

# Parameter Synthesis with IC3

Alessandro Cimatti    Alberto Griggio\*    Sergio Mover    Stefano Tonetta  
Fondazione Bruno Kessler, Trento, Italy  
{cimatti, griggio, mover, tonettas}@fbk.eu

**Abstract**—Parametric systems arise in different application domains, such as software, cyber-physical systems or tasks scheduling. A key challenge is to estimate the values of parameters that guarantee the desired behaviours of the system.

In this paper, we propose a novel approach based on an extension of the IC3 algorithm for infinite-state transition systems. The algorithm finds the feasible region of parameters by complement, incrementally finding and blocking sets of “bad” parameters which lead to system failures. If the algorithm terminates we obtain the precise region of feasible parameters of the system.

We describe an implementation for symbolic transition systems with linear constraints and perform an experimental evaluation on benchmarks taken from the domain of hybrid systems. The results demonstrate the potential of the approach.

## I. INTRODUCTION

Parametric systems arise in many application domains from real-time systems to software to cyber-physical systems. In these applications, the system is often part of a larger environment, and the designer has to define the system relative to some unknown parameters of the environment. The design of a robust system requires the verification to not rely on concrete values for the parameters but to prove the correctness of the system for a certain region of values. The use of parameters is fundamental in the early phases of the development, giving the possibility to explore different design choices. In fact, a parametric system represents a set of (non-parametric) systems, one for each valuation of the parameters.

A key challenge for the design of parametric systems is the estimation of the parameter valuations that guarantee the correct behavior of the system. Manual estimation of these values is time consuming and does not find optimal solutions for specific design problems. Therefore, a fundamental problem is to automatically synthesize the maximal region of parameter valuations for which the system satisfies some properties.

In this paper, we focus on the verification of invariant properties and how to extend the SMT-based algorithms to solve the synthesis problem. The general approach works by complement, building the set of “bad” parameter valuations. It relies on the enumeration of counterexamples violating the properties, extracting from the counterexample a region of bad parameter valuations by quantification of the state variables.

The novel contribution of this paper is a new synthesis algorithm based on IC3, one of the major recent breakthroughs in SAT-based model checking, and lately extended to the SMT

case. The key idea of the synthesis algorithm is to exploit the features of IC3. First, IC3 may find a set of counterexamples consisting of a sequence of set of states  $s_0, \dots, s_k$ , where each state in  $s_i$  is guaranteed to reach some of the bad states in  $s_k$  in  $k - i$  steps; this is exploited in the expensive quantification of the state-variables, that can be performed on shortest, and thus more amenable, counterexamples. Second, the internal structures of IC3 allows our extension to be integrated in a fully incremental fashion, never restarting the search from scratch to find a new counterexample.

Various approaches already solve the parameter synthesis problem for several kind of systems, like infinite-state transition systems [4], timed and hybrid automata [10], [12], [9], [7], [1], [2]. The advantages of the new algorithm are that it synthesizes an optimal region of parameters (unlike [9], [1]), it is incremental and applies quantifier elimination only to small formulas (unlike [9], [7]), and it avoids computing the whole set of the reachable states (unlike [10], [12]).

We implemented the algorithm for symbolic transition systems with linear constraints and performed an experimental evaluation on benchmarks on timed and hybrid systems. We compared the approach with similar SMT-based techniques and with techniques based on the computation of the reachable states. The results show the potential of the approach.

## II. BACKGROUND

### A. Transition Systems

A *transition system*  $S$  is a tuple  $S = \langle X, I, T \rangle$  where  $X$  is a set of (state) variables,  $I(X)$  is a formula representing the initial states, and  $T(X, X')$  is a formula representing the transitions. In this paper, we shall deal with *linear rational arithmetic* formulas, that is, Boolean combinations of propositional variables and linear inequalities over rational variables. A *state* of  $S$  is an assignment to the variables  $X$ . A *path* of  $S$  is a finite sequence  $s_0, s_1, \dots, s_k$  of states such that  $s_0 \models I$  and for all  $i$ ,  $0 \leq i < k$ ,  $s_i, s'_{i+1} \models T$ . Given a formula  $P(X)$ , the *verification problem* denoted with  $S \models P$  is the problem to check if for all paths  $s_0, s_1, \dots, s_k$  of  $S$ , for all  $i$ ,  $0 \leq i \leq k$ ,  $s_i \models P$ . The dual problem is the *reachability problem*, which is the problem to find a path  $s_0, s_1, \dots, s_k$  of  $S$  such that  $s_k \models \neg P$ .  $P(X)$  represents the “good” states, while  $\neg P$  represents the “bad” states.

### B. Parameter Synthesis

In parametric systems, besides the standard constants, the formulas can include also *parameters*, which are rigid symbols

\* Supported by Provincia Autonoma di Trento and the European Community’s FP7/2007-2013 under grant agreement Marie Curie FP7 - PCOFUND-GA-2008-226070 “progetto Trentino”, project ADAPTATION.

with “unknown” values. Let  $U$  be the set of parameters. A *parameter valuation* is an assignment to the parameters. Given a formula  $\phi$  and a parameter valuation  $\gamma$ , we denote with  $\gamma(\phi)$  the formula obtained from  $\phi$  by replacing each parameter in  $U$  with the assignment given by  $\gamma$ .

A *parametric transition system*  $S$  is a tuple  $S = \langle U, X, I, T \rangle$  where  $U$  is the set of parameters,  $X$  is the set of variables,  $I(U, X)$  is the initial formula, and  $T(U, X, X')$  is the transition formula. Each parameter valuation  $\gamma$  induces a transition system  $S_\gamma = \langle X, \gamma(I), \gamma(T) \rangle$ .

Given a parametric transition system  $S = \langle U, X, I, T \rangle$  and a formula  $P(U, X)$ , we say that a parameter valuation  $\gamma$  is *feasible* iff  $S_\gamma \models \gamma(P)$ . The *parameter synthesis problem* is the problem of finding a set  $\rho(U)$  of feasible parameter valuations (i.e., for every  $\gamma \in \rho$ ,  $S_\gamma \models \gamma(P)$ ). A set of feasible parameter valuations  $\rho(U)$  is *optimal* if it contains all the feasible parameter valuations.

### C. IC3 with SMT

IC3 [3] is an efficient algorithm for the verification of finite-state systems, with Boolean state variables and propositional logic formulas. IC3 was subsequently extended to the SMT case in [5], [11]. In the following, we present its main ideas, following the description of [5]. For brevity, we have to omit several important details, for which we refer to [3], [5], [11].

Let  $S$  and  $P$  be a transition system and a set of good states as in §II-A. The IC3 algorithm tries to prove that  $S \models P$  by finding a formula  $F(X)$  such that: (i)  $I(X) \models F(X)$ ; (ii)  $F(X) \wedge T(X, X') \models F(X')$ ; and (iii)  $F(X) \models P(X)$ .

In order to construct an inductive invariant  $F$ , IC3 maintains a sequence of formulas (called *trace*)  $F_0(X), \dots, F_k(X)$  such that: (i)  $F_0 = I$ ; (ii)  $F_i \models F_{i+1}$ ; (iii)  $F_i(X) \wedge T(X, X') \models F_{i+1}(X')$ ; (iv) for all  $i < k$ ,  $F_i \models P$ .

The algorithm proceeds incrementally, by alternating two phases: a blocking phase, and a propagation phase. In the *blocking* phase, the trace is analyzed to prove that no intersection between  $F_k$  and  $\neg P(X)$  is possible. If such intersection cannot be disproved on the current trace, the property is violated and a counterexample can be reconstructed. During the blocking phase, the trace is enriched with additional formulas, that can be seen as strengthening the approximation of the reachable state space. At the end of the blocking phase, if no violation is found,  $F_k \models P$ .

The *propagation* phase tries to extend the trace with a new formula  $F_{k+1}$ , moving forward the clauses from preceding  $F_i$ 's. If, during this process, two consecutive elements of the trace (called *frames*) become identical (i.e.  $F_i = F_{i+1}$ ), then a fixpoint is reached, and IC3 terminates with  $F_i$  being an inductive invariant proving the property.

In the *blocking* phase IC3 maintains a set of pairs  $(s, i)$ , where  $s$  is a set of states that can lead to a bad state, and  $i > 0$  is a position in the current trace. New formulas (in the form of clauses) to be added to the current trace are derived by (recursively) proving that a set  $s$  of a pair  $(s, i)$  is unreachable starting from the formula  $F_{i-1}$ . This is done by checking the

satisfiability of the formula:

$$F_{i-1} \wedge \neg s \wedge T \wedge s'. \quad (1)$$

If (1) is unsatisfiable, and  $s$  does not intersect the initial states  $I$  of the system, then  $\neg s$  is *inductive relative to*  $F_{i-1}$ , and IC3 strengthens  $F_i$  by adding  $\neg s$  to it<sup>1</sup>, thus *blocking* the bad state  $s$  at  $i$ . If, instead, (1) is satisfiable, then the overapproximation  $F_{i-1}$  is not strong enough to show that  $s$  is unreachable. In this case, let  $p$  be a subset of the states in  $F_{i-1} \wedge \neg s$  such that all the states in  $p$  lead to a state in  $s'$  in one transition step. Then, IC3 continues by trying to show that  $p$  is not reachable in one step from  $F_{i-2}$  (that is, it tries to block the pair  $(p, i-1)$ ). This procedure continues recursively, possibly generating other pairs to block at earlier points in the trace, until either IC3 generates a pair  $(q, 0)$ , meaning that the system does not satisfy the property, or the trace is eventually strengthened so that the original pair  $(s, i)$  can be blocked.

A key difference between the original Boolean IC3 and its SMT extensions in [5], [11] is in the way sets of states to be blocked or generalized are constructed. In the blocking phase, when trying to block a pair  $(s, i)$ , if the formula (1) is satisfiable, then a new pair  $(p, i-1)$  has to be generated such that  $p$  is a cube in the *preimage of*  $s$  wrt.  $T$ . In the propositional case,  $p$  can be obtained from the model  $\mu$  of (1) generated by the SAT solver, by simply dropping the primed variables occurring in  $\mu$ . This cannot be done in general in the first-order case, where the relationship between the current state variables  $X$  and their primed version  $X'$  is encoded in the theory atoms, which in general cannot be partitioned into a primed and an unprimed set. The solution proposed in [5] is to compute  $p$  by existentially quantifying (1) and then applying an *under-approximated* existential elimination algorithm for linear rational arithmetic formulas. Similarly, in [11] a theory-aware generalization algorithm for linear rational arithmetic (based on interpolation) was proposed, in order to strengthen  $\neg s$  before adding it to  $F_i$  after having successfully blocked it.

## III. PARAMETER SYNTHESIS WITH IC3

### A. Solving the synthesis problem with reachability

A naive approach to synthesize the set of parameters  $\rho(U)$  is to incrementally find the complement set  $\beta(U)$  (thus,  $\rho = \neg\beta$ ) of unfeasible parameter valuations, rephrasing the problem as a reachability problem for a transition system  $S_\rho$  and iteratively removing the counterexamples to  $S_\rho \models P$ .

More specifically, given the parametric transition system  $S = \langle U, X, I, T \rangle$ , the algorithm keeps an over-approximation  $\rho(U)$ , initially true, of the safe region. The encoding of  $S$  is the transition system  $S_\rho = \langle X \cup P, I_\rho, T_\rho \rangle$  where  $T_\rho = T \wedge \bigwedge_{p \in U} p' = p$  forces parameters to not change their value in the evolution of the system and  $I_\rho = I \wedge \rho$  restricts the parameter valuations to the over-approximation.

At every iteration, a new parameter valuation is removed from  $\rho$ . The algorithm terminates if it proves that  $S_\rho \models P$ , and  $\rho$  is the solution to the synthesis problem.

<sup>1</sup> $\neg s$  is actually *generalized* before being added to  $F_i$ . Although this is fundamental for the IC3 effectiveness, we do not discuss it for simplicity.

This simple approach does not work in the context of infinite-state transition systems, where the possible number of counterexamples and the values of the parameters are infinite. For this reason, we need an algorithm that removes a set of parameters, instead of a single point.

### B. Description of the synthesis algorithm with IC3

We embed a reasoning similar to the naive algorithm in IC3, exploiting its generalization of counterexamples and incrementality. The generalization avoids to explicitly enumerate the counterexamples, while incrementality allows to reuse all the clauses learned by IC3 across different safety checks.

Therefore, IC3 is used to prove that  $S_\rho \models P$ . If it is successful (recall that in the SMT extension, the problem is undecidable), we can conclude that  $\rho$  is a set of feasible parameters and, in particular, is optimal. Instead, if there exists a set of parameters such that  $S \not\models P$ , IC3 might find a counterexample to  $P$ . The counterexample is found in the blocking phase as a sequence  $\pi := (s_0, 0), \dots, (s_n, n)$ , where  $s_0 \models I'$ ,  $s_n \models \neg P$  and for  $0 < i < n - 1$ ,  $s_i \wedge T' \models s_{i+1}$ . Possibly,  $\pi$  does not represent a single path of the system that reaches a violation, but a set of paths that reach  $\neg P$ . This is an intrinsic feature of IC3, which generalizes the counterexamples to induction found in the blocking phase, trying to block set of states rather than a single state<sup>2</sup>. The state  $s_o$  represents a set of states that will eventually reach  $\neg P$ . Thus, we compute from  $s_o$  a set of bad parameters  $\beta_{s_o}(U)$  that will eventually reach  $s_n$ :  $\beta_{s_o}(U) := \exists X. s_o(U, X)$ . We rely on a quantifier elimination procedure to get a quantifier-free formula for  $\beta_{s_o}$ .

The algorithm refines its conjecture about the unfeasible parameters of the system. Let  $\beta' := \beta \vee \beta_{s_o}$  and  $\rho' := \rho \wedge \neg \beta_{s_o}$  be the new approximations of unfeasible and feasible regions of parameters. We have to prove that  $S_{\rho'} \models P$ . We perform the verification incrementally, reusing all the frames of IC3. Since  $\rho' := \rho \wedge \neg \beta_{s_o}$ , we have that  $S_{\rho'} = \langle X \cup P, I_\rho \wedge \neg \beta_{s_o}, T_\rho \rangle$ <sup>3</sup>. Thus, we incrementally encode  $S_{\rho'}$  strengthening the initial condition and the transition relation used in the algorithm, and also strengthening the first frame kept by the IC3 algorithm (i.e.  $F_0 := F_0 \wedge \neg(\beta_{s_o})$ ). The strengthening of  $F_0$  removes the state  $s_o$  from  $I$  (possibly blocking also other bad states).

Since  $S_\rho$  is an overapproximation of  $S_{\rho'}$ , the invariant maintained by IC3 ( $F_0 = I$ ,  $F_i \models F_{i+1}$ ,  $F_i \models P$ ,  $F_i(X) \wedge T(X, X') \models F_{i+1}(X')$ ) holds for the problem  $S_{\rho'} \models P$ .

From this point, we rely on the usual behaviour of IC3, which tries to block  $(s_1, 1)$  with the strengthened frame  $F_0$ . The algorithm terminates if either  $P$  is proved or the  $F_0$  becomes unsatisfiable, showing that  $\rho$  is empty.

*Theorem 1:* Given a parametric transition system  $S = \langle U, X, I, T \rangle$  and a formula  $P(X)$ ,  $\rho(U) := \text{PARAMIC3}(U, I, T, P)$  is the *optimal set* of feasible parameter valuations.

<sup>2</sup>We follow the IC3 formulation of [8], which shows that IC3 can find a set of counterexamples, improving its performance. Moreover, in the SMT-based IC3 [5] the approximate pre-image computes a set of states.

<sup>3</sup>We add  $\neg \beta_{s_o}$  also to  $T_\rho$ , since it is an inductive invariant of  $S_{\rho'}$ .

For lack of space, the proof is available in a technical report ([http://es.fbk.eu/people/mover/fmcad13\\_ext.pdf](http://es.fbk.eu/people/mover/fmcad13_ext.pdf)).

### C. Optimizations

We presented a version of the algorithm which computes a region of bad states  $\beta_{s_o}$  only from the initial states of  $\pi := (s_0, 0), \dots, (s_n, n)$ . However, this is only one of the possible choices, since more general regions of bad parameters can be found considering each  $s_i$  in  $\pi$ . In fact,  $\beta_{s_o}$  is one of the extreme cases, while the other one is  $\beta_n(U) := \exists X. (BMC_n)$ , which encodes the set of all the parameters that may reach  $\neg P$  in  $n$  steps, where  $BMC_n$  denotes  $I^0 \wedge \bigwedge_{i=0}^{n-1} T^i \wedge \neg P^n$ . However, the cost of eliminating the quantifiers grows as well, and it might in fact become impractical. In principle, one may consider the intermediate cases  $\beta_{s_i}$  (that is, the reachability of one of the intermediate states  $s_i$  in  $\pi$ ) to trade the generality of the result with the cost of the quantifier elimination. Furthermore, we notice that for soundness we do not need the precise set  $\beta_{s_i}$ , but we can consider its under-approximations, since this still guarantees to remove only bad parameters valuations. As an advantage, in this case the quantifier elimination problems are easier to solve and are more general than  $\beta_{s_o}$ . In practice, we use an heuristic, which we describe in the next Section, that combines the precise and the under-approximated approach, enabling us to find a trade-off between generality and the cost of quantifier elimination.

## IV. EXPERIMENTS

We have implemented the algorithm described in the previous section on top of the fully symbolic SMT-based IC3 of [5]. The tool uses MATHSAT [6] as backend SMT engine, and works on transition systems with linear arithmetic constraints.

*Evaluation.* Our evaluation consists of three parts. In the first, we compare our implementation (called PARAMIC3 in what follows) with the approach described in [7], in order to evaluate the viability of our technique when compared to other SMT-based solutions. For this, we have implemented the algorithm described in [7] using our “regular” SMT-based IC3 implementation as the backend engine for reachability checking (called ITERATIVE-BLOCK-PATH(IC3) in what follows). We remark that the tool of [7] was based only on Bounded Model Checking (BMC), and exploited domain-specific information for computing the maximum needed bound, which is not available in our more general context.

In the second part, we evaluate the effectiveness of the optimizations described in the previous section, by comparing the default heuristic used by PARAMIC3, using both the full counterexample path  $\pi$  and its initial state  $(s_0, 0)$  for blocking bad regions of parameters, with the basic strategy using only  $(s_0, 0)$  (called PARAMIC3-basic in the following). In particular, the default heuristic used by PARAMIC3 works as follows. At the beginning, only initial states  $(s_0, 0)$  of counterexample paths are used to block bad regions of parameters. If the algorithm starts enumerating too many bad regions, it starts exploiting also full paths  $\pi$ , by computing the bad region  $\beta_k^\pi(U) = \exists X. BMC_k^\pi$ , where  $k$  is the length of  $\pi$ , and  $BMC_k^\pi$




Fig. 1. Run time comparison (sec.) between PARAMIC3 and ITERATIVE-BLOCK-PATH(IC3).




Fig. 2. Run time comparison (sec.) between PARAMIC3 and PARAMIC3-basic.




Fig. 3. Run time comparison (sec.) between PARAMIC3 and RED.

is the formula encoding all the counterexample traces of length  $k$  where the values for the Boolean variables are the same as in  $\pi$ , similarly to what is done in [7]. The computation of  $\beta_k^\pi$  is aborted if it becomes too expensive<sup>4</sup>, in order to control the tradeoff between the quality of the obtained bad region and the cost of performing quantifier elimination.

Finally, in the third part of our evaluation, we compare PARAMIC3 against RED [12], a state-of-the-art tool for parameter synthesis for linear-hybrid automata.

**Benchmarks.** We have selected benchmarks used in previous work on parameter synthesis for hybrid systems. Most of them come from the suite of RED. We have a total of 92 instances from 13 different families. All the instances, the scripts and the tools used for reproducing our experiments are available at <http://es.fbk.eu/people/mover/fmcad13.tar.gz>. For the first two parts of our evaluation, we have experimented with two different ways of encoding linear hybrid automata into symbolic transition systems, resulting in a set of 192 instances. For the comparison with RED, we picked the encoding giving the best overall performance for PARAMIC3.

**Results.** We have run our experiments on a cluster of Linux machines with a 2.27GHz Xeon CPU, using a timeout of 600 seconds and a memory limit of 3Gb for each instance. Figures 1–3 show the scatter plots that compare the total run time (in seconds) of the different techniques. From the plots, we can make the following observations. (i) Our new algorithm is clearly superior to the technique of [7], both in number of completed instances and in execution time. Overall, PARAMIC3 successfully solves 5 more instances than ITERATIVE-BLOCK-PATH(IC3), and it is almost always faster. We remark that both algorithms use the same implementation of IC3 as backend, run with the same options. (ii) Our heuristic for using full counterexample paths  $\pi$  for blocking bad regions of parameters pays off for harder problems. With it, PARAMIC3 solves 6 more instances which were previously out of reach, without any overhead for the other instances. (iii) The comparison with RED shows that our technique is very promising. Although there is no clear winner, there are more instances for which PARAMIC3 outperforms RED than the converse. In general, the

two tools seem to be somewhat complementary. We remark that RED is specialized for timed and linear-hybrid automata and that most of the benchmarks we used come from its suite, whereas PARAMIC3 works for arbitrary transition systems and it is not tuned for linear hybrid systems in any way.

## V. CONCLUSIONS AND FUTURE WORK

We proposed a new algorithm based on IC3 for synthesizing an optimal region of parameter valuations guaranteeing the satisfaction of an invariant property. The algorithm exploits the features of IC3 to incrementally remove sets of bad parameter valuations and to reduce the cost of expensive quantifier elimination operations by performing them on small formulas. Our experimental results show that the new synthesis algorithm performs better than similar SMT-based techniques and is complementary to other techniques based on the computation of the reachable states. In the future, we plan to improve the algorithm by better exploiting the structure of the problem, to evaluate it in other domains such as software, and to apply it in the context of modular component-based verification.

## REFERENCES

- [1] É. André and U. Kühne. Parametric analysis of hybrid systems using HyMITATOR. In *IFM*, 2012.
- [2] G. Behrmann, K. Guldstrand Larsen, Jacob Illum Rasmussen. Beyond liveness: Efficient parameter synthesis for time bounded liveness. In *FORMATS*, 2005.
- [3] A. Bradley. Sat-based model checking without unrolling. In *VMCAI*, 2011.
- [4] R. Bruttomesso, A. Carioni, S. Ghilardi, and S. Ranise. Automated analysis of parametric timing-based mutual exclusion algorithms. In *NFM*, 2012.
- [5] A. Cimatti and A. Griggio. Software Model Checking via IC3. In *CAV*, 2012.
- [6] A. Cimatti, A. Griggio, B. Schaafsma, and R. Sebastiani. The MathSAT5 SMT Solver. In *TACAS*, 2013.
- [7] A. Cimatti, L. Palopoli, and Y. Ramadian. Symbolic computation of schedulability regions using parametric timed automata. In *RTSS*, 2008.
- [8] N. Eén, A. Mishchenko, and Robert K. Brayton. Efficient implementation of property directed reachability. In *FMCAD*, pages 125–134, 2011.
- [9] G. Frehse, S. Jha, and B. Krogh. A counterexample-guided approach to parameter synthesis for linear hybrid automata. In *HSCC*, 2008.
- [10] T. Henzinger and P. Ho. Hytech: The cornell hybrid technology tool. In *Hybrid Systems*, 1994.
- [11] K. Hoder and N. Bjørner. Generalized property directed reachability. In *SAT*, 2012.
- [12] F. Wang. Symbolic parametric safety analysis of linear hybrid systems with bdd-like data-structures. *IEEE Trans. Software Eng.*, 31(1), 2005.

<sup>4</sup>We currently use a cutoff value on the number of elementary operations in the quantifier elimination module of MATHSAT for this.