

# Parameterized Model-Checking of Timed Systems with Conjunctive Guards

Luca Spalazzi and Francesco Spegni<sup>(✉)</sup>

DII - Università Politecnica Delle Marche, Ancona, Italy  
{spalazzi, spegni}@dii.univpm.it

**Abstract.** In this work we extend the Emerson and Kahlon’s cutoff theorems for process skeletons with conjunctive guards to Parameterized Networks of Timed Automata, i.e. systems obtained by an *a priori* unknown number of Timed Automata instantiated from a finite set  $U_1, \dots, U_n$  of Timed Automata templates. In this way we aim at giving a tool to universally verify software systems where an unknown number of software components (i.e. processes) interact with continuous time temporal constraints. It is often the case, indeed, that distributed algorithms show an heterogeneous nature, combining dynamic aspects with real-time aspects. In the paper we will also show how to model check a protocol that uses special variables storing identifiers of the participating processes (i.e. PIDs) in Timed Automata with conjunctive guards. This is non-trivial, since solutions to the parameterized verification problem often relies on the processes to be symmetric, i.e. indistinguishable. On the other side, many popular distributed algorithms make use of PIDs and thus cannot directly apply those solutions.

## 1 Introduction

Software model-checking emerged as a natural evolution of applying model checking to verify hardware systems. Some factors, among several ones, that still make software model checking challenging are: the inherently dynamic nature of software components, the heterogeneous nature of software systems and the relatively limited amount of modular tools (both theoretical and practical) for verifying generic software systems.

Software systems definable as an arbitrary number of identical copies of some process template, are called parameterized systems, and are an example of infinite state systems [17]. Sometimes the nature of a software system is heterogeneous, meaning that it combines several “characteristics” (e.g. a clock synchronization algorithm is supposed to work with an arbitrary number of processes but also to terminate within a certain time). The scarcity of modular tools is witnessed by the fact that almost everyone trying to model check a software system, has to build his/her own toolchain that applies several intermediate steps (usually translations and abstractions) before building a model that can be actually model checked.

Despite such obstacles, several industries already apply model checking as part of their software design and/or software testing stages. (e.g., Microsoft [8], NASA [25], Bell Labs [20], IBM [9], UP4ALL<sup>1</sup>). In the aerospace industry, the DO178C international standard [27] even consider software model checking (or more generally, software verification) an alternative to software testing, under suitable assumptions.

The core of our work is an extension of the Emerson and Kahlon’s Cutoff Theorem [15] to *parameterized and timed systems*. Assuming a parameterized system based on Timed Automata  $U_1, \dots, U_m$  that synchronize using conjunctive Boolean guards, the cutoff theorem allows to compute a list of positive numbers  $(c_1, \dots, c_m)$  such that, let  $\phi$  be a given specification, then:

$$\begin{aligned} \forall i \in [1, m]. (\forall n_i \in [0, \infty) . (U_1, \dots, U_m)^{(n_1, \dots, n_m)} \models \phi \text{ iff} \\ \forall n_i \in [0, c_i] . (U_1, \dots, U_m)^{(n_1, \dots, n_m)} \models \phi) \end{aligned}$$

Intuitively, the proof shows that the cutoff configuration is *trace equivalent* to each “bigger” system.

The contribution of this work is multifold, w.r.t. the aforementioned factors: it reduces the problem of model checking an *infinite state* real-time software system to model checking a finite number of finite state systems; it shows a concrete example of how to combine verification algorithms from distinct domains, to verify what we call a *heterogeneous* software systems; the cutoff theorem for real-time systems is a theoretical tool that can be applied as a first step when verifying a parameterized and real-time algorithm. A second contribution is methodological: this paper describes how to exploit the cutoff theorem to model variables that store process identifiers (PIDs) of processes participating to the distributed algorithm. This is non trivial, since the former relies on the fact that processes should be symmetrical, thus indistinguishable. In order to show this, we will use a popular benchmark protocol, viz. the Fischer’s protocol for mutual exclusion. To the best of our knowledge, this is the first time that the Fischer’s protocol has been verified using model checking techniques, for an *a priori* unknown number of processes.

## 2 Related Work

**Infinite State System.** Timed Automata and Parameterized Systems are two examples of infinite state systems [17]. In general, the problem of model checking infinite state systems is undecidable [6]. A classic approach to overcome this limitation, is to find suitable subsets of infinite state systems that can be reduced to model checking of finitely many finite state systems, e.g. identifying a precise abstraction (e.g. clock-zones for Timed Automata [10]). Other approaches are based on the idea of finding a finite-state abstraction that is correct but not complete, such that a property verified for the abstract system holds for the

<sup>1</sup> <http://www.uppaal.com/index.php?sida=203&&rubrik=92> URL visited on April’14.

original system as well [7, 14, 19, 30]. Some other approaches are based on the idea of building an invariant representing the common behaviors exhibited by the system [24]. When a given relation over the invariant is satisfied, then the desired property is satisfied by the original system. Its limitation is that building the abstraction or the invariant is usually not automatic.

**Cutoffs for Parameterized Systems.** Concerning the use of cutoff for model checking parameterized systems, there exists two main approaches: computing the cutoff number of process replications or the cutoff length of paths. The former consists in finding a finite number of process instances such that if they satisfy a property then the same property is satisfied by an arbitrary number of such processes. Emerson and Kahlon [15] established a cutoff value of about the number of template states, for a clique of interconnected process skeletons. In the case of rings, a constant between 2 and 5 is enough [18]. For shared resources management algorithms [11], the cutoff value is the number of resources plus the quantified processes (in the decidable fragment of processes with equal priority). Other works proved that one process per template is enough, for certain grids [26]. Recently, in [5] it has been showed that certain parameterized systems may admit a cutoff which is not computable, while Hanna *et al.* [22] proposed a procedure to compute a cutoff for Input-Output Automata that is independent of the communication topology. On the other hand, computing the cutoff length of paths of a parameterized system consists in finding an upper bound on the number of nodes in its longest computation path. When a property is satisfied within the bounded path, then the property holds for a system with unbound paths, i.e., with an arbitrary number of process instances. The classic work from German and Sistla [19], Emerson and Namjoshi [16] proved that such a cutoff exists for the verification of parameterized systems composed of a control process and an arbitrary number of user processes against indexed LTL properties. Yang and Li [29] proposed a sound and complete method to compute such a cutoff for parameterized systems with only rendezvous actions. In that work, the property itself is represented as an automaton. Lately it has been also showed that parameterized systems on pairwise rendezvous do not admit, in general, a cutoff [7]. To the best of our knowledge, cutoff theorems have not been stated previously for timed systems. Surprisingly enough, extending Emerson and Kahlon cutoff theorems [15] to timed systems does not increase the cutoff value.

**Parameterized Networks of Timed or Hybrid Automata.** The realm of real-time systems (timed automata and, more in general, hybrid automata) with a finite but unknown number of instances has been explored. Abdulla and Jonsson [1] proposed in their seminal work to reduce safety properties to reachability properties. They worked with a network composed by an arbitrary set of identical timed automata controlled by a controller (i.e. a finite timed automaton as well). Abdulla *et al.* show also that checking safety properties in networks of timed automata with multiple clocks is an undecidable problem [2], as well as the problem of determining if a state is visited infinitely often, in the continuous time model (in the discrete time model, instead, it is decidable) [3]. It should be remarked that in their undecidability proof, the network of

timed automata must rely on synchronous rendezvous in order to prove the undecidability results. This motivated us to explore timed automata with different synchronization mechanisms in this work. Ghilardi et al. [13], reduced model checking safety properties to reachability problem. Similarly to Abdulla and Jonsson, they applied their approach to networks composed by an arbitrary set of timed automata interacting with a controller. Their original contribution consisted in the usage of Satisfiability Modulo Theories techniques. Göthel and Glesner [21] proposed a semi-automatic verification methodology based on finding network invariants and using both theorem proving and model checking. Along the same line, Johnson and Mitra [23] proposed a semi-automatic verification of safety properties for parameterized networks of hybrid automata with rectangular dynamics. They based their approach on a combination of invariant synthesis and inductive invariant proving. Their main limitation is that specifications are often not inductive properties (e.g. the mutual exclusion property it is not an inductive property). In this case one must show that a set of inductive invariants can imply the desired property. This last step is often not fully automatic.

We consider systems composed of a finite number of templates, each of which can be instantiated an arbitrary number of times. We limit Timed Automata to synchronize using Conjunctive Guards, instead of the classic Pairwise Rendezvous [10], because, as already mentioned, parameterized systems with pairwise rendezvous do not admit, in general, a cutoff [7]. Finally, the verification proposed in this paper is completely automatic.

### 3 Parameterized Networks of Timed Automata

This work introduces Parameterized Networks of Timed Automata (PNTA), an extensions of Timed Automata that synchronize using conjunctive Boolean guards. We also introduce Indexed-Timed CTL\*, a temporal logic that integrates TCTL and MTL [12], for reasoning about timed processes, together with Indexed-CTL\* $\setminus$ X [15], for reasoning about parametric networks of processes. In the following definition we will make use of a set of *temporal constraints*  $TC(C_l)$ , defined as:

$$TC(C) ::= \top \mid \neg TC(C) \mid TC(C) \vee TC(C) \mid \\ C \sim C \mid C \sim \mathbb{Q}^{\geq 0}$$

where  $\sim \in \{<, \leq, >, \geq, =\}$ ,  $C$  is a set of clock variables and  $\mathbb{Q}$  denotes the set of rational numbers.

**Definition 1 (Timed Automaton Template).** A Timed Automaton (TA) Template  $U_l$  is a tuple  $\langle S_l, \hat{s}_l, C_l, \Gamma_l, \tau_l, I_l \rangle$  where:

- $S_l$  is a finite set of states, or locations;
- $\hat{s}_l \in S_l$  is a distinguished initial state;
- $C_l$  is a finite set of clock variables;

- $\Gamma_l$  is a finite set of Boolean guards built upon  $S_l$ ;
- $\tau_l \subseteq S_l \times TC(C_l) \times 2^{C_l} \times \Gamma_l \times S_l$  is a finite set of transitions;
- $I_l : S_l \rightarrow TC(C_l)$  maps a state to an invariant, such that  $I_l(\hat{s}_l) = \top$ ;

We will denote with  $|U_l| = |S_l|$  the size of the timed automaton. A network of timed automata can be defined as a set of  $k$  TA templates, where each TA template (say  $U_l$ ) is instantiated an arbitrary number (say  $n_l$ ) of times.

**Definition 2 (PNTA).** Let  $(U_1, \dots, U_k)$  be a set of Timed Automaton templates. Let  $(n_1, \dots, n_k)$  be a set of natural numbers. Then

$$(U_1, \dots, U_k)^{(n_1, \dots, n_k)}$$

is a Parameterized Network of Timed Automata denoting the asynchronous parallel composition of timed automata  $U_1^1 || \dots || U_1^{n_1} || \dots || U_k^1 || \dots || U_k^{n_k}$ , such that for each  $l \in [1, k]$  and  $i \in [1, n_l]$ , then  $U_l^i$  is the  $i$ -th copy of  $U_l$ .

Let us remark that every component of  $U_l^i$  is a disjoint copy of the corresponding template component. In the following will be described how every process  $U_l^i$ , also called instance, can take a local step after having checked that the neighbors' states satisfy the transition (conjunctive) Boolean guard. In such system a process can check it is "safe" to take a local step, but it cannot induce a move on a different instance. A PNTA based on conjunctive guards is defined as follows.

**Definition 3 (PNTA with Conjunctive Guards).** Let  $(U_1, \dots, U_k)^{(n_1, \dots, n_k)}$  be a PNTA. Then, it is a PNTA with Conjunctive Guards iff every  $\gamma \in \Gamma_l^i$  is a Boolean expression with the following form:

$$\bigwedge_{\substack{m \in [1, n_1] \\ m \neq i}} (\hat{s}_l(m) \vee s_l^1(m) \vee \dots \vee s_l^p(m)) \wedge \bigwedge_{\substack{h \in [1, k] \\ h \neq l}} \left( \bigwedge_{j \in [1, n_j]} (\hat{s}_h(j) \vee s_h^1(j) \vee \dots \vee s_h^q(j)) \right)$$

where, for all  $l \in [1, k]$ ,  $i \in [1, n_l]$  and  $p > 0$ ,  $\{s_l^1, \dots, s_l^p\} \subseteq S_l$ ,  $s_l(i) \in S_l^i$  and  $\hat{s}_l$  is the initial states of  $U_l$ . The initial states  $\hat{s}_l(m)$  and  $\hat{s}_h(j)$  must be present.

We remark that our definitions of Timed Automaton template, PNTA and PNTA with Conjunctive Guards are variants of the notion of *timed automata* and *networks of timed automata* found in literature (e.g. [10]).

The *operational semantics* of PNTA with conjunctive guards is expressed as a transition system over *PNTA configurations*.

**Definition 4 (PNTA Configuration).** Let  $(U_1, \dots, U_k)^{(n_1, \dots, n_k)}$  be a PNTA. Then a configuration is a tuple:

$$\mathbf{c} = (\langle \bar{s}_1, \bar{u}_1 \rangle, \dots, \langle \bar{s}_k, \bar{u}_k \rangle)$$

where, for each  $l \in [1, k]$ :

- $\bar{s}_l : [1, n_l] \rightarrow S_l$  maps an instance to its current state, and
- $\bar{u}_l : [1, n_l] \rightarrow (C_l \rightarrow \mathbb{R}^{\geq 0})$ , maps an instance to its clock function, s.t.

$$\forall i . \bar{u}_l(i) \models I_l^i(\bar{s}_l(i)) \quad (1)$$

$\mathfrak{C}$  is the set of all the configurations.

Intuitively, let  $(\dots, \langle \bar{s}_l, \bar{u}_l \rangle, \dots)$  be a configuration, then  $\bar{s}_l(i) \in S_l$  denotes the state where instance  $U_l^i$  is in that configuration.  $\bar{u}_l(i)$  is the clock assignment function (i.e.,  $\bar{u}_l(i) : C_l \rightarrow \mathbb{R}^{\geq 0}$ ) of instance  $U_l^i$  in that configuration. In other words, for each  $c \in C_l$ ,  $\bar{u}_l(i)(c)$  is the current value that the clock variable  $c$  assumes for instance  $U_l^i$ . Any assignment to such clock variables must satisfy the invariant for the corresponding state (see Eq. (1)). The notion of transition requires some auxiliary notations. Let  $l \in [1, k]$ , and let  $i \in [1, n_l]$ , then we call:

- *initial configuration*  
 $\hat{\mathbf{c}} \in \mathfrak{C}$  such that, for each  $l \in [1, k]$ , for each  $i \in [1, n_k]$ :

$$\begin{aligned} \bar{s}_l(i) &= \hat{s}_l^i, \text{ and} \\ \forall c \in C_l, \bar{u}_l(i)(c) &= 0. \end{aligned}$$

- *projection*  
 $\forall \mathbf{c} = (\langle \bar{s}_1, \bar{u}_1 \rangle, \dots, \langle \bar{s}_l, \bar{u}_l \rangle, \dots, \langle \bar{s}_k, \bar{u}_k \rangle) \in \mathfrak{C}$ ,  
 $\mathbf{c}(l) = \langle \bar{s}_l, \bar{u}_l \rangle$ , and  
 $\mathbf{c}(l, i) = \langle \bar{s}_l(i), \bar{u}_l(i) \rangle$ .

- *state-component*  
 $\forall \mathbf{c} = (\langle \bar{s}_1, \bar{u}_1 \rangle, \dots, \langle \bar{s}_l, \bar{u}_l \rangle, \dots, \langle \bar{s}_k, \bar{u}_k \rangle) \in \mathfrak{C}$ ,  
 $state(\mathbf{c}) = (\bar{s}_1, \dots, \bar{s}_l, \dots, \bar{s}_k)$ ,  
 $state(\mathbf{c}(l)) = \bar{s}_l$ , and  
 $state(\mathbf{c}(l, i)) = \bar{s}_l(i)$ .

- *clock-component*  
 $\forall \mathbf{c} = (\langle \bar{s}_1, \bar{u}_1 \rangle, \dots, \langle \bar{s}_l, \bar{u}_l \rangle, \dots, \langle \bar{s}_k, \bar{u}_k \rangle) \in \mathfrak{C}$ ,  $\forall c \in C_l$ ,  
 $clock(\mathbf{c}) = (\bar{u}_1, \dots, \bar{u}_l, \dots, \bar{u}_k)$ ,  
 $clock(\mathbf{c}(l)) = \bar{u}_l$ ,  
 $clock(\mathbf{c}(l, i)) = \bar{u}_l(i)$ , thus  
 $clock(\mathbf{c}(l, i))(c) = \bar{u}_l(i)(c)$ .

- *time increase*  
 $\forall c \in C_l. \forall d \in \mathbb{R}^{\geq 0}. (\bar{u}_l + d)(i)(c) = \bar{u}_l(i)(c) + d$   
 $(clock(\mathbf{c}) + d) = (\bar{u}_1 + d, \dots, \bar{u}_l + d, \dots, \bar{u}_k + d)$ ,  
 $(clock(\mathbf{c}(l)) + d) = (\bar{u}_l + d)$ , and  
 $(clock(\mathbf{c}(l, i)) + d) = (\bar{u}_l + d)(i)$ .

- *clock reset*  
 $\forall c \in C_l . \forall r \subseteq C_l . \forall j .$

$$\bar{u}_l[(i, r) \mapsto 0](j)(c) = \begin{cases} 0 & \text{if } i = j \text{ and } c \in r \\ \bar{u}_l(j)(c) & \text{otherwise} \end{cases}$$

- *clock constraint evaluation*  
 $\bar{u}_l(i) \models g$  iff the clock values of instance  $U_l^i$  denoted by  $\bar{u}_l(i)$  satisfy the clock constraint  $g$ ; the semantics  $\models$  is defined as usual by induction on the structure of  $g$ ;

– *guard evaluation*

$state(\mathbf{c}) \models \gamma$  iff the set of states of  $(U_1, \dots, U_k)^{(n_1, \dots, n_k)}$  denoted by  $state(\mathbf{c})$  satisfies the Boolean guard  $\gamma$ ; this predicate as well can be defined by induction on the structure of  $\gamma$ .

**Definition 5 (PNTA Transitions).** *The transitions among PNTA configurations are governed by the following rules:*

(*delay*)

$$\mathbf{c} \xrightarrow{d} \mathbf{c}' \quad \text{if} \quad d \in \mathbb{R}^{\geq 0}$$

$$\begin{aligned} & state(\mathbf{c}') = state(\mathbf{c}) \\ & clock(\mathbf{c}') = (clock(\mathbf{c}) + d) \\ & \forall l, i, d' \in [0, d]. clock(\mathbf{c}(l, i)) + d' \models I_l^i(state(\mathbf{c}(l, i))) \end{aligned}$$

(*synchronization*)

$$\mathbf{c} \xrightarrow{\gamma} \mathbf{c}' \quad \text{if} \quad \exists l \in [1, k]. \exists i \in [1, n_l] :$$

$$\begin{aligned} & s \xrightarrow{g, r, \gamma} t \in \tau_l^i. \\ & state(\mathbf{c}(l, i)) = s, \\ & clock(\mathbf{c}(l, i)) \models g, \\ & state(\mathbf{c}) \models \gamma, \\ & \mathbf{c}'(h) = \mathbf{c}(h) \text{ for each } h \neq l, \\ & \mathbf{c}'(l, j) = \mathbf{c}(l, j) \text{ for each } j \neq i, \\ & \mathbf{c}'(l, i) = \langle t, clock(\mathbf{c}(l, i))[(r, i) \mapsto 0] \rangle \\ & clock(\mathbf{c}'(l, i)) \models I_l^i(state(\mathbf{c}'(l, i))) \end{aligned}$$

Let us define what is a *timed-computation* for PNTA.

**Definition 6 (Timed Computation).** *Let  $\hat{\mathbf{c}}_0$  be an initial configuration, a timed-computation  $x$  is a finite or infinite sequence of pairs:*

$$x = (\mathbf{c}_0, t_0) \dots (\mathbf{c}_v, t_v) \dots$$

*s.t.*  $t_0 = 0$  and  $\forall v \geq 0 . (\exists d > 0 . \mathbf{c}_v \xrightarrow{d} \mathbf{c}_{v+1} \wedge t_{v+1} = t_v + d) \vee (\exists \gamma . \mathbf{c}_v \xrightarrow{\gamma} \mathbf{c}_{v+1} \wedge t_{v+1} = t_v)$

In other words, a timed computation can be seen as a sequence of *snapshots* of the transition system configurations taken at successive times. It should be noticed that, according to Emerson and Kahlon [15], in this work, it has been adopted the so-called *interleaving semantics*. This means that in a transition between two configurations, only one instance can change its state (see the *synchronization* rule in Definition 5). For the sake of conciseness, let us extend the notion of *projection*, *state-component*, and *clock-component* to timed computations. Let  $x = (\mathbf{c}_0, t_0) \dots (\mathbf{c}_v, t_v) \dots$  be a timed computation, let  $x_v = (\mathbf{c}_v, t_v)$  be the  $v$ -th element of  $x$ , then

$$\begin{array}{llll}
x(l) & = (\mathbf{c}_0(l), t_0) \dots (\mathbf{c}_v(l), t_v) \dots & \text{clock}(x_v) & = \text{clock}(\mathbf{c}_v) \\
x(l, \hat{i}) & = (\mathbf{c}_0(l, \hat{i}), t_0) \dots (\mathbf{c}_v(l, \hat{i}), t_v) \dots & \text{clock}(x_v(l)) & = \text{clock}(\mathbf{c}_v(l)) \\
x_v(l) & = (\mathbf{c}_v(l), t_v) & \text{clock}(x_v(l, \hat{i})) & = \text{clock}(\mathbf{c}_v(l, \hat{i})) \\
x_v(l, \hat{i}) & = (\mathbf{c}_v(l, \hat{i}), t_v) & & \\
\text{state}(x_v) & = \text{state}(\mathbf{c}_v) & \text{time}(x_v) & = t_v \\
\text{state}(x_v(l)) & = \text{state}(\mathbf{c}_v(l)) & \text{time}(x_v(l)) & = t_v \\
\text{state}(x_v(l, \hat{i})) & = \text{state}(\mathbf{c}_v(l, \hat{i})) & \text{time}(x_v(l, \hat{i})) & = t_v
\end{array}$$

$x(l, \hat{i})$  is called the *local computation* of the  $i$ -th instance of automaton template  $l$ .  $\text{time}(x_v)$ ,  $\text{time}(x_v(l))$ , and  $\text{time}(x_v(l, \hat{i}))$  are the *time-components* of  $x_v$ ,  $x_v(l)$ , and  $x_v(l, \hat{i})$  respectively.

**Definition 7 (Idle Local Computation).** Let  $U_l^i = \langle S_l^i, \hat{s}_l^i, C_l^i, \tau_l^i, I_l^i \rangle$  be the  $i$ -th instance of the timed automaton template  $U_l$ . An idle local computation  $\hat{s}(l, \hat{i})$  is a timed local computation such that, for all  $v \geq 0$ :

$$\begin{aligned}
\hat{s}(l, \hat{i}) &= (\langle \hat{s}_l^i, \bar{u}_l(i) \rangle, t_0) \dots (\langle \hat{s}_l^i, \bar{u}_l(i) + t_v \rangle, t_v) \dots \\
\hat{s}_v(l, \hat{i}) &= (\langle \hat{s}_l^i, \bar{u}_l(i) + t_v \rangle, t_v)
\end{aligned}$$

where  $t_0 = 0$  and for each  $c \in C_l$ ,  $\bar{u}_l(i)(c) = 0$ .

It should be noticed that for each  $v$ , it must be  $\bar{u}_l(i) + t_v \models I_l^i(\hat{s}_l^i)$ , since  $I_l^i(\hat{s}_l^i) = \top$  according to Definition 1. Intuitively, an idle local computation is an instance of the automaton template  $U_l$  that stutters in its initial state.

**Definition 8 (Stuttering).** Let  $x$  and  $y$  be two timed computations. Let  $x = x_0 \cdot \dots \cdot x_v \cdot x_{v+1} \dots$ . The timed computation  $y$  is a stuttering of the timed computation  $x$  iff for all  $v \geq 0$ , there exists  $r \geq 0$ , such that

$$y = x_0 \cdot \dots \cdot x_v \cdot x_{v, \delta_1} \cdot x_{v, \delta_2} \cdot \dots \cdot x_{v, \delta_r} \cdot x_{v+1} \dots$$

where  $\delta_1, \delta_2, \dots, \delta_r \in \mathbb{R}^{\geq 0}$ ,  $\delta_1 \leq \delta_2 \leq \dots \leq \delta_r$ ,  $t_v + \delta_r \leq t_{v+1}$ , and

$$\begin{aligned}
x_{v, \delta_1} &= (\langle \text{state}(x_v), \text{clock}(x_v) + \delta_1 \rangle, t_v + \delta_1) \\
x_{v, \delta_2} &= (\langle \text{state}(x_v), \text{clock}(x_v) + \delta_2 \rangle, t_v + \delta_2) \\
&\dots \\
x_{v, \delta_r} &= (\langle \text{state}(x_v), \text{clock}(x_v) + \delta_r \rangle, t_v + \delta_r)
\end{aligned}$$

Intuitively, the above definition means that a stuttering of a given timed computation  $x$  can be generated by inserting an arbitrary number of *delay transitions* (see Definition 5) short enough to not alter the validity of temporal conditions of the original computation  $x$ . It only represents a more detailed view (i.e. a finer sampling) of the interval between a configuration and the next one without changing the original sequence of states.

For the purpose of this work, timed computations conforming to Definition 6 (i.e. each configuration complies with Eq. (1)) can be classified in three different kinds of computation:



- *Infinite Timed Computation*:  $x$  is a timed computation of infinite length.
- *Deadlocked Timed Computation*:  $x$  is a maximal finite timed computation, i.e. in it reaches a final configuration where all transitions are disabled.
- *Finite Timed Computation*:  $x$  is a (not necessarily maximal) final timed computation, i.e. it is either a deadlocked computation or a finite prefix of an infinite one.

## 4 A Temporal Logic for PNTA

A dedicated logic is needed in order to specify behaviors of a PNTA. This logic, named Indexed-Timed-CTL<sup>\*</sup>, allows to reason about real-time intervals and temporal relations (until, before, after, ...) in systems of arbitrary size. While its satisfiability problem is undecidable, the problem of model checking a PNTA is proved to be decidable, under certain conditions.

**Definition 9 (Indexed-Timed-CTL<sup>\*</sup>).** *Let  $\{P_l\}_{l \in [1, k]}$  be finite sets of atomic propositions. Let  $p(l, i)$  be any atomic proposition such that  $l \in [1, k]$ ,  $i \in \mathbb{N}^{>0}$ , and  $p \in P_l$ . Then, the set of ITCTL<sup>\*</sup> formulae is inductively defined as follows:*

$$\begin{aligned} \phi &::= \top \mid p(l, i) \mid \phi \wedge \phi \mid \neg\phi \mid \bigwedge_{i_l} \phi \mid A\Phi \mid A_{fin}\Phi \mid A_{inf}\Phi \\ \Phi &::= \phi \mid \Phi \wedge \Phi \mid \neg\Phi \mid \Phi \mathcal{U}_{\sim_q} \Phi \end{aligned}$$

where  $\sim \in \{<, \leq, \geq, >, =\}$  and  $q \in \mathbb{Q}^{\geq 0}$ .

As usual for branching-time temporal logics, the terms in  $\phi$  denote *state* formulae, while terms in  $\Phi$  denote *path* formulae. For the purpose of this work it is enough to assume the set of atomic propositions coincides with the set of states of a given PNTA, i.e.  $P_l = S_l$ , for every  $l$ .

The path quantifier  $A_{fin}$  (resp.  $A_{inf}$ ) is a variant of the usual universal path quantifier  $A$ , restricted to paths that are of finite length (resp. infinite length). Such variants are inspired by [15]. Missing Boolean ( $\vee, \rightarrow, \dots$ ) operators, temporal operators ( $\mathcal{G}, \mathcal{F}, \mathcal{W}, \dots$ ), as well as path quantifiers ( $E, E_{fin}, E_{inf}$ ) can be defined as usual. The semantics of ITCTL<sup>\*</sup> is defined w.r.t. a Kripke Structure integrating the notions of parametric system size and continuous time semantics [12]. The continuous time model requires that between any two configurations it always exists a third state. It is possible, though, introduce *continuous time computation trees* [4]. Let us call *s-path* a function  $\rho : \mathbb{R}^{\geq 0} \rightarrow \mathfrak{C}$  that intuitively maps a time  $t$  with the current system configuration at that time. The mapping  $\rho_{\lfloor t' } : [0, t') \rightarrow \mathfrak{C}$  is a *prefix* of  $\rho$  iff  $\forall t < t'. \rho_{\lfloor t' }(t) = \rho(t)$ . The mapping  $\rho_{\lceil t' } : [t', \infty) \rightarrow \mathfrak{C}$  is a *suffix* of  $\rho$  iff  $\forall t \geq t'. \rho_{\lceil t' }(t) = \rho(t)$ . Let us take a prefix  $\rho_{\lfloor t' }$  and an s-path  $\rho'$ , then their *concatenation* is defined as:

$$(\rho_{\lfloor t' } \cdot \rho')(t) = \begin{cases} \rho_{\lfloor t' }(t) & \text{if } t < t' \\ \rho'(t - t') & \text{else} \end{cases}$$

Let  $\Pi$  be a set of s-paths, then  $\rho_{\downarrow t'} \cdot \Pi = \{\rho_{\downarrow t'} \cdot \rho' : \rho' \in \Pi\}$ . A *continuous time computation tree* is a mapping  $f : \mathfrak{C} \rightarrow 2^{[\mathbb{R}^{\geq 0} \rightarrow \mathfrak{C}]}$  such that:

$$\forall \mathfrak{c} \in \mathfrak{C}. \forall \rho \in f(\mathfrak{c}). \forall t \in \mathbb{R}^{\geq 0} . \rho_{\downarrow t} \cdot f(\rho(t)) \subseteq f(\mathfrak{c}).$$

For the purpose of this work, here only *s-paths* defined over *timed computations* will be considered.

**Definition 10 (PNTA s-paths).** For each timed computation  $x = (\mathfrak{c}_0, t_0) \dots (\mathfrak{c}_v, t_v) \dots$ , let us call PNTA s-path the s-path  $\rho : \mathbb{R}^{\geq 0} \rightarrow \mathfrak{C}$  satisfying:

$$\forall v. \forall t \in [t_v, t_{v+1}) . \rho(t) = \langle s, c \rangle$$

where  $s = \text{state}(\mathfrak{c}_v)$  and  $c = \text{clock}(\mathfrak{c}_v) + t - t_v$ .

It should be noticed that, according to the above construction, an infinite set of timed computations can generate the same s-path  $\rho$ ; let us denote such set by  $tcomp(\rho)$ . As a consequence, for each  $y \in tcomp(\rho)$ , there exists  $x \in tcomp(\rho)$  such that  $y$  is a stuttering of  $x$  (see Definition 8). The continuous semantics of ITCTL\* can be defined as follows.

**Definition 11 (Satisfiability of ITCTL\*).** Let  $(U_1, \dots, U_k)^{(n_1, \dots, n_k)}$  be a PNTA and  $\mathfrak{c}$  be the current configuration. Let  $\phi$  denote an ITCTL\* state formula, then the satisfiability relation  $\mathfrak{c} \models \phi$  is defined by structural induction as follows:

$$\begin{aligned} \mathfrak{c} &\models \top \\ \mathfrak{c} &\models p(l, i) && \text{iff } p = \text{state}(\mathfrak{c}(l, i)) \\ \mathfrak{c} &\models \phi_1 \wedge \phi_2 && \text{iff } \mathfrak{c} \models \phi_1 \text{ and } \mathfrak{c} \models \phi_2 \\ \mathfrak{c} &\models \neg \phi_1 && \text{iff } \mathfrak{c} \not\models \phi_1 \\ \mathfrak{c} &\models A\phi_1 && \text{iff } \rho \models \phi_1, \quad \text{for all } \rho \in f(\mathfrak{c}) \text{ and} \\ &&& (|\rho| = \omega \text{ or } \text{deadlock}(\rho)) \\ \mathfrak{c} &\models A_{inf}\phi_1 && \text{iff } \rho \models \phi_1, \quad \text{for all } \rho \in f(\mathfrak{c}) \text{ and } |\rho| = \omega \\ \mathfrak{c} &\models A_{fin}\phi_1 && \text{iff } \rho \models \phi_1, \quad \text{for all } \rho \in f(\mathfrak{c}) \text{ and } |\rho| < \omega \\ \mathfrak{c} &\models \bigwedge_{i_l} \phi(i_l) && \text{iff } \mathfrak{c} \models \phi_1(i_l), \quad \text{for each } i_l \in [1, n_l] \\ \\ \rho &\models \phi_1 && \text{iff } \rho(0) \models \phi_1 \\ \rho &\models \phi_1 \wedge \phi_2 && \text{iff } \rho \models \phi_1 \text{ and } \rho \models \phi_2 \\ \rho &\models \neg \phi_1 && \text{iff } \rho \not\models \phi_1 \\ \rho &\models \phi_1 \mathcal{U}_{\sim q} \phi_2 && \text{iff for some } t' \sim q, \text{ where } \sim \in \{<, \leq, \geq, >, =\} \\ &&& \rho_{\downarrow t'} \models \phi_2, \text{ and } \rho_{\downarrow t} \models \phi_1 \text{ for all } t \in [0, t') \end{aligned}$$

where  $|\rho| = \omega$  (resp.  $|\rho| < \omega$ , resp.  $\text{deadlock}(\rho)$ ) denotes that the s-path  $\rho$  has infinite length (resp. has finite length, resp. is deadlocked).

Note that a finite s-path is not necessarily deadlocked, since it can be a finite prefix of some infinite s-path. When a given PNTA  $(U_1, \dots, U_k)^{(n_1, \dots, n_k)}$  satisfies an ITCTL\* state-formula  $\phi$  at its initial configuration  $\hat{\mathfrak{c}}$ , this is denoted by

$$(U_1, \dots, U_k)^{(n_1, \dots, n_k)} \models \phi$$

**Theorem 1 (Undecidability of ITCTL\*).** *The satisfiability problem for ITCTL\* is undecidable.*

*Proof.* The satisfiability problem for TCTL is undecidable [4]. TCTL is included in ITCTL\*, therefore the latter is undecidable.

In the next section we will call *IMTL* the fragment of ITCTL\* having formulae with the following forms:  $\bigwedge_{i_l} Qh(i_l)$ , where  $Q \in \{A, A_{fin}, A_{inf}\}$  and in  $h$  only Boolean ( $\wedge$  and  $\neg$ ) and temporal ( $\mathcal{U}_{\sim q}$ ) operators are allowed. We will call *IMITL* the subset of IMTL where equality constraints (i.e.  $\mathcal{U}_{=q}$ ) are excluded.

## 5 Cutoff Theorem for PNTA with Conjunctive Guards

In this section we prove that a cutoff can be computed to make the PMCP of PNTAs with conjunctive guards decidable, for a suitable set of formulae. The system in which every template is instantiated as many times as its cutoff, will be called the *cutoff system*. Given two instantiations  $I = (U_1, \dots, U_k)^{(c_1, \dots, c_k)}$  and  $I' = (U_1, \dots, U_k)^{(c'_1, \dots, c'_k)}$ , such that all  $c'_i \geq c_i$  and at least one  $c'_j > c_j$ , it can be said that  $I'$  is *bigger* than  $I$ , written  $I' > I$ . The cutoff theorem states that given a cutoff system  $I$ , for each  $I' > I$ , both  $I'$  and  $I$  satisfy the same subset of ITCTL\* formulae.

**Theorem 2 (Conjunctive Cutoff Theorem).** *Let  $(U_1, \dots, U_k)$  be a set of TA templates with conjunctive guards. Let  $\phi = \bigwedge_{i_{l_1}, \dots, i_{l_h}} Q\Phi(i_{l_1}, \dots, i_{l_h})$  where  $Q \in \{A, A_{inf}, A_{fin}, E, E_{inf}, E_{fin}\}$  and  $\Phi$  is an IMTL formula and  $\{l_1, \dots, l_h\} \subseteq [1, k]$ . Then*

$$\begin{aligned} \forall (n_1, \dots, n_k). (U_1, \dots, U_k)^{(n_1, \dots, n_k)} \models \phi \text{ iff} \\ \forall (d_1, \dots, d_k) \preceq (c_1, \dots, c_k). (U_1, \dots, U_k)^{(d_1, \dots, d_k)} \models \phi \end{aligned}$$

where the cutoff  $(c_1, \dots, c_k)$  can be computed as follows:

- In case  $Q \in \{A_{inf}, E_{inf}\}$  (i.e., deadlocked or finite timed computations are ignored). Then  $c_l = 2$  if  $l \in \{l_1, \dots, l_h\}$ , and  $c_l = 1$  otherwise (i.e.  $l \in [1, k] \setminus \{l_1, \dots, l_h\}$ ).
- In case  $Q \in \{A_{fin}, E_{fin}\}$  (i.e. finite timed computations, either deadlocked or finite prefixes of infinite computations). Then  $c_l = 1$  for each  $l$ .
- In case  $Q \in \{A, E\}$  (i.e., infinite and deadlocked). Then  $c_l = 2|U_l| + 1$  if  $l \in \{l_1, \dots, l_h\}$ ;  $c_l = 2|U_l|$  otherwise (i.e.  $l \in [1, k] \setminus \{l_1, \dots, l_h\}$ ).

The proof of the Cutoff Theorem consists of three steps. The first step (*Conjunctive Monotonicity Lemma*) shows that adding instances to the system does not alter the truth of logic formulae. The second step (*Conjunctive Bounding Lemma*) proves that removing an instance beyond the cutoff number, does not alter the truth of logic formulae either. The third step (*Conjunctive Truncation Lemma*) generalizes the Conjunctive Bounding Lemma to a system that has two automaton templates with an arbitrary number of instances. The given proofs can be generalized to systems with an arbitrary number of templates.

**Theorem 3 (Conjunctive Monotonicity Lemma).** *Let  $U_1$  and  $U_2$  be two TA templates with conjunctive guards. Let  $\Phi(1_l)$  be an IMTL formula, with  $l \in \{1, 2\}$ . Then for any  $n \in \mathbb{N}$  such that  $n \geq 1$  we have:*

$$\begin{aligned} (i) \quad & (U_1, U_2)^{(1,n)} \models Q\Phi(1_2) \Rightarrow (U_1, U_2)^{(1,n+1)} \models Q\Phi(1_2) \\ (ii) \quad & (U_1, U_2)^{(1,n)} \models Q\Phi(1_1) \Rightarrow (U_1, U_2)^{(1,n+1)} \models Q\Phi(1_1) \end{aligned}$$

where  $Q \in \{E, E_{inf}, E_{fin}\}$ .

A detailed proof of the theorem is in the extended version of this paper [28]. Intuitively, from any time computation  $x$  one can build a new time computation  $y$  where each instance behaves as in  $x$ , except for a new instance of  $U_2$  that halts in its initial state (remember that by definition the initial states don't falsify any conjunctive guard).

**Theorem 4 (Conjunctive Bounding Lemma).** *Let  $U_1$  and  $U_2$  be two TA templates with conjunctive guards. Let  $\Phi(1_l)$  be an IMTL formula, with  $l \in \{1, 2\}$ . Then for any  $n \in \mathbb{N}$  such that  $n \geq 1$  we have:*

$$\begin{aligned} (i) \quad & \forall n \geq c_2. (U_1, U_2)^{(1,n)} \models Q\Phi(1_2) \Rightarrow (U_1, U_2)^{(1,c_2)} \models Q\Phi(1_2) \\ (ii) \quad & \forall n \geq c_1. (U_1, U_2)^{(1,n)} \models Q\Phi(1_1) \Rightarrow (U_1, U_2)^{(1,c_1)} \models Q\Phi(1_1) \end{aligned}$$

where  $Q \in \{E, E_{inf}, E_{fin}\}$  and:

- $c_1 = 1$  and  $c_2 = 2$ , when  $Q = E_{inf}$ ;
- $c_1 = c_2 = 1$ , when  $Q = E_{fin}$ ;
- $c_1 = 2|U_2|$  and  $c_2 = 2|U_2| + 1$ , when  $Q = E$ .

**Theorem 5 (Truncation Lemma).** *Let  $U_1$  and  $U_2$  be two TA templates with conjunctive guards. Let  $\Phi(1_l)$  be an IMTL formula, with  $l \in \{1, 2\}$ , then:*

$$\forall n_1, n_2 \geq 1. (U_1, U_2)^{(n_1, n_2)} \models Q\Phi(1_2) \text{ iff } (U_1, U_2)^{(n'_1, n'_2)} \models Q\Phi(1_2)$$

where  $Q \in \{E, E_{inf}, E_{fin}\}$ ,  $n'_1 = \min(n_1, c_1)$ ,  $n'_2 = \min(n_2, c_2)$ , and:

- $c_1 = 1$  and  $c_2 = 2$ , when  $Q = E_{inf}$ ;
- $c_1 = c_2 = 1$ , when  $Q = E_{fin}$ ;
- $c_1 = 2|U_2|$  and  $c_2 = 2|U_2| + 1$ , when  $Q = E$ .

The detailed proofs of Theorems 4 and 5 are given in the extended version [28]. Thanks to the Truncation Lemma and the duality between operators  $A$  and  $E$ , the Conjunctive Cutoff Theorem can be easily proved. The Cutoff Theorem together with the known decidability and complexity results of the model checking problems for various timed temporal logics [12] justify the following decidability theorem.

**Theorem 6 (Decidability Theorem).** *Let  $(U_1, \dots, U_k)$  be a set of TA templates with conjunctive guards and let  $\phi = \bigwedge_{i_1, \dots, i_h} Q\Phi(i_{l_1}, \dots, i_{l_h})$  where  $Q \in \{A, A_{inf}, A_{fin}, E, E_{inf}, E_{fin}\}$  and  $\{l_1, \dots, l_h\} \in [1, k]$ . The parameterized model checking problem (under the continuous time semantics)*

$$\forall(n_1, \dots, n_k) \succeq (1, \dots, 1) \cdot (U_1, \dots, U_k)^{(n_1, \dots, n_k)} \models \phi$$

is:

- UNDECIDABLE when  $\Phi$  is an *IMTL* formula;
- DECIDABLE and 2-EXPSpace when  $\Phi$  is an *IMITL* formula;
- DECIDABLE and EXPSpace when  $\phi$  is a *TCTL* formula.

*Proof.* For the first two results, consider that the Cutoff Theorem reduces the parameterized model checking problem to an ordinary model checking problem. The latter is undecidable for MTL and is decidable and EXPSpace-Complete (i.e.  $DSPACE(2^{O(n)})$ ), for MITL [12]. Since the model has an exponential number of states (i.e.  $n = 2^{|U| \log(|U|)}$ , where  $U$  is the “biggest” template), the problem is at most 2-EXPSpace. Concerning the third statement, the TCTL model checking problem is PSPACE-Complete [12]. Again, since the model has an exponential number of states, the parameterized model checking problem is at most EXPSpace. A more detailed proof can be found in the extended version [28].

## 6 Case Study

We use the Fischer’s protocol for mutual exclusion to show how to model-check a parameterized and timed systems. The protocol uses a single timed automaton template, instantiated an arbitrary number of time. Figure 1 depicts such template, where  $inv(b_1) = (c \leq k)$  [13]. In Fischer’s protocol every process (a) reads and writes a PID from and into a shared variable, and (b) waits a constant amount of time between when it asks to enter the critical section, and when it actually does so. The Fischer’s protocol cannot be directly modeled in our framework because of the shared variable. We will first abstract the variable into a finite state system with conjunctive guards, and subsequently we will present the results of our verification.

**Abstracting Process Identifier.** A variable can be modeled naively as an automaton with the structure of a completely connected graph, whose vertices denote possible assigned values (let us call  $V$  such model, see Fig. 2). The state space can thus be infinite or finite, but even in the latter case it is usually too big and makes the verification task unfeasible.






Fig. 1. Process in Fischer’s protocol as a Timed Automaton with integer variables



**Fig. 2.**  $V$ : a shared variable





**Fig. 3.**  $W$ : a process-centric view of a shared PID variable

An abstract shared variable for PIDs can be defined, under the assumptions:

- the variable only stores PID values;
- the variable is shared among all processes;
- every PID value overwrites the previous values of the variable itself;
- every process can compare the variable value only with its own PID value.

As in a predicate abstraction, we replace the shared variable with its *process-centric view*. The latter has only two relevant states: it is either the same PID as the process, or it stores a different one. We use  $W$  to denote such process (see Fig. 3). Every process  $P$  is in a one-to-one relation with its own view of the variable. We introduce a process template  $P' = P \times W$  that results from the synchronous product of the  $P$  and  $W$ . We could then model check a system  $P'^{(n)}$ . Doing this, we would probably obtain many spurious counter-examples, since two processes could have their copy of  $W$  in state  $*\_Mypid$ . Since no variable can store multiple values, this is impossible. Conjunctive guards, though, allow to constraint the system in such a way that no two processes can be in a state of the  $*\_Mypid$  group. This solution rules out the undesired spurious behaviors, and is very convenient since it can be applied whenever an algorithm uses a shared variable. We thus define  $P''$  to be the refined version of  $P'$  represented in Fig. 4 using the Uppaal notation. It is possible to show that the abstract system simulates the concrete system, namely  $(P \times V)^{(1,n)} \preceq (P \times W)^{(1,n)}$ , for any positive  $n$ .

Figure 4 depicts template  $P''$ . Some of the eight states resulting from the product are not reached by any transition, and can thus be removed from the model, implying a smaller cutoff. The model manipulation up to this point can be completely automatized. We notice that it is safe to remove state `b2_diff` and connect directly state `b2_Mypid` with `Init_Diff`, obtaining the reduced system in Fig. 5. Finally, let us remark that variable `mypid` in Figs. 4 and 5 is added to overcome Uppaal syntax limitations that cannot refer directly to process states in guards and specifications. The reduced system has 4 states, and thus the cutoff is 9.


**Fig. 4.**  $P'' = (P \times W) + CG$  template

**Fig. 5.** Reduced  $P''$ 

**Verification Results.** Below are the formulae that have been model checked, together with the required time and memory.<sup>2</sup>

Formula	Outcome	Time (s)	Mem. (MB)
(1) $\bigwedge_i E\mathcal{F}_{\geq 0}(CS\_mypid(i))$	true	0.01	155.2
(2) $\bigwedge_{i \neq j} AG_{\geq 0}!(CS\_mypid(i) \wedge CS\_mypid(j))$	true	30.1	155.2
(3) $\bigwedge_i A\mathcal{F}_{\geq 0}(CS\_mypid(i))$	false	0.59	155.2

Formula (1) checks that a process can enter its critical section, while (2) checks the actual mutual exclusion property. Finally (3) states that a process will always be able to enter its critical section. It is well known that while the Fischer's protocol ensures the mutual exclusion property (i.e. formulae (1) and (2)), it also suffers from the problem of processes to possibly starve (i.e. formula (3)).

## 7 Conclusions

In this work we presented the combined study of timed and parameterized systems. We proved that a cutoff exists for PNTA with conjunctive guards and a subset of ITCTL\* formulae. Moreover, the cutoff value is equal to the value computed in Emerson and Kahlon's work for untimed systems [15]. This proves that the parameterized model checking problem is decidable for networks of timed automata with disjunctive guards, for a suitable logic. We remark that for timed systems, applying Theorem 2 one obtains a considerably smaller cutoff than applying the (untimed) Emerson and Kahlon's cutoff theorem after reducing the original timed system to a finite state system by means of the traditional region or zone abstractions.

Finally, we used the Fischer's protocol for mutual exclusion as a benchmark for showing how to apply the cutoff theorem. We claim that the use of conjunctive

<sup>2</sup> The experiments were run on an Intel Core2 Duo CPU T5870 @ 2.0 Ghz with 4GB RAM, OS Linux 3.13-1-amd64.

guards is convenient for verifying systems based of shared variables, since they naturally express the constraint that a variable can store only one value at any time. As a follow-up of this work, we aim at two main goals: (a) finding more algorithms for real-time and distributed systems that can be model checked using our framework, and (b) extending the Emerson and Kahlon cutoff theorem also to PNTA with Disjunctive Guards.

## References

1. Abdulla, P.A., Jonsson, B.: Verifying networks of timed processes (extended abstract). In: Steffen, B. (ed.) TACAS 1998. LNCS, vol. 1384, pp. 298–312. Springer, Heidelberg (1998)
2. Abdulla, P.A., Deneux, J., Mahata, P.: Multi-clock timed networks. In: Proceedings of the 19th IEEE Symposium on Logic in Computer Science, pp. 345–354 (2004)
3. Abdulla, P.A., Jonsson, B.: Model checking of systems with many identical timed processes. *Theoret. Comput. Sci.* **290**(1), 241–264 (2003)
4. Alur, R., Courcoubetis, C., Dill, D.: Model-checking for real-time systems. In: Proceedings of the Fifth Symposium on Logic in Computer Science, pp. 414–425 (1990)
5. Aminof, B., Jacobs, S., Khalimov, A., Rubin, S.: Parameterized model checking of token-passing systems. In: McMillan, K.L., Rival, X. (eds.) VMCAI 2014. LNCS, vol. 8318, pp. 262–281. Springer, Heidelberg (2014)
6. Apt, K., Kozen, D.: Limits for automatic verification of finite-state concurrent systems. *Inf. Process. Lett.* **22**, 307–309 (1986)
7. Aminof, B., Kotek, T., Rubin, S., Spegni, F., Veith, H.: Parameterized model checking of rendezvous systems. In: Baldan, P., Gorla, D. (eds.) CONCUR 2014. LNCS, vol. 8704, pp. 109–124. Springer, Heidelberg (2014)
8. Ball, T., Levin, V., Rajamani, S.: A decade of software model checking with SLAM. *Commun. ACM* **54**(7), 68–76 (2011)
9. Ben-David, S., Eisner, C., Geist, D., Wolfsthal, Y.: Model checking at IBM. *Formal Methods Sys. Des.* **22**(2), 101–108 (2003)
10. Bengtsson, J., Yi, W.: Timed Automata: Semantics, Algorithms and Tools. Technical report 316, UNU-IIST (2004)
11. Bouajjani, A., Habermehl, P., Vojnar, T.: Verification of parametric concurrent systems with prioritised FIFO resource management. *Formal Methods Syst. Des.* **32**, 129–172 (2008)
12. Bouyer, P.: Model-checking timed temporal logics. *Electron. Notes Theor. Comput. Sci.* **231**, 323–341 (2009)
13. Carioni, A., Ghilardi, S., Ranise, S.: MCMT in the land of parameterized timed automata. In: Proceedings of VERIFY@IJCAR 2010, pp. 1–16 (2010)
14. Clarke, E., Grumberg, O., Browne, M.: Reasoning about networks with many identical finite-state processes. In: Proceedings of the 5th Annual ACM Symposium on Principles of Distributed Computing, pp. 240–248 (1986)
15. Emerson, A., Kahlon, V.: Reducing model checking of the many to the few. In: McAllester, D. (ed.) CADE-17. LNCS, vol. 1831, pp. 236–254. Springer, Heidelberg (2000)
16. Emerson, A., Namjoshi, K.: Automatic verification of parameterized synchronous systems. In: Alur, R., Henzinger, T.A. (eds.) CAV 1996. LNCS, vol. 1102, pp. 87–98. Springer, Heidelberg (1996)



17. Emerson, E., Namjoshi, K.: On model checking for non-deterministic infinite-state systems. In: Proceedings of 13th IEEE Symposium on Logic in Computer Science, pp. 70–80 (1998)
18. Emerson, E.A., Namjoshi, K.: On reasoning about rings. *Int. J. Found. Comput. Sci.* **14**(4), 527–550 (2003)
19. German, S.M., Sistla, A.P.: Reasoning about systems with many processes. *J. ACM* **39**(3), 675–735 (1992)
20. Godefroid, P.: Software model checking: The Verisoft approach. *Formal Methods Syst. Des.* **26**(2), 77–101 (2005)
21. Gonthier, T., Glesner, S.: Towards the semi-automatic verification of parameterized real-time systems using network invariants. In: 8th IEEE International Conference on Software Engineering and Formal Methods (SEFM), pp. 310–314 (2010)
22. Hanna, Y., Samuelson, D., Basu, S., Rajan, H.: Automating Cut-off for Multi-parameterized systems. In: Dong, J.S., Zhu, H. (eds.) ICFEM 2010. LNCS, vol. 6447, pp. 338–354. Springer, Heidelberg (2010)
23. Johnson, T.T., Mitra, S.: A small model theorem for rectangular hybrid automata networks. In: Giese, H., Rosu, G. (eds.) FORTE/FMOODS 2012. LNCS, vol. 7273, pp. 18–34. Springer, Heidelberg (2012)
24. Kurshan, R., McMillan, K.: A structural induction theorem for processes. In: ACM Symposium on Principles of Distributed Computing, pp. 239–247 (1989)
25. Mansouri-Samani, M., Mehlitz, P., Pasareanu, C., Penix, J., Brat, G., Markosian, L., O’Malley, O., Pressburger, T., Visser, W.: Program model checking—a practitioners guide. Technical report NASA/TM-2008-214577, NASA (2008)
26. Pagliarecci, F., Spalazzi, L., Spegni, F.: Model checking grid security. *Future Gener. Comput. Syst.* **29**(3), 811–827 (2013)
27. RTCA. Software Considerations in Airborne Systems and Equipment Certification. Technical report DO-178C, RTCA Inc. (2011)
28. Spalazzi, L., Spegni, F.: Parameterized model-checking for timed systems with conjunctive guards (extended version) (2014). arxiv:1407.7305[cs.Lo]
29. Yang, Q., Li, M.: A cut-off approach for bounded verification of parameterized systems. In: Proceedings of the International Conference on Software Engineering, pp. 345–354. ACM (2010)
30. Zuck, L., Pnueli, A.: Model checking and abstraction to the aid of parameterized systems (a survey). *Comp. Lang. Syst. Struct.* **30**(3–4), 139–169 (2004)