STORMED hybrid systems

Vladimeros Vladimerou, Pavithra Prabhakar, Mahesh Viswanathan, and Geir Dullerud

University of Illinois at Urbana-Champaign
Champaign, Illinois, USA

Abstract. We introduce STORMED hybrid systems, a decidable class which is similar to o-minimal hybrid automata in that the continuous dynamics and constraints are described in an o-minimal theory. However, unlike o-minimal hybrid automata, the variables are not initialized in a memoryless fashion at discrete steps. STORMED hybrid systems require flows which are monotonic with respect to some vector in the continuous space and can be characterised as bounded-horizon systems in terms of their discrete transitions. We demonstrate that such systems admit a finite bisimulation, which can be effectively constructed provided the o-minimal theory used to describe the system is decidable. As a consequence, many verification problems for such systems have effective decision algorithms.

1 Introduction

Embedded processors and electronic controllers are seeing increasingly ubiquitous use and in critical cases require extremely accurate and predictable functionality. Such devices compute discrete steps while interacting with an environment with continuous dynamics and meeting real-time constraints. *Hybrid automata* [1] are a popular formal model used to describe such systems. They have (finitely many) discrete states, and continuous states evolving with time. The discrete and continuous states dictate when discrete transitions take place as well as what the effect of the transition is on the continuous part. Once such a system is modeled, the verification problem asks whether the formal model meets certain correctness requirements.

While the problem of verifying a general hybrid automaton against even simple properties (like invariants) is known to be undecidable, important decidable classes have been identified. *Timed automata* [2], certain special kinds of *rectangular hybrid automata* [9], and *o-minimal hybrid automata* [10] are important classes of general hybrid automata for which many verification problems are decidable. The decidability in all these cases is proved by demonstrating the existence of a finite, computable partition of the state space that is *bisimilar* to the original system. However, all these classes of decidable automata suffer from serious drawbacks — timed and rectangular hybrid automata have very simple dynamics for the way the continuous variables evolve, while o-minimal systems have strong reset conditions on discrete transitions, that decouples the

discrete dynamics from the continuous one, leaving the continuous state largely unaffected by the discrete transitions. The many undecidability results in the area [1, 9, 3, 4, 12] have reinforced the folklore belief that one must either restrict the continuous dynamics or the discrete dynamics to something simple, in order to achieve decidability. Notable exceptions like dynamical systems with piecewise constant derivatives [3] and polygonal hybrid systems [8] are however restricted to very low dimensions (only 2 variables are allowed to obtain decidability).

In this paper we introduce a new class of hybrid automata that we call *STORMED* hybrid systems (STORMED h.s.). These adhere to the following constraints. First the guards of any two transitions are **s**eparable in space by some minimum, non-zero distance. Next, all the constraints (i.e, the guards, invariants, and flows) must be definable in a *order-minimal* (or **o**-minimal) theory. Further we require the existence of a vector ϕ such that the flows in all the control states have a positive projection on ϕ, and the projections of the guards to have **d**elimited-**e**nds on this vector ϕ. These automata also have monotonic resets, which either leave the continuous state unchanged or advance its projection along ϕ. A form of monotonicity was also captured in [5].

Our main result in this paper is that STORMED h.s. can be shown to be *bisimilar* to a finite state transition system. Moreover the finite transition system can be effectively constructed provided the o-minimal theory in which the automaton is defined is decidable. Thus, STORMED h.s. can be verified against rich branching time properties expressed in logics such as CTL and μ-calculus [7].

STORMED h.s. are both more general in some respects, and more restrictive in some ways, when compared with other subclasses of hybrid automata investigated before. They allow for a richer continuous dynamics than timed automata and rectangular hybrid automata, and the discrete transitions can affect the continuous dynamics in non-trivial ways unlike o-minimal systems. However, they are required to have separable guards, monotonic flows/resets and delimited ends on guard constraints. In spite of these restrictions, we believe STORMED h.s. can be conveniently used to model interesting systems. For example, monotonicity is implicitly present in terms of a depleting resource, like fuel or time, while separability of guards translates to infrequency of discrete steps.

Finally we look at some relaxations of the STORMED model, and prove that removal of any single constraint cannot be tolerated. Such an investigation demonstrates that our model is reasonably tight; most relaxations of the constraints yield undecidable models.

2 Preliminaries

Equivalence Relations and Partitions. A binary relation R on a set A is a subset of $A \times A$. We will say aRb to denote $(a, b) \in R$. An *equivalence relation* on a set A is a binary relation R that is reflexive, symmetric and transitive. An equivalence relation partitions the set A into *equivalence classes*: $[a]_R = \{b \in A \mid aRb\}$. A partition Π of the set A defines a natural equivalence relation $\equiv_\Pi$, where $a \equiv_\Pi b$ iff a and b belong to the same partition in Π. In this paper, we will use

the partition Π to mean both the partition, as well as the equivalence relation associated with it. Finally, we will say an equivalence relation R_1 *refines* another equivalence relation R_2 iff $R_1 \subseteq R_2$.

Transition Systems and Bisimulation. A *transition system* is given by $\mathcal{S} = (Q, Q^0, \rightarrow)$, where Q is a set of states, Q^0 $subseteqQ$ is the set of initial states, and $\rightarrow \subseteq Q \times Q$ is the transition relation. For a transition system $\mathcal{S} = (Q, Q^0, \rightarrow)$, a *simulation relation* is a binary relation $R \subseteq Q \times Q$ such that if $(q_1, q_1') \in R$ and $q_1 \rightarrow q_2$ then there is q_2' such that $q_1' \rightarrow q_2'$ and $(q_2, q_2') \in R$. A binary relation R is said to be a *bisimulation* iff both R and R^{-1} are simulation relations. q_1 is said to be *bisimilar* to q_2 when there is a bisimulation R such that $(q_1, q_2) \in R$, and we denote this by $q_1 \cong q_2$. Bisimilarity $\cong$ is an equivalence relation and a bisimulation [11]. It is said to be of *finite* index if it has finitely many equivalence classes. A bisimulation R is said to *respect* a partition Π iff R refines the equivalence relation defined by Π.

First Order Logic. In this paper we will consider first order vocabularies consisting of only relation symbols and constant symbols; we will call $\mathcal{A}$ to be a τ-structure if it is a structure over the vocabulary τ. Recall that a k-ary relation $S \subseteq A^k$, where A is the domain of $\mathcal{A}$, is said to be *definable* in the structure $\mathcal{A}$ if there is a formula $\varphi(x_1, x_2, \ldots x_k)$, with free variables $x_1, \ldots x_k$, such that $S = \{(a_1, \ldots, a_k) \mid \mathcal{A} \models \varphi[x_i \mapsto a_i]_{i=1}^k\}$. A k-ary function f will be said to be definable if its graph, i.e., the set of all $(x_1, \ldots, x_k, f(x_1, \ldots x_k))$, is definable. A *theory* $T(\mathcal{A})$ of a structure $\mathcal{A}$ is the set of all sentences that hold in $\mathcal{A}$. $T(\mathcal{A})$ (or sometimes simply) is said to be *decidable* if there is an effective procedure to decide membership in the set $T(\mathcal{A})$. One consequence of this is that it is also decidable to check the emptiness of a definable relation, and whether two definable relations are equal.

O-minimality. A binary relation $\leq$ on a set A is said to be a *total ordering* if it is reflexive, transitive, antisymmetric $((a \leq b \ \wedge \ b \leq a) \ \Rightarrow a = b)$, and total $(a \leq b \vee b \leq a)$. The set A is said to be totally ordered if there is a total order on it. An *interval* is a set defined in a totally order set, using one or two bounds as follows: $\{x : a \leq x \leq b\}$, $\{x : x \leq a\}$, and $\{x : a \leq x\}$. Trivially, $\{x : a \leq x \leq b\}$ with $a = b$, is an interval consisting of a single point. We write $\mathcal{A} = (A, \leq, \ldots)$ to convey that the τ-structure $\mathcal{A}$ has an ordering relation $\leq$ and other elements in its structure. A totally ordered first-order structure $\mathcal{A} = (A, \leq, \ldots)$ is *o-minimal* (order-minimal) if every definable set is a finite union of intervals [16]. The theory of this structure is also called o-minimal. Examples of o-minimal structures include $(\mathbb{R}, <, +, -, \cdot, \texttt{exp})$ and $(\mathbb{R}, <, +, -, \cdot)$, where $+, -, \cdot, \texttt{exp}$ are the addition, subtraction, multiplication and exponentiation operations on reals, respectively. Additional examples can be found in [15, 16]. The theory of $(\mathbb{R}, <, +, -, \cdot)$ is known to be decidable [14].

3 Hybrid Systems and Special Subclasses

Hybrid systems mix discrete events with continuous dynamics. One formal representation that has been found to conveniently model the behavior of such systems is *hybrid automata* [9]. In this section, we recall the basic definition and introduce special classes of such systems, as a prelude to STORMED hybrid systems that we define in the next section and is the main object of study in this paper.

Definition 1 *A* hybrid automaton $\mathcal{H}$ *is a tuple* $(Q, \Delta, X, X_0, q_0, \mathcal{I}, \mathcal{F}, \mathcal{R}, \mathcal{G})$ *where*

- Q *is a finite set of* (discrete) control states,
- $\Delta \subseteq Q \times Q$ *is the set of edges between control states,*
- $X = \mathbb{R}^n$, *is the domain of the* continuous (part of the) state,
- $X_0 \subseteq X$ *is the set of* initial continuous states,
- $q_0 \in Q$ *is* initial control state,
- $\mathcal{I} : Q \to 2^X$, *is the function that associates with every control state an* invariant,
- $\mathcal{F} : Q \times X \to (\mathbb{R}_+ \to X)$ *is the function that associates with each* $(q, x) \in Q \times X$, *a flow function that describes how the continuous state changes with time,*
- $\mathcal{G} : \Delta \to 2^X$ *is the function that assigns to each edge a* guard, *which is a condition on the continuous state that must hold in order to take the discrete transition,*
- $\mathcal{R} : \Delta \to 2^{X \times X}$ *is the function that associates with each edge a* reset, *which is a binary relation that describes how the continuous state changes when a discrete transition is taken.*

In the above hybrid automaton, we call n the *dimension* of $\mathcal{H}$.
Notation: In order to make the text more readable, we will often write the argument of a function as a subscript. In particular, $\mathcal{I}_q$ will be used to denote the invariant associated with control state q instead of $\mathcal{I}(q)$, and similarly $\mathcal{G}_{(p,q)}$ and $\mathcal{R}_{(p,q)}$ to denote the guard and reset conditions associated with an edge (p, q) instead of $\mathcal{G}(p, q)$ and $\mathcal{R}(p, q)$. We will use $\mathcal{F}_{(q,x)}$ for the flow associated with (q, x) instead of $\mathcal{F}(q, x)$. Also, we call members of $Q \times X$ locations.

Before defining the semantics of the hybrid automata, we observe some conditions that the flow function must satisfy for it to define "reasonable continuous dynamics"; we call this *time-independent spatially consistent.*

Definition 2 *The flow function* $\mathcal{F} : Q \times X \to (\mathbb{R}_+ \to X)$ *is said to be* time-independent spatially-consistent *(TISC) if for every* $q \in Q$ *and* $x \in X$, $\mathcal{F}_{(q,x)}$ *satisfies the following conditions:*

1. $\mathcal{F}_{(q,x)}$ *is continuous and* $\mathcal{F}_{(q,x)}(0) = x$.
2. *It satisfies the following "semi-group" property: for every* $t \geq 0$ *and* $x' \in X$, *if* $\mathcal{F}_{(q,x)}(t) = x'$ *then for every* $t' \geq 0$, $\mathcal{F}_{(q,x)}(t + t') = \mathcal{F}_{(q,x')}(t')$.

Henceforth, we will assume all flows in the hybrid automata to be TISC flows.

Remark 3 TISC flows are a very basic requirement on the continuous dynamics satisfied by most definitions of hybrid automata in the literature (except in [6]). Typically the requirement is ensured by specifying the continuous dynamics in a control state by a differential equation which gives the derivative with respect to time of the continuous state evolution. The dynamics of the flow itself is then described in terms of the "integrals" of these differential constraints. In this paper, we find it convenient to instead directly talk about the flows themselves, rather than the differentials. Notice that a TISC flow is not required to be differentiable and therefore it allows for more general dynamics than is typically considered.

The semantics of a hybrid automaton $\mathcal{H}$ is defined in terms of a transition system $[\![\mathcal{H}]\!] = (C, C_0, \rightarrow)$, where

- $C = Q \times X$ is the set of states,
- $C_0 = q_0 \times X_0$ is the set of initial states, and
- the transition relation $\rightarrow$ is the union of *time transitions* $\rightarrow_t$ and discrete transitions $\rightarrow_d$ given by:
  - $(q_1, x_1) \rightarrow_t (q_2, x_2)$ iff $q_1 = q_2$ and there is a $t \in \mathbb{R}_+$ such that $x_2 = \mathcal{F}_{(q,x_1)}(t)$ and for all $t' \in [0, t]$, $\mathcal{F}_{(q,x_1)}(t') \in \mathcal{I}_{q_1}$.
  - $(q_1, x_1) \rightarrow_d (q_2, x_2)$ iff there is an edge $(q_1, q_2) \in \Delta$ such that $x_1 \in \mathcal{I}_{q_1}$, $x_2 \in \mathcal{I}_{q_2}$, $x_1 \in \mathcal{G}_{(q_1,q_2)}$, and $(x_1, x_2) \in \mathcal{R}_{(q_1,q_2)}$.

In a time transition, the discrete part q_1 of the state does not change but the continuous part changes according to the flow $\mathcal{F}_{q_1}$ while remaining within the invariant $\mathcal{I}_{q_1}$. On the other hand, in a discrete transition, control state changes according to an edge in the automaton, the continuous part of the state before the transition is required to satisfy the guard associated with the edge, and the result of taking the transition changes the continuous state according to the reset conditions associated with the edge.

An *execution* starting from state (q, x) is a sequence of states $(q_1, x_1), (q_2, x_2), \ldots, (q_k, x_k)$ such that $(q_1, x_1) = (q, x)$, and for all $i < k$, $(q_i, x_i) \rightarrow (q_{i+1}, x_{i+1})$. (q_k, x_k) is said to be *reachable* from (q, x). For a hybrid automaton $\mathcal{H}$, we say a control state q is *reachable*, if for some $x \in X$, $x_0 \in X_0$, (q, x) is reachable from an initial state (q_0, x_0). For a hybrid automaton $\mathcal{H}$, the *reachability problem* is to determine if a given control state is reachable.

3.1 Special Definitions

In this subsection we look at some special restrictions on hybrid automata that will be relevant for defining STORMED hybrid system that we consider in this paper.

3.2 Separable guards

A hybrid system $\mathcal{H} = (Q, \Delta, X, X_0, q_0, \mathcal{I}, \mathcal{F}, \mathcal{R}, \mathcal{G})$ is said to have *separable guards* if there exists $d_{min} > 0$ such that for every pair of distinct edges $(p_1, q_1), (p_2, q_2) \in$

Δ, $\min\{\|x_1 - x_2\| \mid x_1 \in \mathcal{G}_{(p_1,q_1)} \text{ and } x_2 \in \mathcal{G}_{(p_2,q_2)}\} \geq d_{min}$. The guards of $\mathcal{H}$ are said to be d_{min}-separable.

Here $\|\cdot\|$ denotes euclidean distance. Also, we will be using the dot product $x \cdot y$, where $x, y \in X$, to denote the real value of the length of the projection of y onto x as it is commonly used.

The guard separability removes the so-called Zeno behavior, i.e., it avoids infinite number of discrete steps in finite time.

Thus far, our discussion on hybrid automata did not address the issue of how the automaton is formally presented. The general definition presented does not give an effective presentation. We will consider automata where all the conditions, guards, invariants, etc. are described in a logical theory, and even more specifically in an o-minimal theory.

3.3 O-minimal Definability

A hybrid system $\mathcal{H}$ is said to be *definable in an o-minimal structure* $\mathcal{A} = \{A, \leq, \ldots\}$ (or simply called o-minimal), if all its initial conditions, invariants, flows, resets and guards are definable in $\mathcal{A}$.

Remark 4 In the literature, o-minimal hybrid automata [10] refer to hybrid automata as defined above with the additional restriction that all resets are *strong*. In other words, for any edge (p, q) the reset $\mathcal{R}_{(p,q)}$ is of the form $\mathcal{G}_{(p,q)} \times X'$ for some $X' \subseteq X$. This allows one to decouple the system into separate dynamical systems, with the discrete transitions "resetting" the continuous state on each discrete step. We do not need this decoupling in STORMED, but we do make use of o-minimality.

The subclass of hybrid automata that we will consider in this paper will have monotonicity requirements on the flow. We define these next.

3.4 Monotonic Flows

The set of flows $\mathcal{F}$ of $\mathcal{H}$ is *monotonic* with respect to a vector $\phi \in X$, if there exists an $\epsilon > 0$ such that for every $q \in Q, x \in X$, and $t, \tau \geq 0$,

$$\phi \cdot (\mathcal{F}_{(q,x)}(t + \tau) - \mathcal{F}_{(q,x)}(t)) \geq \epsilon \|\mathcal{F}_{(q,x)}(t + \tau) - \mathcal{F}_{(q,x)}(t)\|,$$

where $a \cdot b$ refers to the dot-product between the vectors. We call such a set of flows (ϵ, ϕ)-monotonic.

The above monotonicity requirement says that as the continuous state evolves with time according to any flow, the projection on the vector ϕ increases at a minimum rate ϵ. This guarantees that the projection on ϕ will never decrease.

Some obvious examples of monotonic flows are:

1. Linear flows of the form $\mathcal{F}_{(q,x)}(t) = x + \alpha_q(t)$, where $x \in \mathbb{R}^n$, and $\alpha_q \in (\mathbb{R}_+ - \{0\})^n$.
2. Analytic flows with their time-derivatives ranging on only one half-space, i.e., there exists a ϕ such that for all $q \in Q$ and $x \in X$, we have $\nabla_t \mathcal{F}_{(q,x)}(t) \cdot \phi > \epsilon \|\nabla_t \mathcal{F}_{(q,x)}(t)\|$.

3.5 Monotonic Resets

The collection of reset sets $\mathcal{R}$ of $\mathcal{H}$ is said to be *monotonic* with respect to some $\phi \in X$, if there exist $\epsilon, \zeta > 0$ such that for every $(p, q) \in \Delta$ and $x_1, x_2 \in X$ s.t. $(x_1, x_2) \in \mathcal{R}_{(p,q)}$, we have:

(i) if $p = q$, then either $x_1 = x_2$ or $\phi \cdot (x_2 - x_1) \geq \zeta$, and
(ii) if $p \neq q$, then $\phi \cdot (x_2 - x_1) \geq \epsilon \|x_2 - x_1\|$.

We call such a collection of resets (ϵ, ζ, ϕ)-monotonic.

Remark 5 Notice that in the case when the discrete state changes, we do not require the reset to move the continuous state along ϕ by a minimum value. It only requires the change in the continuous state along ϕ is lower bounded by the actual change in the continuous state. In particular, it forbids resets that take the continuous state back along ϕ. Also our definition allows for identity resets.

Our definition guarantees that a minimum distance of $\min\{\zeta, \epsilon d_{min}\}$ is traveled along ϕ between two successive discrete transitions when the flow of the hybrid systems is (ϵ, ϕ)-monotonic and its guards are d_{min}-separable. The only exception is when the discrete state does not change and the reset is the identity map. However in this case we can behave as if the transition was never taken. In all other cases, condition (i) avoids Zeno behaviors in a discrete self-loop, and condition (ii) ensures that we cannot have infinitely fast switching along ϕ when the guards are separable. To see the last remark, if the reset itself changes the value of the continuous state enough to move it to another guard, then $\|x_2 - x_1\|$ will be at least d_{min}. Hence the distance traveled along ϕ would be at least ϵd_{min}. Otherwise, suppose $\|x_2 - x_1| < d_{min}$, it moves at least $\phi \cdot (x_2 - x_1)$ along ϕ which is at least $\epsilon \|x_2 - x_1\|$, and it needs to travel a minimum of $(d_{min} - \|x_2 - x_1\|)$ before taking the next transition. But since the flow is (ϵ, ϕ)-monotonic, it moves another $\epsilon(d_{min} - \|x_2 - x_1\|)$ at least along ϕ. Hence it moves at least ϵd_{min} in total.

4 STORMED Hybrid Systems

In this section we formally introduce the special class of hybrid systems that we study in this paper, and show that they admit a finite bisimulation.

Definition 6 (STORMED Hybrid Systems) *A* STORMED hybrid system *is a tuple* $(\mathcal{H}, \mathcal{A}, \phi, b_-, b_+, d_{min}, \epsilon, \zeta)$ *where* $\mathcal{H} = (Q, \Delta, X, X_0, q_0, \mathcal{I}, \mathcal{F}, \mathcal{R}, \mathcal{G})$ *is a hybrid automaton,* $\mathcal{A}$ *is an o-minimal structure,* $b_-, b_+, d_{min} \in \mathbb{R}$*, and* $\phi \in X$ *is a vector such that the following conditions are satisfied:*

(S) The guards of $\mathcal{H}$ *are* d_{min}*-***S***eparable.*
(T) The flows of $\mathcal{H}$ *are* **T***ISC.*
(O) $\mathcal{H}$ *is definable in the* **O***-minimal structure* $\mathcal{A}$*.*
(RM) **R***esets and flows* $\mathcal{F}_{(\cdot,\cdot)}(\cdot)$ *are* **M***onotonic:* (ϵ, ζ, ϕ)*-monotonic and* (ϵ, ϕ)*-monotonic respectively.*

(ED) **E***nds are* **D***elimited: for all* $(p,q) \in \Delta$ *we have* $\{\phi \cdot x : x \in \mathcal{G}_{(p,q)} \in (b_-, b_+)$ *meaning that the projection of each of the guard sets on* ϕ *is bounded below by (or is greater than)* b_- *and bounded above by (or is less than)* b_+*.*

Before we turn to proving our main result on the existence of a bisimulation for the STORMED systems, we will introduce a few definitions and a lemma to aid the proof.

Definition 7 *Given a partition* $\mathcal{V}$ *of* $Q \times X$*, define* $F_t^\star(\mathcal{V})$ *to be the coarsest bisimulation*[1] *with respect to* $\rightarrow_t$ *that respects* $\mathcal{V}$*. Further, define* $F_d(\mathcal{V}) := \{(s_1, s_2) | \big(\exists s_1' \ . \ s_1 \rightarrow_d s_1'\big) \Rightarrow \big(\exists s_2' \ . \ s_2 \rightarrow_d s_2' \wedge s_1' \mathcal{V} s_2'\big)\} \cap \mathcal{V}$*.*

It can be easily observed that (a) The functionals $F_t^\star(\cdot)$ and $F_d(\cdot)$ are monotonic; (b) $F_t^\star(\mathcal{V})$ is a refinement of $\mathcal{V}$ and so is $F_d(\mathcal{V})$, i.e. $F_t^\star(\mathcal{V}) \subseteq \mathcal{V}$ and $F_d(\mathcal{V}) \subseteq \mathcal{V}$; (c) $F_t^\star(\cdot)$ is idempotent, i.e. $F_t^\star(F_t^\star(\mathcal{V})) = F_t^\star(\mathcal{V})$

Definition 8 *For a hybrid system, we define the* i*-th neighborhood* $N_i \in Q \times X$ *to be the set of all locations starting from which there is no execution that can have more than* i *non-trivial*[2] *discrete transitions. Note that* $N_{i+1} \supseteq N_i$*.*

Lemma 9 *Given a STORMED Hybrid System* $(\mathcal{H}, \mathcal{A}, \phi, b_-, b_+, d_{min}, \epsilon, \zeta)$ *and a partition* $\mathcal{P} = \{P_1, P_2, ..., P_k\}$ *of its state space* $Q \times X$*, let* $\cong$ *to be a bisimulation relation on* $\mathcal{H}$ *refining* $\mathcal{P}$*. Define a sequence of partitions* $\{W_0, W_1, \ldots\}$ *inductively by setting* $W_0 = F_t^\star(\mathcal{P})$ *and* $W_{i+1} = F_t^\star(F_d(W_i))$*. The following hold for all* $i \geq 0$*:*

(a) W_i *is a finite partition definable in the o-minimal theory,*
(b) $\cong \subseteq W_i$*, and*
(c) W_i *is a bisimulation on locations in the* i*-th neighborhood* N_i *that respects* $\mathcal{P}$*.*

Proof: The proof follows by an induction on i. We show that if $(q_1, x_1) W_i (q_2, x_2)$, then (q_1, x_1) simulates (q_2, x_2) and vice versa. Details are given in the Appendix A.

Lemma 10 *Given a STORMED Hybrid System* $(\mathcal{H}, \mathcal{A}, \phi, b_-, b_+, d_{min}, \epsilon, \zeta)$*, any execution of the system can have at most* $i^\star = \lceil \frac{b_+ - b_-}{\eta} \rceil$ *non-trivial discrete transitions, where* $\eta := min\{\zeta, \epsilon d_{min}\}$*.*

Proof: The proof follows from some simple observations extending Remark 5. Details are given in the Appendix B.

[1] The coarsest bisimulation with respect to a subset of the transition relation $\rightarrow' \subseteq \rightarrow$ is the coarsest partition $\mathcal{P} = \{P_i\}$ of the state space $Q \times X$ such that $\mathcal{P}$ is a bisimulation relation of the transition system given by $(Q \times X, q_0 \times X_0, \rightarrow')$.

[2] We ignore the non-trivial (identity) discrete transitions, i.e. $(q, x) \rightarrow_d (q, x)$, which are allowed by monotonic resets because they are trivial and can be omitted for our purposes.

Theorem 11 (Finite Bisimulation) *The transition system induced by a STORMED hybrid system* $(\mathcal{H}, \mathcal{A}, \phi, b_-, b_+, d_{min}, \epsilon, \zeta)$ *has a finite bisimulation that respects any* $\mathcal{A}$*-definable partition* $\mathcal{P}$*. Moreover, if* $\mathcal{A}$ *is decidable, then there is an effective algorithm for constructing that bisimulation.*

Proof: Again, let $\eta := \min\{\zeta, \epsilon d_{min}\}$ and $i^\star := \lceil \frac{b_+ - b_-}{\eta} \rceil$. We can simply observe that since, by Lemma 10, any execution in a STORMED system can go through at most $i^\star$ discrete transitions, all reachable states belong to $N_{i^\star}$. Therefore, by Lemma 9, $W_{i^\star}$ is a bisimulation for all reachable states in $Q \times X$, it respects $\mathcal{P}$ and it is definable in $\mathcal{A}$. Therefore, if $\mathcal{A}$ is decidable, there exists an effective algorithm for constructing $W_{i^\star}$. ■

Corollary 12 (Reachability) *Given a STORMED hybrid system* $(\mathcal{H}, \mathcal{A}, \phi, b_-, b_+, d_{min}, \epsilon, \zeta)$*,*

1. *the set-to-set reachability problem, which is, given two sets* $S_1, S_2 \subseteq Q \times X$*, if there is a point in* S_1 *that can reach some point in* S_2*, is decidable, if* $\mathcal{A}$ *is.*
2. *Claim 1 is true even if the guards are not delimited, as long as the initial conditions satisfy* $\{\phi \cdot x : \exists q \in Q \ .(q, x) \in S_1\} \in [b_-, \infty]$ *and the final set satisfies* $\{\phi \cdot x : \exists q \in Q \ .(q, x) \in S_2\} \in [-\infty, b_+]$*.*

Proof: First note that Claim 2 reduces to Claim 1 since there can be no discrete transitions outside the set of states $\{(q, x) : x \in [b_-, b_+], q \in Q\}$ that can reach the set S_2. Therefore we can restrict all guards along ϕ to $[b_-, b_+]$ and be able to answer the same question. To check reachability of a set $S_2 \subseteq Q \times X$ from a non-intersecting set S_1, we can partition the state space to $\mathcal{P} = \{S_1, S_2, Q \times X \setminus (S_1 \cup S_2)\}$ and get a finite bisimulation that respects $\mathcal{P}$. This is possible because of Theorem 11. The reachability problem then reduces to the reachability problem of a finite automaton which is constructible if $\mathcal{A}$ is decidable, and hence the reachability problem for STORMED hybrid systems is decidable. ■

5 Examples of STORMED Hybrid Systems

We believe that STORMED hybrid system model will be useful in modeling many system models. The constraints imposed by STORMED hybrid systems are realized in some physical systems as follows.

- Monotonicity can be associated with energy or time depletion, or in vehicle control problems, with non-decreasing trajectories.
- The Ends-Delimited property can be present as a deadline on the monotonic direction or a spatial confinement.
- Separability of guards represents infrequency in making control decisions, also based on location or time.
- TISC flows arise naturally, whereas o-minimality is not necessarily a common property, but can be used as an approximation most of the time. Linearization and other model reductions may also result to o-minimal realizations.

In Appendix C, we give a toy example illustrating how the characteristics of a physical system map to the constraints imposed by a STORMED hybrid systems.

6 Relaxations of the STORMED model

In this section we show that relaxing the various constraints of the STORMED model makes the reachability problem undecidable, and thus justify the tightness of our definition of STORMED model. We consider TISC property of the flows and o-minimal definablility of the system as intrinsic to our model. The theorem below identifies relaxations which render the model undecidable.

Theorem 13 *1. The reachability problem of the STORMED model with the constraint on the monotonicity of resets removed is undecidable.*
2. The reachability problem of the STORMED model with the constraint on the ends being limited removed is undecidable.

Proof. We first present a proof of the undecidability of the reachability problem of multi-rate timed automata along the lines of [1] , and then describe how it can be modified to serve our purpose. Multi-rate timed automata can simulate two counter-machines thus reducing the reachability problem for two counter-machines to that of multi-rate automata. Consider a 2 counter-machine M with counters C and D. In the multi-rate automaton A simulating it, there are two variables x and y which store the values corresponding to the values of the counters. A counter value of n is stored in the corresponding variable as $1/2^n$. Hence an increment will halve the value of the variable and similarly a decrement will double the value. The execution of A will synchronize with that of M every two time units in the sense that if the i-th configuration of M points to location p with the two counter values m and n, then A at time instant $2i$ will be in state p with values of counters $1/2^m$ and $1/2^n$. The parts of the automaton corresponding to the operations increment, decrement and test for 0 is given in Figure 1. Here g is a variable which keeps track of the global time. All variables not shown are assumed to have a flow of 0.

Observe that automaton A satisfies all the STORMED constraints except monotonic resets and separable guards. In order to prove part 1 of the above theorem we modify A to obtain A_1 such that A_1 simulates M but has separable guards. With every state q we associate a distinct even number h_q. We introduce a new variable v, and include in the transition going out of p a constraint $v \in (h_p, h_p + 1]$. If there is only one transition going out of p' we add to its guard the constraint $v \in (h_{p'}, h_{p'} + 1]$, otherwise we add to the transition going from p' to q the constraint $v = h_{p'} + 1$, and to the transition going from p' to r the constraint $v \in (h_{p'}, h_{p'} + 1/2]$. We have three more variables g', x' and y' whose values equal that of g, x and y, respectively, while entering any state. However the values of x' and y' do not change while in state p and the value of g' does not change in state p'. It is easy to see that this can be ensured by treating the variables x', y' and g' similar to x, y and g respectively everywhere,

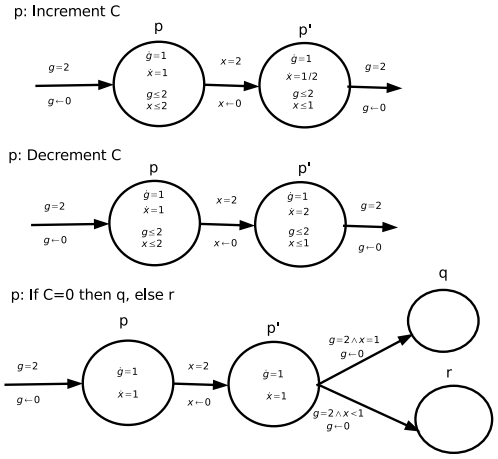

Fig. 1. The parts of the multi-rate automaton A corresponding to the operations increment, decrement and test for zero of the 2-counter machine M.

except that in state p, $\dot{x}' = 0$ and $\dot{y}' = 0$ and in state p', $\dot{g}' = 0$. Finally we set $\dot{v} = h_p/(2 - x') + x'/(2 - x')$ in state p corresponding to an operation on C. In state p' we set $\dot{v} = h_{p'}/(2 - g') + g'/(2 - g')$. Hence the value of v upon exiting p would be $h_p + v_1$ and that upon exiting p' would be $h_{p'} + v_1$ where v_1 is the value of x when entering p. At any point of time the transitions that are enabled in A_1 is the same as that of A.

Now returning to part 2 of the theorem, we show how we can construct the automaton A_2 which restores the monotonicity of resets. However the ends will no more be delimited. A_2 is obtained from A_1 by adding a new variable n which increases monotonically at rate 1. The monotonicity is now along the flow of n. This proves the above theorem. ■

Relaxing combinations of the STORMED constraints causes undecidability at very low dimensions. Without separability of guards and ends-delimited we have undecidability in 4 dimensions. This follows from the results of [3] where piecewise constant derivatives (PCD) with delimited ends in 3 dimensions is shown undecidable. PCD flows are not monotonic but they can be made monotonic by introducing a fourth dimension along which the flows are monotonic. The results in [3] also imply that the reachability problem for STORMED h.s. without guard separability or monotonicity is undecidable in 3 dimensions. With just the relaxation on separability of guards, it follows from the results in [13] that finite bisimulation does not exist even in two dimensions.

7 Conclusions

We introduced a new class of hybrid automata called STORMED hybrid systems and showed that they admit a finite bisimulation. Further, the bisimulation is constructible if the o-minimal theory in which the elements of the system are defined is decidable. STORMED automata allow the continuous variables to have rich dynamics, while at the same time not decoupling the discrete states. However, STORMED hybrid systems require monotonic flows/resets and separable guards. But such constraints are often present in real systems, for example, monotonicity appears in the form of a depleting resource. We also demonstrated that the relaxation of certain constraints from the STORMED hybrid system model results in a model that is undecidable. In the future it would be useful to build a tool to algorithmically analyze systems described as STORMED hybrid system, and evaluate its performance on models of embedded systems.

References

1. R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138(1):3–34, 1995.
2. Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
3. Eugene Asarin, Oded Maler, and Amir Pnueli. Reachability analysis of dynamical systems having piecewise-constant derivatives. *Theoretical Computer Science*, 138(1):35–65, 1995.
4. Vincent D. Blondel, Olivier Bournez, Pascal Koiran, Christos H. Papadimitriou, and John N. Tsitsiklis. Deciding stability and mortality of piecewise affine dynamical systems. *Theoretical Computer Science*, 255(1–2):687–696, 2001.
5. P. Bouyer, T. Brihaye, and F. Chevalier. Weighted o-minimal hybrid systems are more decidable than weighted timed automata! In *Proceedings of the Symposium on Logical Foundations of Computer Science (LFCS'07)*, volume 4514 of *LNCS*, pages 69–83, New-York, NY, USA, June 2007. Springer.
6. Thomas Brihaye and Christian Michaux. On the expressiveness and decidability of o-minimal hybrid systems. *J. Complexity*, 21(4):447–478, 2005.
7. E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 2000.
8. E.Asarin, G.Schneider, and S.Yovine. Algorithmic analysis of polygonal hybrid systems, part i: Reachability. *Theor. Comput. Sci.*, 379(1-2):231–265, 2007.
9. Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. What's decidable about hybrid automata? In *Proc. 27th Annual ACM Symp. on Theory of Computing (STOC)*, pages 373–382, 1995.
10. G. Lafferriere, G. Pappas, and S. Sastry. O-minimal hybrid systems, 1998.
11. Robin Milner. *Communication and Concurrency*. Prentice-Hall, Inc, 1989.
12. V. Mysore and A. Pnueli. Refining the undecidability frontier of hybrid automata. In *Proceedings of the International Conference on the Foundations of Software Technology and Theoretical Computer Science*, pages 261–272, 2005.
13. P. Prabhakar, V. Vladimerou, M. Viswanathan, and Dullerud E. A decidable class of planar linear hybrid systems. Technical report, University of Illinois at Urbana-Champaign, 2008. UIUCDCS-R-2008-2927.

14. Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry.* University of California Press, 2nd edition, 1951.
15. L. van den Dries and C. Miller. On the real exponential field with restricted analytic functions. *Israel Journal of MAthematics*, (85):19–56, 1994.
16. Lou van den Dries. *Tame Topology and O-minimal Structures.* Cambridge Univesity Press, 1998.

A Proof of Lemma 9

Claim (a): Proof by induction on i. We know that for any finite partition $\mathcal{P}$ definable in the o-minimal structure $\mathcal{A}$, $F_t^\star(\mathcal{P})$ has only finitely many equivalence classes and is definable in the o-minimal theory [6]. Therefore W_0 is o-minimal definable, hence the claim is true for $i = 0$. It is easy to see from the definition of F_d that it is also definable. In addition, for any partition $\mathcal{P}$ which is definable in the o-minimal theory, $F_d(\mathcal{P})$ has finitely many equivalence classes because there are finitely many discrete transitions possible from each part of $\mathcal{P}$. Thus, from these observations we have W_{i+1} is definable in the o-minimal theory and has finitely many equivalence classes, if W_i has, since $W_{i+1} = F_t^\star(F_d(\mathcal{P}))$.

Claim (b): Proof by induction on i. Case $i = 0$: $\cong \subseteq \mathcal{P}$. Since $F_t^\star(\cdot)$ is monotonic, $F_t^\star(\cong) \subseteq F_t^\star(\mathcal{P})$. But since $F_t^\star(\cong) = \cong$, we have $\cong \subseteq F_t^\star(\mathcal{P})$. Hence $\cong \subseteq W_0$.

Case $i \geq 1$: By induction hypothesis $\cong \subseteq W_{i-1}$. By monotonicity of the functionals $F_t^\star(\cdot)$ and $F_d(\cdot)$, we have $F_t^\star(F_d(\cong)) \subseteq F_t^\star(F_d(W_{i-1}))$. But since $\cong$ is a bisimulation, $F_t^\star(F_d(\cong)) = \cong$. Hence $\cong \subseteq F_t^\star(F_d(W_{i-1}))$ and therefore $\cong \subseteq W_i$.

Claim (c): We will prove the claim by induction on i. Case $i = 0$: Let $(q, x), (p, y) \in N_0$, and $(q, x)W_0(p, y)$. Suppose $(q, x) \to (q_1, x_1)$. Since $(q, x) \in N_0$, it cannot take a discrete transition, hence $(q, x) \to_t (q_1, x_1)$. But since $W_0 = F_t^\star(\mathcal{P})$, $(q, x)F_t^\star(\mathcal{P})(\bigvee, \dagger)$ and hence there exist (p_1, y_1) such that $(p, y) \to_t (p_1, y_1)$. Therefore (p, y) simulates (q, x). We can argue similarly that (q, x) simulates (p, y). Hence W_0 is a bisimulation on N_0 refining $\mathcal{P}$.

Case $i > 1$: By induction hypothesis W_i is a bisimulation on all locations in N_i refining $\mathcal{P}$. We need to prove that W_{i+1} is a bisimulation relation on all locations in N_{i+1}. W_{i+1} is a refinement of $\mathcal{P}$ since W_i is a refinement of $\mathcal{P}$ and $W_{i+1} = F_t^\star(F_d(\mathcal{P}))$. Given two locations (q, x) and (p, y) in N_{i+1} that satisfy $(q, x)W_{i+1}(p, y)$. We will prove that (p, y) simulates (q, x) and by symmetry, the reverse will also be true.

(1) Suppose $(q, x) \to_t (q_1, x_1)$. Since $W_{i+1} = F_t^\star(F_d(W_i))$ we know that there is a (p_1, y_1) such that $(p, y) \to_t (p_1, y_1)$ and $(q_1, x_1)W_{i+1}(p_1, y_1)$. In addition both (q_1, x_1) and (p_1, y_1) will still be in N_{i+1} since no discrete transition has occurred.
(2) Suppose $(q, x) \to_d (q_1, x_1)$. Since $W_{i+1} = F_t^\star(F_d(W_i)) \subseteq F_d(W_i)$ we know that there is a (p_1, y_1) such that $(p, y) \to_d (p_1, y_1)$ and $(q_1, x_1)F_d(W_i)(p_1, y_1)$. Note that both (q_1, x_1) and (p_1, y_1) will now be in N_i. We also know that since $F_d(W_i) \subseteq W_i$ we have $(q_1, x_1)F_d(W_i)(p_1, y_1) \Rightarrow (q_1, x_1)W_i(p_1, y_1)$. By

the induction hypothesis, we then have $(q_1, x_1) \cong (p_1, y_1)$. Further from Claim (b), we have $\cong \subseteq W_{i+1}$, and hence $(q_1, x_1) W_{i+1} (p_1, y_1)$.

Therefore W_{i+1} is a bisimulation relation on all locations in N_{i+1} which respects $\mathcal{P}$. This concludes the induction proof of Claim (c). ■

B Proof of Lemma 10

We can prove this by showing there is a minimum distance the continuous part of the state travels along ϕ between two consecutive discrete transitions. Indeed, after a non-trivial discrete transition the reset advances the continuous state along ϕ by ζ or some distance proportional by ϵ to the euclidean distance of the total change. This is true by the definition of monotonic resets. In the latter case, the discrete state changes and there cannot be another discrete transition until the continuous part travels enough to reach another guard (separability). By monotonicity of the flow, the distance traveled along ϕ will be ϵd_{min}. Therefore, the minimum distance traveled along ϕ between two discrete jumps is $\eta := \min\{\zeta, \epsilon d_{min}\}$. Since, by the ends-delimited property we can only have discrete transitions along $\{(q, x) : \phi \cdot x \in (b_-, b_+)\}$ we can conclude that we can have at most $i^\star = \lceil \frac{b_+ - b_-}{\eta} \rceil$ discrete transitions in any execution. ■

C Storm Landing Problem

Consider a situation where a small airplane is to land on an isolated airstrip, and a moving storm system is obstructing the landing. The storm is simulated by 3 cloud groups, as in fig. C. There are 3 airports around the airstrip and controllers from each suggest trajectories to the airplane if it lies within their communication range. The suggested trajectories change discretely 3 times for each airport as time goes by. The question is if the small airplane can land before it reaches a low-fuel limit. Not having enough fuel will force it return to an airport away from the storm and miss its original destination.

There are 10 discrete states in Q. States 1, 2, and 3 represent the configuration where the airplane is following different directions from tower 1. States 4,5,6 directions from tower 2, and states 7,8,9 from tower 3. The initial state is $Q_0 = 1$. The final state $F = 10$ represents the situation where the airplane has landed at his destination. We have a fully interconnected graph for out discrete transitions, i.e. $\Delta = \{(q_1, q_2) : q_1, q_2 \in \{1, \ldots, 10\}\}$.

The continuous part of the state X lies in $\mathbb{R}^{14}$. It Consists of:

- the location of the airplane $x_a \in \mathbb{R}^3$,
- the location of the 3 cloud groups x_{s1}, x_{s2}, x_{s3} each in $\mathbb{R}^3$, and
- the gasoline fuel remaining in the airplane $x_g \in \mathbb{R}$
- time $x_t \in \mathbb{R}$

We will be using x as a symbol for the continuous state vector $(x_a, x_{s1}, x_{s2}, x_{s3}, x_g, x_t)$. Our X_0 has $x_t = 0$ and the other components initialized to some other constants.

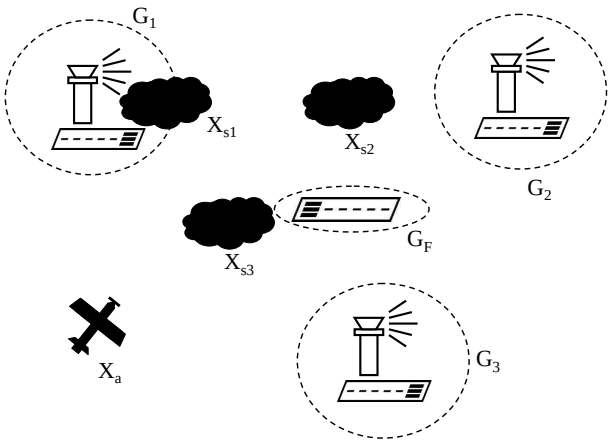

Fig. 2. An illustration of the STORMED storm system showing guards and continuous states.

Each component of x has a different expression for its respective flow. The airplane location component x_a follows airplane dynamics (to the degree simulated by polynomial trajectories), x_t will always follow time $\mathcal{F}_{(\cdot,x_t)}(t) = x_t + t$, gasoline fuel will always decrease at rate ρ liters per second $\mathcal{F}_{(\cdot,x_g)}(t) = x_g - \rho t$, starting at full tank capacity C at time 0, and x_{s1}, x_{s2}, x_{s3} will follow a TISC o-minimal flow according to weather prediction, again independent of the discrete state (only time affects the weather).

The o-minimal structure for this STORMED hybrid system is $\mathcal{M} = \{\mathbb{R}, +, \cdot, \leq\}$. All suggested flows $\mathcal{F}_{(i,\cdot)}(\cdot), i = 1, \ldots, 9$ are TISC polynomial trajectories and lead to the final guard $G_{\star,F} = \{x : (x_a = X_A) \wedge (x_g > L)\}$, where X_A is a constant the location of the destination airstrip and it is separated from all other guards. The rest of the guards are separable by half-planes in time or space. Recall that we have 3 discrete states for each tower, so their guards are only separated in the x_t component. This establishes separability. More specifically:
$G(\star, i) = \{x : (x_t \in T_i) \wedge (x_a \in R_i)\}$
Where $T_i = \begin{cases} [0, 1800) \text{ for } i = 1, 4, 7 \\ [1801, 3600) \text{ for } i = 2, 5, 8 \\ [360, 5400) \text{ for } i = 3, 6, 9 \end{cases}$, with values in seconds are the times when the control towers change their flow directions and
$R_i = \begin{cases} \{x : \|x_a - S_1\| < 10000\} \text{for } i = 1, 2, 3 \\ \{x : \|x_a - S_2\| < 10000\} \text{for } i = 4, 5, 6 \\ \{x : \|x_a - S_3\| < 10000\} \text{for } i = 7, 8, 9 \end{cases}$, where $S_1, S_2, S_3 \in \mathbb{R}^3$ are the lo-

cations of the control towers of the surrounding airports with a range of 10km each. The invariants are $I_1 = \ldots = I_F = \{x : (x_g > L) \wedge (\|x_a - x_{s1}\| > D) \wedge (\|x_a - x_{s2}\| > D) \wedge (\|x_a - x_{s3}\| > D)\}$, where the constant L is the lowest fuel level allowed and the constant D is the minimum safe distance from the airplane to a cloud center.

By construction all constraints of a STORMED h.s. are satisfied, with the exception of delimited guards: TISC flows, o-minimality, reset-free trajectories, monotonicity along the decreasing fuel component x_g, and finally the end guard is delimited by the minimum fuel constant L. Since the initial and final states' projections on the monotonic direction are bounded on the right and left respectively and since the the theory of $\mathcal{M}$ is decidable, by cor. 12 there is an effective algorithm to determine the destiny of the aircraft.