



Compositional Invariant Generation for Timed Systems

*L. Aştefănoaei, S. Ben Rayana, S. Bensalem, M. Bozga,
J. Combaz*

Verimag Research Report n° TR-2014-14

November 12, 2014

Reports are downloadable at the following address

<http://www-verimag.imag.fr>

Unité Mixte de Recherche 5104 CNRS - Grenoble INP - UJF

Centre Equation
2, avenue de VIGNATE
F-38610 GIERES
tel : +33 456 52 03 40
fax : +33 456 52 03 50
<http://www-verimag.imag.fr>



Compositional Invariant Generation for Timed Systems

L. Aştefănoaei, S. Ben Rayana, S. Bensalem, M. Bozga, J. Combaz

November 12, 2014

Abstract

We address the state space explosion problem inherent to model-checking timed systems with a large number of components. The main challenge is to obtain pertinent global timing constraints from the timings in the components alone. To this end, we make use of auxiliary clocks to automatically generate new invariants which capture the constraints induced by the synchronisations between components. The method has been implemented as an extension of the D-Finder tool and successfully experimented on several benchmarks.

Keywords: compositional verification, timed automata, invariants, component invariants, interaction invariants, interactions

Reviewers:

Notes: Research supported by the European Integrated Project 257414 ASCENS, ICT Collaborative Project 288175 CERTAINTY and STREP 318772 D-MILS.

How to cite this report:

```
@techreport {TR-2014-14,  
  title = {Compositional Invariant Generation for Timed Systems},  
  author = {L. Aştefănoaei, S. Ben Rayana, S. Bensalem, M. Bozga, J. Combaz},  
  institution = {{Verimag} Research Report},  
  number = {TR-2014-14},  
  year = {}  
}
```

1 Introduction

Compositional methods in verification have been developed to cope with state space explosion. Generally based on divide et impera principles, these methods attempt to break monolithic verification problems into smaller sub-problems by exploiting either the structure of the system or the property or both. Compositional reasoning can be used in different manners e.g., for deductive verification, assume-guarantee, contract-based verification, compositional generation, etc.

The development of compositional verification for timed systems remains however challenging. State-of-the-art tools [7, 14, 29, 22] for the verification of such systems are mostly based on symbolic state space exploration, using efficient data structures and particularly involved exploration techniques. In the timed context, the use of compositional reasoning is inherently difficult due to the synchronous model of time. Time progress is an action that synchronises continuously all the components of the system. Getting rid of the time synchronisation is necessary for analysing independently different parts of the system (or of the property) but becomes problematic when attempting to re-compose the partial verification results. Nonetheless, compositional verification is actively investigated and several approaches have been recently developed and employed in timed interfaces [2] and contract-based assume-guarantee reasoning [16, 25].

In this paper, we propose a different approach for exploiting compositionality for analysis of timed systems using invariants. In contrast to exact reachability analysis, invariants are symbolic approximations of the set of reachable states of the system. We show that rather precise invariants can be computed compositionally, from the separate analysis of the components in the system and from their composition glue. This method is proved to be sound for the verification of safety state properties. However, it is not complete.

The starting point is the verification method of [10], summarised in Figure 1. The method exploits compositionality as explained next. Consider a system consisting of components B_i interacting by means of a set γ of multi-party interactions, and let φ be a system property of interest. Assume that all B_i as well as the composition through γ can be independently characterised by means of component invariants $CI(B_i)$, respectively interaction invariants $II(\gamma)$. The connection between the invariants and the system property φ can be intuitively understood as follows: if φ can be proved to be a logical consequence of the conjunction of components and interaction invariants, then φ holds for the system.

$$\frac{\vdash \bigwedge_i CI(B_i) \wedge II(\gamma) \rightarrow \varphi}{\|\gamma B_i \models \square \varphi} \quad (VR)$$

Figure 1: Compositional Verification

In the rule (VR) the symbol “ \vdash ” is used to underline that the logical implication can be effectively proved (for instance with an SMT solver) and the notation “ $\|\gamma B_i \models \square \varphi$ ” is to be read as “ φ holds in every reachable state of $\|\gamma B_i$ ”.

The verification rule (VR) in [10] has been developed for untimed systems. Its direct application to timed systems may be weak as interaction invariants do not capture global timings of interactions between components. The key contribution of this paper is to improve the invariant generation method so to better track such global timings by means of auxiliary *history clocks* for actions and interactions. At component level, history clocks expose the local timing constraints relevant to the interactions of the participating components. At composition level, extra constraints on history clocks are enforced due to the simultaneity of interactions and to the synchrony of time progress.

As an illustration, let us consider as running example the timed system in Figure 2 which depicts a “controller” component serving n “worker” components, one at a time. The interactions between the controller and the workers are defined by the set of synchronisations $\{(a \mid b_i), (c \mid d_i) \mid i \leq n\}$. Periodically, after every 4 units of time, the controller synchronises its action a with the action b_i of any worker i whose clock shows at least $4n$ units of time. Initially, such a worker exists because the controller waits for $4n$ units of time before interacting with workers. The cycle repeats forever because there is always a worker “willing” to do b , that is, the system is deadlock-free. Proving deadlock-freeness of the system requires to establish that when the controller is at location lc_1 there is at least one worker such that $y_i - x \geq 4n - 4$.

Unfortunately, this property cannot be shown if we use (VR) as it is in [10]. Intuitively, this is because the proposed invariants are too weak to infer cross constraints relating the clocks of the controller and those of the workers: interaction invariants $II(\gamma)$ relates only locations of components and thus at most eliminates unreachable configurations like $(lc_1, \dots, l_{2i}, \dots)$, while the component invariants can only state local conditions on clocks such as that $x \leq 4$ at lc_1 . Using history clocks allows to recover additional constraints. For example, after the first execution of the loop, each time when the controller is at location lc_1 , there exists a worker i whose clock has an equal value as that of the controller. Similarly, history clocks allow to infer that different $(a \mid b_i)$ interactions are separated by at least 4 time units. These constraints altogether are sufficient to prove the deadlock freedom property.




Figure 2: A Timed System

Related Work

Automatic generation of invariants for concurrent systems is a long-time studied topic. Yet, to our knowledge, specific extensions or applications for timed systems are rather limited. As an exception, the papers [5, 21] propose a monolithic, non-compositional method for finding invariants in the case of systems represented as a single timed automaton.

Compositional verification for timed systems has been mainly considered in the context of timed interface theories [2] and contract-based assume guarantee reasoning [16, 25]. These methods usually rely upon choosing a “good” decomposition structure and require individual abstractions for components to be deterministic timed I/O automata. Finding the abstractions is in general difficult, however, their construction can be automated by using learning techniques [25] in some cases. In contrast to the above, we are proposing a fully automated method generating, in a compositional manner, an invariant approximating the reachable states of a timed system.

Abstractions serve also for compositional minimisation, for instance [11] minimises by constructing timed automata quotients with respect to simulation; these quotients are in turn composed for model-checking. Our approach is orthogonal in that we do not compose at all. Compositional deductive verification as in [17] is also orthogonal on our work in that, by choosing a particular class of local invariants to work with, we need not focus on elaborate proof systems but reason at a level closer to intuition.

The use of additional clocks has been considered, for instance, in [8]. There, extra reference clocks are added to components to faithfully implement a partial order reduction strategy for symbolic state space exploration. Time is allowed to progress desynchronised for individual components and re-synchronised only when needed, i.e., for direct interaction within components. Clearly, the history clocks in our work behave in a similar way, however, our use of clocks is as a helper construction in the generation of invariants and we are totally avoiding state space exploration. Finally, another successful application of extra clocks has been provided in [27] for timing analysis of asynchronous circuits. There, specific history clocks are reset on input signals and used to provide a new time basis for the construction of an abstract model of output signals of the circuit.

Finally, this paper is essentially an extended version of the conference paper [4]. The extension is with respect to incorporating proofs, detailing some technicalities about handling initial states, and formalising some heuristics implicitly used in the conference paper.

Organisation of the paper

Section 2 recalls the needed definitions for modelling timed systems and their properties. Section 3 presents our method for compositional generation of invariants. Section 4 describes three heuristics to speed up and simplify invariant generation. Section 5 describes the prototype implementing the method and presents concrete results obtained on case studies. Section 6 concludes.

2 Timed Systems and Properties

In the framework of the present paper, components are timed automata and systems are compositions of timed automata with respect to multi-party interactions. The timed automata we use are essentially the ones from [3], however, slightly adapted to embrace a uniform notation throughout the paper.

Definition 1 (Syntax). *A component is a timed automaton $(L, A, \mathcal{X}, T, \text{tpc}, s_0)$ where L is a finite set of locations, A a finite set of actions, \mathcal{X} is a finite set of local¹ clocks, $T \subseteq L \times (A \times \mathcal{C} \times 2^{\mathcal{X}}) \times L$ is a set of edges labelled with an action, a guard, and a set of clocks to be reset, $\text{tpc} : L \rightarrow \mathcal{C}$ assigns a time progress condition² to each location. \mathcal{C} is the set of clock constraints and $s_0 \in L \times \mathcal{C}$ provides the initial configuration. A clock constraint is defined by the grammar:*

$$C ::= \text{true} \mid x \# ct \mid x - y \# ct \mid C \wedge C$$

with $x, y \in \mathcal{X}$, $\# \in \{<, \leq, =, \geq, >\}$ and $ct \in \mathbb{Z}$. Time progress conditions are restricted to conjunctions of constraints as $x \leq ct$.

Before recalling the semantics of a component, we first fix some notation. Let \mathbf{V} be the set of all clock valuation functions $\mathbf{v} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$. For a clock constraint C , $C(\mathbf{v})$ denotes the evaluation of C in \mathbf{v} . The notation $\mathbf{v} + \delta$ represents a new \mathbf{v}' defined as $\mathbf{v}'(x) = \mathbf{v}(x) + \delta$ while $\mathbf{v}[r]$ represents a new \mathbf{v}' which assigns any x in r to 0 and otherwise preserves the values from \mathbf{v} .

Definition 2 (Semantics). *The semantics of a component $B = (L, A, \mathcal{X}, T, \text{tpc}, s_0)$ is given by the labelled transition system (Q, A, \rightarrow, Q_0) where $Q \subseteq L \times \mathbf{V}$ denotes the states of B , $\rightarrow \subseteq Q \times (A \cup \mathbb{R}_{\geq 0}) \times Q$ denotes the transitions according to the rules:*

- $(l, \mathbf{v}) \xrightarrow{\delta} (l, \mathbf{v} + \delta)$ if $(\forall \delta' \in [0, \delta]). (\text{tpc}(l)(\mathbf{v} + \delta'))$ (time progress);
- $(l, \mathbf{v}) \xrightarrow{a} (l', \mathbf{v}[r])$ if $(l, (a, g, r), l') \in T, g(\mathbf{v}) \wedge \text{tpc}(l')(\mathbf{v}[r])$ (action step).

and $Q_0 = \{(l_0, \mathbf{v}_0) \mid s_0 = (l_0, c_0) \wedge c_0(\mathbf{v}_0)\}$ denotes the initial states.

Because the semantics defined above is in general infinite, we work with the so called zone graph [23] as a finite symbolic representation. The symbolic states in a zone graph are pairs (l, ζ) where l is a location of B and ζ is a zone, a set of clock valuations defined by clock constraints. The initial configuration $s_0 = (l_0, c_0)$ corresponds trivially to a symbolic state (l_0, ζ_0) . Given a symbolic state (l, ζ) , its successor with respect to a transition t of B is denoted as $\text{succ}(t, (l, \zeta))$ and defined by means of its timed and its discrete successor:

- $\text{time_succ}((l, \zeta)) = (l, \nearrow \zeta \cap \text{tpc}(l))$
- $\text{disc_succ}(t, (l, \zeta)) = (l', (\zeta \cap g)[r] \cap \text{tpc}(l'))$ if $t = (l, (_, g, r), l')$

¹Locality is essential for avoiding side effects which would break compositionality and local analysis.

²To avoid confusion with invariant properties, we prefer to adopt the terminology of “time progress condition” from [12] instead of “location invariants”.

- $\text{succ}(t, (l, \zeta)) = \text{norm}(\text{time_succ}(\text{disc_succ}(t, (l, \zeta))))$

where $\nearrow, [r], \text{norm}$ are usual operations on zones: $\nearrow \zeta$ is the forward diagonal projection of ζ , i.e., it contains any valuation \mathbf{v}' for which there exists a real δ such that $\mathbf{v}' - \delta$ is in ζ ; $\zeta[r]$ is the set of all valuations in ζ after applying the resets in r ; $\text{norm}(\zeta)$ corresponds to normalising ζ such that all bounds on clocks and clock differences are either bounded by some finite value or infinite. Since we are seeking component invariants which are over-approximations of the reachable states, a more thorough discussion on normalisation is not relevant for the present paper. The interested reader may refer to [9, 13] for more precise definitions.

A symbolic execution of B is a sequence of symbolic states s_0, \dots, s_i, \dots ³ such that for any $i > 0$, there exists a transition t for which s_i is $\text{succ}(t, s_{i-1})$. The set of reachable symbolic states of B is $\text{Reach}_B(s_0)$ where Reach_B is defined recursively as:

$$\text{Reach}_B(s) = \{s\} \cup \bigcup_{t \in T} \text{Reach}_B(\text{succ}(t, s)).$$

for an arbitrary s and T the set of transitions in B . We remind that the set $\text{Reach}_B(s_0)$ can be shown finite knowing that the number of normalised zones is finite. In general, the symbolic zone graph provides an over-approximation of the set of reachable states. This over-approximation is exact only for timed automata without diagonal constraints [9, 13].

In our framework, components communicate by means of *interactions*, which are synchronisations between actions. Given n components $(B_i)_{i=1,n}$, with disjoint sets of actions A_i , an interaction is a subset $\alpha \subseteq \cup_i A_i$ containing at most one action per component. We denote interactions α as sets $\{a_i\}_{i \in I}$, with $a_i \in A_i$ for all $i \in I \subseteq \{1, \dots, n\}$. For readability, in examples, we use the alternative notation $(a_1 \mid a_2 \mid \dots \mid a_i)$ instead. Given a set of interactions γ , we denote by $\text{Act}(\gamma)$ the set of actions involved in γ , that is, $\text{Act}(\gamma) = \cup_{\alpha \in \gamma} \alpha$.

Definition 3 (Timed System). *Let n components $(B_i)_{i=1,n}$ with disjoint sets of actions, where $B_i = (L_i, A_i, \mathcal{X}_i, T_i, \text{tpc}_i, s_{0i})$, $s_{0i} = (l_{0i}, c_{0i})$. Let γ be a set of interactions constructed from $\cup_i A_i$. The timed system $\parallel_\gamma B_i$ is defined as the component $(L, \gamma, \mathcal{X}, T_\gamma, \text{tpc}, s_0)$ where $L = \times_i L_i$, $\mathcal{X} = \cup_i \mathcal{X}_i$, $\text{tpc}(\bar{l}) = \bigwedge_i \text{tpc}(l_i)$, $s_0 = ((l_{01}, \dots, l_{0n}), \bigwedge_i c_{0i})$ and*

$$T_\gamma = \left\{ (\bar{l}, (\alpha, g, r), \bar{l}') \mid \begin{array}{l} \bar{l} = (l_1, \dots, l_n) \in L, \bar{l}' = (l'_1, \dots, l'_n) \in L \\ \alpha = \{a_i\}_{i \in I} \in \gamma, \forall i \in I. (l_i, (a_i, g_i, r_i), l'_i) \in T_i, \forall i \notin I. l_i = l'_i \\ g = \bigwedge_{i \in I} g_i, r = \bigcup_{i \in I} r_i \end{array} \right\}$$

In the timed system $\parallel_\gamma B_i$, a component B_i can execute an action a_i only as part of an interaction α , $a_i \in \alpha$, that is, along with the execution of all other actions $a_j \in \alpha$. This corresponds to the usual notion of multi-party interaction. We note that interactions can only restrict the behaviour of components, i.e., the states reached by B_i in $\parallel_\gamma B_i$ belong to $\text{Reach}_{B_i}(s_{0i})$. This is a property which is exploited in the verification rule (VR).

To give a logical characterisation of components and their properties, we use invariants. An invariant Φ is a state predicate which holds in every reachable state of B , in symbols, $B \models \Box \Phi$. We use $CI(B)$ and $II(\gamma)$, to denote **component**, respectively **interaction invariants**. For component invariants, our choice is to work with their reachable symbolic set. More precisely, for component B , its associated component invariant $CI(B)$ is the disjunction of $(l \wedge \zeta)$ for all symbolic states (l, ζ) in $\text{Reach}_B(s_0)$. To ease the reading, we abuse of notation and use l as a place holder for a state predicate “ $at(l)$ ” which holds in any symbolic state with location l , that is, the semantics of $at(l)$ is given by $(l, \zeta) \models at(l)$. As an example, the component invariants for the example in Figure 2 with one worker are:

$$\begin{aligned} CI(\text{Controller}) &= (lc_0 \wedge x \geq 0) \vee (lc_1 \wedge x \leq 4) \vee (lc_2 \wedge x \geq 0) \\ CI(\text{Worker}_1) &= (l_{11} \wedge y_1 \geq 0) \vee (l_{21} \wedge y_1 \geq 4). \end{aligned}$$

³We tacitly assume that s_0 is such that $s_0 = \text{time_succ}(s_0)$. If this is not the case, one can always consider $\text{time_succ}(s_0)$ instead of s_0 for the definition of symbolic executions and reachable states.

The interaction invariants are computed by the method explained in [10]. Interaction invariants are over-approximations of the global state space allowing us to disregard certain tuples of local states as unreachable. As an illustration, consider the interaction invariant for the running example when the controller is interacting with one worker:

$$II(\{(a \mid b_1), (c \mid d_1)\}) = (l_{11} \vee lc_2) \wedge (l_{21} \vee lc_0 \vee lc_1).$$

Let us remark that component and interaction invariants proposed⁴ above are inductive invariants. A state predicate is called *inductive* for a component or system B if, whenever it holds for a state s of B it equally holds for any of its successors s' . That is, the validity of an inductive predicate is preserved by executing any transition, timed or discrete. An inductive predicate which moreover holds at initial states is an (inductive) invariant. Trivially, such a predicate holds in all reachable states.

As for **component properties**, we are interested in arbitrary invariant state properties that can be expressed as boolean combinations of “*at*(l)” predicates and clock constraints. Invariant properties include generic properties such as mutual exclusion, absence of deadlock, unreachability of “bad” states, etc. As a simple illustration consider the property $lc_1 \rightarrow \bigvee_i (y_i - x \geq 4n - 4)$, discussed for our running example introduced in Section 1. As a more sophisticated example, consider *absence of deadlock*. Intuitively, a timed system with a set of interactions γ is *deadlocked* when no interaction in γ is enabled. Absence of deadlock is therefore expressed as the disjunction $\bigvee_{\alpha \in \gamma} \text{enabled}(\alpha)$. As for the enabledness predicate, we borrow it from [28] where it is essentially constructed from the syntactic definition of the timed system. More precisely, for an interaction α , $\text{enabled}(\alpha)$ is $\bigvee_t \text{enabled}(t)$, with t being a global transition triggered by α . In turn, for $t = (\bar{l}, (\alpha, g, r), \bar{l}')$, $\text{enabled}(t)$ is defined using elementary operations on zones as $\bar{l} \wedge \swarrow (g \cap [r] \text{tpc}(\bar{l}') \cap \text{tpc}(\bar{l}))$, where $\swarrow \zeta$ is the backward diagonal projection of ζ , $[r]\zeta$ is the set of valuations \mathbf{v} such that $\mathbf{v}[r]$ is in ζ .

3 Timed Invariant Generation

As explained in the introduction, a direct application of the compositional verification rule (VR) may not be useful in itself in the sense that the component and the interaction invariants alone are usually not enough to prove global properties, especially when such properties involve relations between clocks in different components. More precisely, though component invariants encode timings of local clocks, there is no direct way – the interaction invariant is orthogonal on timing aspects – to constrain the bounds on the differences between clocks in different components. To give a concrete illustration, consider the property $\varphi_{\text{Safe}} = (lc_1 \wedge l_{11} \rightarrow x \leq y_1)$ that holds in the running example with one worker. We note that if this property is satisfied, it is guaranteed that the global system is not deadlocked when the controller is at location lc_1 and the worker is at location l_{11} . It is not difficult to see that φ_{Safe} cannot be deduced from $CI(\text{Controller}) \wedge CI(\text{Worker}_1) \wedge II(\{(a \mid b_1), (c \mid d_1)\})$ as no relation can be established between x and y_1 .

3.1 History Clocks for Actions

In this section, we show how we can, by means of some auxiliary constructions, apply (VR) more successfully. To this end, we “equip” components (and later, interactions) with *history clocks*, a clock per action; then, at interaction time, the clocks corresponding to the actions participating in the interaction are reset. This basic transformation allows us to automatically compute a new invariant of the system with history clocks. This new invariant, together with the component and interaction invariants, is shown to be, after projection of history clocks, an invariant of the initial system.

Definition 4 (Components with History Clocks). *Given component $B = (L, A, \mathcal{X}, T, \text{tpc}, s_0)$, its extension with history clocks is the component $B^h = (L, A, \mathcal{X} \cup \mathcal{H}_A, T^h, \text{tpc}, s_0^h)$ where*

- $\mathcal{H}_A = \{h_0\} \cup \{h_a \mid a \in A\}$ is the set of history clocks,

⁴The rule (VR) is generic enough to work with other types of invariants. For example, one could use any over-approximation of the reachable set in the case of component invariants, however, this comes at the price of losing precision.

- $T^h = \{(l, (a, g, r \cup \{h_a\}), l') \mid (l, (a, g, r), l') \in T\}$,
- $s_0^h = (l_0, c_0^h)$, where $c_0^h = (c_0 \wedge h_0 = 0 \wedge \bigwedge_{a \in A} h_a > 0)$, given $s_0 = (l_0, c_0)$.

The clock h_0 measures the time from the initialisation. This clock equals 0 in s_0^h and is never tested or reset. Due to this very restricted use, the same clock h_0 can be consistently used (shared) by all components B^h and consequently, allows to capture clock constraints derived from the common system initialisation time.

Every history clock h_a measures the time passed from the last occurrence of action a . These history clocks are initially strictly greater than 0 and are reset when the corresponding action is executed. As a side effect, whenever h_a is strictly bigger than h_0 , we can infer that the action a has not been (yet) executed.

Since there is no timing constraint involving history clocks, these have no influence on the behaviour. The extended model is, in fact, bisimilar to the original model. Moreover, any invariant of the extended model of B^h corresponds to an invariant of original component. By abuse of notation, given set of actions $A = \{a_1, \dots, a_m\}$ use $\exists \mathcal{H}_A$ to stand for $\exists h_{a_1} \exists h_{a_2} \dots \exists h_{a_m} \exists h_0$.

Proposition 1.

1. If Φ^h is an invariant of B^h then $\Phi = \exists \mathcal{H}_A. \Phi^h$ is an invariant of B .
2. If Φ^h is an invariant of B^h and Ψ^h an inductive assertion of B^h expressed on history clocks $\mathcal{H}_A \setminus \{h_0\}$ then $\Phi = \exists \mathcal{H}_A. (\Phi^h \wedge \Psi^h)$ is an invariant of B .

Proof. (1) It suffices to notice that any symbolic state (l, ζ^h) in the reachable set $Reach_{B^h}(s_0^h)$ corresponds to a symbolic state (l, ζ) in the reachable set $Reach_B(s_0)$ such that moreover ζ is the projection of ζ^h to clocks in \mathcal{X} , that is $\zeta \equiv \exists \mathcal{H}_A. \zeta^h$. Henceforth, $\exists \mathcal{H}_A. Reach_{B^h}(s_0^h) \equiv Reach_B(s_0)$. Moreover, for any invariant Φ^h of B^h it holds $\exists \mathcal{H}_A. Reach_{B^h}(s_0^h) \subseteq \exists \mathcal{H}_A. \Phi^h$. By combining the two facts, we obtain that Φ is an invariant of B .

(2) Consider the modified component with history clocks B_Ψ^h defined as B^h but with initial configuration $(l_0, c_0^h \wedge \Psi^h)$. This initial configuration is valid, as Ψ^h constrain exclusively clocks in \mathcal{H}_A whereas c_0^h leaves all of them unconstrained. Now, it can be easily shown that $\Phi^h \wedge \Psi^h$ is an invariant of B_Ψ^h . Then, following the same reasoning as for point (1) we obtain that $\exists \mathcal{H}_A. (\Phi^h \wedge \Psi^h)$ is an invariant of B . \square

The only operation acting on history clocks is reset. Its effect is that immediately after an interaction takes place, all history clocks involved in the interaction are equal to zero. All the remaining ones preserve their previous values, thus they are at least greater in value than those being reset. This basic observation is exploited in the following definition, which builds, recursively, all the inequalities that could hold given an interaction set γ .

Definition 5 (Interaction Inequalities for History Clocks). *Given an interaction set γ , we define the following interaction inequalities $\mathcal{E}(\gamma)$:*

$$\mathcal{E}(\gamma) = \bigvee_{\alpha \in \gamma} \left(\left(\bigwedge_{\substack{a_i, a_j \in \alpha \\ a_k \in Act(\gamma \ominus \alpha)}} h_{a_i} = h_{a_j} \leq h_{a_k} \right) \wedge \mathcal{E}(\gamma \ominus \alpha) \right).$$

where $\gamma \ominus \alpha = \{\beta \setminus \alpha \mid \beta \in \gamma \wedge \beta \not\subseteq \alpha\}$.

We can use the interpreted function “min” as syntactic sugar to have a slightly more compact expression for $\mathcal{E}(\gamma)$ as follows:

$$\mathcal{E}(\gamma) = \bigvee_{\alpha \in \gamma} \left(\bigwedge_{a_i, a_j \in \alpha} h_{a_i} = h_{a_j} \leq \min_{a_k \in Act(\gamma \ominus \alpha)} h_{a_k} \wedge \mathcal{E}(\gamma \ominus \alpha) \right).$$

As an example, for $\gamma = \{(a \mid b_1), (c \mid d_1)\}$ corresponding to the interactions between the controller and one worker in Figure 2, the compact form is:

$$(h_a = h_{b_1} \leq \min(h_c, h_{d_1}) \wedge h_c = h_{d_1}) \vee (h_c = h_{d_1} \leq \min(h_a, h_{b_1}) \wedge h_a = h_{b_1}).$$

$\mathcal{E}(\gamma)$ characterises the relations between history clocks during any possible execution. It can be shown that this characterisation is, in fact, an inductive predicate of the extended system with history clocks.

Proposition 2. $\mathcal{E}(\gamma)$ is an inductive predicate of $\|\gamma B_i^h$.

Proof. Assume $\mathcal{E}(\gamma)$ holds in some arbitrary state s of $\|\gamma B_i^h$. We have two categories of successor states for s , namely time successors and discrete successors. Obviously $\mathcal{E}(\gamma)$ holds for all time successors s' , as all clocks progress uniformly and henceforth all the relations between them are preserved. Let now s' be a discrete successor of s by an arbitrary interaction α . As all the history clocks for actions in α have just been reset, s' satisfies

$$\bigwedge_{\substack{a_i, a_j \in \alpha \\ a_k \in Act(\gamma \ominus \alpha)}} 0 = h_{a_i} = h_{a_j} \leq h_{a_k}$$

To conclude the proof, we need to show that moreover, for the remaining clocks of actions in $Act(\gamma \ominus \alpha)$, they satisfy $\mathcal{E}(\gamma \ominus \alpha)$. Actually, we can show the more general fact that for any set of interactions γ and for any interaction α the implication $\mathcal{E}(\gamma) \rightarrow \mathcal{E}(\gamma \ominus \alpha)$ is valid. This fact can be simply proven by induction on the size of the set interactions γ following the definition of \mathcal{E} . \square

By using Proposition 2 and Proposition 1, we can safely combine the component and interaction invariants of the system with history clocks with the interaction inequalities. We can eliminate the history clocks from $\bigwedge_i CI(B_i^h) \wedge II(\gamma) \wedge \mathcal{E}(\gamma)$ and obtain an invariant of the original system. This invariant is usually stronger than $\bigwedge_i CI(B_i) \wedge II(\gamma)$ and yields more successful applications of the rule (VR).

Corrolary 1. $\Phi = \exists \mathcal{H}_A. (\bigwedge_i CI(B_i^h) \wedge II(\gamma) \wedge \mathcal{E}(\gamma))$ is an invariant of $\|\gamma B_i$.

Example 1. We reconsider the model of a controller and a worker from Figure 2. We illustrate how the safety property $\varphi_{Safe} = (lc_1 \wedge l_{11} \rightarrow x \leq y_1)$ introduced in the beginning of the section can be shown to hold by using the newly generated invariant. The invariants for the components with history clocks are:

$$\begin{aligned} CI(\text{Controller}^h) = & (lc_0 \wedge x = h_0 < h_a \wedge h_0 < h_c) \vee \\ & (lc_1 \wedge x \leq h_0 - 4 \wedge x \leq 4 \wedge h_0 < h_a \wedge h_0 < h_c) \vee \\ & (lc_1 \wedge x \leq 4 \wedge x = h_c \leq h_a \leq h_0 - 8) \vee \\ & (lc_2 \wedge x \leq h_0 - 8 \wedge h_a = x \wedge h_0 < h_c) \vee \\ & (lc_2 \wedge x = h_a \wedge h_c = h_a + 4 \leq h_0 - 8) \\ CI(\text{Worker}_1^h) = & (l_{11} \wedge y_1 = h_0 < h_{d_1} \wedge h_0 < h_{b_1}) \vee \\ & (l_{11} \wedge y_1 = h_{d_1} \leq h_{b_1} \leq h_0 - 4) \vee \\ & (l_{21} \wedge h_{b_1} + 4 \leq y_1 = h_0 < h_{d_1}) \vee \\ & (l_{21} \wedge y_1 = h_{d_1} \leq h_0 - 4 \wedge h_{b_1} \leq h_{d_1} - 4) \end{aligned}$$

By using the interaction invariant described in Section 2 and the inequality constraints $\mathcal{E}((a \mid b_1), (c \mid d_1))$, after the elimination of the existential quantifiers in

$$(\exists h_a. \exists h_{b_1}. \exists h_c. \exists h_{d_1}. \exists h_0) CI(\text{Controller}^h) \wedge CI(\text{Worker}_1^h) \wedge II(\gamma) \wedge \mathcal{E}(\gamma)$$

we obtain the following invariant Φ :

$$\begin{aligned} \Phi = & (l_{11} \wedge lc_0 \wedge \mathbf{x = y_1}) \vee \\ & (l_{11} \wedge lc_1 \wedge (\mathbf{y_1 = x} \vee \mathbf{x + 4 \leq y_1})) \vee \\ & (l_{21} \wedge lc_2 \wedge (\mathbf{y_1 = x + 4} \vee \mathbf{x + 8 \leq y_1})). \end{aligned}$$

We used bold fonts in Φ to highlight relations between x and y_1 which are not in $CI(\text{Controller}) \wedge CI(\text{Worker}_1) \wedge II(\gamma)$. It can be easily checked now that that $\Phi \rightarrow \varphi_{Safe}$ holds and consequently, this proves that φ_{Safe} holds for the system.

To sum up, the basic steps of our invariant generation method described so far are:

1. compute the interaction invariant $II(\gamma)$

2. extend the components B_i to components with history clocks B_i^h
3. compute component invariants $CI(B_i^h)$
4. compute inequality constraints $\mathcal{E}(\gamma)$ for interactions γ
5. finally, eliminate the history clocks in $\bigwedge_i CI(B_i^h) \wedge II(\gamma) \wedge \mathcal{E}(\gamma)$.

We conclude the section with a remark on the size of $\mathcal{E}(\gamma)$. Due to the combination of recursion and disjunction, $\mathcal{E}(\gamma)$ can be large. Much more compact formulae can be obtained by exploiting non-conflicting interactions, i.e., interactions that do not share actions.

Proposition 3. *If $\gamma = \gamma_1 \cup \gamma_2$ such that $Act(\gamma_1) \cap Act(\gamma_2) = \emptyset$ then $\mathcal{E}(\gamma) \equiv \mathcal{E}(\gamma_1) \wedge \mathcal{E}(\gamma_2)$.*

Proof. By induction on the number of interactions in γ . In the base case, γ has a single interaction and the property trivially holds. For the induction step, for the ease of reading, we introduce $eq(\alpha)$ and $leq(\alpha, \gamma)$ to denote respectively $\bigwedge_{a_i, a_j \in \alpha} h_{a_i} = h_{a_j}$ and $\bigwedge_{\substack{a_i \in \alpha \\ a_k \in Act(\gamma \ominus \alpha)}} h_{a_i} \leq h_{a_k}$. $\mathcal{E}(\gamma)$ can be rewritten as follows:

$$\begin{aligned}
 \mathcal{E}(\gamma) &= \bigvee_{\alpha \in \gamma_1} eq(\alpha) \wedge leq(\alpha, \gamma) \wedge \mathcal{E}((\gamma_1 \cup \gamma_2) \ominus \alpha) \vee \bigvee_{\alpha \in \gamma_2} eq(\alpha) \wedge leq(\alpha, \gamma) \wedge \mathcal{E}((\gamma_1 \cup \gamma_2) \ominus \alpha) \\
 &\quad (\text{using } \gamma_2 \ominus \alpha = \gamma_2 \text{ for } \alpha \in \gamma_1 \text{ and by ind. for } \gamma' = (\gamma_1 \ominus \alpha) \cup \gamma_2) \\
 &\equiv \bigvee_{\alpha \in \gamma_1} eq(\alpha) \wedge leq(\alpha, \gamma) \wedge \mathcal{E}(\gamma_1 \ominus \alpha) \wedge \mathcal{E}(\gamma_2) \vee \bigvee_{\alpha \in \gamma_2} eq(\alpha) \wedge leq(\alpha, \gamma) \wedge \mathcal{E}(\gamma_1) \wedge \mathcal{E}(\gamma_2 \ominus \alpha) \\
 &\quad (\text{using } \bigvee_{\alpha \in \gamma_1} eq(\alpha) \wedge leq(\alpha, \gamma_1) \wedge \mathcal{E}(\gamma_1 \ominus \alpha) = \mathcal{E}(\gamma_1)) \\
 &\equiv \mathcal{E}(\gamma_1) \wedge \mathcal{E}(\gamma_2) \wedge \left(\bigvee_{\alpha \in \gamma_1} leq(\alpha, \gamma_2) \vee \bigvee_{\alpha \in \gamma_2} leq(\alpha, \gamma_1) \right) \\
 &\quad (\text{using totality of "}\leq\text{" and disjointness of } \gamma_i) \\
 &\equiv \mathcal{E}(\gamma_1) \wedge \mathcal{E}(\gamma_2)
 \end{aligned}$$

□The following corollary is an immediate consequence of Proposition 3.

Corollary 2. *If the interaction model γ has only disjoint interactions, i.e., for any $\alpha_1, \alpha_2 \in \gamma$, $\alpha_1 \cap \alpha_2 = \emptyset$, then $\mathcal{E}(\gamma) \equiv \bigwedge_{\alpha \in \gamma} \left(\bigwedge_{a_i, a_j \in \alpha} h_{a_i} = h_{a_j} \right)$.*

The two interactions in $\gamma = \{(a \mid b_1), (c \mid d_1)\}$ are disjoint. Thus, we can simplify the expression of $\mathcal{E}(\gamma)$ to $(h_a = h_{b_1}) \wedge (h_c = h_{d_1})$.

3.2 History Clocks for Interactions

The equality constraints on history clocks allow to relate the local constraints obtained individually on components. In the case of non-conflicting interactions, the relation is rather “tight”, that is, expressed as conjunction of equalities on history clocks. In contrast, the presence of conflicts lead to a significantly weaker form. Intuitively, every action in conflict can be potentially used in different interactions. The uncertainty on its exact use leads to a disjunctive expression as well as to more restricted equalities and inequalities amongst history clocks.

Nonetheless, the presence of conflicts themselves can be additionally exploited for the generation of new invariants. That is, in contrast to equality constraints obtained from interaction, the presence of conflicting actions enforce disequalities (or separation) constraints between all interactions using them. In what follows, we show a generic way to automatically compute such invariants enforcing differences between the timings of the interactions themselves. To effectively implement this, we proceed in a similar manner as in the previous section: we again make use of history clocks and corresponding resets but this time we associate them to interactions, at the system level.

Definition 6 (System with Interaction History Clocks). *Given a timed system $\|_{\gamma} B_i$, its extension with history clocks for interactions is the timed system $B^* \|_{\gamma^h} B_i^h$ where:*

- B^* is an auxiliary component $(\{l^*\}, A_{\gamma}, \mathcal{H}_{\gamma}, T, (l^* \mapsto true), (l^*, true))$ where:
 - the set of actions $A_{\gamma} = \{a_{\alpha} \mid \alpha \in \gamma\}$
 - the set of interaction history clocks $\mathcal{H}_{\gamma} = \{h_{\alpha} \mid \alpha \in \gamma\}$
 - the set of transitions $T = \{(l^*, (a_{\alpha}, true, \{h_{\alpha}\}), l^*) \mid \alpha \in \gamma\}$
- $\gamma^h = \{(a_{\alpha} \mid \alpha) \mid \alpha \in \gamma\}$ with $(a_{\alpha} \mid \alpha)$ denoting $\{a_{\alpha}\} \cup \{a \mid a \in \alpha\}$.

As before, it can be shown that any invariant of $B^* \|_{\gamma^h} B_i^h$ corresponds to an invariant of $\|_{\gamma} B_i$. The history clocks for interactions do not impact the behaviour and henceforth the two systems are bisimilar.

Proposition 4.

1. If Φ^h is an invariant of $B^* \|_{\gamma^h} B_i^h$, then $\Phi = \exists \mathcal{H}_A \exists \mathcal{H}_{\gamma}. \Phi^h$ is an invariant of $\|_{\gamma} B_i$.
2. If Φ^h is an invariant of $B^* \|_{\gamma^h} B_i^h$ and Ψ^h an inductive predicate of $B^* \|_{\gamma^h} B_i^h$ expressed on history clocks for actions and interactions $\mathcal{H}_{\gamma} \cup \mathcal{H}_A \setminus \{h_0\}$ then $\Phi = \exists \mathcal{H}_A \exists \mathcal{H}_{\gamma}. (\Phi^h \wedge \Psi^h)$ is an invariant of $\|_{\gamma} B_i$.

Proof. Similar to Proposition 1. □

We use history clocks for interactions to express additional constraints on their timing. The starting point is the observation that when two conflicting interactions compete for the same action a , no matter which one is first, the latter must wait until the component which owns a is again able to execute a . This is referred to as a “separation constraint” for conflicting interactions.

Definition 7 (Separation Constraints for Interaction Clocks). *Given an interaction set γ , the induced separation constraints, $\mathcal{S}(\gamma)$, are defined as follows:*

$$\mathcal{S}(\gamma) = \bigwedge_{a \in \text{Act}(\gamma)} \bigwedge_{\substack{\alpha \neq \beta \in \gamma \\ a \in \alpha \cap \beta}} |h_{\alpha} - h_{\beta}| \geq k_a$$

where $|x|$ denotes the absolute value of x and k_a is a constant computed locally on the component executing a , and representing the minimum elapsed time between two consecutive executions of a .

In our running example the only conflicting actions are a and c within the controller, and both k_a and k_c are equal to 4. The expression of the separation constraints reduces to:

$$\mathcal{S}((a \mid b_i)_i, (c \mid d_i)_i) \equiv \bigwedge_{i \neq j} |h_{c|d_i} - h_{c|d_j}| \geq 4 \wedge \bigwedge_{i \neq j} |h_{a|b_i} - h_{a|b_j}| \geq 4.$$

Proposition 5. *Let*

$$\mathcal{S}^*(\gamma) = \bigwedge_{a \in \text{Act}(\gamma)} \bigwedge_{\substack{\alpha \neq \beta \in \gamma \\ a \in \alpha \cap \beta}} (h_a \leq h_{\alpha} \leq h_{\beta} - k_a \vee h_a \leq h_{\beta} \leq h_{\alpha} - k_a)$$

Then

1. $\mathcal{S}^*(\gamma)$ is an inductive predicate of $B^* \|_{\gamma^h} B_i^h$.
2. The equivalence $\mathcal{S}(\gamma) \equiv \exists \mathcal{H}_A. \mathcal{S}^*(\gamma)$ is a valid formula.

Proof. (1) Let us fix an arbitrary term $S(a, \alpha, \beta)$ defined as

$$S(a, \alpha, \beta) = (h_a \leq h_\alpha \leq h_\beta - k_a \vee h_a \leq h_\beta \leq h_\alpha - k_a)$$

Assume $S(a, \alpha, \beta)$ holds in an arbitrary state s of $B^* \parallel_{\gamma^h} B_i^h$. Then, it obviously holds for any time successors as well as for any discrete successors by interactions not containing the action a . For an interaction involving a , but different than α and β , h_a is reset to zero whereas h_α and h_β are unchanged. Henceforth, $S(a, \alpha, \beta)$ remains valid as only h_a changes to 0. Let consider the situation α is executed (the case of β is perfectly dual). In this case, both h_a and h_α are reset to 0, whereas h_β is unchanged. Two situations can happen, either

1. $h_a \leq h_\alpha \leq h_\beta - k_a$ holds in s . Then, obviously, the same holds in s' where h_a and h_α are reset.
2. $h_a \leq h_\beta \leq h_\alpha - k_a$ holds in s . This is the interesting case where we need the assumption about the separation time k_a . As consecutive executions of a are separated by k_a , to execute α it must actually hold that $h_a \geq k_a$ in s . Consequently, $h_\beta \geq k_a$ in s , as well as in s' (because h_β does not change from s to s'). Then, knowing that $h_a = h_\alpha = 0$ in s' we have that $h_a \leq h_\alpha \leq h_\beta - k_a$ in s' .

(2) We can equivalently write

$$\begin{aligned} \mathcal{S}^*(\gamma) &\equiv \bigwedge_{a \in \text{Act}(\gamma)} \bigwedge_{\substack{\alpha \neq \beta \in \gamma \\ a \in \alpha \cap \beta}} (h_a \leq h_\alpha \wedge h_a \leq h_\beta \wedge |h_\alpha - h_\beta| \leq k_a) \\ &\equiv \mathcal{S}(\gamma) \wedge \bigwedge_{a \in \text{Act}(\gamma)} \bigwedge_{\substack{\alpha \neq \beta \in \gamma \\ a \in \alpha \cap \beta}} (h_a \leq h_\alpha \wedge h_a \leq h_\beta) \end{aligned}$$

which leads immediately to the result. □

The predicate $\mathcal{S}(\gamma)$ is expressed over history clocks for interactions. Component invariants $CI(B_i^h)$ are however expressed using history clocks for actions. In order to “glue” them together in a meaningful way, we need some tighter connection between action and interaction history clocks. This connection is formally addressed by the constraints \mathcal{E}^* defined below.

Definition 8 (\mathcal{E}^*). *Given an interaction set γ , we define $\mathcal{E}^*(\gamma)$ as follows:*

$$\mathcal{E}^*(\gamma) = \bigwedge_{a \in \text{Act}(\gamma)} h_a = \min_{\alpha \in \gamma, a \in \alpha} h_\alpha.$$

By a similar argument as the one in Proposition 2, it can be shown that $\mathcal{E}^*(\gamma)$ is an inductive predicate of the extended system $B^* \parallel_{\gamma^h} B_i^h$. Moreover, there exists a tight connection between \mathcal{E} and \mathcal{E}^* as given in Proposition 6.

Proposition 6.

1. $\mathcal{E}^*(\gamma)$ is an inductive predicate of $B^* \parallel_{\gamma^h} B_i^h$.
2. The equivalence $\exists \mathcal{H}_\gamma. \mathcal{E}^*(\gamma) \equiv \mathcal{E}(\gamma)$ is a valid formula.

Proof. (1) To see that $\mathcal{E}^*(\gamma)$ is an inductive predicate it suffices to note that, the predicate is preserved by time progress transitions and for any discrete action a , there is always an interaction α containing a such that h_a and h_α are both reset in the same time.

(2) The proof follows directly from the definitions of $\mathcal{E}(\gamma)$ and $\mathcal{E}^*(\gamma)$. Consider that $\gamma = \{\alpha_1, \alpha_2, \dots, \alpha_m\}$. We have the following equivalences:

$$\begin{aligned}
 \exists \mathcal{H}_\gamma. \mathcal{E}^*(\gamma) &\equiv \exists \mathcal{H}_\gamma. \bigvee_{\alpha_{k_1} \prec \alpha_{k_2} \prec \dots \prec \alpha_{k_m}} (h_{\alpha_{k_1}} \leq h_{\alpha_{k_2}} \leq \dots \leq h_{\alpha_{k_m}} \wedge \mathcal{E}^*(\gamma)) \\
 &\text{(by choosing an arbitrary ordering } \prec \text{ on interactions)} \\
 &\equiv \exists \mathcal{H}_\gamma. \bigvee_{\alpha_{k_1} \prec \alpha_{k_2} \prec \dots \prec \alpha_{k_m}} (h_{\alpha_{k_1}} \leq h_{\alpha_{k_2}} \leq \dots \leq h_{\alpha_{k_m}} \wedge \\
 &\quad \bigwedge_{a \in \alpha_{k_1}} (h_a = h_{\alpha_{k_1}}) \wedge \bigwedge_{a \in \alpha_{k_2} \setminus \alpha_{k_1}} (h_a = h_{\alpha_{k_2}}) \wedge \dots \bigwedge_{a \in \alpha_{k_m} \setminus \alpha_{k_1} \dots \alpha_{k_{m-1}}} (h_a = h_{\alpha_{k_m}})) \\
 &\text{(by expanding the definition of } \mathcal{E}^*(\gamma) \text{ along the chosen order)} \\
 &\equiv \exists \mathcal{H}_\gamma. \bigvee_{\alpha_{k_1} \prec \alpha_{k_2} \prec \dots \prec \alpha_{k_m}} (h_{\alpha_{k_1}} \leq h_{\alpha_{k_2}} \leq \dots \leq h_{\alpha_{k_m}} \wedge \bigwedge_{\ell=1}^m \bigwedge_{a \in \alpha_{k_\ell} \setminus \alpha_{k_1} \dots \alpha_{k_{\ell-1}}} (h_a = h_{\alpha_{k_\ell}})) \\
 &\text{(by rewriting to a more compact form)} \\
 &\equiv \bigvee_{\alpha_{k_1} \prec \alpha_{k_2} \prec \dots \prec \alpha_{k_m}} \exists \mathcal{H}_\gamma. (h_{\alpha_{k_1}} \leq h_{\alpha_{k_2}} \leq \dots \leq h_{\alpha_{k_m}} \wedge \bigwedge_{\ell=1}^m \bigwedge_{a \in \alpha_{k_\ell} \setminus \alpha_{k_1} \dots \alpha_{k_{\ell-1}}} (h_a = h_{\alpha_{k_\ell}})) \\
 &\text{(by distributing the existential quantifiers over the disjunction)} \\
 &\equiv \bigvee_{\alpha_{k_1} \prec \alpha_{k_2} \prec \dots \prec \alpha_{k_m}} \bigwedge_{\ell=1}^m \bigwedge_{\substack{a_i, a_j \in \alpha_{k_\ell} \setminus \alpha_{k_1} \dots \alpha_{k_{\ell-1}} \\ a_k \notin \alpha_{k_1} \dots \alpha_{k_\ell}}} (h_{a_i} = h_{a_j} \leq h_{a_k}) \equiv \mathcal{E}(\gamma) \\
 &\text{(by eliminating the existential quantifiers)}
 \end{aligned}$$

□

From Propositions 6, 4, and 5, it follows that $\exists \mathcal{H}_A \exists \mathcal{H}_\gamma. (\bigwedge_i CI(B_i^h) \wedge II(\gamma) \wedge \mathcal{E}^*(\gamma) \wedge \mathcal{S}(\gamma))$ is an invariant of $\|_\gamma B_i$. This new invariant is in general stronger than $\exists \mathcal{H}_A. (\bigwedge_i CI(B_i^h) \wedge II(\gamma) \wedge \mathcal{E}(\gamma))$ and it provides better state space approximations for timed systems with conflicting interactions.

Corollary 3. $\Phi = \exists \mathcal{H}_A \exists \mathcal{H}_\gamma. (\bigwedge_i CI(B_i^h) \wedge II(\gamma) \wedge \mathcal{E}^*(\gamma) \wedge \mathcal{S}(\gamma))$ is an invariant of $\|_\gamma B_i$

Example 2. To get some intuition about the invariant generated using separation constraints, let us reconsider the running example with two workers. The subformula which we emphasise here is the conjunction of \mathcal{E}^* and \mathcal{S} . The interaction invariant is:

$$II(\gamma) = (l_{11} \vee lc_1 \vee lc_2) \wedge (l_{12} \vee lc_1 \vee lc_2) \wedge (lc_2 \vee l_{11} \vee l_{12}) \wedge (lc_0 \vee lc_1 \vee l_{21} \vee l_{22})$$

The components invariants are:

$$\begin{aligned}
 CI(\text{Controller}^h) &= (lc_0 \wedge x = h_0 \wedge h_0 < h_a \wedge h_0 < h_c) \vee \\
 &\quad (lc_1 \wedge x \leq h_0 - 8 \wedge x \leq 4 \wedge h_0 < h_a \wedge h_0 < h_c) \vee \\
 &\quad (lc_1 \wedge x \leq 4 \wedge x = h_c \leq h_a \leq h_0 - 12) \vee \\
 &\quad (lc_2 \wedge x \leq h_0 - 12 \wedge h_a = x \wedge h_0 < h_c) \vee \\
 &\quad (lc_2 \wedge x = h_a \wedge h_c = h_a + 4 \leq h_0 - 12)
 \end{aligned}$$

$$\begin{aligned}
 CI(\text{Worker}_i^h) &= (l_{1i} \wedge y_i = h_0 \wedge h_0 < h_{d_i} \wedge h_0 < h_{b_i}) \vee \\
 &\quad (l_{1i} \wedge y_i = h_{d_i} \leq h_{b_i} \leq h_0 - 8) \vee \\
 &\quad (l_{2i} \wedge y_i \geq h_{b_i} + 8 \leq h_0 < h_{d_i}) \vee \\
 &\quad (l_{2i} \wedge y_i = h_{d_i} \leq h_0 - 8 \wedge h_{b_i} \leq h_{d_i} - 8)
 \end{aligned}$$

The inequalities for action and interaction history clocks are:

$$\begin{aligned} \mathcal{E}^*(\gamma) = & (h_{b_1} = h_{a|b_1}) \wedge (h_{b_2} = h_{a|b_2}) \wedge (h_a = \min_{i=1,2}(h_{a|b_i})) \wedge \\ & (h_{d_1} = h_{c|d_1}) \wedge (h_{d_2} = h_{c|d_2}) \wedge (h_c = \min_{i=1,2}(h_{c|d_i})) \end{aligned}$$

By recalling the expression of $\mathcal{S}(\gamma)$ from above we obtain that:

$$\exists \mathcal{H}_\gamma. \mathcal{E}^*(\gamma) \wedge \mathcal{S}(\gamma) = (|h_{b_2} - h_{b_1}| \geq 4 \wedge |h_{d_2} - h_{d_1}| \geq 4)$$

and thus, after quantifier elimination in

$$\exists \mathcal{H}_A \exists \mathcal{H}_\gamma. (CI(\text{Controller}^h) \wedge \bigwedge_i CI(\text{Worker}_i^h) \wedge II(\gamma) \wedge \mathcal{E}^*(\gamma) \wedge \mathcal{S}(\gamma))$$

After simplification, we obtain the following invariant Φ :

$$\begin{aligned} \Phi = & (l_{11} \wedge l_{12} \wedge lc_0 \wedge x = y_1 = y_2) \vee \\ & (l_{11} \wedge l_{12} \wedge lc_1 \wedge x \leq 4 \wedge (y_1 = y_2 \geq x + 8 \vee \\ & \quad (y_1 = x \wedge \mathbf{y_2} - \mathbf{y_1} \geq 4) \vee \\ & \quad (y_1 \geq x + 8 \wedge y_1 - y_2 \geq 8) \vee \\ & \quad (y_2 = x \wedge \mathbf{y_1} - \mathbf{y_2} \geq 4) \vee \\ & \quad (y_2 \geq x + 8 \wedge y_2 - y_1 \geq 8))) \vee \\ & (l_{21} \wedge l_{12} \wedge lc_2 \wedge y_1 \geq x + 8 \wedge ((y_2 \geq x + 4 \wedge |\mathbf{y_1} - \mathbf{y_2}| \geq 4) \vee \\ & \quad y_2 \geq x + 12)) \vee \\ & (l_{11} \wedge l_{22} \wedge lc_2 \wedge y_2 \geq x + 8 \wedge ((y_1 \geq x + 4 \wedge |\mathbf{y_1} - \mathbf{y_2}| \geq 4) \vee \\ & \quad y_1 \geq x + 12)) \end{aligned}$$

We emphasised in the expression of Φ the newly discovered constraints. All in all, Φ is strong enough to prove that the system is deadlock free.

We conclude the section with a discussion about the computation of the separation constants k_a . A simple but incomplete heuristics to test that a given value k_a is a correct separation constraint for an action a is as follows. Consider all elementary paths connecting two transitions (not necessarily distinct) labelled by a . If on every such path, there exists a clock x which is reset and then tested in a guard $x \geq ct$, with $ct \geq k_a$ then, it is safe to conclude that actually k_a is a correct separation value. Nonetheless, alternative methods to exactly compute k_a have been already proposed in the literature. For details, the interested reader can refer, for instance, to [15] which reduces this computation to finding a shortest path in a weighted graph built from the zone graph associated to the component.

4 Improving (VR) - Three Heuristics

We describe and elaborate on heuristics allowing to strengthen the generated invariants and to reduce the generation time. These heuristics have been mandatory for handling our case studies considered later in section 5.

4.1 More about conflicting interactions

The use of initialising the history clock h_0 provides a convenient way to express and reason about invariants relating occurrences of various actions and interactions at execution. The assertion $h_\alpha \leq h_0$ has the intuitive meaning that “ α has been executed”. We describe below a new family of invariants providing a finer characterisation for the execution of conflicting interactions and related actions.

We fix a as a potential conflicting action within some component $B = (L, A, T, \mathcal{X}, \text{tpc})$. We define the set of preceding actions $\text{Prec}(a)$ as all actions of B that can immediately precede a at execution, formally $\text{Prec}(a) = \{a' \in A \mid \exists l, l', l'' \in L. l \xrightarrow{a'} l', l' \xrightarrow{a} l''\}$. For any two conflicting interactions α_1, α_2 involving a the following assertion:

$$h_{\alpha_1} \leq h_0 \wedge h_{\alpha_2} \leq h_0 \Rightarrow \bigvee_{a' \in \text{Prec}(a)} h_{a'} \leq h_0$$

is an invariant. Intuitively, the assertion states that whenever α_1 and α_2 have both been executed (implying that a has also been executed twice or more), at least one of the preceding actions of a must also have been executed. We remark that the invariant above is rather weak and can be implied by the component invariant $CI(B)$ and the glue invariant \mathcal{E}^* in many situations. In fact, whenever a is an action which is *not enabled* at the initial location of B , the component invariant $CI(B)$ implies that

$$h_a \leq h_0 \Rightarrow \bigvee_{a' \in \text{Prec}(a)} h_{a'} \leq h_0.$$

This states that whenever a has been executed, at least one of its preceding actions has been executed as well. Knowing moreover that $h_a = \min_{\alpha \in \alpha} h_\alpha$, we can then infer the invariant above.

Nonetheless, if a is an action that is enabled at the initial location, the newly proposed invariant is stronger and cannot be derived as shown before. In this case, a can be actually executed once while none of its predecessors has been executed yet. The component invariant alone does not relate anymore the execution of a to the execution of its preceding actions. Moreover, the component invariant considers always the last occurrence of a and has no means to distinguish cases where a has been executed only once or more. This information can sometimes be re-discovered when interaction history clocks $h_{\alpha_1}, h_{\alpha_2}$ are taken into account, henceforth, leading to the proposed invariant. A concrete illustration is provided later in Section 5.

4.2 Invariant computation using regular expressions

There exist situations where the computation of component invariants can be extremely costly. In particular, for untimed components extended with history clocks, their zone graphs will most likely have an exponential size. In fact, due to history clocks, the zones will record the order of (the last) occurrences of actions, and there could be exponentially many of them, reachable at different locations. Obviously, the same issue exists for timed components as well, however, the timing constraints expressed on local clocks bring more information in a zone graph, thus restricting the number of possible successors from a given zone.

The above observation suggests (and was confirmed by our experiments) that applying the same methodology for computing component invariants (based on the reachability graph of the corresponding components with history clocks) regardless of the components being timed or not leads to large characterisations (in the size of the formulae) when possibly shorter ones exist.

Example 3. Consider the untimed component presented in Figure 3 (left) and its extension with history clocks (right). The entire zone graph reachable from $\langle l_0, \zeta_0 \rangle$, with $\zeta_0 = (h_0 = 0, h_{a,b,c} > 0)$ has 6 symbolic states. Therefore, the component invariant is expressed as a disjunction of 16 terms, 9 of them are related to location l_0 and 7 are related to location l_1 .

We recall that untimed automata have elegant and compact encodings as regular expressions. This basic fact can be exploited in order to provide an alternative computation method for component invariants. More concretely, given an untimed component $B = (L, A, T)$ we show how to automatically compute the invariant describing the relations between the history clocks of B^h at some location ℓ , from the language accepted by B at some designated location ℓ . The first key observation is that only the last occurrence of each action should be retained. This implies that it is safe to abstract with respect to a last occurrence retention operation the regular expression characterising the language accepted at the chosen control location. The second key observation is that, regular expressions in some restricted form, can be used to directly



Figure 3: An untimed component (left) and its extension with history clocks (right).

generate the associated constraint on the history clocks, in a much compact form. Our regular expression based method can be therefore summarised as follows:

1. construct the regular expression E_ℓ representing the language accepted by B at location ℓ ,
2. abstract E_ℓ with respect to the last occurrence retention towards some restricted form $E_\ell^\# = \sum_i e_i^\#$ where, every $e_i^\#$ contains each action at most once, and does not contain nested $*$ -operators,
3. generate from every $e_i^\#$ a characteristic formula on history clocks $\phi(e_i^\#)$ and obtain as invariant for B the assertion $\ell \Rightarrow \bigvee_i \phi(e_i^\#)$.

The first step is well known for finite automata and will not be detailed here. For the second abstraction step, the key ingredients are the simplification rules in Figure 4.

$$\begin{array}{ll} \textbf{Rule 1 [Last Occurrence Retention]:} & E \cdot a \longrightarrow (E \setminus a) \cdot a \\ \textbf{Rule 2 [Back-unfolding]:} & E^* \longrightarrow (E^* \cdot E) + \varepsilon \end{array}$$

Figure 4: Simplification Rules

Rule 1 eliminates all but the last occurrence of the trailing a symbol from a regular expression of the form $E \cdot a$. The “ \setminus ” denotes a syntactic *elimination operator* defined structurally on expressions as follows. Let a and x be two symbols and E , E_1 and E_2 be arbitrary regular expressions.

$$\begin{aligned} \epsilon \setminus a &= \epsilon \\ x \setminus a &= \begin{cases} \epsilon & \text{if } x = a \\ x & \text{if } x \neq a \end{cases} \\ (E_1 + E_2) \setminus a &= (E_1 \setminus a) + (E_2 \setminus a) \\ (E_1 \cdot E_2) \setminus a &= (E_1 \setminus a) \cdot (E_2 \setminus a) \\ E^* \setminus a &= (E \setminus a)^* \end{aligned}$$

Rule 2 simply unfolds $*$ -expressions once. By using this rule and other basic manipulation of regular expressions, further simplification opportunities for Rule 1 are enabled.

Example 4. Let consider again the example presented in Figure 3. The language accepted at l_1 is defined as $(a + bc^*b)^*bc^*$. This expression is progressively abstracted into the restricted form as follows:

$$\begin{aligned} (a + bc^*b)^*bc^* &\rightsquigarrow (a + c^*)^*bc^* && \text{(by Rule 1)} \\ &\equiv (a + c^*)^*b(c^*c + \epsilon) && \text{(by Rule 2)} \\ &\equiv (a + c^*)^*bc^*c + (a + c^*)^*b && \text{(by splitting the last +)} \\ &\rightsquigarrow (a + \epsilon)^*bc + (a + c^*)^*b && \text{(by Rule 1)} \\ &\equiv a^*bc + (a + c)^*b && \text{(by standard transformation)} \end{aligned}$$

In the example above, we have applied the iterative strategy consisting of (1) choosing symbols from right to left and applying Rule 1 until no longer possible and then (2) applying Rule 2 to unfold the right-most *-expression and split the incoming +. It can be shown that such a strategy always terminates with expressions in the restricted form. Intuitively, what happens is that Rule 2 splits larger expressions into smaller ones and, further, for each of these Rule 1 eliminates repetitions of symbols.

For the third step, we construct from regular expressions in restricted form an equivalent formula on history clocks. This formula represents *exactly* the set of orders on actions (the strings) encoded by the regular expression, formally:

$$\phi(e) \equiv \bigvee_{a_1 \dots a_n \in L(e)} (h_0 \geq h_{a_1} \geq \dots \geq h_{a_n} \wedge \bigwedge_{c \neq a_1, \dots, a_n} h_c > h_0)$$

We introduce the principle of the construction with the help of a generic example. Consider the regular expression $(b_1 + \dots + b_m)^* a_1 \dots a_n$ in restricted form (whenever $a_1, \dots, a_n, b_1, \dots, b_m$ are distinct). The corresponding formula on history clocks is

$$h_0 \geq h_{a_1} \geq \dots \geq h_{a_n} \wedge h_{b_1} \geq h_{a_1} \wedge \dots \wedge h_{b_m} \geq h_{a_1} \wedge \bigwedge_{c \neq a_i, b_j} h_c > h_0.$$

The first part encodes the ordering constraints on the *mandatory* string $a_1 \dots a_n$. These actions all occur (consequently, their history clocks are smaller than h_0) in this precise order. The second part considers the constraints on occurrences of b_j actions, which are *optional*. If some of them occur at execution, they are unconstrained each other, but should however take place before a_1 . Finally, the last part deals with actions c which do not appear in the regular expression. For all of them, their history clocks should be strictly greater than h_0 . We remark that, for this particular example, the obtained formula has linear size with respect to the size of the regular expression. In contrast, the number of strings encoded (i.e., whenever restricted to last occurrences of symbols) is exponential, with respect to the number of b actions. The construction above can be generalised for arbitrary restricted regular expressions without much difficulty. The resulting formula remains of polynomial size (at worse quadratic) with respect to the size of the restricted regular expression provided as input.

Example 5. *The restricted regular expression $a^*bc + (a+c)^*b$ obtained in the previous example translates into the following expression:*

$$(h_0 \geq h_b \geq h_c \wedge h_a \geq h_b) \vee (h_0 \geq h_b \wedge h_a \geq h_b \wedge h_c \geq h_b)$$

We note this expression is significantly smaller, yet logically equivalent to the disjunction of 7 distinct terms corresponding to symbolic zones reached at l_1 as initially presented in Example 3.

To sum up, we described an heuristic which can be applied to untimed components to automatically compute an invariant with a reasonable enough size to be handled by existing SMT solvers. Given an untimed component B , our heuristic makes use of the regular expressions characterizing the language accepted by B to avoid a direct construction of the zone graph of B^h which would result in considerably large invariants.

Proposition 7. *Given B an untimed component, E_l the regular expression characterising the language accepted by B at location l . Let $E_l^\#$ be the result of applying the simplification rules. We have that $\bigvee_l (l \wedge \phi(E_l^\#))$ is an invariant of B^h .*

4.3 Exploiting Symmetry

At a closer examination of the definition of separation constraints in Section 3.2, it can be noticed that it characterises all possible orderings of conflicting interactions with respect to permutations. The size of the corresponding search space is exponential in the number of conflicting interactions and this, in turn, may be a bottleneck for the solver. Such situations are to (and can) be avoided especially in the case of symmetric

systems. What we show next is how the inherent symmetry in the formula can be eliminated such that the search space becomes considerably smaller.

The use of symmetry has long been addressed, mostly with the intention of making model-checking more feasible and especially in the context of parameterised systems [19, 20, 18, 26]. There the goal is to show the existence of a small cutoff bound which allows the reduction of the verification problem from an arbitrary number of instances to a small, fixed one. Our context is different, that is, breaking the symmetry in some of the generated invariants, for an a priori known number of components.

The types of systems we consider next are formed of a fixed number, be it n , of isomorphic components interacting with a controller, thus the interactions are binary. Isomorphic components are obtained from a generic component B by attaching an index i (from 1 to n) to all symbols in B . The resulting component is denoted by P_i . For any i, j , B_i and B_j are isomorphic⁵. For the ease of reference, we denote systems like $C \parallel_{\gamma} B_i$ by the letter M and we use $Exec$ to denote the set of their global executions.

In this framework, the notion of symmetry is intrinsically related to permutations. Let Π_n denote the group of permutations of n . The application of permutations is defined on the structure of systems and properties. For a system M as $C \parallel_{\gamma}^n B_i$, and a permutation π , $\pi(M)$ is defined as $C \parallel_{\pi(\gamma)}^n \pi(B_i)$ where $\pi(B_i)$ is defined as $B_{\pi(i)}$ and $\pi(\gamma)$ as $\{\pi(\alpha) \mid \alpha \in \gamma\}$ with $\pi(a_c \mid a_i) = a_c \mid a_{\pi(i)}$ for α an arbitrary binary interaction between an action a_c of C and an action a_i of a P_i . For an execution $\sigma = \alpha_1, \dots, \alpha_i, \dots, \alpha_k$, $\pi(\sigma)$ is defined as $\pi(\alpha_1), \pi(\alpha_2), \dots, \pi(\alpha_i), \dots, \pi(\alpha_k)$. For a global state $s = (s_c, s_1, \dots, s_n)$, $\pi(s)$ is defined as $(s_c, s_{\pi(1)}, \dots, s_{\pi(n)})$. As for system properties φ , we restrict to those built (with the usual logical connectors) from clock constraints and locations, and define:

$$\pi(\varphi) = \begin{cases} x_{\pi(i)} \text{ rop } x_{\pi(j)} & \text{if } \varphi = x_i \text{ rop } x_j \text{ and } \text{rop} \in \{<, \leq, =, >, \geq\} \\ l_{\pi(i)} & \text{if } \varphi = l_i \\ \neg\pi(\varphi_1) & \text{if } \varphi = \neg\varphi_1 \\ \pi(\varphi_1) \text{ op } \pi(\varphi_2) & \text{if } \varphi = \varphi_1 \text{ op } \varphi_2 \text{ and } \text{op} \in \{\wedge, \vee\} \end{cases}$$

where l_i, x_i denote a location, respectively, a clock in P_i .

The symmetric systems we consider are symmetric in a “strong” sense, i.e., they are *fully symmetric*. A system M is fully symmetric if for any $\pi \in \Pi_n$, $\pi(M)$ is syntactically identical to M . Similarly, a property φ is fully symmetric if for any permutation π , $\pi(\varphi)$ is equivalent to φ . A property like $l_1 \wedge l_2 \wedge \dots \wedge l_n$ is symmetric. On the contrary, $G = x_1 \leq x_2$ is not as for the permutation $\pi(1) = 2, \pi(2) = 1, \pi(G) = x_{\pi(1)} \leq x_{\pi(2)} = x_2 \leq x_1$ which is not equivalent to G .

Symmetric systems have the convenient property that, whenever started in a symmetric state, for any of its executions $\sigma \in Exec$, $\pi(\sigma)$ is itself an execution, that is, $\pi(\sigma) \in Exec$. To see why this is indeed the case, let γ be the interaction set and $\alpha = (a_c \mid a_i)$ an interaction in γ . It suffices to note that if α is possible after σ , then it is also the case for $\pi(\alpha)$ after $\pi(\sigma)$. Note also that, thanks to symmetry, $\pi(\alpha)$ is in γ .

The idea behind simplifying the separation constraints \mathcal{S} is to break the symmetry by replacing the constraints on absolute values $|h_{\alpha_i} - h_{\alpha_j}|$. More precisely, given a conflicting (controller) action a_c , in an execution where interaction $\alpha_i = a_c \mid a_i$ executes before $\alpha_j = a_c \mid a_j$ for $j > i$, we can naturally replace $|h_{\alpha_i} - h_{\alpha_j}|$ by $h_{\alpha_i} - h_{\alpha_j}$. As for an execution which violates this natural ordering (or “canonicity”), we show that we can make use of symmetry to rearrange it. First, we formalise what we mean more precisely by *canonicity*. Given an execution σ and an interaction $\alpha_i = a_c \mid a_i$ we denote by $lpos(\sigma, \alpha_i)$ the last position of α_i in σ . An execution σ is canonical with respect to a_c if $lpos(\sigma, \alpha_i) < lpos(\sigma, \alpha_j)$ for any $i < j$. Let $Exec^c$ be the set of canonical executions. Thanks to symmetry, any execution has a corresponding canonical execution. Assume σ is such that there is a conflicting a_c and for $i > j$ the last occurrence of $\alpha_i = a_c \mid a_i$ appears latter than that of $\alpha_j = a_c \mid a_j$. Let π be such that $\pi(i) = j$ and $\pi(j) = i$. Then $\pi(\sigma)$ is itself an execution and is canonical.

For a canonical execution with a_c being the action of interest \mathcal{S} simplifies to:

$$\mathcal{S}^c(\gamma) = \bigwedge_{\substack{i < j \\ a_c \in \alpha_i \cap \alpha_j}} h_{\alpha_i} - h_{\alpha_j} \geq k_{a_c} \wedge \bigwedge_{\substack{b \neq a_c \\ b \in \beta_i \cap \beta_j}} |h_{\beta_i} - h_{\beta_j}| \geq k_b$$

⁵We note that, by construction, isomorphic components cannot have clock constraints involving indices: any constraint in a worker B_i is obtained from those in B which are oblivious to indices i .

We note that \mathcal{S}^c reduces \mathcal{S} by $n!$. This is the best we can get in general. However, under particular conditions, \mathcal{S} can be further reduced. For instance, if the controller is such that it ensures mutual exclusion and moreover, requires the use of some designated action a_c , then \mathcal{S} further reduces to:

$$\bigwedge_{a \in Act(C)} \bigwedge_{\substack{i < j \\ a \in \alpha_i \cap \alpha_j}} h_{\alpha_i} - h_{\alpha_j} \geq k_{a_c}$$

This is because by mutual exclusion, all conflicting interactions involving the controller follow the same order as defined for the designated action a_c .

Finally, we show that for symmetric systems and properties it is correct to consider \mathcal{S}^c instead of \mathcal{S} .

Proposition 8. *Let M be a symmetric system and φ be a symmetric property. We have that if $\vdash GI[\mathcal{S} \leftarrow \mathcal{S}^c] \rightarrow \varphi$ then $M \models \Box\varphi$.*

Proof. (sketch) It suffices to show that $\vdash GI[\mathcal{S} \leftarrow \mathcal{S}^c] \rightarrow \varphi$ iff $GI \rightarrow \varphi$.

“ \Leftarrow ”: trivial. “ \Rightarrow ”: It boils down to show that if φ is an invariant of $Exec^c$ then it is also an invariant of the remaining executions σ in $Exec \setminus Exec^c$. If σ does not have a , we are done, as \mathcal{S}^c is an invariant by default. Else, we make use of the fact that σ has a canonical representation and that φ is symmetric. \square

An immediate application of the above reduction results in the simplification we make use of in the Temperature controller example from Section 5. Naturally, the results can be extended also to systems with less symmetry by adapting the standard constructions of automorphisms from, for example, [20]. More precisely, for a system M for which $Aut(M) = \{\pi \mid \pi(M) = M\}$ is a proper subgroup of Π_n , we need to restrict to canonical executions which are consistent with the permutations in $Aut(M)$. However, though such a generalisation is possible, it is not clear if it is also useful: as it is well pointed out in the literature about symmetries, determining $Aut(M)$ is, in itself, is a hard problem. This, together with the goal of keeping the presentation as clear as possible, were the reasons why we strictly considered only *fully* symmetric systems.

5 Implementation and Experiments

The method has been implemented in a Scala (scala-lang.org/) prototype (www-verimag.imag.fr/~lastefan/tas) which is currently being integrated with the D-Finder tool [10] for verification of Real-Time BIP systems [1]. The prototype takes as input components B_i , an interaction set γ and a global safety property φ and checks whether the system satisfies φ . Internally, it uses PPL (bugseng.com/products/ppl) to manipulate zones (essentially polyhedra) and to compute component invariants. It generates Z3 (z3.codeplex.com) Python code to check the satisfiability of the formula $\bigwedge_i CI(B_i) \wedge II(\gamma) \wedge \Phi^* \wedge \neg\varphi$ where Φ^* , depending on whether γ is conflicting, stands for $\mathcal{E}(\gamma)$ or $\mathcal{E}^*(\gamma) \wedge \mathcal{S}(\gamma)$. If the formula is not satisfiable, the prototype returns `no solution`, that is, the system is guaranteed to satisfy φ . Otherwise, it returns a substitution for which the formula is satisfiable, that is, the conjunction of invariants is true while φ is not. This substitution may correspond to a false positive in the sense that the state represented by the substitution could be unreachable.

For experiments, we chose three classical benchmarks which we discuss below.

5.1 Train gate controller (TGC)

This is a classical example from [3]. The system is composed of a controller, a gate and a number of trains. For simplicity, Figure 5 depicts only one train interacting with the controller and the gate. The controller lowers and raises the gate when a train enters, respectively exits. The safety property of interest is that when a train is at location `in`, the gate has been lowered: $\bigwedge_i (in_i \rightarrow g_2)$. The results are presented in Table 1. When there is only one train in the system, $\mathcal{E}(\gamma)$ is enough to show safety. When there are more trains, be it n , the interactions $approach_i \mid approach$ (respectively $exit_i \mid exit$), for $1 \geq i \geq n$ are in conflict on $approach$ (respectively $exit$) of the controller. In this case, in addition to the separation constraints,

we made use of the first heuristic presented in Section 4.1. More precisely, the invariant generated by the heuristic is as follows:

$$\bigwedge_{i \neq j} ((h_{\text{approach}_i} \leq h_0 \wedge h_{\text{approach}_j} \leq h_0) \rightarrow h_{\text{raise}} \leq h_0)$$

Thanks to its addition, we were able to show that the system is safe also for an arbitrary number of trains.

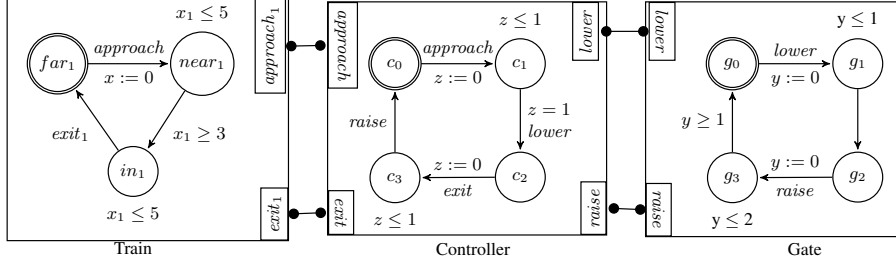


Figure 5: A Controller Interacting with a Train and a Gate

5.2 Fischer protocol

This is a well-studied protocol for mutual exclusion [24]. The protocol specifies how processes can share a resource one at a time by means of a shared variable to which each process assigns its own identifier number. After θ time units, the process with the id stored in the variable enters the critical state and uses the resource. We use an auxiliary component `Id Variable` to mimic the role of the shared variable. The system with two concurrent processes is represented in Figure 6. The property of interest is mutual exclusion: $(cs_i \wedge cs_j) \rightarrow i = j$.

The component `Id Variable` has combinatorial behavior and a large number of actions $(2n + 1)$, thus the generated invariant is huge except for very small values of n . To overcome this issue, we made use of the second heuristic presented in Section 4.2. To simplify, we write s_i instead of set_i and e_i instead of eq_i . We construct the regular expression corresponding to location l_i and project it for actions e_i, e_j, s_i, s_j , respectively e_i, e_0, s_i, s_0 . The latter projection leads to the following regular expression in restricted form:

$$\mathbf{R}_i = (e_0 + s_0)^* e_i . s_i + (e_0 + s_0)^* s_i . e_i + (e_0 + e_i)^* s_0 s_i + (e_i + s_0)^* e_0 s_i + s_i$$

This regular expression translates into the following constraint on history clocks:

$$\begin{aligned} \phi(\mathbf{R}_i) = & (h_{e_0} \geq h_{e_i} \wedge h_{s_0} \geq h_{e_i} \wedge h_{e_i} \geq h_{s_i} \wedge h_{e_i} \leq h_0) \vee \\ & (h_{e_0} \geq h_{s_i} \wedge h_{s_0} \geq h_{s_i} \wedge h_{e_i} \leq h_{s_i} \wedge h_{s_i} \leq h_0) \vee \\ & (h_{e_0} \geq h_{s_0} \wedge h_{e_i} \geq h_{s_0} \wedge h_{s_0} \geq h_{s_i} \wedge h_{s_0} \leq h_0) \vee \\ & (h_{s_0} \geq h_{e_0} \wedge h_{e_i} \geq h_{e_0} \wedge h_{e_0} \geq h_{s_i} \wedge h_{e_0} \leq h_0) \vee \\ & (h_{s_i} \leq h_0 \wedge h_{s_0}, h_{e_0}, h_{e_i} > h_0) \end{aligned}$$

We deduce that $at(l_i) \rightarrow \phi(R_i)$ is an invariant of the `Id Variable`, for any i . These invariants in addition to component invariants of processes and inequality constraints $\mathcal{E}(\gamma)$ are sufficient to show that mutual exclusion holds.

5.3 Temperature controller (TC)

This example is an adaptation from [10]. It represents a simplified model of a nuclear plant. The system consists of a controller interacting with an arbitrary number n of rods (two, in Figure 7) in order to maintain the temperature between the bounds 450 and 900: when the temperature in the reactor reaches 900 (resp. 450), a rod must be used to cool (resp. heat) the reactor. The rods are enabled to cool only after $900n$ units

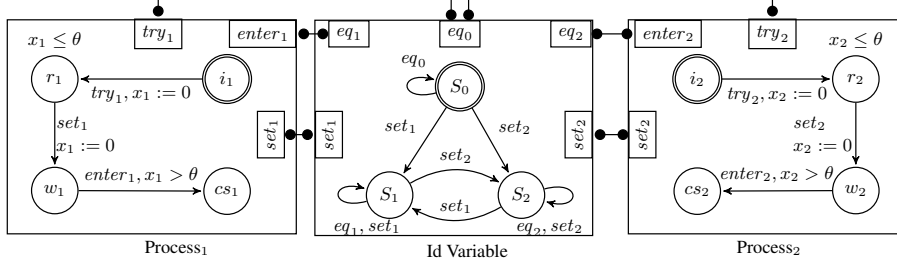


Figure 6: The Fischer Protocol

of time. The global property of interest is the absence of deadlock, that is, the system can run continuously and keep the temperature between the bounds. To express this property in our prototype, we adapt from [28] the definition of *enabled* states, while in Uppaal, we use the query $A[]$ not deadlock. For one rod, $\mathcal{E}(\gamma)$ is enough to show the property. For more rods, because interactions are conflicting, we need the separation constraints which basically bring as new information conjunctions as $\bigwedge_i (h_{rest_{\pi(i)}} - h_{rest_{\pi(i-1)}} \geq 1350)$ for π an ordering on rods. Recalling the discussion from Section 4.3, such a reduction is correct because the system enjoys the particularly helpful property of being symmetric.




Figure 7: A Controller Interacting with 2 Rods

5.4 Results

The experiments were run on a Dell machine with Ubuntu 12.04, an Intel(R) Core(TM)i5-2430M processor of frequency 2.4GHz \times 4, and 5.7GiB memory. The results, synthesised in Table 1, show the potential of our method in terms of accuracy (no false positives) and scalability. For larger numbers of components, the size of the resulting invariants was not problematic for Z3. However, it may be the case that history clocks considerably increase the size of the generated formulae. It can also be observed that Uppaal being highly optimised, it has better scores on the first example in particular and on smaller systems in general. The timings for our prototype are obtained with the Unix command `time` while the results for Uppaal come from the command `verifyta` which comes with the Uppaal 4.1.14 distribution.

6 Conclusion and Future Work

We presented a fully automated compositional method to generate global invariants for timed systems described as parallel compositions of timed automata components using multi-party interactions. The soundness of the method proposed has been proven. In addition, it has been implemented and successfully tested on several benchmarks. The results show that our method may outperform existing exhaustive exploration-based techniques for large systems, thanks to the use of compositionality and over-approximations.

Model & Property	Size	Time/Space	
		$\mathcal{E}^* \wedge \mathcal{S}$	Uppaal
TGC & mutual exclusion	1*	0m0.156s/2.6kB+140B	0ms/8 states
	2	0m0.176s/3.2kB+350B	0ms/13 states
	64	0m4.82s/530kB+170kB	0m0.210s/323 states
	124	0m17.718s/700kB+640kB	0m1.52s/623 states
Fischer & mutual exclusion	2*	0m0.144/3kB	0m0.008s/14 states
	4*	0m0.22s/6.5kB	0m0.012s/156 states
	6*	0m0.36s/12.5kB	0m0.013s/1714 states
	14*	0m2.840s/112kB	no results in 4 hours
TC & absence of deadlock	1*	0m0.172s/840B+60B	0m0.01s/4 states
	8	0m0.5s/23kB+2.4kB	11m0.348s/57922 states
	16	0m2.132s/127kB+9kB	no results in 6 hours
	124	0m19.22s/460kB+510kB	idem

Table 1: Results from Experiments. The marking “*” highlights the cases when \mathcal{E} alone was enough to prove the property. The expressions $x + y$ are to be read as the formula $\wedge_i CI(B_i) \wedge II(\gamma) \wedge \mathcal{E}(\gamma)$, respectively $\mathcal{E}^*(\gamma) \wedge \mathcal{S}(\gamma)$, has length x , respectively y .

This work is currently being extended in several directions. First, we work on integrating it within D-Finder tool [10] and the Real-Time BIP framework [1]. In order to achieve a better integration, we are working on handling *urgencies* [6] on transitions. Actually, urgencies provide an alternative way to constrain time progress, which is more intuitive to use by programmers but much difficult to handle in a compositional way. A second direction of research which is potentially interesting for systems containing identical, replicated components and closely related to the symmetry-based reduction is the application of our method to the verification of parameterised timed systems. Finally, we are considering specific extensions to particular classes of timed systems and properties, in particular, for schedulability analysis of systems with mixed-critical tasks.

References

- [1] T. Abdellatif, J. Combaz, and J. Sifakis. Model-based implementation of real-time applications. In *EMSOFT*, 2010. 5, 6
- [2] L. D. Alfaro, T. A. Henzinger, and M. Stoelinga. Timed interfaces. In *EMSOFT*, 2002. 1, 1
- [3] R. Alur and D. L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 1994. 2, 5.1
- [4] L. Aștefănoaei, S. B. Rayana, S. Bensalem, M. Bozga, and J. Combaz. Compositional invariant generation for timed systems. In *TACAS*, 2014. 1
- [5] B. Badban, S. Leue, and J.-G. Smaus. Automated invariant generation for the verification of real-time systems. In *WING@ETAPS/IJCAR*, 2010. 1
- [6] A. Basu, M. Bozga, and J. Sifakis. Modeling heterogeneous real-time components in BIP. In *SEFM*, 2006. 6
- [7] G. Behrmann, A. David, K. G. Larsen, J. Håkansson, P. Pettersson, W. Yi, and M. Hendriks. UPPAAL 4.0. In *QEST*, 2006. 1
- [8] J. Bengtsson, B. Jonsson, J. Lilius, and W. Yi. Partial order reductions for timed systems. In *CONCUR*, 1998. 1
- [9] J. Bengtsson and W. Yi. On clock difference constraints and termination in reachability analysis of timed automata. In *ICFEM*, 2003. 2

- [10] S. Bensalem, M. Bozga, J. Sifakis, and T.-H. Nguyen. Compositional verification for component-based systems and application. In *ATVA*, 2008. [1](#), [1](#), [2](#), [5](#), [5.3](#), [6](#)
- [11] J. Berendsen and F. W. Vaandrager. Compositional abstraction in real-time model checking. In *FORMATS*, 2008. [1](#)
- [12] S. Bornot and J. Sifakis. An algebraic framework for urgency. *Information and Computation*, 1998. [2](#)
- [13] P. Bouyer. Forward analysis of updatable timed automata. *Form. Methods Syst. Des.*, 2004. [2](#)
- [14] M. Bozga, C. Daws, O. Maler, A. Olivero, S. Tripakis, and S. Yovine. Kronos: A model-checking tool for real-time systems. In *CAV*, 1998. [1](#)
- [15] C. Courcoubetis and M. Yannakakis. Minimum and maximum delay problems in real-time systems. *Formal Methods in System Design*, 1992. [3.2](#)
- [16] A. David, K. G. Larsen, A. Legay, M. H. Møller, U. Nyman, A. P. Ravn, A. Skou, and A. Wasowski. Compositional verification of real-time systems using Ecdar. *STTT*, 2012. [1](#), [1](#)
- [17] F. S. de Boer, U. Hannemann, and W. P. de Roever. Hoare-style compositional proof systems for reactive shared variable concurrency. In *FSTTCS*, 1997. [1](#)
- [18] E. A. Emerson and V. Kahlon. Reducing model checking of the many to the few. In *CADE*, 2000. [4.3](#)
- [19] E. A. Emerson and K. S. Namjoshi. Reasoning about rings. In *POPL*, 1995. [4.3](#)
- [20] E. A. Emerson and A. P. Sistla. Symmetry and model checking. *Formal Methods in System Design*, 9(1/2), 1996. [4.3](#), [4.3](#)
- [21] A. Fietzke and C. Weidenbach. Superposition as a decision procedure for timed automata. *Mathematics in Computer Science*, 2012. [1](#)
- [22] G. Gardey, D. Lime, M. Magnin, and O. H. Roux. ROMEO: A tool for analyzing time petri nets. In *CAV*, 2005. [1](#)
- [23] T. A. Henzinger, X. Nicollin, J. Sifakis, and S. Yovine. Symbolic model checking for real-time systems. *Inf. Comput.*, 1994. [2](#)
- [24] L. Lamport. A fast mutual exclusion algorithm. *ACM Trans. Comput. Syst.*, 1987. [5.2](#)
- [25] S.-W. Lin, Y. Liu, P.-A. Hsiung, J. Sun, and J. S. Dong. Automatic generation of provably correct embedded systems. In *ICFEM*, 2012. [1](#), [1](#)
- [26] K. S. Namjoshi. Symmetry and completeness in the analysis of parameterized systems. In *VMCAI*, 2007. [4.3](#)
- [27] R. B. Salah, M. Bozga, and O. Maler. Compositional timing analysis. In *EMSOFT*, 2009. [1](#)
- [28] S. Tripakis. Verifying progress in timed systems. In *ARTS*, 1999. [2](#), [5.3](#)
- [29] F. Wang. Redlib for the formal verification of embedded systems. In *ISoLA*, 2006. [1](#)

References

- [1] T. Abdellatif, J. Combaz, and J. Sifakis. Model-based implementation of real-time applications. In *EMSOFT*, 2010. 5, 6
- [2] L. D. Alfaro, T. A. Henzinger, and M. Stoelinga. Timed interfaces. In *EMSOFT*, 2002. 1, 1
- [3] R. Alur and D. L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 1994. 2, 5.1
- [4] L. Astefanoaei, S. B. Rayana, S. Bensalem, M. Bozga, and J. Combaz. Compositional invariant generation for timed systems. In *TACAS*, 2014. 1
- [5] B. Badban, S. Leue, and J.-G. Smaus. Automated invariant generation for the verification of real-time systems. In *WING@ETAPS/IJCAR*, 2010. 1
- [6] A. Basu, M. Bozga, and J. Sifakis. Modeling heterogeneous real-time components in BIP. In *SEFM*, 2006. 6
- [7] G. Behrmann, A. David, K. G. Larsen, J. Håkansson, P. Pettersson, W. Yi, and M. Hendriks. UPPAAL 4.0. In *QEST*, 2006. 1
- [8] J. Bengtsson, B. Jonsson, J. Lilius, and W. Yi. Partial order reductions for timed systems. In *CONCUR*, 1998. 1
- [9] J. Bengtsson and W. Yi. On clock difference constraints and termination in reachability analysis of timed automata. In *ICFEM*, 2003. 2
- [10] S. Bensalem, M. Bozga, J. Sifakis, and T.-H. Nguyen. Compositional verification for component-based systems and application. In *ATVA*, 2008. 1, 1, 2, 5, 5.3, 6
- [11] J. Berendsen and F. W. Vaandrager. Compositional abstraction in real-time model checking. In *FORMATS*, 2008. 1
- [12] S. Bornot and J. Sifakis. An algebraic framework for urgency. *Information and Computation*, 1998. 2
- [13] P. Bouyer. Forward analysis of updatable timed automata. *Form. Methods Syst. Des.*, 2004. 2
- [14] M. Bozga, C. Daws, O. Maler, A. Olivero, S. Tripakis, and S. Yovine. Kronos: A model-checking tool for real-time systems. In *CAV*, 1998. 1
- [15] C. Courcoubetis and M. Yannakakis. Minimum and maximum delay problems in real-time systems. *Formal Methods in System Design*, 1992. 3.2
- [16] A. David, K. G. Larsen, A. Legay, M. H. Møller, U. Nyman, A. P. Ravn, A. Skou, and A. Wasowski. Compositional verification of real-time systems using Ecdar. *STTT*, 2012. 1, 1
- [17] F. S. de Boer, U. Hannemann, and W. P. de Roever. Hoare-style compositional proof systems for reactive shared variable concurrency. In *FSTTCS*, 1997. 1
- [18] E. A. Emerson and V. Kahlon. Reducing model checking of the many to the few. In *CADE*, 2000. 4.3
- [19] E. A. Emerson and K. S. Namjoshi. Reasoning about rings. In *POPL*, 1995. 4.3
- [20] E. A. Emerson and A. P. Sistla. Symmetry and model checking. *Formal Methods in System Design*, 9(1/2), 1996. 4.3, 4.3
- [21] A. Fietzke and C. Weidenbach. Superposition as a decision procedure for timed automata. *Mathematics in Computer Science*, 2012. 1
- [22] G. Gardey, D. Lime, M. Magnin, and O. H. Roux. ROMEO: A tool for analyzing time petri nets. In *CAV*, 2005. 1

- [23] T. A. Henzinger, X. Nicollin, J. Sifakis, and S. Yovine. Symbolic model checking for real-time systems. *Inf. Comput.*, 1994. [2](#)
- [24] L. Lamport. A fast mutual exclusion algorithm. *ACM Trans. Comput. Syst.*, 1987. [5.2](#)
- [25] S.-W. Lin, Y. Liu, P.-A. Hsiung, J. Sun, and J. S. Dong. Automatic generation of provably correct embedded systems. In *ICFEM*, 2012. [1](#), [1](#)
- [26] K. S. Namjoshi. Symmetry and completeness in the analysis of parameterized systems. In *VMCAI*, 2007. [4.3](#)
- [27] R. B. Salah, M. Bozga, and O. Maler. Compositional timing analysis. In *EMSOFT*, 2009. [1](#)
- [28] S. Tripakis. Verifying progress in timed systems. In *ARTS*, 1999. [2](#), [5.3](#)
- [29] F. Wang. Redlib for the formal verification of embedded systems. In *ISoLA*, 2006. [1](#)