Analysis of Hybrid Systems: An ounce of realism can save an infinity of states*

Martin Fränzle**

Department of Computer Science
Carl-von-Ossietzky Universität Oldenburg
P.O. Box 2503, D-26111 Oldenburg, Germany

Abstract. Hybrid automata have been introduced in both control engineering and computer science as a formal model for the dynamics of hybrid discrete-continuous systems. In the case of so-called linear hybrid automata this formalization supports semi-decision procedures for state reachability, yet no decision procedures due to inherent undecidability [4]. Thus, unlike finite or timed automata, already linear hybrid automata are out-of-scope of fully automatic verification.
In this article, we devise a new semi-decision method for safety of linear and polynomial hybrid systems which may only fail on pathological, practically uninteresting cases. These remaining cases are such that their safety depends on the *complete* absence of noise, a situation unlikely to occur in real hybrid systems. Furthermore, we show that if low probability effects of noise are ignored akin to the way they are suppressed in digital modelling then safety becomes fully decidable.

Keywords: Hybrid Systems, Verification, Decision Procedures

1 Introduction

Hybrid systems consist of interacting discrete and continuous components. Most embedded systems belong to this class of systems, as they operate within tightly coupled networks of both types of components. Consequently, integration of discrete and continuous reasoning within a single formal model has recently attracted much interest. The hope is that such combined formalisms may ultimately help in developing real embedded systems.

Among such formalisms, the automata-based ones provide the most immediate prospect for mechanization. Roughly, hybrid automata can be characterized as a *combination of finite automata* whose transitions are triggered by predicates on the continuous plant state *with a description of the evolution of the continuous plant.* The latter consists of a set of real-valued variables that are governed by sets of syntactically restricted differential (in-)equations from which a currently active set is selected depending on the current automaton state. Mechanization

* This article reflects work that has been partially funded by the Deutsche Forschungsgemeinschaft under contract DFG Da 206/5-2.

** Email: `Martin.Fraenzle@Informatik.Uni-Oldenburg.De`,
Phone/Fax: `+49-441-798-2412/2145`.

of hybrid-automaton verification has partly become true, as e.g. the HyTech tool [3] supports among others a semi-decision procedure for state reachability in linear hybrid automata, which can — at least in principle, if complexity does not become prohibitive — be used for effectively *falsifying* safety properties. However, as state reachability has been shown to be undecidable for linear hybrid automata [4], the complementary problem, i.e. *verifying* a safety property by showing that no undesirable state may ever be reached, cannot always be done effectively.

However, it is illustrating to take a closer look at the proof technique used in [4] for showing undecidability of state reachability in linear hybrid automata. The core machinery is an instantiation of the following proof pattern:

> *Effectively encode two-counter machines by hybrid automata, representing the counter values by two continuous variables of bounded range. E.g., by variables of range* $[0, 1]$ *through the embedding* ε *defined as* $\varepsilon(k) = 2^{-k}$.

Although the results thus obtained are formally correct and absolutely well-done, their relevance to the practical design problems hybrid automata are intended to cover is questionable. The encodings used (e.g. ε) encode infinite information, namely the set of natural numbers, within a compact interval of continuous states, whereas the ubiquity of noise limits the information content encodable by any bounded continuous variable encountered in real hybrid systems to a finite value. Hence, on simple information-theoretic grounds, the undecidability results thus obtained can be said to be artefacts of an overly idealized formalization.

However, while this implies that the particular proof pattern sketched above lacks physical interpretation, it does not yield any insight as to whether the state reachability problem for hybrid systems featuring noise is decidable or not. We conjecture that there is a variety of realistic noise models for which the problem is indeed decidable. Within this article, we demonstrate this on a very simple model of noise, which we combine with a pragmatic attitude towards thresholds for noise being considered relevant. In Sect. 3 we devise a new decision method that is able to decide safety for those hybrid automata where safety does not depend on the complete absence of noise, i.e. which, if not unsafe, can tolerate some amount of noise without becoming unsafe. Furthermore, Sect. 4 shows that the aforementioned decision method can cope with arbitrary hybrid systems whenever low-probability effects of noise are neglected. A technical side-condition for both results is that the safety region or the continuous state space of the hybrid automaton be bounded or — which is slightly more general — has strongly finite diameter (cf. Def. 1) wrt. an arbitrary metrics.

Related work. To the best of our knowledge, dynamics of hybrid automata under noise has so far only been analyzed for the subclass of timed automata, where two fundamentally different models of the "robust", noise-resistant behaviour of timed automata have been proposed by Gupta, Henzinger, and Jagadeesan [2] and by Puri [6] (both reports do, however, sketch generalizations to hybrid automata). While in the former line of work, the idealized behavioural model of timed automata is essentially kept and only filtered a posteriori by removing accepted and adding rejected trajectories that are isolated wrt. some topology on trajectories, the latter is more akin to our approach in that it models the impact of noise (i.e. clock drift, if only timed automata are dealt with) as a widening of the transition relation and defines the reachable states as those that are reachable under arbitrarily small positive noise.

2 Hybrid automata

We start our investigation by providing a formalization of hybrid automata. The class of hybrid automata we will deal with goes beyond linear hybrid automata in that we will allow polynomial (instead of only linear) activities and polynomial (instead of linear) predicates for state invariants, transition guards, and transition effects. The reasons for adopting a more general class than usual are twofold. First, all our results can be shown for this more general class without any extra effort. Second, all definitions, statements, and proofs concerning this class of hybrid automata are more compact as no need to keep state invariants separate from activity predicates and transition guards separate from transition effects arises. Instead, every state and transition can be described by just one polynomial predicate, formalised through the first-order logic over the real-closed field, denoted $\mathrm{FOL}(\mathbb{R},+,\times)$ in the remainder.

Therefore, within this article, a *(polynomial) hybrid automaton of dimensionality d* ($d \in \mathbb{N}$) is a six-tuple

$$(\Sigma, \mathbf{x}, (act_\sigma)_{\sigma\in\Sigma}, (trans_{\sigma\to\sigma'})_{\sigma,\sigma'\in\Sigma}, (initial_\sigma)_{\sigma\in\Sigma}, (safe_\sigma)_{\sigma\in\Sigma}) \ ,$$

where Σ is a finite set, representing the discrete states, and $\mathbf{x} = (x_1, \ldots, x_d)$ is a vector of length d of variable names, the continuous variables of the hybrid system.[1] $(act_\sigma)_{\sigma\in\Sigma}$ is a Σ-indexed family of formulae from $\mathrm{FOL}(\mathbb{R},+,\times)$ with free variables $\overleftarrow{\mathbf{x}}, \mathbf{x}$, representing the continuous activities and corresponding state constraints, and $(trans_{\sigma\to\sigma'})_{\sigma,\sigma'\in\Sigma}$ a doubly Σ-indexed family of formulae from $\mathrm{FOL}(\mathbb{R},+,\times)$ with free variables $\overleftarrow{\mathbf{x}}, \mathbf{x}$, representing the discrete transitions and their guarding conditions. Finally, $(initial_\sigma)_{\sigma\in\Sigma}$ and $(safe_\sigma)_{\sigma\in\Sigma}$ are Σ-indexed families of formulae from $\mathrm{FOL}(\mathbb{R},+,\times)$ with free variables $\mathbf{x}$ representing the initial and the safe states of the hybrid automaton.

The interpretation is as follows:

- An *activity predicate* act_σ defines the possible evolution of the continuous state while the system is in discrete state σ. Hooked variable names in the predicate refer to the values of the corresponding system variables beforethe activity, while undecorated variable names refer to the values thereafter. A satisfying valuation $I_{\overleftarrow{\mathbf{x}},\mathbf{x}}$ of its free variables $\overleftarrow{\mathbf{x}}, \mathbf{x}$ is interpreted as: if the system is in state σ and its continuous variables have values $I_{\overleftarrow{\mathbf{x}}}$ then the continuous variables may evolve to $I_{\mathbf{x}}$ while staying in state σ.

[1] Here and in the following, we use the convention to print vectors of variables or constants in boldface. All these vectors have length d.

Note that this single predicate thus generalizes both the state invariants and the activity functions of classical hybrid automata. E.g. the activity predicate

$$\exists\, \Delta t\,.\ \overleftarrow{x} + 3 \times \Delta t \leq x \wedge x < \overleftarrow{x} + 5 \times \Delta t \wedge 4 < x \wedge x \leq 10$$

encodes both the linear activity $\dot{x} \in [3, 5)$ and the state invariant $x \in (4, 10]$.

– A *transition predicate* $trans_{\sigma \to \sigma'}$ defines when the system may evolve from state σ to state σ' by a discrete transition (i.e., it specifies the transition guard) and what the effect of that transition on the continuous variables is. A satisfying valuation $I_{\overleftarrow{\mathbf{x}},\mathbf{x}}$ of its free variables $\overleftarrow{\mathbf{x}}, \mathbf{x}$ is interpreted as: if the system is in state σ and its continuous variables have values $I_{\overleftarrow{\mathbf{x}}}$ then the system may evolve to state σ' with its continuous variables taking the new values $I_{\mathbf{x}}$.
  Assignments as present in hybrid automata are simple to represent in this framework: e.g. for a hybrid system with continuous variables $(x_1, \ldots, x_8)$ the assignment $x_7 := 5$ (where all other continuous variables are left unchanged) is encoded by the predicate $x_7 = 5 \wedge \bigwedge_{i \in \{1,\ldots,6,8\}} x_i = \overleftarrow{x_i}$. Accordingly, guards are simply encoded through predicates over the hooked variables: e.g. the guard $x_3 > 11$ is represented by $\overleftarrow{x_3} > 11$. A transition with guard $x_3 > 11$ and assignment $x_7 := 5$ is thus encoded by the predicate

$$\overleftarrow{x_3} > 11 \wedge x_7 = 5 \wedge \bigwedge_{i \in \{1,\ldots,6,8\}} x_i = \overleftarrow{x_i} \quad .$$

  Multiple different transitions between the same state pair can be represented by disjunction of their encodings.

It should be noted that there is no strict need to distinguish between activities and transitions in $\mathrm{FOL}(\mathbb{R}, +, \times)$. This is just a matter of convenience for the further development.

Dynamic behaviour. During execution, hybrid automata engage in an alternating sequence of discrete transitions and evolution phases, where the continuous variables evolve according to an activity. Hence a (partial) execution containing $n \in \mathbb{N}$ transitions comprises n transitions interspersed between $n+1$ evolution phases, where the final states (wrt. both discrete and continuous state components) of the evolution phases meet the initial states of the following transition and vice versa the final states of the transitions meet the initial states of the following evolution phase. Thus reachability of a final discrete state σ' and a final continuous state $I_{\mathbf{x}}$ from an initial discrete state σ and a initial continuous state $I_{\overleftarrow{\mathbf{x}}}$ through an execution containing n transitions can be formalised through the inductively defined predicate $\phi^n_{\sigma \to \sigma'}$, where

$$\phi^0_{\sigma \to \sigma'} = \begin{cases} \texttt{false}\,, & \text{if } \sigma \neq \sigma' \,, \\ act_\sigma \,, & \text{if } \sigma = \sigma' \,, \end{cases}$$

$$\phi^{n+1}_{\sigma \to \sigma'} = \bigvee_{\tilde{\sigma} \in \Sigma} \exists\, \mathbf{x}_1, \mathbf{x}_2\,. \begin{pmatrix} \phi^n_{\sigma \to \tilde{\sigma}}[\mathbf{x}_1/\mathbf{x}] \wedge \\ trans_{\tilde{\sigma} \to \sigma'}[\mathbf{x}_1, \mathbf{x}_2/\overleftarrow{\mathbf{x}}, \mathbf{x}] \wedge \\ act_{\sigma'}[\mathbf{x}_2/\overleftarrow{\mathbf{x}}] \end{pmatrix}$$

A hybrid automaton is called *safe* iff it may never reach an unsafe state. As usual, an *unsafe state* is considered *reachable* iff there is a partial execution (of arbitrary length) from some initial state to a state not belonging to the set of safe states. An unsafe state is *reachable in at most n steps* iff there is such a partial execution containing at most n discrete transitions.

The latter property can be easily formalised using the formalization of partial run: an *unsafe state is reachable within n steps* iff $\phi^i_{\sigma\to\sigma'} \wedge initial_\sigma[\overleftarrow{\mathbf{x}}/\mathbf{x}] \wedge \neg safe_{\sigma'}$ is satisfiable for some discrete states σ and σ' and some $i \leq n$. This is equivalent to satisfiability of the formula

$$Unsafe_n \stackrel{\text{def}}{=} \bigvee_{\sigma' \in \Sigma} Reach^{\leq n}_{\sigma'} \wedge \neg safe_{\sigma'} \quad ,$$

where

$$Reach^{\leq n}_{\sigma'} \stackrel{\text{def}}{=} \bigvee_{i \in \mathbb{N}_{\leq n}} \bigvee_{\sigma \in \Sigma} \phi^i_{\sigma\to\sigma'} \wedge initial_\sigma[\overleftarrow{\mathbf{x}}/\mathbf{x}] \tag{1}$$

characterizes the continuous states reachable in at most n steps within discrete state σ'. Consequently, an *unsafe state is reachable* iff there is some $n \in \mathbb{N}$ for which $Unsafe_n$ is satisfiable.

Note that $Unsafe_n$ is a formula of $\mathrm{FOL}(\mathbb{R}, +, \times)$ such that reachability of an unsafe state within at most n steps is decidable due to the decidability of $\mathrm{FOL}(\mathbb{R}, +, \times)$ [7]. By successively testing increasing n, this does immediately yield a (well-known) semi-decision procedure for reachability of unsafe states:

Lemma 1 (Semi-decidability of safety). *It is semi-decidable whether a given polynomial hybrid automaton is unsafe.* □

However, this semi-decision procedure does not generalize to a decision procedure: state reachability and thus safety is known to be undecidable even for hybrid automata featuring just two clocks and a single stop-watch [4], where a clock is a continuous variable having constant slope 1 within any activity, and a stop-watch is a continuous variable alternating between slopes 0 and 1 only. Thus, undecidability applies already to hybrid automata with just three continuous variables x_1, x_2, x_3, and activity predicates of only the two forms $(x_1 = \overleftarrow{x_1}) \wedge \bigwedge_{i=\{2,3\}}(x_i = \overleftarrow{x_i} + \Delta t)$ and $\bigwedge_{i=\{1,2,3\}}(x_i = \overleftarrow{x_i} + \Delta t)$.

Of course, there are lucky cases where the set of reachable hybrid states stabilizes finitely, i.e. where $n \in \mathbb{N}$ exists s.t. for each $\sigma' \in \Sigma$, the formula $Reach^{\leq n}_{\sigma'}$ characterizing the continuous state set reachable in at most n steps within discrete state σ' is logically equivalent to its successor $Reach^{\leq n+1}_{\sigma'}$. In such cases, these finite approximations of the reachable hybrid state set cannot only be used for falsifying safety properties (as in the semi-decision procedure of Lemma 1), but also for verifying safety. However, as the undecidability result shows, such stabilization need not occur. Therefore, this analysis, which is supported e.g. by the HyTech tool [3], is only a partial remedy.

3 Noise in hybrid automata

Let us step back for a moment and take a philosophical perspective on the problem just encountered. Unlike finite automata, where the full reach set can always be constructed in a finite number of steps, namely at most the size of the state space, the reach set of hybrid automata need not converge finitely. This suggests that hybrid automata are indeed infinite state systems, where the infinite memory is provided by the continuous state components. However, real hybrid systems are always subject to noise, and thus one might suspect that their continuous components can provide only finite memory. If so, they would in fact be finite state system, with the size of the state space being the product of the size of the discrete state space and the effective size of the continuous state space modulo noise.

But then, the reach set computation of hybrid automata modeling real systems should also converge finitely, just as with finite automata, yielding decidability of state reachability and safety. And, vice versa, any hybrid automaton for which safety cannot be determined finitely would then be an *unrealistic* one which crucially relies on complete absence of noise for realizing an infinite state set. We conjecture that the reach set of such a hybrid automaton is practically uninteresting in that it changes drastically under even the slightest disturbance.

It is the theme of the remainder of this article to make these ideas operational. In particular, we devise a new procedure for determining state reachability in polynomial hybrid automata which may fail only on extremely noise-sensitive borderline cases. Non-termination of this procedure may only occur with hybrid automata that cannot reach the questioned states, yet may reach them under even the slightest disturbance. Thus, safety (in the sense of not reaching an undesirable state) of these automata is practically uninteresting as it crucially depends on complete absence of noise, and therefore is just a fiction.

3.1 A simple model of noise

We begin by formalizing a simple model of noise in hybrid automata. Within this model we assume that noise will leave the discrete state set and transitions unaffected, but will make the evolution phases more nondeterministic. I.e., given a hybrid automaton

$$A = (\Sigma, \mathbf{x}, (act_\sigma)_{\sigma\in\Sigma}, (trans_{\sigma\to\sigma'})_{\sigma,\sigma'\in\Sigma}, (initial_\sigma)_{\sigma\in\Sigma}, (safe_\sigma)_{\sigma\in\Sigma}) \ ,$$

a *disturbed variant* of A is any hybrid automaton

$$\tilde{A} = (\Sigma, \mathbf{x}, (\widetilde{act}_\sigma)_{\sigma\in\tilde{\Sigma}}, (trans_{\sigma\to\sigma'})_{\sigma,\sigma'\in\tilde{\Sigma}}, (initial_\sigma)_{\sigma\in\tilde{\Sigma}}, (safe_\sigma)_{\sigma\in\tilde{\Sigma}})$$

with $act_\sigma \Rightarrow \widetilde{act}_\sigma$ for each $\sigma \in \Sigma$.

Now assume that the continuous state space comes equipped with a first-order definable metrics $dist$, with $dist(\mathbf{x}, \mathbf{y})$ being its definition in $\mathrm{FOL}(\mathbb{R}, +, \times)$.

Furthermore, let ε be a first-order definable constant in $\mathbb{R}_{>0}$. We say that a disturbance $\tilde{A}$ of A is a *disturbance of noise level ε or more* wrt. *dist* iff

$$\exists \mathbf{y} \,.\, \big(act_\sigma[\mathbf{y}/\mathbf{x}] \wedge dist(\mathbf{x},\mathbf{y}) < \varepsilon\big) \Rightarrow \widetilde{act}_\sigma \tag{2}$$

for each $\sigma \in \Sigma$. Thus, evolution phases can drift away within at least an open ball of radius ε under such disturbances. It seems reasonable to assume that within any realistic noise field such disturbances exist for some sufficiently small ε.[2]

Obviously, a disturbance will yield additional possible behaviours, possibly leading to otherwise unreachable states. In particular, the disturbed system yields an overapproximation of the undisturbed system. It comes as a surprise that, under fairly mild extra conditions, such an overapproximation can be constructed within a finite number of steps, as the following lemma shows.

Before we can turn to this central lemma, we have to define the crucial notion of sets of strongly finite diameter.

Definition 1 (Strongly finite diameter). *Let $S \subseteq \mathbb{R}^d$. We say that S has* finite diameter *wrt. the metrics dist iff* $\sup\{dist(x,y) \mid x,y \in S\} < \infty$*. We say that S has* strongly finite diameter *wrt. the metrics dist iff for any $\varepsilon > 0$, each subset $P \subseteq S$ containing only points that have a mutual distance of at least ε is finite, i.e.* $\forall \varepsilon > 0, P \subseteq S \,.\, \big(\forall x,y \in P \,.\, \big(x \neq y \Rightarrow dist(x,y) \geq \varepsilon\big) \Rightarrow |P| < \infty\big)$.

Note that the notions of finite diameter and strongly finite diameter coincide for most of the "natural" metrics on $\mathbb{R}^d$, like Euclidean distance, maximum-norm, and taxi-cab metrics, yet differ for some others, like discrete metrics or the so-called radar-screen metrics.

In the following, let A be a hybrid automaton, $\varepsilon > 0$, and $\tilde{A}$ be a disturbance of A of noise level ε or more. By $\widetilde{Reach}^{\leq n}_{\sigma'}$ we denote the predicate obtained from the defining equation (1) when applied to $\tilde{A}$ instead of A. Thus, $\widetilde{Reach}^{\leq n}_{\sigma'}$ is a formula in $\mathrm{FOL}(\mathbb{R},+,\times)$ formalizing the continuous states reachable by the disturbed automaton $\tilde{A}$ within discrete state σ'. Finally, we denote in the remainder for any $\mathrm{FOL}(\mathbb{R},+,\times)$-formula ϕ by $[\![\phi]\!]_{\mathbf{x}}$ the subset

$$\left\{ (c_1,\ldots,c_d) \in \mathbb{R}^d \,\middle|\, I \oplus [x_1 \mapsto c_1, \ldots, x_d \mapsto c_d] \models \phi \text{ for some valuation } I \right\}$$

of $\mathbb{R}^d$.

[2] Concerning the range of applicability of the model, it is worth noting that for the theory exposed in the remainder, the fact that *every* activity is subject to drift of at least ε for some $\varepsilon > 0$ is not essential. While we have chosen such a model in order to simplify exposition, the theory itself can be easily extended to the more general situation that only every cycle in the discrete state space of the hybrid automaton (i.e. every alternating sequence of activities and transitions going from some discrete state σ via other discrete states back to σ) contains at least one activity with drift of at least ε.

Lemma 2 (Finite overapproximation). *Assume that the reach set* $RS_\sigma = \bigcup_{n\in\mathbb{N}} [\![\widetilde{Reach}_\sigma^{\le n}]\!]_{\mathbf{x}}$ *of* $\tilde{A}$ *has strongly finite diameter wrt. the metrics dist. Then there is* $i \in \mathbb{N}$ *s.t. the dynamics of the undisturbed hybrid automaton A is contracting on* $[\![\widetilde{Reach}_\sigma^{\le i}]\!]_{\mathbf{x}}$*, i.e.*

$$\bigvee_{\sigma'\in\Sigma} \exists \mathbf{x}_1, \mathbf{x}_2 \,.\, \begin{pmatrix} \widetilde{Reach}_{\sigma'}^{\le i}[\mathbf{x}_1/\mathbf{x}] \wedge \\ trans_{\sigma'\to\sigma}[\mathbf{x}_1, \mathbf{x}_2/\overleftarrow{\mathbf{x}}, \mathbf{x}] \wedge \\ act_\sigma[\mathbf{x}_2/\overleftarrow{\mathbf{x}}] \end{pmatrix} \Rightarrow \widetilde{Reach}_\sigma^{\le i} \tag{3}$$

holds for each $\sigma \in \Sigma$*. I.e., there is some* $i \in \mathbb{N}$ *s.t. the state space reachable in the disturbed automaton* $\tilde{A}$ *within i steps is closed under any possible evolution of A.*

Proof. We use contraposition and show that RS_σ does not have strongly finite diameter if (3) is invalid for all $i \in \mathbb{N}$. Therefore, we start from the assumption that $out_i \overset{\text{def}}{=} \neg(3)$ is satisfiable for each $i \in \mathbb{N}$. We will show that this implies existence of an infinite set $P \subset RS_\sigma$ of ε-separated points, which implies that RS_σ does not have strongly finite diameter. An appropriate P is defined as $\bigcup_{i\in\mathbb{N}}\{p_i\}$ with p_i being an arbitrary element of $[\![out_i]\!]_{\mathbf{x}}$. Note that for each $i \in \mathbb{N}$ such a p_i exists due to satisfiability of out_i. It remains to be shown that

(a) $P \subseteq RS_\sigma$, (b) P is infinite, and (c) P contains only ε-separated points.

For both (b) and (c) it suffices to show that $dist(p_i, p_j) \ge \varepsilon$ for $i < j$. The key argument towards this is illustrated in Fig. 1. For a formal proof, we show that $p \in [\![out_i]\!]_{\mathbf{x}}$ implies that $p' \in [\![\widetilde{Reach}_\sigma^{\le i+1}]\!]_{\mathbf{x}}$ for each p' with $dist(p, p') < \varepsilon$. Therefore observe that by definition

$$\begin{aligned}
&[\![\widetilde{Reach}_\sigma^{\le i+1}]\!]_{\mathbf{x}} \\
&= [\![\bigvee_{\sigma'\in\Sigma} \exists \mathbf{x}_1, \mathbf{x}_2 \,.\, \begin{pmatrix} \widetilde{Reach}_{\sigma'}^{\le i}[\mathbf{x}_1/\mathbf{x}] \wedge \\ \widetilde{trans}_{\sigma'\to\sigma}[\mathbf{x}_1, \mathbf{x}_2/\overleftarrow{\mathbf{x}}, \mathbf{x}] \wedge \\ \widetilde{act}_\sigma[\mathbf{x}_2/\overleftarrow{\mathbf{x}}] \end{pmatrix}]\!]_{\mathbf{x}} && [\text{Def. of } \widetilde{Reach}_\sigma^{\le i+1}] \\
&\supseteq [\![\bigvee_{\sigma'\in\Sigma} \exists \mathbf{x}_1, \mathbf{x}_2, \mathbf{y} \,.\, \begin{pmatrix} \widetilde{Reach}_{\sigma'}^{\le i}[\mathbf{x}_1/\mathbf{x}] \wedge \\ trans_{\sigma'\to\sigma}[\mathbf{x}_1, \mathbf{x}_2/\overleftarrow{\mathbf{x}}, \mathbf{x}] \wedge \\ act_\sigma[\mathbf{x}_2, \mathbf{y}/\overleftarrow{\mathbf{x}}, \mathbf{x}] \wedge \\ dist(\mathbf{x}, \mathbf{y}) < \varepsilon \end{pmatrix}]\!]_{\mathbf{x}} && [\text{Property (2) of } \widetilde{act}] \\
&\supseteq [\![\exists \mathbf{y} \,.\, \begin{pmatrix} out_i[\mathbf{y}/\mathbf{x}] \wedge \\ dist(\mathbf{x}, \mathbf{y}) < \varepsilon \end{pmatrix}]\!]_{\mathbf{x}} && [\text{Def. } out_i] \\
&= \left\{ p' \in \mathbb{R}^d \,\middle|\, dist(p, p') < \varepsilon \text{ for some } p \in [\![out_i]\!]_{\mathbf{x}} \right\} \;.
\end{aligned}$$

Hence, $p \in [\![out_i]\!]_{\mathbf{x}}$ and $dist(p, p') < \varepsilon$ implies $p' \in [\![\widetilde{Reach}_\sigma^{\le i+1}]\!]_{\mathbf{x}}$. In particular,

$$dist(p_i, p') < \varepsilon \text{ implies } p' \in [\![\widetilde{Reach}_\sigma^{\le i+1}]\!]_{\mathbf{x}} \;. \tag{4}$$

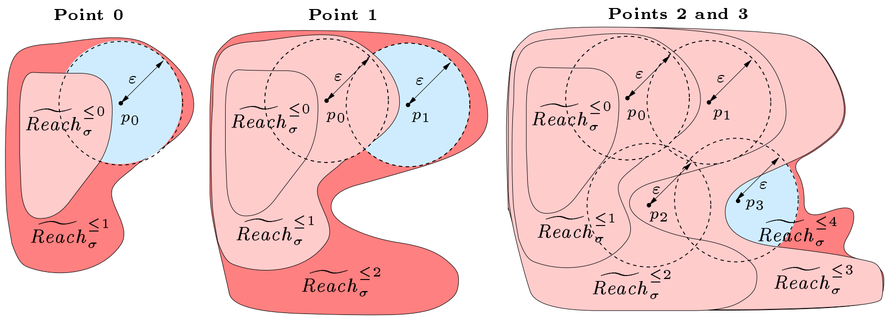

Fig. 1. *Constructing an infinite set of ε-separated points in phase-space.* Points p_i are selected s.t. $p_i \notin [\![\widetilde{Reach}_\sigma^{\leq i}]\!]_{\mathbf{x}}$, yet is reachable from $[\![\widetilde{Reach}_\sigma^{\leq i}]\!]_{\mathbf{x}}$ under a transition followed by an activity of the *undisturbed* automaton A. Consequently, the disturbed automaton $\tilde{A}$ may under the same transition and the corresponding activity reach any point within the ε-ball around p_i s.t. this neighborhood is fully covered by $[\![\widetilde{Reach}_\sigma^{\leq i+1}]\!]_{\mathbf{x}}$. Hence, for $j > i$, point p_j has a minimum distance of ε from p_i.

As $dist(p_i, p_i) = 0$, this shows that $p_i \in [\![\widetilde{Reach}_\sigma^{\leq i+1}]\!]_{\mathbf{x}} \subseteq RS_\sigma$ and thus proves (a). Furthermore, if $i < j$ then $p_j \notin [\![\widetilde{Reach}_\sigma^{\leq i+1}]\!]_{\mathbf{x}}$ can be inferred from the fact that by definition $p_j \in [\![out_j]\!]_{\mathbf{x}}$ and out_j entails $\neg\widetilde{Reach}_\sigma^{\leq j}$, which in turn entails $\neg\widetilde{Reach}_\sigma^{\leq i+1}$ as $i + 1 \leq j$. Therefore, (4) yields $dist(p_i, p_j) \geq \varepsilon$ for $i < j$, which proves (b) and (c). □

Furthermore, it is easy to see that such an $i \in \mathbb{N}$ with the property that the state space reachable in the disturbed automaton $\tilde{A}$ within i steps is closed under any possible dynamic evolution of A can be determined effectively.

Corollary 1. *If the reach set RS_σ of $\tilde{A}$ has strongly finite diameter for each $\sigma \in \Sigma$, then an $i \in \mathbb{N}$ can be* effectively *determined s.t. property (3) holds for each $\sigma \in \Sigma$.*

Proof. For each $i \in \mathbb{N}$, property (3) is a formula of $\mathrm{FOL}(\mathbb{R}, +, \times)$ and thus decidable. Hence, by e.g. μ-recursion, an $i \in \mathbb{N}$ s.t. property (3) holds for each $\sigma \in \Sigma$ can be determined effectively iff there is such an i. However, existence of an $i \in \mathbb{N}$ s.t. property (3) holds is guaranteed by Lemma 2. □

Now assume that we have determined an $i \in \mathbb{N}$ such that the state space reachable in the disturbed automaton $\tilde{A}$ within i steps is closed under any possible evolution of A. As the state space reachable by $\tilde{A}$ within i steps trivially covers the initial state set of $\tilde{A}$ and thus of A, closure of this state space under the possible dynamic evolutions of A implies that this state space covers all states reachable by A. This yields the following corollary:

Corollary 2. *If the reach set RS_σ of $\tilde{A}$ has strongly finite diameter for each $\sigma \in \Sigma$, then an $i \in \mathbb{N}$ can be effectively determined s.t. for each $\sigma \in \Sigma$*

$$\bigcup_{n \in \mathbb{N}} [\![Reach_\sigma^{\leq n}]\!]_{\mathbf{x}} \subseteq [\![\widetilde{Reach}_\sigma^{\leq i}]\!]_{\mathbf{x}} \ .$$

Proof. By Cor. 1 an $i \in \mathbb{N}$ satisfying property (3) can be effectively determined if the reach set RS_σ of $\tilde{A}$ has strongly finite diameter for each $\sigma \in \Sigma$. Now, property (3) for each $\sigma \in \Sigma$ implies that $R \stackrel{\text{def}}{=} \bigcup_{\sigma \in \Sigma} \{\sigma\} \times [\![\widetilde{Reach}_\sigma^{\leq i}]\!]_{\mathbf{x}}$ is closed under the possible evolutions of A. As R contains any initial state of $\tilde{A}$ and thus of A, this implies that R covers all reachable states of A. Therefore, $\bigcup_{n \in \mathbb{N}} [\![Reach_\sigma^{\leq n}]\!]_{\mathbf{x}} \subseteq [\![\widetilde{Reach}_\sigma^{\leq i}]\!]_{\mathbf{x}}$ for each $\sigma \in \Sigma$. □

This implies that if $\tilde{A}$ is safe and its reach set has strongly finite diameter then we obtain within a finite number of iterations — and thus effectively — a witness for safety of A.

Now assume that the safety condition is such that it implies strongly finite diameter of the reachable state set, i.e. that $[\![safe_\sigma]\!]_{\mathbf{x}}$ has strongly finite diameter for each $\sigma \in \Sigma$. This gives rise to the peculiar situation that falsification of $\tilde{A}$ and verification of A become tightly coupled:

Corollary 3. *If safety implies strongly finite diameter, i.e. if $[\![safe_\sigma]\!]_{\mathbf{x}}$ has strongly finite diameter for each $\sigma \in \Sigma$,[3] then either the disturbed automaton $\tilde{A}$ is unsafe (which can be determined effectively) or safety of the undisturbed automaton A can be determined effectively.*

Proof. We distinguish the two cases that either the reach set RS_σ of $\tilde{A}$ has strongly finite diameter for each $\sigma \in \Sigma$, or not. In the first case there is an $i \in \mathbb{N}$ satisfying property (3) for each $\sigma \in \Sigma$ due to Cor. 1. In the second case, the premiss that safety implies strongly finite diameter implies that $\exists\ \overleftarrow{\mathbf{x}}, \mathbf{x} \,.\, \widetilde{Unsafe}_i$ holds for some $i \in \mathbb{N}$. As both (3) and $\exists\ \overleftarrow{\mathbf{x}}, \mathbf{x} \,.\, \widetilde{Unsafe}_i$ are decidable formulae, the minimum $i \in \mathbb{N}$ such that

$$\bigwedge_{\sigma \in \Sigma} (3) \ \vee\ \exists\ \overleftarrow{\mathbf{x}}, \mathbf{x} \,.\, \widetilde{Unsafe}_i \tag{5}$$

can be determined effectively in either case.

Once such an i has been determined, it remains to check the (decidable) property $\exists\ \overleftarrow{\mathbf{x}}, \mathbf{x} \,.\, \widetilde{Unsafe}_i$. If it holds then $\tilde{A}$ is unsafe. Otherwise (5) implies that (3) holds for each $\sigma \in \Sigma$. But then, according to Cor. 2, $\bigcup_{n \in \mathbb{N}} [\![Reach_\sigma^{\leq n}]\!]_{\mathbf{x}} \subseteq [\![\widetilde{Reach}_\sigma^{\leq i}]\!]_{\mathbf{x}}$, which implies

$$\bigcup_{n \in \mathbb{N}} [\![Reach_\sigma^{\leq n}]\!]_{\mathbf{x}} \subseteq [\![safe_\sigma]\!]_{\mathbf{x}} \ ,$$

as $\exists\ \overleftarrow{\mathbf{x}}, \mathbf{x} \,.\, \widetilde{Unsafe}_i$ does not hold. I.e., A has then been shown to be safe. □

An illustration of the verification procedure outlined in above corollary can be found in Fig. 2.

Now, let us take a look at what is a pragmatically reasonable correctness criterion for hybrid systems. It is pragmatically clear that a system should not be called safe if no more than the slightest disturbance is necessary to render it unsafe. This motivates the following definition of robustness.

Definition 2 (Robustness). *A hybrid system is called* fragile *if it is safe, yet any disturbance of arbitrarily small positive noise level is unsafe. All other hybrid systems are called* robust*. I.e., a hybrid automaton A is robust iff it is unsafe or there is an $\varepsilon > 0$ and a safe disturbance $\tilde{A}$ of noise level ε or more.*

[3] Note that for Euclidean distance or equivalent metrics, this condition is equivalent to the safety region being finitely bounded, a case frequently encountered in practice.

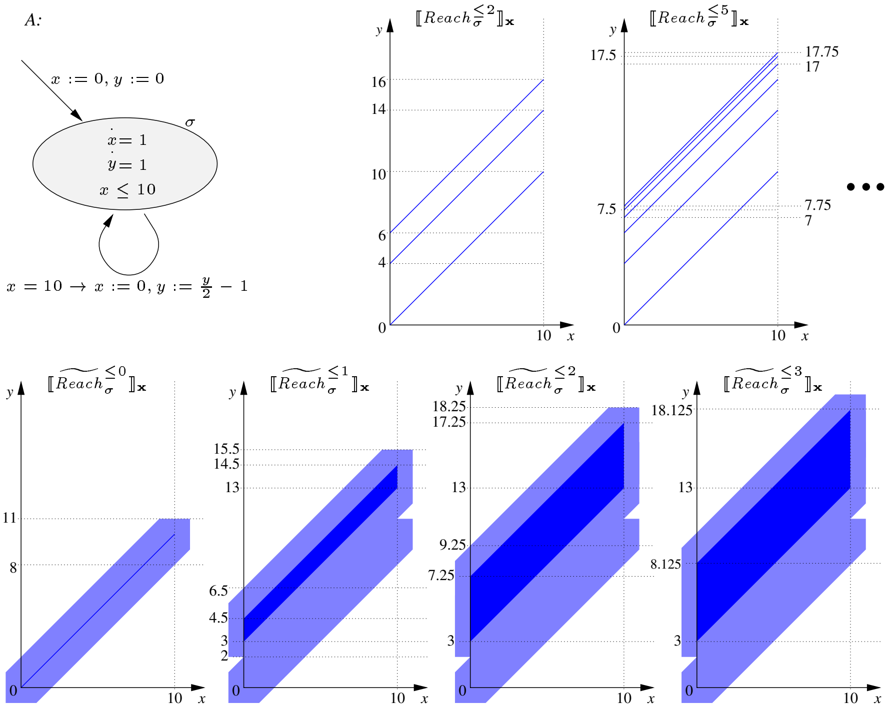

Fig. 2. *Verification using the technique of Cor. 3.* Top left: A hybrid automaton A. Top right: Some steps of its reach set computation. Bottom: Reach set computation for a disturbed variant of noise-level 1 under the max. norm. Shaded parts denote $[\![\widetilde{Reach}_\sigma^{\leq i}]\!]_{\mathbf{x}}$ while solid colour denotes the states reachable by the *undisturbed* automaton from $[\![Reach_\sigma^{\leq i-1}]\!]_{\mathbf{x}}$. Note that reach set computation neither terminates for A nor for its disturbed variant, yet A is contracting on $[\![\widetilde{Reach}_\sigma^{\leq 2}]\!]_{\mathbf{x}}$.

As practical interest clearly is in robust systems only, the following theorem, which states that robust systems can (at least in principle) be automatically verified, is a strong result.

Theorem 1 (Decidability of reachability for robust systems). *If A is a robust hybrid automaton, and if safety implies strongly finite diameter, then it is decidable whether A is safe.*

Proof. A semi-decision procedure for A being unsafe has already been devised in Sect. 2. Hence, it remains to establish a semi-decision procedure for safety of robust automata. This can be done as follows:

1. Select some $n \in \mathbb{N}$.
2. Build the hybrid automaton $\tilde{A} \overset{\text{def}}{=} A_{\frac{1}{n}}$, where for arbitrary $\delta \geq 0$,
$$A_\delta \overset{\text{def}}{=} (\Sigma, I, \mathbf{x}, (\widetilde{act}_\sigma)_{\sigma\in\Sigma}, (trans_{\sigma\to\sigma'})_{\sigma,\sigma'\in\Sigma})$$
with $\widetilde{act}_\sigma \overset{\text{def}}{=} \exists\, \mathbf{y} \,.\, act_\sigma[\mathbf{y}/\mathbf{x}] \wedge dist(\mathbf{x},\mathbf{y}) < \delta$.

3. Using Cor. 3, determine whether $\tilde{A}$ is unsafe or A is safe. In the former case select a strictly greater $n \in \mathbb{N}$ than before and redo from step 2. Otherwise a witness for safety of A has been found. Terminate successfully. □

Now, given the fact that *robust* systems can be automatically verified, it remains the question whether robustness can be automatically detected.

Corollary 4 (Undecidability of robustness). *It is undecidable whether a hybrid automaton is robust or fragile, even when we restrict interest to safety predicates that imply strongly finite diameter. It is, however, semi-decidable in the latter case.*

Proof. Let us restrict attention to safety predicates that imply strongly finite diameter. All fragile systems are by definition safe (cf. Def. 2). Hence, safety of a hybrid automaton is trivially decidable once it is known to be fragile. Likewise, by the extra condition on safety predicates, safety of a hybrid automaton is decidable due to theorem 1 once it is known to be robust. Consequently, decidability of being robust would imply decidability of safety for arbitrary hybrid automata. However, state reachability and thus safety is undecidable even for linear hybrid automata with finitely bounded state space (where safety trivially implies strongly finite diameter wrt. e.g. Euclidean distance) [4], and therefore robustness is undecidable.

It is, however, semi-decidable by a minor variation of the decision procedure of theorem 1: a hybrid automaton A is robust iff it is either unsafe or there is some noise level $\varepsilon > 0$ and some disturbed variant $\tilde{A}$ of noise level ε or more that is safe. However, in the latter case A also has a disturbed variant of positive noise level that is not only safe, but also robust itself, e.g. the automaton $A_{\frac{\varepsilon}{2}}$. Therefore, a hybrid automaton A is robust iff it is either unsafe or there is some noise level $\varepsilon > 0$ such that $A_{\frac{\varepsilon}{2}}$ is robust and safe. The first case is semi-decidable by the semi-decision procedure for being unsafe devised in Sect. 2, while the latter case is semi-decidable by applying the semi-decision procedure for safety of robust automata of Theorem 1 to the automata $A_{\frac{1}{2n}}$ for successively smaller n. □

The undecidability result, while disappointing, should nevertheless be no serious obstacle in practice. It is good engineering practice to make systems tolerant against noise. Thus, any well-engineered system should be robust s.t. the decision procedure of theorem 1 comes to an answer. A simple sufficient criterion for robustness is nevertheless currently unknown. This comes as no surprise, as the same applies even for the related notion of robustness of timed automata [6]. However, robustness of timed automata is decidable.

Even without decidability of robustness, it is in principle possible to compute the exact noise margin that a hybrid automaton can cope with. Here, the noise margin that a hybrid automaton A can cope with is defined as

$$\mathit{Noisemargin}(A) \stackrel{\text{def}}{=} \sup \left\{ \varepsilon > 0 \;\middle|\; \begin{array}{l} A \text{ has a } \textit{safe} \text{ disturbance } \tilde{A} \\ \text{of noise level } \varepsilon \text{ or more} \end{array} \right\} ,$$

with $\sup \emptyset \stackrel{\text{def}}{=} 0$, for the sake of completeness.

Theorem 2 (Computability of robustness margins). *For any hybrid automaton A where safety implies strongly finite diameter, the maximum noise margin $\mathit{Noisemargin}(A)$ that A can cope with can be computed effectively (in the sense of computable real numbers).*

Proof. We use the following definition of computable real numbers: a real number x is computable iff there is an algorithm that given as input a positive rational number denoting the desired precision yields an approximation of x of at least that precision. I.e., given any rational number $\delta > 0$ it yields a rational number $\tilde{x}$ with $x - \delta \leq \tilde{x} \leq x + \delta$. To achieve this, it suffices to have an effective procedure that determines for every pair $a < b$ of rational numbers whether $x \leq b$ or $x \geq a$ in the following sense: if $a \leq x \leq b$ then it may give an arbitrary answer; if however $x < a$ ($x > b$, resp.) then it gives the definitive answer $x \leq b$ ($x \geq a$, resp.).

Such a procedure exists for the noise margin: It suffices to apply Cor. 3 to the two automata A_a and A_b instead of A and $\tilde{A}$. The decision procedure outlined there will either prove A_a to be safe, in which case $Noisemargin(A) \geq a$ follows, or proves A_b to be unsafe, in which case $Noisemargin(A) \leq b$ follows. □

4 Practical modelling issues

Let us finally take a pragmatic attitude towards hybrid system modelling. To this end, we would like to abstract from the impact of noise on the continuous components in a similar way as is generally done for the discrete components in finite-state modelling. The digital model used for modelling digital components as finite state systems is inherently approximative, as low probability deviations — e.g., due to noise — from the ideal digital behaviour are simply neglected. Thereby, no identifiable threshold probability is used for distinguishing effects that are to be modelled from those which may be neglected. Instead, a wide variety of thresholds is accepted in favor of a uniform model, ranging from the extremely low probability of noise-induced error within synchronously clocked subsystems to the far higher probability of error at the interface to an asynchronous environment occurring due to metastable states — known as synchronization failure [1]. This leads to the effect that some unmodelled behavioural aspects may well have higher overall probability than some of the modelled effects.

There is no reason to be more demanding wrt. modelling of the continuous behaviour than wrt. discrete behaviour. I.e., just as with discrete transitions the model-builder may freely include or exclude some of the low-probability effects. However, instead of letting the model-builder do the selection we may as well let the verification algorithm do so itself. I.e., the model builder is asked to devise both a hybrid automaton A excluding all (or most) low-probability effects of noise on the continuous components and a hybrid automaton $\tilde{A}$ including all (or most) low-probability effects. This would apparently lead to $\tilde{A}$ being a disturbance of A of noise level ε or more for some $\varepsilon > 0$. As we do not really care for the low-probability effects, the verification algorithm is then deliberately free to construct an arbitrary intermediate model and to decide its safety. I.e., we insist that our verification procedure gives a positive answer whenever both A and $\tilde{A}$ (and hence all intermediate models) are safe and a negative answer whenever both A and $\tilde{A}$ (and hence all intermediate models) are unsafe. But we will accept any answer if A is safe and $\tilde{A}$ is unsafe.

This might look like a complication, but in fact ideally fits the verification procedures outlined in Sect. 3: if safety implies strongly finite diameter then above proof obligation can be discharged automatically according to Cor. 3. I.e., given such a pragmatic modelling discipline, the safety verification problem for linear or polynomial hybrid automata can be solved fully automatically in the important case of a bounded safety region.

Practitioners may object that practical problems with tools like HyTech [3] are not primarily caused by nontermination of the verification procedure due to the inherent undecidability of state reachability. Often, problems arise due to the extremely large number of iterations needed until the reachable state set stabilizes and due to the numerical precision needed for representing those intermediate state sets. However, the pragmatic modelling discipline sketched here, together with the verification technology of Cor. 3, helps overcome these other problems also. First of all, less applications of the transition relation and the evolution predicates are generally needed for constructing an overapproximation with Cor. 3 than for reaching the fixed point in an incremental calculation of the reachable state space. Second, we may reduce the precision needed in the calculations as we may freely replace $\tilde{A}$ by a slightly less disturbed variant B of A that upon each activity applies some rounding discipline to the corner points of the reachable state sets in order to reduce the necessary numerical precision.

5 Discussion

We have been able to show that safety of linear and polynomial hybrid automata can be decided algorithmically in most practically interesting cases featuring a bounded (or at least strongly finite wrt. some metrics) safety region. The remaining cases are such that their safety depends on the complete absence of noise and, furthermore, apply dissimilar approximation schemes for modelling the discrete and continuous parts of the system, as explained in Sect. 4.

Instrumental to that success has been the insight that a common proof pattern of undecidability results for hybrid-systems formalisms does rely on artefacts of the formalization rather than on an encoding of inherent complexity of the design problem. In a nutshell, those proofs rely on storing infinite information, namely an encoding of the state set of a counter automaton, by continuous variables of bounded range, whereas only a finite part thereof can become effective in any embedded control application due to the ubiquity of noise. We have shown that already with a simplistic model of noise, combined with a pragmatic attitude towards thresholds for noise being considered relevant, the problem of undecidability can be overcome. In this sense, an ounce of realism can save an infinity of states in the analysis of hybrid systems.

However, it may be argued that the model of noise employed in Sect. 3 is too simplistic. Indeed, a realistic model of noise should better be quantitative, representing noise as a probabilistic process. Within this article we have deliberately refrained from this, as such a model would hardly yield any practical verification procedure due to the large computational overhead caused by calculating the density function when iterating the probabilistic transition and evolution steps. I.e., in contrast to the positive effects on the practical complexity of the verification problem that the simplistic model has, as outlined in Sect. 4, such a probabilistic view would certainly yield an intractable model.

However, from a theoretical point of view, it is still interesting what impact on decidability such a detailed modelling of noise has. In a detailed analysis performed by the author it turned out that a probabilistic model of noise has similar impact on hybrid automata as on analog neural nets. In [5], Maass and Sontag analyzed the language acceptance capabilities of analog neural nets that are subject to noise with a wide support, e.g. Gaussian noise, where noise was modelled as a probabilistic process. Under a probabilistic acceptance condition featuring two probability thresholds $a < b$, where a word $w \in A^*$ is rejected if the net reaches an accepting state with probability less than a and w is accepted if the net reaches an accepting state with probability greater b, thus giving a don't-care region between a and b, Maass and Sontag showed that such noisy analog neural nets can only recognize definite languages, i.e. languages of the form $U \cup (A^*.V)$ with finite $U, V \subset A^*$. Under the very same conditions, i.e. noise with a wide support and the acceptance condition featuring a don't-care region, we have shown hybrid automata with finitely bounded continuous state space to recognize[4] regular languages. While this implies that hybrid automata with finitely bounded continuous state space are finite state devices if subject to probabilistic noise with a wide support, this does not immediately imply decidability of e.g. safety, as computability of the noise distribution plays a crucial role.

References

1. T. J. Chaney and C. E. Molnar. Anomalous behavior of synchronizer and arbiter circuits. *IEEE Transactions on Computing*, 22:421–422, 1973.
2. V. Gupta, T. A. Henzinger, and R. Jagadeesan. Robust timed automata. In O. Maler, editor, *Proceedings of the First International Workshop on Hybrid and Real-Time Systems (HART 97)*, volume 1201 of *Lecture Notes in Computer Science*, pages 331–345. Springer-Verlag, 1997.
3. T. A. Henzinger, P.-H. Ho, and H. Wong-Toi. HyTech: The next generation. In *16th Annual IEEE Real-time Systems Symposium (RTSS 1995)*, pages 56–65. IEEE Computer Society Press, 1995.
4. T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What's decidable about hybrid automata. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on the Theory of Computing*, pages 373–382. ACM, 1995.
5. W. Maass and E. D. Sontag. Analog neural nets with Gaussian or other common noise distributions cannot recognize arbitrary regular languages. In *Electronic Colloquium on Computational Complexity, technical reports*. 1997.
6. A. Puri. Dynamical properties of timed automata. In A. P. Ravn and H. Rischel, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'98)*, volume 1486 of *Lecture Notes in Computer Science*, pages 210–227. Springer-Verlag, 1998.
7. A. Tarski. A decision method for elementary algebra and geometry. RAND Corporation, Santa Monica, Calif., 1948.

[4] Under an appropriate extension of the notion of hybrid automaton used within this article, featuring transition labels and a set of accepting states.