

# Exact State Set Representations in the Verification of Linear Hybrid Systems with Large Discrete State Space<sup>\*</sup>

Werner Damm<sup>2,3</sup>, Stefan Disch<sup>1</sup>, Hardi Hungar<sup>3</sup>, Swen Jacobs<sup>4</sup>, Jun Pang<sup>2</sup>, Florian Pigorsch<sup>1</sup>, Christoph Scholl<sup>1</sup>, Uwe Waldmann<sup>4</sup>, and Boris Wirtz<sup>2</sup>

<sup>1</sup> Albert-Ludwigs-Universität Freiburg

Georges-Köhler-Allee 51, 79110 Freiburg, Germany

<sup>2</sup> Carl von Ossietzky Universität Oldenburg

Ammerländer Heerstraße 114-118, 26111 Oldenburg, Germany

<sup>3</sup> OFFIS e.V., Escherweg 2, 26121 Oldenburg, Germany

<sup>4</sup> Max-Planck-Institut für Informatik

Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany

**Abstract.** We propose algorithms significantly extending the limits for maintaining exact representations in the verification of linear hybrid systems with large discrete state spaces. We use AND-Inverter Graphs (AIGs) extended with linear constraints (LinAIGs) as symbolic representation of the hybrid state space, and show how methods for maintaining compactness of AIGs can be lifted to support model-checking of linear hybrid systems with large discrete state spaces. This builds on a novel approach for eliminating sets of redundant constraints in such rich hybrid state representations by a suitable exploitation of the capabilities of SMT solvers, which is of independent value beyond the application context studied in this paper. We used a benchmark derived from an Airbus flap control system (containing  $2^{20}$  discrete states) to demonstrate the relevance of the approach.

## 1 Introduction

We target the verification of safety properties for embedded control applications in the transportation domain. Typical for such applications is a ratio of between 1:5 to 1:10 between the core control algorithms and diagnostic and fault-tolerance measures integrated into the controller, leading to a blow up of the discrete state space against pure control applications often reaching some  $10^6$  discrete states. As an example, we analyze a model derived from an Airbus flap controller [13], which on top of its control-loop for flap extraction and retraction is performing envelope protection to prevent loads on flaps possibly causing physical ruptures, and offers extensive monitoring of the health of its sub-systems to e. g. react on

---

<sup>\*</sup> This work was partly supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR 14 AVACS, <http://www.avacs.org/>).

loss of hydraulic pressure, rupture of the transmission shaft, or hardware failures. To prove safety of such controllers, we must combine methods for analyzing the pure control part (typically using linear dynamics for design models serving as reference for subsequent implementation steps) with state-space exploration methods dealing with large discrete state spaces. Such applications are out of reach for existing hybrid verification tools such as CheckMate [26], PHAVer [11], HyTech [15],  $d/dt$  [5]: while their strength rests in being able to address complex dynamics, they do not scale in the discrete dimension, since modes – the only discrete states considered – are represented explicitly when performing reachability analysis. On the other hand, hardware verification tools such as SMV [20] and VIS [27] scale to extremely large discrete systems, but clearly fail to be applicable to systems with continuous dynamics. To achieve a compact representation of such hybrid state-spaces, we enrich AND-Inverter Graphs (AIGs) with linear constraints. Previous work [23] demonstrated advantages of AIGs over BDDs for representing large discrete state-spaces compactly, due to their higher robustness in handling broad classes of Boolean functions in exhaustive state-space exploration. We lift methods such as test vector generation and SAT checking to detect equivalent (and thus redundant) nodes to the LinAIG level, providing a suite of heuristics including precise checks for equivalent LinAIG nodes using the SMT solver HySAT [10]. Moreover, we provide efficient methods for detecting and eliminating redundant linear constraints from LinAIGs, which are basically arbitrary boolean combinations of boolean variables and linear constraints. This extends results for eliminating redundant linear constraints from convex polyhedra used by Wang [28] and Frehse [11]. Our approach can be applied to perform backward reachability both for discrete time models (such as reference models for embedded controller implementation) and linear hybrid automata enriched with large discrete state spaces. In the latter case, we exploit the fact that the number of modes of a single controller is typically small (in the order of tens of modes) – this allows us to co-factor the LinAIG representation along modes. For each mode, we use the Loos-Weispfenning quantifier elimination technique for backward evaluation of the symbolic state-space representation along continuous flows. We counteract a worst-case quadratic blow up of linear constraints by tightly integrating redundancy elimination into the quantifier elimination process. Jointly, the presented techniques allow to achieve preciseness while maintaining sufficiently compact representations for the targeted application class.

This paper significantly extends our previous work [7] in adding quantifier elimination and redundancy elimination. The introduction of quantifier elimination was originally motivated by the wish to reduce the diameter of discrete time models. In allowing to fold the effect of large sequences of discretized flows into a single substitution, we accelerate hybrid system verification. This is different from the acceleration by folding hybrid control loops as in [6] which is performed in the world of few discrete states.

The presented methods are orthogonal to and may in the future be combined with abstraction techniques (such as bounding the degree of precision or loosening constraints as in [11]), incorporating robustness [12,8] or slackness [1,2]

in models allowing precise abstractions by finite grids under robustness respectively slackness assumptions, counter-example guided abstraction refinement as in [24,17,25] and techniques such as hybridization [4] for approximate linearization of richer dynamics.

The paper is organized as follows: Sections 2 and 3 give the formal mathematical model and present the backward-reachability algorithm. Sections 4 and 5 are dedicated to flow extrapolation and redundancy elimination. Evaluation results on the flap controller case study are presented in Section 6.

## 2 System Model

### 2.1 An Informal Description

This section elaborates on the characteristics of the systems to be analyzed, and motivates particular choices incorporated in the formal definition given in the following section. Our definition of hybrid systems can be seen as an extension of linear hybrid automata (LHA) [14] with a set of discrete variables. The state space is spanned by three classes of variables:

- *Continuous variables* represent sensor values, actuator values, plant states, and other real-valued variables used for the modeling of control-laws and plant dynamics.
- *Mode variables* represent a finite (small) set of *modes*, corresponding to the discrete states of an LHA; each mode is uniquely associated with a constant slope for each of the continuous variables, determining how the continuous valuation evolves over time as long as the system is in the given mode.
- *Discrete variables* code states from state-machines, switches, counters, sanity bits of sensor values, etc., and appear in modeling tools typically as bits, range types, or integer sub-ranges. In this paper we will assume some Boolean encoding of these variables. There are additional discrete input variables to our system.

Our models are closed-loop models without continuous input variables, combining controller and its controlled plant, hence sensors and actuators are internal continuous variables. Interactions of the environment are only possible through discrete input variables, allowing e.g. to select set-points, and to react to protocol messages. Non-deterministic choices are also modeled using discrete input variables. We remark that employing the construction from [3] permits us to extend our procedure to cope with slope sets bounded by constants which allow for non-determinism in plant dynamics, though we will not provide technical details of this extension in this paper.

The system evolves in alternating between continuous flows, in which time passes and only continuous variables are changed according to their slopes associated with the currently active mode of the system, and sequences of discrete transitions, which happen in zero time. Such discrete transitions update both discrete and continuous variables, and finally select the next active mode. All

(discrete) transitions are urgent, eliminating the need to associate state invariants with modes, as in other models of hybrid systems. Discrete inputs enter only in assignments to other discrete variables, i.e. they are disregarded during continuous evolutions. To allow e.g. for periodic sampling of discrete inputs, one can explicitly encode a (continuous) clock within one mode, and test for expiration of the clock-cycle within a transition guard.

## 2.2 Formal Model

We assume disjoint sets of variables  $C$ ,  $D$  and  $I$ . The elements of  $C$  are continuous variables, which are interpreted over the reals  $\mathbb{R}$ . The elements of  $D$  and  $I$  are discrete variables, where  $I$  will be used for inputs. For simplicity, we assume that they are of type boolean and range over the domain  $\mathbb{B} = \{0, 1\}$ . In the same way we assume that modes are encoded by a set  $\mathbf{M} \subseteq \{0, 1\}^l$  of boolean vectors of some fixed length  $l$ , leading to a set  $M$  of  $l$  (boolean) mode variables. We denote a valuation of (a subset of) these variables by  $(\mathbf{d}, \mathbf{i}, \mathbf{c}, \mathbf{m})$ .

A set of valuations (or states) can be represented symbolically using a suitable (quantifier-free) logic formula over  $D \cup I \cup C \cup M$ . We denote by  $\mathcal{B}(D \cup I)$  the set of boolean expressions over  $D \cup I$  and by  $\mathcal{B}(M)$  the set of boolean expressions over  $M$ . Here we restrict terms over  $C$  to the class of linear terms of the form  $\sum \alpha_i c_i + \alpha_0$  with rational constants  $\alpha_i$  and  $c_i \in C$ . Predicates are given by the set  $\mathcal{L}(C)$  of linear constraints, they have the form  $t \sim 0$ , where  $\sim \in \{=, <, \leq\}$  and  $t$  is a linear term. Finally,  $\mathcal{P}(D, C)$  is the set of all boolean combinations of variables from  $D$  and linear constraints over  $C$ .

In the following we use  $\xi$  for formulas in  $\mathcal{P}(D, C)$ ,  $\theta$  for boolean expressions from  $\mathcal{B}(M)$ ,  $g$  for boolean expressions from  $\mathcal{B}(D \cup I)$ ,  $t$  for linear terms over  $C$ , and  $\ell$  for linear constraints over  $C$ .

**Definition 1 (Syntax of CTHSs).** A continuous-time hybrid system *CTHS* contains six components:

- $D = \{d_1, \dots, d_n\}$  is a finite set of discrete variables,  $I = \{d_{n+1}, \dots, d_p\}$ , ( $p \geq n$ ) is a finite set of discrete inputs.
- $C = \{c_1, \dots, c_f\}$  is a finite set of continuous variables.
- $M = \{m_1, \dots, m_l\}$  is a finite set of mode variables,  $\mathbf{M} = \{\mathbf{m}_1, \dots, \mathbf{m}_k\} \subseteq \{0, 1\}^l$  is a finite set of modes, each value  $\mathbf{m}_i$  is associated with a vector  $\mathbf{v}_i \in \mathbb{R}^f$  of slopes for the variables in  $C$ .
- $GC$  is a global constraint in the form  $g_{gc}(D) \wedge \bigwedge_i \ell_i$  with linear constraints  $\ell_i$ .
- *Init* is a set of initial states, given in the form of  $\xi_0 \wedge \theta_0$ , where  $\xi_0 \in \mathcal{P}(D, C)$  and  $\theta_0 \in \mathcal{B}(M)$ .
- *DTrs* is the set of discrete transitions; each discrete transition is given as a guarded assignment  $ga_i$  ( $i = 1, \dots, u$  and  $u \geq 1$ ) in the form

$$\begin{aligned} \xi_i \wedge \theta_i \rightarrow (d_1, \dots, d_n) &:= (g_{i,1}, \dots, g_{i,n}); \\ (c_1, \dots, c_f) &:= (t_{i,1}, \dots, t_{i,f}); \\ (m_1, \dots, m_l) &:= \mathbf{m}_{j_i}. \end{aligned}$$

The typical usage of  $GC$  is to specify lower and upper bounds for continuous variables in runs to be considered. For simplicity, we assume that discrete inputs appear only on the right-hand side of assignments, but not in conditions.<sup>1</sup>

We add the following derived notions and restrictions to the CTHSs we consider:

**Definition 2 (Restrictions on CTHSs)**

- The guards of the discrete transitions must be mutually exclusive, i. e.  $(\xi_i \wedge \theta_i) \Rightarrow \neg(\xi_j \wedge \theta_j)$  for  $i \neq j$ .
- For each mode  $\mathbf{m}_i$  its boundary condition  $\beta_i$  is given by the cofactor of the disjunction of all discrete transition guards wrt.  $\mathbf{m}_i$ .<sup>2</sup> The boundary conditions have to form closed subsets of  $\mathbb{R}^f$  for each valuation of variables in  $D$ .

**Definition 3 (Semantics of CTHSs)**

- A state of a CTHS is a valuation  $s = (\mathbf{d}, \mathbf{c}, \mathbf{m})$  of  $D$ ,  $C$  and  $M$ .
- A discrete transition  $ga_i$  relates two states  $s \rightarrow_i s'$  iff the guard  $\xi_i \wedge \theta_i$  is true in  $s$  and the values in  $s'$  result from executing the assignments for some valuation  $\mathbf{i}$  of the input variables.
- A state  $s = (\mathbf{d}, \mathbf{c}, \mathbf{m}_i)$  evolves in time  $\lambda \in \mathbb{R}_{>0}$  into  $s' = (\mathbf{d}, \mathbf{c} + \lambda \mathbf{v}_i, \mathbf{m}_i)$ , written as  $s \rightsquigarrow^\lambda s'$ .  $s'$  is a  $\lambda$ -time successor of  $s$  ( $s \rightarrow^\lambda s'$ ), if  $s \rightsquigarrow^\lambda s'$  and for all  $s''$  with  $s = s''$  or  $s \rightsquigarrow^{\lambda''} s''$  for some  $\lambda'' < \lambda$ , we have  $s'' \models GC$  and  $s'' \not\models \beta_i$  (i. e. neither we violate the global constraints nor hit a discrete transition guard along the way).
- $\rightarrow =_{\text{df}} (\bigcup_{i=1}^u \rightarrow_i) \cup (\bigcup_{\lambda>0} \rightarrow^\lambda)$  is the transition relation of the CTHS. A trajectory is a finite or infinite sequence of states  $(s^j)_{j \geq 0}$  with  $s^0 \in \text{Init}$ , all  $s^j \models GC$ , and  $s^{j-1} \rightarrow s^j$  for each  $j > 0$ . A state is reachable if there is a trajectory ending in that state.

Note that the definition of a time successor makes the discrete transitions *urgent*: they fire once they become enabled. This explains why we do not need invariants of modes while on the other hand we have to require closed sets for boundary conditions.

### 3 Approach

In this section, we describe the main structure of our algorithm. We recall the ingredients which it shares with its predecessor from [7] and point to the new constituents which are detailed in the ensuing sections.

*Overview.* Our algorithm checks whether all reachable states are within a given set of (safe) states  $S_0$ . To establish this, a backwards fixpoint computation is performed. Starting with the set  $S_0$  enriched by all states violating the global

<sup>1</sup> The usual solution is to check discrete inputs only when a timer has expired.

<sup>2</sup> The cofactor is the partial evaluation of the disjunction wrt.  $(m_1, \dots, m_l) = \mathbf{m}_i$ . It does not depend on  $M$  anymore.

constraints, repeatedly the (safe) pre-image is computed until a fixpoint is reached or some initial state is removed from the fixpoint approximant. In the latter case, a state outside of  $S_0$  is reachable (while observing the global constraints). So we employ repeatedly

$$\text{Safepre}(S) =_{\text{df}} \{ s \in S \mid \forall s'. s \rightarrow s' \Rightarrow s' \in S \},$$

which corresponds to the temporal operator **AX**. We have chosen the backwards direction, because for discrete transitions the pre-image is expressed essentially by a substitution (see Hoare’s program logic [16]).

*Step computation.* We split the computation of *Safepre* into a discrete (*Safepre<sub>D</sub>*) and a continuous (*Safepre<sub>C</sub>*) part. The computation of *Safepre<sub>D</sub>* using boolean operations, substitutions (both for boolean and real variables), and boolean quantification has been already described in [7]. We will explain our new method to cope with continuous-time evolutions (which did not occur in the discrete-time models of the precursor paper) in detail in Sect. 4.

*Termination.* Since the equivalence of state sets (we deal with boolean combinations of linear constraints, as detailed in the following) is decidable, termination of the algorithms enables us to answer the reachability question. However, it should be noted that termination is not guaranteed – otherwise our algorithm would constitute a solution to an undecidable problem<sup>3</sup>. We expect that the algorithm terminates – in theory – for the great majority of problems coming from applications. We consider complexity the much more relevant challenge in practice. Let us also remark that the implemented fixpoint computation is more elaborated in detail than the somewhat simplified version described here (due to lack of space).

*Representation of state sets.* Our algorithm operates on a specific data structure efficiently implementing formulas from  $\mathcal{P}(D, C) \cup \mathcal{B}(M)$ . These can be seen as boolean combinations over  $D$ ,  $M$  and linear constraints  $\mathcal{L}(C)$ . We use a set of new (boolean) *constraint variables*  $Q$  as encodings for the linear constraints, where each occurring  $\ell \in \mathcal{L}(C)$  is encoded by some  $q_\ell \in Q$ . An important characteristic of our procedure is that the set of constraint variables may grow as the step computation continues, so that new variables are introduced continuously.

For the boolean structure we employ Functionally Reduced AND-Inverter Graphs (FRAIGs) [21,23]. These are a semi-canonical variant of AND-Inverter Graphs (AIGs) [22,18]. Basically, they are boolean circuits consisting only of AND gates and inverters. Semi-canonical means that no two nodes represent the same boolean function. In the presence of atoms encoding linear constraints, we call them linear constraint AIGs, or shortly LinAIGs. Their structure is illustrated in Fig. 1.

<sup>3</sup> Even if the global constraints define a bounded region, one can straightforwardly encode arithmetic on integers represented as fractions  $1/2^n$  of continuous values. This is a common integer representation used in the literature for showing undecidabilities in related domains.

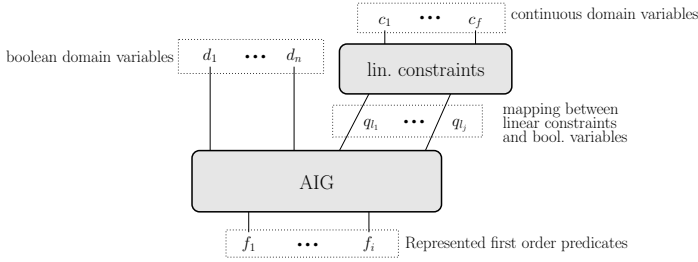


Fig. 1. The LinAIG structure

*Efficiency measures.* We have put much effort into the efficiency of our implementation, in particular into the time efficiency of the routines which keep the representations as small as possible. We briefly summarize some techniques described in more detail in [7], while Sec. 5 presents important improvements. Basically, functional reducedness (generalized from FRAIGs to LinAIGs) can be achieved by checking all pairs of nodes for equivalence, taking the interpretation of constraint variables  $q_\ell$  by the corresponding linear constraints  $\ell$  into account. This task can be performed by an SMT (SAT modulo theories) solver such as HySAT [10], which combines DPLL with linear programming as a decision procedure. However, it would be much too costly to call HySAT every time a new node is introduced. Instead, a hierarchy of approximate techniques is used to factor out “easy” problem instances. In first steps purely boolean approximations are employed: If two nodes represent equivalent boolean formulas, we do not need to refer to the definition of the constraint variables. Here we make use of capabilities of FRAIGs, which include local boolean normalization rules, simulation, and SAT checks. Additionally, boolean reasoning is supported by (approximate) knowledge on linear constraints such as implications between constraints. For identifying non-equivalent LinAIG nodes we use test vectors with valuations  $\mathbf{c} \in \mathbb{R}^f$ , and it proved to be worthwhile to use not only randomly generated test vectors, but also test vectors extracted from failed exact checks done by HySAT (*learning* test vectors). All of these techniques are arranged in a carefully designed and tested strategy of when to apply which technique.

## 4 Flow Extrapolation

*Continuous transitions.* In our system model, the time steps only concern the evolutions of continuous variables and leave the discrete part unchanged. For each mode, the continuous safe pre-image  $Safepre_C$  can be expressed as a formula with one quantified real variable (time). We will show how to eliminate this quantifier to arrive at a formula which can again be represented by a LinAIG.

Let  $\phi(D, M, Q)$  be a representation of a state set. Each valuation  $\mathbf{m}_i$  of the mode variables in  $M$  encodes a concrete mode with an associated evolution  $\mathbf{v}_i$  of  $C$  and boundary condition  $\beta_i$ . Let  $\phi_i$  be the cofactor of  $\phi$  w. r. t. mode  $\mathbf{m}_i$ . Thus we have  $\phi \Leftrightarrow \bigvee_{i=1}^k \phi_i \wedge (m_1, \dots, m_l) = \mathbf{m}_i$ , where each  $\phi_i$  is a boolean

formula over  $D$  and  $Q$ . For each mode  $\mathbf{m}_i$ , we must now determine the set of all valuations for which every (arbitrarily long) evolution along  $\mathbf{v}_i$  remains in the set of valuations satisfying  $\phi_i$ , either forever or until it meets a point that satisfies the boundary condition  $\beta_i$  or violates the global constraints  $GC$ . We denote this set by  $\text{Safe}_{pre_C}(\phi_i, \mathbf{v}_i, \beta_i)$ . Logically, it can be described by the formula

$$\forall \lambda. (\lambda < 0 \vee \phi_i(\mathbf{c} + \lambda \mathbf{v}_i) \vee \neg GC(\mathbf{c} + \lambda \mathbf{v}_i) \vee \exists \lambda'. (\lambda' \geq 0 \wedge \lambda' < \lambda \wedge (\beta_i(\mathbf{c} + \lambda' \mathbf{v}_i) \vee \neg GC(\mathbf{c} + \lambda' \mathbf{v}_i))))).$$

Under the assumption that the set described by  $GC$  is convex, and using the fact that we are only interested in states satisfying  $GC$ , this formula can be simplified (modulo  $GC$ ) to

$$\forall \lambda. (\lambda < 0 \vee \phi_i(\mathbf{c} + \lambda \mathbf{v}_i) \vee \neg GC(\mathbf{c} + \lambda \mathbf{v}_i) \vee \exists \lambda'. (\lambda' \geq 0 \wedge \lambda' < \lambda \wedge \beta_i(\mathbf{c} + \lambda' \mathbf{v}_i))).$$

Our task is now to convert this formula over  $\lambda, \lambda', C$ , and  $D$  into an equivalent formula over the original variables in  $C$  and  $D$ . If the variables in  $C$  occur in  $\phi$  and  $\beta$  only within linear constraints, then this amounts to variable elimination for linear real arithmetic.<sup>4</sup>

*Test points.* The Loos-Weispfenning test point method [19,9] eliminates universal quantifiers by converting them into finite conjunctions (and dually, existential quantifiers into finite disjunctions). The method is based on the following observation: Assume that a formula  $\psi(x, \vec{y})$  is written as a positive boolean combination of linear constraints  $x \sim_i t_i(\vec{y})$  and  $0 \sim'_j t'_j(\vec{y})$ , where  $\sim_i, \sim'_j \in \{=, \neq, <, \leq, >, \geq\}$ . Let us keep the values of  $\vec{y}$  fixed for a moment. If the set of all  $x$  such that  $\psi(x, \vec{y})$  does not hold is non-empty, then it can be written as a finite union of (possibly unbounded) intervals, whose boundaries are among the  $t_i(\vec{y})$ . To check whether  $\forall x. \psi(x, \vec{y})$  holds, it is therefore sufficient to test  $\psi(x, \vec{y})$  for either all upper or all lower boundaries of these intervals. The test values may include  $+\infty, -\infty$ , or a positive infinitesimal  $\varepsilon$ , but these can easily be eliminated from the substituted formula. For instance, if  $x$  is substituted by  $t_j(\vec{y}) - \varepsilon$ , then both the linear constraints  $x \leq t_i(\vec{y})$  and  $x < t_i(\vec{y})$  are turned into  $t_j(\vec{y}) \leq t_i(\vec{y})$ , and both  $x \geq t_i(\vec{y})$  and  $x > t_i(\vec{y})$  are turned into  $t_j(\vec{y}) > t_i(\vec{y})$ .

There are two possible sets of test points, depending on whether we consider upper or lower boundaries:

$$\begin{aligned} TP_1 &= \{+\infty\} \cup \{t_i(\vec{y}) \mid \sim_i \in \{\neq, >\}\} \cup \{t_i(\vec{y}) - \varepsilon \mid \sim_i \in \{=, \geq\}\} \\ TP_2 &= \{-\infty\} \cup \{t_i(\vec{y}) \mid \sim_i \in \{\neq, <\}\} \cup \{t_i(\vec{y}) + \varepsilon \mid \sim_i \in \{=, \leq\}\}. \end{aligned}$$

Let  $TP$  be the smaller one of the two sets and let  $T$  be the set of all symbolic substitutions  $x/t$  for  $t \in TP$ . Then the formula  $\forall x. \psi(x, \vec{y})$  can be replaced by an equivalent finite conjunction  $\bigwedge_{\sigma \in T} \psi(x, \vec{y})\sigma$ . The size of  $TP$  is in general linear in

<sup>4</sup> The variables in  $D$  are assumed to remain constant during mode  $\mathbf{m}_i$ , so boolean expressions over  $D$  behave like propositional variables. For simplicity, we will ignore them in the rest of this section.



the size of  $\psi$ , so the size of the resulting formula is quadratic in the size of  $\psi$ . This is independent of the boolean structure of  $\psi$  – conversion to CNF is not required. On the other hand, if  $\psi$  is a conjunction  $\bigwedge \psi_i$ , then the test point method can also be applied to each of the formulas  $\psi_i$  individually, leading to a smaller number of test points. Moreover, when the test point method transforms each  $\psi_i$  into a finite conjunction  $\bigwedge \psi_i^j$ , then each  $\psi_i^j$  contains at most as many linear constraints as the original  $\psi_i$ , and only the length of the outer conjunction increases.

*Applying the test point method to flow extrapolation.* We have demonstrated above that the safe pre-image  $\text{Safepre}_C(\phi_i, \mathbf{v}_i, \beta_i)$  of the formula  $\phi_i$  is

$$\forall \lambda. (\lambda < 0 \vee \phi_i(\mathbf{c} + \lambda \mathbf{v}_i) \vee \neg GC(\mathbf{c} + \lambda \mathbf{v}_i) \vee \exists \lambda'. (\lambda' \geq 0 \wedge \lambda' < \lambda \wedge \beta_i(\mathbf{c} + \lambda' \mathbf{v}_i))).$$

Assuming that  $\phi_i$  equals  $\bigwedge_k \phi_{ik}$  and that  $\beta_i$  equals  $\bigvee_j \beta_{ij}$ , we obtain

$$\bigwedge_k \forall \lambda. (\lambda < 0 \vee \phi'_{ik}(\mathbf{c} + \lambda \mathbf{v}_i) \vee \bigvee_j \exists \lambda'. (\lambda' \geq 0 \wedge \lambda' < \lambda \wedge \beta_{ij}(\mathbf{c} + \lambda' \mathbf{v}_i))).$$

where  $\phi'_{ik}$  abbreviates  $\phi_{ik} \vee \neg GC$ . Applying the test point method, we replace the universal and the existential quantifier by a finite conjunction or disjunction using a set of symbolic substitutions  $T'_j$  for  $\lambda'$  (which depends on  $\beta_{ij}$  and  $\mathbf{v}_i$ ) and a set of symbolic substitutions  $T_k$  for  $\lambda$  (which depends on  $\phi_{ik}$ , the  $\beta_{ij}$ , and  $\mathbf{v}_i$ ):

$$\begin{aligned} \text{Safepre}_C(\phi_i, \mathbf{v}_i, \beta_i) &= \bigwedge_k \bigwedge_{\sigma \in T_k} ((\lambda < 0 \vee \phi'_{ik}(\mathbf{c} + \lambda \mathbf{v}_i))\sigma \\ &\quad \vee \bigvee_j \bigvee_{\tau \in T'_j} (\lambda' \geq 0 \wedge \lambda' < \lambda \wedge \beta_{ij}(\mathbf{c} + \lambda' \mathbf{v}_i))\tau\sigma). \end{aligned}$$

Note that the test point method can work directly on the internal formula representation of LinAIGs – in contrast to the classic Fourier-Motzkin algorithm, there is no need for a costly CNF or DNF conversion before eliminating quantifiers. Moreover, the resulting formulas preserve most of the boolean structure of the original ones: the method behaves largely like a generalized substitution.

*Convexity.* It should be noted that some of the complexity of the general case disappears automatically if the complement of the boundary conditions is convex, that is, if every  $\beta_{ij}$  is a single linear inequation. Consider the formula  $\bigvee_{\tau \in T'_j} (\lambda' \geq 0 \wedge \lambda' < \lambda \wedge \beta_{ij}(\mathbf{c} + \lambda' \mathbf{v}_i))\tau$ . If  $\beta_{ij}$  is a single linear inequation, then two test points are always sufficient:<sup>5</sup> (a) If  $\beta_{ij}(\mathbf{c} + \lambda' \mathbf{v}_i)$  has the form  $\lambda' \leq t(\mathbf{c})$  or  $\lambda' < t(\mathbf{c})$ , then the test points are  $-\infty$  and 0, (b) otherwise, if  $\beta_{ij}(\mathbf{c} + \lambda' \mathbf{v}_i)$  has the form  $\lambda' \geq t(\mathbf{c})$  or  $\lambda' > t(\mathbf{c})$ , or if  $\lambda'$  is cancelled out completely in  $\beta_{ij}(\mathbf{c} + \lambda' \mathbf{v}_i)$ , then the test points are  $+\infty$  and  $\lambda - \varepsilon$ . Moreover, if  $+\infty$  or  $-\infty$  is substituted for  $\lambda'$ , the conjunction becomes trivially false, so the whole formula is reduced to  $0 < \lambda \wedge \beta_{ij}(\mathbf{c})$  in case (a) and to  $\lambda > 0 \wedge \beta_{ij}(\mathbf{c} + (\lambda - \varepsilon)\mathbf{v}_i)$  in case (b).

---

<sup>5</sup> Since we want to eliminate an existential quantifier, we have to use the dual form of the method described above.

## 5 Redundancy Elimination

Our earlier experiments demonstrated that LinAIGs form an efficient data structure for boolean combinations of boolean variables and linear constraints over real variables [7]. However, in connection with flow extrapolation using Loos-Weispfennig quantifier elimination, one observes that the number of “redundant” linear constraints grows rapidly during the fixpoint iteration of the model checker. For illustration see Fig. 2 and 3, which show a typical example from a model checking run representing a small state set based on two real variables: Lines in Figures 2 and 3 represent linear constraints, and the gray shaded area represents the space defined by some boolean combination of these constraints. While the representation depicted in Fig. 2 contains 24 linear constraints, a closer analysis shows that an optimized representation can be found using only 15 linear constraints as depicted in Fig. 3.




Fig. 2. Before redundancy removal




Fig. 3. After redundancy removal

*Removing redundant constraints* from our representations turned out to be a crucial task for the success of our methods. It should be noted that, since we represent arbitrary boolean combinations of linear constraints (and boolean variables), this task is not as straightforward as for other approaches such as [14,11] which represent sets of convex polyhedra, i. e., sets of conjunctions  $\ell_1 \wedge \dots \wedge \ell_n$  of linear constraints. If one is restricted to convex polyhedra, the question whether a linear constraint  $\ell_1$  is redundant in the representation reduces to the question whether  $\ell_2 \wedge \dots \wedge \ell_n$  represents the same polyhedron as  $\ell_1 \wedge \dots \wedge \ell_n$ , or equivalently, whether  $\overline{\ell_1} \wedge \ell_2 \wedge \dots \wedge \ell_n$  represents the empty set. This question can simply be answered by a linear constraint solver.

For redundancy elimination in our context consider a predicate  $F(b_1, \dots, b_k, \ell_1, \dots, \ell_n)$  (represented by a LinAIG) where  $b_1, \dots, b_k$  are boolean variables,  $\ell_1, \dots, \ell_n$  are linear constraints over  $C$ , and  $F$  is a boolean function.

**Definition 4 (Redundancy of linear constraints).** *The linear constraints  $\ell_1, \dots, \ell_r$  ( $1 \leq r \leq n$ ) are called redundant in the representation of  $F(b_1, \dots, b_k, \ell_1, \dots, \ell_n)$  iff there is a boolean function  $G$  with the property that  $F(b_1, \dots, b_k, \ell_1, \dots, \ell_n)$  and  $G(b_1, \dots, b_k, \ell_{r+1}, \dots, \ell_n)$  represent the same predicates.*

In order to be able to check for redundancy, we need a disjoint copy  $C' = \{c'_1, \dots, c'_f\}$  of the continuous variables  $C = \{c_1, \dots, c_f\}$ . Moreover, for each

linear constraint  $\ell_i$  ( $1 \leq i \leq n$ ) we introduce a corresponding linear constraint  $\ell'_i$  which coincides with  $\ell_i$  up to replacement of variables  $c_j \in C$  by variables  $c'_j \in C'$ . Our check for redundancy is based on the following theorem:

**Theorem 5 (Redundancy check).** *The linear constraints  $\ell_1, \dots, \ell_r$  ( $1 \leq r \leq n$ ) are redundant in the representation of  $F(b_1, \dots, b_k, \ell_1, \dots, \ell_n)$  iff the predicate*

$$F(b_1, \dots, b_k, \ell_1, \dots, \ell_n) \oplus F(b_1, \dots, b_k, \ell'_1, \dots, \ell'_n) \wedge \bigwedge_{i=r+1}^n (\ell_i \equiv \ell'_i) \quad (1)$$

(where  $\oplus$  denotes exclusive-or) is not satisfiable by any assignment of boolean values to  $b_1, \dots, b_k$  and real values to the variables  $c_1, \dots, c_f, c'_1, \dots, c'_f$ .

Note that the check from Thm. 5 can be performed by an SMT solver such as HySAT [10]. By lack of space we just give a sketch of the intuition behind Thm. 5.

According to Def. 4 linear constraints  $\ell_1, \dots, \ell_n$  are redundant iff there is a boolean function  $G$  such that  $G(b_1, \dots, b_k, \ell_{r+1}, \dots, \ell_n)$  and  $F(b_1, \dots, b_k, \ell_1, \dots, \ell_n)$  represent the same predicates. Now let us look at  $F(b_1, \dots, b_k, \ell_1, \dots, \ell_n)$  as a boolean function  $F(b_1, \dots, b_k, q_{\ell_1}, \dots, q_{\ell_n})$  with (new) boolean constraint variables  $q_{\ell_1}, \dots, q_{\ell_n}$  and a mapping connecting  $q_{\ell_i}$  to  $\ell_i$  (just as in our definition of LinAIGs). In comparison to  $F$  the required boolean function  $G$  must depend only on variables  $b_1, \dots, b_k, q_{\ell_{r+1}}, \dots, q_{\ell_n}$ .

If formula (1) is satisfied by some assignment  $\mathbf{d} \in \{0, 1\}^k$  to the boolean variables  $b_1, \dots, b_k$ ,  $\mathbf{c} \in \mathbb{R}^f$  to the real variables  $c_1, \dots, c_f$  (which are inputs of linear constraints  $\ell_i$ ), and  $\mathbf{c}' \in \mathbb{R}^f$  to the copied real variables  $c'_1, \dots, c'_f$  (which are inputs of copied linear constraints  $\ell'_i$ ), then the first part of formula (1), i. e.  $F(b_1, \dots, b_k, \ell_1, \dots, \ell_n) \oplus F(b_1, \dots, b_k, \ell'_1, \dots, \ell'_n)$  enforces that the predicate  $F$  changes its value if input  $\mathbf{c}$  is replaced by input  $\mathbf{c}'$  in the corresponding linear constraints. On the other hand, the second part  $\bigwedge_{i=r+1}^n (\ell_i \equiv \ell'_i)$  enforces that the truth assignment to linear constraints  $\ell_{r+1}, \dots, \ell_n$  does not change when replacing  $\mathbf{c}$  by  $\mathbf{c}'$ . However, since  $G$  only depends on variables  $b_1, \dots, b_k, q_{\ell_{r+1}}, \dots, q_{\ell_n}$  (whose truth assignments are not changed), function  $G$  “cannot see” the effect of changing  $\mathbf{c}$  to  $\mathbf{c}'$ . Thus  $G$  is not able to change its value like  $F$  when replacing  $\mathbf{c}$  by  $\mathbf{c}'$  and therefore it is not able to represent the same predicate as  $F$ .

Conversely, it can be seen that an appropriate function  $G$  can be constructed, when formula (1) is unsatisfiable. When constructing  $G$ , we use the notion of the *don't care set DC induced by linear constraints  $\ell_1, \dots, \ell_n$* : This don't care set  $DC := \{(v_{b_1}, \dots, v_{b_k}, v_{\ell_1}, \dots, v_{\ell_n}) \mid \exists (v_{c_1}, \dots, v_{c_f}) \in \mathbb{R}^f \text{ with } \ell_i(v_{c_1}, \dots, v_{c_f}) = v_{\ell_i} \forall 1 \leq i \leq n\}$  contains all boolean combinations that can not occur due to inconsistent assignments to boolean constraint variables. While for all  $(\mathbf{d}, \mathbf{c}) \in \overline{DC} := \{0, 1\}^{k+n} \setminus DC$  we have to postulate  $G(\mathbf{d}, \mathbf{c}) = F(\mathbf{d}, \mathbf{c})$ , the value of  $G$  may be chosen arbitrarily for all  $(\mathbf{d}, \mathbf{c}) \in DC$ , since these values can not occur due to inconsistencies between linear constraints. A closer analysis shows that – under assumption of unsatisfiability of formula (1) – it is indeed possible to define the function values of  $G(\mathbf{d}, \mathbf{c})$  for  $(\mathbf{d}, \mathbf{c}) \in DC$  in such a way that  $G$  will not depend on variables  $q_{\ell_1}, \dots, q_{\ell_r}$ . This proves that linear constraints  $\ell_1, \dots, \ell_r$  are then redundant.

A straightforward realization of this approach would need a (compact) representation of the don't care set  $DC$  in order to compute an appropriate boolean

function  $G$ . However, two interesting observations turn the basic idea into a feasible approach:

1. In general, we do not need the complete set  $DC$  for the definition of the boolean function  $G$ .
2. A representation of a subset of  $DC$  which is needed for removing the redundant constraints  $\ell_1, \dots, \ell_r$  is already computed by an SMT solver when checking satisfiability of formula (1).

Again, more details on how the SMT solver internally computes a representation of a sufficient subset of  $DC$  and on the method for actually removing redundant constraints from our representations are omitted due to lack of space. Our ideas for redundancy detection and removal have been implemented based on the SMT solver HySAT. Experiments given in Section 6 show that integrating redundancy removal is crucial for the success of our methods.

## 6 Experimental Results

Our sample application is derived from a case study for Airbus, a controller for the flaps of an aircraft [13]. The flaps are extended during take-off and landing to generate more lift at low velocity. They are not robust enough for high velocity, so they must be retracted for other periods. It is the controller’s task to correct the pilot’s commands if he endangers the flaps. Additionally, there is also an extensive monitoring of the health of its sub-systems, checking for instance for hardware failures. The health monitoring system interacts with the flap control by enforcing a more conservative behavior of the control when errors are supposed to be in the system.

The benchmark used here is a simplified version of the full system including the flap controller and a health monitoring system, which is triggered by a timer. The model has three continuous variables: the velocity, the flap angle, and the timer value. Discrete states of the controller and of the health monitoring system contribute to the discrete state space. The discrete state space contains  $2^{20}$  discrete states. This size is clearly out of reach for hybrid verification tools known from the literature, which do not scale in the discrete dimension, since modes – the only discrete states considered – are represented *explicitly* when performing reachability analysis.


The safety property to be established for our model is “For the current flap setting, the aircraft’s velocity shall not exceed the nominal velocity (w. r. t. the flap position) plus 7 knots”. Whether this requirement holds for our model depends on a “race” between flap retraction and speed increase. The controller is correct, if it initiates flap retraction (by correcting the pilot) early enough.

Based on the ideas presented in the previous sections we implemented a prototype model checker using LinAIGs for representing sets of states. Our experiments were run on an AMD Opteron with 2.6 GHz and 16 GB RAM.


Our model checker was able to prove the given safety invariant for the case study in 888.6 CPU seconds. The LinAIG representation had a maximum number

of 30887 nodes and a maximum number of 80 linear constraints. The number of flow extrapolation steps using Loos-Weispfennig quantifier elimination was 6, the number of discrete image computation steps performed until reaching the fixpoint was 20. This result clearly demonstrates that our approach is able to successfully verify hybrid systems including discrete parts with state spaces of considerable sizes.

In the following we analyze how the individual ingredients of our method contribute to its overall success. Redundancy elimination turned out to be absolutely necessary to make flow extrapolation using Loos-Weispfennig quantifier elimination feasible. Fig. 4 illustrates the difference between the model checking runs for our case study with and without redundancy removal by plotting the numbers of linear constraints used during the model checking run. Without redundancy removal (dotted line), the number of linear constraints is rapidly increasing up to a number of 1000 linear constraints and 150000 LinAIG nodes in the fourth flow extrapolation.<sup>6</sup> On the other hand, redundancy elimination detects many of the linear constraints to be redundant in our LinAIG representations. Having a closer look at the solid line in Fig. 4 one can identify six groups of three peaks in the number of linear constraints corresponding to six flow extrapolations for three modes, respectively. One notices that redundancy elimination is able to keep the numbers of linear constraints small after Loos-Weispfennig quantifier elimination, so that the number of linear constraints does not exceed 80 during the model checking run. Redundancy elimination removes redundant constraints early and has thus the additional effect that the number of constraints does not blow up due to a series of further substitutions into the removed constraints in following flow extrapolation steps.



**Fig. 4.** Comparison of the LinAIG evolution with and without redundancy removal



**Fig. 5.** Discrete time example with different time steps

Finally, we want to compare the results for our current system model, which includes continuous evolution of variable length in one operation based on flow

<sup>6</sup> Without redundancy removal the remaining two flow extrapolations could not be performed within our timeout of 24 hours.

extrapolation, to results for a corresponding model with discrete time semantics as presented in [7]. The system model from [7] has no continuous evolution, and discrete steps take fixed time  $\delta$ . We emphasize that, although time discretized models are widely used in practical applications, they have the problem that unsafe states may be reachable from the initial state, but reachability of these states is not observed due to the time discrete nature of steps. Reducing the width of the discrete time steps can alleviate this problem, but it comes at the cost of a larger number of steps for fixpoint iterations and a larger number of LinAIG nodes for representing sets of states. Our continuous time approach does not show this problem. An analysis of this issue (here done for a flap controller without health monitoring system) is given in Fig. 5. It shows the peak numbers of LinAIG nodes (dotted line) and the run times (solid line) for the example. Here the width of the discrete time step varies between 400 ms and 50 ms. Our analysis clearly shows that run times in the discrete time model largely depend on the width of the time discretization step. At a time step of 400 ms the fixpoint iteration took 5.4 CPU seconds for 9 steps, at 100 ms 763.8 CPU seconds for 33 steps, and at 50 ms 22497 CPU seconds for 65 steps.

This demonstrates a dilemma of the time discretized version: We have to keep the time step small both to be sure not to miss relevant reachable states and to be able to model the system correctly (of course, with discrete time steps of 400 ms we are not able to model realistic controllers sampling every 100 ms). However, decreasing the time step too much may turn the model checking problem intractable. In contrast, in our novel approach we do not work with time discretizations, but we are able to compute continuous evolutions of variable lengths in one operation based on flow extrapolation. Sequences of discrete steps of the previous version [7] where no mode switches are triggered are collapsed into a single symbolic substitution in this way. Note that in the example without health monitoring system only five flow extrapolation steps are needed to reach the fixpoint within a runtime of 27.7 s (whereas for the discrete time model with a time step of 50 ms, e. g., the number of steps amounts to 65 with a run time of 22497 s).

## 7 Conclusion

We consider the tight integration of LinAIGs and HySAT in backward reachability analysis a core technology to address scalability of hybrid system verification methods with large discrete state spaces, and have demonstrated the relevance of the approach using a benchmark derived from an Airbus flap controller. The redundancy elimination technique presented in Section 5 is of independent value and could be integrated in other hybrid verification tools. Next imminent extensions of our approach cover differential inclusions and continuous inputs. We will experiment with incorporating orthogonal extensions to our approach such as exploiting robustness, over-approximation, and counterexample guided abstraction refinement to address richer dynamics and achieve further scalability.

## References

1. Agrawal, M., Thiagarajan, P.S.: Lazy rectangular hybrid automata. In: Alur, R., Pappas, G.J. (eds.) HSCC 2004. LNCS, vol. 2993, pp. 1–15. Springer, Heidelberg (2004)
2. Agrawal, M., Thiagarajan, P.S.: The discrete time behavior of lazy linear hybrid automata. In: Morari, M., Thiele, L. (eds.) HSCC 2005. LNCS, vol. 3414, pp. 55–69. Springer, Heidelberg (2005)
3. Alur, R., Courcoubetis, C., Halbwachs, N., Henzinger, T.A., Ho, P.-H., Nicollin, X., Olivero, A., Sifakis, J., Yovine, S.: The algorithmic analysis of hybrid systems. *Theoretical Computer Science* 138(1), 3–34 (1995)
4. Asarin, E., Dang, T., Girard, A.: Hybridization methods for the analysis of non-linear systems. *Acta Informatica* 43(7), 451–476 (2007)
5. Asarin, E., Dang, T., Maler, O.: The d/dt tool for verification of the hybrid systems. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 365–370. Springer, Heidelberg (2002)
6. Boigelot, B., Herbretreau, F.: The power of hybrid acceleration. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 438–451. Springer, Heidelberg (2006)
7. Damm, W., Disch, S., Hungar, H., Pang, J., Pigorsch, F., Scholl, C., Waldmann, U., Wirtz, B.: Automatic verification of hybrid systems with large discrete state space. In: Graf, S., Zhang, W. (eds.) ATVA 2006. LNCS, vol. 4218, pp. 276–291. Springer, Heidelberg (2006)
8. Damm, W., Pinto, G., Ratschan, S.: Guaranteed termination in the verification of LTL properties of non-linear robust discrete time hybrid systems. *Journal of Foundations of Computer Science* 18(1), 63–86 (2007)
9. Dolzmann, A.: Algorithmic Strategies for Applicable Real Quantifier Elimination. PhD thesis, Universität Passau (2000)
10. Fränzle, M., Herde, C.: HySAT: An efficient proof engine for bounded model checking of hybrid systems. *Formal Methods in System Design* 30(3), 179–198 (2007)
11. Frehse, G.: Compositional Verification of Hybrid Systems using Simulation Relations. PhD thesis, Radboud Universiteit Nijmegen (2005)
12. Girard, A., Pappas, G.J.: Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control* 52(5), 782–798 (2007)
13. H3 FOMC Team. The flap controller description.  
<http://www.avacs.org/Benchmarks/flapcontroller.pdf>
14. Henzinger, T.A.: The theory of hybrid automata. In: 11th IEEE Symposium on Logic in Computer Science, pp. 278–292. IEEE Press, Los Alamitos (1996)
15. Henzinger, T.A., Ho, P.-H., Wong-Toi, H.: HyTech: A model checker for hybrid systems. *Software Tools for Technology Transfer* 1(1–2), 110–122 (1997)
16. Hoare, C.A.R.: An axiomatic basis for computer programming. *Communication of the ACM* 12, 576–583 (1969)
17. Jha, S., Brady, B., Seshia, S.: Symbolic reachability analysis of lazy linear hybrid automata. Technical report, EECS Dept. UC Berkeley (2007)
18. Kuehlmann, A., Paruthi, V., Krohm, F., Ganai, M.K.: Robust boolean reasoning for equivalence checking and functional property verification. *IEEE Transactions on Computer-Aided Design* 21(12), 1377–1394 (2002)
19. Loos, R., Weispfenning, V.: Applying linear quantifier elimination. *The Computer Journal* 36(5), 450–462 (1993)
20. McMillan, K.L.: Symbolic Model Checking. Kluwer Academic Publishers, Dordrecht (1993)

21. Mishchenko, A., Chatterjee, S., Jiang, R., Brayton, R.K.: FRAIGs: A unifying representation for logic synthesis and verification. Technical report, EECS Dept. UC Berkeley (2005)
22. Paruthi, V., Kuehlmann, A.: Equivalence checking combining a structural SAT-solver, BDDs, and simulation. In: 18th IEEE Conference on Computer Design, pp. 459–464. IEEE Press, Los Alamitos (2000)
23. Pigorsch, F., Scholl, C., Disch, S.: Advanced unbounded model checking by using AIGs, BDD sweeping and quantifier scheduling. In: 6th Conference on Formal Methods in Computer Aided Design, pp. 89–96. IEEE Press, Los Alamitos (2006)
24. Platzer, A., Clarke, E.: The image computation problem in hybrid systems model checking. In: 10th Workshop on Hybrid Systems: Computation and Control. LNCS, vol. 4416, pp. 473–486. Springer, Heidelberg (2007)
25. Segelken, M.: Abstraction and counterexample-guided construction of  $\omega$ -automata for model checking of step-discrete linear hybrid models. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 433–448. Springer, Heidelberg (2007)
26. Silva, B.I., Richeson, K., Krogh, B.H., Chutinan, A.: Modeling and verification of hybrid dynamical system using CheckMate. In: 4th Conference on Automation of Mixed Processes (2000)
27. The VIS Group. VIS: A system for verification and synthesis. In: Alur, R., Henzinger, T.A. (eds.) CAV 1996. LNCS, vol. 1102, pp. 428–432. Springer, Heidelberg (1996)
28. Wang, F.: Symbolic parametric safety analysis of linear hybrid systems with BDD-like data-structures. IEEE Transactions on Software Engineering 31(1), 38–52 (2005)