

LTL Parameter Synthesis of Parametric Timed Automata

Peter Bezděk^(✉), Nikola Beneš, Jiří Barnat, and Ivana Černá

Faculty of Informatics, Masaryk University, Brno, Czech Republic
bezdek@mail.muni.cz, {xbenes3,barnat,cerna}@fi.muni.cz

Abstract. The parameter synthesis problem for parametric timed automata is undecidable in general even for very simple reachability properties. In this paper we introduce restrictions on parameter valuations under which the parameter synthesis problem is decidable for LTL properties. The investigated bounded integer parameter synthesis problem could be solved using an explicit enumeration of all possible parameter valuations. We propose an alternative symbolic zone-based method for this problem which results in a faster computation. Our technique extends the ideas of the automata-based approach to LTL model checking of timed automata. To justify the usefulness of our approach, we provide experimental evaluation and compare our method with explicit enumeration technique.

1 Introduction

Model checking [1] is a formal verification technique applied to check for logical correctness of discrete distributed systems. While it is often used to prove the unreachability of a bad state (such as an assertion violation in a piece of code), with a proper specification formalism, such as the *Linear Temporal Logic* (LTL), it can also check for many interesting liveness properties of systems, such as repeated guaranteed response, eventual stability, live-lock, etc.

Timed automata have been introduced in [2] and have emerged as a useful formalism for modelling time-critical systems as found in many embedded and cyber-physical systems. The formalism is built on top of the standard finite automata enriched with a set of real-time clocks and allowing the system actions to be guarded with respect to the clock valuations. In the general case, such a timed system exhibits infinite-state semantics (the clock domains are continuous). Nevertheless, when the guards are limited to comparing clock values with integers only, there exists a bisimilar finite state representation of the original infinite-state real-time system referred to as the region abstraction. A practically efficient abstraction of the infinite-state space came with the so called zones [3]. The zone-based abstraction is much coarser and the number of zones *reachable*

N. Beneš—The author has been supported by the Czech Science Foundation grant no. GA15-11089S.

from the initial state is significantly smaller. This in turns allows for an efficient implementation of verification tools for timed automata, see e.g. UPPAAL [4].

Very often the correctness of a time-critical system relates to a proper timing, i.e. it does not only depend on the logical result of the computation, but also on the time at which the results are produced. To that end the designers are not only in the need of tools to verify correctness once the system is fully designed, but also in the need of tools that would help them derive proper time parameters of individual system actions that would make the system as a whole satisfy the required specification. After all this problem of *parameter synthesis* is more urgent in practice than the verification as such.

Related Work. The problem of the existence of a parameter valuation for a reachability property of a parametric timed automaton in continuous time has been shown to be undecidable in [5,6] for a parametric timed automaton with as few as 3 clocks. This problem remains undecidable even for integer-valued parameters [7]. A solution for the parameter synthesis problem and reachability properties is presented in [8] where the authors provide a semi-decision algorithm which is not guaranteed to terminate in all cases. Authors also introduce a subclass of parametric timed automata, called L/U automata for which the emptiness problem is decidable. Decidability results for the class of L/U automata are further extended in [9]. In particular, the authors show that emptiness, finiteness and universality problems of the set of parameter valuations for which there is an infinite accepting run are decidable.

To obtain a decidable version of parameter synthesis problem for parametric timed automata we need to restrict parameter valuations to bounded integers. When modelling a real-time system, designers can usually provide practical bounds on time parameters of individual system actions. Therefore, introducing a parameter synthesis method with such a restriction is still reasonable. In [10] the authors show that the problem of existence of bounded integer parameter value such that a given property is satisfied is PSPACE-complete for a significant number of properties, which include Timed Computational Tree Logic. They give symbolic algorithms only for reachability and unavailability properties.

Contribution. The main contribution of this paper is a symbolic method that solves the parameter synthesis problem for specifications given in the Linear Time Logic (LTL) and parametric timed automata with bounded integer parameters. To this end, we introduce a finite abstraction of parametric timed automata with bounded integer parameters and provide an algorithm working over this abstraction. To evaluate our technique we implemented both a symbolic approach and explicit enumeration technique in a proof-of-concept tool and compare the techniques on a case study. The finite abstraction does not provide a unique representation of states and therefore we design an efficient state storage mechanism that deals with this problem. The experiments demonstrate the strength of the symbolic approach which may be faster by an order of magnitude.

Outline. The rest of the paper is organised as follows. The problem definition is given in Sect. 2 that also introduces the basic notions. We then define the symbolic semantics of a parametric timed Büchi automaton and its finite abstraction in Sect. 3. Section 4 describes the parameter synthesis algorithm itself. Section 5 describes the implementation and used heuristics. Then, in Sect. 6 we experimentally evaluate the proposed algorithm and compare it with explicit enumeration. Finally, Sect. 7 concludes the paper.

2 Preliminaries and Problem Statement

In order to state our main problem formally, we need to describe the notion of a parametric timed automaton. We start by describing some basic notation.

Let P be a finite set of *parameters*. An *affine expression* is an expression of the form $z_0 + z_1p_1 + \dots + z_np_n$, where $p_1, \dots, p_n \in P$ and $z_0, \dots, z_n \in \mathbb{Z}$. We use $E(P)$ to denote the set of all affine expressions over P . A *parameter valuation* is a function $v : P \rightarrow \mathbb{Z}$ which assigns an integer number to each parameter. Let $lb : P \rightarrow \mathbb{Z}$ be a lower bound function and $ub : P \rightarrow \mathbb{Z}$ be an upper bound function. For an affine expression e , we use $e[v]$ to denote the integer value obtained by replacing each p in e by $v(p)$. We use $max_{lb,ub}(e)$ to denote the maximal value obtained by replacing each p with a positive coefficient in e by $ub(p)$ and replacing each p with a negative coefficient in e by $lb(p)$. We say that the parameter valuation v respects lb and ub if for each $p \in P$ it holds that $lb(p) \leq v(p) \leq ub(p)$. We denote the set of all parameter valuations respecting lb and ub by $Val_{lb,ub}(P)$. In the following, we only consider parameter valuations from $Val_{lb,ub}(P)$.

Let X be a finite set of *clocks*. We assume the existence of a special *zero clock*, denoted by x_0 , that has always the value 0. A *guard* is a finite conjunction of expressions of the form $x_i - x_j \sim e$ where $x_i, x_j \in X$, $e \in E(P)$ and $\sim \in \{\leq, <\}$. We use $G(X, P)$ to denote the set of all guards over a set of clocks X and a set of parameters P . A *simple guard* is a guard containing only expressions of the form $x_i - x_j \sim e$ where $x_i, x_j \in X$, $e \in E(P)$, $\sim \in \{\leq, <\}$, and $x_i = x_0$ or $x_j = x_0$. We also use $\bar{G}(X, P)$ to denote the set of all simple guards over a set of clocks X and a set of parameters P . A *clock valuation* is a function $\eta : X \rightarrow \mathbb{R}_{\geq 0}$ assigning non-negative real numbers to each clock such that $\eta(x_0) = 0$. We denote the set of all clock valuations by $Val(X)$. Let $g \in G(X, P)$ and v be a parameter valuation and η be a clock valuation. Then $g[v, \eta]$ denotes a boolean value obtained from g by replacing each parameter p with $v(p)$ and each clock x with $\eta(x)$. A pair (v, η) *satisfies* a guard g , denoted by $(v, \eta) \models g$, if $g[v, \eta]$ evaluates to true. The *semantics* of a guard g , denoted by $\llbracket g \rrbracket$, is a set of all valuation pairs (v, η) such that $(v, \eta) \models g$. For a given parameter valuation v we write $\llbracket g \rrbracket_v$ for the set of clock valuations $\{\eta \mid (v, \eta) \models g\}$.

We define two operations on clock valuations. Let η be a clock valuation, d a non-negative real number and $R \subseteq X$ a set of clocks. We use $\eta + d$ to denote the clock valuation that adds the delay d to each clock, i.e. $(\eta + d)(x) = \eta(x) + d$ for all $x \in X \setminus \{x_0\}$. We further use $\eta\langle R \rangle$ to denote the clock valuation that resets clocks from the set R , i.e. $\eta\langle R \rangle(x) = 0$ if $x \in R$, $\eta\langle R \rangle(x) = \eta(x)$ otherwise.

Definition 2.1 (PTA). A parametric timed automaton (PTA) is a tuple $M = (L, l_0, X, P, \Delta, Inv)$ where

- L is a finite set of locations,
- $l_0 \in L$ is the initial location,
- X is a finite set of clocks,
- P is a finite set of parameters,
- $\Delta \subseteq L \times \overline{G}(X, P) \times 2^X \times L$ is a finite transition relation, and
- $Inv : L \rightarrow \overline{G}(X, P)$ is an invariant function.

We use $q \xrightarrow{g, R} \Delta q'$ to denote $(q, g, R, q') \in \Delta$. The semantics of a PTA is given as a labelled transition system. A *labelled transition system* (LTS) over a set of symbols Σ is a triple (S, s_0, \rightarrow) , where S is a set of states, $s_0 \in S$ is an initial state and $\rightarrow \subseteq S \times \Sigma \times S$ is a transition relation. We use $s \xrightarrow{a} s'$ to denote $(s, a, s') \in \rightarrow$.

Definition 2.2 (PTA semantics). Let $M = (L, l_0, X, P, \Delta, Inv)$ be a PTA and v be a parameter valuation. The semantics of M under v , denoted by $\llbracket M \rrbracket_v$, is an LTS $(\mathbb{S}_M, s_0, \rightarrow)$ over the set of symbols $\{act\} \cup \mathbb{R}_{\geq 0}$, where

- $\mathbb{S}_M = L \times Val(X)$ is a set of all states,
- $s_0 = (l_0, \mathbf{0})$, where $\mathbf{0}$ is a clock valuation with $\mathbf{0}(x) = 0$ for all x , and
- the transition relation \rightarrow is specified for all $(l, \eta), (l', \eta') \in \mathbb{S}_M$ as follows:
 - $(l, \eta) \xrightarrow{d} (l', \eta')$ if $l = l'$, $d \in \mathbb{R}_{\geq 0}$, $\eta' = \eta + d$, and $(v, \eta') \models Inv(l')$,
 - $(l, \eta) \xrightarrow{act} (l', \eta')$ if $\exists g, R : l \xrightarrow{g, R} \Delta l'$, $(v, \eta) \models g$, $\eta' = \eta \langle R \rangle$, and $(v, \eta') \models Inv(l')$.

The transitions of the first kind are called delay transitions, the latter are called action transitions.

We write $s_1 \xrightarrow{act}_d s_2$ if there exists $s' \in \mathbb{S}_M$ and $d \in \mathbb{R}_{\geq 0}$ such that $s_1 \xrightarrow{act} s' \xrightarrow{d} s_2$. A proper run π of $\llbracket M \rrbracket_v$ is an infinite alternating sequence of delay and action transitions that begins with a delay transition $\pi = (l_0, \eta_0) \xrightarrow{d_0} (l_0, \eta_0 + d_0) \xrightarrow{act} (l_1, \eta_1) \xrightarrow{d_1} \dots$. A proper run is called *Zeno* if the sum of all its delays is finite.

Let M be a PTA, $\mathcal{L} : L \rightarrow 2^{Ap}$ be a labelling function that assigns a set of atomic propositions to each location of M , v be a parameter valuation, and φ be an LTL formula. We say that M under v with \mathcal{L} satisfies φ , denoted by $(M, v, \mathcal{L}) \models \varphi$ if for all proper runs π of $\llbracket M \rrbracket_v$, π satisfies φ where atomic propositions are determined by \mathcal{L} .

Given a parametric timed automaton M , a labelling function \mathcal{L} , and an LTL property φ , the *parameter synthesis problem* is to compute the set of all parameter valuations v such that $(M, v, \mathcal{L}) \models \varphi$. Unfortunately, it is known that the parameter synthesis problem for a PTA is undecidable even for very simple (reachability) properties [5]. Instead of solving the general problem, we thus focus on a more constrained version which is still reasonable for practical purposes.

Problem Formulation. Given an LTL property φ , a parametric timed automaton $M = (L, l_0, X, P, \Delta, Inv)$, a labelling function \mathcal{L} , a lower bound function lb and an upper bound function ub , the *bounded integer parameter synthesis problem* is to compute the set of all parameter valuations v such that $(M, v, \mathcal{L}) \models \varphi$ and $lb(p) \leq v(p) \leq ub(p)$ for each $p \in P$.

This problem is trivially decidable using the standard zone-based abstraction and explicit enumeration of all parameter valuations. In order to avoid the necessity of the explicit enumeration of all parameter valuations we use a combination of the zone-based abstraction and a symbolic representation of parameter valuation sets. Our algorithmic framework which solves this problem consists of three steps.

As the first step, we apply the standard automata-based LTL model checking of timed automata [2] to parametric timed automata. We employ this approach in the following way. From a PTA M and an LTL formula φ we produce a product parametric timed Büchi automaton (PTBA) A . The accepting runs of the automaton A correspond to the runs of M violating the formula φ .

As the second step, we employ a symbolic semantics of a PTBA A with a suitable extrapolation. From the symbolic state space of a PTBA A we finally produce a Büchi automaton B in which each state is associated symbolic information about parameter valuations. This transformation is described in Sect. 3.

As the last step, we need to detect all parameter valuations for which there exists an accepting run in Büchi automaton B . To that end, we employ a new algorithm, which we call the Cumulative NDFS. The algorithm is described in detail in Sect. 4.

We now proceed with the definitions of a Büchi automaton, a parametric timed Büchi automaton and its semantics.

Definition 2.3 (BA). A Büchi automaton (BA) is a tuple $B = (Q, q_0, \Sigma, \rightarrow, F)$, where Q is a finite set of states, $q_0 \in Q$ is an initial state, Σ is a finite set of symbols, $\rightarrow \subseteq Q \times \Sigma \times Q$ is a set of transitions, and $F \subseteq Q$ is a set of accepting states (acceptance condition). An ω -word $w = a_0a_1a_2 \dots \in \Sigma^\omega$ is accepting if there is an infinite sequence of states $q_0q_1q_2 \dots$ such that $q_i \xrightarrow{a_i} q_{i+1}$ for all $i \in \mathbb{N}$, and there exist infinitely many $i \in \mathbb{N}$ such that $q_i \in F$.

Definition 2.4 (PTBA). A parametric timed Büchi automaton (PTBA) is a pair $A = (M, F)$ where $M = (L, l_0, X, P, \Delta, Inv)$ is a PTA, and $F \subseteq L$ is a set of accepting locations.

Zeno runs represent non-realistic behaviours and it is desirable to ignore them in analysis. Therefore, we are interested only in non-Zeno accepting runs of a PTBA. There is a syntactic transformation to the so-called strongly non-Zeno form [11] of a PTBA, which guarantees that each accepting run is non-Zeno. For the rest of the paper, we thus assume that there are no Zeno accepting runs in the PTBA.

Definition 2.5 (PTBA semantics). Let $A = (M, F)$ be a PTBA and v be a parameter valuation. The semantics of A under v , denoted by $\llbracket A \rrbracket_v$, is defined as $\llbracket M \rrbracket_v = (\mathbb{S}_M, s_0, \rightarrow)$.

We say a state $s = (l, \eta) \in \mathbb{S}_M$ is accepting if $l \in F$. A proper run $\pi = s_0 \xrightarrow{d_0} s'_0 \xrightarrow{act} s_1 \xrightarrow{d_1} s'_1 \xrightarrow{act} \dots$ of $\llbracket A \rrbracket_v$ is accepting if there exists an infinite set of indices i such that s_i is accepting.

3 Symbolic Semantics

In this section we show the construction of a finite system which represents the semantics of a given PTBA. First, we describe a parametric extension of the zone abstraction. This extension is based on constrained parametric difference bound matrices, described in [8]. However, this abstraction itself does not guarantee finiteness in our setting. To solve this problem we further introduce a finite parametric extrapolation.

3.1 Constrained Parametric Difference Bound Matrix

A *constraint* is an inequality of the form $e \sim e'$ where $e, e' \in E$ and $\sim \in \{>, \geq, \leq, <\}$. We define $c[v]$ as the boolean value obtained by replacing each p in c by $v(p)$. A valuation v satisfies a constraint c , denoted $v \models c$, if $c[v]$ evaluates to true. The *semantics* of a constraint c , denoted $\llbracket c \rrbracket$, is the set of all valuations that satisfy c . A finite set of constraints C is called a *constraint set*. A valuation satisfies a constraint set C if it satisfies each $c \in C$. The *semantics* of a constraint set C is given by $\llbracket C \rrbracket = \bigcap_{c \in C} \llbracket c \rrbracket$. A constraint set C is *satisfiable* if $\llbracket C \rrbracket \neq \emptyset$. A constraint c covers a constraint set C , denoted $C \models c$, if $\llbracket C \rrbracket \subseteq \llbracket c \rrbracket$.

As in [8], we identify the relation symbol \leq with the boolean value true and $<$ with the boolean value false. Then, we treat boolean connectives on relation symbols $\leq, <$ as operations with boolean values. For example, $(\leq \implies <) = <$.

We now define the parametric difference bound matrix, the constrained parametric difference bound matrix, several operations on them, and the symbolic semantics of a PTBA.

Definition 3.1. A parametric difference bound matrix (PDBM) over P and X is a set D which contains for all $0 \leq i, j \leq |X|$ a guard of the form $x_i - x_j \prec_{ij} e_{ij}$ where $x_i, x_j \in X$ and $e_{ij} \in E(P) \cup \{\infty\}$ and $i = j \implies e_{ii} = 0$. We denote by D_{ij} a guard of the form $x_i - x_j \prec_{ij} e_{ij}$ contained in D . Given a parameter valuation v , the semantics of D is given by $\llbracket D \rrbracket_v = \llbracket \bigwedge_{i,j} D_{ij} \rrbracket_v$. A PDBM D is satisfiable with respect to v if $\llbracket D \rrbracket_v$ is non-empty.

Definition 3.2. A constrained parametric difference bound matrix (CPDBM) is a pair (C, D) , where C is a constraint set and D is a PDBM and for each $0 \leq i \leq |X|$ it holds that $C \models e_{0i} \geq 0$. The semantics of (C, D) is given by $\llbracket C, D \rrbracket = \{(v, \eta) \mid v \in \llbracket C \rrbracket \wedge \eta \in \llbracket D \rrbracket_v\}$. We call (C, D) satisfiable if $\llbracket C, D \rrbracket$ is non-empty. A CPDBM (C, D) is said to be in the canonical form if and only if for all i, j, k , $C \models e_{ij} (\prec_{ik} \wedge \prec_{kj}) e_{ik} + e_{kj}$.

Resetting a Clock. Suppose (C, D) is a CPDBM in the canonical form. The reset of the clock x_r in (C, D) , denoted by $(C, D)\langle x_r \rangle$, is given as $(C, D\langle x_r \rangle)$ where:

$$D\langle x_r \rangle_{ij} = \begin{cases} D_{0j} & \text{if } i \neq j \text{ and } i = r, \\ D_{i0} & \text{if } i \neq j \text{ and } j = r, \\ D_{ij} & \text{else.} \end{cases}$$

We can again generalise this definition to a set of clocks:

$$(C, D)\langle x_{i_0}, x_{i_1}, \dots, x_{i_k} \rangle \stackrel{def}{\iff} (C, D)\langle x_{i_0} \rangle \langle x_{i_1} \rangle \dots \langle x_{i_k} \rangle.$$

Applying a Guard. Suppose g is a guard of the form $x_i - x_j \prec e$, (C, D) is a CPDBM in the canonical form and $D_{ij} = (e_{ij}, \prec_{ij})$. The application of the guard g on (C, D) generally results in a set of CPDBMs and is defined as follows:

$$(C, D)[g] = \begin{cases} \{(C, D[g])\} & \text{if } C \models \neg(e_{ij}(\prec_{ij} \implies \prec)e), \\ \{(C, D)\} & \text{if } C \models e_{ij}(\prec_{ij} \implies \prec)e, \\ \{(C \cup \{e_{ij}(\prec_{ij} \implies \prec)e\}, D)\}, & \text{otherwise,} \\ \{(C \cup \{\neg e_{ij}(\prec_{ij} \implies \prec)e\}, D[g])\}, & \end{cases}$$

where $D[g]$ is defined as follows:

$$D[g]_{kl} = \begin{cases} (e, \prec) & \text{if } k = i \text{ and } l = j, \\ D_{kl} & \text{else.} \end{cases}$$

We can generalise this definition to conjunctions of guards as follows:

$$D[g_{i_0} \wedge g_{i_1} \wedge \dots \wedge g_{i_k}] \stackrel{def}{\iff} D[g_{i_0}][g_{i_1}] \dots [g_{i_k}].$$

Time Successors. Suppose (C, D) is a CPDBM in the canonical form. The time successor of (C, D) , denoted by $(C, D)^\uparrow$, represents a CPDBM with all upper bounds on clocks removed and is given as (C, D^\uparrow) where:

$$D^\uparrow_{ij} = \begin{cases} (\infty, \prec) & \text{if } i \neq 0 \text{ and } j = 0, \\ D_{ij} & \text{else.} \end{cases}$$

The reset and time successor operations preserve the canonical form of a CPDBM. After the application of a guard the CPDBM may no longer be in the canonical form and thus a transformation to the canonical form needs to be performed. However, due to the presence of parameters the standard canonisation [12] process can be ambiguous. The canonisation procedure is therefore extended to cope with this ambiguity. As a consequence, the result of the canonisation is not a single CPDBM, but may generally be a set containing potentially more CPDBMs in the canonical form with mutually disjoint constraint sets.

To canonise the given CPDBM we need to derive the tightest constraint on each clock difference. Deriving the tightest constraint on a clock difference can be

seen as finding the shortest path in the graph interpretation of the CPDBM. In [8] the authors implement the canonisation using a nondeterministic extension of the Floyd-Warshall algorithm where on each relaxation a split into two different CPDBMs can occur.

Canonisation. First, we define a relation \longrightarrow_{FW} on constrained parametric bound matrices as follows, for all $0 \leq k, i, j \leq |X|$:

- $(k, i, j, C_1, D_1) \longrightarrow_{FW} (k, i, j + 1, C_2, D_2)$
if $(C_2, D_2) \in (C_1, D_1)[x_i - x_j (\prec_{ik} \wedge \prec_{kj}) e_{ik} + e_{kj}]$
- $(k, i, |X| + 1, C_1, D_1) \longrightarrow_{FW} (k, i + 1, 0, C_1, D_1)$
- $(k, |X| + 1, 0, C_1, D_1) \longrightarrow_{FW} (k + 1, 0, 0, C_1, D_1)$

The relation \longrightarrow_{FW} can be seen as a representation of the computation steps of the extended Floyd-Warshall algorithm.

Suppose now (C, D) is a CPDBM. The canonical set of (C, D) , denoted as $(C, D)_c$, represents a set of CPDBMs with the tightest constraint on each clock difference in D and is defined as follows:

$$(C, D)_c = \{(C', D') \mid (0, 0, 0, C, D) \longrightarrow_{FW}^* (|X| + 1, 0, 0, C', D')\}$$

Example 3.3. Let $x, y \in X$ and $p, q \in P$. For a CPDBM $(C, D) = (\emptyset, \{x \leq p, y \leq q, y \leq x, y \leq x\})$ we obtain by canonisation $(C, D)_c = \{(\{p \leq q\}, \{x \leq p, y \leq p, y \leq x, y \leq x\}), (\{q < p\}, \{x \leq q, y \leq q, y \leq x, y \leq x\})\}$.

Definition 3.4. (PTBA symbolic semantics). Let $A = ((L, l_0, X, P, \Delta, Inv), F)$ be a PTBA. Let lb and ub be a lower bound function and an upper bound function on parameters. The symbolic semantics of A with respect to lb and ub is a transition system $(\mathbb{S}_A, \mathbb{S}_{init}, \Longrightarrow)$, denoted as $\llbracket A \rrbracket_{lb, ub}$, where

- $\mathbb{S}_A = L \times \{\llbracket C, D \rrbracket \mid (C, D) \text{ is a CPDBM}\}$ is the set of all symbolic states,
- the set of initial states $\mathbb{S}_{init} = \{(l_0, \llbracket C, D \rrbracket) \mid (C, D) \in (\emptyset, E^\dagger)[Inv(l_0)]\}$, where
 - E is a PDBM with $E_{i,j} = (0, \leq)$ for each i, j , and
 - for each $p \in P$, the constraints $p \geq lb(p)$ and $p \leq ub(p)$ are in C .
- There is a transition $(l, \llbracket C, D \rrbracket) \Longrightarrow (l', \llbracket C', D'_c \rrbracket)$ if
 - $l \xrightarrow{g, R}_\Delta l'$ and
 - $(C'', D'') \in (C, D)[g]$ and
 - $(C''_c, D''_c) \in (C'', D'')_c$ and
 - $(C', D') \in (C''_c, D''_c \langle R \rangle^\dagger)[Inv(l')]$ and
 - $(C'_c, D'_c) \in (C', D')_c$.

We say that a state $S = (l, \llbracket C, D \rrbracket) \in \mathbb{S}_A$ is accepting if $l \in F$. We say that $\pi = S_0 \Longrightarrow S_1 \Longrightarrow \dots$ is a run of $\llbracket A \rrbracket_{lb, ub}$ if $S_0 \in \mathbb{S}_{init}$ and for each i , $S_i \in \mathbb{S}_A$ and $S_{i-1} \Longrightarrow S_i$. A run respects a parameter valuation v if for each state $S_i = (l_i, \llbracket C_i, D_i \rrbracket)$ it holds that $v \in \llbracket C_i \rrbracket$. A run π is accepting if there exists an infinite set of indices i such that S_i is accepting. For the rest of the paper we fix lb, ub and use $\llbracket A \rrbracket$ to denote $\llbracket A \rrbracket_{lb, ub}$.

3.2 Finite Abstraction

Similarly to the nonparametric case, the symbolic transition system $\llbracket A \rrbracket$ may be infinite. In order to obtain a finite transition system we need to apply a finite abstraction over $\llbracket A \rrbracket$. In the standard case of timed automata without parameters we use one of the extrapolation techniques [13, 14]. In our parametric setup we define a new finite abstraction called the *pk-extrapolation* which is a parametric extension of the widely used *k-extrapolation* [13]. The *k-extrapolation* identifies states which are identical except for the clock values which exceeds the maximal constant from guards and invariants.

In our parametric setup, we need to define the maximal constant with which each clock within a PTBA is compared. We define $M(x)$ as the maximal value in $\{max_{lb,ub}(e) \mid e \text{ is compared with } x \text{ in a guard or an invariant of the considered PTBA}\}$. The core idea of *pk-extrapolation* is the same as the idea of *k-extrapolation*. We substitute each bound on clock difference in the CPDBM whenever this bound exceeds the maximal constant. The precise description of this substitution process is given in the Definition 3.5. Contrary to the nonparametric case, due to the occurrence of parameters in the CPDBM bounds, the substitution process may be ambiguous. In these situations we restrict the parameter values in order to obtain an unambiguous situation. This solution is similar to the constraint set splitting that is done in the application of a guard and in the canonisation procedure. Therefore, the result of *pk-extrapolation* is a set of CPDBMs instead of a single CPDBM.

Definition 3.5. *Let A be a PTBA, $(l, \llbracket C, D \rrbracket)$ be a symbolic state of $\llbracket A \rrbracket$ and $D_{ij} = x_i - x_j \prec_{ij} e_{ij}$ for each $0 \leq i, j \leq |X|$. We define the *pk-extrapolation* α_{pk} in the following way. $\alpha_{pk}(l, \llbracket C, D \rrbracket)$ is the set of all $(l, \llbracket C', D' \rrbracket)$ such that for each i, j , $0 \leq i, j \leq |X|$ one of the following conditions holds:*

- $D'_{ij} = x_i - x_j \prec_{ij} e_{ij}$ and the constraint $(e_{ij} \leq M(x_i)) \in C'$,
- $D'_{ij} = x_i - x_j < \infty$ and the constraint $(e_{ij} > M(x_i)) \in C'$,
- $D'_{ij} = x_i - x_j \prec_{ij} e_{ij}$ and the constraint $(e_{ij} \geq -M(x_j)) \in C'$,
- $D'_{ij} = x_i - x_j < -M(x_j)$ and the constraint $(e_{ij} < -M(x_j)) \in C'$.

Example 3.6. Consider $x, y \in X$, $p \in P$, $p \in [0, 7]$, $M(x) = M(y) = 10$, and the symbolic state $(l, \llbracket C, D \rrbracket)$ where $C = \emptyset$ and $D = \{x \leq y, y \leq x, y \leq 2p\}$. Now, $\alpha_{pk}(l, \llbracket C, D \rrbracket)$ contains two symbolic states: $(l, \llbracket C_1, D_1 \rrbracket)$ and $(l, \llbracket C_2, D_2 \rrbracket)$ where $C_1 = \{2p \leq 10\}$, $D_1 = \{x \leq y, y \leq x, y \leq 2p\}$, $C_2 = \{2p > 10\}$, $D_2 = \{x \leq y, y \leq x, y < \infty\}$.

Theorem 3.7. *Let A be a PTBA. The *pk-extrapolation* is a finite abstraction that preserves all accepting runs of $\llbracket A \rrbracket_v$ for each parameter valuation v .*

Proof Idea. We can transform the proof of Theorem 1 of [15] as well as the corresponding lemmata and definitions into our parametric setup. Due to space constraints, we did not include the full technically detailed proof and we kindly refer the reader to [16].

4 Parameter Synthesis Algorithm

We recall that our main objective is to find all parameter valuations for which the parametric timed automaton satisfies its specification. In the previous sections we have described the standard automata-based method employed under a parametric setup which produces a Büchi automaton. For the rest of this section we use $s.\llbracket C \rrbracket$ to denote the set $\llbracket C \rrbracket$ where $s = (l, \llbracket C, D \rrbracket)$ is a state of the input Büchi automaton. We say that a sequence of states $s_1 \Longrightarrow s_2 \Longrightarrow \dots \Longrightarrow s_n \Longrightarrow s_1$ is a cycle under the parameter valuation v if each state s_i in the sequence satisfies $v \in s_i.\llbracket C \rrbracket$. A cycle is called accepting if there exists $0 \leq i \leq n$ such that s_i is accepting.

The standard automata-based LTL model checking checks the emptiness of the produced Büchi automaton. The emptiness check can be performed using the Nested Depth First Search (NDFS) algorithm [17]. The NDFS algorithm is a modification of the depth first search algorithm which allows a detection of an accepting cycle in the given Büchi automaton.

Contrary to the standard LTL model checking, it is not enough to check the emptiness of the produced Büchi automaton. Our objective is to check the emptiness of the produced Büchi automaton for each considered parameter valuation. To solve this objective, we introduce a new algorithm called the Cumulative NDFS algorithm which is an extension of the NDFS algorithm. The pseudocode of Cumulative NDFS is given in Algorithm 1. Our modification is based on the set *Found* which accumulates all detected parametric valuations such that an accepting cycle under these valuations was found. In contrast to the NDFS algorithm, whenever Cumulative NDFS detects an accepting cycle, parameter valuations are saved to the set *Found* and the computation continues with a search for another accepting cycle. Note the fact that whenever we reach a state s' with $s'.\llbracket C \rrbracket \subseteq \text{Found}$ we already have found an accepting cycle under all valuations from $s'.\llbracket C \rrbracket$ and there is no need to continue with the search from s' . Therefore, we are able to speed up the computation whenever we reach such a state.

The crucial property the algorithm is based on is that of monotonicity. The set of parameter valuations $s.\llbracket C \rrbracket$ can not grow along any run of the input automaton. Lemma 4.1 states this observation formally. The observation follows from the definition of successors in $\llbracket A \rrbracket^\alpha$ and the definition of operations on CPDBMs. The clear corollary of Lemma 4.1 is the fact that each state s on a cycle has the same set $s.\llbracket C \rrbracket$.

Lemma 4.1. *Let A be a PTBA, α be an abstraction and s be a state in $\llbracket A \rrbracket^\alpha$. For every state s' reachable from s it holds that $s'.\llbracket C \rrbracket \subseteq s.\llbracket C \rrbracket$.*

Theorem 4.2. *Let A be a PTBA and α an abstraction over $\llbracket A \rrbracket$. A parameter valuation v is contained in the output of the CumulativeNDFS($\llbracket A \rrbracket^\alpha$) if and only if there exists an accepting run respecting v in $\llbracket A \rrbracket^\alpha$.*

Due to space constraints, we did not include the full technically detailed proof and we kindly refer the reader to [16].

<p>Algorithm <i>CumulativeNDFS</i>(G)</p> <p>1 $Found \leftarrow \emptyset$; $Stack \leftarrow \emptyset$ $Outer \leftarrow \emptyset$; $Inner \leftarrow \emptyset$</p> <p>2 $OuterDFS(s_{init})$</p> <p>3 return $Accepted \leftarrow Found$</p> <p>Procedure <i>OuterDFS</i>(s)</p> <p>4 $Stack \leftarrow Stack \cup \{s\}$</p> <p>5 $Outer \leftarrow Outer \cup \{s\}$</p> <p>6 foreach s' such that $s \rightarrow s'$ do</p> <p>7 if $s' \notin Outer \wedge s' \notin Stack \wedge$ $s'.\llbracket C \rrbracket \not\subseteq Found$ then</p> <p>8 $OuterDFS(s')$</p> <p>9 if $s \in Accepting \wedge s.\llbracket C \rrbracket \not\subseteq Found$ then</p> <p>10 $InnerDFS(s)$</p> <p>11 $Stack \leftarrow Stack \setminus \{s\}$</p>	<p>Procedure <i>InnerDFS</i>(s)</p> <p>12 $Inner \leftarrow Inner \cup \{s\}$</p> <p>13 foreach s' such that $s \rightarrow s'$ do</p> <p>14 if $s' \in Stack$ then</p> <p>15 “Cycle detected”</p> <p>16 $Found \leftarrow Found \cup s'.\llbracket C \rrbracket$</p> <p>17 return</p> <p>18 if $s' \notin Inner \wedge$ $s'.\llbracket C \rrbracket \not\subseteq Found$ then</p> <p>19 $InnerDFS(s')$</p>
--	---

Algorithm 1. Cumulative NDFS

As the last step in the solution to our problem, we need to complement the set *Accepted*. Thus, the solution is the complement of the set *Accepted*, more precisely the set $Val_{lb,ub}(X, P) \setminus Accepted$. To conclude this section, we state that Theorem 4.2 together with Theorem 3.7 imply the correctness of our solution.

5 Implementation

We have implemented our approach in a proof-of-concept tool. We are able to process models given as networks of parametric timed automata. A network represents a product of several parametric timed automata where handshake synchronization of two components at a time is allowed. We also extend the parametric timed automata with data variables which enable the usage of guards on data values and transition effects on data values. Such model is considered standard in the field and is used as the modelling language in the tool UPPAAL.

Deadlocks. Cumulative NDFS algorithm returns all parameter valuations for which LTL property does not hold. However, state space can contain deadlock states which also need to be detected and reported. In the nonparametric setting a state is a deadlock state if there are no enabled outgoing transitions. In a parametric setting the deadlock status of a state depends on the parameter valuation. To decide for which parameter valuations a state $(l, \llbracket C, D \rrbracket)$ is a deadlock we need to consider all guards g_1, \dots, g_n of the outgoing transitions of l . The state $(l, \llbracket C, D \rrbracket)$ is a deadlock for all parameter valuations in $\llbracket C, D \rrbracket[\neg g_1 \wedge \dots \wedge \neg g_n]$. Applying this detection to each reachable state, all parameter valuations leading to deadlock are detected during computation.

<pre> Procedure <i>InitializeStorage</i>() 1 <i>Storage</i> $\leftarrow \emptyset$; $M_1 \leftarrow \emptyset$; $M_2 \leftarrow \emptyset$ Procedure <i>SetData</i>($l, C, D, data$) 2 if $M_2(C, D) \neq \emptyset$ then 3 $(C', D') \leftarrow M_2(C, D)$ 4 <i>Storage</i>(l, C', D') $\leftarrow data$ 5 else 6 $IH \leftarrow integerHull(C, D)$ 7 foreach (C', D') <i>in</i> $M_1(IH)$ do 8 if $\llbracket C', D' \rrbracket = \llbracket C, D \rrbracket$ then 9 $M_2(C, D) \leftarrow (C', D')$ 10 <i>Storage</i>(l, C', D') $\leftarrow data$ 11 $M_2(C, D) \leftarrow (C, D)$ 12 $M_1(IH) \leftarrow M_1(IH) \cup \{(C, D)\}$ 13 <i>Storage</i>(l, C, D) $\leftarrow data$ </pre>	<pre> Procedure <i>GetData</i>(l, C, D) 14 if $M_2(C, D) \neq \emptyset$ then 15 $(C', D') \leftarrow M_2(C, D)$ 16 return <i>Storage</i>(l, C', D') 17 else 18 $IH \leftarrow integerHull(C, D)$ 19 foreach (C', D') <i>in</i> $M_1(IH)$ do 20 if $\llbracket C', D' \rrbracket = \llbracket C, D \rrbracket$ then 21 $M_2(C, D) \leftarrow (C', D')$ 22 return <i>Storage</i>(l, C', D') 23 $M_2(C, D) \leftarrow (C, D)$ 24 $M_1(IH) \leftarrow M_1(IH) \cup \{(C, D)\}$ 25 <i>Storage</i>(l, C, D) $\leftarrow initialData$ 26 return <i>initialData</i> </pre>
---	--

Algorithm 2. State space storage operations

State space storage. One of the performance critical parts of the implementation is the state space storage. We use the state space storage to look up and store information about presence of each state in the sets *Inner*, *Outer*, and *Stack*. We refer to this information as *data*. A straightforward implementation would simply store each state together with its data. Such a solution is only efficient when a unique representation of states is available. Without such a unique representation the storage operations have to perform expensive equivalence checks with each stored state in the worst case scenario. In [10] the authors introduce unique representation based on a computation of an integer hull. The integer hull of a given set is a convex hull of all integer elements of a given set.

The solution of [10] assumes the existence of an upper bound for each clock. We do not have such an upper-bound assumption and therefore this solution is not directly applicable in our technique. However, we use the integer hull as a heuristic approximation of a unique representation of a CPDBM instead. This way we obtain a practically efficient solution that deals with the non-existence of a unique representation of a state.

The solution is based on two mappings. The first mapping, denoted by M_1 maps a given integer hull to a list of CPDBM representations. Each such list contains the representations of semantically different CPDBMs with the same integer hull. Thanks to M_1 we can quickly distinguish states with different integer hulls. However, each storage operation still needs to perform the expensive computation of the integer hull. In order to reduce number of the integer hull computations, we introduce the second mapping, denoted by M_2 . This second mapping serves as a cache which maps a given CPDBM to its unique representative in the storage. Once a CPDBM representative is resolved, it is saved in M_2 .

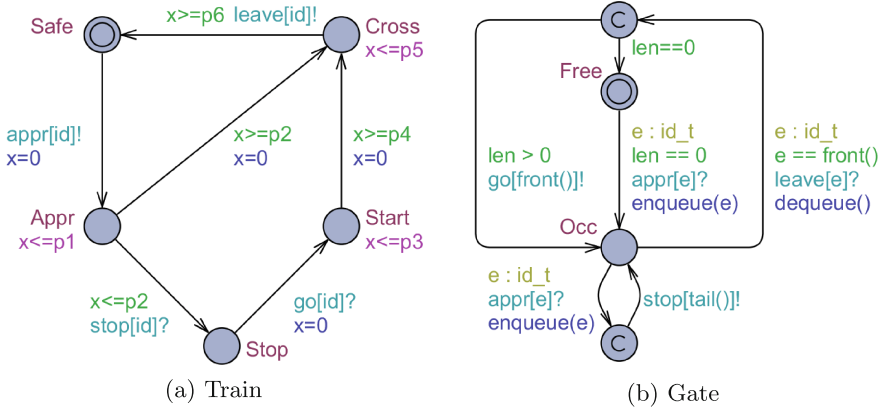


Fig. 1. Parametric TrainGate model

The pseudo code of state space storage operations is given in Algorithm 2. Note that the procedures *SetData* and *GetData* are analogous. In our prototype tool, the two mappings as well as the storage itself are implemented using hash tables. Checking whether two states are semantically equivalent is implemented using Parma Polyhedra Library [18]. The library is also used to check parametric constraint satisfaction in the CPDBM operations.

6 Experimental Evaluation

We have implemented the proposed technique for integer parameter synthesis in our proof-of-concept tool. Our goal is to compare our method with the explicit enumeration technique. To be able to compare performance of both techniques under similar conditions we also implemented the standard DBM-based LTL model checker for timed automata. Both tools use the same LTL to BA translation method [19] and analogous extrapolation techniques.

Our evaluation was performed on a parametric extension of the case study TrainGate [20] provided with the tool UPPAAL. In the TrainGate model we substitute all 6 integer bounds with separate parameters and consider two trains. This model is presented in Fig. 1. We checked two LTL properties. The first property *prop1* states that the two trains can not cross the bridge simultaneously ($G!(Train_1.cross \text{ and } Train_2.cross)$). The second property *prop2* states that whenever the first train is approaching the bridge it will cross the bridge eventually ($G Train_1.appr \implies F Train_1.cross$). For all considered parameter valuations which do not lead to the deadlock, *prop1* and *prop2* are satisfied.

Experiments were performed on a PC with CPU i5-4690 and 16 GB RAM. We considered a timeout of 12 h for each task. We provide percentage of solved parameter valuations if the timeout was reached by explicit enumeration.

Table 1 shows the impact of the number of parameters used in the model. For models with a small number of parameters and small value ranges the explicit

Table 1. Impact of model parameter count

TrainGate model 2 trains	3 params	4 params	5 params	6 params
	$p_1 \in [20, 50]$	$p_1 \in [20, 50]$	$p_1 \in [20, 50]$	$p_1 \in [20, 50]$
	$p_2 \in [10, 50]$	$p_2 \in [10, 50]$	$p_2 \in [10, 50]$	$p_2 \in [10, 50]$
	$p_3 \in [15, 50]$	$p_3 \in [15, 50]$	$p_3 \in [15, 50]$	$p_3 \in [15, 50]$
	$p_4 = 7$	$p_4 \in [7, 50]$	$p_4 \in [7, 50]$	$p_4 \in [7, 50]$
	$p_5 = 5$	$p_5 = 5$	$p_5 \in [5, 50]$	$p_5 \in [5, 50]$
	$p_6 = 3$	$p_6 = 3$	$p_6 = 3$	$p_6 \in [3, 50]$
prop1 explicit enumeration	0:01:03	0:44:50	Timeout(51%)	Timeout(2%)
prop1 cumulative algorithm	0:08:16	0:54:39	3:20:25	7:58:42
prop2 explicit enumeration	0:01:21	0:58:17	Timeout(42%)	Timeout(1%)
prop2 cumulative algorithm	0:12:20	1:23:37	5:11:01	10:48:16

Table 2. Impact of parameter range size

TrainGate model	$p_1 \in [20, 50]$	$p_1 \in [20, 100]$	$p_1 \in [10, 100]$
2 trains	$p_2 \in [10, 50]$	$p_2 \in [10, 100]$	$p_2 \in [10, 100]$
4 parameters	$p_3 \in [15, 50]$	$p_3 \in [15, 100]$	$p_3 \in [10, 100]$
$p_5 = 5$ $p_6 = 3$	$p_4 \in [7, 50]$	$p_4 \in [7, 100]$	$p_4 \in [10, 100]$
prop1 explicit enumeration	0:44:50	Timeout(68%)	Timeout(63%)
prop1 cumulative algorithm	0:54:39	7:39:43	6:56:49
prop2 explicit enumeration	0:58:17	Timeout(56%)	Timeout(53%)
prop2 cumulative algorithm	1:23:37	10:25:28	8:59:11

enumeration can be more efficient. However, higher parameter count significantly favours the cumulative algorithm. Table 2 shows the impact of the parameter range size on the execution times. Note that for larger parameter ranges the cumulative algorithm is faster than explicit enumeration.

7 Conclusion and Future Work

We have presented an algorithmic framework for the bounded integer parameter synthesis for parametric timed automata with an LTL specification. The proposed framework allows the avoidance of the explicit enumeration of all possible parameter valuations.

Our symbolic technique is based on the zone abstraction and uses a parametric extension of difference bound matrices. To be able to employ the zone-based method successfully we have introduced a finite abstraction called the pk-extrapolation. To be able to synthesize all violating parameter valuations we have introduced the Cumulative NDFS algorithm which is an extension of the NDFS algorithm.

We have implemented the proposed technique in an experimental tool and our experiments confirm that this technique can be significantly faster than the explicit enumeration technique.

As for future work we plan to introduce different finite abstractions based on different extrapolations and compare their influence on the state space size. We also plan to introduce a parallel version of the cumulative algorithm. Other area that can be investigated is the employment of different linear specification logics, e.g. Clock-Aware LTL [21] which enables the use of clock-valuation constraints as atomic propositions.

References

1. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press, Cambridge (1999)
2. Alur, R., Dill, D.L.: A theory of timed automata. *Theoret. Comput. Sci.* **126**(2), 183–235 (1994)
3. Daws, C., Tripakis, S.: Model checking of real-time reachability properties using abstractions. In: Steffen, B. (ed.) TACAS 1998. LNCS, vol. 1384, pp. 313–329. Springer, Heidelberg (1998)
4. Behrmann, G., David, A., Larsen, K.G., Hakansson, J., Petterson, P., Yi, W., Hendriks, M.: Uppaal 4.0. In: QEST, pp. 125–126. IEEE (2006)
5. Alur, R., Henzinger, T.A., Vardi, M.Y.: Parametric real-time reasoning. In: Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, pp. 592–601. ACM (1993)
6. Miller, J.S.: Decidability and complexity results for timed automata and semi-linear hybrid automata. In: Lynch, N.A., Krogh, B.H. (eds.) HSCC 2000. LNCS, vol. 1790, p. 296. Springer, Heidelberg (2000)
7. Beneš, N., Bezděk, P., Larsen, K.G., Srba, J.: Language emptiness of continuous-time parametric timed automata. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015. LNCS, vol. 9135, pp. 69–81. Springer, Heidelberg (2015)
8. Hune, T., Romijn, J., Stoelinga, M., Vaandrager, F.: Linear parametric model checking of timed automata. *J. Logic Algebraic Programm.* **52**, 183–220 (2002)
9. Bozzelli, L., La Torre, S.: Decision problems for lower/upper bound parametric timed automata. *Formal Methods Syst. Des.* **35**(2), 121–151 (2009)
10. Jovanovic, A., Lime, D., Roux, O.H.: Integer parameter synthesis for real-time systems. *IEEE Trans. Softw. Eng.* **41**(5), 445–461 (2015)
11. Tripakis, S., Yovine, S., Bouajjani, A.: Checking timed büchi automata emptiness efficiently. *Formal Methods Syst. Des.* **26**(3), 267–292 (2005)
12. Dill, D.L.: Timing assumptions and verification of finite-state concurrent systems. In: Sifakis, J. (ed.) Automatic Verification Methods for Finite State Systems. LNCS, vol. 407, pp. 197–212. Springer, Heidelberg (1990)
13. Bouyer, P.: Forward analysis of updatable timed automata. *Formal Methods Syst. Des.* **24**(3), 281–320 (2004)
14. Behrmann, G., Bouyer, P., Larsen, K.G., Pelánek, R.: Lower and upper bounds in zone-based abstractions of timed automata. *Int. J. Softw. Tools Technol. Transf.* **8**(3), 204–215 (2006)

15. Li, G.: Checking timed Büchi automata emptiness using LU-abstractions. In: Ouaknine, J., Vaandrager, F.W. (eds.) FORMATS 2009. LNCS, vol. 5813, pp. 228–242. Springer, Heidelberg (2009)
16. Bezděk, P., Beneš, N., Barnat, J., Černá, I.: LTL parameter synthesis of parametric timed automata. CoRR abs/1409.3696 (2016)
17. Courcoubetis, C., Vardi, M., Wolper, P., Yannakakis, M.: Memory-efficient algorithms for the verification of temporal properties. In: Clarke, E.M., Kurshan, R.P. (eds.) CAV. LNCS, vol. 531, pp. 233–242. Springer, Heidelberg (1992)
18. Bagnara, R., Hill, P.M., Zaffanella, E.: The parma polyhedra library: toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems. *Sci. Comput. Programm.* **72**(1–2), 3–21 (2008)
19. Gastin, P., Oddoux, D.: Fast LTL to Büchi automata translation. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, p. 53. Springer, Heidelberg (2001)
20. Behrmann, G., David, A., Larsen, K.G.: A tutorial on UPPAAL. In: Bernardo, M., Corradini, F. (eds.) SFM-RT 2004. LNCS, vol. 3185, pp. 200–236. Springer, Heidelberg (2004)
21. Bezděk, P., Beneš, N., Havel, V., Barnat, J., Černá, I.: On clock-aware LTL properties of timed automata. In: Ciobanu, G., Méry, D. (eds.) ICTAC 2014. LNCS, vol. 8687, pp. 43–60. Springer, Heidelberg (2014)