

Template-based Controller Synthesis for Timed Systems

Bernd Finkbeiner and Hans-Jörg Peter

Reactive Systems Group
Department of Computer Science
Universität des Saarlandes
66123 Saarbrücken, Germany

Abstract. We present an effective controller synthesis method for real-time systems modeled as timed automata with safety requirements. Under the realistic assumption of partial observability, the problem is undecidable in general, and prohibitively expensive (2EXPTIME-complete) if a bound on the granularity of the controller is set in advance. We investigate the synthesis of controllers from templates, given as timed automata with parametric control structure. Template-based synthesis is significantly cheaper (PSPACE-complete) than standard synthesis and produces much simpler controllers. We present an efficient symbolic synthesis algorithm based on automatic abstraction refinement and report on encouraging experimental results from an implementation in the timed verification and synthesis tool SYNTHIA.

1 Introduction

In controller synthesis, we automatically transform a given model of a plant and a safety requirement into a finite-state controller that monitors and affects the ongoing behavior of the plant to ensure the safety of its operation. There has been a lot of recent progress [AT02,CDF⁺05,BCD⁺07,PM09,PEM11] in synthesizing controllers for real-time systems, where the plant is given as a timed automaton. Notably, the UPPAAL-TiGA tool [BCD⁺07], which is based on the popular timed model checker UPPAAL, has extended the highly efficient state-space traversal based on symbolic zone representations from verification to synthesis.

Unfortunately, timed controller synthesis quickly turns into an intractable problem if one makes the realistic assumption that the controller does not have access to the full state of the plant, but rather only sees a subset of the events. Under *partial observability*, the controller synthesis problem is undecidable in general, and remains prohibitively expensive (2EXPTIME-complete) if the problem is made decidable by fixing a bound on the granularity of the controller in advance [BDMP03]. Furthermore, since the size of the controller may be doubly-exponential in the size of the plant, it is often infeasible to actually construct the controller.




Fig. 1. Full template with one, two, and three locations

In this paper, we propose a new synthesis approach, where the size and general shape of the controller is fixed in advance in the form of a *template*. A template is a timed or untimed automaton with parametric control structure. Figures 1 and 2 show two example template families. In the *full template*, shown in Figure 1, every pair of locations is connected by an edge for every possible action. In the *cyclic-executive template*, shown in Figure 2, the controller implements some schedule according to which the actions are handled; the controller alternates between waiting for an uncontrollable action and responding with some controllable action. The template families are organized according to the number of locations. Typically, we start the synthesis process with small templates and then iteratively increase the size until an optimal controller is found. A controller *matches* a template if it can be obtained by removing a subset of the edges. We encode the presence of edges using Boolean parameter variables and combine the resulting parametric timed automaton with the plant automaton: any valuation of the Boolean parameters under which only safe states are reachable represents a correct controller.

Template-based synthesis has several attractive features: Since the observations of the controller are limited by the template, template-based synthesis naturally solves the controller synthesis problem with partial observability. The size of the controller is also limited by the size of the template. Because the templates model standard types of controllers, the synthesized controllers are well-structured, resembling a manually built controller.

In terms of complexity, it is not surprising that template-based synthesis is much simpler than standard synthesis. The problem is PSPACE-complete, matching the complexity of model checking. Template-based synthesis can in fact be understood as parametric model checking, where we verify a timed automaton that is parametrized with Boolean variables.

The technical challenge in developing fast algorithms for template-based synthesis thus lies in the efficient manipulation of the potential valuations of the parameters. In the paper, we present a solution to this challenge based on auto-




Fig. 2. Cyclic-executive template with two and four locations

matic abstraction refinement. Starting with an initial abstraction that considers all parameter valuations, a refinement loop incrementally focuses the search towards smaller and smaller sets of parameter valuations. The loop terminates as soon as, for the remaining parameter valuations, only safe states are reachable. New refinements are computed by identifying situations where an unsafe state is reachable for a subset of the parameter valuations.

Our experimental results indicate that template-based synthesis is not only an effective solution to the controller synthesis problem with partial observability, it is an attractive alternative to standard synthesis also in the simpler case of complete observability, outperforming tools like UPPAAL-TIGA on benchmarks where structurally simple controllers, corresponding to the available templates, exist.

Related work. The basic timed controller synthesis problem for timed automata [AD94] with complete observability was defined by Maler et al. [MPS95,AMPS98] in the setting of turn-based timed games. In their fundamental work, the decidability of the problem was shown by demonstrating that the standard discrete attractor construction [Tho95] on the region graph suffices to obtain winning strategies. Henzinger and Kopke showed that this construction is theoretically optimal by proving that the problem is EXPTIME-complete [HK99]. D’Souza and Madhusudan investigated the complexity of timed controller synthesis against external specifications [DM02]. Bouyer et al. continued this line of research and also investigated the impact of partial information [BDMP03].

A first more practical approach to timed controller synthesis, implemented in the tool SYNTHKRO, was proposed by Altisen and Tripakis [AT02]. The approach requires, however, an expensive preprocessing step. Cassez et al. presented a symbolic algorithm, implemented in UPPAAL-TIGA, that avoids the upfront state explosion by combining the backward attractor construction with a forward zone graph exploration [CDF⁺05,BCD⁺07]. Our timed verification and synthesis tool SYNTHIA [PEM11] is based on abstraction refinement techniques that combine symbolic representations for the discrete and the continuous state components [EMP10] and exploit the compositional structure of the timed system [PM09]. We have implemented the approach presented in this paper as an extension of SYNTHIA.

Compared to the significant body of work on timed controller synthesis with complete observability, there has been comparatively little work on the more re-

alistic setting of partial observability. Fundamental results on the undecidability of the general problem and the complexity for fixed granularity are due to Bouyer et al. [BDMP03]. Cassez et al. proposed a pragmatic approach to handle partial information, which restricts the choices and the observability of the controller so that a zone-based synthesis algorithm remains possible [CDL⁺07]. An extension of this work uses *alternating timed simulation relations* to efficiently control partially observable systems [CDL09]. This approach has also been implemented in UPPAAL-TIGA.

Template-based synthesis is related to the *bounded synthesis* approach [SF07], where one fixes the size (but not the structure) of the controller. Bounded synthesis has so far been limited to purely discrete systems. There are efficient algorithms for bounded synthesis based on SMT-solving [FS07], antichains [FJR09], and BDDs [Ehl10], which, however, unfortunately do not seem to have straightforward extensions to the timed case. Another interesting restriction on the type of controllers to be considered has been proposed by Lustig et al.: *synthesis from component libraries* [LV09] attempts to construct a controller by assembling routines from a given library. The difference to template-based synthesis is that the synthesized controller is a combination of predefined components rather than an instantiation of a parametric template. Currently, this approach is also limited to discrete systems.

2 Timed Systems

Timed automata. The components of a timed system are represented by *timed automata*. A timed automaton [AD94] is a tuple $\mathcal{A} = (L, l_0, \Sigma, \Delta, X, I)$, where L is a finite set of (control) locations, $l_0 \in L$ is the initial location, Σ is a finite set of actions, $\Delta \subseteq (L \times \Sigma \times \mathcal{C}(X) \times 2^X \times L)$ is an edge relation, X is a finite set of real valued clocks, $I : L \rightarrow \mathcal{C}(X)$ maps each location to an invariant, and $\mathcal{C}(X)$ is the set of clock constraints over X . A (rectangular) clock constraint $\varphi \in \mathcal{C}(X)$ is of the form

$$\varphi = \mathbf{true} \mid x \leq c \mid c \leq x \mid x < c \mid c < x \mid \varphi_1 \wedge \varphi_2,$$

where x is a clock in X and c is a constant in \mathbb{N}_0 . A *clock valuation* $\mathbf{t} : X \rightarrow \mathbb{R}_{\geq 0}$ assigns a nonnegative value to each clock and can also be represented by a $|X|$ -dimensional vector $\mathbf{t} \in \mathcal{R}$, where $\mathcal{R} = \mathbb{R}_{\geq 0}^X$ denotes the set of all clock valuations.

The states of a timed automaton are pairs (l, \mathbf{t}) of locations and clock valuations. Timed automata have two types of transitions: *timed transitions*, where only time passes and the location remains unchanged, and *discrete transitions*, where no time passes, the current location can be changed and some clocks can be reset to zero. In a timed transition, denoted by $(l, \mathbf{t}) \xrightarrow{d} (l, \mathbf{t} + d \cdot \mathbf{1})$, the same nonnegative value $d \in \mathbb{R}_{\geq 0}$ is added to all clocks such that, for each $0 \leq d' \leq d$, $\mathbf{t} + d'$ satisfies the location invariant $I(l)$. A discrete transition, denoted by $(l, \mathbf{t}) \xrightarrow{a} (l', \mathbf{t}')$ for some action $a \in \Sigma$, corresponds to an edge $\delta = \langle l, a, \varphi, \lambda, l' \rangle$ of Δ such that \mathbf{t} satisfies the clock constraint φ of δ , and $\mathbf{t}' = \mathbf{t}[\lambda := 0]$ is obtained

from \mathbf{t} by setting the clocks in λ to 0 and satisfies the location invariant $I(l')$. For two states s and s' , we write $s \xrightarrow{\delta} s'$ if there is a delay $d \in \mathbb{R}_{\geq 0}$ and an edge δ with action a such that there is an s'' with $s \xrightarrow{d} s''$ and $s'' \xrightarrow{a} s'$.

We say that a state s is *reachable* if there is a finite sequence of transitions of the form $s_0 \xrightarrow{\delta_0} s_1 \dots s_{n-1} \xrightarrow{\delta_{n-1}} s$ such that $\delta_0, \dots, \delta_{n-1} \in \Delta$ are edges in Δ , $s_0 = (l_0, \mathbf{0})$ is the initial state (where $\mathbf{0}$ is the zero vector), and for all $1 \leq i \leq n$, the individual $s_i = (l_i, \mathbf{t}_i)$ are states of the automaton. We define $\text{Reach}(\mathcal{A})$ as the set of all forward reachable states of a timed automaton \mathcal{A} .

Composition. Timed automata can be composed to networks, in which the automata run in parallel and synchronize on shared actions. For two timed automata $\mathcal{A} = (L_1, l_0^1, \Sigma_1, \Delta_1, X_1, I_1)$ and $\mathcal{A}' = (L_2, l_0^2, \Sigma_2, \Delta_2, X_2, I_2)$ with disjoint clock sets $X_1 \cap X_2 = \emptyset$, the *parallel composition* $\mathcal{A}_1 \parallel \mathcal{A}_2$ is the timed automaton $(L_1 \times L_2, (l_0^1, l_0^2), \Sigma_1 \cup \Sigma_2, \Delta, X_1 \cup X_2, I)$, where $I(l_1, l_2) = I_1(l_1) \wedge I_2(l_2)$ for all $l_1 \in L_1$ and $l_2 \in L_2$, and Δ is the smallest set that contains

- for $a \in \Sigma_1 \cap \Sigma_2$, $\langle (l_1, l_2), a, \varphi_1 \wedge \varphi_2, \lambda_1 \cup \lambda_2, (l'_1, l'_2) \rangle$ if $\langle l_1, a, \varphi_1, \lambda_1, l'_1 \rangle \in \Delta_1$ and $\langle l_2, a, \varphi_2, \lambda_2, l'_2 \rangle \in \Delta_2$,
- for $a \in \Sigma_1 \setminus \Sigma_2$, $\langle (l_1, l_2), a, \varphi_1, \lambda_1, (l'_1, l_2) \rangle$ if $\langle l_1, a, \varphi_1, \lambda_1, l'_1 \rangle \in \Delta_1$, and
- for $a \in \Sigma_2 \setminus \Sigma_1$, $\langle (l_1, l_2), a, \varphi_2, \lambda_2, (l_1, l'_2) \rangle$ if $\langle l_2, a, \varphi_2, \lambda_2, l'_2 \rangle \in \Delta_2$.

Finite Semantics. The decidability of the reachability problem of timed automata relies on the existence of the *region equivalence relation* [AD94] on \mathcal{R} which has a finite index.

For a timed automaton $\mathcal{A} = (L, l_0, \Sigma, \Delta, X, I)$, we call the value of a clock $x \in X$ *maximal* if it is strictly greater than the highest constant c_{max} any clock is compared to. We say that two clock valuations $\mathbf{t}_1, \mathbf{t}_2 \in \mathcal{R}$ are in the same *clock region*, denoted $\mathbf{t}_1 \sim_R \mathbf{t}_2$, if

- the set of clocks with maximal value is the same in \mathbf{t}_1 and in \mathbf{t}_2 ($\forall x \in X : \mathbf{t}_1(x) > c_{max} \Leftrightarrow \mathbf{t}_2(x) > c_{max}$), and
- \mathbf{t}_1 and \mathbf{t}_2 agree (1) on the integer parts of the clock values, (2) on the relative order of the noninteger parts of the clock values, and (3) on the equality of the noninteger parts of the clock values with 0. That is, for all clocks x and y with nonmaximal value, it holds that (1) $\lfloor \mathbf{t}_1(x) \rfloor = \lfloor \mathbf{t}_2(x) \rfloor$, (2) $\widehat{\mathbf{t}}_1(x) \leq \widehat{\mathbf{t}}_1(y) \Leftrightarrow \widehat{\mathbf{t}}_2(x) \leq \widehat{\mathbf{t}}_2(y)$, and (3) $\widehat{\mathbf{t}}_1(x) = 0$ if, and only if, $\widehat{\mathbf{t}}_2(x) = 0$, where $\widehat{\mathbf{t}}_i(x) = \mathbf{t}_i(x) - \lfloor \mathbf{t}_i(x) \rfloor$ for $i \in \{1, 2\}$.

We denote with $[\mathbf{t}]_R = \{\mathbf{t}' \in \mathcal{R} \mid \mathbf{t} \sim_R \mathbf{t}'\}$ the clock region \mathbf{t} belongs to. We say that two states $s_1 = (l_1, \mathbf{t}_1)$ and $s_2 = (l_2, \mathbf{t}_2)$ of \mathcal{A} are *region-equivalent*, denoted by $s_1 \sim_R s_2$, if their locations are the same ($l_1 = l_2$) and the clock valuations are in the same clock region ($\mathbf{t}_1 \sim_R \mathbf{t}_2$), and denote with $[(l, \mathbf{t})]_R = \{(l, \mathbf{t}') \in L \times \mathcal{R} \mid \mathbf{t} \sim_R \mathbf{t}'\}$ the equivalence class of region-equivalent states that (l, \mathbf{t}) belongs to.

Regions are a suitable semantics for the abstraction of timed automata because they essentially preserve the language: if there is a discrete transition $s \xrightarrow{a} s'$ from a state s to a state s' of a timed automaton, then there is, for all states r with $r \sim_R s$, a state r' with $r' \sim_R s'$ such that $r \xrightarrow{a} r'$ is a discrete

transition with the same label. For timed transitions, a slightly weaker property holds: if there is a timed transition $s \xrightarrow{t} s'$ from a state s to a state s' , then there is, for all states r with $r \sim_R s$, a state r' with $r' \sim_R s'$ such that there is a timed transition $r \xrightarrow{t'} r'$ (but possibly with $t' \neq t$).

The *finite semantics* of a timed automaton $\mathcal{A} = (L, l_0, \Sigma, \Delta, X, I)$ is the finite graph $\llbracket \mathcal{A} \rrbracket = (Q, q_0, T)$ where

- the symbolic state set $Q = \{[(l, \mathbf{t})]_R \mid (l, \mathbf{t}) \in L \times \mathcal{R}\}$ of $\llbracket \mathcal{A} \rrbracket$ is the set of equivalence classes of region-equivalent states of \mathcal{A} , with
- the initial state $q_0 = [(l_0, \mathbf{t}_0)]_R$, and
- the set $T = \{(q, q') \in Q \times Q \mid \exists r \in q, r' \in q', a \in \Sigma \cup \mathbb{R}_{\geq 0}. r \xrightarrow{a} r'\}$ of transitions.

The finite semantics is reachability-preserving:

Lemma 1. [AD94] *For a timed automaton $\mathcal{A} = (L, l_0, \Sigma, \Delta, X, I)$ there is a finite path from a state (l, \mathbf{t}) to a state (l', \mathbf{t}') if, and only if, there is a finite path from $[(l, \mathbf{t})]_R$ to $[(l', \mathbf{t}')]_R$ in $\llbracket \mathcal{A} \rrbracket$.*

Assuming a binary encoding of the constants in the clock constraints, the number of states of the finite semantics is exponential in the number of clocks and in the magnitude of the constants:

Lemma 2. [AD94] *For a timed automaton $\mathcal{A} = (L, l_0, \Sigma, \Delta, X, I)$, with c_x as the maximal constant appearing in any constraint of \mathcal{A} , the number of states of $\llbracket \mathcal{A} \rrbracket$ is bounded by*

$$|L| \cdot |X|! \cdot 2^{|X|-1} \cdot \prod_{x \in X} O(c_x) = |L| \cdot |X|! \cdot O(c_x)^{|X|}.$$

As it turns out, the finite semantics is a theoretically optimal state space representation for deciding reachability:

Theorem 1. [AD94] *For a timed automaton \mathcal{A} and a set of states B , testing whether $\text{Reach}(\mathcal{A}) \cap B = \emptyset$ is PSPACE-complete.*

In practice, instead of deciding $\text{Reach}(\mathcal{A}) \cap B = \emptyset$ based on an explicit construction of the finite semantics, tools like SYNTHIA or UPPAAL use the much coarser *clock zones* as the fundamental representation of clock values.

3 Template-based Controller Synthesis

In this section, we formalize controller templates and the template instantiation problem. A *controller template* is a tuple (\mathcal{T}, P, Π) consisting of a timed automaton $\mathcal{T} = (L, l_0, \Sigma, \Delta, X, I)$, a finite set of Boolean parameters P , and a total function $\Pi : \mathcal{P} \rightarrow 2^\Delta$ defining which edges are enabled for a given parameter valuation, where $\mathcal{P} = P \rightarrow \mathbb{B}$ is the set of all parameter valuations. In the following, we will assume that the timed automaton modeling the environment (or plant) is already integrated (by parallel composition) in \mathcal{T} . As usual, we assume that the controller does neither reset plant clocks, inhibit plant actions, nor introduce timelocks.

Definition 1. For a controller template (\mathcal{T}, P, Π) with $\mathcal{T} = (L, l_0, \Sigma, \Delta, X, I)$ and a set of bad states B , the instantiation problem asks for a parameter valuation $\mathbf{p} \in \mathcal{P}$ such that $\mathcal{I} = (L, l_0, \Sigma, \Pi(\mathbf{p}), X, I)$ and $\text{Reach}(\mathcal{I}) \cap B = \emptyset$.

We call an instantiation of the template that satisfies the condition of the definition *feasible*. Synthesizing a template-based controller corresponds to *statically* finding a feasible instantiation. This is in contrast to the classical formulation of the timed controller synthesis problem [MPS95, AMPS98], where the controller is an arbitrary timed automaton whose behavior depends *dynamically* on the observed events of the plant.

The complexity of the template instantiation problem is the same as the complexity of standard timed model checking: the exponential size of the region graph dominates the size of the search space (the possible valuations of the parameters).

Theorem 2. The instantiation problem for a controller template (\mathcal{T}, P, Π) and a set of bad states B is PSPACE-complete.

We note that the synthesis of controllers with full observability is already EXPTIME-complete [HK99]. The case where the controller can only observe a subset of the events of the plant, is even undecidable in general, or 2EXPTIME-complete if the granularity (number and precision of the clocks) of the controller is bounded in advance [BDMP03]. Furthermore, the size of the controller can be exponential or even, in case of partial observability, doubly-exponential in the size of the plant. Template-based synthesis is not only much cheaper (PSPACE-complete), it also has the advantage that the size of the controller is fixed in advance.

Template-based synthesis thus provides a much more promising setting for effective controller synthesis than the standard approach. The remainder of the paper is devoted to the development of an efficient template-based synthesis algorithm and an experimental evaluation.

4 Symbolic Parameter Synthesis

We now present a symbolic algorithm for finding feasible instantiations for a given controller template (\mathcal{T}, P, Π) with $\mathcal{T} = (L, l_0, \Sigma, \Delta, X, I)$ and a set of bad states $B \subseteq S$, where S is the set of states of \mathcal{T} . In the rest of this section, we assume that \mathcal{T} , P , Π , S , and B are fixed.

We develop the algorithm in three steps: first, we describe the immediate, exact, computation of the set of feasible instantiations based on forward and backward propagation; then we give an approximate computation based on an abstraction of the template; finally, we describe an abstraction refinement procedure, which increases the precision of the approximate computation until either a feasible instantiation has been found, or it has been shown that no feasible instantiation exists.

4.1 Precise Computation of the Feasible Instantiations

The precise set of feasible instantiations can be computed in a standard fixed point construction that either starts from the initial state and propagates, in a forward manner, the reachable combinations of states and parameter valuations, or starts with the bad states and propagates, in a backward manner, those combinations of states and parameter valuations that have a path to the bad states.

To accommodate both directions, we define a successor and a predecessor propagation function $\text{Succ}, \text{Pred} : 2^{S \times \mathcal{P}} \rightarrow 2^{S \times \mathcal{P}}$ with

$$\begin{aligned} \text{Succ}(Y) &= \{(s', \mathbf{p}) \in S \times \mathcal{P} \mid \exists \delta \in \Pi(\mathbf{p}) : \exists s \in S : (s, \mathbf{p}) \in Y \wedge s \xrightarrow{\delta} s'\} \text{ and} \\ \text{Pred}(Y') &= \{(s, \mathbf{p}) \in S \times \mathcal{P} \mid \exists \delta \in \Pi(\mathbf{p}) : \exists s' \in S : (s', \mathbf{p}) \in Y' \wedge s \xrightarrow{\delta} s'\}. \end{aligned}$$

The set FR of forward-reachable states and parameter valuations and the set BR of backward-reachable states and parameter valuations are obtained by the following fixed point computations (the index identifies the round of the fixpoint iteration):

$$\begin{aligned} \text{FR}_0 &= \{(l_0, \mathbf{0})\} \times \mathcal{P} & \text{BR}_0 &= B \times \mathcal{P} \\ \text{FR}_{i+1} &= \text{Succ}(\text{FR}_i) \cup \text{FR}_i & \text{BR}_{i+1} &= \text{Pred}(\text{BR}_i) \cup \text{BR}_i \\ \text{FR} &= \lim_i \text{FR}_i & \text{BR} &= \lim_i \text{BR}_i. \end{aligned}$$

Clearly, if there is some $(s, \mathbf{p}) \in \text{FR}_i$ then this means that state s is reached after $i \in \mathbb{N}$ forward steps for parameter valuation \mathbf{p} , which corresponds to a path $s_0 \xrightarrow{\delta_1} s_1 \xrightarrow{\delta_2} \dots \xrightarrow{\delta_i} s$, where each $\delta_1, \delta_2, \dots, \delta_i$ is in $\Pi(\mathbf{p})$. Dually, if there is some state $(s, \mathbf{p}) \in \text{BR}_i$ then this means that state s is reached after $i \in \mathbb{N}$ backward steps for parameter valuation \mathbf{p} , which corresponds to a path $s \xrightarrow{\delta_1} s_1 \xrightarrow{\delta_2} \dots \xrightarrow{\delta_i} b$, where $b \in B$ and each $\delta_1, \delta_2, \dots, \delta_i$ is in $\Pi(\mathbf{p})$.

We can obtain the feasible instantiations either by looking for parameter valuations in FR that are not paired up with bad states, or by looking for parameter valuations in BR that are not paired up with the initial state. Both constructions identify the same set of feasible instantiations.

Theorem 3. *The set*

$$G = \{\mathbf{p} \in \mathcal{P} \mid (B \times \{\mathbf{p}\}) \cap \text{FR} = \emptyset\} = \{\mathbf{p} \in \mathcal{P} \mid ((l_0, \mathbf{0}), \mathbf{p}) \notin \text{BR}\}$$

consists of exactly the feasible instantiations.

In practice, neither construction performs well. The problem is that it is difficult and expensive to maintain the correlation between parameter valuations and reachable states; typically, each parameter valuation results in a different set of states.

Instead of directly computing the precise set of parameter valuations, in the next subsection, we will present an abstraction technique that allows us to reason about approximations of parameter valuations.

4.2 The Focus Abstraction

We now consider an abstraction of the template based on a given set $P \subseteq \mathcal{P}$ of parameter valuations, which we call *focus*. We use the parameter valuations in P to obtain an over- or underapproximation of the sets FR and BR, by considering P as an equivalence class: we require that a transition must exist for some or all parameter valuations in P , respectively. In the following, we use an overapproximation for the forward construction and an underapproximation for the backward construction; obviously, all constructions can also be dualized. We obtain the following approximate successor and predecessor functions: $\overline{\text{Succ}}^P, \underline{\text{Pred}}^P : 2^S \rightarrow 2^S$ with

$$\begin{aligned} \overline{\text{Succ}}^P(Y) &= \{s' \in S \mid \exists \mathbf{p} \in P : \exists \delta \in \Pi(\mathbf{p}) : \exists s \in Y : s \xrightarrow{\delta} s'\} \text{ and} \\ \underline{\text{Pred}}^P(Y') &= \{s \in S \mid \forall \mathbf{p} \in P : \exists \delta \in \Pi(\mathbf{p}) : \exists s' \in Y' : s \xrightarrow{\delta} s'\}. \end{aligned}$$

Replacing the precise Succ and Pred operators in the fixed point construction from Subsection 4.1, we obtain two new fixed point constructions for the approximations $\overline{\text{FR}}^P$ and $\underline{\text{BR}}^P$:

$$\begin{aligned} \overline{\text{FR}}_0^P &= \{(l_0, \mathbf{0})\} & \underline{\text{BR}}_0^P &= B \\ \overline{\text{FR}}_{i+1}^P &= \overline{\text{Succ}}^P(\overline{\text{FR}}_i^P) \cup \overline{\text{FR}}_i^P & \underline{\text{BR}}_{i+1}^P &= \underline{\text{Pred}}^P(\underline{\text{BR}}_i^P) \cup \underline{\text{BR}}_i^P \\ \overline{\text{FR}}^P &= \lim_i \overline{\text{FR}}_i^P & \underline{\text{BR}}^P &= \lim_i \underline{\text{BR}}_i^P. \end{aligned}$$

Clearly, if there is some state $s \in \overline{\text{FR}}_i^P$ then this means that state s is reached after $i \in \mathbb{N}$ forward steps for a set of parameter valuations P , which corresponds to a path $s_0 \xrightarrow{\delta_1} s_1 \xrightarrow{\delta_2} \dots \xrightarrow{\delta_i} s$, where, for each δ_i , there is a $\mathbf{p}_i \in P$ such that δ_i in $\Pi(\mathbf{p}_i)$. Dually, if there is some state $s \in \underline{\text{BR}}_i^P$ then this means that state s is reached after $i \in \mathbb{N}$ backward steps for a set of parameter valuations P , which corresponds to a path $s \xrightarrow{\delta_1} s_1 \xrightarrow{\delta_2} \dots \xrightarrow{\delta_i} b$, where $b \in B$ and, for each δ_i and each $\mathbf{p} \in P$, we have δ_i in $\Pi(\mathbf{p})$.

The following lemma clarifies the relationships between the approximate and precise versions of FR and BR: $\overline{\text{FR}}^P$ overapproximates FR on P , $\underline{\text{BR}}^P$ underapproximates BR on P .

Lemma 3. *For every set $P \subseteq \mathcal{P}$ of parameter valuations, it holds that*

$$\begin{aligned} \overline{\text{FR}}^P &\supseteq \{s \in S \mid \exists \mathbf{p} \in P : (s, \mathbf{p}) \in \text{FR}\} \text{ and} \\ \underline{\text{BR}}^P &\subseteq \{s \in S \mid \exists \mathbf{p} \in P : (s, \mathbf{p}) \in \text{BR}\}. \end{aligned}$$

Combining Lemma 3 with Theorem 3, we obtain that the focus abstraction allows us to approximate the set of feasible instantiations: A set of parameter valuations P definitely represents feasible instantiations if no bad states appear in $\overline{\text{FR}}^P$. Dually, the parameter valuations in P definitely represent infeasible instantiations if the initial state appears in $\underline{\text{BR}}^P$. Hence, we obtain the following lower and upper bounds for the set of feasible instantiations.

Theorem 4. *Let G be the precise set of feasible instantiations. For every set $P \subseteq \mathcal{P}$, it holds that*

$$\{\mathbf{p} \in P \mid B \cap \overline{\text{FR}}^P = \emptyset\} \subseteq G \subseteq \{\mathbf{p} \in \mathcal{P} \mid \mathbf{p} \in P \Rightarrow (l_0, \mathbf{0}) \notin \underline{\text{BR}}^P\}.$$

In the next subsection, we will describe an automatic refinement algorithm for the Focus abstraction.

4.3 Abstraction Refinement

We now describe a refinement procedure that computes an increasingly precise approximation of the set of feasible instantiations. The procedure starts with the set \mathcal{P} of all parameter valuations, and then splits the set into smaller and smaller subsets, until either a feasible instance is found, or it is established that no feasible instance exists.

Algorithm 1 $\text{Solve}(P)$: The algorithm computes a safe subset of a given set P of parameter valuations, or returns `fail` if no safe subset exists.

```

1: if  $P = \emptyset$  then
2:   return fail
3: else if  $(l_0, \mathbf{0}) \in \underline{\text{BR}}^P$  then
4:   return fail
5: else if  $\overline{\text{FR}}^P \cap \underline{\text{BR}}^P = \emptyset$  then
6:   return  $P$ 
7: else
8:    $P_1 := \text{Refine}(P)$ 
9:    $R_1 := \text{Solve}(P_1)$ 
10:  if  $R_1 \neq \text{fail}$  then
11:    return  $R_1$ 
12:  else
13:     $P_2 := P \setminus P_1$ 
14:    return  $\text{Solve}(P_2)$ 

```

The procedure is shown as Algorithm 1. The input to the procedure is the current focus P , for which we initially use \mathcal{P} . Unless the (un)reachability of some bad state can be surely established, after each refinement step, `Solve` recurs on the refined focus. In each call of `Solve`, the set of bad states are augmented with the states in $\underline{\text{BR}}^P$. This is justified by the following lemmas, which state that the old underapproximation $\underline{\text{BR}}^P$ is a subset of the new underapproximation $\underline{\text{BR}}^{P'}$ for a refinement $P' \subset P$, and that excluding $\underline{\text{BR}}^P$ from $\overline{\text{FR}}^P$ does not affect the resulting upper bound on the feasible instantiations.

Lemma 4. *For two sets $P, P' \subseteq \mathcal{P}$ of parameter valuations such that $P' \subset P$, it holds that $\underline{\text{BR}}^P \subseteq \underline{\text{BR}}^{P'}$.*

Lemma 5. *For every set $P \subseteq \mathcal{P}$ of parameter valuations, it holds that*

$$\{\mathbf{p} \in P \mid B \cap \overline{\text{FR}}^P = \emptyset\} = \{\mathbf{p} \in P \mid \underline{\text{BR}}^P \cap \overline{\text{FR}}^P = \emptyset\}.$$

It remains to specify the function `Refine`, which is called in procedure `Solve` to find an appropriate subset of \mathcal{P} to split on. Since \mathcal{P} is finite, we could, in principle, choose any strict (and non-empty) subset of \mathcal{P} during the refinement step. In the following we describe a heuristic choice that has proved useful in practice: we choose a set of parameter valuations that are guaranteed to increase $\underline{\text{BR}}^P$ in the next iteration.

Suppose the termination conditions of procedure `Solve` are not true yet, i.e., the initial state is not in $\underline{\text{BR}}^P$ and there are still states in $\overline{\text{FR}}^P \cap \underline{\text{BR}}^P$. We choose a state $s \in \overline{\text{FR}}^P \setminus \underline{\text{BR}}^P$ and a state $s' \in \underline{\text{BR}}^P$, such that there exists a transition $\delta \in \Pi(\mathbf{p})$ that leads from s to s' for some $\mathbf{p} \in P$, but not for all $\mathbf{p} \in P$. The refinement proceeds with the parameter valuations that allow a transition from s to s' :

$$\text{Refine}(P) = \{\mathbf{p} \in P \mid \exists \delta \in \Pi(\mathbf{p}) : s \xrightarrow{\delta} s'\}$$

Since such a pair s, s' of states can be found until the termination conditions of procedure `Solve` become true, we obtain that `Refine` always ensures progress of our refinement algorithm.

Lemma 6. *For every set of parameter valuations $P \subseteq \mathcal{P}$, if $P \neq \emptyset$, $(l_0, \mathbf{0}) \notin \underline{\text{BR}}^P$, and $\overline{\text{FR}}^P \cap \underline{\text{BR}}^P \neq \emptyset$, then there is a state $s \in \overline{\text{FR}}^P \setminus \underline{\text{BR}}^P$ and a state $s' \in \underline{\text{BR}}^P$ such that*

$$\emptyset \subset \text{Refine}(P) \subset P.$$

Putting everything together, we obtain the following correctness theorem for `Solve`(\mathcal{P}), where Lemma 6 guarantees termination and Theorem 4 guarantees soundness of the result.

Theorem 5. *Called with the set \mathcal{P} of parameter valuations, Procedure `Solve`(\mathcal{P}) terminates after at most $|\mathcal{P}|$ refinement steps and either computes a feasible template instantiation or reports failure, in which case no feasible template instantiation exists.*

5 Experimental Results

In this section, we report on experimental results based on a prototype implementation of the symbolic instantiation algorithm from Section 4.

Implementation. We have implemented the symbolic instantiation algorithm from Section 4 in the SYNTHIA tool [PEM11]. SYNTHIA is a verification and synthesis tool for timed automata extended with bounded integer variables. The tool provides facilities for automatic abstraction refinement and combines reduced ordered binary decision diagrams (ROBDDs) [Bry86] with difference bound matrices (DBMs) [Dil89] to obtain symbolic state representations for both discrete and continuous state components.

While standard SYNTIA already includes a game-based synthesis algorithm for timed controllers with complete observability, we only use the verification functionality of SYNTIA for template-based synthesis. For a template instantiation problem given by a controller template (\mathcal{T}, P, Π) and a set of bad states, we encode the Boolean parameters P by global integer variables whose initial values are left undefined.

The algorithm from Section 4 is realized as a specialized refinement procedure for SYNTIA’s standard (location-based) abstraction refinement loop. Whenever an edge for refinement is found, we identify the parameter valuations associated with that edge and split the global abstraction with these valuations. In the subsequent refinement step, we focus (i.e., we restrict the forward exploration) on the identified parameters of the last refinement.

Benchmarks. In the *Chinese Juggler* benchmark [LBS], a performer needs to stabilize spinning plates to prevent them from falling. After a certain amount of time has passed since a plate was stabilized, it can nondeterministically become unstable. If no restabilization takes place, it ultimately falls down. The plates have different sizes, and hence, different times to become unstable. It takes the performer one time unit to stabilize a certain plate. During that time, he cannot stabilize another plate. The controller synthesis task consists in finding a safe strategy for the performer such that no plate will ever fall down. The benchmark size is parametrized in the number of plates n . For the template-based synthesis, we use a generic cyclic-executive template (cf. Section 1) with n locations. In each step of the cyclic execution, the controller decides which plate should be stabilized next.

In the *Dam* benchmark, a controller is to be synthesized that determines the speed of the inflow to a dam. The controller can either stop the inflow or choose between a slow or a fast inflow speed. The bounded reachability requirement is that the fill level should reach a certain value between a minimal and maximal bound. While a fast speed might reach the desired fill level more quickly, the variance of the actual inflow is larger so that the maximal level might be exceeded. On the other hand, being in slow mode, it takes longer to reach the desired fill level, but the variance is not so high so that it is always possible to exactly reach a desired fill level. Thus, being in one mode all the time is not feasible, since a feasible controller must alternate between fast and slow at least once to fulfill the requirement. The benchmark size is parametrized in the degree of precision in which the fill level and the inflow amount is digitized. For the template-based synthesis, we use a controller template that models a parametric two-phase program: in the first phase, a certain inflow speed is set until a threshold of the current fill level is passed. Then, the controller enters the second phase with a possibly different speed. The controller stops as soon as a desired fill level is reached. The first and the second speed, as well as the phase-switching threshold are parameters, for which feasible instantiations are to be found.

Results. We compare the performance of the template-based extension of SYNTIA against standard SYNTIA and UPPAAL-TIGA 0.16 [BCD⁺07]. In Table 1,

Table 1. Experimental evaluation of template-based synthesis. We compare the performance of the template-based extension of SYNTHIA against standard SYNTHIA and UPPAAL-TIGA.

Benchmark	Template-based SYNTHIA				Standard SYNTHIA				UPPAAL-TIGA		
	Steps	Abs	Time	Mem	Steps	Abs	Time	Mem	States	Time	Mem
Juggler 2	6	19	0	53	2	5	1	53	57	0	6
Juggler 3	38	136	0	61	6	9	1	61	477	0	6
Juggler 4	110	421	2	87	TIMEOUT				6755	6	57
Juggler 5	423	1899	59	247	TIMEOUT				81292	1095	79
Juggler 6	1445	8335	1932	1335	TIMEOUT				TIMEOUT		
Juggler 7	TIMEOUT				TIMEOUT				TIMEOUT		
Dam 5	58	100	1	80	230	149	4	80	88592	2	65
Dam 25	268	380	13	87	1115	718	1182	91	3114648	307	443
Dam 50	530	730	87	105	TIMEOUT				13545848	5018	2355
Dam 75	793	1080	329	111	TIMEOUT				TIMEOUT		
Dam 100	1055	1430	927	143	TIMEOUT				TIMEOUT		
Dam 125	1318	1780	1949	149	TIMEOUT				TIMEOUT		
Dam 150	1580	2130	3483	153	TIMEOUT				TIMEOUT		
Dam 175	1843	2480	5127	213	TIMEOUT				TIMEOUT		
Dam 200	TIMEOUT				TIMEOUT				TIMEOUT		

the columns show, from left to right, the benchmark instance, the number of refinement steps and abstract locations in the final abstraction of the parameter synthesis algorithm, the running time and memory consumption of our template-based implementation, the performance of SYNTHIA’s standard controller synthesis algorithm, UPPAAL-TIGA’s number of explored states, running time and memory consumption. For UPPAAL-TIGA, we measured the performance for various parameters and always selected the best results. Running times are given in seconds, memory consumption in MB, the time limit was set to 2 hours, and the memory limit was set to 4 GB. All experiments were conducted on a 2.6 GHz AMD Opteron computer running Ubuntu 10.04.

For both benchmarks, template-based SYNTHIA clearly outperforms the game-based synthesis techniques implemented in standard SYNTHIA and UPPAAL-TIGA. A closer look at the *Chinese Juggler* example reveals that a major source of complexity results from the subtraction operation that occurs in the backwards computation of the winning states. Subtraction is expensive because it does not preserve convexity, and therefore requires a split into multiple zones. The much better performance of template-based synthesis is due to the fact that template-based synthesis is based on model checking, rather than game solving, and model checking does not require such nonconvex operations. In the *Dam* example, we observe that the size of the abstraction, and, thus, the running time, of the template-based approach increases polynomially in the size of the benchmark, while both standard SYNTHIA and UPPAAL-TIGA suffer from an exponential blow-up.

6 Conclusion

Our results demonstrate that template-based synthesis is an attractive alternative to the standard game-based approach to timed synthesis. Template-based synthesis has the better worst-case complexity, is easier to implement with symbolic data structures such as DBMs, and produces nicely structured controllers with a small number of locations.

In future work, we plan to expand the class of templates considered by the synthesis algorithm. Particularly interesting is the introduction of parameters in the clock constraints. Results from parametric timed model checking [AHV93] indicate that the analysis of such templates is in general undecidable. However, subclasses of parametric timed automata, such as L/U automata [HRSV01], for which the emptiness problem is decidable, are promising candidates for a more expressive and yet computationally feasible class of templates.

The long-term goal is to obtain a succinct but comprehensive library of standard templates that serves as a basis for a fully automatic template-based synthesis approach.

Acknowledgments. The authors would like to thank Alexandre David for pointing out the *Chinese Juggler* benchmark and Christoph Scholl for helpful comments on an early draft of the paper.

This work was supported by the German Research Foundation (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS).

References

- AD94. Rajeev Alur and David L. Dill, *A theory of timed automata*, Theoretical Computer Science **126** (1994), no. 2, 183–235.
- AHV93. Rajeev Alur, Thomas A. Henzinger, and Moshe Y. Vardi, *Parametric real-time reasoning*, STOC, 1993, pp. 592–601.
- AMPS98. Eugene Asarin, Oded Maler, Amir Pnueli, and Joseph Sifakis, *Controller synthesis for timed automata*, Proc. 5th IFAC Conference on System Structure and Control (J.-F. Lafay, ed.), Elsevier, 1998, pp. 469–474.
- AT02. Karine Altisen and Stavros Tripakis, *Tools for controller synthesis of timed systems*, 2nd Workshop on Real-Time Tools (RT-TOOLS), 2002.
- BCD⁺07. Gerd Behrmann, Agnès Cougnard, Alexandre David, Emmanuel Fleury, Kim Guldstrand Larsen, and Didier Lime, *UPPAAL-Tiga: Time for playing games!*, Proc. 19th International Conference on Computer Aided Verification (CAV’07) (Werner Damm and Holger Hermanns, eds.), Lecture Notes in Computer Science, vol. 4590, Springer, 2007, pp. 121–125.
- BDMP03. Patricia Bouyer, Deepak D’Souza, P. Madhusudan, and Antoine Petit, *Timed control with partial observability*, CAV (Warren A. Hunt Jr. and Fabio Somenzi, eds.), Lecture Notes in Computer Science, vol. 2725, Springer, 2003, pp. 180–192.
- Bry86. Randal E. Bryant, *Graph-based algorithms for boolean function manipulation*, IEEE Trans. Computers **35** (1986), no. 8, 677–691.

- CDF⁺05. Franck Cassez, Alexandre David, Emmanuel Fleury, Kim Guldstrand Larsen, and Didier Lime, *Efficient on-the-fly algorithms for the analysis of timed games*, Proc. 16th International Conference on Concurrency Theory (CONCUR'05) (Martín Abadi and Luca de Alfaro, eds.), Lecture Notes in Computer Science, vol. 3653, Springer, 2005, pp. 66–80.
- CDL⁺07. Franck Cassez, Alexandre David, Kim Guldstrand Larsen, Didier Lime, and Jean-François Raskin, *Timed control with observation based and stuttering invariant strategies*, in Namjoshi et al. [NYHO07], pp. 192–206.
- CDL09. Thomas Chatain, Alexandre David, and Kim G. Larsen, *Playing games with timed games*, Proceedings of the 3rd IFAC Conference on Analysis and Design of Hybrid Systems (ADHS'09) (Zaragoza, Spain) (Alessandro Giua, Manuel Silva, and Janan Zaytoon, eds.), September 2009.
- Dil89. David L. Dill, *Timing assumptions and verification of finite-state concurrent systems*, Automatic Verification Methods for Finite State Systems (Joseph Sifakis, ed.), Lecture Notes in Computer Science, vol. 407, Springer, 1989, pp. 197–212.
- DM02. Deepak D'Souza and P. Madhusudan, *Timed control synthesis for external specifications*, STACS (Helmut Alt and Afonso Ferreira, eds.), Lecture Notes in Computer Science, vol. 2285, Springer, 2002, pp. 571–582.
- Ehl10. Rüdiger Ehlers, *Symbolic bounded synthesis*, CAV (Tayssir Touili, Byron Cook, and Paul Jackson, eds.), Lecture Notes in Computer Science, vol. 6174, Springer, 2010, pp. 365–379.
- EMP10. Rüdiger Ehlers, Robert Mattmüller, and Hans-Jörg Peter, *Combining symbolic representations for solving timed games*, FORMATS (Krishnendu Chatterjee and Thomas A. Henzinger, eds.), Lecture Notes in Computer Science, vol. 6246, Springer, 2010, pp. 107–121.
- FJR09. Emmanuel Filiot, Naiyong Jin, and Jean-François Raskin, *An antichain algorithm for ltl realizability*, CAV (Ahmed Bouajjani and Oded Maler, eds.), Lecture Notes in Computer Science, vol. 5643, Springer, 2009, pp. 263–277.
- FS07. Bernd Finkbeiner and Sven Schewe, *SMT-based synthesis of distributed systems*, Proceedings of the 2nd Workshop on Automated Formal Methods (AFM 2007), 6 November, Atlanta, Georgia, USA, ACM Press, 2007, pp. 69–76.
- HK99. Thomas A. Henzinger and Peter W. Kopke, *Discrete-time control for rectangular hybrid automata*, Theoretical Computer Science **221** (1999), no. 1-2, 369–392.
- HRSV01. Thomas Hune, Judi Romijn, Mariëlle Stoelinga, and Frits W. Vaandrager, *Linear parametric model checking of timed automata*, TACAS (Tiziana Margaria and Wang Yi, eds.), Lecture Notes in Computer Science, vol. 2031, Springer, 2001, pp. 189–203.
- LBS. Kim G. Larsen, Gerd Behrmann, and Arne Skou, *Exercises for Uppaal*, <http://www.cs.aau.dk/~bnielsen/TOV08/ESV04/exercises>.
- LV09. Yoad Lustig and Moshe Y. Vardi, *Synthesis from component libraries*, FOSACS (Luca de Alfaro, ed.), LNCS, vol. 5504, 2009, pp. 395–409.
- MPS95. Oded Maler, Amir Pnueli, and Joseph Sifakis, *On the synthesis of discrete controllers for timed systems (an extended abstract)*, Proc. 12th Annual Symposium on Theoretical Aspects of Computer Science (STACS'95) (Ernst W. Mayr and Claude Puech, eds.), Lecture Notes in Computer Science, vol. 900, Springer, 1995, pp. 229–242.

- NYHO07. Kedar S. Namjoshi, Tomohiro Yoneda, Teruo Higashino, and Yoshio Okamura (eds.), *Automated technology for verification and analysis, 5th international symposium, atva 2007, tokyo, japan, october 22-25, 2007, proceedings*, Lecture Notes in Computer Science, vol. 4762, Springer, 2007.
- PEM11. Hans-Jörg Peter, Rüdiger Ehlers, and Robert Mattmüller, *Synthia: Verification and synthesis for timed automata*, CAV (Ganesh Gopalakrishnan and Shaz Qadeer, eds.), Lecture Notes in Computer Science, vol. 6806, Springer, 2011, pp. 649–655.
- PM09. Hans-Jörg Peter and Robert Mattmüller, *Component-based abstraction refinement for timed controller synthesis*, IEEE Real-Time Systems Symposium (Theodore P. Baker, ed.), IEEE Computer Society, 2009, pp. 364–374.
- SF07. Sven Schewe and Bernd Finkbeiner, *Bounded synthesis*, in Namjoshi et al. [NYHO07], pp. 474–488.
- Tho95. Wolfgang Thomas, *On the synthesis of strategies in infinite games*, STACS, 1995, pp. 1–13.

A Proofs

Theorem 2. The controller template instantiation problem for a controller template (\mathcal{T}, P, Π) and a set of bad states B is PSPACE-complete.

Proof. We assume $\mathcal{T} = (L, l_0, \Sigma, \Delta, X, I)$. PSPACE-hardness follows immediately from Theorem 1 as the original reachability problem for timed automata is just a special case of the instantiation problem: choose Π in such a way that $\Pi(\mathbf{p}) = \Delta$, for any \mathbf{p} .

Membership in $\text{NPTIME}^{\text{PSPACE}} \subseteq \text{NPSpace}$, which is known to coincide with PSPACE, is established through the following nondeterministic algorithm. First, observe that the number of possible controllers is bounded by the number of distinct parameter valuations $|\mathcal{P}| = 2^{|P|}$. Hence, some controller can be represented using only $|P|$ bits. Our algorithm now proceeds as follows. It first guesses the bits of a feasible controller instantiation \mathcal{I} in linear time and then verifies the guess by model checking \mathcal{I} , which, according to Theorem 1, is in PSPACE. This concludes the proof of the upper bound. \square

Theorem 3. The set

$$G = \{\mathbf{p} \in \mathcal{P} \mid B \times \{\mathbf{p}\} \cap \text{FR} = \emptyset\} = \{\mathbf{p} \in \mathcal{P} \mid ((l_0, \mathbf{0}), \mathbf{p}) \notin \text{BR}\}$$

consists of exactly the feasible instantiations.

Proof. It is easy to see that $G_f = \{\mathbf{p} \in \mathcal{P} \mid B \times \{\mathbf{p}\} \cap \text{FR} = \emptyset\}$ comprises exactly those parameter valuations for which no bad state is forward reachable. Hence, by definition, $G = G_f$.

On the other hand, because for every forward reachable state s , the initial state s_0 is also backward reachable from s , we have that $G_b = \{\mathbf{p} \in \mathcal{P} \mid ((l_0, \mathbf{0}), \mathbf{p}) \notin \text{BR}\}$ coincides with G_f . Hence, $G = G_f = G_b$. \square

Lemma 3. For every set $P \subseteq \mathcal{P}$ of parameter valuations, it holds that

$$\begin{aligned} \overline{\text{FR}}^P &\supseteq \{s \in S \mid \exists \mathbf{p} \in P : (s, \mathbf{p}) \in \text{FR}\} \text{ and} \\ \underline{\text{BR}}^P &\subseteq \{s \in S \mid \exists \mathbf{p} \in P : (s, \mathbf{p}) \in \text{BR}\}. \end{aligned}$$

Proof. We first show $\overline{\text{FR}}^P \supseteq Y = \{s \in S \mid \exists \mathbf{p} \in P : (s, \mathbf{p}) \in \text{FR}\}$. By definition, for every state s contained in Y , there is a $\mathbf{p} \in P$ such that there is a path

$$s_0 \xrightarrow{\delta_0} s_1 \xrightarrow{\delta_1} \dots \xrightarrow{\delta_n} s,$$

where, for each $0 \leq i \leq n$, $\delta_i \in \Pi(\mathbf{p})$. On the other hand, by definition, for every state s contained in $\overline{\text{FR}}^P$, there is a path

$$s_0 \xrightarrow{\delta_0} s_1 \xrightarrow{\delta_1} \dots \xrightarrow{\delta_n} s,$$

where, for each δ_i , $0 \leq i \leq n$, there is a $\mathbf{p}_i \in P$ such that δ_i in $\Pi(\mathbf{p}_i)$. Now, when we fix each $\mathbf{p}_i = \mathbf{p}$, for some $\mathbf{p} \in P$, then we get Y . Hence, $\overline{\text{FR}}^P \supseteq Y$.

Secondly, we show $\underline{\text{BR}}^P \subseteq Y = \{s \in S \mid \exists \mathbf{p} \in P : (s, \mathbf{p}) \in \text{BR}\}$. By definition, for every state s contained in Y , there is a $b \in B$ and a $\mathbf{p} \in P$ such that there is a path

$$s \xrightarrow{\delta_0} s_1 \xrightarrow{\delta_1} \dots \xrightarrow{\delta_n} b,$$

where, for each $0 \leq i \leq n$, $\delta_i \in \Pi(\mathbf{p})$. On the other hand, by definition, for every state s contained in $\underline{\text{BR}}^P$, there is a $b \in B$ such that there is a path

$$s \xrightarrow{\delta_0} s_1 \xrightarrow{\delta_1} \dots \xrightarrow{\delta_n} b,$$

where, for each δ_i , $0 \leq i \leq n$, and each $\mathbf{p} \in P$, we have $\delta_i \in \Pi(\mathbf{p})$. Now, it is easy to see that every state s in $\underline{\text{BR}}^P$, is backward reachable for an arbitrary $\mathbf{p} \in P$, and thus, also contained in Y . Hence, $\underline{\text{BR}}^P \subseteq Y$. \square

Theorem 4. Let G be the precise set of feasible instantiations. For every set $P \subseteq \mathcal{P}$,

$$\{\mathbf{p} \in P \mid B \cap \overline{\text{FR}}^P = \emptyset\} \subseteq G \subseteq \{\mathbf{p} \in \mathcal{P} \mid \mathbf{p} \in P \Rightarrow (l_0, \mathbf{0}) \notin \underline{\text{BR}}^P\}.$$

Proof. We first show that $Y = \{\mathbf{p} \in P \mid B \cap \overline{\text{FR}}^P = \emptyset\} \subseteq G$. By Lemma 3, $\overline{\text{FR}}^P$ overapproximates those states that are precisely forward reachable for P . By definition of Y , either (1) $Y = \emptyset \subseteq G$, if $B \cap \overline{\text{FR}}^P \neq \emptyset$, or (2) $Y = P$, if $B \cap \overline{\text{FR}}^P = \emptyset$, i.e., surely no bad state is reachable for P . By Theorem 3, a parameter valuation is feasible if no bad state is forward reachable. Hence, also in case (2), $Y = P \subseteq G$.

Secondly, we prove that $G \subseteq Y = \{\mathbf{p} \in \mathcal{P} \mid \mathbf{p} \in P \Rightarrow (l_0, \mathbf{0}) \notin \underline{\text{BR}}^P\}$ by showing that every $\mathbf{p} \in G$ is also contained in Y . We distinguish two cases.

Case 1: assume $\mathbf{p} \in G \cap P$. By Theorem 3, since $\mathbf{p} \in G$, the initial state is not backward reachable for \mathbf{p} . By Lemma 3, $\underline{\text{BR}}^P$ underapproximates those states that are precisely backward reachable for P . Since $\mathbf{p} \in P$, we have that the premise and the implicant of the definition of Y are true, and thus, \mathbf{p} is also contained in Y .

Case 2: assume $\mathbf{p} \in G \setminus P$. \mathbf{p} violates the premise in the implication of the definition of Y , and thus, is trivially contained in Y .

Hence, $G \subseteq Y$. \square

Lemma 4. For two sets $P, P' \subseteq \mathcal{P}$ of parameter valuations such that $P' \subset P$, it holds that $\underline{\text{BR}}^P \subseteq \underline{\text{BR}}^{P'}$.

Proof. We assume $P' \subset P$ and prove $\underline{\text{BR}}^P \subseteq \underline{\text{BR}}^{P'}$ by showing that every state in $\underline{\text{BR}}^P$ is also contained in $\underline{\text{BR}}^{P'}$. By definition, for every state $s \in \underline{\text{BR}}^P$, there is a $b \in B$ such that there is a path

$$s \xrightarrow{\delta_0} s_1 \xrightarrow{\delta_1} \dots \xrightarrow{\delta_n} b,$$

where, for each δ_i , $0 \leq i \leq n$, and each $\mathbf{p} \in P$, we have $\delta_i \in \Pi(\mathbf{p})$. Now, it is easy to see that the same path also exists for the smaller set $P' \subset P$. \square

Lemma 5. For every set $P \subseteq \mathcal{P}$ of parameter valuations, it holds that

$$\{\mathbf{p} \in P \mid B \cap \overline{\text{FR}}^P = \emptyset\} = \{\mathbf{p} \in P \mid \underline{\text{BR}}^P \cap \overline{\text{FR}}^P = \emptyset\}.$$

Proof. Let $Y = \{\mathbf{p} \in P \mid B \cap \overline{\text{FR}}^P = \emptyset\}$ and $Y' = \{\mathbf{p} \in P \mid \underline{\text{BR}}^P \cap \overline{\text{FR}}^P = \emptyset\}$. We distinguish two cases:

Case 1: assume $B \cap \overline{\text{FR}}^P \neq \emptyset$. In this case, since $B \subseteq \underline{\text{BR}}^P$, $Y = Y' = \emptyset$.

Case 2: assume $B \cap \overline{\text{FR}}^P = \emptyset$. Observe that, for each state $s \in \underline{\text{BR}}^P$, there is a path from s to some state $b \in B$, for all parameter valuations in P . On the other hand, by assumption, $\overline{\text{FR}}^P$ does not comprise any states leading to some bad state for some parameter valuation in P . Hence, $\underline{\text{BR}}^P \cap \overline{\text{FR}}^P = \emptyset$, and thus, $Y = Y'$. \square

Lemma 6. For every set of parameter valuations $P \subseteq \mathcal{P}$, if $P \neq \emptyset$, $(l_0, \mathbf{0}) \notin \underline{\text{BR}}^P$, and $\overline{\text{FR}}^P \cap \underline{\text{BR}}^P \neq \emptyset$, then there is a state $s \in \overline{\text{FR}}^P \setminus \underline{\text{BR}}^P$ and a state $s' \in \underline{\text{BR}}^P$ such that

$$\emptyset \subset \text{Refine}(P) \subset P.$$

Proof. By assumption, since $\overline{\text{FR}}^P \cap \underline{\text{BR}}^P \neq \emptyset$, there is a path from $s_0 = (l_0, \mathbf{0})$ to some state $b \in \underline{\text{BR}}^P$ of the form

$$s_0 \xrightarrow{\delta_0} s_1 \xrightarrow{\delta_1} \dots s_n \xrightarrow{\delta_n} b,$$

where, for each δ_i , $0 \leq i \leq n$, there is a $\mathbf{p}_i \in P$ such that $\delta_i \in \Pi(\mathbf{p}_i)$. Since $s_0 \notin \underline{\text{BR}}^P$, there is surely some $s_i \in \overline{\text{FR}}^P \setminus \underline{\text{BR}}^P$, whose successor s_{i+1} is in $\underline{\text{BR}}^P$. In fact, if b is the first state from $\underline{\text{BR}}^P$ which appears on the path, then a candidate edge for refinement is $s_n \xrightarrow{\delta_n} b$, i.e., $s = s_n$ and $s' = b$. Thus, $\emptyset \subset \text{Refine}(P)$ as $\text{Refine}(P)$ contains those $\mathbf{p} \in P$, for which $\delta_n \in \Pi(\mathbf{p})$. On the other hand, if $\text{Refine}(P)$ subsumes exactly those $\mathbf{p} \in P$, for which $\delta_n \in \Pi(\mathbf{p})$, we can conclude that $\text{Refine}(P) \subset P$, as s_n would be contained in $\underline{\text{BR}}^P$ in the first place if δ_n would be available for all $\mathbf{p} \in P$. \square

Theorem 5. Called with the set \mathcal{P} of parameter valuations, Procedure $\text{Solve}(\mathcal{P})$ terminates after at most $|\mathcal{P}|$ refinement steps and either computes a feasible template instantiation or reports failure, in which case no feasible template instantiation exists.

Proof. Termination of $\text{Solve}(\mathcal{P})$ is easily established by applying Lemma 6 which ensures progress of each refinement step. Since there are only finitely many parameter valuations in \mathcal{P} , the focus P will eventually, after at most $|\mathcal{P}|$ refinement steps, end up with \emptyset or some parameter valuations for which one of the termination conditions becomes true.

Soundness of $\text{Solve}(\mathcal{P})$ follows immediately from Theorem 4.

Completeness of $\text{Solve}(\mathcal{P})$ is established by the fact that Solve does not miss any parameter valuation: in the lines 9 and 14 of Algorithm 1, Solve recurs on the refinement P_1 and on $P \setminus P_1$. \square