On Hybrid Systems and the Modal μ-Calculus (extended abstract)*

J. M. Davoren
Center for Foundations of Intelligent Systems
626 Rhodes Hall, Cornell University
Ithaca, NY 14853–3801, USA
davoren@cornell.edu

Introduction

It is hardly controversial to claim that the *μ-calculus* is a formal logic of central import for the analysis and verification of hybrid automata and related classes of systems. The fundamental concepts of *reachability* and *invariance* for hybrid trajectories are expressible in terms of fixed-points of operators mapping sets of states to sets of states, and thus definable in the μ-calculus. The iterative computation of the denotation of such fixed point formulas lies at the heart of symbolic model checking tools for hybrid and real-time systems such as HyTech (Alur, Henzinger, & Ho 1996), (Henzinger 1996) and Kronos (Daws *et al.* 1996). More generally, the propositional μ-calculus is well-recognized as a richly expressive logic over transition system models: the power of its fixed-point quantifiers is such that it subsumes virtually all temporal, modal and dynamic logics (Emerson 1997), (Janin & Walukiewicz 1996).

However, the current practice, within the allied field of automated verification of (discrete) reactive systems as well as within the hybrid systems community, is to treat the μ-calculus not as a working or usable logic but rather as a logic of the substratum. It provides a common "machine" language and semantics for verification by model checking over transition system models, with user-input specifications written in the more "natural" languages of temporal logics, and then translated into that of the μ-calculus (Kupferman & Vardi 1998), (Henzinger 1996).

In the hybrid and real-time systems literature, the bulk of the current work on logics and formal methods is an expansion of, but firmly anchored in, the framework of temporal logic verification of finite state systems (see, for example, (Manna & Pnueli 1993), (Henzinger 1996), (Manna & Sipma 1998)). The core computational model is that of a hybrid automaton, which is represented formally as a *labeled transition system* over a hybrid state space $X \subseteq Q \times \mathbb{R}^n$, where Q is a finite set of discrete modes, and $\mathbb{R}^n$ is Euclidean space. Both sorts of system dynamics – both continuous evolution within a control mode, and the effects of discrete jumps between control modes – are uniformly represented as binary transition relations $r \subseteq X \times X$. Within the temporal logic framework, the principal focus has been on the formal specification of dynamic properties of classes of trajectories of hybrid automata such as safety/invariance, or its dual of reachability (Henzinger 1996), (Manna & Sipma 1998).

In this paper, we take the basic hybrid automaton and its standard transition system model from (Henzinger 1996), (Lafferriere, Pappas, & Sastry 1998), and examine them afresh.

We work in the *modal* rather than the temporal variant of the μ-calculus (see (Stirling 1992) §4), which includes in its formal language a pair of modal operators $\langle a \rangle$ and $[a]$ for each of the component transition relations, interpreted by the familiar relational *pre-image* (or *post-image*) operators on sets. The modal framework provides a *modular* specification language. Classes of hybrid trajectories can be simply described by sequences of alternating compositions of evolution and jump relations. We show how to write clean *and* "human-readable" formulas of μ-calculus expressing safety and liveness properties of hybrid trajectories.

The shift of mindset to that of modal logic opens up a wealth of new possibilities. The key move is to view a transition system model not merely as some form of "discrete abstraction" (Henzinger 1996), but rather as a skeleton which can be fleshed out by imbuing the state space with *topological*, *metric tolerance* or other structure; we then represent such structure by enriching the language of the μ-calculus with special-purpose modal operators. In the resulting logical formalisms, we can simply and clearly express what we mean by *continuous* and *discrete* dynamics, and hybrids of the two. We can formally express *topological* concepts, such as the topological *interior*, *closure* or *boundary* of a set, or notions of *imprecision* or *metric tolerance*, such as the property of "being within distance ϵ" of a set, for a given $\epsilon > 0$. By viewing transition relations $r \subseteq X \times X$ in their equivalent form as *set-valued maps* $r : X \rightsquigarrow X$, i.e. functions $r : X \to \mathcal{P}(X)$, and drawing on the resources of set-valued analysis and dynamical systems

*The full version of this paper appears in P. Antsaklis *et al.* (eds.), *Hybrid Systems V*, LNCS **1567**. Springer-Verlag, Berlin, 1999. 38–69. The paper is available online at: `ftp://cam.cornell.edu/pub/davoren/davoren.html`

theory (Aubin & Frankowska 1990), (Akin 1993), we open the way to a richer formal analysis of robustness and stability properties for hybrid automata and related classes of systems.

A further advantage of the modal framework is that it supports not only the specification and verification of single properties, but the larger task of representing and building up a *knowledge base* of properties of a system, starting with structural properties assumed in the modeling, and then adding new facts as they are verified by either model-checking or deductive means. Building on the work of (Kozen 1983), (Walukiewicz 1996) and (Ambler, Kwiatkowska, & Measor 1995), we show that the modal μ-calculus and various of its normal polymodal extensions have sound and complete axiomatic proof systems.

We can also provide a clean account of the relationship between propositional modal (and thus in general second-order) *specification languages* for expressing system properties, and first-order *system description languages*. From the very recent work of (Lafferriere, Pappas, & Sastry 1998), we can assume each of components of a hybrid automaton have explicit first-order definitions in the language $\mathcal{L}(\overline{\mathbb{R}})$ of an *o-minimal structure* $\overline{\mathbb{R}} = (\mathbb{R}; <, +, -, \cdot, 0, 1, ...)$ expanding (or a substructure of) the real-closed field (van den Dries 1998). (In the hybrid automata literature, the adjective "linear" means the components are all definable in the first-order language $\mathcal{L}(<, +, -, 0, 1)$ of the reals with only order, addition and integer constants.) Clarifying the work on *finite bisimulation quotients* and decidability of temporal verification in (Lafferriere, Pappas, & Sastry 1998), we show how and when modal μ-calculus sentences can be translated into first-order formulas.

This work is one installment of a larger project. An analysis of the concept of bisimulation, and its relation to the algebraic semantics for the μ-calculus, is given in (Davoren 1998b), and (Davoren 1999) is a full treatment of completeness of deductive proof systems for normal polymodal extensions of the μ-calculus. Related logics and earlier versions of some of the ideas are found in (Davoren 1998a).

The remainder of this extended abstract is a brief tour through the full paper.

Labeled transition systems

Modulo notational variations, the labeled transition system is the common basic model for all propositional temporal and polymodal logics.

Definition 1 *A* modal signature *is a pair* (Φ, Σ), *where* Φ *is a set of* propositional constants (*observation or event labels*), *and* Σ *is a set of* transition labels. *A* labeled transition system (*LTS*), (*or* generalized Kripke model) *of signature* (Φ, Σ) *is a structure*

$$\mathfrak{M} = \left(X, \{a^{\mathfrak{M}}\}_{a \in \Sigma}, \{||p||^{\mathfrak{M}}\}_{p \in \Phi}\right)$$

where:

- $X \neq \varnothing$ *is the state space* (*of* arbitrary *cardinality*);
- *for each transition label* $a \in \Sigma$, $a^{\mathfrak{M}} : X \rightsquigarrow X$ *is a binary relation on* X; *and*
- *for each propositional constant* $p \in \Phi$, $||p||^{\mathfrak{M}} \subseteq X$ *is a fixed subset of* X.

An LTS model $\mathfrak{M}$ is a clean and simple abstraction of a finite automaton. It is an *abstract machine* over state space X, with input or action alphabet Σ, and additionally equipped with an observation alphabet Φ, and the output relation which maps a state x to the set of all atomic propositions $p \in \Phi$ such $x \in ||p||^{\mathfrak{M}}$. Sets of initial or final states can be identified by specific labels in Φ.

Over a topological or metric space X, an LTS $\mathfrak{M}$ is a generalized (set-valued) dynamical system.

Syntax and semantics of the modal μ-calculus

The μ-calculus originated in the 1960's as a formal logic of digital programs, and is formalized in (Kozen 1983). Contemporary introductions to the μ-calculus can be found in (Stirling 1992), (Emerson 1997).

Definition 2 *Fix a modal signature* (Φ, Σ), *and let* PVar *denote a fixed set of propositional (second-order or set-valued) variables. The collection* $\mathcal{F}_\mu(\Phi, \Sigma)$ *of* formulas *of the propositional modal* μ*-calculus is generated by the grammar:*

$$\varphi ::= \mathbf{ff} \mid p \mid Z \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \langle a \rangle \varphi \mid \mu Z.\varphi$$

for $p \in \Phi$, $Z \in$ PVar, *and* $a \in \Sigma$, *with the proviso that in* $\mu Z.\varphi$, *the variable* Z *occur* positively, *i.e. each occurrence of* Z *within the scope of an even number of negations.*

The other (classical) propositional connectives, modalities and greatest fixed point quantifier are defined in the usual way; in particular, $[a]\varphi \doteq \neg\langle a \rangle \neg\varphi$ and $\nu Z.\varphi \doteq \neg\mu Z.\neg\varphi[Z := \neg Z]$.

Let $\mathcal{S}_\mu(\Phi, \Sigma)$ denote the set of all *sentences* of the μ-calculus; that is, formulas containing no free propositional variables Z. And let $\mathcal{S}(\Phi, \Sigma)$ denote the set of all *purely modal* sentences in the signature (Φ, Σ); that is, without any fixed point quantifiers or propositional variables.

The (standard) relational Kripke semantics of the labeled modalities $[a]$ and $\langle a \rangle$ are given by the *universal* and *existential pre-image operators* of the corresponding relations $r = a^{\mathfrak{M}}$. For relations $r : X \rightsquigarrow Y$, and sets $A \subseteq Y$, define:

$$\begin{aligned} \tau(r)(A) &\doteq \{x \in X \mid (\forall y \in Y)[x \xrightarrow{r} y \Rightarrow y \in A]\} \\ \sigma(r)(A) &\doteq \{x \in X \mid (\exists y \in Y)[x \xrightarrow{r} y \wedge y \in A]\} \end{aligned}$$

In the notation of (Henzinger, Kupferman, & Qadeer 1998), $\sigma(r) = \mathrm{pre}[r]$ and $\tau(r) = \widetilde{\mathrm{pre}}[r]$. The semantic readings of the modalities are *forward*-looking:

$$\begin{aligned} [a]\varphi &\rightleftharpoons \text{"}\textit{All } a\text{-successors satisfy } \varphi\text{"} \\ \langle a \rangle \varphi &\rightleftharpoons \text{"}\textit{Some } a\text{-successor satisfies } \varphi\text{"} \end{aligned}$$

In temporal logic, one usually works with the *global* transition relation $R^{\mathfrak{M}} = \bigcup_{a\in\Sigma} a^{\mathfrak{M}}$ (standardly assumed to be *total*), and the labeled modalities are replaced by global temporal "next" operators, written $\forall X$ or $\forall\bigcirc$, and $\exists X$ or $\exists\bigcirc$.

In an LTS model $\mathfrak{M}$, sentences $\varphi \in \mathcal{S}_\mu(\Phi,\Sigma)$ denote *sets of states* $\|\varphi\|^{\mathfrak{M}} \subseteq X$. A sentence is *true* in $\mathfrak{M}$, written $\mathfrak{M} \models \varphi$, iff $\|\varphi\|^{\mathfrak{M}} = X$, or equivalently, $\|\neg\varphi\|^{\mathfrak{M}} = \varnothing$. The propositional connectives $\neg$, $\wedge$ and $\vee$ are interpreted by set-theoretic complement, intersection and union, and the labeled modalities are interpreted by the pre-image operators. In particular, $\|\mathtt{tt}\|^{\mathfrak{M}} = X$ for all $\mathfrak{M}$, and for implications, $\mathfrak{M} \models \varphi \to \psi$ exactly when $\|\varphi\|^{\mathfrak{M}} \subseteq \|\psi\|^{\mathfrak{M}}$.

Formulas $\varphi \in \mathcal{F}_\mu(\Phi,\Sigma)$ with free variables denote sets $\|\varphi\|^{\mathfrak{M}}_\xi \subseteq X$, relative to a *variable assignment* $\xi : \mathrm{PVar} \to \mathcal{P}(X)$. The semantics of the μ and ν quantifiers are given in terms of the least and greatest fixed points of operators $\varphi^{\mathfrak{M}}_{\xi,Z} : \mathcal{P}(X) \to \mathcal{P}(X)$ defined by:

$$\left(\varphi^{\mathfrak{M}}_{\xi,Z}\right)(A) \triangleq \|\varphi\|^{\mathfrak{M}}_{\xi(A/Z)}$$

where $\xi(A/Z)$ is the variant assignment which is the same as ξ except for assigning the set A to Z. The syntactic restriction on formulas $\mu Z.\varphi$ ensures that the operator $\varphi^{\mathfrak{M}}_{\xi,Z}$ is $\subseteq$-monotone. The completeness of $\mathcal{P}(X)$ as a lattice ensures (by the Hitchcock-Park fixed-point theorem) that the set $\|\mu Z.\varphi\|^{\mathfrak{M}}_\xi$ may also be characterized as a transfinite union of an $\subseteq$-chain of approximations $\|\mu Z.\varphi\|^{\mathfrak{M}}_{\xi,\alpha}$ for ordinals α (of cardinality $\leq Card(X)$), beginning with the empty set, applying $\varphi^{\mathfrak{M}}_{\xi,Z}$ at successor ordinals and taking unions at limits. When this operator is *ω-chain-additive*, i.e. distributes over unions of countable $\subseteq$-increasing chains of sets, the ordinal of convergence is at worst ω. In such cases, we have a sequence of approximation formulas $\varphi^0 \triangleq \mathtt{ff}$ and $\varphi^{n+1} \triangleq \varphi[Z := \varphi^n]$, and $\|\mu Z.\varphi\|^{\mathfrak{M}}_\xi = \bigcup_{n<\omega} \|\varphi^n\|^{\mathfrak{M}}_\xi$.

In the hybrid systems context, the classes of LTS models of particular interest are those $\mathfrak{M}$ such that the state space $X \subseteq \mathbb{R}^n$, the transition relations $a^{\mathfrak{M}} \subseteq \mathbb{R}^{2n}$, and the observation sets $\|p\|^{\mathfrak{M}} \subseteq \mathbb{R}^n$, are all first-order definable in some structure $\overline{\mathbb{R}}$ over the reals. For such $\mathfrak{M}$, it is immediate that for all purely modal sentences $\varphi \in \mathcal{S}(\Phi,\Sigma)$, the denotation set $\|\varphi\|^{\mathfrak{M}}_\xi \subseteq \mathbb{R}^n$ is also first-order definable in $\overline{\mathbb{R}}$, based on the straight-forward modal translation using the definitions of the pre-image operators. For μ-sentences $\mu Z.\varphi$, a first-order translation is available provided that the sequence $\|\varphi^n\|^{\mathfrak{M}}$ for $n < \omega$ of denotations of approximation formulas is guaranteed to converge at a *finite* stage. Such is the case when $\mathfrak{M}$ has a *finite bisimulation quotient*: see (Lafferriere, Pappas, & Sastry 1998). Symbolic model-checking tools such as HyTech are predicated on such convergence.

Hybrid automata and their LTS models

We base our discussion on the systems considered in (Lafferriere, Pappas, & Sastry 1998), typically depicted by a graph of the form of Figure 1.

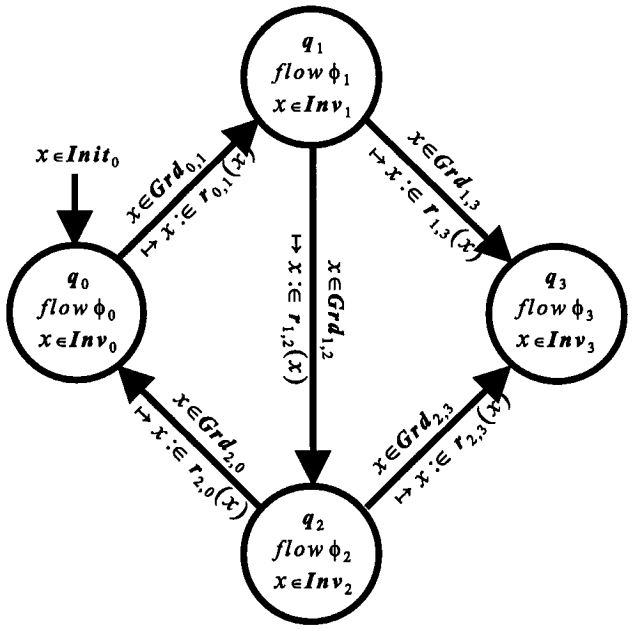

Figure 1: Basic hybrid automaton

Definition 3 *A (basic, evolution time-deterministic)* hybrid system *is a structure $\mathcal{H}$ consisting of the following components:*

- *a finite set Q of* discrete states *or* control modes*;*
- *a* control graph *$G \subseteq Q \times Q$ of* discrete transitions*;*
- *for each $q \in Q$,*
  - *a state space $X_q \subseteq \mathbb{R}^n$ for mode q;*
  - *the continuous (semi-) flow $\phi_q : X_q \times \mathbb{R}^+ \to X_q$ of a vector field on X_q;*
  - *a set $Inv_q \subseteq X_q$ of* invariant *states for mode q – the* domain of permitted evolution *within mode q;*
  - *a set $Init_q \subseteq Inv_q$ of* initial *states for mode q (possibly empty, but not for all q);*
- *for each discrete transition $(q,q') \in G$,*
  - *a set $Grd_{q,q'} \subseteq X_q$, the* guard set *for the jump from q to q';*
  - *a* reset relation *$r_{q,q'} : X_q \rightsquigarrow X_{q'}$; for each $x \in X_q$, $r_{q,q'}(x) \subseteq X_{q'}$ is the set of possible reassignment states after the jump from q to q'.*
- *the* hybrid state space *of the system is the set*

$$X_{\mathcal{H}} = \bigcup_{q\in Q} \{q\} \times X_q$$

For definiteness, we take a *hybrid automaton* to be a hybrid system $\mathcal{H}$ with a *concrete syntactic description*: each of the component sets X_q, $Init_q$, Inv_q,

$Grd_{q,q'} \subseteq \mathbb{R}^n$, semi-flows ϕ_q, and reset relations $r_{q,q'}$ have explicit first-order definitions in the language $\mathcal{L}(\overline{\mathbb{R}})$ of some specified structure over the reals.

Definition 4 (Henzinger 1996), (Lafferriere, Pappas, & Sastry 1998). *Given a hybrid system $\mathcal{H}$, the ("time-abstract") LTS model $\mathfrak{M}_{\mathcal{H}}$ determined by $\mathcal{H}$ has the following components:*

- *the state space $X \triangleq X_{\mathcal{H}}$;*
- *for each discrete state $q \in Q$, the* constrained evolution *("time-step") relation $e_q : X_q \rightsquigarrow X_q$ defined by:*

$$x \xrightarrow{e_q} x' \;\triangleq\; (\exists t \in \mathbb{R}^+) \quad [\; x' = \phi_q(x,t) \wedge (\forall s \in [0,t])\ \phi_q(x,s) \in Inv_q \;]$$

- *for each discrete transition $(q,q') \in G$, the* controlled jump *("discrete-step") relation $c_{q,q'} : X_q \rightsquigarrow X_{q'}$ defined by:*

$$x \xrightarrow{c_{q,q'}} x' \;\triangleq\; x \in Grd_{q,q'} \;\wedge\; x' \in Inv_{q'} \;\wedge\; x \xrightarrow{r_{q,q'}} x'$$

- *the observation sets X_q, $Init_q$, Inv_q, $Grd_{q,q'}$.*

We adopt the notational convention of identifying, when convenient, sets $A_q \subseteq X_q$ and $\{q\} \times A_q \subseteq X$; moreover, the relations $e_q : X_q \rightsquigarrow X_q$ and $c_{q,q'} : X_q \rightsquigarrow X_{q'}$ can be "lifted" to relations $X \rightsquigarrow X$ in the unique obvious way. It is immediate that whenever $\mathcal{H}$ is a hybrid automaton, in the sense above, the LTS model $\mathfrak{M}_{\mathcal{H}}$ is also first-order definable in the same structure.

Figure 2 is an illustration of the operation of a hybrid automaton.

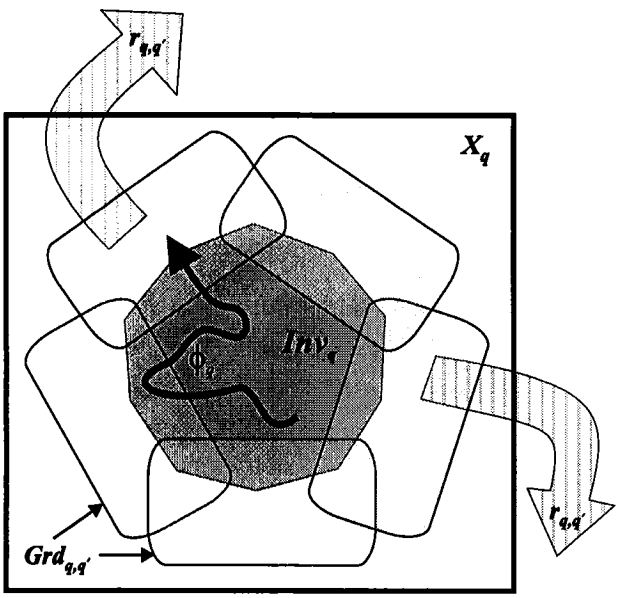

Figure 2: Operation of basic hybrid automaton

The transition alphabet Σ of an LTS model $\mathfrak{M}_{\mathcal{H}}$ will include symbols $\mathbf{e}_q$ for $q \in Q$ and $\mathbf{c}_{q,q'}$ for each edge $(q,q') \in G$, and the observation alphabet Φ will include propositional constants $\mathbf{Init}_q$ and $\mathbf{Inv}_q$ for $q \in Q$, and $\mathbf{Grd}_{q,q'}$ for $(q,q') \in G$.

A trajectory of $\mathcal{H}$ is a finite or infinite sequence $\langle \delta_i, q_i, \gamma_i \rangle_{i \in I}$ such that for each $i \in I$: the duration $\delta_i \geq 0$; the curve $\gamma_i : [0, \delta_i] \to X_{q_i}$ is such that $(q_i, \gamma_i(0)) \xrightarrow{e_{q_i}} (q_i, \gamma_i(t))$ for all $t \in [0, \delta_i]$; $(q_i, q_{i+1}) \in G$; and $(q_i, \gamma_i(\delta_i)) \xrightarrow{c_{q_i,q_{i+1}}} (q_{i+1}, \gamma_{i+1}(0))$. When I is finite, with largest element N, it is allowed that $\delta_N = \infty$. When a hybrid automaton is thought of as a discrete controller interacting with a physical plant, the class of trajectories, so defined, are founded on implicit operational assumptions of temporally continuous and perfect precision sensing, and instantaneous control switches (Henzinger 1996).

Using the modal μ-calculus

The modal sentences:

$$\psi \to [\mathbf{c}_{q,q'}]\varphi \quad \text{and} \quad \psi \to [\mathbf{e}_q]\varphi$$

with the semantic readings "If ψ holds, then all $c_{q,q'}$-successors satisfy φ", and likewise for e_q, correspond precisely to the two types of (temporal logic) safety verification conditions for hybrid systems in (Manna & Pnueli 1993) §4.1. Their Hoare-triple notation is: $\{\psi\}\tau\{\varphi\}$ and $\{\psi\}cont\{\varphi\}$ respectively, where τ ranges over jump transitions and "*cont*" denotes the union of all the evolution relations.

In the LTS model $\mathfrak{M} = \mathfrak{M}_{\mathcal{H}}$, a modal sentence (1):

$$\langle \mathbf{e}_{q_0} \rangle \langle \mathbf{c}_{q_0,q_1} \rangle \langle \mathbf{e}_{q_1} \rangle \langle \mathbf{c}_{q_1,q_2} \rangle \langle \mathbf{e}_{q_2} \rangle \cdots \langle \mathbf{e}_{q_{k-1}} \rangle \langle \mathbf{c}_{q_{k-1},q_k} \rangle \langle \mathbf{e}_{q_k} \rangle \varphi$$

denotes the set of states (q_0, x) from which *some* trajectory with discrete trace $(q_0, q_1, \ldots, q_k)$ *reaches* the set $\|\varphi\|^{\mathfrak{M}} \subseteq X$. Dually, a modal sentence (2):

$$[\mathbf{e}_{q_0}][\mathbf{c}_{q_0,q_1}][\mathbf{e}_{q_1}][\mathbf{c}_{q_1,q_2}][\mathbf{e}_{q_2}] \cdots [\mathbf{e}_{q_{k-1}}][\mathbf{c}_{q_{k-1},q_k}][\mathbf{e}_{q_k}] \varphi$$

denotes the set of states from which *all* $(q_0, q_1, \ldots, q_k)$-trajectories *reach* the set $\|\varphi\|^{\mathfrak{M}}$, upon the last jump c_{q_{k-1},q_k} and remain in $\|\varphi\|^{\mathfrak{M}}$ throughout the last evolution e_{q_k}.

Defining $\mathbf{e}$ and $\mathbf{c}$ to denote the relational sum (union) of, respectively, the relations for the $\mathbf{e}_q$'s for $q \in Q$, and the relations for the $\mathbf{c}_{q,q'}$'s for $(q,q') \in G$, the dynamics of the class of all hybrid trajectories with finite discrete traces are captured by the dual fixed-point definable modalities:

$$\langle \mathbf{h} \rangle \varphi \;\triangleq\; \mu Z.\, \langle \mathbf{e} \rangle \varphi \vee \langle \mathbf{e} \rangle \langle \mathbf{c} \rangle Z$$
$$[\mathbf{h}]\, \varphi \;\triangleq\; \nu Z.\, [\mathbf{e}]\varphi \wedge [\mathbf{e}][\mathbf{c}] Z$$

Since the corresponding semantic operator is ω-chain-additive, the sentence $\langle \mathbf{h} \rangle \varphi$ "unwinds" to the infinite union of all sentences of the form (1); dually, $[\mathbf{h}]\, \varphi$ corresponds to the intersection of all sentences of the form (2). Semantically, $\langle \mathbf{h} \rangle$ and $[\mathbf{h}]$ correspond to the dual $\exists$ and $\forall$ pre-image operators of the *reachability relation* h of the system under the control of $\mathcal{H}$; that is,

$(q,x) \xrightarrow{h} (q',x')$ iff some trajectory $\langle \delta_i, q_i, \gamma_i \rangle_{i \in I}$ with $q_0 = q$ and $\gamma_0(0) = x$ passes through the point (q',x').

We now have the formal linguistic machinery to succinctly express various system specifications. The *safety* sentence

$$\mathbf{Init} \rightarrow [\mathbf{h}]\,\varphi$$

is *true* in the model $\mathfrak{M} = \mathfrak{M}_{\mathcal{H}}$ exactly when every trajectory that starts in the set $\|\mathbf{Init}\|^{\mathfrak{M}}$ *always* remains within $\|\varphi\|^{\mathfrak{M}}$. From the fixed-point rules of Kozen's axiomatization of the μ-calculus, one readily derives an obvious *invariance rule* for hybrid trajectories:

$$\frac{\psi \rightarrow \varphi \qquad \varphi \rightarrow [\mathbf{e}_q]\varphi \qquad \varphi \rightarrow [\mathbf{c}_{q,q'}]\varphi}{\psi \rightarrow [\mathbf{h}]\varphi}$$

the hypotheses holding for all $q \in Q$ and $(q,q') \in G$; c.f. (Manna & Pnueli 1993), (Manna & Sipma 1998).

To express liveness properties, we use the "box-diamond" construct, as in temporal logic. For example, the sentence

$$\varphi \rightarrow [\mathbf{h}]\langle \mathbf{e} \rangle \langle \mathbf{c} \rangle \langle \mathbf{e} \rangle\, \mathbf{tt}$$

is true in $\mathfrak{M} = \mathfrak{M}_{\mathcal{H}}$ exactly when every maximal $\mathcal{H}$ trajectory from a state in $\|\varphi\|^{\mathfrak{M}}$ has an *infinite* discrete trace. This is so because $[\mathbf{h}]\langle \mathbf{e} \rangle \langle \mathbf{c} \rangle \langle \mathbf{e} \rangle\, \mathbf{tt}$ denotes the set of states from which every trajectory with a finite discrete trace can be properly extended. Similarly, the sentence $\varphi \rightarrow [\mathbf{h}]\langle \mathbf{e} \rangle \langle \mathbf{c} \rangle \langle \mathbf{e} \rangle \varphi$ is true in $\mathfrak{M}$ exactly when every trajectory from $\|\varphi\|^{\mathfrak{M}}$ returns to $\|\varphi\|^{\mathfrak{M}}$ via a controlled jump *infinitely often*. And $[\mathbf{h}]\langle \mathbf{h} \rangle \varphi$ denotes the set of states from which every hybrid trajectory *eventually* reaches $\|\varphi\|^{\mathfrak{M}}$.

The modularity of the modal μ-calculus allows us to succinctly express not only *desired properties* – i.e. those to be verified, but also various of the structural properties of an LTS model $\mathfrak{M}_{\mathcal{H}}$ that it will typically possess *by assumption*. The full paper proposes a list of 16 such axioms. As one example, it is immediate from Definition 4 that the controlled jumps relations $c_{q,q'} : X_q \rightsquigarrow X_{q'}$ have the relational decompositions $c_{q,q'} = r_{q,q'} \cap (Grd_{q,q'} \times Inv_{q'})$; this is expressed modally by the sentence scheme:

$$\langle \mathbf{c}_{q,q'} \rangle Z \leftrightarrow \mathbf{Grd}_{q,q'} \wedge \langle \mathbf{r}_{q,q'} \rangle (Z \wedge \mathbf{Inv}_{q'})$$

The constrained evolution relation $e_q : X_q \rightsquigarrow X_q$ may be characterized as a restriction of the (positive) *orbit relation* $f_q : X_q \rightsquigarrow X_q$ of the semi-flow ϕ_q (Akin 1993) given by:

$$x \xrightarrow{f_q} x' \;\triangleq\; (\exists t \in \mathbb{R}^+)\;\; x' = \phi_q(x,t)$$

When the set Inv_q is *convex* with respect to the semi-flow ϕ_q – in the sense that no curve segment of ϕ_q with both endpoints in Inv_q ever leaves Inv_q at an intermediate point – we have the relational decomposition $e_q = f_q \cap (Inv_q \times Inv_q)$, which is likewise expressible in a modal sentence scheme. With the addition of *relational converse*, we can also modally express the convexity property.

Adding topological structure

Within modal logic, there is a well-known way of representing a *topology* $\mathcal{T}$ on the state space X of an LTS or Kripke model. From McKinsey and Tarski's work in the 1940's (McKinsey & Tarski 1944), (Rasiowa & Sikorski 1963), the axioms for the box modality of the modal logic **S4** correspond exactly to those of the Kuratowski axioms for the topological interior operator, and dually, the **S4** diamond corresponds to topological closure. The logic **S4** is better known by its relational Kripke semantics in terms of *pre-orders*: reflexive and transitive relations $\preccurlyeq \subseteq X \times X$. In showing that the Kripke semantics are a special case of the topological semantics, one is lead to a study of *Alexandroff* topologies (Davoren 1998b).

Let $\mathcal{F}_{\mu,\Box}(\Phi,\Sigma)$ denote the collection of formulas defined as in Definition 2 with an additional clause for a plain (unlabeled) $\Box$ modality. The diamond is defined by the usual duality: $\Diamond\varphi \triangleq \neg\Box\neg\varphi$. For LTS models $\mathfrak{M} = (X, \mathcal{T}, \{a^{\mathfrak{M}}\}_{a\in\Sigma}, \{\|p\|^{\mathfrak{M}}\}_{p\in\Phi})$ additionally equipped with a topology $\mathcal{T}$ on X, the extra semantic clauses for formulas $\varphi \in \mathcal{F}_{\mu,\Box}(\Phi,\Sigma)$ are:

$$\|\Box\varphi\|^{\mathfrak{M}}_{\xi} \triangleq int_{\mathcal{T}}\left(\|\varphi\|^{\mathfrak{M}}_{\xi}\right) \quad \text{and} \quad \|\Diamond\varphi\|^{\mathfrak{M}}_{\xi} \triangleq cl_{\mathcal{T}}\left(\|\varphi\|^{\mathfrak{M}}_{\xi}\right)$$

In the enriched language, we can simply express topological properties of sets of states, such as being *open*, *closed*, *dense* or *nowhere dense*.

The appropriate topological notions of *continuity* are those for relations/set-valued maps, as introduced by Kuratowski and Bouligand in the 1930's. Instead of the functional continuity "the inverse-image of every open set is open", there are two distinct notions of *semi-continuity*: the *u.s.c.* property is "the $\forall$-pre-image of every open set is open", while the *l.s.c.* property substitutes the $\exists$-pre-image operator. The two properties are expressible by the sentence schemes:

$$[a]\Box Z \rightarrow \Box[a]Z \quad \text{and} \quad \langle a\rangle\Box Z \rightarrow \Box\langle a\rangle Z$$

In the setting of compact metric spaces and relations with closed set-values, there is an analog of the familiar ϵ-δ characterization of continuity of functions for each of the two semi-continuity properties (Akin 1993). For the orbit relation of a semi-flow, the picture is that of an *ϵ-tube*, as illustrated in Figure 3.

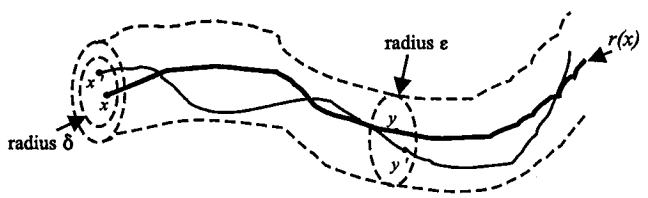

Figure 3: The u.s.c. property in the compact metric setting.

Adding metric tolerance structure

Metric structure on the state space of an LTS model can be used to define explicit metric tolerance relations. For X a metric space with metric d_X, and $\epsilon > 0$, define a relation of *ϵ-tolerance* or *ϵ-indiscernability* $(\epsilon) : X \rightsquigarrow X$ by:

$$x \ (\epsilon) \ x' \quad \text{iff} \quad d_X(x, x') < \epsilon$$

Such a relation is *reflexive* and *symmetric*, but not transitive. Formally, we extend the transition alphabet Σ with a new symbol ϵ, and interpret the new modalities $\langle\epsilon\rangle$ and $[\epsilon]$ in the standard way with the corresponding pre-image operators. The sentence $\langle\epsilon\rangle\varphi$ thus denotes the *ϵ-ball* around the set $\|\varphi\|^{\mathfrak{M}}$. In the full paper, we explore ways of relaxing the definition of "perfect precision" hybrid trajectories using metric tolerance relations.

Conclusion

We demonstrate that the modal μ-calculus and various of its polymodal extensions provide an expressively rich yet highly usable logical framework for the formal analysis of hybrid automata and related hybrid dynamical systems.

Acknowledgments.

I would like to thank Prof. Anil Nerode, David Cook, Joe Miller, Suman Ganguli and Prof. Dexter Kozen for valuable conversations, and Xi Krump for graphic artistry. This research was supported by the ARO under the MURI program "An Integrated Approach to Intelligent Systems", grant no. DAA H04-96-1-0341.

References

Akin, E. 1993. *The General Topology of Dynamical Systems.* Providence, R.I.: AMS.

Alur, R.; Henzinger, T.; and Ho, P.-H. 1996. Automatic symbolic verification of embedded systems. *IEEE Transactions on Software Engineering* 22:181–201.

Ambler, S.; Kwiatkowska, M.; and Measor, N. 1995. Duality and completeness of the modal μ-calculus. *Theoretical Computer Science* 151:3–27.

Aubin, J.-P., and Frankowska, H. 1990. *Set-Valued Analysis.* Boston: Birkhäuser.

Davoren, J. 1998a. *Modal Logics for Continuous Dynamics.* Ph.D. Dissertation, Department of Mathematics, Cornell University.

Davoren, J. 1998b. Topologies, continuity and bisimulations. Technical Report 98-13, CFIS, Cornell University.

Davoren, J. 1999. On continuous dynamics and modal logics. Forthcoming.

Daws, C.; Olivero, A.; Tripakis, S.; and Yovine, S. 1996. The tool KRONOS. In Alur, R.; Henzinger, T.; and Sontag, E. D., eds., *Hybrid Systems III*, LNCS 1066. Berlin: Springer-Verlag. 208–219.

Emerson, E. A. 1997. Modal checking and the mu-calculus. In Immerman, N., and Kolaitis, P. G., eds., *Descriptive Complexity and Finite Models.* Providence, R.I.: AMS. 185–208.

Henzinger, T.; Kupferman, O.; and Qadeer, S. 1998. From pre-historic to post-modern symbolic model checking. In *Proc. of 10th International Conference on Computer-aided Verification (CAV'98)*, LNCS 1427, 195–206. Berlin: Springer-Verlag.

Henzinger, T. 1996. The theory of hybrid automata. In *Proc. of 11th Annual IEEE Symposium on Logic in Computer Science (LICS'96)*, 278–292. IEEE Computer Society Press.

Janin, D., and Walukiewicz, I. 1996. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *Proc. of 7th International Conference on Concurrency Theory (CONCUR'96)*, LNCS 1119, 263–277. Berlin: Springer-Verlag.

Kozen, D. 1983. Results on the propositional μ-calculus. *Theoretical Computer Science* 27:333–354.

Kupferman, O., and Vardi, M. 1998. Freedom, weakness, and determinism: From linear-time to branching-time. In *Proc. of 13th Annual IEEE Symposium on Logic in Computer Science (LICS'98).* IEEE Computer Society Press.

Lafferriere, G.; Pappas, G.; and Sastry, S. 1998. O-minimal hybrid systems. Technical Report UCB/ERL M98/29, Dept. EECS, UC Berkeley.

Manna, Z., and Pnueli, A. 1993. Verifying hybrid systems. In Grossman, R.; Nerode, A.; Ravn, A.; and Rischel, H., eds., *Hybrid Systems*, LNCS 736. Berlin: Springer-Verlag. 4–35.

Manna, Z., and Sipma, H. B. 1998. Deductive verification of hybrid systems using step. In Henzinger, T. A., and Sastry, S., eds., *Hybrid Systems-Computation and Control (HSCC '98)*, LNCS 1386. Berlin: Springer-Verlag.

McKinsey, J., and Tarski, A. 1944. The algebra of topology. *Annals of Mathematics* 45:141–191.

Rasiowa, H., and Sikorski, R. 1963. *The Mathematics of Metamathematics.* Warsaw: PWN.

Stirling, C. 1992. Modal and temporal logics. In Abramsky, S.; Gabbay, D.; and Maibaum, T., eds., *Handbook of Logic in Computer Science*, volume 2. Clarendon Press, Oxford: Oxford University Press. 477–563.

van den Dries, L. 1998. *Tame Topology and O-minimal Structures.* London Math. Soc. Lecture Notes 248. CUP.

Walukiewicz, I. 1996. A note on the completeness of Kozen's axiomatization of the propositonal μ-calculus. *Bulletin of Symbolic Logic* 2:349–366.