



HAL
open science

Abstractions for the local-time semantics of timed automata: a foundation for partial-order methods

R. Govind, Frédéric Herbreteau, B. Srivathsan, Igor Walukiewicz

► To cite this version:

R. Govind, Frédéric Herbreteau, B. Srivathsan, Igor Walukiewicz. Abstractions for the local-time semantics of timed automata: a foundation for partial-order methods. 37th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2022, Aug 2022, Haifa, Israel. hal-03644039v2

HAL Id: hal-03644039

<https://hal.science/hal-03644039v2>

Submitted on 2 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Abstractions for the local-time semantics of timed automata: a foundation for partial-order methods

R. Govind, Frédéric Herbreteau, B. Srivathsan and Igor Walukiewicz

Abstract

A timed network is a parallel composition of timed automata synchronizing on common actions. We develop a methodology that allows to use partial-order methods when solving the reachability problem for timed networks. It is based on a local-time semantics proposed by [Bengtsson et al. 1998]. A new simulation based abstraction of local-time zones is proposed. The main technical contribution is an efficient algorithm for testing subsumption with respect to this abstraction operator. The abstraction is not finite for all networks. It turns out that, under relatively mild conditions, there is no finite abstraction for local-time zones that works for arbitrary timed networks. To circumvent this problem, we introduce a notion of a bounded-spread network. The spread of a network is a parameter that says how far the local times of individual processes need to diverge. For bounded-spread networks, we show that it is possible to use subsumption and partial-order methods at the same time.

1 Introduction

The reachability problem for timed automata [3] is to decide if a given automaton has an execution from an initial to a final state. Very frequently a model is given as a network of timed automata working in parallel and synchronizing on common actions. It is tempting to exploit the concurrency information provided by such a representation to speed up reachability testing. For untimed systems, partial-order methods [2, 15, 19, 20, 22, 35, 37, 38] can provide exponential improvements. The presence of time greatly complicates the picture because individual automata may synchronize implicitly via time. In this work we extend the classical zone based approach to the reachability problem so that partial-order reduction methods become applicable.

Let us explain the challenge on a simple example. Figure 1a shows a network of two processes. The first process does a local action b , the second a local action c , and then they synchronize on action $\$$. If we ignore the timing constraints, the graph of all executions of this system has a *diamond*: since b and c are executed on different processes they are independent, so the sequence bc leads to the same state as cb as shown in Figure 1b. The timing constraints break this diamond: the sequence cb is impossible since doing c requires to wait at least 2 time units, and then it is too late for doing b that needs to be executed within 1 time unit from the




Figure 1. A network of two processes. Timing constraints break the diamond formed by two independent actions.

start, as in Figure 1c. This is a major obstacle for applying partial-order methods in the timed automata setting.

Before addressing this obstacle let us review how the reachability problem of a single timed automaton is solved. Most efficient solutions to the reachability problem construct an explicit graph, that we call here an *abstract zone graph with subsumptions*. *Zones* [17] are special convex sets of clock valuations with the property that a set of valuations reachable from a zone is once again a zone. The nodes of the graph represent the zones that are reachable by the transitions of the automaton. For some timed automata, there could be infinitely many reachable zones. This is why *abstractions* [6, 17], such as $Extra_{LU}$ or $\alpha_{\leq LU}$, are used to "abstract" a finite number of representative sets. Finally, only zones whose abstractions are maximal with respect to inclusion are kept during exploration. This technique is called *subsumption* [17], and it is essential for efficiency.

When applying this method to networks of timed automata, the state explosion problem occurs. For untimed systems this problem can be alleviated either by partial-order methods, or by symbolic methods based on BDDs or SAT-solving. For timed systems BDD and SAT based solutions [4, 5, 7, 18, 33, 34, 39] have been tried with mixed results. Here we pursue a partial-order approach to tackle the state explosion problem.

Partial-order methods limit the search space by using the diamonds present in the graph of executions. In our example from Figure 1b, it is enough to explore the sequence $bc\$$ as thanks to the diamond we are sure that the sequence $cb\$$ leads to the same state. For more complicated cases, this approach can give exponential gains in time as well in the size of the graph to be explored. In order to apply partial-order




Figure 2. Diamonds are recovered in local-time semantics. In the rightmost path the local times of the two process differ.

methods for timed systems it is essential to recover the diamonds lost due to implicit synchronization caused by time constraints. Otherwise, choosing $cb\$$ as a representative for $bc\$$ in Figure 1c would lead to incompleteness. Solutions proposed in the literature consider only diamonds where time does not elapse [9, 10, 29], or try to deduce which diamonds are still bound to stay despite time constraints [16, 24]. Here we develop a set of theoretical results permitting a much wider use of partial-order methods in constructing abstract zone graphs with subsumption.

Our starting point is the local-time semantics [8] for networks of timed automata, that addresses exactly the diamond problem by making time local to each process. The processes are required to synchronize their times when performing common actions. As a result, local-time semantics can be actually used to solve the reachability problem despite allowing more behaviours than the standard global-time semantics. Moreover, actions executing on different processes are independent as there are no implicit synchronizations on time.

Let us revisit our example to see how diamonds are recovered thanks to local-time. Figure 2 illustrates the graph of executions under the local-time semantics. The two processes have their local and independent times represented by clocks t_1 and t_2 , respectively. The path $bc\$$ is still feasible as before, by keeping the local times of process P_1 and P_2 synchronized. But now, $cb\$$ becomes feasible as well. P_2 may delay by 2 time units and do c while P_1 does not delay at all. Then, P_1 can do b , and then delay 2 time units to synchronize its time with P_2 and enable the common action $\$$. As in the standard (global-time) setting, there is a notion of a *local-zone*, and one can try to use the local-zone graph for checking reachability. However, this graph may be infinite.

What is lacking to make the local-time approach algorithmically interesting, is an efficient abstraction operator that would guarantee finiteness of abstract local-zone graphs. An abstraction operator has been proposed in [8], but as we show here, the associated decision problem is PSPACE-hard (Proposition 2), so there is little hope that it can be used to give an efficient solution.

To sum up, to be able to use partial-order methods with the help of local-time semantics, we need to find an abstraction operator for local zones that:

1. preserves reachability,
2. leads to a finite abstract local-zone graph,
3. is efficient algorithmically,
4. preserves diamonds of actions from distinct processes.

The first two conditions are required for correctness and termination of an exploration algorithm. The third is essential to be competitive with existing solutions: computing an inclusion between two abstracted zones should be easier than solving the reachability problem in the first place. The fourth condition is needed to apply partial-order methods. The formalization of the fourth condition is actually weaker than requiring diamonds to exist in the abstract local-zone graph with subsumptions. The latter property would be much too strong to demand; c.f. Figure 3.

Our first result is an extension of the well-known $\alpha_{\leq LU}$ abstraction for global-time semantics [6, 26] to the local-time setting (Theorem 5). We call it $\alpha_{\leq LU}^*$. The main technical result is an efficient algorithm for testing inclusion $\alpha_{\leq LU}^*(Z) \subseteq \alpha_{\leq LU}^*(Z')$ in time $O((|X| + n)^2)$, where $|X|$ is the number of clocks, and n is the number of processes in the network (Theorem 6). This complexity is essentially the same as in the global-time setting, with the factor n coming due to extra clocks added by the local-time semantics.

Unfortunately, the $\alpha_{\leq LU}^*$ abstraction is not finite, that is, it does not satisfy property (2). Actually, we observe a strong negative result: there is no simulation based abstraction operator satisfying properties (1), (2) and (4) at the same time (Theorem 3). This is a serious obstacle because we do not know how to guarantee (4) for abstractions that are not simulation based. To the best of our knowledge, all abstractions used in timed automata verification algorithms are simulation based. The main hindrance to get finiteness is that the local times of processes can drift from each other by arbitrary amounts, but this quantity cannot be abstracted away.

Given this roadblock, we propose a restricted setting of *bounded-spread networks*. These are networks where the drift between processes can be controlled: every sequence of actions can be realized while maintaining a bounded drift between processes. For such networks, a suitable adaptation of the $\alpha_{\leq LU}^*$ abstraction becomes finite, and has all the required four properties (Theorem 7).

The final step is to apply partial-order methods to the finite abstract local-zone graph with subsumptions. This is slightly delicate because the abstract local-zone graph with subsumptions does not have diamonds, precisely due to subsumptions (cf. Figure 3). Yet, we do not want to disallow subsumptions as they are essential to get an effective and an algorithmically efficient solution. We show that every partial-order method that works on graphs without subsumptions, intuitively for untimed systems, can be used for bounded-spread networks (Theorem 2). Abstractly, we see a partial-order method as computing a function src indicating for every state a subset of its outgoing transitions, such that exploring the smaller

set of transitions is sufficient to verify reachability. In our example from Figure 1b we may have $src(0, 0) = \{b\}$ indicating that it is sufficient to explore only the transition b from the initial node. Since the $\alpha_{\leq LU}^*$ abstraction is based on a simulation, we can show that if the src function is correct for the local-zone graph (without subsumptions), it is also correct to use it for the abstract local-zone graph with subsumptions, even though the latter does not have diamonds.

Putting these results together we obtain a methodology allowing to apply existing partial-order methods to timed-systems. The methodology is not general because it applies only to networks of bounded spread, while in general networks could have unbounded spread. Moreover, computing a spread of a given network is at least as difficult as testing reachability. On a positive side, we give examples of some types of networks that are guaranteed to have a bounded spread. We also propose a method to convert an arbitrary network into a bounded-spread network by introducing synchronizations between processes. We conclude with simple examples where our method can bring exponential gains.

Related work. Local-time semantics has been considered by three groups. Bengtsson et al. in their paper introducing local-time semantics [8] propose an algorithm for reachability checking. For this they introduce an abstraction called catch-up equivalence. It is rather improbable that an algorithm using this equivalence can be competitive against standard solutions because, as we show here, checking if two valuations are catch-up equivalent is PSPACE-hard. In [32] another equivalence is proposed, but it turns out to be not sound [23]. Paper [23] introduces sync-subsumption, but this subsumption does not preserve diamonds, hence it is not suitable for partial-order reduction.

An alternative to local-time zones was proposed by Lugiez et al. [31]. In that approach zones maintain a partial-order between clocks. To get finiteness an abstraction similar to sync-subsumption of [23] is used. Once again, this does not preserve diamonds and hence this approach is not suitable for partial-order reduction on the control states.

Other works have proposed partial-order methods for timed automata, while keeping the standard semantics. For example, limiting partial-order methods only to parts where independent actions occur in zero-time [11, 29, 33]. Some works propose ways to discover which actions remain independent despite time constraints, either statically [16] or dynamically [24]. Two works [12, 13] apply unfolding techniques to bounded timed automata which admit a finite representation of their state space without abstraction.

Partial-order methods have been introduced in the 90s [22, 35, 37] as a method to speed up verification of transition systems. Later the accent shifted to program verification, and in particular to stateless model-checking [21]. The subject has become very active since the work of Abdulla et al. [1, 2] introducing a notion of optimal partial order reduction

(see [14, 28, 40] and references within). In this paper we take an abstract view of partial-order methods and do not focus on any concrete methods. The most recent works need some adaptation to be applicable in our setting. One reason is that they consider only straight-line processes, i.e., without branching. This is too restrictive in our setting.

Synopsis. In the Preliminaries, we introduce local-time semantics, local-zone graphs and their most important properties. We also present a succinct description of partial-order methods that is sufficient for this work. In Section 3 we introduce a notion of abstraction for local-zone graphs, and study conditions under which an abstraction can be used for reachability. In Section 4 we show how partial-order methods can be used in the presence of abstractions. Unfortunately, under mild assumptions, abstractions of local-zone graphs compatible with partial-order cannot be finite (Section 5). Our solution is to put a restriction on timed networks, but before this we develop in Section 6 an abstraction operator $\alpha_{\leq LU}^*$, which is a generalization of the well-known $\alpha_{\leq LU}$ operator. We show that inclusion testing with respect to the new operator $\alpha_{\leq LU}^*$ can be done efficiently. In Section 7 we introduce bounded-spread networks, give examples of such networks, as well as a general construction transforming a timed network into a bounded-spread network. We show that a modification of $\alpha_{\leq LU}^*$ is finite on bounded-spread networks. We conclude with some examples where our method gives exponential gains and discuss where they come from.

2 Preliminaries

In this section we introduce networks of timed automata, local-time semantics, and partial-order reduction. We present the standard global-time semantics as a special case of the local-time semantics. This clearly shows the differences between the two. Our approach allows to transfer any partial-order method from the untimed setting to the timed setting. In this paper, a partial-order method is given as an oracle that tells which transitions need to be explored. The only constraint is that the method keeps at least one execution from each trace equivalence class. At the end of the section, we introduce local-zone graphs, and state their properties.

We use \mathbb{N} for the set of natural numbers, \mathbb{R} for the set of reals and $\mathbb{R}_{\geq 0}$ for 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 $\sim \in \{<, \leq, =, \geq, >\}$. The base constraints $x \sim c$ will be called *atomic constraints*.

A *network of timed automata* is a collection of timed automata communicating with each other via shared actions. We have seen an example of a network in Figure 1a. Each automaton participating in the network is called a *process*. Formally, a timed network is a k -tuple of processes $\mathcal{N} = \langle A_1, \dots, A_k \rangle$. 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 . 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 $p_1 \neq p_2$. The sets of labels need not be disjoint - a label shared by two processes represents an action synchronizing the processes. The remaining components are an initial state q_p^{init} and a set of transitions $T_p \subseteq (Q_p \times \Sigma_p \times \phi(X_p) \times 2^{X_p} \times Q_p)$. A transition $(q, b, g, R, q') \in T_p$ has a label b , a *guard* g , and a set R of clocks to be *reset*. We write $Proc$ for the set of all processes. We will use some abbreviations: $Q = \prod_{p=1}^k Q_p$, $\Sigma = \bigcup_{p=1}^k \Sigma_p$ and $X = \bigcup_{p=1}^k X_p$. For a tuple of states $q \in Q$, we write $q(p)$ to be the state of process p in the tuple q . Every action b has its domain $dom(b) = \{p : b \in \Sigma_p\}$. The execution of action b requires participation of all processes in the domain. We denote by q^{init} the tuple of initial states q_p^{init} for each process p .

Local-time semantics. We introduce the local-time semantics of timed automata [8], and then the standard global-time semantics as a particular case. Fix a timed network \mathcal{N} .

In the local-time semantics, each process p has its local time represented by a clock t_p . The processes synchronize their times when doing a common action. The clock t_p , called the *reference clock* of process p , is never tested in a guard nor reset by the process. We will denote by X^t the set $\{t_p \mid p \in Proc\}$ of reference clocks. The other clock variables will store the local-time when the clock was last reset. Thus the value of t_p cannot be smaller than values of clocks of process p . More formally, a *local valuation* assigns a value, a real number, to each clock in $X \cup X^t$:

$$v : (X \cup X^t) \rightarrow \mathbb{R} \quad \text{provided } v(t_p) \geq v(x) \text{ for } x \in X_p$$

With this intuition, the difference $v(t_p) - v(x)$ gives the time since the last reset of clock x . This is what is considered as the value of x in the standard semantics. In the local-time semantics we allow negative values for $v(t_p), v(x)$ since we will always work with the difference $v(t_p) - v(x)$ and allowing for negative values offers some simplicity later while handling zones of local valuations. We use $LocalVal(X, Proc)$ for the set of local valuations, but mostly we will just write $LocalVal$ as X and $Proc$ will be clear from the context.

Operations of clock reset for local valuations as well as local time elapse are defined accordingly, based on the interpretation given above. For a set of clocks R , let $v[R]$ denote the local valuation obtained by resetting R in v . That is: $v[R](x) = v(t_p)$ if $x \in R \cap X_p$ for some $p \in Proc$, and $v[R](x) = v(x)$ otherwise. For a tuple of non-negative reals $\Delta = \{\delta_p \in \mathbb{R}_{\geq 0}\}_{p \in Proc}$ we define $v \xrightarrow{\Delta} v'$ when $v'(t_p) = v(t_p) + \delta_p$ for all $p \in Proc$, and $v'(x) = v(x)$ for all $x \in X$. This denotes a *local delay* of Δ from the valuation v . The notion of a local valuation satisfying a guard is also adapted to this interpretation. For x a clock of process p , i.e. $x \in X_p$, we define $v \models x \sim c$ if $v(t_p) - v(x) \sim c$ for $\sim \in \{<, \leq, =, \geq, >\}$.

Remark: One may wonder why not just keep the value of the clock in $v(x)$. This interpretation gives a big problem

later when we consider zones of local valuations. It turns out that in this interpretation the set of valuations reachable by a transition from a zone may not be a zone. Quite remarkably the interpretation presented above avoids this problem [8].

A configuration of the network is a pair (q, v) where q is a tuple of control states of all processes, and v is a local valuation. An *initial valuation* v_0 associates the same real to each clock: $v_0(r - s) = 0$ for all $r, s \in X \cup X^t$. Let V_0 denote the set of initial local valuations. The initial configurations are $\{q^{init}\} \times V_0$. For an action b , the network \mathcal{N} can execute a transition $(q, v) \xrightarrow{b} (q', v')$ if there is a tuple of b -transitions $\{(q_p, b, g_p, R_p, q'_p)\}_{p \in dom(b)}$ such that:

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

A *local run* of \mathcal{N} is a sequence of local delay and action transitions from an initial configuration (q_0, v_0) :

$$(q_0, v_0) \xrightarrow{\Delta_0} (q_0, v'_0) \xrightarrow{b_1} (q_1, v_1) \xrightarrow{\Delta_1} \dots \xrightarrow{b_n} (q_n, v_n) \xrightarrow{\Delta_n} (q_n, v'_n)$$

We write $(q_0, v_0) \xrightarrow{u} (q_n, v'_n)$ to say that there is a sequence as above for $u = b_1 \dots b_n$ and adequate delays.

Global-time semantics and reachability. The standard semantics of a network, which we refer to as global-time semantics or just global semantics in short, is given by the semantics of the monolithic timed automaton obtained as the “synchronized product” of the individual processes. There is a common time for all processes: their reference clocks are always equal. In other words global semantics uses only *synchronized valuations* v where $v(t_p) = v(t_q)$ for all $p, q \in Proc$. In consequence, in the global semantics, we only allow *global delays* $v \xrightarrow{\Delta} v'$ which are local delays such that $\delta_p = \delta_q$ for any two processes $p, q \in Proc$. We use δ for global delays to distinguish from local delays Δ . A *global run* of \mathcal{N} is an alternating sequence of global delay and action transitions starting from a synchronized valuation: $(q_0, v_0) \xrightarrow{\delta_0} (q_0, v'_0) \xrightarrow{b_1} (q_1, v_1) \xrightarrow{\delta_1} (q_1, v'_1) \xrightarrow{b_2} \dots \xrightarrow{b_n} (q_n, v_n) \xrightarrow{\delta_n} (q_n, v'_n)$. Observe that all the valuations on a global run are synchronized.

The *reachability problem* asks if a state q_f is reachable in the global-time semantics. In other words, does there exist a global run from an initial configuration to a configuration (q_f, v) for some synchronized valuation v . This problem is known to be PSPACE-complete [3]. Most algorithms solving the reachability problem use the global semantics [6, 17, 25, 26]. The state q_f that we check for reachability is called the *final state* of the network \mathcal{N} in the sequel.

Partial-order reduction (POR). We give a general outline of partial-order reductions that is sufficient for this work. The main idea is to use information about concurrency to avoid

exploring equivalent interleavings of actions. In Figure 1, we have seen that the order of execution between b and c is irrelevant: starting from $(0, 0)$ both sequences end in $(1, 1)$. We say that b and c are *independent* in a transition system \mathcal{S} if this property holds for every state of \mathcal{S} . The notion of independence leads to trace equivalence on sequences: two sequences are trace equivalent, denoted $u \sim_{\mathcal{S}} v$, if one can be obtained from the other by permuting adjacent independent actions. This is an equivalence relation on sequences of actions. Moreover, if u leads from an initial to a final state in \mathcal{S} then so does v . A POR method aims at exploring at least one path from every trace-equivalence class, but preferably not much more. For instance in Figure 1 we may only explore the sequence $bc\$,$ and ignore the sequence $cb\$.$ This avoids visiting state $(0, 1)$. In some cases this optimization may lead to exponential reductions in the number of visited states.

We think of a POR method as a way of computing for a given transition system \mathcal{S} a *source function*, $src : Q \rightarrow \mathcal{P}(\Sigma)$ assigning to every state of \mathcal{S} a set of relevant actions. A path in \mathcal{S} is a *source path* if it is a path in the restriction of \mathcal{S} where from every state q we eliminate transitions on actions that are not in $src(q)$. A source function should be *trace-faithful* meaning that for every state q and every path $q \xrightarrow{u} q_f$ to a final state q_f , there must be a trace-equivalent sequence $v \sim_{\mathcal{S}} u$ such that $q \xrightarrow{v} q_f$ is a source path. In the example from Figure 1 we may take $src(0, 0) = src(0, 1) = \{b\}$, $src(1, 0) = \{c\}$ and $src(1, 1) = \{\$\}$. The goal is to find a trace-faithful source function without exploring the transition system.

A common way to get a src function is to look at the parallelism in a given system. In a network of automata without timing constraints, two actions with disjoint domains are independent in the sense of the previous paragraph. Stubborn sets [37], ample sets [35], persistent sets [22], faithful decompositions [27], stamper sets [36], source sets [2], are different ways of computing a source function in this setting.

Our goal in this work is to develop a theory allowing to use the same approach for networks of timed automata. As seen in Figure 1c, in timed networks, two domain-disjoint actions may not be independent. Global time destroys diamonds, making it difficult to find out which actions are independent. Local-time semantics allows to recover diamonds, cf. Figure 2. Reachability can be solved using local-time, as we see next.

Reachability and diamonds in local-time. Observe that every global run is a local-time run. Conversely, for every local-time run there is a trace equivalent global run.

Lemma 1. [23] *Let v, v' be synchronized local valuations, and let $(q, v) \xrightarrow{u} (q', v')$ be a local run. Then there exists a global run $(q, v) \xrightarrow{w} (q', v')$ such that $u \sim w$.*

Since the initial valuations are all synchronized, the above lemma ensures that a control state q is reachable in the local-time semantics iff it is reachable in the global-time semantics.

This is particularly true of the final state q_f . Given this correspondence, we will henceforth work completely with the local-time semantics. Additionally, the local-time semantics offers the diamond property which is essential for POR.

Lemma 2 (Diamond property). *Suppose $dom(a) \cap dom(b) = \emptyset$. If $(q, v) \xrightarrow{ab} (q', v')$ then $(q, v) \xrightarrow{ba} (q', v')$.*

Local-zone graphs. To make the local-time semantics feasible for use in algorithms, a notion of local-zones, analogous to the zones in the global-time setting [17], is employed. A *local-zone* is a set of local valuations given by conjunctions of constraints: $x - y \leq c$ where $x, y \in X \cup X^t$ and $\leq \in \{<, \leq\}$. For a set of local valuations W , define:

- $local\text{-elapse}(W) := \{v + \Delta \mid v \in W, \Delta \in \mathbb{R}_{\geq 0}^k\}$,
- $W[R] := \{v[R] \mid v \in W\}$, for a set of clocks $R \subseteq X$.
- $W \cap g := \{v \mid v \models g\}$ for a guard g .

It can be shown that for a local-zone Z , the sets $local\text{-elapse}(Z)$ (local-time delay), $Z[R]$ (clock reset) and $Z \cap g$ (intersection with guard) are local-zones [8, 23].

Local-zones can be implemented using Difference Bound Matrices (DBMs), similar to the case of standard zones. Hence, they can be computed and stored as efficiently as standard zones. Before defining the local zone graph, we lift the local semantics from configurations to sets of configurations.

Definition 1 (Symbolic transition relation). Let W be a set of local valuations. We write $(q, W) \xrightarrow{b} (q', W')$ if there exists a tuple of b -transitions $\{(q_p, b, g_p, R_p, q'_p)\}_{p \in dom(b)}$ such that

- $q(p) = q_p$ and $q'(p) = q'_p$ for all $p \in dom(b)$, and $q(p) = q'(p)$ for all $p \notin dom(b)$;
- $W' = local\text{-elapse}(W_2)$ is not empty, where W_2 is defined as follows: $W_2 = W_1[\bigcup_{p \in dom(b)} R_p]$ and $W_1 = W \cap (\bigwedge_{p \in dom(b)} g_p \wedge \bigwedge \{t_p = t_q \mid p, q \in dom(b)\})$

We write $(q, W) \xrightarrow{b_1 \dots b_n} (q_n, W_n)$ if there is a sequence of symbolic transitions $(q, W) \xrightarrow{b_1} (q_1, W_1) \dots \xrightarrow{b_n} (q_n, W_n)$.

The following lemma states the relation between transitions on zones and on valuations. Its proof follows from the definition of symbolic transitions. We say that a local zone Z is *time-elapsd* if $Z = local\text{-elapse}(Z)$.

Lemma 3 (Pre and post properties). *For every network of timed automata and every action b :*

pre-property: *If $(q, v) \xrightarrow{b} (q', v')$ and $v \in Z$ for some time-elapsd local-zone Z then $(q, Z) \xrightarrow{b} (q', Z')$ and $v' \in Z'$ for some local-zone Z' .*

post-property: *If $(q, Z) \xrightarrow{b} (q', Z')$ and $v' \in Z'$ for local-zones Z, Z' , then $(q, v) \xrightarrow{b} (q', v')$ for some $v \in Z$.*

Definition 2 (Local-zone graph $LZG(\mathcal{N})$). The local-zone graph $LZG(\mathcal{N})$ of a network \mathcal{N} is a transition system whose

nodes are of the form (q, Z) where q is a state of the network, and Z is a local-zone. The initial node is (q_0, Z_0) with $Z_0 = \text{local-elapsed}(V_0)$ where V_0 is the set of initial valuations and $q_0 = q^{init}$. The transitions are given by the symbolic transition relation $(q, Z) \xRightarrow{b} (q', Z')$.

The initial zone is time-elapsed. This entails that every zone reachable by \Rightarrow transitions is also time-elapsed, due to Definition 1. Using this observation along with the pre- and post-properties of Lemma 3, we get the following theorem.

Theorem 1. [8, 23] *For a given network \mathcal{N} , there is a run of \mathcal{N} reaching a state q iff there is a path in $\text{LZG}(\mathcal{N})$ from the initial node to a node (q, Z) .*

This theorem suggests that the local-zone graph $\text{LZG}(\mathcal{N})$ could potentially be used to analyze reachability. The local-zone graph is an untimed transition system and we are interested in applying partial-order methods on it. As desired, domain-disjoint actions are independent in the local zone graph. This is a consequence of Lemmas 2 and 3.

Proposition 1 (Diamond property of $\text{LZG}(\mathcal{N})$). *Let $\text{dom}(a) \cap \text{dom}(b) = \emptyset$. If $(q, Z) \xRightarrow{ab} (q', Z')$ then $(q, Z) \xRightarrow{ba} (q', Z')$.*

Let us remark that the so called *enabledness* property [15] may not hold in a local-zone graph: it is possible to construct a network, a local-zone Z and two independent actions a, b such that from (q, Z) there are both \xRightarrow{a} and \xRightarrow{b} transitions but neither \xRightarrow{ab} nor \xRightarrow{ba} are feasible from (q, Z) [32]. Enabledness is however true at the level of configurations.

Although the local-zone graph is sound and complete for reachability, and has the diamond property, there are networks for which the local-zone graph is infinite. Hence a finite abstraction of the local-zone graph is required for analysis. This is the subject for the next section.

3 Abstract local-zone graphs

The goal of this section is to study finite abstractions of local-zone graphs that can be used to answer the reachability question. We introduce a general definition of an abstraction and of an abstract local-zone graph. Then we put restrictions on abstractions that make the abstract local-zone graph sound and complete for reachability. We fix a timed network \mathcal{N} . This allows us to omit indexing every notion with \mathcal{N} .

Definition 3. A *quasi-abstraction operator* $\alpha : \mathcal{P}(\text{LocalVal}) \rightarrow \mathcal{P}(\text{LocalVal})$ is a function from sets of local valuations to sets of local valuations such that $\alpha(\alpha(W)) = \alpha(W)$ for all sets of local valuations W . If the operator additionally satisfies $W \subseteq \alpha(W)$ for all sets W , we call it an *abstraction operator*. A quasi-abstraction operator is *finite* if its co-domain is finite: there are finitely many sets $\alpha(W)$.

The definition of the abstraction operator is the same as in the global-time semantics [6, 26], except that now we work

with local valuations. We will use the weaker notion which we have called a quasi-abstraction to get finite abstractions in our setting.

A quasi-abstraction operator allows to compute an abstract local-zone graph. An exploration of a local-zone graph is stopped when a node with a bigger abstraction is already in the graph. The smaller node is said to be *subsumed* by the bigger node. If the quasi-abstraction is finite, then we can have a finite abstract graph.

Definition 4 ($\text{LZG}^\alpha(\mathcal{N})$). Suppose α is a quasi-abstraction operator. An *abstract local-zone graph* is a subset of nodes and edges of $\text{LZG}(\mathcal{N})$ together with some new edges called *subsumption edges*. Each node is labeled either *covered* or *uncovered*. The graph must satisfy the following conditions:

- The initial node of $\text{LZG}(\mathcal{N})$ belongs to the graph.
- For every uncovered node (q, Z) , all its successors together with associated transitions $(q, Z) \xRightarrow{b} (q', Z')$ in $\text{LZG}(\mathcal{N})$ should be in the graph.
- For every covered node (q, Z) there is an uncovered node (q, Z') with $\alpha(Z) \subseteq \alpha(Z')$; moreover there is an explicit subsumption edge $(q, Z) \rightsquigarrow (q, Z')$.
- Every node of the graph must be reachable from the initial node by a path of \Rightarrow edges.

We denote by $\text{LZG}^\alpha(\mathcal{N})$ some abstract zone graph for \mathcal{N} . One can imagine that we take the first one in some fixed order on graphs.

A point worth noting is that the algorithm stores zones and not abstract sets. Indeed we do not assume that an abstraction of a zone is a zone, and therefore we do not know a priori how to store and manipulate an abstract set directly.

The question now is when it is correct to examine the abstract local-zone graph instead of the network itself: when can we say that a given state is reachable by a run in a network iff it is reachable in its abstract local-zone graph. Since every node of $\text{LZG}^\alpha(\mathcal{N})$ is reachable by a sequence of \Rightarrow transitions, we have:

Lemma 4. *Every abstract local-zone graph is sound: if a final state is reachable in $\text{LZG}^\alpha(\mathcal{N})$ then it is reachable in \mathcal{N} .*

We now study the converse implication.

Definition 5. A quasi-abstraction operator α is *complete* when reachability of a state q in \mathcal{N} implies its reachability in $\text{LZG}^\alpha(\mathcal{N})$.

The challenge is to get complete and finite quasi-abstraction operators for which the test $\alpha(Z) \subseteq \alpha(Z')$ is efficient. Abstractions for the global semantics are based on simulation relations [6, 17, 26]. Our next direction would be to consider abstractions based on simulations for the local semantics.

Definition 6. A (*time-abstract*) *simulation* relation \preceq on the local semantics is a reflexive and transitive relation $(q, v) \preceq$

(q, v') between configurations having the same discrete state that satisfies two conditions:

1. for every local delay transition $(q, v) \xrightarrow{\Delta} (q, v_1)$, there exists a local delay Δ' such that $(q, v') \xrightarrow{\Delta'} (q, v'_1)$ and $(q, v_1) \preceq (q, v'_1)$,
2. for every transition $(q, v) \xrightarrow{b} (q_1, v_1)$ there is a transition $(q, v') \xrightarrow{b} (q_1, v'_1)$ with $(q_1, v_1) \preceq (q_1, v'_1)$.

We say $v \preceq v'$ if $(q, v) \preceq (q, v')$ for all states q . When $\Delta' = \Delta$ in the first condition above, the relation is called a *strong-timed simulation*.

Definition 7. A quasi-abstraction operator α is *simulation based* if there is a simulation \preceq such that $\alpha(W) \subseteq \{v : \exists v' \in W. v \preceq v'\}$.

In particular, there is the biggest abstraction operator based on a simulation \preceq . It is simply the downward closure operator with respect to \preceq .

Lemma 5. A simulation based abstraction operator is complete.

This lemma (whose proof is in Appendix A) is not true in general for quasi-abstractions. The proof of the lemma crucially uses $Z \subseteq \alpha(Z)$, a property which may not hold in a quasi-abstraction. We propose an additional condition for quasi-abstractions that requires the abstraction $\alpha(Z)$ to keep some of the “good” valuations from the local-zone Z .

Definition 8. A quasi-abstraction α *keeps runs* if for every node (q, Z) in $\text{LZG}(\mathcal{N})$ that is reachable from the initial node, and every path $(q, Z) \xrightarrow{u} (q_f, Z_f)$ to the final state q_f there is a valuation $v \in \alpha(Z)$ and a run $(q, v) \xrightarrow{u} (q_f, v_f)$.

This property means that α should keep all paths leading to a final state. Observe that every abstraction operator keeps runs since $Z \subseteq \alpha(Z)$. The property of keeping runs, along with the operator being simulation based, gives a complete quasi-abstraction.

Lemma 6. A simulation based quasi-abstraction operator that keeps runs is complete for reachability.

The aim now is to come up with a concrete quasi-abstraction α that satisfies the properties of the above lemma and for which the test $\alpha(Z) \subseteq \alpha(Z')$ is efficient. We make a short digression into one of the first quasi-abstractions proposed for the local-time semantics.

Catch-up equivalence. A quasi-abstraction operator based on a relation between configurations called *catch-up equivalence* has been defined in [8]. However, as we show below, deciding whether two configurations are catch-up equivalent is PSPACE-hard.

We start with a definition of the equivalence. A delay $(q, v) \xrightarrow{\Delta} (q, v')$ is a *catch-up delay* if $\max(\{v'(t)\}_{t \in T}) =$




Figure 3. A diamond that is destroyed by a subsumption.

$\max(\{v(t)\}_{t \in T})$. So catch-up delays only allow the processes that are behind in time to join the most advanced processes. Two local-time configurations (q, v) and (q', v') are *catch-up equivalent* if the two can reach the same synchronized regions (i.e. Alur & Dill’s regions [3]) through catch-up delays and discrete transitions.

Proposition 2. The problem of deciding if two given configurations (q, v) , (q', v') of a given timed network are catch-up equivalent is PSPACE-hard.

The proof is by a reduction from the language emptiness of intersection of finite automata. It is presented in Appendix B. This hardness result makes it very unlikely that catch-up equivalence is suitable in practice.

To summarize this section, we have seen the properties we need of a quasi-abstraction to get a correct abstract local-zone graph (Lemma 6). In Section 6 we will present an efficient simulation based abstraction for local-zone graphs. Before that, we talk about partial-order reduction.

4 POR on abstract local-zone graphs

We discuss how to use partial-order methods on abstract zone graphs. At this point, we have a local-zone graph of a network $\text{LZG}(\mathcal{N})$ that has diamonds but may be infinite. We suppose that we have some quasi-abstraction α giving a finite abstract local-zone graph $\text{LZG}^\alpha(\mathcal{N})$. We would like to use partial-order methods on $\text{LZG}^\alpha(\mathcal{N})$, but this graph may not have diamonds as we illustrate in Figure 3. Due to subsumption there are no transitions from (q_b, Z_b) . So, $\text{LZG}(\mathcal{N})$ has diamonds but may be infinite, and $\text{LZG}^\alpha(\mathcal{N})$ is finite but has no diamonds. We show that when α satisfies the conditions given by Lemma 6, every partial-order method for $\text{LZG}(\mathcal{N})$ can be used on $\text{LZG}^\alpha(\mathcal{N})$.

In this section we will assume that we have a source function src for $\text{LZG}(\mathcal{N})$ given by a partial-order method as described in Section 2. In $\text{LZG}(\mathcal{N})$, we have diamonds and so we can use any partial-order method to calculate a source function. Recall that the nodes of $\text{LZG}(\mathcal{N})$ are pairs (q, Z) . The graph $\text{LZG}(\mathcal{N})$ may be infinite since there are infinitely many local-zones. As we want the source function to be given by some finite description, we assume that it does not depend on the local-zone, and instead depends only on the state q and the set of actions enabled from (q, Z) , denoted as $\text{enabled}(q, Z)$.

Definition 9. A *source function* for a timed network \mathcal{N} is a function $src : Q \times \mathcal{P}(\Sigma) \rightarrow \mathcal{P}(\Sigma)$. A source function of \mathcal{N} is trace-faithful if for every node (q, Z) and a path u from (q, Z) to a final state there is a source path $w \sim u$ from (q, Z) .

The concept of trace-faithful source function is directly inspired by partial-order methods. Indeed, they always compute trace-faithful source functions as they guarantee that every path has at least one equivalent source path.

Remark. Partial-order methods in general require both the diamond and enabledness properties [15]. In our case $LZG(\mathcal{N})$ has diamonds, but not necessarily the enabledness property. The latter property is not needed if, for example, final states are reached by a global synchronization action, or final states are determined by a state of one of the processes. The definition above of the source function hides this problem. When applying some existing partial-order methods, some precaution, or transformation of a system, should be done to ensure that the source function is indeed trace-faithful.

We can now combine abstraction and partial-order reduction.

Definition 10. For a timed network \mathcal{N} and a source function $src : Q \times \mathcal{P}(\Sigma) \rightarrow \mathcal{P}(\Sigma)$, the graph $LZG^{src}(\mathcal{N})$ is obtained from $LZG(\mathcal{N})$ by keeping only the edges allowed by the src function: $(q, Z) \xrightarrow{b} (q', Z')$ such that $b \in src(q, \text{enabled}(q, Z))$.

Then, $LZG^{a,src}(\mathcal{N})$ is a graph obtained from $LZG^{src}(\mathcal{N})$ that satisfies the conditions in Definition 4.

We now have a graph $LZG^{a,src}(\mathcal{N})$ on which both subsumption and POR have been applied. Used separately, both of them yield transition systems that are sound and complete for reachability. The next theorem, proved in Appendix C says that even the combination is correct.

Theorem 2. *If src is a trace-faithful source function and \mathbf{a} is a simulation based quasi-abstraction that keeps runs, then a final state is reachable in $LZG(\mathcal{N})$ iff it is reachable in $LZG^{a,src}(\mathcal{N})$.*

5 No finite abstractions for local-zone graphs

Theorem 2 gives a sufficient condition for a quasi-abstraction to be compatible with POR. There is one ingredient missing to get an algorithm. We need a finite quasi-abstraction. Unfortunately, we show that this is impossible under the assumptions made on the quasi-abstraction in Theorem 2. In the argument below, we do not really need that the quasi-abstraction keeps all runs. It would be enough to keep for every path a run with the same Parikh image.

Theorem 3. *There is a network \mathcal{N}^- such that $LZG^a(\mathcal{N}^-)$ is infinite for every simulation based quasi-abstraction operator that keeps runs.*

Proof. We present a network \mathcal{N}^- such that $LZG^a(\mathcal{N}^-)$ is infinite for every simulation based quasi-abstraction \mathbf{a} that

keeps runs. The same example appears in Lugiez et al. [31] in a similar context. The network \mathcal{N}^- , presented in Figure 4, consists of two processes A_1 and A_2 . It is easy to see that any accepting run of the network executes an equal number of b 's and c 's followed by the global synchronizing action $\$$.




Figure 4. A network of two processes without a finite abstract zone graph that contains all runs.

Consider $LZG(\mathcal{N}^-)$, the local zone graph of \mathcal{N}^- . For every $m, n \geq 0$, the network has a run on $b^m c^n$. Let $(p_0, q_0, Z_{m,n})$ be the node in $LZG(\mathcal{N}^-)$ reached from the initial node after the sequence $b^m c^n$: $(p_0, q_0, Z_0) \xrightarrow{b^m c^n} (p_0, q_0, Z_{m,n})$. Pick $i > j \geq 0$. We claim that:

- $\mathbf{a}(Z_{i,j}) \not\subseteq \mathbf{a}(Z_{k,l})$ for any $k, l \geq 0$ with $(i-j) \neq (k-l)$.

Suppose to the contrary that $\mathbf{a}(Z_{i,j}) \subseteq \mathbf{a}(Z_{k,l})$ for some i, j, k, l with $(i-j) \neq (k-l)$. Consider an execution $(p_0, q_0, v_0) \xrightarrow{b^i c^j} (p_0, q_0, v_{i,j}) \xrightarrow{c^{i-j} \$} (p_1, q_1, v)$ of \mathcal{N}^- . We have $v_{i,j} \in Z_{i,j}$. Hence by pre-property from Lemma 3, there is a path $(p_0, q_0, Z_{i,j}) \xrightarrow{c^{i-j}} (p_1, q_1, Z)$ in $LZG(\mathcal{N}^-)$. Now, as the operator \mathbf{a} keeps runs, there is $v'_{i,j} \in \mathbf{a}(Z_{i,j})$ and a run $(p_0, q_0, v'_{i,j}) \xrightarrow{c^{i-j}} (p_1, q_1, v')$. By $\mathbf{a}(Z_{i,j}) \subseteq \mathbf{a}(Z_{k,l})$ we have $v'_{i,j} \in \mathbf{a}(Z_{k,l})$. Since \mathbf{a} is simulation based there is $v_{k,l} \in Z_{k,l}$ with configuration $(p_0, q_0, v_{k,l})$ simulating $(p_0, q_0, v'_{i,j})$. Hence, we have $(p_0, q_0, v_{k,l}) \xrightarrow{c^{i-j} \$} (p_1, q_1, u)$ in \mathcal{N}^- . From the fact that $v_{k,l} \in Z_{k,l}$ and the post-property (Lemma 3), there is an execution $(p_0, q_0, v_0) \xrightarrow{b^k c^l} (p_0, q_0, v_{k,l})$. Combining the last two executions we obtain: $(p_0, q_0, v_0) \xrightarrow{b^k c^l} (p_0, q_0, v_{k,l}) \xrightarrow{c^{i-j} \$} (p_1, q_1, u)$. This is impossible for $(i-j) \neq (k-l)$.

By the diamond property of $LZG(\mathcal{N}^-)$, any sequence containing k occurrences of b , and l occurrences of c ends in $(p_0, q_0, Z_{k,l})$. From $\mathbf{a}(Z_{i,j}) \not\subseteq \mathbf{a}(Z_{k,l})$, the node $(p_0, q_0, Z_{i,j})$ (reached by any sequence containing i occurrences of b and j occurrences of c) cannot be subsumed by any other node. This shows that there are infinitely many nodes in $LZG^a(\mathcal{N})$ as there is at least one for every difference $(i-j)$. \square

In [23], a simulation based quasi-abstraction is defined which is shown to be finite and complete. This operator however does not keep runs, which is in accordance with the above result. Due to this reason, this operator is not amenable for partial-order reduction.

In the following sections we propose a way out from the apparent deadlock created by Theorems 2 and 3. One direction could be to find an abstraction operator not satisfying the hypothesis of Theorem 3, that is, either not simulation based or not keeping runs. We do not know how to do this while still preserving some form of Theorem 2. Our solution is to put some restrictions on the timed networks we consider. We will first generalize the $\alpha_{\leq LU}$ abstraction [6] for global-time semantics to the local-time semantics. Then we will show sufficient conditions under which it is finite.

6 LU-simulation for the local semantics

We will present a concrete strong-timed simulation that generalizes of the LU -simulation [6] known in the global semantics to the local semantics. It is parameterized by two functions L and U that keep for each clock the maximum constant among lower bound constraints $x \geq c, x > c$ and upper bound constraints $x \leq c, x < c$ respectively. The simulation induces an abstraction operator $\alpha_{\leq LU}^*$ which is sound, complete and keeps runs for networks with bounds L and U . The impossibility result from the previous section still applies though. Indeed the operator is not finite. In Section 7, we will present a restriction on timed networks and modify the abstraction operator to a quasi-abstraction operator that will be finite for the restricted class of networks.

Definition 11. An LU -bounds is a pair of functions $L : X \rightarrow \mathbb{N} \cup \{-\infty\}$ and $U : X \rightarrow \mathbb{N} \cup \{-\infty\}$, each of which maps process clocks to a natural number or $-\infty$. An atomic constraint $x \sim c$ is an LU -constraint if $c \leq L(x)$ when $\sim \in \{\geq, >\}$ (lower bound constraint) and if $c \leq U(x)$ when $\sim \in \{<, \leq\}$ (upper bound constraint). A network \mathcal{N} is an LU -network if every guard in \mathcal{N} is a conjunction of LU -constraints.

We next lift the LU -preorder [6], written as \leq_{LU} and defined for the global-time semantics to the local-time setting. Here, when we relate v and v' , we require that the difference between reference clocks is the same for both v and v' .

Definition 12 (\leq_{LU}^* -preorder). Given LU -bounds L and U . For two local valuations v, v' , we say $v \leq_{LU}^* v'$ if:

- $v(t_p - t_q) = v'(t_p - t_q)$ for all $p, q \in Proc$
- for all $p \in Proc$ and all $x \in X_p$
 - $v(t_p - x) \leq U_x \Rightarrow v'(t_p - x) \leq v(t_p - x)$
 - $v(t_p - x) \leq L_x \Rightarrow v'(t_p - x) \geq v(t_p - x)$
 - $v(t_p - x) > L_x \Rightarrow v'(t_p - x) > L_x$

Intuitively, the relation $v \leq_{LU}^* v'$ ensures the following: (1) whenever $v + \Delta$ synchronizes t_p and t_q , $v' + \Delta$ also synchronizes them, (2) whenever $v + \Delta$ satisfies an LU -constraint, $v' + \Delta$ also satisfies the same constraint. This is the basis for \leq_{LU}^* to induce a simulation over the local semantics. When v, v' are synchronized valuations, the \leq_{LU}^* preorder is identical to the \leq_{LU} preorder of the global-time semantics.

We overload the notation \leq_{LU}^* to a relation between configurations: we define $(q, v) \leq_{LU}^* (q, v')$ whenever $v \leq_{LU}^* v'$.

The next theorem (proved in Appendix D) states that \leq_{LU}^* relation is a strong-timed simulation on \mathcal{N} . We illustrate the theorem on an example. Consider a transition $q \xrightarrow{b} q'$ with guard $x_1 > c \wedge x_2 \leq d$ and a reset $\{x_1\}$. Action b is shared between processes 1 and 2. Suppose $(q, v) \xrightarrow{b} (q_1, v_1)$. Then $v(t_1) = v(t_2)$, and v satisfies the guard. Let $v \leq_{LU}^* v'$. We will see that $(q, v') \xrightarrow{b} (q_1, v'_1)$ and $v_1 \leq_{LU}^* v'_1$. Firstly, we have $v'(t_1) = v'(t_2)$ by the first item in the \leq_{LU}^* definition. Next, we have $v(t_1 - x_1) > c$ and $v(t_2 - x_2) \leq d$. If $v(t_1 - x_1) \leq L(x_1)$, then $v'(t_1 - x_1) \geq v(t_1 - x_1)$ by the second sub-item in the second condition; else $v'(t_1 - x_1) > L(x_1)$ by third sub-item. Since $L(x_1) \geq c$, we get $v'(t_1 - x_1) > c$ in both cases. Similarly, we can argue that $v'(t_2 - x_2) \leq d$ using the first sub-item with $U(x_2)$. Moreover, after resetting x_1 , all conditions of \leq_{LU}^* are still satisfied in the resulting valuations v_1 and v'_1 .

Theorem 4. Let \mathcal{N} be an LU -network. The relation \leq_{LU}^* is a strong-timed simulation on the local semantics of \mathcal{N} .

Definition 13. The abstraction operator $\alpha_{\leq LU}^*$ is defined as $\alpha_{\leq LU}^*(W) := \{v \mid v \leq_{LU}^* v' \text{ for some } v' \in W\}$ for every set of local valuations W . This is the downward closure of W with respect to the \leq_{LU}^* relation.

Theorem 5. For every LU -network \mathcal{N} , the abstraction operator $\alpha_{\leq LU}^*$ is sound and complete. It also keeps runs.

Unfortunately, despite this theorem we still miss two pieces to analyze timed networks with local semantics:

- We need an efficient test for $\alpha_{\leq LU}^*(Z) \subseteq \alpha_{\leq LU}^*(Z')$ because it is used in the definition of $LZG^{\alpha_{\leq LU}^*}(\mathcal{N})$.
- We need $LZG^{\alpha_{\leq LU}^*}(\mathcal{N})$ to be finite.

We discuss an efficient inclusion test in Section 6.1. The impossibility result from Theorem 3 tells us that $LZG^{\alpha_{\leq LU}^*}(\mathcal{N})$ cannot be always finite. To address this, we introduce the concept of a bounded-spread network in Section 7 and show that a variant of $LZG^{\alpha_{\leq LU}^*}(\mathcal{N})$ is finite there.

6.1 An algorithm for $\alpha_{\leq LU}^*(Z) \subseteq \alpha_{\leq LU}^*(Z')$

The counterpart of $\alpha_{\leq LU}^*$ in the global-semantics is the abstraction operator $\alpha_{\leq LU}$ [6]. It is well known that the $\alpha_{\leq LU}$ abstraction of a zone need not result in a zone, in fact, it may not even be convex [6, 26]. The current abstraction operator $\alpha_{\leq LU}^*$ is a generalization of $\alpha_{\leq LU}$ which is identical to $\alpha_{\leq LU}$ over zones that contain only synchronized valuations. Therefore, $\alpha_{\leq LU}^*$ is not convex. As in the global setting, the challenge is to decide the inclusion $\alpha_{\leq LU}^*(Z) \subseteq \alpha_{\leq LU}^*(Z')$ by looking at zones Z and Z' . We start with some simplification steps. Since $\alpha_{\leq LU}^*$ is the downward closure operator with respect to \leq_{LU}^* , we make the first simplification below.

Lemma 7. For every pair of zones Z, Z' : $\alpha_{\leq LU}^*(Z) \subseteq \alpha_{\leq LU}^*(Z')$ iff $Z \subseteq \alpha_{\leq LU}^*(Z')$.

The test $Z \subseteq \alpha_{\leq LU}^*(Z')$ can be seen as checking whether for every $v \in Z$ there exists a $v' \in Z'$ such that $v \leq_{LU}^* v'$.

Define $\langle v \rangle^* := \{v' \mid v \leq_{LU}^* v'\}$. The next lemma shows that we can reduce inclusion to intersection.

Lemma 8. *Let Z, Z' be non-empty zones. Then, $Z \not\subseteq \mathbf{a}_{\leq LU}^*(Z')$ iff there exists $v \in Z$ satisfying $\langle v \rangle^* \cap Z' = \emptyset$.*

As mentioned before, when Z, Z' contain only synchronized valuations, we have $\mathbf{a}_{\leq LU}^*(Z) = \mathbf{a}_{\leq LU}(Z)$, $\mathbf{a}_{\leq LU}^*(Z') = \mathbf{a}_{\leq LU}(Z')$ and the test boils down to checking $Z \subseteq \mathbf{a}_{\leq LU}(Z')$, which is studied in [26] for the global semantics. In the local semantics we need to consider valuations that are desynchronized. However, by definition of \leq_{LU}^* , for $v \leq_{LU}^* v'$, we require $v(t_p - t_q) = v'(t_p - t_q)$. This property allows us to lift the technique used in [26] to our setting.

For our analysis, we will make use of a graph representation of local-zones, called *distance graphs* [26, 30]. A distance graph has vertices $X \cup X'$. For every $x, y \in X \cup X'$ there is an edge $x \rightarrow y$ with a *weight* that is either (\prec, ∞) or of the form (\prec, c) with $c \in \mathbb{R}$ and \prec standing for \leq or $<$. The edge $x \xrightarrow{(\prec, c)} y$ represents the constraint $y - x < c$. For example, the zone $Z_1 := t_1 - x \geq 5 \wedge t_2 - y \leq 2$ can be represented as a graph with edges: $t_1 \xrightarrow{(\leq, -5)} x$, $t_2 \xrightarrow{(\leq, 0)} y$ and $y \xrightarrow{(\leq, 2)} t_2$. To reason about cumulative constraints of a path in this graph representation, an arithmetic over weights is defined.

Order: for $c_1, c_2 \in \mathbb{R}$, we say $(\prec_1, c_1) < (\prec_2, c_2)$ if $c_1 < c_2$, or $c_1 = c_2$, \prec_1 is $<$ and \prec_2 is \leq ; secondly, we have $(\prec, c) < (\prec, \infty)$ for every $c \in \mathbb{R}$. *Addition:* for $c_1, c_2 \in \mathbb{R}$, we have $(\prec_1, c_1) + (\prec_2, c_2)$ to be equal to (\prec, d) where $d = c_1 + c_2$ and \prec is $<$ if one of \prec_1 or \prec_2 is $<$, and \prec is \leq otherwise; secondly, $(\prec, c) + (\prec, \infty)$ is defined to be (\prec, ∞) for every weight (\prec, c) .

The addition allows us to define the weight of a path in a distance graph, as the sum of weights of the edges. A distance graph is *canonical* if for all pairs of vertices $x \neq y$, the smallest weight of a path from x to y is given by the weight of the edge $x \rightarrow y$. For a zone Z we denote by Z_{xy} the weight of the $x \rightarrow y$ edge in the canonical distance graph representing Z . We now have all the notation to state our inclusion test. Details of arriving at this test are in Appendix E.

Theorem 6. *Let Z, Z' be non-empty local zones. We have $Z \not\subseteq \mathbf{a}_{\leq LU}^*(Z')$ iff there exist two variables $x, y \in X \cup X_t$ s.t.*

- $Z'_{yx} < Z_{yx}$, and
- $(\leq, U_x) + Z_{t_p x} \geq (\leq, 0)$ if $x \in X_p$ for a process p , and
- $(\prec, -L_y) + Z'_{yx} < Z_{t_q x}$, if $y \in X_q$ for some process q .

The test runs over pairs of variables x, y and uses weights $Z_{yx}, Z'_{yx}, Z_{t_p x}$ and $Z_{t_q x}$ to check the conditions given by the theorem. This procedure can be implemented in time $\mathcal{O}(|X \cup X_t|^2)$. When we look at local-zones consisting of only synchronized valuations, we can add constraints $t_p = t_q$ and derive the test $Z \subseteq \mathbf{a}_{\leq LU}(Z')$ in the global-setting as a special case of the above theorem.

7 Bounded-spread networks

The impossibility result for local time semantics (Theorem 3) says that no simulation based abstraction can ensure finiteness of an abstract zone graph. Even if we go to quasi-abstractions, it is impossible to get a finite abstraction that keeps runs. As we do not know how to obtain abstractions that would go around this problem, we need to look for subclasses of timed networks where abstraction guarantees finiteness. In the example from the proof of Theorem 3, the local times of the two processes can differ by an arbitrary amount, and moreover this difference influences future behavior. We give a sufficient condition to avoid this situation.

We introduce the notion of bounded-spread networks and show how we can adapt the $\mathbf{a}_{\leq LU}^*$ abstraction (Definition 13) to get a finite quasi-abstraction that keeps runs for bounded-spread networks. This gives an algorithm for bounded-spread networks that can use both subsumption and POR at the same time. We also discuss some cases when a network is guaranteed to be of bounded spread, as well as present a method of converting any network into an equivalent bounded-spread network by adding some synchronizations.

Definition 14. The *spread* between processes A_p, A_q in a local valuation v is the absolute value of the difference between their reference clocks: $|v(t_p) - v(t_q)|$. Let $D \geq 0$ be a natural number. We say that a *valuation v has spread D* if the spread between every pair of processes in v is at most D .

Definition 15. A run in the local time semantics

$$(q_0, v_0) \xrightarrow{\Delta_0} (q_0, v'_0) \xrightarrow{b_1} (q_1, v_1) \xrightarrow{\Delta_1} \dots \xrightarrow{b_n} (q_n, v_n) \xrightarrow{\Delta_n} (q_n, v'_n)$$

is said to be *D -spread* if all $v_0, v'_0, \dots, v_n, v'_n$ have spread D .

Definition 16. A network \mathcal{N} is said to be *D -spread* if every local run of \mathcal{N} can be converted to a D -spread run by adjusting the delays: that is, for every run $(q_0, v_0) \xrightarrow{\Delta_0} (q_0, v'_0) \xrightarrow{b_1} (q_1, v_1) \xrightarrow{\Delta_1} \dots \xrightarrow{b_n} (q_n, v_n) \xrightarrow{\Delta_n} (q_n, v'_n)$ there exists a D -spread run $(q_0, \hat{v}_0) \xrightarrow{\Delta'_0} (q_0, \hat{v}'_0) \xrightarrow{b_1} (q_1, \hat{v}_1) \xrightarrow{\Delta'_1} \dots \xrightarrow{b_n} (q_n, \hat{v}_n) \xrightarrow{\Delta'_n} (q_n, \hat{v}'_n)$ where $\hat{v}_0 = v_0$.

Example. Consider the network in Figure 1a. We have two processes one with clock x and the other with clock y . There are also two reference clocks t_1 and t_2 . Let $v_{i,j}$ stand for a valuation $t_1 = i, t_2 = j, x = y = 0$. In particular, $v_{0,0}$ is an initial valuation. In the local semantics we have a run $(0, 0, v_{0,0}) \xrightarrow{(0,9)} (0, 0, v_{0,9}) \xrightarrow{c} (0, 1, v_{0,9}) \xrightarrow{b} (1, 1, v_{0,9})$. Valuation $v_{0,9}$ has spread 9. Yet the run has a spread 1 because we can adjust the delays: $(0, 0, v_{0,0}) \xrightarrow{(1,2)} (0, 0, v_{1,2}) \xrightarrow{c} (0, 1, v_{1,2}) \xrightarrow{b} (1, 1, v_{1,2})$. If we did not allow for adjusting delays in the definition of D -spread, this network would have an unbounded spread. With the adjustment, it is 1-spread.

7.1 When is a network bounded-spread

We give some examples where it is easy to check that a network is of bounded spread. For such classes we can apply the verification method presented in this work. In general, checking if the spread of a network is bounded by a given D is at least as hard as checking reachability. So an approach consisting of taking an arbitrary network, calculating its spread, and then applying our method, would not work. Observe that every network can be made 0-spread if one makes all the processes synchronize on all actions. However, this removes all parallelism in the network and any possibility of applying POR. We show a less radical method of converting any network to a D -spread network. The method introduces some new synchronizations, but still leaves some parallelism where partial-order methods can be applied. We start with some sufficient conditions for a network to be bounded spread and later describe the general construction.

Acyclic systems. We claim that acyclic systems are bounded-spread where the bound depends on the size of the network description. The local-zone graph of an acyclic network is finite and no abstraction is needed. But, making use of “cross” subsumptions reduces the state-space when there are multiple ways to reach a state. Showing that an acyclic system is bounded-spread allows to use both subsumption and POR.

Lemma 9. *Suppose M is a maximal constant in guards. The spread of a run of length n is bounded by $nM + 1$.*

Here is an example where we get the maximal spread:

$$\xrightarrow{b, x > M, x := 0} \dots \xrightarrow{b, x > M, x := 0} \xrightarrow{a, y = 0}$$

Actions a and b are local actions of two processes. At the beginning the two clocks are at 0. The clock of b process gets to some $nM + \epsilon'$ for $0 < \epsilon' < 1$, while the clock of a process is still 0. Lemma 9 gives an upper bound for the spread of any run in an acyclic system (see Appendix F for proofs).

Corollary 1. *An acyclic timed network $\mathcal{N} = \langle A_1, \dots, A_k \rangle$ is $(\sum_{i=1}^k |T_i|) \times M + 1$ -spread bounded where $|T_i|$ is the number of transitions in A_i , and M is the maximum constant in \mathcal{N} .*

Frequently communicating systems. More interesting examples of bounded-spread systems are frequently communicating client/server systems. In such systems we have one server process S , and a number of client processes C_1, \dots, C_n . The only communication actions are between the server and clients: the domain of an action can be either a singleton or $\{S, C_i\}$ for some i . Such a network is frequently communicating if there is a bound D such that every client communicates with server in every time interval of length D . It is not difficult to see that in this case the network is $2D$ -spread.

Another example is a network with barriers that can be modeled as global synchronizations. Assume there are no other communication actions: each action is either local to a process or it is a global synchronization. If we know that

there is a synchronizing action on every loop of every process then the system is bounded-spread thanks to Lemma 9.

This idea of frequent communication resulting in bounded-spread brings us to the next construction. It is possible to convert an arbitrary system to a bounded-spread system, at the price of reducing concurrency. Lemma 9 suggests adding global synchronizations, say on every loop. This would indeed bound the spread as the length of runs between two global configurations would be bounded. This transformation is unfortunately not correct: we miss some behaviors of the original system. Another solution is to synchronize everybody every D units of time, which we formalize below.

Definition 17. Let \mathcal{N} be an arbitrary timed network and $D \geq 1$ a natural number. Define \mathcal{N}^D to be the timed network obtained from \mathcal{N} as follows. Add a fresh clock z_p to every process p , and a new synchronization action s whose domain is the set of all processes. To every state of every process we add a self-loop on s with a guard $z_p = D$ and reset of z_p . To every other transition add $z_p < D$ to the existing guard.

The construction ensures that in every D units of time every process needs to do the s transition. Hence the resulting system is D -spread bounded. Moreover reachability is preserved as every state that is reachable is reachable by a global run (Lemma 1) and global runs are 0-spread. Hence all global runs of \mathcal{N} appear in \mathcal{N}^D , with embedded s actions.

Proposition 3. *For every network \mathcal{N} and natural number $D \geq 1$, the system \mathcal{N}^D is D -spread. A final state q_f is reachable in \mathcal{N} iff it is reachable in \mathcal{N}^D .*

The methodology that we develop in the subsequent section can be applied to \mathcal{N}^D . While independence between actions of the original network \mathcal{N} is preserved in \mathcal{N}^D , there may be more traces due to new synchronization actions. In Section 8, we will see an example where these extra traces get compensated by POR, in fact by an exponential factor.

7.2 Abstraction for bounded-spread networks

We come back to the crucial point of obtaining finite abstractions of bounded-spread networks. For a given bound D we use $\alpha_{\leq LU}^*$ abstraction restricted to D -spread valuations. We show that this guarantees finiteness of the abstract graph.

Definition 18. For a set of valuations W define $\text{spread}_D(W)$ to be $\{v \in W \mid v \text{ has spread } D\}$. The quasi-abstraction operator $\alpha_{\leq LU}^D$ is defined as $\alpha_{\leq LU}^D(W) := \alpha_{\leq LU}^*(\text{spread}_D(W))$ for every set of valuations W .

Theorem 7. *Quasi-abstraction $\alpha_{\leq LU}^D$ is sound, complete, finite, and keeps runs for D -spread networks. The inclusion $\alpha_{\leq LU}^D(Z) \subseteq \alpha_{\leq LU}^D(Z')$ can be checked in time $O(|X \cup X_t|^2)$ for time-elapsed local-zones Z, Z' .*


Proof. Soundness follows from Lemma 4. We now show that $\alpha_{\leq LU}^D$ keeps runs. Consider a D -spread network \mathcal{N} and a node

(q, Z) reachable from the initial node in $\text{LZG}(\mathcal{N})$: there exists a path $(q_0, Z_0) \xrightarrow{u'} (q, Z)$. Let $(q, Z) \xrightarrow{u} (q_f, Z_f)$ be a path to a final state. By post-property there is a local run $(q_0, v_0) \xrightarrow{u'} (q, v) \xrightarrow{u} (q_f, v_f)$ with $v_0 \in Z_0, v \in Z$ and $v_f \in Z_f$. Since \mathcal{N} is D -spread we can assume $v \in \text{spread}_D(Z)$, and hence by Definition 18 we have $v \in \mathfrak{a}_{\leq LU}^D(Z)$. This proves that $\mathfrak{a}_{\leq LU}^D$ keeps runs, as per Definition 8. Lemma 6 then entails that $\mathfrak{a}_{\leq LU}^D$ is complete. By using the inclusion test from Theorem 6, we show that an order between zones defined as $Z \leq Z'$ if $\mathfrak{a}_{\leq LU}^D(Z) \subseteq \mathfrak{a}_{\leq LU}^D(Z')$ is a well-quasi order (Lemma 14 in Appendix F.1). This proves finiteness of $\mathfrak{a}_{\leq LU}^D$. Complexity of the inclusion test is discussed in Lemma 13 of Appendix F. Since the local-zone graph has only time-elapsing zones, it is sufficient to consider such zones for the inclusion test. \square

8 Examples with exponential gain

Theorems 7 and 2, along with Definitions 10 and 9 give an algorithm for testing reachability in bounded-spread networks: explore the local zone graph restricted to the successors given by the src function, and for each fresh node (q, Z) that is discovered, do not explore further if it is subsumed, that is $\mathfrak{a}_{\leq LU}^D(Z) \subseteq \mathfrak{a}_{\leq LU}^D(Z')$ for an already visited node (q, Z') . We now present two examples on which this method gives exponential gain.

The first example shows advantages of local-time semantics together with partial-order methods. Consider a network of N timed automata \mathcal{A}_i as depicted below, where a_i and b_i are local actions, whereas $\$$ is a synchronized action:




Notice that not all sequences of actions are feasible in global time. For instance, b_2 requires a delay of 4 time units, hence it cannot happen before a_3 which only allows a delay of 3 time units. Still, all 3^N combinations of states q_0, q_1 and q_2 are reachable, although some of them are deadlocks. POR cannot be applied in algorithms using global-time semantics since the guards remove most diamonds.

In the local-time semantics, all sequences of actions are feasible. The spread is bounded by $2N$. We can apply a very simple partial-order technique: if there is a local action in $\text{enabled}(q, Z)$ then keep only the action of the smallest process, otherwise only $\$$ action is enabled and keep this action. This source function is complete for reachability of the final state (q_3, \dots, q_3) . There is only one source path and it follows the sequence of actions $a_1 b_1 a_2 b_2 \dots a_N b_N \$$. Recall that without POR at least 3^N states are visited.

The second example illustrates the general construction given in Definition 17 converting any network to a bounded-spread network (Proposition 3). Recall the example of a network with unbounded spread from Figure 4. We consider an extension of it to n processes, \mathcal{N}_n^- . We apply to it the construction for bounding the spread to 1, obtaining a network

$\mathcal{N}_n^+ = \langle A_1, \dots, A_n \rangle$, where A_i are as in the figure below. The actions s and $\$$ are global actions and b_i is local to A_i .



Consider a src function that gives for each (p, Z_σ) , the action s and the action b_i with the least index that is enabled at Z_σ . Denote by $\text{LZG}_{LU}^{1, \text{src}}(\mathcal{N}_n^+)$ the abstract local-zone graph over the $\mathfrak{a}_{\leq LU}^D$ operator with $D = 1$. The uncovered nodes in $\text{LZG}_{LU}^{1, \text{src}}(\mathcal{N}_n^+)$ are those accessible by paths $sb_1 \dots b_i$ for $0 \leq i < n$ and by $sb_1 \dots b_i sb_1 \dots b_j$ where $0 \leq i < n$ and $0 \leq j < i$. The covered nodes are the ones accessible by paths $sb_1 \dots b_n$ and by $sb_1 \dots b_i sb_1 \dots b_j s$ where $0 \leq i < n, 0 \leq j \leq i$. This gives an $\mathcal{O}(n^2)$ bound on the number of nodes present in $\text{LZG}_{LU}^{1, \text{src}}(\mathcal{N}_n^+)$. Without partial-order reduction, there will be a node (p, Z_{su}) for each u that is a sequence of b actions without repetitions. This is because Z_{su_1} and Z_{su_2} with u_1 and u_2 not being interleavings of each other cannot be covered with respect to each other by the $\mathfrak{a}_{\leq LU}^D$ quasi-abstraction. A detailed analysis, presented in Appendix G, shows that without a partial-order method, an exploration needs to visit exponentially many zones, be it in local or global-time semantics.

9 Conclusion

We have introduced a framework for applying partial-order methods to the analysis of timed automata. It uses local-time semantics in order to regain commutativity of independent actions. However, the resulting local-time zone graph is usually infinite, and prior finite abstractions were either impractical or incompatible with partial-order methods. We have introduced a new abstraction $\mathfrak{a}_{\leq LU}^*$ that is simulation based, and hence compatible with partial-order methods. The abstraction $\mathfrak{a}_{\leq LU}^*$ is generally not finite. Even worse, as we have shown here, there does not exist an abstraction that is finite, and simulation based. To circumvent this obstacle, we have introduced bounded-spread timed networks, for which the $\mathfrak{a}_{\leq LU}^*$ abstraction can be made finite. This requires the introduction of quasi-abstractions. We have given examples of subclasses of timed networks that are naturally bounded-spread, and we have shown that every timed network can be made bounded-spread, at the cost of reducing concurrency. We have illustrated the benefits of our framework on two examples.

Our next steps will be designing concrete partial-order methods that provide exponential gains for a wide class of timed networks. We hope that our framework can be extended to other verification problems, like liveness or solving timed games, as well as to richer timed models, like push-down timed automata, or weighted timed automata.

A Appendix for Section 3

► **Lemma 5.** A simulation based abstraction operator is complete.

Proof. Let \leq be the simulation on which the abstract local zone graph is based on. Let $(q_0, v_0) \xrightarrow{\Delta_0} (q_0, v'_0) \xrightarrow{b_1} (q_1, v_1) \cdots \xrightarrow{b_n} (q_n, v_n) \xrightarrow{\Delta_n}$ be a local run. For every (q_i, v_i) we will identify an uncovered node (q_i, Z_i) and a valuation $u_i \in Z_i$ such that $v_i \leq u_i$.

Base case is easy since the initial node (q_0, Z_0) contains the initial valuation v_0 . Assume we have identified (q_i, Z_i) and u_i . By property of simulations, there is a local run $(q_i, u_i) \xrightarrow{\Delta'_i} (q_i, u'_i) \xrightarrow{b_{i+1}} (q_{i+1}, u_{i+1})$ such that $v'_i \leq u_i$ and $v_{i+1} \leq u_{i+1}$. By pre-property (Lemma 3) there is a symbolic transition $(q_i, Z_i) \xrightarrow{b_i} (q_{i+1}, Z_{i+1})$ with $u_{i+1} \in Z_{i+1}$. If (q_{i+1}, Z_{i+1}) is uncovered, we are done. If not, we have a node (q_{i+1}, \hat{Z}_{i+1}) such that $\mathbf{a}(Z_{i+1}) \subseteq \mathbf{a}(\hat{Z}_{i+1})$. Since \mathbf{a} is an abstraction operator, we have $Z_{i+1} \subseteq \mathbf{a}(Z_{i+1})$. Hence there exists a $\hat{u}_{i+1} \in \hat{Z}_{i+1}$ such that $u_{i+1} \leq \hat{u}_{i+1}$. This node (q_{i+1}, \hat{Z}_{i+1}) and valuation \hat{u}_{i+1} give the required conclusion. \square

► **Lemma 6.** A simulation based quasi-abstraction operator that keeps runs is complete for reachability.

Proof. We show a more general result in Theorem 2. This lemma follows by taking src in Theorem 2 to be the set of all enabled actions for every (q, Z) . \square

B Catch-up equivalence is PSPACE-complete

A delay $(l, v) \xrightarrow{\Delta} (l, v')$ is a *catch-up delay* if $\max(\{v'(t)\}_{t \in T}) \leq \max(\{v(t)\}_{t \in T})$. So catch-up delays only allow the processes that are behind in time to join the most advanced processes.


Two local-time configurations (q, v) and (q', v') are *catch-up equivalent* if (q, v) and (q', v') can reach the same synchronized regions (i.e. Alur&Dill's regions) through catch-up delays and discrete transitions.

In the sequel we consider the following decision problem:

INPUT: A network of timed automata \mathcal{N} and two local-time configurations (q, v) and (q', v') of \mathcal{N} .

QUESTION: are (q, v) and (q', v') catch-up equivalent?

To warm-up we consider a simple network:




Process 0 has no transitions, and process 1 has one transition guarded with $x = 0$. We claim that the following

configurations are not catch-up equivalent:

$$(q_0^0, q_1^0, [t_0 = t_1 = 1, x = 0]) \not\sim (q_0^0, q_0^1, [t_0 = 1, t_1 = x = 0])$$

Indeed, in the first configuration no catchup transition is possible. In the second process 1 can do transition c immediately, and then wait in q_1^1 reaching $(q_0^0, q_1^1, [t_0 = 1, t_1 = x = 1])$. On the other hand if there is no c transition then the two configurations are equivalent, because the only thing possible is that process 1 lets the time pass to catch-up with process 0.

We use the same idea to reduce the language emptiness of the intersection of n finite automata.



We are given automata $\mathcal{A}_1, \dots, \mathcal{A}_n$ with initial states q_i^i and final states q_f^i . We guard every transition of every automaton with $x = 0$ (technically we need x_i for every process but this just makes notation worse). Moreover on a new letter b we add transitions from the initial state of automaton \mathcal{A}_i to every state of \mathcal{A}_i . Finally, from the tuple of finite states (q_f^1, \dots, q_f^n) we add a transition c with the guard $x = 0$, the transition synchronizes process $1, \dots, n$. We claim that if $L(\mathcal{A}_1) \cap \dots \cap L(\mathcal{A}_n) \neq \emptyset$ then the two configurations are not equivalent:

$$\begin{aligned}
 \text{conf}_{1,1} &= (q_0^0, q_0^1, \dots, q_0^n, [t_0 = t_1 = \dots = t_n = 1, x = 0]) \\
 \text{conf}_{1,0} &= (q_0^0, q_0^1, \dots, q_0^n, [t_0 = 1, t_1 = \dots = t_n = x = 0])
 \end{aligned}$$

If intersection is empty then transition c cannot be taken. The synchronized configurations reachable from $\text{conf}_{1,1}$ are all combinations of states of $\mathcal{A}_1, \dots, \mathcal{A}_n$ thanks to the added b transitions. Similarly, from $\text{conf}_{1,0}$, as the processes can just wait in the initial state to $\text{conf}_{1,1}$.

If the intersection is non-empty then from $\text{conf}_{1,0}$ processes can get to q_f^1, \dots, q_f^n in 0-time, and then do c transition reaching $(q_0^0, q_r^1, \dots, q_r^n, [t_0 = t_1 = \dots = t_n = 1, x = 0])$. This synchronized state is not possible to reach from $\text{conf}_{1,1}$.

C Appendix for Section 4

► **Theorem 2.** If src is a trace-faithful source function and \mathbf{a} is a simulation based quasi-abstraction that keeps runs, then a final state is reachable in $\text{LZG}(\mathcal{N})$ iff it is reachable in $\text{LZG}^{\mathbf{a}, \text{src}}(\mathcal{N})$.

Proof. If a final state is reachable in $\text{LZG}^{\text{a},\text{src}}(\mathcal{N})$ then it is reachable by a sequence of \Longrightarrow transitions by definition (c.f. Definitions 10, 4). This gives a path in $\text{LZG}(\mathcal{N})$.

Consider left-to-right direction. Since *src* is trace-faithful, it is sufficient to show that each source path in $\text{LZG}(\mathcal{N})$ that leads to a final state has a representative source path in $\text{LZG}^{\text{a}}(\mathcal{N})$, potentially with subsumption edges, that goes to a final state. The latter is a path in $\text{LZG}^{\text{a},\text{src}}(\mathcal{N})$ by definition. Suppose w_0 is a source path in $\text{LZG}(\mathcal{N})$ from (q_0, Z_0) to (q_n, Z_n) , with q_n an accepting state. Let $n = |w_0|$ be the length of w_0 . By induction on i we show that there are paths u_i, w_i such that:

- $(q_0, Z_0) \xrightarrow{u_i} (q_i, Z_i)$ is a source path in $\text{LZG}^{\text{a},\text{src}}(\mathcal{N})$,
- $(q_i, Z_i) \xrightarrow{w_i} (q_n, Z_n)$ is a source path in $\text{LZG}(\mathcal{N})$,
- $|w_i| = n - i$

The initial step is trivial. The induction step is easy if (q_i, Z_i) is uncovered in $\text{LZG}^{\text{a},\text{src}}(\mathcal{N})$. In this case, let b be the first letter of w_i : $w_i = bw_{i+1}$. We have $(q_i, Z_i) \xrightarrow{b} (q_{i+1}, Z_{i+1})$ and $b \in \text{src}(q_i, \text{enabled}(Z_i))$. Hence taking $u_{i+1} = u_i b$ and w_{i+1} we obtain the induction step.

It remains to check what happens when (q_i, Z_i) is covered in $\text{LZG}^{\text{a},\text{src}}(\mathcal{N})$. Say (q_i, Z_i) is subsumed by (q_i, Z'_i) , meaning $\mathfrak{a}(Z_i) \subseteq \mathfrak{a}(Z'_i)$. Since $(q_i, Z_i) \xrightarrow{w_i} (q_n, Z_n)$ and \mathfrak{a} keeps runs, there exists a valuation $v_i \in \mathfrak{a}(Z_i)$ and an execution $(q_i, v_i) \xrightarrow{w_i} (q_n, v_n)$. From $\mathfrak{a}(Z_i) \subseteq \mathfrak{a}(Z'_i)$, we have $v_i \in \mathfrak{a}(Z'_i)$. Secondly, as \mathfrak{a} is simulation based, there exists $v'_i \in Z'_i$ such that $v_i \leq v'_i$, where \leq is the simulation on which \mathfrak{a} is based on. Hence we also have an execution $(q_i, v'_i) \xrightarrow{w_i} (q_n, v'_n)$. By pre-property of zones $(q_i, Z'_i) \xrightarrow{w_i} (q_n, Z'_n)$ in $\text{LZG}(\mathcal{N})$. Since w_i is a path reaching a final state, and *src* function is trace faithful, there is a source path $(q_i, Z'_i) \xrightarrow{w'_i} (q_n, Z'_n)$ in $\text{LZG}(\mathcal{N})$ with $w_i \sim w'_i$. In particular, $|w_i| = |w'_i|$. Let b be the first letter of w'_i , i.e., $w'_i = bw'_{i+1}$. We claim that $u_{i+1} = u_i b$ and w_{i+1} satisfy the induction conditions. The path $u_i b$ contains a subsumption edge. \square

D Appendix for Section 6

► Theorem 4. Let \mathcal{N} be an *LU*-network. The relation \leq_{LU}^* is a strong-timed simulation on the local semantics of \mathcal{N} .

Proof. It is easy to see that \leq_{LU}^* is reflexive and transitive. Moreover, notice that if $v \leq_{\text{LU}}^* v'$ then $(v + \Delta) \leq_{\text{LU}}^* (v' + \Delta)$ for every local delay Δ . Hence \leq_{LU}^* is a reflexive and transitive relation that satisfies condition (1) required for a simulation as in Definition 6.

We now show condition (2). Let $v \leq_{\text{LU}}^* v'$ and suppose $(q, v) \xrightarrow{b} (q_1, v_1)$. This first means that all processes in $\text{dom}(b)$ are synchronized in v , that is, $v(t_p - t_q) = 0$ for all $t_p, t_q \in \text{dom}(b)$. From the first item of Definition 12, we have $v'(t_p - t_q) = 0$ as well, for $p, q \in \text{dom}(b)$. Hence the

processes in $\text{dom}(b)$ are synchronized in v' too. Next, from the definition of the local step $(q, v) \xrightarrow{b} (q_1, v_1)$, there is a tuple of b -transitions $\{(q_p, g_p, R_p, q'_p)\}_{p \in \text{dom}(b)}$ such that $v \models g_p$ for all $p \in \text{dom}(b)$. Since b is enabled at v , valuation v satisfies all the constraints occurring in all the guards g_p . We will now show that v' satisfies all these constraints by invoking the second item of Definition 12.

Since \mathcal{N} is an *LU*-network, for every $x < c$ (resp. $y > d$) occurring in the tuple, we have $c \leq U(x)$ (resp. $d \leq L(y)$). Consider a constraint $x < c$ from some g_p . As $v(t_p - x) < c$, we have $v(t_p - x) \leq U_x$. Hence $v'(t_p - x) \leq v(t_p - x)$ from the first sub-item. This implies $v'(t_p - x) < c$. Consider a constraint $y > d$. As $v \models y > d$, we have $v(t_p - y) > d$. If $v(t_p - y) \leq L_y$, then $v'(t_p - y) \geq v(t_p - y)$ from the second sub-item. This implies $v' \models y > d$. Otherwise, from the third sub-item, $v'(t_p - y) > L_y \geq d$, which again implies that $v' \models y > d$.

Therefore there is a transition $(q, v') \xrightarrow{b} (q_1, v'_1)$. It remains to show that $v_1 \leq_{\text{LU}}^* v'_1$. Since the transition b is instantaneous, $v_1(t_p - t_q) = v(t_p - t_q)$ and $v'_1(t_p - t_q) = v'(t_p - t_q)$ for all p, q . As $v(t_p - t_q) = v'(t_p - t_q)$ we get $v_1(t_p - t_q) = v'_1(t_p - t_q)$ for all $p, q \in \text{Proc}$. This gives the first item in the \leq_{LU}^* definition. Secondly, notice that $v_1 = [R]v$ and $v'_1 = [R]v'$. Therefore, for all clocks $x \notin R$, the second item is already satisfied for valuations v_1 and v'_1 . For all $x \in R$, we have $v_1(t_p - x) = v'_1(t_p - x) = 0$, when $x \in X_p$. Hence the second item is true for such clocks as well. \square

► Theorem 5. For every *LU*-network \mathcal{N} , the abstraction operator $\mathfrak{a}_{\leq_{\text{LU}}^*}^*$ is sound and complete. It also keeps runs.

Proof. Soundness and completeness follow from Lemmas 4, 5. Since $Z \subseteq \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z)$, the abstraction also keeps runs. \square

E Appendix for Section 6.1

► Lemma 7. For every pair of zones Z, Z' : $\mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z) \subseteq \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z')$ iff $Z \subseteq \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z')$.

Proof. The left-to-right direction is immediate since $Z \subseteq \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z)$. For the right-to-left direction, suppose $Z \subseteq \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z')$. Pick $v \in \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z)$. There exists $v_1 \in Z$ such that $v \leq_{\text{LU}}^* v_1$. From $Z \subseteq \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z')$, we have $v_1 \in \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z')$, and hence there is $v'_1 \in Z'$ such that $v_1 \leq_{\text{LU}}^* v'_1$. This also implies that $v \leq_{\text{LU}}^* v'_1$, and hence $v \in \mathfrak{a}_{\leq_{\text{LU}}^*}^*(Z)$. \square

E.1 Representing local zones

Local zones can be represented using Difference Bound Matrices (DBMs). For our analysis, we will make use of a graph representation of local zones, called *distance graphs*. A distance graph has vertices $X \cup X^t$. For every $x, y \in X \cup X^t$ there is an edge $x \rightarrow y$ with a *weight* is either (∞) or of the form (\leq, c) with $c \in \mathbb{R}$ and \leq standing for \leq or $<$. The edge $x \xrightarrow{(\leq, c)} y$ represents the constraint $y - x < c$. For a graph G , we will write $\llbracket G \rrbracket$ for the set of valuations satisfying

all the constraints given by G . To reason about cumulative constraints of a path in this graph representation, we define an arithmetic over weights:

Order. for $c_1, c_2 \in \mathbb{R}$, we say $(\ll_1, c_1) < (\ll_2, c_2)$ if $c_1 < c_2$, or $c_1 = c_2$, \ll_1 is $<$ and \ll_2 is \leq ; secondly, we have $(\ll, c) < (<, \infty)$ for every $c \in \mathbb{R}$,

Addition. for $c_1, c_2 \in \mathbb{R}$, we have $(\ll_1, c_1) + (\ll_2, c_2)$ to be equal to (\ll, d) where $d = c_1 + c_2$ and \ll is $<$ if one of \ll_1 or \ll_2 is $<$, and \ll is \leq otherwise; secondly, $(\ll, c) + (<, \infty)$ is defined to be (\ll, ∞) for every weight (\ll, c) .

The addition allows us to define the weight of a path in a distance graph, as the sum of weights of the edges. A distance graph is *canonical* if for all pairs of vertices $x \neq y$, the smallest weight of a path from x to y is given by the weight of the edge $x \rightarrow y$. Given two distance graphs G_1, G_2 we define $\min(G_1, G_2)$ to be the graph obtained by replacing the weight of every edge by the minimum of the corresponding weights from G_1 and G_2 . Finally, we will often reason about cycles in a distance graph. A cycle in a distance graph is *positive* if the sum of the weights of its edges is greater than or equal to $(\leq, 0)$. Otherwise, it is *negative*. It is well-known that $\ll[G]$ is non-empty iff there are no negative cycles in G .

We end this section with an observation about the local zones present in the local zone graph. This says that each local zone in the local zone graph can be described by difference constraints that use only integers.

Lemma 10. *Let \mathcal{N} be a network. For every node (q, Z) in $\text{LZG}(\mathcal{N})$, the canonical distance graph of Z has weight either (\ll, ∞) or (\ll, c) with $c \in \mathbb{Z}$, in each of its edges.*

Proof. This is true of the initial zone. We show that this property is preserved during successor computation.

Suppose G_1, G_2 are distance graphs with only integral weights. Then their intersection $\min(G_1, G_2)$ will have only integral weights since the canonicalization procedure only adds weights. This observation is sufficient to show the required property since each operation in the successor computation either involves removing edges or doing the intersection as above. \square

E.2 Steps to the final test

For convenience of presentation, we define two sets of clocks for a given local valuation v :

$$L\text{-bounded}(v) := T \cup \bigcup_{p \in \text{Proc}} \{x \in X_p \mid v(t_p - x) \leq L_x\}$$

$$U\text{-bounded}(v) := T \cup \bigcup_{p \in \text{Proc}} \{x \in X_p \mid v(t_p - x) \leq U_x\}$$

Notice that the reference clocks T are present in both $L\text{-bounded}(v)$ and $U\text{-bounded}(v)$.

Define $\langle v \rangle^* := \{v' \mid v \preceq_{LU}^* v'\}$.

► **Lemma 8.** Let Z, Z' be non-empty zones. Then, $Z \not\subseteq \alpha_{\preceq_{LU}^*}^*(Z')$ iff there exists $v \in Z$ satisfying $\langle v \rangle^* \cap Z' = \emptyset$.

Definition 19 (Distance graph H^v). Let $x, y \in X \cup X^t$ be two clocks, possibly reference clocks. Assume that $y \neq x$ and $y \in X_q \cup \{t_q\}$ for some process q . The weight of the edge $x \rightarrow y$ in the distance graph H^v is given by:

$$\left\{ \begin{array}{ll} (\leq, v(y-x)) & \text{if } x \in U\text{-bounded}(v), \\ & y \in L\text{-bounded}(v) \\ (\leq, v(t_q-x)) + (<, -L_y) & \text{if } x \in U\text{-bounded}(v), \\ & y \notin L\text{-bounded}(v), L_y \neq -\infty \\ (\leq, v(t_q-x)) & \text{if } x \in U\text{-bounded}(v), \\ & y \notin L\text{-bounded}(v), L_y = -\infty \\ (<, \infty) & \text{otherwise} \end{array} \right.$$

Lemma 11. *For every valuation v , we have $\ll[H^v] = \langle v \rangle^*$. Furthermore, the distance graph H^v is in canonical form.*

Proof. Proving $\ll[H^v] \subseteq \langle v \rangle^*$: Pick $v' \in H^v$. We will prove $v \preceq_{LU}^* v'$, by showing that v' satisfies the conditions of Definition 12.

For the first item of Definition 12, suppose $x = t_p$ and $y = t_q$. We have edges $t_p \rightarrow t_q$ and $t_q \rightarrow t_p$ with weights $(\leq, v(t_q - t_p))$ and $(\leq, v(t_p - t_q))$ in H^v . Hence $v'(t_p - t_q) = v(t_p - t_q)$ as required.

For the second item of Definition 12 take some process p and a clock $x \in X_p$. We have three conditions to check.

For the first condition, suppose $v(t_p - x) \leq U_x$. The edge $x \rightarrow t_p$ gives the constraint $(\leq, v(t_p - x))$ in H^v since x is U -bounded in v , and t_p is L -bounded in v . As v' satisfies this constraint, we get the desired $v'(t_p - x) \leq v(t_p - x)$.

For the second condition, suppose $v(t_p - x) \leq L_x$. The edge $t_p \rightarrow x$ has weight $v(x - t_p)$. Thus v' satisfies $v'(x - t_p) \leq v(x - t_p)$ equivalent to the required $v'(t_p - x) \geq v(t_p - x)$.

The third condition assumes $v(t_p - x) > L_x$. We have two cases. The first one is when $L_x \neq -\infty$. In this case the weight of the edge $t_p \rightarrow x$ is $(\leq, v(t_p - t_p)) + (<, -L_x)$. So $v'(x - t_p)$ satisfies $(\ll, -L_x)$, giving $v'(t_p - x) > L_x$. The second case is when $L_x = -\infty$. The constraint on the edge $t_p \rightarrow x$ is 0, giving the constraint $v'(t_p - x) \geq 0$. This constraint always holds as v' is a local valuation.

Proving $\langle v \rangle^* \subseteq \ll[H^v]$: Pick $v' \in \langle v \rangle^*$. We will show that v' satisfies every edge constraint $x \rightarrow y$ in H^v . Let us start with the case when x is a process clock in X_p and y is a process clock in X_q . Then, rewrite $v'(y - x)$ as:

$$v'(y - x) = v'(y - t_q) + v'(t_q - t_p) + v'(t_p - x) \quad (1)$$

We restrict to the situation when $x \in U\text{-bounded}(v)$, because if not, the constraint $x \rightarrow y$ in H^v is (\ll, ∞) . We now make some conclusions from the definition of \preceq_{LU}^* preorder. Since $x \in U\text{-bounded}(v)$, we have $v'(t_p - x) \leq v(t_p - x)$. Further, we have $v'(t_q - t_p) = v(t_q - t_p)$. Therefore:

$$v'(y - x) \leq v'(y - t_q) + v(t_q - t_p) + v(t_p - x) \quad (2)$$

When $y \in L\text{-bounded}(v)$, we have $v'(t_q - y) \geq v(t_q - y)$ from Definition 12. Hence $v'(y - t_q) \leq v(y - t_q)$. Plugging this to (2) gives $v'(y - x) \leq v(y - x)$. Hence the $x \rightarrow y$

constraint of H^v is satisfied. When $y \notin L$ -bounded(v) and $L_y \neq -\infty$, we have $v'(t_q - y) > L_y$, which gives $v'(y - t_q) < -L_y$. Plugging this to (2) gives $v'(y - x) < -L_y + v(t_q - x)$. Hence v' satisfies the constraint when the edge weight comes from the second item. When $L_y = -\infty$, we still have a trivial constraint that $v'(t_q - y) \geq 0$ as v' is a local valuation. This can be rewritten as $v'(y - t_q) \leq 0$. Plugging this in (2) gives $v'(y - x) \leq v(t_q - x)$, therefore satisfying the $x \rightarrow y$ edge constraint when the weight comes from the third item.

When x is a reference clock, the third term of (1) is 0. When y is a reference clock, the first term of (1) is 0. The rest of the argument follows similarly.

Proving that H^v is in canonical form. To show that H^v is canonical, we will show that the weight of $x \rightarrow y$ is smaller than or equal to weight of the path $x \rightarrow s \rightarrow y$ for every variable s . When $x \notin U$ -bounded(v), edge $x \rightarrow s$ has weight $(<, \infty)$ and the claim is trivially true. Similarly, if $s \notin U$ -bounded(v), the claim is true as $s \rightarrow y$ has weight $(<, \infty)$. Let us therefore assume $x, s \in U$ -bounded(v).

Suppose to the contrary that the sum of constraints on $x \rightarrow s \rightarrow y$ is strictly smaller than the constraint on $x \rightarrow y$. Let $w_{x \rightarrow s}$ stand for the constraint on the edge $x \rightarrow s$. If $y \in L$ -bounded(v) then we get $w_{x \rightarrow s} + (\leq, v(y - s)) < (\leq, v(y - x))$. After a simplification this gives $w_{x \rightarrow s} < (\leq, v(s - x))$. If $y \notin L$ -bounded(v) and $L_y \neq -\infty$ then we get $w_{x \rightarrow s} + (\leq, v(t_q - x)) + (<, L_y) < (\leq, v(t_q - s)) + (<, L_y)$. Once again this simplifies to $w_{x \rightarrow s} < (\leq, v(s - x))$. The same happens when $L_y = -\infty$.

It remains to show that $w_{x \rightarrow s} < (\leq, v(s - x))$ is impossible. If $s \in L$ -bounded(v) then $w_{x \rightarrow s}$ is $(\leq, v(s - x))$. If $s \notin L$ -bounded(v) we have $v(t_p - s) > L_s$, where p is the process of the clock s . If $L_s \neq -\infty$ then $w_{x \rightarrow s}$ is $(\leq, v(t_p - x)) + (<, -L_s)$. This is strictly bigger than $(\leq, v(t_p - x)) + (\leq, v(s - t_p)) = (\leq, v(s - x))$. If $L_s = -\infty$ then $w_{x \rightarrow s}$ is $(\leq, v(t_p - x))$. Observe that $v(s - x) = v(s - t_p) + v(t_p - x)$. Since $v(s - t_p) \leq 0$ as v is a local valuation, we get $v(s - x) \leq v(t_p - x) = w_{x \rightarrow s}$. \square

Proposition 4. *The intersection $\langle v \rangle^* \cap Z'$ is empty iff there are two variables $x, y \in X \cup T$ s.t. $x \in U$ -bounded(v), $L_y \neq -\infty$ when y is a process clock, and $H_{xy}^v + Z'_{yx} < (\leq, 0)$.*

Proof. Let $H_{Z'}$ be the canonical distance graph of Z' . The intersection $\langle v \rangle^* \cap Z'$ is empty iff there is a negative cycle in $H_{\min} := \min(H^v, H_{Z'})$. Suppose $H_{xy}^v + Z'_{yx} < (\leq, 0)$, then there is a negative cycle in H_{\min} . This gives the right-to-left direction of the proposition. We will now show the left-to-right direction.

Suppose $\langle v \rangle^* \cap Z'$ is empty. Then H_{\min} has a negative cycle N . Note that some of the edges of H_{\min} come from H^v and the others come from $H_{Z'}$. We will now reduce N to the form given in the right-hand-side of the proposition.

Step 1. Since H^v and $H_{Z'}$ are canonical, we can replace consecutive edges $x \rightarrow y \rightarrow u$ coming from the same graph with the edge $x \rightarrow u$ from that graph. Hence we can assume

that the edges in N alternate between edges from H^v and $H_{Z'}$.

Step 2. We transform N so that every edge coming from H^v has a weight given by either Item 1 or 2 of Definition 19. Clearly Item 4 does not apply as the sum of weights in N is a finite negative value, and hence we cannot have edges with $(<, \infty)$ weight in N . Suppose there is an edge $x \rightarrow y$ falling under Item 3. This edge can be replaced with the sequence $x \rightarrow t_q \rightarrow y$ from H^v with weight $(\leq, v(t_q - x))$ to edge $x \rightarrow t_q$ (due to Item 1) and weight $(\leq, 0)$ to edge $t_q \rightarrow y$ (due to Item 3). The weight of the edge $t_q \rightarrow y$ in $H_{Z'}$ is lesser than or equal to $(\leq, 0)$: this is because in local valuations the value of the corresponding reference clock is always greater than or equal to the value of a process clock, hence $t_q \geq y$ in all valuations of Z' , reflecting that $y - t_q \leq 0$. Therefore replacing the edge $t_q \rightarrow y$ in N with the corresponding edge from $H_{Z'}$ gives another negative cycle with weight at most that of N . This way we remove all edges coming from Item 3. Eventually, we apply once again Step 1 to collapse consecutive edges from Z' , so we have a cycle N with edges alternating between those of H^v and $H_{Z'}$.

Step 3. Consider an edge $x \rightarrow y$ in N coming from H^v and having weight due to Item 2, that is, $x \in U$ -bounded(v), $y \notin L$ -bounded(v) and $L_y \neq -\infty$. The weight of the edge is $(\leq, v(t_q - x)) + (<, -L_y)$, where t_q is the reference clock of y .

Replace this edge with two edges $x \xrightarrow{(\leq, v(t_q - x))} t_q \xrightarrow{(<, -L_y)} y$, both from H^v . This keeps the same value of the negative cycle. Perform this change for every edge coming from Item 2. We now have a negative cycle N where blocks of edges alternate between H^v and $H_{Z'}$: each H^v block either has a single edge $x \rightarrow y$ with weight $(\leq, v(y - x))$ from Item 1, or two edges $x \rightarrow t_q \rightarrow y$ with $x \rightarrow t_q$ having weight $(\leq, v(t_q - x))$ from Item 1, and $t_q \rightarrow y$ having weight $(<, -L_y)$ from Item 2.

Step 4. Suppose N has two H^v edges $x_1 \rightarrow y_1$ and $x_2 \rightarrow y_2$ with weights $u_1 := (\leq, v(y_1 - x_1))$ and $u_2 := (\leq, v(y_2 - x_2))$ due to Item 1. Therefore, $x_1, x_2 \in U$ -bounded(v) and $y_1, y_2 \in L$ -bounded(v). Let w_1 be the weight of the path in N from y_1 to x_2 and w_2 the weight from y_2 to x_1 . The cycle N can be broken into four parts as depicted below:

$$x_1 \xrightarrow{u_1} y_1 \cdots w_1 \cdots x_2 \xrightarrow{u_2} y_2 \cdots w_2 \cdots x_1$$

From Definition 19, the weights of edges $x_1 \rightarrow y_2$ and $x_2 \rightarrow y_1$ come due to Item 1. Let the weight of the edge $x_1 \rightarrow y_2$ be $u := (\leq, v(y_2 - x_1))$. If $u \leq u_1 + w_1 + u_2$, then the path from $x_1 \cdots y_2$ in N can be replaced with the edge $x_1 \xrightarrow{u} y_2$. Else, $u_1 + w_1 + u_2 < u$. Expanding this inequality, we get: $(\leq, v(y_1 - x_1)) + w_1 + (\leq, v(y_2 - x_2)) < (\leq, v(y_2 - x_1))$. Adding $(\leq, v(x_1 - y_2))$ on both sides gives $w_1 + (\leq, v(y_1 - x_2)) < (\leq, 0)$. As mentioned in the beginning of this paragraph, we know that the weight of the edge $x_2 \rightarrow y_1$ in H^v is $(\leq, v(y_1 - x_2))$. Therefore, the situation $u_1 + w_1 + u_2 < u$ gives a different negative cycle $y_1 \cdots w_1 \cdots x_2 \rightarrow y_1$, with the last edge $x_2 \rightarrow y_1$ from H^v . In both the cases, we replace

two Item 1 edges with a single Item 1 edge. Moreover, we still have the property that between two consecutive Z' edges on the cycle there is either a single Item 1 edge or a single Item 1 edge followed by an Item 2 edge. This is because transitions entering x_1 and x_2 on the cycle must come from Z' . Moreover, transitions from y_1 and y_2 must be either from Z' or Item 2 transitions followed by a transition from Z . After the reduction we can have two consecutive transitions from Z' , but then we can apply Step 1 to shorten the cycle.

Applying the transformation repeatedly, we are left with a negative cycle having a single Item 1 edge. This edge can be followed by an Item 2 edge. There can only be one Z' edge on the cycle. This means that N is either $x \xrightarrow{(\leq, v(y-x))} y \xrightarrow{Z'_{yx}} x$ or $x \xrightarrow{(\leq, v(t_q-x))} t_q \xrightarrow{(<, -L_y)} y \xrightarrow{Z'_{yx}} x$. In the latter case, we can replace $x \rightarrow t_q \rightarrow y$ with the H^v edge $x \rightarrow y$ and get a negative cycle $x \xrightarrow{(\leq, v(t_q-x)) + (<, -L_y)} y \xrightarrow{Z'_{yx}} x$. This finally gives us a negative cycle in the required form. \square

► **Theorem 6.** Let Z, Z' be non-empty local zones. We have $Z \not\subseteq \mathbf{a}_{\leq LU}^*(Z')$ iff there exist two variables $x, y \in X \cup T$ such that

- $Z'_{yx} < Z_{yx}$, and
- if $x \in X_p$ for some process p , then $(\leq, U_x) + Z_{t_p x} \geq (\leq, 0)$, and
- if $y \in X_q$ for some process q , then $(<, -L_y) + Z'_{yx} < Z_{t_q x}$.

Proof of Theorem 6

Left-to-right direction. Suppose $Z \not\subseteq \mathbf{a}_{\leq LU}^*(Z')$. Then there is a $v \in Z$ such that $\langle v \rangle^* \cap Z' = \emptyset$ (Lemma 8). From Proposition 4, there exist two clocks $x \in U$ -bounded(v) and if y is a process clock, then $L_y \neq -\infty$, and $H_{xy}^v + Z'_{yx} < (\leq, 0)$. We will use these conclusions to show the right hand side of the theorem.

We start with the second item, that is, to show that if $x \in X_p$ then $(\leq, U_x) + Z_{t_p x} \geq (\leq, 0)$. As $x \in U$ -bounded(v), we have $v(t_p - x) \leq U_x$. Secondly, since $v \in Z$, replacing the $x \rightarrow t_p$ edge of Z with $(\leq, v(t_p - x))$ will give no negative cycles. Hence, in particular: $(\leq, v(t_p - x)) + Z_{t_p x} \geq (\leq, 0)$. Plugging $v(t_p - x) \leq U_x$ into this inequality gives $(\leq, U_x) + Z_{t_p x} \geq (\leq, 0)$.

For the first and third items, we make use of a preliminary lemma.

Lemma 12. Let \bar{Z} be a non-empty local zone and let $\bar{v} \in \bar{Z}$. Let $r, s \in X \cup T$ be arbitrary clocks, and let (\leq, d) be a weight with $d \in \mathbb{Z}$. If $(\leq, \bar{v}(r-s)) + (\leq, d) < (\leq, 0)$ then $(\leq, d) < \bar{Z}_{rs}$.

Proof. Since $\bar{v} \in \bar{Z}$, \bar{v} satisfies all constraints of \bar{Z} . Consider the weight \bar{Z}_{rs} , which is the weight of the $r \rightarrow s$ edge in the canonical distance graph of \bar{Z} . This weight gives an upper bound for $v(s-r)$.

Suppose $\bar{Z}_{rs} = (\leq, c)$ (with a weak inequality in the weight). Then $\bar{v}(s-r) \leq c$, which implies $-c \leq \bar{v}(r-s)$. Since $(\leq, \bar{v}(r-s)) + (\leq, d) < (\leq, 0)$, we also have $(\leq, -c) + (\leq, d) < (\leq, 0)$. From this inequality, we can infer that either $d-c < 0$, or $d=c$ and $\leq < <$. Either way, we get $(\leq, d) < (\leq, c)$.

Suppose $\bar{Z}_{rs} = (<, c)$ (with a strict inequality in the weight). Then $v(s-r) < c$, which implies $-c < v(r-s)$. Let $v(r-s) = -c + \varepsilon$ for some $\varepsilon > 0$. We then have $(\leq, -c + \varepsilon) + (\leq, d) < (\leq, 0)$. By the previous argument, we have $(\leq, d) < (\leq, c - \varepsilon)$. Since d is an integer, this implies $(\leq, d) < (<, c)$.

In both cases, we can infer $(\leq, d) < \bar{Z}_{rs}$. \square

Thanks to Lemma 12, it is sufficient to show

$$(\leq, v(y-x)) + Z'_{yx} < (\leq, 0) \quad (3)$$

to conclude the first item, and

$$(\leq, v(t_q-x)) + (<, -L_y) + Z'_{yx} < (\leq, 0) \quad (4)$$

to conclude the third item.

We consider the case when $y \in L$ -bounded(v). In this case H_{xy}^v is $(\leq, v(y-x))$. The assumption $H_{xy}^v + Z'_{yx} < (\leq, 0)$ gives immediately (3). For equation (4) we use the fact that y is L -bounded, giving us $v(y-t_q) \geq -L_y$. Substituting this inequality into $v(y-x) = v(y-t_q) + v(t_q-x)$ we obtain $v(y-x) \geq -L_y + v(t_q-x)$. Then the hypothesis $H_{xy}^v + Z'_{yx} < (\leq, 0)$ gives the desired $(\leq, -L_y) + (\leq, v(t_q-x)) + Z'_{yx} < (\leq, 0)$.

When $y \notin L$ -bounded(v), we have $H_{xy}^v = (\leq, v(t_q-x)) + (<, -L_y)$. Therefore, $(\leq, v(t_q-x)) + (<, -L_y) + Z'_{yx} < (\leq, 0)$. Hence (4) is true. Moreover, $v(t_q-y) > L_y$ as $y \notin L$ -bounded(v). Now, $v(y-x) = v(y-t_q) + v(t_q-x)$ which is strictly lesser than $-L_y + v(t_q-x)$. In terms of weights, $(\leq, v(y-x)) \leq (<, -L_y) + (\leq, v(t_q-x))$. This implies that $(\leq, v(y-x)) + Z'_{yx} < (\leq, 0)$ is also true, proving (3).

Right to left direction. We will show that if the right hand side is true, there is a valuation $v \in Z$ satisfying the left hand side of Proposition 4 with clocks x and y . The third item of the right hand side already shows that $L_y \neq -\infty$ when y is a process clock. We now need to get a valuation $v \in Z$ such that $x \in U$ -bounded(v) and $H_{xy}^v + Z'_{yx} < (\leq, 0)$. Let G_Z be the canonical distance graph of Z .

Step 1. Consider the graph G_1 obtained from G_Z by replacing weight of edge $x \rightarrow t_p$ with $\min((\leq, U_x), Z_{x t_p})$ where p is the process of clock x , $x \in X_p$. Adding this edge causes no negative cycles, since $(\leq, U_x) + Z_{t_p x} \geq (\leq, 0)$. Therefore $\llbracket G_1 \rrbracket \neq \emptyset$ and contains the set of all valuations $v \in Z$ such that $x \in U$ -bounded(v). Let G_1^* be the canonical graph derived from G_1 . The shortest path from any variable s to x in G_1 does not involve the edge $x \rightarrow t_p$ since any path from s to x containing edge $x \rightarrow t_p$ will have a cycle, and we have seen that cycles in G_1 have non-negative weight. Therefore, weight of $s \rightarrow x$ in G_1^* is Z_{sx} .

Step 2. Suppose $Z_{yx} = (<_{yx}, c_{yx})$ and $Z_{t_q x} = (<_{t_q x}, c_{t_q x})$, where q is the process clock of y . Define $\eta_{yx} = c_{yx}$ if $<_{yx}$ equals \leq , otherwise $\eta_{yx} = c_{yx} - 0.5$. Since $Z'_{yx} < Z_{yx}$ and Z'

has integer weights, we also have $Z'_{yx} < (\leq, \eta_{yx})$. Similarly, define $\eta_{t_p x}$. From the third item of the rhs, we have $(\leq, -L_y) + Z'_{yx} < (\leq, \eta_{t_q x})$.

Let G_2 be the distance graph obtained from G_1^* by replacing $x \rightarrow y$ with $(\leq, -\eta_{yx})$ and $x \rightarrow t_q$ with $(\leq, -\eta_{t_q x})$. This gives the set of valuations simultaneously satisfying $v(y - x) \leq -\eta_{yx}$ and $v(t_q - x) \leq -\eta_{t_q x}$. If indeed $\llbracket G_2 \rrbracket$ is non-empty and there is such a valuation v , then we are done: we will have $(\leq, v(y - x)) + Z'_{yx} < (\leq, 0)$ and $(\leq, v(t_q - x)) + (\leq, -L_y) + Z'_{yx} < (\leq, 0)$. Notice that the value of H_{xy}^v comes from either case 1 or 2 of Definition 19. Hence we get $H_{xy}^v + Z'_{yx} < (\leq, 0)$.

It remains to show that G_2 has no negative cycles. The only two edges that are modified from G_1^* are $x \rightarrow y$ and $x \rightarrow t_q$. If there is a negative cycle, it should contain at least one of these two edges. If it contains both then the cycle can be broken down into two, with one of them being negative and containing exactly one of the above two edges. Therefore, we can assume without loss of generality that the negative cycle contains exactly one edge $x \rightarrow s$ where s is either y or t_q . As G_1^* is canonical, and the edges of G_2 other than these two come from G_1^* , we can conclude that the shortest path from $s \rightarrow x$ is the weight of $s \rightarrow x$ in G_1^* , which we have seen in the end of Step 1 to be Z_{sx} . Therefore, the possible negative cycle is of the form $x \rightarrow s \rightarrow x$ with weight $(\leq, -\eta_{sx}) + Z_{sx}$. By construction of η_{sx} this cycle cannot be negative for both the cases, when $s = y$ and $s = t_q$.

F Appendix for Section 7

► **Lemma 9.** Suppose M is a maximal constant in guards. The spread of a run of length n is bounded by $nM + 1$.

Proof. Consider a timed automaton and let M be its maximal constant. We claim that the minimal time for executing n actions in the automaton is at most $nM + 1$ in the global semantics. Indeed, in the global-time semantics there is no point of waiting more than M time units in a state, since after waiting M time units the valuation is already in the biggest region and valuations within a region simulate each other (see [3] for the definition and properties of regions). This intuition is less evident in local-time semantics but we can transfer this observation from the global-time to local-time.

Consider a local run on a sequence $b_1 \dots b_n$. By Lemma 1 there is a global run on a sequence $c_1 \dots c_n$ such that $b_1 \dots b_n$ is trace equivalent to $c_1 \dots c_n$. This means that there is a bijection $f : [n] \rightarrow [n]$ with $b_i = c_{f(i)}$ and respecting order of actions on each process: if b_i and b_j are two actions of process p , and $i < j$ then $f(i) < f(j)$.

The global run has the form:

$$(q_0, v_0) \xrightarrow{\delta_1} (q_0, v'_0) \xrightarrow{c_1} (q_1, v_1) \xrightarrow{\delta_2} \dots \xrightarrow{c_n} (q_n, v_n) .$$

We can assume that the cumulated time of this run is at most $nM + 1$; this is because if some δ_i is strictly bigger than M then we can shorten it to $M + \epsilon$ for a $0 < \epsilon < \frac{1}{n}$ as anyway the resulting valuation is the maximal region. Let θ_i be the

cumulated time before the i -th action: $\theta_i = \delta_1 + \dots + \delta_i$. We construct a local-time execution

$$(q_0, v_0) \xrightarrow{\Delta_1} (q_0, v'_0) \xrightarrow{b_1} (q'_1, v_1) \xrightarrow{\Delta_2} \dots \xrightarrow{b_n} (q_n, v_n) .$$

such that for every i and process $p \in \text{dom}(b_i)$ we have $v'_i(t_p) = \theta_{f(i)}$. This means that we execute action b_i exactly the time when the corresponding action $c_{f(i)}$ was executed in the global run. This constraint determines Δ_i , hence determines the run completely. It can be checked that it is indeed a run: all Δ 's are positive, and all guards are satisfied. By definition we have $v_i(t_p) - v_0(t_p) \leq nM + 1$ for all reference clocks and for all i . Since all reference clocks are equal in v_0 , we get that the spread is at most $nM + 1$. \square

► **Corollary 1.** An acyclic timed network $\mathcal{N} = \langle A_1, \dots, A_k \rangle$ is $(\sum_{i=1}^k |T_i|) \times M + 1$ -spread bounded where $|T_i|$ gives the number of transitions in process A_i , and M is the maximum constant used in \mathcal{N} .

Proof. No transition repeats in an acyclic system. Hence the length of a run is bounded by $(\sum_{i=1}^k |T_i|)$. Lemma 9 gives the bound $(\sum_{i=1}^k |T_i|) \times M + 1$ for each run, and hence the system is spread bounded with this constant. \square

► **Proposition 3.** For every network \mathcal{N} and natural number $D \geq 1$, the system \mathcal{N}^D is D -spread bounded. A final state q is reachable in \mathcal{N} iff it is reachable in \mathcal{N}^D .

Proof. Consider a local run in \mathcal{N}^D . Valuations reached after the s action are synchronized. Between any two such valuations, the run can elapse at most D time units in each process. As the initial valuation is synchronized, the prefix of the run upto the first s action can also elapse at most D time units. Similarly, after the last s action, the run cannot elapse more than D time units in each process. This shows that the run is D -spread.

Pick a global run in \mathcal{N} . At every delay of D time units, insert the action s . This gives a run in \mathcal{N}^D . By Lemma 1, every state q reachable in the local semantics is reachable in the global semantics. As \mathcal{N}^D contains a representative for every global run, we get that \mathcal{N}^D is complete for reachability. \square

Lemma 13. Let Z, Z' be time-elapsing zones. The test $\mathbf{a}_{\leq LU}^D(Z) \subseteq \mathbf{a}_{\leq LU}^D(Z')$ can be done in time $O(|X \cup X_t|^2)$.

Proof. The test involves two steps: (1) computing $Z_1 := \text{spread}_D(Z)$ and $Z'_1 := \text{spread}_D(Z')$ and then (2) checking $\mathbf{a}_{\leq LU}^*(Z_1) \subseteq \mathbf{a}_{\leq LU}^*(Z'_1)$. The second step can be done in time $O(|X \cup X_t|^2)$ thanks to Theorem 6. We show that $\text{spread}_D(Z)$ can also be computed in the same complexity.

Let G_Z be the canonical distance graphs of Z . Let (\ll_{xy}, c_{xy}) be the weight of $x \rightarrow y$ in G_Z . Since Z is time-elapsing, there are no constraints that give an upper bound on the reference clocks, that is, there are no constraints of the form $t_p - x < c$. This implies that every edge of the form $x \rightarrow t_p$

with $x \in X \cup X_t$ and $t_p \in X_t$ has weight (\langle, ∞) . Same is the case with $G_{Z'}$ as Z' is time-elapsd.

Computing $\text{spread}_D(Z)$ involves taking G_Z , adding edges $t_p \xrightarrow{(\leq, D)} t_q$ between every pair of reference clocks t_p, t_q and canonicalizing the resulting graph. Call this resulting graph \bar{G} . Since G_Z had no incoming edges to reference clocks, the only incoming edges to t_p in \bar{G} are from other reference clocks. In particular, there are no edges in \bar{G} of the form $x \rightarrow t_p$ where $x \in X$. Therefore, the only shortest paths that can change are of the form $t_p \rightarrow y$, where y is a process clock. The shortest path from t_p to y is given by the minimum of $(\langle_{t_p y}, c_{t_p y})$ and $(\leq, D) + (\langle_{t_q y}, c_{t_q y})$ over all q . This can be computed in $O(|X|_t)$ for each $t_p \rightarrow y$. Overall, ranging over all such edges, we get a complexity $O((|X_t| \cdot |X|) \cdot |X|_t)$. This gives an $O(|X \cup X_t|^2)$ procedure for computing $\text{spread}_D(Z)$. The same complexity holds for computing $\text{spread}_D(Z')$. \square

F.1 $\mathfrak{a}_{\leq LU}^D$ is finite.

Since $\mathfrak{a}_{\leq LU}^D$ first restricts to D -spread valuations, let us restrict to zones containing only D -spread valuations. For such zones $\mathfrak{a}_{\leq LU}^D(Z) \subseteq \mathfrak{a}_{\leq LU}^D(Z')$ iff $Z \subseteq \mathfrak{a}_{\leq LU}^*(Z')$. Define an order $Z \preceq Z'$ if $Z \subseteq \mathfrak{a}_{\leq LU}^*(Z')$.

Lemma 14. *The order \preceq on D -spread zones is a wqo.*

Proof. Suppose \preceq is not a wqo. Then there is an infinite antichain for \preceq . If $Z \not\preceq Z'$ there is a pair of clocks x, y witnessing the non-inclusion. Since the number of such pairs is finite, by standard Ramsey arguments, there are two clocks x, y and infinite sequence Z^1, Z^2, \dots such that $Z^i \not\preceq Z^{i+1}$ for all $i \geq 1$ is witnessed by x, y . This by Theorem 6 means that for all $i \geq 1$ we have:

1. $Z_{yx}^{i+1} < Z_{yx}^i$
2. if $x \in X_p$, then $(\leq, U_x) + Z_{t_p x}^i \geq (\leq, 0)$,
3. if $y \in X_q$ then $(\langle, -L_y) + Z_{yx}^{i+1} < Z_{t_q x}^i$

Let us assume that x and y are process clocks, with $x \in X_p$ and $y \in X_q$. The argument below can be adapted to the cases when one of them is a reference clock.

From 2 and from the fact that in all valuations we have $x - t_p \leq 0$, we get $(\leq, -U_x) \leq Z_{t_p x}^i \leq (\leq, 0)$. From 1, we see that Z_{yx}^{i+1} keeps decreasing as i increases and since the zones we get have only integer weights, for every $K \leq 0$, we can find an i such that $Z_{yx}^i < (\leq, K)$.

Now, due to canonicity, we have $Z_{t_p x}^{i+1} \leq Z_{t_p t_q}^{i+1} + Z_{t_q y}^{i+1} + Z_{yx}^{i+1}$. We have $Z_{t_p t_q}^{i+1} \leq (\leq, D)$ as Z^{i+1} is a D -spread zone and $Z_{t_q y}^{i+1} \leq 0$ as $y - t_q \leq 0$ for all local valuations. By choosing a sufficiently large i , we can get a small value for Z_{yx}^{i+1} and hence the sum $Z_{t_p t_q}^{i+1} + Z_{t_q y}^{i+1} + Z_{yx}^{i+1}$ to be strictly smaller than $(\leq, -U_x)$. This is a contradiction with 2. \square

G Appendix for Section 8

We describe the second example from section 8 in more detail here.

Since \mathcal{N}^+ is obtained by applying the construction of Definition 17 to an n -process extension of \mathcal{N}^- , we have \mathcal{N}_n^+ to be 1-spread by Proposition 3. We show that network \mathcal{N}_n^+ is in fact 0-spread. We will subsequently work with the graph $\text{LZG}_{LU}^{D,src}(\mathcal{N}_n^+)$ by taking $D = 0$ instead of $D = 1$ as this makes the discussion simpler. The same results hold when we consider $D = 1$ too. Consider an arbitrary local run of \mathcal{N}_n^+ . Since s is a global synchronization the valuations reached after doing s are synchronized, hence 0-spread. The final action $\$$ is also a global synchronization which results in a 0-spread valuation. Between two s actions, we can only have some b -sequence happening in 0 time. Hence each intermediate valuation is 0-spread, making the run 0-spread.

Next, we look at some of the valuations reached. From the above paragraph, every valuation obtained after an action transition in a run has $z_i = 0$ for all i : action s resets z_i , and between two s there is no time elapse. The z_i clocks will not play a role in deciding subsumption and hence we will ignore z_i clocks for our analysis. For a sequence of actions σ not containing $\$$, we let (p, v_σ) be the configuration reached after executing σ from the initial configuration (p, v_0) . Let u, u_1, u_2 be b -sequences. We say $i \in u$ if u contains b_i . Our sequences of interest are sus, su, su_1su_2s and su_1su_2 . We tabulate the values of the valuations reached after such sequences. For simplicity, we will write $v(x_i)$ for $v(t_i - x_i)$, where t_i is the reference clock of A_i . All valuations below are synchronized.

$$v_{sus}(x_i) = \begin{cases} 1 & i \in u \\ 2 & i \notin u \end{cases} \quad v_{su}(x_i) = \begin{cases} 0 & i \in u \\ 1 & i \notin u \end{cases}$$

$$v_{su_1su_2s}(x_i) = \begin{cases} 1 & i \in u_2 \\ 2 & i \in u_1, i \notin u_2 \\ 3 & \text{otherwise} \end{cases}$$

$$v_{su_1su_2}(x_i) = \begin{cases} 0 & i \in u_2 \\ 1 & i \in u_1, i \notin u_2 \\ 2 & \text{otherwise} \end{cases}$$

The zone Z_σ reached after each sequence σ as above is given by $\text{local-elapse}(v_\sigma)$. The 0-spread valuations in Z_σ are those obtained by elapsing the same local delay on each process from valuation v_σ . By property of simulations, $v_\sigma \preceq_{LU}^* v_{\sigma'}$ implies $v_\sigma + \Delta \preceq_{LU}^* v_{\sigma'} + \Delta$. Therefore, $v_\sigma \preceq_{LU}^* v_{\sigma'}$ implies $\mathfrak{a}_{\leq LU}^0(Z_\sigma) \subseteq \mathfrak{a}_{\leq LU}^0(Z_{\sigma'})$ (in fact, $v_\sigma \preceq_{LU}^* v_{\sigma'}$ implies $\mathfrak{a}_{\leq LU}^D(Z_\sigma) \subseteq \mathfrak{a}_{\leq LU}^D(Z_{\sigma'})$ for all $D \geq 0$ and hence the analysis that follows will hold for $D = 1$ too). Notice that we have $L = U = 1$ for every x_i . Definition 12 then gives $v_\sigma \preceq_{LU}^* v_{\sigma'}$ iff either $v_\sigma(x_i) = v_{\sigma'}(x_i)$ or both $v_\sigma(x_i), v_{\sigma'}(x_i) > 1$.

Based on the valuations above, we have: $v_{su_1su_2s} \preceq_{LU}^* v_{su_2s}$. Secondly we have $v_{sb_1b_2\dots b_n} \preceq_{LU}^* v_0$ where v_0 is an initial valuation that is synchronized and has $v_0(x_i) = 0$ for all i . This

gives $\mathbf{a}_{\leq LU}^0(Z_{su_1su_2s}) \subseteq \mathbf{a}_{\leq LU}^0(Z_{su_2s})$ and $\mathbf{a}_{\leq LU}^0(Z_{sb_1b_2\dots b_n}) \subseteq \mathbf{a}_{\leq LU}^0(Z_0)$. Now consider v_{su} where $u \neq b_1b_2\dots b_n$. Valuation v_{su} can neither be simulated by $v_{su'}$ nor any of its time successors, when $u' \neq u$. Similarly, valuation v_{su} cannot be simulated by valuations of the form v_{su_1s} or its time successors. Therefore, Z_{su} cannot be subsumed by any node of the form $Z_{su'}$ or Z_{su_1s} . An analogous reasoning gives that $Z_{su_1su_2}$ cannot be subsumed by zones Z_{su} , Z_{su_s} , $Z_{su'_1su}$, $Z_{su_1su'_2}$ where $u_1 \neq u'_1$ and $u_2 \neq u'_2$. This gives an abstract zone graph (without partial-order reduction) where sequences with three s actions are covered by a sequence with two s actions, and a sequence su where u contains all actions b_i is covered by Z_0 . There is an uncovered node for every Z_{su} where u has at most $n-1$ b_i actions. Hence, without partial-order reduction, the abstract local zone graph has at least 2^{n-1} nodes.

We now look at what happens with partial-order reduction. Here is a trace-faithful source function: at (p, Z_σ) the source function gives s , the b_i with the least index that is enabled at Z_σ and the $\$$ action if it is enabled. Observe that $\text{enabled}(Z_{su}) = \{s\} \cup \{b_i \mid i \notin u\} \cup \{\$\mid u \text{ is empty}\}$ and $\text{enabled}(Z_{su_1su_2}) = \{s\} \cup \{b_i \mid i \in u_1, i \notin u_2\}$. Therefore, at Z_{su} the source function picks the smallest b_i that is not present in u , and at $Z_{su_1su_2}$ it picks the smallest b_i that is present in u_1 , but not yet in u_2 .

The zone graph for four processes, $\text{LZG}_{LU}^{0,src}(N_4^+)$ is depicted in Figure 5. In the general case, the uncovered nodes in $\text{LZG}_{LU}^{0,src}(N_n^+)$ are $sb_1\dots b_i$ for $0 \leq i < n$ and $sb_1\dots b_isb_1\dots b_j$ where $0 \leq i < n$ and $0 \leq j < i$. The covered nodes are $sb_1\dots b_n$ (since $v_{sb_1\dots b_n} \leq_{LU}^* v_0$) and $sb_1\dots b_isb_1\dots b_js$ where $0 \leq i < n, 0 \leq j \leq i$. This gives an $\mathcal{O}(n^2)$ bound on the number of nodes present in $\text{LZG}_{LU}^{0,src}(N_n^+)$.

References

- [1] Parosh Aziz Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. Optimal dynamic partial order reduction. In Suresh Jagannathan and Peter Sewell, editors, *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, pages 373–384. ACM, 2014.
- [2] Parosh Aziz Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. Source sets: A foundation for optimal dynamic partial order reduction. *J. ACM*, 64(4):25:1–25:49, 2017.
- [3] Rajeev Alur and David L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 126(2):183–235, 1994.
- [4] Gilles Audemard, Alessandro Cimatti, Artur Kornilowicz, and Roberto Sebastiani. Bounded model checking for timed systems. In Doron A. Peled and Moshe Y. Vardi, editors, *Formal Techniques for Networked and Distributed Systems - FORTE 2002, 22nd IFIP WG 6.1 International Conference Houston, Texas, USA, November 11-14, 2002, Proceedings*, volume 2529 of *Lecture Notes in Computer Science*, pages 243–259. Springer, 2002.
- [5] Bahareh Badban and Martin Lange. Exact incremental analysis of timed automata with an smt-solver. In Uli Fahrenberg and Stavros Tripakis, editors, *Formal Modeling and Analysis of Timed Systems - 9th International Conference, FORMATS 2011, Aalborg, Denmark, September 21-23, 2011. Proceedings*, volume 6919 of *Lecture Notes in Computer Science*, pages 177–192. Springer, 2011.




Figure 5. Illustration of the abstract local zone graph $\text{LZG}_{LU}^{0,src}(N_4^+)$. The label of the node reached after a sequence σ represents the value of x_1, x_2, x_3, x_4 in the valuation v_σ . Covered nodes are shown in blue. Nodes su_1su_2s are covered by su_2s and node $sb_1b_2b_3b_4$ is covered by the initial node. Subsumption edges have not been shown explicitly.

- [6] Gerd Behrmann, Patricia Bouyer, Kim Guldstrand Larsen, and Radek Pelánek. Lower and upper bounds in zone-based abstractions of timed automata. *Int. J. Softw. Tools Technol. Transf.*, 8(3):204–215, 2006.
- [7] Gerd Behrmann, Kim Guldstrand Larsen, Justin Pearson, Carsten Weise, and Wang Yi. Efficient timed reachability analysis using clock difference diagrams. In Nicolas Halbwachs and Doron A. Peled, editors, *Computer Aided Verification, 11th International Conference, CAV '99, Trento, Italy, July 6-10, 1999, Proceedings*, volume 1633 of *Lecture Notes in Computer Science*, pages 341–353. Springer, 1999.
- [8] 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.
- [9] Frederik M. Bønneland, Peter Gjø Jensen, Kim G. Larsen, Marco Muñoz, and Jiri Srba. Stubborn set reduction for timed reachability and safety games. In Catalin Dima and Mahsa Shirmohammadi, editors, *Formal Modeling and Analysis of Timed Systems - 19th International Conference, FORMATS 2021, Paris, France, August 24-26, 2021, Proceedings*, volume 12860 of *Lecture Notes in Computer Science*, pages 32–49. Springer, 2021.
- [10] Frederik M. Bønneland, Peter Gjø Jensen, Kim Guldstrand Larsen, Marco Muñoz, and Jiri Srba. Start pruning when time gets urgent: Partial order reduction for timed systems. In Hana Chockler and Georg Weissenbacher, editors, *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part I*, volume 10981 of *Lecture Notes in Computer Science*, pages 527–546. Springer, 2018.
- [11] Frederik Meyer Bønneland, Peter Gjø Jensen, Kim Guldstrand Larsen, Marco Muñoz, and Jiri Srba. Stubborn set reduction for two-player reachability games. *Log. Methods Comput. Sci.*, 17(1), 2021.
- [12] Patricia Bouyer, Serge Haddad, and Pierre-Alain Reynnier. Timed unfoldings for networks of timed automata. In Susanne Graf and Wenhui Zhang, editors, *Automated Technology for Verification and Analysis, 4th International Symposium, ATVA 2006, Beijing, China, October 23-26, 2006*, volume 4218 of *Lecture Notes in Computer Science*, pages 292–306.

- Springer, 2006.
- [13] Franck Cassez, Thomas Chatain, and Claude Jard. Symbolic unfoldings for networks of timed automata. In Susanne Graf and Wenhui Zhang, editors, *Automated Technology for Verification and Analysis, 4th International Symposium, ATVA 2006, Beijing, China, October 23-26, 2006*, volume 4218 of *Lecture Notes in Computer Science*, pages 307–321. Springer, 2006.
- [14] Krishnendu Chatterjee, Andreas Pavlogiannis, and Viktor Toman. Value-centric dynamic partial order reduction. *Proc. ACM Program. Lang.*, 3(OOPSLA):124:1–124:29, 2019.
- [15] Edmund M. Clarke, Orna Grumberg, Marius Minea, and Doron A. Peled. State space reduction using partial order techniques. *Int. J. Softw. Tools Technol. Transf.*, 2(3):279–287, 1999.
- [16] Dennis Dams, Rob Gerth, Bart Knaack, and Ruurd Kuiper. Partial-order reduction techniques for real-time model checking. *Formal Aspects Comput.*, 10(5-6):469–482, 1998.
- [17] 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.
- [18] Rüdiger Ehlers, Daniel Fass, Michael Gerke, and Hans-Jörg Peter. Fully symbolic timed model checking using constraint matrix diagrams. In *Proceedings of the 31st IEEE Real-Time Systems Symposium, RTSS 2010, San Diego, California, USA, November 30 - December 3, 2010*, pages 360–371. IEEE Computer Society, 2010.
- [19] Cormac Flanagan and Patrice Godefroid. Dynamic partial-order reduction for model checking software. In Jens Palsberg and Martin Abadi, editors, *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, pages 110–121. ACM, 2005.
- [20] Patrice Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems - An Approach to the State-Explosion Problem*, volume 1032 of *Lecture Notes in Computer Science*. Springer, 1996.
- [21] Patrice Godefroid. Model checking for programming languages using verisoft. In Peter Lee, Fritz Henglein, and Neil D. Jones, editors, *Conference Record of POPL '97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the Symposium, Paris, France, 15-17 January 1997*, pages 174–186. ACM Press, 1997.
- [22] Patrice Godefroid and Pierre Wolper. A partial approach to model checking. *Inf. Comput.*, 110(2):305–326, 1994.
- [23] R. Govind, Frédéric Herbretreau, B. Srivathsan, and Igor Walukiewicz. Revisiting local time semantics for networks of timed automata. In *CONCUR*, volume 140 of *LIPICs*, pages 16:1–16:15. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2019.
- [24] Henri Hansen, Shang-Wei Lin, Yang Liu, Truong Khanh Nguyen, and Jun Sun. Diamonds are a girl's best friend: Partial order reduction for timed automata with abstractions. In *CAV*, volume 8559 of *Lecture Notes in Computer Science*, pages 391–406. Springer, 2014.
- [25] Frédéric Herbretreau, B. Srivathsan, and Igor Walukiewicz. Lazy abstractions for timed automata. In *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 990–1005. Springer, 2013.
- [26] Frédéric Herbretreau, B. Srivathsan, and Igor Walukiewicz. Better abstractions for timed automata. *Inf. Comput.*, 251:67–90, 2016.
- [27] Shmuel Katz and Doron A. Peled. Verification of distributed programs using representative interleaving sequences. *Distributed Comput.*, 6(2):107–120, 1992.
- [28] Michalis Kokologiannakis, Iason Marmanis, Vladimir Gladstein, and Viktor Vafeiadis. Truly stateless, optimal dynamic partial order reduction. *Proc. ACM Program. Lang.*, 6(POPL), jan 2022.
- [29] Kim G. Larsen, Marius Mikucionis, Marco Muñoz, and Jiri Srba. Urgent partial order reduction for extended timed automata. In Dang Van Hung and Oleg Sokolsky, editors, *Automated Technology for Verification and Analysis - 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings*, volume 12302 of *Lecture Notes in Computer Science*, pages 179–195. Springer, 2020.
- [30] Kim Guldstrand Larsen, Fredrik Larsson, Paul Pettersson, and Wang Yi. Compact data structures and state-space reduction for model-checking real-time systems. *Real Time Syst.*, 25(2-3):255–275, 2003.
- [31] 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.
- [32] 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.
- [33] Jesper B. Møller, Jakob Lichtenberg, Henrik Reif Andersen, and Henrik Hulgaard. Fully symbolic model checking of timed systems using difference decision diagrams. *Electron. Notes Theor. Comput. Sci.*, 23(2):88–107, 1999.
- [34] Peter Niebert, Moez Mahfoudh, Eugene Asarin, Marius Bozga, Oded Maler, and Navendu Jain. Verification of timed automata via satisfiability checking. In Werner Damm and Ernst-Rüdiger Olderog, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems, 7th International Symposium, FTRTFT 2002, Co-sponsored by IFIP WG 2.2, Oldenburg, Germany, September 9-12, 2002, Proceedings*, volume 2469 of *Lecture Notes in Computer Science*, pages 225–244. Springer, 2002.
- [35] Doron A. Peled. All from one, one for all: on model checking using representatives. In Costas Courcoubetis, editor, *Computer Aided Verification, 5th International Conference, CAV '93, Elounda, Greece, June 28 - July 1, 1993, Proceedings*, volume 697 of *Lecture Notes in Computer Science*, pages 409–423. Springer, 1993.
- [36] Doron A. Peled, Antti Valmari, and Ilkka Kokkarinen. Relaxed visibility enhances partial order reduction. *Formal Methods Syst. Des.*, 19(3):275–289, 2001.
- [37] Antti Valmari. Stubborn sets for reduced state space generation. In Grzegorz Rozenberg, editor, *Advances in Petri Nets 1990 [10th International Conference on Applications and Theory of Petri Nets, Bonn, Germany, June 1989, Proceedings]*, volume 483 of *Lecture Notes in Computer Science*, pages 491–515. Springer, 1989.
- [38] Antti Valmari. A stubborn attack on state explosion. *Formal Methods Syst. Des.*, 1(4):297–322, 1992.
- [39] Farn Wang. Symbolic verification of complex real-time systems with clock-restriction diagram. In Myungchul Kim, Byoungmoon Chin, Sungwon Kang, and Danhyung Lee, editors, *Formal Techniques for Networked and Distributed Systems, FORTE 2001, IFIP TC6/WG6.1 - 21st International Conference on Formal Techniques for Networked and Distributed Systems, August 28-31, 2001, Cheju Island, Korea*, volume 197 of *IFIP Conference Proceedings*, pages 235–250. Kluwer, 2001.
- [40] Naling Zhang, Markus Kusano, and Chao Wang. Dynamic partial order reduction for relaxed memory models. In David Grove and Stephen M. Blackburn, editors, *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, pages 250–259. ACM, 2015.