such that $\alpha \frown \beta \in L$. Since $L \subseteq S$, we have that $\alpha \frown \beta \in S \cap L$. Thus, (S, L) is machine-closed. ■

The fact that two properties are machine-closed can be formalized by using other conditions equivalent to those we used in our formal definition above. The first property in the following theorem states that a pair (S, L) is machine closed if S is the same as the prefix and limit closure of the intersection of S and L . The second property states that if the intersection of S and L is contained in a safety property, it must be the case that S itself is contained in the same safety property. That is, L does not add new safety constraints to those already defined by S .

Theorem 6.7 *Let S be a safety property and L be a liveness property for an automaton \mathcal{A} . The pair (S, L) is machine closed iff either of the following holds:*

1. $S = \text{safe}(S \cap L)$.
2. If S' is a safety property and $S \cap L \subseteq S'$ then $S \subseteq S'$.

Proof: We show the following three implications: (1) if (S, L) is machine-closed then $S = \text{safe}(S \cap L)$, (2) if $S = \text{safe}(S \cap L)$, then for any safety property S' , $S \cap L \subseteq S'$ implies $S \subseteq S'$, and (3) if for every safety property S' , $S \cap L \subseteq S'$ implies $S \subseteq S'$, then (S, L) is machine-closed.

1. Suppose (S, L) is machine-closed. In order to show that $S = \text{safe}(S \cap L)$, we need to establish $S \subseteq \text{safe}(S \cap L)$ and $\text{safe}(S \cap L) \subseteq S$. To establish $S \subseteq \text{safe}(S \cap L)$ we take some $\alpha \in S$ and consider the following two cases:
 - (a) α is a closed execution fragment.
By the machine-closure assumption there exists β such that $\alpha \frown \beta \in S \cap L$. Since $\text{safe}(S \cap L)$ contains all the prefixes of elements of $S \cap L$, $\alpha \in \text{safe}(S \cap L)$, as needed.
 - (b) α is an infinite sequence or a finite sequence ending with a right-open trajectory. Then α must be the limit of a chain of closed execution fragments $\alpha_0 \alpha_1 \dots$ in S . Since S is a safety property, every prefix of α is in S . Therefore for each i , we have $\alpha_i \in S$. By case (a), for each i , $\alpha_i \in \text{safe}(S \cap L)$. By definition of $\text{safe}(S \cap L)$ the limit α is also in $\text{safe}(S \cap L)$, as needed.

To show $\text{safe}(S \cap L) \subseteq S$, take some $\alpha \in \text{safe}(S \cap L)$. We consider two cases:

- (a) α is a closed execution fragment.
Then, by Lemma 6.3, α is a prefix of some element in $S \cap L$. That is to say, $\alpha \frown \beta \in (S \cap L)$ for some β and it follows that $\alpha \frown \beta \in S$. Since S is a safety property we have $\alpha \in S$, as needed.
 - (b) α is an infinite sequence or a finite sequence ending with a right-open trajectory.
Then α must be the limit of a chain of closed execution fragments $\alpha_0 \alpha_1 \dots$ in $\text{safe}(S \cap L)$. We have established in case (a) that each closed execution fragment α_i is in S . Since S is a safety property, the limit α must also be in S , as needed.
2. Suppose $S = \text{safe}(S \cap L)$. Let S' be a safety property such that $S \cap L \subseteq S'$. Let $\alpha \in S$ and show that $\alpha \in S'$.
- (a) α is a closed execution fragment.
Since $S = \text{safe}(S \cap L)$ by assumption, $\alpha \in \text{safe}(S \cap L)$, and since α is closed, by Lemma 6.3, α is a prefix of some element in $S \cap L$. Since $(S \cap L) \subseteq S'$ we have that α is a prefix of some element of S' . Since S' is a safety property, $\alpha \in S'$.
 - (b) α is an infinite sequence or a finite sequence ending with a right-open trajectory.
Then α must be the limit of a chain of closed execution fragments $\alpha_0 \alpha_1 \dots$ in $\text{safe}(S \cap L)$. We have established in case (a) that each closed execution fragment α_i is in S' . Since S' is a safety property, the limit α must also be in S' , as needed.
3. Suppose that for every safety property S' , $S \cap L \subseteq S'$ implies $S \subseteq S'$. We must show that for every closed execution fragment $\alpha \in S$, there exists β such that $\alpha \frown \beta \in S \cap L$. Let α be a closed execution fragment in S . By Lemma 6.2 we have that $\text{safe}(S \cap L)$ is a safety property. Since $S \cap L \subseteq \text{safe}(S \cap L)$, by assumption $S \subseteq \text{safe}(S \cap L)$. Since $\alpha \in S$, we have that $\alpha \in \text{safe}(S \cap L)$. Since α is closed, by Lemma 6.3, α is a prefix of some element of $S \cap L$. That is to say, there exists β such that $\alpha \frown \beta \in S \cap L$, as needed.

■

6.1.3 Special kinds of properties

Fairness properties: Proving interesting liveness properties requires some assumptions saying that certain activities in a concurrent system get “enough” chances to make progress. Fairness properties are special kinds of liveness properties that express such assumptions. We define two types of fairness: weak fairness and strong fairness.

Let \mathcal{A} be a TA and let C be a subset of the actions of \mathcal{A} . Let α be an execution fragment of \mathcal{A} . Then:

1. α is *weakly fair* for C if (at least) one of the following conditions holds:
 - (a) α contains infinitely many events from C .
 - (b) There is no suffix β of α such that C is enabled in all states of β .
2. α is *strongly fair* for C if (at least) one of the following conditions holds:
 - (a) α contains infinitely many events from C .
 - (b) There is some suffix β of α such that C is disabled in all states of β .

Consider a finite execution fragment α . If α ends with a closed trajectory, the definition above says that for α to be weakly fair or strongly fair for C , C must be disabled in $\alpha.lstate$. On the other hand, if α ends with a right-open trajectory, α is weakly fair provided that there are state occurrences with C disabled, at times arbitrarily close to $\alpha.ltime$ and α is strongly fair provided that C is continuously disabled from some point on in α .

Theorem 6.8 *Let \mathcal{A} be a TA, C a subset of actions of \mathcal{A} and α be an execution fragment of \mathcal{A} . If α is strongly fair for C then α is weakly fair for C .*

Proof: Follows from the definitions of strong and weak fairness. ■

Theorem 6.9 *For any timed automaton \mathcal{A} and any subset C of its actions, the set of strongly fair execution fragments for C is a liveness property for \mathcal{A} .*

Proof: Fix \mathcal{A} a TA, C a subset of the actions of \mathcal{A} and let α be a closed execution fragment of \mathcal{A} . We are required to show that for some β , $\alpha \frown \beta$ is strongly fair for C . Construct an execution fragment $\beta = \alpha_0 \frown \alpha_1 \frown \dots$ as follows:

- $\alpha_0 = \wp(\alpha.lstate)$,
- For each $i \geq 1$, if there exists $(\alpha_{i-1}.lstate, b, \mathbf{y}) \in \mathcal{D}_{\mathcal{A}}$ for some $b \in C$ and some $\mathbf{y} \in Q_{\mathcal{A}}$, then choose some such b and \mathbf{y} and define $\alpha_i = \wp(\alpha_{i-1}.lstate) b \wp(\mathbf{y})$; otherwise, $i - 1$ is the final index in the sequence.

It follows that, if β is a finite sequence then C is disabled in its last state. Therefore, for some suffix of β , C is disabled in all states and $\alpha \frown \beta$ is strongly fair with respect to C . On the other hand, if β is an infinite sequence then $\alpha \frown \beta$ has infinitely many events from C , as needed. ■

Corollary 6.10 *For any timed automaton \mathcal{A} and any subset C of its actions, the set of weakly fair execution fragments for C is a liveness property for \mathcal{A} .*

Proof: Follows from Theorem 6.8 and Theorem 6.9. ■

Admissibility: Admissibility is another notion that is fundamental to any useful formal model for timed systems. It is hard to think about executions such as those that arise from Zeno behavior, yet they make formal sense. Admissibility conditions help one to avoid considering such executions in reasoning about properties. The formal definition of admissibility is given in 3.4.1. Formally, an execution fragment α is *admissible* if $\alpha.ltime = \infty$.

Theorem 6.11 *A timed automaton \mathcal{A} is feasible if and only if its set of admissible execution fragments is a liveness property for \mathcal{A} .*

Proof: Immediate from the definitions of feasibility and liveness property. ■

History-independence: History-independence is an important characteristic of properties that simplifies the analysis of the behavior of an automaton. A property P of a timed automaton \mathcal{A} is said to be *history-independent* provided that the following holds: For every execution fragment α , if α' is a suffix of α , then $\alpha \in P$ if and only if $\alpha' \in P$. In other words, whether or not α satisfies P is determined only by what happens in its suffixes—it is not affected by what happens in any initial portion of α . If a property P is known to be history-independent, then one can prove that an execution fragment α satisfies P by considering the portion of α from some point onward.

The liveness properties that are typically used are history-independent. Fairness and admissibility properties defined earlier in the section constitute examples of history-independent properties, as shown in the following theorems.

Theorem 6.12 *For any timed automaton \mathcal{A} , and any subset C of its actions, the set of weakly fair execution fragments for C is history-independent.*

Proof: Fix \mathcal{A} a TA, C a subset of actions of \mathcal{A} and let $\alpha = \alpha' \frown \alpha''$ with $\alpha'.lstate = \alpha''.fstate$ be an execution fragment of \mathcal{A} .

First, suppose that α is weakly fair for C . We are required to show that α'' is also weakly fair with respect to C . By definition of weak fairness, either α contains infinitely many events from C , or it has no suffix in which C is enabled in all states. Since α'' is a suffix of α , in either case we conclude that α'' is weakly fair with respect to C by using the definition of weak fairness.

Now, suppose that α'' is weakly fair for C . We are required to show that α is also weakly fair with respect to C . Similar to the case above, this is easy to show by using the definition of weak fairness and the fact that α'' is a suffix of α . ■

Theorem 6.13 *For any timed automaton \mathcal{A} , and a subset C of its actions, the set of strongly fair execution fragments for C is history-independent.*

Theorem 6.14 *For any timed automaton \mathcal{A} , the set of admissible execution fragments is history-independent.*

6.2 Implementation Relationships

We define another preorder for automata with properties:

- $(\mathcal{A}_1, P_1) \leq (\mathcal{A}_2, P_2)$ provided that $traces_{(\mathcal{A}_1, P_1)} \subseteq traces_{(\mathcal{A}_2, P_2)}$.

If P_1 is a liveness property for a TA \mathcal{A}_1 and P_2 is any property for a TA \mathcal{A}_2 , and (\mathcal{A}_1, P_1) and (\mathcal{A}_2, P_2) are related by the preorder defined above, then every closed trace of \mathcal{A}_1 is also a trace of \mathcal{A}_2 . This is shown in the following theorem.

Theorem 6.15 *Suppose that P_1 is a liveness property for \mathcal{A}_1 and P_2 is any property for \mathcal{A}_2 . If $(\mathcal{A}_1, P_1) \leq (\mathcal{A}_2, P_2)$ then every closed trace of \mathcal{A}_1 is a trace of \mathcal{A}_2 .*

Proof: Assume $(\mathcal{A}_1, P_1) \leq (\mathcal{A}_2, P_2)$ and let β be a closed trace of \mathcal{A}_1 . Let α be a closed execution of \mathcal{A}_1 with $trace(\alpha) = \beta$. Since P_1 is a liveness property of \mathcal{A}_1 , there exists an execution fragment α' of \mathcal{A}_1 such that $\alpha \frown \alpha' \in P_1$.

Let $\beta' = trace(\alpha \frown \alpha')$; then clearly $\beta' \in traces_{(\mathcal{A}_1, P_1)}$. Then because $(\mathcal{A}_1, P_1) \leq (\mathcal{A}_2, P_2)$, we have that $\beta' \in traces_{(\mathcal{A}_2, P_2)}$. Since β is a prefix of β' and the set of traces of \mathcal{A}_2 is prefix-closed, it follows that β is a trace of \mathcal{A}_2 , as needed. ■

6.3 Simulation Relations

As we have seen in Section 4.5, simulation relations provide a useful tool for reasoning about implementation relationships between automata at multiple levels of abstraction. The existence of a forward or a backward simulation relation, or a history or a prophecy relation, from one timed automaton \mathcal{A} to another, \mathcal{B} , is sufficient to establish that each trace of \mathcal{A} is also a trace of \mathcal{B} .

For any TA \mathcal{A} the set of all execution fragments of \mathcal{A} , $frags_{\mathcal{A}}$, constitutes a safety property. This follows from the definition of a safety property in Section 6.1.1 by using the fact that $frags_{\mathcal{A}}$ is prefix and limit closed. Suppose we define a safety property S_1 for an automaton \mathcal{A} to be $frags_{\mathcal{A}}$ and a safety property S_2 for an automaton \mathcal{B} to be $frags_{\mathcal{B}}$. The existence of a forward simulation relation from \mathcal{A} to \mathcal{B} would imply that for any execution $\alpha \in S_1$, there is an execution $\beta \in S_2$ such that $trace(\alpha) = trace(\beta)$. However, the same implication does not in general hold, if we replace safety properties S_1 and S_2 with arbitrary liveness properties $L_1 \subseteq S_1$ and $L_2 \subseteq S_2$ for \mathcal{A} and \mathcal{B} , respectively. In [9] Attie addresses this issue in an untimed setting and proposes several notions of “liveness-preserving” simulation relations. The liveness properties that he considers are

of a special form that are analogous to the acceptance condition of a complemented-pairs automaton [7].

In the two theorems below, we consider the special classes of weak and strong fairness properties and state some extra constraints on forward simulation relations. The existence of a forward simulation relation from an automaton \mathcal{A} to another \mathcal{B} that satisfies these additional constraints allows us to conclude that the trace of each fair execution of \mathcal{A} is also a trace of a fair execution of \mathcal{B} . The constraints that we impose on the forward simulation relation for discrete steps turn out to be special cases of Attie's constraints[9].

Let \mathcal{A} and \mathcal{B} be comparable TAs. Let $C_{\mathcal{A}}$ be a set of actions of \mathcal{A} and $C_{\mathcal{B}}$ be a set of actions of \mathcal{B} . A *fair forward simulation* from \mathcal{A} to \mathcal{B} with respect to $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ is a relation $R \subseteq Q_{\mathcal{A}} \times Q_{\mathcal{B}}$ satisfying the following conditions, for all states $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ of \mathcal{A} and \mathcal{B} , respectively:

1. If $\mathbf{x}_{\mathcal{A}} \in \Theta_{\mathcal{A}}$ then there exists a state $\mathbf{x}_{\mathcal{B}} \in \Theta_{\mathcal{B}}$ such that $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$. Moreover, if $C_{\mathcal{A}}$ is disabled in $\mathbf{x}_{\mathcal{A}}$, then $C_{\mathcal{B}}$ is disabled in $\mathbf{x}_{\mathcal{B}}$.
2. If $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$ and α is an execution fragment of \mathcal{A} consisting of an action a surrounded by two point trajectories, with $\alpha.fstate = \mathbf{x}_{\mathcal{A}}$, then \mathcal{B} has a closed execution fragment β such that $\beta.fstate = \mathbf{x}_{\mathcal{B}}$, $trace(\beta) = trace(\alpha)$, and $\alpha.lstate R \beta.lstate$. Moreover,
 - (a) If $a \in C_{\mathcal{A}}$ then β contains some event in $C_{\mathcal{B}}$.
 - (b) If $C_{\mathcal{A}}$ is disabled in $\alpha.lstate$ then
 - i. If $\beta = \wp(\mathbf{x}_{\mathcal{B}})$ then $C_{\mathcal{B}}$ is disabled in $\mathbf{x}_{\mathcal{B}}$.
 - ii. If $\beta \neq \wp(\mathbf{x}_{\mathcal{B}})$ then $C_{\mathcal{B}}$ is disabled in all states in β except possibly in $\mathbf{x}_{\mathcal{B}}$.
3. If $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$ and α is an execution fragment of \mathcal{A} consisting of a single closed trajectory, with $\alpha.fstate = \mathbf{x}_{\mathcal{A}}$, then \mathcal{B} has a closed execution fragment β such that $\beta.fstate = \mathbf{x}_{\mathcal{B}}$, $trace(\beta) = trace(\alpha)$, and $\alpha.lstate R \beta.lstate$. Moreover,
 - (a) If $\beta.ltime = 0$ and $C_{\mathcal{A}}$ is disabled in $\mathbf{x}_{\mathcal{A}}$ then $C_{\mathcal{B}}$ is disabled in all states in β .
 - (b) If $\beta.ltime > 0$ then for all t such that $0 < t \leq \alpha.ltime$, if $C_{\mathcal{A}}$ is disabled in $\alpha(t)$ then for each closed prefix β' of β such that $\beta'.ltime = t$, $C_{\mathcal{B}}$ is disabled in $\beta'.lstate$.

We say that R is a fair forward simulation from \mathcal{A} to \mathcal{B} , without mentioning $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ explicitly, when those sets are clear from the context.

Now, we define a construction that, given two automata \mathcal{A} and \mathcal{B} , two sets of actions $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$, a fair forward simulation R from \mathcal{A} to \mathcal{B} , and an execution α of \mathcal{A} , generates an execution β of \mathcal{B} by using the definition of a fair forward simulation.

Let \mathcal{A} and \mathcal{B} be two TAs, $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ be sets of actions for \mathcal{A} and \mathcal{B} , respectively, and R be a fair forward simulation from \mathcal{A} to \mathcal{B} with respect to $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. Let α be an execution of \mathcal{A} . The construction consists of the following steps:

1. Using axioms **T1** and **T2**, write α as a concatenation $\alpha_0 \frown \alpha_1 \frown \alpha_2 \cdots (\alpha_0 \frown \alpha_1 \frown \cdots \frown \alpha_n$ if α is a finite sequence ending with a closed trajectory), in which each execution fragment α_i consists of either a single closed trajectory or one action surrounded by two point trajectories. Without loss of generality, we can assume that for each $i \geq 0$, $\alpha_i.lstate = \alpha_{i+1}.fstate$.
2. Define inductively a sequence $\beta_0 \beta_1 \dots$ of closed execution fragments of \mathcal{B} , such that $\beta_0.fstate = \mathbf{x}_B$ for some $\mathbf{x}_B \in \Theta_B$ and, for each i , $\alpha_i.lstate R \beta_i.lstate$, $\beta_i.lstate = \beta_{i+1}.fstate$, and $trace(\alpha_i) = trace(\beta_i)$. We use Properties 1 and 3 of a fair forward simulation in the construction of β_0 , Property 2 in the construction of β_i consisting of one action surrounded by two point trajectories, and Property 3 in the construction of β_i consisting of a single closed trajectory.
3. Let β be the concatenation $\beta_0 \frown \beta_1 \frown \cdots$.

For such β , we say that β *corresponds to* α with respect to R, C_A and C_B . When R, C_A and C_B are clear from the context, we do not state their names explicitly.

Lemma 6.16 *Let \mathcal{A} and \mathcal{B} be two TAs, C_A and C_B be sets of actions for \mathcal{A} and \mathcal{B} , respectively, and R be a fair forward simulation from \mathcal{A} to \mathcal{B} with respect to C_A and C_B . Let α be an execution of \mathcal{A} and β be an execution of \mathcal{B} that corresponds to α . Suppose that α is expressed as $\alpha_0 \frown \alpha_1 \frown \cdots$ and β is expressed as $\beta_0 \frown \beta_1 \frown \cdots$ in the construction of β . Then, β satisfies the following properties:*

1. If C_A is disabled in $\alpha_0.fstate$, then C_B is disabled in $\beta_0.fstate$.
2. For each α_i of the form $\wp(\mathbf{x}_A)$ a $\wp(\mathbf{x}'_A)$ let $\mathbf{x}_B = \beta_i.fstate$. Then,
 - If $a \in C_A$ then β_i contains some event in C_B .
 - If C_A is disabled in \mathbf{x}'_A then
 - If $\beta_i = \wp(\mathbf{x}_B)$ then C_B is disabled in \mathbf{x}_B .
 - If $\beta_i \neq \wp(\mathbf{x}_B)$ then C_B is disabled in all states in β_i except possibly in \mathbf{x}_B .
3. For each α_i consisting of a single closed trajectory:
 - If $\beta_i.ltime = 0$ and C_A is disabled in $\alpha_i.fstate$ then C_B is disabled in all states in β_i .
 - If $\beta_i.ltime > 0$ then for all t such that $0 < t \leq \alpha_i.ltime$, if C_A is disabled in $\alpha_i(t)$ then for each closed prefix β'_i of β_i such that $\beta'_i.ltime = t$, C_B is disabled in $\beta'_i.lstate$.
4. β is an execution of \mathcal{B} such that $trace(\beta) = trace(\alpha)$.

Proof: Properties 1, 2 and 3 follow from the construction of β and the definition of a fair forward simulation relation. We show property 4 as follows. By Lemma 4.7, β is an execution fragment of \mathcal{B} . By the construction of β , $\beta_0.fstate = \mathbf{x}_{\mathcal{B}}$ for some $\mathbf{x}_{\mathcal{B}} \in \Theta_{\mathcal{B}}$. Therefore, that β is an execution of \mathcal{B} . By Lemma 3.9 applied to both α and β , $trace(\beta) = trace(\alpha)$. ■

Lemma 6.17 *Let \mathcal{A} and \mathcal{B} be two TAs, $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ be sets of actions for \mathcal{A} and \mathcal{B} , respectively, and R be a fair forward simulation from \mathcal{A} to \mathcal{B} with respect to $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. Let α be an execution of \mathcal{A} , and let β be an execution of \mathcal{B} that corresponds to α . Then, if α contains infinitely many events from $C_{\mathcal{A}}$ it must be the case that β contains infinitely many events from $C_{\mathcal{B}}$.*

Proof: We know that, in the construction of β , α is expressed as $\alpha_0 \frown \alpha_1 \frown \dots$ in which each execution fragment α_i consists of either a single closed trajectory or one action surrounded by two point trajectories, and β is expressed as $\beta_0 \frown \beta_1 \frown \dots$. Suppose that α contains infinitely many events from $C_{\mathcal{A}}$. By property 2 of Lemma 6.16 in the construction of β , we have that for each α_i consisting of one action surrounded by two point trajectories, if α_i contains a $C_{\mathcal{A}}$ event, then β_i contains a $C_{\mathcal{B}}$ event. Since there are infinitely many $C_{\mathcal{A}}$ events in α , there must be infinitely many $C_{\mathcal{B}}$ events in β , as needed. ■

Lemma 6.18 *Let \mathcal{A} and \mathcal{B} be two TAs, $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ be sets of actions for \mathcal{A} and \mathcal{B} , respectively, and R be a fair forward simulation from \mathcal{A} to \mathcal{B} with respect to $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. Let α be an execution of \mathcal{A} that is a finite sequence ending with a closed trajectory, and let β be an execution of \mathcal{A} that corresponds to α . Then, if $C_{\mathcal{A}}$ is disabled in $\alpha.lstate$ it must be the case that $C_{\mathcal{B}}$ is disabled in $\beta.lstate$.*

Proof: We know that, in the construction of β , α is expressed as $\alpha_0 \frown \alpha_1 \frown \dots \frown \alpha_n$ in which each execution fragment α_i consists of either a single closed trajectory or one action surrounded by two point trajectories and β is expressed as $\beta_0 \frown \beta_1 \frown \dots \frown \beta_n$. Suppose that $C_{\mathcal{A}}$ is disabled in $\alpha.lstate$. Since $\alpha.lstate = \alpha_n.lstate$, we have that $C_{\mathcal{A}}$ is disabled in $\alpha_n.lstate$. Now, consider the following cases:

1. α_n is a single closed trajectory.
Since $C_{\mathcal{A}}$ is disabled in $\alpha_n.lstate$, by using property 3 in Lemma 6.16, we have that $C_{\mathcal{B}}$ is disabled in $\beta_n.lstate$. Since $\beta.lstate = \beta_n.lstate$, we have that $C_{\mathcal{B}}$ is disabled in $\beta.lstate$, as needed.
2. α_n is one action surrounded by point trajectories.
Since $C_{\mathcal{A}}$ is disabled in $\alpha_n.lstate$, by using property 2 in Lemma 6.16, we have that $C_{\mathcal{B}}$ is disabled in $\beta_n.lstate$. Since $\beta.lstate = \beta_n.lstate$, we have that $C_{\mathcal{B}}$ is disabled in $\beta.lstate$, as needed.

■

Lemma 6.19 *Let \mathcal{A} and \mathcal{B} be two TAs, $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ be sets of actions for \mathcal{A} and \mathcal{B} , respectively, and R be a fair forward simulation from \mathcal{A} to \mathcal{B} with respect to $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. Let α be an execution of \mathcal{A} such that α is an infinite sequence or a finite sequence ending with an open trajectory, and let β be an execution of \mathcal{B} that corresponds to α . Then, if for some suffix α' of α , $C_{\mathcal{A}}$ is disabled in all states in α' , it must be the case that for some suffix β' of β , $C_{\mathcal{B}}$ is disabled in all states in β' .*

Proof: We know that, in the construction of β , α is expressed as $\alpha_0 \frown \alpha_1 \frown \dots$ in which each execution fragment α_i consists of either a single closed trajectory or one action surrounded by two point trajectories, and β is expressed as $\beta_0 \frown \beta_1 \frown \dots$. Suppose that for some suffix α' of α , $C_{\mathcal{A}}$ is disabled in all states in α' . Consider the following cases:

1. For infinitely many $i \geq 0$, α_i is an execution fragment consisting of an action surrounded by point trajectories.

Without loss of generality we can assume that $\alpha' = \alpha_i \frown \alpha_{i+1} \frown \dots$ for some $i \geq 0$ and α' is an infinite sequence starting with a discrete action surrounded by two point trajectories. Now, consider the corresponding execution fragment $\beta' = \beta_i \frown \beta_{i+1} \frown \dots$ of \mathcal{B} . Let β'' be the suffix $\beta_{i+1} \frown \beta_{i+2} \frown \dots$ of β' . Since $C_{\mathcal{A}}$ is disabled in all states in α' , $C_{\mathcal{A}}$ is disabled in $\alpha_i.lstate$. By property 2 of Lemma 6.16 we know that $C_{\mathcal{B}}$ is disabled in $\beta_i.lstate$. Then for each $j > i$, by properties 2 and 3 of Lemma 6.16, we know that $C_{\mathcal{B}}$ is disabled in all states in β_j , except possibly in $\beta_j.fstate$. Since for each $j > i$, $\beta'_j.fstate = \beta'_{j-1}.lstate$ by the construction of β , we know that $C_{\mathcal{B}}$ is disabled in all states of β'' , which is a suffix of β .

2. For only finitely many $i \geq 0$, α_i consists of an action surrounded by point trajectories. Then for all sufficiently large $i \geq 0$, α_i consists of a single closed trajectory. Without loss of generality we can assume that $\alpha' = \alpha_i \frown \alpha_{i+1} \frown \dots$ for some sufficiently large $i \geq 0$ and for each $j \geq i$, α_j is a single closed trajectory. Now consider the corresponding execution fragment $\beta' = \beta_i \frown \beta_{i+1} \frown \dots$. Let β'' be the suffix $\beta_{i+1} \frown \beta_{i+2} \frown \dots$ of β' . Since $C_{\mathcal{A}}$ is disabled in all states in α' , $C_{\mathcal{A}}$ is disabled in $\alpha_i.lstate$. Then, by property 3 of Lemma 6.16, we know that $C_{\mathcal{B}}$ is disabled in $\beta_i.lstate$ and for each $j > i$, $C_{\mathcal{B}}$ is disabled in all states in β_j , except possibly in $\beta_j.fstate$. Since for each $j > i$, $\beta'_j.fstate = \beta'_{j-1}.lstate$ by the construction of β , we know that $C_{\mathcal{B}}$ is disabled in all states of β'' , which is a suffix of β .

■

Lemma 6.20 *Let \mathcal{A} and \mathcal{B} be two TAs, $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ be sets of actions for \mathcal{A} and \mathcal{B} , respectively, and R be a fair forward simulation from \mathcal{A} to \mathcal{B} with respect to $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$.*

Let α be an execution of \mathcal{A} such that α is an infinite sequence or a finite sequence ending with an open trajectory, and let β be an execution of \mathcal{A} that corresponds to α . Then, if there is no suffix α' of α such that $C_{\mathcal{A}}$ is enabled in all states in α' it must be the case that there is no suffix β' of β such that $C_{\mathcal{B}}$ is enabled in all states in β' .

Proof: We know that, in the construction of β , α is expressed as $\alpha_0 \frown \alpha_1 \frown \dots$ in which each execution fragment α_i consists of either a single closed trajectory or one action surrounded by two point trajectories, and β is expressed as $\beta_0 \frown \beta_1 \frown \dots$. Suppose that there is no suffix α' of α such that $C_{\mathcal{A}}$ is enabled in all states in α' . This means that for infinitely many $i \geq 0$, $C_{\mathcal{A}}$ is disabled in some state of α_i . Then, by properties 2 and 3 in Lemma 6.16, we know that for infinitely many $i \geq 0$, $C_{\mathcal{B}}$ is disabled in some state of β_i . This implies that β has no suffix in which $C_{\mathcal{B}}$ is enabled in all states. ■

The following lemma states that a fair forward simulation from \mathcal{A} to \mathcal{B} yields a correspondence for open trajectories.

Lemma 6.21 *Let \mathcal{A} and \mathcal{B} be comparable TAs, $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$ be sets of actions of \mathcal{A} and \mathcal{B} respectively, and R be a fair forward simulation from \mathcal{A} to \mathcal{B} with respect to $C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. Let $\mathbf{x}_{\mathcal{A}}$ and $\mathbf{x}_{\mathcal{B}}$ be states of \mathcal{A} and \mathcal{B} , respectively, such that $\mathbf{x}_{\mathcal{A}} R \mathbf{x}_{\mathcal{B}}$. Let α be an execution fragment of \mathcal{A} from state $\mathbf{x}_{\mathcal{A}}$ consisting of a single open trajectory τ . Then \mathcal{B} has an execution fragment β with $\beta.fstate = \mathbf{x}_{\mathcal{B}}$ and $trace(\beta) = trace(\alpha)$. Moreover, β satisfies the following condition: for all t such that $0 < t \leq \tau.ltime$, if $C_{\mathcal{A}}$ is disabled in $\tau(t)$ then for each prefix β' of β such that $\beta'.ltime = t$, $C_{\mathcal{B}}$ is disabled in $\beta'.lstate$.*

Proof: Let τ be the single open trajectory in α . Using axioms **T1** and **T2**, we construct an infinite sequence $\tau_0 \tau_1 \dots$ of closed trajectories of \mathcal{A} such that $\tau = \tau_0 \frown \tau_1 \frown \dots$. Then, working recursively, we construct a sequence $\beta_0 \beta_1 \dots$ of closed execution fragments of \mathcal{B} such that $\beta_0.fstate = \mathbf{x}_{\mathcal{B}}$ and, for each i , $\tau_i.lstate R \beta_i.lstate$, $\beta_i.lstate = \beta_{i+1}.fstate$, $trace(\tau_i) = trace(\beta_i)$, and the following fairness condition holds: for all t such that $0 < t \leq \tau_i.ltime$, if $C_{\mathcal{A}}$ is disabled in $\tau_i(t)$ then for each prefix β'_i of β_i such that $\beta'_i.ltime = t$, $C_{\mathcal{B}}$ is disabled in $\beta'_i.lstate$. This construction uses induction on i , using Property 3 of the definition of a fair forward simulation in the induction step. Now let $\beta = \beta_0 \frown \beta_1 \frown \dots$. By Lemma 4.7, β is an execution fragment of \mathcal{B} . Clearly, $\beta.fstate = \mathbf{x}_{\mathcal{B}}$. By Lemma 3.9 applied to both α and β , $trace(\beta) = trace(\alpha)$. Using Property 3 for each β_i , and the inductive hypothesis $\beta_i.lstate = \beta_{i+1}.fstate$, we have that for all t such that $0 < t \leq \tau.ltime$, if $C_{\mathcal{A}}$ is disabled in $\tau(t)$ then for each prefix β' of β such that $\beta'.ltime = t$, $C_{\mathcal{B}}$ is disabled in $\beta'.lstate$. Thus β has the required properties. ■

Theorem 6.22 *Suppose that R is a fair forward simulation relation from \mathcal{A} to \mathcal{B} with respect to a set $C_{\mathcal{A}}$ of actions of \mathcal{A} and a set $C_{\mathcal{B}}$ of actions of \mathcal{B} . Let $L_{\mathcal{A}}$ be the set of strongly fair executions of \mathcal{A} for $C_{\mathcal{A}}$ and let $L_{\mathcal{B}}$ be the set of strongly fair executions of \mathcal{B} for $C_{\mathcal{B}}$. Then $(\mathcal{A}, L_{\mathcal{A}}) \leq (\mathcal{B}, L_{\mathcal{B}})$.*

Proof: Let α be an execution of \mathcal{A} such that $\alpha \in L_{\mathcal{A}}$ and let β be an execution fragment of \mathcal{B} that corresponds to α with respect to $R, C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. By property 4 in Lemma 6.16 we know that β is an execution of \mathcal{B} such that $trace(\alpha) = trace(\beta)$. We show that $\beta \in L_{\mathcal{B}}$ by considering the following cases:

1. α contains infinitely many events from $C_{\mathcal{A}}$.
By Lemma 6.17, we know that β has infinitely many events from $C_{\mathcal{B}}$. Then, by definition of strong fairness $\beta \in L_{\mathcal{B}}$, as needed.
2. For some suffix α' of α , $C_{\mathcal{A}}$ is disabled in all states in α' .
 - (a) α is either an infinite sequence or a finite sequence ending with an open trajectory.
Then, by Lemma 6.19, we have that $C_{\mathcal{B}}$ is disabled in all states in some suffix of β . Then, by definition of strong fairness $\beta \in L_{\mathcal{B}}$, as needed.
 - (b) α is a finite sequence ending with a closed trajectory.
By Lemma 6.18, we have that $C_{\mathcal{B}}$ is disabled in $\beta.lstate$. Since $\beta.lstate$ is a suffix of β , by definition of strong fairness $\beta \in L_{\mathcal{B}}$, as needed.

■

Theorem 6.23 *Suppose that R is a fair forward simulation relation from \mathcal{A} to \mathcal{B} with respect to a set $C_{\mathcal{A}}$ of actions of \mathcal{A} and a set $C_{\mathcal{B}}$ of actions of \mathcal{B} . Let $L_{\mathcal{A}}$ be the set of weakly fair executions of \mathcal{A} for $C_{\mathcal{A}}$ and let $L_{\mathcal{B}}$ be the set of weakly fair executions of \mathcal{B} for $C_{\mathcal{B}}$. Then $(\mathcal{A}, L_{\mathcal{A}}) \leq (\mathcal{B}, L_{\mathcal{B}})$.*

Proof: Let α be an execution of \mathcal{A} such that $\alpha \in L_{\mathcal{A}}$ and let β be an execution fragment of \mathcal{B} that corresponds to α with respect to $R, C_{\mathcal{A}}$ and $C_{\mathcal{B}}$. By property 4 in Lemma 6.16 we know that β is an execution of \mathcal{B} such that $trace(\alpha) = trace(\beta)$. We show that $\beta \in L_{\mathcal{B}}$ by considering the following cases:

1. α contains infinitely many events from $C_{\mathcal{A}}$.
By Lemma 6.17, we know that β has infinitely many events from $C_{\mathcal{B}}$. Then, by definition of weak fairness $\beta \in L_{\mathcal{B}}$, as needed.
2. There is no suffix α' of α such that $C_{\mathcal{A}}$ is enabled in all states in α' .
 - (a) α is either an infinite sequence or a finite sequence ending with an open trajectory.
Then, by Lemma 6.20, we have that there is no suffix β' of β such that $C_{\mathcal{B}}$ is enabled in all states in β' . By definition of weak fairness $\beta \in L_{\mathcal{B}}$, as needed.

(b) α is a finite sequence ending with a closed trajectory.

By Lemma 6.18, we have that $C_{\mathcal{B}}$ is disabled in $\beta.lstate$. Therefore, β cannot have any suffix in which $C_{\mathcal{B}}$ is enabled in all states. Then, by definition of weak fairness $\beta \in L_{\mathcal{B}}$, as needed. ■

It would have been possible to prove Theorem 6.23 for a slightly different notion of fair forward simulation obtained by weakening Property 3 of the current definition. The current definition requires that the disabling is carried over from the low-level automaton to the high-level one for all states in a trajectory, except for the first state of trajectories with limit time greater than zero. For proving Theorem 6.23, it would have been sufficient to require that the disabling be carried over for some states only.

6.4 Composition

This section includes results that are essential for compositional reasoning about timed automata with properties. They are specializations of the similar results in Section 5.1.

6.4.1 Definitions and Basic Results

If \mathcal{A}_1 and \mathcal{A}_2 are two compatible timed automata and P_1 and P_2 are properties for \mathcal{A}_1 and \mathcal{A}_2 , respectively, then we define $P_1 \parallel P_2$ to be $\{\alpha \in frags_{\mathcal{A}_1 \parallel \mathcal{A}_2} \mid \alpha \upharpoonright (A_i, X_i) \in P_i, i \in \{1, 2\}\}$. Using this, we define composition of automata with properties $(\mathcal{A}_1, P_1) \parallel (\mathcal{A}_2, P_2)$ as $(\mathcal{A}_1 \parallel \mathcal{A}_2, P_1 \parallel P_2)$.

Theorem 6.24 *Let \mathcal{A}_1 and \mathcal{A}_2 be two compatible TAs and P_1 and P_2 be properties for \mathcal{A}_1 and \mathcal{A}_2 , respectively. Then $traces_{(\mathcal{A}_1 \parallel \mathcal{A}_2, P_1 \parallel P_2)}$ is exactly the set of (E, \emptyset) -sequences whose restrictions to \mathcal{A}_1 and \mathcal{A}_2 are $traces_{(\mathcal{A}_1, P_1)}$ and $traces_{(\mathcal{A}_2, P_2)}$, respectively. That is, $traces_{(\mathcal{A}_1 \parallel \mathcal{A}_2, P_1 \parallel P_2)} = \{\beta \mid \beta \text{ is an } (E, \emptyset)\text{-sequence and } \beta \upharpoonright (E_i, \emptyset) \in traces_{(\mathcal{A}_i, P_i)}, i \in \{1, 2\}\}$.*

Proof: Follows from definition of composition of automata with properties and Theorem 5.4. ■

6.4.2 Substitutivity Results

Theorem 6.25 *Suppose that \mathcal{A}_1 , \mathcal{A}_2 , and \mathcal{B} are TAs, \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with \mathcal{B} . Suppose that P_1 , P_2 , and Q are properties for \mathcal{A}_1 , \mathcal{A}_2 , and \mathcal{B} , respectively. If $(\mathcal{A}_1, P_1) \leq (\mathcal{A}_2, P_2)$ then $(\mathcal{A}_1, P_1) \parallel (\mathcal{B}, Q) \leq (\mathcal{A}_2, P_2) \parallel (\mathcal{B}, Q)$.*

This theorem can be strengthened with two corollaries.

Corollary 6.26 *Suppose $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1,$ and \mathcal{B}_2 are TAs, \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, \mathcal{B}_1 and \mathcal{B}_2 have the same external actions, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with each of \mathcal{B}_1 and \mathcal{B}_2 . Suppose that P_i and Q_i are properties for \mathcal{A}_i and \mathcal{B}_i , respectively for $i \in \{1, 2\}$. If $(\mathcal{A}_1, P_1) \leq (\mathcal{A}_2, P_2)$ and $(\mathcal{B}_1, Q_1) \leq (\mathcal{B}_2, Q_2)$ then $(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_1, Q_1) \leq (\mathcal{A}_2, P_2) \parallel (\mathcal{B}_2, Q_2)$.*

Corollary 6.27 *Suppose $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1,$ and \mathcal{B}_2 are TAs, \mathcal{A}_1 and \mathcal{A}_2 have the same external actions, \mathcal{B}_1 and \mathcal{B}_2 have the same external actions, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with each of \mathcal{B}_1 and \mathcal{B}_2 . Suppose that P_i and Q_i are properties for \mathcal{A}_i and \mathcal{B}_i , respectively for $i \in \{1, 2\}$. If $(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_2, Q_2) \leq (\mathcal{A}_2, P_2) \parallel (\mathcal{B}_2, Q_2)$ and $(\mathcal{B}_1, Q_1) \leq (\mathcal{B}_2, Q_2)$ then $(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_1, Q_1) \leq (\mathcal{A}_2, P_2) \parallel (\mathcal{B}_2, Q_2)$.*

7 Timed I/O Automata

In this section we refine the timed automaton model of Section 4 by distinguishing between input and output actions. Typically, an interaction between a system and its environment is modeled by using output and input actions to represent, respectively, the external events under the control of the system and the environment. We extend the results on simulation relations and composition from Sections 4 and 5 to this new setting. We also introduce special kinds of timed I/O automata: I/O feasible, progressive, and receptive TIOAs.

7.1 Definition of Timed I/O Automata

A *timed I/O automaton (TIOA)* \mathcal{A} is a tuple (\mathcal{B}, I, O) where

- $\mathcal{B} = (X, Q, \Theta, E, H, \mathcal{D}, \mathcal{T})$ is a timed automaton.
- I and O partition E into *input* and *output actions*, respectively. Actions in $L \triangleq H \cup O$ are called *locally controlled*; as before we write $A \triangleq E \cup H$.
- The following additional axioms are satisfied:

E1 (Input action enabling)

For every $\mathbf{x} \in Q$ and every $a \in I$, there exists $\mathbf{x}' \in Q$ such that $\mathbf{x} \xrightarrow{a} \mathbf{x}'$.

E2 (Time-passage enabling)

For every $\mathbf{x} \in Q$, there exists $\tau \in \mathcal{T}$ such that $\tau.lstate = \mathbf{x}$ and either

1. $\tau.ltime = \infty$, or
2. τ is closed and some $l \in L$ is enabled in $\tau.lstate$.

Input action enabling is the input enabling condition of ordinary I/O automata; it says that a TIOA is able to perform an input action at any time. The time-passage enabling condition says that a TIOA either allows time to advance forever, or it allows time to advance for a while, up to a point where it is prepared to react with some locally controlled action. Because TIOAs have no external variables, **E1** and **E2** are slightly simpler than the corresponding axioms for HIOAs.

Notation: As we did for TAs, we often denote the components of a TIOA \mathcal{A} by $\mathcal{B}_{\mathcal{A}}, I_{\mathcal{A}}, O_{\mathcal{A}}, X_{\mathcal{A}}, Q_{\mathcal{A}}, \Theta_{\mathcal{A}}$, etc., and those of a TIOA \mathcal{A}_i by $H_i, I_i, O_i, \dots, X_i, Q_i, \Theta_i$, etc. We sometimes omit these subscripts, where no confusion is likely. We abuse notation slightly by referring to a TIOA \mathcal{A} as a TA when we intend to refer to $\mathcal{B}_{\mathcal{A}}$.

Example 7.1 (TAs viewed as TIOAs) The automaton $TimedChannel(b, M)$ described in Example 4.1 can be turned into a TIOA by classifying the *send* actions as inputs, and the *receive* actions as outputs. Since there is no precondition for *send* actions, they are enabled in each state, so clearly the input enabling condition **E1** holds. It is also easy to see that axiom **E2** holds: in each state either *queue* is nonempty, in which case a *receive* output action is enabled after a point trajectory, or *queue* is empty, in which case time can advance forever.

The automaton $ClockSync(u, \rho)_i$ of Example 4.6 can be turned into a TIOA by classifying the *send* actions as outputs, and the *receive* actions as inputs. Axiom **E1** then holds trivially. Axiom **E2** holds since from each state either time can advance forever, or we have an outgoing trajectory (possibly of length 0) to a state in which $physclock = nextsend$, and from there a *send* output action is enabled. ■

7.2 Executions and Traces

An *execution fragment*, *execution*, *trace fragment*, or *trace* of a TIOA \mathcal{A} is defined to be an execution fragment, execution, trace fragment, or trace of the underlying TA $\mathcal{B}_{\mathcal{A}}$, respectively.

We say that an execution fragment of a TIOA is *locally-Zeno* if it is Zeno and contains infinitely many locally controlled actions, or equivalently, if it has finite limit time and contains infinitely many locally controlled actions.

7.3 Special Kinds of Timed I/O Automata

7.3.1 Feasible and I/O Feasible TIOAs

A TIOA $A = (\mathcal{B}, I, O)$ is defined to be feasible provided that its underlying TA \mathcal{B} is feasible according to the definition given in Section 4.3.1. As noted in Section 4.3.1, feasibility is a

basic requirement that any TA (or TIOA) should satisfy. I/O feasibility is a strengthened version of feasibility that take inputs into account. It says that the automaton is capable of providing some response from any state, for any sequence of input actions and any amount of intervening time-passage. In particular, it should allow time to pass to infinity if the environment does not submit any input actions. Formally, we define a TIOA to be *I/O feasible* provided that, for each state \mathbf{x} and each (I, \emptyset) -sequence β , there is some execution fragment α from \mathbf{x} such that $\alpha \upharpoonright (I, \emptyset) = \beta$. That is, an I/O feasible TIOA accommodates arbitrary input actions occurring at arbitrary times. The given (I, \emptyset) -sequence β describes the inputs and the amounts of intervening times.

7.3.2 Progressive TIOAs

A progressive TIOA never generates infinitely many locally controlled actions in finite time. Formally, a TIOA \mathcal{A} is *progressive* if it has no locally-Zeno execution fragments.

The following lemma says that any progressive TIOA is capable of advancing time forever.

Lemma 7.2 *Every progressive TIOA is feasible.*

Proof: Let \mathcal{A} be a progressive TIOA and let \mathbf{x} be a state of \mathcal{A} . Since \mathcal{A} is a TIOA it satisfies axiom **E2**. We construct an admissible execution fragment $\alpha = \alpha_0 \frown \alpha_1 \frown \alpha_2 \cdots$ from \mathbf{x} as follows.

1. $\alpha_0 = \wp(\mathbf{x})$.
2. For each $i > 0$,
 - (a) If there exists a trajectory τ from $\alpha_{i-1}.lstate$ such that $\tau.ltime = \infty$ then α_i is the final execution fragment in the sequence and $\alpha_i = \tau$.
 - (b) Otherwise, let τ_i be a closed execution fragment from $\alpha_{i-1}.lstate$ such that $l \in L$ is enabled in $\tau_i.lstate$. Define $\alpha_i = \tau_i \frown \tau_{i+1}$ where $\tau_{i+1} = \wp(\mathbf{y})$ and $\tau_i.lstate \xrightarrow{l} \mathbf{y}$.

The above construction either ends after finitely many stages such that the last trajectory of α is admissible, or goes through infinitely many stages such that α contains infinitely many local actions. In the former case, we know that α is admissible since it ends with an admissible trajectory. In the latter case, since \mathcal{A} is progressive, the fact that α has infinitely many local actions implies that α is admissible, as needed. ■

The following lemma says that a progressive TIOA is capable of allowing any amount of time to pass from any state.

Lemma 7.3 *Let \mathcal{A} be a progressive TIOA, let \mathbf{x} be a state of \mathcal{A} , and let $\tau \in \text{trajs}(\emptyset)$. Then there exists an execution fragment α of \mathcal{A} such that $\alpha.\text{fstate} = \mathbf{x}$ and $\alpha \upharpoonright (I, \emptyset) = \tau$.*

Proof: The result follows from the construction used in the proof of Lemma 7.2. Let α be an admissible execution fragment from \mathbf{x} constructed as in the proof of Lemma 7.2. Let α' be a prefix of α such that $\alpha' \upharpoonright (\emptyset, \emptyset) = \tau$. Since our construction uses no actions from I , we have $\alpha' \upharpoonright (I, \emptyset) = \alpha' \upharpoonright (\emptyset, \emptyset) = \tau$, as needed. ■

The following theorem says that a progressive TIOA is capable not just of allowing arbitrary amounts of time to pass, but of allowing arbitrary input actions at arbitrary times.

Theorem 7.4 *Every progressive TIOA is I/O feasible.*

Proof: Let \mathcal{A} be a progressive TIOA, let \mathbf{x} be a state of \mathcal{A} , and let $\beta = \tau_0 a_1 \tau_1 a_2 \tau_2 \dots$ be an (I, \emptyset) -sequence. We construct a finite or infinite sequence $\alpha_0 \alpha_1 \dots$ of execution fragments such that:

1. $\alpha_0.\text{fstate} = \mathbf{x}$.
2. For each nonfinal index i , $\alpha_i.\text{lstate} = \alpha_{i+1}.\text{fstate}$.
3. For each i , $(\alpha_0 \frown \alpha_1 \frown \dots \frown \alpha_i) \upharpoonright (I, \emptyset) = \tau_0 a_1 \tau_1 \dots \tau_i$.

The construction is carried out recursively. To define α_0 , we start with \mathbf{x} and use Lemma 7.3 to span τ_0 . For $i > 0$, we define α_i by starting with $\alpha_{i-1}.\text{lstate}$, using axiom **E1** to perform the input action a_i and move to a new state and then using Lemma 7.3 to span τ_i .

Let $\alpha = \alpha_0 \frown \alpha_1 \frown \dots$. By Lemma 3.8, α is an execution fragment of \mathcal{A} from \mathbf{x} such that $\alpha \upharpoonright (I, \emptyset) = \beta$, as needed. ■

7.3.3 Receptive Timed I/O Automata

In this section, we define the notion of *receptiveness* for TIOAs. A TIOA will be defined to be receptive provided that it admits a *strategy* for resolving its nondeterministic choices that never generates infinitely many locally controlled actions in finite time. This notion has an important consequence: A receptive TIOA provides some response from any state, for any sequence of discrete input actions at any times. This implies that the automaton has a nontrivial set of execution fragments, in fact, it has execution fragments that accommodate any inputs from the environment. The automaton cannot simply stop at

some point and refuse to allow time to elapse; it must allow time to pass to infinity if the environment does so. Previous studies of receptiveness properties include [12, 1, 36, 24]. The notion of receptiveness for TIOAs as discussed here is a special case of the same notion for HIOAs [22].

We build our definition of receptiveness on our earlier definition of progressive TIOAs. Namely, we define a *strategy* for resolving nondeterministic choices, and define receptiveness in terms of the existence of a progressive strategy.

We define a *strategy* for a TIOA \mathcal{A} to be a TIOA \mathcal{A}' that differs from \mathcal{A} only in that $\mathcal{D}' \subseteq \mathcal{D}$ and $\mathcal{T}' \subseteq \mathcal{T}$. That is, we require:

- $\mathcal{D}' \subseteq \mathcal{D}$.
- $\mathcal{T}' \subseteq \mathcal{T}$.
- $X = X'$, $Q = Q'$, $\Theta = \Theta'$, $E = E'$, $H = H'$, $I = I'$, and $O = O'$.

Our strategies are nondeterministic and memoryless. They provide a way of choosing some of the evolutions that are possible from each state \mathbf{x} of \mathcal{A} . The fact that the state set Q' of \mathcal{A}' is the same as the state set Q of \mathcal{A} implies that \mathcal{A}' chooses evolutions from every state of \mathcal{A} .

Notions of strategy have been used also in previous studies of receptiveness [12, 1, 36, 24]. However, in these earlier works, strategies have been formalized using two-player games rather than automata. Defining strategies using automata allows us to avoid introducing extra mathematical machinery.

Lemma 7.5 *If \mathcal{A}' is a strategy for \mathcal{A} , then every execution fragment of \mathcal{A}' is also an execution fragment of \mathcal{A} .*

We define a TIOA to be *receptive* if it has a progressive strategy. The following theorem says that any receptive TIOA can respond to any inputs from the environment.

Theorem 7.6 *Every receptive TIOA is I/O feasible.*

Proof: The proof is similar to that of the corresponding theorem for HIOAs [22]. ■

Example 7.7 (Progressive and receptive TIOAs) The time-bounded channel automaton described in Example 4.1 is not progressive since it allows for an infinite execution in which *send* and *receive* actions alternate without any passage of time in between. The time-bounded channel automaton is receptive, however, as we may construct a progressive strategy for it by adding a condition $u = \text{now}$ to the precondition of the *receive* action.

In this way we enforce that the channel operates maximally slow and messages are only delivered at their delivery deadline. The clock synchronization automaton of Example 4.6 is progressive (and therefore receptive) since it can only generate a locally controlled action each time its physical clock advances by u time units and the real time that elapses between two locally produced actions is at least $u(1 - \rho)$ time units. ■

7.4 Implementation Relationships

Two TIOAs \mathcal{A}_1 and \mathcal{A}_2 are *comparable* if their inputs and outputs coincide, that is, if $I_1 = I_2$ and $O_1 = O_2$. If \mathcal{A}_1 and \mathcal{A}_2 are comparable, then $\mathcal{A}_1 \leq \mathcal{A}_2$ is defined to mean that the traces of \mathcal{A}_1 are included among those of \mathcal{A}_2 : $\mathcal{A}_1 \leq \mathcal{A}_2 \triangleq \text{traces}_{\mathcal{A}_1} \subseteq \text{traces}_{\mathcal{A}_2}$.

Lemma 7.8 *Let $\mathcal{A}_1, \mathcal{A}_2$ be two comparable TIOAs and let $\mathcal{B}_1, \mathcal{B}_2$ be, respectively, the underlying TAs for \mathcal{A}_1 and \mathcal{A}_2 . Then \mathcal{B}_1 and \mathcal{B}_2 are comparable and $\mathcal{A}_1 \leq \mathcal{A}_2$ iff $\mathcal{B}_1 \leq \mathcal{B}_2$.*

Proof: Immediate from the definitions. ■

7.5 Simulation Relations

The definition of forward simulation for TIOAs is the same as for TAs. Formally, if $\mathcal{A}_1 = (\mathcal{B}_1, I_1, O_1)$ and $\mathcal{A}_2 = (\mathcal{B}_2, I_2, O_2)$ are two comparable TIOAs, then a forward simulation from \mathcal{A}_1 to \mathcal{A}_2 is a forward simulation from \mathcal{B}_1 to \mathcal{B}_2 .

Theorem 7.9 *If \mathcal{A}_1 and \mathcal{A}_2 are comparable TIOAs and there is a forward simulation from \mathcal{A}_1 to \mathcal{A}_2 , then $\mathcal{A}_1 \leq \mathcal{A}_2$.*

The definitions and results about backward simulations, history and prophecy relations for timed automata from Section 4 carry over to timed automata with input and output distinction in a similar fashion.

8 Operations on Timed I/O Automata

8.1 Composition

In this section, we define the operations of composition and hiding and present projection, pasting and substitutivity results for TIOAs. We revisit the special kinds of TIOAs introduced in Section 7 and show that the classes of progressive and receptive timed I/O automata are closed under composition, while this is not true for the class of I/O feasible automata.

8.1.1 Definitions and Basic Results

The definition of composition for TIOAs is based on the corresponding definition for TAs, but also takes the input/output structure into account. We say that TIOAs \mathcal{A}_1 and \mathcal{A}_2 are *compatible* if, for $i \neq j$, $X_i \cap X_j = H_i \cap A_j = O_i \cap O_j = \emptyset$.

Lemma 8.1 *If $\mathcal{A}_1 = (\mathcal{B}_1, I_1, O_1)$ and $\mathcal{A}_2 = (\mathcal{B}_2, I_2, O_2)$ are compatible TIOAs, then \mathcal{B}_1 and \mathcal{B}_2 are compatible TAs.*

If \mathcal{A}_1 and \mathcal{A}_2 are compatible TIOAs then their *composition* $\mathcal{A}_1 \parallel \mathcal{A}_2$ is defined to be the tuple $\mathcal{A} = (\mathcal{B}, I, O)$ where

- $\mathcal{B} = \mathcal{B}_1 \parallel \mathcal{B}_2$,
- $I = (I_1 \cup I_2) - (O_1 \cup O_2)$
- $O = O_1 \cup O_2$.

Thus, an external action of the composition is classified as an output if it is an output of one of the component automata, and otherwise it is classified as an input. The composition of two TIOAs is guaranteed to be a TIOA:

Theorem 8.2 *If \mathcal{A}_1 and \mathcal{A}_2 are TIOAs then $\mathcal{A}_1 \parallel \mathcal{A}_2$ is a TIOA.*

Proof: The proof is straightforward except for showing that Axiom **E2** is satisfied by the composition. Let \mathbf{x} be a state of $\mathcal{A}_1 \parallel \mathcal{A}_2$. We need to show the existence of a trajectory from \mathbf{x} that satisfies **E2**.

By definition of $\mathcal{A}_1 \parallel \mathcal{A}_2$, $\mathbf{x} \upharpoonright X_1$ is a state of \mathcal{A}_1 and $\mathbf{x} \upharpoonright X_2$ is a state of \mathcal{A}_2 . We know that both \mathcal{A}_1 and \mathcal{A}_2 satisfy **E2**. Let τ_1 be a trajectory of \mathcal{A}_1 with $\tau_1.fstate = \mathbf{x} \upharpoonright X_1$ that satisfies **E2**, let τ_2 be a trajectory of \mathcal{A}_2 with $\tau_2.fstate = \mathbf{x} \upharpoonright X_2$ that satisfies **E2**, and consider the following cases:

1. $\tau_1.ltime = \infty$ and $\tau_2.ltime = \infty$.
Then, define τ such that $\tau \downarrow X_1 = \tau_1$ and $\tau \downarrow X_2 = \tau_2$.
2. $\tau_1.ltime = \infty$ and τ_2 is closed where some $l \in L_2$ is enabled in $\tau_2.lstate$.
Then, define τ such that $\tau \downarrow X_1 = \tau_1 \upharpoonright \text{dom}(\tau_2)$ and $\tau \downarrow X_2 = \tau_2$.
3. τ_1 is closed where some $l \in L_1$ is enabled in $\tau_1.lstate$ and $\tau_2.ltime = \infty$.
Then, define τ such that $\tau \downarrow X_1 = \tau_1$ and $\tau \downarrow X_2 = \tau_2 \upharpoonright \text{dom}(\tau_1)$.

4. τ_1 is closed where some $l \in L_1$ is enabled in $\tau_1.lstate$ and τ_2 is closed where some $l \in L_2$ is enabled in $\tau_2.lstate$.
 If $dom(\tau_1) \subseteq dom(\tau_2)$, then define τ such that $\tau \downarrow X_1 = \tau_1$ and $\tau \downarrow X_2 = \tau_2 \upharpoonright dom(\tau_1)$. Otherwise, define τ such that $\tau \downarrow X_1 = \tau_1 \upharpoonright dom(\tau_2)$ and $\tau \downarrow X_2 = \tau_2$.

In all the cases, by definition of trajectories for a TIOA, τ is a trajectory of $\mathcal{A}_1 \parallel \mathcal{A}_2$ from \mathbf{x} , which satisfies **E2** by construction. ■

Note that this theorem is stronger than the corresponding theorem (Theorem 6.12 in [22]) for general HIOAs. Two HIOAs \mathcal{A}_1 and \mathcal{A}_2 are required to be “strongly compatible” for their composition to be a hybrid I/O automaton. This extra condition is needed to rule out dependencies among external variables that may prevent the component automata from evolving together. The absence of external variables in TIOA eliminates this kind of problematic behavior. Thus, for the timed case, we do not require the notion of strong compatibility that was needed for the hybrid case.

Composition of TIOAs satisfies the following projection and pasting result, which follows from Theorem 5.4.

Theorem 8.3 *Let \mathcal{A}_1 and \mathcal{A}_2 be comparable TIOAs, and let $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$. Then $traces_{\mathcal{A}}$ is exactly the set of (E, \emptyset) -sequences whose restrictions to \mathcal{A}_1 and \mathcal{A}_2 are traces of \mathcal{A}_1 and \mathcal{A}_2 , respectively. That is, $traces_{\mathcal{A}} = \{\beta \mid \beta \text{ is an } (E, \emptyset)\text{-sequence and } \beta \upharpoonright (E_i, \emptyset) \in traces_{\mathcal{A}_i}, i = \{1, 2\}\}$.*

8.1.2 Substitutivity Results

The following theorem is analogous to Theorem 5.8 for TAs without input/output distinction. It shows that the introduction of the input/output distinction does not cause any changes to the substitutivity results we obtained for general TAs.

Theorem 8.4 *Suppose \mathcal{A}_1 and \mathcal{A}_2 are comparable TIOAs with $\mathcal{A}_1 \leq \mathcal{A}_2$. Suppose that \mathcal{B} is a TIOA that is compatible with each of \mathcal{A}_1 and \mathcal{A}_2 . Then $\mathcal{A}_1 \parallel \mathcal{B} \leq \mathcal{A}_2 \parallel \mathcal{B}$.*

The corollaries below follow from the Corollaries 5.9 and 5.10 of Theorem 5.8.

Corollary 8.5 *Suppose $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1$, and \mathcal{B}_2 are TIOAs, \mathcal{A}_1 and \mathcal{A}_2 are comparable, \mathcal{B}_1 and \mathcal{B}_2 are comparable, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with each of \mathcal{B}_1 and \mathcal{B}_2 . If $\mathcal{A}_1 \leq \mathcal{A}_2$ and $\mathcal{B}_1 \leq \mathcal{B}_2$ then $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.*

Corollary 8.6 *Suppose $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1$, and \mathcal{B}_2 are TAs, \mathcal{A}_1 and \mathcal{A}_2 are comparable, \mathcal{B}_1 and \mathcal{B}_2 are comparable, and each of \mathcal{A}_1 and \mathcal{A}_2 is compatible with each of \mathcal{B}_1 and \mathcal{B}_2 . If $\mathcal{A}_1 \parallel \mathcal{B}_2 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ and $\mathcal{B}_1 \leq \mathcal{B}_2$ then $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.*

The basic substitutivity theorem, Theorem 8.4, is desirable for any formalism for interacting processes. For design purposes, it enables one to refine individual components without violating the correctness of the system as a whole. For verification purposes, it enables one to prove that a composite system satisfies its specification by proving that each component satisfies its specification, thereby breaking down the verification task into more manageable pieces. However, it might not always be possible or easy to show that each component \mathcal{A}_1 (resp. \mathcal{B}_1) satisfies its specification \mathcal{A}_2 (resp. \mathcal{B}_2) without using any assumptions about the environment of the component. *Assume-guarantee* style results such as those presented in [19, 33, 38, 1, 2, 18, 39] are special kinds of substitutivity results that state what *guarantees* are expected from each component in an environment constrained by certain *assumptions*. Since the environment of each component consists of the other components in the system, assume-guarantee style results need to break the circular dependencies between the assumptions and guarantees for components. We present below two assume-guarantee style theorems Theorem 8.7 and Corollary 8.8, which can be used for proving that a system specified as a composite automaton $\mathcal{A}_1 \parallel \mathcal{B}_1$ implements a specification represented by a composite automaton $\mathcal{A}_2 \parallel \mathcal{B}_2$.

The main idea behind Theorem 8.7 is to assume that \mathcal{A}_1 implements \mathcal{A}_2 in a context represented by \mathcal{B}_2 , and symmetrically that \mathcal{B}_1 implements \mathcal{B}_2 in a context represented by \mathcal{A}_2 where \mathcal{A}_2 and \mathcal{B}_2 are automata whose trace sets are closed under limits. The requirement about limit-closure implies that \mathcal{A}_2 and \mathcal{B}_2 specify trace safety properties. Moreover, we assume that the trace sets of \mathcal{A}_2 and \mathcal{B}_2 are closed under time-extension. That is, the automata allow arbitrary time-passage. This is the most general assumption one could make to ensure that $\mathcal{A}_2 \parallel \mathcal{B}_2$ does not impose stronger constraints on time-passage than $\mathcal{A}_1 \parallel \mathcal{B}_1$. Note that the definitions of limit and time extension of a hybrid sequence can be found in Section 9.2.

Theorem 8.7 *Suppose $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1, \mathcal{B}_2$ are TIOAs such that \mathcal{A}_1 and \mathcal{A}_2 are comparable, \mathcal{B}_1 and \mathcal{B}_2 are comparable, and \mathcal{A}_i is compatible with \mathcal{B}_i for $i \in \{1, 2\}$. Suppose further that:*

1. *The sets $\text{traces}_{\mathcal{A}_2}$ and $\text{traces}_{\mathcal{B}_2}$ are closed under limits.*
2. *The sets $\text{traces}_{\mathcal{A}_2}$ and $\text{traces}_{\mathcal{B}_2}$ are closed under time-extension.*
3. *$\mathcal{A}_1 \parallel \mathcal{B}_2 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ and $\mathcal{A}_2 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.*

Then $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.

Proof: We first prove by induction on the length of traces of $\mathcal{A}_1 \parallel \mathcal{B}_1$ that every closed trace of $\mathcal{A}_1 \parallel \mathcal{B}_1$ is a trace of $\mathcal{A}_2 \parallel \mathcal{B}_2$.

For the base case, let β be a trace of $\mathcal{A}_1 \parallel \mathcal{B}_1$ such that $\beta \in \text{trajs}(\emptyset)$ (a single trajectory over the empty set of variables). By Axiom **T0** in the definition of a TA, we know that

\mathcal{A}_2 and \mathcal{B}_2 have traces α_1 and α_2 such that $\alpha_1.ltime = \alpha_2.ltime = 0$. By Assumption 2 we have $\alpha_1 \frown \beta \in traces_{\mathcal{A}_2}$ and $\alpha_2 \frown \beta \in traces_{\mathcal{B}_2}$. Since, $\alpha_1 \frown \beta = \beta$ and $\alpha_2 \frown \beta = \beta$, it follows that $\beta \in traces_{\mathcal{A}_2}$ and $\beta \in traces_{\mathcal{B}_2}$. By pasting using Theorem 8.3, $\beta \in traces_{\mathcal{A}_2 \parallel \mathcal{B}_2}$, as needed.

For the inductive step we consider the following cases:

1. $\beta = \beta' a \tau$, where a is an output action of \mathcal{A}_1 and τ is a point trajectory.

Then $\beta \uparrow (E_{\mathcal{A}_1}, \emptyset) \in traces_{\mathcal{A}_1}$ by projection using Theorem 8.3. By inductive hypothesis, $\beta' \in traces_{\mathcal{A}_2 \parallel \mathcal{B}_2}$. So $\beta' \uparrow (E_{\mathcal{B}_2}, \emptyset) \in traces_{\mathcal{B}_2}$, by projection using Theorem 8.3. Let α be an execution of \mathcal{B}_2 such that $trace(\alpha) = \beta' \uparrow (E_{\mathcal{B}_2}, \emptyset)$. Since \mathcal{A}_1 and \mathcal{B}_1 are compatible TIOAs, \mathcal{B}_1 and \mathcal{B}_2 are comparable, and a is an output action of \mathcal{A}_1 , we know that either a is an input action of \mathcal{B}_2 or the action set of \mathcal{B}_2 does not contain a . In the former case, by the input-enabling axiom (**E1**) we know that there exists \mathbf{x}' such that $(\alpha.lstate, a, \mathbf{x}')$ is a discrete transition of \mathcal{B}_2 . It follows that $\beta \uparrow (E_{\mathcal{B}_2}, \emptyset) \in traces_{\mathcal{B}_2}$. In the latter case, since $\beta \uparrow (E_{\mathcal{B}_2}, \emptyset) = \beta' \uparrow (E_{\mathcal{B}_2}, \emptyset)$ and $\beta' \uparrow (E_{\mathcal{B}_2}, \emptyset) \in traces_{\mathcal{B}_2}$ we get $\beta \uparrow (E_{\mathcal{B}_2}, \emptyset) \in traces_{\mathcal{B}_2}$. By pasting using Theorem 8.3, $\beta \in traces_{\mathcal{A}_1 \parallel \mathcal{B}_2}$. Then by Assumption 3, $\beta \in traces_{\mathcal{A}_2 \parallel \mathcal{B}_2}$.

2. $\beta = \beta' b \tau$, where b is an output action of \mathcal{B}_1 and τ is a point trajectory.

This case is symmetric with the previous one.

3. $\beta = \beta' c \tau$, where c is an input action of both \mathcal{A}_1 and \mathcal{B}_1 and τ is a point trajectory.

By inductive hypothesis, $\beta' \in traces_{\mathcal{A}_2 \parallel \mathcal{B}_2}$. By projection using Theorem 8.3 we get $\beta' \uparrow (E_{\mathcal{A}_2}, \emptyset) \in traces_{\mathcal{A}_2}$ and $\beta' \uparrow (E_{\mathcal{B}_2}, \emptyset) \in traces_{\mathcal{B}_2}$. Let α be an execution of \mathcal{A}_2 such that $trace(\alpha) = \beta' \uparrow (E_{\mathcal{A}_2}, \emptyset)$. Since \mathcal{A}_1 and \mathcal{A}_2 are comparable and a is an input action of \mathcal{A}_1 we know that a is an input action of \mathcal{A}_2 . By the input-enabling axiom (**E1**) we know that there exists \mathbf{x}' such that $(\alpha.lstate, a, \mathbf{x}')$ is a discrete transition of \mathcal{A}_2 . It follows that $\beta \uparrow (E_{\mathcal{A}_2}, \emptyset) \in traces_{\mathcal{A}_2}$. Similarly, let α' be an execution of \mathcal{B}_2 such that $trace(\alpha') = \beta' \uparrow (E_{\mathcal{B}_2}, \emptyset)$. Since \mathcal{B}_1 and \mathcal{B}_2 are comparable and a is an input action of \mathcal{B}_1 we know that a is an input action of \mathcal{B}_2 . By the input-enabling axiom (**E1**) we know that there exists \mathbf{y}' such that $(\alpha'.lstate, a, \mathbf{y}')$ is a discrete transition of \mathcal{B}_2 . It follows that $\beta \uparrow (E_{\mathcal{B}_2}, \emptyset) \in traces_{\mathcal{B}_2}$. By pasting using Theorem 8.3, we get $\beta \in traces_{\mathcal{A}_2 \parallel \mathcal{B}_2}$.

4. $\beta = \beta' d \tau$, where d is an input action of \mathcal{A}_1 but not an action of \mathcal{B}_1 and τ is a point trajectory.

By inductive hypothesis, $\beta' \in traces_{\mathcal{A}_2 \parallel \mathcal{B}_2}$. By projection using Theorem 8.3, we have $\beta' \uparrow (E_{\mathcal{A}_2}, \emptyset) \in traces_{\mathcal{A}_2}$ and $\beta' \uparrow (E_{\mathcal{B}_2}, \emptyset) \in traces_{\mathcal{B}_2}$. Let α be an execution of \mathcal{A}_2 such that $trace(\alpha) = \beta' \uparrow (E_{\mathcal{A}_2}, \emptyset)$. Since \mathcal{A}_1 and \mathcal{A}_2 are comparable TIOAs and a is an input action of \mathcal{A}_1 , a must be an input action of \mathcal{A}_2 . By the input-enabling axiom (**E1**) we know that there exists \mathbf{x}' such that $(\alpha.lstate, a, \mathbf{x}')$ is a discrete transition of \mathcal{A}_2 . It follows that $\beta \uparrow (E_{\mathcal{A}_2}, \emptyset) \in traces_{\mathcal{A}_2}$. Since \mathcal{B}_1 and

\mathcal{B}_2 are comparable and a is not an action of \mathcal{B}_1 , a cannot be an external action of \mathcal{B}_2 . Therefore, $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) = \beta' \upharpoonright (E_{\mathcal{B}_2}, \emptyset)$. Since $\beta' \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$ we get $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. By pasting using Theorem 8.3, we get $\beta \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_2}$.

5. $\beta = \beta' d\tau$, where d is an input action of \mathcal{B}_1 but not an action of \mathcal{A}_1 and τ is a point trajectory.

This case is symmetric with the previous one.

6. $\beta = \beta' \frown \beta''$, where β'' is a hybrid sequence consisting of a single trajectory τ .

By inductive hypothesis, $\beta' \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_2}$. By projection using Theorem 8.3, we get $\beta' \upharpoonright (E_{\mathcal{A}_2}, \emptyset) \in \text{traces}_{\mathcal{A}_2}$ and $\beta' \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. By Assumption 2, we have $\beta' \upharpoonright (E_{\mathcal{A}_2}, \emptyset) \frown \beta'' \upharpoonright (E_{\mathcal{A}_2}, \emptyset) \in \text{traces}_{\mathcal{A}_2}$ and $\beta' \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \frown \beta'' \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. Then by pasting using Theorem 8.3, $\beta \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_2}$, as needed.

We have thus shown that every closed trace of $\mathcal{A}_1 \parallel \mathcal{B}_1$ is a trace of $\mathcal{A}_2 \parallel \mathcal{B}_2$. Now consider any non-closed trace β of $\mathcal{A}_1 \parallel \mathcal{B}_1$. This β can be written as the limit of a sequence $\beta_1 \beta_2 \dots$ of closed traces of $\mathcal{A}_1 \parallel \mathcal{B}_1$. By the first part of the proof we know that each $\beta_i \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_2}$, and by projection using Theorem 8.3 each $\beta_i \upharpoonright (E_{\mathcal{A}_2}, \emptyset)$ is a closed trace of \mathcal{A}_2 , and $\beta_i \upharpoonright (E_{\mathcal{B}_2}, \emptyset)$ is a closed trace of \mathcal{B}_2 . We know that $\beta \upharpoonright (E_{\mathcal{A}_2}, \emptyset)$ is the limit of the $\beta_i \upharpoonright (E_{\mathcal{A}_2}, \emptyset)$ and similarly $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset)$ is the limit of the $\beta_i \upharpoonright (E_{\mathcal{B}_2}, \emptyset)$. Since the sets $\text{traces}_{\mathcal{A}_2}$ and $\text{traces}_{\mathcal{B}_2}$ are limit-closed by Assumption 1, we get $\beta \upharpoonright (E_{\mathcal{A}_2}, \emptyset) \in \text{traces}_{\mathcal{A}_2}$ and $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. Finally, by pasting using Theorem 8.3, we get $\beta \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_2}$. ■

Note that automata with FIN and timing-independence (see Section 4.3.1 for definitions) constitute examples for context automata \mathcal{A}_2 and \mathcal{B}_2 that satisfy Assumptions 1 and 2. The property FIN implies Assumption 1 (Lemma 4.18) and timing-independence implies Assumption 2.

Theorem 8.7 has a corollary, Corollary 8.8 below, which can be used in the decomposition of proofs even when \mathcal{A}_2 and \mathcal{B}_2 neither admit arbitrary time-passage nor have limit-closed trace sets. The main idea behind this corollary is to assume that \mathcal{A}_1 implements \mathcal{A}_2 in a context \mathcal{B}_3 that is a variant of \mathcal{B}_2 , and symmetrically that \mathcal{B}_1 implements \mathcal{B}_2 in a context that is a variant of \mathcal{A}_2 . That is, the correctness of implementation relationship between \mathcal{A}_1 and \mathcal{A}_2 does not depend on all the environment constraints, just on those expressed by \mathcal{B}_3 (symmetrically for $\mathcal{B}_1, \mathcal{B}_2$, and \mathcal{A}_3). In order to use this corollary to prove $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ one needs to be able to find appropriate variants of \mathcal{A}_2 and \mathcal{B}_2 that meet the required closure properties. This corollary prompts one to pin down what is essential about the behavior of the environment in proving the intended implementation relationship, and also allows one to avoid the unnecessary details of the environment in proofs. In Section 9 we extend this corollary to the case where properties, typically liveness properties, are added to automaton specifications.

Corollary 8.8 *Suppose $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ are TIOAs such that $\mathcal{A}_1, \mathcal{A}_2$, and \mathcal{A}_3 are comparable, $\mathcal{B}_1, \mathcal{B}_2$, and \mathcal{B}_3 are comparable, and \mathcal{A}_i is compatible with \mathcal{B}_i for $i \in \{1, 2, 3\}$. Suppose further that:*

1. *The sets $\text{traces}_{\mathcal{A}_3}$ and $\text{traces}_{\mathcal{B}_3}$ are closed under limits.*
2. *The sets $\text{traces}_{\mathcal{A}_3}$ and $\text{traces}_{\mathcal{B}_3}$ are closed under time-extension.*
3. *$\mathcal{A}_2 \parallel \mathcal{B}_3 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$ and $\mathcal{A}_3 \parallel \mathcal{B}_2 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$.*
4. *$\mathcal{A}_1 \parallel \mathcal{B}_3 \leq \mathcal{A}_2 \parallel \mathcal{B}_3$ and $\mathcal{A}_3 \parallel \mathcal{B}_1 \leq \mathcal{A}_3 \parallel \mathcal{B}_2$.*

Then $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.

Proof: Since $\mathcal{A}_2 \leq \mathcal{A}_3$ by Assumption 3 and $\mathcal{A}_1 \parallel \mathcal{B}_3 \leq \mathcal{A}_2 \parallel \mathcal{B}_3$ by Assumption 4, we get $\mathcal{A}_1 \parallel \mathcal{B}_3 \leq \mathcal{A}_2 \parallel \mathcal{B}_3 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$, by Theorem 8.4. Similarly, we have $\mathcal{A}_3 \parallel \mathcal{B}_1 \leq \mathcal{A}_3 \parallel \mathcal{B}_2 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$. Since $\mathcal{A}_1 \parallel \mathcal{B}_3 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$ and $\mathcal{A}_3 \parallel \mathcal{B}_1 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$, by using Assumptions 1 and 2, and Theorem 8.7 we have $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$.

Let β be a trace of $\mathcal{A}_1 \parallel \mathcal{B}_1$. By projection using Theorem 8.3, $\beta \upharpoonright (E_{\mathcal{A}_1}, \emptyset) \in \text{traces}_{\mathcal{A}_1}$ and $\beta \upharpoonright (E_{\mathcal{B}_1}, \emptyset) \in \text{traces}_{\mathcal{B}_1}$. Since $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_3 \parallel \mathcal{B}_3$, we know that $\beta \in \text{traces}_{\mathcal{A}_3 \parallel \mathcal{B}_3}$. By projection using Theorem 8.3, $\beta \upharpoonright (E_{\mathcal{A}_3}, \emptyset) \in \text{traces}_{\mathcal{A}_3}$ and $\beta \upharpoonright (E_{\mathcal{B}_3}, \emptyset) \in \text{traces}_{\mathcal{B}_3}$. By pasting using Theorem 8.3, we have $\beta \in \text{traces}_{\mathcal{A}_1 \parallel \mathcal{B}_3}$ and $\beta \in \text{traces}_{\mathcal{A}_3 \parallel \mathcal{B}_1}$. By Assumption 4, we get $\beta \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_3}$ and $\beta \in \text{traces}_{\mathcal{A}_3 \parallel \mathcal{B}_2}$. Then, by projection using Theorem 8.3, $\beta \upharpoonright (E_{\mathcal{A}_2}, \emptyset) \in \text{traces}_{\mathcal{A}_2}$ and $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. Finally, by pasting using Theorem 8.3 we have $\beta \in \text{traces}_{\mathcal{A}_2 \parallel \mathcal{B}_2}$, as needed. \blacksquare

Example 8.9 (Using environment assumptions to prove safety)

This example illustrates that, in cases where specifications \mathcal{A}_2 and \mathcal{B}_2 satisfy certain closure properties, it is possible to decompose the proof of $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ by using Theorem 8.7, even if it is not the case that $\mathcal{A}_1 \leq \mathcal{A}_2$ or $\mathcal{B}_1 \leq \mathcal{B}_2$.

The automata *AlternateA* and *AlternateB* in Figure 16 are timing-independent automata in which no consecutive outputs occur without inputs happening in between. *AlternateA* and *AlternateB* perform a handshake, outputting an alternating sequence of a and b actions when they are composed. The automata *CatchUpA* and *CatchUpB* in Figure 17 are timing-dependent automata that do not necessarily alternate inputs and outputs as *AlternateA* and *AlternateB*. *CatchUpA* can perform an arbitrary number of b actions, and can perform an a provided that $\text{count}_a \leq \text{count}_b$. It allows count_a to increase to one more than count_b . *CatchUpB* can perform an arbitrary number of a actions, and can perform a b provided that $\text{count}_a \geq \text{count}_b + 1$. It allows count_b to reach count_a . Timing constraints require each output to occur exactly one time unit after the last action. *CatchUpA* and *CatchUpB* perform an alternating sequence of a actions and b actions when they are composed.

Automaton *AlternateA*

Variables X : **discrete** $myturn \in Bool$ **initially** true
States Q : $val(X)$
Actions A : **input b , output a**
Transitions \mathcal{D} : **input b** **output a**
 effect **precondition**
 $myturn := true$ $myturn$
 effect
 $myturn := false$
Trajectories \mathcal{T} : **satisfies**
 constant($myturn$)

Automaton *AlternateB*

Variables X : **discrete** $myturn \in Bool$ **initially** false
States Q : $val(X)$
Actions A : **input a , output b**
Transitions \mathcal{D} : **input a** **output b**
 effect **precondition**
 $myturn := true$ $myturn$
 effect
 $myturn := false$
Trajectories \mathcal{T} : **satisfies**
 constant($myturn$)

Figure 16: Example automata for \mathcal{A}_2 and \mathcal{B}_2 in Theorem 8.7

Suppose that we want to prove that $CatchUpA \parallel CatchUpB \leq AlternateA \parallel AlternateB$. We cannot apply the basic substitutivity theorem Theorem 8.7, in particular Corollary 8.5, since the assertions $CatchUpA \leq AlternateA$ and $CatchUpB \leq AlternateB$ are not true. Consider the trace $\tau_0 b \tau_1 a \tau_2 a \tau_3$ of $CatchUpA$ where τ_0, τ_1, τ_2 and τ_3 are trajectories with limit time 1. After having performed one b and one a , $CatchUpA$ can perform another a . But, this is impossible for $AlternateA$ which needs an input to enable the second a . $AlternateA$ and $CatchUpA$ behave similarly only when put in a context that imposes alternation.

It is easy to check that $AlternateA$ and $AlternateB$ satisfy the closure properties required by Assumptions 1 and 2 of Theorem 8.7 and, hence can be substituted for \mathcal{A}_2

and \mathcal{B}_2 respectively. Similarly, we can easily check that Assumption 3 is satisfied if we substitute $CatchUpA$ for \mathcal{A}_1 and $CatchUpB$ for \mathcal{B}_1 . ■

Example 8.10 (Extracting essential environment assumptions with auxiliary automata) This example illustrates that it may be possible to decompose verification, using Corollary 8.8, in cases where Theorem 8.7 is not applicable. If the aim is to show $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$ where \mathcal{A}_2 and \mathcal{B}_2 do not satisfy the assumptions of Theorem 8.7, then we find appropriate context automata \mathcal{A}_3 and \mathcal{B}_3 that abstract from those details of \mathcal{A}_2 and \mathcal{B}_2 that are not essential in proving $\mathcal{A}_1 \parallel \mathcal{B}_1 \leq \mathcal{A}_2 \parallel \mathcal{B}_2$.

Consider the automata $UseOldInputA$ and $UseOldInputB$ in Figure 18. $UseOldInputA$ keeps track of whether or not it is $UseOldInputA$'s turn, and when it is $UseOldInputA$'s turn, it keeps track of the next time it is supposed to perform an output. The number of outputs that $UseOldInputA$ can perform is bounded by a natural number. In the case of repeated b inputs, it is the oldest input that determines when the next output will occur. The automaton $UseOldInputB$ is the same as $UseOldInputA$ (inputs and outputs reversed) except that the turn variable of $UseOldInputB$ is set to false initially. Note that $UseOldInputA$ and $UseOldInputA$ are not timing-independent and their trace sets are not limit-closed. For each automaton, there are infinitely many start states, one for each natural number. We can build an infinite chain of traces, where each element in the chain corresponds to an execution starting from a distinct start state. The limit of such a chain, which contains infinitely many outputs, cannot be a trace of $UseOldInputA$ or $UseOldInputA$ since the number of outputs they can perform is bounded by a natural number. The automaton $UseNewInputA$ in Figure 19 behaves similarly to $UseOldInputA$ except for the handling of inputs. In the case of repeated b inputs, it is the most recent input that determines when the next output will occur. The automaton $UseNewInputB$ in Figure 19 is the same as $UseNewInputA$ (inputs and outputs reversed) except that the turn variable of $UseNewInputB$ is set to false initially.

Suppose that we want to prove that:

$$UseNewInputA \parallel UseNewInputB \leq UseOldInputA \parallel UseOldInputB.$$

Theorem 8.7 is not applicable here because the high-level automata $UseOldInputA$ and $UseOldInputB$ do not satisfy the required closure properties. However, we can use Corollary 8.8 to decompose verification. It requires us to find auxiliary automata that are less restrictive than $UseOldInputA$ and $UseOldInputB$ but that are restrictive enough to express the constraints that should be satisfied by the environment, for $UseNewInputA$ to implement $UseOldInputA$ and for $UseNewInputB$ to implement $UseOldInputB$.

The automata $AlternateA$ and $AlternateB$ in Figure 16 can be used as auxiliary automata in this example. They satisfy the closure properties required by Corollary 8.8 and impose alternation, which is the only additional condition to ensure the needed trace inclusion.

We can define a forward simulation relation from $UseNewInputA \parallel UseNewInputB$ to $UseOldInputA \parallel UseOldInputB$, which is based on the equality of the turn variables of the implementation and the specification automata. The fact that this simulation relation only uses the equality of turn variables reinforces the idea that the auxiliary contexts, which only keep track of their turn, capture exactly what is needed for the proof of $UseNewInputA \parallel UseNewInputB \leq UseOldInputA \parallel UseOldInputB$. We can observe that a direct proof of this assertion would require one to deal with state variables such as *maxout* and *next* of both $UseOldInputA$ and $UseOldInputB$, which do not play any essential role in the proof. On the other hand, by decomposing the proof along the lines of Corollary 8.8 some of the unnecessary details can be avoided. Even though, this is a toy example with an easy proof it should not be hard to observe how this simplification would scale to large proofs. ■

8.1.3 Composition of Special Kinds of TIOAs

The following example illustrates that the set of I/O feasible TIOAs is not closed under composition:

Example 8.11 (Two I/O feasible TIOAs whose composition is not I/O feasible)

Consider two I/O feasible TIOAs \mathcal{A} and \mathcal{B} , where $O_{\mathcal{A}} = I_{\mathcal{B}} = \{a\}$ and $O_{\mathcal{B}} = I_{\mathcal{A}} = \{b\}$. Suppose that \mathcal{A} performs its output a at time 0 and then waits, allowing time to pass, until it receives input b . If and when it receives b , it responds with output a without allowing any time to pass (and ignoring any inputs that occur before it has a chance to perform its output). On the other hand, \mathcal{B} starts out waiting, allowing time to pass, until it receives input a . If and when it receives a , it responds with output b without allowing time to pass.

It is not difficult to see that \mathcal{A} and \mathcal{B} are individually I/O feasible. We claim that the composition $\mathcal{A} \parallel \mathcal{B}$ is not I/O feasible. To see this, consider the start state of $\mathcal{A} \parallel \mathcal{B}$ and the unique input sequence β with $\beta.time = \infty$; β simply allows time to pass to infinity. The composition $\mathcal{A} \parallel \mathcal{B}$ has no way of accommodating this input, since it will never allow time to pass beyond 0. ■

On the other hand, the following theorems say that the classes of progressive and receptive TIOAs are closed under composition:

Theorem 8.12 *If \mathcal{A}_1 and \mathcal{A}_2 are compatible progressive TIOAs, then their composition is also progressive.*

Proof: The proof is similar to the proof of Theorem 7.4 in [22]. The main idea behind the proof is that a Zeno execution of $\mathcal{A}_1 \parallel \mathcal{A}_2$ with infinitely many locally controlled contains

infinitely many locally controlled actions of either \mathcal{A}_1 or \mathcal{A}_2 . Suppose without loss of generality that the automaton that contributes infinitely many locally controlled actions is \mathcal{A}_1 . Then the projection onto \mathcal{A}_1 violates progressiveness for \mathcal{A}_1 . ■

Theorem 8.13 *Let \mathcal{A}_1 and \mathcal{A}_2 be two compatible TIOAs with strategies \mathcal{A}'_1 and \mathcal{A}'_2 , respectively. Then $\mathcal{A}'_1 \parallel \mathcal{A}'_2$ is a strategy for $\mathcal{A}_1 \parallel \mathcal{A}_2$.*

Proof: The proof is similar to the proof of Theorem 7.7 in [22]. ■

Now, we can state the main result of this section, which follows easily from the previous two theorems. It shows that the class of receptive TIOAs is closed under composition.

Theorem 8.14 *Let \mathcal{A}_1 and \mathcal{A}_2 be two compatible receptive TIOAs with progressive strategies \mathcal{A}'_1 and \mathcal{A}'_2 , respectively. Then $\mathcal{A}_1 \parallel \mathcal{A}_2$ is a receptive TIOA with progressive strategy $\mathcal{A}'_1 \parallel \mathcal{A}'_2$.*

Example 8.15 (Composition of receptive TIOAs) Theorem 8.14 implies that the composition of clock synchronization automata with channel automata described in Example 5.7 (viewed as TIOAs as explained in Example 7.1) is receptive. By Theorem 7.6 we also have that it is I/O feasible. ■

In fact, the fact that the set of I/O feasible TIOAs is not closed under composition motivated the definition of the more restrictive class of receptive TIOAs. That is, receptiveness is a reasonable sufficient condition that implies I/O feasibility, and that also is preserved by composition.

The special case of the HIOA model, represented by the TIOA model, has simpler and stronger composition theorems than the general HIOA model. In particular, the main compositionality result for receptive HIOAs (Theorem 7.12 in [22]) has a more intricate proof than ours. It makes an assumption about the existence of strongly compatible strategies (discussed briefly at the end of Section 8.1.1) and needs an additional lemma that shows that if two HIOAs \mathcal{A}_1 and \mathcal{A}_2 which may not be strongly compatible have strongly compatible strategies \mathcal{A}'_1 and \mathcal{A}'_2 , then \mathcal{A}_1 and \mathcal{A}_2 are also strongly compatible.

8.2 Hiding

We extend the definition of action hiding to any TIOA \mathcal{A} . For TIOAs, we consider hiding outputs only (but not inputs), by converting them to internal actions. Namely, if $O \subseteq O_{\mathcal{A}}$, then $\text{ActHide}(O, \mathcal{A})$ is the TIOA \mathcal{B} that is equal to \mathcal{A} except that $O_{\mathcal{B}} = O_{\mathcal{A}} - O$ and $H_{\mathcal{B}} = H_{\mathcal{A}} \cup O$.

Lemma 8.16 *If \mathcal{A} is a TIOA and $O \subseteq O_{\mathcal{A}}$ then $\text{ActHide}(O, \mathcal{A})$ is a TIOA.*

Lemma 8.17 *If \mathcal{A} is a TIOA and $O \subseteq O_{\mathcal{A}}$ then $\text{traces}_{\text{ActHide}(O, \mathcal{A})} = \{\beta \upharpoonright (O_{\mathcal{A}} - O, V_{\mathcal{A}}) \mid \beta \in \text{traces}_{\mathcal{A}}\}$.*

Theorem 8.18 *Suppose \mathcal{A} and \mathcal{B} are TIOAs with $\mathcal{A} \leq \mathcal{B}$, and suppose $O \subseteq O_{\mathcal{A}}$. Then $\text{ActHide}(O, \mathcal{A}) \leq \text{ActHide}(O, \mathcal{B})$.*

9 Properties for Timed I/O Automata

In this section, we present some definitions and results for timed I/O automata with properties. We focus on the definitions and results, such as those that involve receptiveness for properties, that become of interest with the introduction of input, output distinction to the model.

9.1 Definitions and Basic Results

A *property* for a timed I/O automaton $\mathcal{A} = (\mathcal{B}, I, O)$ is defined to be a property of its underlying timed automaton, that is, it is a subset of the execution fragments of \mathcal{B} .

Now, we introduce a notion of liveness property that takes into account how a system responds to inputs from its environment. A property P for a TIOA \mathcal{A} is defined to be an *I/O liveness* property provided that for each closed execution fragment α of \mathcal{A} and each (I, \emptyset) -sequence β , there is some execution fragment α' such that $\alpha' \upharpoonright (I, \emptyset) = \beta$ and $\alpha \frown \alpha' \in P$. In other words, no matter how \mathcal{A} behaves for a finite period of time, and no matter what inputs arrive, it is still possible for \mathcal{A} to continue in some way and satisfy P .

The following theorem relates I/O feasibility and I/O liveness. An I/O feasible TIOA can be characterized by the fact that its set of execution fragments form an I/O liveness property.

Theorem 9.1 *A TIOA is I/O feasible if and only if its set of execution fragments is an I/O liveness property.*

Proof: Fix \mathcal{A} , a TIOA. First, assume that \mathcal{A} is I/O feasible. Let α be a closed execution fragment of \mathcal{A} with $\alpha.lstate = \mathbf{x}$ and let β be an (I, \emptyset) -sequence. I/O feasibility of \mathcal{A} implies that there is some α' from \mathbf{x} such that $\alpha' \upharpoonright (I, \emptyset) = \beta$. Since $\alpha \frown \alpha' \in \text{frags}_{\mathcal{A}}$, we can conclude that the set of execution fragments $\text{frags}_{\mathcal{A}}$ of \mathcal{A} is an I/O liveness property.

For the converse, suppose that the set of execution fragments of \mathcal{A} is an I/O liveness property. Let \mathbf{x} be a state of \mathcal{A} and β be an (I, \emptyset) -sequence. Since the set of execution fragments of \mathcal{A} is an I/O liveness property, there must be some α' such that $\wp(\mathbf{x}) \frown \alpha' \in \text{frags}_{\mathcal{A}}$ and $\alpha' \upharpoonright (I, \emptyset) = \beta$. Clearly, $(\wp(\mathbf{x}) \frown \alpha') \upharpoonright (I, \emptyset) = \beta$, and therefore \mathcal{A} is I/O feasible. ■

9.2 Composition

The following projection and pasting theorem for TIOAs with properties follows from a similar theorem, Theorem 6.24, for TAs with properties.

Theorem 9.2 *Let \mathcal{A}_1 and \mathcal{A}_2 be two compatible TAs and P_1 and P_2 be properties for \mathcal{A}_1 and \mathcal{A}_2 , respectively. Then $\text{traces}_{(\mathcal{A}_1 \parallel \mathcal{A}_2, P_1 \parallel P_2)}$ is exactly the set of (E, \emptyset) -sequences whose restrictions to \mathcal{A}_1 and \mathcal{A}_2 are $\text{traces}_{(\mathcal{A}_1, P_1)}$ and $\text{traces}_{(\mathcal{A}_2, P_2)}$, respectively. That is, $\text{traces}_{(\mathcal{A}_1 \parallel \mathcal{A}_2, P_1 \parallel P_2)} = \{\beta \mid \beta \text{ is an } (E, \emptyset)\text{-sequence and } \beta \upharpoonright (E_i, \emptyset) \in \text{traces}_{(\mathcal{A}_i, P_i)}, i \in \{1, 2\}\}$.*

Theorem 8.7 and its corollary presented in Section 8 assume specification automata whose trace sets are closed under limits, and hence express safety constraints. In this section we present a theorem that can be used in the decomposition of verification where the specification automata may also express liveness properties.

The decomposition of a proof of the assertion $(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_1, Q_1) \leq (\mathcal{A}_2, P_2) \parallel (\mathcal{B}_2, Q_2)$ can be viewed as consisting of two parts. The first part involves the decomposition of the proof that (\mathcal{A}_1, P_1) and (\mathcal{B}_1, Q_1) satisfy their safety properties and the second part involves the decomposition of the proof that (\mathcal{A}_1, P_1) and (\mathcal{B}_1, Q_1) satisfy their liveness properties. Theorem 9.3 uses Corollary 8.8 for the safety part of proofs; the first four hypotheses of Theorem 9.3 imply those of Corollary 8.8. The remaining two hypotheses involve the liveness part of proofs. It requires one to find auxiliary automata with properties, (\mathcal{A}_3, P_3) and (\mathcal{B}_3, Q_3) , such that (\mathcal{A}_1, P_1) implements (\mathcal{A}_3, P_3) in the context of \mathcal{B}_3 without relying on the liveness property of \mathcal{B}_3 , and (\mathcal{B}_1, Q_1) implements (\mathcal{B}_3, Q_3) in the context of \mathcal{A}_3 without relying on the liveness property of \mathcal{A}_3 . Moreover, (\mathcal{A}_1, P_1) must implement (\mathcal{A}_2, P_2) in the context of (\mathcal{B}_3, Q_3) and (\mathcal{B}_1, Q_1) must implement (\mathcal{B}_2, Q_2) in the context of (\mathcal{A}_3, P_3) . That is, the implementation relation between (\mathcal{A}_1, P_1) and (\mathcal{A}_2, P_2) depend on the liveness property Q_3 of the auxiliary context, and the implementation relation between (\mathcal{B}_1, Q_1) and (\mathcal{B}_2, Q_2) depend on the liveness property P_3 of the auxiliary context.

Theorem 9.3 *Suppose $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ are TIOAs such that $\mathcal{A}_1, \mathcal{A}_2$, and \mathcal{A}_3 are comparable, $\mathcal{B}_1, \mathcal{B}_2$, and \mathcal{B}_3 are comparable, and \mathcal{A}_i is compatible with \mathcal{B}_i for $i \in \{1, 2, 3\}$. Suppose that P_i is a property for \mathcal{A}_i and Q_i is a property for \mathcal{B}_i for $i \in \{1, 2, 3\}$. Suppose further that:*

1. *The sets $\text{traces}_{\mathcal{A}_3}$ and $\text{traces}_{\mathcal{B}_3}$ are closed under limits.*
2. *The sets $\text{traces}_{\mathcal{A}_3}$ and $\text{traces}_{\mathcal{B}_3}$ are closed under time-extension.*
3. *$\mathcal{A}_2 \leq \mathcal{A}_3$ and $\mathcal{B}_2 \leq \mathcal{B}_3$.*
4. *$\mathcal{A}_1 \parallel \mathcal{B}_3 \leq \mathcal{A}_2 \parallel \mathcal{B}_3$ and $\mathcal{A}_3 \parallel \mathcal{B}_1 \leq \mathcal{A}_3 \parallel \mathcal{B}_2$.*
5. *$(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_3, \text{frags}_{\mathcal{B}_3}) \leq (\mathcal{A}_3, P_3) \parallel (\mathcal{B}_3, \text{frags}_{\mathcal{B}_3})$ and $(\mathcal{A}_3, \text{frags}_{\mathcal{A}_3}) \parallel (\mathcal{B}_1, Q_1) \leq (\mathcal{A}_3, \text{frags}_{\mathcal{A}_3}) \parallel (\mathcal{B}_3, Q_3)$.*

6. $(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_3, Q_3) \leq (\mathcal{A}_2, P_2) \parallel (\mathcal{B}_3, Q_3)$ and
 $(\mathcal{A}_3, P_3) \parallel (\mathcal{B}_1, Q_1) \leq (\mathcal{A}_3, P_3) \parallel (\mathcal{B}_2, Q_2)$.

Then $(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_1, Q_1) \leq (\mathcal{A}_2, P_2) \parallel (\mathcal{B}_2, Q_2)$.

Proof: Let $\beta \in \text{traces}_{(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_1, Q_1)}$. By definition of composition for automata with properties, $\beta \in \text{traces}_{(\mathcal{A}_1 \parallel \mathcal{B}_1)}$. By Assumptions 1, 2, 3 and 4 and Theorem 8.8, we have $\beta \in \text{traces}_{(\mathcal{A}_2 \parallel \mathcal{B}_2)}$. By projection using Theorem 8.3, $\beta \upharpoonright (E_{\mathcal{A}_2}, \emptyset) \in \text{traces}_{\mathcal{A}_2}$ and $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_2}$. By Assumption 3, $\beta \upharpoonright (E_{\mathcal{A}_2}, \emptyset) \in \text{traces}_{\mathcal{A}_3}$ and $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) \in \text{traces}_{\mathcal{B}_3}$. Since \mathcal{A}_2 and \mathcal{A}_3 are comparable, $\beta \upharpoonright (E_{\mathcal{A}_2}, \emptyset) = \beta \upharpoonright (E_{\mathcal{A}_3}, \emptyset)$ and $\beta \upharpoonright (E_{\mathcal{B}_2}, \emptyset) = \beta \upharpoonright (E_{\mathcal{B}_3}, \emptyset)$. Therefore, $\beta \upharpoonright (E_{\mathcal{A}_3}, \emptyset) \in \text{traces}_{\mathcal{A}_3}$ and $\beta \upharpoonright (E_{\mathcal{B}_3}, \emptyset) \in \text{traces}_{\mathcal{B}_3}$.

By projection using Theorem 9.2, we have $\beta \upharpoonright (E_{\mathcal{A}_1}, \emptyset) \in \text{traces}_{(\mathcal{A}_1, P_1)}$ and $\beta \upharpoonright (E_{\mathcal{B}_1}, \emptyset) \in \text{traces}_{(\mathcal{B}_1, Q_1)}$. By pasting using Theorem 9.2, we have $\beta \in \text{traces}_{(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_3, \text{frags}_{\mathcal{B}_3})}$ and $\beta \in \text{traces}_{(\mathcal{B}_1, Q_1) \parallel (\mathcal{A}_3, \text{frags}_{\mathcal{A}_3})}$. By Assumption 5, we have $\beta \in \text{traces}_{(\mathcal{A}_3, P_3) \parallel (\mathcal{B}_3, \text{frags}_{\mathcal{B}_3})}$ and $\beta \in \text{traces}_{(\mathcal{B}_3, Q_3) \parallel (\mathcal{A}_3, \text{frags}_{\mathcal{A}_3})}$. By projection using Theorem 9.2, we get $\beta \upharpoonright (E_{\mathcal{A}_3}, \emptyset) \in \text{traces}_{(\mathcal{A}_3, P_3)}$ and $\beta \upharpoonright (E_{\mathcal{B}_3}, \emptyset) \in \text{traces}_{(\mathcal{B}_3, Q_3)}$. Since $\beta \upharpoonright (E_{\mathcal{A}_1}, \emptyset) \in \text{traces}_{(\mathcal{A}_1, P_1)}$, by pasting using Theorem 9.2, we have $\beta \in \text{traces}_{(\mathcal{A}_1, P_1) \parallel (\mathcal{B}_3, Q_3)}$, similarly since $\beta \upharpoonright (E_{\mathcal{B}_1}, \emptyset) \in \text{traces}_{(\mathcal{B}_1, Q_1)}$, we have $\beta \in \text{traces}_{(\mathcal{B}_1, Q_1) \parallel (\mathcal{A}_3, P_3)}$. \blacksquare

Example 9.4 (Using environment assumptions to prove liveness) This example illustrates the use of Theorem 9.3 in decomposing the proof of an implementation relationship where the implementation and specification are not merely composition of automata but composition of automata that satisfy some liveness property.

Let $UseOldInputA'$, $UseOldInputB'$, $UseNewInputA'$, and $UseNewInputB'$ be automata which are defined exactly as $UseOldInputA$, $UseOldInputB$, $UseNewInputA$, and $UseNewInputB$ from Example 8.10 except that there is no bound on the number of outputs that the automata can perform. That is, $maxout$ is removed from their sets of state variables. Let P_1, P_2, Q_1 and Q_2 be properties for, respectively, $UseNewInputA'$, $UseOldInputA'$, $UseNewInputB'$ and $UseOldInputB'$ defined as follows:

- P_1 consists of the admissible execution fragments of $UseNewInputA'$.
- Q_1 consists of the admissible execution fragments of $UseNewInputB'$.
- P_2 consists of the execution fragments of $UseOldInputA'$ that contain infinitely many a actions.
- Q_2 consists of the execution fragments of $UseOldInputB'$ that contain infinitely many b actions.

Suppose that we want to prove that:

$$(UseNewInputA', P_1) \parallel (UseNewInputB', Q_1) \leq (UseOldInputA', P_2) \parallel (UseOldInputB', Q_2).$$

The automata $UseNewInputA' \parallel UseNewInputB'$ and $UseOldInputA' \parallel UseOldInputB'$ perform an alternating sequence of a and b actions. The properties express the additional condition that as time goes to infinity the composite automaton $UseNewInputA' \parallel UseNewInputB'$ performs infinitely many a and infinitely many b actions where a and b actions alternate.

As in Example 8.10 automata $AlternateA$ and $AlternateB$ from Figure 16 satisfy the required closure properties for auxiliary automata and capture what is essential about the safety part of the proof, namely that the environments of $UseNewInputA'$ and $UseNewInputB'$ impose alternation. The essential point in the proof of the liveness part is that each automaton responds to each input it receives from its environment. Therefore, we need to pair $AlternateA$ and $AlternateB$ with properties that eliminate non-responding behavior. The properties P_3 and Q_3 defined below satisfy this condition:

- P_3 consists of execution fragments α of $AlternateA$ that satisfy the following condition: if α has finitely many actions then the last action in α is a .
- Q_3 consists of execution fragments α of $AlternateB$ that satisfy the following condition: if α has finitely many actions and contains at least one a then the last action in α is b .

In order to see why the first part of Assumption 5 is satisfied we can inspect the definition of $UseNewInputA$ and observe that $UseNewInputA$ performs an output a one time unit after each input b , when it is composed with $AlternateB$. This implies that in any admissible execution fragment of $UseNewInputA \parallel AlternateB$ with finitely many actions the last action must be a . This is exactly the liveness constraint expressed by P_3 . The second part of Assumption 5 can be seen to hold using a symmetric argument.

In order to see why the first part of Assumption 6 holds consider any execution fragment β of $UseNewInputA \parallel AlternateB$. For β to satisfy P_1 and Q_3 at the same time, it must consist of an infinite sequence in which a and b actions alternate. It is not possible for $UseNewInputA \parallel AlternateB$ to have an admissible execution fragment with finitely many actions because the definition of $UseNewInputA$ requires such a sequence to end in a while this is ruled out by Q_3 , which requires $AlternateB$ to respond to a . The second part of Assumption 6 can be seen to hold using a symmetric argument.

Note that in our explanations we refer to execution fragments rather than traces of execution fragments. This is because our examples do not include any internal actions and our arguments for execution fragments extend to trace fragments in a straightforward way. ■

9.3 Receptiveness for Properties

If we would define a live TIOA to be a pair (\mathcal{A}, L) of a TIOA \mathcal{A} coupled with an I/O liveness property L then the resulting class of systems would not be closed under composition. The

problem, and this was noted already in previous studies of liveness properties for timed I/O automata such as [36], is that this definition allows a system to choose its relative speed with respect to the environment, and to base its decisions on the future behavior of the environment. As a result, the live preorder is not substitutive for parallel composition. To solve these problems, previous studies have introduced notions of *receptive strategies* to guarantee that a system does not constrain its environment. The TIOA framework incorporates a simpler (although less general) notion of strategy than those considered in previous work on timed I/O automata [36].

We begin with a definition of receptiveness for a property. Let \mathcal{A} be a TIOA and let P be a property for \mathcal{A} , that is, a subset of the execution fragments of \mathcal{A} . Then we say that \mathcal{A} is *receptive for P* provided that there exists a strategy \mathcal{A}' for \mathcal{A} such that every execution fragment of \mathcal{A}' is in P . That is, \mathcal{A} has a strategy that can always ensure that P is satisfied (regardless of the behavior of the environment).

The following theorem shows that if \mathcal{A} is receptive for P and P is history-independent, then we can conclude that P is a liveness property for \mathcal{A} . Theorem 9.6 strengthens this result: if we also know that P consists of non-locally-Zeno execution fragments, then P must be an I/O liveness property.

Theorem 9.5 *If a TIOA \mathcal{A} is receptive for P and P is history-independent then P is a liveness property for \mathcal{A} .*

Proof: Suppose that \mathcal{A} is receptive for P . That is, \mathcal{A} has a strategy \mathcal{A}' such that $\text{frags}_{\mathcal{A}'} \subseteq P$. Let α be a closed execution fragment of \mathcal{A} with $\alpha.\text{lstate} = \mathbf{x}$. Since $Q_{\mathcal{A}} = Q_{\mathcal{A}'}$, we know that $\mathbf{x} \in Q_{\mathcal{A}'}$. Now, we need to show that there exists some α' such that $\alpha \cap \alpha' \in P$. Let $\alpha' = \wp(\mathbf{x})$. We know that $\wp(\mathbf{x}) \in \text{frags}_{\mathcal{A}'}$ by axiom **T0**. Since $\text{frags}_{\mathcal{A}'} \subseteq P$, $\alpha' \in P$. Since P is history-independent $\alpha \cap \alpha' \in P$, as needed. ■

Theorem 9.6 *If a TIOA \mathcal{A} is receptive for P and P is a history-independent property for \mathcal{A} consisting of non-locally-Zeno execution fragments, then P is an I/O liveness property for \mathcal{A} .*

Proof: Suppose \mathcal{A} is receptive for P . Then there exists a strategy \mathcal{A}' for \mathcal{A} such that $\text{frags}_{\mathcal{A}'} \subseteq P$. Since all elements of P are non-locally-Zeno, it follows that every element in $\text{frags}_{\mathcal{A}'}$ is non-locally-Zeno, equivalently, \mathcal{A}' is progressive. By Theorem 7.4, we know that any progressive strategy is I/O feasible.

Now, let α be a closed execution fragment of \mathcal{A} with $\alpha.\text{lstate} = \mathbf{x}$ and let β be an (I, \emptyset) -sequence. Since $Q_{\mathcal{A}} = Q_{\mathcal{A}'}$, we have $\mathbf{x} \in Q_{\mathcal{A}'}$, and since \mathcal{A}' is I/O feasible, there exists some execution fragment α' of \mathcal{A}' from \mathbf{x} such that $\alpha' \upharpoonright (I, \emptyset) = \beta$. Since $\alpha' \in P$ and P is history-independent we have that $\alpha \cap \alpha' \in P$. Hence, P is an I/O liveness property for \mathcal{A} . ■

The need for the history-independence assumption for the two theorems above stems from the fact that strategies of our framework are memoryless whereas liveness properties are defined in terms of the possibility of extending every closed execution fragment to a live execution fragment. The history-independence assumption might become unnecessary if we defined strategies to have memory while keeping the liveness property definition as is. Alternatively, we could change the definition of a liveness property to a non-standard one such that a property P for \mathcal{A} is defined to be a liveness property provided that for any state \mathbf{x} of \mathcal{A} , there is some execution fragment α from \mathbf{x} that is in P .

The following is a basic theorem that has nice consequences for composition of automata with liveness properties. Together with Theorems 9.5 and 9.6, it can be used for compositional reasoning about TIOAs with liveness properties.

Theorem 9.7 *Let A_1 and A_2 be two compatible TIOAs. If \mathcal{A}_1 is receptive for P_1 and \mathcal{A}_2 is receptive for P_2 then $\mathcal{A}_1 \parallel \mathcal{A}_2$ is receptive for $P_1 \parallel P_2$.*

Proof: The proof follows from Theorem 8.13 and the definition of composition of properties $P_1 \parallel P_2$ from Section 6. ■

10 Conclusions

In this paper, we have defined a new timed I/O automaton modeling framework for describing and analyzing the behavior of timed systems. This framework is a special case of the recently presented hybrid I/O automaton modeling framework [22]. We used what we have learned in developing the HIOA framework to revise the earlier work on timed I/O automaton models. Our main motivation was to have a timed I/O automaton model that is compatible with the new HIOA model. We sought to benefit from the new style used in describing hybrid behavior in simplifying the prior definitions and results on timed I/O automata. Moreover, we extended the work on the HIOA model by investigating safety and liveness properties and receptiveness for general liveness, not only for feasibility as in the HIOA framework. The results presented in this paper suggest that we are not that far from having a unified framework for timed and hybrid systems in which we can collect and summarize previous results of our own work. We have also established formal relationships with other models that are comparable to ours, showing that the TIOA framework is general enough to express previous results from other frameworks, such as [29, 28, 6, 27, 25, 36].

Designers of real-time systems or timing-based algorithms can use the TIOA framework to describe complex systems and to decompose them into manageable pieces. In particular, they can use the TIOA framework to describe their systems at multiple levels of abstraction, to establish implementation relationships between these levels and to decompose their systems into more primitive, interacting components.

The TIOA framework supports precise statement and verification of safety, liveness, and performance properties of timing-dependent systems. Since the TIOA framework is purely mathematical, proofs are generally done by hand at present. However, the TIOA framework provides a natural basis for computer support tools, which will be developed in the future as an extension to the IOA toolkit [13]. These tools include a syntax and static semantics checker for TIOA specifications, a simulator and partially automated proof tools that employ dynamic invariant detection techniques. There is also work in progress toward a tool to automatically translate TIOA specifications into the input language of UPPAAL [32, 21], which is discussed in more detail in Section 1.2. This would allow us to benefit from fully automated methods in verifying TIOAs that are expressible in UPPAAL.

A Notational Conventions

a, b	action
f, g, h	function
i, j	index
l	locally controlled action
t	time point
v, x	variable
A	set of actions
C	task
E	set of external actions
F	set of functions
H	set of internal (hidden) actions
I	set of input actions
J	interval
K	set of time points
L	set of locally controlled actions
O	set of output actions
P	set of elements in cpo
Q	set of automaton states
R	(simulation) relation
S	set
T	set of trajectories
V	set of variables
X	set of internal variables
\mathbf{x}	state
\mathbf{v}	valuation
$\mathcal{A}, \mathcal{B}, \mathcal{C}$	timed (I/O) automaton
\mathcal{D}	set of discrete transitions
\mathcal{T}	set of trajectories
\mathbb{N}	the natural numbers
\mathbb{R}	the real numbers
\mathbb{T}	the time axis
\mathbb{Z}	the integers
\mathbb{V}	the universe of variables
α, β, δ	(A, V) -sequence
γ	sequence
λ	the empty sequence
π	projection function
σ, ρ	sequence
τ, ν	trajectory
Θ	set of start states

References

- [1] M. Abadi and L. Lamport. Composing specifications. *ACM Transactions on Programming Languages and Systems*, 1(15):73–132, 1993.
- [2] M. Abadi and L. Lamport. Conjoining specifications. *ACM Transactions on Programming Languages and Systems*, 17(3):507–534, 1995.
- [3] B. Alpern and F.B. Schneider. Defining liveness. *Information Processing Letters*, 21:181–185, 1985.
- [4] R. Alur. Timed automata. In *Proc. of 11th International Conference on Computer-Aided Verification (CAV)*, volume 1633 of *LNCS*, pages 8–22. Springer-Verlag, 1999. An earlier and longer version appears in NATO-ASI Summer School on Verification of Digital and Hybrid Systems.
- [5] R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P. H. Ho, X. Nicolin, A. Olivero, J. Sifakis, and Yovine S. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138:3–34, 1995.
- [6] R. Alur and D.L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.
- [7] R. Alur and T. A. Henzinger. Local liveness for compositional modeling of fair reactive systems. In P. Wolper, editor, *Proceedings of CAV 95: Computer Aided Verification*, volume 939 of *Lecture Notes in Computer Science*, pages 166–179. Springer-verlag, 1995.
- [8] K.R. Apt, N. Francez, and S. Katz. Appraising fairness in languages for distributed programming. *Distributed Computing*, 2:226–241, 1988.
- [9] Paul C. Attie. Liveness-preserving simulation relations. In *Proc. of Principles of Distributed Computing*, 1999.
- [10] F. Dederichs and R. Weber. Safety and liveness from a methodological point of view. *Information Processing Letters*, 36(1):25–30, 1990.
- [11] Roberto DePrisco, Butler Lampson, and Nancy Lynch. Revisiting the Paxos algorithm. In Marios Mavronicolas and Philippas Tsigas, editors, *Distributed Algorithms 11th International Workshop, WDAG'97, Saarbrücken, Germany, September 1997 Proceedings*, volume 1320 of *Lecture Notes in Computer Science*, pages 111–125, Berlin-Heidelberg, 1997. Springer-Verlag.
- [12] D. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-Independent Circuits*. ACM Distinguished Dissertations. MIT Press, 1988.

- [13] S. Garland, N. Lynch, and M. Vaziri. *IOA: A Language for Specifying, Programming, and Validating Distributed Systems*. MIT Laboratory for Computer Science, Cambridge, MA, 2001. URL <http://theory.lcs.mit.edu/tds/ioa.html>.
- [14] R. Gawlick, R. Segala, J.F. Sogaard-Andersen, and N.A. Lynch. Liveness in timed and untimed systems. In S. Abiteboul and E. Shamir, editors, *Proceedings 21th ICALP*, Jerusalem, volume 820 of *Lecture Notes in Computer Science*. Springer-Verlag, 1994. A full version appears as MIT Technical Report number MIT/LCS/TR-587.
- [15] C.A. Gunter. *Semantics of Programming Languages: Structures and Techniques*. MIT Press, Cambridge, Massachusetts, 1992.
- [16] M. Hennessy. *Algebraic Theory of Processes*. MIT Press, Cambridge, Massachusetts, 1988.
- [17] T. A. Henzinger, P.-H. Ho, and H. Wong-Toi. HYTECH: A Model Checker for Hybrid Systems. In O. Grumberg, editor, *Proceedings of the 9th International Conference on Computer Aided Verification*, volume 1254 of *Lecture Notes in Computer Science*, pages 460–463. Springer-Verlag, 1997.
- [18] T. A. Henzinger, S. Qadeer, and S. K. Rajamani. Decomposing refinement proofs using assume-guarantee reasoning. In *Proceedings of the International Conference on Computer-Aided Design (ICCAD)*, pages 245–252. IEEE Computer Society Press, 2000.
- [19] C. B. Jones. Specification and design of parallel programs. In R. E. A. Mason, editor, *Information Processing 83: Proceedings of the IFIP 9th World Congress*, pages 321–332. North-Holland, 1983.
- [20] D. E. Knuth. *Fundamental Algorithms*, volume 1 of *The Art of Computer Programming*. Addison-Wesley, Reading, Massachusetts, second edition, 1973.
- [21] Kim G. Larsen, Paul Pettersson, and Wang Yi. Uppaal in a nutshell. *Journal of Software Tools for Technology Transfer*, 1–2:134–152, 1997.
- [22] N.A. Lynch, R. Segala, and F.W. Vaandrager. Hybrid I/O automata. *Information and Computation*, 185(1):105–157, 2003. Also Technical Report MIT-LCS-TR-827d, MIT Laboratory for Computer Science.
- [23] N.A. Lynch, R. Segala, F.W. Vaandrager, and H.B. Weinberg. Hybrid I/O automata. In R. Alur, T.A. Henzinger, and E.D. Sontag, editors, *Hybrid Systems III*, volume 1066 of *Lecture Notes in Computer Science*, pages 496–510. Springer-Verlag, 1996.
- [24] N.A. Lynch, R. Segala, F.W. Vaandrager, and H.B. Weinberg. Hybrid I/O automata. Report CSI-R9907, Computing Science Institute, University of Nijmegen, April 1999.

- [25] N.A. Lynch and F.W. Vaandrager. Action transducers and timed automata. *Formal Aspects of Computing*, 8(5):499–538, 1996.
- [26] Nancy Lynch and Alex Shvartsman. RAMBO: A reconfigurable atomic memory service for dynamic networks. In D. Malkhi, editor, *Distributed Computing (Proceedings of the 16th International Symposium on Distributed Computing (DISC), Toulouse, France, October 2002)*, volume 2508 of *Lecture Notes in Computer Science*, pages 173–190. Springer-Verlag, 2002. Also, Technical Report MIT-LCS-TR-856.
- [27] Nancy Lynch and Frits Vaandrager. Forward and backward simulations — Part II: Timing-based systems. *Information and Computation*, 128(1):1–25, July 1996.
- [28] O. Maler, Z. Manna, and A. Pnueli. From timed to hybrid systems. In J.W. de Bakker, C. Huizing, W.P. de Roever, and G. Rozenberg, editors, *Proceedings REX Workshop on Real-Time: Theory in Practice*, Mook, The Netherlands, June 1991, volume 600 of *Lecture Notes in Computer Science*, pages 447–484. Springer-Verlag, 1992.
- [29] M. Merritt, F. Modugno, and M. Tuttle. Time constrained automata. In J.C.M. Baeten and J.F. Groote, editors, *Proceedings CONCUR 91*, Amsterdam, volume 527 of *Lecture Notes in Computer Science*, pages 408–423. Springer-Verlag, 1991.
- [30] R. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer-Verlag, 1980.
- [31] S. Mitra, Y. Wang, N. Lynch, and E. Feron. Safety verification of pitch controller for model helicopter. In O. Maler and A. Pnueli, editors, *Proc. of Hybrid Systems: Computation and Control*, volume 2623 of *Lecture Notes in Computer Science*, pages 343–358, Prague, the Czech Republic April 3-5, 2003.
- [32] Paul Petterson. *Modelling and Verification of Real-Time Systems Using Timed Automata: Theory and Practice*. PhD thesis, Department of Computer Systems, Uppsala University, 1999. Technical Report DoCs 99/101.
- [33] A. Pnueli. In transition from global to modular temporal reasoning about programs. In K. R. Apt, editor, *Logis and Models of Concurrent Systems*, NATO ASI, pages 123–144. Springer-Verlag, 1984.
- [34] A. Pnueli. Development of hybrid systems. In H. Langmaack, W.-P. de Roever, and J. Vytupil, editors, *Proceedings of the Third International School and Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'94)*, Lübeck, Germany, September 1994, volume 863 of *Lecture Notes in Computer Science*, pages 77–85. Springer-Verlag, 1994.
- [35] J.W. Polderman and J.C. Willems. *Introduction to Mathematical Systems Theory: A Behavioural Approach*, volume 26 of *Texts in Applied Mathematics*. Springer-Verlag, 1998.

- [36] R. Segala, R. Gawlick, J.F. Sogaard-Andersen, and N.A. Lynch. Liveness in timed and untimed systems. *Information and Computation*, 141(2):119–171, March 1998.
- [37] E.D. Sontag. *Mathematical Control Theory — Deterministic Finite Dimensional Systems*, volume 6 of *Texts in Applied Mathematics*. Springer-Verlag, 1990.
- [38] E. W. Stark. A proof technique for relt/guarantee properties. In S. N. Maheshwari, editor, *Foundations of Software Technology and Theoretical Computer Science*, volume 206 of *LNCS*, pages 369–391. Springer-Verlag, 1985.
- [39] S. Tasiran, R. Alur, R.P. Kurshan, and R.K. Brayton. Verifying abstractions of timed systems. In *Proceedings of the Seventh Conference on Concurrency Theory (CONCUR 1996)*, volume 1119 of *LNCS*.

Index

- (A, V)-restriction, 19
- (A, V)-sequence, 17

- abstraction, 5
- admissible, 17, 19, 88
- Alur-Dill timed automaton, 8, 37, 38, 77
- analog, 14
- analog variable, 22
- assume-guarantee, 105

- backward simulation, *see* simulation relation, 48

- chain, 11
- clock synchronization, 27, 44
- ClockSync*, 29, 58, 98
- comparable
 - TA, 39
 - TIOA, 102
- compatible
 - TA, 55
 - TIOA, 103
- complete partial order, 11
- composition, 5, 55, 103
- congruence, 73
- cpo, *see* complete partial order

- discrete
 - variable, 14
- discrete action, 20
- discrete transition, 20
- discrete variable, 14, 22
- dynamic type, 13

- effect, 22
- enabled, 20
- execution, 30, 98
 - PeriodicSend*, 32
 - Timeout*, 32
- execution fragment, 30, 31, 98

- fair forward simulation, *see* simulation relation
- fairness property, *see* property
- feasible, 37
- FIN, *see* finite internal nondeterminism, 107
- finite internal nondeterminism, 35
- Fischer's mutual exclusion, 26, 33, 71
- FischerME*, 27
- FischerME2*, 71
- forward simulation, *see* simulation relation
 - clock synchronization, 44
 - time-bounded channels, 43

- hiding, 62
- HIOA, 6, 104
- history relation, 50, 51, 103
 - time-bounded channels, 53
- history variable, 50, 51
 - time-bounded channels, 50
- history-independent property, *see* property
- hybrid automaton, 21, 55
- Hybrid I/O Automaton modeling framework, 6, 122
- hybrid sequence, 16
 - admissible, 17
 - closed, 17
 - concatenation, 18
 - limit time, 17
 - prefix, 18
 - time-bounded, 17
 - Zeno, 17
- HyTech, 8

- I/O feasible, 99, 114
- I/O liveness property, *see* property
- implementation, 5, 39
- invariant, 31
 - clock agreement, 60
 - clock validity, 59, 60
 - ClockSync*, 59, 60

- failure and timeout, 58
 - FischerME*, 33, 34
 - TimedChannel*, 33
 - timeout, 57
- isomorphism, 46
- limit, 11
- linear hybrid automaton, 8
- liveness property, *see* property
- locally Zeno, 98
- machine-closed, 83–85
- machine-closure, 6
- non-Zeno, 17, 19
- parallel composition, *see* composition
- partial order, 11
 - complete partial order, 11
- periodic sending process, 24, 32
- periodic sending process with failures, 24
- PeriodicSend*, 24, 56
- PeriodicSend2*, 57
- PeriodicSend2*, 25
- point trajectory, *see* trajectory
- precondition, 22
- progressive, 99, 102
- property, 81, 116
 - fairness, 6, 86
 - history-independent, 88
 - I/O liveness, 116
 - liveness, 6, 82, 89, 116, 118
 - safety, 6, 81, 89
- prophecy relation, 53, 103
- prophecy variable, 53
- reachable, 31
- receptive, 102, 115
- receptiveness, 6, 100
- receptiveness for a property, 120
- refinement, 46
- safety property, *see* property
- sequence, 10
- simulation relation, 5, 41
 - backward simulation, 41, 47, 103
 - forward simulation, 41, 102
 - refinement, 46
- static type, 13
- stopping condition, 23
- strategy, 100, 101
- strongly fair, 87
- substitutivity, 61, 62, 104, 105
- suffix, 31
- TA, *see* timed automaton
- TA with bounds, 65
- task, 65
 - lower bound, 66
 - upper bound, 66
- time axis, 13
- time interval, 13
- time-bounded channel, 23, 33, 43, 50, 53
- timed automaton, 20
- timed automaton model, 20
- Timed I/O automaton, 5, 97
- Timed Input/Output Automaton modeling
 - framework, 5
- TimedChannel*, 23, 56, 57, 98
- Timeout*, 26, 56, 57
- timeout process, 25, 32
- timing-independent, 37, 107
- TIOA, *see* Timed I/O automaton
- trace, 5, 31, 98
 - PeriodicSend*, 32
 - Timeout*, 33
- trace fragment, 31, 98
- trajectory, 14, 20
 - concatenation, 15
 - limit time, 15
 - point trajectory, 14, 17
 - prefix, 15
- untimed automaton, 12
- untiming, 71, 73
- UPPAAL, 8, 123
- variables, 13, 14, 20

analog, 14
discrete, 14
dynamic types, *see* static type
static type, *see* static type

weakly fair, 87

Zeno, 6, 17, 34, 88