Centre Fédéré en Vérification

Technical Report number 2005.46

Systematic Implementation of Real-Time Models

Martin De Wulf, Laurent Doyen, Jean-François Raskin



This work was partially supported by a FRFC grant: 2.4530.02

http://www.ulb.ac.be/di/ssd/cfv

Systematic Implementation of Real-Time Models *[†]

Martin De Wulf

Laurent Doyen[‡]

Jean-François Raskin

Département d'Informatique Faculté des Sciences Université Libre de Bruxelles

Abstract

Recently we have proposed the "almost ASAP" semantics as an alternative semantics for timed automata. This semantics is useful when modeling realtime controllers : control strategies modeled with this semantics are robust and implementable (without making the synchrony hypothesis). We show in this paper how to effectively encode this semantics using timed automata along with their classical semantics. We have implemented a tool set that allows us to verify, using HYTECH and UPPAAL, the almost ASAP behavior of controllers and generate automatically provably correct code from verified models. To illustrate the applicability of our results, we show how we have synthesized the code for the Philips Audio Control Protocol for LEGO MINDSTORMSTM.

1 Introduction

Timed automata are an important formal model for the specification and analysis of real-time systems. Formalisms like timed automata and hybrid automata are central in the so-called *model-based development* methodology for embedded controllers. The steps underlying that methodology can be summarized as follows: (*i*) construct a (timed/hybrid automaton) model Env of the environment in which the controller will be embedded; (*ii*) make clear what is the control objective: for example, prevent the environment to enter a set of Bad states; (*iii*) design a (timed automaton) model Cont of the control strategy; (*iv*) verify that Reach([[Env || Cont]]) \cap Bad = \emptyset . When Cont has been proven correct, it would be valuable to ensure that an implementation Impl of that model can be obtained in a systematic way in order to ensure the conservation of correctness, that is to ensure that Reach([[Env || Impl]]) \cap Bad = \emptyset is obtained by construction.

Unfortunately, this is often not possible for several *fundamental* and/or *technical* reasons. First, the notion of time used in the traditional semantics of timed automata is *continuous* and defines *perfect clocks* with *infinite precision* while implementations can only access time through *digital* and *finitely precise* clocks. Second,

^{*}This is the extended version of the paper with the same title to appear in Proc. of Formal Methods 2005 (FM'05), Newcastle upon Tyne, UK, Jul. 2004, LNCS, Springer, 2004

 $^{^\}dagger Supported$ by the FRFC project "Centre Fédéré en Vérification" funded by the Belgian National Science Foundation (FNRS) under grant nr2.4530.02

[‡]Research fellow supported by the Belgian National Science Foundation (FNRS)

timed automata react *instantaneously* to events and time-outs while implementations can only react within a given, usually small but not zero, *reaction delay*. Third, timed automata may describe control strategies that are *unrealistic*, like *zeno-strategies* or strategies that ask the controller to act faster and faster [CHR02]. For one of those three reasons, a model for a digital controller that has been proven correct may not be implementable (at all) or it may not be possible to turn it systematically into an implementation that is proven correct w.r.t. this model.

To overcome those problems, we recently proposed an alternative semantics to timed automata in [DDR04]. This semantics is called the Almost ASAP semantics, AASAP for short. The AASAP semantics of a timed automaton A, noted $[\![A]\!]_{\Delta}^{AAsap}$, is a parametric semantics that leaves as a parameter Δ , which takes value $\delta \in \mathbb{Q}^{\geq 0}$, the *reaction delay* of the controller. This semantics relaxes the classical semantics of timed automata in that it does not impose on the controller to react instantaneously but imposes on the controller to react within δ time units. We have proven that a timed controller is implementable with a sufficiently fast hardware if there exists $\delta \in \mathbb{Q}^{>0}$ such that Reach($[[Env]] \parallel [[Cont]]_{\delta}^{AAsap} \cap Bad = \emptyset$.

To use the AASAP semantics in practice, we need tool support. In [DDR04], we have shown that the AASAP semantics of a controller can be encoded using a single parameter timed automaton. Unfortunately, this construction is exponential in all cases, which makes it useless for all but the toy examples. In this paper, we define a new compositional construction that avoids the exponential blow-up. The exponential behavior can still appear during the verification phase but only in the worst case. Thanks to this new construction, we have implemented a tool set in order to manipulate the AASAP semantics on top of HyTECH [HHWT95] and UPPAAL [PL00]. We show the practical interest of our construction by applying our tool set to a non-trivial example: the Philips Audio Control Protocol [BPV94]. We show how the AASAP semantics can be used to produce provably correct executable code for this protocol. The code that we have produced automatically can be run on LEGO MINDSTORMSTM. With this case study, we believe that we have shown that the AASAP semantics is useful when supported by computer aided verification tools and that it can be used to produce correct code for non-trivial embedded controllers without making the synchrony hypothesis. To the best of our knowledge, this is the first time that provably correct (without making the synchrony hypothesis) real-time code is produced for a non-trivial case study.

The rest of the paper is organized as follows. In Section 2, we recall some preliminary notions. In Section 3, we review the syntax and classical semantics of timed automata. In Section 4, we recall the AASAP semantics and summarize its properties. In Section 5, we present our compositional construction. In Section 6, we present our tool set. In Section 7, we show how to apply the AASAP semantics to synthesize provably correct code for a real-time protocol.

2 Preliminaries

Definition 1 [STTS] A structured timed transition system \mathcal{T} is a tuple $\langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau}, \rightarrow \rangle$, where S is a (possibly infinite) set of states, $\iota \in S$ is the initial state, the set of labels is structured in three disjoint components: Σ_{in} is the finite set of incoming labels, Σ_{out} is the finite set of outgoing labels, Σ_{τ} is the finite set of internal labels, and $\rightarrow \subseteq S \times \Sigma_{in} \cup \Sigma_{out} \cup \Sigma_{\tau} \cup \mathbb{R}^{\geq 0} \times S$ is the transition relation.

A state $s \in S$ of a STTS $\mathcal{T} = \langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau}, \rightarrow \rangle$ is *reachable* if there exists a finite sequence $s_0 s_1 \dots s_n$ of states such that $s_0 = \iota, s_n = s$ and for any $i, 0 \leq i < n$,

there exists $\sigma \in \Sigma_{in} \cup \Sigma_{out} \cup \Sigma_{\tau} \cup \mathbb{R}^{\geq 0}$ such that $(s_i, \sigma, s_{i+1}) \in \rightarrow$. The set of reachable states of \mathcal{T} is noted Reach (\mathcal{T}) .

Some more notions. We only present intuitively other notions that are useful in the sequel. The formal definitions can be found in appendix. We use a natural definition of the *composition* $\mathcal{T}^1 || \mathcal{T}^2$ of two STTS \mathcal{T}^1 and \mathcal{T}^2 with synchronizations similar to the ones in the input-output automata framework [LT87]: a common label must be an output label (sending) in one of the STTS and an input label (receiving) in the other. The composition $\mathcal{T}^1 || \mathcal{T}^2$ is a STTS where synchronized labels are considered as internal.

Such syncronizations is a blocking communication mechanism. This may be problematic as on one hand we want to verify that the controller does not control the environment by refusing to synchronize on its output, and on the other hand, we do not want our controller to issue outputs that can not be accepted by the environment. To avoid such problems we impose *input enabledness* of the STTS that we compose, which means that input labels have the property of being enabled in every state. In this point, the presentation differs from [DDR04].

Finally, given two *input enabled* STTS \mathcal{T}^1 (the controller) with state space S^1 , \mathcal{T}^2 (the environment) with state space S^2 and a set $B \subseteq S^2$ of bad states, we say that \mathcal{T}^1 controls \mathcal{T}^2 to avoid B if Reach $(\mathcal{T}^1 || \mathcal{T}^2) \cap S^1 \times B$ is empty.

3 Timed Automata and Urgency

Let X be a finite set of real-valued variables. A valuation for X is a function v: $X \to \mathbb{R}$. We write $[X \to \mathbb{R}]$ for the set of all valuations for X. Let Δ be a *parameter*. Define the set of *terms* to be $\mathsf{T} = \mathbb{Q} \cup \{+\infty\}$, and the set of *parametric terms* to be $\mathsf{PT} = \mathsf{T} \cup \{c + \Delta, c - \Delta \mid c \in \mathbb{Q}\}$. A *(parametric) rectangular constraint* over X is a formula of the form $\varphi \equiv a \sim_1 x \sim_2 b$ where $x \in X, \sim_1, \sim_2 \in \{<, \le\}$ and a, b are (parametric) terms. Let $lb(\varphi) = a$ and $rb(\varphi) = b$ denote the left (resp. right) bound of φ . A *(parametric) rectangular predicate* is a finite set of (parametric) rectangular constraints interpreted as a conjunction. A (parametric) multirectangular predicate is a finite set of (parametric) rectangular predicates interpreted as a disjunction. Given $\delta \in \mathbb{Q}$ and a parametric term a, let $\llbracket a \rrbracket_{\delta} = a$ if $a \in \mathsf{T}$ and $\llbracket a \rrbracket_{\delta} = c + \delta$ (resp. $c - \delta$) if $a = c + \Delta$ (resp. $c - \Delta$). For a parametric rectangular predicate p, a valuation v and a rational $\delta \in \mathbb{Q}$, we write $v \models_{\delta} p$ iff $[a]_{\delta} \sim_1 v(x) \sim_2 [b]_{\delta}$ for all " $a \sim_1 x \sim_2 b$ " in p. For a parametric multirectangular predicate q, we write $v \models_{\delta} q$ iff there exists $p \in q$ such that $v \models_{\delta} p$. For a parametric (multi)rectangular predicate p, let $[\![p]\!]_{\delta}$ denote the set $\{v \mid v \models_{\delta} p\}$. We sometimes write $v \models p$ instead of $v \models_0 p$.

We say that a rectangular predicate over X is in normal form if it contains at most one rectangular constraint for each variable $x \in X$, with the convention that the empty predicate p (such that $\llbracket p \rrbracket = \emptyset$) is represented by $\{x \in [+\infty, +\infty] \mid x \in X\}$; any rectangular predicate can be put in that normal form. Let g be a rectangular predicate in normal form, then g(x) denotes the rectangular constraint " $a \sim_1 x \sim_2 b$ " if it is the constraint over x in g and true if there is no constraint over x in g. We defined predicates as sets because it is useful in the sequel for manipulating the predicates that appear in timed automata. However, some operations are easier to represent with classical boolean operations (\wedge, \vee, \neg) . It is easy to extend the definition of such operators to our set-predicates. For example, for two multirectangular predicates q and r, the multirectangular predicate $q \wedge r$ is the set $\bigcup_{p_1 \in q, p_2 \in r} \{ p_1 \cup p_2 \}.$

We note Rect(X) the set of rectangular predicates built using variables in X, PRect(X) the set of parametric rectangular predicates, MultPRect(X) the set of parametric multirectangular predicates and $\text{Rect}_{c}(X)$ the set of rectangular predicates containing only closed rectangular constraints ($\sim_1, \sim_2 \in \{\leq\}$).

Let $v: E_1 \to \mathbb{R}$ be a valuation, let $E_2 \subseteq E_1$, and $c \in \mathbb{R}$, then $v[E_2 := c]$ denotes the valuation v' such that v'(e) = c if $e \in E_2$ and v'(e) = v(e) if $e \notin E_2$. In the sequel, we sometimes write v[e := c] instead of $v[\{e\} := c]$. Let $v: X \to \mathbb{R}$ be a valuation, for any $t \in \mathbb{R}^{\geq 0}$, v - t is a valuation such that for any $x \in X$, (v-t)(x) = v(x) - t. We define v + t in a similar way. We extend this definition to valuations v in $[X \to \mathbb{R}^{\geq 0} \cup \{\bot\}]$ as follows: (v + t)(x) = v(x) + t, if $v(x) \in \mathbb{R}^{\geq 0}$, and $(v + t)(x) = \bot$ otherwise. We are now equipped to define our flavor of timed automata [AD94] (with one parameter and a urgency flag Asap) and their classical semantics.

Definition 2 [Single parametric timed automata] A single parametric timed automaton¹ is a tuple $\langle \text{Loc}, l_0, \text{Var}, \text{Inv}, \text{Lab}, \text{Edg}, \text{Asap} \rangle$ where (i) Loc is a finite set of locations representing the discrete states of the automaton. (ii) $l_0 \in \text{Loc}$ is the initial location. (iii) $\text{Var} = \{x_1, \ldots, x_n\}$ is a finite set of real-valued clocks whose values continuously increase as time passes with first derivative equal to one. (iv) $\text{Inv} : \text{Loc} \rightarrow \text{MultPRect}(\text{Var})$ is the invariant condition. The automaton can stay in location l as long as the tuple of values of the variables x_1, \ldots, x_n lies in Inv(l). To ensure the existence of an initial state, we require that the valuation v_0 such that $v_0(x) = 0$ for every $x \in \text{Var}$ lies in $\text{Inv}(l_0)$. (v) $\text{Lab} = \text{Lab}_{in} \cup \text{Lab}_{out} \cup \text{Lab}_{\tau}$ is a structured finite alphabet of labels, partitioned into input labels Lab_{in} , output labels Lab_{out} , and internal labels Lab_{τ} . (vi) $\text{Edg} \subseteq \text{Loc} \times \text{Loc} \times \text{PRect}(\text{Var}) \times \text{Lab} \times 2^{\text{Var}}$ is a set of edges. An edge (l, l', g, σ, R) represents a discrete transition from location l to location l' with guard g, event σ and a subset $R \subseteq \text{Var}$ of the variables to be reset. The guard g is a rectangular predicate. (vii) $\text{Asap} : \text{Edg} \rightarrow \{\top, \bot\}$ is a special flag used to model urgency.

Definition 3 [Semantics of single parametric timed automata] Let $A = \langle \mathsf{Loc}, l_0, \mathsf{Var}, \mathsf{Inv}, \mathsf{Lab}, \mathsf{Edg}, \mathsf{Asap} \rangle$ be a timed automaton and $\delta \in \mathbb{Q}^{\geq 0}$. The semantics of A is the STTS $\llbracket A \rrbracket_{\delta} = (S, \iota, \Sigma_{\mathsf{in}}, \Sigma_{\mathsf{out}}, \Sigma_{\tau}, \rightarrow)$ where: (i) $S = \{(l, v) \mid l \in \mathsf{Loc} \land v \in \llbracket \mathsf{Inv}(l) \rrbracket_{\delta}\}$. (ii) $\iota = (l_0, v_0)$ such that for any $x \in \mathsf{Var} : v_0(x) = 0$. (iii) $\Sigma_{\mathsf{in}} = \mathsf{Lab}_{\mathsf{in}}, \Sigma_{\mathsf{out}} = \mathsf{Lab}_{\mathsf{out}}, \text{ and } \Sigma_{\tau} = \mathsf{Lab}_{\tau}$. (iv) the transition relation \rightarrow is defined as follows:

- (a) For the discrete transitions, $((l, v), \sigma, (l', v')) \in \to$ iff there exists an edge $(l, l', g, \sigma, R) \in \mathsf{Edg}$ such that $v \models_{\delta} g, v' = v[R := 0]$.
- (b) For the continuous transitions $((l, v), t, (l', v')) \in \rightarrow$ iff first l = l', second for any edge $e = (l_1, l_2, g, \sigma, R) \in \mathsf{Edg}$: if $l_1 = l$ then $\mathsf{Asap}(e) = \bot$ and third $\forall x \in \mathsf{Var}, t \in \mathbb{R}^{\geq 0} : v'(x) = v(x) + t$ and $\forall t' \in [0, t] : v + t' \in [\mathsf{Inv}(l)]_{\delta}$

For simplicity, we often say timed automaton instead of single parametric timed automaton. We use the classical definition of the synchronized product $A_1 \times A_2$ of two timed automata. For the urgency flag, an edge in the product is flagged Asap if one of the corresponding edges in A_1 or A_2 is flagged. This is the semantics used in the HYTECH tool for the Asap flag [HHWT95]. Notice that (in the final product only), the Asap flag can be replaced by a clock which is reset on every transition, and forced by an invariant to stay nil in every location with an outgoing

¹In this paper, single parametric timed automata always use the parameter Δ .

Asap edge, showing that Asap is a feature that does not add expressive power to timed automata, but just allows us to design timed automata in a modular way.

4 ELASTIC Controllers and AASAP Semantics

Controllers are specified using a subclass of timed automata, called ELASTIC², without invariants and with only closed guards. In general, invariants are used to express urgency but in ELASTIC urgency is implicit : a controller shall make an action (almost) as soon as it becomes possible. Formally, this almost urgency is defined in the AASAP semantics of the controller by allowing some delay (bounded by a parameter Δ) before forcing an enabled transition.

Definition 4 [ELASTIC Controllers] An ELASTIC controller A is a tuple $\langle \text{Loc}, l_0, \text{Var}, \text{Lab}, \text{Edg} \rangle$ where Loc is a finite set of locations, $l_0 \in \text{Loc}$ is the initial location, $\text{Var} = \{x_1, \ldots, x_n\}$ is a finite set of clocks, Lab is a finite structured alphabet of labels, partitioned into input labels Lab_{in}, output labels Lab_{out}, and internal labels Lab_{τ}, Edg is a set of edges of the form (l, l', g, σ, R) where $l, l' \in \text{Loc}$ are locations, $\sigma \in \text{Lab}$ is a label, $g \in \text{Rect}_c(\text{Var})$ is a guard and $R \subseteq \text{Var}$ is a set of clocks to be reset.

Notations. We define the function TrueSince: $[Var \to \mathbb{R}^{\geq 0}] \times \text{Rect}_{c}(Var) \to \mathbb{R}^{\geq 0} \cup \{-\infty\}$, noted TS, as follows: either $v \models g$ and $\mathsf{TS}(v,g) = t$ where t is s.t. $v - t \models g \land \forall t' > t : v - t' \not\models g$, or $v \not\models g$ and $\mathsf{TS}(v,g) = -\infty$.

Let $p \equiv a \sim_1 x \sim_2 b$ be a rectangular constraint. Given $\Delta_1, \Delta_2 \in \mathsf{PT}$, the symbol \langle standing either for [or (and the symbol \rangle standing either for] or), we define the notation $\Delta_1 \langle p \rangle_{\Delta_2}$ for the parametric rectangular constraint:

$$a - \Delta_1 \sim_1' x \sim_2' b + \Delta_2$$

where \sim'_1 stands either for \leq if \langle is [, or for \langle if \langle is (, and \sim'_2 is interpreted symetrically. For example, let $p \equiv 2 \leq x \leq 5$, then $_{-\frac{1}{3}}(p]_{\Delta} \equiv 2 + \frac{1}{3} < x \leq 5 + \Delta$. The notation is naturally extended to rectangular predicates.

With those two additional notations we are now ready to define the AASAP semantics [DDR04].

Definition 5 [AASAP semantics] Given an ELASTIC controller $A = \langle Loc, l_0, Var, Lab_{in}, Lab_{out}, Lab_{\tau}, Edg \rangle$ and $\delta \in \mathbb{Q}^{\geq 0}$, the AASAP semantics of A is the STTS $[A]_{\delta}^{AAsap} = \langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau}, \rightarrow \rangle$ where:

- (A1) S is the set of tuples (l, v, I, d) where $l \in \mathsf{Loc}, v \in [\mathsf{Var} \to \mathbb{R}^{\geq 0}], I \in [\Sigma_{\mathsf{in}} \to \mathbb{R}^{\geq 0} \cup \{\bot\}]$ and $d \in \mathbb{R}^{\geq 0}$;
- (A2) $\iota = (l_0, v, I, 0)$ where v is such that for any $x \in Var : v(x) = 0$, and I is such that for any $\sigma \in \Sigma_{in}$, $I(\sigma) = \bot$;
- (A3) $\Sigma_{in} = \mathsf{Lab}_{in}, \Sigma_{out} = \mathsf{Lab}_{out}, \text{ and } \Sigma_{\tau} = \mathsf{Lab}_{\tau} \cup \overline{\mathsf{Lab}_{in}} \cup \{\epsilon\};$
- (A4) The transition relation is defined as follows:

- for the discrete transitions, we distinguish five cases:

(A4.1) let $\sigma \in \mathsf{Lab}_{\mathsf{out}}$. We have $((l, v, I, d), \sigma, (l', v', I, 0)) \in \rightarrow$ iff there exists $(l, l', g, \sigma, R) \in \mathsf{Edg}$ such that $v \models_{\delta \Delta}[g]_{\Delta}$ and v' = v[R := 0];

²Event-based LAnguage for Simple TImed Controllers.

- (A4.2) let $\sigma \in \mathsf{Lab}_{\mathsf{in}}$. We have $((l, v, I, d), \sigma, (l, v, I', d)) \in \to$ iff \cdot either $I(\sigma) = \bot$ and $I' = I[\sigma := 0];$ \cdot or $I(\sigma) \neq \bot$ and I' = I.
- (A4.3) let $\bar{\sigma} \in \overline{\mathsf{Lab}_{in}}$. We have $((l, v, I, d), \bar{\sigma}, (l', v', I', 0)) \in \rightarrow$ iff there exists $(l, l', g, \sigma, R) \in \mathsf{Edg}, v \models_{\delta} \Delta[g]_{\Delta}, I(\sigma) \neq \bot, v' = v[R := 0]$ and $I' = I[\sigma := \bot];$
- (A4.4) let $\sigma \in \mathsf{Lab}_{\tau}$. We have $((l, v, I, d), \sigma, (l', v', I, 0)) \in \to$ iff there exists $(l, l', g, \sigma, R) \in \mathsf{Edg}, v \models_{\delta \Delta} [g]_{\Delta}$, and v' = v[R := 0];
- (A4.5) let $\sigma = \epsilon$. We have for any $(l, v, I, d) \in S : ((l, v, I, d), \epsilon, (l, v, I, d)) \in \rightarrow$.
- for the continuous transitions:
- (A4.6) for any $t \in \mathbb{R}^{\geq 0}$, we have $((l, v, I, d), t, (l, v + t, I + t, d + t)) \in \rightarrow$ iff the two following conditions are satisfied:
 - for any edge $(l, l', g, \sigma, R) \in \mathsf{Edg}$ with $\sigma \in \mathsf{Lab}_{\mathsf{out}} \cup \mathsf{Lab}_{\tau}$, we have that: $\forall t' : 0 \leq t' \leq t : (d + t' \leq \delta \lor \mathsf{TS}(v + t', g) \leq \delta)$
 - for any edge $(l, l', g, \sigma, R) \in \mathsf{Edg}$ with $\sigma \in \mathsf{Lab}_{in}$, we have that: $\forall t' : 0 \le t' \le t : (d + t' \le \delta \lor \mathsf{TS}(v + t', g) \le \delta \lor (I + t')(\sigma) \le \delta)$

Comments on the AASAP semantics. Rule (A1) defines the states that are tuples of the form $\langle l, v, I, d \rangle$. The first two components, location l and valuation v, are the same as in the classical semantics; I and d are new. The function I records, for each input event σ , the time elapsed since its oldest "untreated" occurrence. The treatment of an event σ happens when a transition labelled with $\bar{\sigma}$ is fired. Once this oldest occurence is treated, the function returns \perp for σ until a new occurence of σ , forgetting about the σ 's that happened between the oldest occurrence and the treatment. The time elapsed since the last location change in the controller is recorded by d. Rule (A2) and (A3) are straightforward. Rules (A4.1-6) require more explanations. Rule (A4.1) defines when it is allowed for the controller to emit an output event. The only difference with the classical semantics is that we enlarge the guard by the parameter Δ . Rules (A4.2) defines how inputs from the environment are received by the controller. The controller maintains, through the function I, a list of events that have occurred and are not treated yet. An input event σ can be received at any time, but only the age of the oldest untreated σ is stored in the I function. Note that the rule ensures input enabledness of the controller. Rule (A4.3) defines when inputs are treated by the controller. An input σ is treated when a transition with an enlarged guard and labelled with $\bar{\sigma}$ is fired. Once σ has been treated, the value of $I(\sigma)$ goes back to \perp . Rule (A4.4) is similar to (A4.1). Rule (A4.5) expresses that the ϵ event can always be emitted. Rule (A4.6) specifies how much time can elapse. Intuitively, time can pass as long as no transition starting from the current location is *urgent*. A transition labeled with an output or an internal event is urgent in a location l when the control has been in l for more than δ time units $(d + t' \geq \delta)$ and the guard of the transition has been true for more than δ time units $(\mathsf{TS}(v+t',g) \geq \delta)$. A transition labeled with an input event σ is urgent in a location l when the control has been in l for more than δ time units $(d + t' \geq \delta)$, the guard of the transition has been true for more that δ time units $(\mathsf{TS}(v+t', g) \geq \delta)$ and the last untreated occurrence of σ event has been emitted by the environment at least δ time units ago $(I + t'(\sigma) \geq \delta)$ (we define \perp to be smaller than any rational value). This notion of urgency parameterized by Δ is the main difference between the AASAP semantics and the usual ASAP semantics.

Properties We informally recall the main properties of the AASAP semantics which have been established in [DDR04].

First, the AASAP semantics has the desirable property that "faster is better": if a controller with reaction time bounded by δ_1 safely controls an environment, then so does the same controller with a reaction time bounded by any $\delta_2 < \delta_1$.

Second, we can implement a controller that has been proven correct (that is, such that for some $\delta > 0$ its AASAP semantics safely controls the environment). The correctness of the controller is preserved by the implementation provided the hardware is sufficiently fast and has a sufficiently precise digital clock. This has been formally proven by showing that the AASAP semantics can simulate (in the formal sense) a program semantics which defines what is an implementation of an ELASTIC controller. Intuitively, it is a procedure that repeats forever *execution* rounds defined as follows: (i) first, the current time is read in the clock register of the CPU and stored in a variable, say T; (ii) the list of input events to treat is updated: the input sensors are checked for new events issued by the environment; (*iii*) guards of the edges of the current locations are evaluated with the value stored in T. If at least one guard evaluates to true then take nondeterministically one of the enabled transitions; (iv) the next round is started. All we require from the hardware is to respect the following two requirements: (i) the clock register of the CPU is incremented every Δ_P time units and (ii) the time spent in one loop is bounded by a fixed value Δ_L . We choose this semantics for its simplicity and also because it is obviously implementable. The condition for the preservation of the correctness is that $\delta > 3\Delta_L + 4\Delta_P$.

Third, the AASAP semantics can be encoded by a classical single parameter timed automaton, so that it can by analyzed automatically by timed automata model-checkers like HYTECH or UPPAAL. However, this encoding has a limited interest in practice because its size is always exponential in $|Lab_{in}|$, the number of input labels of the controller. We solve this problem in the next section by giving a new translation which is compositional and at most quadratic in the size of the controller.

5 Compositional Construction for the AASAP Semantics

The main idea underlying our compositional construction is to treat the incoming events (issued by the environment) independently of the control structure of the ELASTIC controller, with a network of automata. This leads to technical difficulties we explain and address in this section.

Following the rule (A4.6) defining almost urgency of the AASAP semantics, there are essentially three reasons for allowing time to pass: (i) either the controller has been in its current location for less than Δ time units, (ii) or all last untreated occurences of an event have been issued by the environment less than Δ time units ago, (iii) or finally the guard of the outgoing transitions have not been enabled for more than Δ time units. Roughly, those conditions will be checked in our compositional construction by respectively A^2 , which is a transformation of the ELASTIC controller A, and two types of widgets: the event-watchers and the guardwatchers.

In timed automata, there is essentially one way for modeling urgency: invariants on locations. Roughly, if we have a transition guarded by a lower bound constraint g, it can be forced as soon as it is enabled by adding as invariant in its source



Figure 1: Event-Watcher W_{α} .

Figure 2: Guard-Watcher $W^l_{\alpha}(\bar{\varphi}_{\mathsf{evt}}(l, \alpha)).$

location the closure of $\neg g$. E.g. for a guard $x \ge 3$ we can add the invariant $x \le 3$. This way, time is blocked when the guard is satisfied and the discrete transition is forced. If we enlarge the invariant by Δ ($x \le 3 + \Delta$), we get the *almost urgency* we need. To formalize this idea, we will need to introduce some more notations:

Additional notations (i) Given an ELASTIC controller $A = \langle \text{Loc}, l_0, \text{Var}, \text{Lab}, \text{Edg} \rangle$ and a location $l \in \text{Loc}$, let $G_{\text{act}}(l) = \{g \mid (l, l', g, \sigma, R) \in \text{Edg} \land \sigma \in \text{Lab}_{out} \cup \text{Lab}_{\tau}\}$ be the set of guards labelling output transitions or internal transitions, and for $\alpha \in \text{Lab}_{in}^1$, let $G_{\text{evt}}(l, \alpha) = \{g \mid (l, l', g, \alpha, R) \in \text{Edg}\}$ be the set of guards labelling event transitions. (ii) Then define $\bar{\varphi}_{a}(l) = \bigwedge_{g \in G_{\text{act}}(l)} \neg (-\Delta(g)_0)$ and $\bar{\varphi}_e(l, \alpha) = \bigwedge_{g \in G_{\text{evt}}(l,\alpha)} \neg (-\Delta(g)_0)$. For example, let $G_{\text{act}}(l) = \{2 \leq x \leq 5, 0 \leq y \leq 1\}$, then $\bar{\varphi}_{a}(l) \equiv (x \leq 2 + \Delta \lor x \geq 5) \land (y \leq \Delta \lor y \geq 1)$.

Those constraints will be used as invariant to match the third part of rule (A4.6). The constraint $\bar{\varphi}_{a}(l)$ will be used as an invariant for location l in A^{2} to force an output transition when it becomes possible. The constraint $\bar{\varphi}_{e}(l, \alpha)$ will be used in the *guard-watchers*, to ensure that when a guard has been true for enough time, the corresponding transition becomes urgent (as long as it is allowed by other parts of rule (A4.6)).

Those invariants are central to our construction, but if we want a compositional construction (a product of automata), invariants are too restrictive to express urgency since urgency also depends on the current state of the other automata offering enabled synchronizations in the product. Hence, we should not block time simply when a transition is enabled in *one* automaton but only when it is enabled in *every* automaton of the product. Therefore, some compositional mechanism is needed to model urgency in a product: we will use the Asap flag. Remember that this flag expresses the fact that a transition is urgent as soon as it is enabled in the whole product.

The formal definition of our construction is given in Definition 6. From an ELASTIC controller A and a parameter Δ we construct $\mathcal{F}(A, \Delta)$ as a product of three types of components: event-watchers, guard-watchers and A^2 directly obtained from A. We omit the formal definitions of event-watchers and guard watchers which should be clear from the figures and anyway can be found in the appendix.

Event-Watcher Associated to an event $\alpha \in \Sigma_{in}$, we define W_{α} (see Figure 1) that records the event α . It has a clock z_{α} encoding the value of $I(\alpha)$ in the AASAP



Figure 4: The timed automaton A^2 associated to the ELASTIC controller A of Figure 3.

semantics. z_{α} records the time elapsed since the last untreated event α was issued by the environment. When $I(\alpha) \neq \bot$, the value of the clock z_{α} is equal to $I(\alpha)$.

This widget is intended to record the occurrence of the events α (as expressed by rule (A4.2) in the definition of the AASAP semantics), and then to propose a synchronization on $\bar{\alpha}$ with an Asap flag in location W_2 . Remember that the notation $\bar{\alpha}$ corresponds to the detection of event α by the controller. From the invariant of location W_1 , this synchronization will not become urgent before Δ time units.

Guard-Watchers. We introduce *Guard-Watchers* (see Figure 2) to monitor the truth value of a set of guards. They are associated to an event $\alpha \in \Sigma_{in}$ and a location $l \in \mathsf{Loc}$. When the controller is not in location l, the guard-watchers $W^l_{\alpha}(G)$ do not influence the execution, being in location U_0 and offering a self-loop synchronization on $\bar{\alpha}$. When location l is reached, the synchronization on in_l forces $W^l_{\alpha}(G)$ enter location U_1 and to become active. The watcher get back in U_0 as soon as l is exited by out_l . Thus, it is active when it is not in U_0 . Its role is then to prevent the label $\bar{\alpha}$ to become urgent whenever there is no transition labeled with $\bar{\alpha}$ that has been enabled for more than Δ units of time. Hence, we use $W^l_{\alpha}(G)$ with the set of guards $G = \bar{\varphi}_e(l, \alpha)$.

Controller transformation We illustrate the transformation of the ELASTIC controller with an example. The timed automaton A^2 corresponding to the ELASTIC controller A of Figure 3 is depicted on Figure 4. The automaton A^2 has a similar structure to A. It is used to (*i*) guarantee a maximum delay of Δ when location changes (as modeled by the variable d in the AASAP semantics) (*ii*) make transitions labeled with actions $\sigma \in \mathsf{Lab}_{out} \cup \mathsf{Lab}_{\tau}$ urgent when their guard has been satisfied for more than Δ time units (through invariant of Out_l) and (*iii*) enlarge the guards of the controller's transitions (as expressed by rules (A4.1), (A4.3) and (A4.4)).

Definition 6 [Compositional construction \mathcal{F}] Let $A = \langle \mathsf{Loc}^1, l_0^1, \mathsf{Var}^1, \mathsf{Lab}^1, \mathsf{Edg}^1 \rangle$

be an ELASTIC controller. The compositional construction $\mathcal{F}(A, \Delta)$ is the synchronized product of the following timed automata:

- the event-watchers W_{α} for every $\alpha \in \mathsf{Lab}_{in}^1$,
- the guard-watchers $W^l_{\alpha}(G_{\mathsf{evt}}(l,\alpha))$ for every $\alpha \in \mathsf{Lab}^1_{\mathsf{in}}, l \in \mathsf{Loc}^1$,
- and the timed automaton $A^2 = \langle \mathsf{Loc}^2, l_0^2, \mathsf{Var}^2, \mathsf{Inv}^2, \mathsf{Lab}^2, \mathsf{Edg}^2, \mathsf{Asap}^2 \rangle$ where:
 - $-\operatorname{Loc}^{2} = \{\operatorname{PreIn}_{l}, \operatorname{In}_{l}, \operatorname{Out}_{l}, \operatorname{PostOut}_{l,l'} \mid l, l' \in \operatorname{Loc}^{1}\};$
 - $l_0^2 = In_{l_0^1};$
 - $\operatorname{Var}^2 = \operatorname{Var}^1 \cup \{d\};$
 - $\mathsf{Lab}_{\mathsf{out}}^2 = \mathsf{Lab}_{\mathsf{out}}^1, \mathsf{Lab}_{\mathsf{in}}^2 = \emptyset \text{ and } \mathsf{Lab}_{\tau}^2 = \mathsf{Lab}_{\tau}^1 \cup \overline{\mathsf{Lab}_{\mathsf{in}}^1} \cup \{\epsilon, in_l, out_l\};$
 - $\begin{array}{l} \ \mathsf{Edg}^2 \ \text{contains} \ \ (a) \ \text{the edges} \ (Out_l, PostOut_{l,l'}, \vartriangle[g]_{\Delta}, \sigma, R \cup \{d\}) \ \text{such} \\ \text{that there exists} \ (l,l',g,\sigma,R) \in \mathsf{Edg}^1 \ \text{with} \ \sigma \in \mathsf{Lab}^1_{\mathsf{out}} \cup \mathsf{Lab}^1_{\tau} \ (b) \ \text{the edges} \\ (Out_l, PostOut_{l,l'}, \vartriangle[g]_{\Delta}, \bar{\alpha}, R \cup \{d\}) \ \text{such that there exists} \ (l,l',g,\alpha,R) \in \mathsf{Edg}^1 \ \text{with} \ \alpha \in \mathsf{Lab}^1_{\mathsf{in}} \ \text{and} \ (c) \ \text{the edges} \ (PostOut_{l,l'}, PreIn_{l'}, \emptyset, out_l, \emptyset) \ \text{for} \\ \text{each} \ l,l' \in Loc^1, \ \text{and the edges} \ (PreIn_l, In_l, \emptyset, in_l, \emptyset) \ \text{and} \ (In_l, Out_l, \emptyset, \epsilon, \emptyset) \\ \text{for each} \ l' \in Loc^1. \end{array}$
 - $\mathsf{Asap}^2(e) = \bot$ for every $e \in \mathsf{Edg}^2$;
 - The function Inv^2 is defined as follows. For each $l, l' \in Loc^1$, (a) $\operatorname{Inv}^2(In_l) = \{\{d \leq \Delta\}\}\ (b)\ \operatorname{Inv}^2(Out_l) = \{\{d \leq \Delta \lor \bar{\varphi}_{a}(l)\}\}\$ and (c) $\operatorname{Inv}^2(PreIn_l) = \operatorname{Inv}^2(PostOutl, l') = \{\{d = 0\}\}.$

In summary, $\mathcal{F}(A, \Delta) = A^2 \times \prod_{\alpha \in \mathsf{Lab}_{i\alpha}^1} W_\alpha \times \prod_{\alpha \in \mathsf{Lab}_{i\alpha}^1, l \in \mathsf{Loc}^1} W_\alpha^l(G_{\mathsf{evt}}(l, \alpha)).$

The correctness of our compositional construction is established by the following theorem.

Theorem 1 For any ELASTIC controller A, for any environment STTS E and a set Bad of its states, for any $\delta \in \mathbb{Q}^{>0}$, $[\![A]\!]^{AAsap}_{\delta}$ controls E to avoid Bad iff $[\![\mathcal{F}(A, \Delta)]\!]_{\delta}$ controls E to avoid Bad.

The proof can be found in appendix. Since the correctness of AASAP semantics of A implies its implementability, we can verify the compositional construction with an automatic tool and generate systematically the implementation code. In the second part of this paper, we show how we have applied this methodology in practice on a real-world protocol.

6 Tool Set

We briefly describe the tool set that we have implemented. The structure of the tool set is depicted in Figure 5 and it consists of three tools: (i) ELASTIC2HYTECH, (ii) HYTECH2UPPAAL, and (iii) ELASTIC2BRICK.

ELASTIC2HYTECH is the main component of the tool set: it implements the compositional construction of Section 5. Given an ELASTIC controller Cont (expressed in an HYTECH like syntax), it produces a one-parameter HYTECH specification $Cont'(\Delta)$ following the construction defined in the previous section. To obtain a model of the entire system, this specification of the controller has to be composed with a model of the environment (in which the controller is embedded). This is



Figure 5: Structure of our tool set.

given as a product of rectangular automata (in HYTECH syntax). The environment is noted Env in the sequel. We can then use HYTECH to reason about the system. The following three correctness problems can be formulated and answered with HYTECH (if the analysis terminates):

- [Fixed] Given a set of bad states Bad, a value δ ∈ Q^{≥0}, does the controller, when reacting within δ, control the environment to avoid Bad, that is: Reach([Cont'(Δ)||Env]_Δ) ∩ Bad = Ø
- [Existence] Given a set of bad states Bad, does there exist a value δ ∈ Q such that when the controller reacts within δ, it controls the environment to avoid Bad, that is: ∃δ > 0 : Reach([Cont'(Δ)||Env]]_δ) ∩ Bad = Ø
- [Maximization] Given a set of Bad, what is the largest value for δ ∈ Q^{≥0} such that when the controller reacts within δ, it controls the environment to avoid Bad, that is: max{δ > 0 : Reach([Cont'(Δ)||Env]]_δ) ∩ Bad = Ø}

To tackle large examples, we also use UPPAAL. The tool HYTECH2UPPAAL translates HYTECH specifications into UPPAAL specifications. As UPPAAL is restricted to the analysis of timed automata (and it does that very efficiently), it is only applicable if the environment can be modeled as a product of timed automata and the parameter Δ is fixed. Obviously, this allows us to answer the [fixed] version of the correctness problem. Thanks to the "faster is better" property of the AASAP semantics, by doing a binary search on the value space of the parameter δ , we can approximate the maximal value of δ for which the controller is correct up to any precision.

The main purpose of the AASAP semantics is to give a way to synthesize executable code for a controller from its model and to ensure that the properties that have been proved on the model are preserved on the code (without making the synchrony hypothesis). To obtain executable code from the ELASTIC model of a controller, we use the tool ELASTIC2BRICK that produces C-code from an annotated ELASTIC specification. The annotations assign to each transition a piece of code that has to be executed when the transition is fired. The translation is very simple: we assign to each edge of the ELASTIC controller a thread that is ran when the associated input event is perceived or when the associated output event has to be produced. We have chosen to produce code for LEGO MINDSTORMSTM running BRICKOS³. LEGO MINDSTORMSTM are toys but the internals are a fully functional micro-computer linked with sensor and actuators. When running BRICKOS, we can use priorities to ensure real-time properties of the code that is executed on the Brick. Details can be found in [Doy03].

7 Case Study: the "Philips Audio Control Protocol"

Introduction Bosscher et al study in [BPV94] "a simple protocol for the physical layer of an interface bus that connects the devices of a stereo equipment". This protocol was proposed by Philips engineers. The protocol is based on Manchester encoding to transmit binary sequences on a wire between a single sender and a single receiver.

In our case study, we will use LEGO MINDSTORMSTM Bricks to implement the sender and the receiver. To connect the two Bricks, we use a wire plugged to an output gate of the sender and to an input gate of the receiver. The difficulties here to implement the protocol are similar to the ones that the engineers in Philips were facing: (i) although the receiver knows the length of a time slot, it does not know when it begins (the two Bricks are running asynchronously); (ii) a receiver does not know the length of the bit string it is receiving; (iii) only UP signals can be reliably detected by our sensors (this constraint is taken to fit with the case study of [BPV94]; (iv) the sender and the receiver have digital clocks that have finite granularity, so there will be imprecision in both sending and receiving times; (v) in BRICKOS sensors are polled periodically. As a consequence, the moment at which a bit is perceived can be substantially later than the moment it has been sent. The first three difficulties should be solved by the logic of the protocol. The last two difficulties are much lower level and we would like to forget them when designing a high level version of the protocol. This is exactly what the AASAP semantics allows us to do.

Next, we present the idealized version of the protocol and how we modeled it with two ELASTIC controllers: one for the sender and one for the receiver. Here, the environment is an observer that compares the sequence of bits sent by the sender with the sequence of bits decoded by the receiver. The observer reaches the location *error* whenever the two sequences do not match.

Afterwards, we explain how we can use the AASAP semantics during the verification process and verify the robustness of the protocol. The verification phase allows us to generate code that is correct by construction.

ELASTIC **models.** An idealized version of the protocol uses evenly spaced time slots. To transmit a 1, the sender must let the signal go from low voltage to high in the middle of a slot and from high to low for a 0. To repeat a bit, the sender is thus forced between two slots to turn the signal off for a 1 or on for a 0. The receiver is not able to detect precisely moments when the signal goes down and then only relies on the UP signals to decode the messages. This implies that a message has to begin by a 1 and that messages ending in 10 or in 1 are not distinguishable without adding information bits. Rather than adding bits, the protocol restricts messages to be either odd in length or to end in 00.

³http://brickos.sourceforge.net/



Figure 6: The Sender automaton.



Figure 7: The Receiver automaton.



Figure 8: The Observer automaton.

Our modelisation of the protocol can be found in Figure 6 for the sender and Figure 7 for the receiver. There is an additional *observer automaton* playing the role of the environment on Figure 8 that allows us to verify the correct transmission of the bits (this observer was proposed by Ho and Wong-Toi in [HWT95]). The unit of time of the model, noted U, is a quarter of the time slot. This unit is not written in the constraints, to alleviate the presentation.

One can easily check that the sender automaton can send any sequence conforming to the protocol restrictions. Arrival in location OneSent (ZeroSent) means the signal for a 1 (a 0) has just been sent. The clock x is used for the timing of the sequence. The discrete variable i is non-deterministically set to 1 or 0 each time a bit is sent (not shown on the figures). Its value determines which shall be the next bit. The discrete variables p and doublezero encode respectively if the current sequence is odd in length and if it ends in 00. Finally, the discrete variables c and leng are used to encode the bits that have been sent but not decoded by the receiver yet. c simply encodes in an integer the binary word composed of the last such bits and leng is the number of those bits. The decrementing of c and leng is done by the observer automaton every time it succeeds in matching a sent bit with a received bit.

The receiver automaton decodes its incoming UP signals by rounding its local time for when it received the signal to the nearest possible time it expects a signal. This is what makes the protocol robust. If no signal is received in due time, the sequence is interpreted as being complete. The discrete variable m is used to encode parity of the received sequence. It allows the receiver to know if it has to complete a sequence with an additional 0 to conform to the protocol restrictions. The discrete variable r encodes the one or two bits that were last received. This variable is checked by the observer automaton against c and *leng* of the sender to verify if the sent bits are the same as the received ones. The label FINALZERO does not correspond to an event. It is an internal action done when the receiver understands it must add a 0 to the sequence to end it. The observer automaton then synchronizes on this label to know a new bit has been decoded. As said before, the receiver does not synchronize on DOWN signals.

This modelisation uses finite range discrete variables, which are not present in the formal definitions. This is not a problem since all those discrete variables are bounded and thus could be encoded in locations. For the sake of clarity, we did not do this. Furthermore, the tools that we are using allow the use of such finite range discrete variables.

Parametric verification Let us now turn to the use of the AASAP semantics during the verification phase. We take the opportunity here to present some methodological aspects too.

Using ELASTIC2HYTECH, we generate for the sender and the receiver the HYTECH specification of their AASAP semantics following Definition 6. Those two semantics are noted $[Sender(\Delta)]_{\delta_1}^{AAsap}$ and $[Receiver(\Delta)]_{\delta_2}^{AAsap}$.

We can first check that if the protocol executed in an idealized setting, that is for $\delta_1 = 0$ and $\delta_2 = 0$, is correct. This is formalized by the following question: Reach([Sender(Δ)]]^{AAsap}₀||[Receiver(Δ)]]^{AAsap}₀||[Observer]]) \cap Bad $\neq \emptyset$, where Bad are the states in which the observer is in location *error*. With HYTECH (or UPPAAL), we can easily show that this test is passed successfully by our modelisation of the protocol. If this verification had failed then we should have concluded that the protocol was flawed in its logic.

To continue the study of the protocol and determine if it can be implemented, we should check its robustness. In our context, we must determine what are the maximum values of δ_1 and δ_2 which ensure that the system $[\![Sender(\Delta)]\!]_{\delta_1}^{AAsap}|\!|$ $[\![Receiver(\Delta)]\!]_{\delta_2}^{AAsap}|\!|[\![Observer]\!] \cap Bad = \emptyset$. Those maximal value will be expressed in the unit of time U of the system that we have not fixed so far. Remember U is a quarter of a timeslot. By tuning this value, we can then maximize the throughput of the protocol. We should then look for the smallest implementable U on our implementation platform. For BRICKOS, the value $\Delta_L U$ (length of the loop in the execution procedure) and $\Delta_P U$ (precision of the clocks) can be set to as low as 6 ms and 1 ms. To guarantee a correct implementation of Sender(Δ) (and Receiver(Δ)), we need to have $\Delta > 3\Delta_L + 4\Delta_P$, and so $\Delta U > 22ms$.

So, we know that $\delta_1 U$ and $\delta_2 U$ should be srictly below 22 ms. If $\delta_1 \leq \delta_2$, the infimum for U is $\frac{22 \text{ ms}}{\delta_1}$ else it is $\frac{22 \text{ ms}}{\delta_2}$. Now if we increase the value of one of the parameters δ_i , the correct value for the other decreases. This is because increasing the parameter value for the AASAP semantics of a controller strictly increases its looseness, forcing the other to be more precise as compensation, which corresponds to a smaller value for its parameter. Using this fact, we can conclude that the best U for the system will be obtained when δ_1 and δ_2 are equal.

Guiding HYTECH with this information, by a parametric search, we found that, for ensuring correctness, the parameters must be strictly less than $\frac{1}{4}U$. In fact, we proved that a sufficient condition to avoid the *error* state is that $\delta_1 + \delta_2 < \frac{1}{2}$. Execution times of different analysis are given in Figure 9. Note that to make HYTECH terminate, we needed to give some initial constraints. Execution times with UPPAAL are very encouraging: the problems solved are simpler as the models are not parametric but this problems are those to be solved in practice as a precise parametric analysis is nice in theory but not required in practice (if the target platform is fixed).

Implementation From annotated models of the sender and the receiver, we have generated, using **ELASTIC2BRICK** the C-code for the sender and the receiver. The C files are about 500 lines long for each controller. The annotations of the models are very natural. Here are some examples of annotations. Assume that we want to use the protocol to exchange variable length strings of bits that are stored in an array, say **A** in the sender and **B** in the receiver. Instead of assigning the bit variable *i* non-deterministically, we should execute the annotation $\{i := A[j]; j++;\}$, and in the

Tool	Constraint	Result	Time
НүТесн	$\delta_1 + \delta_2 < 1/2$	Safe	55s
	$\delta_1 = \delta_2 = 1/5$	Safe	50s
	$\delta_1 = \delta_2 = 1/4$	Unsafe	90s
Uppaal	$\delta_1 = \delta_2 = 1/5$	Safe	< 1s
	$\delta_1 = \delta_2 = 1/4$	Unsafe	< 1s

Figure 9: Execution times for the different models.

Receiver automaton, we add the code {B[k]:= α ; k++;} to transitions setting r to $\alpha \in \{0, 1\}$, and the code {B[k]:=0; B[k+1]:=1; k+=2;} to the transition setting r to 2.

Evaluation The code that we have obtained is correct by construction and can safely be executed on LEGO MINDSTORMSTM Brick as an alternative communication mean with real-time guarantees. For that, it suffices to give the highest level of priority to the protocol to ensure its real-time behavior. This should not spoil the behavior of other applications running on the Brick as the resources needed by the protocol are very low. Now, let us look at the performance of the protocol in our implementation. The throughput obtained, when the length of the sequence goes to infinity, is around 2.84 bits per seconds. This may look quite low and we could think that far better throughput could be obtained by a hand-made implementation. But this is not the case. Indeed, we can show using the results of Ho and Wong-Toi [HWT95] and by taking into account only the imprecision due to reading on digital clocks every time slice, that the throughput of the protocol on LEGO MINDSTORMSTM is bounded from above by around 4.16 bits per seconds. So, the price in term of performance loss to obtain automatically generated and correct code is not too high in our opinion. Let us also note that we were only able to find error by testing when the throughput was set around 7 bits per seconds. That shows the limit of testing at least when done in a naive way.

References

- [AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- [BPV94] D. Bosscher, I. Polak, and F. Vaandrager. Verification of an Audio Control Protocol. In H. Langmaack, W.-P. de Roever, and J. Vytopil, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems*, volume 863, pages 170–192, Lübeck, Germany, 1994. Springer-Verlag.
- [CHR02] F. Cassez, T.A. Henzinger, and J.-F. Raskin. A comparison of control problems for timed and hybrid systems. In HSCC 02: Hybrid Systems— Computation and Control, Lecture Notes in Computer Science 2289, pages 134–148. Springer-Verlag, 2002.
- [DDR04] M. De Wulf, L. Doyen, and J.-F. Raskin. Almost ASAP semantics: From timed models to timed implementations. In HSCC 04: Hybrid Systems—Computation and Control, Lecture Notes in Computer Science 2993, pages 296–310. Springer-Verlag, 2004.

- [Doy03] Laurent Doyen. A systematic implementation of simple timed controllers. Technical Report 504, U.L.B., 2003.
- [HHWT95] Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. Hytech: The next generation. In 16th Annual Real-Time Systems Symposium (RTSS), pages 56–65. IEEE Computer Society Press, 1995.
- [HWT95] P.-H. Ho and H. Wong-Toi. Automated analysis of an audio control protocol. In P. Wolper, editor, *Proceedings of the 7th International Conference On Computer Aided Verification*, volume 939, pages 381– 394, Liege, Belgium, 1995. Springer Verlag.
- [LT87] N. Lynch and M. Tuttle. Hierarchical correctness proofs for distributed algorithms. In 6th ACM Symp. on Principles of Distributed Computing, pages 137–151, 1987.
- [PL00] Paul Pettersson and Kim G. Larsen. UPPAAL2k. Bulletin of the European Association for Theoretical Computer Science, 70:40–44, February 2000.

Preliminary Definitions Α

Below, we give the definitions of input enabledness, composition and safety control which we have informally introduced in Section 2. We also introduce weak simulation and hiding, which we will need to prove Theorem 1.

Definition 7 [Input enabled STTS] A STTS $\mathcal{T} = \langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau}, \rightarrow \rangle$ is input enabled if for all $\sigma \in \Sigma_{in}$, for all $s_1 \in S$ there exists $s_2 \in S$ such that $(s_1, \sigma, s_2) \in \rightarrow$.

Definition 8 [Weak simulation relation for STTS] Given two STTS $\mathcal{T}^1 = \langle S^1, \iota^1, \Sigma_{\text{in}}^1, \Sigma_{\text{out}}^1, \Sigma_{\tau}^1, \rightarrow^1 \rangle$ and $\mathcal{T}^2 = \langle S^2, \iota^2, \Sigma_{\text{in}}^2, \Sigma_{\text{out}}^2, \Sigma_{\tau}^2, \rightarrow^2 \rangle$, let $\Sigma = \Sigma_{\text{out}}^1 \cup \Sigma_{\text{in}}^1 \cup \Sigma_{\tau}^1$, we say that \mathcal{T}^1 is weakly simulable by \mathcal{T}^2 , noted $\mathcal{T}^1 \sqsubseteq \mathcal{T}^2$, if there exists a relation $\mathcal{R} \subseteq S^1 \times S^2$ (called a *weak simulation relation*) such that :

- $(\iota^1, \iota^2) \in \mathcal{R}$
- for any $(s_1^1, s_1^2) \in \mathcal{R}$, for any $\sigma \in \Sigma \setminus \{\epsilon\} \cup \mathbb{R}^{\geq 0}$, for any s_2^1 s.t. $(s_1^1, \sigma, s_2^1) \in \to^1$, there exists $s_2^2, s_3^2, \ldots, s_n^2 \in S^2$ s.t. for some j, $(s_j^2, \sigma, s_{j+1}^2) \in \to^2$ and for any $i \neq j$, $(s_i^2, \epsilon, s_{i+1}^2) \in \to^2$ and $(s_2^1, s_n^2) \in \mathcal{R}$;

In definitions and theorems of [DDR04], we used a classical simulation relation (without the possibility of using ϵ transitions). The results we have obtained are not affected by this new definition because weak simulation preserve observable labels as simulation does.

Definition 9 [Hiding] Given a STTS $\mathcal{T} = \langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau}, \rightarrow \rangle$ and a subset $H \subseteq$ Σ_{τ} of its internal labels, define the *hiding* of H in \mathcal{T} by $\mathcal{T} \downarrow H = \langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau} \backslash H, \to'$ \rangle where $(s, \sigma, s') \in \to'$ iff either $\sigma \in \Sigma_{in} \cup \Sigma_{out} \cup \Sigma_{\tau} \setminus H$, $\sigma \neq \epsilon$ and $(s, \sigma, s') \in \to$ or $\sigma = \epsilon$ and there exists $\sigma' \in H$ such that $(s, \sigma', s') \in \rightarrow$.

Definition 10 [Composition of STTS] Let $\mathcal{T}^1 = \langle S^1, \iota^1, \Sigma_{\mathsf{in}}^1, \Sigma_{\mathsf{out}}^1, \Sigma_{\tau}^1, \to^1 \rangle$ and $\mathcal{T}^2 = \langle S^2, \iota^2, \Sigma_{\mathsf{in}}^2, \Sigma_{\mathsf{out}}^2, \Sigma_{\tau}^2, \to^2 \rangle$ be two STTS. Let $\Sigma^1 = \Sigma_{\mathsf{in}}^1 \cup \Sigma_{\mathsf{out}}^1 \cup \Sigma_{\tau}^1$ and $\Sigma^2 = \Sigma_{\mathsf{in}}^2 \cup \Sigma_{\mathsf{out}}^2 \cup \Sigma_{\tau}^2$. We say that \mathcal{T}^1 and \mathcal{T}^2 are *composable* if whenever $\sigma \in \Sigma^1 \cap \Sigma^2$ then either $\sigma \in \Sigma_{in}^1 \cap \Sigma_{out}^2$ or $\sigma \in \Sigma_{out}^1 \cap \Sigma_{in}^2$. Their composition, noted $\mathcal{T}^1 \| \mathcal{T}^2$ is the STTS $\mathcal{T} = \langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau}, \rightarrow \rangle$ such

that:

let
$$\Sigma_{in}^{12} = \Sigma_{in}^1 \cup \Sigma_{in}^2$$
, $\Sigma_{out}^{12} = \Sigma_{out}^1 \cup \Sigma_{out}^2$ and $\Sigma_{\tau}^{12} = \Sigma_{\tau}^1 \cup \Sigma_{\tau}^2$

$$\begin{split} S &= S^1 \times S^2, \ \iota = (\iota^1, \iota^2), \ \Sigma_{\mathsf{in}} = \Sigma_{\mathsf{in}}^{12} \backslash \Sigma_{\mathsf{out}}^{12}, \ \Sigma_{\mathsf{out}} = \Sigma_{\mathsf{out}}^{12} \backslash \Sigma_{\mathsf{in}}^{12}, \ \Sigma_{\tau} = \Sigma_{\tau}^{12} \cup \left(\Sigma_{\mathsf{in}}^{12} \backslash \Sigma_{\mathsf{in}} \right), \\ \text{and} \to \text{is such that} \ \left((s_1^1, s_1^2), \sigma, (s_2^1, s_2^2) \right) \in \to \text{ iff one of the following three assertions} \end{split}$$
holds:

- $\sigma \in (\Sigma_{in}^{12} \setminus \Sigma_{in}) \cup \mathbb{R}^{\geq 0}$ and $(s_1^1, \sigma, s_2^1) \in \to^1$ and $(s_1^2, \sigma, s_2^2) \in \to^2$
- $\sigma \in \Sigma^1 \cap (\Sigma_{in} \cup \Sigma_{out} \cup \Sigma^{12}_{\tau})$ and $(s_1^1, \sigma, s_2^1) \in \to^1$ and $s_1^2 = s_2^2$
- $\sigma \in \Sigma^2 \cap (\Sigma_{in} \cup \Sigma_{out} \cup \Sigma^{12}_{\tau})$ and $(s_1^2, \sigma, s_2^2) \in \to^2$ and $s_1^1 = s_2^1$

Definition 11 [Safety Control] Let $\mathcal{T}^1 = \langle S^1, \iota^1, \Sigma_{in}^1, \Sigma_{out}^1, \Sigma_{\tau}^1, \rightarrow^1 \rangle$ and $\mathcal{T}^2 = \langle S^2, \iota^2, \Sigma_{in}^2, \Sigma_{out}^2, \Sigma_{\tau}^2, \rightarrow^2 \rangle$ be two composable STTS. Let $B \subseteq S^2$, we say that \mathcal{T}^1 controls \mathcal{T}^2 to avoid B if the following condition holds:

 $\mathsf{Reach}(\mathcal{T}^1 || \mathcal{T}^2) \cap \{(s^1, s^2) \mid s^1 \in S^1 \land s^2 \in B\}$ is empty.

Theorem 2 ([DDR04]) Let \mathcal{T}^1 and \mathcal{T}^2 be two composable STTS, let \mathcal{T}^3 be a STTS such that $\mathcal{T}^3 \sqsubseteq \mathcal{T}^1$, and let $B \subseteq S^2$, if \mathcal{T}^1 controls \mathcal{T}^2 to avoid B then \mathcal{T}^3 controls \mathcal{T}^2 to avoid B.

Β Proof of Theorem 1

Before proving Theorem 1, we need some preliminary lemmas and definitions presented below.

First, we add to the notations introduced in Section 4 the following shortcuts: given an ELASTIC controller $A = (Loc, l_0, Var, Lab, Edg)$, a location $l \in Loc$, and an input label $\alpha \in \mathsf{Lab}_{in}^1$, let $\varphi_{\mathsf{a}}(l) = \bigwedge_{g \in G_{\mathsf{act}}(l)} \neg(-\Delta(g]_0)$; and $\varphi_{\mathsf{e}}(l, \alpha) = \varphi_{\mathsf{e}}(l, \alpha)$ $\bigwedge_{g \in G_{\mathsf{evt}}(l,\alpha)} \neg (-\Delta(g]_0). \text{ Thus, } \bar{\varphi}_{\mathsf{a}}(l) \text{ and } \bar{\varphi}_{\mathsf{e}}(l,\alpha), \text{ defined in section 5, are the closure of the section for the se$ of $\varphi_{a}(l)$ and $\varphi_{e}(l, \alpha)$ respectively.

Lemma 1 Let $g \in \text{Rect}(\text{Var})$ be a guard and let $\delta \in \mathbb{Q}^{\geq 0}$. For any valuation v, we have:

- (a) $\mathsf{TS}(v,g) \leq \delta$ iff $v \models_{\delta} \neg (-\Delta(g]_0)$
- (b) $v \models_{\delta} \neg (-\Delta(g)_0)$ iff $v \models_{\delta} \neg (-\Delta(g)_0) \lor \forall t \in \mathbb{R}^{>0} : v + t \models_{\delta} \neg (-\Delta(g)_0)$

Proof. (a) See also Figure 10. For $x \in Var$, let $a_x = lb(g(x))$ and $b_x = rb(g(x))$). $\mathsf{TS}(v,g) \leq \delta$

- $(v \models_{\delta} g) \to \exists x \in \mathsf{Var} : v(x) \delta \leq a_x$ \leftrightarrow
- $(v \models_{\delta} g) \rightarrow (\exists x \in \mathsf{Var} : a_x \leq v(x) \leq b_x \land v(x) \delta \leq a_x)$
- $\leftrightarrow \quad (v \not\models_{\delta} g) \lor (\exists x \in \mathsf{Var} : a_x \le v(x) \le \min(b_x, a_x + \delta))$
- $(\exists x \in \mathsf{Var} : v(x) < a_x \lor v(x) > b_x)$ \leftrightarrow
- $\begin{array}{l} \forall (\exists x \in \mathsf{Var}: a_x \leq v(x) \leq \min(b_x, a_x + \delta)) \\ \leftrightarrow \quad \exists x \in \mathsf{Var}: v(x) \leq a_x + \delta \lor v(x) > b_x \end{array}$
- $\leftrightarrow \quad \exists x \in \mathsf{Var} : v(x) \notin (a_x + \delta, b_x]$
- $\leftrightarrow v \not\models_{\delta} -\Delta(g]_0$
- $\leftrightarrow \quad v \models_{\delta} \neg (-\Delta(g]_0)$

(b) The second part of the Lemma can be rewritten: $v \models_{\delta} -\Delta(g)_0$ iff $v \models_{\delta}$ $(-\Delta(g]_0) \wedge \exists t \in \mathbb{R}^{>0} : v + t \models_{\delta} -\Delta(g]_0$. If $v \models_{\delta} -\Delta(g)_0$, then $\forall x \in \mathsf{Var} : a + d = 0$ $\delta < v(x) < b$ which implies that $\forall x \in Var$ we have both $a + \delta < v(x) \leq b$ and $\exists t \in \mathbb{R}^{>0} : a + \delta < v(x) + t \leq b$. The converse holds obviously.

Definition 12 [Right-closed predicates] A (parametric) rectangular constraint " $a \sim_1$ $x \sim_2 b$ " is right-closed if $\sim_2 = \leq \leq$. A (parametric) rectangular predicate is rightclosed if it is a finite set of right-closed (parametric) rectangular constraints. A (parametric) multirectangular predicate is *right-closed* if it is a finite set of rightclosed (parametric) rectangular predicates.

Notice that $\varphi_{a}(l)$, $\bar{\varphi}_{a}(l)$, $\varphi_{e}(l, \alpha)$ and $\bar{\varphi}_{e}(l, \alpha)$ are right-closed parametric multirectangular predicates.

Lemma 2 Let g be a right-closed parametric multirectangular predicate and let $\delta \in \mathbb{Q}^{\geq 0}$. For any valuation v, if for some $t_1, t_2 \in \mathbb{R}$, $t_1 < t_2$, we have $\forall t \in [t_1, t_2)$: $v + t \models_{\delta} g$, then $v + t_2 \models_{\delta} g$.

Proof. The proof of this lemma is left to the reader.

To prove Theorem 1, we will need to formally define what is a watcher and what is a guard-watcher. A graphic representation of those two definitions are respectively given in figure 1 and 2



Figure 10: The function $\mathsf{TS}(v, a \le x \le b)$ is strictly above δ iff $v(x) \in (a + \delta, b]$.

Definition 13 [Event-Watcher] Given an ELASTIC controller $A = (Loc^1, l_0, Var^1, Lab^1, Edg^1)$, for $\alpha \in Lab_{in}$ define the timed automaton *Event-Watcher* $W_{\alpha} = (Loc, W_0, Var, Inv, Lab, Edg, Asap)$ where:

- Loc = { W_0, W_1, W_2 }
- Var = $\{z_{\alpha}\}$
- $\operatorname{Inv}(W_0) = \operatorname{Inv}(W_2) = \{\emptyset\} \text{ and } \operatorname{Inv}(W_1) = \{\{z_\alpha \leq \Delta\}\}$
- $\mathsf{Lab}_{\mathsf{in}} = \{\alpha\}, \, \mathsf{Lab}_{\mathsf{out}} = \emptyset \text{ and } \mathsf{Lab}_{\tau} = \{\epsilon, \bar{\alpha}\}$
- $\mathsf{Edg} = \{e_i | 1 \le i \le 5\}$ where:
 - $e_1 = (W_0, W_1, \emptyset, \alpha, \{z_{\alpha}\})$
 - $-e_2 = (W_1, W_2, \emptyset, \epsilon, \emptyset)$
 - $-e_3 = (W_2, W_0, \emptyset, \overline{\alpha}, \emptyset)$
 - $-e_4 = (W_1, W_1, \emptyset, \alpha, \emptyset)$
 - $e_5 = (W_2, W_2, \emptyset, \alpha, \emptyset).$
- $\mathsf{Asap}(e_1) = \mathsf{Asap}(e_2) = \mathsf{Asap}(e_4) = \mathsf{Asap}(e_5) = \bot$ and $\mathsf{Asap}(e_3) = \top$.

Definition 14 [Guard-Watchers] Given an ELASTIC controller $A = \langle \mathsf{Loc}^1, l_0, \mathsf{Var}^1, \mathsf{Lab}^1, \mathsf{Edg}^1 \rangle$, for an event $\alpha \in \mathsf{Lab}_{in}$, a location $l \in \mathsf{Loc}^1$ and a set of guards $G \subseteq \mathsf{Rect}_{\mathsf{c}}(\mathsf{Var})$, define the timed automaton *Guard-Watcher* $W^l_{\alpha}(G) = \langle \mathsf{Loc}, U_{\mathsf{init}}, \mathsf{Var}, \mathsf{Inv}, \mathsf{Lab}, \mathsf{Edg}, \mathsf{Asap} \rangle$ where:

- Loc = { U_0, U_1, U_2, U_3 };
- $U_{\text{init}} = U_1$ if $l = l_0$ and $U_{\text{init}} = U_0$ otherwise;
- Var = Var¹ \cup {u};

- $\operatorname{Inv}(U_0) = \operatorname{Inv}(U_1) = \{\emptyset\}, \operatorname{Inv}(U_2) = \bigwedge_{g \in G} \neg(-\Delta(g)_0) \text{ and } \operatorname{Inv}(U_3) = \{\{u = 0\}\}, v \in U_1\}$ 0};
- $\mathsf{Lab}_{\mathsf{in}} = \emptyset$, $\mathsf{Lab}_{\mathsf{out}} = \emptyset$ and $\mathsf{Lab}_{\tau} = \{\bar{\alpha}, in_l, out_l, \epsilon\};$
- $\mathsf{Edg} = \{e_i | 1 \le i \le 7\}$ where:

$$- e_1 = (U_0, U_1, \emptyset, in_l \emptyset)$$

 $- e_2 = (U_1, U_0, \emptyset, out_l, \emptyset)$ $e_0 = (U_1, U_2, \emptyset, \epsilon, \emptyset)$

$$- e_3 = (U_1, U_2, \emptyset, \epsilon, \emptyset)$$
$$- e_4 = (U_2, U_1, \emptyset, \epsilon, \emptyset)$$

$$- e_5 = (U_1, U_3, \emptyset, \bar{\alpha}, \{u\})$$

- $-e_6 = (U_3, U_0, \emptyset, out_l, \emptyset)$
- $-e_7 = (U_0, U_0, \emptyset, \overline{\alpha}, \emptyset);$
- $\mathsf{Asap}(e) = \bot$ for every $e \in \mathsf{Edg}$.

For the simplicity of the presentation, we define the meta-transitions associated to an STTS so that there will be a 1-to-1 correspondence between the steps in the AASAP semantics and the compositional construction. This is correct since we consider weak simulation relations.

Definition 15 [Meta-transition] Given a STTS $\langle S, \iota, \Sigma_{in}, \Sigma_{out}, \Sigma_{\tau}, \rightarrow \rangle$ such that $\epsilon \in$ Σ_{τ} , define the meta-transition relation $\Rightarrow \subseteq S \times \Sigma_{\mathsf{in}} \cup \Sigma_{\mathsf{out}} \cup \Sigma_{\tau} \cup \mathbb{R}^{\geq 0} \times S$ by:

- $(s,\epsilon,s') \in \Rightarrow$ iff s = s' or there exist some states $s_1,\ldots,s_n \in S$ such that $s_1 = s, s_n = s'$ and for every $i, 1 \le i < n$, we have $(s_i, \epsilon, s_{i+1}) \in \rightarrow$;
- For $\sigma \in \Sigma_{in}$, $(s, \sigma, s') \in \Rightarrow$ iff $(s, \sigma, s') \in \rightarrow$;
- For $\sigma \in \Sigma_{\mathsf{out}} \cup \Sigma_{\tau} \setminus \{\epsilon\}, (s, \sigma, s') \in \Rightarrow$ iff there exist some states $s_1, s_2, s_3 \in S$ such that $(s, \epsilon, s_1) \in \Rightarrow$, $(s_1, \sigma, s_2) \in \rightarrow$, $(s_2, \epsilon, s_3) \in \rightarrow$ and $(s_3, \epsilon, s') \in \rightarrow$.
- For $t \in \mathbb{R}^{\geq 0}$, $(s, t, s') \in \Rightarrow$ iff there exist some states $s_1, s'_1, s_2, s'_2, \ldots, s_n \in S$ such that $s_1 = s$ and $s_n = s'$ and some real numbers $t_1, \ldots, t_n \in \mathbb{R}^{\geq 0}$ such that $t = \sum_{i=1}^{n} t_i$ and for every $i, 1 \le i < n, (s_i, t_i, s'_i) \in \rightarrow$ and $(s'_i, \epsilon, s_{i+1}) \in \Rightarrow$;

The third item of the definition above imposes two ϵ transitions after an action because after having output σ , the automaton A^2 changes its location, and doing so it issues two ϵ transitions for (i) getting out of its current location l (synchronizing on out_l) and (ii) getting in its target location $l'(synchronizing on in_{l'})$.

Proof of Theorem 1.

Let us first remind to the reader that, at the beginning of Section 5, we have tried to give intuitions to facilitate the understanding of this proof.

The proof is by establishing that the two STTS $\mathcal{F}(A, \Delta)$ and $\llbracket A \rrbracket^{\mathsf{AAsap}}_{\delta}$ are mutually similar. Since we have used in $\mathcal{F}(A, \Delta)$ the labels in_l, out_l $(l \in \mathsf{Loc})$ for synchronizing the location changes, we hide those labels in the simulation's proof since they do not appear in the AASAP semantics.

Further, we use the following notations: for a location $\ell = (l_1, \ldots, l_n)$ of a synchronized product $A_1 \times \cdots \times A_n$ of timed automata, we define $\ell(A_1) = l_1$, $\ell(A_2) = l_2, \ldots, \ell(A_n) = l_n$. Let $\hat{l} \in \mathsf{Loc}^i$ be a location of A_i , we define $\ell[A_i := \hat{l}]$ to be the location ℓ' such that $\ell'(A_i) = \hat{l}$ and $\ell'(A_j) = l_j$ for $j \neq i$. Let A be the tuple $(\mathsf{Loc}^1, l_0^1, \mathsf{Var}^1, \mathsf{Lab}^1, \mathsf{Edg}^1)$. To prove that the compositional

construction is correct, we have to show that:

- $(A) \ \llbracket A \rrbracket_{\delta}^{\mathsf{AAsap}} \sqsubseteq \llbracket \mathcal{F}(A, \Delta) \rrbracket_{\delta} \downarrow \{ in_l, out_l \mid l \in \mathsf{Loc}^1 \}$
- $(B) \ \llbracket \mathcal{F}(A,\Delta) \rrbracket_{\delta} \downarrow \{ in_l, out_l \mid l \in \mathsf{Loc}^1 \} \sqsubseteq \llbracket A \rrbracket_{\delta}^{\mathsf{AAsap}}$

Let $\llbracket A \rrbracket_{\delta}^{\mathsf{AAsap}}$ be the STTS $(S^1, \iota^1, \Sigma_{\mathsf{in}}^1, \Sigma_{\mathsf{out}}^1, \Sigma_{\tau}^1, \to^1)$ and $\llbracket \mathcal{F}(A, \Delta) \rrbracket_{\delta} \downarrow \{in_l, out_l \mid l \in \mathsf{Loc}^1\}$ be the STTS $(S^2, \iota^2, \Sigma_{\mathsf{in}}^2, \Sigma_{\mathsf{out}}^2, \Sigma_{\tau}^2, \to^2)$. In the proof, we consider instead of \to^2 the transition relation \Rightarrow^2 defined as follows: let $\Rightarrow^{2'}$ be the meta-transition relation associated to $\llbracket \mathcal{F}(A, \Delta) \rrbracket_{\delta} \downarrow \{in_l, out_l \mid l \in \mathsf{Loc}^1\}$ (in the sense of Definition 15). Then $(s, \sigma, s') \in \Rightarrow^2$ iff either $\sigma \neq \epsilon$ and $(s, \sigma, s') \in \Rightarrow^{2'}$ or $\sigma = \epsilon$ and s = s'.

To establish (A), we use the relation $\mathcal{R} \subseteq S^1 \times S^2$ such that $((l^1, v^1, I^1, d^1), (\ell^2, v^2)) \in \mathcal{R}$ iff:

1. $\ell^2(A^2) = \begin{cases} In_{l^1} & \text{if } d^1 \leq \delta \\ Out_{l^1} & \text{otherwise} \end{cases}$

2. For any
$$\alpha \in \mathsf{Lab}_{\mathsf{in}}^1$$
: $\ell^2(W_\alpha) = \begin{cases} W_0 & \text{if } I^1(\alpha) = \bot \\ W_1 & \text{if } I^1(\alpha) \neq \bot \land I^1(\alpha) \leq \delta \\ W_2 & \text{otherwise} \end{cases}$

3. For any $\alpha \in \mathsf{Lab}_{\mathsf{in}}^1$, $l \in \mathsf{Loc}^1$: $\ell^2(W_\alpha^l) = \begin{cases} U_1 & \text{if } l = l_1 \\ U_0 & \text{otherwise} \end{cases}$

- 4. $\forall x \in \mathsf{Var}^1 : v^2(x) = v^1(x);$
- 5. $v^2(d) = d^1;$
- 6. For every $\alpha \in \mathsf{Lab}_{\mathsf{in}}^1$, if $I^1(\alpha) \neq \bot$ then $v^2(z_\alpha) = I^1(\alpha)$.

Let us show that \mathcal{R} is a weak simulation relation, in the sense of Definition 8.

- It appears obviously that $(\iota^1, \iota^2) \in \mathcal{R}$.
- Assume that $(s_1^1, s_1^2) = ((l_1^1, v_1^1, I_1^1, d_1^1), (\ell_1^2, v_1^2)) \in \mathcal{R}$ and that $(s_1^1, \sigma, s_2^1) \in \to^1$ (with $s_2^1 = (l_2^1, v_2^1, I_2^1, d_2^1)$).

We must prove that the following assertion holds:

(A1) There exists a state $s_2^2 \in S^2$ such that $(s_1^2, \sigma, s_2^2) \in \Rightarrow^2$ and $(s_2^1, s_2^2) \in \mathcal{R}$.

To prove (A1) we work case by case on the different possible types of σ :

- If $\sigma \in \Sigma^1_{\text{in}}$.

By definition of the AASAP semantics, rule (A4.2), since $(s_1^1, \sigma, s_2^1) \in \to^1$ we have two possible cases :

* $I_1^1(\sigma) = \bot$ and $s_2^1 = (l_1^1, v_1^1, I_1^1[\sigma := 0], d_1^1)$. Then, by definition of the relation \mathcal{R} , we know from $(s_1^1, s_1^2) \in \mathcal{R}$ that $\ell_1^2(W_\alpha) = W_0$. Therefore, there is a transition $(s_1^2, \sigma, s_2^2) \in \Rightarrow^2$ with $s_2^2 = (\ell_2^2, v_2^2)$ where $\ell_2^2 = \ell_1^2[W_\sigma := W_1]$ and $v_2^2 = v_1^2[z_\sigma := 0]$. This is because there is an edge $(W_0, W_1, \emptyset, \sigma, \{z_\sigma\})$ in W_σ : the guard is trivially satisfied and the variable z_σ is reset, so that the invariant $z_\sigma \leq \Delta$ of W_1 is satisfied by v_2^2 . Note that W_σ is the only automaton in $\mathcal{F}(A, \Delta)$ synchronizing on σ .

Finally, it is easy to see from the definition of \mathcal{R} that $(s_2^1, s_2^2) \in \mathcal{R}$ (in particular, $v_2^2(z_{\sigma}) = I_2^1(\sigma) = 0$).

- * $I_1^1(\sigma) \neq \bot$ and $s_2^1 = s_1^1$. Then, by definition of the relation \mathcal{R} , we know from $(s_1^1, s_1^2) \in \mathcal{R}$ that $\ell_1^2(W_\alpha) = W_1 \lor \ell_1^2(W_\alpha) = W_2$. In both cases there exists a transition $(s_1^2, \sigma, s_2^2) \in \Rightarrow^2$. Those are the self-loop on the location W_1 and W_2 : $(W_1, W_1, \emptyset, \sigma, \emptyset)$ and $(W_2, W_2, \emptyset, \sigma, \emptyset)$. After any of those transitions we get $s_2^2 = s_1^2$, which implies that $(s_2^1, s_2^2) \in \mathcal{R}$ since $(s_1^1, s_1^2) \in \mathcal{R}$ and $s_1^1 = s_2^1$.
- If $\sigma \in \Sigma^1_{\mathsf{out}}$.

By definition of the AASAP semantics, rule (A4.1), since $(s_1^1, \sigma, s_2^1) \in \to^1$ we know that there exists an edge $(l_1^1, l_2^1, g, \sigma, R) \in \mathsf{Edg}^1$ such that $v_1^1 \models_{\delta} \Delta[g]_{\Delta}$ and $v_2^1 = v_1^1[R := 0]$. Also $I_2^1 = I_1^1$ and $d_2^1 = 0$. Since $(s_1^1, s_1^2) \in \mathcal{R}$, we have either $\ell_1^2(A^2) = In_{l_1^1}$ or $\ell_1^2(A^2) = Out_{l_1^1}$. If $\ell_1^2(A^2) = In_{l_1^1}$, we use the ϵ -edge $(In_{l_1^1}, Out_{l_1^1}, \emptyset, \epsilon, \emptyset)$ of A^2 to get into $Out_{l_1^1}$. This can be achieved since $[\mathsf{Inv}^2(In_{l_1^1})]_{\delta} \subseteq [\mathsf{Inv}^2(Out_{l_1^1})]_{\delta}$. From $Out_{l_1^1}$, we can take the edge $(Out_{l_1^1}, PostOut_{l_1^1, l_2^1}, \Delta[g]_{\Delta}, \sigma, R \cup \{d\})$. This is because the guard is satisfied by v_1^2 (v_1^1 satisfies the guard and agrees with v_1^2 on the variables in Var^1), and d is reset so that the invariant of $PostOut_{l_1^{1,l_2^1}}$ is also satisfied. Then we can go to $PreIn_{l_2^1}$ (invariant d = 0) and $In_{l_2^1}$ (invariant containing $d \leq \Delta$ disjunctively) by synchronizing on $out_{l_1^1}$ and

 $in_{l_2^1}$. This makes the location of the guard-watchers $W_{\alpha}^{l_1^1}$ and $W_{\alpha}^{l_2^1}$ (for each $\alpha \in \mathsf{Lab}_{in}$) change to U_0 and U_1 respectively.

Finally, it is easy to see from the definition of \mathcal{R} that $(s_2^1, s_2^2) \in \mathcal{R}$ (in particular, $v_2^2(d) = d_2^1 = 0$ and for every $x \in R : v_2^2(x) = v_2^1(x) = 0$).

- If $\sigma \in \Sigma^1_{\tau}$. Recall that $\Sigma^1_{\tau} = \mathsf{Lab}^1_{\tau} \cup \overline{\mathsf{Lab}^1_{\mathsf{in}}} \cup \{\epsilon\}$.
 - * If $\sigma \in \mathsf{Lab}^1_{\tau}$, the proof is identical to the previous case.
 - * If $\sigma = \bar{\alpha} \in \mathsf{Lab}^1_{\mathsf{in}}$.

By definition of the AASAP semantics, rule (A4.3), since $(s_1^1, \sigma, s_2^1) \in \rightarrow^1$ we know that there exists an edge $(l_1^1, l_2^1, g, \sigma, R) \in \mathsf{Edg}^1$ such that $v_1^1 \models_{\delta \Delta} [g]_{\Delta}$ and $v_2^1 = v_1^1[R := 0]$. Also $I_1^1(\sigma) \neq \bot$, $I_2^1 = I_1^1[\sigma := \bot]$ and $d_2^1 = 0$.

We show below that there exists a transition $(s_1^2, \sigma, s_2^2) \in \Rightarrow^2$ such that $s_2^2 = (\ell_2^2, v_2^2)$ with $v_2^2 = v_1^2[R \cup \{d\} := 0]$ and ℓ_2^2 is such that $\ell_2^2(A^2) = In_{l_2^1}, \ell_2^2(W_\alpha) = W_0$, and for each $\beta \in \mathsf{Lab}_{in}^1, \ell_2^2(W_\beta^{l_1^1}) = U_0$ and $\ell_2^2(W_\beta^{l_2^1}) = U_1$. Finally, ℓ_2^2 agrees with ℓ_1^2 on location of other automata in $\mathcal{F}(A, \Delta)$. Note that the guard-watchers W_α^l associated with locations $l \neq l_1^1$ are in location U_0 so that they can synchronize on $\bar{\alpha}$ without changing location.

First, we show that A^2 can move: since $(s_1^1, s_1^2) \in \mathcal{R}$, we have either $\ell_1^2(A^2) = In_{l_1^1}$ or $\ell_1^2(A^2) = Out_{l_1^1}$. If $\ell_1^2(A^2) = In_{l_1^1}$, we use the ϵ -edge $(In_{l_1^1}, Out_{l_1^1}, \emptyset, \epsilon, \{d\})$ of A^2 to get into $Out_{l_1^1}$. This can be achieved since $[\![\operatorname{Inv}^2(In_{l_1^1})]\!]_{\delta} \subseteq [\![\operatorname{Inv}^2(Out_{l_1^1})]\!]_{\delta}$. From $Out_{l_1^1}$, we can take the edge $(Out_{l_1^1}, PostOut_{l_1^1, l_2^1}, \Delta[g]_{\Delta}, \bar{\alpha}, R \cup \{d\})$. This is because the guard is satisfied by v_1^2 (v_1^1 satisfies the guard and agrees with v_1^2 on the variables in Var^1), and d is reset so that the invariant of $PostOut_{l_1^1, l_2^1}$ is also satisfied. Then we can go to $PreIn_{l_2^1}$ (invariant d = 0) and $In_{l_2^1}$ (invariant containing $d \leq \Delta$ disjunctively) by synchronizing on $out_{l_1^1}$ and $in_{l_2^1}$.

Second, the synchronizations in A^2 (on $\bar{\alpha}$, $out_{l_1^1}$ and $in_{l_2^1}$) make the location of the guard-watchers $W_{\beta}^{l_1^1}$ and $W_{\beta}^{l_2^1}$ (for each $\beta \in \mathsf{Lab}_{\mathsf{in}}^1$) change to U_0 and U_1 respectively. In particular, for $\beta = \alpha$, the guard-watcher $W_{\beta}^{l_1^1}$ passes through U_3 .

Third, we show that the event-watcher W_{α} can move: since $(s_1^1, s_1^2) \in \mathcal{R}$ and $I_1^1(\sigma) \neq \bot$, we have either $\ell_1^2(W_{\alpha}) = W_1$ or $\ell_1^2(W_{\alpha}) = W_2$. In the case $\ell_1^2(W_{\alpha}) = W_1$, we use the ϵ -edge $(W_1, W_2, \emptyset, \epsilon, \emptyset)$ of W_{α} to get into W_2 which has a trivial invariant. From W_2 , we take the edge $(W_2, W_0, \emptyset, \bar{\alpha}, \emptyset)$, synchronizing with A^2 on $\bar{\alpha}$. Those edges are enabled since their guard and the invariant of their target location are trivial.

Finally, notice that the labels $out_{l_1^1}$ and $in_{l_2^1}$ being hidden, we have two ϵ transition after $\bar{\alpha}$. It is easy to see from the definition of \mathcal{R} that $(s_2^1, s_2^2) \in \mathcal{R}$ (in particular, $v_2^2(d) = d_2^1 = 0$ and for every $x \in R : v_2^2(x) = v_2^1(x) = 0$).

* If $\sigma = \epsilon$. Then $s_1^1 = s_2^1$. This case is trivial since $(s_1^2, \epsilon, s_1^2) \in \Rightarrow^2$. - If $\sigma = t \in \mathbb{R}^{\geq 0}$.

By definition of the AASAP semantics, rule (A4.6), since $(s_1^1, t, s_2^1) \in \to^1$ we have:

(H1) for any edge $(l_1^1, l', g, \sigma, R) \in \mathsf{Edg}^1$ with $\sigma \in \mathsf{Lab}_{\mathsf{out}}^1 \cup \mathsf{Lab}_{\tau}^1$: $\forall t': 0 < t' < t: (d_1^1 + t' \le \delta \lor \mathsf{TS}(v_1^1 + t', g) \le \delta)$

(H2) for any edge
$$(l_1^1, l', g, \sigma, R) \in \mathsf{Edg}^1$$
 with $\sigma \in \mathsf{Lab}_{in}^1$:

$$\forall t': 0 \le t' \le t: d_1^1 + t' \le \delta \lor \mathsf{TS}(v_1^1 + t', g) \le \delta \lor (I_1^1 + t')(\sigma) \le \delta$$

From the definition of \mathcal{R} , we have for any $t' \in \mathbb{R}^{\geq 0}$:

- $(P1) \ d_1^1 + t' \leq \delta \leftrightarrow v_1^2(d) + t' \leq \delta$
- (P2) For any $\alpha \in \mathsf{Lab}_{in}^1$: $(I_1^1 + t')(\alpha) \leq \delta \leftrightarrow v_1^2(z_\alpha) + t' \leq \delta$ (provided $(I_1^1 + t')(\alpha) \neq \bot$)

(P3) For any
$$g \in \mathsf{Rect}_{\mathsf{c}}(\mathsf{Var}^1)$$
: $\mathsf{TS}(v_1^1 + t', g) \leq \delta \leftrightarrow v_1^2 + t' \models_{\delta} \neg(-\Delta(g]_0)$

The third proposition is established by applying Lemma 1(a) since v_1^1 and v_1^2 agree on the variables in Var¹.

We must show that $\mathcal{F}(A, \Delta)$ can let time pass up to t. For this, we show that (i) each automaton in $\mathcal{F}(A, \Delta)$ can let time pass and that (ii) the synchronized product of those automata can avoid staying in urgent locations.

(i) Let us successively consider each automaton in $\mathcal{F}(A, \Delta)$:

- * Event-Watcher W_{α} , for $\alpha \in \mathsf{Lab}_{\mathsf{in}}^1$. If $\ell_1^2(W_{\alpha}) = W_0$ or $\ell_1^2(W_{\alpha}) = W_2$, then time can pass since the invariant is trivial. Assume $\ell_1^2(W_{\alpha}) = W_1$. If $t \leq \delta - v_1^2(z_{\alpha})$, then we can stay waiting in this location. Otherwise, we wait $\delta - v_1^2(z_{\alpha})$ time units in W_1 so that the value of z_{α} is δ . Using an ϵ -edge, we jump to W_2 where the invariant is trivial.
- * Guard-Watchers W_{α}^{l} , for $l \in \mathsf{Loc}^{1}$ and $\alpha \in \mathsf{Lab}_{\mathsf{in}}^{1}$. We have $\ell_{1}^{2}(W_{\alpha}^{l_{1}^{1}}) = U_{1}$ and for $l \neq l_{1}^{1}$, $\ell_{1}^{2}(W_{\alpha}^{l}) = U_{0}$. Both U_{0} and U_{1} have a trivial invariant.
- * A^2 . Assume $\ell_1^2(A^2) = In_{l_1^1}$. If $t \leq \delta v_1^2(d)$, then we can stay waiting in this location. Otherwise, we wait $\delta v_1^2(d)$ time units, and then we

take the ϵ -edge in A^2 to location $Out_{l_1^1}$. This is possible since now the value of d is δ and the expression " $d \leq \Delta$ " appears disjunctively in the invariant of $Out_{l_1^1}$. We claim that in this location, we can let the remaining time $t - (\delta - v_1^2(d))$ pass. This is because (H1) is equivalent to $\forall g \in G_{\mathsf{act}}(l_1^1) \forall t' : 0 \leq t' \leq t : (v_1^2(d) + t' \leq \delta \lor v_1^2 + t' \models_{\delta} \neg (-\Delta(g]_0))$ from the (P1) and (P3). Therefore, we have $v_1^2 + t' \models_{\delta} d \leq \Delta \lor \varphi_{\mathsf{a}}(l_1^1)$, which implies $v_1^2 + t' \models_{\delta} d \leq \Delta \lor \varphi_{\mathsf{a}}(l_1^1)$ that is $v_1^2 + t' \models_{\delta} \mathsf{Inv}^2(Out_{l_1^1})$ for every $0 \leq t' \leq t$. The proof is similar for when $\ell_1^2(A^2) = Out_{l_1^1}$.

(*ii*) It remains to check that we do not let time pass in a urgent location. The urgent locations we could enter by letting time pass in $\mathcal{F}(A, \Delta)$ are the locations ℓ_{α} (for each $\alpha \in \mathsf{Lab}_{\mathsf{in}}^1$) such that $\ell_{\alpha}(A^2) = Out_{l_1^1}$,

$$\ell_{\alpha}(W_{\alpha}) = W_2 \text{ and } \ell_{\alpha}(W_{\alpha}^{\iota_1}) = U_1.$$

Assume that after t' < t time units, we get into such a urgent location ℓ_{α} (for some $\alpha \in \mathsf{Lab}_{\mathsf{in}}^1$). Then, we necessarily have $v_1^2(d) + t' \geq \delta$ and $v_1^2(z_{\alpha}) + t' \geq \delta$. From (P1) and (P2), for every t'' > t' we have $d_1^1 + t'' > \delta$ and $(I_1^1 + t'')(\alpha) > \delta$. Therefore, (H2) where $\sigma = \alpha$ reduces to $\forall g \in G_{\mathsf{evt}}(l_1^1, \alpha) \forall t'' : t' < t'' \leq t : \mathsf{TS}(v_1^1 + t'', g) \leq \delta$, which entails by (P3) that $v_1^2 + t'' \models_{\delta} \neg(_{-\Delta}(g]_0)$ and finally $v_1^2 + t'' \models_{\delta} \varphi_{\mathsf{e}}(l_1^1, \alpha)$ for $t' < t'' \leq t$. Since $\bar{\varphi}_{\mathsf{e}}(l_1^1, \alpha)$ is the closure of $\varphi_{\mathsf{e}}(l_1^1, \alpha)$, we also have $v_1^2 + t' \models_{\delta} \bar{\varphi}_{\mathsf{e}}(l_1^1, \alpha)$. Hence, we can jump to U_2 in $W_{\alpha}^{l_1^1}$ and continue waiting there during the remaining t - t' time units.

Since the number of urgent locations ℓ_{α} is equal to $|\mathsf{Lab}_{\mathsf{in}}^1|$ and we do not get back into a location ℓ_{α} we have left earlier, there will be a finite number of 'jumps to U_2 ' (at most $|\mathsf{Lab}_{\mathsf{in}}^1|$). Hence, by applying a finite number of times the above reasoning we can finally wait the whole t time units.

The reader can easily check that we have $(s_2^1, s_2^2) \in \mathcal{R}$.

To establish (B), we use the relation $\mathcal{R} \subseteq S^2 \times S^1$ such that $((\ell^2, v^2), (l^1, v^1, I^1, d^1)) \in \mathcal{R}$ iff:

1. One of the following assertions holds for every $\alpha \in \mathsf{Lab}_{in}^1$:

(a)
$$\ell^2(A^2) = In_{l^1}, \ \ell^2(W^{l^1}_{\alpha}) \neq U_0 \text{ and } \ell^2(W^{l}_{\alpha}) = U_0 \text{ for any } l \neq l^1;$$

- (b) or $\ell^2(A^2) = Out_{l^1}, \, \ell^2(W^{l^1}_{\alpha}) \neq U_0 \text{ and } \ell^2(W^l_{\alpha}) = U_0 \text{ for any } l \neq l^1.$
- 2. For any $\alpha \in \mathsf{Lab}_{\mathsf{in}}^1$: $I^1(\alpha) = \begin{cases} \bot & \text{if } \ell^2(W_\alpha) = W_0 \\ v^2(z_\alpha) & \text{otherwise} \end{cases}$
- 3. $\forall x \in Var^1 : v^1(x) = v^2(x);$
- 4. $d^1 = v^2(d)$.

Let us show that \mathcal{R} is a weak simulation relation, in the sense of Definition 8.

- It appears obviously that $(\iota^1, \iota^2) \in \mathcal{R}$.
- Assume that $(s_1^2, s_1^1) = ((\ell_1^2, v_1^2), (l_1^1, v_1^1, I_1^1, d_1^1)) \in \mathcal{R}$ and that $(s_1^2, \sigma, s_2^2) \in \Rightarrow^2$ (with $s_2^2 = (\ell_2^2, v_2^2)$).

We must prove that the following assertion holds:

(B1) There exists a state $s_2^1 \in S^1$ such that $(s_1^1, \sigma, s_2^1) \in \to^1$ and $(s_2^2, s_2^1) \in \mathcal{R}$.

To prove (B1) we work case by case on the different possible types of σ :

- If $\sigma \in \Sigma^2_{\text{in}}$.

Since W_{σ} is the only automaton in $\mathcal{F}(A, \Delta)$ that synchronizes on σ , we know from $(s_1^2, \sigma, s_2^2) \in \Rightarrow^2$ that there are two possible cases :

- * $\ell_1^2(W_{\sigma}) = W_0$, $\ell_2^2(W_{\sigma}) = W_1$, and $v_2^2 = v_1^2[z_{\sigma} := 0]$. Then, by definition of \mathcal{R} , we have $I_1^1(\sigma) = \bot$. Therefore, there is a transition $(s_1^1, \sigma, s_2^1) \in \to^1$ with $s_2^1 = (l_1^1, v_1^1, I_1^1[\sigma := 0], d_1^1)$ (from rule (A4.2) of the AASAP semantics). Hence, we have $(s_2^2, s_2^1) \in \mathcal{R}$ (in particular, $I_2^1(\sigma) = v_2^2(z_{\sigma}) = 0$).
- * $\ell_1^2(W_{\sigma}) = W_1 \vee \ell_1^2(W_{\sigma}) = W_2$. In this case, $s_1^2 = s_2^2$. Furthermore, since $(s_1^1, s_1^2) \in \mathcal{R}$, we know that $I_1^1 \neq \bot$, which implies from rule (A4.2) that there exists $(s_1^1, s_2^1) \in \Rightarrow^1$ and that $s_1^1 = s_2^1$. It is then immediate that $(s_2^1, s_2^2) \in \mathcal{R}$.
- If $\sigma \in \Sigma^2_{\mathsf{out}}$.

In this case, A^2 is the only automaton synchronizing on σ . Clearly, there exists an edge $(l_1^1, l_2^1, g, \sigma, R) \in \mathsf{Edg}^1$ such that $v_1^2 \models_{\delta \Delta} [g]_{\Delta}$ and $v_2^2 = v_1^2[R \cup \{d\} := 0]$. Also, $\ell_2^2(A^2) = In_{l_2^1}$.

Then, by definition of \mathcal{R} , $v_1^1 \models_{\delta} \Delta[g]_{\Delta}$, and there is a transition $(s_1^1, \sigma, s_2^1) \in \to^1$ with $s_2^1 = (l_2^1, v_1^1[R := 0], I_1^1, 0)$ (from rule (A4.1) of the AASAP semantics). Hence, we have $(s_2^2, s_2^1) \in \mathcal{R}$ (in particular, $d_2^1 = v_2^2(d) = 0$ and for every $x \in R : v_2^1(x) = v_2^2(x) = 0$).

- If $\sigma \in \Sigma^2_{\tau}$. Recall that $\Sigma^2_{\tau} = \mathsf{Lab}^1_{\tau} \cup \overline{\mathsf{Lab}^1_{\mathsf{in}}} \cup \{\epsilon\}$.
 - * If $\sigma \in \mathsf{Lab}^{1}_{\tau}$, the proof is identical to the previous case.
 - * If $\sigma = \bar{\alpha} \in \overline{\mathsf{Lab}^1_{\mathsf{in}}}$.

The automata in $\mathcal{F}(A, \Delta)$ that are synchronizing on $\bar{\alpha}$ are: A^2 , W_{α} and W_{α}^l for every $l \in \mathsf{Loc}^1$. By definition of the semantics of timed automata and by construction of A^2 , since $(s_1^2, \sigma, s_2^2) \in \Rightarrow^2$ we know that there exists an edge $(l_1^1, l_2^1, g, \bar{\alpha}, R) \in \mathsf{Edg}^1$ such that $v_1^2 \models_{\delta} \Delta[g]_{\Delta}$ and $v_2^2 = v_1^2[R \cup \{d\} := 0]$. Also, $\ell_2^2(A^2) = In_{l_2^1}$. By construction of W_{α} , it must be that $\ell_1^2(W_{\alpha}) = W_2$ and $\ell_2^2(W_{\alpha}) = W_0$. Hence, by definition of \mathcal{R} , we have $v_1^1 \models_{\delta} \Delta[g]_{\Delta}$ and $I_1^1(\alpha) \neq \bot$. Thus there is a transition $(s_1^1, \sigma, s_2^1) \in \rightarrow^1$ with $s_2^1 = (l_2^1, v_1^1[R := 0], I_1^1[\alpha := \bot], 0)$ (from rule (A4.3) of the AASAP semantics) and we have $(s_2^2, s_2^1) \in \mathcal{R}$ (in particular, $d_2^1 = v_2^2(d) = 0$ and for every $x \in R : v_2^1(x) = v_2^2(x) = 0$)..

* If $\sigma = \epsilon$. Then $s_1^2 = s_2^2$. This case is trivial since $(s_1^1, \epsilon, s_1^1) \in \to^1$ (from rule (A4.5) of the AASAP semantics).

- If $\sigma = t \in \mathbb{R}^{\geq 0}$.

Consider the possible cases for the current location ℓ_1^2 :

1. Assume $\ell_1^2(A^2) = In_{l_1^1}$. By definition of the semantics of timed automata, we have: $\forall t' \in [0, t] : v_1^2 + t' \in \llbracket \operatorname{Inv}^2(In_{l_1^1}) \rrbracket_{\delta}$. Therefore, we have $v_1^2(d) + t' \leq \delta$ and hence $d_1^1 + t' \leq \delta$ for any $t' \in [0, t]$. Hence, by definition of the AASAP semantics, there is a transition $(s_1^1, t, s_2^1) \in \to^1$ and $(s_2^2, s_2^1) \in \mathcal{R}$. 2. Assume $\ell_1^2(A^2) = Out_{l_1^1}$

We have to show that the rule (A4.6) of the definition of the AASAP semantics holds.

(*i*) Consider the edges with a label in $\mathsf{Lab}_{\mathsf{out}} \cup \mathsf{Lab}_{\tau}$. We know that $v_1^2 + t' \in \llbracket \mathsf{Inv}^2(Out_{l_1^1}) \rrbracket_{\delta}$ for any $t' \in [0, t]$. Since $\mathsf{Inv}^2(Out_{l_1^1}) = d \leq \delta \vee \bar{\varphi}_{\mathsf{a}}(l_1^1)$, from Lemma 1(*b*) this implies that for any $t' \in [0, t]$, we have either:

 $(E1) \ v_1^2(d) + t' \le \delta$

(E2) or $v_1^2 + t' \models_{\delta} \varphi_{\mathsf{a}}(l_1^1)$

(E3) or $\forall t'' > t' : v_1^2 + t'' \models_{\delta} \varphi_{\mathsf{a}}(l_1^1)$

Below, we show that $(E1) \lor (E2)$ holds for any $t' \in [0, t]$. From the preliminaries (P1) and (P3), this implies that for any edge $(l^1, l', g, \sigma, R) \in \mathsf{Edg}^1$ with $\sigma \in \mathsf{Lab}_{\mathsf{out}} \cup \mathsf{Lab}_{\tau}$ we have that:

 $\forall t': 0 \le t' \le t: (d_1^1 + t' \le \delta \lor \mathsf{TS}(v_1^1 + t', g) \le \delta)$

which completes the proof of (A4.6) for edges with a label in $\mathsf{Lab}_{\mathsf{out}} \cup \mathsf{Lab}_{\tau}$.

To show that $(E1) \lor (E2)$ holds for any $t' \in [0, t]$, let us assume that (E1) and (E2) are falsified for some $\hat{t} \in [0, t]$ and let show that this leads to a contradiction. First, (E1) and (E2) cannot be both falsified for any other $t^* \neq \hat{t}$. Indeed, if it was the case, assuming wlog that $t^* > \hat{t}$, we would have a contradiction from (E3) when $t'' = t^*$ ((E3) saying that $v_1^2 + t^* \models_{\delta} \varphi_{\mathbf{a}}(l_1^1)$ which is (E2) for $t' = t^*$).

It cannot be that $\hat{t} > 0$ since the predicate $p := (E1) \lor (E2)$ is right-closed, by Lemma 2 (with $t_1 = 0$ and $t_2 = \hat{t}$), we would have p satisfied in \hat{t} , a contradiction.

It is also impossible that $\hat{t} = 0$ because of the following: consider the last continuous transition that occurred in $[\![\mathcal{F}(A, \Delta)]\!]_{\delta}$. Since that last continuous transition, there cannot have been a discrete jump in A^2 (since then we would have $v_1^2(d) \leq \delta$, that is (E1) satisfied for $t' = \hat{t} = 0$). Hence, since only A^2 could reset d and the variables appearing in $\varphi_a(l_1^1)$, we must conclude that $(E1) \lor (E2) \lor (E3)$ was in fact true for every $t' \in [-\tilde{t}, t]$ (where $\tilde{t} > 0$ is the duration of the last continuous transition).

Clearly, for $t' \in [-\tilde{t}, 0)$, (E3) cannot be true (if it was true, then (E2) would be true for $t' = \hat{t} = 0$). Therefore, we have for $t' \in [-\tilde{t}, 0)$: (E1) \lor (E2). Since (E1) and (E2) are both right-closed, their disjunction must also hold for t' = 0 (by Lemma 2). This contradicts our assumption.

- (*ii*) Now, consider edges labeled by $\alpha \in \mathsf{Lab}_{in}^1$.
- (a) Assume $\ell_1^2(W_\alpha) = W_0$. By definition of \mathcal{R} , we have $I_1^1(\alpha) = \bot$. Hence, for the edges labeled with α , we have trivially that $\forall t' : 0 \le t' \le t : (d_1^1 + t' \le \delta \lor \mathsf{TS}(v_1^1 + t', g) \le \delta \lor (I_1^1 + t')(\alpha) \le \delta)$. Hence, there is a transition $(s_1^1, t, s_2^1) \in \to^1$ and $(s_2^2, s_2^1) \in \mathcal{R}$.
- (b) Assume $\ell_1^2(W_\alpha) = W_1$. Since $v_1^2 \models_{\delta} (z_\alpha \leq \delta)$ and $I_1^1(\alpha) = v_1^2(z_\alpha)$, we get the same conclusion as in the previous subcase.
- (c) Assume $\ell_1^2(W_\alpha) = W_2$ and $\ell_1^2(W_\alpha^{l^1}) = U_2$. By definition of the semantics of timed automata, the invariant of ℓ_1^2 is satisfied by $v_1^2 + t'$ for every $t' \in [0, t]$. In particular, in U_2 : $v_1^2 + t' \models_{\delta} \bar{\varphi}_{\mathbf{e}}(l_1^1, \alpha)$ for any $t' \in [0, t]$. Therefore, similarly as in the proof

for edges with a label in $\mathsf{Lab}_{\mathsf{out}} \cup \mathsf{Lab}_{\tau}$, we have for any edge $(l^1, l', g, \alpha, R) \in \mathsf{Edg}^1$:

 $\begin{array}{l} \forall t': 0 \leq t' \leq t: (d_1^1 + t' \leq \delta \lor \mathsf{TS}(v_1^1 + t',g) \leq \delta \lor (I_1^1 + t')(\alpha) \leq \delta) \\ \text{Hence, time can pass in the AASAP semantics: there is a transition } (s_1^1,t,s_2^1) \in \rightarrow^1 \text{ and } (s_2^2,s_2^1) \in \mathcal{R}. \end{array}$

(d) Assume $\ell_1^2(W_\alpha) = W_2$ and $\ell_1^2(W_\alpha^{l^1}) = U_1$. Since time can pass in ℓ_1^2 , it cannot be a urgent location. Hence, there is no outgoing edge labeled by $\bar{\alpha}$ in location $Out_{l_1^1}$ and therefore, there is no edge labeled by α in Edg¹. Consequently, the conditions for having a transition $(s_1^1, t, s_2^1) \in \to^1$ are trivially satisfied and $(s_2^2, s_2^1) \in \mathcal{R}$.

The other cases are either impossible or do not allow time passing.

The insightful reader has probably noticed that the use of $\bar{\varphi}_a$ and $\bar{\varphi}_e$ instead of instead of φ_a and φ_e somewhat complicates the proof. The reason for this follows: in Definition 2, invariants of timed automata can be a disjunction of rectangular predicates. It allows to define any boolean combination of rectangular constraint as invariant by converting it into DNF. However, in general the verification tools (like HYTECH or UPPAAL) accept only rectangular predicates as invariant. In practice, we have to split the locations: a location l with invariant $\operatorname{Inv}(l) = \{p_1, \ldots, p_n\}$ is replaced by n locations l_i with invariants $\operatorname{Inv}(l_i) = \{p_i\}$ for $1 \leq i \leq n$. Those locations are connected to each other by empty edges. The initial location and the edges of the automaton are modified as expected.

In the compositional construction $\mathcal{F}(A, \Delta)$, the obvious would be to model the fact that $\mathsf{TS}(v, g) \leq \delta$ by an invariant $\neg(\neg_{\Delta}(g]_0)$ (as in Lemma 1(*a*)). However, the verification tools do not accept disjunctions, and hence we must split the disjunctive invariants. This can be done safely (with conservation of the behaviors) when predicates in disjunctions are all *closed*. Otherwise, correctness is not guaranteed (*e.g.*, splitting the invariant $x < 1 \lor x \geq 1$ is not safe because from the location with invariant x < 1, it is impossible to reach the location with invariant $x \geq 1$ and thus time is blocked, which was not the case in the unsplitted location). To guarantee safe splitting, we use the closure of the expression $\neg(\neg_{\Delta}(g]_0)$ in invariants (that is $\bar{\varphi}_a$ instead of φ_a and $\bar{\varphi}_e$ instead of φ_e).