Partial Order Reduction for Verification of Real-Time Components

John Håkansson¹ and Paul Pettersson^{1,2}

¹ Dept. of Information Technology, Uppsala University, P.O. Box 337,
SE-751 05 Uppsala, Sweden. E-mail: {johnh,paupet}@it.uu.se.

² Dept. of Computer Science and Electronics, Mälardalen University, P.O. Box 883,
SE-721 23, Västerås, Sweden. E-mail: Paul.Pettersson@mdh.se.

Abstract. We describe a partial order reduction technique for a real-time component model. Components are described as timed automata with data ports, which can be composed in static structures of unidirectional control and data flow. Compositions can be encapsulated as components and used in other compositions to form hierarchical models. The proposed partial order reduction technique uses a local time semantics for timed automata, in which time may progress independently in parallel automata which are resynchronized when needed. To increase the number of independent transitions and to reduce the problem of re-synchronizing parallel automata we propose, and show how, to use information derived from the composition structure of an analyzed model. Based on these ideas, we present a reachability analysis algorithm that uses an ample set construction to select which symbolic transitions to explore. The algorithm has been implemented as a prototype extension of the real-time model-checker UPPAAL. We report from experiments with the tool that indicate that the technique can achieve substantial reduction in the time and memory needed to analyze a real-time system described in the studied component model.

1 Introduction

Component-based development has been successfully used for desktop and e-business applications, and it is currently being introduced in many embedded systems. The resource constrained nature of these systems has motivated the development of specific component models [1,13,15,20] and formal frameworks, e.g. [8,9,11].

In general, a component based system is a composition of components, where a component is an open system that accepts input from its environment and produces output. The internal behaviour of a component can be described by a composition, thereby forming a hierarchy of compositions. Components interact with their environment through ports, according to interfaces defined for the ports. Figure 1 shows three components A, B and C. The two components A and B each have an output port, while component C has three input ports and an output port. Components can be composed into more complex functional units

Fig. 1. an example of composition where components A, B and C are composed by connecting port p_1 to p_3 , and p_2 to p_4 .

with well defined interfaces. A *horizontal composition* is a set of components with their ports connected, as in Fig. 1. The connections define how components can interact within a composition. A *vertical composition* is a component with its internal behaviour defined by a horizontal composition.

We use a model for components and composition based on the SAVE³ component model [1,5], and designed for vehicle applications with analysability and safety in mind. The component model is similar to IEC 1131 [15] and Rubus [15]. In our model a component is either idle or executing, and data is transferred from a component when its execution has finished. Some input ports are called trigger ports, and are used to trigger the transition of a component's state from idle to executing. Control flow is specified by means of trigger ports: when one component becomes idle, it can trigger other components so that they become executing. As timeliness is an important property for many embedded systems, we model the execution of components as timed automata [2].

Model checking is an well-established and popular approach for analysis of models, although it is inherently complex and suffers from the so-called state-space explosion problem [12]. Partial order reduction [10,19,17,7] has been suggested as a technique to reduce the state-space explosion caused by parallelism. The idea is to explore representative traces — a property preserving subset of the full model based on independence of transitions. In this paper we present a partial order reduction technique for real-time systems, which is guided by the structure the component based system being analyzed. As in [3,16] we use local time semantics to increase independence. For timed automata the implicit synchronization of global time restricts independence of actions. For our component model we note that the separation of communication from internal computations makes internal transitions independent of actions in other components. We also note that we have extensive information on how components communicate, which is useful for our ample set construction. To increase independence further we relax the synchronization, so that we abstract from the exact time of non-triggering write operations while preserving the order of writes. We present an

³ SAVE is a project supported by Swedish Foundation for Strategic Research. See <http://www.mrtc.mdh.se/SAVE/> for more information.

algorithm for partial order reduction that takes advantage of these ideas, and provide some experimental results from a prototype implementation. The experiments indicate that that our component model is well suited for partial order reduction, and they show good additional reductions when the further relaxed synchronization is used.

Related work includes partial order reduction techniques for timed systems, in particular the local time semantics for timed automata first introduced by Bengtsson et. al., [3]. They also apply the partial order reduction in reachability analysis of timed automata. This work is extended to Timed LTL model-checking by Minea in [16]. We adopt the local time semantics to develop a reachability analysis algorithm for a component model and study how the particular semantics and the static structure of the model can be used to improve previous results. A more recent approach to symbolic model checking of timed automata based on partial order semantics is presented by Lugiez et. al., in [14]. It relies on constraints over event occurrences, instead of clock constraints. In [18], Salah et. al., show that the union of zones reached by interleavings of the same set of transitions is convex. Concurrent semantics for networks of timed automata are investigated in [6,4], by a symbolic unfolding into petri nets with read arcs (to support urgency and invariants).

The rest of this paper is organized as follows: the component model is described in Section 2. In Section 3 we describes our approach to partial order reduction, and in Section 4 we give an algorithm for checking reachability and presents results from an experiment. Section 5 concludes the paper.

2 The Component Model

We introduce timed behaviours to model the execution of components as timed automata, and go on to define syntax and semantics for our component model.

Example 1 (Running Example). Figure 1 shows a horizontal composition of components A, B and C. Assume A is a timer, C a controller, and B a component generating setpoint for the controller. The timer A is connected to the input trigger port p_3 to periodically activate C. The port p_5 is used to read sensor input, which is compared to the setpoint when the controller computes its output to the actuator, port p_6 .

2.1 Timed Behaviour

We define a timed behaviour as a timed automaton, extended with data variables and a final location such that no edges are leading out from this location. For a timed behaviour we have two sets of variables, the set V_C of clock variables, and V_D of data variables. The domain of variables in V_C is the non-negative real numbers $\mathbb{R}_{\geq 0}$, and for variables in V_D the domain is a bounded set of integers INT. We denote by $\mathcal{P}(V_C)$ the power-set of V_C , i.e. the set of all subsets of V_C .

A term t is generated by the grammar $t ::= m | y | t_1 \otimes t_2$, where m is a natural number, $y \in V_D$ is a data variable, $\otimes \in \{+, -, \times, /\}$, and t_1, t_2 are terms. Let U

Fig. 2. timed behaviours: **(a)** a timer with period T and jitter J , **(b)** a computation updating data variable a after between Min and Max time units.

be the set of *variable updates*, each in the form $y := t$ for a data variable $y \in V_D$ and a term t . An atomic clock constraint is of the form $y \sim m$, for $y \in V_C$, m a natural number, and $\sim \in \{<, \leq, =, \geq, >\}$. Similarly, an atomic data constraint is of the form $t_1 \sim t_2$, with terms t_1 and t_2 . We denote by $\text{conj}(V_D, V_C)$ the set of conjunctions of atomic constraints. For $g \in \text{conj}(V_D, V_C)$ we have $g_D \in \text{conj}(V_D)$ the atomic data constraints of g , and $g_C \in \text{conj}(V_C)$ the atomic clock constraints of g .

A *timed behaviour* is a timed automaton $\mathcal{B} = \langle N, l_0, l_f, V_D, V_C, r_0, r_f, E, I \rangle$, where N is a set of locations, l_0 is the initial location, l_f is the final location, V_D and V_C are sets of variables, $r_0 \subseteq V_C$ and $r_f \subseteq V_C$ are sets of clocks (initial and final resets), E is a set of edges so that $E \subseteq N \times \text{conj}(V_D, V_C) \times U \times \mathcal{P}(V_C) \times N$, and I maps each location l in $N \setminus \{l_f\}$ to its invariant $I(l)$, a conjunction of upper bounds on clocks ($y \leq m$ or $y < m$). We write $l \xrightarrow{g, e, r} l'$ iff $\langle l, g, e, r, l' \rangle \in E$ to denote an edge from location l to l' with a guard g , variable update e , and reset clocks $r \subseteq V_C$.

2.2 Component

A component in our setting is defined by its interface and a timed behaviour. The interface of a component consists of data ports and trigger ports, where connected data ports define data flow between components, and connected trigger ports define control flow. The ports are either input or output. An input data port has an associated data variable holding the current data item for this port.

A component is initially *idle*, and it remains in this state until all input trigger ports have been activated, at which point it is *triggered* and switches to the executing state. The internal computation of a component starts with a *read* phase, where all the input data ports are stored internally. The internal copies of input data are used together with internal state variables during the *execute* phase, where the internal behaviour of the component is executed. When the execute phase is over the *write* phase writes output to the output data ports. Finally, the input trigger ports are reset and all outgoing trigger ports are activated, after which the component returns to the idle state.

Each component C is a tuple $\langle P_{\text{in}}, P_{\text{out}}, P_{\text{trig}}, \mathcal{B}, \pi \rangle$, where P_{in} is a set of input ports, P_{out} is a set of output ports, $P_{\text{trig}} \subseteq P_{\text{in}}$ are the trigger input ports, $\mathcal{B}$ is

the timed behaviour, and $\pi : P \mapsto V_D$ is a mapping from ports to variables. We denote by P the set of ports $P_{\text{in}} \cup P_{\text{out}}$ of a component.

Example 2. Component A of Example 1 has no input ports so it is spontaneously triggered. Figure 2 (a) shows the timed behaviour $\mathcal{B}^A$ of A, with period $T = 10$ and jitter $J = 1$. When A is triggered the read phase leads to the initial location l_0 . The automaton uses a clock y to ensure that l_f is reached every T time units, with a non-deterministic offset J . Reaching the final location l_f starts the write phase, after which A is spontaneously triggered again. Figure 2 (b) shows the timed behaviour $\mathcal{B}^B$ of component B, with response time between $\text{Min} = 5$ and $\text{Max} = 20$. The clock z is used to ensure that l_f is reached after between Min and Max time units. The setpoint value a is updated to $1 - a$. The port mapping π^B for B is such that $\pi^B(p_2) = a$, meaning that a is copied to p_2 in the write phase.

Semantics In order to define a state of a component we first introduce clock and data valuations. For a set of clocks V_C a clock valuation is a map $u : V_C \mapsto \mathbb{R}_{\geq 0}$. Similarly, for a set of data variables V_D and ports P a data valuation is a map $v : (V_D \cup P) \mapsto \text{INT}$. Operations on valuations are:

$$\begin{aligned} u' &= [r := 0]u \quad \text{iff} \quad u'(y) = 0 \text{ for clocks } y \in r, \text{ and} \\ &\quad u'(y') = u(y') \text{ for } y' \notin r. \\ v' &= [y := t]v \quad \text{iff} \quad v'(y) = v(t) \text{ for } y, \text{ and} \\ &\quad v'(y') = v(y') \text{ for } y' \neq y. \\ u' &= u \oplus \delta \quad \text{iff} \quad \delta \in \mathbb{R}_{\geq 0} \text{ and } u'(y) = u(y) + \delta \text{ for any clock } y. \end{aligned}$$

We introduce the *idle* location $l_{\perp} \notin N$, and denote by $N_{\perp}$ the set $N \cup \{l_{\perp}\}$. A *state* of a component is a tuple $\langle l, v, u \rangle$, where l is a location in $N_{\perp}$, v is a data valuation, and u is a clock valuation. We introduce values active and inactive for trigger ports, and define a component as *triggered* for a data valuation v , $\text{triggered}(v)$, iff for each $p \in P_{\text{trig}}$ we have $v(p) = \text{active}$. A function $\text{input}(v)$ is used to copy values from input ports to corresponding internal variables, similarly $\text{output}(v)$ copies internal variables to output ports, and $\text{idle}(v)$ inactivates trigger inputs:

$$\begin{aligned} \text{input}(v) &= [y := p \mid p \in P_{\text{in}}, y = \pi(p)]v \\ \text{output}(v) &= [p := y \mid p \in P_{\text{out}}, y = \pi(p)]v \\ \text{idle}(v) &= [p := \text{inactive} \mid p \in P_{\text{trig}}]v \end{aligned}$$

The transition rules for a component C are:

- *delay transition*: $\langle l, v, u \rangle \xrightarrow{\delta} \langle l, v, u \oplus \delta \rangle$ if $\delta \in \mathbb{R}_{\geq 0}$, $u \oplus \delta \models I(l)$, $l \neq l_f$, and if $l = l_{\perp}$ then $\neg \text{triggered}(v)$.
- *internal transition*: $\langle l, v, u \rangle \xrightarrow{\tau} \langle l', v', u' \rangle$ along an edge $l \xrightarrow{g, e, r} l'$ with e in the form $y := t$ if $v \models g_D, u \models g_C, u' \models I(l')$, $v' = [y := t]v$, and $u' = [r := 0]u$.
- *read transition*: $\langle l_{\perp}, v, u \rangle \xrightarrow{r} \langle l_0, \text{input}(v), [r_0 := 0]u \rangle$ if $\text{triggered}(v)$.
- *write transition*: $\langle l_f, v, u \rangle \xrightarrow{w} \langle l_{\perp}, \text{idle}(\text{output}(v)), [r_f := 0]u \rangle$.

We use the following restriction on the internal behaviour of components to avoid spurious local time traces:

Definition 1 (Time Divergence). *We require for a timed behaviour that time diverges, i.e. that there is no time-stop due to invariants, and within any finite time bound only a finite number of transitions can be taken (non-zenoness).*

2.3 Composition

A composition is a set of interconnected components, also known as a horizontal composition. We define a composition as a tuple $\langle P_{\text{in}}^x, P_{\text{out}}^x, \mathcal{C}, \mathcal{X} \rangle$, where P_{in}^x and P_{out}^x are external ports connecting the composition to its environment, $\mathcal{C}$ is a set of components and $\mathcal{X}$ is a set of connections. A connection $x = \langle p, p' \rangle$ connects port $p \in (P_{\text{out}}^i \cup P_{\text{in}}^x)$ to port $p' \in (P_{\text{in}}^j \cup P_{\text{out}}^x)$ for C^i and C^j in $\mathcal{C}$. We do not allow conflicting connections, i.e. connecting output ports of the same component with the same port.

Example 3. The composition of Fig. 1 has an external input port p_5 for sensor input, an external output port p_6 for actuation, three components A, B and C, and two connections $\langle p_1, p_3 \rangle$ and $\langle p_2, p_4 \rangle$.

Semantics A state of a composition is a triple $\langle l, v, u \rangle$, where l is a location vector, v is a data valuation and u is a clock valuation. For a state s we denote by $s[i]$ the state $\langle l[i], v[i], u[i] \rangle$ of a component $C^i \in \mathcal{C}$. The local valuations $v[i]$ and $u[i]$ for a component C^i are such that $v[i](y) = v(y)$ for $y \in (P^i \cup V_D^i)$, and $u[i](y) = u(y)$ for $y \in V_C^i$. In addition to the local valuations, the data valuation v also maps external ports P^x to their values. The transfer of data and triggering introduced by writing to ports Q :

$$\begin{aligned} \text{writedata}(Q, v) &= [p' := p \mid \langle p, p' \rangle \in \mathcal{X}, p \in Q, p' \in P_{\text{out}}^x \cup P_{\text{in}}^j \setminus P_{\text{trig}}^j]v \\ \text{writetrig}(Q, v) &= [p' := \text{active} \mid \langle p, p' \rangle \in \mathcal{X}, p \in Q, p' \in P_{\text{trig}}^j]v \end{aligned}$$

The transition rules for a composition are then:

- *delay transition:* $s \xrightarrow{\delta} s'$ if $s[i] \xrightarrow{\delta} s'[i]$ for each component $C^i \in \mathcal{C}$.
- *internal transition:* $s \xrightarrow{\tau^i} s'$ in the behaviour of C^i if $s[i] \xrightarrow{\tau} s'[i]$, and $s[j] = s'[j]$ for $j \neq i$.
- *read transition:* $s \xrightarrow{r^i} s'$ if $s[i] \xrightarrow{r} s'[i]$ and $s'[j] = s[j]$ for $j \neq i$.
- *write transition:* $s \xrightarrow{w^i} s'$ where either $C^i \in \mathcal{C}$ for internal component C^i writing to ports $Q = P_{\text{out}}^i$ or $i = Q$ for external write to ports $Q \subseteq P_{\text{in}}^x$ if
 - internal state of writer is updated: $s[i] \xrightarrow{w} s_1[i]$ if $C^i \in \mathcal{C}$, $s_1[j] = s[j]$ for $j \neq i$ (for external writes $s_1 = s$), and
 - data or triggering is transferred from ports Q : $s' = \langle l_1, v', u_1 \rangle$ with $v' = \text{writetrig}(Q, \text{writedata}(Q, v_1))$.

2.4 Composite Component

In our component model [1,5] we introduce composite components to support hierarchical composition, by allowing the behaviour of a component to be described by a composition. As any other component, a composite component is defined by its interface (ports) and its timed behaviour. Unlike other components the behaviour of a composite component is described as an internal composition. This is sometimes referred to as vertical composition. Composite components can be constructed from compositions that are time divergent (Definition 1). We also require that internal components have at least one input trigger port, to avoid spontaneous triggering.

For a composite component C , an internal transition is either an internal, read or write transition of some internal component. The read operation of C correspond to a write to the external input ports of the internal composition. Component C can write when all internal components are idle. The port values are already updated by internal writes, so the write operation only need to inactivate input trigger ports.

3 Partial Order Reduction

The idea of partial order reduction is to explore representative sequences of independent transitions, instead of examining all possible sequences. However, the implicit synchronization of global time restricts independence for transitions of timed automata. As in [3,16] we use local time semantics to increase independence. It essentially allow us to analyse components of a composition in isolation, and then synchronize the components to a shared state whenever one writes to the others. To increase independence further than [3,16] we relax the synchronization, so that we abstract from the exact time of non-triggering write operations.

3.1 Representatives and Local Time Traces

To describe the concept of representative traces, we first need a notion of independent transitions. Two transitions are considered independent if they can be reordered within a trace without affecting the final state of the trace, or the validity of the trace. We denote by $\text{enabled}(\sigma)$ the set of transitions that can immediately follow a finite trace σ , and define independent transitions as in [16]:

Definition 2. *Two transitions α_1 and α_2 are independent iff for any trace σ such that $\alpha_1, \alpha_2 \in \text{enabled}(\sigma)$:*

- Enabledness: $\alpha_2 \in \text{enabled}(\sigma\alpha_1)$ and $\alpha_1 \in \text{enabled}(\sigma\alpha_2)$.
- Commutativity: *Any state reachable by the trace $\sigma\alpha_1\alpha_2$ can also be reached by the trace $\sigma\alpha_2\alpha_1$.*

Independence is a sufficient condition for reordering transitions within a trace so that the same state is reached, however it is not a sufficient condition for reordering of transitions during reachability analysis. There could for example be a transition α_3 in $\text{enabled}(\sigma\alpha_1)$ which is not in $\text{enabled}(\sigma\alpha_2)$, which we would miss if we only considered the trace $\sigma\alpha_2\alpha_1$. For analysis of a composition we need a strategy to make sure that some representative trace is explored for each possible trace of the full state graph. We examine conditions for reordering further in Sect. 4.1.

Independence can be concluded from the structure of a composition, using a set $\text{active}(\alpha)$ of components that participate in a transition α such that $\text{active}(\alpha^i) = \{C^i\}$ for $\alpha \in \{\tau, \delta, r\}$ and $\text{active}(w^i) = \{C^i\} \cup \{C^j \mid \langle p, p' \rangle \in \mathcal{X}, p \in P_{\text{out}}^i, p' \in P_{\text{in}}^j\}$. We then restate a theorem found in e.g. [3,16], i.e. that two local time actions are independent if no automata participate in both actions:

Theorem 1. $\text{active}(\alpha_1) \cap \text{active}(\alpha_2) = \emptyset \Rightarrow \text{independent}(\alpha_1, \alpha_2)$

Proof. See [16]. □

We reduce the independence relation further (for internal transitions) in Sect. 3.2, where we introduce local time semantics. We say that two transitions α_1 and α_2 are dependent, and write $\text{dependent}(\alpha_1, \alpha_2)$, whenever $\text{independent}(\alpha_1, \alpha_2)$ cannot be concluded from the structure of a composition. The dependency relation is thus a safe approximation of transitions that are not independent.

We define a local time trace to be a representative of some timed trace. A trace σ is a representative of a trace σ' iff independent transitions of σ can be reordered to construct σ' . A timed trace (as defined in [2]) is a pair $\langle \sigma, t \rangle$ such that $\sigma(i)$ is the i th transition of the trace, $t(i)$ is the time of this transition, and the timepoints $t(i)$ are monotonic, i.e. $i \leq j$ implies $t(i) \leq t(j)$. Local time traces are then defined as:

Definition 3. *A local time trace is a trace $\langle \sigma, t \rangle$ such that dependent transitions are monotonic, i.e. for any $\sigma(i)$ and $\sigma(j)$ that are dependent we have $i \leq j$ implies $t(i) \leq t(j)$.*

3.2 Local Time Semantics

To keep track of the local time within components we introduce a reference clock $c^j \in V_C^i$ for each component $C^i \in \mathcal{C}$. We define a local delay transition for component C^i , where other components do not need to delay correspondingly.

We relax the synchronization of our local time semantics compared to [3] by completing the internal computation of a component before synchronizing with other components. This can be done since values written to input ports of a component are not used during its internal computation. We run internal computations ahead implicitly in the rule for a write operation w^i of C^i by requiring that any C^j such that $\text{dependent}(w^i, w^j)$ is either idle ($l[j] = l_\perp^j$) or finished ($l[j] = l_f^j$).

Time is synchronized so that the local time in all components dependent on a write operation is ensured to be later than the local time of the writer, i.e. $c^i \leq c^j$ for a writer C^i and any C^j such that $\text{dependent}(w^i, w^j)$. This is to make traces of the local time semantics be *local time traces*, according to Definition 3. The set of components that are triggered by a write, given a location vector l and data valuation v , is:

$$\text{trig}(l, v) = \{C^j \mid l[j] = l_{\perp}^j \wedge \text{triggered}^j(v[j])\}$$

The clock synchronization constraint $\text{sync}^i(l, v)$ also preserves the exact time of triggering by requiring that $c^i = c^j$ for triggered components C^j :

$$\text{sync}^i(l, v) = \left(\bigwedge_{C^j \in \text{trig}(l, v)} c^i = c^j \right) \wedge \left(\bigwedge_{\text{dependent}(w^i, w^j)} c^i \leq c^j \right)$$

We use compositions to describe component based systems. We define local time transitions $s \longrightarrow_t s'$ for a composition using the corresponding rules for transitions $s \longrightarrow s'$:

- *delay transition*: $s \xrightarrow{\delta^i}_t s'$ if $s[i] \xrightarrow{\delta} s'[i]$ and $s'[j] = s[j]$ for $j \neq i$.
- *internal transition*: $s \xrightarrow{\tau^i}_t s'$ if $s \xrightarrow{\tau^i} s'$.
- *read transition*: $s \xrightarrow{r^i}_t s'$ if $s \xrightarrow{r^i} s'$.
- *write transition*: $s \xrightarrow{w^i}_t s'$ for $s = \langle l, v, u \rangle$ if $s \xrightarrow{w^i} s'$, $u \models \text{sync}^i(l, v)$, and either $l[j] = l_{\perp}^j$ or $l[j] = l_f^j$ for C^j such that $\text{dependent}(w^i, w^j)$.

Lemma 1 shows that internal transitions in the local time semantics are independent of write transitions. This independence makes our model suited for partial order reduction, and enables the weak synchronization (otherwise we would need $c^i = c^j$ also for non-triggering participants, instead we use $c^i \leq c^j$ to preserve order of write operations).

Lemma 1. *For the local time semantics we have:*

- $\text{independent}(\alpha_1^i, \alpha_2^j)$ if $i \neq j$ and α_1, α_2 in $\{\tau, \delta\}$
- $\text{independent}(w^i, \alpha^j)$ if $i \neq j$ and α in $\{\tau, \delta\}$
- $\text{independent}(w^i, r^j)$ if $C^j \notin \text{active}(w^i)$
- $\text{independent}(w^i, w^j)$ if $\text{active}(w^i) \cap \text{active}(w^j) = \emptyset$

Proof. by Theorem 1, and for $\text{independent}(w^i, \tau^j)$ we note that w^i is not enabled if τ^j is enabled for $C^j \in \text{active}(w^i)$, similarly for δ^j . $\square$

Theorem 2 (Correctness of Local Time Semantics). *Assume local time states s_0 and $s_f = \langle l_f, v_f, u_f \rangle$, global time states s'_0 and $s'_f = \langle l'_f, v'_f, u'_f \rangle$, and a component C^k . Let $s_0 = s'_0$ except for reference clocks which are zero in s_0 but not included in s'_0 .*

- (Soundness) whenever $s_0 \xrightarrow{*}_t s_f$ then $s'_0 \xrightarrow{*}_t s'_f$ so that $l_f[k] = l'_f[k]$, $v_f[k](y) = v'_f[k](y)$ for $y \in V_D^k$ (i.e. not for $y \in P^k$), and $u_f[k] = u'_f[k]$.
- (Completeness) whenever $s'_0 \xrightarrow{*}_t s'_f$ then $s_0 \xrightarrow{*}_t s_f$ so that $l_f[k] = l'_f[k]$, $v_f[k](y) = v'_f[k](y)$ for $y \in V_D^k$ (i.e. not for $y \in P^k$), and $u_f[k] = u'_f[k]$.

Proof. Soundness is shown by induction over a local time trace, and completeness by construction of the corresponding local time trace (see Appendix A.1). $\square$

3.3 Symbolic Local Time Semantics

We define a zone Z as a set of clock valuations, and use symbolic states $\langle l, v, Z \rangle$ to represent all states $\langle l, v, u \rangle$ such that $u \in Z$. For a zone Z and a clock constraint g we define *conjunction* $Z \wedge g$ as the set of valuations $u \in Z$ such that $u \models g$, *reset* $r(Z)$ as u' such that $u' = [r := 0]u$ for $u \in Z$, and *symbolic local delay* $Z^{\uparrow i}$ as u' such that for $u \in Z$, $\delta \in \mathbb{R}_{\geq 0}$ we have $u'[i] = u[i] \oplus \delta$ and $u'[j] = u[j]$ for $i \neq j$.

The initial symbolic state is $\langle l_0, v_0, Z_0 \rangle$, where $Z_0 = (\{u_0\})^{\uparrow} \wedge I(l_0)$ incorporates an initial delay in all components from a zone with a single solution u_0 . To improve the presentation we incorporate the semantics of a component into the transition rules for a composition:

- *internal transition*: $\langle l, v, Z \rangle \xrightarrow{\tau^i}_t \langle [l^{i'}/l^i]l, v', Z' \rangle$ along an edge $l^i \xrightarrow{g,e,r} l^{i'}$ with e in the form $y := t$ if $v \models g_D$, $v' = [y := t]v$, and if $l^{i'} = l_f^i$ then $Z' = r(Z \wedge g_C)$, otherwise $Z' = (r(Z \wedge g_C))^{\uparrow i} \wedge I^i(l^{i'})$.
- *read transition*: $\langle l, v, Z \rangle \xrightarrow{r^i}_t \langle [l_0^i/l_{\perp}^i]l, v', Z' \rangle$ if $l[i] = l_{\perp}^i$ and triggered ^{i} ($v[i]$), with $v'[i] = \text{input}(v[i])$, $v'[j] = v[j]$ for $j \neq i$, and if $l_0^i = l_f^i$ then $Z' = r_0^i(Z)$, otherwise $Z' = (r_0^i(Z))^{\uparrow i} \wedge I^i(l_0^i)$.
- *write transition*: $\langle l, v, Z \rangle \xrightarrow{w^i}_t \langle [l_{\perp}^i/l_f^i]l, v', Z' \rangle$ where either $C^i \in \mathcal{C}$ writing to ports $Q = P_{\text{out}}^i$, or $i = Q$ for external write to $Q \subseteq P_{\text{in}}^x$, if $l[i] = l_f^i$ and:
 - $l[j] = l_{\perp}^j$ or $l[j] = l_f^j$ for C^j such that dependent(w^i, w^j),
 - $v_1[i] = \text{idle}(\text{output}(v[i]))$ if $C^i \in \mathcal{C}$, $v_1[j] = v[j]$ for $j \neq i$,
 - $v' = \text{writetrig}(Q, \text{writedata}(Q, v_1))$, and
 - if triggered ^{i} ($v'[i]$) then $Z' = r_f^i(Z \wedge \text{sync}^i(l, v))$, otherwise $Z' = (r_f^i(Z \wedge \text{sync}^i(l, v)))^{\uparrow i}$.

For global time semantics a zone can be represented as a conjunction of clock difference constraints. Constraints on two clocks are preserved by global time delay because both clocks grow equally, but for local time we need a different zone representation that is preserved by local time delay. Local time zones can be efficiently represented [3,16] as difference constraints on reference clocks c^i and timestamps t_y for the latest reset of a clock y .

Example 4. We explore a symbolic trace of the composition in Fig. 1. In the initial state $\langle [l_{\perp}^A, l_{\perp}^B, l_{\perp}^C], v_0, Z_0 \rangle$ both A and B can read, since they have no trigger

input ports (and so all their triggers are trivially active). If A reads first (r^A) we get to a state $\langle [l_0^A, l_\perp^B, l_\perp^C], v_0, Z_1 \rangle$ with $Z_1 = Z_0^{\uparrow A} \wedge (y \leq 10)$. We continue the trace with r^B to $\langle [l_0^A, l_0^B, l_\perp^C], v_0, Z_2 \rangle$ with $Z_2 = r_z(Z_1)^{\uparrow B} \wedge (z \leq 20)$ for $r_z = \{z\}$. From this state internal transitions τ_A and τ_B are enabled. By τ^B we get to a state $\langle [l_0^A, l_f^B, l_\perp^C], v_3, Z_3 \rangle$ with $v_3 = [a := 1]v_0$ and $Z_3 = Z_2 \wedge (z \geq 5)$. Component B cannot write from this state, because w^B depends on w^A and A is neither idle or in its final location, so we take τ^A to the state $\langle [l_1^A, l_f^B, l_\perp^C], v_3, Z_4 \rangle$ with $Z_4 = r_y(Z_3 \wedge y = 10)^{\uparrow A} \wedge y \leq 1$ for $r_y = \{y\}$. Another τ^A leads to $\langle [l_f^A, l_f^B, l_\perp^C], v_3, Z_4 \rangle$. From this state both A and B can write, A when $10 \leq c^A \leq 11$ and B when $5 \leq c^B \leq 20$, but it cannot be determined if C will get data from B before being triggered by A.

Theorem 3 (Correctness of Symbolic Local Time Semantics). *Assume location vectors l_0 , l_f , variable valuations v_0 , v_f , clock valuations u_0 , u_f , and local time zones Z_0 , Z_f .*

- (Soundness) *whenever $\langle l_0, v_0, Z_0 \rangle \Longrightarrow_t^* \langle l_f, v_f, Z_f \rangle$ then for any $u_f \in Z_f$ $\langle l_0, v_0, u_0 \rangle \longrightarrow_t^* \langle l_f, v_f, u_f \rangle$.*
- (Completeness) *whenever $\langle l_0, v_0, u_0 \rangle \longrightarrow_t^* \langle l_f, v_f, u_f \rangle$ then $\langle l_0, v_0, Z_0 \rangle \Longrightarrow_t^* \langle l_f, v_f, Z_f \rangle$ so that $u_f \in Z_f$.*

Proof. Symbolic transition rules are constructed from local time semantics and definitions of zone operations. Preservation of zone representation is shown in [16].

(See Appendix A.2) □

4 Reachability Analysis

We perform reachability analysis by exploring a subset of enabled transitions from each explored state, in order to reach a target location denoted $l_\odot^k$ for a component C^k . In the analysis we use the symbolic local time semantics, to get the independence introduced in our local time semantics and to get a finite state space.

4.1 The Ample Set Method

An ample set [17] is a subset of the enabled transitions that is sufficient to explore when model checking. The ample set method reduces a state graph G to a subgraph R such that correctness of model checking is preserved, i.e. checking R gives the same result as checking G . In general we need to select an ample set so that the checked property is preserved, i.e. a property holds in a representative trace σ in R iff the same property holds in all traces σ' in G represented by σ , and so that all traces in G have a representative in R .

For local reachability we note first that we only need to consider traces of G that can actually reach the target location. We also note that local reachability

in a component C^k is preserved by representative traces, since the reordering of independent transitions does not affect which local states are reachable. The following describes how we can construct an ample set:

Definition 4 (Ample Set Construction). *From the static structure of a composition we compute $\mathcal{P}^k$ the enabled transitions that can give progress in C^k , i.e. transitions of C^k or of some C^j producing data or triggering to C^k (possibly via other components). From $\mathcal{P}^k$ we define the ample set as follows:*

- *The ample set is empty iff $\mathcal{P}^k$ is empty, otherwise some $\alpha_0 \in \mathcal{P}^k$ is in the ample set. This must hold for each valuation of a zone, as discussed in [16].*
- *If α is in the ample set, α' is enabled, and $\text{dependent}(\alpha, \alpha')$ then α' must also be in the ample set.*
- *If α is in the ample set, α' is not enabled, $\text{dependent}(\alpha, \alpha')$, and there is a transition α'' that can lead to $\alpha' \in \mathcal{P}^k$ before α is taken, then α'' must also be in the ample set. This can be determined from the static structure of a composition.*
- *For a local cycle, at least one transition in $\mathcal{P}^k$ and outside of the cycle is in the ample set.*

Theorem 4. *For any trace σ in G reaching the target location, the subgraph R induced by the ample set construction in Definition 4 contains a representative of an extension $\sigma\rho$ in G .*

Proof. By a construction similar to that in [7], with simplifications for local reachability, and cycle closing due partly to construction of $\mathcal{P}^k$ (see Appendix A.3). $\square$

As mentioned in [17] constructing an optimal ample set with respect to state space reduction is NP hard, so we suggest a heuristic. The construction in Definition 4 starts from a transition α_0 in $\mathcal{P}^k$, according to the first rule. Once a transition has been selected the other rules are used to find the least fixpoint, which is an ample set. To reduce the size of the ample set we select α_0 as an internal transition τ^i , if possible. Otherwise select α^i with minimal upper bound on the reference clock of component C^i . This reduces the possibility for other components to interfere with the execution of α^i . We also prefer read operations over writes when selecting α_0 .

4.2 Model Checking Algorithm

An algorithm for symbolic reachability analysis based on the symbolic local time semantics and ample set construction is shown in Fig. 3. It is a standard reachability algorithm, with the exception that the *ample* transitions are explored instead of all enabled transitions. The normalisation $\text{max}(Z)$ is used to ensure termination, as the symbolic semantics is not finite. In [3] it is shown that there is a finite partitioning of the state space, and [16] suggests a method for constructing $\text{max}(Z)$. Efficient representations of local time zones are also discussed in [3,16].

```

PASSED :=  $\emptyset$ 
WAITING :=  $\{\langle l_0, v_0, Z_0 \rangle\}$ 
repeat for some  $\langle l, v, Z \rangle \in \text{WAITING}$ 
  WAITING := WAITING  $\setminus \{\langle l, v, Z \rangle\}$ 
  if  $l[k] = l_{\odot}^k$  then return “REACHABLE”
  else if  $\max(Z) \not\subseteq Z'$  for all  $\langle l, v, Z' \rangle \in \text{PASSED}$  then
    PASSED := PASSED  $\cup \langle l, v, \max(Z) \rangle$ 
    SUCC :=  $\{\langle l', v', Z' \rangle \mid \langle l, v, Z \rangle \xrightarrow{\alpha}_t \langle l', v', Z' \rangle, \alpha \in \text{ample}(l, v, Z)\}$ 
    WAITING := WAITING  $\cup \text{SUCC}$ 
until WAITING =  $\emptyset$ 
return “NOT REACHABLE”

```

Fig. 3. an algorithm for symbolic reachability analysis, exploring a selected subset of enabled transitions from each state until component C^k reaches $l_{\odot}^k$.

We expect the algorithm to perform well: the partial order reduction is a subset of the symbolic local time semantics, by exploring only the ample set: a subset of enabled transitions. The local time semantics is sound (Theorem 2) with respect to the global time semantics. The partitioning of the state space induced by $\max(Z)$ is however incomparable with the normalisation of global time zones, so we cannot conclude any strict improvements. For timed automata with local time semantics [3,16] soundness is shown only for synchronized states, and in general for a local time state there might not exist a corresponding synchronized state. This means that some local time traces lead outside the global time semantics, giving a larger state graph to search. The components of our model are time divergent, can always accept input, and never require input to be available. Because of this components can always catch up, which is why we can show soundness for local states.

Example 5. A symbolic trace is described in Example 4. The local traces for A and B are $r^A \tau^A \tau^A$ and $r^B \tau^B$, respectively. For global time semantics there are six possible interleavings, although probably more interesting nine states are passed. When searching for a location in C using the ample set construction in Definition 4 we select either r^A and r^B first. The selected component will reach its final location before the other components read operation is selected. Even though we have not specified which transition to select only one trace is explored, because r^A and r^B are independent. The trace explored using the ample set construction passes through six states, instead of the nine for global time semantics.

4.3 Experimental Results

We have developed a prototype implementation of our method as an extension of the UPPAAL tool⁴. Figure 4 illustrates the synthetic benchmarks we use to

⁴ See the web site www.uppaal.com for more information.

Fig. 4. the synthetic benchmark 5x3 (with N=5 columns, M=3 rows).

Method	3x2	3x3	3x4	4x3	4x4
global	264/0.16s	499/0.22s	838/0.31s	2553/0.65s	4778/1.6s
strict	19/0.14s	105/0.19s	443/0.47s	105/0.26s	990/1.7s
local	19/0.14s	65/0.18s	275/0.43s	93/0.25s	336/1.0s
	4x5	5x4	5x5	6x5	6x6
global	8146/3.8s	26108/12s	48096/36s	481318/11m16s	$\perp$
strict	7549/21s	990/3.5s	15892/33s	15892/2m59s	$\perp$
local	2380/10s	598/2.8s	2156/20s	4316/1m08s	15101/7m36s

Table 1. benchmark results: we use $\perp$ to denote that the experiment did not terminate in 30 minutes.

evaluate our implementation (similar to the benchmarks of Salah et.al. [18]). A synthetic benchmark $N \times M$ is a grid of components in N columns and M rows. Each component $C_{n,m}$ is connected to $C_{n,m+1}$ and to $C_{n+1,m+1}$. The components in the first row are triggered once, and the timed behaviour of each component is a delay by at least 4 time units. As target component C^k for our ample set construction we use $C_{N,M}$ (i.e. $C_{5,3}$ for Fig. 4). Table 1 shows the computation time and number of explored states for *global* time semantics (Section 2), local time with *strict* synchronization (using constraints $c^i = c^j$ also for w^i, w^j dependent) and *local* time semantics (Section 3). The prototype does not implement the normalisation step $\max(Z)$. Normalisation is not required for the benchmark models, and is turned off in the global time implementation. We note that the algorithms with partial order reduction performs much better than the algorithm without partial order reduction, and that weak synchronization performs better than the strict. We also note that *global* cover more states per second, this is because of the overheads when synchronizing components and constructing ample sets.

5 Conclusion

In this paper, we have applied and improved existing symbolic partial order reduction techniques for timed automata to develop an efficient model-checking

technique for real-time components. The behavior of components are internally described as timed automata that can be hierarchical in the sense that a component can be described as a composition of components. To compose components explicit data and control flow is modeled, a property that is exploited in order to increase the independence between components, and thus to reduce the growth of the state-space caused by interleavings. We give a concrete and symbolic local time semantics for the component model, as well as a symbolic reachability analysis algorithm that uses an ample set construction to select symbolic transitions to explore. We also describe a heuristics that can be used for accelerating the analysis of local reachability properties (e.g., reachability of a location in a single component).

Our experiments suggest that our technique can achieve substantial reduction in the time and space needed to analyze a real-time system described in the studied component model. As future work we plan to further evaluate the reduction in a case study for the component model. We will also complete our implementation of the proposed reachability analysis and evaluate the achieved reduction with respect to existing techniques, such as the *event zones* of [14]. We also plan to further enrich the component model with more complex interaction structures, and support for modeling of other non-functional properties than real-time.

References

1. M. Åkerholm, J. Carlson, J. Fredriksson, H. Hansson, J. Håkansson, A. Möller, P. Pettersson, and M. Tivoli. The SAVE approach to component-based development of vehicular systems. *Journal of Systems and Software*, 2006.
2. R. Alur and D. L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
3. J. Bengtsson, B. Jonsson, J. Lilius, and W. Yi. Partial order reductions for timed systems. In *Proceedings, Ninth International Conference on Concurrency Theory*, volume 1466 of *Lecture Notes in Computer Science*, pages 485–500. Springer–Verlag, 1998.
4. P. Bouyer, S. Haddad, and P.-A. Reynier. Timed unfoldings for networks of timed automata. In S. Graf and W. Zhang, editors, *Proc. of the 4th International Symposium on Automated Technology for Verification and Analysis, ATVA 2006*, number 4218 in *Lecture Notes in Computer Science*, pages 292–306. Springer–Verlag, 2006.
5. J. Carlson, J. Håkansson, and P. Pettersson. SaveCCM: An analysable component model for real-time systems. In *Proceedings of the 2nd Workshop on Formal Aspects of Components Software (FACS 2005)*, *Electronic Notes in Theoretical Computer Science*. Elsevier, 2005.
6. F. Cassez, T. Chatain, and C. Jard. Symbolic unfoldings for networks of timed automata. In S. Graf and W. Zhang, editors, *Proc. of the 4th International Symposium on Automated Technology for Verification and Analysis, ATVA 2006*, number 4218 in *Lecture Notes in Computer Science*, pages 307–321. Springer–Verlag, 2006.
7. E. M. Clarke, O. Grumberg, M. Minea, and D. Peled. State space reduction using partial order techniques. *International Journal on Software Tools for Technology Transfer*, 2(3):279–287, November 1999.

8. L. de Alfaro and T. A. Henzinger. Interface automata. In *Proceedings of the Ninth Annual Symposium on Foundations of Software Engineering*, pages 109–120. ACM Press, 2001.
9. L. de Alfaro, T. A. Henzinger, and M. Stoelinga. Timed interfaces. In *Proceedings of the Second International Workshop on Embedded Software*, Lecture Notes in Computer Science. Springer-Verlag, 2002.
10. P. Godefroid and P. Wolper. Using partial orders for the efficient verification of deadlock freedom and safety properties. In *Proceedings, 3rd International Conference on Computer Aided Verification*, volume 575 of *Lecture Notes in Computer Science*, pages 332–342. Springer-Verlag, 1991.
11. Gregor Gössler and Joseph Sifakis. Composition for component-based modelling. *Science of Computer Programming*, 55(1-3):161–183, 2005.
12. Gerard Holzmann. *The Design and Validation of Computer Protocols*. Prentice Hall, 1991.
13. IEC. International standard IEC 1131: Programmable controllers, 1992.
14. D. Lugiez, P. Niebert, and S. Zennou. A partial order semantics approach to the clock explosion problem of timed automata. *Theoretical Computer Science*, 345(1):27–59, 2005.
15. Kurt-Lennart Lundbäck, John Lundbäck, and Mats Lindberg. Development of dependable real-time applications. Arcticus Systems, December 2004.
16. Marius Minea. Partial order reduction for model checking of timed automata. In *Proceedings, 10th International Conference on Concurrency Theory*, volume 1664 of *Lecture Notes in Computer Science*, pages 431–446. Springer-Verlag, 1999.
17. D. Peled. All from one, one for all: on model checking using representatives. In *Proceedings, 5th International Conference on Computer Aided Verification*, volume 697 of *Lecture Notes in Computer Science*, pages 409–423. Springer-Verlag, 1993.
18. R. B. Salah, M. Bozga, and O. Maler. On interleaving in timed automata. In *Proceedings, 17th International Conference on Concurrency Theory*, volume 4137 of *Lecture Notes in Computer Science*, pages 465–476. Springer-Verlag, 2006.
19. Antti Valmari. A stubborn attack on state explosion. In *Proceedings, Advances in Petri Nets*, volume 483 of *Lecture Notes in Computer Science*. Springer-Verlag, 1990.
20. Rob van Ommering, Frank van der Linden, Jeff Kramer, and Jeff Magee. The Koala component model for consumer electronics software. *IEEE Computer*, pages 78–85, 2000.

A Proofs

A.1 Correctness of Local Time Semantics

According to Theorem 2: assume local time states s_0 and $s_f = \langle l_f, v_f, u_f \rangle$, global time states s'_0 and $s'_f = \langle l'_f, v'_f, u'_f \rangle$, and a component C^k . Let $s_0 = s'_0$ except for reference clocks which are zero in s_0 but not included in s'_0 .

- (*Soundness*) whenever $s_0 \xrightarrow{*}_t s_f$ then $s'_0 \xrightarrow{*} s'_f$ so that $l_f[k] = l'_f[k]$, $v_f[k](y) = v'_f[k](y)$ for $y \in V_D^k$ (i.e. not for $y \in P^k$), and $u_f[k] = u'_f[k]$.
- (*Completeness*) whenever $s'_0 \xrightarrow{*} s'_f$ then $s_0 \xrightarrow{*}_t s_f$ so that $l_f[k] = l'_f[k]$, $v_f[k](y) = v'_f[k](y)$ for $y \in V_D^k$ (i.e. not for $y \in P^k$), and $u_f[k] = u'_f[k]$.

Notation We define $\sigma[i]$ as the part of a trace σ involving component C^i :

- $\epsilon[i] = \epsilon$ (where ϵ is the empty trace),
- $(\sigma\alpha)[i] = \sigma[i]$ if $C^i \notin \text{active}(\alpha)$, and
- $(\sigma\alpha)[i] = (\sigma[i])\alpha$ otherwise.

We can reconstruct global time traces by combining traces $\sigma_1 \parallel \sigma_2$:

- $(\sigma_1 \parallel \epsilon) = \sigma_1$ and $(\sigma_1 \parallel \sigma_2) = (\sigma_2 \parallel \sigma_1)$,
- $(\sigma_1\alpha \parallel \sigma_2\alpha) = (\sigma_1 \parallel \sigma_2)\alpha$ when both traces participate in α , and
- $(\sigma_1\alpha \parallel \sigma_2) = (\sigma_1 \parallel \sigma_2)\alpha$ when α does not involve components in both traces, and α occurs after (non-strict) the last transition in σ_2 .

Note that the construction is non-deterministic due to non-strict interpretation of *after*, and that (σ, t) is a possible construction from $(\sigma[1] \parallel \sigma[2], t)$ for a two component system: the time-points of events are not altered. A similar construction for local time traces, $\sigma_1 \parallel_t \sigma_2$:

- $(\sigma_1 \parallel_t \epsilon) = \sigma_1$ and $(\sigma_1 \parallel_t \sigma_2) = (\sigma_2 \parallel_t \sigma_1)$,
- $(\sigma_1\alpha \parallel_t \sigma_2\alpha) = (\sigma_1 \parallel_t \sigma_2)\alpha$ when both traces participate in α , and
- $(\sigma_1\alpha \parallel_t \sigma_2) = (\sigma_1 \parallel_t \sigma_2)\alpha$ when α does not involve components in both traces, and α occurs after (non-strict) the last transition in σ_2 *dependent* on α (holds trivially for internal transitions).

Proof For soundness we assume a set of components $\mathcal{K}$ such that for all $C^i, C^j \in \mathcal{K}$ we have $u_f(c^i) = u_f(c^j)$, implying the case $\mathcal{K} = \{k\}$, and continue to show soundness by induction over number of write actions in the local time trace (σ, t) leading to $\langle l_f, v_f, u_f \rangle$:

- (*Soundness, no writes*) Each component C^i can get to $\langle l'_f[i], v'_f[i], u'_f[i] \rangle$ by a modification σ_i of $\sigma[i]$. We may need to extend $\sigma[i]$ with delays and local transitions if $u_f(c^i) < u_f(c^k)$ for $C^k \in \mathcal{K}$, and we may need to remove delays and local transitions if $u_f(c^i) > u_f(c^k)$ for $C^k \in \mathcal{K}$. We get from $\langle l_0, v_0, u'_0 \rangle$ to $\langle l'_f, v'_f, u'_f \rangle$ by $(\sigma_1 \parallel \dots \parallel \sigma_n)$ for a composition of n components.
- (*Soundness, some write*) Let w^i be the latest write operation in σ , and $\mathcal{K}' = \{C^i\} \cup \mathcal{K}$. Partition $\sigma[i]$ into $\sigma_1^i w^i \sigma_2^i$ and $\sigma[k]$ into $\sigma_1^k \sigma_2^k$ for $C^k \in \mathcal{K}$, $C^j \notin \mathcal{K}'$ so that $(\sigma_1^k \parallel \sigma_1^i \parallel \sigma[j])$ reaches $\langle l_n, v_n, u_n \rangle$ and $u_n(c^k) = u_n(c^i)$. By the induction hypothesis with $\mathcal{K}'$ there is $\langle l_0, v_0, u'_0 \rangle \xrightarrow{\sigma'} \langle l'_n, v'_n, u'_n \rangle$ so that for any $C^k \in \mathcal{K}'$ we have $l_n[k] = l'_n[k]$, $v_n[k](y) = v'_n[k](y)$ for any $y \in V_D^k$, and $u_n[k] = u'_n[k]$.

We get to $\langle l'_f, v'_f, u'_f \rangle$ by $(\sigma'[j]\sigma^j \parallel \sigma[k]\sigma_2^k)$ where σ^i starts with w^i if $C^i \notin \mathcal{K}$ and σ^j is a sequence of internal, delay and read transitions so that $u'_f(c^j) = u'_f(c^k)$. Here the time divergence (Definition 1) comes in: we know that we can run these components until $u_f(c^k)$ for some $k \in \mathcal{K}$ (their results are insignificant, as they will not write to any component in $\mathcal{K}$).

We have the same local states for components in $\mathcal{K}$ at $\langle l'_f, v'_f, u'_f \rangle$ as for $\langle l_f, v_f, u_f \rangle$, because we have the same local states for components in $\mathcal{K}'$

at $\langle l'_n, v'_n, u'_n \rangle$ as for $\langle l_n, v_n, u_n \rangle$, the values for components in $\mathcal{K}$ after w^i depends only on the local state of components in $\mathcal{K}'$ before w^i , and we run the same local transitions, at the same global time, after w^i for components in $\mathcal{K}$.

- (*Completeness*) Let σ be the global time transition sequence. We construct a transition sequence $\sigma' = (\sigma^1 \parallel_t \dots \parallel_t \sigma^n)$, $\sigma^k = \sigma[k]$ and σ^i is constructed from $\sigma[i]$. If there is a first write operation w^i in $\sigma[i] = \sigma^i w^i \rho^i$ that require C^k to reach its final location after $\sigma[k]$ then σ^i is the prefix of $\sigma[i]$ before the write. If some writes in σ' require σ^i to reach its final location then $\sigma^i = \sigma[i] \rho^i$ with ρ^i any sequence of internal transitions reaching l_f^i ⁵. We have the same state for C^k since we have the same trace, except for some independent transitions added or removed for components $C^i \neq C^k$. $\square$

A.2 Correctness of Symbolic Local Time Semantics

According to Theorem 3: assume location vectors l_0, l_f , variable valuations v_0, v_f , clock valuations u_0, u_f , and local time zones Z_0, Z_f .

- (*Soundness*) whenever $\langle l_0, v_0, Z_0 \rangle \Longrightarrow_t^* \langle l_f, v_f, Z_f \rangle$ then for any $u_f \in Z_f$ $\langle l_0, v_0, u_0 \rangle \longrightarrow_t^* \langle l_f, v_f, u_f \rangle$.
- (*Completeness*) whenever $\langle l_0, v_0, u_0 \rangle \longrightarrow_t^* \langle l_f, v_f, u_f \rangle$ then $\langle l_0, v_0, Z_0 \rangle \Longrightarrow_t^* \langle l_f, v_f, Z_f \rangle$ so that $u_f \in Z_f$.

Proof We construct symbolic transitions $\langle l, v, Z \rangle \xRightarrow{\alpha}_t \langle l', v', Z' \rangle$ from transitions $\langle l, v, u \rangle \xrightarrow{\alpha}_t \langle l', v', u' \rangle$. The locations l' and data valuation v' are exactly the same in symbolic semantics as in concrete, and Z' is constructed so that for all $u \in Z$ we have $u' \in Z'$:

- For $\alpha = \delta^i$ with $l[i] \neq l_f^i$, and either $l[i] \neq l_\perp$ or $\neg \text{triggered}^i(v[i])$ we have:
$$\begin{aligned} Z' &= \{u' \mid u \in Z, \delta \in \mathbb{R}_{\geq 0}, \langle l, v, u \rangle \xrightarrow{\delta^i}_t \langle l, v, u' \rangle\} = \\ &= \{u' \mid u \in Z, \delta \in \mathbb{R}_{\geq 0}, \langle l[i], v[i], u[i] \rangle \xrightarrow{\delta} \langle l[i], v[i], u'[i] \rangle, u'[j] = u[j], j \neq i\} = \\ &= \{u' \mid u \in Z, \delta \in \mathbb{R}_{\geq 0}, u'[i] = u[i] \oplus \delta, u'[i] \models I^i(l[i]), u'[j] = u[j], j \neq i\} = \\ &= \{u' \mid u \in Z, \delta \in \mathbb{R}_{\geq 0}, u'[i] = u[i] \oplus \delta, u'[j] = u[j], j \neq i\} \wedge I^i(l[i]) = \\ &= \{u \mid u \in Z\}^{\uparrow i} \wedge I^i(l[i]) = Z^{\uparrow i} \wedge I^i(l[i]). \end{aligned}$$
- For $\alpha = \tau^i$ along an edge $l^i \xrightarrow{g, e, r}_t l'^i$ with e in the form $y := t$ and $v \models g_D$ we have:
$$\begin{aligned} Z' &= \{u' \mid u \in Z, \langle l, v, u \rangle \xrightarrow{\tau^i}_t \langle l', v', u' \rangle\} = \\ &= \{u' \mid u \in Z, \langle l[i], v[i], u[i] \rangle \xrightarrow{\tau^i}_t \langle l'[i], v'[i], u'[i] \rangle, u'[j] = u[j], j \neq i\} = \\ &= \{u' \mid u \in Z, u[i] \models g_C, u[i]' = [r := 0](u[i]), u'[j] = u[j], j \neq i\} = \\ &= \{u' \mid u \in Z, u \models g_C, u' = [r := 0]u\} = \\ &= r(\{u \mid u \in Z, u \models g_C\}) = r(\{u \mid u \in Z\} \wedge g_C) = r(Z \wedge g_C) \\ \text{If } C^i \text{ can delay: } &Z' = (r(Z \wedge g_C))^{\uparrow i} \wedge I^i(l[i]). \end{aligned}$$

⁵ This suggests that we require termination for timed behaviour, but instead we can allow writes if lower bound of a targets reference clock is greater than upper bound of writers (in the symbolic semantics).

- For $\alpha = r^i$ and $\text{triggered}^i(v)$ we have:
$$Z' = \{u' \mid u \in Z, \langle l, v, u \rangle \xrightarrow{r^i}_t \langle l', v', u' \rangle\} =$$

$$\{u' \mid u \in Z, \langle l[i], v[i], u[i] \rangle \xrightarrow{r^i}_t \langle l'[i], v'[i], u'[i] \rangle, u'[j] = u[j], j \neq i\} =$$

$$\{u' \mid u \in Z, u'[i] = [r_0^i := 0](u[i]), u'[j] = u[j], j \neq i\} =$$

$$\{u' \mid u \in Z, u' = [r_0^i := 0](u)\} = r_0^i(\{u \mid u \in Z\}) = r_0^i(Z)$$
 If C^i can delay: $Z' = (r_0^i(Z))^{\uparrow i} \wedge I^i(l[i])$.
- For $\alpha = w^i$ with $C^i \in \mathcal{C}$ and $l[i] = l_f^i$ we have:
$$Z' = \{u' \mid u \in Z, \langle l, v, u \rangle \xrightarrow{w^i}_t \langle l', v', u' \rangle\} =$$

$$\{u' \mid u \in Z, \langle l[i], v[i], u[i] \rangle \xrightarrow{w^i}_t \langle l'[i], v'[i], u'[i] \rangle, u'[j] = u[j], j \neq i\} =$$

$$\{u' \mid u \in Z, u'[i] = [r_f^i := 0](u[i]), u'[j] = u[j], j \neq i\} =$$

$$\{u' \mid u \in Z, u' = [r_f^i := 0]u\} = r_f^i(\{u \mid u \in Z\}) = r_f^i(Z).$$
 If C^i can delay: $Z' = (r_f^i(Z))^{\uparrow i} \wedge I^i(l[i])$.

Using these constructions, showing soundness and completeness by induction on the length of a transition sequence is easy. $\square$

A.3 Correctness of Ample Set Construction

According to Theorem 4: for any trace σ in G reaching the target location, the subgraph R induced by the ample set construction in Definition 4 contains a representative of an extension $\sigma\rho$ in G . To show that the construction is finite we will use the following lemma:

Lemma 2. *Assume local computations terminate. If some w^i is in ample infinitely often and C^i is connected to C^k via at most n components, then for all C^j connected to C^k via at most n components w^j is in ample infinitely often.*

Proof. By induction on n (distance to C^k).

- (Base case, $n = 0$): If w^i is in ample infinitely often then w^j is, since for $n = 0$ we have $i = j = k$.
- (Induction step, $n > 0$): Assume w^i is in ample infinitely often and C^i is connected to C^k via n components. A write operation w^i is disabled if components C^j such that $\text{dependent}(w^i, w^j)$ are not idle or finished. If $w^i \in \mathcal{P}^k$ then one of the components C^j is finished and w^j triggers a component connected to C^k via $n - 1$ components. By induction, for all C^j connected to C^k via at most $n - 1$ components w^j is in ample infinitely often. Because a write operation is dependent on its target also writes from components connected to C^k via n components are in ample infinitely often (every read is eventually followed by a write as local computations terminate). $\square$

Proof Similar to [7], by constructing a representative σ' in R for some $\sigma = \alpha_0, \alpha_1 \dots$ in G , starting from some state s_0 in G .

- If $\alpha_0 \in \text{ample}(s_0)$ then $s_0 \xrightarrow{\alpha_0} s_1$ is in R and we construct $\alpha_1, \alpha_2 \dots$ from s_1 inductively.

- If $\alpha_0 \notin \text{ample}(s_0)$ and α_k is the first transition from $\text{ample}(s_0)$ in σ then α_k is independent of all transitions α_i with $0 \leq i < k$, otherwise some $\alpha_j \in \text{enabled}(s_0)$ with $0 \leq j \leq i$ would enable α_i and so α_j must be in $\text{ample}(s_0)$ by ample set construction. By independence we know that $\alpha_k \alpha_0 \alpha_1 \dots \alpha_{k-1} \alpha_{k+1} \dots$ is also in G , so we can construct its representative from s_1 after $s_0 \xrightarrow{\alpha_k} s_1$.
- If $\alpha_0 \notin \text{ample}(s_0)$ and σ contains no transitions from $\text{ample}(s_0)$, let β be an arbitrary transition in $\text{ample}(s_0)$. Since σ does not contain β , it is independent of all transitions in σ , otherwise some α_i would be in $\text{ample}(s_0)$ by ample set construction. The construction of a representative for σ continues from s_1 after $s_0 \xrightarrow{\beta} s_1$. However, since β is not in the trace σ we must ensure that we eventually take some transition in σ in order to get a finite construction.

Selecting all of $\mathcal{P}^k$ somewhere in case of a local loop ensures that some transition in σ is taken in this case. By Lemma 2 a loop involving write operations (i.e. a non-local loop) will have some transition in σ in ample, since σ leads to $l_{\odot}^k$ in C^k . $\square$