

PHAVer: Algorithmic Verification of Hybrid Systems past HyTech*

Goran Frehse

VERIMAG,
2, ave de Vignate, 38610 Gieres, France, e-mail: goran.frehse@imag.fr,
WWW home page: <http://www-verimag.imag.fr/~frehse>

Received: June 20, 2006 / Revised version: date

Abstract. In 1995, HyTech broke new ground as a potentially powerful tool for verifying hybrid systems – yet its applicability remains limited to relatively simple systems. We address the main problems of HyTech with PHAVer, a new tool for the exact verification of safety properties of hybrid systems with piecewise constant bounds on the derivatives. Affine dynamics are handled by on-the-fly overapproximation, partitioning the state space based on user-definable constraints and the dynamics of the system. PHAVer features exact arithmetic in a robust implementation that, based on the Parma Polyhedra Library, supports arbitrarily large numbers. To manage the complexity of the polyhedral computations, we propose methods to conservatively limit the number of bits and constraints of polyhedra. Experimental results for a navigation benchmark and a tunnel diode circuit show the effectiveness of the approach.

Key words: hybrid systems – verification – tools – polyhedra

1 Introduction

Systems with interacting discrete as well as continuous dynamics, so-called hybrid systems, are known to be notoriously complex to analyze. Their algorithmic verification remains a challenging problem, from a theoretical point of view because of decidability problems, and from the implementation side because the slightest numerical errors can void the results. We present PHAVer (Polyhedral Hybrid Automaton Verifier), a verification tool that

aims at overcoming the limitations of predecessors such as as HyTech [1]. Although PHAVer uses in principle the same algorithm as HyTech we not merely present a new implementation of known algorithms, but propose heuristics to overcome inherent and fundamental problems of reachability computations, namely excessive complexity, slow convergence and loss of accuracy due to overapproximations. As a result, we are able to analyze systems previously beyond the reach of verification tools.

In this paper we consider the reachability problem, i.e., computing the set of reachable states of a system or a conservative overapproximation thereof. While reachability is a relatively simple property in the context of verification, more complex safety or bounded liveness properties can be formulated as reachability problems by adding monitoring components to the system [2]. Computations in PHAVer are based on linear hybrid automata (LHA) [3], which are defined by linear predicates and piecewise constant bounds on the derivatives. They stand out from other classes of hybrid automata because there are known algorithms for computing exact reachability over an infinite time horizon [4]. PHAVer is based on such reachability algorithms for LHA, but can also compute conservative overapproximations for hybrid automata with more complex dynamics thanks to an on-the-fly implementation of phase-portrait approximation [5]. In practise, reachability computations face three fundamental problems: excessive complexity, slow convergence and insufficient accuracy. The *complexity* of the set of reachable states can be prohibitively large. In addition to the state explosion problem, well-known from the domain of discrete systems, one has to deal with expensive polyhedral computations whose complexity increases, often exponentially, with each iteration of the analysis. Inevitably, one has to resort to overapproximations for all but the most simple problems. Although termination is not guaranteed in general, it may be induced by overapproximation. However, *convergence* might be

* A preliminary version of this paper appeared in the *Proceedings of Hybrid Systems: Computation and Control (HSCC 2005)*, *Lecture Notes in Computer Science 3414*, Springer-Verlag, 2005, pp. 258–273.

prohibitively slow, and overapproximations may have a positive or negative effect that is hard to predict. Overapproximations are essential to handle complexity and convergence problems, but they forcibly decrease the *accuracy* of the results. E.g., the obtained set of reachable states might contain forbidden states that are not actually reachable by the system.

In applications, complexity and convergence problems manifest themselves in terms of computation time, memory consumption and degradation of accuracy, and the challenge is to find an acceptable trade-off between them. We propose the following heuristics to enable such a trade-off. We manage the complexity of polyhedral computations by imposing a limit on the number of bits and constraints in polyhedral representations, overapproximating them in a manner that is guaranteed to be conservative. These overapproximations immediately imply termination if the invariants of the hybrid automaton are bounded. However, they may lead to severe convergence problems, which we counter by applying them only when triggered by a given complexity threshold. Experimental results indicate that the proposed overapproximations are sufficiently accurate in practical applications.

In addition to the reachability algorithm, PHAVer includes a separate engine for computing simulation relations between hybrid automata. It can be used to verify equivalence or refinement between different models, and in assume-guarantee-style proofs. The reader is referred to [6,7] for a description of the approach and some experimental results. Using PHAVer, we successfully verified analog and mixed-signal circuits that were previously beyond the reach of verification tools, see [8,9]. PHAVer has been used by others to verify safety properties of hybrid Chi specifications [10], and to verify a type of stability of hybrid systems [11]. It was also used in abstraction-refinement schemes based on rectangular automata [12] and based on LHA [9]. A performance comparison between PHAVer and HSOLVER in [13] turned out in PHAVer's favor.

Hybrid Automata Our hybrid automata are based on the classic model in [14]. We added a distinction between input and output variables, comparability and compatibility definitions from [15]. In addition, we adopted from [3] that the set of variables may be different for each automaton in a composition, and there is a set of initial states.

From HyTech to PHAVer In 1995, Henzinger et al. presented the tool HyTech [1]. It features a powerful input language permitting to specify fixed-point algorithms based on post- and pre-operators, but suffers from a major flaw: It uses exact arithmetic with a data structure of limited digits, which can quickly lead to overflow errors. It was successfully used to analyze a number of relatively small examples [16,17,18,19,20,21], but

the overflow problem prohibits any application to more complex systems. Nonetheless, the numerous valuable experiences with HyTech have spawned important suggestions for improvement [17], most of which we incorporated in PHAVer. The basic operations in PHAVer, i.e., discrete and timed post-operators based on polyhedral computations, are identical to those in HyTech, and use exact arithmetic with unlimited precision. Our implementation is immune to overflow errors thanks to its use of the Parma Polyhedra Library (PPL) [22] and the GNU Multiple Precision Arithmetic Library (GMP) [23], which can handle computations with hundreds of thousands of bits on a standard PC. The first HyTech prototype was implemented in Mathematica and did not have any numerical restrictions, but it was also 50–1000 times slower than the later version written in C++ [24].

Overapproximating Dynamics The on-the-fly overapproximation in PHAVer is a variation of the phase-portrait approximation from [5]. It differs in that the invariants of the partitioned locations do not have an open cover, because we split along hyperplanes and the invariants only overlap on the plane. It has been shown in [12] that this is valid for affine dynamics. Similar variants of partitioning the state space are used throughout literature, e.g., in [25,26], although we're not aware of any work using the angular spread of the derivatives. Earlier attempts to improve over HyTech include the use of interval arithmetic [27], which can quickly lead to prohibitively large overapproximations. Recently, interval arithmetic was combined with abstraction refinement in a tool called HSOLVER, which can deal conservatively with nonlinear dynamics [28]. For a survey of verification tools for hybrid automata, see [29].

Managing complexity In the context of timed systems, zones are *rounded* in [30] by dropping constraints or rounding to a constant lower bound. An algorithm specialized on rectangular automata was proposed in [31] and implemented based on the HyTech engine. It is able to use a limited number of bits through component-wise conservative rounding of the coefficients. However, the rectangular over-approximation can become prohibitively large. An improvement was proposed in [32] by allowing arbitrary convex polyhedra. It also incorporates a strategy to reduce the number of bits by component-wise overapproximation, but is based on a vertex representation of polyhedra and its complexity is exponential in the number of variables. For the simplification of polyhedra it has been suggested to use bounding boxes or oriented rectangular hulls [33]. Instead, we propose to simply drop the least significant of the constraints, as this seems a good compromise in terms of accuracy and speed.

In the next section we introduce the hybrid automaton model used in PHAVer. In Sect. 3 we present the reachability analysis algorithm and its on-the-fly over-

approximation of affine dynamics. Experimental results are provided for a navigation benchmark. Methods to manage the complexity of polyhedra by limiting the bits and constraints are proposed in Sect., 4, and illustrated with a tunnel diode circuit. We draw some conclusions in Sect. 5.

2 Hybrid Automata with Controlled Variables

Hybrid automata [34] are widely recognized as an intuitive and expressive modeling paradigm for hybrid systems. Out of several different definitions present in literature we choose a slight variation of the hybrid automata in [14] because they admit compositional reasoning. To reason about equivalence between automata in a practical manner, we enrich them with a distinction between input and output variables, as in [15], although we forego such a distinction between synchronization labels for the sake of simplicity.

Preliminaries Given a set $X = \{x_1, \dots, x_n\}$ of variables, a *valuation* is a function $v : X \rightarrow \mathbb{R}$. We use \dot{X} to denote the set $\{\dot{x}_1, \dots, \dot{x}_n\}$ of dotted variables, and X' to denote the set $\{x'_1, \dots, x'_n\}$ of primed variables. Let $V(X)$ denote the set of valuations over X . The *projection* of v is $v \downarrow_{\bar{X}} = \{x \rightarrow v(x) | x \in \bar{X}\}$. The *embedding* of a set $U \subseteq V(X)$ into variables $\bar{X} \supseteq X$ is $U \uparrow^{\bar{X}} = \{v \in V(\bar{X}) | v \downarrow_X \in U\}$. When a valuation u over X and a valuation v over \bar{X} *agree*, i.e., $u \downarrow_{X \cap \bar{X}} = v \downarrow_{X \cap \bar{X}}$, we use $u \sqcup v$ to denote the valuation w defined by $w \downarrow_X = u$ and $w \downarrow_{\bar{X}} = v$. Arithmetic operations on valuations are defined in the straightforward way. An *activity* over X is a function $f : \mathbb{R}^{\geq 0} \rightarrow V(X)$. Let $Acts(X)$ denote the set of activities over X . The *derivative* \dot{f} of an activity f is an activity over \dot{X} , defined analogously to the derivative in \mathbb{R}^n . The extension of operators from valuations to activities is done pointwise. Let $const_X(Y) = \{(v, v') | v, v' \in V(X), v \downarrow_Y = v' \downarrow_Y\}$.

A *hybrid automaton* $H = (Loc, (X, O, C), Lab, Edg, Flow, Inv, Init)$ consists of the following:

- A set $X = \{x_1, \dots, x_n\}$ of continuous, real-valued *variables*. They are divided into *controlled* variables C and *input* variables $I = X \setminus C$, and a subset O of C is designated as the *output* variables.
- A set Loc of *locations*. A *state* is a pair (l, v) of a location l and a valuation v , which attributes a value to each of the variables.
- A set Lab of *synchronization labels* including the *stutter* label τ .
- A set Edg of *transitions* that describe instantaneous changes of location, in the course of which variables may change their values. Each transition $(l, a, \mu, l') \in Edg$ has a *source location* l , a *target location* l' . It is associated with a *synchronization label* a and with a *jump relation* μ that relates values of variables before the transition to values they take after the transition.

The projection of the jump relation to the variables before the transition describes for which variable values the transition is enabled; this is often referred to as a *guard*. Each location l has a self-loop *stutter transition* $(l, \tau, const_X(C), l)$ that lets the input variables change arbitrarily within the invariant, while the controlled variables remain constant.

- A mapping *Flow* attributes to each location a set of valuations over the variables and their derivatives, which determines how variables can change over time.
- A set of states Inv called *invariant*. All behavior is constrained to the invariant at all times. In this paper we consider only invariants that are convex in each location.
- A set Ini of *initial states*, contained in the invariants, from which all behavior of the automaton originates.

A *linear constraint* over a set of variables X has the form

$$\sum_i \alpha_i x_i + \beta \bowtie 0, \quad (1)$$

where α_i and β are integer constants and \bowtie is a sign \bowtie is either $<$ or \leq . A *convex linear predicate* is a conjunction of linear constraints. We write $\dot{X} = \{\dot{x}_1, \dots, \dot{x}_n\}$ to denote derivatives, and using $X' = \{x'_1, \dots, x'_n\}$ describe a relation $R \in X \times X'$ with a predicate over $X \cup X'$. In a *Linear Hybrid Automaton* (LHA) [3], invariants and initial states are given by linear predicates over X , flow predicates by convex linear predicates over \dot{X} , and jump relations by linear predicates over $X \cup X'$. PHAVer can handle *affine hybrid automata*, which are defined like LHA except that flows may be given by linear constraints over $X \cup \dot{X}$. The use of inequalities in flow predicates allows us to model dynamics of the form $\dot{x} = Ax + b$ with the elements of A and b given by intervals.¹ Note that LHA can model discrete-time dynamics of the form $x_{k+1} = Ax_k + b$ by updating the variables in transitions with a jump relation $x'_i = \sum_j a_{ij} x_j + b_i$.

Complex systems can be constructed in a modular fashion using *parallel composition*. Our hybrid automata are *compositional* with respect to reachability properties [7], which means that a state that is not reachable in a given automaton H will also not be reachable in a composition of H with any other automaton. This allows one to analyze isolated subsystems and conclude that any non-reachability property found will also hold in the entire system. We use the standard definition of parallel composition:

Definition 1 (Parallel Composition). [14] Hybrid automata H_1, H_2 are *compatible* if $C_1 \cap C_2 = \emptyset$, $X_1 \cap C_2 \subseteq O_2$ and $X_2 \cap C_1 \subseteq O_1$. The *parallel composition* of compatible hybrid automata H_1, H_2 is the hybrid automaton H with

¹ In literature, affine hybrid automata are sometimes also referred to as a *linear* hybrid automata, e.g., in [35], which may lead to confusion with LHA.

- $Loc = Loc_1 \times Loc_2$,
- $X = X_1 \cup X_2$, $C = C_1 \cup C_2$, $O = O_1 \cup O_2$, $Lab = Lab_1 \cup Lab_2$
- $((l_1, l_2), a, \mu, (l'_1, l'_2)) \in Edg$ iff
 - $(l_1, a_1, \mu_1, l'_1) \in Edg_1$ and $(l_2, a_2, \mu_2, l'_2) \in Edg_2$
 - either $a = a_1 = a_2$, or $a = a_1 \notin Lab_2$ and $a_2 = \tau$, or $a_1 = \tau$ and $a = a_2 \notin Lab_1$,
 - $\mu = \{(v, v') | (v \downarrow_{X_i}, v' \downarrow_{X_i}) \in \mu_i\}$;
- $Flow(l_1, l_2) = Flow_1(l_1)|^{X \cup \dot{X}} \cap Flow_2(l_2)|^{X \cup \dot{X}}$;
- $Inv(l_1, l_2) = Inv_1(l_1)|^X \cap Inv_2(l_2)|^X$;
- $Init(l_1, l_2) = Init_1(l_1)|^X \cap Init_2(l_2)|^X$.

Example 1. Consider the model of a timing based mutual-exclusion protocol shown in Fig. 1, adapted from [4]. In every location l of P_i , there is a transition (l, τ, μ, l) with $\mu = \{(v, v') | v(x_i) = v'(x_i), v(k), v'(k) \in \mathbb{R}\}$ (omitted from the figure). Variables stay constant in transitions unless there is an explicit assignment. The system is considered *safe* if there are never two or more processes in the critical section at the same time. It is a compositional adaptation of the model given in [4], and parameterized to n processes with time constants c_i and d_i that represent the minimal, respectively maximal, skew of their clocks. The processes P_i have a controlled variable x_i to model their local clock and an input variable k that models a semaphore. Because none of the processes controls k , it is modeled separately in an automaton S we call a *shared variable model*. S has k as a controlled variable and fixes its derivative to zero. It gives the processes access to k by synchronizing on transitions labeled with set_i or $release_i$. It has no restrictions on the jump relations of these transitions, and thus allows the processes to change k to an arbitrary value.

Semantics The behavior of a hybrid automaton is based on two types of transitions: *Discrete* transitions are manifestations of the transitions in Edg , and change the location and variables instantaneously. *Timed* transitions describe the change of the variables over time with an admissible activity. An activity $f(t) \in Acts(X)$ is called *admissible* over an interval $[0, \delta]$ in a location l if $\delta = 0$, or $\forall t, 0 \leq t \leq \delta : f(t) \in Inv(l), f(t) \sqcup \dot{f}(t) \in Flow(l)$. In weak runs, we consider τ -transitions unobservable as long as they do not change the variables.

Definition 2 (Run, Reachability). For a given hybrid automaton H , there is a *discrete transition* $\sigma = p \xrightarrow{\delta, a} p'$ with *source and target states* p, p' and label $a \in Lab$ iff $p, p' \in Inv$ and there is a transition $loc(p) \xrightarrow{a, \mu} loc(p')$ with $(val(p), val(p')) \in \mu$. There is a *timed transition* $\sigma = p \xrightarrow{\delta, f} p'$ with *duration* $\delta \in \mathbb{R}^{\geq 0}$ and activity f over X called *witness* iff $p, p' \in Inv$, and f is differentiable and admissible over $[0, \delta]$ in $loc(p)$ and $f(0) = val(p)$. A *run* of a hybrid automaton H is a finite or infinite sequence

$$\sigma = p_0 \xrightarrow{\delta_0, f_0, a_0} p_1 \xrightarrow{\delta_1, f_1, a_1} \dots \xrightarrow{\delta_{n-1}, f_{n-1}, a_{n-1}} p_n$$

such that for all $i \geq 0$ $p_i \xrightarrow{\delta_i, f_i} (loc(p_i), f_i(\delta_i)) \xrightarrow{a_i} p_{i+1}$. A state p' is *reachable* if there exists a run σ with $p_0 \in Init$ and $p_n = p'$.

3 Reachability Analysis in PHAVer

We compute the set of reachable states as the fixed-point of two *Post*-operators, which yield the image states of discrete, respectively timed, transitions. If the flow in a location is affine, it is overapproximated with LHA-dynamics, i.e., with constant bounds on the derivatives. The accuracy of this overapproximation depends on the size of the invariant, so locations are partitioned on the fly depending on several criteria. We give an overview of our algorithm, and then provide detailed descriptions of the operations involved.

The implementation in PHAVer processes a shared passed and waiting list as shown in Fig.2. First, the locations of the initial states are partitioned, time elapse is applied, and they are put on the passed and waiting list. In the main loop, a set of states is taken from the waiting list, and the successor states of discrete transitions are computed. Then the locations of these successors are partitioned, and in each partition the dynamics are overapproximated with LHA-dynamics. The set of successors is subjected to complexity management, which may overapproximate it with a set of lower complexity. Time elapse is applied to this new set, and it is added to the passed list. The new states, i.e., those that were not previously on the passed list, are put onto the waiting list. The loop repeats until the waiting list is empty.

Remark 1. Note that to be sound, time elapse must be applied after any overapproximation, such as convex hull or complexity management, unless the overapproximation is guaranteed to be invariant with respect to time elapse.

The basic algorithm, shown on the left side of Fig.2, is similar to the one used in HyTech, and summarized in the next section. Partitioning, overapproximation and complexity management are specific to PHAVer, and will be discussed in detail in the remainder of the paper.

3.1 Reachability of LHA

We summarize the standard reachability algorithm from [4, 36] for LHA using polyhedra. To compute the set of reachable states, one repeatedly computes the successors of timed and discrete transitions. Let

$$Post_t(P) = \{p' | \exists \delta, f, p \in P : p \xrightarrow{\delta, f} p'\}$$

be the set of timed successors and

$$Post_d(P) = \{p' | \exists a, p \in P : p \xrightarrow{a} p'\}$$

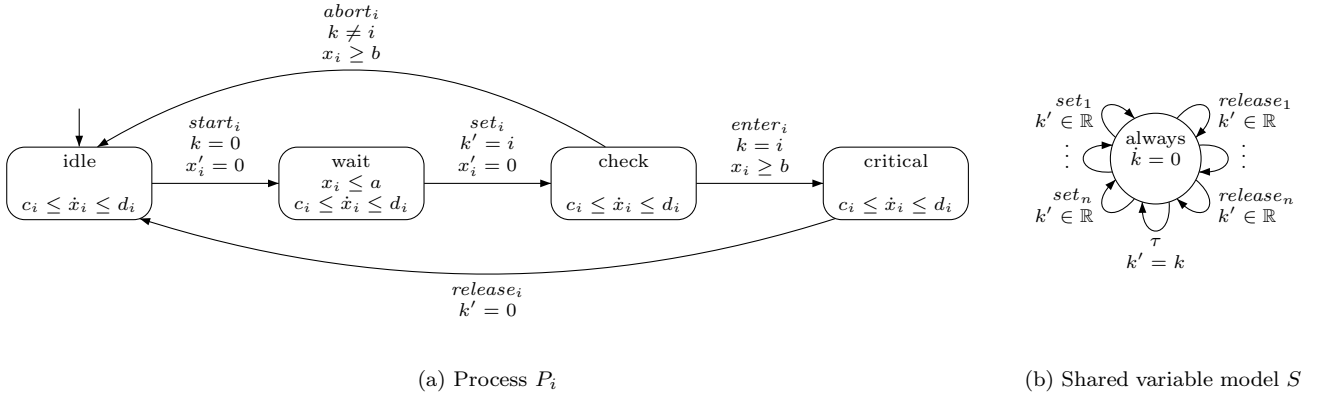


Fig. 1. Compositional model of timing based mutual-exclusion protocol in [4]

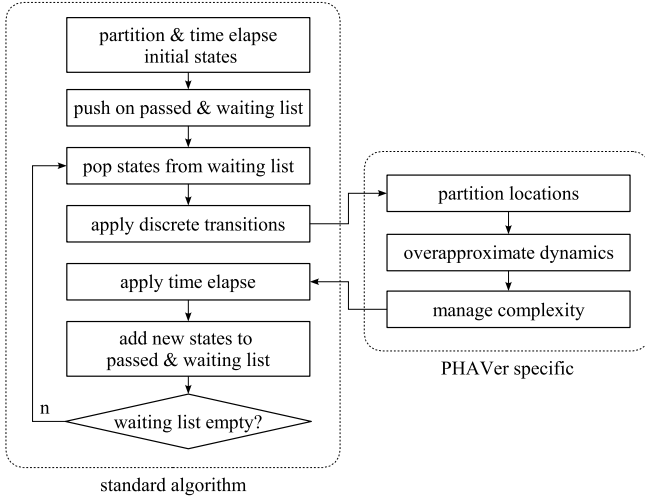


Fig. 2. Reachability algorithm in PHAVer

be the set of discrete successors. The set of reachable states, denoted by $Reach$, is the smallest fixed-point of the sequence given by $P_0 = Init$ and

$$P_{k+1} = Post_d(Post_t(P_k)). \quad (2)$$

We recall the basic definitions and operations on polyhedra, and use them to define the $post$ -operators for LHA.

Definition 3. A *convex polyhedron* S is a set in \mathbb{R}^n that can be represented as the conjunction of a finite number of linear constraints, i.e.,

$$S = \{x \mid \bigwedge_i a_i^T x \bowtie_i b_i\}, \quad (3)$$

where $a_i \in \mathbb{R}^n$, $b_i \in \mathbb{R}$ and $\bowtie_i \in \{<, \leq\}$. This is referred to as the *constraint representation* of the polyhedron. Strict inequalities $a_i^T x < b_i$ can be replaced with nonstrict inequalities $a_i^T x + \varepsilon \leq b_i$ by introducing an auxiliary variable ε that may take an arbitrary positive

value [36]. In the following, we will therefore only consider nonstrict inequalities, which define closed polyhedra. Alternatively, a polyhedron can be described by a *generator representation* (V, R) as the convex hull of a finite set $V \subseteq \mathbb{R}^n$ of *vertices* and a finite set $R \subseteq \mathbb{R}^n$ of *rays*, i.e.,

$$S = \left\{ \sum_{v_i \in V} \lambda_i v_i + \sum_{r_i \in R} \mu_i r_i \mid \lambda_i, \mu_i \geq 0, \sum_i \lambda_i = 1 \right\}. \quad (4)$$

One may switch between constraint and generator representations, although at a potentially exponential cost. Some operations are easier in constraint form, e.g., intersection, while others are easier in generator form, e.g., testing for emptiness. Therefore modern polyhedral libraries, such as the PPL, keep both representations in what is referred to as the *double description method*. Projection and embedding operators are easily implemented in constraint form by removing elements or adding zeros to the vectors a_i .

Recall that in a discrete transition $(l, a, \mu, l') \in Edg$, the jump relation μ is a predicate over $X \cup X'$, where the primed variables specify the new values after the transition. The discrete successors of a set of states P are

$$Post_d(P) = \bigcup_{(l, a, \mu, l') \in Edg} l' \times \left((P(l) \upharpoonright^{X \cup X'} \cap \mu) \downarrow_{X' \rightarrow X} \cap Inv(l') \right), \quad (5)$$

where $\downarrow_{X' \rightarrow X}$ denotes that we project onto the primed variables X' , and rename the variables as unprimed. The timed successor operator is slightly more involved. For LHA with convex invariants there is an activity from one state in location l to another iff they are connected by a straight line whose derivative is in $Flow(l)$ [4]. This allows one to replace the passage of time by simple existential quantification as follows. Let $S \nearrow F$ be defined as

$$S \nearrow F = \{v' \mid \exists \delta \in \mathbb{R}^{\geq 0}, w \in F, v \in S : v' = v + w\delta\}. \quad (6)$$




Fig. 3. Time elapse for LHA

This operation is easily implemented using the generator representation [36]. Let (V, R) be the generator representation of S , and (V', R') that of F . Then $S \nearrow F = (V, R \cup V' \cup R')$. For convex invariants, the timed successors are then given by

$$Post_t(P) = \bigcup_{l \in loc(P)} l \times (P(l) \nearrow Flow(l)) \cap Inv(l). \quad (7)$$

Example 2. Consider the mutual exclusion protocol from Ex. 1 for two processes, with $a = 3$, $c_1 = c_2 = 1$, $d_1 = 3$, $d_2 = 2$. We compute the time elapse for the set $S = \{x_1 = x_2 = 0\}$, with both processes in location *wait*. The flow in the location is given by $F = Flow((wait, wait)) = \{1 \leq \dot{x}_1 \leq 3 \wedge 1 \leq \dot{x}_2 \leq 2\}$, which translates in generator form to $F = (\{(1, 1), (3, 1), (3, 2), (1, 2)\}, \{\})$, see Fig. 3(a). To apply the time elapse, the vertices of F in the derivative space are reinterpreted as rays in the state space, shown as dashed arrows in Fig. 3(b). These rays are added to the generator representation of $S = (\{(0, 0)\}, \{\})$ and we obtain

$$S \nearrow F = (\{(0, 0), (1, 1), (3, 1), (3, 2), (1, 2)\}, \{\}).$$

Finally, one has to restrict the time elapse to the invariant $Inv((wait, wait)) = \{x_1 \leq 3, x_2 \leq 3\}$. The resulting set $Post_t((wait, wait) \times S) = (S \nearrow F) \cap Inv((wait, wait))$ is shown in Fig. 3(b).

Remark 2. The time elapse operator implemented in the PPL (v0.9) computes the smallest convex set containing $S \nearrow F$. Strictly speaking, this is an overapproximation as the actual set may be nonconvex. Consider $S = \{x = 0 \wedge y = 0\}$ and $F = \{y > 0\}$. Then $S \nearrow F = \{(x = 0 \wedge y = 0) \vee y > 0\}$, while the PPL returns $S \nearrow_{PPL} F = \{y \geq 0\}$.

In the following example we compare the performance of HyTech and PHAVer. The results indicate that the overhead introduced by using a multiprecision number representation is compensated by the efficiency of the PPL, and improvements such as the shared passed and waiting list. The comparison is limited to simple parameters and systems to avoid overflow problems with HyTech.

Example 3. Consider the mutual exclusion protocol from Ex. 1 for n processes, with $a = 1$, $b = 4$, $c_i = 1$, $d_i = 2$. The experimental results are shown for MEX1 with $n = 1$ to MEX5 with $n = 5$ in Table 1, as well as parametric versions MEXP1–4 in which the allowed range of a and b is computed. In these instances, PHAVer outperforms HyTech for nontrivial systems. The downside is its memory consumption, which is about twice that of HyTech. In the instance of a parametric analysis of an audio protocol AUDSP from [37], HyTech outperforms PHAVer. The experiments in this paper were carried out on a Pentium IV, 3.2 GHz with 1 GB RAM running Linux.

3.2 Overapproximation of Dynamics

To deal with non-LHA dynamics, we apply a simple variant of *phase-portrait approximation* [5]. This method utilizes the fact that overapproximating any of the sets that define an automaton, in particular $Flow$, results in an overapproximation of its behavior. We can therefore apply our reachability algorithm to any hybrid automaton by overapproximating it with LHA-style sets, and obtain a conservative overapproximation of its reachable states. Because the accuracy of the overapproximation depends on the size of the invariant, the locations are adequately partitioned before applying this transformation, as will be discussed in the next section.

Currently, this is implemented for affine dynamics, which have the form $M\dot{x} = Ax + b$, or more precisely a relaxed version thereof that may consist of inequalities. We describe two methods for overapproximating them with LHA-dynamics, i.e., for obtaining constant bounds on the derivatives in the form of linear constraints. In the following, we consider $Flow$ to be given as *relaxed affine dynamics*, i.e., a conjunction of constraints

$$a_i^T \dot{x} + \hat{a}_i^T x \bowtie_i b_i, \quad (8)$$

where $a_i, \hat{a}_i \in \mathbb{Z}^n$, $b_i \in \mathbb{Z}$, $\bowtie_i \in \{<, \leq, =\}$, $i = 1, \dots, m$.

Projection Given a set of states S that constrains the possible values of x , we can obtain bounds on the derivatives by intersecting $Flow$ with S and projecting onto the derivatives:

$$Flow_{pr}(S) = \left(Flow \cap S \Big|_{x \cup \dot{x}} \right) \Big|_{\dot{x}}. \quad (9)$$

We may always choose $S = Inv$, but further refinement is possible, as discussed below. PHAVer's default time elapse operator is obtained by substituting (9) for the flow in (7):

$$Post_{t,pr}(P) = \bigcup_{l \in loc(P)} l \times (P(l) \nearrow Flow_{pr}(Inv(l)) \cap Inv(l). \quad (10)$$

When $Flow$ is given as a conjunction of constraints of the form (8) and S is given by linear constraints, (9) involves only standard operations on polyhedra in the

Table 1. Tool performance, HyTech vs. PHAVer

Instance	HyTech			PHAVer			Reachable Set	
	Time [s]	Mem. [MB]	Iter.	Time [s]	Mem. [MB]	Iter.	Loc.	Poly.
MEX2	0.0	3.1	7	0.1	5.2	9	13	21
MEX3	0.7	6.0	17	1.0	10.3	15	39	139
MEX4	17.5	28.8	23	10.0	58.6	21	113	1169
MEX5	928.3	437.8	29	238.9	814.5	27	323	12001
MEX2P	0.1	4.5	12	0.2	5.9	9	16	37
MEX3P	10.5	27.4	23	2.9	17.9	19	64	564
MEX4P	2999.0	1660.6	34	114.0	443.1	29	256	12408
AUDSP	8.8	62.7	39	20.1	143.6	39	193	5849

space of $X \cup \dot{X}$, and is therefore straightforward to implement. Sometimes, $Flow_{pr}(S)$ can be prohibitively complex, e.g., when S is iteratively refined, as discussed below, or in abstraction-refinement schemes such as the one in [9]. If this is the case, the following solution may be advantageous.

Constraint-Based The following approach has the advantage that it results in the same number of constraints as $Flow$, independent on how complex the invariant or other restrictions might be. We assume that equalities in (8) are modeled using conjuncts of pairs of inequalities. We transform each constraint (8) into a linear constraint over \dot{X} by finding a lower bound on $\hat{a}_i^T x$. If no such bound exists, we drop the constraint entirely.

Given a set of states S that constrains the possible values of x , let


$$p/q = \inf_{x \in S} \hat{a}_i^T x, \quad p, q \in \mathbb{Z}. \quad (11)$$

If the infimum exists, the set of \dot{x} that satisfy (8) is contained in the set defined by $\hat{a}_i^T \dot{x} \bowtie_i b_i - p/q$. Multiplying both sides with q yields the LHA-style constraint

$$q\hat{a}_i^T \dot{x} \bowtie_i qb_i - p. \quad (12)$$

Let $Flow_{con}(S)$ be obtained from $Flow(S)$ by replacing each constraint of the form (8) by its corresponding constraint (12) if the infimum in (11) exists, and otherwise removing the constraint.

A one-to-one correspondence of the original and the LHA-style constraints allows the user to specify the type of constraints he wishes to obtain on the dynamics. E.g., let the dynamics be given by $\dot{x} = f(x, y) \wedge \dot{y} = g(x, y)$, f and g being affine functions. Specifying the constraints in this form in PHAVer results in rectangular LHA-style constraints, i.e., of the form $\dot{x} \in [.,.] \wedge \dot{y} \in [.,.]$. Octagonal constraints, i.e., of the form $\dot{x} \in [.,.] \wedge \dot{y} \in [.,.] \wedge \dot{x} - \dot{y} \in [.,.] \wedge \dot{x} + \dot{y} \in [.,.]$ can be obtained by specifying $\dot{x} = f(x, y) \wedge \dot{y} = g(x, y) \wedge \dot{x} - \dot{y} = f(x, y) - g(x, y) \wedge \dot{x} + \dot{y} = f(x, y) + g(x, y)$. Octagonal constraints are useful because they preserve some of the correlation

**Fig. 4.** Overapproximating affine dynamics with LHA dynamics

between variables, and thus may lead to drastically improved accuracy.

As the following example shows, the constraint based overapproximation can be very inaccurate, and does not even have to be tight, i.e., touch the original set.

Example 4. Let $Inv(l) = \{x_l \leq x \leq x_u\}$ and $Flow(l) = \{\dot{x} \leq ax + b_u, \dot{x} \leq -ax + b, \dot{x} \geq ax + b_l\}$ as shown in Fig. 4. With projection we obtain $Flow_{pr}(l) = \{\dot{x} \leq (b - b_u)/(2a)\}$. In the constraint-based approach, we compute the infimum for each constraint separately, obtaining $Flow_{con}(l) = \{ax_l + b_l \leq \dot{x} \leq -ax_l + b\}$. As Fig. 4 illustrates, the overapproximation of the constraint-based approach can be considerable.

Iterative Refinement The set of states S containing all possible values of x is initially chosen to be the invariant. $Flow_{pr}(Inv)$, respectively $Flow_{con}(Inv)$, yields the set of all derivatives possible inside this invariant, which is independent from the set that is actually reachable. We can further refine the set of derivatives – and the set of reachable states – by restricting S to the set of states that are reachable with $Flow_{pr}(Inv)$, possibly reiterating the process until convergence. This yields the following fixed-point computation for the operator $Post_t(P)$, where $Post_{t,pr}^0 = P$:

$$Post_{t,pr}^{k+1} = \bigcup_{l \in loc(P)} l \times (P(l) \nearrow Flow_{pr}(Post_{t,pr}^k)) \cap Inv(l). \quad (13)$$

Note that $Post_{t,pr}^1 = Post_{t,pr}(P)$. We denote the fixed-point of the above sequence as $Post_{t,pr}^\infty(P)$, and the corresponding result of the constraint-based operator with $Post_{t,con}^\infty(P)$. The following example shall illustrate the qualitative difference that a single iteration of (13) can make.

*Example 5.*² Consider $Flow(l) = \{\dot{x} = y \wedge \dot{y} = 0\}$ and $Inv^l(l) = \{-1 \leq y \leq 1\}$. With $Flow_{pr}(Inv) = l \times \{\dot{x} \in [-1, 1], \dot{y} = 0\}$, the timed successors of a set of states $P = l \times \{x = 0 \wedge y = 1\}$ are $Post_{t,pr}(P) = l \times \{y = 1\}$. The overapproximation of $Flow$ with $Flow_{pr}$ results in negative values of x being reachable, even though $\dot{x} = y = 1$. Applying the fixed-point computation (13), we obtain using $Post_{t,pr}^1 = Post_{t,pr}$ and $Flow_{pr}(Post_{t,pr}^1) = l \times \{\dot{x} = 1, \dot{y} = 0\}$ the fixed-point $Post_{t,pr}^2(P) = l \times \{x \geq 0 \wedge y = 1\}$, which is equal to the actual reachable set $Post_t(P)$.

In practise, only the first 2–3 iterations yield a significant improvement, so PHAVer limits the number of iterations to a user-defined value.

3.3 Partitioning Locations

When the dynamics are overapproximated as in the previous section, the error depends on the size of the invariant of the location. We partition each invariant into sufficiently small parts with a recursive splitting operator. Each splitting cuts the invariant in two along a hyperplane, and replaces the original location with two copies, each assigned a part of the original invariant. The hyperplanes for splitting are chosen from a user-defined set. The splitting recursion terminates if the size of the invariant or the spread of the derivatives fall below a given threshold. By using small enough partitions, one is able to approximate the behavior of the original hybrid automaton arbitrarily close [5]. Note that this requires the introduction of the derivatives as auxiliary variables, as will be discussed at the end of this section.

The splitting of a location consists of duplicating the location, including incoming and outgoing transitions as well as flow, invariant and initial states. The invariants of the location duplicates may be restricted to subsets as long as they cover the original invariants. Formally, this corresponds to the following operation:

Definition 4 (Invariant split). (modified from [5]³) A *split* \mathcal{S} for a hybrid automaton H maps each location l to a finite set $\{S_1^l, \dots, S_k^l\}$ of sets of valuations over X such that there exists a finite cover $\mathcal{O}^l = \{O_1^l, \dots, O_k^l\}$ of $Inv(l)$ with $S_i^l = Inv(l) \cap O_i^l$ for $i = 1, \dots, k$. The *split* of H along \mathcal{S} is the hybrid automaton $split(H, \mathcal{S}) = (Loc_{\mathcal{S}}, (X, C, O), Lab, Edg_{\mathcal{S}}, Flow_{\mathcal{S}}, Inv_{\mathcal{S}}, Init_{\mathcal{S}})$ with

- $Loc_{\mathcal{S}} = \{(l, S) \mid l \in Loc, S \in \mathcal{S}(l)\}$,
- $Edg_{\mathcal{S}} = \{((l, S), a, \mu, (l', S')) \mid (l, a, \mu, l') \in Edg\}$,
- $Flow_{\mathcal{S}}((l, S)) = Flow(l)$,
- $Inv_{\mathcal{S}}((l, S)) = Inv(l) \cap S$, and
- $Init_{\mathcal{S}}((l, S)) = Init(l) \cap S$.

The behavior of the split automaton $H_{\mathcal{S}}$ is identical, i.e., bisimilar, to the behavior of the original automaton if the cover \mathcal{O} is open [5], or the dynamics are affine [12].

Recall that a hyperplane h is defined by an equation $a_h^T x = b_h$, where the normal vector a_h determines its direction and the inhomogeneous term b_h its position. Let the *slack* of h in a location l be defined by

$$\Delta(a_h) = \sup_{x \in Inv(l)} a_h^T x - \inf_{x \in Inv(l)} a_h^T x.$$

In PHAVer, we recursively split one location l at a time along a hyperplane $a_h^T x = b_h$, i.e., we apply an invariant split with $S_1^l = \{a_h^T x \leq b_h\}$, $S_2^l = \{a_h^T x \geq b_h\}$. Note that only reachable locations are split. The user provides a list of candidate normal vectors $a_{h,i}$ together with a minimum and maximum value for the slack of each hyperplane:

$$Cand = \{(a_{h,1}, \Delta_{min,1}, \Delta_{max,1}), \dots, (a_{h,m}, \Delta_{min,m}, \Delta_{max,m})\}.$$

This allows the user to include expert knowledge by choosing planes and location sizes suitable for the system. A trivial choice for the hyperplanes are the directions of the axes. Let P be the set of reachable states at the time of the splitting. The position b_h of the hyperplane is chosen to be the center of the location, i.e.,

$$b_h = (\sup_{x \in Inv(l)} a_h^T x + \inf_{x \in Inv(l)} a_h^T x)/2,$$

if both supremum and infimum exist. Otherwise, let

$$b_h = \Delta_{min} + \sup_{x \in P} a_h^T x, \quad \text{or} \quad b_h = -\Delta_{min} + \inf_{x \in P} a_h^T x,$$

depending on which lies inside the invariant. In related work, a method for positioning of the hyperplane based on optimizing the set of derivatives was proposed in [12].

In each split, the best hyperplane is chosen according to a number of criteria. We provide an overview of the criteria, followed by a detailed description of the selection process. In principle, each location is split until the slack of every candidate hyperplane h_i satisfies $\Delta(a_{h,i}) \leq \Delta_{min,i}$. We account for the dynamics of the system using the spatial angle that is spanned by the derivatives in a location. Let the *spread* of a set of valuations X be

$$\angle(X) = \inf_{x,y \in X} x^T y / |x||y|.$$

The spread of the derivatives in a location l , constrained to the states S is then

$$\angle_{deriv}(l, S) = \angle((Flow(l) \cap S)^{X \cup \dot{X}} \downarrow \dot{X}).$$

² Many thanks to R. J. M. Theunissen for originating the example.

³ We do not require the cover to be open.

The spread of the derivatives is used in two ways: The partitioning of a location is stopped once the spread is smaller than a given minimum \triangleleft_{min} , and the constraints are sorted according to the spread of the derivatives in the location after the splitting. The slack $\Delta_{max,i}$ may be specified to split at least to that slack without regard to \triangleleft_{min} .

For each splitting, the candidate hyperplanes $a_h^T x = b_h$ are sorted according to the following criteria:

1. slack

$$c_1 = \begin{cases} \Delta(a_h)/\Delta_{min,h} & \text{if } \Delta(a_h) > \Delta_{min,h}, \\ \infty & \text{otherwise.} \end{cases}$$

2. reachable states only on one side

$$c_2 = \begin{cases} 1 & \text{if } \exists x, x' \in P : a^T x < b \wedge a^T x' > b \\ 0 & \text{otherwise.} \end{cases}$$

3. spread of the derivatives (discard constraint if a minimum spread \triangleleft_{min} is reached and the slack is smaller than $\Delta_{max,h}$)

$$c_3 = \begin{cases} -\triangleleft_{deriv}(l, Inv) & \text{if } \triangleleft_{deriv}(l, Inv) \geq \triangleleft_{min} \\ & \vee \Delta(a_h) > \Delta_{max,h}, \\ \infty & \text{otherwise.} \end{cases}$$

4. derivative spread after the constraint is applied

$$c_4 = -\max\{\triangleleft_{deriv}(l, \{(l, x) \in Inv \mid a_h^T x \leq b_h\}), \\ \triangleleft_{deriv}(l, \{(l, x) \in Inv \mid a_h^T x \geq b_h\})\}.$$

By default, the candidate hyperplanes are sorted by evaluating the tuple (c_1, c_2, c_3) lexicographically. The best one is chosen that does not have a criterion evaluating to ∞ . If no such hyperplane exists, the splitting recursion terminates. Note that c_3 can be deactivated by specifying $\triangleleft_{min} = 1$. When c_4 is activated, the sorting is according to (c_4, c_1, c_2, c_3) .

According to [5], the overapproximation of a hybrid automaton H with a linear hybrid automaton can be arbitrarily close to the original if the partition size is chosen sufficiently small. However, the splitting method in [5] relies on splitting not only the invariant, but also the flow predicate. This is necessary to approximate the flow arbitrarily close: Consider, e.g., a nonconvex flow. Any overapproximation by a convex flow will contain the convex hull no matter how small the invariant is partitioned, and can consequently not be arbitrarily close. When splitting the invariant is not sufficient, one may introduce the derivatives as auxiliary variables to implement the method of [5] by invariant split. Let $Y = \{y_1, \dots, y_n\}$ be the auxiliary variables representing the derivatives. The transformed automaton is $H' = (Loc, (X \cup Y, C \cup Y, O), Lab, Edg', Flow', Inv', Init')$, where

- $Edg' = \{(l, a, \mu', l') \mid (l, a, \mu, l') \in \rightarrow, \mu' = \mu|^{X \cup Y \cup X' \cup Y'}\}$,
- $Flow'(l) = Flow(l)|^{X \cup \dot{X} \cup Y \cup \dot{Y}} \cap \{\bigwedge_{i=1, \dots, n} \dot{x}_i = y_i\}$,

- $Inv'(l) = Inv(l)|^{X \cup Y}$, and
- $Init'(l) = Init(l)|^{X \cup Y}$.

The auxiliary variables are not restricted in transitions, invariants or initial states, and their derivatives are not restricted in the flow. They do not modify the behavior of the hybrid automaton except for the constraint $\dot{x}_i = y_i$ in the flow, which has no effect on \dot{x}_i since y_i is unrestricted. However, splitting the invariant will now also split the derivatives, therefore implementing the method in [5].

3.4 Example: Navigation Benchmark

We illustrate the reachability analysis of PHAVer with a navigation benchmark proposed in [38]. It models an object moving in a plane, and it must be shown that a set of bad states is not reachable, and that a set of target states is eventually reached. It is a challenging 4-dimensional system because of two integrators (turning velocity into position), whose marginal stability contributes to the accumulation of errors.

The dynamics of the object is defined by “desired”, or equilibrium, velocities that depend on its position in a map. The set of desired velocities is given by $v_d(i) = (\sin(i\pi/4), \cos(i\pi/4))^T$, $i = 0, \dots, 7$, where i is attributed to each unit square in the plane by a given map M . A special symbol **A** denotes the set of target states, and **B** denotes the set of forbidden states for the object. We verified that the forbidden states are not reachable for the instances shown in Fig. 5, whose maps are given by:

$$M_{NAV01-03} = \begin{pmatrix} \mathbf{B} & 2 & 4 \\ 2 & 3 & 4 \\ 2 & 2 & \mathbf{A} \end{pmatrix}, M_{NAV04} = \begin{pmatrix} \mathbf{B} & 2 & 4 \\ 2 & 2 & 4 \\ 1 & 1 & \mathbf{A} \end{pmatrix}.$$

The dynamics of the state vector $(x_1, x_2, v_1, v_2)^T$ are given by

$$\begin{pmatrix} \dot{x} \\ \dot{v} \end{pmatrix} = \begin{pmatrix} 0 & I \\ 0 & A \end{pmatrix} \begin{pmatrix} x \\ v \end{pmatrix} - \begin{pmatrix} 0 \\ A \end{pmatrix} \begin{pmatrix} 0 \\ v_d(i) \end{pmatrix}, \\ A = \begin{pmatrix} -1.2 & 0.1 \\ 0.1 & -1.2 \end{pmatrix}.$$

The initial states for for NAV01–NAV03 are defined by $x_0 \in [2, 3] \times [1, 2]$, for NAV04 by $x_0 \in [0, 1] \times [0, 1]$, and

$$\begin{aligned} v_{0,NAV01} &\in [-0.3, 0.3] \times [-0.3, 0], \\ v_{0,NAV02} &\in [-0.3, 0.3] \times [-0.3, 0.3], \\ v_{0,NAV03} &\in [-0.4, 0.4] \times [-0.4, 0.4], \\ v_{0,NAV04} &\in [0.1, 0.5] \times [0.05, 0.25]. \end{aligned}$$

As splitting constraints we use

$$Cand = \{(v_1, \delta_1, \infty), (v_2, \delta_2, \infty)\},$$

where appropriate δ_i were established by some trial-and-error runs, and (c_1) as splitting criterion. Note that x_1, x_2 need not be partitioned, since they depend only on

Table 2. PHAVer performance, navigation benchmark

Instance	Time [s]	Mem. [MB]	Iter.	Checks	Automaton		Reachable Set			
					Loc.	Trans.	Loc.	Poly.	Con.	Bits
NAV01	8.7	29.0	11	386	72	1123	45	386	20	62
NAV02	14.5	38.2	10	803	72	1126	45	803	20	33
NAV03	14.4	38.0	10	795	72	1126	45	795	20	33
NAV04	13.6	47.6	8	562	122	2057	85	562	18	32

v. The other analysis parameters were left at their default setting. While we need to specify bounds for the analysis region, we can handle the unbounded case by checking that the reachable state space is strictly contained in the analysis region. All instances shown were obtained with a-priori bounds of $[-2, 2]$ on the velocities, and the reachable velocities remained within an interval $[-1.1, 1.1]$, which confirms our a-priori bounds as valid. Figure 5 shows the set of reachable states computed by PHAVer as a result. For the instances NAV01–NAV04, the analysis was fairly straightforward, with $\delta_i = 1$. Higher instances require a much higher level of accuracy, and lead to large numbers of polyhedra. Applying a convex hull overapproximation would remedy this problem, but without complexity management the analysis did not terminate. We present results after introducing complexity management in the next section. In comparison, for a predicate abstraction tool the following times were reported in [39]: For NAV01–NAV03 34s, 153s (68MB) and 152s (180MB), respectively, on a Sun Enterprise 3000 (4 x 250 MHz UltraSPARC) with 1 GB RAM. In [13] it is reported that HSOLVER was not able to show safety of instances similar to NAV01–NAV03.

4 Managing the Complexity of Polyhedra

In exact fixed-point computations with polyhedra, the size of numbers in the formula as well as the number of constraints typically increases unless the structure of the hybrid system imposes boundaries, e.g., with resets or invariants. To keep the complexity manageable, we propose to simplify polyhedra in a strictly conservative fashion by limiting the number of bits, i.e., the size of coefficients, and the number of constraints. While showing good performance in practice, both methods have the disadvantage that the approximation error may be unbounded, and is not localized. In practice, both simplifications are applied when the number of bits or constraints exceeds a given threshold that is significantly higher than the reduction level. The resulting hysteresis between exact computations and overapproximations gives cyclic dependencies time to stabilize.

4.1 Limiting the Number of Bits

We consider a convex polyhedron in constraint representation, i.e., given as the conjunction of constraints $a_i^T x \bowtie_i b_i$, where a_i is a vector of the coefficients $a_{ij} \in \mathbb{Z}$, $i = 1, \dots, m$, $j = 1, \dots, n$, $\bowtie_i \in \{<, \leq\}$, and inhomogeneous coefficients $b_i \in \mathbb{Z}$. We assume that the a_{ij} and b_i have no common factor and that there are no redundant constraints. Note that we usually do not simplify equalities to preserve the affine dimension, i.e., the inhabited subspace, of the polyhedron, but they can be simplified by converting each equality into the conjunction of two inequalities. The goal is to find, with as little overapproximation as possible, a new constraint $\alpha_i^T x \bowtie_i \beta_i$ that contains all solutions of $a_i^T x \bowtie_i b_i$, and whose coefficients α_{ij}, β_i have less than z bits, i.e., $|\alpha_{ij}|, |\beta_i| \leq 2^z - 1$.

Our approach is to reduce scale down the a_{ij} , and then find a β_i that makes the constraint conservative. There is no a-priori estimate of how large β_i will have to be, so we obtain an initial estimate of a scaling factor, and reduce it if it results in a β_i that is too large. Because the new coefficients need to be integers, we need to account for rounding errors. We can represent the new coefficients as

$$\begin{aligned}\alpha_{ij} &= f a_{ij} + r_{ij}, \\ \beta_i &= f b_i + r_i,\end{aligned}$$


with a scaling factor $f > 0$, rounding errors r_{ij} , $|r_{ij}| \leq 0.5$ and an unbounded error r_i for the inhomogeneous term. In this representation, our goal is to find a scaling factor f that is close to 1 and results in a small error r_i . There is no a-priori bound on r_i because it depends on the new direction α_i and the other constraints that define the polyhedron.

We obtain an initial estimate of f based on the assumption that β_i is close to $f b_i$, say the closest integer. Since β_i must be rounded upwards to guarantee conservativeness, we get $|r_i| \leq 1$ as an optimistic estimate. This gives us the following upper bounds for f :

$$f \leq (2^z - 3/2)/|a_{ij}|, \text{ and} \quad (14)$$

$$f \leq (2^z - 2)/|b_i|. \quad (15)$$

To predict the effects of rounding on the new coefficients is difficult and would lead to a mixed integer linear program. We employ a heuristic algorithm, shown in Fig. 6.


 Fig. 5. Reachable states in the x_1, x_2 -plane (initial states darkest)

(a) NAV01

procedure *LimitConstraintBits*
Input: Polyhedron as a set of constraints

$$P = \{a_k^T x \bowtie_k b_k \mid k = 1, \dots, m\},$$

 index i to constraint to be limited,

 desired number of bits z
Output: new constraint $\alpha_i^T x \bowtie_i b_i$
 $success := false;$
 $f := \min\{(2^z - 3/2)/|a_{ij}|, (2^z - 2)/|b_i| \mid j = 1, \dots, n\};$
while $\neg success$ **do**

 for $j = 1, \dots, n$ **do** $\alpha_{ij} := round(fa_{ij})$ **od**;

 $q := \inf_{x \in P} \alpha_i^T x;$

 if $\alpha_i = 0$ **or** $q = -\infty$ **then abort fi**;

 $\beta_i := ceil(q);$

 if $|\beta_i| \leq 2^z - 1$ **then** $success := true$

 else $f := \min\{f/2 - 3/(4|a_{ij}|), (2^z - 2)/|$
 $\beta_i| \mid j = 1, \dots, n\}$ **fi**;

od.

(b) NAV02

(c) NAV03

(d) NAV04

Fig. 6. Limiting the number of bits of a constraint




Fig. 7. Limiting the number of bits of a constraint

Let $round(x)$ be a function that returns the next integer between x and zero, and $ceil(x)$ be a function that rounds to the next larger integer. First, we estimate f based on (14),(15), then we compute a new β_i using linear programming. If β_i has more than z bit, we decrease f and start over. The procedure is repeated until all $\alpha_{ij} = 0$, in which case the problem is infeasible. Note that even if it is feasible, the new polyhedron may not be bounded because the reduced constraints may be parallel. Figure 7 illustrates the basic scheme. The normal vector a_i of the constraint, shown in (a), is approximated by α_i , as shown in (b). Linear programming yields the inhomogeneous term q that makes the constraint tangent to the polyhedron, as in (c). Rounding of q yields β_i , and the polyhedron outlined in (d).

4.2 Limiting the Number of Constraints

A set of symbolic states is described by a linear predicate, the convex sub-predicates of which define convex polyhedra, which in turn are described by a set of constraints. The number of constraints usually increases with the number of iterations of a fixpoint computation, which forces us to limit this number to remain within a computationally feasible complexity. This reduction has been implemented in literature by means of, e.g., bounding boxes [40], or oriented rectangular hulls [33]. Instead, we propose to simply drop the least significant of the constraints, as this seems a good compromise in terms of accuracy and speed. In addition, the constraints in our applications are, as a whole, invariant with respect to the fixpoint computation. This invariance has a greater chance of being preserved as good as possible if we keep constraints instead of drawing up an entirely new set. As with limiting the number of bits, we usually chose to not limit equalities in order to preserve the affine dimension of the polyhedron. If an equality is to be limited, it must be replaced by two inequalities.

We measure the significance of a constraint based on a criterion *crit* that measures the the difference between the polyhedron with and without the constraint. Let P be a set of linear constraints describing a convex polyhedron, and $P^{\setminus i} = P \setminus \{a_i^T x \bowtie_i b_i\}$ be the polyhedron without it's i th constraint. Then the difference between the points contained P and $P^{\setminus i}$ is the polyhedron $P^{-i} = P^{\setminus i} \cup \{-a_i^T x \bowtie_i -b_i\}$, where $(\bowtie_i, \bar{\bowtie}_i) \in \{(<, \leq), (\leq, <)\}$, obtained by simply replacing the i th constraint with its complement. It has less non-redundant constraints than P and is therefore preferable in the formulations below. We consider three methods:

1. volumetric: Let $V(P)$ be the volume of the points contained in P . Then $crit = V(P^{\setminus i}) - V(P) = V(P^{-i})$. Requires P^{-i} to be bounded.
2. slack: Let $b_{max} = \max_x a_i^T x$ s.t. $x \in P^{-i}$. Then $crit = (b_{max} - b_i) / \|a_i\|$, i.e., the distance, measured in the direction of the constraint, between the points farthest apart in P^{-i} . Requires P^{-i} to be bounded in the direction of a_i .
3. angle: $crit = -\max_{j \neq i} a_j^T a_i$. Measures the negative cosine of the closest angle between the normal vector of the i th constraint and all others.

Our goal is to select the z most important out of m original constraints. It would be prohibitively expensive to examine all $\binom{m}{z}$ possible combinations of constraints. Instead, we consider two heuristics:

1. deconstruction: Starting from the entire set of constraints, drop the $m - z$ constraints with the least effect according to *crit*.
2. reconstruction: Starting from an empty set of constraints, add the z constraints with the greatest effect according to *crit*.

The criteria based on volume and slack require the initial polyhedron to be bounded, for which one could use, e.g., the invariant of the location. The following example shall illustrate the difference between volumetric and angle criteria, and its potential unboundedness.

Example 6. Consider the polyhedron shown in Fig. 8(a). It has 6 constraints A–F, whose angles with the neighbors are noted in the graph. In a volume based deconstruction with 5 constraints, constraint A is removed since that causes the smallest change in volume. The resulting polyhedron is shown hashed in Fig. 8(b). The angle based reconstruction with 5 constraints results in the shaded polyhedron in Fig. 8(b), where the constraints are labeled in the order they are chosen: First, an arbitrary initial constraint is chosen, say constraint C. The second choice is the constraint that has the largest angle with C, i.e., that is most opposed to it. In this case, this is constraint F, since it has an angle of 180° with C. The third choice is the one that is most opposed to both C and F, here constraint B because it has an angle of 90° with both C and F. The fourth constraint is A, with minimum angles of 45° , and the fifth is D with a minimum angle of 30° . Figure 8(c) shows the reduction to 4 constraints. Here the angle based method results in an unbounded polyhedron because constraint D is not chosen. An algorithm should take this possibility into account and test for boundedness.

The construction method with an angle criterion was the fastest in our experiments. The angle calculations can be sped up by using a look-up table $\alpha(i, j)$ that maps an angle to every pair of constraints. This yields an algorithm of complexity $O(nm^2 + m^3)$, shown in Fig. 9, where C is the set of candidate constraints and H is the set of chosen constraints. It includes a test that preserves the boundedness of P . H is initialized with the set of equalities, which are not reduced to preserve the affine dimension of the polyhedron, and an arbitrary initial constraint. Here we choose the one with the smallest coefficients. In a while-loop, the constraint is chosen based on the best of the worst-cases, i.e., the smallest angle with the constraints in H . Since $a_j^T a_i$ is the cosine of the angle, choosing the smallest angle translates into maximizing $a_j^T a_i$. The constraint is added to H and removed from the candidates C , and the procedures is repeated until $|H| \geq z$ and the boundedness of P implies boundedness of H . This algorithm is in our implementation $\sim 1000\times$ faster than a slack based deconstruction for limiting 400 constraints down to 32 in 4 dimensions.

The following example shall illustrate that complexity management is useful to induce termination, and that the constraint-based overapproximation of dynamics may be preferable when a projection-based overapproximation leads to prohibitively complex sets of states. Still, the relationship between overapproximation method, complexity and quick convergence is not a trivial one.

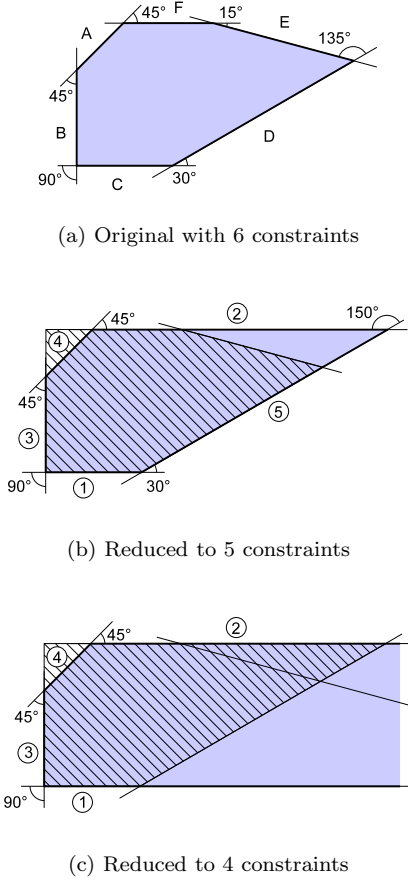


Fig. 8. Example for limiting the number of constraints by volumetric deconstruction (hashed) and angle based reconstruction (shaded)

procedure *LimitConstraintsByAngle*

Input: Polyhedron P as a set of constraints

$$a_i^T x \bowtie_i b_i, \quad i = 1, \dots, m,$$

desired number of constraints z

Output: Polyhedron H

for $i = 1, \dots, m, j = 1, \dots, m, j > i$ **do**

$$\alpha(i, j) := a_i^T a_j$$

od;

$$H := \{a_k^T x \bowtie_k b_k \mid k = \operatorname{argmin}_k (\max_j |a_{kj}|)\} \cup \{a_i^T x \bowtie_i b_i \mid \bowtie_i \in \{=\}\};$$

$$C := P \setminus H;$$

while $(|C| > 0 \wedge (|H| < z \vee (\operatorname{bounded}(P) \wedge \neg \operatorname{bounded}(H))))$ **do**

$$j = \operatorname{argmin}_j (\max_i \alpha(i, j)) \text{ s.t.}$$

$$a_i^T x \bowtie_i b_i \in H, a_j^T x \bowtie_j b_j \in C;$$

$$H := H \cup \{a_j^T x \bowtie_j b_j\};$$

$$C := C \setminus \{a_j^T x \bowtie_j b_j\}$$

od.

Fig. 9. Reconstructing a polyhedron with a limited number of constraints by angle prioritization

Example 7. Consider the navigation benchmark given in Sect. 3.4. Instances NAV01–NAV04 can be analyzed without complexity management, but the computed reachable set consists of a large number of polyhedra. This can be remedied by including a convex hull operation in the successor computation, for each location individually. This computation, however, does not terminate as the solution converges towards nonrational numbers. Limiting the number of bits remedies this problem, as the results in Table 3 show. The constraints were limited to 24 bits, with a triggering threshold of 300, and to 48 constraints with a threshold of 96. This reduction cut the computation time and memory in half, and reduced the number of polyhedra by almost a factor of 20. This should allow us to analyze instances of the navigation benchmark that are larger. The analysis of NAV07, however, does not finish within reasonable time. The computed polyhedra consist of up to 944 constraints at 311 bits. This is a result of using the projection based overapproximation of the affine dynamics, which has the potential to drastically increase number of constraints. If instead we use the constraint based overapproximation with constraints in the direction of v_1 and v_2 , see NAV07 in the bottom half of Table 3, the analysis successfully terminates after 25 s, using less than 40 constraints in most iterations. On the other hand, NAV04 does not terminate when using constraint based overapproximation, and requires a smaller partition size to show that the forbidden states are not reachable.

4.3 Example: Tunnel-Diode Oscillator Circuit

We present experimental results for a tunnel-diode oscillator circuit taken from [26] with the parameters used in [41]. We model the current I_L through the inductor and the voltage drop V_d of a tunnel diode in parallel with the capacitor of a serial RLC circuit, which are in stable oscillation for the given parameters. The state equations are given by

$$\dot{V}_d = 1/C(-I_d(V_d) + I_L), \quad (16)$$

$$\dot{I}_L = 1/L(-V_d - 1/G \cdot I_L + V_{in}), \quad (17)$$

where $C = 1 \text{ pF}$, $L = 1 \text{ } \mu\text{H}$, $G = 5 \text{ m}\Omega^{-1}$, $V_{in} = 0.3 \text{ V}$, and the diode current

$$I_d = \begin{cases} 6.01V_d^3 - 0.992V_d^2 + 0.0545V_d & \text{if } V_d \leq 0.055, \\ 0.0692V_d^3 - 0.0421V_d^2 + 0.004V_d + 8.96 \cdot 10^{-4} & \text{if } 0.055 \leq V_d \leq 0.35, \\ 0.263V_d^3 - 0.277V_d^2 + 0.0968V_d - 0.0112 & \text{if } 0.35 \leq V_d. \end{cases}$$

Following the procedure outlined in the previous section, a piecewise affine envelope was constructed for the tunnel diode characteristic $I_d(V)$. We choose 32 intervals for the range $V_d \in [-0.1, 0.6]$ to yield sufficient accuracy and so obtain a piecewise affine model for (16). It

Table 3. PHAVer performance, navigation benchmark with convex hull overapproximations and complexity management

Instance	Time [s]	Mem. [MB]	Iter.	Checks	Automaton		Reachable Set			
					Loc.	Trans.	Loc.	Poly.	Con.	Bits
<i>Projection-based overapproximation of dynamics</i>										
NAV01	5.2	25.6	22	185	70	1090	43	43	30	308
NAV02	6.4	25.6	20	230	72	1126	45	45	42	282
NAV03	6.5	25.6	20	232	72	1126	45	45	48	290
NAV04	6.7	25.9	52	330	72	1116	54	54	18	32
NAV07 ^a	>250k	~100.0	>11k	>738k	224	4283	224	224	994	311
<i>Constraint-based overapproximation of dynamics</i>										
NAV01	4.6	24.0	7	135	72	1126	45	45	16	33
NAV02	4.7	26.2	6	137	72	1126	45	45	16	37
NAV03	4.8	26.2	6	137	72	1126	45	45	16	40
NAV04 ^b	>1k	~75.0	>100	>5k	220	4350	124	124	777	340
NAV07	24.9	85.5	32	898	220	4211	192	192	95	369

^a no termination, computation aborted

^b forbidden states only unreachable when partition size reduced from $\delta_i = 1$ to $\delta_i = 0.5$

is modeled as a hybrid automaton with V_d as an output- and I_L as an input-variable, and consists of 32 locations, one for each interval. Equation (17) is affine, and is modeled as a hybrid automaton with V_d as an input and a single location. Both models are composed and analyzed in PHAVer. Figure 10(a) shows the computed reachable states for initial states given by $V_d \in [0.42V, 0.52V]$, $I_L = 0.6mA$. It also shows the invariants (dashed) generated by the partitioning algorithm using the constraints $Cand = \{(V_d, 0.7/128, \infty), (I_L, 1.5/128, \infty)\}$, i.e., max. 128 partitions in both directions, and splitting criterion (c_3, c_1) with $\triangleleft_{min} = 0.99$. The analysis with PHAVer takes 17.1s and 68.0MB RAM, with the largest coefficient taking up 2508 bits and at most 5 constraints per polyhedron.

We now obtain bounds on the cycle time of the oscillator by composing the circuit model with a monitor automaton. The cycle time is the maximum time it takes any state to cross the threshold $I = 0.6\mu A$, $V > 0.25V$ twice. For the clocked circuit, the number of bits and constraints grows rapidly and a more precise analysis, such as shown in Fig. 10(b) is only possible with limits on both. We compare the unlimited and limited analysis for constraint candidates $Cand = \{(V_d, 0.7/64, \infty), (I_L, 1.5/64, \infty)\}$ and using convex hull overapproximations for the flows. The bits are limited to 24 when a threshold of 300 bits is reached, and the constraints to 32 with a threshold of 200. Figures 11(a) and 11(b) show a polynomial increase in the number of constraints, and an exponential increase of the number of bits in the new polyhedra found at each iteration. The analysis takes 1412 s (381 MB) when unlimited, and 90 s (132 MB) when limited and yields the reachable states shown in Fig. 10(b). The bounds for on the cycle time are $[12.75, 14.90]\mu s$ when unlimited. The relative error of the limited analysis is 0.006% for the lower bound and 0.08%


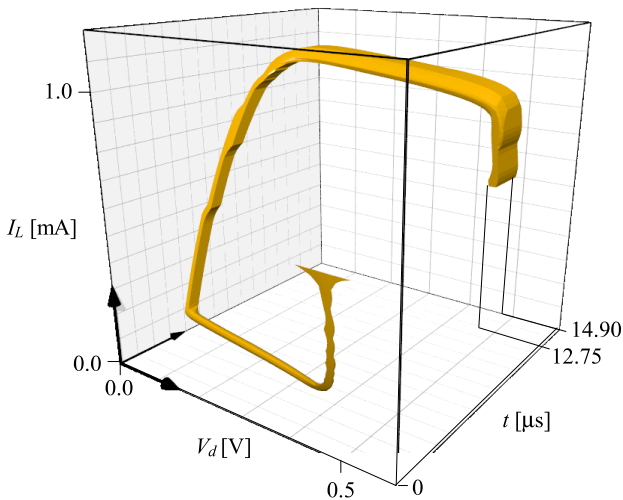
for the upper bound. At a more than fifteenfold increase in speed, the overapproximation is negligible and results in a cycle time estimate that is practically identical. Note that a comparison was only possible at the lowest level of accuracy. Higher accuracy can not be achieved at all without limiting due to an drastic increase in computation time.

In comparison with CheckMate [41], PHAVer is able to analyze the circuit at higher accuracy, and obtains results for parameters where CheckMate does not [8].

5 Conclusions


PHAVer, a new tool for verifying safety properties of linear hybrid automata, provides infinite precision arithmetic in a robust implementation, on-the-fly overapproximation of affine dynamics, and supports compositional and assume/guarantee-reasoning. To manage the complexity of the underlying polyhedral computations, we propose heuristics to conservatively overapproximate polyhedra by limiting the number of bits and constraints. Experimental results for a navigation benchmark and a tunnel diode circuit demonstrate the effectiveness of the approach, with PHAVer outperforming other current verification tools. The results indicate that the benefits of exact polyhedral computations, such as exact containment testing, can outweigh the costs incurred by the indispensable complexity management. Future research will focus on accelerating convergence and termination, e.g., using widening [36], and using the transition topology to improve the search algorithm. PHAVer is available at <http://www.cs.ru.nl/~goranf/>.

Acknowledgements. The author is most grateful for the numerous inspiring discussions with Prof. Bruce Krogh, whose


(a) V_d - I_L -Plane, invariants dashed(b) V_d and I_L over time**Fig. 10.** Reachable states of Tunnel Diode Circuit

insightful guidance was indispensable in this work, and to Prof. Frits W. Vaandrager and Prof. Sebastian Engell for their generous support and supervision. This research was supported in part by the US ARO contract no. DAAD19-01-1-0485, the US NSF contract no. CCR-0121547, and the Semiconductor Research Corporation under task ID 1028.001.

Special thanks go to R.J.M. Theunissen, Gabriela Marin, Laurent Doyen, Li Hong and Scott Little for bug reports and helpful suggestions, and to Flavio Lerda for his invaluable compilation scripts. Finally, a big thanks goes to the PPL development team and in particular to Roberto Bagnara for his wonderful help.



(a) Number of bits



(b) Number of constraints

Fig. 11. Clocked Tunnel Diode Circuit, exact (dashed) and with limits on bits and constraints (solid)

References

1. Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. HYTECH: A model checker for hybrid systems. *Int. Journal on Software Tools for Technology Transfer*, 1(1–2):110–122, December 1997.
2. Moshe Y. Vardi and Pierre Wolper. An automata-theoretic approach to automatic program verification (preliminary report). In *LICS*, pages 332–344. IEEE Computer Society, 1986.
3. Thomas A. Henzinger. The theory of hybrid automata. In *Proc. 11th Annual IEEE Symposium on Logic in Computer Science, LICS'96, New Brunswick, New Jersey, 27–30 July 1996*, pages 278–292. IEEE Computer Society Press, 1996. An extended version appeared in *Verification of Digital and Hybrid Systems* (M.K. Inan, R.P. Kurshan, eds.), NATO ASI Series F: Computer and Systems Sciences, Vol. 170, Springer-Verlag, 2000, pp. 265–292.
4. Rajeev Alur, Thomas A. Henzinger, and Pei-Hsin Ho. Automatic symbolic verification of embedded systems. *IEEE Transactions on Software Engineering*, 22:181–201, 1996. A preliminary version appeared in the Proceedings of the 14th Annual Real-Time Systems Sympos-

- sium (RTSS), IEEE Computer Society Press, 1993, pp. 2-11.
5. Thomas A. Henzinger, Pei-Hsin Ho, and Howard Wong-Toi. Algorithmic analysis of nonlinear hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):540–554, 1998.
 6. Goran Frehse, Zhi Han, and Bruce H. Krogh. Assume-guarantee reasoning for hybrid i/o-automata by over-approximation of continuous interaction. In *Proc. 43rd IEEE Conf. Decision and Control (CDC'04), December 14–17, 2004, Atlantis, Bahamas, 2004*.
 7. Goran Frehse. *Compositional Verification of Hybrid Systems Using Simulation Relations*. PhD thesis, Radboud University Nijmegen, 2005. available at http://webdoc.ubn.ru.nl/mono/f/frehse_g/compveofh.pdf.
 8. Goran Frehse, Bruce H. Krogh, Rob A. Rutenbar, and Oded Maler. Time domain verification of oscillator circuit properties. 153(3):9–22, 2006.
 9. G. Frehse, B. H. Krogh, and R. A. Rutenbar. Verifying analog oscillator circuits using forward/backward refinement. In *Proceedings of the 9th Conference on Design, Automation and Test in Europe (DATE 06)*. ACM SIGDA, Munich, Germany, 2006.
 10. D.A. van Beek, K.L. Man, M.A. Reniers, J.E. Rooda, and R.R.H. Schiffelers. Formal verification of chi models using phaver. In I. Troch and F. Breitenacker, editors, *Proc. MathMod 2006, Vienna*, ARGESIM Reports, February 2006.
 11. Andreas Podelski and Silke Wagner. Model checking of hybrid systems: From reachability towards stability. In João P. Hespanha and Ashish Tiwari, editors, *HSCC*, volume 3927 of *Lecture Notes in Computer Science*, pages 507–521. Springer, 2006.
 12. Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Automatic rectangular refinement of affine hybrid systems. In *Proc. FORMATS'05*, volume 3829 of *LNCS*, pages 144–161. Springer, 2005.
 13. I. Ben Makhoul and S. Kowalewski. An evaluation of two recent reachability analysis tools for hybrid systems. In *Proc. IFAC Conf. Analysis and Design of Hybrid Systems (ADHS'06)*, 2006.
 14. Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, Pei-Hsin Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138(1):3–34, 1995. A preliminary version appeared in the Proceedings of the 11th International Conference on Analysis and Optimization of Systems: Discrete-Event Systems (ICAOS), Lecture Notes in Control and Information Sciences 199, Springer-Verlag, 1994, pp. 331-351.
 15. Nancy A. Lynch, Roberto Segala, and Frits W. Vaandrager. Hybrid I/O automata. *Information and Computation*, 185(1):105–157, 2003. URL <http://www.cs.kun.nl/ita/publications/papers/fvaan/HIOA.html>.
 16. Darren D. Cofer, Eric Engstrom, Robert P. Goldman, David J. Musliner, and Steve Vestal. Applications of model checking at Honeywell Laboratories. In Matthew B. Dwyer, editor, *Model Checking Software, 8th International SPIN Workshop, Toronto, Canada, May 19–20, 2001, Proceedings*, volume 2057 of *LNCS*, pages 296–303. Springer, 2001.
 17. Thomas A. Henzinger, Joerg Preussig, and Howard Wong-Toi. Some lessons from the hytech experience. In *Proceedings of the 40th Annual Conference on Decision and Control (CDC'01)*, pages pp. 2887–2892. IEEE Press, 2001.
 18. Stefan Kowalewski, Olaf Stursberg, Martin Fritz, Holger Graf, Ingo Hoffmann, Joerg Preussig, Manuel Remelhe, Silke Simon, and Heinz Treseler. A case study in tool-aided analysis of discretely controlled continuous systems: The two tanks problem. In Panos J. Antsaklis, Wolf Kohn, Michael D. Lemmon, Anil Nerode, and Shankar Sastry, editors, *Hybrid Systems V*, volume 1567 of *LNCS*, pages 163–185. Springer, 1999.
 19. Claire Tomlin. Verification of an air traffic management protocol using hytech. Course Project for EE290A, taught by Prof. T. A. Henzinger, Spring 1996, Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, 1996.
 20. Thomas A. Henzinger and Howard Wong-Toi. Using HyTech to synthesize control parameters for a steam boiler. In *Formal Methods for Industrial Applications: Specifying and Programming the Steam Boiler Control*, number 1165 in *LNCS*, pages 265–282. Springer Verlag, 1996.
 21. Olaf Mller and Thomas Stauner. Modelling and verification using linear hybrid automata - a case study. *Mathematical and Computer Modelling of Dynamical Systems*, 6(1):71–89, 2000.
 22. Roberto Bagnara, Elisa Ricci, Enea Zaffanella, and Patricia M. Hill. Possibly not closed convex polyhedra and the Parma Polyhedra Library. In Manuel V. Hermenegildo and German Puebla, editors, *Static Analysis: Proc. of the 9th Int. Symposium*, volume 2477 of *LNCS*, pages 213–229. Springer, Madrid, Spain, 2002.
 23. Torbjörn Granlund and Kevin Ryde. The GNU Multiple Precision arithmetic library version 4.0, 2001. <http://www.swox.com/gmp/>.
 24. T. A. Henzinger, Pei-Hsin Ho, and H. Wong-Toi. HYTECH: the next generation. In *Proceedings of the 16th IEEE Real-Time Systems Symposium (RTSS '95)*, page 56. IEEE Computer Society, 1995.
 25. Olaf Stursberg and Stefan Kowalewski. Approximating switched continuous systems by rectangular automata. In *Proc. 5th European Control Conference, Karlsruhe, 1999*.
 26. Walter Hartong, Lars Hedrich, and Erich Barke. On discrete modeling and model checking for nonlinear analog systems. In Ed Brinksma and Kim Guldstrand Larsen, editors, *Computer Aided Verification, 14th International Conference, CAV 2002, Copenhagen, Denmark, July 27–31, 2002, Proceedings*, volume 2404 of *LNCS*, pages 401–413. Springer, 2002.
 27. Thomas A. Henzinger, Benjamin Horowitz, Rupak Majumdar, and Howard Wong-Toi. Beyond HYTECH: Hybrid systems analysis using interval numerical methods. In Nancy A. Lynch and Bruce H. Krogh, editors, *Hybrid Systems: Computation and Control, Third International Workshop, HSCC 2000, Pittsburgh, PA, USA, March 23–25, 2000, Proceedings*, volume 1790 of *LNCS*, pages 130–144. Springer, 2000.
 28. Stefan Ratschan and Zhikun She. Safety verification of hybrid systems by constraint propagation based abstraction refinement. In Manfred Morari and Lothar Thiele,

- editors, *Proc. of the 8th International Workshop on Hybrid Systems: Computation and Control*, number 3414 in LNCS, pages 573–589. Springer, 2005.
29. B. I. Silva, O. Stursberg, B. H. Krogh, and S. Engell. An assessment of the current status of algorithmic approaches to the verification of hybrid systems. In *Proc. 40th Conference on Decision and Control (CDC'01)*, December 2001.
 30. Howard Wong-Toi. Symbolic approximations for verifying real-time systems, December 1994.
 31. Jörg Preußig, Stephan Kowalewski, Howard Wong-Toi, and Thomas A. Henzinger. An algorithm for the approximative analysis of rectangular automata. In *Proceedings of the Fifth International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT)*, number 1486 in LNCS, pages 228–240. Springer-Verlag, 1998.
 32. Jörg Preußig. *Formale Überprüfung der Korrektheit von Steuerungen mittels rektangulärer Automaten*. PhD thesis, Schriftenreihe des Lehrstuhls für Anlagensteuerungstechnik Band 4/2000, Universität Dortmund, Shaker Verlag, 2000. (in German).
 33. Olaf Stursberg and Bruce H. Krogh. Efficient representation and computation of reachable sets for hybrid systems. In Oded Maler and Amir Pnueli, editors, *Hybrid Systems: Computation and Control, 6th International Workshop, HSCC 2003 Prague, Czech Republic, April 3-5, 2003, Proceedings*, volume 2623 of LNCS, pages 482–497. Springer, 2003.
 34. Rajeev Alur, Costas Courcoubetis, Thomas A. Henzinger, and Pei-Hsin Ho. Hybrid automata: An algorithmic approach to the specification and verification of hybrid systems. In Robert L. Grossman, Anil Nerode, Anders P. Ravn, and Hans Rischel, editors, *Hybrid Systems*, volume 736 of LNCS, pages 209–229. Springer, 1993.
 35. Gerardo Lafferriere, George J. Pappas, and Sergio Yovine. Symbolic reachability computation for families of linear vector fields. *Journal of Symbolic Computation*, 32:231–253, 2001.
 36. Nicolas Halbwachs, Yann-Erick Proy, and Patrick Roumanoff. Verification of real-time systems using linear relation analysis. *Formal Methods in System Design: An International Journal*, 11(2):157–185, August 1997.
 37. Pei-Hsin Ho and Howard Wong-Toi. Automated analysis of an audio control protocol. In *Proc. Conf. on Computer-Aided Verification*, volume 939 of LNCS, pages 381–394. Springer, Liege, Belgium, 1995.
 38. Ansgar Fehnker and Franjo Ivancic. Benchmarks for hybrid systems verification. In Rajeev Alur and George J. Pappas, editors, *Hybrid Systems: Computation and Control, 7th International Workshop, HSCC 2004, Philadelphia, PA, USA, March 25-27, 2004, Proceedings*, volume 2993 of LNCS, pages 326–341. Springer, 2004.
 39. Franjo Ivancic. *Modeling and Analysis of Hybrid Systems*. PhD thesis, University of Pennsylvania, Philadelphia, PA, December 2003.
 40. Alberto Bemporad and Manfred Morari. Verification of hybrid systems via mathematical programming. In Frits W. Vaandrager and Jan H. van Schuppen, editors, *Hybrid Systems: Computation and Control, Second International Workshop, HSCC'99, Berg en Dal, The Netherlands, March 29-31, 1999, Proceedings*, volume 1569 of LNCS, pages 31–45. Springer, 1999.
 41. Smriti Gupta, Bruce H. Krogh, and Rob A. Rutenbar. Towards formal verification of analog designs. In *Proc. IEEE Intl. Conf. on Computer-Aided Design (ICCAD-2004), Nov. 7-11, 2004, San Jose CA (USA)*, 2004.