Learning-based compositional parameter synthesis for event-recording automata^{*}

Étienne André¹ and Shang-Wei Lin²

¹ Université Paris 13, LIPN, CNRS, UMR 7030, France

² SCSE, Nanyang Technological University, Singapore

Abstract. We address the verification of timed concurrent systems with unknown or uncertain constants considered as parameters. First, we introduce parametric event-recording automata (PERAs), as a new subclass of parametric timed automata (PTAs). Although in the non-parametric setting event-recording automata yield better decidability results than timed automata, we show that the most common decision problem remains undecidable for PERAs. Then, given one set of components with parameters and one without, we propose a method to compute an abstraction of the non-parametric set of components, so as to improve the verification of reachability properties in the full (parametric) system. We also show that our method can be extended to general PTAs. We implemented our method, which shows promising results.

1 Introduction

Verifying distributed systems involving timing constraints is notoriously difficult, especially when timing constants may be uncertain. This problem becomes even more difficult (often intractable) in the presence of timing parameters, i. e., unknown timing constants. Parametric reachability synthesis aims at synthesizing timing parameter valuations for which a set of (usually bad) states is reachable. Parametric timed automata (PTAs) [2] is a parametric extension of timed automata (TAs) to model and verify models involving (possibly parametric) timing constraints and concurrency. Its high expressiveness comes with the drawback that most interesting problems are undecidable [3].

Related work Despite undecidability of the theoretical problems, several monolithic (non-compositional) techniques for parametric reachability synthesis in PTAs have been proposed in the past, either in the form of semi-algorithms (a procedure that is correct but may not terminate), or using approximations. In [2], a basic semi-algorithm (called *EFsynth* in [14]) has been proposed: it explores the symbolic state space until bad states are found, and gathers the associated parameter constraints. In [12], approximated parametric reachability synthesis is performed using counter-example guided abstraction refinement (CEGAR) techniques for parametric linear hybrid automata, a class of models more expressive

^{*} This work is partially supported by the ANR national research program “PACS” (ANR-14-CE28-0002).

than PTAs. In [7], we proposed a point-based technique: instead of attacking the reachability synthesis in a brute-force manner, we iterate on (some) integer parameter valuations, and derive for each of them a constraint around this valuation that preserves the (non-)reachability of the bad locations. Although numerous iterations may be needed, each of them explores a much smaller part of the state space than the brute-force exploration of **EFsynth**, often resulting in a faster execution than **EFsynth**.

Distributed systems are often made of a set of components interacting with each other; taking advantage of the compositionality is a goal often desired to speed up verification. In [11], a learning-based approach is proposed to automate compositional verification of untimed systems modeled by labeled transition systems (LTS). For timed systems, we proposed a learning-based compositional verification framework [15] for event-recording automata (ERAs), a subclass of TAs for which language inclusion is decidable [1]. This approach showed to be much faster than monolithic verification.

The recent work [9] is close to our goal, as it proposes an approach for compositional parameter synthesis, based on the derivation of interaction and component invariants. The method is implemented in a prototype in Scala, making use of **IMITATOR** [5]. Whereas both [9] and our approach address reachability or safety properties, the class of PTAs of [9] is larger; conversely, we add no further restrictions on the models, whereas in [9] all clocks and (more problematically) parameters must be local to a single component and cannot be shared.

Contribution In this work, we propose an approach relying on a point-based technique for parametric reachability synthesis, combined with learning-based abstraction techniques, for a subclass of PTAs, namely parametric event-recording automata. We propose this subclass due to the decidability of the language inclusion in the non-parametric setting. We consider a set of parametric components A (where parameters are dense in a bounded parameter domain D_0) and a set of non-parametric components B , with their parallel composition denoted by $A \parallel B$. For each integer parameter valuation v not yet covered by a good or bad constraint, we try to compute, by learning, an abstraction $\tilde{B}$ of B s.t. $v(A) \parallel B$ does not reach the bad locations. We then “enlarge” the valuation v using the abstract model $A \parallel \tilde{B}$, which yields a dense good constraint; we prove the correctness of this approach. If the learning fails to compute an abstraction, we derive a counter-example, and we then replay it in the fully parametric model $A \parallel B$, which allows us to derive very quickly a bad dense constraint. We iterate until (at least) all integer points in D_0 are covered. In practice, we cover not only all rational-valued in D_0 , but in fact the entire parameter space (except for one benchmark for which we fail to compute a suitable abstraction).

We propose the following technical contributions:

1. we introduce a parametrization of event-recording automata (PERAs);
2. we show that the reachability emptiness problem is undecidable for PERAs;
3. we then introduce our approach that combines iteration-based synthesis with learning-based abstraction;

4. we implement our approach into a toolkit using IMITATOR and CV, and we demonstrate its efficiency on several case studies.

Outline Section 2 introduces the necessary preliminaries. Section 3 recalls the parametric reachability preservation [7]. Section 4 introduces parametric event-recording automata, and proves the undecidability of the reachability emptiness problem. Section 5 introduces our main contribution, and Section 6 evaluates it on benchmarks. Section 7 concludes the paper.

2 Preliminaries

2.1 Clocks, parameters and constraints

Let $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}_+$ and $\mathbb{R}_+$ denote the sets of non-negative integers, integers, non-negative rational and non-negative real numbers respectively.

Throughout this paper, we assume a set $X = \{x_1, \dots, x_H\}$ of *clocks*, i.e., real-valued variables that evolve at the same rate. A clock valuation is a function $\mu : X \rightarrow \mathbb{R}_+$. We write $\mathbf{0}$ for the clock valuation that assigns 0 to all clocks. Given $d \in \mathbb{R}_+$, $\mu + d$ denotes the valuation such that $(\mu + d)(x) = \mu(x) + d$, for all $x \in X$. Given $R \subseteq X$, we define the *reset* of a valuation μ , denoted by $[\mu]_R$, as follows: $[\mu]_R(x) = 0$ if $x \in R$, and $[\mu]_R(x) = \mu(x)$ otherwise.

We assume a set $P = \{p_1, \dots, p_M\}$ of *parameters*, i.e., unknown rational-valued constants. A parameter *valuation* (or *point*) v is a function $v : P \rightarrow \mathbb{Q}_+$.

In the following, we assume $\triangleleft \in \{<, \leq\}$ and $\bowtie \in \{<, \leq, \geq, >\}$. Throughout this paper, lt denotes a linear term over $X \cup P$ of the form $\sum_{1 \leq i \leq H} \alpha_i x_i + \sum_{1 \leq j \leq M} \beta_j p_j + d$, with $\alpha_i, \beta_j, d \in \mathbb{Z}$. Similarly, plt denotes a parametric linear term over P , that is a linear term without clocks ($\alpha_i = 0$ for all i). A *constraint* C (i.e., a convex polyhedron) over $X \cup P$ is a conjunction of inequalities of the form $lt \bowtie 0$. Given a parameter valuation v , $v(C)$ denotes the constraint over X obtained by replacing each parameter p in C with $v(p)$. Likewise, given a clock valuation μ , $\mu(v(C))$ denotes the Boolean value obtained by replacing each clock x in $v(C)$ with $\mu(x)$.

A *guard* g is a constraint over $X \cup P$ defined by a conjunction of inequalities of the form $x \bowtie plt$.

A parameter constraint K is a constraint over P . We write $v \models K$ if $v(K)$ evaluates to true. $\perp$ (resp. $\top$) denotes the special parameter constraint containing no (resp. all) parameter valuations. We will sometime manipulate *non-convex* constraints over P , i.e., finite unions of parameter constraints. Such non-convex constraints can be implemented using finite lists of constraints, and therefore all definitions extend in a natural manner to non-convex constraints.

A *parameter domain* is a box parameter constraint, i.e., a conjunction of inequalities of the form $p \bowtie d$, with $d \in \mathbb{N}$. A parameter domain D is *bounded* if, for each parameter, there exists in D an inequality $p \triangleleft d$ (recall that, additionally, all parameters are bounded below from 0 as they are non-negative). Therefore D can be seen as a hypercube in M dimensions.

2.2 Parametric Timed Automata

Definition 1 (PTA). A parametric timed automaton (hereafter PTA) A is a tuple $(\Sigma, L, l_0, X, P, I, E)$, where: i) Σ is a finite set of actions, ii) L is a finite set of locations, iii) $l_0 \in L$ is the initial location, iv) X is a finite set of clocks, v) P is a finite set of parameters, vi) I is the invariant, assigning to every $l \in L$ a guard $I(l)$, vii) E is a finite set of edges $e = (l, g, a, R, l')$ where $l, l' \in L$ are the source and target locations, $a \in \Sigma$, $R \subseteq X$ is a set of clocks to be reset, and g is a guard.

Given a PTA A and a parameter valuation v , we denote by $v(A)$ the non-parametric timed automaton where all occurrences of a parameter p_i have been replaced by $v(p_i)$.

As usual, PTAs can be composed by performing their parallel composition, i.e., their synchronized product on action names.

Definition 2 (Concrete semantics). Given a PTA $A = (\Sigma, L, l_0, X, P, I, E)$, and a parameter valuation v , the concrete semantics of $v(A)$ is given by the timed transition system $(S, s_0, \rightarrow)$, with $S = \{(l, \mu) \in L \times \mathbb{R}_+^H \mid \mu(v(I(l))) \text{ is true}\}$, $s_0 = (l_0, \mathbf{0})$, and $\rightarrow$ consists of the discrete and delay transition relations:

- discrete transitions: $(l, \mu) \xrightarrow{e} (l', \mu')$, if $(l, \mu), (l', \mu') \in S$, there exists $e = (l, g, a, R, l') \in E$, $\mu' = [\mu]_R$, and $\mu(v(g))$ is true.
- delay transitions: $(l, \mu) \xrightarrow{d} (l, \mu + d)$, with $d \in \mathbb{R}_+$, if $\forall d' \in [0, d], (l, \mu + d') \in S$.

A (concrete) run is a sequence $\rho = s_0 \gamma_0 s_1 \gamma_1 \dots s_n \gamma_n \dots$ such that $\forall i, (s_i, \gamma_i, s_{i+1}) \in \rightarrow$. We consider as usual that concrete runs strictly alternate delays d_i and discrete transitions e_i and we thus write concrete runs in the form $\rho = s_0 \xrightarrow{(d_0, e_0)} s_1 \xrightarrow{(d_1, e_1)} \dots$. The corresponding *timed word* is $(a_0, t_0), (a_1, t_1), \dots$ where a_i is the action of e_i and $t_i = \sum_{j=0}^i d_j$. Given a state $s = (l, \mu)$, we say that s is reachable (or that $v(A)$ reaches s) if s belongs to a run of $v(A)$. By extension, we say that l is reachable in $v(A)$, if there exists a state (l, μ) that is reachable. Given $L^\odot \subseteq L$, we say that $L^\odot$ is reachable in $v(A)$ if $\exists l \in L^\odot$ s.t. l is reachable.

Let $\rho = (l_0, \mu_0) \xrightarrow{(d_0, e_0)} (l_1, \mu_1) \xrightarrow{(d_1, e_1)} \dots (l_n, \mu_n) \xrightarrow{(d_n, e_n)} \dots$ be a run of $v(A)$. The *trace* of this run (denoted by $\text{trace}(\rho)$) is the sequence $e_0 e_1 \dots e_n \dots$, and the *untimed word* of this run is $a_0 a_1 \dots a_n \dots$, where a_i is the action of e_i for all i . The *trace set* of $v(A)$ is the set of traces associated with all runs of A .

Symbolic semantics Let us recall the symbolic semantics of PTAs (as in e.g., [4, 14]). We define the *time elapsing* of a constraint C , denoted by $C^\nearrow$, as the constraint over X and P obtained from C by delaying all clocks by an arbitrary amount of time. That is, $C^\nearrow = \{(\mu, v) \mid \mu \models v(C) \wedge \forall x \in X : \mu'(x) = \mu(x) + d, d \in \mathbb{R}_+\}$. Given $R \subseteq X$, we define the *reset* of C , denoted by $[C]_R$, as the constraint obtained from C by resetting the clocks in R , and keeping the other clocks unchanged. We denote by $C \downarrow_P$ the projection of C onto P , i.e., obtained by eliminating the clock variables (e.g., using Fourier-Motzkin).

A *parametric zone* is a convex polyhedron over $X \cup P$ in which constraints are of the form $x \bowtie plt$, or $x_i - x_j \bowtie plt$, where $x_i, x_j \in X$ and plt is a parametric linear term over P .

A symbolic state is a pair $\mathbf{s} = (l, C)$ where $l \in L$ is a location, and C its associated parametric zone. The initial symbolic state of A is $\mathbf{s}_0 = (l_0, (\{\mathbf{0}\} \wedge I(l_0))^{\nearrow} \wedge I(l_0))$.

The symbolic semantics relies on the **Succ** operation. Given a symbolic state $\mathbf{s} = (l, C)$ and an edge $e = (l, g, a, R, l')$, $\text{Succ}(\mathbf{s}, e) = (l', C')$, with $C' = (([C \wedge g])_R \wedge I(l'))^{\nearrow} \wedge I(l')$. The **Succ** operation is effectively computable, using polyhedra operations; also note that the successor of a parametric zone C is a parametric zone (see e. g., [14]).

A symbolic run of a PTA is an alternating sequence of symbolic states and edges starting from the initial symbolic state, of the form $\mathbf{s}_0 \xrightarrow{e_0} \mathbf{s}_1 \xrightarrow{e_1} \dots \xrightarrow{e_{m-1}} \mathbf{s}_m$, such that for all $i = 0, \dots, m-1$, we have $e_i \in E$, and $\mathbf{s}_{i+1} = \text{Succ}(\mathbf{s}_i, e_i)$.

Given a symbolic run $\mathbf{s}_0 \xrightarrow{e_0} \mathbf{s}_1 \xrightarrow{e_1} \dots$, its *trace* is the sequence $e_0 e_1 \dots$. Two runs (symbolic or concrete) are *equivalent* if they have the same trace.

3 Parametric reachability preservation

Let us briefly recall the parametric reachability preservation algorithm PRP [7]. Given a set of locations $L^\ominus$, $\text{PRP}(A, v, L^\ominus)$ synthesizes a dense (convex) constraint K containing at least v and such that, for all $v' \in K$, $v'(A)$ preserves the reachability of $L^\ominus$ in $v(A)$. By preserving the reachability of $L^\ominus$ in $v(A)$, we mean that some locations of $L^\ominus$ are reachable in $v'(A)$ iff they are in $v(A)$. That is, if $v(A)$ is safe (i. e., it does not reach $L^\ominus$), then $v'(A)$ is safe too. Conversely, if $v(A)$ is unsafe (i. e., $L^\ominus$ is reachable for some runs), then $v'(A)$ is unsafe too.

Lemma 1 (Soundness of PRP [7]). *Let A be a PTA, v a parameter valuation, and $L^\ominus$ a subset of locations. Let $K = \text{PRP}(A, v, L^\ominus)$.*

For all $v' \models K$, $v'(A)$ reaches $L^\ominus$ iff $v(A)$ reaches $L^\ominus$.

A specificity of PRP is that it does *not* aim at completeness; instead, it focuses on behaviors “similar” to that of $v(A)$ so as not to explore a too large part of the state space, and outputs valuations neighboring v . A sort of completeness can be achieved by iterating PRP on various parameter valuations: when $v(A)$ has computed K , the algorithm can be called again on a valuation v_2 “neighbor” of the result K , and so on until either the entire parameter space has been covered, or when a certain coverage of a bounded parameter domain has been achieved (e. g., 99%). This iterated version is called PRPC (for PRP cartography), takes as input a PTA A and a bounded parameter domain D_0 , and iteratively calls PRP on parameter valuations of D_0 with a given precision (e. g., at least all integer-valued). This gives a cartography of D_0 with a union K_{good} of safe constraints (valuations for which $L^\ominus$ is unreachable) and a union K_{bad} of unsafe constraints (for which $L^\ominus$ is reachable). Although only the coverage of the discrete points (e. g., integer-valued) can be theoretically guaranteed, PRPC often covers most (if not all) of the dense state space within D_0 , and often outside too.

Fig. 1: An example of a PERA

4 Parametric event-recording automata

Event-recording automata (ERAs) [1] are a subclass of timed automata, where each action label is associated with a clock such that, for every edge with a label, the associated clock is reset. We propose here a parametric extension of ERAs, following the parameterization of TAs into PTAs.

Formally, let Σ be a set of actions: we denote by X_Σ the set of clocks associated with Σ , i. e., $\{x_a \mid a \in \Sigma\}$. A Σ -guard is a guard on $X_\Sigma \cup P$.

Definition 3 (PERAs). A parametric event-recording automaton (PERA) is a tuple $(\Sigma, L, l_0, P, I, E)$, where: *i)* Σ is a finite set of actions, *ii)* L is a finite set of locations, *iii)* $l_0 \in L$ is the initial location, *iv)* P is a finite set of parameters, *v)* I is the invariant, assigning to every $l \in L$ a Σ -guard $I(l)$, *vi)* E is a finite set of edges $e = (l, g, a, x_a, l')$ where $l, l' \in L$ are the source and target locations, $a \in \Sigma$, x_a is the clock to be reset, and g is a Σ -guard.

Just as for ERAs, PERAs can be seen as a syntactic subclass of PTAs: a PERA is a PTA for which there is a one-to-one matching between clocks and actions and such that, for each edge, the clock corresponding to the action is the only clock to be reset.

Following the conventions used for ERAs, we do not explicitly represent graphically the clock x_a reset along an edge labeled with a : this is implicit.

Example 1. Fig. 1 depicts an example of PERA with 3 actions (and therefore 3 clocks x_a , x_b and x_c), and one parameter p . Only clock x_a is used in a guard.

It is well-known that the EF-emptiness problem (“is the set of parameter valuations for which it is possible to reach a given location empty?”) is undecidable for PTAs [2,6]. Reusing the proof of [6], we show below that this remains undecidable for PERAs.

Theorem 1. *The EF-emptiness problem is undecidable for PERAs, even with bounded parameters.*

Proof. The proof works by adapting to PERAs the proof of [6, Theorem 1].

This negative result rules out the possibility to perform exact synthesis for PERAs. Still, in the next section, we propose an approach that is sound, though maybe not complete: the synthesized valuations are correct, but some may be missing. More pragmatically, we aim at improving the synthesis efficiency.

Fig. 2: AGR proof rule (left) and TL* (right)

5 Compositional parameter synthesis for PERAs

Fig. 2a recalls the common proof rule used in Assume-Guarantee Reasoning (AGR), which is one of the compositional verification techniques. Given two components A , B and a safety property φ , the proof rule tells us that if A can satisfy the property φ under an assumption $\tilde{B}$ and B can guarantee this assumption $\tilde{B}$, then we can conclude that $A \parallel B$ satisfies φ .

5.1 Partitioning the system

The proof rule is presented in the context of two components. If a system consists of more than two components, an intuitive way is to partition the components into two groups to fit the proof rule. For example, if we have four components M_1 , M_2 , M_3 , and M_4 , we could partition them as $A = M_1 \parallel M_2$ and $B = M_3 \parallel M_4$. However, the number of possible partitions is exponential to the number of components. In addition, an investigation [10] showed that a good partition is very critical to AGR because it affects the verification performance significantly. In this work, we adopt the following heuristics:

1. If a component has timing parameters, it is collected in group A ;
2. If a component shares common action labels with the property, the component is collected in group A .

Other components are collected in group B .

Heuristics 1 is required for our approach to be sound. Concerning heuristics 2, in AGR, the ideal case is when A satisfies the property with the weakest assumption $\tilde{B}$ that allows everything, i. e., A itself is sufficient to prove the property no matter how B behaves. Based on this observation, the rationale behind heuristics 2 is that if a component shares common action labels with the property, it is very likely to be necessary to prove the property. We will show that heuristics 2 indeed yields good performance in practice.

5.2 Computing an abstraction via learning

Let us explain how to automatically generate $\tilde{B}$ by learning for non-parametric timed systems. We adopt the TL* algorithm [15], which is a learning algorithm to infer ERAs. The TL* algorithm has to interact with a teacher. The interaction

between them is shown in Fig. 2b. Notice that only the teacher knows about the ERA (say U) to be learned. During the learning process, the TL^* algorithm makes two types of queries: membership and candidate queries.

A *membership query* asks whether a word is accepted by U . After several membership queries, TL^* constructs a candidate ERA C , and makes a candidate query for it. A *candidate query* asks whether an ERA accepts the same timed language as U . If the teacher answers “yes”, then the learning process is finished, and C is the ERA learned by TL^* . If the candidate C accepts more (or less) timed words than U , the teacher answers “no” with a counterexample run ρ . TL^* will refine the candidate ERA based on the counterexamples provided by the teacher until the answer to the candidate query is “yes”. See [15] for details.

The two condition checkings in Fig. 3 ($A \parallel C \models \varphi$ and $B \models C$) can be done by model checking, and counterexamples given by model checking can also serve as counterexamples to the TL^* algorithm. Fig. 3 shows our overall procedure $LearnAbstr(B, A, \varphi)$ that returns either an assumption (denoted by $Abstraction(\tilde{B})$) when it is proved that $A \parallel B \models \varphi$ holds, or a counterexample (denoted by $Counterex(\tau)$) otherwise. $Counterex$ and $Abstraction$ are “tags” containing a value, in the spirit of data exchanged in distributed programming or types in functional programming; these tags will be used later on to differentiate between the two kinds of results output by $LearnAbstr$. Also note that, in our setting, we need a counterexample in the form of a trace τ , which is why $LearnAbstr$ returns $Counterex(trace(\rho))$.

Lemma 2. *Let A, B be two ERAs. Assume $LearnAbstr(B, A, \varphi)$ terminates with result $Abstraction(\tilde{B})$. Then $A \parallel \tilde{B} \models \varphi$ and $A \parallel B \models \varphi$.*

Proof. $Abstraction(\tilde{B})$ is returned only if $A \parallel \tilde{B} \models \varphi$ and $B \models \tilde{B}$. Thus, $A \parallel \tilde{B} \models \varphi$ holds. In addition, according to Fig. 2a, we can conclude that $A \parallel B \models \varphi$.

5.3 Replaying a trace

In this section, we explain how to synthesize the exact set of parameter valuations for which a finite trace belongs to the trace set.

Fig. 3: $LearnAbstr(B, A, \varphi)$

Algorithm 1: $\text{ReplayTrace}(\mathbf{A}, \tau)$

input : PTA $\mathbf{A}$, finite trace $\tau = e_0, e_1, \dots, e_{n-1}$
output : Constraint over the parameters

1 $\mathbf{s} = \mathbf{s}_0$
2 **for** $i = 0$ **to** $n - 1$ **do** $\mathbf{s} \leftarrow \text{Succ}(\mathbf{s}, e_i)$;
3 **return** $\mathbf{s} \downarrow_P$

Replaying a trace is close to two undecidable problems for PTAs: *i*) the reachability of a location is undecidable for PTAs [2], and therefore this result trivially extends to the reachability of a single edge; *ii*) the emptiness of the set of valuations for which the set of untimed words is the same as a given valuation is undecidable for PTAs [8] (where a proof is provided even for a *unique* untimed word). Nevertheless, computing the set of parameter valuations for which a given *finite* trace belongs to the trace set can be done easily by exploring a small part of the symbolic state space as follows.

We give our procedure $\text{ReplayTrace}(\mathbf{A}, \tau)$ in Algorithm 1. Basically, ReplayTrace computes the symbolic run equivalent to τ , and returns the projection onto P of the last symbolic state of that run. The correctness of ReplayTrace comes from the following results (proved in, e. g., [13]):

Lemma 3. *Let $\mathbf{A}$ be a PTA, and let ρ be a run of $\mathbf{A}$ reaching (l, C) . Let v be a parameter valuation. There exists an equivalent run in $v(\mathbf{A})$ iff $v \models C \downarrow_P$.*

Proof. From [13, Propositions 3.17 and 3.18].

Lemma 4. *Let $\mathbf{A}$ be a PTA, let v be a parameter valuation. Let ρ be a run of $v(\mathbf{A})$ reaching (l, μ) . Then there exists an equivalent symbolic run in $\mathbf{A}$ reaching (l, C) , with $v \models C \downarrow_P$.*

Proof. From [13, Proposition 3.18].

Proposition 1. *Let $\mathbf{A}$ be a PTA, let τ a trace of $v_0(\mathbf{A})$ for some v_0 . Let $K = \text{ReplayTrace}(\mathbf{A}, \tau)$. Then, for all v , τ is a trace of $v(\mathbf{A})$ iff $v \models K$.*

Proof. τ is a trace of $v_0(\mathbf{A})$ for some v_0 , and therefore it corresponds to some run ρ of $v_0(\mathbf{A})$. Then from Lemma 4 there exists an equivalent symbolic run in $\mathbf{A}$ reaching (l, C) , with $v_0 \models C \downarrow_P$. Now, from Lemma 3, for all v , there exists an equivalent run in $v(\mathbf{A})$ iff $v \models C \downarrow_P$. As $\text{ReplayTrace}(\mathbf{A}, \tau)$ returns exactly $K = C \downarrow_P$ therefore τ is a trace of $v(\mathbf{A})$ iff $v \models K$.

5.4 Exploiting the abstraction and performing parameter synthesis

We give our procedure in Algorithm 2: it takes as arguments a set of PERA components $\mathbf{A}$, a set of ERA components $\mathbf{B}$, a bounded parameter domain D_0 and a set of locations to be avoided. We maintain a safe non-convex parameter constraint K_{good} and an unsafe non-convex parameter constraint K_{bad} , both

Algorithm 2: CompSynth($A, B, D_0, L^\odot$)

input : PERA A , ERA B , parameter domain D_0 , subset $L^\odot$ of locations
output : Good and bad constraint over the parameters

```

1  $K_{bad} \leftarrow \perp$ ;  $K_{good} \leftarrow \perp$ 
2 while  $D_0 \cap \mathbb{N} \cap (K_{bad} \cup K_{good}) \neq \emptyset$  do
3   Pick  $v$  in  $D_0 \cap \mathbb{N} \cap (K_{bad} \cup K_{good})$ 
4   switch LearnAbstr( $B, v(A), AG \neg L^\odot$ ) do
5     case Abstraction( $\tilde{B}$ )
6        $K_{good} \leftarrow K_{good} \cup \text{PRP}(A \parallel \tilde{B}, v, L^\odot)$ 
7     case Counterex( $\tau$ )
8        $K_{bad} \leftarrow K_{bad} \cup \text{ReplayTrace}(A \parallel B, \tau)$ 
9 return ( $K_{good}, K_{bad}$ )

```

initially containing no valuations (line 1). Then CompSynth iterates on integer points: while not all integer points in D_0 are covered, i.e., do not belong to $K_{bad} \cup K_{good}$ (line 2), such an uncovered point v is picked (line 3). Then, we try to learn an abstraction of B w.r.t. $v(A)$ (line 5) so that $L^\odot$ is unreachable (“ $AG \neg L^\odot$ ” stands for “no run should ever reach $L^\odot$ ”). If an abstraction is successfully learned, then PRP is called on v and the abstract model $A \parallel \tilde{B}$ (line 6); the constraint K_{good} is then refined. Note that K_{good} is refined because, if an abstraction is computed, then necessarily the property is satisfied and therefore the (abstract) system is safe. Alternatively, if LearnAbstr fails to compute a valid abstraction, then a counterexample trace τ is returned (line 7); then this trace is replayed using ReplayTrace (line 8), and the constraint K_{bad} is updated.

5.5 Soundness

Proposition 2 (soundness). *Let $A \parallel B$ be a PERA and D_0 be a bounded parameter domain. Assume CompSynth($A, B, D_0, L^\odot$) terminates with result (K_{good}, K_{bad}) .*

Then, for all v i) if $v \models K_{good}$ then $v(A \parallel B)$ does not reach $L^\odot$; ii) if $v \models K_{bad}$ then $v(A \parallel B)$ reaches $L^\odot$.

Proof. i) Assume $v \models K_{good}$. From Algorithm 2, K_{good} is a finite union of convex constraints, each of them being the result of a call to PRP. Necessarily, $v \models K$, where K is one of these convex constraints, resulting from a call to $(A \parallel \tilde{B}, v')$, for some v' . From Lemma 2, $v'(A) \parallel \tilde{B} \models (AG \neg L^\odot)$. Since B and $\tilde{B}$ are non-parametric, we can write $v'(A \parallel \tilde{B}) \models (AG \neg L^\odot)$, i.e., $v'(A \parallel \tilde{B})$ does not reach $L^\odot$. From Lemma 1, for all $v'' \models K$, $v''(A \parallel \tilde{B})$ does not reach $L^\odot$. Now, since $\tilde{B}$ is a valid abstraction of B (i.e., $B \models \tilde{B}$), therefore $\tilde{B}$ contains more behaviors than B . Therefore for all $v'' \models K$, $v''(A \parallel B)$ does not reach $L^\odot$ either. Since $v \models K$, therefore $v(A \parallel B)$ does not reach $L^\odot$.

ii) Assume $v \models K_{bad}$. From Algorithm 2, K_{bad} is a finite union of convex constraints, each of them being the result of a call to ReplayTrace. Necessarily,

$v \models K$, where K is one of these convex constraints, resulting from a call to `ReplayTrace(A || B,  $\tau$ )` for some trace τ reaching $L^\ominus$. This trace was generated by `LearnAbstr` for some v' and is a valid counter-example, i. e., this trace τ reaches $L^\ominus$ in $v'(A) || B$. From [Lemma 4](#), this trace is also a trace reaching $L^\ominus$ in $A || B$. Then, from [Proposition 1](#), for all $v'' \models K$, τ is a valid trace of $v''(A || B)$ which reaches $L^\ominus$ and therefore $v''(A || B)$ reaches $L^\ominus$. Since $v \models K$, then $v(A || B)$ reaches $L^\ominus$.

Proposition 3 (integer-completeness). *Let A be a PERA and D_0 be a bounded parameter domain. Assume `CompSynth`($A, B, D_0, L^\ominus$) terminates with result (K_{good}, K_{bad}) . Then, for all $v \in D_0 \cap \mathbb{N}$, $v \in K_{good} \cup K_{bad}$.*

Proof. From [Algorithm 2](#) (line 2).

Remark 1. Note that the integerness can be scaled down to, e. g., multiples of 0.1, or in fact arbitrarily small numbers. The time needed to perform the verification might grow, but the coverage of all these discrete points is still guaranteed.

6 Experiments

6.1 Handling general PTAs

So far, we showed that our framework is sound for PERAs. We now show that, since we address only reachability, any PTA can be transformed into an equivalent PERA, and therefore our framework is much more general. The idea is that, since we are interested in reachability properties, we can rename some of the actions so that the PTA becomes a PERA.

Basically, we remove any action labels along the edges, and we add them back as follows: 1) if clock x is reset along an edge, the action label will be a_x ; 2) if no clock is reset along an edge, the action label will be na , where na is a (unique) label, the clock associated to which (say x_{na}) is never used (in guards and invariants) in the PERA; note that, by definition, x_{na} is reset along each edge labeled with na (although this has no impact in the PERA); 3) if more than one clock is reset along the edge, we split the edge into 2 consecutive edges in 0-time, where each clock is reset after the other, following the mechanism described above. Note that the 0-time can be ensured using an invariant $x \leq 0$, where x is the first clock to be reset.

Basically, our transformation leaves the structure of the PTA unchanged (with the exception of a few transitions in 0-time to simulate multiple simultaneous clock resets). For each parameter valuation, the resulting PERA has the same timed language as the original PTA – up to action renaming and with the introduction of some 0-time transitions (that could be considered as silent transitions if the language really mattered). Therefore, reachability is preserved.

Note that this construction provides an alternative proof for [Theorem 1](#).

Example 2. [Fig. 4a](#) shows a PTA, and [Fig. 4b](#) its translation into an equivalent PERA. (Recall that clock resets are implicit in PERAs.)

Fig. 4: General PTA and its translation to a PERA

Remark 2. In our benchmarks, although we only address reachability, action labels are not entirely useless: they are often used for action synchronization between components. Therefore, renaming all actions is not a valid transformation, as components may not synchronize anymore the way it was expected. In fact, we ensured that our models either only work using interleaving (no action synchronization) or, when various components of a PTA synchronize on an action label, at most one clock is reset along that transition for all PTAs synchronizing on this action label.

6.2 Experiments

We implemented our method in a toolkit made of the following components:

- IMITATOR [5] is a state of the art tool for verifying real-time systems modeled by an extension of PTAs with stopwatches, broadcast synchronization and integer-valued shared variables. IMITATOR is implemented in OCaml, and the polyhedra operations rely on the Parma Polyhedra Library (PPL).
- CV (Compositional Verifier) is a prototype implementation (in C++) of the proposed learning-based compositional verification framework for ERAs.

The architecture of our toolkit is shown in Fig. 5. The leading tool is IMITATOR, that takes the input model (in the IMITATOR input format), and eventually outputs the result. IMITATOR implements both algorithms CompSynth and ReplayTrace, while CV implements LearnAbstr. The interface between both tools is handled by a Python script, that is responsible for retrieving the abstraction of B computed by CV and re-parameterizing the components A. We used IMITATOR 2.9-alpha1, build 2212.³ Experiments were run on a MacBook Pro with an i7 CPU 2.67GHz and 3,7 GiB memory running Kubuntu 14.04 64 bits.

Benchmarks We evaluated our approach using several benchmarks, with various (reachability) properties. We give in Table 1 the case studies, with the numbers of PERAs in parallel, of clocks (equal to the number of actions, by definition) and of parameters, followed by the specification number; then, we

³ Sources, binaries, models and results are available at imitator.fr/static/FORTE17

Fig. 5: Architecture of our toolkit

Case study	#A	#X	#P	Spec	EFsynth	PRPC		CompSynth			
						#iter	total	#abs	#c.-ex.	learning	total
FMS-1	6	18	2	1	0.299	2	0.654	1	1	0.074	0.136
				2	0.010	1	0.372	0	1	0.038	0.046
				3	0.282	1	0.309	1	0	0.090	0.242
FMS-2	11	37	2	1	T.O.	-	T.O.	1	1	84.2	88.9
				2	T.O.	-	T.O.	1	0	81.4	85.2
				3	0.051	-	T.O.	0	2	1.10	2.44
				4	0.062	-	T.O.	0	1	1.42	1.53
				5	T.O.	-	T.O.	1	0	31.4	40.8
				6	T.O.	-	T.O.	1	0	37.2	42.4
AIP	11	46	2	1	0.551	-	T.O.	0	1	0.086	0.114
				2	2.11	-	T.O.	0	1	1.22	1.25
				3	3.91	-	T.O.	0	1	8.50	8.54
				4	0.235	-	T.O.	1	1	8.39	8.42
				5	T.O.	-	T.O.	1	0	0.394	0.871
				6	T.O.	-	T.O.	1	0	5.32	9.58
				7	T.O.	-	T.O.	1	0	1.76	3.19
				8	T.O.	-	T.O.	1	0	1.13	4.35
				9	T.O.	-	T.O.	1	1	0.762	1.84
				10	0.022	-	T.O.	0	1	0.072	0.094
Fischer-3	5	12	2		2.76	4	14.0	0	1	-	T.O.
Fischer-4	6	16	2		T.O.	-	T.O.	0	1	-	T.O.

Table 1: Experiments: comparison between algorithms

compare the computation time (in s) for EFsynth, PRPC, and CompSynth (for which we also give the number of abstractions and counter-examples generated by LearnAbstr, and the learning time required by LearnAbstr). “T.O.” denotes a timeout ($> 600s$). FMS-1 and -2 are two versions of a flexible manufacturing system [15] (Fig. 1 depicts the conveyor component of FMS-1). AIP is a manufacturing system producing two products from two different materials [15]. Fischer-3 (resp. 4) is a PERA version of the mutual exclusion protocol with 3 (resp. 4) processes; it was obtained using the transformation in Section 6.1.

Comparison Although reachability synthesis is intractable for PERAs (Theorem 1), CompSynth always terminates for our case studies (except for Fischer, for which the abstraction computation is too slow). In contrast, EFsynth does often not terminate. In addition, CompSynth always gives a complete (dense) result not only within D_0 but in fact in the entire parameter domain ($\mathbb{Q}_+^M$).

First, CompSynth outperforms PRPC for all but one benchmark: this suggest to use CompSynth instead of PRPC in the future.

Second, CompSynth is faster than EFsynth in 13/20 cases. In addition, whereas EFsynth often does not terminate, CompSynth always outputs a result (except for Fischer). In some cases (FMS-2:3, FMS-2:4, AIP:4), EFsynth is much faster because it immediately derives $\perp$, whereas CompSynth has to compute the ab-

Case study	#A	#X	#P	Spec	D_0	CompSynth				
						#abs	#c.-ex.	find next point	learning	total
FMS-2	11	37	2	1	2,500	1	1	0.0	81.0	85.7
					10,000	1	1	0.1	82.5	87.3
					250,000	1	1	2.2	82.0	89.0
					1,000,000	1	1	8.9	83.1	96.7
					25,000,000	1	1	221.2	83.1	309.0
					100,000,000	1	1	888.1	83.5	976.4

Table 2: Experiments: scalability w.r.t. the reference domain

straction first. Even in these unfavorable cases, **CompSynth** is never much behind **EFsynth**: the worst case is AIP:4, with 8 seconds slower. This suggests that **CompSynth** may be preferred to **EFsynth** for PERAs benchmarks.

Interestingly, in almost all benchmarks, at most one abstraction (for good valuations) and one counter-example (for bad valuations) is necessary for **CompSynth**. In addition, most of the computation time of **CompSynth** (71 % in average) comes from **LearnAbstr**; this suggests to concentrate our future optimization efforts on this part. Perhaps an on-the-fly composition mixed with synthesis could help speeding-up this part; this would also solve the issue of constraints $\perp$ synthesized only after the abstraction phase is completed (FMS-2:3, FMS-2:4, AIP:4).

For Fischer, our algorithm is very inefficient: this comes from the fact that the model is strongly synchronized, and the abstraction computation does not terminate within 600s. In fact, in both cases, **LearnAbstr** successfully derives very quickly a counter-example that is used by **CompSynth** to immediately synthesize all “bad” valuations; but then, as **LearnAbstr** fails in computing an abstraction, the good valuations are not synthesized. Improving the learning phase for strongly synchronized models is on our agenda.

We were not able to perform a comparison with [9]; the prototype of [9] always failed to compute a result. In addition, our Fischer benchmark does not fit in [9] as Fischer makes use of shared parameters.

Size of the parameter domain [Algorithm 2](#) is based on an enumeration of integer points: although we could use an SMT solver to find the next uncovered point, in our implementation we just enumerate *all* points, and therefore the size of D_0 may have an impact on the efficiency of **CompSynth**. [Table 2](#) shows the impact of the size of D_0 w.r.t. **CompSynth**. “find next point” is the time to find the next uncovered point (and therefore includes the enumeration of all points). The overhead is reasonable up to 1,000,000 points, but then becomes very significant. Two directions can be taken to overcome this problem for very large parameter domains: 1) using an SMT solver to find the next uncovered point; or 2) using an on-the-fly refinement of the precision (e.g., start with multiples of 100, then 10 for uncovered subparts of D_0 , then 1... until $D_0 \subseteq K_{bad} \cup K_{good}$).

Partitioning Finally, although the use of heuristic [2](#) is natural, we still wished to evaluate it. Results show that our partitioning heuristic yields always the best execution time, or almost the best execution time.

7 Conclusion and perspectives

We proposed a learning-based approach to improve the verification of parametric distributed timed systems, that turns to be globally efficient on a set of benchmarks; most importantly, it outputs an exact result for most cases where the monolithic procedure `EFsynth` fails.

Among the limitations of our work is that the input model must be a PERA (although we provide an extension to PTAs), and that all parametric ERAs must be in the same component `A`. How to lift these assumptions is on our agenda.

Another perspective is the theoretical study of PERAs, i. e., their expressiveness and decidability (beyond EF-emptiness, that we proved to be undecidable).

Finally, addressing other properties than reachability is also on our agenda.

Acknowledgment We warmly thank Lăcrămioara Aștefănoaei for her appreciated help with installing and using the prototype tool of [9].

References

1. Alur, R., Fix, L., Henzinger, T.A.: Event-clock automata: A determinizable class of timed automata. *Theoretical Computer Science* 211(1-2), 253–273 (1999)
2. Alur, R., Henzinger, T.A., Vardi, M.Y.: Parametric real-time reasoning. In: *STOC*. pp. 592–601. ACM (1993)
3. André, É.: What’s decidable about parametric timed automata? In: *FTSCS. CCIS*, vol. 596, pp. 1–17. Springer (2015)
4. André, É., Chatain, T., Encrenaz, E., Fribourg, L.: An inverse method for parametric timed automata. *IJFCS* 20(5), 819–836 (2009)
5. André, É., Fribourg, L., Kühne, U., Soulat, R.: *IMITATOR 2.5: A tool for analyzing robustness in scheduling problems*. In: *FM. LNCS*, vol. 7436. Springer (2012)
6. André, É., Lime, D., Roux, O.H.: Decision problems for parametric timed automata. In: *ICFEM. LNCS*, vol. 10009, pp. 400–416. Springer (2016)
7. André, É., Lipari, G., Nguyen, H.G., Sun, Y.: Reachability preservation based parameter synthesis for timed automata. In: *NFM. LNCS*, vol. 9058, pp. 50–65. Springer (2015)
8. André, É., Markey, N.: Language preservation problems in parametric timed automata. In: *FORMATS. LNCS*, vol. 9268, pp. 27–43. Springer (2015)
9. Aștefănoaei, L., Bensalem, S., Bozga, M., Cheng, C., Ruess, H.: Compositional parameter synthesis. In: *FM. LNCS*, vol. 9995, pp. 60–68 (2016)
10. Cobleigh, J.M., Avrunin, G.S., Clarke, L.A.: Breaking up is hard to do: An evaluation of automated assume-guarantee reasoning. *TOSEM* 17(2), 7:1–7:52 (2008)
11. Cobleigh, J.M., Giannakopoulou, D., Pasareanu, C.S.: Learning assumptions for compositional verification. In: *TACAS. LNCS*, vol. 2619, pp. 331–346 (2003)
12. Frehse, G., Jha, S.K., Krogh, B.H.: A counterexample-guided approach to parameter synthesis for linear hybrid automata. In: *HSCC. LNCS*, vol. 4981 (2008)
13. Hune, T., Romijn, J., Stoelinga, M., Vaandrager, F.W.: Linear parametric model checking of timed automata. *JLAP* 52-53, 183–220 (2002)
14. Jovanović, A., Lime, D., Roux, O.H.: Integer parameter synthesis for timed automata. *Transactions on Software Engineering* 41(5), 445–461 (2015)
15. Lin, S.W., André, É., Liu, Y., Sun, J., Dong, J.S.: Learning assumptions for compositional verification of timed systems. *TSE* 40(2), 137–153 (2014)