Verification of Hybrid Systems: Formalization and Proof Rules in PVS*

Erika Ábrahám-Mumm
Inst. für Informatik und Praktische Mathematik
Christian-Albrechts-Universität zu Kiel
Preußerstraße 1–9
D-24105 Kiel, Germany

Ulrich Hannemann
Computing Science Department
University of Nijmegen
P.O. Box 9010
NL – 6500 GL Nijmegen, The Netherlands

Martin Steffen
Inst. für Informatik und Praktische Mathematik
Christian-Albrechts-Universität zu Kiel
Preußerstraße 1–9, D-24105 Kiel, Germany

E-mail: {eab|ms}@informatik.uni-kiel.de,ulrichh@cs.kun.nl

Abstract

Combining discrete state-machines with continuous behavior, hybrid systems are a well-established mathematical model for discrete systems acting in a continuous environment. As a priori infinite *state systems, their computational properties are undecidable in the general model and the main line of research concentrates on model checking of finite abstractions of restricted subclasses of the general model. In our work, we use* deductive methods*, falling back upon the general-purpose theorem prover PVS.*

To do so we extend the classical approach for the verification of state-based programs by developing an inductive proof method to deal with the parallel composition of hybrid systems. It covers shared variable communication, label-synchronization, and especially the common continuous activities in the parallel composition of hybrid automata. Besides hybrid systems and their parallel composition, we formalized their operational step semantics and a number of proof-rules within PVS, for one of which we give also a rigorous completeness proof. Moreover, the theory is applied to the verification of a number of examples.

Keywords: *hybrid systems, deductive methods, machine-assisted verification.*

*The work was supported by the technology foundation STW, project EIF 3959, "Formal Design of Industrial Safety-Critical Systems" and further by the German Research Council (DFG) within the special program KONDISK under grant LA 1012/5-1.

1 Introduction

Combining discrete state-machines with continuous behavior, hybrid systems [2] have been successfully used to model a large number of applications in areas such as real-time software, embedded systems, and others. Basically, it is a state-based formalism augmented by real-valued variables which may continuously evolve over time. The discrete behavior is given as a labeled transition system, typically in guarded-command notation, allowing shared-variable communication and synchronization over transition labels. The continuous behavior — the activities, as it is called — is typically specified per control-state by differential (in-)equations.

By its continuous part, hybrid automata are a priori *infinite* state systems. Moreover, their computational properties are undecidable in the general model (this is already true for timed-automata, an important subclass). Depending on various restrictions on the form of the invariants, the guards, the activities, etc., a score of variants and simplifications of the general model have been investigated, especially to obtain decidable and automatically checkable subclasses of the general definition (cf. for instance [2] [3] [18] [12] [23]). The main line of research concentrated on model checking of finite abstractions of restricted subclasses of the general model. Besides the drawback of limited expressive power, fully-automatic approaches suffer from the usual state-space explosion problem, when dealing with the parallel composition of subsystems.

Hence in our work, we pursue an alternative route, using

deductive methods and falling back upon a general-purpose theorem prover. To assure rigorous formal reasoning, we employ the interactive theorem prover PVS [22], based on higher-order logic, extensive libraries of data-structures and theories, powerful strategies to assist in routine verification tasks, and modularization facilities.

A classical approach for the verification of state-based programs is that of *inductive assertions:* to prove invariance of a property for all reachable states, it suffices to give an inductive proof, i.e., to prove initial satisfaction and preservation under computational steps. To cope with the verification of parallel systems, it is advantageous to exploit the system's parallel structure (cf. for instance [7] for an extensive monograph on the topic). In the present paper we develop an inductive proof method to deal with the parallel composition of hybrid systems, which we prove to be complete. The method covers the shared variable communication, label-synchronization, and especially the common continuous activities in the parallel composition of hybrid automata. Besides hybrid systems and their parallel composition, we formalize the operational step semantics and a number of proof-rules within PVS, and apply the theory to the verification of a number of examples.

The rest of the paper is organized as follows. We start in Section 2, defining hybrid systems and their transition semantics. Section 3 develops a proof method for verifying safety properties of hybrid systems, based on assertion networks. After defining the parallel composition of hybrid systems (Section 4), we generalize the proof rules in Sections 5. Finally, in Section 6 we conclude with related and future work. The library of PVS-theories formalizing the hybrid system model, together with the proof methods and the examples is available via `http://www.informatik.uni-kiel.de/~eab`.

2 Hybrid systems

Hybrid systems [2] are a well-known formal model for discrete systems acting in a continuous environment. The system's discrete part is represented as a finite set of locations or modes Loc, connected by discrete transitions or edges. The continuous part is given by a finite set Var of variables ranging over the real numbers $\mathbb{R}$. To be able to reason about the parallel composition of hybrid systems, the variable set Var of a hybrid system is a finite subset of a common countably-infinite variable set Var_g. A mapping $\nu : Var \to \mathbb{R}$ of variables to real values is called a *valuation*; the set of all valuations is denoted by V. A location-valuation pair $\sigma = (l, \nu) \in Loc \times V$ constitutes a *state* of a hybrid system. Let $\Sigma = Loc \times V$ denote the set of all states. A state set $Ini \subseteq \Sigma$ characterizes the initial states of the system.

As states consist of a discrete and a continuous part, so do the transitions of a hybrid system. Discrete change of state is captured by the edges of the graph: leading from one location to another, a transition changes the discrete part of the state; besides that, in going from one location to the next, it may alter non-deterministically the values of the variables.

To cater for synchronization between parallel components, the edges come decorated with a synchronization label from a finite label set Lab. The set of labels contains a specific stutter label τ denoting internal moves, not eligible for synchronization. Each location l is assumed to be able to perform a *stutter* transition labelled by τ. Such a transition stands, as usual, for a "do-nothing" step and denotes that other hybrid systems involved in the parallel composition take some discrete transitions. To distinguish between variables the component has under its control in a stutter transition and those it cannot actively influence, the variable set is split into *control* and *non-control* variables. The distinction is drawn per location by a function $Con : Loc \to 2^{Var}$. Stutter transitions leave the valuations for control variables of the given location unchanged, while putting no restriction on the effect concerning the non-control variables, as they are considered as being influenced solely by the outside.

For the continuous part, the values of the variables may evolve over time, where the corresponding behavior is described, per location, by a set of *activities.* An activity is a continuous function, describing the variables' change starting from the moment the location is entered. Since the specific entrance point in time should not influence the behavior relative to that moment, the set of activities for a location is required to be insensitive against shift in time, or time-invariant. Let $\mathcal{F}$ denote the set of all continuous functions in $\mathbb{R}^{\geq 0} \to V$. A set $F \subseteq \mathcal{F}$ of activities is called *time-invariant*, if for all $f \in F$ and $t \in \mathbb{R}^{\geq 0}$, also $f + t \in F$, where $f + t$ denotes the function which assigns to each $t' \in \mathbb{R}^{\geq 0}$ the value $f(t + t')$. An *invariant* finally is attributed to each location, i.e., a predicate over the set of valuations V, where the system is allowed to enter or stay in a location only as long as the invariant evaluates to true.

Before giving the formal definition of a hybrid system, let us fix some notations. We write $f|_{A'} : A' \to B$ for the restriction of a function $f : A \to B$ to a sub-domain $A' \subseteq A$; the same notation is used for the extension of the restiction operator to sets of functions, as well. For $f \in \mathbb{R}^{\geq 0} \to V$ and $x \in Var$, we denote by f^x the function in $\mathbb{R}^{\geq 0} \to \mathbb{R}$ such that $f^x(t) = f(t)(x)$ for all $t \in \mathbb{R}^{\geq 0}$. We call a function $f \in \mathbb{R}^{\geq 0} \to V$ continuous, if for all $x \in Var$, f^x is continuous. The following definition corresponds to the one encoded in PVS; to avoid overly baroque notation, we elide type declarations present in PVS within the definitions in the paper, whenever the type can unambigously be inferred from the context. This convention applies to all the following definitions.

Definition 1 (Hybrid system) *A* hybrid system H *is a tuple* $(Loc, Var, Con, Ini, Lab, Edg, Act, Inv)$*, where* Loc *is a finite, non-empty set of* locations *and* Var *a finite, non-empty set of* variables*. The function* $Con \in Loc \to 2^{Var}$ *defines the* control variables *in each state, the set* $Ini \subseteq \Sigma = Loc \times V$ *the* initial states*. The* transitions *are given by* $Edg \subseteq Loc \times Lab \times (2^V, V \to 2^V) \times Loc$*, where* Lab *denotes a finite set of* labels *containing the stutter label* τ*. For all* $l \in Loc$ *there is a stutter transition* $(l, \tau, (\phi, f), l) \in Edg$ *such that* $\phi = V$ *and* $f(\nu) = \{\nu' \mid \nu|_{Con(l)} = \nu'|_{Con(l)}\}$*. The* activities *are given by* $Act : Loc \to 2^{\mathcal{F}}$ *such that* $Act(l)$ *is time-invariant for each location* $l \in Loc$*. The function* $Inv : Loc \to 2^V$ *specifies the* invariants*.*

For a discrete transition $(l_1, \alpha, (\phi, f), l_2) \in Edg$, $\phi \subseteq V$ is called the *guard* and $f : V \to 2^V$ its *effect*. Depending on various restrictions on the form of the invariants, the guards, the activities etc., a score of variants and simplifications of this model have been investigated, especially to obtain decidable and automatically checkable subclasses of the general definition (cf. for instance [2] [3] [18] [12] [23]). As in this paper we are concerned with formulating a proof method within a deductive framework, we will stick to the general definition.

Let's illustrate the definitions so far on a simple, well-tried example, the thermostat.

Example 2 (Thermostat) *The temperature* x *of a room is controlled by a thermostat, which continuously senses the temperature and turns a heater on and off if the threshold values* x^{min} *and* x^{max} *are reached, where* $x^{min} < x^{max}$ *and* $x^{min}, x^{max} \in \mathbb{R}^{>0}$*. When the heater is off, the temperature decreases according to the function* $x(t) = x_0 e^{-Kt}$*, where* x_0 *is the initial temperature,* t *the time, and* $K \in \mathbb{R}^{>0}$ *a room constant. With the heater turned on, the temperature follows the function* $x(t) = (x_0 - h)e^{-Kt} + h$*, where* $h > x^{min} + x^{max}$ *is a real-valued constant which depends on the power of the heater. The initial temperature is* x^{max} *degrees and the heater is off initially. The two modes of the thermostat are represented by two control locations,* l_{off} *and* l_{on}*. In both locations, the time-invariant activity sets are specified as differential equations. The invariants represent the conditions concerning the extremal temperatures for each location. Together with the guards on the edges between the two locations, these invariants force the system two switch between the two modes iff. the extremal temperatures are reached. Two variables* y *and* z *serve to record the duration of the time spent in the heating and the non-heating mode. The resulting hybrid system is shown in Figure 1. By convention, trivial components of an edge* $(l, \alpha, (\phi, f), l')$*, i.e.,* $\alpha = \tau$*,* $\phi = true$*, or* $f = Id$ *are not shown, and neither are stutter transitions. The same simplification is done for trivial invariants in locations.*

As mentioned before, a system's state can change in two ways: either by time delay or by discrete transitions. Hence there are two kinds of transitions between states: *time steps,* written $\to^{f,t}$, describe the evolution of the values of the variables in a given location and according to an activity in that location. An instantaneous, discrete step, written $\to^{\alpha}$, follows an edge $(l_1, \alpha, (\phi, f), l_2)$ of the system, thereby moving from location l_1 to l_2 and possibly changing the values of the variables according to (ϕ, f). For both relations, control may stay in a location (i.e., time can progress in a location), resp. enter a location through a discrete state change, only if the invariant is not violated. The resulting transition semantics is summarized in the following definition.

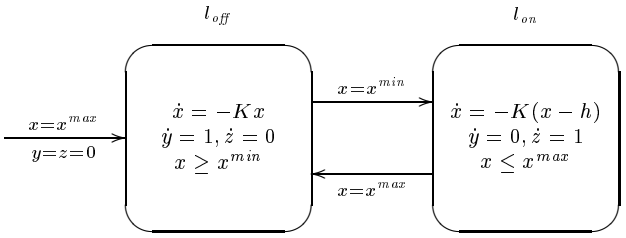

Figure 1. Thermostat

Definition 3 (Step semantics) *Let* $H = (Loc, Var, Con, Ini, Lab, Edg, Act, Inv)$ *be a hybrid system with set of states* $\Sigma = Loc \times V$*. For an activity* f*, a non-negative time delay* t*, and a discrete transition* $(l_1, \alpha, (\phi, f), l_2)$*, the discrete step relation* $\to^{\alpha} \subseteq \Sigma \times \Sigma$ *and the time step relation* $\to^{f,t} \subseteq \Sigma \times \Sigma$ *are given by the following two rules:*

$$\frac{\nu_2 \in f(\nu_1) \quad \nu_1 \in Inv(l_1) \quad \nu_2 \in Inv(l_2) \quad (l_1, \alpha, (\phi, f), l_2) \in Edg \quad \phi(\nu_1) = true}{(l_1, \nu_1) \to^{\alpha} (l_2, \nu_2)}$$

$$\frac{f(0) = \nu_1 \quad f(t) = \nu_2 \quad \forall 0 \leq t' \leq t.\ f(t') \in Inv(l) \quad t \in \mathbb{R}^{\geq 0} \quad f \in Act(l)}{(l, \nu_1) \to^{f,t} (l, \nu_2)}$$

The one-step *relation* $\to$ *is defined by* $\to^{\alpha} \cup \to^{f,t}$*. A* run *of the hybrid system* H *is a (finite or infinite) sequence* $\rho = \sigma_0 \to \sigma_1 \to \sigma_2 \to \cdots$*, with* $\sigma_0 = (l_0, \nu_0) \in Ini$ *and* $\nu_0 \in Inv(l_0)$*. We denote the set of runs of* H *by* $[\![H]\!]$*. A state* $\sigma \in \Sigma$ *is* reachable *in* H*, if there exists a run* $\rho = \sigma_0 \to \sigma_1 \to \sigma_2 \to \cdots \to \sigma_n$ *of* H *with* $\sigma_n = \sigma$*. We write* $R(H)$ *for the set of all reachable states of* H*.*

We use $\to^n$, $\to^*$, and $\to^+$ to denote respectively the n-step relation, the reflexive-transitive closure, and the transitive closure of the one step relation.

The semantics allows a system to evolve by a arbitrarily mixing discrete and continuous steps. It will sometimes be convenient later to assume a more constrained form of trajectories, disallowing time steps of duration zero and disallowing two consecutive time steps. We will call such trajectories to be in *normal form.* Since discrete stutter steps are always possible, for each trajectory there exists an equivalent one in normal form, i.e., $\rightarrow^* = (\rightarrow^\alpha \rightarrow^{f,t>0} \cup \rightarrow^\alpha)^*$.

Before giving an example, let us fix some conventions to specify the components of the hybrid system. The standard way to describe the activities is as solutions of differential equations and differential inclusions. A differential equation is written $\dot{\vec{x}} = g(\vec{x})$, where $\vec{x} = (x_1, \ldots, x_n)$ are variables from the variable set Var of a given hybrid system, and with g a function from V to V. The solution set of a differential equation $\dot{\vec{x}} = g(\vec{x})$ is the set of all functions $f \in \mathcal{F}$ such that for all $i \in \{1, \ldots, n\}$, f^{x_i} is differentiable and $\dot{f}^{x_i}(t) = g(f(t))(x_i)$ for all $t \in \mathbb{R}^{>0}$. Similarly, a differential inclusion is an expression of the form $\dot{\vec{x}} \in g(\vec{x})$, where g is a function from V to 2^V. The solution set of a differential inclusion $\dot{\vec{x}} \in g(\vec{x})$ is the set of all functions $f \in \mathcal{F}$ such that for all $i \in \{1, \ldots, n\}$, f^{x_i} is differentiable and $\dot{f}^{x_i}(t) \in g(f(t))(x_i)$ for all $t \in \mathbb{R}^{>0}$. We will write subsets of valuations V, like the invariants of the locations, in form of boolean predicates $\varphi : V \rightarrow Bool$. Such a predicate φ defines the set of all valuations ν, such that $\varphi(\nu) = true$. In such formulas we write short x for the evaluation $\nu(x)$, where $x \in Var$. In a transition relation (ϕ, f), the non-deterministic change of valuation associated with an edge of the system expressed by f will be written in the form of a simultaneous, non-deterministic assignment $x_1, \ldots, x_n := g_1, \ldots, g_n$, where $x_1, \ldots, x_n \in Var$, and $g_1, \ldots, g_n$ are set-valued functions from V to $2^{\mathbb{R}}$. The relation f is then defined as the set of all valuation pairs $(\nu, \nu') \in V^2$ such that $\nu'(x_i) \in g_i(\nu)$ for all $i = 1, \ldots, n$, and $\nu(y) = \nu'(y)$ for all $y \in Var \setminus \{x_1, \ldots, x_n\}$.

3 Inductive assertional proofs

Our approch and formalization to analyze the behaviour of hybrid systems is based on Floyd's *inductive assertion method* [8]. In this classical state-based verification method one associates an assertion, i.e., a predicate over the current values of variables, with each control location of the underlying program. This gives a finite number of verification conditions to check for proving the given correctness criteria of that program. While originally developed in the context of sequential programs, the inductive assertion method serves also as fundamental technique in the analysis of concurrent programs [7]. We extend the inductive assertion method to hybrid systems.

Let $(Loc, Var, Con, Ini, Lab, Edg, Act, Inv)$ be a hybrid system. An *assertion* on a location l is a boolean predicate over V, or equivalently a subset of V, and an *assertion network* is a subset of the global set $\Sigma = Loc \times V$ of states. For a given assertion network Q of H and a location l, let the assertion $Q_l \subseteq V$ be defined by $Q_l = \{\nu \mid (l, \nu) \in Q\}$, i.e., $\nu \in Q_l$ iff. $(l, \nu) \in Q$. Considering subsets of states as predicates on or properties of the states, we say Q_l holds for a valuation ν, in case $\nu \in Q_l$, and correspondingly for states and assertion networks. By the same token, we will speak of an assertion network implying a property etc. In connection with the system's transition semantics, an assertion network is *invariant,* if it holds for all reachable states, it is called *inductive,* if it holds for the initial states and is preserved under the transition relation.

Definition 4 *An assertion network Q of H is called* inductive, *if*

1. $\{(l, \nu) \in Ini \mid \nu \in Inv(l)\} \subseteq Q$,
2. *if* $\sigma \in Q$ *and* $\sigma \rightarrow^\alpha \sigma'$, *then* $\sigma' \in Q$, *and*
3. *if* $\sigma \in Q$ *and* $\sigma \rightarrow^{f,t} \sigma'$, *then* $\sigma' \in Q$,

for arbitrary states σ and σ' from Σ. We call Q an invariant *of H, if $R(H) \subseteq Q$.*

Obviously, each inductive network is invariant, while the converse will, in general, not hold. Therefore, when interested in verifying a property φ to hold for all reachable states, one can do so by finding a stronger invariant, i.e., an inductive assertion network Q which implies φ. This proof principle, known as *inductive assertion method,* is summarized in the following rule:

$$\frac{Q \rightarrow \varphi \qquad Q \text{ inductive for } H}{R(H) \rightarrow \varphi} \text{ IND}$$

It is standard to show that the rule is sound and complete. We have to refer to the technical report [1] resp. the PVS-code for details. We illustrate the approach, continuing with the thermostat example.

Example 5 (Thermostat) *The thermostat of Example 2 is expected to keep the temperature between x^{min} and x^{max} degrees. Additionally, we show that at any point in time the sum of duration of the heating periods is less or equal than the overall non-heating time, provided that $h > x^{min} + x^{max}$. The duration of the two periods is recorded in the two variables y and z. Both requirements are expressed by the following predicate:*

$$\varphi = x^{min} \leq x \leq x^{max} \wedge y - z \geq 0.$$

We define the network Q as follows:

$$Q_{l_{off}} = x^{min} \leq x \leq x^{max} \wedge y - z \geq -\frac{1}{K} \ln \frac{x}{x^{max}},$$

$$Q_{l_{on}} = x^{min} \leq x \leq x^{max} \wedge y - z \geq -\frac{1}{K} \ln \frac{x^{min}(x^{min} - h)}{x^{max}(x - h)}.$$

Since logarithms and exponentials are not predefined in PVS, we introduced axiomatically a number of algebraic and mathematical facts about these.

4 The parallel composition of hybrid systems

Complex systems are often built from smaller components working in parallel. Before we can start to reason about the composition of systems, we need to introduce their parallel composition [2]. The parallel composition of two hybrid systems H_1 and H_2 is given by a standard product construction and written as $H_1 \times H_2$. Locations are paired and the set of variables combined. The two partners can take a common discrete step, either by synchronizing on the same label, or in that one of the contributors performs a discrete non-synchronizing transition while its partner stutters. Besides synchronizing on the label in a common discrete step, the conjunction of the actions on the variables is taken, i.e., a common step is possible only if both guards are true and if the outcome on the variables coincides. On variables it does not control, a component cannot block non-synchronizing transitions of its partner, since stutter transitions, available at each location, don't restrict the behavior of non-controlled variables. On control variables, on the other hand, stuttering is allowed only without changing the variables' values. Time transitions of the composed systems are time transitions in both systems, i.e., the activities of the composed system, restricted to the local variables, are activities of the component systems. Invariants of the composition finally are conjunctions of the component invariants.

Definition 6 (Parallel composition) *Let H_1 and H_2 be two hybrid systems of the forms $(Loc_i, Var_i, Con_i, Ini_i, Lab_i, Edg_i, Act_i, Inv_i)$. The* product $H_1 \times H_2$ *is the hybrid system $H = (Loc_1 \times Loc_2, Var_1 \cup Var_2, Con, Ini, Lab_1 \cup Lab_2, Edg, Act, Inv)$ such that for all $l_1, l_1' \in Loc_1$, $l_2, l_2' \in Loc_2$, $\alpha \in Lab$, $\phi \subseteq V$, $f : V \to 2^V$ and $g \in \mathcal{F}$:*

1. *$((l_1, l_2), \nu) \in Ini$ iff. $(l_i, \nu|_{Var_i}) \in Ini_i$, for $i = 1, 2$;*
2. *$Con((l_1, l_2)) = Con_1(l_1) \cup Con_2(l_2)$;*
3. *$(l_1, l_2) \longrightarrow^{\alpha}_{(\phi, f)} (l_1', l_2') \in Edg$, iff. there exist $l_i \longrightarrow^{\alpha_i}_{(\phi_i, f_i)} l_i' \in Edg_i$, such that*
   (a) *$\alpha = \alpha_1 = \alpha_2$, or $\alpha_1 = \alpha \notin Lab_2$ and $\alpha_2 = \tau$, or $\alpha_1 = \tau$ and $\alpha = \alpha_2 \notin Lab_1$,*
   (b) *$\phi(\nu) = \phi_1(\nu|_{Var_1}) \wedge \phi_2(\nu|_{Var_2})$, and*
   (c) *$\nu' \in f(\nu)$, iff. $\nu'|_{Var_1} \in f_1(\nu|_{Var_1})$ and $\nu'|_{Var_2} \in f_2(\nu|_{Var_2})$;*
4. *$g \in Act((l_1, l_2))$, iff. for both $i = 1$ and $i = 2$, there exist $g_i \in Act_i(l_i)$, such that $\lambda t.g(t)|_{Var_i} = g_i$;*
5. *$Inv((l_1, l_2))|_{Var_i} = Inv_i(l_i)$ for $i = 1, 2$.*

Note that by construction the set of activities $Act((l_1, l_2))$ for a composed location is time invariant, since $Act_1(l_1)$ and $Act_2(l_2)$ are. It is routine albeit tedious to show that parallel composition is associative and commutative. For a parallel composition $H_1 \times \ldots \times H_n$ with $n > 0$ and $j \in \{1, \ldots, n\}$, we call the composition system without H_j the *context* of H_j. Let Σ_H denote the state space of H. For the product system $H = H_1 \times H_2$, and a state $\sigma = ((l_1, l_2), \nu)$ of H, we write $\sigma\downarrow_{H_1} = (l_1, \nu|_{Var_1})$ and $\sigma\downarrow_{H_2} = (l_2, \nu|_{Var_2})$ for the projections on the respective components; we will use the same notation for sets of states. For a run $\rho = \sigma_0 \to_0 \sigma_1 \to_1 \sigma_2 \to_2 \cdots$ of the product system H, the projection $\rho\downarrow_{H_1}$ is the sequence $\sigma_0' \to_0' \sigma_1' \to_1' \sigma_2' \to_2' \cdots$, where for all $i = 0, 1, 2, \ldots$, $\sigma_i' = \sigma_i\downarrow_{H_1}$, and if $\to_i$ is a time-step of the form $\to^{f_i, t_i}$, then $\to_i'$ is a time-step $\to^{f_i', t_i}$ where f_i' is defined by $f_i'(t) = f_i(t)|_{Var_1}$ for all $t \in \mathbb{R}^{\geq 0}$, and if $\to_i$ is a discrete-step $\to^{\alpha_i}$ then $\to_i'$ is also a discrete-step $\to^{\alpha_i'}$ where $\alpha_i' = \alpha_i$ if $\alpha_i \in Lab_1$, and $\alpha_i' = \tau$ otherwise. The projection $\rho\downarrow_{H_2}$ is defined analogous.

A basic property of the product system is that all runs of the product projected to one of the component systems are runs of that component system:

Lemma 7 *Let $H = H_1 \times H_2$ be the parallel composition of two hybrid systems H_1 and H_2. Then $[\![H]\!]\downarrow_{H_i} \subseteq [\![H_i]\!]$ and $R(H)\downarrow_{H_i} \subseteq R(H_i)$, for $i = 1, 2$.*

Example 8 (Thermostat revisited) *We extend Example 2, considering a finite number of heating systems $H_1, \ldots, H_n$ running in parallel and sharing a common fuel tank. An additional variable u_i for each thermostat records its fuel consumption; it is a control variable in both the on- and the off-locations, with values $u_i = 0$ in location l_{off} and $u_i = c_i$ in l_{on}, and where the constant c_i represents the consumption per time unit. We abbreviate the system's overall fuel consumption $\sum_{i=1}^{n} u_i$ by the variable u. Heating is possible only as long as the tank is not empty and a heater can be turned on only if there's a certain amount of fuel in the tank, represented by the threshold value $\epsilon \in \mathbb{R}^{>0}$. The discrete transitions and the invariants from Example 2 are modified accordingly.*

The fuel tank H_{fuel} is emptied by the heaters in their on-mode, and can be refilled with a constant rate d to keep

the filling level v between Vol_1 and Vol_2 with $\mathrm{Vol}_1, \mathrm{Vol}_2 \in \mathbb{R}^{>0}$ and $\epsilon < \mathrm{Vol}_1 < \mathrm{Vol}_2$. Initially, the tank is full at level Vol_2. We assume that the filling rate d exceeds the sum $\sum_{i=1}^{n} c_i$ of the outflows in case of maximal heating, i.e., when all heaters are on. The resulting composed system is shown in Figure 2.

5 Inductive assertional proofs for parallel composition

A first approach to the verification of the parallel composition of hybrid systems is immediately given by the observation that the parallel composition of two hybrid systems as defined in Definition 6 results in a hybrid system for which one can construct an inductive assertion network again. The number of resulting verification obligations turns out to reflect the state-explosion problem.

The basic idea for an improvement over the plain product of assertions for the classical programming concepts (cf. to [7] for an exhaustive treatment) is a two-level approach, where first *local* assertion networks are checked for local consistency, and then some global consistency test relates these local networks, reducing the amount of verification conditions. In the sequel, we improve the inductive assertion method for hybrid systems following those very ideas. An important technique in this extension of the inductive method is the introduction of new, otherwise unused variables, which allow to speak about the peer processes of a component. These variable are commonly called *auxiliary variables*. The extension of a system by introducing (finitely many) auxiliary variables is called *augmentation.* Auxiliary variables allow the formulation of a sound and complete proof rule for the parallel composition of hybrid systems which we apply to the above example.

5.1 Augmentation

As auxiliary variables are added for the sole purpose of verification, their addition must not influence the system's behaviour in any way, i.e., the unaugmented part of the system must be left untouched. The additional variables may be control variables (per location) but the augmentation must not change the "control"-status of the old variables. Furthermore, projected onto the domain belonging to the original set of variables, the state space and the set of initial states remains unchanged (the set of locations is not altered, anyway). As for the discrete transition, each one in the original system is mirrored in the augmented one, and vice versa. Especially, the transition relation for each edge must result in the same change on the old variables and conversely as well: each transition from the original system must be reflected by some of the augmented one. Analogously, augmenting the system with auxiliary variables must not constrain its continuous behavior: each activity on the original variables must have a counterpart in the extended system and vice versa. To be compliant with the definition of a hybrid system, the new set of activities in each location must be closed under time-shift, i.e., be time-invariant. The invariants, finally, must not restrain the new variables at all.

Definition 9 (Augmented hybrid system) *Given a set of global variables Var_g, a hybrid system $H = (Loc, Var, Con, Ini, Lab, Edg, Act, Inv)$, and a set of fresh auxiliary variables $Var_{aux} \subseteq Var_g$, i.e., $Var_{aux} \cap Var = \emptyset$. We call H' an* augmentation *of H with Var_{aux}, if $H' = (Loc, Var', Con', Ini', Lab, Edg', Act', Inv')$ with $Var' = Var \dot\cup Var_{aux}$ and if the following conditions hold:*

1. *$Con(l) \subseteq Con'(l)$ and $Con'(l) \setminus Con(l) \subseteq Var_{aux}$, for all locations $l \in Loc$;*
2. *$Ini = \{(l, \nu) \mid \exists (l, \nu') \in Ini'.\ \nu = \nu'|_{Var}\}$;*
3. *$Inv'(l) = \{\nu' \mid \nu'|_{Var} \in Inv(l)\}$, for all locations $l \in Loc$;*
4. *(a) for all $(l_1, a, (\phi, f), l_2) \in Edg$ there exists a $(l_1, a, (\phi', f'), l_2) \in Edg'$ such that for all $\nu' \in V'$: $\phi(\nu'|_{Var})$ implies $\phi'(\nu')$ as well as $f(\nu'|_{Var}) \subseteq f'(\nu')|_{Var}$;*
   (b) for all $(l_1, a, (\phi', f'), l_2) \in Edg'$ exists $(l_1, a, (\phi, f), l_2) \in Edg$ such that $\phi'|_{Var} \subseteq \phi$, and for all valuations $\nu' \in V'$, $f'(\nu')|_{Var} \subseteq f(\nu'|_{Var})$;
   (c) for all locations $l \in Loc$ there is a stutter transition $(l, \tau, (\phi', f'), l) \in Edg'$ such that $\phi' = V'$ and $f = \lambda\nu' : V'.\lambda\hat\nu' : V'.\hat\nu'|_{Con'(l)} = \nu'|_{Con'(l)}$;
5. *for all locations $l \in Loc$*
   (a) $Act'(l)$ is time-invariant;
   (b) for all $f \in Act(l)$ and $\nu' \in V'$ with $\nu'|_{Var} = f(0)$, there exists $f' \in Act'(l)$ such that $f = \lambda t.f'(t)|_{Var}$;
   (c) for all $f' \in Act'(l)$, there exists $f \in Act(l)$ such that $f = \lambda t.f'(t)|_{Var}$.

In the following, we will write $H' \geq H$, when H' is an augmentation of H. As the control flow and the activities for variables of H are not influenced by the auxiliary variables, the set of reachable states of H' restricted to Var in the valuation component equals the reachable states of the original system.

Lemma 10 *Let H be a hybrid system with variable set Var and H' an augmentation of H by auxiliary variables Var_{aux}. Then $R(H')\downarrow_{Var} = R(H)$.*

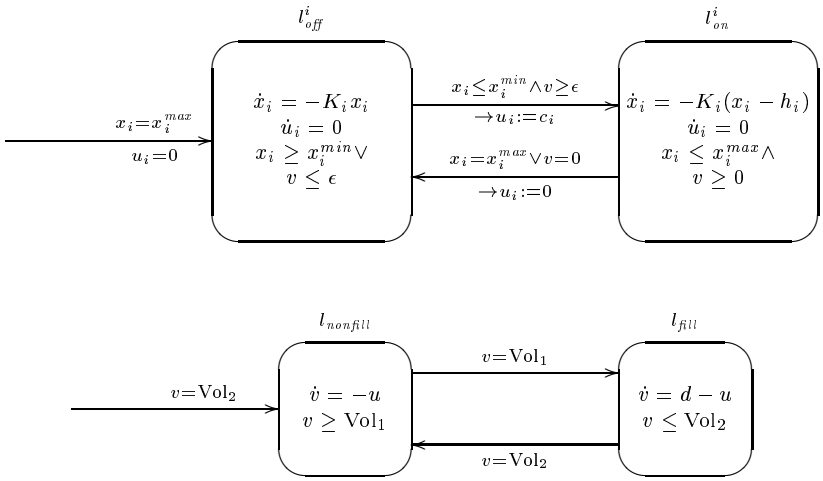

Figure 2. Thermostats with common fuel tank: $H_1 \times \ldots H_n \times H_{fuel}$

Thus, a property whose satisfaction does not depend on the values for the auxiliary variables, holds for all reachable states of H', iff. it holds for all reachable states of H.

5.2 A complete proof method

Two hybrid systems running parallel are related by the time. We allow to refer to the global time in the assertions as a variable t. Referring to the global time is equivalent to having a common (auxiliary) variable in all hybrid systems running in parallel, whose initial value is 0, whose derivative is always 1, and whose is not affected by any transition. Allowing to refer to t as a pre-defined common auxiliary variable simplifies the verification.

Definition 11 *Let H_1 and H_2 be hybrid systems, $H = H_1 \times H_2$ their parallel composition, and Q_1 and Q_2 assertion networks for H_1 and H_2, respectively. We define the composition of the local assertion networks as $Q_1 \times Q_2 = \{\sigma \in \Sigma_H | \sigma\downarrow_{H_1} \in Q_1 \wedge \sigma\downarrow_{H_2} \in Q_2\}$.*

Note that $Q_1 \times Q_2$ is an assertion network of H. So let H_1 and H_2 be two hybrid systems, $H = H_1 \times H_2$ with variable set Var their parallel composition, and $\varphi \subseteq \Sigma_H$ a predicate on the set of H's states. Then φ is an invariant of H if and only if there exists an auxiliary variable set Var_{aux}, hybrid systems H'_1 and H'_2, such that $H' = H'_1 \times H'_2$ is an augmentation of H with Var_{aux}, and inductive assertion networks Q'_1 and Q'_2 of H'_1 and H'_2, respectively, such that $Q = \{(l,\nu) \in \Sigma_H \mid \exists (l,\nu') \in Q'_1 \times Q'_2.\ \nu = \nu'|_{Var}\} \subseteq \varphi$. With these conventions, we can formulate the proof rule to deal with the parallel composition of systems.

$$\frac{\begin{array}{c} H'_1 \times H'_2 \geq H_1 \times H_2 \\ Q'_1 \textit{ inductive for } H'_1 \quad Q'_2 \textit{ inductive for } H'_2 \\ \forall (l,\nu') : \Sigma_{H'}.\ Q'_1 \times Q'_2(l,\nu') \rightarrow \varphi(l,\nu'|_{Var}) \end{array}}{R(H) \rightarrow \varphi}\ \textsc{Comp}$$

Proposition 12 *The proof rule* (COMP) *is sound and complete.*

Proof sketch: Soundness is shown with the help of Lemma 7 and Lemma 10 by a straightforward inductive argument. The soundness proof is formalized in PVS.

The proof of (semantical) completeness hinges on the fact that for each trajectory there exists an equivalent one in normal form (cf. Section 2). We introduce three auxiliary variables $Var'_i = Var_i \dot{\cup} \{h_t, h_\rho, h_f\}$ to H_1 and H_2, where h_t measures the elapsed time, the history variable h_ρ records the sequence of past pairs of locations together with the common sojourn times of the continuous part of the transition steps, and h_f is used to enforce the existence of common continuous behaviour. These auxiliary variables allow to construct H'_1 and H'_2 forming a canonical augmentation of $H_1 \times H_2$ which satisfies the requirements of (COMP). Then the key step for completeness is to show that the history variables record enough information to ensure that each state reachable separately in H'_1 and H'_2 is commonly reachable, as well, i.e.,

$$\begin{array}{c} \text{if } \sigma'\downarrow_{H'_1} \in R(H'_1) \text{ and } \sigma'\downarrow_{H'_2} \in R(H'_2), \\ \text{then } \sigma' \in R(H'_1 \times H'_2). \end{array}$$

For details we refer to the full technical report [1]. □

5.3 Examples

Besides formalizing the proof rules in PVS, we applied the method to a number of examples, e.g., non-linear variations of the water level monitor [2], or a modified clock synchronization of the MPEG4 standard.

Example 13 (Thermostat) *For the extended thermostat system of Example 8, we want to verify that the temperature of all rooms stays between the specified limits, and that the tank never gets below the minimum level needed for turning a heater on, i.e.,*

$$\varphi = \epsilon \leq v \wedge \bigwedge_i x_i^{min} \leq x_i \leq x_i^{max}$$

is estabilied invariant of $H_1 \times \ldots \times H_n \times H_{fuel}$ (cf. Figure 2). Using the proof method from above, $Q_i^{cons} = u_i \leq c_i$ can be established as invariant of H_i, and thus the parallel composition H satisfies the conjunction $\bigwedge_i Q_i^{cons}$. Furthermore, under the assumption that the context of H_{fuel} satisfies the invariant $u \leq \sum_i c_i$, the fuel system H_{fuel} enjoys $Q_{fuel} = \mathrm{Vol}_1 \leq v$ as invariant. Finally, assuming the context of H_i satisfies the invariant Q_{fuel}, the property $Q_i^{temp} = x_i^{min} \leq x_i \leq x_i^{max}$ is an invariant of H_i. Since $Q_{fuel} \wedge \bigwedge_i Q_i^{temp}$ implies φ, the desired property φ holds for the composed system $H_1 \times \ldots \times H_n \times H_{fuel}$.

The PVS formalization of these examples and the verified properties is available on the web-site and in [1].

6 Conclusion

In this paper we present an assertional, deductive proof method for the verification of hybrid systems. For the verification of composed systems, we give a complete proof rule to reduce the complexity introduced by the parallel composition. To facilitate the tedious verification of hybrid systems without restricting the model artificially, we embedded the proof system into the PVS theorem prover.

Beside offering the full power of higher-order logic, a further advantage of such a deductive verification environment is that it allows a straightforward rigorous formalization of the mathematical definitions, without the need to resort to any specific logic. Furthermore, PVS comes equipped with a wide range of automated proof-strategies and heuristics. Also its modularization facilities are helpful in structuring the work.

Of course, to perform the verification requires some expertise in hybrid systems, PVS, and the nature of inductive proofs, but the goal-directed manner of applying the proof-methods and rules gives valuable guidance. Sometimes, the theorem prover surprises the user by taking a long time for seemingly routing task (for instance, type-checking the collected theorems). In those cases, rephrasing the formalization into a mathematically equivalent one but using different data representations (e.g. using lists instead of sets) can cut down the times noticeably. The examples we analyzed in PVS demonstrate how to apply our mathematical framework in praxis. Despite their relative small size, they point out the relevance of our method for the verification of *arbitrary* hybrid systems (for instance hybrid systems with non-linear properties or parameterized hybrid systems).

As the main line of research on hybrid systems focuses on model checking techniques for appropriately *restricted* subclasses, there are less investigations on deductive methods for their verification. Closest in spirit to our work is [5], which embed timed automata into PVS and apply their approach on the steam boiler example. The same example is treated in [25], with the goal of deriving an implementation of a real-time program in a number of refinement steps [16]. The PVS theorem prover is also used in [14] in combination with model checking using HyTech [4] for the reachability analysis for various classes of linear hybrid automata. For the verification of safety properties of hybrid systems, [17] employ hybrid temporal logic HTL, an extension of interval temporal logic. They give a number of proof-rules which they prove sound. Likewise building upon temporal logic, [21] use the Stanford theorem prover STeP as proof environment. See [19] for an overview over deductive and algorithmic approaches for verification of hybrid systems.

For future work, we intend to apply our method to larger case studies, especially to extend the control example based on MPEG4 of [6], and further a laser steering system for mass spectroscopy. For the verification of complex case studies the addition of a tool which allows for a graphical specification of hybrid systems is essential, as is a parser that automatically generates the PVS verification conditions for a given specification. This also allows for a clean separation between type checking (i.e., checking that the specified system is indeed a hybrid system) and the actual verification task. Moreover, the representation of the verification conditions for each control mode and each discrete transition as separate lemmas improves the structure of the proofs and increases the level of automation. To improve the specification structure of hybrid systems, the interface information can be extended, for instance separating the variable set into input and output variables as in [20]. Such a cleaner separation is a necessary prerequisite for the development of an assume-guarantee reasoning scheme (cf. [24] or in the context of hybrid systems [13, 9]). Especially we expect that the verification will benefit from an alternative semantics allowing for compositional proofs [11].

References

[1] E. Ábrahám-Mumm, U. Hannemann, and M. Steffen. Verification of hybrid systems: Formalization and proof rules in PVS. Technical Report TR-ST-01-1, Lehrstuhl für Software-Technologie, Institut für Informatik und praktische Mathematik, Christian-Albrechts-Universität Kiel, Jan. 2001.

[2] R. Alur, C. Courcoubetis, T. Henzinger, P. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138:3–34, 1995. A preliminary version appeared in the proceedings of 11th. International Conference on Analysis and Optimization of Systems: Discrete Event Systems (LNCI 199).

[3] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:252–235, 1994.

[4] R. Alur, T. A. Henzinger, and P. Ho. Automatic symbolic verification of embedded systems. In *Proc. 14th Annual Real-Time Systems Symposium*, pages 2–11. IEEE Computer Society Press, 1993.

[5] M. Archer and C. Heitmeyer. Verifying hybrid systems modeled as timed automata: A case sudy. In O. Maler, editor, *Hybrid and Real-Time Systems (HART'97)*, number 1201 in Lecture Notes in Computer Science, pages 171–185. Springer-Verlag, 1997.

[6] J. B. de Meer and E. Ábrahám-Mumm. Formal methods for reflective system specification. In J. Grabowski and S. Heymer, editors, *Formale Beschreibungstechniken für verteilte Systeme*, pages 51–57. Universität Lübeck/Shaker Verlag, Aachen, Juni 2000 2000.

[7] W.-P. de Roever, F. de Boer, U. Hannemann, J. Hooman, Y. Lakhnech, M. Poel, and J. Zwiers. *Concurrency Verification: Introduction to Compositional and Noncompositional Proof Methods*. Cambridge University Press, 2001. to appear.

[8] R. W. Floyd. Assigning meanings to programs. In J. T. Schwartz, editor, *Proc. Symp. in Applied Mathematics*, volume 19, pages 19–32, 1967.

[9] G. F. Frehse, O. Stursberg, S. Engel, R. Huuck, and B. Lukoschus. Modular analysis of discrete controllers for distributed hybrid systems. 2001. Technical Report, submitted for publication.

[10] R. L. Grossman, A. Nerode, A. P. Ravn, and H. Rischel, editors. *Hybrid Systems*, volume 736 of *Lecture Notes in Computer Science*. Springer-Verlag, 1993.

[11] U. Hannemann. *Semantic Analysis of Compositional Proof Methods for Concurrency*. PhD thesis, Utrecht Universiteit, Oct. 2000.

[12] T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What's decidable about hybrid automata. In *27th Annual ACM Symposium on Theory of Computing*. ACM Press, 1995.

[13] T. A. Henzinger, M. Minea, and V. Prabhu. Assume-guarantee reasoning for hierarchical hybrid systems. In M. D. D. Benedetto and A. Sangiovanni-Vincentelli, editors, *Proceedings of the 4th International Wokrshop on Hybrid Systems: Computation and Control (HSCC 2001), Rome*, volume 2034 of *Lecture Notes in Computer Science*, pages 275–290. Springer-Verlag, 2001.

[14] T. A. Henzinger and V. Rusu. Reachability verification for hybrid automata. In Henzinger and Sastry [15], pages 190–204.

[15] T. A. Henzinger and S. Sastry, editors. *Proceedings of the First International Workshop on Hybrid Systems: Computation and Control (HSCC'98)*, volume 1386 of *Lecture Notes in Computer Science*. Springer, 1998.

[16] J. Hooman. A compositional approach to the design of hybrid systems. In Grossman et al. [10], pages 121–148.

[17] A. Kapur, T. A. Henzinger, Z. Manna, and A. Pnueli. Proving safety properties of hybrid systems. In H. Langmaack, W.-P. de Roever, and J. Vytopil, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems 1994*, volume 863 of *Lecture Notes in Computer Science*, Kiel, Germany, September 1994. Springer-Verlag. 3rd International School and Symposium.

[18] Y. Kesten, A. Pnueli, J. Sifakis, and S. Yovine. Integration graphs: a class of decidable hybrid systems. In Grossman et al. [10], pages 179–208.

[19] S. Kowalewski, P. Herrmann, S. Engell, H. Krumm, H. Treseler, Y. Lakhnech, R. Huuck, and B. Lukoschus. Approaches to the formal verification of hybrid systems. *at-Automatisierungstechnik. Special Issue: Hybrid Systems II: Analysis, Modeling, and Verification*, 49(2):66–74, Feb. 2001.

[20] N. A. Lynch, R. Segala, and F. W. Vaandrager. Hybrid I/O automata revisited. In *Proceedings of HSCC'01*, 2001. to appear.

[21] Z. Manna and H. B. Sipma. Deductive verification of hybrid systems using STeP. In Henzinger and Sastry [15].

[22] S. Owre, J. M. Rushby, and N. Shankar. PVS: A prototype verification system. In D. Kapur, editor, *Automated Deduction (CADE-11)*, volume 607 of *Lecture Notes in Computer Science*, pages 748–752. Springer-Verlag, 1992.

[23] O. Roux and V. Rusu. Uniformity for the decidability of hybrid automata. In R. Cousot and D. A. Schmidt, editors, *Proceedings of SAS '96*, volume 1145 of *Lecture Notes in Computer Science*, pages 301–316. Springer-Verlag, 1996.

[24] N. Shankar. Lazy compositional verification. In W.-P. de Roever, H. Langmaack, and A. Pnueli, editors, *Compositionality: The Significant Difference (Compos '97)*, volume 1536 of *Lecture Notes in Computer Science*, pages 541–564. Springer, 1998.

[25] J. Vitt and J. Hooman. Assertional specification and verification using PVS of the steam boiler system. In *Formal Methods for Industrial Applications: Specifying and Programming the Steam Boiler Control System*, volume 1165 of *Lecture Notes in Computer Science*, pages 453–472. Springer, 1996.