

Revisiting Local Time Semantics for Networks of Timed Automata

R. Govind

Chennai Mathematical Institute, India

Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France

govindr@cmi.ac.in

Frédéric Herbreteau

Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France

fh@labri.fr

B. Srivathsan

Chennai Mathematical Institute, India

sri@cmi.ac.in

Igor Walukiewicz

Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France

igw@labri.fr

Abstract

We investigate a zone based approach for the reachability problem in timed automata. The challenge is to alleviate the size explosion of the search space when considering networks of timed automata working in parallel. In the timed setting this explosion is particularly visible as even different interleavings of local actions of processes may lead to different zones. Salah *et al.* in 2006 have shown that the union of all these different zones is also a zone. This observation was used in an algorithm which from time to time detects and aggregates these zones into a single zone.

We show that such aggregated zones can be calculated more efficiently using the local time semantics and the related notion of local zones proposed by Bengtsson *et al.* in 1998. Next, we point out a flaw in the existing method to ensure termination of the local zone graph computation. We fix this with a new algorithm that builds the local zone graph and uses abstraction techniques over (standard) zones for termination. We evaluate our algorithm on standard examples. On various examples, we observe an order of magnitude decrease in the search space. On the other examples, the algorithm performs like the standard zone algorithm.

2012 ACM Subject Classification Theory of computation → Verification by model checking

Keywords and phrases Timed automata, verification, local-time semantics, abstraction

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2019.16

Related Version A full version of this paper is available at [10], <http://arxiv.org/abs/1907.02296>.

Funding Work supported by UMI 2000 ReLaX, ANR project TickTac, and project IoTTTA – CEFIPRA Indo-French program in ICST – DST/CNRS ref. 2016-01. Author B. Srivathsan is partially funded by grants from Infosys Foundation, India and Tata Consultancy Services, India.

1 Introduction

Timed automata [1] are a popular model for real-time systems. They extend finite state automata with real valued variables called *clocks*. Constraints on clock values can be used as guards for transitions, and clocks can be reset to zero during transitions. Often, it is more natural to use a network of timed automata which operate concurrently and synchronize on joint actions. We study the reachability problem for networks of timed automata: given a state of the timed automaton network, is there a run from the initial state to the given state.



© R. Govind, Frédéric Herbreteau, B. Srivathsan, and Igor Walukiewicz;
licensed under Creative Commons License CC-BY

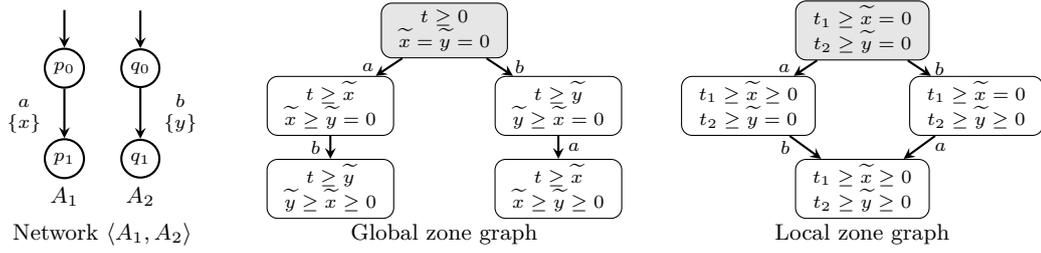
30th International Conference on Concurrency Theory (CONCUR 2019).

Editors: Wan Fokkink and Rob van Glabbeek; Article No. 16; pp. 16:1–16:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Illustration of commutativity in the local zone graph.

A widely used technique to solve the reachability problem constructs a *zone graph* [7] whose nodes are $(state, zone)$ pairs consisting of a state of the automaton and a *zone* representing a set of clock valuations [8]. This graph may not be finite, so in order to guarantee termination of an exploration algorithm, various sound and complete abstraction techniques are used [7, 2, 12, 9].

Dealing with automata operating in parallel poses the usual state-space explosion problem arising due to different interleavings. Consider an example of a network with two processes in Figure 1. Actions a and b are local to each process. Variables x, y are clocks and $\{x\}$ denotes that clock x is reset in transition a . The (global) zone graph maintains the set of configurations reached after each sequence of actions. Although a and b are local actions, there is an intrinsic dependence between the two processes happening due to time. Hence, the zone reached after executing sequence ab contains configurations where x is reset before y and the zone after ba contains configurations where y is reset before x . So the sequences ab and ba lead to different zones. The number of different interleavings of sequences of local actions increases exponentially when their length grows, or when more processes get involved. Every interleaving can potentially lead to a different zone.

Salah et al. [17] have shown a surprising property that the *union* of all zones reached by the interleavings of a sequence of actions is also a zone. We call it an *aggregated zone*. Their argument is based on the fact that for a given sequence one can write a zone-like constraint defining all the runs on the interleavings of the sequence. Then the aggregated zone is obtained simply by projecting this big constraint on relevant components. This approach requires to work with sets of constraints whose size grows with the length of a sequence. This is both inefficient and limited to finite sequences. They use this observation in an algorithm where, when all the interleavings of a sequence σ have been explored, the resulting zones are aggregated to a single zone and further exploration is restricted to this aggregated zone. This requires detecting from time to time whether aggregation can happen. This is an obstacle in using aggregated zones in efficient reachability-checking algorithms. Another limitation of this approach is that it works only for acyclic automata.

Another approach by Bengtsson et al. [4] involves making time local to each process. This local time approach is based on a very elegant idea: make time in every process progress independently, and synchronize local times of processes when they need to perform a common action. In consequence, the semantics has the desired property: two actions whose process domains are disjoint are commutative. In the example above, depending on the local time in processes A_1 and A_2 , the sequence ab may result, on a global time scale, in a occurring before b , as well as b occurring before a . As a result, the set of valuations reached after ab does not remember the order in which a and b occurred. Thus, sequences ab and ba lead to the same set of configurations: those obtained after doing a and b concurrently. Similar to standard zones, we now have local zones and a local zone graph having $(state, local\ zone)$ pairs. Due

to the above argument, the local zone graph has a nice property – if a run σ from an initial zone reaches a local zone, then all the runs equivalent to σ lead to the same local zone. Local zone graphs are therefore ideal for handling interleavings. However, substantially more involved abstraction techniques are needed for local zones to make the local zone graph finite. A *subsumption relation* between local zones was defined in [4] but no algorithm was proposed and there was no effective way to use local time semantics. Later Minea [15] proposed a widening operator on local zones to construct finite local zone graphs.

Summary of results in the paper:

- We show that the aggregated zone of interleavings of σ in the standard zone graph is obtained by synchronizing valuations in the local zone obtained after σ (Theorem 18). Hence computing the local zone graph gives a more direct and efficient algorithm than Salah et al. to compute aggregated zones. Moreover, this algorithm is not restricted to acyclic timed automata.
- We point out a flaw in the abstraction procedure of Minea to get a finite local zone graph (Section 5).
- We propose a different algorithm to get a finite local zone graph. This gives a new reachability algorithm for networks of timed automata, which works with local zones but uses subsumption on standard zones (Definition 19 and Theorem 20). Instead of subsumptions between local zones, we use subsumptions between the synchronized valuations inside these local zones. This helps us to exploit the (well-studied) subsumptions over standard zones [7, 2, 12, 9]. Moreover, this subsumption is much more aggressive than the standard one since, thanks to local-time semantics, the (aggregated) zone used in the subsumption represents all the valuations reachable not only by the execution that we are exploring but also by all the executions equivalent to it.
- We report on experiments performed with a prototype implementation of the algorithm (Table 1). The algorithm performs surprisingly well on some examples, and it is never worse than the standard zone graph algorithm.

Related work. The basis of this work is local-time semantics and local zones developed by Bengtsson et. al. [4]. The authors have left open how to use this semantics to effectively compute the local zone graph since no efficient procedure was provided to ensure its finiteness. To that purpose, Minea [16, 15] has proposed a subsumption operation on local zones, and an algorithm using this operation. Unfortunately, as we exhibit here, the algorithm has a flaw that is not evident to repair. Lugiez et. al. [14] also use local time but their method is different. They use constraints to check if local clocks of an execution can be synchronized sufficiently to obtain a standard run.

Aggregated zones are crucial to obtain an efficient verification procedure for networks of timed automata. Coming to the same state with different zones inflicts a huge blowup in the zone graph, since the same paths are explored independently from each of these zones. This has also been observed in the context of multi-threaded program verification in [18]. Solving this problem in the context of program analysis requires to over approximate the aggregated state. Fortunately, in the context of timed automata, the result of aggregating these zones is still a zone [17]. Efficient computation of aggregated zones is thus an important advance in timed automata verification as demonstrated by our experimental results in Section 6. A full version of this paper is available at [10].

2 Networks of timed automata

We start by defining networks of timed automata and two semantics for them: a global-time semantics (the usual one) and a local-time semantics (introduced in [4]). Then, we recall the fact that they are equivalent, and state some interesting properties of local-time semantics w.r.t. concurrency that were observed in [4]. Let \mathbb{N} denote the set of natural numbers and $\mathbb{R}_{\geq 0}$ the set of non-negative reals. Let X be a finite set of variables called *clocks*. Let $\phi(X)$ denote a set of clock constraints generated by the following grammar: $\phi := x \sim c \mid \phi \wedge \phi$ where $x \in X$, $c \in \mathbb{N}$, and $\sim \in \{<, \leq, =, \geq, >\}$.

► **Definition 1** (Network of timed automata). *A network of timed automata with k processes, is a k -tuple of timed automata A_1, \dots, A_k . Each process $A_p = \langle Q_p, \Sigma_p, X_p, q_p^{init}, T_p \rangle$ has a finite set of states Q_p , a finite alphabet of actions Σ_p , a finite set of clocks X_p , an initial state q_p^{init} , and transitions $T_p \subseteq \Sigma_p \times Q_p \times \phi(X_p) \times 2^{X_p} \times Q_p$. We require that the sets of states, and the sets of clocks are pairwise disjoint: $Q_{p_1} \cap Q_{p_2} = \emptyset$, and $X_{p_1} \cap X_{p_2} = \emptyset$ for $p_1 \neq p_2$. We write *Proc* for the set of all processes.*

A network is a parallel composition of timed automata. Its semantics is that of the timed automaton obtained as the “synchronized product” of the processes. For an action b , a b -transition of a process p is an element of T_p with b in the first component. Synchronization happens on two levels: (i) via time that advances the same way in all the processes, and (ii) via common actions, for example if $b \in \Sigma_1 \cap \Sigma_2$, then processes 1 and 2 need to synchronize by doing a b -transition. We define the domain of an action b : $dom(b) = \{p : b \in \Sigma_p\}$ as the set of processes that must synchronize to do b . We will use some abbreviations: $Q = \prod_{p=1}^k Q_p$, $\Sigma = \bigcup_{p=1}^k \Sigma_p$ and $X = \bigcup_{p=1}^k X_p$.

The semantics of a network is governed by the value of clocks at each instant. We choose to represent these values using offsets as this allows a uniform presentation of the global-time semantics below and the local-time semantics in Section 2.1. For every clock x , we introduce an offset variable \tilde{x} . The value of \tilde{x} is the time-stamp at which x was last reset. In addition, we consider a variable t which tracks the global time: essentially t is a clock that is never reset. We now make this notion precise. Let $\tilde{X}_p = \{\tilde{x} \mid x \in X_p\}$ and $\tilde{X} = \bigcup_{p=1}^k \tilde{X}_p$. A *global valuation* v is a function $v : \tilde{X} \cup \{t\} \mapsto \mathbb{R}_{\geq 0}$ such that $v(t) \geq v(\tilde{x})$ for all variables \tilde{x} . In this representation, the value of clock x corresponds to $v(t) - v(\tilde{x})$, denoted $\bar{v}(x)$ ¹.

Recall that offset variable \tilde{x} stores the last time-stamp at which clock x has been reset. Hence, a delay in offset representation increases the value of reference clock t and leaves offset variables \tilde{x} unchanged. Formally: for $\delta \in \mathbb{R}_{\geq 0}$, we denote by $v + \delta$ the valuation defined by: $(v + \delta)(t) = v(t) + \delta$ and $(v + \delta)(\tilde{x}) = v(\tilde{x})$ for all \tilde{x} . Similarly, resetting the clocks in $R \subseteq X$, yields a global valuation $[R]v$ defined by: $([R]v)(t) = v(t)$, and $([R]v)(\tilde{x})$ is $v(t)$ if $x \in R$, and $v(\tilde{x})$ otherwise. Given a global valuation v and a clock constraint g , we write $v \models g$ if every constraint in g holds after replacing x with its value $\bar{v}(x) = v(t) - v(\tilde{x})$.

A configuration of the network is a pair (q, v) where $q \in Q$ is a global state, and v is a global valuation. We will write $q(p)$ to refer to the p -th component of the state q .

► **Definition 2** (Global-time semantics). *The semantics of a network \mathcal{N} is given by a transition system whose states are configurations (q, v) . The initial configuration is (q^{init}, v^{init}) where $q^{init}(p) = q_p^{init}$ is the tuple of initial states, and $v^{init}(y) = 0$ for $y \in \tilde{X} \cup \{t\}$.*

¹ Usually, semantics of timed automata is described using valuations of the form \bar{v} . Here, we have chosen v which uses offsets since it extends naturally to the local-time setting. Translations between these two kinds of valuations is straightforward, as shown.

There are two kinds of transitions, which we call steps: global delay, and action steps. A global delay by the amount $\delta \in \mathbb{R}_{\geq 0}$ gives a step $(q, v) \xrightarrow{\delta}_{st} (q, v + \delta)$. An action step on action b gives $(q, v) \xrightarrow{b}_{st} (q', v')$ if there is a set of b -transitions $\{(q_p, g_p, R_p, q'_p)\}_{p \in \text{dom}(b)}$ of the respective processes such that:

- processes from $\text{dom}(b)$ change states: $q_p = q(p)$, $q'_p = q'(p)$, for $p \in \text{dom}(b)$, and $q(p) = q'(p)$ for $p \notin \text{dom}(b)$;
- all the guards are satisfied: $v \models g_p$, for $p \in \text{dom}(b)$;
- all resets are performed: $v' = [\bigcup_{p \in \text{dom}(b)} R_p]v$;

A run of an automaton from a configuration (q_0, v_0) is a sequence of steps starting in (q_0, v_0) . For a sequence $u = b_1 \dots b_n$ of actions, we write $(q_0, v_0) \xrightarrow{u} (q_n, v'_n)$ if there is a run

$$(q_0, v_0) \xrightarrow{\delta_0}_{st} (q_0, v'_0) \xrightarrow{b_1}_{st} (q_1, v_1) \xrightarrow{\delta_1}_{st} (q_1, v'_1) \dots \xrightarrow{b_n}_{st} (q_n, v_n) \xrightarrow{\delta_n}_{st} (q_n, v'_n)$$

for some delays $\delta_0, \dots, \delta_n \in \mathbb{R}_{\geq 0}$.

► **Definition 3** (Reachability problem). *The reachability problem is to decide, given a network \mathcal{N} and a state q , whether there is a run reaching q ; or in other words, whether there exists a sequence of transitions u such that $(q^{init}, v^{init}) \xrightarrow{u} (q, v)$ for some valuation v .*

The reachability problem for networks of timed automata is PSPACE-complete [1].

2.1 Local-time semantics

The definition of a network of automata suggests an independence relation between actions: a pair of actions with disjoint domains (i.e. involving distinct processes) should commute. We say that two sequences of actions are *equivalent*, written $u \sim w$ if one can be obtained from the other by repeatedly permuting adjacent actions with disjoint domains.

► **Lemma 4.** *For two equivalent sequences $u \sim w$: if there are two runs $(q, v) \xrightarrow{u} (q_u, v_u)$, and $(q, v) \xrightarrow{w} (q_w, v_w)$ then $q_u = q_w$.*

Observe that in the above lemma we cannot assert that $v_u = v_w$. Even further, the existence of $(q, v) \xrightarrow{u} (q_u, v_u)$ does not imply that a run from (q, v) on w is feasible. This happens due to global time delays, i.e., delays that involve all the processes. For example, consider actions a and b on disjoint processes with a having guard $x \leq 1$ and b having guard $y \geq 2$. Then from the initial valuation one can execute ab but not ba .

One solution to get commutativity between actions with disjoint domains is to consider local-time semantics [4]. In this semantics, time elapses independently in every process, and time elapse is synchronized before executing a synchronized action. This way, two actions with disjoint domains become commutative. In the example from the previous paragraph, while the process executing b elapses 2 time units, the other process is allowed to not elapse time at all and hence ba becomes possible. Moreover, for the reachability problem, local-time semantics is equivalent to the standard one.

In the local-time semantics, we replace the clock t which was tracking the global time, with individual *reference clocks* t_p for each process A_p which track the local time of each process. We set $\tilde{X}'_p = \tilde{X}_p \cup \{t_p\}$ and $\tilde{X}' = \bigcup_p \tilde{X}'_p$. A *local valuation* v is a valuation over the set of clocks \tilde{X}' such that $v(t_p) \geq v(\tilde{x})$ for all processes $p \in \text{Proc}$ and all clocks $\tilde{x} \in \tilde{X}_p$. This restriction captures the intuition that t_p is a reference clock for process p , and it is never reset. In this setting, the value $v(t_p) - v(\tilde{x})$ of clock x is defined relative to the reference clock t_p of process p that owns x , i.e. $x \in X_p$. We will use the notation v for local valuations to distinguish from global valuations v .

We introduce a new operation of local time elapse. For a process $p \in Proc$ and $\delta \in \mathbb{R}_{\geq 0}$, operation $\mathbf{v} +_p \delta$ adds δ to $\mathbf{v}(t_p)$, the value of the reference clock t_p of process p , and leaves the other variables unchanged. Formally, $(\mathbf{v} +_p \delta)(t_p) = \mathbf{v}(t_p) + \delta$ and $(\mathbf{v} +_p \delta)(y) = \mathbf{v}(y)$ for all $y \in \tilde{X}' \setminus \{t_p\}$. A local valuation \mathbf{v} satisfies a clock constraint g , denoted $\mathbf{v} \models g$ if every constraint in g holds after replacing x by its value $\mathbf{v}(t_p) - \mathbf{v}(\tilde{x})$ where p is the process such that $x \in X_p$. We denote by $[R]\mathbf{v}$ the valuation obtained after resetting the clocks in $R \subseteq X$ and defined by: $([R]\mathbf{v})(t_p) = \mathbf{v}(t_p)$ for every reference clock t_p , $([R]\mathbf{v})(\tilde{x}) = \mathbf{v}(\tilde{x})$ if $x \notin R$, and $([R]\mathbf{v})(\tilde{x}) = \mathbf{v}(t_p)$ if $x \in R$ and p is the process such that $x \in X_p$.

► **Definition 5** (Local steps of a timed automata network). *There are two kinds of local steps in a network \mathcal{N} : local delay, and local action. A local delay $\delta \in \mathbb{R}_{\geq 0}$ in process $p \in Proc$ is a step $(q, \mathbf{v}) \xrightarrow{p, \delta}_{st} (q, \mathbf{v} +_p \delta)$. For an action b , we have a step $(q, \mathbf{v}) \xrightarrow{b}_{st} (q', \mathbf{v}')$ if there is a set of b -transitions of respective processes $\{(q_p, g_p, R_p, q'_p)\}_{p \in dom(b)}$ such that:*

- $q_p = q(p)$, $q'_p = q'(p)$, for $p \in dom(b)$, and $q(p) = q'(p)$ for $p \notin dom(b)$;
- start times are synchronized: $\mathbf{v}(t_{p_1}) = \mathbf{v}(t_{p_2})$, for every $p_1, p_2 \in dom(b)$;
- guards are satisfied: $\mathbf{v} \models g_p$, for every $p \in dom(b)$;
- resets are performed: $\mathbf{v}' = [\bigcup_{p \in dom(b)} R_p]\mathbf{v}$;

The main difference with respect to global semantics is the presence of local time delay. As a result, every process can be in a different local time as emphasized by the reference clocks in each process. In consequence, in local action steps we require that when processes do a common action their local times should be the same. Of course a standard delay δ on all processes can be simulated by a sequence of delays on every process separately, as $\xrightarrow{1, \delta}_{st} \dots \xrightarrow{k, \delta}_{st}$. For a sequence of local delays $\Delta = (p_1, \delta_1) \dots (p_n, \delta_n)$ we will write $(q, \mathbf{v}) \xrightarrow{\Delta}_{st} (q, \mathbf{v}')$ to mean $(q, \mathbf{v}) \xrightarrow{p_1, \delta_1}_{st} (q, \mathbf{v}_1) \xrightarrow{p_2, \delta_2}_{st} \dots \xrightarrow{(p_n, \delta_n)}_{st} (q, \mathbf{v}')$.

► **Definition 6** (Local run). *A local run from a configuration (q_0, \mathbf{v}_0) is a sequence of local steps. For a sequence of actions $u = b_1 \dots b_n$, we write $(q_0, \mathbf{v}_0) \xrightarrow{u} (q_n, \mathbf{v}'_n)$ if for some sequences of local delays $\Delta_0, \dots, \Delta_n$ there is a local run*

$$(q_0, \mathbf{v}_0) \xrightarrow{\Delta_0}_{st} (q_0, \mathbf{v}'_0) \xrightarrow{b_1}_{st} (q_1, \mathbf{v}_1) \xrightarrow{\Delta_1}_{st} \dots \xrightarrow{b_n}_{st} (q_n, \mathbf{v}_n) \xrightarrow{\Delta_n}_{st} (q_n, \mathbf{v}'_n)$$

Observe that a run may start and end with a sequence of delays. In the next section we will make a link between local and global runs. For this we will first examine independence properties of local runs which are much better than for global runs (cf. Lemma 4).

► **Lemma 7** (Independence). *Suppose $dom(a) \cap dom(b) = \emptyset$. If $(q, \mathbf{v}) \xrightarrow{ab} (q', \mathbf{v}')$ then $(q, \mathbf{v}) \xrightarrow{ba} (q', \mathbf{v}')$. If $(q, \mathbf{v}) \xrightarrow{a} (q_a, \mathbf{v}_a)$ and $(q, \mathbf{v}) \xrightarrow{b} (q_b, \mathbf{v}_b)$ then $(q, \mathbf{v}) \xrightarrow{ab} (q_{ab}, \mathbf{v}_{ab})$ for some q_{ab} and \mathbf{v}_{ab} .*

Recall that two sequences of actions are equivalent, written $u \sim w$ if one can be obtained from the other by repeatedly permuting adjacent actions with disjoint domains. Directly from the previous lemma we obtain.

► **Lemma 8.** *If $(q_0, \mathbf{v}_0) \xrightarrow{u} (q_n, \mathbf{v}_n)$ and $u \sim w$ then $(q_0, \mathbf{v}_0) \xrightarrow{w} (q_n, \mathbf{v}_n)$.*

With local-time semantics two equivalent sequences not only reach the same state q_n , but also the same local valuation \mathbf{v}_n (in contrast with Lemma 4 for global-time semantics).

2.2 Comparing local and global runs

We have presented two semantics for networks of timed automata: global-time and local-time. Local runs have more freedom as time can elapse independently in every process. Yet, with respect to state reachability the two concepts turn out to be equivalent.

► **Definition 9.** A local valuation \mathbf{v} is synchronized if for every pair of processes p_1, p_2 , the values of their reference clocks are equal: $\mathbf{v}(t_{p_1}) = \mathbf{v}(t_{p_2})$.

For a synchronized local valuation \mathbf{v} , let $\mathbf{global}(\mathbf{v})$ be the global valuation v such that $\mathbf{v}(\tilde{x}) = v(\tilde{x})$ and $v(t) = \mathbf{v}(t_1) = \dots = \mathbf{v}(t_k)$. Conversely, to every global valuation v , we associate the synchronized local valuation $\mathbf{local}(v) = \mathbf{v}$ where $\mathbf{v}(\tilde{x}) = v(\tilde{x})$ and $\mathbf{v}(t_p) = v(t)$ for every reference clock t_p .

► **Lemma 10.** If $(q, \mathbf{v}) \xrightarrow{u} (q', \mathbf{v}')$ is a local run where \mathbf{v} and \mathbf{v}' are synchronized local valuations, there exists a global run $(q, \mathbf{global}(\mathbf{v})) \xrightarrow{w} (q', \mathbf{global}(\mathbf{v}'))$ for some $w \sim u$. Conversely, if $(q, v) \xrightarrow{u} (q', v')$ is a global run, then there is a local run $(q, \mathbf{local}(v)) \xrightarrow{u} (q', \mathbf{local}(v'))$.

The reachability problem with respect to local semantics is defined as before: q is reachable if there is a local run $(q^{init}, \mathbf{v}^{init}) \xrightarrow{u} (q, \mathbf{v})$ for some \mathbf{v} where $\mathbf{v}^{init} = \mathbf{local}(v^{init})$. By adding some local delays at the end of the run we can always assume that \mathbf{v} is synchronized. Lemma 10 thus implies that the reachability problem in local semantics is equivalent to the standard one in global semantics.

3 Zone graphs

We introduce zones, a standard approach for solving reachability in timed automata. Zones are sets of valuations that can be represented efficiently using simple constraints.

Let us fix a network \mathcal{N} of timed automata with k processes. Recall that each process p has a set of clocks X_p and corresponding offset variables \tilde{X}_p . The set of clocks in \mathcal{N} is $X = \bigcup_{i=1}^k X_p$. Similarly, the set of offset variables in \mathcal{N} is $\tilde{X} = \bigcup_{i=1}^k \tilde{X}_p$.

3.1 Standard zone-based algorithm for reachability

Recall that in the global-time semantics, \mathcal{N} has a reference clock t . A *global zone* is a set of global valuations described by a conjunction of constraints of the form $y_1 - y_2 \triangleleft c$ where $y_1, y_2 \in \tilde{X} \cup \{t\}$, $\triangleleft \in \{<, \leq\}$ and $c \in \mathbb{Z}$. Since global valuations need to satisfy $v(\tilde{x}) \leq v(t)$ for every offset variable $\tilde{x} \in \tilde{X}$, a global zone satisfies $\tilde{x} \leq t$ for every $\tilde{x} \in \tilde{X}$.

Let g be a guard and R a set of clocks. We define the following operations on zones: $Z_g = \{v \mid v \models g\}$ is the set of global valuations satisfying g , $[R]Z := \{[R]v \mid v \in Z\}$ and $\overrightarrow{Z} := \{v \mid \exists v' \in Z, \exists \delta \in \mathbb{R}_{\geq 0} \text{ s. t. } v = v' + \delta\}$. From [4], Z_g , $[R]Z$ and \overrightarrow{Z} are all zones. We say that a zone is *time-elapsd* if $Z = \overrightarrow{Z}$.

The semantics of a network of timed automata can be described in terms of global zones. For an action b , consider a set of b -transitions of respective processes $\{(q_p, g_p, R_p, q'_p)\}_{p \in \text{dom}(b)}$. Let $R = \bigcup_{p \in \text{dom}(b)} R_p$, and $g = \bigwedge_{p \in \text{dom}(b)} g_p$. Then we have a transition $(q, Z) \xrightarrow{b} (q', Z')$ where $Z' = \overrightarrow{[R](Z \cap Z_g)}$ provided that $q(p) = q_p$, $q'(p) = q'_p$ if $p \in \text{dom}(b)$, and $q'(p) = q(p)$ otherwise, and Z' is not empty. We write \Rightarrow for the union over all \xrightarrow{b} .

The *global zone graph* $ZG(\mathcal{N})$ of a timed automaton network \mathcal{N} is a transition system whose nodes are of the form (q, Z) where $q \in Q$ and Z is a time-elapsd global zone. The transition relation is given by \Rightarrow . The initial node is (q^{init}, Z^{init}) where q^{init} is the tuple of

initial states and $Z^{init} = \overrightarrow{\{v^{init}\}}$ where v^{init} is the initial global valuation. The zone graph $ZG(\mathcal{N})$ is known to be sound and complete with respect to reachability. This means that a state q is reachable by a run of \mathcal{N} iff a node (q, Z) for some non-empty Z is reachable from (q^{init}, Z^{init}) in the zone graph. As zone graphs may be infinite, an abstraction operator is used to obtain a finite quotient.

An *abstraction operator* $\mathbf{a} : P(\mathbb{R}_{\geq 0}^X) \rightarrow P(\mathbb{R}_{\geq 0}^X)$ [2] is a function from sets of valuations to sets of valuations such that $W \subseteq \mathbf{a}(W)$ and $\mathbf{a}(\mathbf{a}(W)) = \mathbf{a}(W)$. Simulation relations between valuations are a convenient way to construct abstraction operators that are correct for reachability. A *time-abstract simulation* is a relation between valuations that depends on a given network \mathcal{N} . We say that v_1 can be simulated by v_2 , denoted $v_1 \preceq v_2$ if for every state q of \mathcal{N} , and every delay-action step $(q, v_1) \xrightarrow{\delta_1} \xrightarrow{b} (q', v'_1)$ there is a delay $\delta_2 \in \mathbb{R}_{\geq 0}$ such that $(q, v_2) \xrightarrow{\delta_2} \xrightarrow{b} (q', v'_2)$ and $v'_1 \preceq v'_2$. The simulation relation can be lifted to global zones: we say that Z is simulated by Z' , written as $Z \preceq Z'$ if for all $v \in Z$ there exists a $v' \in Z'$ such that $v \preceq v'$. An abstraction \mathbf{a} based on \preceq is defined as $\mathbf{a}(W) = \{v \mid \exists v' \in W \text{ with } v \preceq v'\}$. The abstraction \mathbf{a} is finite if its range is finite. Given two nodes (q, Z) and (q', Z') of $ZG(\mathcal{N})$, (q, Z) is *subsumed* by (q', Z') , denoted $(q, Z) \sqsubseteq^{\mathbf{a}} (q', Z')$, if $q = q'$ and $Z \subseteq \mathbf{a}(Z')$.

► **Remark 11.** Our definition of zones slightly differs from the standard definition in the literature (e.g. [5]) since we use offset variables to represent clock valuations. Yet, finite time-abstract simulations from the literature [2, 12] can be adapted to our settings as a simulation over standard valuations \bar{v} can be expressed as a simulation over global valuations v since for every clock x , $\bar{v}(x) = v(t) - v(\tilde{x})$, and zones over valuations v can be translated to zones over standard valuations \bar{v} .

A finite abstraction \mathbf{a} allows to construct a finite *global zone graph with subsumption* for a network of timed automata \mathcal{N} . The construction starts from the initial node of $ZG(\mathcal{N})$. Using, say, a breadth-first-search (BFS), for every constructed node we examine all its successors in $ZG(\mathcal{N})$, and keep only those that are maximal w.r.t. to $\sqsubseteq^{\mathbf{a}}$ relation. Computing such a zone graph with subsumption gives an algorithm for the reachability problem. However, the global zone graph, and hence the algorithm above, are sensitive to the combinatorial explosion arising from parallel composition. Global time makes any two actions potentially dependent – the same is still true on the level of zones. Zone graphs based on the local-time semantics, as presented next, solve this problem.

3.2 Local zone graphs

The goal of this section is to introduce a concept similar to global zones and global zone graphs for local-time semantics. Recall that in the local-time semantics, each process p has a reference clock t_p .

A *local zone* is a zone over local valuations: a set of local valuations defined by constraints $y_1 - y_2 \triangleleft c$ where $y_1, y_2 \in \tilde{X} \cup \{t_1, \dots, t_k\}$. Recall that a local valuation \mathbf{v} needs to satisfy: $\mathbf{v}(\tilde{x}) \leq \mathbf{v}(t_p)$ for every process p and every $\tilde{x} \in \tilde{X}_p$. This means that a local zone satisfies $\tilde{x} \leq t_p$ for every process p and every $\tilde{x} \in \tilde{X}_p$. We will use Z , eventually with subscripts, to range over local zones, and distinguish from global zones Z . Local zones are closed under all basic operations involved in a local step of a network of timed automata, namely: local time elapse, intersection with a guard, and reset of clocks [4]. That is, for every local zone Z :

- the set $\text{local-elapse}(Z) = \{\mathbf{v} +_1 \delta_1 +_2 \dots +_k \delta_k \mid \mathbf{v} \in Z, \delta_1, \dots, \delta_k \in \mathbb{R}_{\geq 0}\}$ is a local zone.
- for every guard g the set $Z_g = \{\mathbf{v} \mid \mathbf{v} \models g\}$ is a local zone.
- for every set of clocks $R \subseteq X$, the set $[R]Z = \{[R]\mathbf{v} \mid \mathbf{v} \in Z\}$ is a local zone.

Local zones can be implemented using DBMs. Hence, they can be computed and stored as efficiently as standard zones.

The operations of local time elapse, guard intersection, and reset, enable us to describe a local step $(q, Z) \xrightarrow{b} (q', Z')$ on the level of local zones. This is done in the same way as for global zones. Observe that a local step is indexed only by an action, as the time is taken care of by the local time elapse operation. Formally, for an action b consider a set of b -transitions $\{(q_p, g_p, R_p, q'_p)\}_{p \in \text{dom}(b)}$ of respective processes. Then we have a transition $(q, Z) \xrightarrow{b} (q', Z')$ for $Z' = \text{local-elapse}([R](Z \cap Z_g \cap Z_{\text{sync}}))$ where $Z_g = \bigcap_{p \in \text{dom}(b)} Z_{g_p}$, $Z_{\text{sync}} = \{v \mid v(t_{p_1}) = v(t_{p_2}) \text{ for } p_1, p_2 \in \text{dom}(b)\}$ and $R = [\bigcup_{p \in \text{dom}(b)} R_p]$. Intuitively, Z' is the set of valuations obtained through reset and then local-time elapse, from valuations in Z that satisfy the guard and such that the processes involved in action b are synchronised. We extend \xrightarrow{b} to $(q, Z) \xrightarrow{u} (q', Z')$ for a sequence of actions u in the obvious way.

Using local zones, we construct a local zone graph and show that it is sound and complete for reachability testing. The only missing step is to verify the pre/post properties of runs on local zones. We say a local zone is *time-elapsed* if $Z = \text{local-elapse}(Z)$.

► **Lemma 12** (Pre and post properties of runs on local zones). *Let u be a sequence of actions.*

- *If $(q, v) \xrightarrow{u} (q', v')$ and $v \in Z$ for some time-elapsed local zone Z then $(q, Z) \xrightarrow{u} (q', Z')$ and $v' \in Z'$ for some local zone Z' .*
- *If $(q, Z) \xrightarrow{u} (q', Z')$ and $v' \in Z'$ then $(q, v) \xrightarrow{u} (q', v')$, for some $v \in Z$.*

► **Definition 13** (Local zone graph). *For a network of timed automata \mathcal{N} the local zone graph of \mathcal{N} , denoted $\text{LZG}(\mathcal{N})$, is a transition system whose nodes are of the form (q, Z) where Z is a time elapsed local zone, and whose transitions are steps $(q, Z) \xrightarrow{b} (q', Z')$. The initial node $(q^{\text{init}}, Z^{\text{init}})$ consists of the initial state q^{init} of the network and the local zone $Z^{\text{init}} = \text{local-elapse}(\{v^{\text{init}}\})$.*

Directly from Lemma 12 we obtain the main property of local zone graphs stated in [4]. This allows us to use local zone graphs for reachability testing.

► **Theorem 14.** *For a given network of timed automata \mathcal{N} , there is a run of the network reaching a state q iff for some non-empty local zone Z , node (q, Z) is reachable in $\text{LZG}(\mathcal{N})$ from its initial node.*

Notice that $\text{LZG}(\mathcal{N})$ may still be infinite and it cannot be used directly for reachability checking. The solution in Remark 11 does not apply to local zones due to the multiple reference clocks. This problem will be addressed in Section 5. We first focus on important properties of the local-time zone graph w.r.t. concurrency.

4 Why are local zone graphs better than global zone graphs?

The important feature about local zone graphs, as noticed in [4], is that two transitions on actions with disjoint domains commute (see Figure 1).

► **Lemma 15** (Commutativity on local zones [4]). *Suppose $\text{dom}(a) \cap \text{dom}(b) = \emptyset$. If $(q, Z) \xrightarrow{ab} (q', Z')$ then $(q, Z) \xrightarrow{ba} (q', Z')$.*

From the above lemma, we get the following property.

► **Corollary 16.** *If $(q, Z) \xrightarrow{u} (q', Z')$ and $u \sim w$, then $(q, Z) \xrightarrow{w} (q', Z')$*

Thus starting from a local zone, all equivalent interleavings of a sequence of actions u end up in the same local zone. This is in stark contrast to the global zone graph, where each interleaving results in a possibly different global zone. Let

$$MZ(q, Z, u) = \{v' \mid \exists v \in Z, \exists w, w \sim u \text{ and } (q, v) \xrightarrow{w} (q', v')\}$$

denote the union of all these global zones. Salah et al. [17] have shown that, surprisingly, $MZ(q, Z, u)$ is always a global zone. We call it *aggregated zone*, and the notation MZ is in the memory of Oded Maler. In the same work, this observation was extended to an algorithm for *acyclic timed automata* that from time to time merged zones reached by equivalent paths to a single global zone. We prove below that this aggregated zone can, in fact, be obtained directly in the local zone graph: the aggregated (global) zone is exactly the set of synchronized valuations obtained after executing u in the local zone semantics. Here we need some notation: let Z be a global zone and Z a local zone; define $\text{sync}(Z) = \{v \in Z \mid v \text{ is synchronized}\}$; $\text{local}(Z) = \{\text{local}(v) \mid v \in Z\}$ and $\text{global}(\text{sync}(Z)) = \{\text{global}(v) \mid v \in \text{sync}(Z)\}$.

► **Lemma 17.** *For every global zone Z and local zone Z : $\text{sync}(Z)$ and $\text{local}(Z)$ are local zones and $\text{global}(\text{sync}(Z))$ is a global zone.*

Proof. $\text{sync}(Z)$ is the local zone $Z \wedge \bigwedge_{i,j}(t_i = t_j)$; $\text{local}(Z)$ is the local zone obtained by replacing t with some t_i in each constraint, and adding the constraints $\bigwedge_{i,j}(t_i = t_j)$; $\text{global}(\text{sync}(Z))$ is obtained by replacing each t_i with t in each constraint of Z . ◀

► **Theorem 18.** *Consider a state q , a sequence of actions u and a time elapsed global zone Z . Consider the local zone $Z = \text{local-elapse}(\text{local}(Z))$. If $(q, Z) \xrightarrow{u} (q', Z')$, we have $MZ(q, Z, u) = \text{global}(\text{sync}(Z'))$, otherwise $MZ(q, Z, u) = \emptyset$.*

Proof. Pick $v' \in MZ(q, Z, u)$. There exists $w \sim u$, $v \in Z$ and a global run $(q, v) \xrightarrow{w} (q', v')$. From Lemma 10, there exists a local run $(q, \text{local}(v)) \xrightarrow{w} (q', \text{local}(v'))$. By assumption, $\text{local}(v) \in Z$. Hence from the pre property of local zones (Lemma 12), there exists $(q, Z) \xrightarrow{w} (q', Z_w)$ such that $\text{local}(v') \in Z_w$. As $\text{local}(v')$ is synchronized, we get $\text{local}(v') \in \text{sync}(Z_w)$. But, by Corollary 16, $Z_w = Z'$. This proves $\text{local}(v') \in \text{sync}(Z')$ and hence $v' \in \text{global}(\text{sync}(Z'))$.

For the other direction take $v' \in \text{global}(\text{sync}(Z'))$. As $(q, Z) \xrightarrow{u} (q', Z')$, by post property of local zones (Lemma 12) there is a local run $(q, v_u) \xrightarrow{u} (q', \text{local}(v'))$ for some $v_u \in Z$. Since $v_u \in Z$, it is obtained by a local time elapse from some $v \in \text{local}(Z)$. Hence v is synchronized and $\text{global}(v) \in Z$. From Lemma 10 we get that for some $w \sim u$ there is a global run $(q, \text{global}(v)) \xrightarrow{w} (q', v')$. Hence $v' \in MZ(q, Z, u)$. ◀

Theorem 18 gives an efficient way to compute aggregated zones: it is sufficient to compute local zone graphs. Computing local zone graphs is not more difficult than computing global zone graphs. But, surprisingly, the combinatorial explosion due to interleaving does not occur in local zone graphs, thanks to the theorem above. Hence, this gives an incentive to work with local zone graphs instead of global zone graphs.

This contrasts with the aggregation algorithm in [17] which requires to store all the paths to a global zone and detect situations where zones can be merged, that is, when all the equivalent permutations have been visited. Another important limitation of the algorithm in [17] is that it can only be applied to acyclic zone graphs. If local zone graphs can be computed for general timed automata (which contain cycles), we can get to use the aggregation feature for all networks (and not only acyclic ones). To do this, there is still a major problem left: local zone graphs could be infinite when the automata contain cycles.

5 Making local zone graphs finite

The standard approach to make a global zone graph finite is to use a subsumption operation between global zones [2, 12]. Such a subsumption operation is usually based on a finite index simulation between (global) valuations parameterized by certain maximum constants

occurring in guards. We first discuss technical problems that arise when we lift these simulations to local valuations. In the paper introducing local time semantics and local zone graphs [4] a notion of a *catch-up equivalence* between local valuations is defined. This equivalence is a finite index simulation. So one could, in theory, construct a finite local zone graph using catch-up subsumption as a finite abstraction. Unfortunately, the question of effective algorithms for catch-up subsumption was left open in [4], and, to the best of our knowledge, there is no efficient procedure for catch-up subsumption.

Another finite abstraction of the local zone graph was proposed by Minea [15, 16]. We however believe that Minea’s approach carries a flaw, and a different idea is needed to get finiteness. Minea’s approach is founded on an equivalence between local valuations in the lines of the region equivalence [1]. Let c_{\max} be the maximum constant appearing among the guards of a timed automata network. Two local valuations \mathbf{v}_1 and \mathbf{v}_2 are said to be equivalent, written as $\mathbf{v}_1 \simeq_{\text{reg}} \mathbf{v}_2$ if for all variables $\alpha, \beta \in \tilde{X} \cup \{t_1, \dots, t_k\}$ (note that all reference clocks are included):

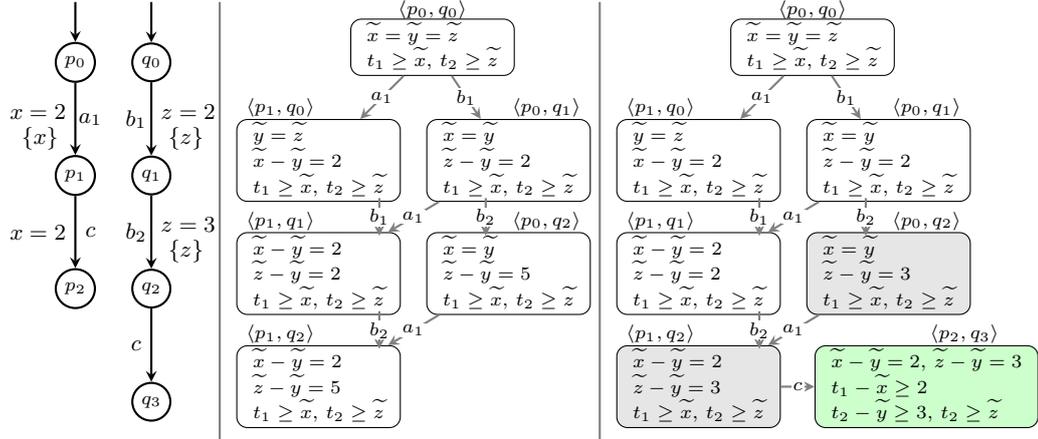
- either $\lfloor \mathbf{v}_1(\alpha) - \mathbf{v}_1(\beta) \rfloor = \lfloor \mathbf{v}_2(\alpha) - \mathbf{v}_2(\beta) \rfloor$,
- or $\lfloor \mathbf{v}_1(\alpha) - \mathbf{v}_1(\beta) \rfloor > c_{\max}$ and $\lfloor \mathbf{v}_2(\alpha) - \mathbf{v}_2(\beta) \rfloor > c_{\max}$,
- or $\lfloor \mathbf{v}_1(\alpha) - \mathbf{v}_1(\beta) \rfloor < -c_{\max}$ and $\lfloor \mathbf{v}_2(\alpha) - \mathbf{v}_2(\beta) \rfloor < -c_{\max}$.

It is claimed (in Proposition 6 of [15]) that this equivalence is preserved over local time elapse: for every process p , and for $\delta \geq 0$ there exists $\delta' \geq 0$ such that $\mathbf{v}_1 +_p \delta \simeq_{\text{reg}} \mathbf{v}_2 +_p \delta'$. However, this is not true, as we now exhibit a counter-example. Consider 2 processes with clocks $X_1 = \{x\}$, $X_2 = \{y\}$. This gives $\tilde{X} = \{\tilde{x}, \tilde{y}\}$ and $T = \{t_1, t_2\}$. Let $c_{\max} = 3$. Define local valuations $\mathbf{v}_1 : \tilde{x} = 0, t_1 = 0, \tilde{y} = 0, t_2 = 4$ and $\mathbf{v}_2 : \tilde{x} = 0, t_1 = 0, \tilde{y} = 0, t_2 = 5$.

Note that the differences in \mathbf{v}_1 are either 0, 4 or -4 and the corresponding differences in \mathbf{v}_2 are 0, 5 or -5 . Hence by definition, $\mathbf{v}_1 \simeq_{\text{reg}} \mathbf{v}_2$. Consider valuation $\mathbf{v}_1 +_1 2$ obtained from \mathbf{v}_1 by local delay of 2 units in component 1, that is $\mathbf{v}_1 +_1 2 : \tilde{x} = 0, t_1 = 2, \tilde{y} = 0, t_2 = 4$. Observe that in $\mathbf{v}_1 +_1 2$, the difference $t_1 - \tilde{x} = 2$ and $t_2 - t_1 = 2$ which are both smaller than c_{\max} . We claim there is no local delay δ' such that $\mathbf{v}_1 +_1 2 \simeq_{\text{reg}} \mathbf{v}_2 +_1 \delta'$. Valuation $\mathbf{v}_2 +_1 \delta'$ is given by $\tilde{x} = 0, t_1 = \delta', \tilde{y} = 0, t_2 = 5$. If $\mathbf{v}_1 +_1 2 \simeq_{\text{reg}} \mathbf{v}_2 +_1 \delta'$, we need $\delta' = 2$ (due to difference $t_1 - \tilde{x}$) and $5 - \delta' = 2$ (due to difference $t_2 - t_1$). This is not possible.

The main problem is that the above equivalence “forgets” actual values when the difference between reference clocks is above c_{\max} . Even if this difference is bigger than the maximum constant, local delays can bring them within the constant c_{\max} . Such a situation does not arise in the global semantics, as there is a single reference clock.

In [15] a widening operator on local zones based on \simeq_{reg} is used for finiteness: given a canonical representation of a zone Z , the *maximized zone* with respect to c_{\max} is obtained by changing every constraint $y_1 - y_2 < c$ to $y_1 - y_2 < \infty$ if $c > c_{\max}$, and to $y_1 - y_2 < -c_{\max}$ if $c < -c_{\max}$. In the local zone graph construction, each local zone is maximized and inclusion between maximized zones is used for termination. Figure 2 gives an example of a network $\langle A_1, A_2 \rangle$ where the maximized local zone graph is unsound. This is shown by making use of the valuations \mathbf{v}_1 and \mathbf{v}_2 above. We also add an extra clock z in component A_2 for convenience. Note that $c_{\max} = 3$. Although clock y does not appear in A_2 , one can assume that there are other transitions from q_0 that deal with y (we avoid illustrating these transitions explicitly). In the discussion below, $\mathbf{v}_1, \mathbf{v}_2$ are valuations restricted to $\tilde{x}, \tilde{y}, t_1$ and t_2 . In order to reach the state p_2 , the synchronization action c needs to be taken: transition sequence a_1c requires c to be taken at global time 4, and transition sequence b_1b_2c requires c at global time 5. Hence c is not enabled in the network. This is witnessed by c not being enabled in the local zone graph (middle picture). Valuation \mathbf{v}_2 is present in the zone reached after b_1b_2 .



■ **Figure 2** *Left*: network $\langle A_1, A_2 \rangle$; *Middle*: local zone graph; *Right*: maximized local zone graph of [15]. State (p_2, q_3) is not reachable in local zone graph, but becomes reachable after maximization.

The maximized local zone graph is shown on the right. Zones where maximization makes a difference are shaded gray. In particular, the zone $b_1 b_2$ on maximization adds valuation v_1 , from which $a_1 c$ is enabled, giving a zone in the maximized local zone graph with state p_2 .

5.1 Synchronized valuations for subsumption

A finite abstraction of the differences between reference clocks constitutes the main challenge in obtaining a finite local zone graph. We propose a different solution which bypasses the need to worry about such differences: *restrict to synchronized valuations for subsumption*.

Subsumptions over global zones are well studied [1, 2, 12]. Taking an off-the-shelf finite abstraction α and a subsumption \sqsubseteq^α between global zones, we want to do the following: given two local zones Z_1 and Z_2 we perform a subsumption test $\text{global}(\text{sync}(Z_1)) \sqsubseteq^\alpha \text{global}(\text{sync}(Z_2))$. By finiteness of the abstraction, we will get only finitely many local zones Z with incomparable $\text{global}(\text{sync}(Z))$. In order to do this, we need the crucial fact that $\text{global}(\text{sync}(Z))$ is a zone (shown in Lemma 17). Given two configurations $s := (q, Z)$ and $s' := (q', Z')$ of the local zone graph, we write $s \sqsubseteq_{\text{sync}}^\alpha s'$ if $q = q'$ and $\text{global}(\text{sync}(Z)) \sqsubseteq^\alpha \text{global}(\text{sync}(Z'))$.

► **Definition 19** (Local sync graph). *Let α be a finite abstraction over global zones. A local sync graph G of a network of timed automata \mathcal{N} based on α , is a tree and a subgraph of $\text{LZG}(\mathcal{N})$ satisfying the following conditions:*

- C0** every node of G is labeled either covered or uncovered;
- C1** the initial node of $\text{LZG}(\mathcal{N})$ belongs to G and is labeled uncovered;
- C2** every node is reachable from the initial node;
- C3** for every uncovered node s , all its successor transitions $s \xrightarrow{a} s'$ occurring in $\text{LZG}(\mathcal{N})$ should be present in G ;
- C4** for every covered node $s \in G$ there is an uncovered node $s' \in G$ such that $s \sqsubseteq_{\text{sync}}^\alpha s'$. A covered node has no successors.

The above definition essentially translates to this algorithm: explore the local zone graph say in a BFS fashion, and subsume (cover) using $\sqsubseteq_{\text{sync}}^\alpha$ (similar to algorithm for global zone graph as described in page 8). Local sync graphs are not unique, since the final graph depends on the order of exploration. Every local sync graph based on a finite abstraction α is finite. Theorem 20 below states that local sync graphs are sound and complete for reachability, and this algorithm is correct.

► **Theorem 20** (Soundness and completeness of local sync graphs). *A state q is reachable in a network of timed automata \mathcal{N} iff a node (q, Z) , with Z non-empty, is reachable from the initial node in a local sync graph for \mathcal{N} .*

Proof. If (q, Z) is reachable in a local sync graph then it is trivially reachable in the local zone graph and the (backward) implication follows from soundness of local zone graphs (Theorem 14). For the other direction which proves completeness of local sync graphs, let us take a global run: $(q_{init}, v_{init}) \xrightarrow{\delta_1, b_1} (q_1, v_1) \cdots \xrightarrow{\delta_n, b_n} (q_n, v_n)$. Take a local sync graph G , based on abstraction \mathbf{a} coming from a simulation \preceq . By induction on i , for every (q_i, v_i) we will find an uncovered node (q_i, Z_i) of G , and a synchronized local valuation $\mathbf{v}_i \in Z_i$ such that $v_i \preceq \mathbf{global}(\mathbf{v}_i)$. This proves completeness, since every reachable (q_i, v_i) will have a representative node in the local sync graph.

The induction base is immediate, so let us look at the induction step. Consider the global step $(q_i, v_i) \xrightarrow{\delta_i, b_i} (q_{i+1}, v_{i+1})$. Since $v_i \preceq \mathbf{global}(\mathbf{v}_i)$, there is a delay δ'_i such that $(q_i, \mathbf{global}(\mathbf{v}_i)) \xrightarrow{\delta'_i, b_i} (q_{i+1}, v'_{i+1})$ and $v_{i+1} \preceq v'_{i+1}$. As the global delay δ'_i can be thought of a sequence of local delays, we have local run $(q_i, \mathbf{v}_i) \xrightarrow{b_i} (q_{i+1}, \mathbf{v}'_{i+1})$, where $\mathbf{v}'_{i+1} = \mathbf{local}(v'_{i+1})$. Note that \mathbf{v}'_{i+1} is synchronized and $v_{i+1} \preceq \mathbf{global}(\mathbf{v}'_{i+1})$. From the pre-property of local zones (Lemma 12) there exists a transition $(q_i, Z_i) \xrightarrow{b_i} (q_{i+1}, Z'_{i+1})$ with $\mathbf{v}'_{i+1} \in Z'_{i+1}$, in fact, $\mathbf{v}'_{i+1} \in \mathbf{sync}(Z'_{i+1})$. If (q_{i+1}, Z'_{i+1}) is uncovered, take \mathbf{v}'_{i+1} for \mathbf{v}_{i+1} and Z'_{i+1} for Z_{i+1} (needed by the induction step). Otherwise, from condition C4, there is an uncovered node (q_{i+1}, Z''_{i+1}) such that $\mathbf{global}(\mathbf{sync}(Z'_{i+1})) \preceq \mathbf{global}(\mathbf{sync}(Z''_{i+1}))$. This gives $\mathbf{v}''_{i+1} \in \mathbf{sync}(Z''_{i+1})$ such that $\mathbf{global}(\mathbf{v}'_{i+1}) \preceq \mathbf{global}(\mathbf{v}''_{i+1})$. Now take \mathbf{v}''_{i+1} for \mathbf{v}_{i+1} and Z''_{i+1} for Z_{i+1} . ◀

6 Experiments

We have implemented the construction of local sync graphs in our prototype TChecker [11] and compared it with two implementations of the usual global zone graph method: TChecker and UPPAAL [13, 3], the state-of-the-art verification tool for timed automata. The three implementations use a breadth-first search with subsumption, and the \sqsubseteq_{sync}^a subsumption in the case of local sync graph. Table 1 presents results of our experiments on standard models from the literature (except “Parallel” that is a model we have introduced).

Local sync graphs yield no gain on 3 standard examples (which are not given in Table 1): “CSMA/CD”, “FDDI” and “Fischer”. In these models, the three algorithms visit and store the same number of nodes. The reason is that for “CSMA/CD” and “FDDI”, replacing each local zone Z in the local zone graph by its set of synchronized valuations $\mathbf{sync}(Z)$ yields exactly the zone graph. In the third model, “Fischer”, every control state appears at most once in the global zone graph. So there is no hope to achieve any gain with our technique. This is due to the fact that doing ab or ba results in two different control states in the automaton. So “Fischer” is out of the scope of our technique.

In contrast, we observe significant improvements on other standard models (Table 1). Observe that due to subsumption, the order in which nodes are visited impacts the total number of visited nodes. UPPAAL and our prototype TChecker (Global ZG column) may not visit the same number of nodes despite the fact that they implement the same algorithm. In our prototype we use the same order of exploration for Global ZG and Local ZG. “CorSSO” and “Critical region” are standard examples from the literature. “Dining Philosophers” is a modification of the classical problem where a philosopher releases her left fork if she cannot take her right fork within a fixed amount of time [14]. “Parallel” is a model we have introduced, where concurrent processes compete to access a resource in mutual exclusion.

■ **Table 1** Experimental results obtained by running UPPAAL and our prototype TChecker (Global ZG and Local ZG) on a MacBook Pro 2013 with 4 2.4GHz Intel Core i5 and 16 GB of memory. The timeout is 90 seconds. For each model we report the number of concurrent processes.

Models (# processes)	UPPAAL			Global ZG			Local ZG		
	visited	stored	sec.	visited	stored	sec.	visited	stored	sec.
CorSSO 3	64378	61948	1.48	64378	61948	1.41	1962	1962	0.05
CorSSO 4		timeout			timeout		23784	23784	0.69
CorSSO 5		timeout			timeout		281982	281982	16.71
Critical reg. 4	78049	53697	1.45	75804	53697	2.27	44490	28400	2.40
Critical reg. 5		timeout			timeout		709908	389614	75.55
Dining Phi. 7	38179	38179	34.61	38179	38179	7.28	2627	2627	0.32
Dining Phi. 8		timeout			timeout		8090	8090	1.65
Dining Phi. 9		timeout			timeout		24914	24914	7.10
Dining Phi. 10		timeout			timeout		76725	76725	30.20
Parallel 6	11743	11743	4.82	11743	11743	1.09	256	256	0.02
Parallel 7		timeout			timeout		576	576	0.04
Parallel 8		timeout			timeout		1280	1280	0.11

We observe an order of magnitude gains for most of these four models. The reason is that in most states when two processes can perform actions a and b , doing ab or ba leads to the same control-state of the automaton. Hence, a difference between ab and ba (if any) is encoded in distinct zones Z_{ab} and Z_{ba} obtained along these two paths in the global zone graph. In contrast, the two paths result in the same zone Z (containing both Z_{ab} and Z_{ba}) in the local sync graph. In consequence, our approach that combines the local zone graph and abstraction using synchronized zones is very efficient in this situation.

7 Conclusions

We have revisited local-time semantics of timed automata and local zone graphs. We have discovered a very useful fact that local zones calculate aggregated zones: global zones that are unions of all the zones obtained by equivalent executions [17]. We have used this fact as a theoretical foundation for an algorithm constructing local zone graphs and using subsumption on aggregated zones at the same time.

We have shown that, unfortunately, subsumption operations on local zones proposed in the literature do not work. We have proposed a new subsumption for local zone graphs based on standard abstractions for timed automata, applied on synchronized zones. The restriction to synchronized zones is crucial as standard abstractions cannot handle multiple reference clocks. A direction for future work is to find abstractions for local zones.

Our algorithm is the first efficient implementation of local time zone graphs and aggregated zones. Experimental results show an order of magnitude gain with respect to state-of-the-art algorithms on several standard examples.

As future work, we plan to develop partial-order techniques taking advantage of the high level of commutativity in local zone graphs. Existing methods are not directly applicable in the timed setting. In particular, contrary to expectations, actions with disjoint domains may not be independent (in a partial order sense) in a local zone graph [15]. Thus, it will be very interesting to understand the structure of local zone graphs better. A recent partial-order method proposed for timed-arc Petri nets [6] gives a hope that such obstacles can be overcome. For timed networks with cycles, the interplay of partial-order and subsumption adds another level of difficulty.

References

- 1 Rajeev Alur and David L. Dill. A Theory of Timed Automata. *Theoretical Computer Science*, 126:183–235, 1994.
- 2 Gerd Behrmann, Patricia Bouyer, Kim G. Larsen, and Radek Pelánek. Lower and Upper Bounds in Zone-Based Abstractions of Timed Automata. *International Journal on Software Tools for Technology Transfer*, 8(3):204–215, 2006.
- 3 Gerd Behrmann, Alexandre David, Kim Guldstrand Larsen, John Hakansson, Paul Pettersson, Wang Yi, and Martijn Hendriks. UPPAAL 4.0. In *QEST*, pages 125–126. IEEE Computer Society, 2006.
- 4 Johan Bengtsson, Bengt Jonsson, Johan Lilius, and Wang Yi. Partial Order Reductions for Timed Systems. In *CONCUR*, volume 1466 of *Lecture Notes in Computer Science*, pages 485–500, 1998.
- 5 Johan Bengtsson and Wang Yi. Timed Automata: Semantics, Algorithms and Tools. In *ACPN 2003*, volume 3098 of *Lecture Notes in Computer Science*, pages 87–124. Springer, 2003.
- 6 Frederik M. Bønneland, Peter Gjøøl Jensen, Kim Guldstrand Larsen, Marco Muñoz, and Jiri Srba. Start Pruning When Time Gets Urgent: Partial Order Reduction for Timed Systems. In *CAV*, volume 10981 of *Lecture Notes in Computer Science*, pages 527–546. Springer, 2018.
- 7 Conrado Daws and Stavros Tripakis. Model Checking of Real-Time Reachability Properties Using Abstractions. In *TACAS*, volume 1384 of *Lecture Notes in Computer Science*, pages 313–329. Springer, 1998.
- 8 D. L. Dill. Timing Assumptions and Verification of Finite-state Concurrent Systems. In *Proceedings of the International Workshop on Automatic Verification Methods for Finite State Systems*, pages 197–212, 1990.
- 9 Paul Gastin, Sayan Mukherjee, and B. Srivathsan. Fast algorithms for handling diagonal constraints in timed automata. *CoRR*, abs/1904.08590, 2019. [arXiv:1904.08590](https://arxiv.org/abs/1904.08590).
- 10 R Govind, Frédéric Herbreteau, B. Srivathsan, and Igor Walukiewicz. Revisiting local time semantics for networks of timed automata. *CoRR*, abs/1907.02296, 2019. [arXiv:1907.02296](https://arxiv.org/abs/1907.02296).
- 11 F. Herbreteau and G. Point. TChecker. <https://github.com/fredher/tchecker>, v0.2 - April 2019.
- 12 Frédéric Herbreteau, B. Srivathsan, and Igor Walukiewicz. Better Abstractions for Timed Automata. In *LICS*, pages 375–384. IEEE Computer Society, 2012.
- 13 Kim Guldstrand Larsen, Paul Pettersson, and Wang Yi. UPPAAL in a nutshell. *STTT*, 1(1-2):134–152, 1997.
- 14 Denis Lugiez, Peter Niebert, and Sarah Zennou. A partial order semantics approach to the clock explosion problem of timed automata. *Theor. Comput. Sci.*, 345(1):27–59, 2005.
- 15 Marius Minea. Partial Order Reduction for Model Checking of Timed Automata. In *CONCUR*, volume 1664 of *Lecture Notes in Computer Science*, pages 431–446. Springer, 1999.
- 16 Marius Minea. *Partial Order Reduction for Verification of Timed Systems*. PhD thesis, School of Computer Science, Carnegie Mellon University Pittsburgh, PA 15213, 1999.
- 17 Ramzi Ben Salah, Marius Bozga, and Oded Maler. On Interleaving in Timed Automata. In *CONCUR*, volume 4137 of *Lecture Notes in Computer Science*, pages 465–476. Springer, 2006.
- 18 Marcelo Sousa, César Rodríguez, Vijay D’Silva, and Daniel Kroening. Abstract Interpretation with Unfoldings. In *CAV*, pages 197–216, 2017.