

Input Synthesis for Sampled Data Systems by Program Logic

Takumi Akazaki Ichiro Hasuo

Department of Computer Science

The University of Tokyo, Japan

ultraredrays@is.s.u-tokyo.ac.jp ichiro@is.s.u-tokyo.ac.jp

Kohei Suenaga

Graduate School of Informatics

Kyoto University, Japan

ksuenaga@fos.kuis.kyoto-u.ac.jp

Inspired by a concrete industry problem we consider the *input synthesis* problem for hybrid systems: given a hybrid system that is subject to *input* from outside (also called *disturbance* or *noise*), find an input sequence that steers the system to the desired postcondition. In this paper we focus on *sampled data systems*—systems in which a digital controller interrupts a physical plant in a periodic manner, a class commonly known in control theory—and furthermore assume that a controller is given in the form of an imperative program. We develop a structural approach to input synthesis that features forward and backward reasoning in *program logic* for the purpose of reducing a search space. Although the examples we cover are limited both in size and in structure, experiments with a prototype implementation suggest potential of our program logic based approach.

1 Introduction

Cyber-physical systems (CPS)—integration of digital control with physical environments—are gaining yet more and more importance, with cars, airplanes and all others controlled by computers. *Hybrid systems* capture one of the crucial aspects of CPS, by focusing on the combination of continuous *flow* dynamics and discrete *jump* dynamics. Quality assurance of hybrid systems is therefore a big concern in industry as well as in academia.

In this paper we study the *input synthesis* problem of hybrid systems: given a hybrid system that is subject to *input* from outside (also commonly called *disturbance* or *noise*), we aim to find an input sequence that steers the system to the desired postcondition. Our interest in input synthesis stems from the following concrete problem; it was provided by our research partner in car manufacturing industry as a prototype of the problems they often encounter in their design process.

Example 1.1. In the system below in Fig. 1, the controller interrupts the plant (a car) once every second and manages the velocity v of the car. The controller chooses one mode m_1 and the plant operates in that mode for one second, after which the value of v is fed back to the controller via the sensor. The problem is to come up with an initial state of the whole system together with an input sequence $i_0 \cdots i_{999}$, such that:

- **(precondition)** the initial state satisfies $\text{cnt} = 0$ and $x \in [-0.1, 0.1]$; and
- **(postcondition)** after 1000 seconds, the system satisfies $\text{cnt} = 100$.

The input synthesis problem can arise in many different contexts in quality assurance of hybrid systems. One example is *testing*: the desired postcondition is the trigger for some countermeasure (e.g. a fuse) against certain extremity (the countermeasure is outside the model); and we seek for input (i.e.

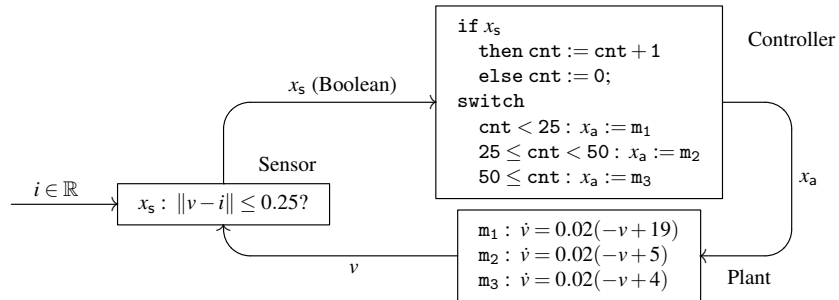


Figure 1: A hybrid system

a test case) that drives the system to activating the countermeasure. The input sequence thus discovered in the model can be fed to the physical realization of the system to see if the countermeasure works properly.

This paper contributes an algorithm for solving the input synthesis problem. Its novelty is the use of *program logic*: we make the most of the structures expressed in the digital controller given in the form of a program. In fact, a likely human effort for the problem in Example 1.1 is:

(*) “for the system to have $\text{cnt} = 100$ at time $k = 1000$, the Boolean value x_s must be true from $k = 900$ through $k = 999$, and ...”;

this is nothing but reasoning in program logic and is included in our proposed algorithm.

More specifically, we restrict our attention to a class of hybrid systems commonly called *sampled data systems*. One such system consists of a physical *plant*, a digital *controller* that periodically interrupts the plant (for simplicity we assume a fixed interval), and a *sensor* that feeds the state of the plant back to the controller. This structural assumption—restrictive yet realistic—allows us to think of the behaviors of such systems quite much as the semantics of programs, and enables forward and backward reasoning in program logic. In our algorithm for solving the input synthesis problem, reasoning in program logic (like the above (*)) contributes to the reduction of the search space. Indeed our prototype implementation successfully solves the problem in Example 1.1.

Related Work The closest to the current work is one by Zutshi, Sankaranarayanan and Tiwari [17], where they verify safety properties of sampled data systems. Their model is more expressive, in that a plant can autonomously change its modes without interruption by a controller. While their goal is reachability analysis and is different from the current paper’s, their relational abstraction technique can be useful in our algorithm, too, in particular for the forward approximation phase.

SMT-solver based approaches [6, 8] to hybrid system analysis are related, too, especially in their emphases on discrete jump dynamics rather than continuous flow. Their effectivity in the input synthesis problem is not yet clear, though: the only available implementation (that of dReal [8]) returned ‘unsat’ to Example 1.1.

More generally, an important feature of our modeling is that a digital controller is given in the form of a program, unlike an automaton used in a majority of existing work (including [8, 17]). The contrast is comparable to the difference between the *theorem proving* (or *type-based*) approach and *software model checking* in program verification. While there have been results [11, 12] that suggest these two approaches are equivalent on a fundamental level, differences do remain especially in applications. In our proposed algorithm it is an advantage that we can exploit rich structural information that is explicit in a program in inferring impossibility (`false`) more quickly.

The backward search phase of our algorithm resembles a *membership question* addressed in the seminal work by Alur et al. [2]. Since our plant (flow) dynamics is not necessarily linear, it is not easy to see how the results in [2] can be used in our problem. They could nevertheless be applied to meta-properties of the problem such as complexity.

Fainekos and his colleagues have developed several techniques for analyzing *robustness* of hybrid systems. Among them is a tool called S-Taliro [3]: it searches for a trajectory by optimization that relies on the continuous nature of the system dynamics. It is possible to encode the input synthesis problem into an input to S-Taliro. However our leading example (Example 1.1), of a jump-heavy nature, seems to fall out of the tool’s focus (it timed out with a smaller problem of 15, not 1000, time units).

Several techniques for testing hybrid systems have been proposed [1, 4, 5, 7, 10]. Although they synthesize test cases and therefore seem similar to what we do here, their goal is to meet certain coverage criteria (such as star discrepancy in [5]) and not to come up with input that steers the system to a specific desired postcondition.

The current work is on logical analysis of hybrid systems; and in that respect it is close to Platzer’s recent series of work (see e.g. [13]) where dynamic logic is extended in a systematic way so that it encompasses continuous dynamics too. Also related is the work [14, 15] by some of the authors where: flow is turned into jump with the help of *nonstandard analysis*; and (discrete) program logic is applied as it is to hybrid systems.

Future Work In this paper we applied program logic to the specific problem of input synthesis. We believe the technique have a greater potential and plan to look at other applications.

The current implementation can only handle continuous plants of dimension 1. Its extension to larger dimensions seems feasible. Specifically, the forward approximation phase of our algorithm will be unproblematic, while in the backward search phase we will have to give up completeness.

Currently our modeling of a sampled data system has a fixed clock cycle. It does not seem hard to accommodate variable intervals; such extension as well as its use is a topic of our future work.

Our modeling benefits a lot from the assumption that the controller communicates with the plant and the sensor using finite datatypes. Some hybrid systems do call for relaxation of this assumption in their modeling; it is our future work to see how the current input synthesis algorithm carries over to such relaxation.

Organization of the Paper In §2 we introduce our modeling of sampled data systems and formalize the input synthesis problem. In §3 we describe our algorithm, explaining its three phases one by one. In §4 our implementation is described, together with the experimental results. The proofs are deferred to the appendix.

Acknowledgments We are grateful to the reviewers of an earlier version for their useful comments and suggestions. T.A. and I.H. are supported by Grants-in-Aid for Young Scientists (A) No. 24680001, and by Aihara Innovative Mathematical Modeling Project, FIRST Program, JSPS/CSTP; K.S. is supported by Grants-in-Aid for Young Scientists (B) No. 70633692 and The Hakubi Project of Kyoto University.


Notations \mathbb{R} is the set of real numbers; $\mathbb{B} = \{\text{tt}, \text{ff}\}$ is the set of Boolean values. We let $f[x_0 \mapsto y_0]$ denote *function update*: it carries x_0 to y_0 and acts as f on the other input.

2 Modeling Sampled Data Systems

2.1 Overview

Sampled data systems are a class of hybrid systems commonly known in control theory. In those systems a physical plant is interrupted by a digital controller in a periodic manner. In the current paper where our interests are in input synthesis, it is convenient to explicitly separate the third component called a *sensor*. The three components are then organized in a loop, as shown on the right in (1).

In the execution of sampled data systems thus modeled, we refer to the three stages in which the sensor, the controller, and the plant operates, respectively, as the *sense*, *think*, and *act* stages. Note that the sensor also takes input from outside the system.



(1)

For simplification we further assume the following.

1. A (digital) controller is written in an imperative programming language.
2. In the execution of a sampled data system, the sense-think-act loop is executed at fixed intervals—once every one second.¹
3. The sense and control stages take no time for their execution.
4. The controller governs the plant by picking a *mode*, from a finite set $\{m_1, \dots, m_M\}$. In particular, the controller cannot feed the plant with a continuous value r .
5. In the act stage the plant operates according to (the ODE associated with) the mode m_i picked by the controller. The act stage lasts for one second (a fact that follows from 2. and 3).
6. The data sent from the sensor to the controller is finitely many Boolean values.


While there are many actual systems that fall out of the realm of this modeling, it does cover fairly many—among which are fixed interval digital controllers, a class of hybrid systems ubiquitous in industry. Sampled data systems, especially under the above assumptions, come to exhibit pleasant structural properties: its behaviors are much like those of programs and we can apply forward and backward reasoning in program logic. Assumptions 2. and 3. are common (see e.g. [17]). For example, Assumption 3. is reasonable considering the speed of digital circuits and typical sensing intervals ($\Delta \approx 1\text{ms}$). Assumptions 4. and 6.—that the controller communicates via finite datatypes—are essential in reducing the input synthesis problem to a search problem.

2.2 The Language IMP_{Ctrl}

We start with defining an imperative programming language IMP_{Ctrl} that is used to describe the (digital) controller of a sampled data system. It is a standard one and is much like IMP in [16], but lacks the `while` construct. It is indeed unrealistic to have `while` loops in real-time applications like cyber-physical systems. Moreover, without `while` loops we can succinctly express weakest preconditions and strongest preconditions—the latter are fully exploited in our algorithm for input synthesis.

¹The clock cycle can be an arbitrary number Δ ; in this paper we assume $\Delta = 1$ for simplicity.

In $\mathbf{IMP}_{\text{Ctrl}}$ the set $\mathbf{Var} = \mathbf{Var}_t \cup \mathbf{Var}_s \cup \mathbf{Var}_a$ of variables is divided into three classes: the *think*, *sense* and *act* variables. The distinction is for the purpose of communicating with the other two components (plant and sensor) of a system. As we will see, a think variable $x_t \in \mathbf{Var}_t$ stores a real number (which will be a floating-point number in an actual implementation); a sense variable $x_s \in \mathbf{Var}_s$ represents a Boolean value sent from the sensor; and the (only) act variable x_a in $\mathbf{Var}_a = \{x_a\}$ tells the plant which mode m_i the plant should take in the coming interval.



Definition 2.1 (the language $\mathbf{IMP}_{\text{Ctrl}}$). Let $\text{Modes} = \{m_1, \dots, m_M\}$ be a fixed finite set of *modes*; \mathbf{Var}_t be a countable set of *think variables*; \mathbf{Var}_s be a finite set of *sense variables*; and $\mathbf{Var}_a = \{x_a\}$. The syntax of $\mathbf{IMP}_{\text{Ctrl}}$ is as follows.

$$\begin{array}{lll}
\mathbf{AExp} \ni & a ::= r \mid x_t \mid a_1 \text{ aop } a_2 & \text{arithmetic expr.} \\
\mathbf{BExp} \ni & b ::= \text{true} \mid \text{false} \mid x_s \mid a_1 \text{ rop } a_2 \mid \neg b \mid b_1 \vee b_2 \mid b_1 \wedge b_2 & \text{Boolean expr.} \\
\mathbf{Cmd} \ni & c ::= \text{skip} \mid x_t := a \mid x_a := m_i \mid c_1; c_2 \mid \text{if } b \text{ then } c_1 \text{ else } c_2 & \text{commands}
\end{array}$$

Here $r \in \mathbb{R}$, $m_i \in \text{Modes}$, $x_t \in \mathbf{Var}_t$, $x_s \in \mathbf{Var}_s$, $\text{aop} \in \{+, -, \times\}$ and $\text{rop} \in \{=, <, \leq, >, \geq\}$.

The semantics of $\mathbf{IMP}_{\text{Ctrl}}$ is as usual, like in [16]. See Def. A.1 for details.

2.3 Assertions for $\mathbf{IMP}_{\text{Ctrl}}$

We now introduce an assertion language for $\mathbf{IMP}_{\text{Ctrl}}$. Its formulas are used to express pre- and postconditions in the input synthesis problem, as well as in program logic. The semantics of the first-order language $\mathbf{Assn}_{\text{Ctrl}}$ is as usual. See Def. A.2.

Definition 2.2 (the assertion language $\mathbf{Assn}_{\text{Ctrl}}$). We fix a set \mathbf{Var}' of “logical” variables such that $\mathbf{Var} \cap \mathbf{Var}' \neq \emptyset$. The assertion language $\mathbf{Assn}_{\text{Ctrl}}$ is defined as follows.

$$\begin{array}{lll}
\mathbf{AExp} \ni & a ::= r \mid x_t \mid v' \mid a_1 \text{ aop } a_2 & \text{arithmetic expressions} \\
\mathbf{MExp} \ni & m ::= m_i \mid x_a & \text{mode expressions} \\
\mathbf{Fml} \ni & \Phi ::= \text{true} \mid \text{false} \mid x_s \mid a_1 \text{ rop } a_2 \mid m = m \mid \neg \Phi \mid & \text{formulas} \\
& \Phi_1 \vee \Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid \forall v' \in \mathbb{R}. \Phi \mid \exists v' \in \mathbb{R}. \Phi
\end{array}$$

Here $r \in \mathbb{R}$, $m_i \in \text{Modes}$, $x_t \in \mathbf{Var}_t$, $x_s \in \mathbf{Var}_s$, and $v' \in \mathbf{Var}'$. Intuitively, $\sigma \in \Sigma$ is a valuation that depends on the state of a sampled data system; and $\gamma \in \mathbb{R}^{\mathbf{Var}'}$ is another valuation of (logical) variables in $\mathbf{Assn}_{\text{Ctrl}}$.

2.4 Calculi for Weakest Preconditions and Strongest Postconditions

We introduce program logic for $\mathbf{IMP}_{\text{Ctrl}}$ in the form of a *weakest precondition calculus* (see e.g. [16]) and a *strongest postcondition calculus* (see e.g. [9]). The calculi will be exploited for the search space reduction in input synthesis.

Definition 2.3 (weakest precondition $w[[c, \Phi]]$; strongest postcondition $s[[c, \Phi]]$). Given $c \in \mathbf{Cmd}$ of $\mathbf{IMP}_{\text{Ctrl}}$ and $\Phi \in \mathbf{Fml}$ of $\mathbf{Assn}_{\text{Ctrl}}$, we define a formula $w[[c, \Phi]] \in \mathbf{Fml}$ inductively on c .

$$\begin{aligned}
w[[\text{skip}, \Phi]] &\equiv \Phi, & w[[c_1; c_2, \Phi]] &\equiv w[[c_1, w[[c_2, \Phi]]]] , \\
w[[x_t := a, \Phi]] &\equiv \Phi[a/x_t], & w[[x_a := m_i, \Phi]] &\equiv \Phi[m_i/x_a], \\
w[[\text{if } b \text{ then } c_1 \text{ else } c_2, \Phi]] &\equiv (b \wedge w[[c_1, \Phi]]) \vee (\neg b \wedge w[[c_2, \Phi]]);
\end{aligned} \tag{2}$$

A formula $s[[c, \Phi]] \in \mathbf{Fml}$ is defined as follows, similarly by induction.

$$\begin{aligned}
s[[\text{skip}, \Phi]] &\equiv \Phi, & s[[c_1; c_2, \Phi]] &\equiv s[[c_2, s[[c_1, \Phi]]]] , \\
s[[x_t := a, \Phi]] &\equiv \exists v' \in \mathbb{R}. (\Phi[v'/x_t] \wedge x_t = a[v'/x_t]) , \\
s[[x_a := m_i, \Phi]] &\equiv (\Phi[m_1/x_a] \vee \dots \vee \Phi[m_M/x_a]) \wedge x_a = m_i , \\
s[[\text{if } b \text{ then } c_1 \text{ else } c_2, \Phi]] &\equiv s[[c_1, b \wedge \Phi]] \vee s[[c_2, \neg b \wedge \Phi]] .
\end{aligned} \tag{3}$$

In our implementation, $\mathbf{Assn}_{\text{Ctrl}}$ is restricted to its propositional fragment for tractability. The quantifier in (3) is thus immediately eliminated using the quantifier elimination mechanism in Mathematica. The third line in (3) is essentially the same as the second; there we can dispense with a quantifier \exists since $\text{Modes} = \{m_1, \dots, m_M\}$ is a finite set.

Proposition 2.4. For any $\sigma \in \Sigma$ and $\gamma \in \mathbb{R}^{\text{Var}'}$,

1. (weakest precondition) $\sigma, \gamma \models w[[c, \Phi]]$ if and only if $[[c]](\sigma), \gamma \models \Phi$;
2. (strongest postcondition) $\sigma, \gamma \models \Phi$ if and only if $[[c]](\sigma), \gamma \models s[[c, \Phi]]$.

2.5 Modeling Sampled Data Systems, Formally

We present the formal definition of our modeling of sampled data systems, under the assumptions in §2.1.

Definition 2.5 (sampled data system). Let n be a natural number, and $I \subseteq \mathbb{R}^n$ be a fixed set called the *input domain*. An n -dimensional sampled data system is a triple $\mathcal{S} = (c, p, s)$ where:

- $c \in \mathbf{Cmd}$ is a command of $\mathbf{IMP}_{\text{Ctrl}}$ (§2.2), called a *controller*;
- $p = (\dot{x} = p_{m_i}(t, x))_{m_i \in \text{Modes}}$ is a family of (explicit, n -dimensional) ODEs indexed by $\text{Modes} = \{m_1, \dots, m_M\}$, called a *plant*; and
- $s : \mathbb{R}^n \times I \rightarrow \mathbb{B}^{\text{Var}'_s}$ is a function, called a *sensor*.

A *state* of a sampled data system is a pair (σ, x) of $\sigma \in \Sigma$ and $x \in \mathbb{R}^n$. In a state (σ, x) , the component σ is called a *controller state* (*C-state*), and x a *plant state* (*P-state*).

The dimension n refers to that of the (continuous) plant, meaning that x and \dot{x} in the plant $p = (\dot{x} = p_{m_i}(t, x))_{m_i \in \text{Modes}}$ are vectors in \mathbb{R}^n .

Example 2.6 (count and brake). In Fig. 2 is a simplification of Example 1.1; this will be our running example. The value v is intended to be the velocity of a car.




Figure 2: A sampled data system (running example)

The example follows a pattern of fixed interval controllers commonly used in industry. Namely, a counter cnt is used to tell if extremity $(v + i \geq 1)$ has continued for a certain critical number of intervals (2 here). If cnt reaches the critical number a countermeasure is taken: the plant is set to the braking

mode (Brk) and the velocity v decreases. Otherwise the plant operates in the acceleration mode (Ac1), which is a first-order lag system where the velocity v approaches towards 2.

The system takes input i —whose domain is assumed to be $[-0.2, 0.2]$ —that models disturbance from outside. For example, the road can be slippery, which can make the actual velocity v different from the value that is used by the controller.

2.6 Semantics of Sampled Data Systems

We formally define the semantics of a sampled data system. Our current concern is not so much on the solution of ODEs as on the interaction between a controller and a plant. Therefore we adopt the following black-box view of a plant.

Definition 2.7 ($\text{execPlant}(p, x)$). In what follows we assume that all the ODEs used for a plant have unique solutions. That is, for any n -dimensional ODE $\dot{x} = p(t, x)$ and an initial value $x_0 \in \mathbb{R}^n$, we assume that there exists a unique function $F : [0, 1] \rightarrow \mathbb{R}^n$ such that: $F(0) = x_0$; and for any $t \in [0, 1]$, $\dot{F}(t) = p(t, F(t))$.

By $\text{execPlant}(p, x_0)$ we denote the state of the plant $\dot{x} = p(t, x)$ at time $t = 1$, assuming that the initial state (at time $t = 0$) is x_0 . That is, $\text{execPlant}(p, x_0) = F(1)$ where F is the function in the above.

In our implementation we actually use the result of numerical calculations (by MATLAB) as the value $\text{execPlant}(p, x)$, ignoring numerical errors.

Definition 2.8 (semantics of a sampled data systems). Let $\mathcal{S} = (c, p, s)$ be a sampled data system. The *one-step transition* is a ternary relation \rightarrow among two states (σ, x) , (σ', x') and input $i \in I$; this is denoted by $(\sigma, x) \xrightarrow{i} (\sigma', x')$. It is defined as follows.

We have $(\sigma, x) \xrightarrow{i} (\sigma', x')$ if $(\sigma', x') = (\text{act}_{\mathcal{S}} \circ \text{think}_{\mathcal{S}} \circ \text{sense}_{\mathcal{S}})(\sigma, x, i)$, where the three functions are defined by:

$$\begin{aligned} \text{sense}_{\mathcal{S}} &: \Sigma \times X \times I \longrightarrow \Sigma \times X, & (\sigma, x, i) &\longmapsto (\sigma[x_s \mapsto s(x, i)(x_s)], x); \\ \text{think}_{\mathcal{S}} &: \Sigma \times X \longrightarrow \Sigma \times X, & (\sigma, x) &\longmapsto (\llbracket c \rrbracket(\sigma), x); \\ \text{act}_{\mathcal{S}} &: \Sigma \times X \longrightarrow \Sigma \times X, & (\sigma, x) &\longmapsto (\sigma, \text{execPlant}(p_{\sigma(x_a)}, x)). \end{aligned} \quad (4)$$

Here $\llbracket c \rrbracket$ is as in Def. A.1. It is clear that, given a state (σ, x) and $i \in I$, the post-state (σ', x') such that $(\sigma, x) \xrightarrow{i} (\sigma', x')$ is uniquely determined. A succession $(\sigma_0, x_0) \xrightarrow{i_0} (\sigma_1, x_1) \xrightarrow{i_1} \dots \xrightarrow{i_{T-1}} (\sigma_T, x_T)$ of one-step transition is called a *run* of the system \mathcal{S} .

A specification of a state of a sampled data system is given by a pair of an assertion formula (on the controller) and a subset of \mathbb{R}^n (on the plant).

Definition 2.9 (CP-condition). Let $\mathcal{S} = (c, p, s)$ be an n -dimensional sampled data system. A *controller-plant condition* (CP-condition) for \mathcal{S} is a pair (Φ, X) of an assertion $\Phi \in \mathbf{Fml}$ called the *controller condition* and a condition $X \subseteq \mathbb{R}^n$ called the *plant condition*. The projection to each component is denoted by π_C and π_P respectively.

Given a state $(\sigma, x) \in \Sigma \times \mathbb{R}^n$ of \mathcal{S} and a CP-condition (Φ, X) , we write $(\sigma, x) \models (\Phi, X)$ if $\sigma \models \Phi$ and $x \in X$. (Φ, X) is *satisfiable* if there is a state that satisfies it.

2.7 The Input Synthesis Problem for Sampled Data Systems

Definition 2.10 (input synthesis problem). The *input synthesis problem* is:

- given:
- $\mathcal{S} = (c, p, s)$, an n -dimensional sampled data system;
 - $(\Phi_{\text{init}}, X_{\text{init}})$ and $(\Phi_{\text{final}}, X_{\text{final}})$, a pre- and a post-CP-condition; and
- return:
- $T \in \mathbb{N}$, the number of steps,
 - an initial state $(\sigma_0, x_0) \in \Sigma \times \mathbb{R}^n$ such that $(\sigma_0, x) \models (\Phi_{\text{init}}, X_{\text{init}})$; and
 - an input sequence $i_0, \dots, i_{T-1} \in I$ such that, for the corresponding run $(\sigma_0, x_0) \xrightarrow{i_0} (\sigma_1, x_1) \xrightarrow{i_1} \dots \xrightarrow{i_{T-1}} (\sigma_T, x_T)$ of \mathcal{S} , we have $(\sigma_T, x_T) \models (\Phi_{\text{final}}, X_{\text{final}})$.

Example 2.11. Let \mathcal{S} be the sampled data system in Example 2.6. Consider

a pre-CP-condition $(\text{cnt} = 0, [0, 1])$ and a post-CP-condition $(\text{true}, [1.5, 2])$

and $T = 4$ as the number of steps. In the input synthesis problem, we seek for an initial state (σ_0, x_0) and an input sequence $i_0, i_1, i_2, i_3 \in [-0.2, 0.2]$ such that

$$(\sigma_0, x_0) \models (\text{cnt} = 0, [0, 1]) \text{ , } (\sigma_0, x_0) \xrightarrow{i_0} (\sigma_1, x_1) \xrightarrow{i_1} \dots \xrightarrow{i_3} (\sigma_4, x_4) \text{ and } (\sigma_4, x_4) \models (\text{true}, [1.5, 2]) \text{ .}$$

3 An Algorithm for Input Synthesis for Sampled Data Systems

In this section we present our algorithm. We identify the core of the input synthesis problem to be the discovery of suitable input and output of the controller at each step. More specifically, we seek for a *successful path*

$$\overrightarrow{(\sigma_s, m)} := \langle (\sigma_s^{(T-1)}, m^{(T-1)}), (\sigma_s^{(T-2)}, m^{(T-2)}), \dots, (\sigma_s^{(0)}, m^{(0)}) \rangle \quad (5)$$

where $\sigma_s^{(k)} : \mathbf{Var}_s \rightarrow \mathbb{B}$ is a valuation of sense variables—which shall be henceforth called *sensor output*—and $m^{(k)} \in \text{Modes}$ is a mode.² Together with an initial state (σ_0, x_0) , the sensor output $\sigma_s^{(k)}$ determines the behavior of the controller, and the mode $m^{(k)}$ determines that of the plant, at each step k . Therefore a path like in (5) determines the behavior of the whole sampled data system from step 0 through step T ; a “successful” path is then one that steers the given precondition to the given postcondition.

Towards the discovery of a successful path, our approach is to exploit the program logic in §2.4—i.e. to make most of the structure of the controller as a program. In our modeling of sampled data systems (§2) we have made assumptions so that the program-logic approach is possible.

Concretely, our algorithm consists of the following three phases.

1. **(Forward approximation)** We overapproximate the set of CP-states that the system can reach, starting from the pre-CP-condition $(\Phi_{\text{init}}, X_{\text{init}})$ and going forward step by step. This first phase is seen as a preparation for the second (main) phase.
2. **(Backward search)** A successful path (5) will be a path in a so-called *backward search tree*. Its branching degree is $2^{|\mathbf{Var}_s|} \times |\text{Modes}|$; its nodes are labeled with CP-conditions; and its root is labeled with the post-CP-condition $(\Phi_{\text{final}}, X_{\text{final}})$. We search for a successful path in the tree, in a depth-first manner.
3. **(Synthesis of actual input)** We choose an initial state (σ, x) and go on to synthesize an input sequence i_0, \dots, i_{T-1} , using the successful path $\overrightarrow{(\sigma_s, m)}$ discovered in the previous phase. This can be done in a straightforward linear manner.

The second phase (backward search) is where an actual (depth-first) search is done. Program logic is used there to prune branches and reduce the search space.

²Note that time is reversed in (5). This is purely for the purpose of presentation.

3.1 Forward Approximation

In this phase of the algorithm, we overapproximate the behavior of the given sampled data system and obtain a sequence $(k\text{-FA}(\Phi_{\text{init}}, X_{\text{init}}))_{0 \leq k \leq T}$ of CP-conditions. These are obtained iteratively as follows.

Notation 3.1 (s^{-1}). Let $\mathcal{S} = (c, p, s)$ be an n -dimensional sampled data system; I be its input domain; and $\sigma_s \in \mathbb{B}^{\text{Var}_s}$ be sensor output. We abuse notation and denote by $s^{-1}(\sigma_s)$ the set of plant states that can be “steered” to σ_s . Precisely, $s^{-1}(\sigma_s) := \{x \in \mathbb{R}^n \mid \exists i \in I. s(x, i) = \sigma_s\}$.

For example, let σ_s such that $\sigma_s(x_s) = \text{tt}$ in the setting of Example 2.6. We have $s^{-1}(\sigma_s) = \{x \in X \mid \exists i \in [-0.2, 0.2]. x + i \geq 1\} = [0.8, \infty)$.

Definition 3.2 (1-FA, k -FA). Let $\mathcal{S} = (c, p, s)$ be a sampled data system. Let us first define the functions $1\text{-FA}_{\text{sense}}^{\text{pre}}$, $1\text{-FA}_{\text{think}}^{\text{pre}}$ and $1\text{-FA}_{\text{act}}^{\text{pre}}$ as follows. Their types should be obvious.

$$\begin{aligned} 1\text{-FA}_{\text{sense}}^{\text{pre}}(\sigma_s)(\Phi, X) &:= (s[x_s := \sigma_s(x_s), \Phi], X \cap s^{-1}(\sigma_s)) \\ 1\text{-FA}_{\text{think}}^{\text{pre}}(\Phi, X) &:= (s[c, \Phi], X) , \\ 1\text{-FA}_{\text{act}}^{\text{pre}}(m)(\Phi, X) &:= (\Phi \wedge x_a = m, \text{execPlant}(p_m, X)) . \end{aligned} \quad (6)$$

Here $s[c, \Phi]$ in the second line is the strongest postcondition (Def. 2.3); $\text{execPlant}(p_m, X)$ in the third line is the direct image of $X \subseteq \mathbb{R}^n$ by the function in Def. 2.7; and $s[x_s := \sigma_s(x_s), \Phi]$ in the first line is defined as follows, similarly to Def. 2.3.

$$s[x_s := \sigma_s(x_s), \Phi] \equiv \begin{cases} (\Phi[\text{true}/x_s] \vee \Phi[\text{false}/x_s]) \wedge x_s & \text{if } \sigma_s(x_s) = \text{true} \\ (\Phi[\text{true}/x_s] \vee \Phi[\text{false}/x_s]) \wedge \neg x_s & \text{if } \sigma_s(x_s) = \text{false} \end{cases}$$

These three functions are composed to yield:

$$1\text{-FA}^{\text{pre}}(\sigma_s, m)(\Phi, X) := 1\text{-FA}_{\text{act}}^{\text{pre}}(m)(1\text{-FA}_{\text{think}}^{\text{pre}}(1\text{-FA}_{\text{sense}}^{\text{pre}}(\sigma_s)(\Phi, X))) ;$$

this is understood as the strongest postcondition after the one-step execution of \mathcal{S} , *assuming that the sensor output σ_s and the mode m have been chosen*.

Finally, the *one-step forward approximation* function is defined as the following disjunction/union over different σ_s and m :

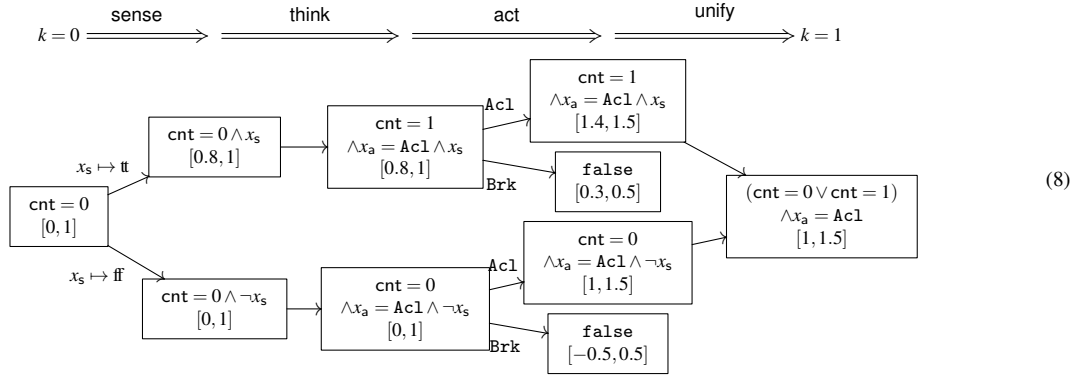
$$1\text{-FA}(\Phi, X) := \left(\bigvee_{(\sigma_s, m) \in \mathcal{M}} \pi_C(1\text{-FA}^{\text{pre}}(\sigma_s, m)(\Phi, X)), \bigcup_{(\sigma_s, m) \in \mathcal{M}} \pi_P(1\text{-FA}^{\text{pre}}(\sigma_s, m)(\Phi, X)) \right) , \\ \text{where } \mathcal{M} := \{(\sigma_s, m) \in \mathbb{B}_s^{\text{Var}} \times \text{Modes} \mid 1\text{-FA}^{\text{pre}}(\sigma_s, m)(\Phi, X) \text{ is satisfiable.}\} \quad (7)$$

The projections π_C and π_P , as well as satisfiability of CP-conditions, are from Def. 2.9.

We write $k\text{-FA}(\Phi, X)$ for $(1\text{-FA})^k(\Phi, X)$. The sequence $(k\text{-FA}(\Phi_{\text{init}}, X_{\text{init}}))_{0 \leq k \leq T}$ of CP-conditions is called the *forward approximation sequence* for \mathcal{S} .

As an example we present forward approximation for Example 2.11. The first one-step approxima-

tion (from $k = 0$ to 1) is shown below, stage by stage.



Observe that we have four CP-conditions in the fourth column from the left. Each of them corresponds to a choice of (σ_s, m) . Two among the four CP-conditions are unsatisfiable and hence discarded (i.e. they are not in \mathcal{M}); the remaining two are unified and yield 1-FA($\text{cnt} = 0, [0, 1]$) in the rightmost column.³

By continuing further we obtain the forward approximation sequence shown on the below in (9), presented pictorially.

For the completeness of our algorithm we need to prove that our forward approximation is indeed an *over-approximation*.

Proposition 3.3. *Let $i_0, \dots, i_{k-1} \in I$ be any input sequence; $(\sigma, x) \xrightarrow{i_0} \dots \xrightarrow{i_{k-1}} (\sigma', x')$ be a run of \mathcal{S} ; and $(\sigma, x) \models (\Phi, X)$. Then $(\sigma', x') \models k\text{-FA}(\Phi, X)$.*
□


3.2 Backward Search

In this phase of the algorithm we search for a successful path (σ_s, m) of sensor output and modes—i.e. one that steers an initial state to a desired postcondition. The search is conducted in a backward depth-first manner in a tree called the *backward search tree*.⁽⁹⁾

For the input synthesis problem, it is not necessary to construct the whole backward search tree: finding a leaf whose CP-condition is compatible with the precondition suffices. We will use program logic (§2.4)—and the forward approximation sequence obtained in the previous phase—in pruning branches and reducing the search space.

Definition 3.4 (backward search tree). Given an input synthesis problem, its *backward search tree* is a tree with branching degree $2^{|\text{Vars}|} \times |\text{Modes}|$ and with height $T + 1$. The nodes of the tree are defined inductively as follows.

- The root of the tree is labeled with the postcondition $(\Phi_{\text{final}}, X_{\text{final}})$.



³ Our approximation can be finer: in (8), in the unification stage, the correlation between a C-condition and a P-condition is forgotten by separately taking the disjunction of C-conditions and the union of P-conditions (see (7)). Finer approximation, however, makes the approximants grow much bigger and slows down the backward search phase of the algorithm.

- Let (Φ, X) be the label at the position $\overrightarrow{(\sigma_s, m)} = \langle (\sigma_s^{(T-1)}, m^{(T-1)}), \dots, (\sigma_s^{(k+1)}, m^{(k+1)}) \rangle$. Its child at the position $\overrightarrow{(\sigma'_s, m')}$ is labeled by

$$(\Phi', X') := \left(\begin{array}{l} \pi_C(k\text{-FA}(\Phi_{\text{init}}, X_{\text{init}})) \wedge \pi_C(1\text{-BS}^{\text{pre}}(\sigma_s, m)(\Phi, X)), \\ \pi_P(k\text{-FA}(\Phi_{\text{init}}, X_{\text{init}})) \cap \pi_P(1\text{-BS}^{\text{pre}}(\sigma_s, m)(\Phi, X)) \end{array} \right), \quad (10)$$

where the function 1-BS^{pre} is defined as follows.

$$\begin{aligned} 1\text{-BS}^{\text{pre}}(\sigma_s, m)(\Phi, X) &:= 1\text{-BS}_{\text{sense}}^{\text{pre}}(\sigma_s) \left(1\text{-BS}_{\text{think}}^{\text{pre}} \left(1\text{-BS}_{\text{act}}^{\text{pre}}(m)(\Phi, X) \right) \right) \text{ where} \\ 1\text{-BS}_{\text{act}}^{\text{pre}}(m)(\Phi, X) &:= (\Phi \wedge x_a = m, \text{execPlant}(\text{rev}(p_m), X)) , \\ 1\text{-BS}_{\text{think}}^{\text{pre}}(\Phi, X) &:= (w \llbracket c, \Phi \rrbracket, X) , \text{ and} \\ 1\text{-BS}_{\text{sense}}^{\text{pre}}(\sigma_s)(\Phi, X) &:= (\Phi[\sigma_s(x_s)/x_s], X \cap s^{-1}(\sigma_s)) . \end{aligned} \quad (11)$$

In the second line, $\text{execPlant}(\text{rev}(p_m), x_1)$ means running the original ODE $\dot{x} = p_m(t, x)$ with time reversed (i.e. from $t = 1$ to $t = 0$) and with the “initial” value x_1 (at time $t = 1$). Concretely, $\text{rev}(p_m)$ is given by: $\text{rev}(p_m)(t, x) = -p_m(1 - t, x)$.

By $\text{BS}(\overrightarrow{(\sigma_s, m)})$ we denote the label in the tree, at the position designated by the path $\overrightarrow{(\sigma_s, m)}$. $\text{BS}(\overrightarrow{(\sigma_s, m)})$ is therefore a CP-condition.

Definition 3.5 (successful path). Let $\overrightarrow{(\sigma_s, m)}$ be a path in the backward search tree. It is *successful* if: 1) it is of length T ; and 2) the label $\text{BS}(\overrightarrow{(\sigma_s, m)})$ at the leaf is satisfiable.

The following establishes that that finding a successful path in the backward search tree is equivalent to solving the input synthesis problem.

Proposition 3.6 (soundness & completeness). Let $\overrightarrow{(\sigma_s, m)}$ be a successful path in the backward search tree. Assume also that $\sigma_0 \in \Sigma$ and $x_0 \in \mathbb{R}^n$ satisfy $(\sigma_0, x_0) \models \text{BS}(\overrightarrow{(\sigma_s, m)})$. Then there exists an input sequence $i_0, \dots, i_{T-1} \in T$ such that an initial state (σ_0, x_0) together with i_0, \dots, i_{T-1} is an answer to the input synthesis problem.

Conversely, assume there is an answer to an input synthesis problem, given by (σ_0, x_0) and i_0, \dots, i_{T-1} . Then there is a successful path $\overrightarrow{(\sigma_s, m)}$. \square

In searching for a successful path in a backward search tree, once we hit an unsatisfiable label, clearly all its offspring are unsatisfiable. We therefore prune such a branch. The use of $k\text{-FA}$ in (10) strengthens the labels and makes more branches pruned.

Lemma 3.7 (pruning is correct). In the backward search tree, assume that the label at the position $\overrightarrow{(\sigma_s, m)}$ is unsatisfiable. Then its child has an unsatisfiable label too. \square

An example is again using Example 2.11. Fig. 3 describes details of one-step generation (from the root $k = 4$ to $k = 3$) of the backward search tree. The rightmost is the root; the four leftmost nodes are the direct children of the root; and the intermediate layers are not present in the backward search tree but are shown for illustration. Each of the four children corresponds to each possible choice of (σ_s, m) . The bottom two children are unsatisfiable—the intuition is that the plant’s mode at time $k = 3$ cannot be Brk for the postcondition to hold. The search for a successful path will therefore be continued from one of the two top children.

Presented in Fig. 4 is a more bird’s-eye view of the backward search tree: it shows one possible trace of the depth-first search. It has found a successful path

$$\overrightarrow{(\sigma_s, m)} = \langle (x_s \mapsto \text{tt}, \text{Ac1}), (x_s \mapsto \text{ff}, \text{Ac1}), (x_s \mapsto \text{tt}, \text{Brk}), (x_s \mapsto \text{tt}, \text{Ac1}) \rangle . \quad (12)$$

In the search shown in Fig. 4, pruning has occurred at the nodes (N1)–(N3).

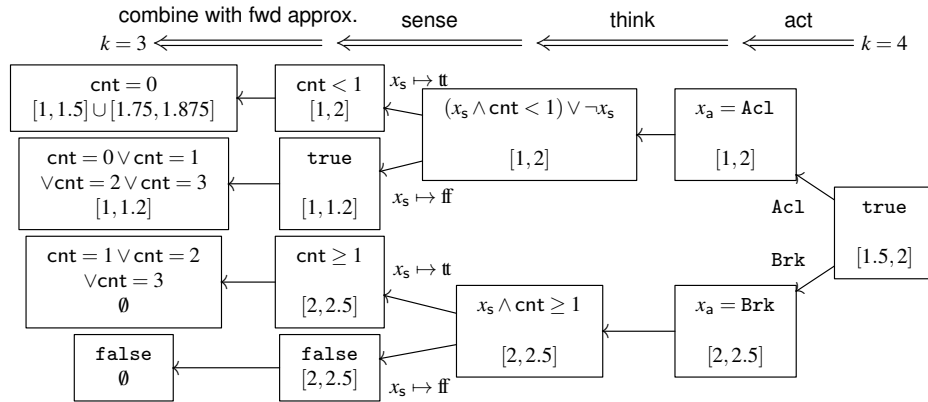


Figure 3: Generation of the backward search tree, in detail




Figure 4: A bird's-eye view of the backward search tree

3.3 Synthesis of Actual Input

The second phase gives us a successful path $\overrightarrow{(\sigma_s, m)}$; as discussed at the beginning of §3, this determines the behavior of the whole sampled data system. We now synthesize an actual answer to the input synthesis problem from the path $\overrightarrow{(\sigma_s, m)}$. Theoretically it is possible (Prop. 3.6); it is moreover computationally cheap, using a CAS like Mathematica.

We describe the procedure by example. For Example 2.11, the second phase gives a successful path in (12), from which we obtain a refinement of the pre-CP-condition $\text{BS}(\overrightarrow{(\sigma_s, m)}) = (\text{cnt} = 0, [0.8, 1])$ (the leftmost node in Fig. 4).

- **(Choosing an initial state)** By Prop. 3.6 any (σ_0, v_0) such that $(\sigma_0, v_0) \models \text{BS}(\overrightarrow{(\sigma_s, m)})$ admits a desired input sequence. Let us say $\sigma_0(\text{cnt}) = 0$ and $v_0 := 0.9$.
- **(Running the plant)** It is crucial that the behavior of the plant is completely determined now, given the initial P-state v_0 and the sequence of modes $\langle m^{(0)}, \dots, m^{(T-1)} \rangle$ extracted from the path $\overrightarrow{(\sigma_s, m)}$. In the current example the plant dynamics is as follows: $0.9 \xrightarrow{\text{Ac1}} 1.45 \xrightarrow{\text{Brk}} 0.95 \xrightarrow{\text{Ac1}} 1.475 \xrightarrow{\text{Ac1}} 1.7375$.
- **(Synthesis of input)** For each moment k , we now know the plant state $v^{(k)}$ and the sensor output $\sigma_s^{(k)} \in \mathbb{B}_s^{\text{var}}$; the latter is extracted from the path $\overrightarrow{(\sigma_s, m)}$. We choose input i_k so that it, combined with $v^{(k)}$, gives the sensor output as specified by $\sigma_s^{(k)}$.

For example let us pick i_2 . Now $\sigma^{(2)}(x_s) = (v + i \geq 1?) = \text{ff}$ and $v^{(2)} = 0.95$; we choose $i_2 \in I = [-0.2, 0.2]$ so that $v^{(2)} + i_2 = 0.95 + i_1 < 1$; say $i_2 = 0$. In implementation we let the FindInstance function of Mathematica do this job.

Overall, we obtain the following run from the pre-CP-condition to the post-CP-condition. This gives an answer $(\text{cnt} \mapsto 0, v = 0.9) \xrightarrow{0.1} (1, 1.45) \xrightarrow{0.0} (2, 0.95) \xrightarrow{0.0} (0, 1.475) \xrightarrow{0.0} (1, 1.7375)$ to the input synthesis problem in Example 2.11.

4 Implementation, Optimization and Experiments

Our prototype implementation has a front-end written in OCaml which in particular implements inferences in program logic (Def. 2.3). Mathematica is used for simplifying arithmetic formulas and inequalities, as well as for picking a value under a certain assumption. We also use MATLAB for numerically solving ODEs.

Our implementation is currently restricted to one-dimensional plants ($n = 1$). From time to time we have to calculate the evolution of an interval according an ODE (like $\text{execPlant}(p_m, X)$ in (6) for a set X); such calculation is done by the method of [6].

Optimization Techniques We further employ the following techniques for speedup. We note that the none of these affect correctness (Prop. 3.6) of our algorithm.

- **(Truncation of forward approximation)** In the forward approximation phase, a problem is that an approximant $k\text{-FA}(\Phi_{\text{init}}, X_{\text{init}})$ can grow exponentially as k grows—as hinted already in (9). Such explosion of approximants slows down not only the forward approximation phase, but also the backward search phase. Moreover such a big approximant tends not to contribute a lot to pruning branches.

To avert this we truncate forward approximation under certain circumstances. Specifically we stop calculating C-conditions when the approximated C-condition has become compatible with any choice of modes—a sign of the C-condition no longer contributing to pruning. Currently such truncation is implemented only for C-conditions; but it should also be possible for P-conditions, e.g. by merging intervals in (9).

- **(Prioritization in search)** In the backward (depth-first) search phase, we can have multiple children from which to pick. Besides randomized picks, we have the following prioritization strategies. In the *by-volume prioritization*, we estimate the volume of the P-condition (i.e. a region in \mathbb{R}^n) of each child, and pick one with the biggest. In the *by-robustness prioritization*, in contrast, we pick the child whose P-condition is the closest to the “center” of the forward-approximated P-condition. In other words, the picked child is the one with a P-condition that intersects with the forward-approximated P-condition *in the most robust manner*. This robustness-driven optimization is much like in S-Talro [3].

Experiments We used Mathematica 9.0.1 and MATLAB 8.1.0 (for Linux x86, 64-bit), on ThinkPad T530 with Intel Core i7-3520M 2.90GHz CPU with 3.7GB memory.

The first table below shows the result of our prototype implementation applied to the problem in Example 2.11, with a varying number T of steps. All the times are in seconds. The rows correspond to different prioritization strategies, and whether truncation of forward approximation is enabled. For random prioritization the experiment was repeated 50 times and the average is shown, together with the standard deviation. From the results we can see that forward approximation truncation is very effective as the problem becomes larger, on the one hand. On the other hand, no clear comparative advantage of any of the three prioritization strategies is observed.

The second table below presents the breakdown of two cases (both $T = 100$) from the first table, into the three phases of the algorithm, together with the number of backtracks in a search. While truncation causes more backtracks (this is because less information is passed to the backward search phase), we see that the speed of both of the first two phases are greatly improved thanks to simpler approximants.

We also applied our implementation to the original problem in Example 1.1. It successfully solved the problem in 638.968 seconds.

Overall, our experiments so far are limited to examples of a specific structure: namely, a counter in the controller, incremented or reset to 0 every second, causes the change of modes of the plant. This structure however is a commonly used one in industry (see Example 2.6); and its discrete nature (the counter takes an integer value that can be fairly large) becomes a challenge in many approaches to verification, testing or input synthesis. The experimental results seem to suggest that our program logic based approach is promising in coping with this kind of challenges.

prioritization	truncation	$T = 10$	$T = 20$	$T = 30$	$T = 100$	$T = 1000$
random		3.228	9.849	18.464	108.437 \pm 37.574	No answer
by volume		3.633	10.197	14.332	115.311	No answer
by robustness		3.072	11.082	22.231	68.464	No answer
random	on	3.409	11.314	20.068	54.132 \pm 31.953	377.361 \pm 80.392
by volume	on	3.689	9.289	12.323	38.425	445.784
by robustness	on	3.552	20.702	41.443	38.803	245.661

prioritization	truncation	fwd. approx.	bwd. search	synthesis	total	num. of backtracks
by volume		36.622	76.568	2.121	115.311	140
by volume	on	15.118	21.565	1.743	38.425	176

References

- [1] Rajeev Alur, Aditya Kanade, S. Ramesh & K. C. Shashidhar (2008): *Symbolic analysis for improving simulation coverage of Simulink/Stateflow models*. In Luca de Alfaro & Jens Palsberg, editors: *EMSOFT*, ACM, pp. 89–98. Available at <http://doi.acm.org/10.1145/1450058.1450071>.
- [2] Rajeev Alur, Robert P. Kurshan & Mahesh Viswanathan (1998): *Membership Questions for Timed and Hybrid Automata*. In: *RTSS*, IEEE Computer Society, pp. 254–263. Available at <http://doi.ieeecomputersociety.org/10.1109/REAL.1998.739751>.
- [3] Yashwanth Annpureddy, Che Liu, Georgios E. Fainekos & Sriram Sankaranarayanan (2011): *S-TaLiRo: A Tool for Temporal Logic Falsification for Hybrid Systems*. In Parosh Aziz Abdulla & K. Rustan M. Leino, editors: *TACAS, Lecture Notes in Computer Science 6605*, Springer, pp. 254–257. Available at http://dx.doi.org/10.1007/978-3-642-19835-9_21.
- [4] Bahareh Badban, Martin Fränzle, Jan Peleska & Tino Teige (2006): *Test automation for hybrid systems*. In Johannes Mayer, Neelam Gupta & Yves Ledru, editors: *SOQUA*, ACM, pp. 14–21. Available at <http://doi.acm.org/10.1145/1188895.1188902>.
- [5] Thao Dang & Tarik Nahhal (2009): *Coverage-guided test generation for continuous and hybrid systems*. *Formal Methods in System Design* 34(2), pp. 183–213. Available at <http://dx.doi.org/10.1007/s10703-009-0066-0>.
- [6] Andreas Eggers, Nacim Ramdani, Nedialko Nedialkov & Martin Fränzle (2011): *Improving SAT Modulo ODE for Hybrid Systems Analysis by Combining Different Enclosure Methods*. In Gilles Barthe, Alberto Pardo & Gerardo Schneider, editors: *SEFM, Lect. Notes in Comp. Sci. 7041*, Springer, pp. 172–187. Available at http://dx.doi.org/10.1007/978-3-642-24690-6_13.
- [7] J.M. Esposito (2004): *Randomized test case generation for hybrid systems: metric selection*. In: *System Theory, 2004. Proceedings of the Thirty-Sixth Southeastern Symposium on*, pp. 236–240, doi:10.1109/SSST.2004.1295655.

- [8] Sicun Gao, Soonho Kong & Edmund M. Clarke (2013): *dReal: An SMT Solver for Nonlinear Theories over the Reals*. In Maria Paola Bonacina, editor: *CADE, Lecture Notes in Computer Science 7898*, Springer, pp. 208–214. Available at http://dx.doi.org/10.1007/978-3-642-38574-2_14.
- [9] Mike Gordon & Hlne Collavizza (2010): *Forward with Hoare*. In A.W. Roscoe, Cliff B. Jones & Kenneth R. Wood, editors: *Reflections on the Work of C.A.R. Hoare*, Springer London, pp. 101–121. Available at http://dx.doi.org/10.1007/978-1-84882-912-1_5.
- [10] A. Agung Julius, Georgios E. Fainekos, Madhukar Anand, Insup Lee & George J. Pappas (2007): *Robust Test Generation and Coverage for Hybrid Systems*. In Alberto Bemporad, Antonio Bicchi & Giorgio C. Buttazzo, editors: *HSCC, Lecture Notes in Computer Science 4416*, Springer, pp. 329–342. Available at http://dx.doi.org/10.1007/978-3-540-71493-4_27.
- [11] Naoki Kobayashi & C.-H. Luke Ong (2009): *A Type System Equivalent to the Modal Mu-Calculus Model Checking of Higher-Order Recursion Schemes*. In: *LICS, IEEE Computer Society*, pp. 179–188. Available at <http://doi.ieeecomputersociety.org/10.1109/LICS.2009.29>.
- [12] Mayur Naik & Jens Palsberg (2008): *A type system equivalent to a model checker*. *ACM Trans. Program. Lang. Syst.* 30(5). Available at <http://doi.acm.org/10.1145/1387673.1387678>.
- [13] André Platzer (2012): *The Complete Proof Theory of Hybrid Systems*. In: *LICS, IEEE*, pp. 541–550. Available at <http://dx.doi.org/10.1109/LICS.2012.64>.
- [14] Kohei Suenaga & Ichiro Hasuo (2011): *Programming with Infinitesimals: A While-Language for Hybrid System Modeling*. In Luca Aceto, Monika Henzinger & Jiri Sgall, editors: *ICALP (2), Lecture Notes in Computer Science 6756*, Springer, pp. 392–403. Available at http://dx.doi.org/10.1007/978-3-642-22012-8_31.
- [15] Kohei Suenaga, Hiroyoshi Sekine & Ichiro Hasuo (2013): *Hyperstream processing systems: nonstandard modeling of continuous-time signals*. In Roberto Giacobazzi & Radhia Cousot, editors: *POPL, ACM*, pp. 417–430. Available at <http://doi.acm.org/10.1145/2429069.2429120>.
- [16] Glynn Winskel (1993): *The Formal Semantics of Programming Languages*. MIT Press.
- [17] Aditya Zutshi, Sriram Sankaranarayanan & Ashish Tiwari (2012): *Timed Relational Abstractions for Sampled Data Control Systems*. In P. Madhusudan & Sanjit A. Seshia, editors: *CAV, Lecture Notes in Computer Science 7358*, Springer, pp. 343–361. Available at http://dx.doi.org/10.1007/978-3-642-31424-7_27.

A Auxiliary Definitions and Lemmas

Definition A.1 (semantics $\llbracket _ \rrbracket$ of $\mathbf{IMP}_{\text{Ctrl}}$). Let Σ be the set of *valuations*, that is,

$$\Sigma = \{ \sigma : \mathbf{Var} \rightarrow \mathbb{R} \cup \mathbb{B} \cup \text{Modes} \mid \sigma(\mathbf{Var}_t) \subseteq \mathbb{R}, \sigma(\mathbf{Var}_s) \subseteq \mathbb{B}, \sigma(x_a) \in \text{Modes} \} . \quad (13)$$

For each expression e of $\mathbf{IMP}_{\text{Ctrl}}$, their *semantics* $\llbracket e \rrbracket$ is defined in the following standard way. For $a \in \mathbf{AExp}$, $\llbracket a \rrbracket : \Sigma \rightarrow \mathbb{R}$ is defined by

$$\llbracket r \rrbracket(\sigma) = r , \quad \llbracket x_t \rrbracket(\sigma) = \sigma(x_t) , \quad \llbracket a_1 \text{ aop } a_2 \rrbracket(\sigma) = \llbracket a_1 \rrbracket(\sigma) \text{ aop } \llbracket a_2 \rrbracket(\sigma) .$$

For $b \in \mathbf{BExp}$, $\llbracket b \rrbracket : \Sigma \rightarrow \{\text{tt}, \text{ff}\}$ is defined by

$$\begin{aligned} \llbracket b_1 \vee b_2 \rrbracket(\sigma) &= \llbracket b_1 \rrbracket(\sigma) \vee \llbracket b_2 \rrbracket(\sigma) \quad \text{and similarly for } \neg, \wedge, \text{true and false;} \\ \llbracket x_s \rrbracket(\sigma) &= \sigma(x_s) ; \quad \text{and} \quad \llbracket a_1 \text{ rop } a_2 \rrbracket(\sigma) = \llbracket a_1 \rrbracket(\sigma) \text{ rop } \llbracket a_2 \rrbracket(\sigma) . \end{aligned}$$

For $c \in \mathbf{Cmd}$, $\llbracket c \rrbracket : \Sigma \rightarrow \Sigma$ is defined by

$$\begin{aligned} \llbracket \text{skip} \rrbracket(\sigma) &= \sigma , \quad \llbracket x_t := a \rrbracket(\sigma) = \sigma[x_t \mapsto \llbracket a \rrbracket(\sigma)] , \quad \llbracket x_a := m_i \rrbracket(\sigma) = \sigma[x_a \mapsto m_i] , \\ \llbracket c_1 ; c_2 \rrbracket(\sigma) &= \llbracket c_2 \rrbracket(\llbracket c_1 \rrbracket(\sigma)) , \quad \llbracket \text{if } b \text{ then } c_1 \text{ else } c_2 \rrbracket(\sigma) = \begin{cases} \llbracket c_1 \rrbracket(\sigma) & \text{if } \llbracket b \rrbracket(\sigma) = \text{tt} \\ \llbracket c_2 \rrbracket(\sigma) & \text{if } \llbracket b \rrbracket(\sigma) = \text{ff} \end{cases} \end{aligned}$$

Here $f[x_0 \mapsto y_0]$ denotes *function update*: the function $f[x_0 \mapsto y_0]$ carries x_0 to y_0 and acts as f on the other input.

Definition A.2 (semantics of $\mathbf{Assn}_{\text{Ctrl}}$). We define the semantics of $a \in \mathbf{AExp}$ of $\mathbf{Assn}_{\text{Ctrl}}$ as a function $\llbracket a \rrbracket : \Sigma \times \mathbb{R}^{\mathbf{Var}'} \rightarrow \mathbb{R}$.

$$\begin{aligned} \llbracket r \rrbracket(\sigma, \gamma) &= r , & \llbracket x_t \rrbracket(\sigma, \gamma) &= \sigma(x_t) , \\ \llbracket v' \rrbracket(\sigma, \gamma) &= \gamma(v') , & \llbracket a_1 \text{ aop } a_2 \rrbracket(\sigma) &= \llbracket a_1 \rrbracket(\sigma, \gamma) \text{ aop } \llbracket a_2 \rrbracket(\sigma, \gamma) \end{aligned}$$

For $m \in \mathbf{MExp}$, its semantics is a function $\llbracket m \rrbracket : \Sigma \times \mathbb{R}^{\mathbf{Var}'} \rightarrow \text{Modes}$ defined by $\llbracket m_i \rrbracket(\sigma, \gamma) = m_i$ and $\llbracket x_a \rrbracket(\sigma, \gamma) = \sigma(x_a)$. Finally for formulas, the semantics of $\Phi \in \mathbf{Fml}$ is given by the relation \models between $\Sigma \times \mathbb{R}^{\mathbf{Var}'}$ and \mathbf{Fml} defined as follows. $\sigma, \gamma \models \text{true}$, $\sigma, \gamma \not\models \text{false}$, and

$$\begin{aligned} \sigma, \gamma \models x_s &\stackrel{\text{def.}}{\iff} \sigma(x_s) = \text{tt} , \\ \sigma, \gamma \models a \text{ rop } a' &\stackrel{\text{def.}}{\iff} \llbracket a \rrbracket(\sigma, \gamma) \text{ rop } \llbracket a' \rrbracket(\sigma, \gamma) , \\ \sigma, \gamma \models m = m' &\stackrel{\text{def.}}{\iff} \llbracket m \rrbracket(\sigma, \gamma) = \llbracket m' \rrbracket(\sigma, \gamma) , \\ \sigma, \gamma \models \Phi_1 \vee \Phi_2 &\stackrel{\text{def.}}{\iff} \sigma, \gamma \models \Phi_1 \text{ or } \sigma, \gamma \models \Phi_2 \quad (\text{similarly for } \neg \text{ and } \wedge), \\ \sigma, \gamma \models \forall v' \in \mathbb{R}. \Phi &\stackrel{\text{def.}}{\iff} \sigma, \gamma[v' \mapsto r] \models \Phi \text{ for any } r \in \mathbb{R} \quad (\text{similarly for } \exists). \end{aligned}$$

We write $\sigma \models \Phi$ if $\sigma, \gamma \models \Phi$ for every $\gamma \in \mathbb{R}^{\mathbf{Var}'}$.

The next observation follows immediately from Def. 2.5 and 2.9.

Lemma A.3. (Φ, X) is unsatisfiable if and only if Φ is an unsatisfiable formula (i.e. logically equivalent to false) or $X = \emptyset$. □

B Omitted Proofs

B.1 Proof of Prop. 3.3

Lemma B.1. *Let $i \in I$, $\sigma \in \Sigma$ and $x \in \mathbb{R}^n$. Assume $(\sigma, x) \models (\Phi, X)$. Then:*

$$\begin{aligned} \text{sense}(\sigma, x, i) &\models 1\text{-FA}_{\text{sense}}^{\text{pre}}(s(i, x))(\Phi, X) , & \text{think}(\sigma, x) &\models 1\text{-FA}_{\text{think}}^{\text{pre}}(\Phi, X) , \\ \text{act}(\sigma, x) &\models 1\text{-FA}_{\text{act}}^{\text{pre}}(\sigma(x_a))(\Phi, X) . \end{aligned} \quad (14)$$

Proof. For sense, it is easily shown that $\sigma[x_s \mapsto s(x, i)] \models s[x_s := \sigma_s(x_s), \Phi]$ by induction on Φ . Moreover $x \in X \cap s^{-1}(s(x, i))$ by Notation 3.1. Therefore we have $\text{sense}(\sigma, x, i) \models 1\text{-FA}_{\text{sense}}^{\text{pre}}(s(i, x))(\Phi, X)$.

For think, the claim is obvious from Prop. 2.4.

For act, we trivially have $\sigma \models x_a = \sigma(x_a)$. Therefore

$$(\sigma, \text{execPlant}(p_{\sigma(x_a)}, x)) \models (\Phi \wedge x_a = \sigma(x_a), \text{execPlant}(p_{\sigma(x_a)}, X))$$

follows immediately from the assumption. \square

Proof. (Of Prop. 3.3) It is sufficient to show for $k = 1$; the general case follows by induction. Let $i \in I$, $(\sigma, x) \xrightarrow{i} (\sigma', x')$ be a run of \mathcal{S} , and $(\sigma, x) \models (\Phi, X)$. We need to show

$$(\sigma', x') \models 1\text{-FA}(\Phi, X) .$$

That is obvious because we obtain the following from Lemma. B.1.

$$(\sigma', x') \models 1\text{-FA}^{\text{pre}}(s(i, x), \sigma'(x_a))(\Phi, X) . \quad \square$$

B.2 Proof of Prop. 3.6

Lemma B.2. *The following three properties hold.*

$$\begin{aligned} (\sigma, x) \models 1\text{-BS}_{\text{sense}}^{\text{pre}}(\sigma_s)(\Phi, X) &\implies \text{sense}(\sigma, x, i) \models (\Phi, X) \text{ for some } i \in I \\ (\sigma, x) \models 1\text{-BS}_{\text{think}}^{\text{pre}}(\Phi, X) &\implies \text{think}(\sigma, x) \models (\Phi, X) \\ (\sigma, x) \models 1\text{-BS}_{\text{act}}^{\text{pre}}(m)(\Phi, X) &\implies \text{act}(\sigma, x) \models (\Phi, X) \end{aligned}$$

It follows that: if $(\sigma, x) \models 1\text{-BS}^{\text{pre}}(\sigma_s, m)(\Phi, X)$, then there exist input $i \in I$ and a CP-state (σ', x') for which we have $(\sigma, x) \xrightarrow{i} (\sigma', x')$ and $(\sigma', x') \models (\Phi, X)$.

Proof. For sense, it follows from the assumption that $\sigma \models \Phi[\sigma_s(x_s)/x_s]$ and $x \in X \cap s^{-1}(\sigma_s)$. Then we have $(\sigma[x_s \mapsto \sigma_s(x_s)], x) \models (\Phi, X)$. Moreover, since $x \in s^{-1}(\sigma_s)$ there exists some $i \in I$ such that $s(x, i) = \sigma_s$. For this choice of i we have $\text{sense}(\sigma, x, i) = (\sigma[x_s \mapsto \sigma_s(x_s)], x)$.

For think, we have $\sigma \models w[[c, \Phi]]$ and $x \in X$ from the assumption. Therefore $([[c]](\sigma), x) \models (\Phi, X)$ from Prop. 2.4.

For act, we have $\sigma \models \Phi \wedge x_a = m$ and $x \in \text{execPlant}(\text{rev}(p_m), X)$ from the assumption. Then we have $\sigma(x_a) = m$, hence $x \in \text{execPlant}(\text{rev}(p_{\sigma(x_a)}), X)$, that is, x is reached from the region X by running $p_{\sigma(x_a)}$ with time reversed. From this the claim $(\sigma, \text{execPlant}(p_{\sigma(x_a)}, x)) \models (\Phi, X)$ follows. \square

Proof. (Of Prop. 3.6) For soundness, first we observe that $(\sigma_0, x_0) \models (\Phi_{\text{init}}, X_{\text{init}})$. This is because $\text{BS}(\overrightarrow{(\sigma_s, m)})$ implies $(\Phi_{\text{init}}, X_{\text{init}}) = 0\text{-FA}(\Phi_{\text{init}}, X_{\text{init}})$ by Def. 3.4; see in particular (10).

Starting from (σ_0, x_0) , we can repeatedly apply Lem. B.2 to obtain input i_0, \dots, i_{T-1} and CP-states $(\sigma_1, x_1), \dots, (\sigma_T, x_T)$ such that: $(\sigma_0, x_0) \xrightarrow{i_0} \dots \xrightarrow{i_{T-1}} (\sigma_T, x_T)$; and $(\sigma_k, x_k) \models \text{BS}(\overrightarrow{(\sigma_s^{(T-1)}, m^{(T-1)})}, \dots, (\sigma_s^{(k)}, m^{(k)})$)

for each $k \in [0, T]$. Then in particular $(\sigma_T, x_T) \models \text{BS}(\varepsilon) = (\Phi_{\text{final}}, X_{\text{final}})$ where ε denotes the empty sequence. This means that (σ_0, x_0) and i_0, \dots, i_{T-1} qualify as an answer.

For completeness, let $(\sigma_0, x_0) \xrightarrow{i_0} \dots \xrightarrow{i_{T-1}} (\sigma_T, x_T)$ be a run of \mathcal{S} . We can take

$$\overrightarrow{(\sigma_s, m)} := \langle (s(x_{T-1}, i_{T-1}), \sigma_T(x_a)), \dots, (s(x_0, i_0), \sigma_1(x_a)) \rangle . \quad \square$$

B.3 Proof of Lem. 3.7

Proof. Assume the node (Φ, X) at the position $\overrightarrow{(\sigma_s, m)}$ has an unsatisfiable label. Consider its child at the position $\overrightarrow{(\sigma'_s, m')}$, for arbitrary $\sigma'_s \in \mathbb{B}_s^{\text{Var}}$ and $m' \in \text{Modes}$.

We know $\Phi = \text{false}$ or $X = \emptyset$ from Lem. A.3. In case $\Phi = \text{false}$, we easily see that

$$\pi_{\text{C}}(1\text{-BS}^{\text{pre}}(\sigma_s, m)(\Phi, X)) = \text{false} .$$

In case $X = \emptyset$, similarly

$$\pi_{\text{P}}(1\text{-BS}^{\text{pre}}(\sigma_s, m)(\Phi, X)) = \emptyset . \quad \square$$