

Chapter 30

Verification of Hybrid Systems

Laurent Doyen, Goran Frehse, George J. Pappas, and André Platzer

Abstract Hybrid systems are models which combine discrete and continuous behavior. They occur frequently in safety-critical applications in various domains such as health care, transportation, and robotics, as a result of interactions between a digital controller and a physical environment. They also have relevance in other areas such as systems biology, in which the discrete dynamics arises as an abstraction of fast continuous processes. One of the prominent models is that of *hybrid automata*, where differential equations are associated with each node, and jump constraints such as guards and resets are associated with each edge.

In this chapter, we focus on the problem of model checking of hybrid automata against reachability and invariance properties, enabling the verification of general temporal logic specifications. We review the main decidability results for hybrid automata, and since model checking is in general undecidable, we present three complementary analysis approaches based on symbolic representations, abstraction, and logic. In particular, we illustrate polyhedron-based reachability analysis, finite quotients, abstraction refinement techniques, and logic-based verification. We survey important tools and application domains of successful hybrid system verification in this vibrant area of research.

L. Doyen
LSV, CNRS & ENS Paris-Saclay, Cachan, France

G. Frehse
Verimag, University Grenoble Alpes, Grenoble, France

G.J. Pappas
University of Pennsylvania, Philadelphia, PA, USA

A. Platzer (✉)
Carnegie Mellon University, Pittsburgh, PA, USA
e-mail: aplatzer@cs.cmu.edu

30.1 Introduction

Information technology (IT) has dramatically changed our lives. The first revolution in information technology led to the birth of the computer. The second information technology revolution led to the creation of the Internet, connecting computing around the world and resulting in the hyper-connected world that we live in. The third revolution that is now taking place is connecting all the computational power to the physical world. Computing powerhouses, such as Intel, are investing in wearable computing and smart watches, Google has invested in self-driving cars and bought Nest, the makers of a learning thermostat, thereby connecting Google to building control and energy markets, and Amazon is investing in robotics and unmanned aerial vehicles. Similarly the most significant innovation in automotive companies recently came from software-intensive car technology, leading from adaptive cruise control to driverless cars. There is a plethora of novel medical devices, either wearable or implantable, that sense patient vitals and use computer algorithms to diagnose medical conditions or even perform life-critical functions. The fundamental aspect of this third revolution is the fusion of IT with physical devices that interact with the physical world.

The marriage of IT with the physical world is known as embedded computing as it consists of computing that is embedded and tightly interacts with the physical world. A major modeling challenge is how to formally capture the interaction between computing and physics so that we can *reason about the effect of physics on computing and vice versa*. This led to the development of *hybrid systems* [8, 34, 35, 91, 132], where both discrete and continuous behaviors of the system are important. Hybrid systems grew out of the necessity to enrich purely digital models of computing with analog models of physics. As a result, hybrid systems contain both digital models of computing (such as automata or programs) as well as analog elements (such as differential equations) integrated in such a way that one can model many embedded computing applications.

The need for formal models of hybrid systems arises from the fact that many embedded computing problems are safety-critical. They arise in collision avoidance protocols in air traffic control [31, 100, 113, 129, 149, 174–177], cruising controllers for automotive vehicles [24, 46, 53, 62, 101, 114, 164, 168], obstacle avoidance algorithms for autonomous ground robots [130, 182], and software-controlled medical devices that actively regulate life-critical functions or help surgeons with surgical robotic systems [104]. Therefore there is not only a need for formal models of embedded computing, but also for rigorous verification approaches to guarantee that the embedded computation, as modeled by a hybrid system, is formally safe. This has resulted in the development of a new paradigm within the formal methods community, namely the formal verification of hybrid models of embedded computing.

There is a range of formal models for hybrid systems [8, 27, 29, 35, 52, 91, 118, 120, 131, 132, 134, 135, 143, 167, 179, 180], each with different advantages for different purposes. This chapter focuses on *hybrid automata* [8, 91], because they directly generalize the timed automata that have been considered in Chap. 29. The basic idea in hybrid automata is to associate differential equations with the nodes of

an automaton. The automaton structure defines how and under which conditions the system switches between the various differential equations and what happens to the state if they switch. Timed automata, which are discussed in Chap. 29, are a special case of hybrid automata, where all differential equations are of the form $\dot{x} = 1$ such that x is a clock variable measuring the progress of time and additional linearity assumptions are met for the switching conditions. Timed automata are an interesting subclass of hybrid automata, because reachability is decidable in this subclass. For systems with more general continuous dynamics, e.g., moving, acceleration, or curving, however, timed automata are not sufficient, and hybrid systems models are needed instead.

In this chapter, we give a survey of model-checking techniques for hybrid systems with an emphasis on the handling of continuous dynamics. It has been proved that continuous dynamics verification is the most fundamental question in hybrid systems verification [135, 141], because discrete dynamics can be verified exactly as well as continuous dynamics.

In this section, we survey a set of complementary verification techniques for hybrid systems, including explicit-state reachability computations with termination criteria like bounded-horizon (Sect. 30.4, which is related to Chap. 5), abstraction techniques and abstraction refinement loops (Sect. 30.5, also see Chap. 13), and logic-based verification approaches (Sect. 30.6, which is related to Chaps. 2, 26, 20, and 15). Other surveys of several aspects of hybrid systems can be found in the literature [7, 12, 39, 80, 103, 143, 163, 171, 172]. A control-theoretic view on hybrid systems verification has been reported in a book by Tabuada [165]. A logic and proofs view on hybrid systems verification can be found in a book by one of the authors [137]. Introductions to embedded systems from a cyber-physical systems perspective have been reported in the literature [111, 122] and in university courses.

Hybrid systems has become a very active and successful area of research with a vibrant community. Giving a complete overview of all relevant approaches is impossible in this chapter, instead we focus on giving an overview of some of the most important representative classes of techniques. By their very nature, hybrid systems tend to be mathematically demanding, but they can also be exceedingly beautiful. The broad applicability and scope of the resulting hybrid systems analysis techniques make hybrid systems a very rewarding area of science with the potential for significant impact on practical applications.

30.2 Basic Definitions

Hybrid systems combine discrete evolutions (namely, mode changes and variable updates) and continuous evolutions through variables whose dynamics is governed by differential equations. Hybrid system models have been introduced to deal with such systems in a uniform way [8, 27, 29, 35, 52, 91, 118, 120, 131, 132, 134–136, 167, 179, 180]. The original definitions are very general. In this chapter, we focus on subclasses of particular interest. Timed automata are an important class of hybrid automata for which safety verification is decidable (see Chap. 29). When

continuous variables are subject to rectangular flow constraints, that is constraints of the form $\dot{x} \in [a, b]$, hybrid automata are called *rectangular*. For that subclass of hybrid automata, there exists a reasonably efficient algorithm to compute the image of a (simple) set. Based on this algorithm, there exists an iterative method that computes the exact set of reachable states when it terminates. This semi-algorithm can be used to establish or refute *safety properties*. On the other hand, if the evolution of the continuous variables is subject to more complicated flow constraints, for example affine dynamics like $\dot{x} = 3x - y$, computing the flow successor is much more difficult and only approximate methods are known.

30.2.1 Predicates

Let $X = \{x_1, \dots, x_n\}$ be a finite set of variables. Given a valuation $v : X \rightarrow \mathbb{R}$ and $Y \subseteq X$, define $v|_Y : Y \rightarrow \mathbb{R}$ by $v|_Y(x) = v(x)$ for every $x \in Y$.

Definition 1 (Polynomial term) A *polynomial term* over a finite set of variables $X = \{x_1, \dots, x_n\}$ is an expression of the form $y \equiv \sum_{i \in \mathbb{N}^n} a_i x_1^{i_1} \dots x_n^{i_n}$ where $a_i \in \mathbb{Q}$ ($i = (i_1, \dots, i_n) \in \mathbb{N}^n$) are rational constants and almost all a_i are zero. Given a valuation v over X , we write $\llbracket y \rrbracket_v$ for the real number $\sum_{i \in \mathbb{N}^n} a_i v(x_1)^{i_1} \dots v(x_n)^{i_n}$ obtained by evaluating the polynomial term at v . We denote by $\text{PTerm}(X)$ the set of all polynomial terms over the variables X .

Definition 2 (Polynomial constraint) A *polynomial constraint* over X (also known as a semi-algebraic constraint) is a finite formula φ defined by the following grammar rule:

$$\varphi ::= \theta \bowtie 0 \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$$

where $\theta \in \text{PTerm}(X)$ and $\bowtie \in \{<, \leq, =, >, \geq\}$. We denote by $\text{PConstr}(X)$ the class of polynomial constraints over the set of variables X .

Definition 3 Given a valuation $v : X \rightarrow \mathbb{R}$ and a polynomial constraint $\varphi \in \text{PConstr}(X)$, we write $v \models \varphi$ and say that v *satisfies* φ , which we define inductively as:

- $v \models \theta \bowtie 0$ if $\llbracket \theta \rrbracket_v \bowtie 0$,
- $v \models \varphi_1 \wedge \varphi_2$ if $v \models \varphi_1$ and $v \models \varphi_2$,
- $v \models \varphi_1 \vee \varphi_2$ if $v \models \varphi_1$ or $v \models \varphi_2$.

We also write $v \in \llbracket \varphi \rrbracket$ when $v \models \varphi$. If $v : X \rightarrow \mathbb{R}$ and $w : Y \rightarrow \mathbb{R}$ are two valuations for disjoint sets of variables X, Y (with $X \cap Y = \emptyset$), we also write $(v, w) \in \llbracket \varphi \rrbracket$ when $u \models \varphi$ where $u : X \cup Y \rightarrow \mathbb{R}$ is defined such that $u|_X = v$ and $u|_Y = w$.

Linear constraints are an important special case of polynomial constraints. The set of solutions of a linear constraint describes a set of polyhedra. This geometric

interpretation is sometimes used for model checking since image computations of polyhedra can be quite efficient.

Definition 4 (Linear constraint) A *linear term* is a polynomial term of the form $y \equiv a_0 + \sum_{x_i \in X} a_i x_i$ with $a_i \in \mathbb{Q}$. We denote the set of all linear terms over X by $\text{LTerm}(X)$. A *linear constraint* is a polynomial constraint where all terms are linear. It is called *conjunctive* if it does not contain any disjunctions. We denote by $\text{LConstr}(X)$ the class of linear constraints over X and by $\text{LConstr}_c(X)$ the class of conjunctive linear constraints. The constraints true and false are defined as abbreviations in a standard way.

Definition 5 (Polyhedron) A set of valuations that can be defined by a conjunctive linear constraint is called a *polyhedron*, and a closed and bounded polyhedron is called a *polytope*. We denote a polyhedron in its *constraint representation* as

$$P = \left\{ x \mid \bigwedge_{i=0}^m a_i^\top x \bowtie_i b_i \right\}, \quad \text{with } \bowtie_i \in \{<, \leq, =, >, \geq\},^1$$

where the $a_i \in \mathbb{Q}^n$ are called *facet normals* and the $b_i \in \mathbb{Q}^n$ constants. In vector-matrix notation, this corresponds to

$$P = \{x \mid Ax \bowtie b\}, \quad \text{with } A = \begin{pmatrix} a_1^\top \\ \vdots \\ a_m^\top \end{pmatrix}, \quad \bowtie = \begin{pmatrix} \bowtie_1 \\ \vdots \\ \bowtie_m \end{pmatrix}, \quad b = \begin{pmatrix} b_1 \\ \vdots \\ b_m \end{pmatrix}.$$

A closed polyhedron $P \subseteq \mathbb{R}^n$ can be represented by a pair (V, R) , called the *generators* of P , where $V \subseteq \mathbb{Q}^n$ is a finite set of *vertices*, and $R \subseteq \mathbb{Q}^n$ is a finite set of *rays*, with:

$$P = \left\{ \sum_{v_i \in V} \lambda_i \cdot v_i + \sum_{r_j \in R} \mu_j \cdot r_j \mid \lambda_i \geq 0, \mu_j \geq 0, \sum_i \lambda_i = 1 \right\}.$$

The representation can be extended with *closure points* to deal with non-closed polyhedra [22].

There are algorithms for transforming one representation into the other, namely the Fourier–Motzkin procedure (or quantifier elimination) for computing the system of inequalities from the generators [55, 63], and Chernikova’s algorithm for computing the generators from a set of predicates [44].

¹ $x^\top y = \sum_{i=1}^n x_i y_i$ is the scalar product of n -dimensional vectors $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$.

30.2.2 Hybrid Automata

We define hybrid automata with polynomial dynamics [6, 96].

Definition 6 (Hybrid automaton with polynomial dynamics) *A hybrid automaton H with polynomial dynamics is a tuple*

$$\langle \text{Loc}, \text{Lab}, \text{Edg}, X, \text{Init}, \text{Inv}, \text{Flow}, \text{Jump}, \text{Final} \rangle$$

where:

- $\text{Loc} = \{\ell_1, \dots, \ell_m\}$ is a finite set of *locations*;
- Lab is a finite set of *labels*, including the *silent* label τ ;
- $\text{Edg} \subseteq \text{Loc} \times \text{Lab} \times \text{Loc}$ is a finite set of *edges*;
- $X = \{x_1, \dots, x_n\}$ is a finite set of *variables*;
- $\text{Init} : \text{Loc} \rightarrow \text{PConstr}(X)$ gives the *initial condition* $\text{Init}(\ell)$ of location ℓ . The automaton can start in ℓ with an initial valuation v lying in $\llbracket \text{Init}(\ell) \rrbracket$;
- $\text{Inv} : \text{Loc} \rightarrow \text{PConstr}(X)$ gives the *evolution domain restriction* $\text{Inv}(\ell)$ (also called the invariant) of location ℓ . The automaton can stay in ℓ as long as the values of its variables lie in $\llbracket \text{Inv}(\ell) \rrbracket$;
- $\text{Flow} : \text{Loc} \rightarrow \text{PConstr}(X \cup \dot{X})$ is the *flow constraint*, which constrains the evolution of the variables in each location. In a location ℓ , if the valuation of the variables is v_0 at time $t = 0$, then at time $t \geq 0$, the value of the variables is $\phi(t)$ where $\phi : \mathbb{R} \rightarrow \mathbb{R}^X$ is such that the flow relation $\text{Flow}(\ell)(\phi(t), \dot{\phi}(t))$ holds for the flow $\phi(t)$ and its time-derivative $\dot{\phi}(t)$, and $\phi(0) = v_0$.²
- $\text{Jump} : \text{Edg} \rightarrow \text{PConstr}(X \cup X^+)$ with $X^+ = \{x_1^+, \dots, x_n^+\}$ gives the *jump condition* $\text{Jump}(e)$ of edge e . The variables in X^+ refer to the updated values of the variables after the edge has been traversed. Jump conditions are often conjunctions of a guard and a reset constraint. There, the constraints purely on variables in X are called *guards*, and the constraints that describe variables in X^+ in terms of variables in X are called *updates* or *resets*.
- $\text{Final} : \text{Loc} \rightarrow \text{PConstr}(X)$ gives the *final condition* $\text{Final}(\ell)$ of location ℓ . Depending on the analysis question at hand, final conditions can either specify the unsafe states of the system or the desired states of the system.

The labels on edges can be used to synchronize hybrid automata in a compositional design. In the rest of this chapter, we assume that a single automaton is to be analyzed.

Example Figure 1 represents an affine automaton modeling a single gas-burner that is shared for heating alternatively two water tanks. It has three locations ℓ_0, ℓ_1, ℓ_2 and two variables x_1 and x_2 , the temperature in the two tanks. The gas-burner can

²Note that the semantics of flow constraints requires some attention, see differential-algebraic constraints [136].

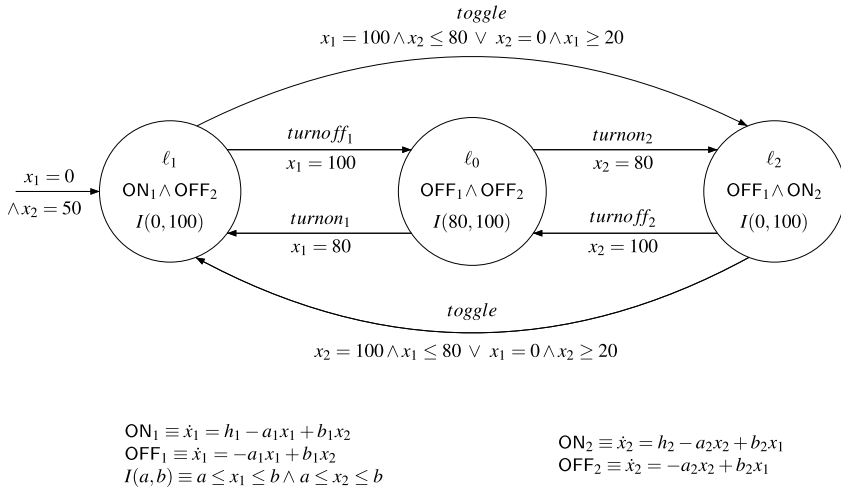


Fig. 1 A shared gas-burner

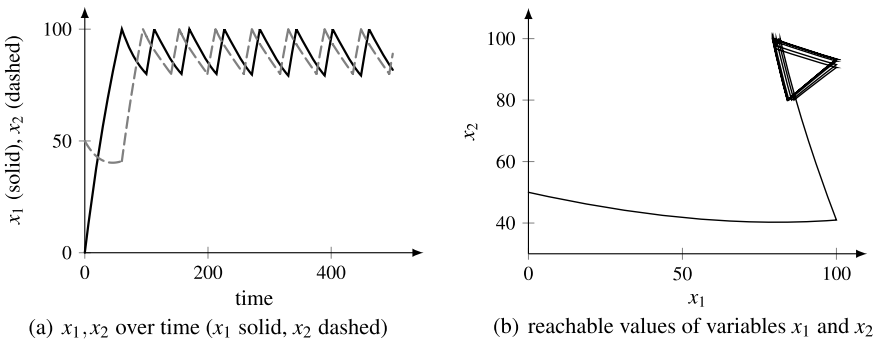


Fig. 2 The evolution of the continuous variables of the shared gas burner example, starting from location ℓ_1 with initial values $x_1 = 0$ and $x_2 = 50$

be either switched off (in ℓ_0) or turned on heating one of the two tanks (in ℓ_1 or ℓ_2). The dynamics in each location is given by a combination of the predicates ON_i and OFF_i ($i = 1, 2$) where the constants a_i model the heat exchange rate of the tank i with the room in which the tanks are located, b_i model the heat exchange rate between the two tanks and h_i depends on the power of the gas-burner. On every edge of the automaton, we have omitted the condition $x_1^+ = x_1 \wedge x_2^+ = x_2$ also written as $stable(x_1, x_2)$ that asks that the values of the variables are maintained when the edge is traversed. In the sequel, we fix the constants $h_1 = h_2 = 2$, $a_1 = a_2 = 0.01$ and $b_1 = b_2 = 0.005$. The evolution of the continuous variables over time is shown in Fig. 2. Starting in location ℓ_1 , the burner heats up tank 1 until it reaches a temperature of 100 degrees. Since the temperature of tank 2 is below 80 degrees, the automaton takes edge $toggle$ to location ℓ_2 . Note that edge $turnoff_1$ cannot be

taken since the evolution $I(80, 100)$ domain of the target location is not satisfied by x_2 . In location ℓ_2 , the burner heats up tank 2 until it reaches 100 degrees. Since x_1 is still above 80 degrees, the automaton takes edge $turnoff_2$ to location ℓ_0 , where the burner is off. It briefly remains here until x_1 falls to 80 degrees, at which point it takes edge $turnon_1$ to location ℓ_1 , where the burner heats up tank 1. The automaton converges towards a limit cycle of heating tank 1, heating tank 2, and briefly turning off the burner.

The definitions above define what a hybrid automaton consists of (flows, jumps, initial regions, ...) but they do not specify the behavior of a hybrid automaton or how its state evolves over time. This is the purpose of defining a semantics for hybrid automata by providing a transition system for each hybrid automaton.

Definition 7 (Semantics of hybrid automata) The *semantics* of a hybrid automaton $H = \langle \text{Loc}, \text{Lab}, \text{Edg}, X, \text{Init}, \text{Inv}, \text{Flow}, \text{Jump}, \text{Final} \rangle$ is the transition system $\llbracket H \rrbracket = \langle S, S_0, S_f, \Sigma, \rightarrow \rangle$ where $S = \{(\ell, v) \in \text{Loc} \times \mathbb{R}^X \mid v \in \llbracket \text{Inv}(\ell) \rrbracket\}$ is the *state space*, $S_0 = \{(\ell, v) \in S \mid v \in \llbracket \text{Init}(\ell) \rrbracket\}$ is the *initial space*, $S_f = \{(\ell, v) \in S \mid v \in \llbracket \text{Final}(\ell) \rrbracket\}$ is the *final space*, the actions are $\Sigma = \text{Lab} \cup \{\text{time}\}$ (we assume that $\text{time} \notin \text{Lab}$) and the transition relation \rightarrow contains all the tuples $((\ell, v), \sigma, (k, w))$ such that:

- (discrete transition) either there exists $e = (\ell, \sigma, k) \in \text{Edg}$ such that $(v, w) \in \llbracket \text{Jump}(e) \rrbracket$, or
- (continuous transition) $\ell = k$, $\sigma = \text{time}$ and there exists an $r \in \mathbb{R}^{\geq 0}$ and a continuously differentiable function $\xi : [0, r] \rightarrow \mathbb{R}^X$ such that $\xi(0) = v$, $\xi(r) = w$ and $(\xi(t), \dot{\xi}(t)) \in \llbracket \text{Flow}(\ell) \rrbracket$ for all $t \in [0, r]$ and $\xi(t) \in \llbracket \text{Inv}(\ell) \rrbracket$ for all $t \in [0, r]$.

We call ξ a *trajectory* from v to w . We also write $(\ell, v) \xrightarrow{r} (k, w)$ to emphasize that the continuous transition is of duration r . Usually $\text{Flow}(\ell)$ is a differential equation, in which case ξ is a solution of that differential equation.

We write $(\ell, v) \xrightarrow{\sigma} (k, w)$ if \rightarrow contains the tuple $((\ell, v), \sigma, (k, w))$.

A state $q = (\ell, v) \in S$ is *reachable* if there exists a finite path $q_0 \sigma_0 q_1 \sigma_1 \dots \sigma_{n-1} q_n$ where $q_0 \in S_0$, $q = q_n$, and $(q_i, \sigma_i, q_{i+1}) \in \rightarrow$ for all $0 \leq i < n$. This path *generates* the word $\bar{\sigma} = \sigma_0 \sigma_1 \dots \sigma_{n-1} \in \Sigma^*$. If $q \in S_f$ is final, we say that the word $\bar{\sigma}$ is *accepted* by H . The set of words that are accepted by H is the *language* of H , denoted $\mathcal{L}(H)$. The set of reachable states of $\llbracket H \rrbracket$ is denoted by $\text{Reach}(\llbracket H \rrbracket)$. The transition system $\llbracket H \rrbracket$ is *safe* if $\text{Reach}(\llbracket H \rrbracket) \cap S_f = \emptyset$.

Safety Verification Problem. Many verification problems for hybrid systems reduce to the *safety problem* for hybrid automata.

Definition 8 (Safety verification problem for hybrid automata) Given a hybrid automaton H , the *safety verification problem for hybrid automata* asks whether $\llbracket H \rrbracket$ is safe.

A *parameter* in a hybrid automaton is a variable which has first derivative 0 in every location and is never modified by discrete transitions. The *parametric safety*

verification problem for hybrid automata asks, given a hybrid automaton H and a parameter p in H , whether there exists a value $v_p \in \mathbb{R}$ such that $\llbracket H_{p=v_p} \rrbracket$ is safe, where $H_{p=v_p}$ is obtained by replacing every constraint φ in H by $\varphi \wedge (p = v_p)$.

Remark There would be no loss of generality in assuming that there is a location ℓ_{bad} such that $\text{Final}(\ell_{\text{bad}}) = \text{true}$ and $\text{Final}(\ell) = \text{false}$ for all $\ell \neq \ell_{\text{bad}}$. Indeed, to reduce any hybrid automaton to this form, it suffices to add transitions $e_\ell = (\ell, \sigma, \ell_{\text{bad}})$ with $\text{Jump}(e_\ell) = \text{Final}(\ell)$ for each $\ell \in \text{Loc}$, and to substitute $\text{Final}(\ell)$ with false for each $\ell \neq \ell_{\text{bad}}$.

30.3 Decidability and Undecidability Results

We review the most important results about the decidability of the safety verification problem for subclasses of hybrid automata. Details and proofs can be found in the given references.

Safety Verification Problem. The safety verification problem is decidable only for restricted classes of hybrid automata. The main classes for which safety verification is decidable are timed automata (see Chap. 29), initialized rectangular automata, and o-minimal hybrid automata [108]. The safety verification problem is undecidable already for the class of rectangular hybrid automata (and therefore also for linear, and affine hybrid automata).

A rectangular predicate over X is an expression of the form $a < x < b$ where $x \in X$, $< \in \{\leq, <\}$, and $a \leq b$ define a nonempty (possibly unbounded) interval with endpoints $a, b \in \mathbb{Q} \cup \{-\infty, \infty\}$. *Rectangular hybrid automata* are hybrid automata where (i) the flow constraint in each location ℓ is a conjunction of rectangular predicates over \dot{X} , (ii) the initial, final, and evolution domain conditions are conjunctions of rectangular predicates over X , and (iii) the jump condition of every edge is a conjunction of rectangular predicates over X^+ and expressions of the form $x^+ = x$ for $x \in X$. A hybrid automaton is *initialized* if for every edge $e = (\ell, \sigma, k)$ and for every variable x such that $\{v(\dot{x}) \mid v \in \llbracket \text{Flow}(\ell) \rrbracket\} \neq \{v(\dot{x}) \mid v \in \llbracket \text{Flow}(k) \rrbracket\}$, it holds that the set $\text{update}_e^x(v) = \{w(x^+) \mid (v, w) \in \llbracket \text{Jump}(e) \rrbracket\}$ does not depend on the valuation v (i.e., $\text{update}_e^x(v) = \text{update}_e^x(v')$ for all valuations v, v'). In words, whenever the flow condition is changed for a variable x by a discrete transition e , then this variable is (nondeterministically) reinitialized to a new value in update_e^x that is independent of the previous value.

The following decidability result is obtained by a translation of initialized rectangular hybrid automata to timed automata, preserving safety (see also Sect. 30.5.1).

Theorem 1 ([96]) *The safety verification problem is decidable for initialized rectangular hybrid automata (and therefore also for timed automata).*

The safety verification problem remains decidable for various extensions of timed automata. For instance, if *diagonal constraints* of the form $x - y \bowtie c$ for

$x, y \in X$, $\bowtie \in \{<, \leq, =, >, \geq\}$, and $c \in \mathbb{Q}$ are allowed in guards, or if *assignments* of the form $x^+ = y$ are allowed in updates, then the safety verification problem is still decidable [11, 33]. The decidability result for safety verification, useful for model checking, can be extended to the controller synthesis problem, solved as a game. We refer to Chap. 29 for games on timed automata, and mention the decidability of discrete-time control for rectangular hybrid automata [95].

The safety verification problem becomes undecidable for automata with rectangular flow constraints.

Theorem 2 ([96]) *The safety verification problem is undecidable for rectangular hybrid automata (and therefore also for linear, and affine hybrid automata).*

Note that the class of initialized rectangular hybrid automata (for which safety verification is decidable) have a finite language-equivalence quotient [94, 156]. The special case of initialized rectangular hybrid automata with only two variables even has a finite similarity quotient [94], and the class of timed automata has finite bisimilarity quotient (see also Chap. 29).

The result of Theorem 2 has been refined in several directions [96]. The problem is undecidable even if there is a single variable x with two different slopes, i.e., there exist $k_1, k_2 \in \mathbb{Q}$ with $k_1 \neq k_2$ such that in every location ℓ , either $\text{Flow}(\ell)$ implies $\dot{x} = k_1$, or $\text{Flow}(\ell)$ implies $\dot{x} = k_2$. The undecidability result holds for all fixed rational constants $k_1 \neq k_2$. The problem is also undecidable if diagonal constraints or assignments of the form $x^+ = y$ are allowed, and one variable has slope $k \neq 1$. There are extremely simple classes of hybrid systems, stopwatch automata, i.e., timed automata with only differential equations of the form $\dot{x} = 1$ and $\dot{x} = 0$, that are already undecidable [40] (see also Chap. 29). The variant of time-bounded safety verification asks whether, given a time bound T , there exists a final state reachable within a total duration of T time units. This problem is also undecidable for general rectangular hybrid automata, but it is decidable for a larger class than plain safety verification, namely for rectangular hybrid automata with monotone dynamics (the rate of every variable is either always non-negative, or always non-positive [36]).

Note that while differential equations define single continuous executions, the safety verification problem has been considered under various perturbed semantics with finite precision where drifting executions or tubes of executions are considered. It turns out that the undecidability result of Theorem 2 is mostly robust [97], but some decidability results can be obtained [23, 57, 64, 155].

Systems between timed and hybrid automata may remain decidable, e.g., weighted timed automata [15, 25]. Even systems with piecewise constant derivatives quickly become undecidable for dimension three [21]. On the other hand, if the discrete and the continuous parts of a hybrid system are completely independent of each other, the system falls apart into separate continuous systems, so that reachability becomes decidable for certain classes of linear differential equations [108].

Parametric Safety Verification Problem. If parameters are allowed only in the jump conditions of the edges, then it can be shown that the parametric safety veri-

fication problem is decidable for timed automata with one clock [14, 124], and undecidable for timed automata with one parameter and (i) three clocks (all of which being possibly constrained by the parameter) [124], or (ii) four clocks, but only one is compared with the parameter [124]. These undecidability results require the use of equalities in jump conditions. An undecidability result is known for *open* timed automata (in which all guards are open sets, thus forbidding equality constraints) with two parameters and five clocks (among which two are compared with the parameters) [59]. If parameters are allowed in the flow constraints, then it can be shown that the parametric safety verification problem is undecidable for rectangular automata with three variables and one parameter [181].

Computability and Polynomial Constraints. A frequent misconception about the definition of hybrid automata is that they should allow an arbitrary subset $\text{Init}(\ell) \subseteq \mathbb{R}^n$ of the real numbers as initial region for each location ℓ , an arbitrary subset $\text{Inv}(\ell) \subseteq \mathbb{R}^n$ as evolution domain restriction, arbitrary relation $\text{Flow}(\ell) \subseteq \mathbb{R}^n \times \mathbb{R}^n$ as flow constraints, arbitrary relation $\text{Jump}(e) \subseteq \mathbb{R}^n \times \mathbb{R}^n$ jump conditions, and an arbitrary subset $\text{Final}(\ell) \subseteq \mathbb{R}^n$ as final conditions. Generalizations like these have been suggested in the literature numerous times. They are useful as mathematical models, but not for any computational or verification purpose. It is important to understand why.

We can only obtain meaningful model-checking results for a hybrid automaton if we can describe the hybrid automaton (e.g., as an input file in a computer for the model checker). There is no way to describe arbitrary sets $\text{Init}(\ell)$, $\text{Inv}(\ell) \subseteq \mathbb{R}^n$, $\text{Jump}(e) \subseteq \mathbb{R}^n \times \mathbb{R}^n$, etc. as inputs, because there are uncountably many such sets, but model checkers accept only finite input files from a countable set of inputs.

Moreover, even for cases where there is some description of those sets, we still need to equip the model checker with a way to decide membership in those sets. Suppose some model-checking algorithm worked hard to find out that the hybrid automaton will be unsafe when started in a particular state $v \in \mathbb{R}^n$. Then, the model checker still needs to find out whether the hybrid automaton allows v as an initial state or not. That is, we need to give the model checker a way of deciding whether $v \in \text{Init}(\ell)$ for any location $\ell \in \text{Loc}$. Mathematically, this is a simple set inclusion and looks trivial. But that does not mean there is a computer program that can decide whether $v \in \text{Init}(\ell)$ or $v \notin \text{Init}(\ell)$. For arbitrary sets $\text{Init}(\ell) \subseteq \mathbb{R}^n$, this is impossible by classical results on the limits of computation due to Turing, Church, Gödel, and others. The Mandelbrot set is an example of such a set $\text{Init}(\ell)$ for which it is impossible to decide membership even in a very strong model of real computation [32].

Similar observations hold for all the other parts of hybrid automata. Consequently, we have to assume more structure on $\text{Init}(\ell)$ and all the parts of the definitions of hybrid automata. This is the reason why it is crucial that Definition 6 requires hybrid automata to be described in a definable way. Definition 6 requires hybrid automata to be described by polynomial constraints with rational coefficients, which are representable on a computer, unlike constraints with arbitrary real coefficients. This also explains why it is critical to restrict polyhedra to rational coefficients in Definition 5.

It should be noted that these observations about the requirements on hybrid automata are crucial for all model checkers, whether they try to decide fragments or semidecide fragments or whether they just strive to approximately answer the reachability problem. Fundamental limits of computation that represent the *numerical analogue of the halting problem* otherwise cause strong undecidabilities even for approximate answers [147], unless additional assumptions are imposed on the hybrid automata [51, 147].

30.4 Set-Based Reachability Analysis

There are two kinds of events that can take place in a hybrid automaton: time can pass with the state evolving according to the flow constraints, or a jump can take the system instantaneously to a new state. Starting from the initial states, *set-based reachability analysis* exhaustively computes the successor states for both time elapse and jumps in alternation until this no longer produces any new states. Since this process might not terminate (see decidability results in Sect. 30.3), an a priori limit on the search depth is sometimes imposed. The search depth is usually counted in the number of jumps and, in analogy to discrete automata, this is referred to as *bounded model checking*.

Reachability computation can be seen as a generalization of *numerical simulation*. In numerical simulation, one picks an initial state and tries to compute a successor state that lies on one of the solutions of the corresponding flow constraint and also satisfies one of the jump conditions (some intermediate points along the trajectory are usually kept as well). Then one picks one of the successor states of the jump and repeats the process. Like numerical simulation, reachability analysis directly follows the transition semantics of hybrid automata (Definition 7), but considers sets of states instead of single states.

Just like numerical simulation, reachability computation has to use approximations if the dynamics of the system are complex. Working with sets instead of points, approximate reachability can be conservative in the sense that the computed sets are sure to cover all solutions. Computation costs generally increase sharply in terms of the number of continuous variables. Scalable approximations are available for certain types of dynamics, as discussed later in this section, but this performance comes at a price in accuracy. The trade-off between runtime and accuracy remains a central problem in reachability analysis. Surveys of reachability techniques for hybrid automata can be found, e.g., in [7, 117, 119, 165].

30.4.1 Reachability Algorithm

The standard method to compute the reachable states is to iterate the following *one-step successor* operators for discrete and continuous transitions. Given a set

of states S , let $\text{Post}_C(S)$ be the set of states reachable by letting time elapse from any of the states in S ,

$$\text{Post}_C(S) = \{(\ell, w) \mid \exists(\ell, v) \in S : (\ell, v) \xrightarrow{\text{time}} (\ell, w)\}.$$

Let $\text{Post}_D(S)$ be the set of states resulting from taking a discrete transition from any of the states in S ,

$$\text{Post}_D(S) = \{(k, w) \mid \exists(\ell, v) \in S, \exists\sigma \in \text{Lab} : (\ell, v) \xrightarrow{\sigma} (k, w)\}.$$

The reachable states are obtained by applying $\text{Post}_C(S)$ and $\text{Post}_D(S)$ repeatedly and recording all states that are obtained. The basic algorithm for *forward reachability* computes the following sequence, starting from the initial states:

$$\begin{aligned} R_0 &= \{(\ell, v) \mid v \in \llbracket \text{Init}(\ell) \rrbracket\}, \\ R_{i+1} &= R_i \cup \text{Post}_C(R_i) \cup \text{Post}_D(R_i) \quad \text{for } i = 0, 1, 2, \dots \end{aligned}$$

The algorithm terminates when a fixed point is reached, i.e., when $R_{i+1} = R_i$ for some $i \geq 0$ (note that $R_i \subseteq R_{i+1}$ for all $i \geq 0$). This simple algorithm does not necessarily terminate, even for systems where reachability is decidable. For example, a system with an (unbounded) counter would enter a new state at each iteration such that the fixed point is never reached. Abstraction techniques such as *widening* [22, 86] are used in program analysis to ensure termination, and while they have been applied to hybrid systems with simple dynamics [92] it is difficult to obtain finite-state abstractions for more general cases.

Reachability with Symbolic States. A semi-algorithm used frequently for reachability of hybrid automata is shown as Algorithm 1. The states of the hybrid automaton H are represented by finite sets of *symbolic states* (ℓ, P) , where $\ell \in \text{Loc}$ and P is a set of continuous states in a suitable set representation such as polyhedra. The set of states corresponding to such a set $R = \{(\ell_1, P_1), (\ell_2, P_2), \dots\}$ is

$$\llbracket R \rrbracket = \{(\ell, v) \mid \exists(\ell, P) \in R : v \in P\}.$$

If H is safe, Algorithm 1 computes the reachable states by iterating one-step successor computations on such a set R , without guarantee of termination. If H is not safe, the procedure will eventually stop when a nonempty intersection of R with the final states is found. A similar semi-algorithm implements the backward approach by iterating a one-step predecessor operator. Other approaches are possible such as mixed forward-backward, where the forward and backward algorithms are executed in an interleaved fashion [92]. All these variations are semi-algorithms since the problem is undecidable.

The one-step successors $\text{Post}_C(S)$ and $\text{Post}_D(S)$ are implemented for symbolic states by enumerating over locations and transitions, respectively, using the following operators. The *continuous successors* of a set of continuous states P in a location

Algorithm 1: A reachability semi-algorithm using symbolic states**Input** : A hybrid automaton $H = \langle \text{Loc}, \text{Lab}, \text{Edg}, X, \text{Init}, \text{Inv}, \text{Flow}, \text{Jump}, \text{Final} \rangle$.**Output:** If H is safe then SAFE else UNSAFE.**begin** $Bad \leftarrow \{(\ell, \llbracket \text{Final}(\ell) \rrbracket) \mid \ell \in \text{Loc}\};$ $R \leftarrow \{(\ell, \text{post}_\ell(\llbracket \text{Init}(\ell) \rrbracket)) \mid \ell \in \text{Loc}\};$ $R_{old} \leftarrow \emptyset;$ **while** $\llbracket R \rrbracket \not\subseteq \llbracket R_{old} \rrbracket$ **do** $R_{old} \leftarrow R;$ $R \leftarrow \{(\ell, \text{post}_\ell(\text{post}_\varepsilon(P))) \mid (\ell, P) \in R \wedge \varepsilon = (\ell, \sigma, k) \in \text{Edg}\};$ **if** $\llbracket R \rrbracket \cap \llbracket Bad \rrbracket \neq \emptyset$ **then return** UNSAFE;**return** SAFE ;**end** ℓ is the set of continuous states

$$\text{post}_\ell(P) = \{x' \mid \exists x \in P : (\ell, x) \xrightarrow{\text{time}} (\ell, x')\}.$$

Similarly, the *discrete successors* of a set of continuous states P for an edge $\varepsilon = (\ell, \sigma, k)$ is the set of continuous states

$$\text{post}_\varepsilon(P) = \{x' \mid \exists x \in P : (\ell, x) \xrightarrow{\sigma} (k, x')\}.$$

Formally, the one-step successors of a set of symbolic states R are expressed using the above operators as

$$\text{Post}_C(\llbracket R \rrbracket) = \llbracket \{(\ell, \text{post}_\ell(P)) \mid \exists (\ell, P) \in R\} \rrbracket,$$

$$\text{Post}_D(\llbracket R \rrbracket) = \llbracket \{(k, \text{post}_\varepsilon(P)) \mid \exists (\ell, P) \in R, \varepsilon = (\ell, \sigma, k) \in \text{Edg}\} \rrbracket.$$

In the following, we discuss the above successor operators $\text{post}_\ell(P)$, $\text{post}_\varepsilon(P)$ for different classes of hybrid automata with increasingly complex continuous dynamics. We will focus mainly on computing time elapse successors, since this operation usually dominates costs. Other operations of the reachability algorithm may also become bottlenecks, e.g., computing the discrete successors, containment checking, and clustering.

30.4.2 Piecewise Constant Dynamics

Hybrid automata with piecewise constant dynamics (PCDA) are a special case of hybrid automata with polynomial dynamics (Definition 6), where all constraints are conjunctivelinear and the flow constraints are linear predicates over dotted variables

only. That is, the derivatives of the variables are independent of the current continuous state. They are also called *linear hybrid automata* (LHA), where the term linear refers to trajectories instead of dynamics (they do not allow the linear dynamics discussed in the next section). In order to avoid possible confusion resulting from this terminology, we prefer the name PCDA.

Definition 9 (Hybrid automaton with piecewise constant dynamics) A hybrid automaton $H = \langle \text{Loc}, \text{Lab}, \text{Edg}, X, \text{Init}, \text{Inv}, \text{Flow}, \text{Jump}, \text{Final} \rangle$ with polynomial dynamics is called a *hybrid automaton with piecewise constant dynamics* iff:

- Init, Inv, Final are conjunctivelinear constraints over X ,
- Flow are conjunctivelinear constraints over \dot{X} , and
- Jump are conjunctivelinear constraints over $X \cup X^+$.

PCDA are of particular interest in formal verification because the one-step successors can be computed exactly, which is not the case for the more complex dynamics discussed in later sections.

Examples of flow constraints of a PCDA include differential inclusions such as $\dot{x} \in [1, 2]$, and conservation laws such as $\dot{x} + \dot{y} = 0$. The jump constraints of a PCDA admit arbitrary linear updates of the variables, which can generate complex behavior. For example, PCDA can model discrete-time affine systems, a widely used class of control systems, by using jump constraints of the form $x^+ = Ax + b$. Chaotic behavior can arise in PCDA due to switching flows [42] or guarded jumps, with which one can model piecewise affine maps such as the tent map [48].

Continuous Successors. In the following, we discuss computing the states reachable by time elapse in a given location ℓ of a PCDA and write x as shorthand for the state (ℓ, x) . By definition, a trajectory can be an arbitrarily curved function as long as it is differentiable and satisfies both flow constraints and evolution domain restrictions. For the purposes of reachability, it suffices to consider only straight-line trajectories of PCDA, as formalized in the following lemma.

Lemma 1 ([13]) *In any given location of a PCDA, there is a trajectory $\xi(t)$ from $x = \xi(0)$ to $x' = \xi(r)$ for some $r > 0$ iff $\eta(t) = x + qt$ with $q = \frac{x' - x}{r}$ is a trajectory from x to x' .*

Using this lemma, we now show that the states reachable by time elapse from a polyhedral set of states P are given by the union of P with a polyhedron that is readily computable [8, 28]. Consider polyhedra P and Q . The states on straight-line trajectories starting in P with constant derivative $\dot{x} = q$ for any $q \in Q$ are the *time successors*

$$P \nearrow Q = \{x' \mid x \in P, q \in Q, t \in \mathbb{R}^{\geq 0}, x' = x + qt\}. \quad (1)$$

We now transform the right-hand term of (1) into a linear constraint. Let P and Q be polyhedra given in vector-matrix form as $P = \{x \mid Ax \bowtie b\}$, $Q = \{q \mid \bar{A}q \bar{\bowtie} \bar{b}\}$.



(a) $P \nearrow Q$ with closed and unbounded $Q = \{q_1 = 1 \wedge q_2 \leq 1\}$ (b) $P \nearrow Q'$ with bounded and non-closed $Q' = \{q_1 = 1 \wedge -1 \leq q_2 < 1\}$

Fig. 3 Given a polyhedron P and a polyhedral set of derivatives Q , the time successors $P \nearrow Q$ can be a convex set that is not a single polyhedron but the union of P with another polyhedron

By separating the case $t = 0$ from $t > 0$ in (1) we have $q = \frac{x' - x}{t}$. Eliminating q and multiplying with t yields

$$P \nearrow Q = P \cup \{x' \mid Ax \bowtie b \wedge \bar{A}(x' - x) \bar{\bowtie} \bar{b} \cdot t \wedge t > 0\}. \tag{2}$$

The right-hand term of the union in (2) is a polyhedron that can be computed by quantifier elimination over $X \cup \{t\}$ using, e.g., Fourier–Motzkin elimination. If Q is closed and bounded, the constraint $t > 0$ in (2) can be replaced by $t \geq 0$, so the right-hand term contains P and $P \nearrow Q$ becomes a single polyhedron. The following example illustrates that $P \nearrow Q$ can be the union of two polyhedra.

Example 1 For $P = \{x_1 = 0 \wedge x_2 = 0\}$, and Q, Q' given in Fig. 3, (2) yields

$$\begin{aligned} P \nearrow Q &= P \cup \{(x'_1, x'_2) \mid x_1 = 0 \wedge x_2 = 0 \wedge x'_1 - x_1 = t \wedge x'_2 - x_2 \leq t \wedge t > 0\} \\ &= P \cup \{(x'_1, x'_2) \mid x'_1 = t \wedge x'_2 \leq t \wedge t > 0\} \\ &= P \cup \{(x'_1, x'_2) \mid x'_1 > 0 \wedge x'_2 \leq x'_1\}. \end{aligned}$$

Here, the closed but unbounded set Q results in a convex set $P \nearrow Q$ that is not a polyhedron but the union of two polyhedra. Similarly, the bounded but non-closed set Q' results in $P \nearrow Q' = P \cup \{x'_1 > 0 \wedge -x'_1 \leq x'_2 < x'_1\}$, which is also convex and not a polyhedron.


The time successor operation can also be carried out using geometrical operations on the polyhedra P and Q as shown in Fig. 4 [86]. The *positive cone* of Q is the polyhedral set $\text{pos}(Q) = \{q \cdot t \mid q \in Q, t > 0\}$. The time successors are given by the Minkowski sum³ of P and the positive cone of Q ,

$$P \nearrow Q = P \cup (P \oplus \text{pos}(Q)). \tag{3}$$

If P and Q are closed with generator representation (V, R) and (V', R') , respectively, then a generator representation of $P \nearrow Q$ is $(V, R \cup V' \cup R')$.

³The *Minkowski sum* is defined as $P \oplus Q = \{p + q \mid p \in P, q \in Q\}$.

Fig. 4 The time successors $P \nearrow Q$ using geometric operations on polyhedra P and Q



It remains to ensure that the time successors are reachable by trajectories that satisfy $\text{Inv}(\ell)$. Assuming that $P \subseteq \llbracket \text{Inv}(\ell) \rrbracket$, this restriction reduces to $x' \in \llbracket \text{Inv}(\ell) \rrbracket$ since $\llbracket \text{Inv}(\ell) \rrbracket$ is convex and only straight line trajectories need to be considered. This leads us to the following discrete successor operator for PCDA.

Lemma 2 ([8]) *The continuous successors of a polyhedron P in a location ℓ of a PCDA H is the set:*

$$\text{post}_\ell(P) = (P \nearrow \llbracket \text{Flow}(\ell) \rrbracket) \cap \llbracket \text{Inv}(\ell) \rrbracket.$$

Discrete Successors. The *discrete successors* of a polyhedron P for an edge $\varepsilon = (\ell, \sigma, k)$ of a PCDA H is the set:

$$\text{post}_\varepsilon(P) = \{x^+ \mid \exists x \in P : (x, x^+) \in \llbracket \text{Jump}(\varepsilon) \rrbracket \wedge x^+ \in \llbracket \text{Inv}(k) \rrbracket\}.$$

This set is defined using existential quantification, and computing it may require costly quantifier elimination. Frequently occurring special cases can be computed more efficiently. As an example, consider $\text{Jump}(e)$ given by a guard $x \in G$ and a reset $x^+ = Cx + d$, with a constant matrix C and a vector d of appropriate dimensions. The discrete successors are

$$\text{post}_\varepsilon(P) = (C(P \cap G) \oplus \{d\}) \cap \llbracket \text{Inv}(k) \rrbracket. \tag{4}$$

If C is invertible and all sets are polyhedra in constraint representation, the computation is straightforward since intersection corresponds to concatenation of constraints, and for any polyhedron $Q = \{x \mid Ax \bowtie b\}$,

$$CQ \oplus \{d\} = \{x \mid AC^{-1}x \bowtie b + C^{-1}d\}.$$

Computational Cost. Computing the continuous successors using (3) involves the cone, Minkowski sum, and intersection operations, for details see [22, 86]. The cone and Minkowski sum are efficient only in the generator representation of a polyhedron (see Definition 5). The intersection operation is efficient only in constraint representation. Translating the polyhedron from constraints to generators and vice versa can produce a number of generators that is exponential in the number of variables. For example, consider that an n -dimensional cube has $2n$ constraints and 2^n vertices. Dually, an n -dimensional *cross-polytope* (hyperoctahedron) has $2n$ vertices and 2^n constraints. In total, the cost of computing the continuous successors is exponential in the number of variables. Tools such as HyTech and PHAVer use the

geometric version (3) of the time successor operator since in practice it is often more efficient than quantifier elimination [87]. The operator is available in computational geometry libraries such as the Parma Polyhedra Library (PPL) [22].

The cost of computing the discrete successors is exponential for polyhedra in constraint representation since it involves quantifier elimination. For some frequently occurring special cases the cost is polynomial, e.g., in the case of (4) with invertible map.

The containment and emptiness tests in Algorithm 1 are carried out pairwise over the elements of sets of symbolic states. The *containment test* $P \subseteq Q$ is solvable with linear programming (and thus in polynomial time) if P, Q are in constraint representation.⁴ The *emptiness test* $P = \emptyset$ is solvable as a linear program if P is in constraint representation and trivial if P is in generator representation.

Path Constraints. A *path* of a hybrid automaton is a sequence of adjacent edges (usually from an initial to a final location). An interesting property of PCDA is that the reachable states along a given path can be encoded by a conjunction of linear constraints, the so-called *path constraints*. The reachability problem for a given path can therefore be solved very efficiently using linear programming. This approach has been implemented in the tool BACH [37]. The number of paths in a PCDA can be infinite if there are cycles, so techniques such as CEGAR have been used to reduce the number of paths to be checked and accelerate termination [162].

30.4.3 Piecewise Affine Dynamics

Hybrid automata with piecewise affine dynamics (PWA) are a special case of hybrid automata with polynomial dynamics (Definition 6), where all constraints are linear and the flow constraints are linear ordinary differential equations (ODEs). We divide the continuous variables into *state variables* $X = \{x_1, \dots, x_n\}$, whose derivative is explicitly defined, and *input variables* $U = \{u_1, \dots, u_m\}$, whose derivative is unconstrained. The input variables can be used to model nondeterminism such as open inputs to the system, approximation errors, disturbances, etc.

In each location of a PWA, the continuous dynamics are *affine*, i.e., given by differential equations of the form

$$\dot{x} = Ax + Bu, \quad u \in \mathcal{U}, \quad (5)$$

where A and B are matrices of appropriate dimension and the *input set* \mathcal{U} is compact and convex. Note that differential inclusions like $\dot{x} \in \mathcal{U}$ and $Ax - b \leq \dot{x} \leq Ax + b$ can be brought into this form by introducing auxiliary variables. Similarly, jump constraints of an edge e define resets of the form

$$x^+ = Cx + Du, \quad (6)$$

⁴Checking $P \subseteq Q$ is polynomial unless P is in constraint representation and Q is in generator representation, in which case it is known to be NP-complete [71].

where x^+ denotes the value of x after the jump, u is defined as above and C and D are matrices of appropriate dimension. The jump constraints also define a set \mathcal{G} called the *guard* of the edge, and a jump can only take place if $x \in \mathcal{G}$. The formal definition of PWA is as follows.

Definition 10 (Hybrid automaton with piecewise affine dynamics) *A hybrid automaton with piecewise affine dynamics is a hybrid automaton $H = \langle \text{Loc}, \text{Lab}, \text{Edg}, X \cup U, \text{Init}, \text{Inv}, \text{Flow}, \text{Jump}, \text{Final} \rangle$ where*

- *Init and Inv are conjunctivelinear constraints over X .*
- *Inv are conjunctivelinear constraints over $X \cup U$, such that each linear term ranges over variables exclusively from either X or U (no correlation between state and input variables). The input set \mathcal{U} of a location ℓ is given by the terms of $\text{Inv}(\ell)$ that range over U and must be closed and bounded.*
- *Flow are constraints over $\dot{X} \cup X \cup U$ of the form $\dot{x} = Ax + Bu$.*
- *Jump are conjunctivelinear constraints over $X^+ \cup X \cup U$ whose terms either range over X or are of the form $x^+ = Cx + Du$. The *guard* set \mathcal{G} of an edge e is given by the terms of $\text{Jump}(e)$ that range over X .*

The reachable states of a PWA can be computed using Algorithm 1 from Sect. 30.4.1, with suitable operators post_ℓ for continuous and post_e for discrete successors that will be presented in the following section.

30.4.3.1 Successor Computations


The successor computations for affine dynamics can be approximated by sequences of geometric set operations. We first present such a sequence for the continuous successors, then give the equation for the discrete successors. Different set representations can be used to implement these operations, and a selection are discussed in the subsequent Sect. 30.4.3.2.

Continuous Successors. In the following, we discuss how to compute the states reachable by time elapse in a given location ℓ . Since ℓ is clear from the context we call x a (continuous) state. We will initially ignore any evolution domain restriction on x and discuss it after the basic construction has been presented. The evolution of the input variables is described by an *input signal* $\zeta : \mathbb{R}^{\geq 0} \rightarrow \mathcal{U}$ that attributes to each point in time a value of the input u . The input signal does not need to be continuous. A trajectory $\xi(t)$ from a state x_0 is the solution of the differential equation (5) for initial condition $\xi(0) = x_0$ and a given input signal ζ . It has the form

$$\xi_{x_0, \zeta}(t) = e^{At} x_0 + \int_0^t e^{A(t-s)} B \zeta(s) ds. \quad (7)$$

It consists of the superposition of the solution of the *autonomous* system, obtained for $\zeta(t) = 0$, and the input integral obtained for $x_0 = 0$. In the following, this decomposition of (7) will be exploited to obtain efficient and accurate approximations.

Fig. 5 A sequence of sets $\Omega_0, \Omega_1, \dots$ that covers \mathcal{X}_t over a finite time horizon T . The choice of set representation for Ω_k (illustrated here by *ellipsoids*) has a substantial impact on accuracy and computational complexity



A state x' is reachable from some initial set of states \mathcal{X}_0 in time t if for some $x_0 \in \mathcal{X}_0$ and some ζ , $x' = \xi_{x_0, \zeta}(t)$. We now describe the reachable states as sets using (7). Let \mathcal{X}_t be the states reachable in time t from any state in \mathcal{X}_0 and let \mathcal{Y}_t be the states reachable from $\mathcal{X}_0 = \{0\}$, then (7) can be written as

$$\mathcal{X}_t = e^{At} \mathcal{X}_0 \oplus \mathcal{Y}_t. \tag{8}$$

The goal is to conservatively approximate the reachable states over some finite time horizon T , i.e., to compute a finite sequence of sets $\Omega_0, \Omega_1, \dots$ such that

$$\bigcup_{0 \leq t \leq T} \mathcal{X}_t \subseteq \Omega_0 \cup \Omega_1 \cup \dots \tag{9}$$

We present the construction of a sequence of Ω_k for a fixed *sampling time* $\delta > 0$ such that Ω_k covers \mathcal{X}_t for $t \in [k\delta, (k + 1)\delta]$, as illustrated in Fig. 5. The so-called *semi-group* property of reachability says that, starting from \mathcal{X}_s , for any $s \geq 0$, and then waiting r time units leads to the same states as starting from \mathcal{X}_0 and waiting $r + s$ time units. Applying this to (8), we obtain that for any $r, s \geq 0$,

$$\mathcal{X}_{r+s} = e^{Ar} \mathcal{X}_s \oplus \mathcal{Y}_r. \tag{10}$$

Substituting $r \leftarrow \delta, s \leftarrow k\delta$, we obtain a recursive time discretization in the form of

$$\mathcal{X}_{(k+1)\delta} = e^{A\delta} \mathcal{X}_{k\delta} \oplus \mathcal{Y}_\delta.$$

It follows that if we have initial approximations Ω_0 and Ψ_δ such that

$$\bigcup_{0 \leq t \leq \delta} \mathcal{X}_t \subseteq \Omega_0, \quad \mathcal{Y}_\delta \subseteq \Psi_\delta, \tag{11}$$


then the sequence

$$\Omega_{k+1} = e^{A\delta} \Omega_k \oplus \Psi_\delta \tag{12}$$

satisfies (9). Note that Ω_0 covers the reachable set over an interval of time $[0, \delta]$, while Ψ_δ covers the values of the input integral at a single time instant δ .

Computing Initial Approximations Ω_0 and Ψ_δ . The set Ω_0 needs to cover \mathcal{X}_t from $t = 0$ to $t = \delta$. A good starting point for such a cover is the convex hull of

Fig. 6 An approximation Ω_0 that covers \mathcal{X}_t for $t \in [0, \delta]$ can be obtained from the convex hull of \mathcal{X}_0 and \mathcal{X}_δ by enlarging it to compensate for the curvature of trajectories



\mathcal{X}_0 and \mathcal{X}_δ . One approach, shown in Fig. 6(a), is to compute the convex hull in constraint representation, and push the facets out far enough to be conservative [81]. The required values can be computed from a Taylor approximation of (7) [19], or by solving an optimization problem [45]. Note that the cost of computing the exact constraints of the convex hull can be exponential in the number of variables, which limits the scalability of this approach.

A scalable way to obtain Ω_0 is to bloat \mathcal{X}_0 and \mathcal{X}_δ enough to compensate for the curvature of trajectories [75], as illustrated in Fig. 6(b). We present the approach from [75], which uses uniform bloating and whose approximation error is asymptotically linear in the time step δ as $\delta \rightarrow 0$. This is asymptotically optimal for any approximation containing the convex hull of \mathcal{X}_0 and \mathcal{X}_δ [110]. The bloating can be made non-uniform in space and time to obtain a more precise approximation [68, 110]. The bloating factor is derived from a Taylor approximation of (7), whose remainder is bounded using norms. To formalize the above statements, we use the following notation. Let $\|\cdot\|$ be a vector norm and let $\|A\|$ be its induced matrix norm.⁵ Let $\mu(\mathcal{X}) = \max_{x \in \mathcal{X}} \|x\|$ and let \mathcal{B} be the unit ball of the norm, i.e., the largest set \mathcal{B} such that $\mu(\mathcal{B}) = 1$. For a scalar c , let $c\mathcal{X} = \{cx \mid x \in \mathcal{X}\}$. The approximation error is measured using the Hausdorff distance between sets \mathcal{X}, \mathcal{Y} ,

$$d_H(\mathcal{X}, \mathcal{Y}) = \max \left\{ \sup_{x \in \mathcal{X}} \inf_{y \in \mathcal{Y}} \|x - y\|, \sup_{y \in \mathcal{Y}} \inf_{x \in \mathcal{X}} \|x - y\| \right\}.$$

Lemma 3 ([75]) *Given a set of initial states \mathcal{X}_0 and affine dynamics (5), let*

$$\begin{aligned} \alpha_\delta &= \mu(\mathcal{X}_0) \cdot (e^{\|A\|\delta} - 1 - \|A\|\delta), \\ \beta_\delta &= \frac{1}{\|A\|} \mu(B\mathcal{U}) \cdot (e^{\|A\|\delta} - 1), \\ \Omega_0 &= \text{chull}(\mathcal{X}_0 \cup e^{A\delta} \mathcal{X}_0) \oplus (\alpha_\delta + \beta_\delta)\mathcal{B}, \\ \Psi_\delta &= \beta_\delta \mathcal{B}. \end{aligned}$$

⁵For example, the infinity norm $\|x\|_\infty = \max\{|x_1|, \dots, |x_n|\}$ induces the matrix norm $\|A\| = \max_{1 \leq i \leq n} \sum_{j=1}^m |a_{ij}|$, where A is of dimension $n \times m$. Its ball \mathcal{B}_∞ is a cube of side length 2.

Then $\bigcup_{0 \leq t \leq \delta} \mathcal{X}_t \subseteq \Omega_0$ and $\mathcal{Y}_\delta \subseteq \Psi_\delta$. Furthermore, if BU is a ball of the norm, i.e., $BU = \mu(BU)\mathcal{B}$, the approximation error is bounded by

$$d_H\left(\bigcup_{0 \leq t \leq \delta} \mathcal{X}_t, \Omega_0\right) \leq \delta e^{\|A\|\delta} (\mu(BU) + (\frac{1}{2} + \delta)\|A\|\mu(\mathcal{X}_0)),$$

$$d_H(\mathcal{Y}_\delta, \Psi_\delta) \leq \delta^2 \|A\| e^{\|A\|\delta} \mu(BU).$$

Propagating the initial approximation Ω_0 forward in time using (12) gives an approximation of \mathcal{X}_t over a bounded horizon. The following theorem gives a bound on the total approximation error.

Theorem 3 ([75]) *Given Ω_0 and Ψ_δ as defined in Lemma 3, let $\Omega_{k+1} = e^{A\delta} \Omega_k \oplus \Psi_\delta$ for $k = 1, \dots, N - 1$. Then $\bigcup_{0 \leq t \leq N\delta} \mathcal{X}_t \subseteq \bigcup_{0 \leq k \leq N-1} \Omega_k$. Furthermore, if BU is a ball of the norm, the approximation error is bounded by*

$$d_H\left(\bigcup_{0 \leq t \leq N\delta} \mathcal{X}_t, \bigcup_{0 \leq k \leq N-1} \Omega_k\right) \leq \delta e^{\|A\|N\delta} \left(2\mu(BU) + \left(\frac{1}{2} + \delta\right)\|A\|\mu(\mathcal{X}_0)\right).$$

Approximations and the Wrapping Effect. The sequence in (12) can be problematic to compute since the complexity of Ω_k may increase sharply with k . We illustrate this for the case where Ω_k is a polytope in generator representation, and a similar argument can be made for constraint representation. Let N_k be the number of vertices of Ω_k and let Ψ_δ have M vertices. Since Ω_{k+1} is the sum of $e^{A\delta} \Omega_k$ with Ψ_δ it can have $N_{k+1} = N_k \cdot M$ vertices. Resolving the recursion, we get the tight upper bound $N_k \leq N_0 \cdot M^k$. To avoid this increase in complexity, we approximate each Ω_k by a simplified set. Let Appr be an *approximation function* such that for any set P , $P \subseteq \text{Appr}(P)$. The sequence (12) then becomes

$$\hat{\Omega}_{k+1} = \text{Appr}(e^{A\delta} \hat{\Omega}_k \oplus \Psi_\delta). \tag{13}$$

For example, if Appr computes the interval hull (bounding box) and Ω_0 is a polytope, then all $\hat{\Omega}_k$ are polytopes with $2n$ facets. However, the recursive application of the approximation function can lead to an exponential increase in the approximation error. This phenomenon is known in numerical analysis as the *wrapping effect* [105] and is illustrated in Fig. 7.

For affine dynamics, the wrapping effect can be avoided by combining two techniques [77]. First, the approximation operator is chosen such that it distributes over Minkowski sum, i.e., $\text{Appr}(P \oplus Q) = \text{Appr}(P) \oplus \text{Appr}(Q)$. This is the case, e.g., for the interval hull (bounding box). Second, the alternation of the map $e^{A\delta}$ with the Minkowski sum in (12) is avoided by splitting it into two sequences

$$\hat{\Psi}_{k+1} = \text{Appr}(e^{A\delta} \Psi_\delta) \oplus \hat{\Psi}_k, \quad \text{with } \hat{\Psi}_0 = \{0\},$$

$$\hat{\Omega}_k = \text{Appr}(e^{A\delta} \Omega_0) \oplus \hat{\Psi}_k. \tag{14}$$




Fig. 7 The wrapping effect can lead to an exponential increase in the approximation error that can be avoided for affine dynamics. This example shows the exact solution $e^{A k \delta} \mathcal{X}_0$ (shaded) and an interval hull approximation (thick), with $e^{A \delta}$ performing a rotation of 45 degrees around the origin. The wrapping effect occurs if the approximation is applied to the map of the previous approximation (dashed). To illustrate the effect more clearly, \mathcal{X}_0 is used here instead of Ω_0

For sequence (14) it holds that $\hat{\Omega}_k = \text{Appr}(\Omega_k)$, which means the resulting approximation is free of the wrapping effect. The total approximation error consists of the bounds of Theorem 3 plus the error introduced by the operator Appr (measured in terms of the Hausdorff distance).


The approach is easily extended to variable time steps by adapting Ω_0 and Ψ_δ to the time step while computing the sequence [68].

Evolution Domain Restriction. So far we have neglected the evolution domain restriction (invariant) $\text{Inv}(\ell)$ of the location. Let $\mathcal{S} = \llbracket \text{Inv}(\ell) \rrbracket$. A simple but efficient heuristic tries to find, if it exists, the smallest K such that Ω_K lies completely outside \mathcal{S} . The search for such a K may be combined with finding a suitable time horizon T and a suitable time step δ (this search obviously might not terminate). Then one computes the sequence $\Omega_0, \dots, \Omega_K$ and obtains the sequence $\bar{\Omega}_k = \Omega_k \cap \mathcal{S}$ as an approximation of the continuous successors over the time horizon $T = K \delta$.

In cases where the above solution is overly conservative, one can improve the approximation using the following approach from [83]. Let \mathcal{S}_t be the states reachable from \mathcal{S} (neglecting the evolution domain restriction), and let $\xi(\tau)$ be a trajectory inside \mathcal{S} for all $0 \leq \tau \leq t$. Then the semi-group property implies that $\xi(\tau + s) \in \mathcal{S}_s$ for all $0 \leq s \leq t - \tau$, so that $\xi(t) \in \bigcap_{0 \leq \tau \leq t} \mathcal{S}_\tau$. We may therefore improve the approximation by intersecting Ω_k with an approximation of the states reachable from \mathcal{S} , which we obtain from the sequence in (14) with $\Omega_0 \leftarrow \mathcal{S}$. This leads to the following sequence $\bar{\Omega}_k$ that approximates the continuous successors, starting with $k = 0$ and $\Psi_0 = \{0\}$:

$$\begin{aligned} \Psi_{k+1} &= \text{Appr}(e^{A k \delta} \Psi_\delta) \oplus \Psi_k, \\ \bar{\Omega}_k &= (\text{Appr}(e^{A k \delta} \Omega_0) \oplus \Psi_k) \cap \bigcap_{0 \leq i \leq k} (\text{Appr}(e^{A i \delta} \mathcal{S}) \oplus \Psi_i). \end{aligned} \quad (15)$$

Fig. 8 A zonotope is a special form of centrally symmetric polytope, as illustrated here with generators v_1, v_2, v_3, v_4 , and center $c = 0$



Discrete Successors. Consider an edge $\varepsilon = (\ell, \sigma, k)$ of a PWA, whose jump constraints define the reset map

$$x^+ = Cx + Du$$

and the guard set \mathcal{G} , which only lets states jump where $x \in \mathcal{G}$. Recall that $u \in \mathcal{U}$, where \mathcal{U} is compact, convex, and given by constraints in $Inv(\ell)$. Let $S^+ = \llbracket Inv(k) \rrbracket$ be the evolution domain restriction of the target location. The discrete successors of a set P can be written using geometric operators as

$$post_\varepsilon(P) = (C(P \cap \mathcal{G}) \oplus DU) \cap S^+.$$

We now turn to representing the individual sets in the sequences Ψ_k and Ω_k , and which approximation operator Appr to use.

30.4.3.2 Set Representations

Several set representations have been proposed in the literature for computing the continuous successors under affine dynamics, using variations of the algorithm presented in the previous section. To be efficient, scalable implementations or approximations need to be available for the operators in the algorithm. Using the initial approximation from Lemma 3 and the recurrence equation (14), the operators are linear map, Minkowski sum, convex hull, and intersection. The following paragraphs summarize the results for a selection of prominent representations.


Ellipsoids. The first scalable reachability algorithms for affine dynamics were obtained for ellipsoids, see [107, 133] and references therein. An approximation of the reachable states using ellipsoids is shown in Fig. 5. A nondegenerate *ellipsoid* $\mathcal{E}(c, Q) \subseteq \mathbb{R}^n$ is represented by a center $c \in \mathbb{Q}^n$ and a positive definite⁶ matrix $Q \in \mathbb{Q}^{n \times n}$,

$$\mathcal{E}(c, Q) = \{x \mid (x - c)^\top Q^{-1}(x - c) \leq 1\}$$

(this can be generalized to degenerate ellipsoids). Deterministic affine transforms can be computed efficiently for ellipsoids. For a matrix $A \in \mathbb{Q}^{n \times n}$ and vector

⁶A matrix Q is positive definite iff it is symmetric and $x^\top Qx > 0$ for all $x \neq 0$.

Fig. 9 A reach set cover $\hat{\Omega}_0, \hat{\Omega}_1, \dots$, computed with zonotopes using the implementation in [3] (solid)



$$b \in \mathbb{Q}^n,$$

$$A\mathcal{E}(c, Q) + b = \mathcal{E}(Ac + b, AQA^T).$$

Ellipsoids are not closed under Minkowski sum, convex hull, or intersection. Using ellipsoids one therefore generally suffers from the wrapping effect unless BU is a singleton. Efficient approximations are available for Minkowski sum, convex hull, and special cases of intersection, but the computation of discrete successors can be problematic in terms of accuracy. For an implementation, see [106].

Zonotopes. Zonotopes are a compact representation for a special form of polytopes that have been used successfully for reachability analysis due to their computationally attractive features [3, 75]. A *zonotope* $P \subseteq \mathbb{R}^n$ is defined by a center $c \in \mathbb{Q}^n$ and a finite number of generators $v_1, \dots, v_k \in \mathbb{Q}^n$ that span the polytope as bounded linear combinations from the center:

$$P = \left\{ c + \sum_{i=1}^k \alpha_i v_i \mid \alpha_i \in [-1, 1] \right\}.$$

A common denotation for this zonotope is $P = (c, \langle v_1, \dots, v_k \rangle)$. A zonotope with k generators is an affine transformation of a k -dimensional unit hypercube. Zonotopes are central-symmetric convex polytopes, see Fig. 8 for an illustration. Affine transformations can be computed efficiently for zonotopes. For a matrix $A \in \mathbb{Q}^{m \times n}$, the image of the linear transformation can simply be computed component-wise:

$$AP = (Ac, \langle Av_1, \dots, Av_k \rangle)$$

The Minkowski sum can be computed efficiently for zonotopes $P = (c, \langle v_1, \dots, v_k \rangle)$ and $Q = (d, \langle w_1, \dots, w_m \rangle)$ by a single vector addition and a single list concatenation:

$$P \oplus Q = (c + d, \langle v_1, \dots, v_k, w_1, \dots, w_m \rangle).$$

Since zonotopes are closed under Minkowski sum, it is straightforward to devise an approximation operator Appr that distributes over Minkowski sum and use the wrapping-free sequence (14). When the list of generators of a zonotope becomes large, one can efficiently compute a smaller list that results in a cover of the original zonotope [75].




Fig. 10 Evaluating the support function in a set of directions gives a polyhedral outer approximation that can be computed very efficiently

Zonotopes are neither closed under convex hull, nor under intersection. Efficient approximations exist, and the accuracy of approximating the convex hull in the above reachability algorithm can be improved by taking smaller time steps. However, the lack of accuracy in intersections can make the computation of discrete successors with zonotopes problematic. In special cases it can be advantageous to use an approach called *continuization* to avoid the intersection operation, see [5]. Instead of intersecting a set of states with the guard set and then applying the dynamics of the successor location to the result, the states suspected to intersect with the guard set (by some approximative measure) are subjected to nondeterministic dynamics that overapproximate the dynamics both before and after the jump. The dynamics of the successor location are used once enough time steps have been carried out to be sure the set no longer intersects with the guard set.

Reachability with zonotopes is extremely scalable for affine dynamics [3, 77]. The approach has been extended to nonlinear differential algebraic equations [2].

Support Functions. A support function represents a closed, bounded, and convex set exactly, somewhat like a characteristic function. Support functions lead to very scalable algorithms since linear map, Minkowski sum, and convex hull correspond to simple operations on vectors and scalars [74, 83, 116].


The *support function* $\varrho_P : \mathbb{R}^n \rightarrow \mathbb{R}$ of a nonempty, closed, bounded, and convex set P is

$$\varrho_P(d) = \max\{d^T x \mid x \in P\}.$$

It attributes to every *direction* $d \in \mathbb{R}^n$ the position of the tangent halfspace in that direction, see Fig. 10(a). The values of the support function over a set of directions $D \subseteq \mathbb{R}^n$ define an *outer approximation*

$$[P]_D = \bigcap_{d \in D} \{d^T x \leq \varrho_P(d)\}.$$

If $D = \mathbb{R}^n$ or D is the ball of a norm, then $[P]_D = P$, which shows that the support function indeed represents the set exactly. If D is a finite set of directions, the outer approximation is a polyhedron, as illustrated in Fig. 10(b) and applied to reachability in Fig. 11. While for a given direction the numerical value of the support function



(a) using the axis directions gives a bounding box approximation of Ω_k from (14) (b) adding more directions, the approximation approaches Ω_k from (14)

Fig. 11 A reach set cover can be computed with support functions and initial approximations Ω_0, Ψ_δ from a variation of Lemma 3 where the bloating is non-uniform [68]. Evaluating the support function in a given set of directions results in the shown outer approximation $\hat{\Omega}_0, \hat{\Omega}_1, \dots$ (solid)

can often be computed very efficiently, one does not escape the curse of dimensionality if the goal is to compute an outer approximation of a given accuracy: To obtain an outer approximation within a Hausdorff distance ε of P in n dimensions, one needs to evaluate the support function in $\mathcal{O}(\frac{1}{\varepsilon^{n-1}})$ directions. Asymptotically optimal algorithms to construct ε -close approximations are described, e.g., in [116]. However, for some examples even a small number of directions can lead to reachability results with an acceptable approximation error [68].

Linear map, Minkowski sum, and convex hull are easily computed with support functions:

$$\begin{aligned} \varrho_{AP}(d) &= \varrho_P(A^\top d), \\ \varrho_{P \oplus Q}(d) &= \varrho_P(d) + \varrho_Q(d), \\ \varrho_{\text{chull}(P \cup Q)}(d) &= \max\{\varrho_P(d), \varrho_Q(d)\}. \end{aligned}$$

The intersection operation is more complex, and can be formulated as an optimization problem [83].

Thanks to the above properties, support functions serve well as a lazy representation for sets that arise from the successor computations described in Sect. 30.4.3.1. Computing the support function of the sequence (14) for a given direction can be done very efficiently even without the approximation operator Appr [83].

Two issues need to be solved to use support functions efficiently in the reachability Algorithm 1. First, the nesting of support functions should be of limited depth, in particular because evaluating the support function of an intersection operation requires multiple evaluations of its operands. Second, deciding containment is hard for support functions. Both problems can be solved by switching the set representation from a support function to its polyhedral outer approximation at appropriate points in the algorithm [68]. Combining support functions and polyhedral computations for a fixed set of directions D is closely related to reachability with *template polyhedra* [160] and both require that a good set of directions D be chosen. The support function representation can be extended to represent the entire (non-convex) reachable set by parameterizing it over time [69].

Polyhedra. The class of polyhedra is closed under all required operations, i.e., linear map, Minkowski sum, convex hull, and intersection. However, not all of them scale well. As mentioned in Sect. 30.4.2, there are no scalable algorithms for computing convex hull and Minkowski sum on polyhedra in constraint representation. For illustration, consider that using the convex hull of n line segments, each given by $2n$ constraints in n dimensions, one can construct a cross-polytope, which has 2^n constraints. Taking the Minkowski sum can lead to a similar explosion in the number of constraints. This is illustrated by the fact that the Minkowski sum can be computed with a convex hull and an intersection operation in $n + 1$ dimensions using the *Cayley Trick* [173]. A polyhedral approximation for the non-scalable operations can be efficiently computed by a priori fixing the facet normals of the result, e.g., using the outer approximation of the support function. The accuracy of the approximation can be increased by including additional directions, leading to a scalable approach [20].

30.4.3.3 Clustering

The accuracy of the approximation in Lemma 3 depends on the size of the time step. This property, common to all approaches cited in Sect. 30.4.3, points to a potential bottleneck: To achieve a desired accuracy, one may end up with a large number of sets to cover the required time horizon. In the next successor computation, each one of these sets may become the initial set of yet another sequence, and so one may easily end up with an exponential increase in the number of sets. If only very few of these sets intersect with the guard sets, the discrete successor computation results in few sets and therefore acts as a filter that might just keep the number of sets manageable. But this is not the case in general; note that these sets necessarily overlap. To prevent an explosion in the number of sets, a common approach is to cluster together all sets that intersect with the same guard [83]. The clustering operation, e.g., taking the convex hull, can itself be costly and adds to the approximation error in a way that is not easy to quantify. An approach to obtain a suboptimal number of clusters for a given error bound is presented in [69].

30.4.4 Nonlinear Dynamics

We give a very brief overview of techniques that deal with nonlinear dynamics

$$\dot{x} = f(x),$$

where f is usually assumed to be globally Lipschitz continuous.

Linearization. One way to deal with nonlinear dynamics is to approximate them with affine dynamics $\dot{x} = Ax + u$, $u \in \mathcal{U}$ and then use reachability algorithms for

affine dynamics. First, the states are confined to a bounded domain \mathcal{S} . This can be the evolution domain restriction in a location, or \mathcal{S} can be derived iteratively by growing suitable bounds around a given set of initial states. Then, a suitable matrix A and vector b are chosen. For example, linearizing $f(x)$ around a point $x_0 \in \mathcal{S}$ gives a matrix A with elements $a_{ij} = \frac{\partial f_i}{\partial x_j} |_{x=x_0}$ and a vector $b = f(x_0) - Ax_0$. Finally, one derives a set \mathcal{U}_ε that bounds the error such that for all $x \in \mathcal{S}$,

$$f(x) - (Ax + b) \in \mathcal{U}_\varepsilon.$$

Such bounds can be obtained using, e.g., interval arithmetic or optimization techniques. The states reachable using the affine dynamics $\dot{x} = Ax + u$, $u \in \mathcal{U}_\varepsilon \oplus \{b\}$ cover those of the original nonlinear dynamics. This approach constructs an *abstraction* of the system. Such abstractions are discussed more formally in Sect. 30.5.2.

The accuracy of the linearization depends on the size of the domain \mathcal{S} . It can be increased by partitioning \mathcal{S} into smaller parts. Each part can then be associated with smaller error bounds \mathcal{U}_ε and consequently gives a more accurate approximation of the reachable set. The switching of the system from one element of the partition to another is straightforward to model with a hybrid automaton. This process is known as *phase-portrait approximation*, see also Sect. 30.5.2. It can be of use even when dealing with purely continuous dynamical systems, in which case it is also referred to as *hybridization* [18]. The abstract model can be simplified by projecting away variables and adding a clock variable to preserve timing properties [17].

Polynomial Approximations. If the dynamics are polynomial, bringing them into Bernstein form allows one to compute conservative approximations of successor sets in polynomial form [54, 152]. Another approach is to use *Taylor models*, which are polynomial approximations of a function that are derived from a higher-order Taylor expansion and an interval bound on the remainder [30]. The resulting ODE can be solved by iterative approximations using the Picard operator. The reachable states are approximated by sets that are polyhedra [160] or polynomial images of intervals [43]. A similar approach uses polynomial images of zonotopes, which are themselves images of intervals [4]. Since polynomial images of intervals are generally not closed under intersection, the accuracy may be diminished when computing discrete successors. It can also be shown that additional assumptions, such as knowledge of a Lipschitz constant, are required in these approaches in order to ensure computable error bounds [147].

30.5 Abstraction-Based Verification

Explicit-state reachability analysis is very easy to use. Its flat and direct representation of the system behavior can, however, cause it to run into scalability issues for bigger systems. One technique that has been very successful for scaling up discrete model checking is that of abstraction (see also Chap. 13).

The basic idea is to replace the actual system by a simpler, abstract system, in which model checking is easier to perform. The verification results about the abstract system, of course, can only be related back to verification results about the original concrete system under certain conditions on how the abstract and concrete system are related and whether the particular property in question survives this abstraction process.

The options for directly constructing discrete abstractions by finite quotients and for which subclasses they work have been examined by Henzinger [88, 93] and Lafferriere et al. [108]. Because of the limited scope of discrete abstractions, more general predicate abstractions [9, 10] and abstraction refinement techniques like Counterexample-Guided Abstraction Refinement (CEGAR) have been developed subsequently [9, 46]; see Chap. 13. These directions have again worked successfully in discrete and, to some extent, real-time systems.

30.5.1 Discrete Abstractions

We present a general notion of abstraction for transition systems based on simulation relations [125] and we illustrate the principle of using abstractions in the verification of hybrid systems for the class of initialized rectangular automata.

Definition 11 (Abstraction) A transition system $T^A = \langle S^A, S_0^A, S_f^A, \Sigma, \rightarrow_A \rangle$ is an *abstraction* of a transition system (with the same alphabet) $T = \langle S, S_0, S_f, \Sigma, \rightarrow \rangle$ (which is then called the *concrete* system) if there exists an *abstraction mapping* $\alpha : S \rightarrow S^A$ such that the following conditions hold:

1. $\alpha(s) \in S_0^A$ for all initial states $s \in S_0$;
2. for all $\sigma \in \Sigma$, for all states $s_1, s_2 \in S$, if $s_1 \xrightarrow{\sigma} s_2$, then $\alpha(s_1) \xrightarrow{\sigma}_A \alpha(s_2)$;
3. $\alpha(s) \in S_f^A$ for all final states $s \in S_f$.

The abstraction mapping α is in fact a particular case of a (time-abstract) simulation relation [126]. It may be convenient to allow the abstraction mapping to map a state $s \in S$ to several abstract states $s_A^1, s_A^2, \dots, s_A^k \in S^A$, that is to consider abstraction mappings $\alpha : S \rightarrow 2^{S^A}$ or equivalently to consider an abstraction relation over $S \times S^A$, rather than a function. We take the simpler definition which is sufficient for the purpose of describing the main principles of abstraction for hybrid automata.

The main property of abstractions which is useful for the safety verification problem of hybrid automata is that they are conservative. Formally, $\{\alpha(s) \mid s \in \text{Reach}(T)\} \subseteq \text{Reach}(T^A)$, which implies the following.

Lemma 4 *Let T^A be an abstraction of T . If T^A is safe, then T is safe.*

By Lemma 4, if we show (e.g., using algorithmic techniques) that the unsafe states are not reachable in an abstraction of a hybrid system, then we can conclude

that the concrete system is safe. Intuitively, this is because abstractions are over-approximations of the original system, and therefore they exhibit (or simulate) all executions of the concrete system, and possibly more. In particular, every path to an unsafe state has a matching path in the abstraction, which is the main argument for proving Lemma 4. The converse of this lemma does not hold simply because abstractions may introduce spurious executions (which have no matching execution in the concrete system) due to over-approximation.

The main purpose of abstraction for hybrid systems is to obtain finite-state transition systems which are amenable to model checking by automated tools, and give useful conclusions about the original system. Remember that the transition systems of hybrid automata have (uncountably) infinite state space, and (uncountably) infinite branching. In the next subsections, we present ideas for practically constructing such abstractions.

Initialized Rectangular Automata. We illustrate abstractions with an informal argument of why the safety verification problem is decidable for initialized rectangular automata. The idea is that for such hybrid automata H , one can construct a timed automaton A such that A is an abstraction of H , and H is an abstraction of A , thus A is safe if and only if H is safe. Note that in this case the constructed abstraction (the timed automaton A) has infinite state space, but since we know that the safety verification problem for timed automata is decidable, we obtain decidability for initialized rectangular hybrid automata by Lemma 4.

We present the main steps behind this construction. In every location, a variable x with flow constraint $k_1 \leq \dot{x} \leq k_2$ is replaced by two variables x_l and x_u with flow constraint $\dot{x}_l = k_1$ and $\dot{x}_u = k_2$ which track the least and greatest possible value of x respectively. An incoming edge with jump condition $a \leq x^+ \leq b$ (an update) is replaced by $x_l^+ = a \wedge x_u^+ = b$. An edge with jump condition $x \leq b$ (a guard) that occurs in conjunction with $x^+ = x$ is replaced by two copies of the edge, one with the constraint $(x_l \leq b \wedge x_u \geq b \wedge x_u^+ = b)$ and the other with the constraint $x_u \leq b$. More complicated jump conditions (strict inequalities, and conjunction of simple jump conditions) are handled analogously, as well as the constraints in initial, final, and evolution domain conditions (invariants).

After this step, the slope of every variable is a singleton in every location. The next step is to scale the nonzero slope of the variables to 1. To do this, in each location we replace flow constraints $\dot{x} = k$ (when $k \neq 0$) by $\dot{x} = 1$ and divide by k the constants in the guards of outgoing edges, and in the updates of incoming edges. This ensures that the value stored in variable x remains k times smaller than the value of x in the original automaton (as long as the flow constraint $\dot{x} = k$ holds). It is therefore important that the rectangular automaton is initialized, as it guarantees that if the constraint $x^+ = x$ occurs in the jump condition of an edge (ℓ, σ, ℓ') , then the slope of variable x is the same in ℓ and in ℓ' . It remains to eliminate variables with slope 0, which can be done easily by storing the lower and upper value of x in the finite control structure of the automaton (these values can be changed only by discrete jumps).

The technical details of how to deal for instance with strict constraints in a jump condition like $(a < x < b)$, or unbounded flow constraints (like $\dot{x} \geq 1$) can be found in [96].

30.5.2 Phase-Portrait Approximation

Phase-portrait approximations are used as abstractions of hybrid automata with complex flow constraints. We discuss the approach for affine flow constraints, but it also applies to flow constraints that are much more general (e.g., given by $\dot{x} = f(x)$ for a continuous function f). Details about the theory and practice of this approach can be found, e.g., in [67, 93] and extensions on hybridization in [18] and other abstractions in [60, 70, 169].


The objective of phase-portrait approximations is to replace complex dynamics by simple rectangular (or sometimes linear) flow constraints on the dotted variables only. For example, the flow constraint $\dot{x} = f(x)$ where $f(x) = x + 1$ in a location with evolution domain (invariant) $0 \leq x \leq 10$ is replaced by $1 \leq \dot{x} \leq 11$, which over-approximates the exact dynamics. In general, bounds on the derivative can be derived from bounds on the variables and computed as optimization problems, where the lower bound should be less than or equal to $\inf_{v \in \llbracket \text{Inv}(\ell) \rrbracket} f(v)$ and symmetrically for the upper bound. Manual or numerical methods can be used as long as the bounds can be proven to hold.

Formally, a *phase-portrait approximation* of a hybrid automaton $H = \langle \text{Loc}, \text{Lab}, \text{Edg}, X, \text{Init}, \text{Inv}, \text{Flow}, \text{Jump}, \text{Final} \rangle$ is a hybrid automaton $H' = \langle \text{Loc}, \text{Lab}, \text{Edg}, X, \text{Init}, \text{Inv}, \text{Flow}', \text{Jump}, \text{Final} \rangle$ in which all components in H and H' are identical, except the flow constraint which is such that $\llbracket \text{Flow}'(\ell) \rrbracket \supseteq \llbracket \text{Flow}(\ell) \rrbracket$ for every location $\ell \in \text{Loc}$.

Lemma 5 *Let H' be a phase-portrait approximation of H . Then $\llbracket H' \rrbracket$ is an abstraction of $\llbracket H \rrbracket$, and if $\llbracket H' \rrbracket$ is safe, then $\llbracket H \rrbracket$ is safe.*

The safety verification problem for phase-portrait approximations can be solved using the algorithms and data structure presented in Sect. 30.4 for reachability analysis. Rectangular phase-portrait approximations are relatively simple to obtain because bounds are computed for each variable separately. However, the quality of the approximation may be too coarse to establish safety. If the bad states are reachable in the phase-portrait approximation, it may be due to lack of accuracy. More precise approximations are obtained by splitting the evolution domains. For example, a location with evolution domain $0 \leq x \leq 10$ can be replaced by two locations with respective evolution domains $0 \leq x \leq 5$ and $5 \leq x \leq 10$, over which the approximation of $\dot{x} = x + 1$ is more precise, namely $1 \leq \dot{x} \leq 6$ and $6 \leq \dot{x} \leq 11$ respectively. Figure 12 shows the states reachable from $x = t = 0$ (assuming $\dot{t} = 1$) in the rectangular phase-portrait approximation before splitting (light gray) and after splitting (dark gray).

Fig. 12 Tighter approximations using evolution domain (invariant) splitting



In general, splitting consists of replacing a location ℓ by k locations ℓ_1, \dots, ℓ_k with the same flow constraint as in ℓ , and with evolution domains that cover the evolution domain of ℓ , i.e., such that $\llbracket \text{Inv}(\ell) \rrbracket \subseteq \bigcup_{i=1}^k \llbracket \text{Inv}(\ell_i) \rrbracket$. For each incoming edge (ℓ', σ, ℓ) , new edges (ℓ', σ, ℓ_i) ($i = 1, \dots, k$) are created with the same jump condition, and similarly for each outgoing edge. The split locations ℓ_1, \dots, ℓ_k are connected by edges with jump condition $\text{stable}(X) = \bigwedge_{x \in X} x' = x$. It can be shown that location splitting results in hybrid automata that are mutually abstractions of each other, implying that one is safe if and only if the other is safe. By splitting locations, rectangular phase-portrait approximation can be made arbitrarily precise in the following sense. Given a hybrid automaton H and $\varepsilon > 0$, an ε -relaxation of H is a hybrid automaton with the same locations and transition structure as in H , and where all predicates ϕ in H are replaced by predicates ϕ' such that $\llbracket \phi \rrbracket \subseteq \llbracket \phi' \rrbracket \subseteq \llbracket \phi \rrbracket_\varepsilon$ where $\llbracket \phi \rrbracket_\varepsilon := \{v \in \mathbb{R}^X \mid \exists u \in \llbracket \phi \rrbracket : \max_{x \in X} |v(x) - u(x)| \leq \varepsilon\}$ is the set of valuations at a distance at most ε from a valuation satisfying ϕ .

It can be shown that for every hybrid automaton H and $\varepsilon > 0$, there exists a rectangular phase-portrait approximation H_ε of a splitting of H such that H_ε is an abstraction of H , and there exists an ε -relaxation of H which is an abstraction of H_ε (see [93]). This ensures that if H robustly satisfies a safety property (i.e., both H and some ε -relaxation satisfy the safety property), then it is possible to establish the property using rectangular phase-portrait approximation and splitting.

In practice, it is often useful to split locations according to specific information we may have about the given hybrid automaton. For example, a flow constraint $\dot{x} = 3 - x$ suggests the evolution domain should be split along lines parallel to $L \equiv 3 - x = 0$. More generally, a common heuristic is to use linear approximations of the flow constraints as support for cutting planes.

30.5.3 Predicate Abstractions

This part is a survey of [9, 10, 46]. We provide general ideas and guidelines about predicate abstraction schemes for hybrid systems. Chapter 13 provides a detailed

presentation of abstraction techniques for program verification (note that imperative programs can be viewed as a subclass of hybrid systems).

Reachability analysis based on predicate abstraction consists of tracking the truth value of a fixed finite set of predicates instead of computing the value of the continuous variables. The continuous part of the state space is replaced by the Boolean truth values of the predicates.

Let H be a hybrid system, and let $\Pi = \{\pi_1, \dots, \pi_k\}$ be a finite set of linear predicates π_i of the form $y \bowtie 0$ where $y \in \text{LTerm}(X)$ and $\bowtie \in \{<, \leq, =, >, \geq\}$. A truth value for Π is a vector $b \in \mathbb{B}^k$ where $\mathbb{B} = \{0, 1\}$ that assigns a truth value b_i to each predicate $\pi_i \in \Pi$. Truth values induce a partition of the continuous state space into finitely many abstract states. To obtain an abstraction we require that whenever there exists a transition between two concrete states, then there is a transition between the corresponding abstract states. Hence, the transition relation satisfies Definition 11 by construction.

We define an abstraction mapping α_Π as follows. For all states (ℓ, v) of the hybrid automaton H , let $\alpha_\Pi(\ell, v) = (\ell, b)$ if $b = (b_1, \dots, b_k) \in \mathbb{B}^k$ is the vector of truth values of the predicates in Π under valuation v , i.e., such that $\pi_i(v) = b_i$ for all $1 \leq i \leq k$. We sometimes omit the location and write $\alpha_\Pi(v) = b$. We denote by γ_Π the *concretization function* such that $\gamma_\Pi(b) = \{v \in \mathbb{R}^X \mid \alpha_\Pi(v) = b\}$ for all $b \in \mathbb{B}^k$.

The *predicate abstraction* of H induced by Π is the finite-state transition system $H_\Pi = \langle S_\Pi, S_0, S_f, \Sigma, \rightarrow_\Pi \rangle$ where:

- $S = \{(\ell, b) \in \text{Loc} \times \mathbb{B}^k \mid \exists v \in \llbracket \text{Inv}(\ell) \rrbracket : \alpha_\Pi(\ell, v) = (\ell, b)\}$
- $S_0 = \{(\ell, b_0) \in S_\Pi \mid \exists v \in \llbracket \text{Init}(\ell) \rrbracket : \alpha_\Pi(\ell, v) = (\ell, b_0)\}$;
- $S_f = \{(\ell, b_f) \in S_\Pi \mid \exists v \in \llbracket \text{Final}(\ell) \rrbracket : \alpha_\Pi(\ell, v) = (\ell, b_f)\}$;
- $\Sigma = \text{Lab} \cup \{\text{time}\}$ where Lab is the alphabet of H ;
- For each $\sigma \in \text{Lab}$, the transition relation \rightarrow_Π contains all tuples $((\ell, b), \sigma, (\ell', b'))$ such that $\exists e = (\ell, \sigma, \ell') \in \text{Edg} \cdot \exists v \in \gamma_\Pi(b) \cdot \exists v' \in \gamma_\Pi(b') : (\ell, v) \xrightarrow{\sigma} (\ell', v')$; and the transition relation \rightarrow_Π contains the tuples $((\ell, b), \text{time}, (\ell', b'))$ such that $\ell' = \ell$ and $\exists r \geq 0 \cdot \exists v \in \gamma_\Pi(b) \cdot \exists v' \in \gamma_\Pi(b') : (\ell, v) \xrightarrow{r} (\ell, v')$.

While predicate abstractions are finite-state, their size can be of prohibitive computational cost. The number of states in H_Π is at most exponential in the number of predicates in Π . In practice though, many truth value vectors are not feasible (i.e., they have an empty concretization). For example, think of a set of $2k$ predicates over two variables x and y , where k predicates define a partition of the values for x (e.g., $x < 0$, $0 \leq x \leq 1$, and $1 < x$) and k predicates define a partition of the values for y . Then the number of feasible abstract states is at most k^2 rather than 2^{2k} . Note that this example would still give a number of abstract states exponential in the number of variables. The dimension of the space is a well-known source of computational complexity. The choice of predicates is thus very important to obtain precise approximations at the least cost. The initial set of predicates is usually chosen manually. Natural candidates are the predicates occurring in the hybrid automaton itself, like the evolution domains and jump conditions. Automatic construction and refinement of predicate abstractions is discussed in Sect. 30.5.4.

For reachability analysis, it is usually not necessary to construct the entire transition systems of the predicate abstractions, because many states may not be reachable. On-the-fly approaches are used to simultaneously construct and explore the abstraction. Starting from the initial states in the abstraction, the transitions to other abstract states are explored as and when they are computed. A classical strategy is to explore the discrete successors first (because they are less expensive to compute), and then the continuous successors for increasing amounts of time, as long as no new discrete transition is enabled.

Computing Discrete Successors. Discrete successors can be computed as follows. Given $b \in \mathbb{B}^k$, let $\Pi(b) = \bigwedge_{i|b_i=1} \pi_i \wedge \bigwedge_{i|b_i=0} \overline{\pi_i}$ be the constraint defining the abstract state b . A transition $e = (\ell, \sigma, \ell')$ is *enabled* in a state (ℓ, b) if $\text{EN}(e) := \llbracket \exists X \cdot \Pi(b) \wedge \text{Inv}(\ell) \wedge \text{Jump}(e) \rrbracket \cap \llbracket \text{Inv}(\ell') \rrbracket \neq \emptyset$. The successor states of (ℓ, b) by enabled transition e are the abstract states (ℓ', b') such that $R_k \neq \emptyset$ where $R_0 = \text{EN}(e)$ and for all $1 \leq i \leq k$, if $b'_i = 1$ then $R_i = R_{i-1} \cap \llbracket \pi_i \rrbracket$, and if $b'_i = 0$ then $R_i = R_{i-1} \cap \llbracket \overline{\pi_i} \rrbracket$. A procedure for computing b' can easily be derived from this definition. Note that it may be that both $R_{i-1} \cap \llbracket \pi_i \rrbracket \neq \emptyset$ and $R_{i-1} \cap \llbracket \overline{\pi_i} \rrbracket \neq \emptyset$ hold, which would lead $b'_i = 1$ and $b'_i = 0$ to be set successively, and both cases to be explored. A simple optimization of this procedure is for each $1 \leq i \leq k$ to set $b'_i = 1$ beforehand if $\text{EN}(e) \cap \llbracket \overline{\pi_i} \rrbracket = \emptyset$, and set $b'_i = 0$ if $\text{EN}(e) \cap \llbracket \pi_i \rrbracket = \emptyset$. If one of the two cases holds, then the corresponding predicate can be skipped in the computation of R_i 's.

Computing Continuous Successors. In general, the continuous successors are not computed exactly, even according to the abstract transition relation. This is due to the lack of exact algorithmic methods for solving differential equations. Note that this is a difficult problem even if the differential equations in the flow constraints have closed-form solutions, like in linear systems. Given $R \subseteq \mathbb{R}^X$ and location ℓ , we want to compute the set $\text{Post}_C(\{\ell\} \times R)$ of continuous successor states as defined in Sect. 30.4.1, but over-approximations are sufficient for our purpose. This is consistent with the framework of abstraction (in the sense of Definition 11), but strictly speaking we are exploring in this way an over-approximation of the transition system H_{Π} defined above.

Optimizations. Various optimizations and heuristics have been defined and evaluated on many examples in the literature, see, e.g., [9, 10, 46]. For example, when we discover that a new abstract state s is reachable as a continuous successor under some flow constraint, we do not need to explore the continuous successors of s under the same flow constraint (unless s is also reachable by some discrete transition). This may significantly prune the search through the abstract state space. The search can also be guided to discover unsafe reachable states as quickly as possible. Various exploration strategies have been defined, based on giving priority to the most promising states, according to some greedy measures. For example, such measures may estimate the distance from the current state to the unsafe state, such as the Euclidean distance between the valuation of the variables in the abstract state

and in the unsafe states, or a discrete distance as the smallest number of discrete transitions necessary to reach an unsafe state, possibly taking into account the jump condition on the edges. Combination thereof are also possible [10]. Finally, as in program verification [90], it may be useful to maintain a set of predicates Π specific to each location, because certain predicates that are relevant in one location may not be useful in other locations.

30.5.4 Abstraction Refinement

The abstraction schemes presented in Sect. 30.5.1 and Sect. 30.5.3 may not be sufficient to establish the safety of a system. In particular, we know that safety of the abstraction implies safety of the original system, but non-safety of the abstraction is inconclusive. The process of refinement consists of constructing abstractions that are tighter (or more detailed) than a given abstraction, in order to prove safety. If the refinement process repeatedly fails in proving safety, then one can reasonably conclude that even if the original system may indeed be safe, it should not be considered as acceptable because its correctness is not robust, a small deviation in the implementation of the system being able to cause violation of the safety requirement [58, 155]. Such considerations are used to stop the refinement process when a specified level of precision is reached [46, 64].

In general, if T^A is an abstraction of T , then a *refinement* T^B of T^A is an abstraction of T which is such that T^A is an abstraction of T^B .

In the case of splitting and phase-portrait approximations, refinements can be obtained by further splitting locations. For predicate abstractions, adding new predicates gives a refinement. We present one of the most popular frameworks to discover new predicates automatically, the counterexample-guided abstraction refinement (CEGAR) [46, 47]. A general framework of abstraction refinement is presented in Chap. 13.

Spurious Counterexamples. When a predicate abstraction fails to establish safety, the analysis usually returns a witness path from an initial abstract state to a final abstract state. Such a path $\rho = q_0 \xrightarrow{\sigma_1} q_1 \dots \xrightarrow{\sigma_n} q_n$ is a *spurious* counterexample if there exists no path $(\ell_0, v_0) \xrightarrow{\sigma_1} (\ell_1, v_1) \dots \xrightarrow{\sigma_n} (\ell_n, v_n)$ in the original system such that $(\ell_i, v_i) \in \gamma_\Pi(q_i)$ for all $0 \leq i \leq n$. Clearly, if a counterexample is not spurious, then we can immediately conclude that the original system is not safe. We present a standard approach to check whether a counterexample is spurious [9].

To simplify the presentation, we assume in this section that every edge has a different label that identifies it uniquely. The successor operator is

$$\text{Post}_\sigma(S) = \{(\ell', v') \mid \exists(\ell, v) \in S : (\ell, v) \xrightarrow{\sigma} (\ell', v')\},$$

where $\text{Post}_{\text{time}}(\cdot) = \text{Post}_C(\cdot)$ is the one-step continuous successor operator. Similarly, the predecessor operator is

$$\text{Pre}_\sigma(S) = \{(\ell, v) \mid \exists(\ell', v') \in S : (\ell, v) \xrightarrow{\sigma} (\ell', v')\}.$$

Let $R_0 = \gamma_\Pi(q_0) \cap \{(\ell, v) \mid v \in \llbracket \text{Inv}(\ell) \rrbracket\}$, and $R_{i+1} = \text{Post}_{\sigma_i}(R_i) \cap \gamma_\Pi(q_{i+1}) \cap \{(\ell, v) \mid v \in \llbracket \text{Inv}(\ell) \rrbracket\}$ for all $i \geq 0$. The counterexample ρ is spurious iff $R_i = \emptyset$ for some $0 \leq i \leq n$. Note that over-approximations of $\text{Post}_{\sigma_i}(\cdot)$ may suffice to show that a counterexample is spurious, but under-approximations are necessary to establish with certainty that a counterexample exists in the original system.

Refinement. Assume that the counterexample is spurious, and let $j \geq 0$ such that $R_j \neq \emptyset$ and $R_{j+1} = \emptyset$. Then it is easy to prove that $R_j \cap \text{Pre}_{\sigma_{j+1}}(\gamma_\Pi(q_{j+1})) = \emptyset$. New predicates should be added to the set Π in order to rule out the counterexample. Since $R_j \cap \text{Pre}_{\sigma_{j+1}}(\gamma_\Pi(q_{j+1})) = \emptyset$, we can search for a set of predicates which separates R_j and $\text{Pre}_{\sigma_{j+1}}(\gamma_\Pi(q_{j+1}))$. A set $\hat{\Pi}$ of predicates *separates* two sets R and Q if for every truth value $b \in \mathbb{B}^{\hat{\Pi}}$, we have either $\gamma_{\hat{\Pi}}(b) \cap R = \emptyset$ or $\gamma_{\hat{\Pi}}(b) \cap Q = \emptyset$.

Note that to separate closed polyhedra, one simple linear constraint is always sufficient, but since reachable states (and in particular states reachable by continuous flow) are approximated by non-convex unions of polyhedra, several simple constraints may be necessary. Several methods have been developed to separate polyhedral sets, which are beyond the scope of this chapter. We refer to [9] for references and discussion.

In some case, spuriousness can be established by analyzing fragments of the counterexample [46], i.e., trying to show that a sub-sequence in the counterexample is not feasible in the original system. Spurious fragments of length 2 are called *locally infeasible* in [9] and defined as follows: $q_{i-1} \xrightarrow{\sigma_i} q_i \xrightarrow{\sigma_{i+1}} q_{i+1}$ is spurious if $\text{Post}_{\sigma_i}(\gamma_\Pi(q_{i-1})) \cap \gamma_\Pi(q_i) \cap \text{Pre}_{\sigma_{i+1}}(\gamma_\Pi(q_{i+1})) = \emptyset$. Refinement is computed as above using separating predicates.

Various forms of robustness have been considered for hybrid systems, which basically work by not distinguishing between almost safe and almost unsafe hybrid systems so that incorrect answers from the analysis procedure are accepted for such borderline cases, but correct answers are required for clear-cut cases. Different notions of robustness have been considered successfully [64, 157, 158].

30.5.5 Approximate Bisimulations

For discrete systems, the relationships between systems can be described by the notions of language inclusion, simulation, and bisimulation. These concepts have been transposed to continuous and hybrid systems [85], and extended to take advantage of metrics over state spaces [78]. While traditional simulation and bisimulation relations require the output traces of related states to be identical, it suffices for metric

relations that they are sufficiently close. It is then possible to construct a discrete bisimilar quotient by discretizing the state space. The quotient is then amenable to verification and controller synthesis techniques for discrete systems [76].

We briefly sketch out the principle of approximate bisimulations in discrete time. Two states x_1, x_2 are in ε -bisimulation relation if their output values are within distance ε and for every successor state x'_1 of x_1 , x_2 has a matching successor state x'_2 so that x'_1 and x'_2 are also in the relation. As a consequence, the output traces of two states in the relation will never be more than ε apart. Note that the definition coincides with classical bisimulation for $\varepsilon = 0$. It is generally hard to compute ε -bisimulation relations exactly, but (under some mild assumptions) one can define a Lyapunov-like *bisimulation function* that maps pairs of states to a non-negative value, and whose sub-level sets are in an approximate bisimulation relation. The existence of bisimulation relations can be tied to certain types of stability (the tendency of the system to go to its equilibrium point). For example, a bisimulation function of a linear continuous system with dynamics $\dot{x} = Ax$ and output signal $y = Cx$ can be computed efficiently even for high-dimensional systems by solving a set of linear matrix inequalities (LMI) of the form $M \geq C^T C$ and $A^T M + M A \leq 0$. The LMI always has a solution if the system is stable. Therefore, two stable linear systems are always ε -bisimilar, and an upper bound on ε can be computed. Note that approximate bisimulations can be used to relate continuous-time to discrete-time systems, continuous-valued to discrete-valued systems, etc.

Verification by Simulation. Bisimulation relations can also be used to verify bounded-horizon properties on bounded regions by computing a finite number of trajectories, a technique called *verification by simulation* [61, 102]. Here, the proximity measure of the bisimulation relation is combined with a robustness measure on temporal logic formulas. Given an initial state x_0 from which a trajectory satisfies a temporal formula to some measure, a bisimulation metric allows one to identify a neighborhood of initial states that all satisfy the same formula. This is possible since the bisimulation metric guarantees that all trajectories from the neighborhood (including all trajectories starting in x_0) remain sufficiently close together to satisfy the formula. Given a (dense) bounded region of initial states, it is, under suitable assumptions, possible to identify a finite subset of initial states whose trajectories are sufficient to show that the system satisfies a temporal logic formula [79]. A similar approach has been developed for embedded control software [112]. Together with the work on robustness mentioned in Sect. 30.5.4, these results demonstrate how stability and robustness can be used to simplify verification tasks.

30.6 Logic-Based Verification

The working principle behind logic-based verification is to use logical formulas to characterizing some parts of the hybrid systems verification problem and to solve this verification problem or subproblem entirely by checking the corresponding logical formulas for validity. There are even verification techniques for hybrid systems

that are entirely based on logic and proof [137, 143], which are beyond the scope of this chapter, however. In this section we survey the basic principles behind these approaches and show what kind of reasoning can be used to verify safety properties of hybrid systems or their parts by showing the validity of logical formulas.⁷

We survey a number of different approaches that represent the verification problem by various logical formulas or logical constraints:

1. Polynomial barrier certificates [154];
2. Equational certificates from templates [159, 161];
3. Differential invariants [136, 142, 144, 148].

These logic-based verification approaches further have in common that they argue by invariance and are based on variations of the work of Sophus Lie, of Jean Gaston Darboux, or of Aleksandr Lyapunov. Differential invariants are based on Sophus Lie's 1867–1873 work on what are now called Lie derivatives and Lie groups. Equational certificates are based on Darboux's 1878 results [56] on a way to use Sophus Lie's approach. Barrier certificates are based on variations of Aleksandr Lyapunov's 1884–1892 work on a criterion for stability, which is used for safety instead [154]. The logic-based verification techniques for hybrid systems are complementary, so barrier certificates, equational templates, and differential invariants can be used together and also combined as abstractions with reachability analysis techniques.

Consider a location ℓ of a hybrid automaton with polynomial dynamics defined by polynomial differential equations. To emphasize that such a differential equation is considered only once even if it occurs in multiple different locations, it is also referred to as *continuous mode*. Let

$$\dot{x}_1 = f_1(x), \dots, \dot{x}_n = f_n(x)$$

be the polynomial differential equation system of the mode, which we abbreviate by the (vectorial) differential equation $\dot{x} = f(x)$. The mode ℓ has an invariant condition $\text{Inv} \in \text{PConstr}(X)$. What we want to understand in model checking of safety properties is whether the system will always stay in a safe region when it follows this continuous evolution mode starting from some initial region. We represent the desired initial region by a constraint $\text{Init} \in \text{PConstr}(X)$. Finally, we consider a constraint $\text{Safe} \in \text{PConstr}(X)$ defining the safe states for which we want to show that our system never leaves the set of states $\llbracket \text{Safe} \rrbracket$ satisfying Safe .

Definition 12 (Continuous mode safety problem) Let $\dot{x} = f(x)$ be a (vectorial) differential equation, i.e., a polynomial differential equation system

$$\dot{x}_1 = f_1(x), \dots, \dot{x}_n = f_n(x)$$

⁷It should be noted that the other verification techniques surveyed in this chapter benefit from logic as well, for example in their representation of big sets of states using simple logical formulas.

for the system variables $X = \{x_1, \dots, x_n\}$. A *continuous system* $(\text{Init}, \dot{x} = f(x), \text{Inv})$ consists of a constraint $\text{Inv} \in \text{PConstr}(X)$ for the invariant condition (or evolution domain restriction), and a constraint $\text{Init} \in \text{PConstr}(X)$ for the initial condition. We say that the continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$ is *safe with respect to* constraint $\text{Safe} \in \text{PConstr}(X)$ iff all $\delta \in \mathbb{R}^{\geq 0}$ and all continuously differentiable functions $\varphi : [0, \delta] \rightarrow \mathbb{R}^X$ with $\varphi(0) \in \llbracket \text{Init} \rrbracket$ also satisfy $\varphi(\delta) \in \llbracket \text{Safe} \rrbracket$ provided that $\dot{\varphi}(t) = f(\varphi(t))$ and $f(t) \in \llbracket \text{Inv} \rrbracket$ for all $t \in [0, \delta]$. We also say that the continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$ *respects* Safe if $(\text{Init}, \dot{x} = f(x), \text{Inv})$ is safe with respect to property Safe .

Logic-based verification techniques provide easily checkable witnesses to verify that a continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$ respects Safe . The immediate significance for model checking is that they induce abstractions that can be used to terminate a reachability computation.

Lemma 6 (Logical abstraction) *Let $(\text{Init}, \dot{x} = f(x), \text{Inv})$ be a continuous system of a mode $\ell \in \text{Loc}$ of a hybrid automaton. If $(\text{Init}, \dot{x} = f(x), \text{Inv})$ respects Safe , then*

$$\text{post}_\ell(\llbracket \text{Init} \rrbracket) \subseteq \llbracket \text{Safe} \rrbracket.$$

If the continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$ of a mode $\ell \in \text{Loc}$ of a hybrid automaton respects the desired safety property Safe , (continuous) reachability computation can be terminated for all states in any subset $P \subseteq \llbracket \text{Init} \rrbracket$, because, by monotonicity, Lemma 6 then implies

$$\text{post}_\ell(P) \subseteq \llbracket \text{Safe} \rrbracket.$$

In particular, notice that it is useful for fast reachability computation if we can identify big sets Init that make $(\text{Init}, \dot{x} = f(x), \text{Inv})$ respect Safe . These sets Init are often much bigger than the original initial sets from Definition 6.

The logic-based verification techniques mentioned above have in common that they provide easily checkable witnesses for the verification. They further enjoy the benefit that they can be used for highly nonlinear dynamics. The primary challenge in all cases is the need to first find the witnesses or their shape, which corresponds to the challenge of finding the right directions for support functions.

An interesting special case of the continuous safety problem from Definition 12 is the case where Init and Safe are the same formula F . If the continuous system $(F, \dot{x} = f(x), \text{Inv})$ is safe with respect to F , then F is called a (safety) *invariant*. In that case, Lemma 6 implies

$$\text{post}_\ell(\llbracket F \rrbracket) \subseteq \llbracket F \rrbracket.$$

That is, the continuous system will never be able to leave F . Thus, without reachability computation, one can conclude that reachable sets that are within $\llbracket F \rrbracket$ will stay there forever.

Observe that, despite the similar name, there is a crucial difference between an invariant condition Inv of a continuous system (or a mode in a hybrid system) and

a safety invariant F . The difference is that we need to verify whether F is a safety invariant, while we just assume that the system obeys the invariant condition Inv . That is why Inv is also called an evolution domain restriction, because it restricts the admissible evolution domain of the continuous system. So, Inv is part of the system model, yet F is part of a safety property that we verify for the system model.

One of many possible approaches to logic-based verification is the one that focuses on showing that a formula F is a global invariant of a hybrid automaton by showing that it is an invariant for each discrete transition and an invariant for each continuous transition of the automaton. The best case is if F is the safety property and turns out to be a global invariant of the system in this manner. This is generally somewhat overly simplistic, because the verification does not necessarily have to work with the same invariant F in all places so that multiple invariants need to be used instead. Nevertheless, having this simple example of a single global invariant in mind is a useful guiding principle for logic-based verification approaches.

Related arguments have also been used for invariant generation techniques for abstract interpretation [169]. Based on decidability results for o-minimal hybrid automata [109], this includes invariant generation techniques for linear systems based on Gröbner basis computations [169] rather than based on quantifier elimination [109]. The case of (hyper-rectangle) box invariants has been discussed in more detail elsewhere [170].

30.6.1 Polynomial Barrier Certificates

The basic idea behind barrier certificates is to find a barrier separating good and bad states that we can easily show to be impenetrable by the continuous system dynamics. Barrier certificates were proposed for safety verification in [154].

Theorem 4 (Weak barrier certificate [154]) *Let $(\text{Init}, \dot{x} = f(x), \text{Inv})$ be a continuous system with safety constraint **Safe**. If B is a (weak) barrier certificate for a continuous safety problem, i.e., a polynomial satisfying*

$$\begin{aligned} B(x) &\leq 0 \quad \text{for all initial states } x \in \llbracket \text{Init} \rrbracket, \\ B(x) &> 0 \quad \text{for all unsafe states } x \notin \llbracket \text{Safe} \rrbracket, \quad \text{and} \\ \frac{\partial B}{\partial x}(x) f(x) &\leq 0 \quad \text{for all states } x \in \llbracket \text{Inv} \rrbracket, \end{aligned}$$

*then the continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$ respects **Safe**.*

Barrier certificates themselves can be defined for more general non-polynomial cases, but the conditions are generally not computable when $f_i(x)$ and B are not polynomials or Inv , Init , and Safe are not polynomial constraints. The purpose of a barrier certificate is to separate safe from unsafe states in such a way that initial

states are safe, and the differential equations can easily be seen to never cross the barrier between safe and unsafe states.

The importance of barrier certificates comes from the fact that they reduce a reachability question (can we ever reach an unsafe state) to a simple check on the directional derivative $\frac{\partial B}{\partial x}(x)f(x)$ along ODE of the Barrier certificate.

It had originally been proposed [153] that barrier certificates only need to be checked on the boundary of the barrier and that it would be sufficient to check the third condition in Theorem 4 for all $x \in \llbracket \text{Inv} \rrbracket$ with $B(x) = 0$:

$$\frac{\partial B}{\partial x}(x)f(x) \leq 0 \quad \text{for all states } x \in \llbracket \text{Inv} \rrbracket \text{ with } B(x) = 0. \quad (16)$$

This condition is generally not strong enough and can lead to soundness issues, as the following example shows.

Example 2 When using condition (16), it looks as if the differential equation $\dot{x} = 1$ always stays in the region $\text{Safe} \equiv x^2 \leq 0$ because condition (16) succeeds as follows:

$$\frac{\partial x^2}{\partial x} 1 = 2x \leq 0 \quad \text{for all states } x \text{ with } x^2 = 0$$

This, however, is counterfactual, because the system $\dot{x} = 1$ will, of course, leave region $x^2 \leq 0$. Thus, the condition (16) is unsound. The same issue occurs for a suggestion on how to extend this approach to Boolean combinations of inequalities [84]. A discussion of the assumptions under which the conditions can be restricted to such subsets without losing soundness can be found in the literature [136, 142, 144, 154].

Checking on the boundary is sound, however, if the condition (16) is modified to a strict inequality, instead of a weak inequality:

Theorem 5 (Strict barrier certificate [154]) *Let $(\text{Init}, \dot{x} = f(x), \text{Inv})$ be a continuous system with safety constraint Safe . If B is a (strict) barrier certificate for a continuous safety problem, i.e., a polynomial satisfying*

$$\begin{aligned} B(x) &\leq 0 \quad \text{for all initial states } x \in \llbracket \text{Init} \rrbracket, \\ B(x) &> 0 \quad \text{for all unsafe states } x \notin \llbracket \text{Safe} \rrbracket, \quad \text{and} \end{aligned}$$

$$\frac{\partial B}{\partial x}(x)f(x) < 0 \quad \text{for all states } x \in \llbracket \text{Inv} \rrbracket \text{ with } B(x) = 0,$$

then the continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$ respects Safe .

Search procedures for barrier certificates include approaches that choose a degree-bound for the barrier certificate $B(x)$ and then turn the conditions from Theorem 4 into a convex optimization problem, which can be solved efficiently [154].

A similar approach has been proposed for Theorem 5, but the optimization problem is then non-convex [154], so optimizers can get stuck in local optima.

Barrier certificates can be extended to systems with disturbances and to switching diffusion systems [154]. We refer to the literature for a discussion of these generalizations and examples [154].

30.6.2 Equational Certificates

Equational certificates [159, 161] serve a purpose that has quite some similarity to barrier certificates. They were introduced [161] at the same time as barrier certificates [154], and later rephrased and generalized [159] similarly to a matrix reformulation of that idea [123]. Equational certificates have been investigated earlier by Darboux in 1878 [56] for continuous systems not in the context of hybrid systems. Like barrier certificates, the conditions of equational certificates make a reachability analysis superfluous, because they give a simple certificate showing a property of the system. One major difference of equational certificates compared to barrier certificates is that an equational certificate consists of a single polynomial *equation* $p(x) = 0$, while a barrier certificate consists of a single polynomial *inequality* $B(x) \leq 0$. The other major difference is the condition itself. It is an equational criterion, not using inequalities. Another minor difference is that an equational certificate $p(x) = 0$ shows invariance of the property $p(x) = 0$ instead of separating initial states from bad states. That is a minor difference, though, because Safe is an invariance property that can be read off from a barrier certificate that separates Init from \neg Safe.

Theorem 6 (Equational certificates [161]) *Let $p(x)$ be a polynomial and let $(p(x) = 0, \dot{x} = f(x), \text{Inv})$ be a continuous system. If there is a polynomial $g(x)$ such that*

$$\frac{\partial p}{\partial x}(x) f(x) = g(x) p(x)$$

for all $x \in \llbracket \text{Inv} \rrbracket$, then $(p(x) = 0, \dot{x} = f(x), \text{Inv})$ respects $p(x) = 0$. In particular, $p(x) = 0$ is an invariant of $(p(x) = 0, \dot{x} = f(x), \text{Inv})$.

The equational template approach for equational certificates [161] works as follows. The user chooses a template for the polynomial equation $p(x) = 0$ and the system then uses linear equation solving and/or Gröbner basis computations [38] to check whether the equational certificate condition from Theorem 6 holds. In general, the approach may use the decision procedures of quantifier elimination in real closed fields [49] to handle the nonlinear real arithmetic.

Common special cases of equational certificates include those where only numbers or only 0 is chosen for the polynomial $g(x)$. It had originally been proposed

informally [161] that it should also be sufficient in Theorem 6 to check

$$\frac{\partial p}{\partial x}(x)f(x) = 0 \quad \text{for all } x \in \llbracket \text{Inv} \rrbracket \text{ with } p(x) = 0. \quad (17)$$

This variation is generally not strong enough and can lead to soundness issues.

Example 3 When using condition (17), it may seem as if $x^2 = 0$ were an invariant of the differential equation $\dot{x} = 1$, because condition (17) succeeds as follows:

$$\frac{\partial x^2}{\partial x} 1 = 2x = 0 \quad \text{for all } x \text{ with } x^2 = 0$$

This, however, is counterfactual, because the system $\dot{x} = 1$ will, of course, falsify the safety condition $x^2 = 0$ right away. Thus, the condition (17) is an unsound variation of Theorem 6. We refer to the literature [136, 142] for a discussion of the conditions under which stronger assumptions can be assumed without losing soundness.

There are additional conditions on the system dynamics and p , however, under which the restriction (17) remains correct [142]. That line of research also identifies under which conditions equational templates and equational differential invariants are complete for verifying equational safety properties [142].

30.6.3 Differential Invariants and Logical Certificates

Differential invariants are a generalized form of logic-based witness techniques for hybrid systems and generalize equational certificates [161] and barrier certificates [153, 154]. Like equational certificates [161], a differential invariant can be an equation $p(x) = 0$. Like barrier certificates [153, 154], differential invariants can be inequalities like $p(x) \leq 0$. Differential invariants can be general logical formulas with propositional combinations of mixed equations, strict inequalities, and weak inequalities, and can be extended to contain quantifiers for distributed hybrid systems [138]. Differential invariants have been introduced in 2008 [136] and later refined to an automatic verification procedure that searches for differential invariants [148]. Further results about the theory of differential invariants can be found in the literature [142, 144].

Given a continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$, we want to check whether it respects *Safe*. As a short notation, we say that the formula $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ is *valid* if the continuous system $(\text{Init}, \dot{x} = f(x), \text{Inv})$ respects the safety condition *Safe*. That is, if that continuous system always stays in the region *Safe* when it follows differential equation $\dot{x} = f(x)$ restricted to the evolution domain region *Inv* and when started in any initial state satisfying *Init*. Even though more complex representations can be used, we assume *Init*, *Safe*, and *Inv* to be (semi-algebraic) polynomial constraints. A simple form corresponds to the case where *Init* and *Safe* are the same

formula F . If $F \rightarrow [\dot{x} = f(x) \& \text{Inv}]F$ is valid, then F is called a *continuous invariant* of the dynamics $\dot{x} = f(x) \& \text{Inv}$. That is, if the continuous system starts in F , then it will always stay in F .

In fact, the notation $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ can be understood as a logical formula. The logical formula $[\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ uses the modal operator $[\dot{x} = f(x) \& \text{Inv}]$ to say that formula Safe holds in all states that are reachable along the differential equation $\dot{x} = f(x)$ within evolution domain Inv . The implication $\text{Init} \rightarrow$ in $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ restricts this to only the set of initial states that satisfy Init . The same principle extends to a logic for hybrid systems [135–137, 141, 143] and to a logic for distributed hybrid systems [140]; see [143] for an overview. Both of these logics are relatively complete (similarly to relative completeness of Hoare calculus). That is, they can prove every valid formula about hybrid systems or (distributed) hybrid systems from elementary properties of differential equations. These results also give a precise construction lifting all verification techniques for continuous systems to hybrid systems [141].

Differential invariants can be equational formulas like equational certificates, they can include inequalities like barrier certificates, but they also include mixed cases, Boolean combinations, and cases with more complicated logical formulas.

Definition 13 (Continuous invariant) Let $(\text{Init}, \dot{x} = f(x), \text{Inv})$ be a continuous system with safety constraint Safe . Constraint F is a *continuous invariant* of $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ iff the following formulas are valid (true in all states):

1. $\text{Init} \wedge \text{Inv} \rightarrow F$ (induction start), and
2. $F \rightarrow [\dot{x} = f(x) \& \text{Inv}]F$ (induction step).

A continuous invariant F is *sufficiently strong* for $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ if, in addition, $F \rightarrow \text{Safe}$ is valid, because $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ is then valid.


It is easy to see that the existence of a sufficiently strong continuous invariant for $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ implies that the property $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ is valid.

Continuous invariants are useful notions, but they are not computational per se, because we still need to find a way to check the induction step. The induction start is reasonable, because it is just a constraint, which is a logical formula of first-order real arithmetic and thus decidable by quantifier elimination in real closed fields [49, 50, 166]. But we need to find a checkable representation of the induction step. A checkable condition is made formally precise using the notion of differential invariants.

Definition 14 (Differential invariant) Let $(\text{Init}, \dot{x} = f(x), \text{Inv})$ be a continuous system with safety constraint Safe . A polynomial constraint F is a *differential invariant* of $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ iff the following formulas are valid:

1. $\text{Init} \wedge \text{Inv} \rightarrow F$ (induction start), and
2. $\text{Inv} \rightarrow \nabla_{\dot{x}=f(x)} F$ (induction step),

Fig. 13 Differential invariant F



where $\nabla_{\dot{x}=f(x)} F$ is the conjunction of all directional derivatives of atomic formulas in F in the direction of the vector field of $\dot{x} = f(x)$ (the partial derivative of b by x_i is $\frac{\partial b}{\partial x_i}$):

$$\nabla_{\dot{x}=f(x)} F \equiv \bigwedge_{(b \sim c) \in F} \left(\sum_{i=1}^n \frac{\partial b}{\partial x_i} f_i(x) \right) \sim \left(\sum_{i=1}^n \frac{\partial c}{\partial x_i} f_i(x) \right)$$

where $\sim \in \{=, \geq, >, \leq, <\}$.

A differential invariant F is *sufficiently strong* for $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ if, in addition, $F \rightarrow \text{Safe}$ is valid (because $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ is then valid by Corollary 2 below).

The respective partial derivatives of terms are well defined in the Euclidean space spanned by the variables and can be computed symbolically [136, 137]. Differential invariants capture the condition showing that the formula F is only becoming more true when following the dynamics, not less true, see Fig. 13.

The central property of differential invariants for verification purposes is that they replace infeasible or impossible reachability analysis with feasible symbolic computation.

Theorem 7 (Principle of differential induction [136]) *All differential invariants are continuous invariants.*


Corollary 1 *If F is a differential invariant for $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$, then $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] F$ is valid.*

Corollary 2 *If F is a differential invariant for $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ that is sufficiently strong, then F is a continuous invariant that is sufficiently strong for $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$. In particular, $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$ is valid.*

Example 4 Consider the dynamics $\dot{x} = x^4, \dot{y} = -2$. We are interested in seeing whether $2x \geq 5y$ is an invariant of this dynamics. With differential invariants it is easy to show that this is an invariant for the dynamics without using any state-based reachability verification. We just compute symbolically:

$$\nabla_{\dot{x}=x^4, \dot{y}=-2} (2x \geq 5y) \equiv \frac{\partial 2x}{\partial x} x^4 + \frac{\partial 2x}{\partial y} (-2) \geq \frac{\partial 5y}{\partial x} x^4 + \frac{\partial 5y}{\partial y} (-2) \equiv 2x^4 \geq -10.$$

Fig. 14 Example of a differential invariant indicated by the thick boundary



Since the latter formula is easily found to be valid, $2x \geq 5y$ is proven to be a differential invariant and thus stays true whenever it holds for the initial state of the dynamics.

Consider the case where $\dot{x} = x^4, \dot{y} = -2$ is the dynamics of one location of a hybrid automaton. Then we know that $2x \geq 5y$ is true after staying in this location arbitrarily long, if only we know that $2x \geq 5y$ is also true initially when entering the location. This is a prototypical scenario where local verification results also need to be combined together in order to verify the whole hybrid automaton.

Example 5 Consider the dynamics

$$\dot{x}_1 = 2x_1^4x_2 + 4x_1^2x_2^3 - 6x_1^2x_2, \quad \dot{x}_2 = -4x_1^3x_2^2 - 2x_1x_2^4 + 6x_1x_2^2.$$

Using differential invariants it is easy to show that $x_1^4x_2^2 + x_1^2x_2^4 - 3x_1^2x_2^2 + 1 \leq c$ is an invariant of this dynamics, as illustrated in Fig. 14. The justification again follows by simple symbolic computation as in Example 4:

$$\begin{aligned} \nabla_{\dot{x}}(x_1^4x_2^2 + x_1^2x_2^4 - 3x_1^2x_2^2 + 1) &= (2x_1^4x_2 + 4x_1^2x_2^3 - 6x_1^2x_2)(x_1^4x_2^2 + x_1^2x_2^4 - 3x_1^2x_2^2 + 1) \\ &\quad + (-4x_1^3x_2^2 - 2x_1x_2^4 + 6x_1x_2^2)(x_1^4x_2^2 + x_1^2x_2^4 - 3x_1^2x_2^2 + 1) \\ &\equiv \frac{\partial(x_1^4x_2^2 + x_1^2x_2^4 - 3x_1^2x_2^2 + 1)}{\partial x_1}(2x_1^4x_2 + 4x_1^2x_2^3 - 6x_1^2x_2) \\ &\quad + \frac{\partial(x_1^4x_2^2 + x_1^2x_2^4 - 3x_1^2x_2^2 + 1)}{\partial x_2}(-4x_1^3x_2^2 - 2x_1x_2^4 + 6x_1x_2^2) \leq 0 \end{aligned}$$

which simplifies to true.

Differential invariants work somewhat like loop invariants but for differential equations instead of loops. When checking a loop invariant F , we can assume it holds before the loop in the induction step. It thus looks as if we should be able to assume F when proving the induction step Case 2 of Definition 14 and prove

$$\text{Inv} \wedge F \rightarrow \nabla_{\dot{x}=f(x)} F \tag{18}$$

instead. Or, better, yet, only check the condition on the boundary of the domain like for barrier certificates. Neither of those would be sound, however, according to the following counterexamples from [136, 148]:

Example 6 When using condition (18), it looks as if $x^2 \leq 0$ were an invariant of the differential equation $\dot{x} = 1$, because condition (18) succeeds as follows:

$$(x^2 \leq 0 \rightarrow \nabla_{\dot{x}=1} x^2 \leq 0) \equiv \left(x^2 \leq 0 \rightarrow \frac{\partial x^2}{\partial x} 1 \leq 0 \right) \equiv (x^2 \leq 0 \rightarrow 2x \leq 0)$$

This, however, is counterfactual, because the system $\dot{x} = 1$ will, of course, leave region $x^2 \leq 0$. Thus, the condition (18) is unsound. The same example shows that checking on the boundary of F is unsound in general. We refer to the original work [136] for a discussion of the conditions under which stronger assumptions can be made without losing soundness.

A further elaboration of these phenomena as well as an identification of the conditions under which such extra assumptions would be sound can be found in the literature [142, 144].

It turns out that some properties cannot be verified using differential invariants alone but that additional verification techniques are needed [144]. Differential saturation (repeated application of differential cuts [136, 144]) has been introduced together with differential invariants in 2008 [136] as a sound alternative that can be used to add conditions iteratively without compromising soundness.

Theorem 8 (Differential saturation [136, 144]) *Assume that F is a continuous invariant (e.g., a differential invariant) of $\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe}$, then*

$$\text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv}] \text{Safe} \quad \text{iff} \quad \text{Init} \rightarrow [\dot{x} = f(x) \& \text{Inv} \wedge F] \text{Safe}.$$

An evolution domain constraint Inv (also confusingly referred to as the invariant of a location) is an entirely different entity than an invariant property F of a system. An automaton model *assumes or prescribes* that the system dynamics can only be followed along traces that do not leave Inv , because the system will stop all executions that leave Inv . In contrast, a differential invariant *proves* that the system will never leave F whether it wants to or not. Nevertheless, Theorem 8 gives a sound way of translating a proved differential invariant into a prescriptive evolution domain constraint. Theorem 8 can be used to strengthen the evolution domain constraints to subregions, which then become available for subsequent verification in a sound way. The differential cut principle underlying Theorem 8 is particularly powerful when used repeatedly until saturation [136, 148]. That is, verification with differential invariants often proceeds in stages, where a number of formulas F are verified to be invariants and then used to constrain the evolution of the system using the right-hand side of Theorem 8. This process repeats until all unsafe states have

been verified to be removable from the state space and so verification becomes trivial. Repeating this process in a fixed-point loop has been shown to work successfully in practice [148].

Differential invariants are computationally attractive concepts, because their induction start and induction step are just polynomial constraints, which are formulas of first-order real arithmetic, and are thus decidable by quantifier elimination in real closed fields [49, 50, 166]. Also the check whether a differential invariant F is sufficiently strong to imply a polynomial safety constraint Safe is decidable. The steps needed to compute the induction step of a differential invariant are simple algebraic computations that can be automated easily.

Differential invariants are always sound. That is, every property that can be verified using a differential invariant is correct. The converse question is that of completeness, whether all relevant properties can be verified. It turns out that differential invariants alone are not complete.

Example 7 $x > 0 \rightarrow [\dot{x} = -x \ \&\text{true}]x > 0$ is valid, but $x > 0$ is not a differential invariant of $\dot{x} = -x$, not a barrier certificate, and does not qualify as an equational template either.

More generally, it can be shown that there are properties like Example 7 that are true but cannot be verified [144], except when using an additional verification technique known as differential auxiliaries (alias differential ghosts) that adds additional variables and additional dynamics for verification purposes [144]. Thus, differential auxiliaries are a fundamental extension that is required for verification.


Search procedures for differential invariants include degree-bounded enumeration and fixed-point loops [148]. For completeness guarantees and numerous provability relationships on classes of differential invariants, see [141, 144]. The case of equational differential invariants is elaborated in [142], in which case differential invariants are a necessary and sufficient criterion for invariant functions according to a corresponding result by Lie.

Theorem 9 (Invariant function characterization) *A (polynomial) function p is an invariant function of $\dot{x} = f(x)$, i.e., the value of p along all solutions is constant, iff $p = 0$ is a differential invariant of $\dot{x} = f(x)$.*

A corresponding necessary and sufficient characterization of all algebraic invariant equations of algebraic differential equations is possible with a higher-order generalization of equational differential invariants called *differential radical invariants* [73].

For hybrid systems, differential invariants are used by allowing separate invariants for the respective locations of the hybrid automaton. Consider the hybrid automaton in Fig. 15 and, for the moment, suppose that there are no discrete jumps, i.e., the reset relations are the identity relation. Then, we need to show that starting in F_1 for dynamics $\dot{x} = \theta_1$ will always stay in the region F_2 . In addition, we need to show that, when starting in F_2 the dynamics $\dot{x} = \theta_2$ will always stay in

Fig. 15 Example of a verification loop for a hybrid automaton



the region F_3 , and so on. That is, in general we need to show that, when starting in F_i , the dynamics $\dot{x} = \theta_i$ will always stay in the region $F_{(i+1)\%5}$. In the presence of non-trivial discrete jump relations, we also need to show that these jump relations preserve the respective invariant. That means, we need to show that the jump relation (including its guard) will always transform every state within the invariant region F_i of its source into the invariant region $F_{(i+1)\%5}$ of its target. Finally, we only know that the reachable states of the hybrid automaton are contained in the respective invariant regions F_i if the automaton also starts in the required invariant region F_0 of the initial location. That is, we need to check that the initial state is contained in F_0 .

To make this principle concrete, consider a flyable roundabout maneuver for air traffic control [149], which is a variation of roundabouts that have been proposed a decade before [175]. Flyable roundabouts follow a hybrid automaton similar to Fig. 15, but with locations that correspond to the various phases of the roundabout as depicted schematically in Fig. 16. The aircraft are initially in free flight (free), then, when a conflict arises, agree on a compatible roundabout collision avoidance maneuver (agree), approach the roundabout with an entry procedure (entry), follow the roundabout (circ), and then leave the roundabout (exit), until they are far enough away to enter free flight again. Such roundabout collision avoidance maneuvers for aircraft can be verified using differential invariants, see elsewhere [149] for details.

For an investigation of the theory of differential invariants, we refer to [73, 136, 142–144]. That line of research studies the theoretical and provability properties of differential invariants. It identifies a dozen relations either equating or separat-





Fig. 16 Phases of flyable roundabout maneuver and protocol cycle



DI_{Ω} : properties verifiable using differential invariants built with operators from Ω

Fig. 17 Differential invariance chart: identifies how classes of differential invariants relate to each other, where the operators in the differential invariants are restricted as indicated in subscript Ω

ing the verification power of various classes of differential invariants (Fig. 17 indicates strict inclusion, equivalence, and incomparability of verification power, respectively). These relations further imply that the inclusion of Boolean operators that differential invariants support makes it possible to verify more systems compared to the single polynomial inequalities of barrier certificates or the single polynomial equations of equational templates [144]. The subclass of systems that have equational systems as invariants, however, already have a single equational invariant [144]. Differential cuts, differential saturation, and differential auxiliaries have been identified as fundamental extensions [136, 144]. The surprisingly close relationship of differential invariants to classical discrete invariants has been explored in the literature [141]. The relationship of differential invariants to Lie’s seminal work, a differential operator view, and partial differential equations has been investigated along with a technique called the inverse characteristic method for generating differential invariants [142]. The generalization of differential radical invariants can be generated efficiently using symbolic linear algebra [73].

For a generalization of differential invariants to systems with disturbances and differential-algebraic equations, we refer to the literature [136, 137]. Differential invariants can be generalized to the case of quantified first-order formulas and to distributed hybrid systems [138]. The approach extends to a relatively complete logic for hybrid systems [135–137, 141, 143] and to a relatively complete logic for distributed hybrid systems [140]. Generalizations to reachability and progress conditions can be found elsewhere [136]. Generalizations to stochastic hybrid systems with stochastic differential equations have been proposed [139].

30.7 Verification Tools

Despite the undecidability of the general case, the safety verification problem has been attacked algorithmically: many of the classical tools (among others D/DT [19], CheckMate [45], HYTECH [89]) and many of the more recent tools (PHAVER [67]) use a symbolic analysis of the hybrid automaton with a forward and/or backward approach: starting from the initial (resp. unsafe) states, iterate the Post operator (resp.

Pre) until a fixed point is reached and then check emptiness of the intersection with the unsafe (resp. initial) states. By Theorem 2, these procedures are not guaranteed to terminate in general. As discussed in Sect. 30.4, a major issue is scalability, as the computational cost increases sharply with the number of continuous variables. Performance is achieved by overapproximating the Post operator, and overapproximation can also be used to force termination of the fixed-point procedure. The challenge is to find methods that scale and are still accurate enough to show safety.

We discuss a selection of hybrid systems verification tools representing different classes of approaches that we survey here. A complete overview of all tools is beyond the scope of this chapter. We focus on a subset of the verification tools for which a dedicated tool paper and at least some documentation is available. A more complete collection of tools can be found on the Web.⁸

HSolver: Interval Constraint Propagation

HSolver⁹ [158] is an open-source software package for the formal verification of safety properties of continuous-time hybrid systems. It allows hybrid systems with nonlinear ordinary differential equations and nonlinear jumps assuming a global compact domain restriction on all variables. Even though HSolver is based on fast machine-precision floating point arithmetic, it uses sound rounding, and hence the correctness of its results cannot be hampered by round-off errors. HSolver not only verifies (unbounded horizon) reachability properties of hybrid systems, but—in addition—it also computes abstractions of the input system. So, even for input systems that are unsafe, or for which exhaustive formal verification is too difficult, it will compute abstractions that can be used by other tools. For example, the abstractions could be used for guiding search for error trajectories of unsafe systems.

HSolver is not optimized for special classes of hybrid systems (e.g., systems such as linear hybrid automata that have very simple continuous dynamics). Moreover it does not yet provide mature support for finding counterexamples for unsafe input systems. The method used by HSolver is abstraction refinement based on interval constraint propagation [158], which incrementally refines an abstraction of the input system. Special care is taken to reflect as much information as possible in the abstraction without increasing its size.

HyTech: The HYbrid TECHnology Tool

HyTech¹⁰ [89] was the first tool for reachability analysis of PCDA (Linear Hybrid Automata). The system is specified as a product of automata that synchronize

⁸<http://wiki.grasp.upenn.edu/>.

⁹<http://hsolver.sourceforge.net/>.

¹⁰<http://embedded.eecs.berkeley.edu/research/hytech/>.

on transitions that share the same label. The tool has a simple command language similar to a basic imperative language, allowing the user to program his own exploration algorithms. The basic data type represents a union of polyhedra associated with each location of the product automaton. Operations such as Boolean operations, existential quantification, emptiness test, and reachability computation (using the Post operation) are provided. Error traces (counterexamples) can be produced in combination with reachability analysis.

HyTech uses polyhedra with the double description method, which combines constraint and generator representations. The post operators are those described in Sect. 30.4.2 and implemented with exact arithmetic. HyTech can be used for parametric analysis by viewing parameters as variables with first derivative equal to zero. For instance, existential quantification on the reachable states can be used to extract a constraint on the parameters such that a given region is reachable. HyTech has been used to model check an audio control protocol [99] and a steam boiler [98]. A main limitation of HyTech lies in its use of standard integer data types, which quickly leads to integer overflow.

KeYmaera: Logic and Differential Invariants for Compositional Verification

KeYmaera¹¹ [135–137, 141, 150] is a hybrid verification tool for hybrid systems that combines deductive, real algebraic, and computer algebraic prover technologies. It is an automated and interactive theorem prover for a natural specification and verification logic for hybrid systems. With this, the verification principle behind KeYmaera is fundamentally different and complementary to tools like HyTech [89], PHAVer [67], and SpaceEx [68]. KeYmaera supports differential dynamic logic ($d\mathcal{L}$) [135, 137, 141], which is a real-valued first-order dynamic logic for hybrid programs [135, 137, 141], a program notation for hybrid systems. KeYmaera also supports hybrid systems with nonlinear discrete jumps, nonlinear differential equations, differential-algebraic equations, differential inequalities, and systems with non-deterministic discrete or continuous input.

For automation, KeYmaera implements a number of automated proof strategies that decompose the hybrid system symbolically and prove properties of the full system by proving properties of its parts [137]. This compositional verification principle helps scale up verification, because KeYmaera verifies large systems by verifying properties of subsystems (also see assume-guarantee reasoning, Chap. 12). KeYmaera implements fixed-point procedures [148] that compute differential invariants and invariants in fixed-point loops, somewhat like the way classical model checkers compute reachable sets in fixed-point loops. KeYmaera is typically more suitable for verifying parametric hybrid systems than systems with a single numerical state, where simulation is more appropriate. KeYmaera has been used successfully for verifying case studies in train control [151], car control [114, 115], air

¹¹<http://symbolaris.com/info/KeYmaera.html>.

traffic management [149], mobile robotics [130] and surgical robotics [104]. The KeYmaera approach is described in a book about Logical Analysis of Hybrid Systems [137].

A comprehensive introduction is provided in a textbook on the logical foundations of cyber-physical systems [146]. This textbook also explains how the successor tool KeYmaera¹² [72] can achieve the same from a minimal soundness-critical core [145].

PHAVer: Polyhedral Hybrid Automaton Verifier

PHAVer¹³ [67] follows the same basic principles as HyTech. PHAVer is a formal verification tool for computing reachability and simulation relations of PCDA (Linear Hybrid Automata) from Sect. 30.4.2.

PHAVer uses standard operations on polyhedra for reachability computations over an infinite time horizon (similar to those used in HyTech), and the algorithm for computing simulation relations is a straightforward extension of these. Using unbounded integer arithmetic, the computations are exact and formally sound. In addition to PCDA reachability, PHAVer can overapproximate piecewise affine dynamics on the fly, computing an overapproximation of the reachable states that is invariant (all trajectories that start within the set stay within the set). While PCDA are undecidable, PHAVer provides formally sound and precise overapproximation and widening operators that can force termination at the cost of reduced precision. These operators also simplify the computed continuous sets and dynamics of the system, and may result in a considerable speed-up without much loss in precision. The checking of abstraction and equivalence with simulation relations can be applied compositionally, and a sound non-circular assume-guarantee rule is implemented [66]. However, since the required exact computations on polyhedra do not scale well, this approach is limited to very small systems.

With its exact computations and controllable overapproximations, PHAVer is suited to verifying formally stringent properties on small systems with simple dynamics, such as communication protocols with drifting clocks or buffer networks. PHAVer's disadvantage is that the employed polyhedra computations are generally exponential in the number of variables, so that scalability is limited. PHAVer has been used to verify oscillation properties of a voltage-controlled oscillator circuit with three state variables [70], and various academic benchmarks with simple dynamics and up to 14 continuous variables. Since 2011, PHAVer is part of the SpaceEx tool platform [68].

¹²<http://keymaeraX.org/>.

¹³http://www-verimag.imag.fr/~frehse/phaver_web/index.html.

SpaceEx: State Space Explorer

SpaceEx¹⁴ [68] is a tool platform for verifying hybrid systems. It can handle hybrid automata whose continuous and jump dynamics are piecewise affine with nondeterministic inputs, i.e., PWA from Sect. 30.4.3. Nondeterministic inputs are particularly useful for modeling the approximation error when nonlinear systems are brought into piecewise affine form. SpaceEx comes with a compositional modeling language. It allows one to specify complex systems in a modular fashion as a network of interacting hybrid automata with templates and nesting. In the SpaceEx model editor, components are connected in block diagrams known from control theory, and the evolution of continuous variables is specified by hybrid automata with differential algebraic equations and inequalities.

Several different algorithms are implemented on the SpaceEx platform, including an exact algorithm for PCDA and a simulator that can handle nonlinear dynamics. The main verification algorithms, called LGG [68] and STC [69], combine explicit set representations (polyhedra), implicit set representations (support functions), and linear programming to achieve high scalability while maintaining high accuracy. The reachable states are overapproximated in the form of template polyhedra, which are polyhedra whose facets are oriented according to a user-provided set of template directions. The algorithms use adaptive time steps to ensure that the approximation error in each template direction remains below a given value. Empirical measurements indicate that the complexity of the image computations is linear in the number of variables, quadratic in the number of template directions, and linear in the number of time-discretization steps.

The accuracy of the overapproximation can be increased arbitrarily by choosing smaller time steps and adding more template directions. To attain a given approximation error (in the Hausdorff sense), the number of template directions is worst-case exponential. In case studies, the developers of SpaceEx observe that a linear number of user-specified directions, possibly augmented by a small set of critical directions, often suffices. The LGG and STC algorithms use floating-point computations that do not formally guarantee soundness. SpaceEx has been used to verify continuous and hybrid systems with more than 100 continuous variables.

ToolboxLS: Level Set Methods

Level set methods are a class of algorithms designed for approximating the solution of the Hamilton–Jacobi partial differential equation (PDE) [127], which arises in many fields including optimal control, differential games, and dynamic implicit surfaces. In particular, dynamic implicit surfaces can be used to compute backward reachable sets and tubes for nonlinear, nondeterministic, continuous dynamic systems with control and/or disturbance inputs; in other words, inputs and parameters to

¹⁴<http://spaceex.imag.fr/>.

the model can be treated in a worst case and/or best case fashion. The strengths and weaknesses of the Hamilton–Jacobi PDE formulation of reachability are very similar to those of viability theory: it can treat very general dynamics with adversarial inputs and can represent very general sets, but the known computational algorithms require resources that grow exponentially with the number of dimensions (number of variables); for example, in ToolboxLS¹⁵ the level set algorithms run on a Cartesian grid of the state space. The ToolboxLS algorithms also do not guarantee the sign of computational errors, but they deliver higher accuracy for a given resolution than that available from typical sound alternatives.

Because ToolboxLS [128] is designed for dynamic implicit surfaces rather than specifically for reachability, it does not include a specialized verification interface; however, it has a 140-page user manual documenting the software and over twenty complete examples including three reachable-set computations. It has been used primarily for reachability of systems with two to four continuous dimensions, including collision avoidance, quadrotor flips, aerial refueling, automated landing, and glide-path recapture.

Acknowledgements The authors thank the anonymous reviewers of this chapter for helpful feedback, as well as David Henriques and João Martins for their proofreading help. We also thank Andrew Sogokon for suggesting the nice illustration shown in Fig. 14.

References

1. Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25–28, 2012. IEEE (2012)
2. Althoff, M., Krogh, B.: Reachability analysis of nonlinear differential-algebraic systems. *IEEE Trans. Autom. Control* **59**, 371–383 (2014)
3. Althoff, M., Krogh, B.H., Stursberg, O.: Analyzing reachability of linear dynamic systems with parametric uncertainties. In: Rauh, A., Auer, E. (eds.) *Modeling, Design, and Simulation of Systems with Uncertainties*. Springer, Heidelberg (2011)
4. Althoff, M.: Reachability analysis of nonlinear systems using conservative polynomialization and non-convex sets. In: *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, pp. 173–182. ACM, New York (2013)
5. Althoff, M., Krogh, B.H.: Avoiding geometric intersection operations in reachability analysis of hybrid systems. In: *Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control*, pp. 45–54. ACM, New York (2012)
6. Alur, R., Courcoubetis, C., Halbwachs, N., Henzinger, T.A., Ho, P.-H., Nicollin, X., Olivero, A., Sifakis, J., Yovine, S.: The algorithmic analysis of hybrid systems. *Theor. Comput. Sci.* **138**, 3–34 (1995)
7. Alur, R.: Formal verification of hybrid systems. In: Chakraborty et al. [41], pp. 273–278
8. Alur, R., Courcoubetis, C., Henzinger, T.A., Ho, P.-H.: Hybrid automata: an algorithmic approach to the specification and verification of hybrid systems. In: Grossman et al. [82], pp. 209–229
9. Alur, R., Dang, T., Ivancic, F.: Counterexample-guided predicate abstraction of hybrid systems. *Theor. Comput. Sci.* **354**(2), 250–271 (2006)

¹⁵<http://www.cs.ubc.ca/~mitchell/ToolboxLS/>.

10. Alur, R., Dang, T., Ivancic, F.: Predicate abstraction for reachability analysis of hybrid systems. *ACM Trans. Embed. Comput. Syst.* **5**(1), 152–199 (2006)
11. Alur, R., Dill, D.L.: A theory of timed automata. *Theor. Comput. Sci.* **126**(2), 183–235 (1994)
12. Alur, R., Henzinger, T., Lafferriere, G., Pappas, G.J.: Discrete abstractions of hybrid systems. *Proc. IEEE* **88**(7), 971–984 (2000)
13. Alur, R., Henzinger, T.A., Ho, P.-H.: Automatic symbolic verification of embedded systems. *IEEE Trans. Softw. Eng.* **22**(3), 181–201 (1996)
14. Alur, R., Henzinger, T.A., Vardi, M.Y.: Parametric real-time reasoning. In: *ACM Symposium on Theory of Computing*, pp. 592–601 (1993)
15. Alur, R., La Torre, S., Pappas, G.J.: Optimal paths in weighted timed automata. In: *Domenica Di Benedetto, M., Sangiovanni-Vincentelli, A.L. (eds.) HSCC. LNCS, vol. 2034*, pp. 49–62. Springer, Heidelberg (2001)
16. Alur, R., Pappas, G.J. (eds.): *Hybrid Systems: Computation and Control, Proceedings of the 7th International Workshop, HSCC 2004, Philadelphia, PA, USA, March 25–27, 2004. LNCS, vol. 2993*. Springer, Heidelberg (2004)
17. Asarin, E., Dang, T.: Abstraction by projection and application to multi-affine systems. In: *Alur and Pappas [16]*, pp. 32–47
18. Asarin, E., Dang, T., Girard, A.: Hybridization methods for the analysis of nonlinear systems. *Acta Inform.* **43**(7), 451–476 (2007)
19. Asarin, E., Dang, T., Maler, O., Bournez, O.: Approximate reachability analysis of piecewise-linear dynamical systems. In: *Proc. HSCC 00: Hybrid Systems—Computation and Control. LNCS, vol. 1790*, pp. 20–31. Springer, Heidelberg (2000)
20. Asarin, E., Dang, T., Maler, O., Testylier, R.: Using redundant constraints for refinement. In: *Automated Technology for Verification and Analysis*, pp. 37–51. Springer, Heidelberg (2010)
21. Asarin, E., Maler, O., Pnueli, A.: Reachability analysis of dynamical systems having piecewise-constant derivatives. *Theor. Comput. Sci.* **138**(1), 35–65 (1995)
22. Bagnara, R., Hill, P.M., Zaffanella, E.: The Parma Polyhedra Library: toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems. *Sci. Comput. Program.* **72**(1–2), 3–21 (2008)
23. Baier, C., Bertrand, N., Bouyer, P., Brihaye, T., Größer, M.: Almost-sure model checking of infinite paths in one-clock timed automata. In: *Proc. of LICS*, pp. 217–226. IEEE, Piscataway (2008)
24. Balluchi, A., Di Natale, F., Sangiovanni-Vincentelli, A.L., van Schuppen, J.H.: Synthesis for idle speed control of an automotive engine. In: *Alur and Pappas [16]*, pp. 80–94
25. Behrmann, G., Fehnker, A., Hune, T., Larsen, K.G., Pettersson, P., Romijn, J., Vaandrager, F.W.: Minimum-cost reachability for priced timed automata. In: *Proc. of HSCC. LNCS, vol. 2034*, pp. 147–161. Springer, Heidelberg (2001)
26. Bemporad, A., Bicchi, A., Buttazzo, G. (eds.): *Hybrid Systems: Computation and Control, Proceedings of the 10th International Conference, HSCC 2007, Pisa, Italy. LNCS, vol. 4416*. Springer, Heidelberg (2007)
27. Bemporad, A., Morari, M.: Verification of hybrid systems via mathematical programming. In: *Vaandrager and van Schuppen [178]*, pp. 31–45
28. Benerecetti, M., Faella, M., Minopoli, S.: Automatic synthesis of switching controllers for linear hybrid systems: safety control. *Theor. Comput. Sci.* **493**, 116–138 (2013)
29. Bergstra, J.A., Middelburg, C.A.: Process algebra for hybrid systems. *Theor. Comput. Sci.* **335**(2–3), 215–280 (2005)
30. Berz, M., Makino, K.: Verified integration of ODEs and flows using differential algebraic methods on high-order Taylor models. *Reliab. Comput.* **4**(4), 361–369 (1998)
31. Bicchi, A., Pallottino, L.: On optimal cooperative conflict resolution for air traffic management systems. *IEEE Trans. Intell. Transp. Syst.* **1**(4), 221–231 (2000)
32. Blum, L., Cucker, F., Shub, M., Smale, S.: *Complexity and Real Computation*. Springer, New York (1998)

33. Bouyer, P., Dufourd, C., Fleury, E., Petit, A.: Updatable timed automata. *Theor. Comput. Sci.* **321**(2–3), 291–345 (2004)
34. Branicky, M.S.: General hybrid dynamical systems: modeling, analysis, and control. In: Alur, R., Henzinger, T.A., Sontag, E.D. (eds.) *Hybrid Systems*, vol. 1066, pp. 186–200. Springer, Heidelberg (1995)
35. Branicky, M.S., Borkar, V.S., Mitter, S.K.: A unified framework for hybrid control: model and optimal control theory. *IEEE Trans. Autom. Control* **43**(1), 31–45 (1998)
36. Brihaye, T., Doyen, L., Geeraerts, G., Ouaknine, J., Raskin, J.-F., Worrell, J.: On reachability for hybrid automata over bounded time. In: *Proceedings of ICALP 2011: International Colloquium on Automata, Languages and Programming (Part II)*. LNCS, vol. 6756, pp. 416–427. Springer, Heidelberg (2011)
37. Bu, L., Li, Y., Wang, L., Li, X.: BACH: bounded reachability checker for linear hybrid automata. In: *Formal Methods in Computer-Aided Design, 2008. FMCAD'08*, pp. 1–4. IEEE, Piscataway (2008)
38. Buchberger, B.: An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal. PhD thesis, University of Innsbruck (1965)
39. Carloni, L.P., Passerone, R., Pinto, A., Sangiovanni-Vincentelli, A.L.: Languages and tools for hybrid systems design. *Found. Trends Electron. Des. Autom.* **1**(1/2), 1–193 (2006)
40. Cassez, F., Larsen, K.G.: The impressive power of stopwatches. In: *CONCUR*, pp. 138–152 (2000)
41. Chakraborty, S., Jerraya, A., Baruah, S.K., Fischmeister, S. (eds.): *Proceedings of the 11th International Conference on Embedded Software, EMSOFT 2011, Part of the Seventh Embedded Systems Week, ESWeek 2011, Taipei, Taiwan, October 9–14, 2011*. ACM, New York (2011)
42. Chase, C., Serrano, J., Ramadge, P.J.: Periodicity and chaos from switched flow systems: contrasting examples of discretely controlled continuous systems. *IEEE Trans. Autom. Control* **38**(1), 70–83 (1993)
43. Chen, X., Ábrahám, E., Sankaranarayanan, S.: Taylor model flowpipe construction for nonlinear hybrid systems. In: *RTSS*, pp. 183–192. IEEE, Piscataway (2012)
44. Chernikova, N.V.: Algorithm for discovering the set of all solutions of a linear programming problem. *USSR Comput. Math. Math. Phys.* **8**(6), 282–293 (1968)
45. Chutinan, A., Krogh, B.H.: Verification of polyhedral-invariant hybrid automata using polygonal flow pipe approximations. In: Vaandrager and van Schuppen [178], pp. 76–90
46. Clarke, E.M., Fehnker, A., Han, Z., Krogh, B.H., Ouaknine, J., Stursberg, O., Theobald, M.: Abstraction and counterexample-guided refinement in model checking of hybrid systems. *Int. J. Found. Comput. Sci.* **14**(4), 583–604 (2003)
47. Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: *Proc. of CAV 2000: Computer Aided Verification*. LNCS, vol. 1855, pp. 154–169. Springer, Heidelberg (2000)
48. Collet, P., Eckmann, J.P.: *Iterated Maps on the Interval as Dynamical Systems*, vol. 1. Springer, Heidelberg (1980)
49. Collins, G.E.: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: Barkhage, H. (ed.) *Automata Theory and Formal Languages*. LNCS, vol. 33, pp. 134–183. Springer, Heidelberg (1975)
50. Collins, G.E., Hong, H.: Partial cylindrical algebraic decomposition for quantifier elimination. *J. Symb. Comput.* **12**(3), 299–328 (1991)
51. Collins, P.: Optimal semicomputable approximations to reachable and invariant sets. *Theory Comput. Syst.* **41**(1), 33–48 (2007)
52. Cuijpers, P.J.L., Reniers, M.A.: Hybrid process algebra. *J. Log. Algebraic Program.* **62**(2), 191–245 (2005)
53. Damm, W., Hungar, H., Olderog, E.-R.: Verification of cooperating traffic agents. *Int. J. Control* **79**(5), 395–421 (2006)
54. Dang, T., Testylier, R.: Reachability analysis for polynomial dynamical systems using the Bernstein expansion. *Reliab. Comput.* **17**(2), 128–152 (2012)

55. Dantzig, G.B., Eaves, B.C.: Fourier-Motzkin elimination and its dual. *J. Comb. Theory* **14**, 288–297 (1973)
56. Darboux, J.-G.: Mémoire sur les équations différentielles algébriques du premier ordre et du premier degré. *Bull. Sci. Math. Astron.* **2**(1), 151–200 (1878)
57. De Wulf, M., Doyen, L., Markey, N., Raskin, J.-F.: Robust safety of timed automata. *Form. Methods Syst. Des.* **33**(1–3), 45–84 (2008)
58. De Wulf, M., Doyen, L., Raskin, J.-F.: Almost ASAP semantics: from timed models to timed implementations. *Form. Asp. Comput.* **17**(3), 319–341 (2005)
59. Doyen, L.: Robust parametric reachability for timed automata. *Inf. Process. Lett.* **102**(5), 208–213 (2007)
60. Doyen, L., Henzinger, T.A., Raskin, J.-F.: Automatic rectangular refinement of affine hybrid systems. In: *Proc. of FORMATS 2005: Formal Modelling and Analysis of Timed Systems*. LNCS, vol. 3829, pp. 144–161. Springer, Heidelberg (2005)
61. Fainekos, G.E., Girard, A., Pappas, G.J.: Temporal logic verification using simulation. In: *Formal Modeling and Analysis of Timed Systems*, pp. 171–186. Springer, Heidelberg (2006)
62. Fehnker, A., Krogh, B.H.: Hybrid system verification is not a sinecure—the electronic throttle control case study. *Int. J. Found. Comput. Sci.* **17**(4), 885–902 (2006)
63. Ferrante, J., Rackoff, C.: A decision procedure for the first order theory of real addition with order. *SIAM J. Comput.* **4**(1), 69–76 (1975)
64. Fränzle, M.: Analysis of hybrid systems: an ounce of realism can save an infinity of states. In: *CSL*. LNCS, vol. 1683, pp. 126–140. Springer, Heidelberg (1999)
65. Frazzoli, E., Grosu, R. (eds.): *Proceedings of the 14th ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2011, Chicago, USA, April 12–14, 2011*. ACM, New York (2011)
66. Frehse, G.: *Compositional verification of hybrid systems using simulation relations*. PhD thesis, Radboud Universiteit Nijmegen (October 2005)
67. Frehse, G.: PHAVer: algorithmic verification of hybrid systems past HyTech. *Int. J. Softw. Tools Technol. Transf.* **10**(3), 263–279 (2008)
68. Frehse, G., Le Guernic, C., Donzé, A., Ray, R., Lebeltel, O., Ripado, R., Girard, A., Dang, T., Maler, O.: SpaceEx: scalable verification of hybrid systems. In: *Gopalakrishnan, G., Qadeer, S. (eds.) CAV*. LNCS, vol. 6806. Springer, Heidelberg (2011)
69. Frehse, G., Kateja, R., Le Guernic, C.: Flowpipe approximation and clustering in space-time. In: *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, pp. 203–212. ACM, New York (2013)
70. Frehse, G., Krogh, B.H., Rutenbar, R.A.: Verifying analog oscillator circuits using forward/backward abstraction refinement. In: *Gielen, G.G.E. (ed.) DATE*, pp. 257–262. European Design and Automation Association, Leuven (2006)
71. Freund, R.M., Orlin, J.B.: On the complexity of four polyhedral set containment problems. *Math. Program.* **33**(2), 139–145 (1985)
72. Fulton, N., Mitsch, S., Quesel, J.-D., Völpl, M., Platzer, A.: KeYmaera X: an axiomatic tactical theorem prover for hybrid systems. In: *Felty, A.P., Middeldorp, A. (eds.) Automated Deduction, CADE-25*. LNCS, vol. 9195, pp. 527–538. Springer, Heidelberg (2015)
73. Ghorbal, K., Platzer, A.: Characterizing algebraic invariants by differential radical invariants. In: *Ábrahám, E., Havelund, K. (eds.) TACAS*. LNCS, vol. 8413, pp. 279–294. Springer, Heidelberg (2014)
74. Ghosh, P.K., Kumar, K.V.: Support function representation of convex bodies, its application in geometric computing, and some related representations. *Comput. Vis. Image Underst.* **72**(3), 379–403 (1998)
75. Girard, A.: Reachability of uncertain linear systems using zonotopes. In: *Morari, M., Thiele, L. (eds.) HSCC*. LNCS, vol. 3414, pp. 291–305. Springer, Heidelberg (2005)
76. Girard, A.: Controller synthesis for safety and reachability via approximate bisimulation. *Automatica* **48**(5), 947–953 (2012)

77. Girard, A., Le Guernic, C., Maler, O.: Efficient computation of reachable sets of linear time-invariant systems with inputs. In: Hespanha, J.P., Tiwari, A. (eds.) HSCC. LNCS, vol. 3927, pp. 257–271. Springer, Heidelberg (2006)
78. Girard, A., Pappas, G.J.: Approximation metrics for discrete and continuous systems. *IEEE Trans. Autom. Control* **52**(5), 782–798 (2007)
79. Girard, A., Zheng, G.: Verification of safety and liveness properties of metric transition systems. *ACM Trans. Embed. Comput. Syst.* **11**(S2), 54 (2012)
80. Goebel, R., Sanfelice, R.G., Teel, A.R.: Hybrid dynamical systems. *IEEE Control Syst. Mag.* **29**(2), 28–93 (2009)
81. Greenstreet, M.R.: Verifying safety properties of differential equations. In: Computer Aided Verification. LNCS, vol. 1102, pp. 277–287. Springer, Heidelberg (1996)
82. Grossman, R.L., Nerode, A., Ravn, A.P., Rischel, H. (eds.): Hybrid Systems. LNCS, vol. 736. Springer, Heidelberg (1993)
83. Le Guernic, C., Girard, A.: Reachability analysis of hybrid systems using support functions. In: Bouajjani, A., Maler, O. (eds.) CAV. LNCS, vol. 5643, pp. 540–554. Springer, Heidelberg (2009)
84. Gulwani, S., Tiwari, A.: Constraint-based approach for analysis of hybrid systems. In: Gupta, A., Malik, S. (eds.) CAV. LNCS, vol. 5123, pp. 190–203. Springer, Heidelberg (2008)
85. Haghverdi, E., Tabuada, P., Pappas, G.J.: Bisimulation relations for dynamical, control, and hybrid systems. *Theor. Comput. Sci.* **342**(2), 229–261 (2005)
86. Halbwachs, N., Proy, Y.-E., Raymond, P.: Verification of linear hybrid systems by means of convex approximations. In: International Static Analysis Symposium, SAS'94, Namur (1994)
87. Henzinger, T.A., Ho, P.-H., Wong-Toi, H.: HyTech: the next generation. In: Proceedings of the 16th IEEE Real-Time Systems Symposium (RTSS '95), p. 56. IEEE, Piscataway (1995)
88. Henzinger, T.A.: Hybrid automata with finite bisimulations. In: ICALP: Automata, Languages, and Programming. LNCS, vol. 944, pp. 324–335. Springer, Heidelberg (1995)
89. Henzinger, T.A., Ho, P.-H., Wong-Toi, H.: HyTech: a model checker for hybrid systems. *Softw. Tools Technol. Transf.* **1**, 110–122 (1997)
90. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy abstraction. In: Proceedings of the 29th Annual Symposium on Principles of Programming Languages, pp. 58–70. ACM, New York (2002)
91. Henzinger, T.A.: The theory of hybrid automata. In: Inan, M.K., Kurshan, R.P. (eds.) Verification of Digital and Hybrid Systems. NATO ASI Series F: Computer and Systems Sciences, vol. 170, pp. 265–292. Springer, Heidelberg (2000)
92. Henzinger, T.A., Ho, P.-H.: A note on abstract interpretation strategies for hybrid automata. In: Hybrid Systems II, pp. 252–264. Springer, Heidelberg (1995)
93. Henzinger, T.A., Ho, P.-H., Wong-Toi, H.: Algorithmic analysis of nonlinear hybrid systems. *IEEE Trans. Autom. Control* **43**, 540–554 (1998)
94. Henzinger, T.A., Kopke, P.W.: State equivalences for rectangular hybrid automata. In: CONCUR: Concurrency Theory. LNCS, vol. 1119, pp. 530–545. Springer, Heidelberg (1996)
95. Henzinger, T.A., Kopke, P.W.: Discrete-time control for rectangular hybrid automata. *Theor. Comput. Sci.* **221**, 369–392 (1999)
96. Henzinger, T.A., Kopke, P.W., Puri, A., Varaiya, P.: What's decidable about hybrid automata? *J. Comput. Syst. Sci.* **57**, 94–124 (1998)
97. Henzinger, T.A., Raskin, J.-F.: Robust undecidability of timed and hybrid systems. In: Proc. of HSCC 00: Hybrid Systems—Computation and Control. LNCS, vol. 1790, pp. 145–159. Springer, Heidelberg (2000)
98. Henzinger, T.A., Wong-Toi, H.: Using HyTech to synthesize control parameters for a steam boiler. In: Abrial, J.-R., Börger, E., Langmaack, H. (eds.) Formal Methods for Industrial Applications. LNCS, vol. 1165, pp. 265–282. Springer, Heidelberg (1995)
99. Ho, P.-H., Wong-Toi, H.: Automated analysis of an audio control protocol. In: Wolper, P. (ed.) CAV. LNCS, vol. 939, pp. 381–394. Springer, Heidelberg (1995)

100. Johnson, T.T., Mitra, S.: Parametrized verification of distributed cyber-physical systems: an aircraft landing protocol case study. In: ICCPS, pp. 161–170. IEEE, Piscataway (2012)
101. Jula, H., Kosmatopoulos, E.B., Ioannou, P.A.: Collision avoidance analysis for lane changing and merging. PATH Research Report UCB-ITS-PRR-99-13, Institute of Transportation Studies, University of California, Berkeley (1999)
102. Agung Julius, A., Fainekos, G.E., Anand, M., Lee, I., Pappas, G.J.: Robust test generation and coverage for hybrid systems. In: Hybrid Systems: Computation and Control, pp. 329–342. Springer, Heidelberg (2007)
103. Kim, K.-D., Kumar, P.R.: Cyber-physical systems: a perspective at the centennial. Proc. IEEE **100**, 1287–1308 (2012). Centennial-Issue
104. Kouskoulas, Y., Renshaw, D.W., Platzer, A., Kazanzides, P.: Certifying the safe design of a virtual fixture control algorithm for a surgical robot. In: Belta, C., Ivancic, F. (eds.) HSCC, pp. 263–272. ACM, New York (2013)
105. Kühn, W.: Rigorously computed orbits of dynamical systems without the wrapping effect. Computing **61**(1), 47–67 (1998)
106. Kurzhanskiy, A.A., Varaiya, P.: Ellipsoidal toolbox (ET). In: 45th IEEE Conference on Decision and Control, pp. 1498–1503. IEEE, Piscataway (2006)
107. Kurzhanskiy, A.A., Varaiya, P.: Ellipsoidal techniques for reachability analysis of discrete-time linear systems. IEEE Trans. Autom. Control **52**(1), 26–38 (2007)
108. Lafferriere, G., Pappas, G.J., Sastry, S.: O-minimal hybrid systems. Math. Control Signals Syst. **13**(1), 1–21 (2000)
109. Lafferriere, G., Pappas, G.J., Yovine, S.: Symbolic reachability computation for families of linear vector fields. J. Symb. Comput. **32**(3), 231–253 (2001)
110. Le Guernic, C.: Reachability analysis of hybrid systems with linear continuous dynamics. PhD thesis, Université Grenoble I, Joseph Fourier (2009)
111. Lee, E.A., Seshia, S.A.: Introduction to Embedded Systems—A Cyber-Physical Systems Approach (2013). Lulu.com
112. Lerda, F., Kapinski, J., Clarke, E.M., Krogh, B.H.: Verification of supervisory control software using state proximity and merging. In: Hybrid Systems: Computation and Control, pp. 344–357. Springer, Heidelberg (2008)
113. Livadas, C., Lygeros, J., Lynch, N.A.: High-level modeling and analysis of TCAS. Proc. IEEE **88**(7), 926–947 (2000)
114. Loos, S.M., Platzer, A., Nistor, L.: Adaptive cruise control: hybrid, distributed, and now formally verified. In: Butler, M., Schulte, W. (eds.) FM. LNCS, vol. 6664, pp. 42–56. Springer, Heidelberg (2011)
115. Loos, S.M., Witmer, D., Steenkiste, P., Platzer, A.: Efficiency analysis of formally verified adaptive cruise controllers. In: Hegyi, A., De Schutter, B. (eds.) ITSC. Springer, Heidelberg (2013)
116. Lotov, A.V., Bushenkov, V.A., Kamenev, G.K.: Interactive Decision Maps. Applied Optimization, vol. 89. Kluwer Academic, Boston (2004)
117. Lunze, J., Lamnabhi-Lagarrigue, F.: Handbook of Hybrid Systems Control: Theory, Tools, Applications. Cambridge University Press, Cambridge (2009)
118. Lynch, N.A., Segala, R., Vaandrager, F.W.: Hybrid I/O automata. Inf. Comput. **185**(1), 105–157 (2003)
119. Maler, O.: Algorithmic verification of continuous and hybrid systems. In: Int. Workshop on Verification of Infinite-State System (Infinity) (2013)
120. Maler, O., Manna, Z., Pnueli, A.: From timed to hybrid systems. In: de Bakker, J.W., Huizing, C., de Roever, W.P., Rozenberg, G. (eds.) REX Workshop, vol. 600, pp. 447–484. Springer, Heidelberg (1991)
121. Maler, O., Pnueli, A. (eds.): Hybrid Systems: Computation and Control, Proceedings of the 6th International Workshop, HSCC 2003, Prague, Czech Republic, April 3–5, 2003. LNCS, vol. 2623. Springer, Heidelberg (2003)
122. Marwedel, P.: Embedded System Design: Embedded Systems Foundations of Cyber-Physical Systems. Springer, Heidelberg (2010)

123. Matringe, N., Vieira Moura, A., Rebiha, R.: Generating invariants for non-linear hybrid systems by linear algebraic methods. In: Cousot, R., Martel, M. (eds.) SAS. LNCS, vol. 6337, pp. 373–389. Springer, Heidelberg (2010)
124. Miller, J.S.: Decidability and complexity results for timed automata and semi-linear hybrid automata. In: Proc. of HSCC 00: Hybrid Systems—Computation and Control. LNCS, vol. 1790, pp. 296–309. Springer, Heidelberg (2000)
125. Milner, R.: An algebraic definition of simulation between programs. In: Cooper, D.C. (ed.) Proc. of the 2nd Int. Joint Conference on Artificial Intelligence, London, UK, September 1971. pp. 481–489. William Kaufmann, British Computer Society, London (1971)
126. Milner, R.: A Calculus of Communicating Systems. LNCS, vol. 92. Springer, Heidelberg (1980)
127. Mitchell, I.M., Bayen, A.M., Tomlin, C.J.: A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE Trans. Autom. Control* **50**(7), 947–957 (2005)
128. Mitchell, I.M., Templeton, J.A.: A toolbox of Hamilton-Jacobi solvers for analysis of non-deterministic continuous and hybrid systems. In: Morari, M., Thiele, L. (eds.) Hybrid Systems: Computation and Control. LNCS, vol. 3414, pp. 480–494. Springer, Heidelberg (2005)
129. Mitra, S., Wang, Y., Lynch, N.A., Feron, E.: Safety verification of model helicopter controller using hybrid input/output automata. In: Maler and Pnueli [121], pp. 343–358
130. Mitsch, S., Ghorbal, K., Platzer, A.: On provably safe obstacle avoidance for autonomous robotic ground vehicles. In: Robotics: Science and Systems (2013)
131. Nerode, A., Yakhnis, A.: Modelling hybrid systems as games. In: Proceedings of the 31st IEEE Conference on Decision and Control, vol. 3, pp. 2947–2952 (1992)
132. Nerode, A., Kohn, W.: Models for hybrid systems: automata, topologies, controllability, observability. In: Grossman et al. [82], pp. 317–356
133. Neumaier, A.: The wrapping effect, ellipsoid arithmetic, stability and confidence regions. In: Validation Numerics, pp. 175–190. Springer, Heidelberg (1993)
134. Nicollin, X., Olivero, A., Sifakis, J., Yovine, S.: An approach to the description and analysis of hybrid systems. In: Grossman et al. [82], pp. 149–178
135. Platzer, A.: Differential dynamic logic for hybrid systems. *J. Autom. Reason.* **41**(2), 143–189 (2008)
136. Platzer, A.: Differential-algebraic dynamic logic for differential-algebraic programs. *J. Log. Comput.* **20**(1), 309–352 (2010)
137. Platzer, A.: Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics. Springer, Heidelberg (2010)
138. Platzer, A.: Quantified differential invariants. In: Frazzoli and Grosu [65], pp. 63–72
139. Platzer, A.: Stochastic differential dynamic logic for stochastic hybrid programs. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE. LNCS, vol. 6803, pp. 431–445. Springer, Heidelberg (2011)
140. Platzer, A.: A complete axiomatization of quantified differential dynamic logic for distributed hybrid systems. *Log. Methods Comput. Sci.* **8**(4), 1–44 (2012). Special issue for selected papers from CSL’10
141. Platzer, A.: The complete proof theory of hybrid systems. In: LICS [1], pp. 541–550
142. Platzer, A.: A differential operator approach to equational differential invariants. In: Beringer, L., Felty, A. (eds.) ITP. LNCS, vol. 7406, pp. 28–48. Springer, Heidelberg (2012)
143. Platzer, A.: Logics of dynamical systems. In: LICS [1], pp. 13–24
144. Platzer, A.: The structure of differential invariants and differential cut elimination. *Log. Methods Comput. Sci.* **8**(4), 1–38 (2012)
145. Platzer, A.: A complete uniform substitution calculus for differential dynamic logic. *J. Autom. Reason.* **59**(2), 219–265 (2017)
146. Platzer, A.: Logical Foundations of Cyber-Physical Systems. Springer, Heidelberg (2018)
147. Platzer, A., Clarke, E.M.: The image computation problem in hybrid systems model checking. In: Bemporad et al. [26], pp. 473–486

148. Platzer, A., Clarke, E.M.: Computing differential invariants of hybrid systems as fixedpoints. *Form. Methods Syst. Des.* **35**(1), 98–120 (2009)
149. Platzer, A., Clarke, E.M.: Formal verification of curved flight collision avoidance maneuvers: a case study. In: Cavalcanti, A., Dams, D. (eds.) *FM. LNCS*, vol. 5850, pp. 547–562. Springer, Heidelberg (2009)
150. Platzer, A., Quesel, J.-D.: KeYmaera: a hybrid theorem prover for hybrid systems. In: Armando, A., Baumgartner, P., Dowek, G. (eds.) *IJCAR. LNCS*, vol. 5195, pp. 171–178. Springer, Heidelberg (2008)
151. Platzer, A., Quesel, J.-D.: European Train Control System: a case study in formal verification. In: Breitman, K., Cavalcanti, A. (eds.) *ICFEM. LNCS*, vol. 5885, pp. 246–265. Springer, Heidelberg (2009)
152. Prabhakar, P., Viswanathan, M.: A dynamic algorithm for approximate flow computations. In: Frazzoli and Grosu [65], pp. 133–142
153. Prajna, S., Jadbabaie, A.: Safety verification of hybrid systems using barrier certificates. In: Alur and Pappas [16], pp. 477–492
154. Prajna, S., Jadbabaie, A., Pappas, G.J.: A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Trans. Autom. Control* **52**(8), 1415–1429 (2007)
155. Puri, A.: Dynamical properties of timed automata. In: *FTRTFT '98. LNCS*, vol. 1486, pp. 210–227. Springer, Heidelberg (1998)
156. Puri, A., Varaiya, P.: Decidability of hybrid systems with rectangular differential inclusion. In: *Proc. of CAV. LNCS*, vol. 818, pp. 95–104. Springer, Heidelberg (1994)
157. Ratschan, S.: Safety verification of non-linear hybrid systems is quasi-semidecidable. In: *TAMC 2010: 7th Annual Conference on Theory and Applications of Models of Computation. LNCS*, vol. 6108, pp. 397–408. Springer, Heidelberg (2010)
158. Ratschan, S., She, Z.: Safety verification of hybrid systems by constraint propagation-based abstraction refinement. *Trans. Embed. Comput. Syst.* **6**(1), 8 (2007)
159. Sankaranarayanan, S.: Automatic invariant generation for hybrid systems using ideal fixed points. In: Johansson, K.H., Yi, W. (eds.) *HSCC*, pp. 221–230. ACM, New York (2010)
160. Sankaranarayanan, S., Dang, T., Ivančić, F.: Symbolic model checking of hybrid systems using template polyhedra. In: *Tools and Algorithms for the Construction and Analysis of Systems*, pp. 188–202. Springer, Heidelberg (2008)
161. Sankaranarayanan, S., Sipma, H.B., Manna, Z.: Constructing invariants for hybrid systems. *Form. Methods Syst. Des.* **32**(1), 25–55 (2008)
162. Segelken, M.: Abstraction and counterexample-guided construction of ω -automata for model checking of step-discrete linear hybrid models. In: Damm, W., Hermanns, H. (eds.) *CAV. LNCS*, vol. 4590, pp. 433–448. Springer, Heidelberg (2007)
163. Sokolsky, O., Lee, I., Heimdahl, M.P.E.: Challenges in the regulatory approval of medical cyber-physical systems. In: Chakraborty et al. [41], pp. 227–232
164. Stursberg, O., Fehnker, A., Han, Z., Krogh, B.H.: Verification of a cruise control system using counterexample-guided search. *Control Eng. Pract.* **12**(10), 1269–1278 (2004)
165. Tabuada, P.: *Verification and Control of Hybrid Systems: A Symbolic Approach*. Springer, Heidelberg (2009)
166. Tarski, A.: *A Decision Method for Elementary Algebra and Geometry*, 2nd edn. University of California Press, Berkeley (1951)
167. Tavernini, L.: Differential automata and their discrete simulators. *Nonlinear Anal.* **11**(6), 665–683 (1987)
168. Tiwari, A.: Approximate reachability for linear systems. In: Maler and Pnueli [121], pp. 514–525
169. Tiwari, A.: Abstractions for hybrid systems. *Form. Methods Syst. Des.* **32**(1), 57–83 (2008)
170. Tiwari, A.: Generating box invariants. In: Egerstedt, M., Mishra, B. (eds.) *HSCC. LNCS*, vol. 4981, pp. 658–661. Springer, Heidelberg (2008)
171. Tiwari, A.: Logic in software, dynamical and biological systems. In: *LICS*, pp. 9–10. IEEE, Piscataway (2011)

172. Tiwari, A., Shankar, N., Rushby, J.M.: Invisible formal methods for embedded control systems. *Proc. IEEE* **91**(1), 29–39 (2003)
173. Tiwary, H.R.: On the hardness of computing intersection, union and Minkowski sum of polytopes. *Discrete Comput. Geom.* **40**(3), 469–479 (2008)
174. Tomlin, C., Pappas, G., Košecák, J., Lygeros, J., Sastry, S.: Advanced air traffic automation: a case study in distributed decentralized control. In: Siciliano, B., Valavanis, K. (eds.) *Control Problems in Robotics and Automation. Lecture Notes in Control and Information Sciences*, vol. 230, pp. 261–295. Springer, Heidelberg (1998)
175. Tomlin, C., Pappas, G.J., Sastry, S.: Conflict resolution for air traffic management: a study in multi-agent hybrid systems. *IEEE Trans. Autom. Control* **43**(4), 509–521 (1998)
176. Umeno, S., Lynch, N.A.: Proving safety properties of an aircraft landing protocol using I/O automata and the PVS theorem prover: a case study. In: Misra, J., Nipkow, T., Sekerinski, E. (eds.) *FM*, vol. 4085, pp. 64–80. Springer, Heidelberg (2006)
177. Umeno, S., Lynch, N.A.: Safety verification of an aircraft landing protocol: a refinement approach. In: Bemporad et al. [26], pp. 557–572
178. Vaandrager, F.W., van Schuppen, J.H. (eds.) *Hybrid Systems: Computation and Control, Proceedings of the Second International Workshop, HSCC'99, Berg en Dal, The Netherlands, March 29–31, 1999. LNCS*, vol. 1569, Springer, Heidelberg (1999)
179. van Beek, D.A., Man, K.L., Reniers, M.A., Rooda, J.E., Schiffelers, R.R.H.: Syntax and consistent equation semantics of hybrid Chi. *J. Log. Algebraic Program.* **68**(1–2), 129–210 (2006)
180. van Beek, D.A., Reniers, M.A., Schiffelers, R.R.H., Rooda, J.E.: Concrete syntax and semantics of the compositional interchange format for hybrid systems. In: *17th IFAC World Congress* (2008)
181. Wong-Toi, H.: Analysis of slope-parametric rectangular automata. In: *Hybrid Systems. LNCS*, vol. 1567, pp. 390–413. Springer, Heidelberg (1997)
182. Wongpiromsarn, T., Mitra, S., Murray, R.M., Lamperski, A.G.: Periodically controlled hybrid systems. In: Majumdar, R., Tabuada, P. (eds.) *HSCC. LNCS*, vol. 5469, pp. 396–410. Springer, Heidelberg (2009)